

Geometric Group Theory

Cornelia Druțu and Michael Kapovich

With an Appendix by Bogdan Nica

To our families: Catalin and Alexandra, Jennifer and Esther

Preface

The goal of this book is to present several central topics in Geometric Group Theory, primarily related to the large scale geometry of infinite groups and of the spaces on which such groups act, and to illustrate them with fundamental theorems such as Gromov's Theorem on groups of polynomial growth, Tits' Alternative Theorem, Mostow's Rigidity Theorem, Stallings' theorem on ends of groups, theorems of Tukia and Schwartz on quasiisometric rigidity for lattices in real-hyperbolic spaces, etc. We give essentially self-contained proofs of all the above mentioned results, and we use the opportunity to describe several powerful tools/toolkits of Geometric Group Theory, such as coarse topology, ultralimits and quasiconformal mappings. We also discuss three classes of groups central in Geometric Group Theory: Amenable groups, hyperbolic groups, and groups with Property (T).

The key idea in Geometric Group Theory is to study groups by endowing them with a metric and treating them as *geometric objects*. This can be done for groups that are *finitely generated*, i.e. that can be reconstructed from a finite subset, *via* multiplication and inversion. Many groups naturally appearing in topology, geometry and algebra (e.g. fundamental groups of manifolds, groups of matrices with integer coefficients) are finitely generated. Given a finite generating set S of a group G , one can define a metric on G by constructing a connected graph, the *Cayley graph* of G , with G serving as set of vertices, and oriented edges joining elements in G that differ by a right multiplication with a generator s from S , and labeled by s . A Cayley graph \mathcal{G} , as any other connected graph, admits a natural metric invariant under automorphisms of \mathcal{G} : Edges are assumed to be of length one, and the distance between two points is the length of the shortest path in the graph joining these points (see Section 2.3). The restriction of this metric to the vertex set G is called the *word metric* dist_S on the group G . The first obstacle to "geometrizing" groups in this fashion is the fact that a Cayley graph depends not only on the group but also on a particular choice of finite generating set. Cayley graphs associated with different generating sets are not isometric but merely *quasiisometric*.

Another typical situation in which a group G is naturally endowed with a (pseudo-)metric is when G acts on a metric space X : In this case the group G maps to X *via* the *orbit map* $g \mapsto gx$. The pull-back of the metric to G is then a pseudo-metric on G . If G acts on X isometrically, then the resulting pseudo-metric on G is G -invariant. If, furthermore, the space X is proper and geodesic and the action of G is *geometric* (i.e. properly discontinuous and cocompact), then G is finitely generated and the resulting (pseudo-)metric is quasiisometric to word metrics on G (Theorem 8.37). For example, if a group G is the fundamental group of a closed Riemannian manifold M , the action of G on the universal cover \widetilde{M} of

M satisfies all these properties. The second class of examples of isometric actions (whose origin lies in functional analysis and representation theory) comes from isometric actions of a group G on a Hilbert space. The square of the corresponding pull-back (pseudo-)metric on G is known in the literature as a *conditionally negative semidefinite kernel*. In this case, if the group is not virtually abelian the action cannot be geometric. (Here and in what follows, when we say that a group has a certain property *virtually* we mean that it has a finite-index subgroup with that property.) On the other hand, the mere existence of a proper action of a group G on a Hilbert space H (i.e. an action so that, as $g \in G$ escapes every compact, the norm $\|gv\|$ diverges to infinity, where v is any vector in H), equivalently the mere existence of a conditionally negative semidefinite kernel on G that is proper as a topological map, has many interesting implications, detailed in Chapter 19.

In the setting of the geometric view of groups, the following questions become fundamental:

- QUESTIONS. (A) *If G and G' are quasiisometric groups, to what extent do G and G' share the same algebraic properties?*
- (B) *If a group G is quasiisometric to a metric space X , what geometric properties (or structures) on X translate to interesting algebraic properties of G ?*

Addressing these questions is the primary focus of this book. Several striking results (like Gromov’s Polynomial Growth Theorem) state that certain algebraic properties of a group can be reconstructed from its loose geometric features, and in particular must be shared by quasiisometric groups.

Closely connected to these considerations are two foundational problems which appeared in different contexts, but both render the same sense of existence of a “demarcation line” dividing the class of infinite groups into “abelian-like” groups and “free-like” groups. The invariants used to draw the line are quite different (existence of a finitely-additive invariant measure in one case and behavior of the growth function in the other); nevertheless, the two problems/questions and the classification results that grew out of these questions, have much in common.

The first of these questions was inspired by work investigating the existence of various types of group-invariant measures, that originally appeared in the context of Euclidean spaces. Namely, the *Banach-Tarski paradox* (see Chapter 17), while denying the existence of such measures on the Euclidean space, inspired John von Neumann to formulate two important concepts: That of *amenable groups* and that of *paradoxical decompositions and groups* [vN28]. In an attempt to connect amenability to the algebraic properties of a group, von Neumann made the observation, in the same paper, that the existence of a free subgroup excludes amenability. Mahlon Day (in [Day50] and [Day57]) extended von Neumann’s work, introduced the terminology *amenable groups*, defined the class of *elementary amenable groups* and proved several foundational results about amenable and elementary amenable groups. In [Day57, p. 520] he also noted¹:

¹Contrary to the common belief, Day neither formulated a conjecture about this issue nor attributed the problem to von Neumann.

- It is not known whether the class of elementary amenable groups equals the class of amenable groups and whether the class of amenable groups coincides with the class of groups containing no free non-abelian subgroups.

This observation later became commonly known as the von Neumann–Day problem (or conjecture):

QUESTION (The von Neumann–Day problem). *Is non-amenability of a group equivalent to the existence of a free non-abelian subgroup?*

The second problem appeared in the context of Riemannian geometry, in connection to attempts to relate, for a compact Riemannian manifold M , the geometric features of its universal cover \widetilde{M} to the behavior of its fundamental group $G = \pi_1(M)$. Two of the most basic objects in Riemannian geometry are the *volume* and the *volume growth rate*. The notion of volume growth extends naturally to discrete metric spaces, such as finitely generated groups. The *growth function* of a finitely generated group G (with a fixed finite generating set S) is the cardinality $\mathfrak{G}(n)$ of the ball of radius n in the metric space (G, dist_S) . While the function $\mathfrak{G}(n)$ depends on the choice of the finite generating set S , the *growth rate* of $\mathfrak{G}(n)$ is independent of S . In particular, one can speak of groups of linear, polynomial, exponential growth, etc. More importantly, the growth rate is preserved by quasi-isometries, which allows to establish a close connection between the Riemannian growth of a manifold \widetilde{M} as above, and the growth of $G = \pi_1(M)$.

One can easily see that every abelian group has polynomial growth. It is a more difficult theorem (proven independently by Hyman Bass [Bas72] and Yves Guivarc’h [Gui70, Gui73]) that all nilpotent groups also have polynomial growth. We provide a proof of this result in Section 14.2. In this context, John Milnor [Mil68c] and Joe Wolf [Wol68] asked the following question:

QUESTION. *Is it true that the growth of each finitely generated group is either polynomial (i.e. $\mathfrak{G}(n) \leq Cn^d$ for some fixed C and d) or exponential (i.e. $\mathfrak{G}(n) \geq Ca^n$ for some fixed $a > 1$ and $C > 0$)?*

Note that Milnor stated the problem in the form of a question, not a conjecture, however, he conjectured in [Mil68c] that each group of polynomial growth is virtually nilpotent.

The answer to the question is positive for *solvable groups*: This is the *Milnor–Wolf Theorem*, which moreover states that *solvable groups of polynomial growth are virtually nilpotent*, see Theorem 14.37 in this book (the theorem is a combination of results due to Milnor and Wolf). This theorem still holds for the larger class of *elementary amenable groups* (see Theorem 18.58); moreover, such groups with non-polynomial growth must contain a free non-abelian subsemigroup.

The proof of the Milnor–Wolf Theorem essentially consists of a careful examination of increasing/decreasing sequences of subgroups in nilpotent and solvable groups. Along the way, one discovers other features that nilpotent groups share with abelian groups, but not with solvable groups. For instance, *in a nilpotent group all finite subgroups are contained in a maximal finite subgroup, while solvable groups may contain infinite strictly increasing sequences of finite subgroups*. Furthermore, *all subgroups of a nilpotent group are finitely generated, but this is no longer true for solvable groups*. One step further into the study of a finitely generated subgroup H in a group G is to compare a word metric dist_H on the subgroup

H to the restriction to H of a word metric dist_G on the ambient group G . With an appropriate choice of generating sets, the inequality $\text{dist}_G \leq \text{dist}_H$ is immediate: All the paths in H joining $h, h' \in H$ are also paths in G , but there might be some other, shorter paths in G joining h, h' . The problem is to find an upper bound on dist_H in terms of dist_G . If G is abelian, the upper bound is linear as a function of dist_G . If dist_H is bounded by a polynomial in dist_G , then the subgroup H is said to be *polynomially distorted* in G , while if dist_H is approximately $\exp(\lambda \text{dist}_G)$ for some $\lambda > 0$, the subgroup H is said to be *exponentially distorted*. It turns out that *all subgroups in a nilpotent group are polynomially distorted, while some solvable groups contain finitely generated subgroups with exponential distortion*.

Both the von Neumann-Day and the Milnor-Wolf questions were answered in the affirmative for linear groups by Jacques Tits:

THEOREM (Tits' Alternative). *Let F be a field of zero characteristic and let Γ be a subgroup of $GL(n, F)$. Then either Γ is virtually solvable or Γ contains a free nonabelian subgroup.*

We prove Tits' Alternative in Chapter 15. Note that this alternative also holds for fields of positive characteristic, provided that Γ is finitely generated.

There are other classes of groups in which both the von Neumann-Day and the Milnor-Wolf questions have positive answers, they include: Subgroups of Gromov-hyperbolic groups ([Gro87, §8.2.F], [GdlH90, Chapter 8]), fundamental groups of closed Riemannian manifolds of nonpositive curvature [Bal95], subgroups of the mapping class groups of surfaces [Iva92], and of the groups of outer automorphisms of free groups [BFH00, BFH05].

The von Neumann-Day question in general has a negative answer: The first counterexamples were given by A. Ol'shanskiĭ in [Ol'80]. In [Ady82] it was shown that the free Burnside groups $B(n, m)$ with $n \geq 2$ and $m \geq 665$, m odd, are also counterexamples. The first finitely presented counterexamples were constructed by A. Ol'shanskiĭ and M. Sapir in [OS02]. Y. Lodha and J.T. Moore later provided, in [LM16], another finitely presented counterexample, a subgroup of the group of piecewise projective homeomorphisms of the real projective line, subgroup which is **torsion-free** (unlike the previous counterexamples, based precisely on the existence of a large torsion), and has an explicit presentation with three generators and nine relators. These papers have lead to the development of certain techniques of constructing "infinite finitely generated monsters". While the negation of amenability (i.e. the paradoxical behavior) is, thus, still not completely understood algebraically, several stronger properties implying nonamenability were introduced, among which are various fixed-point properties, most importantly *Kazhdan's Property (T)* (Chapter 19). Remarkably, amenability (hence paradoxical behavior) is a quasiisometry invariant, while Property (T) is not.

The Milnor-Wolf question, in full generality, likewise has a negative answer: The first groups of *intermediate growth*, i.e. growth which is super-polynomial but subexponential, were constructed by Rostislav Grigorchuk. Moreover, he proved the following:

THEOREM (Grigorchuk's Subexponential Growth theorem). *Let $f(n)$ be an arbitrary sub-exponential function larger than $2\sqrt{n}$. Then there exists a finitely*

generated group Γ with subexponential growth function $\mathfrak{G}(n)$ such that:

$$f(n) \leq \mathfrak{G}(n)$$

for infinitely many $n \in \mathbb{N}$.

Later on, Anna Erschler [Ers04] adapted Grigorchuk’s arguments to improve the above result with the inequality $f(n) \leq \mathfrak{G}(n)$ for all but finitely many n . In the above examples, the exact growth function was unknown. However, Laurent Bartholdi and Anna Erschler [BE12] constructed examples of groups of intermediate growth, where they actually compute $\mathfrak{G}(n)$, up to an appropriate equivalence relation. Note, however, that the Milnor–Wolf Problem is still open for finitely presented groups.

On the other hand, Mikhael Gromov proved an even more striking result:

THEOREM (Gromov’s Polynomial Growth Theorem, [Gro81a]). *Every finitely generated group of polynomial growth is virtually nilpotent.*

This is a typical example of an algebraic property that may be recognized *via* a, seemingly, weak geometric information. A corollary of Gromov’s theorem is *quasiisometric rigidity* for virtually nilpotent groups:

COROLLARY. *Suppose that G is a group quasiisometric to a nilpotent group. Then G itself is virtually nilpotent.*

Gromov’s theorem and its corollary will be proven in Chapter 16. Since the first version of this book was written, Bruce Kleiner [Kle10] and, later, Narutaka Ozawa [Oza15] gave completely different (and much shorter) proofs of Gromov’s polynomial growth theorem, using harmonic functions on graphs (Kleiner) and functional-analytic tools (Ozawa). Both proofs still require the Tits’ Alternative. Kleiner’s techniques provided the starting point for Y. Shalom and T. Tao, who proved the following effective version of Gromov’s Theorem [ST10]:

THEOREM (Shalom–Tao Effective Polynomial Growth Theorem). *There exists a constant C such that for any finitely generated group G and $d > 0$, if for some $R \geq \exp(\exp(Cd^C))$, the ball of radius R in G has at most R^d elements, then G has a finite index nilpotent subgroup of class less than C^d .*

It is also possible to prove Gromov’s Theorem without using the Tits alternative. Indeed, the proofs of either Gromov, Kleiner or Ozawa allow to restrict to the case of linear groups, and from there two different approaches are possible.

The first one is to use the well known remark that groups with subexponential growth are amenable (see Proposition 18.6), and the direct proof of Shalom [Sha98] of the fact that linear amenable groups are virtually solvable. The main ingredient in Shalom’s proof is a version of the Furstenberg lemma stating that, for any local field \mathbb{F} , the stabilizer in $PGL(n, \mathbb{F})$ of a probability measure on the projective space $\mathbb{F}P^{n-1}$ whose support is not included in a finite number of hyperplanes is a compact subgroup of $PGL(n, \mathbb{F})$. See also [Bre14].

The second approach is *via* simple additive combinatorics. E. Breuillard and B. Green have shown in [BG12] that if a finite subset A of the unitary group $U(n)$ satisfies $|A^3| < K|A|$ then A is contained in at most K^C cosets of an abelian subgroup of $U(n)$, where $K > 1$ is an arbitrary constant and $C = C(n)$ is independent of K and A . From this, it can be easily deduced that finitely generated

subgroups of $U(n)$ that have polynomial growth are virtually abelian; see [BG12, Proposition 5.1]. As Kleiner’s proof allows to restrict to the case of subgroups of the unitary group $U(n)$, this concludes the proof of Gromov’s Theorem. The advantage of this approach is that it is elementary: it relies on simple properties of compact Lie groups, and uses neither proximality nor amenability. The result of Breuillard-Green has been further generalized in their joint work with T. Tao [BGT11] to subsets A in Lie groups that are not compact. This improved result can be combined with either of the arguments of Gromov, Kleiner or Ozawa, reducing the problem to linear groups, to provide yet another proof of the Polynomial Growth Theorem avoiding the Tits alternative, less elementary though. Both additive combinatorics proofs have the further advantage that, unlike when using the Tits alternative or the proof of Shalom, one does not need to change field: The entire argument can be carried out in the setting of the real numbers.

We decided to retain, however, Gromov’s original proof since it contains a wealth of ideas that generated in their turn new areas of research. Remarkably, the same piece of logic (a weak version of the axiom of choice) that makes the Banach–Tarski paradox possible also allows to construct *ultralimits*, a powerful tool in the proof of Gromov’s theorem and that of many rigidity theorems (e.g, quasiisometric rigidity theorems of Kapovich, Kleiner and Leeb) as well as in the investigation of fixed point properties.

Regarding Questions (A) and (B), the best one can hope for is that the geometry of a group (up to quasiisometric equivalence) allows to recover, not just some of its algebraic features, but the group itself, *up to virtual isomorphism*. Two groups G_1 and G_2 are said to be *virtually isomorphic* if there exist subgroups

$$F_i \triangleleft H_i \leq G_i, i = 1, 2,$$

so that H_i has finite index in G_i , F_i is a finite normal subgroup in H_i , $i = 1, 2$, and H_1/F_1 is isomorphic to H_2/F_2 . Virtual isomorphism implies quasiisometry but, in general, the converse is false, see Example 8.48. In the situation when the converse implication also holds, one says that the group G_1 is *quasiisometrically rigid*.

An example of quasiisometric rigidity is given by the following theorem proven by Richard Schwartz [Sch96b]:

THEOREM (Schwartz QI rigidity theorem). *Suppose that Γ is a non-uniform lattice of isometries of the hyperbolic space \mathbb{H}^n , $n \geq 3$. Then each group quasiisometric to Γ must be virtually isomorphic to Γ .*

We will present a proof of this theorem in Chapter 24. In the same chapter we will use similar “zooming” arguments to prove the following special case of *Mostow’s Rigidity Theorem*:

THEOREM (The Mostow Rigidity Theorem). *Let Γ_1 and Γ_2 be lattices of isometries of \mathbb{H}^n , $n \geq 3$, and let $\varphi : \Gamma_1 \rightarrow \Gamma_2$ be a group isomorphism. Then φ is given by a conjugation via an isometry of \mathbb{H}^n .*

Note that Schwartz’ Theorem no longer holds for $n = 2$, where non-uniform lattices are virtually free. However, in this case, quasiisometric rigidity still holds as a corollary of Stallings’ Theorem on ends of groups:

THEOREM. *Let Γ be a group quasiisometric to a free group of finite rank. Then Γ is itself virtually free.*

This theorem will be proven in Chapter 20. We also prove:

THEOREM (Stallings “Ends of groups” theorem). *If G is a finitely generated group with infinitely many ends, then G splits as a graph of groups with finite edge-groups.*

In this book we give two proofs of the above theorem, which, while quite different, are both inspired by the original argument of Stallings. In Chapter 20 we prove Stallings’ theorem for *almost finitely presented* groups. This proof follows the ideas of Dunwoody, Jaco and Rubinstein: We will be using *minimal Dunwoody tracks*, where minimality is defined with respect to a certain hyperbolic metric on the presentation complex (unlike the combinatorial minimality used by Dunwoody). In Chapter 21, we will give another proof, which works for all finitely generated groups and follows a proof sketched by Gromov in [Gro87], using least energy harmonic functions. We decided to present both proofs, since they use different machinery (the first is more geometric and the second more analytical) and different (although related) geometric ideas.

In Chapter 20 we also prove the following:

THEOREM (Dunwoody’s Accessibility Theorem). *Let G be an almost finitely presented group. Then G is accessible, i.e. the decomposition process of G as a graph of groups with finite edge groups eventually terminates.*

In Chapter 23 we prove Tukia’s theorem, which establishes quasiisometric rigidity of the class of fundamental groups of compact hyperbolic n -manifolds, and, thus, complements Schwartz’ Theorem above:

THEOREM (Tukia’s QI Rigidity Theorem). *If a group Γ is quasiisometric to the hyperbolic n -space, then Γ is virtually isomorphic to the fundamental group of a compact hyperbolic n -manifold.*

Note that the proofs of the theorems of Mostow, Schwartz and Tukia all rely upon the same analytical tool: Quasiconformal mappings of Euclidean spaces. In contrast, the analytical proofs of Stallings’ theorem presented in the book are mostly motivated by another branch of geometric analysis, namely, the theory of minimal submanifolds and harmonic functions.

In regard to Question (B), we investigate two closely related classes of groups: Hyperbolic and relatively hyperbolic groups. These classes generalize fundamental groups of compact negatively curved Riemannian manifolds and, respectively, complete negatively curved Riemannian manifolds of finite volume. To this end, in Chapters 4 and 11 we cover the basics of hyperbolic geometry and the theory of hyperbolic and relatively hyperbolic groups.

Other sources. Our choice of topics in geometric group theory is far from exhaustive. We refer the reader to [Aea91], [Bal95], [Bow91], [VSCC92], [Bow06a], [BH99], [CDP90], [Dav08], [Geo08], [GdlH90], [dlH00], [NY11], [Pap03], [Roe03], [Sap14], [Väi05], for the discussion of other parts of the theory.

Work on this book started in 2002 and the material which we cover mostly concerns developments in Geometric Group Theory from the 1960s through the 1990s. In the meantime, while we were working on the book, some major exciting developments in the field have occurred which we did not have a chance

to discuss. To name a few, these developments are subgroup separability and its connections with 3-dimensional topology [Ago13, KM12, HW12, Bes14], applications of Geometric Group Theory to higher dimensional and coarse topology [Yu00, MY02, BLW10, BL12], the theory of Kleinian groups [Min10, BCM12, Mj14b, Mj14a], quasiconformal analysis on boundaries of hyperbolic groups and Cannon Conjecture [BK02a, BK05, Bon11, BK13, Mar13, Hai15], the theory of approximate groups [BG08a, Tao08, BGT12, Hru12], the first-order logic of free groups (see [Sel01, Sel03, Sel05a, Sel04, Sel05b, Sel06a, Sel06b, Sel09, Sel13] and [KM98b, KM98a, KM98c, KM05]), the theory of systolic groups [JŠ03, JŠ06, HŠ08, Osa13], probabilistic aspects of Geometric Group Theory [Gro03, Ghy04, Oll04, Oll05, KSS06, Oll07, KS08, OW11, AŁŚ15, DM16].

Requirements. The book is intended as a reference for graduate students and more experienced researchers, it can be used as a basis for a graduate course and as a first reading for a researcher wishing to learn more about Geometric Group Theory. This book is partly based on lectures given at the Oxford University (C.D.), the University of Utah and the University of California, Davis (M.K.). We expect the reader to be familiar with the basics of group theory, algebraic topology (fundamental groups, covering spaces, (co)homology, Poincaré duality) and elements of differential topology and Riemannian geometry. Some of the background material is covered in Chapters 1, 3 and 5. We tried to make the book as self-contained as possible, but some theorems are stated without proof, they are marked as **Theorem**.

Acknowledgments. The work of the first author was supported in part by the ANR projects “Groupe de recherche de Géométrie et Probabilités dans les Groupes”, by the EPSRC grant “Geometric and analytic aspects of infinite groups”, by the ANR-10-BLAN 0116, acronym GGAA, and by the Labex CEMPI. The second author was supported by NSF grants DMS-02-03045, DMS-04-05180, DMS-05-54349, DMS-09-05802, DMS-12-05312 and DMS-16-04241, as well as by the Korea Institute for Advanced Study (KIAS) through the KIAS scholar program. During the work on this book, the authors visited the Max Planck Institute for Mathematics (Bonn) and the Mathematical Sciences Research Institute (Berkeley). We wish to thank these institutions for their hospitality, and the organizers of the program “Geometric Group Theory”, held at MSRI, August to December 2016, for offering us the stimulating working environment of the program while finishing this book. The first author would also like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the program “Non-positive curvature, group actions and cohomology”, funded by EPSRC grant no. EP/KO32208/1, where work on this book was finalized. Our special thanks are to Elia Fioravanti and Federico Vigolo for their careful reading and consistent help with corrections, comments and picture drawing. We are grateful to Thomas Delzant, Bruce Kleiner and Mohan Ramachandran for supplying some proofs, to Jason Behrstock, Emmanuel Breuillard, Pierre-Emmanuel Caprace, Yves de Cornulier, Francois Dahmani, David Hume, Ilya Kapovich, Bernhard Leeb, Andrew Sale, Mark Sapir, Alessandro Sisto and the anonymous referees of this book for the numerous corrections and suggestions. We also thank Yves de Cornulier, Rostislav Grigorchuk, Pierre Pansu and Romain Tessera for useful references and to Bogdan Nica for writing an appendix to the book.

Cornelia Druţu: Mathematical Institute,
Andrew Wiles Building,
University of Oxford,
Woodstock Road,
Oxford OX2 6GG, United Kingdom
`Cornelia.Drutu@maths.ox.ac.uk`

Michael Kapovich: Department of Mathematics,
University of California,
Davis, CA 95616, USA
`kapovich@math.ucdavis.edu`

Contents

Preface	iii
Chapter 1. Geometry and Topology	1
1.1. Set-theoretic preliminaries	1
1.1.1. General notation	1
1.1.2. Growth rates of functions	2
1.1.3. Jensen's inequality	3
1.2. Measure and integral	3
1.2.1. Measures	3
1.2.2. Integrals	5
1.3. Topological spaces. Lebesgue covering dimension	7
1.4. Exhaustions of locally compact spaces	10
1.5. Direct and inverse limits	11
1.6. Graphs	13
1.7. Complexes and homology	17
1.7.1. Simplicial complexes	17
1.7.2. Cell complexes	19
Chapter 2. Metric spaces	23
2.1. General metric spaces	23
2.2. Length metric spaces	25
2.3. Graphs as length spaces	27
2.4. Hausdorff and Gromov–Hausdorff distances. Nets	28
2.5. Lipschitz maps and Banach-Mazur distance	30
2.5.1. Lipschitz and locally Lipschitz maps	30
2.5.2. Bi-Lipschitz maps. The Banach-Mazur distance	33
2.6. Hausdorff dimension	34
2.7. Norms and valuations	35
2.8. Norms on field extensions. Adeles	39
2.9. Metrics on affine and projective spaces	43
2.10. Quasiprojective transformations. Proximal transformations	48
2.11. Kernels and distance functions	51
Chapter 3. Differential geometry	59
3.1. Smooth manifolds	59
3.2. Smooth partition of unity	61
3.3. Riemannian metrics	61
3.4. Riemannian volume	64
3.5. Volume growth and isoperimetric functions. Cheeger constant	67
3.6. Curvature	71

3.7.	Riemannian manifolds of bounded geometry	72
3.8.	Metric simplicial complexes of bounded geometry and systolic inequalities	74
3.9.	Harmonic functions	78
3.10.	Spectral interpretation of the Cheeger constant	81
3.11.	Comparison geometry	82
3.11.1.	Alexandrov curvature and $CAT(\kappa)$ spaces	82
3.11.2.	Cartan's fixed point theorem	86
3.11.3.	Ideal boundary, horoballs and horospheres	87
Chapter 4.	Hyperbolic Space	91
4.1.	Moebius transformations	91
4.2.	Real hyperbolic space	94
4.3.	Classification of isometries	99
4.4.	Hyperbolic trigonometry	103
4.5.	Triangles and curvature of \mathbb{H}^n	105
4.6.	Distance function on \mathbb{H}^n	108
4.7.	Hyperbolic balls and spheres	110
4.8.	Horoballs and horospheres in \mathbb{H}^n	110
4.9.	\mathbb{H}^n as a symmetric space	112
4.10.	Inscribed radius and thinness of hyperbolic triangles	116
4.11.	Existence-uniqueness theorem for triangles	118
Chapter 5.	Groups and their actions	121
5.1.	Subgroups	121
5.2.	Virtual isomorphisms of groups and commensurators	124
5.3.	Commutators and the commutator subgroup	126
5.4.	Semidirect products and short exact sequences	128
5.5.	Direct sums and wreath products	130
5.6.	Geometry of group actions	131
5.6.1.	Group actions	131
5.6.2.	Linear actions	135
5.6.3.	Lie groups	135
5.6.4.	Haar measure and lattices	139
5.6.5.	Geometric actions	141
5.7.	Zariski topology and algebraic groups	142
5.8.	Group actions on complexes	149
5.8.1.	G -complexes	149
5.8.2.	Borel and Haefliger constructions	150
5.8.3.	Groups of finite type	160
5.9.	Cohomology	161
5.9.1.	Group rings and modules	161
5.9.2.	Group cohomology	162
5.9.3.	Bounded cohomology of groups	166
5.9.4.	Ring derivations	167
5.9.5.	Derivations and split extensions	169
5.9.6.	Central coextensions and second cohomology	172
Chapter 6.	Median spaces and spaces with measured walls	177

6.1.	Median spaces	177
6.1.1.	A review of median algebras	178
6.1.2.	Convexity	179
6.1.3.	Examples of median metric spaces	180
6.1.4.	Convexity and gate property in median spaces	182
6.1.5.	Rectangles and parallel pairs	184
6.1.6.	Approximate geodesics and medians; completions of median spaces	187
6.2.	Spaces with measured walls	188
6.2.1.	Definition and basic properties	188
6.2.2.	Relationship between median spaces and spaces with measured walls	191
6.2.3.	Embedding a space with measured walls in a median space	192
6.2.4.	Median spaces have measured walls	194
Chapter 7. Finitely generated and finitely presented groups		201
7.1.	Finitely generated groups	201
7.2.	Free groups	205
7.3.	Presentations of groups	208
7.4.	The rank of a free group determines the group. Subgroups	214
7.5.	Free constructions: Amalgams of groups and graphs of groups	215
7.5.1.	Amalgams	215
7.5.2.	Graphs of groups	216
7.5.3.	Converting graphs of groups into amalgams	218
7.5.4.	Topological interpretation of graphs of groups	218
7.5.5.	Constructing finite-index subgroups	219
7.5.6.	Graphs of groups and group actions on trees	221
7.6.	Ping-pong lemma. Examples of free groups	224
7.7.	Free subgroups in $SU(2)$	227
7.8.	Ping-pong on projective spaces	228
7.9.	Cayley graphs	229
7.10.	Volumes of maps of cell complexes and Van Kampen diagrams	237
7.10.1.	Simplicial, cellular and combinatorial volumes of maps	237
7.10.2.	Topological interpretation of finite-presentability	238
7.10.3.	Presentations of central coextensions	238
7.10.4.	Dehn function and van Kampen diagrams	240
7.11.	Residual finiteness	246
7.12.	Hopfian and cohopfian properties	249
7.13.	Algorithmic problems in the combinatorial group theory	250
Chapter 8. Coarse geometry		253
8.1.	Quasi-isometry	253
8.2.	Group-theoretic examples of quasiisometries	263
8.3.	A metric version of the Milnor-Schwarz Theorem	269
8.4.	Topological coupling	271
8.5.	Quasiactions	273
8.6.	Quasi-isometric rigidity problems	275
8.7.	The growth function	277
8.8.	Codimension one isoperimetric inequalities	282
8.9.	Distortion of a subgroup in a group	285

Chapter 9. Coarse topology	287
9.1. Ends	287
9.1.1. The number of ends	287
9.1.2. The space of ends	290
9.1.3. Ends of groups	295
9.2. Rips complexes and coarse homotopy theory	297
9.2.1. Rips complexes	297
9.2.2. Direct system of Rips complexes and coarse homotopy	299
9.3. Metric cell complexes	300
9.4. Connectivity and coarse connectivity	305
9.5. Retractions	312
9.6. Poincaré duality and coarse separation	314
9.7. Metric filling functions	317
9.7.1. Coarse isoperimetric functions and coarse filling radius	318
9.7.2. Quasi-isometric invariance of coarse filling functions	320
9.7.3. Higher Dehn functions	325
9.7.4. Coarse Besikovitch inequality	329
Chapter 10. Ultralimits of Metric Spaces	333
10.1. The Axiom of Choice and its weaker versions	333
10.2. Ultrafilters and the Stone–Čech compactification	339
10.3. Elements of nonstandard algebra	340
10.4. Ultralimits of families of metric spaces	344
10.5. Completeness of ultralimits and incompleteness of ultrafilters	348
10.6. Asymptotic cones of metric spaces	352
10.7. Ultralimits of asymptotic cones are asymptotic cones	356
10.8. Asymptotic cones and quasiisometries	358
10.9. Assouad–type theorems	359
Chapter 11. Gromov–hyperbolic spaces and groups	363
11.1. Hyperbolicity according to Rips	363
11.2. Geometry and topology of real trees	367
11.3. Gromov hyperbolicity	368
11.4. Ultralimits and stability of geodesics in Rips–hyperbolic spaces	372
11.5. Local geodesics in hyperbolic spaces	376
11.6. Quasiconvexity in hyperbolic spaces	379
11.7. Nearest-point projections	381
11.8. Geometry of triangles in Rips–hyperbolic spaces	382
11.9. Divergence of geodesics in hyperbolic metric spaces	385
11.10. Morse Lemma revisited	387
11.11. Ideal boundaries	390
11.12. Gromov bordification of Gromov–hyperbolic spaces	397
11.13. Boundary extension of quasiisometries of hyperbolic spaces	402
11.13.1. Extended Morse Lemma	402
11.13.2. The extension theorem	404
11.13.3. Boundary extension and quasiactions	406
11.13.4. Conical limit points of quasiactions	406
11.14. Hyperbolic groups	407
11.15. Ideal boundaries of hyperbolic groups	410

11.16.	Linear isoperimetric inequality and Dehn algorithm for hyperbolic groups	414
11.17.	The small cancellation theory	417
11.18.	The Rips construction	417
11.19.	Central coextensions of hyperbolic groups and quasiisometries	419
11.20.	Characterization of hyperbolicity using asymptotic cones	422
11.21.	Size of loops	428
11.21.1.	The minsize	428
11.21.2.	The constriction	430
11.22.	Filling invariants of hyperbolic spaces	432
11.22.1.	Filling area	432
11.22.2.	Filling radius	433
11.22.3.	Orders of Dehn functions of non-hyperbolic groups and higher Dehn functions	436
11.23.	Asymptotic cones, actions on trees and isometric actions on hyperbolic spaces	437
11.24.	Summary of equivalent definitions of hyperbolicity	440
11.25.	Further properties of hyperbolic groups	441
11.26.	Relatively hyperbolic spaces and groups	444
Chapter 12.	Lattices in Lie groups	447
12.1.	Semisimple Lie groups and their symmetric spaces	447
12.2.	Lattices	449
12.3.	Examples of lattices	450
12.4.	Rigidity and superrigidity	452
12.5.	Commensurators of lattices	454
12.6.	Lattices in $PO(n, 1)$	454
12.6.1.	Zariski density	454
12.6.2.	Parabolic elements and non-compactness	456
12.6.3.	Thick-thin decomposition	457
12.7.	Central coextensions	459
Chapter 13.	Solvable groups	463
13.1.	Free abelian groups	463
13.2.	Classification of finitely generated abelian groups	466
13.3.	Automorphisms of \mathbb{Z}^n	469
13.4.	Nilpotent groups	471
13.5.	Polycyclic groups	482
13.6.	Solvable groups: Definition and basic properties	487
13.7.	Free solvable groups and Magnus embedding	489
13.8.	Solvable versus polycyclic	491
Chapter 14.	Geometric aspects of solvable groups	495
14.1.	Wolf's Theorem for semidirect products $\mathbb{Z}^n \rtimes \mathbb{Z}$	495
14.1.1.	Geometry of $H_3(\mathbb{Z})$	497
14.1.2.	Distortion of subgroups of solvable groups	501
14.1.3.	Distortion of subgroups in nilpotent groups	503
14.2.	Polynomial growth of nilpotent groups	511
14.3.	Wolf's Theorem	513

14.4.	Milnor's theorem	515
14.5.	Failure of QI rigidity for solvable groups	518
14.6.	Virtually nilpotent subgroups of $GL(n)$	519
14.7.	Discreteness and nilpotence in Lie groups	522
14.7.1.	Some useful linear algebra	522
14.7.2.	Zassenhaus neighborhoods	523
14.7.3.	Jordan's theorem	526
14.8.	Virtually solvable subgroups of $GL(n, \mathbb{C})$	528
Chapter 15.	The Tits Alternative	535
15.1.	Outline of the proof	536
15.2.	Separating sets	538
15.3.	Proof of existence of free subsemigroup	539
15.4.	Existence of very proximal elements: Proof of Theorem 15.6	539
15.4.1.	Proximality criteria	540
15.4.2.	Constructing very proximal elements	541
15.5.	Finding ping-pong partners: Proof of Theorem 15.7	543
15.6.	The Tits Alternative without finite generation assumption	544
15.7.	Groups satisfying the Tits Alternative	545
Chapter 16.	Gromov's Theorem	547
16.1.	Topological transformation groups	547
16.2.	Regular Growth Theorem	549
16.3.	Consequences of the Regular Growth Theorem	553
16.4.	Weakly polynomial growth	554
16.5.	Displacement function	555
16.6.	Proof of Gromov's theorem	556
16.7.	Quasi-isometric rigidity of nilpotent and abelian groups	559
16.8.	Further developments	560
Chapter 17.	The Banach–Tarski paradox	563
17.1.	Paradoxical decompositions	563
17.2.	Step 1: A paradoxical decomposition of the free group F_2	566
17.3.	Step 2: The Hausdorff paradox	567
17.4.	Step 3: Spheres of dimension ≥ 2 are paradoxical	568
17.5.	Step 4: Euclidean unit balls are paradoxical	569
Chapter 18.	Amenability and paradoxical decomposition.	571
18.1.	Amenable graphs	571
18.2.	Amenability and quasiisometry	576
18.3.	Amenability of groups	581
18.4.	Følner property	586
18.5.	Amenability, paradoxality and the Følner property	589
18.6.	Supramenability and weakly paradoxical actions	593
18.7.	Quantitative approaches to non-amenability and weak paradoxality	599
18.8.	Uniform amenability and ultrapowers	604
18.9.	Quantitative approaches to amenability	606
18.10.	Summary of equivalent definitions of amenability	609
18.11.	Amenable hierarchy	610

Chapter 19. Ultralimits, fixed point properties, proper actions	613
19.1. Classes of Banach spaces stable with respect to ultralimits	613
19.2. Limit actions and point-selection theorem	618
19.3. Properties for actions on Hilbert spaces	622
19.4. Kazhdan's Property (T) and the Haagerup property	624
19.5. Groups acting on trees do not have Property (T)	631
19.6. Property FH, a-T-menability, and group actions on median spaces	634
19.7. Fixed point property and proper actions for L^p -spaces	637
19.8. Groups satisfying Property (T) and the spectral gap	639
19.9. Failure of quasiisometric invariance of Property (T)	641
19.10. Summary of examples	642
Chapter 20. Stallings Theorem and accessibility	645
20.1. Maps to trees and hyperbolic metrics on 2-dimensional simplicial complexes	645
20.2. Transversal graphs and Dunwoody tracks	650
20.3. Existence of minimal Dunwoody tracks	654
20.4. Properties of minimal tracks	657
20.4.1. Stationarity	657
20.4.2. Disjointness of essential minimal tracks	658
20.5. Stallings Theorem for almost finitely presented groups	662
20.6. Accessibility	664
20.7. QI rigidity of virtually free groups and free products	669
Chapter 21. Proof of Stallings' Theorem using harmonic functions	673
21.1. Proof of Stallings' theorem	675
21.2. Nonamenability	679
21.3. An existence theorem for harmonic functions	681
21.4. Energy of minimum and maximum of two smooth functions	683
21.5. A compactness theorem for harmonic functions	684
21.5.1. Positive energy gap implies existence of an energy minimizer	684
21.5.2. Some coarea estimates	688
21.5.3. Energy comparison in the case of a linear isoperimetric inequality	690
21.5.4. Proof of positivity of the energy gap	692
Chapter 22. Quasiconformal mappings	695
22.1. Linear algebra and eccentricity of ellipsoids	696
22.2. Quasisymmetric maps	697
22.3. Quasiconformal maps	698
22.4. Analytical properties of quasiconformal mappings	699
22.4.1. Some notion and results from real analysis	700
22.4.2. Differentiability properties of quasiconformal mappings	703
22.5. Quasisymmetric maps and hyperbolic geometry	709
Chapter 23. Groups quasi-isometric to \mathbb{H}^n	715
23.1. Uniformly quasiconformal groups	716
23.2. Hyperbolic extension of uniformly quasiconformal groups	717
23.3. Least volume ellipsoids	718
23.4. Invariant measurable conformal structure	719
23.5. Quasiconformality in dimension 2	722

23.5.1.	Beltrami equation	722
23.5.2.	Measurable Riemannian metrics	723
23.6.	Proof of Tukia's theorem on uniformly quasiconformal groups	724
23.7.	QI rigidity for surface groups	727
Chapter 24.	Quasiisometries of non-uniform lattices in \mathbb{H}^n	731
24.1.	Coarse topology of truncated hyperbolic spaces	732
24.2.	Hyperbolic extension	736
24.3.	Mostow Rigidity Theorem	737
24.4.	Zooming in	741
24.5.	Inverted linear mappings	743
24.6.	Scattering	746
24.7.	Schwartz Rigidity Theorem	748
Chapter 25.	A survey of quasiisometric rigidity	751
25.1.	Rigidity of symmetric spaces, lattices and hyperbolic groups	751
25.1.1.	Uniform lattices	751
25.1.2.	Non-uniform lattices	752
25.1.3.	Symmetric spaces with Euclidean de Rham factors and Lie groups with nilpotent normal subgroups	754
25.1.4.	QI rigidity for hyperbolic spaces and groups	755
25.1.5.	Failure of QI rigidity	758
25.1.6.	Rigidity of random groups	760
25.2.	Rigidity of relatively hyperbolic groups	760
25.3.	Rigidity of classes of amenable groups	762
25.4.	Bilipschitz vs. quasiisometric	765
25.5.	Various other QI rigidity results and problems	766
Chapter 26.	Appendix by Bogdan Nica: Three theorems on linear groups	775
26.1.	Introduction	775
26.2.	Virtual and residual properties of groups	776
26.3.	Platonov's theorem	776
26.4.	Proof of Platonov's theorem	777
26.5.	The Idempotent Conjecture for linear groups	780
26.6.	Proof of Formanek's criterion	781
26.7.	Notes	783
	Bibliography	785

Geometry and Topology

Treating groups as *geometric objects* is the major theme and defining feature of Geometric Group Theory. In this chapter and Chapter 2 we discuss basics of metric (and topological) spaces, while Chapter 3 will contain a brief overview of Riemannian geometry. For an in-depth discussion of metric geometry, we refer the reader to [BBI01]. We assume basic knowledge of Algebraic Topology as can be found, for instance, in [Hat02] or [Mas91].

1.1. Set-theoretic preliminaries

1.1.1. General notation. Given a set X we denote by $\mathcal{P}(X) = 2^X$ the *power set* of X , i.e. the set of all subsets of X . If two subsets A, B in X have the property that $A \cap B = \emptyset$ then we denote their union by $A \sqcup B$, and we call it the *disjoint union*. For a subset E of a set X we denote the complement of E in X either by $X \setminus E$ or by E^c . A *pointed set* is a pair (X, x) , where x is an element of X . The composition of two maps $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ is denoted either by $g \circ f$ or by gf . The identity map $X \rightarrow X$ will be denoted either by Id_X or simply by Id (when the choice of X is clear). For a map $f : X \rightarrow Y$ and a subset $A \subset X$, we let $f|_A$ denote the restriction of f to A . We use the notation $|E|$ or $\text{card}(E)$ to denote the cardinality of a set E . (Sometimes, however, $|E|$ will denote the Lebesgue measure of a subset of the Euclidean space.) We use the notation \mathbb{D}^n for the closed unit ball centered at the origin in the n -dimensional Euclidean space, and \mathbb{S}^{n-1} for the corresponding unit sphere. In contrast, we use the notation $B(x, r)$ for the open metric ball (in a general metric space) centered at x , of radius r . Accordingly, \mathbb{B}^n will denote the open unit ball in \mathbb{R}^n .

The Axiom of Choice (AC) plays a prominent part in many of the arguments in this book. We discuss it in more detail in section 10.1, where we also list equivalent and weaker forms of AC. Throughout the book we make the following convention:

CONVENTION 1.1. We always assume ZFC: The Zermelo–Fraenkel axioms of set theory and the Axiom of Choice.

Given a non-empty set X , we denote by $\text{Bij}(X)$ the group of bijections $X \rightarrow X$, with composition as the binary operation.

CONVENTION 1.2. Throughout the paper we let $\mathbf{1}_A$ and χ_A denote the characteristic (or indicator) function of a subset A in a set X , i.e. the function $\mathbf{1}_A : X \rightarrow \{0, 1\}$ defined by

$$\mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A. \end{cases}$$

By the *codimension of a subspace X in a space Y* we mean the difference between the dimension of Y and the dimension of X , whatever the notion of dimension that we use.

We will use the notation \cong to denote an isometry of metric spaces and \simeq to denote an isomorphism of groups.

Throughout the book, \mathbb{N} will denote the set of natural numbers and $\mathbb{Z}_+ = \mathbb{N} \cup \{0\}$.

1.1.2. Growth rates of functions. In this book we will be using two different *asymptotic* inequalities and equivalence relations for functions: One is used to compare Dehn functions of groups and the other to compare growth rates of groups.

DEFINITION 1.3. Let X be a subset of \mathbb{R} . Given two functions $f, g : X \rightarrow \mathbb{R}$, we say that *the order of the function f is at most the order of the function g* and we write $f \lesssim g$, if there exist real numbers $a, b, c, d, e > 0$ and x_0 such that for all $x \in X, x \geq x_0$, we have: $bx + c \in X$ and

$$f(x) \leq ag(bx + c) + dx + e.$$

If $f \lesssim g$ and $g \lesssim f$ then we write $f \approx g$ and we say that f and g are *approximately equivalent*.

This definition will be typically used with $X = \mathbb{R}_+$ or $X = \mathbb{N}$, in which case a, b, c, d, e will be natural numbers.

The equivalence class of a function with respect to the relation \approx is called *the order of the function*. If a function f has (at most) the same order as the function x, x^2, x^3, x^d or $\exp(x)$ it is said that *the order of the function f is (at most) linear, quadratic, cubic, polynomial, or exponential*, respectively. A function f is said to have *subexponential* order if it has order at most $\exp(x)$ and is not approximately equivalent to $\exp(x)$. A function f is said to have *intermediate order* if it has subexponential order and $x^n \lesssim f(x)$ for every n .

DEFINITION 1.4. We introduce the following *asymptotic inequality* between functions $f, g : X \rightarrow \mathbb{R}$ with $X \subset \mathbb{R}$: We write $f \preceq g$ if there exist $a, b > 0$ and $x_0 \in \mathbb{R}$ such that for all $x \in X, x \geq x_0$, we have: $bx \in X$ and

$$f(x) \leq ag(bx).$$

If $f \preceq g$ and $g \preceq f$ then we write $f \asymp g$ and we say that f and g are *asymptotically equal*.

Note that this definition is more refined than the *order notion* \approx . For instance, $x \approx 0$ while these functions are not asymptotically equal. This situation arises, for instance, in the case of free groups (which are given free presentation): The Dehn function is zero, while the area filling function of the Cayley graph is $A(\ell) \asymp \ell$. The equivalence relation \approx is more appropriate for Dehn functions than the relation \asymp , because in the case of a free group one may consider either a presentation with no relations, in which case the Dehn function is zero, or another presentation that yields a linear Dehn function.

- EXERCISE 1.5. 1. Show that \approx and \asymp are equivalence relations.
2. Suppose that $x \preceq f, x \preceq g$. Then $f \approx g$ if and only if $f \asymp g$.

1.1.3. Jensen's inequality. Let (X, μ) be a space equipped with a probability measure μ , $f : X \rightarrow \mathbb{R}$ a measurable function and φ a convex function defined on the range of f . Jensen's inequality [Rud87, Theorem 3.3] reads:

$$\varphi\left(\int_X f d\mu\right) \leq \int_X \varphi \circ f d\mu.$$

We will be using this inequality when the function f is strictly positive and $\varphi(t) = t^{-1}$: The function $\varphi(t)$ is convex for $t > 0$. It will be convenient to eliminate the probability measure assumption. We will be working with spaces (X, μ) of finite (but non-zero) measure. Instead of normalizing the measure μ to be a probability measure, we can as well replace integrals $\int_X h d\mu$ with *averages*

$$\int_X h d\mu = \frac{1}{M} \int_X h d\mu,$$

where $M = \int_X d\mu$. With this in mind, Jensen's inequality becomes

$$\left(\int_X f d\mu\right)^{-1} \leq \int_X \frac{1}{f} d\mu.$$

Replacing f with $\frac{1}{f}$ we also obtain:

$$(1.1) \quad \left(\int_X \frac{1}{f} d\mu\right)^{-1} \leq \int_X f d\mu.$$

1.2. Measure and integral

1.2.1. Measures. We recall the relevant definitions from the theory of measure spaces. A reference is [Bau01], whose terminology we adopt here. Let X be a non-empty set.

DEFINITION 1.6. A *ring of subsets* of X is a subset \mathcal{R} of $\mathcal{P}(X)$ containing the empty set, closed with respect to finite unions and differences.

An *algebra of subsets* of X is a non-empty collection \mathcal{A} of subsets of X such that:

- (1) $X \in \mathcal{A}$;
- (2) $A, B \in \mathcal{A} \Rightarrow A \cup B \in \mathcal{A}, A \cap B \in \mathcal{A}$;
- (3) $A \in \mathcal{A} \Rightarrow A^c \in \mathcal{A}$.

A σ -*algebra* of subsets of X is an algebra of subsets closed under countable intersections and countable unions.

Given a topological space X , the smallest σ -algebra of subsets of X containing all open subsets is called the *Borel σ -algebra* of X . Elements of this σ -algebra are called *Borel subsets* of X .

DEFINITION 1.7. A *finitely additive (f. a.) measure* μ on a ring \mathcal{R} is a function $\mu : \mathcal{R} \rightarrow [0, \infty]$ such that $\mu(A \sqcup B) = \mu(A) + \mu(B)$ for all $A, B \in \mathcal{R}$.

An immediate consequence of the f. a. property is that when \mathcal{R} is an algebra \mathcal{A} , for any two sets $A, B \in \mathcal{A}$,

$$\mu(A \cup B) = \mu((A \setminus B) \sqcup (A \cap B) \sqcup (B \setminus A)) = \mu(A \setminus B) + \mu(A \cap B) + \mu(B \setminus A) \leq \mu(A) + \mu(B).$$

In some texts the f. a. measures are called simply ‘measures’. We prefer the terminology above, since in other texts a ‘measure’ is meant to be countably additive as defined below.

DEFINITION 1.8. Given a ring \mathcal{R} , a *countably additive (c. a.) premeasure* on it is a function $\mu : \mathcal{R} \rightarrow [0, +\infty]$ such that

(M_1) for any sequence of pairwise disjoint sets $(A_n)_{n \in \mathbb{N}}$ in \mathcal{R} such that $\bigsqcup_{n \in \mathbb{N}} A_n \in \mathcal{R}$,

$$\mu \left(\bigsqcup_{n \in \mathbb{N}} A_n \right) = \sum_{n \in \mathbb{N}} \mu(A_n).$$

A premeasure is called *σ -finite* if there exists a sequence (A_n) in \mathcal{R} such that $\mu(A_n) < +\infty$ for every n , and $\bigcup_n A_n = X$. A premeasure defined on a σ -algebra is called a *countably additive (c. a.) measure*.

Property (M_1) is equivalent to the following list of two properties:

(M'_1) $\mu(A \sqcup B) = \mu(A) + \mu(B)$;

(M''_1) If $(A_n)_{n \in \mathbb{N}}$ is a non-increasing sequence of sets in \mathcal{R} such that $\bigcap_{n \in \mathbb{N}} A_n = \emptyset$, then $\lim_{n \rightarrow \infty} \mu(A_n) = 0$.

In order to simplify the terminology, we will suppress the dependence of the f.a. (resp. c.a) measure μ on the algebra (resp. σ -algebra) of subsets of X , and will refer to such a μ simply as a f.a. (resp. c.a.) measure on X .

DEFINITION 1.9. If μ is a finitely (resp. countably) additive measure on X , such that $\mu(X) = 1$, then μ is called a f.a. (resp. c.a.) *probability measure* on X , which is abbreviated as f.a.p. measure (resp. c.a.p. measure).

Suppose that G is a group acting on X preserving an algebra (resp. σ -algebra) \mathcal{A} . If μ is a f.a. (resp. c.a.) measure on \mathcal{A} , such that $\mu(\gamma A) = \mu(A)$ for all $\gamma \in G$ and $A \in \mathcal{A}$, then μ is called *G -invariant*.

We will need a precise version of the Carathéodory’s Theorem on the extension of a premeasure μ to a measure, therefore we recall here the notion of an outer measure.

DEFINITION 1.10. Let μ be a c.a. premeasure defined on a ring \mathcal{R} .

For every $Q \subset X$ let $\mathcal{U}(Q)$ designate the set of all sequences (A_n) in \mathcal{R} such that $Q \subset \bigcup_n A_n$. Define $\mu^*(Q) = +\infty$ if $\mathcal{U}(Q) = \emptyset$; if $\mathcal{U}(Q) \neq \emptyset$ then define

$$\mu^*(Q) = \inf \left\{ \sum_{n=1}^{\infty} \mu(A_n) ; (A_n) \in \mathcal{U}(Q) \right\}.$$

The function μ^* is an *outer measure* on the set X .

A subset A of X is called *μ^* -measurable* if for every $Q \in \mathcal{P}(X)$,

$$\mu^*(Q) = \mu^*(Q \cap A) + \mu^*(Q \cap A^c).$$

THEOREM 1.11 (Carathéodory [Bau01], §I.5). (1) *The collection \mathcal{A}^* of*

μ^ -measurable sets is a σ -algebra containing \mathcal{R} , and the restriction of μ^* to \mathcal{A}^* is a measure, while the restriction of μ^* to \mathcal{R} coincides with μ .*

(2) *If μ is σ -finite, then it has a unique extension to a measure on the σ -algebra generated by \mathcal{R} .*

1.2.2. Integrals. We let $B(X)$ denote the vector space of real-valued bounded functions on a set X . In addition to measures we will need the notion of *finitely additive integral*. We discuss integrals of functions $f \in B(X)$, only in the simpler case of finitely additive probability measures μ , defined on the algebra $\mathcal{A} = \mathcal{P}(X)$ (the setting where we will use finitely additive integrals, in Chapter 18). We refer the reader to [DS88] for an exposition of finitely additive integrals in greater generality.

A *finitely additive integral* on $(X, \mathcal{A}, \mu, B(X))$ is a linear functional

$$f \mapsto \int_X f d\mu, \quad f \in B(X), \quad \int_X f d\mu \in \mathbb{R},$$

satisfying the following properties:

- If $f(x) \geq 0$ for all $x \in X$, then $\int_X f d\mu \geq 0$.
- $\int_X \mathbf{1}_A d\mu = \mu(A)$ for all $A \in \mathcal{A}$.

For a subgroup $G \leq \text{Bij}(X)$, the integral \int_X is said to be *G-invariant* if

$$\int_X f \circ \gamma d\mu = \int_X f d\mu$$

for every $\gamma \in G$ and every $f \in B(X)$.

THEOREM 1.12. *If μ is a G-invariant f.a.p. measure on the algebra $\mathcal{P}(X)$, then there exists a G-invariant integral \int_X on $(X, \mathcal{A}, \mu, B(X))$ such that for every A*

$$\int_X \mathbf{1}_A d\mu = \mu(A).$$

PROOF. We let $B_+(X)$ denote the subset of $B(X)$ consisting of all non-negative functions $f \in B(X)$. Observe that the linear span of $B_+(X)$ is the entire $B(X)$. First of all, for each $A \in \mathcal{A}$ we have the integral

$$\int_X \mathbf{1}_A d\mu := \mu(A).$$

We next extend the integral from the set of characteristic functions $\mathbf{1}_A$, $A \in \mathcal{A}$, to the linear subspace $S(X) \subset B(X)$ of *simple functions*, i.e. the linear span of the set of characteristic functions. We also define $S_+(X)$ as $B_+(X) \cap S(X)$. In order to construct an extension of \int_X to $S(X)$, we observe that each $f \in S(X)$ can be written in the form

$$(1.2) \quad f = \sum_{i=1}^n s_i \mathbf{1}_{A_i}$$

where the subsets A_i are pairwise disjoint. Moreover, we can choose the subsets A_i such that either $f|_{A_i} \geq 0$ or $f|_{A_i} < 0$, and this for every i . (Here we are helped by the fact that $\mathcal{A} = 2^X$.) Next, for $s_i \in \mathbb{R}$, $A_i \in \mathcal{A}$, $i = 1, \dots, n$, and $t_j \in \mathbb{R}$, $B_j \in \mathcal{A}$, $j = 1, \dots, m$, finite additivity of μ implies that if

$$\sum_{i=1}^n s_i \mathbf{1}_{A_i} = \sum_{j=1}^m t_j \mathbf{1}_{B_j}$$

then

$$\sum_{i=1}^n s_i \mu(A_i) = \sum_{j=1}^m t_j \mu(B_j).$$

Therefore, we can extend \int to a linear functional on $S(X)$ by linearity:

$$\int_X \left(\sum_{i=1}^n s_i \mathbf{1}_{A_i} \right) d\mu = \sum_{i=1}^n s_i \mu(A_i).$$

Since for every $f \in S_+(X)$ we can assume that in (1.2), each $s_i \geq 0$ and $A_i \cap A_j = \emptyset$ for all $i \neq j$, it follows that

$$\int_X f d\mu \geq 0.$$

Next, given a function $f \in B_+(X)$ we set

$$\int_X f d\mu := \sup \left\{ \int_X g d\mu : g \in S_+(X), g \leq f \right\}.$$

It is clear from this definition that, since μ is G -invariant, so is the map

$$\int_X : B_+(X) \rightarrow \mathbb{R}.$$

Furthermore, it is clear that

$$\int_X a f d\mu = a \int_X f d\mu$$

for all $a \geq 0$ and $f \in B_+(X)$. However, additivity of \int_X thus defined is not obvious. We leave it to the reader to verify the simpler fact that for all functions $f, g \in B_+(X)$ we have

$$\int_X (f + g) d\mu \geq \int_X f d\mu + \int_X g d\mu.$$

We will prove the reverse inequality. For each subset $A \subset X$ and $f \in B_+(X)$ define the integral

$$\int_A f d\mu := \int_X f \mathbf{1}_A d\mu.$$

Given a simple function h , $0 \leq h \leq f + g$, we need to show that

$$\int_X h d\mu \leq \int_X f d\mu + \int_X g d\mu.$$

The function h can be written as

$$h = \sum_{i=1}^n a_i \mathbf{1}_{A_i}$$

with pairwise disjoint $A_i \in \mathcal{A}$ and $a_i > 0$. Therefore, in view of the linearity of \int_X on $S(X)$, it suffices to prove that

$$a_i \mu(A_i) \leq \int_{A_i} f d\mu + \int_{A_i} g d\mu$$

for each i . Thus, the problem reduces to the case $n = 1$, $A_1 = A$ and (by dividing by a_1) to proving the inequality

$$\mu(A) \leq \int_A f d\mu + \int_A g d\mu,$$

for functions $f, g \in B_+(X)$ satisfying

$$(1.3) \quad \mathbf{1}_A \leq f + g.$$

Let c be an integer upper bound for f and g . For each $N \in \mathbb{N}$, consider the following simple functions f_N, g_N :

$$f|_A - \frac{c}{N} \mathbf{1}_A \leq f_N := \sum_{j=0}^N \mathbf{1}_{f^{-1}((cj/N, c(j+1)/N])} \frac{cj}{N} \leq f,$$

$$g|_A - \frac{c}{N} \mathbf{1}_A \leq g_N := \sum_{j=0}^N \mathbf{1}_{g^{-1}((cj/N, c(j+1)/N])} \frac{cj}{N} \leq g.$$

In view of the inequality (1.3) we have

$$\left(1 - \frac{2c}{N}\right) \mathbf{1}_A \leq f_N + g_N.$$

The latter implies (by the definition of \int_X) that

$$\left(1 - \frac{2c}{N}\right) \mu(A) \leq \int_A f_N d\mu + \int_A g_N d\mu \leq \int_A f d\mu + \int_A g d\mu.$$

Since this inequality holds for all $N \in \mathbb{N}$, we conclude that

$$\mu(A) \leq \int_A f d\mu + \int_A g d\mu$$

as required. Thus, \int_X is an additive functional on $B_+(X)$. Since $B_+(X)$ spans $B(X)$, \int_X extends uniquely (by linearity) to a linear functional on $B(X)$. Clearly, the result is a G -invariant integral on $B(X)$. \square

1.3. Topological spaces. Lebesgue covering dimension

In this section we review some topological notions that shall be used in the book.

Notation and terminology. A *neighborhood* of a point in a topological space will always mean an *open neighborhood*. A *neighborhood* of a subset A in a topological space X is an open subset $U \subset X$ containing A .

We will use the notation \bar{A} , $\text{cl} A$ and $\text{cl}(A)$ for the closure of a subset A in a topological space X . We will denote by $\text{int} A$ and $\text{int}(A)$ the interior of A in X . A subset of a topological space X is called *clopen* if it is both closed and open. We will use the notation \mathcal{C}_X and \mathcal{K}_X for the sets of all closed, and of all compact subsets in X , respectively.

A topological space X is said to be *locally compact* if there is a basis of topology of X consisting of relatively compact subsets of X , i.e. subsets of X with compact closures. A space X is called *σ -compact* if there exists a sequence of compact subsets $(K_n)_{n \in \mathbb{N}}$ in X such that $X = \bigcup_{n \in \mathbb{N}} K_n$. A *second countable topological space* is a topological space which admits a countable base of topology (this is sometimes called *the second axiom of countability*). A second countable space is *separable* (i.e. contains a countable dense subset) and *Lindelöf* (i.e. every open cover has a countable sub-cover). A locally compact second countable space is σ -compact.

The *wedge* of a family of pointed topological spaces $(X_i, x_i), i \in I$, denoted by $\vee_{i \in I} X_i$, is the quotient of the disjoint union $\sqcup_{i \in I} X_i$, where we identify all the points x_i . The wedge of two pointed topological spaces is denoted $X_1 \vee X_2$.

If $f : X \rightarrow \mathbb{R}$ is a function on a topological space X , then we will denote by $\text{Supp}(f)$ the *support* of f , i.e. the set

$$\text{cl}(\{x \in X : f(x) \neq 0\}).$$

Given two topological spaces X, Y , we let $C(X; Y)$ denote the space of all continuous maps $X \rightarrow Y$; we also set $C(X) := C(X; \mathbb{R})$. For a function $f \in C(X)$ we define its norm

$$\|f\| = \sup_{x \in X} |f(x)|.$$

We always endow the space $C(X; Y)$ with the *compact-open topology*. A subbasis of this topology consists of the subsets

$$U_{K,V} = \{f : X \rightarrow Y : f(K) \subset V\} \subset C(X; Y),$$

where $K \subset X$ is compact and $V \subset Y$ is open.

If Y is a metric space then the compact-open topology is equivalent to the topology of *uniform convergence on compacts*: A sequence of functions $f_i : X \rightarrow Y$ converges to a function $f : X \rightarrow Y$ if and only if for every compact subset $K \subset X$ the sequence of restrictions $f_i|_K : K \rightarrow Y$ converges to $f|_K$ uniformly.

Given two topological spaces X, Y and two continuous maps

$$f_0, f_1 : X \rightarrow Y,$$

a *homotopy* between f_0 and f_1 is a continuous map $F : X \times [0, 1] \rightarrow Y$ such that $F(x, 0) = f_0(x), F(x, 1) = f_1(x)$, for every $x \in X$. *Tracks* of this homotopy are paths $F(x, t), t \in [0, 1]$, in Y , for various (fixed) points $x \in X$.

A continuous map $f : X \rightarrow Y$ of topological spaces is called *proper* if preimages of compact sets under f are again compact. In line with this, one defines a *proper homotopy* between two maps

$$f_0, f_1 : X \rightarrow Y,$$

by requiring the homotopy F between these maps to be a proper map $F : X \times [0, 1] \rightarrow Y$.

A topological space is called *perfect* if it is non-empty and contains no isolated points, i.e. points $x \in X$ such that the singleton $\{x\}$ is open in X .

DEFINITION 1.13. A topological space X is *regular* if every closed subset $A \subset X$ and a singleton $\{x\} \subset X \setminus A$, have disjoint neighborhoods. A topological space X is called *normal* if every pair of disjoint closed subsets $A, B \subset X$ have disjoint open neighborhoods, i.e. there exist disjoint open subsets $U, V \subset X$ such that $A \subset U, B \subset V$.

- EXERCISE 1.14. 1. Every normal Hausdorff space is regular.
2. Every compact Hausdorff space is normal.

We will also need a minor variation on the notion of normality:

DEFINITION 1.15. Two subsets A, B of a topological space X are said to be *separated by a function* if there exists a continuous function $\rho = \rho_{A,B} : X \rightarrow [0, 1]$ so that

1. $\rho|_A \equiv 0$
2. $\rho|_B \equiv 1$.

A topological space X is called *perfectly normal* if every two disjoint closed subsets of X can be separated by a function.

We will see below (Lemma 2.2) that every metric space is perfectly normal. A much harder result is

THEOREM 1.16 (Tietze–Urysohn extension theorem). *Every normal topological space X is perfectly normal. Moreover, for every closed subset $C \subset X$ and continuous function $f : C \rightarrow \mathbb{R}$, the function f admits a continuous extension to X .*

A proof of this extension theorem can be found in [Eng95]. In view of this theorem, every normal topological space is perfectly normal, since one can take $C = A \cup B$ and let ρ be a continuous extension of the function

$$f : C \rightarrow \mathbb{R}, \quad f|_A \equiv 0, f|_B \equiv 1.$$

COROLLARY 1.17. *In the definition of a normal topological space, one can take U and V to have disjoint closures.*

PROOF. Let $f : C = A \cup B \rightarrow \{0, 1\}$ be the function as above. define $\rho : X \rightarrow \mathbb{R}$ to be a continuous extension of f . Then take

$$U := \rho^{-1}((-\infty, 1/3)), \quad V := \rho^{-1}((2/3, \infty)). \quad \square$$

LEMMA 1.18. [*Extension lemma*] *Suppose that X, Y topological spaces, where Y is regular and X contains a dense subset A .*

1. *If $f : X \rightarrow Y$ is a mapping satisfying the property that for each $x \in X$ the restriction of f to $A \cup \{x\}$ is continuous, then f is continuous.*

2. *Assume now that A is open and set $X \setminus A = Z$. Suppose that $f : X \rightarrow Y$ is such that the restriction $f|_{A \cup \{z\}}$ is continuous at z for every point $z \in Z \subset X$. Then $f : X \rightarrow Y$ is continuous at each point $z \in Z$.*

PROOF. 1. We will verify continuity of f at each $x \in X$. Let $y = f(x)$ and let V be an (open) neighborhood of y in Y ; the complement $C = Y \setminus V$ is closed. Since Y is regular, there exist disjoint open neighborhoods $V_1 \subset V$ of y and V_2 of C . Therefore, the closure W of V_1 is contained in V . By continuity of the map

$$f|_{A \cup \{x\}},$$

there exists an (open) neighborhood U of x in X , such that

$$f(U \cap (A \cup \{x\})) \subset V_1 \subset W.$$

Let us verify that $f(U) \subset W \subset V$. Take $z \in U$. By continuity of

$$f|_{A \cup \{z\}},$$

the preimage

$$D = f^{-1}(W) \cap (A \cup \{z\})$$

is closed in $A \cup \{z\}$. This preimage contains $U \cap A$; the latter is dense in U , since A is dense in X and $U \subset X$ is open. Therefore, D contains the closure of $U \cap A$ in $A \cup \{z\}$. This closure contains the point z since $z \in U$ and $U \cap A$ is dense in U . It follows that z is in D and, hence, $f(z) \in W$. We conclude that $f(U) \subset V$ and, therefore, f is continuous at x .

2. We change the topology \mathcal{T}_X on X to a new topology \mathcal{T}'_X whose basis is the union of \mathcal{T}_X and the power set 2^A . Since $A \in \mathcal{T}'_X$, it follows that the map $f : X \rightarrow Y$ is continuous at each point $a \in A$ with respect to the new topology. Part 1 now

implies continuity of the map $f : (X, \mathcal{T}'_X) \rightarrow (Y, \mathcal{T}_Y)$. It follows that $f : A \cup \{z\} \rightarrow Y$ is continuous at each $z \in Z$ with respect to the original topology. \square

A topological space X is said to be *locally path-connected* if for each $x \in X$ and each neighborhood U of x , there exists a neighborhood $V \subset U$ of x , such that every point $y \in V$ can be connected to x by a path contained in U . In other words, the inclusion $V \hookrightarrow U$ induces the map

$$\pi_0(V) \rightarrow \pi_0(U)$$

whose image is a singleton.

An open covering $\mathcal{U} = \{U_i : i \in I\}$ of a topological space X is called *locally finite* if every subset $J \subset I$ such that

$$\bigcap_{i \in J} U_i \neq \emptyset$$

is finite. Equivalently, every point $x \in X$ has a neighborhood which intersects only finitely many U_i 's.

The *multiplicity* of an open cover $\mathcal{U} = \{U_i : i \in I\}$ of a space X is the supremum of cardinalities of subsets $J \subset I$ so that

$$\bigcap_{i \in J} U_i \neq \emptyset.$$

A cover \mathcal{V} is called a *refinement* of a cover \mathcal{U} if every $V \in \mathcal{V}$ is contained in some $U \in \mathcal{U}$.

DEFINITION 1.19. The (*Lebesgue*) *covering dimension* of a topological space Y is the least number n such that the following holds: Every open cover \mathcal{U} of Y admits a refinement \mathcal{V} which has multiplicity at most $n + 1$.

The following example shows that the covering dimension is consistent with our “intuitive” notion of dimension:

EXAMPLE 1.20. If M is an n -dimensional topological manifold, then n equals the covering dimension of M . See e.g. [Nag83].

1.4. Exhaustions of locally compact spaces

DEFINITION 1.21. A family of compact subsets $\{K_i : i \in I\}$ of a topological space X is said to be an *exhaustion* of X if:

1. $\bigcup_{i \in I} K_i = X$.
2. For each $i \in I$ there exists $j \in I$ such that

$$K_i \subset \text{int}(K_j).$$

PROPOSITION 1.22. *If X is locally compact, Hausdorff and second countable, it admits an exhaustion by a countable collection of compact subsets. Moreover, there exists a countable exhaustion $\{K_n : n \in \mathbb{N}\}$ of X such that $K_n \subset \text{int } K_{n+1}$ for each n .*

PROOF. If X is empty, there is nothing to prove, therefore we will assume that $X \neq \emptyset$. Let \mathcal{B} be a countable basis of X . Define $\mathcal{U} \subset \mathcal{B}$ to be a subset of \mathcal{B} consisting of relatively compact sets.

LEMMA 1.23. *\mathcal{U} is a basis of X .*

PROOF. Let $x \in X$ be a point and V a neighborhood of x . Since X is locally compact, there exists a compact subset $K \subset X$ with

$$x \in \text{int}(K) \subset V.$$

Then the boundary ∂K of K in X is disjoint from $\{x\}$. Since K is regular and \mathcal{B} is a basis, there exists a neighborhood W of ∂K in K and $B \in \mathcal{B}$, a neighborhood of x in X , such that $B \subset \text{int}(K)$ and $B \cap W = \emptyset$. Then the closure \bar{B} of B is compact (and, thus, $B \in \mathcal{U}$) and contained in $K \setminus W \subset \text{int}(K)$. \square

We define an exhaustion of X inductively. We begin by enumerating the elements of the cover:

$$\mathcal{U} = \{U_1, U_2, \dots, U_n, \dots\}.$$

Set K_1 to be the closure of U_1 . Given a compact subset K_n containing $\bigcup_{i=1}^n U_i$, we consider its cover by elements of \mathcal{U} . By compactness of K_n , there exists a finite subcollection $U_{j_1}, \dots, U_{j_k} \in \mathcal{U}$ covering K_n . Set

$$K_{n+1} := \bigcup_{s=1}^k \bar{U}_{j_s} \cup \bar{U}_{n+1}.$$

This is the required exhaustion. \square

1.5. Direct and inverse limits

Let I be a *directed set*, i.e. a partially ordered set, where every two elements i, j have an *upper bound*, which is some $k \in I$ such that $i \leq k, j \leq k$. The reader should think of the set of real numbers, or positive real numbers, or natural numbers, as the main examples of directed sets. A *directed system* of sets (or topological spaces, or groups) indexed by I is a collection of sets (or topological spaces, or groups) $A_i, i \in I$, and maps (or continuous maps, or homomorphisms) $f_{ij} : A_i \rightarrow A_j, i \leq j$, satisfying the following compatibility conditions:

- (1) $f_{ik} = f_{jk} \circ f_{ij}, \forall i \leq j \leq k$,
- (2) $f_{ii} = \text{Id}$.

An *inverse system* is defined similarly, except $f_{ij} : A_j \rightarrow A_i, i \leq j$, and, accordingly, in the first condition we use $f_{ij} \circ f_{jk} = f_{ik}$.

We will use the notation $(A_i, f_{ij}, i, j \in I)$ for direct and inverse systems of sets, spaces and groups.

The *direct limit* of the direct system of sets is the set

$$A = \varinjlim A_i = \left(\prod_{i \in I} A_i \right) / \sim,$$

where $a_i \sim a_j$ whenever $f_{ik}(a_i) = f_{jk}(a_j)$ for some $k \in I$. In particular, we have maps $f_m : A_m \rightarrow A$ given by $f_m(a_m) = [a_m]$, where $[a_m]$ is the equivalence class in A represented by $a_m \in A_m$. Note that

$$A = \bigcup_{m \in I} f_m(A_m).$$

If the sets A_i are groups, then we equip the direct limit with the group operation:

$$[a_i] \cdot [a_j] = [f_{ik}(a_i) \cdot f_{jk}(a_j)],$$

where $k \in I$ is an upper bound for i and j .

If the sets A_i are topological spaces, then we equip the direct limit with the *final topology*, i.e. the topology where $U \subset \varinjlim A_i$ is open if and only if $f_i^{-1}(U)$ is open for every $i \in I$. In other words, this is the quotient topology descending from the disjoint union of A_i 's.

Similarly, the *inverse limit* of an inverse system is

$$\varprojlim A_i = \left\{ (a_i) \in \prod_{i \in I} A_i : a_i = f_{ij}(a_j), \forall i \leq j \right\}.$$

If the sets A_i are groups, we equip the inverse limit with the group operation induced from the direct product of the groups A_i . If the sets A_i are topological spaces, we equip the inverse limit with the *initial topology*, i.e. the subset topology of the Tychonoff topology on the direct product. Explicitly, this is the topology generated by the open sets of the form $f_m^{-1}(U_m)$, where $U_m \subset X_m$ are open subsets and $f_m : \varprojlim A_i \rightarrow A_m$ is the restriction of the coordinate projection.

EXERCISE 1.24. 1. Show that $\varprojlim A_i$ is closed in $\prod_{i \in I} A_i$. 2. Conclude that if each A_i is compact, then so is $\varprojlim A_i$.

Given a subset $J \subset I$, we have the *restriction map*

$$\rho : \prod_{i \in I} A_i \rightarrow \prod_{j \in J} A_j, \quad \lambda \mapsto \lambda|_J$$

where we treat elements of the product spaces as functions $I \rightarrow \bigcup_{i \in I} A_i$ and $J \rightarrow \bigcup_{j \in J} A_j$ respectively. A subset $J \subset I$ is called *cofinal* if for each $i \in I$ there exists $j \in J$ such that $i \leq j$.

EXERCISE 1.25. Show that if $J \subset I$ is cofinal then the restriction map ρ is a bijection between $\varprojlim A_i$ and $\varprojlim A_j$.

EXERCISE 1.26. Assuming that each A_i is a Hausdorff topological space satisfying the first separation axiom (also denoted T_1 , requiring that singletons are closed sets), show that $\varprojlim A_i$ is a closed subset of the product space $\prod_{i \in I} A_i$. Conclude that the inverse limit of a directed system of compact Hausdorff topological spaces is again compact and Hausdorff. Conclude, furthermore, that if each A_i is totally disconnected, then so is the inverse limit.

Suppose that each A_i is a topological space and we are given a subset $A'_i \subset A_i$ with the subspace topology. Then we have a natural continuous embedding

$$\iota : \prod_{i \in I} A'_i \rightarrow \prod_{i \in I} A_i$$

EXERCISE 1.27. Suppose that for each $i \in I$ there exists $j \in I$ such that $f_{ij}(A_j) \subset A'_i$. Verify that the map ι is a bijection.

We now turn from topological spaces to groups.

EXERCISE 1.28. Every group G is the direct limit of a direct system $G_i, i \in I$, consisting of all finitely generated subgroups of G . Here the partial order on I is given by inclusion and homomorphisms $f_{ij} : G_i \rightarrow G_j$ are tautological embeddings.

EXERCISE 1.29. Suppose that G is the direct limit of a direct system of groups $\{G_i, f_{ij} : i, j \in I\}$. Assume also that for every i we are given a subgroup $H_i \leq G_i$ satisfying

$$f_{ij}(H_i) \leq H_j, \quad \forall i \leq j.$$

Then the family of groups and homomorphisms

$$\mathcal{H} = \{H_i, f_{ij}|_{H_i} : i, j \in I\}$$

is again a direct system; let H denote the direct limit of this system. Show that there exists a monomorphism $\phi : H \rightarrow G$, so that for every $i \in I$,

$$f_i|_{H_i} = \phi \circ h_i : H_i \rightarrow G,$$

where $h_i : H_i \rightarrow H$ are the homomorphisms associated with the direct limit of the system \mathcal{H} .

EXERCISE 1.30. 1. Let $H \leq G$ be a subgroup. Then $|G : H| \leq n$ if and only if the following holds: For every subset $\{g_0, \dots, g_n\} \subset G$, there exist $i \neq j$ so that $g_i g_j^{-1} \in H$.

2. Suppose that G is the direct limit of a system of groups $\{G_i, f_{ij}, i, j \in I\}$. Assume also that there exist $n \in \mathbb{N}$ so that for every $i \in I$, the group G_i contains a subgroup H_i of index $\leq n$ and the assumptions of Exercise 1.29 are satisfied. Let the group H be the direct limit of the system

$$\{H_i, f_{ij}|_{H_i} : i, j \in I\}$$

and $\phi : H \rightarrow G$ be the monomorphism as in Exercise 1.29. Show that

$$|G : \phi(H)| \leq n.$$

1.6. Graphs

An *unoriented graph* Γ consists of the following data:

- a set V called the *set of vertices* of the graph;
- a set E called the *set of edges* of the graph;
- a map ι called *incidence map* defined on E and taking values in the set of subsets of V of cardinality one or two.

We will use the notation $V = V(\Gamma)$ and $E = E(\Gamma)$ for the vertex and respectively the edge set of the graph Γ . When $\{u, v\} = \iota(e)$ for some edge e , the two vertices u, v are called the *endpoints* of the edge e ; we say that u and v are *adjacent vertices*.

An unoriented graph can also be seen as a 1-dimensional cell complex (see section 1.7), with 0-skeleton V and with 1-dimensional cells/edges labeled by elements of E , such that the boundary of each 1-cell $e \in E$ is the set $\iota(e)$.

Note that in the definition of a graph we allow for *monogons* (i.e. edges connecting a vertex to itself)¹ and *bigons*² (pairs of distinct edges with the same endpoints). A graph is *simplicial* if the corresponding cell complex is a simplicial complex. In other words, a graph is simplicial if and only if it contains no monogons or bigons³.

¹Not to be confused with *unigons*, which are hybrids of unicorns and dragons.

²Also known as *digons*.

³and, naturally, no unigons, because those do not exist anyway.

The incidence map ι defining a graph Γ is set-valued; converting ι into a map with values in $V \times V$, equivalently into a pair of maps $E \rightarrow V$ is the choice of an *orientation* of Γ : An orientation of Γ is a choice of two maps

$$o : E \rightarrow V, \quad t : E \rightarrow V$$

such that $\iota(e) = \{o(e), t(e)\}$ for every $e \in E$. In view of the Axiom of Choice, every graph can be oriented.

DEFINITION 1.31. An *oriented* or *directed* graph is a graph Γ equipped with an orientation. The maps o and t are called the *head (or origin) map* and the *tail map* respectively.

We will in general denote an oriented graph by $\bar{\Gamma}$, its edge-set by \bar{E} , and oriented edges by \bar{e} .

CONVENTION 1.32. In this book, unless we state otherwise, all graphs are assumed to be unoriented.

The *valency (or valence, or degree) of a vertex v* of a graph Γ is the number of edges having v as an endpoint, where every monogon with both endpoints equal to v is counted twice. The *valency* of Γ is the supremum of valencies of its vertices.

Examples of graphs. Below we describe several examples of graphs which will appear in this book.

EXAMPLE 1.33 (*n-rose*). This graph, denoted R_n , has one vertex and n edges connecting this vertex to itself.

EXAMPLE 1.34. [*i-star or i-pod*] This graph, denoted T_i , has $i + 1$ vertices, v_0, v_1, \dots, v_i . Two vertices are connected by a unique edge if and only if one of these vertices is v_0 and the other one is different from v_0 . The vertex v_0 is the *center* of the star and the edges are called its *legs*.

EXAMPLE 1.35 (*n-circle*). This graph, denoted C_n , has n vertices which are identified with the n -th roots of unity:

$$v_k = e^{2\pi i k/n}.$$

Two vertices u, v are connected by a unique edge if and only if they are adjacent to each other on the unit circle:

$$uv^{-1} = e^{\pm 2\pi i/n}.$$

EXAMPLE 1.36 (*n-interval*). This graph, denoted I_n , has the vertex set equal to $[1, n + 1] \cap \mathbb{N}$, where \mathbb{N} is the set of natural numbers. Two vertices n, m of this graph are connected by a unique edge if and only if

$$|n - m| = 1.$$

Thus, I_n has n edges.

EXAMPLE 1.37 (*Half-line*). This graph, denoted H , has the vertex set equal to \mathbb{N} (the set of natural numbers). Two vertices n, m are connected by a unique edge if and only if

$$|n - m| = 1.$$

The subset $[n, \infty) \cap \mathbb{N} \subset V(H)$ is the vertex set of a subgraph of H also isomorphic to the half-line H . We will use the notation $[n, \infty)$ for this subgraph.

EXAMPLE 1.38 (Line). This graph, denoted L , has the vertex set equal to \mathbb{Z} , the set of integers. Two vertices n, m of this graph are connected by a unique edge if and only if

$$|n - m| = 1.$$

As with general cell complexes and simplicial complexes, we will frequently conflate a graph with its *geometric realization*:

DEFINITION 1.39. The *geometric realization* or *underlying topological space* of an oriented graph $\bar{\Gamma}$ is the quotient space of the topological space

$$U = \bigsqcup_{v \in V} \{v\} \sqcup \bigsqcup_{\bar{e} \in \bar{E}} \{\bar{e}\} \times [0, 1]$$

by the equivalence relation

$$\bar{e} \times \{0\} \sim o(\bar{e}), \quad \bar{e} \times \{1\} \sim t(\bar{e}).$$

One defines the geometric realization of an undirected graph Γ by converting Γ to an oriented graph $\bar{\Gamma}$; the topology of the resulting space is independent of the orientation.

A *morphism* of graphs $f : \Gamma \rightarrow \Gamma'$ is a pair of maps $f_V : V(\Gamma) \rightarrow V(\Gamma')$, $f_E : E(\Gamma) \rightarrow E(\Gamma')$ such that

$$\iota' \circ f_E = f_V \circ \iota$$

where ι and ι' are the incidence maps of the graphs Γ and Γ' respectively. Thus, every morphism of graphs induces a (nonunique) continuous map $f : \Gamma \rightarrow \Gamma'$ of geometric realizations, sending vertices to vertices and edges to edges. A *monomorphism* of graphs is a morphism such that the corresponding maps f_V, f_E are injective. The image of a monomorphism $\Gamma \rightarrow \Gamma'$ is a *subgraph* of Γ' . In other words, a subgraph in a graph Γ' is defined by subsets $V \subset V(\Gamma'), E \subset E(\Gamma')$ such that

$$\iota'(e) \subset V$$

for every $e \in E$. A subgraph Γ' of Γ is called *full* if every $e = [v, w] \in E(\Gamma)$ connecting vertices of Γ' , is an edge of Γ' .

A morphism $f : \Gamma \rightarrow \Gamma'$ of graphs which is invertible (as a morphism) is called an *isomorphism* of graphs: More precisely, we require that the maps f_V, f_E are invertible and the inverse maps define a morphism $\Gamma' \rightarrow \Gamma$. In other words, an isomorphism of graphs is an isomorphism of the corresponding cell complexes.

EXERCISE 1.40. 1. For every isomorphism of graphs there exists a (nonunique) homeomorphism $f : \Gamma \rightarrow \Gamma'$ of geometric realizations, such that the images of the edges of Γ are edges of Γ' and images of vertices are vertices.

2. Isomorphisms of graphs are morphisms such that the corresponding vertex and edge maps are bijective.

We use the notation $\text{Aut}(\Gamma)$ for the group of automorphisms of a graph Γ .

An edge connecting two vertices u, v of a graph Γ will sometimes be denoted by $[u, v]$: This is unambiguous if Γ is simplicial. A finite ordered set of edges of the form $[v_1, v_2], [v_2, v_3], \dots, [v_n, v_{n+1}]$ is called an *edge-path* in Γ . The number n is called the *combinatorial length* of the edge-path. An edge-path in Γ is a *cycle* if $v_{n+1} = v_1$. A *simple cycle* (or a *circuit*) is a cycle with all vertices $v_i, i = 1, \dots, n$, pairwise distinct. In other words, a simple cycle is a subgraph isomorphic to the

n -circle for some n . A graph Γ is *connected* if any two vertices of Γ are connected by an edge-path. Equivalently, the topological space underlying Γ is path-connected.

A subgraph $\Gamma' \subset \Gamma$ is called a *connected component* of Γ if Γ' is a maximal (with respect to the inclusion) connected subgraph of Γ .

A *simplicial tree* is a connected graph without circuits.

EXERCISE 1.41. Simple cycles in a graph Γ' are precisely subgraphs whose underlying spaces are homeomorphic to the circle.

Maps of graphs. Sometimes, it is convenient to consider maps of graphs which are not morphisms. A *map of graphs* $f : \Gamma \rightarrow \Gamma'$ consists of a pair of maps (g, h) :

1. A map $g : V(\Gamma) \rightarrow V(\Gamma')$ sending adjacent vertices to adjacent or equal vertices;
2. A *partially defined* map of the edge-sets:

$$h : E_o \rightarrow E(\Gamma'),$$

where E_o consists only of edges e of Γ whose endpoints $v, w \in V(\Gamma)$ have distinct images by g :

$$g(v) \neq g(w).$$

For each $e \in E_o$, we require the edge $e' = h(e)$ to connect the vertices $g(o(e)), g(t(e))$. In other words, f amounts to a morphism of graphs $\Gamma_o \rightarrow \Gamma'$, where the vertex set of Γ_o is $V(\Gamma)$ and the edge-set of Γ_o is E_o .

Collapsing a subgraph. Given a graph Γ and a (non-empty) subgraph Λ of it, we define a new graph, $\Gamma' = \Gamma/\Lambda$, by “collapsing” the subgraph Λ to a vertex. Here is the precise definition. Define the partition $V(\Gamma) = W \sqcup W^c$,

$$W = V(\Lambda), \quad W^c = V(\Gamma) \setminus V(\Lambda).$$

The vertex set of Γ' equals

$$W^c \sqcup \{v_o\}.$$

Thus, we have a natural surjective map $V(\Gamma) \rightarrow V(\Gamma')$ sending each $v \in W^c$ to itself and each $v \in W$ to the vertex v_o . The edge-set of Γ' is in bijective correspondence to the set of edges in Γ which *do not* connect vertices of Λ to each other. Each edge $e \in E(\Gamma)$ connecting $v \in W^c$ to $w \in W$ projects to an edge, also called e , connecting v to v_o . If an edge e connects two vertices in W^c , it is also retained and connects the same vertices in Γ' .

The map $V(\Gamma) \rightarrow V(\Gamma')$ extends to a *collapsing* map of graphs $\kappa : \Gamma \rightarrow \Gamma'$.

EXERCISE 1.42. If Γ is a tree and Λ is a subtree, then Γ' is again a tree.

DEFINITION 1.43. Let $F \subset V = V(\Gamma)$ be a set of vertices in a graph Γ . The *vertex-boundary* of F , denoted by $\partial_V F$, is the set of vertices in F each of which is adjacent to a vertex in $V \setminus F$. Similarly, the *exterior vertex-boundary* of F is

$$\partial^V F = \partial_V F^c.$$

The *edge-boundary* of F , denoted by $E(F, F^c)$, is the set of edges e such that the set of endpoints $\iota(e)$ intersects both F and its complementary set F^c in exactly one element.

Unlike the vertex-boundary, the edge boundary is the same for F , as for its complement F^c . There is a natural surjective map from the edge-boundary $E(F, F^c)$ to the vertex-boundary $\partial_V F$, sending each edge $e \in E(F, F^c)$ to the vertex v such that $\{v\} = \iota(e) \cap F$. This map is at most C -to-1, where C is the valency of Γ . Hence, assuming that C is finite, the cardinalities of the two types of boundaries are *comparable*:

$$(1.4) \quad |\partial_V F| \leq |E(F, F^c)| \leq C|\partial_V F|,$$

$$(1.5) \quad |\partial^V F| \leq |E(F, F^c)| \leq C|\partial^V F|,$$

DEFINITION 1.44. A simplicial graph Γ is *bipartite* if the vertex set V splits as $V = Y \sqcup Z$, so that each edge $e \in E$ has one endpoint in Y and one endpoint in Z . In this case, we write $\Gamma = \text{Bip}(Y, Z; E)$.

EXERCISE 1.45. Let W be an n -dimensional vector space over a field \mathbb{K} , where $n \geq 3$. Let Y be the set of 1-dimensional subspaces of W and let Z be the set of 2-dimensional subspaces of W . Define the bipartite graph $\Gamma = \text{Bip}(Y, Z, E)$, where $y \in Y$ is adjacent to $z \in Z$ if and only if, as subspaces in W , $y \subset z$.

- (1) Compute (in terms of \mathbb{K} and n) the valency of Γ , the (combinatorial) length of the shortest circuit in Γ , and show that Γ is connected.
- (2) Estimate from above the length of the shortest path between any pair of vertices of Γ . Can you get a bound independent of \mathbb{K} and n ?

1.7. Complexes and homology

Complexes are higher-dimensional generalizations of graphs. In this book, we will primarily use two types of complexes:

- Simplicial complexes.
- Cell complexes.

As we expect the reader to be familiar with basics of algebraic topology, we will discuss simplicial and cell complexes only briefly.

1.7.1. Simplicial complexes.

DEFINITION 1.46. A simplicial complex X consists of a set $V = V(X)$, called the *vertex set* of X , and a collection $S(X)$ of finite non-empty subsets of V ; members of $S(X)$ of cardinality $n + 1$ are called *n -dimensional simplices* or *n -simplices*. A simplicial complex is required to satisfy the following axioms:

- (1) For every simplex $\sigma \in S(X)$, every non-empty subset $\tau \subset \sigma$ is also a simplex. The subset τ is called a *face* of σ . *Vertices* of σ are the 0-faces of σ .
- (2) Every singleton $\{v\} \subset V$ is an element of $S(X)$.

A *simplicial map* or *morphism* of two simplicial complexes $f : X \rightarrow Y$ is a map $f : V(X) \rightarrow V(Y)$ which sends simplices to simplices, where the dimension of a simplex might decrease under f .

Products of simplicial complexes. Let X, Y be simplicial complexes. We order all the vertices of X and Y . The *product* $Z = X \times Y$ is defined as a simplicial

complex whose vertex set is $V(X) \times V(Y)$. Let σ, τ be simplices in X and Y of dimensions m and n respectively with vertex sets

$$\sigma^{(0)} = \{v_0, \dots, v_m\}, \quad \tau^{(0)} = \{w_0, \dots, w_n\}.$$

The product $\sigma \times \tau$, of course, is not a simplex (unless $nm = 0$), but it admits a standard triangulation, whose vertex set is

$$\sigma^{(0)} \times \tau^{(0)}.$$

This triangulation is defined as follows. Pairs $u_{ij} = (v_i, w_j)$ are the vertices of $\sigma \times \tau$. Distinct vertices

$$(u_{i_0, j_0}, \dots, u_{i_k, j_k})$$

span a k -simplex in $\sigma \times \tau$ if and only if $i_0 \leq \dots \leq i_k$ and $j_0 \leq \dots \leq j_k$. These will compose the set of simplices of Z , for arbitrary choices of σ and τ . The product complex Z depends on the orderings of $V(X)$ and $V(Y)$; however, which orderings to choose will be irrelevant for our purposes.

We will use the notation $X^{(i)}$ for the i -th *skeleton* of the simplicial complex X , i.e. the simplicial subcomplex with the same set of vertices but having as collection of simplices only the simplices of dimension $\leq i$ in X .

The *geometric realization* of an n -simplex is the *standard simplex*

$$\Delta^n = \{(x_0, \dots, x_n) \in \mathbb{R}^{n+1} : \sum_{i=0}^n x_i = 1, x_i \geq 0, i = 0, \dots, n\}.$$

Faces of Δ^n are intersections of the standard simplex with coordinate subspaces of \mathbb{R}^{n+1} . Given a simplicial complex X , by gluing copies of standard simplices, one obtains a topological space, which is a *geometric realization* of X .

We define the *interior* of the standard simplex Δ^n as

$$\text{int}(\Delta^n) = \{(x_0, \dots, x_n) \in \mathbb{R}^{n+1} : \sum_{i=0}^n x_i = 1, x_i > 0, i = 0, \dots, n\}.$$

We refer to interiors of simplices in a simplicial complex X as *open simplices* in X .

A *gallery* in an n -dimensional simplicial complex X is a chain of n -simplices $\sigma_1, \dots, \sigma_k$ such that $\sigma_i \cap \sigma_{i+1}$ is an $(n-1)$ -simplex for every $i = 1, \dots, k-1$.

A *homotopy between simplicial maps* $f_0, f_1 : X \rightarrow Y$ is a simplicial map $F : X \times I \rightarrow Y$ which restricts to f_i on $X \times \{i\}$, $i = 0, 1$. The *tracks* of the homotopy F are the paths $\mathfrak{p}(t) = F(x, t)$, $x \in X$.

Cohomology with compact support. Let X be a simplicial complex. Recall that besides the usual cohomology groups $H^*(X; A)$ (with coefficients in a ring A that the reader can assume to be \mathbb{Z} or \mathbb{Z}_2), we also have *cohomology with compact support* $H_c^*(X, A)$, defined as follows. Consider the usual cochain complex $C^*(X; A)$. We say that a cochain $\sigma \in C^*(X; A)$ has *compact support* if it vanishes outside a finite subcomplex in X . Thus, in each chain group $C^k(X; A)$ we have the subgroup $C_c^k(X; A)$ consisting of compactly supported cochains. Then the usual coboundary operator δ satisfies

$$\delta : C_c^k(X; A) \rightarrow C_c^{k+1}(X; A).$$

The cohomology of the new cochain complex $(C_c^*(X; A), \delta)$ is denoted $H_c^*(X; A)$ and is called *cohomology of X with compact support*. Maps of simplicial complexes

no longer induce homomorphisms of $H_c^*(X; A)$ since they do not preserve the compact support property of cochains; however, *proper* maps of simplicial complexes do induce natural maps on H_c^* . Similarly, maps which are *properly* homotopic induce equal homomorphisms of H_c^* and *proper homotopy equivalences* induce isomorphisms of H_c^* . In other words, H_c^* satisfies the functoriality property of the usual cohomology groups as long as we restrict to the category of proper maps.

Bounded cohomology. Another variation on this construction, which has many applications in Geometric Group Theory, is the concept of *bounded cohomology*.

Let A be a subgroup in \mathbb{R} (the groups $A = \mathbb{Z}$ and $A = \mathbb{R}$ will be the main examples here). One defines the group of *bounded cochains* $C_b^k(X; A) \subset C^k(X; A)$ as the group consisting of cochains which are bounded as functions on $C_k(X)$. It is immediate that the usual coboundary operator satisfies

$$\delta_k : C_b^k(X; A) \rightarrow C_b^{k+1}(X; A)$$

for every k . This allows one to define groups of bounded cocycles $Z_b^k(X; A)$ and coboundaries $B_b^{k+1}(X; A)$ as the kernel and image of the coboundary operator restricted to $C_b^k(X; A)$. Hence, one defines the *bounded cohomology groups*

$$H_b^k(X; A) = Z_b^k(X; A) / B_b^k(X; A).$$

The inclusion $C_b^*(X) \rightarrow C^*(X)$ induces the group homomorphism

$$H_b^k(X; A) \rightarrow H^k(X; A).$$

This is sometimes called the *comparison map*.

1.7.2. Cell complexes. A CW complex X is defined as the increasing union of subspaces denoted $X^{(n)}$ (or X^n), called *n-skeleta* of X . The 0-skeleton $X^{(0)}$ of X is a set with discrete topology. Assume that $X^{(n-1)}$ is defined. Let

$$U_n := \coprod_{\alpha \in J} \mathbb{D}_\alpha^n,$$

be a (possibly empty) disjoint union of closed n -balls \mathbb{D}_α^n . Suppose that for each \mathbb{D}_α^n we have a continuous *attaching map* $e_\alpha : \partial \mathbb{D}_\alpha^n \rightarrow X^{(n-1)}$. This defines a map

$$e^n : \partial U_n \rightarrow X^{(n-1)}$$

and an equivalence relation $x \sim y = e^n(x)$, $x \in U, y \in X^{(n-1)}$. The space $X^{(n)}$ is the quotient space of $X^{(n-1)} \sqcup U_n$ with the quotient topology with respect to this equivalence relation. Each attaching map e_α extends to the map $\hat{e}_\alpha : \mathbb{D}_\alpha^n \rightarrow X^{(n)}$. We will use the notation $\sigma = \mathbb{D}_\alpha^n / e_\alpha$ for the image of \mathbb{D}_α^n in X^n , it is homeomorphic to the quotient $\mathbb{D}_\alpha^n / \sim$. We will also use the notation $e_\alpha = \partial \hat{e}_\alpha$ and refer to the image of e_α as the *boundary* of the cell $\sigma = \hat{e}_\alpha(\mathbb{D}_\alpha^n)$, $\partial \sigma = e_\alpha(\partial \mathbb{D}_\alpha^n)$. The set

$$X := \coprod_{n \in \mathbb{N}} X_n$$

is equipped with the *weak topology*, where a subset $C \subset X$ is closed if and only if the intersection of C with each skeleton is closed (equivalently, the intersection of C with each $\hat{e}_\alpha(\mathbb{D}_\alpha^n)$ in X is closed). The space X , together with the collection of maps e_α , is called a *CW complex*. By abuse of terminology, the maps \hat{e}_α , the balls \mathbb{D}_α^n , and their projections to X are called *n-cells* in X . Similarly, we will conflate the cell complex X and its underlying topological space $|X|$. We will also refer

to CW complexes simply as *cell complexes*, even though the usual notion of a cell complex is less restrictive than the one of a CW complex.

We use the terminology of *open n -cell in X* for the open ball $\mathbb{D}_\alpha^n \setminus \partial\mathbb{D}_\alpha^n$, as well as for the restriction of the map \hat{e}_α to the open ball $\mathbb{D}_\alpha^n \setminus \partial\mathbb{D}_\alpha^n$, and for the image of this restriction. We refer to this open n -cell as the *interior* of the corresponding n -cell.

We will use the notation σ for cells and σ° for their interiors.

The *dimension* of a cell complex X is the supremum of n 's such that X has an n -cell. Equivalently, the dimension of X is its topological dimension.

EXERCISE 1.47. A subset $K \subset X$ is compact if and only if it is closed and contained in a finite union of cells.

Regular and almost regular cell complexes. A cell complex X is said to be *regular* if every attaching map e_α is one-to-one. For instance, every simplicial complex is a regular cell complex. If $D \subset \mathbb{R}^n$ is a bounded convex polyhedron, then D has a natural structure of regular cell complex X , where the faces of D are the cells of X .

A regular cell complex is called *triangular* if every cell is naturally isomorphic to a simplex. (Note that a triangular cell complex need not be simplicial since a non-empty intersections of two cells need not be a single cell.)

A slightly more general notion is the one of *almost regular* cell complex. (We could not find this notion in the literature and the terminology is ours.) A cell complex X is *almost regular* if the boundary \mathbb{S}^{n-1} of every n -cell \mathbb{D}_α^n is given a structure of regular cell complex K_α , so that the attaching map e_α is one-to-one on every open cell in \mathbb{S}^{n-1} .

EXAMPLES 1.48. 1. Consider the 2-dimensional complex X constructed as follows. The complex X has a single vertex and a single 1-cell, thus $|X^{(1)}|$ is homeomorphic to \mathbb{S}^1 . Let $e : \mathbb{S}^1 \rightarrow \mathbb{S}^1$ be a k -fold covering. Attach the 2-cell \mathbb{D}^2 to $X^{(1)}$ via the map e . The result is an almost regular (but not regular) cell complex X .

2. Let X be the 2-dimensional cell complex obtained by attaching a single 2-cell to a single vertex by the constant map. Then X is not an almost regular cell complex.

Almost regular 2-dimensional cell complexes (with a single vertex) appear naturally in the context of group presentations, see Definition 7.91. For instance, suppose that X is a simplicial complex and $Y \subset X^{(1)}$ is a *forest*, i.e. a subcomplex isomorphic to a disjoint union of simplicial trees. Then the quotient X/Y is an almost regular cell complex.

Barycentric subdivision of an almost regular cell complex. Our goal is to (canonically) subdivide an almost regular cell complex X so that the result is a triangular regular cell complex $X' = Y$. We define Y as an increasing union of regular subcomplexes Y_n (where $Y_n \subset Y^{(n)}$ but, in general, is smaller).

First, set $Y_0 := X^{(0)}$. Suppose that $Y_{n-1} \subset Y^{(n-1)}$ is defined, so that $|Y_{n-1}| = |X^{(n-1)}|$. Consider the attaching maps $e_\alpha : \partial\mathbb{D}_\alpha^n \rightarrow X^{(n-1)}$. We take the preimage of the regular cell complex structure of Y_{n-1} under e_α to be a *refinement* L_α of the regular cell complex structure K_α on \mathbb{S}^{n-1} . We then define a regular cell complex

M_α on D_α^n by coning off every cell in L_α from the origin $o \in \mathbb{D}_\alpha^n$. Then the cells in M_α are the cones $Cone_{o_\alpha}(s)$, where s is a cell in L_α .




FIGURE 1.1. Barycentric subdivision of a 2-cell.

Since, by the induction assumption, every cell in Y_{n-1} is a simplex, its preimage s in \mathbb{S}^{n-1} is also a simplex, thus $Cone_o(s)$ is a simplex as well. We then attach each cell \mathbb{D}_α^n to Y_n by the original attaching map e_α . It is clear that the new cells $Cone_{o_\alpha}(s)$ are embedded in Y_n and each is naturally isomorphic to a simplex. Lastly, we set

$$Y := \bigcup_n Y_n.$$

Second barycentric subdivision. Note that the complex X' constructed above may not be a simplicial complex. The problem is that if x, y are distinct vertices of L_j , their images under the attaching map e_α could be the same (a point z). Thus the edges $[o_j, x], [o_j, y]$ in Y_{n+1} will intersect in the set $\{o_j, z\}$. However, if the complex X was regular, this problem does not arise and X' is a simplicial complex. Thus, in order to promote X to a simplicial complex (whose geometric realization is homeomorphic to $|X|$), we take the *second barycentric subdivision* X'' of X : Since X' is a regular cell complex, the complex X'' is naturally isomorphic to a simplicial complex.

Mapping cylinders. Let $f : X \rightarrow Y$ be a map of topological spaces. The *mapping cylinder* M_f of f is the quotient space of

$$X \times [0, 1] \sqcup Y$$

by the equivalence relation:

$$(x, 1) \sim f(x).$$

Similarly, given two maps $f_i : X \rightarrow Y_i, i = 0, 1$, we form the *double mapping cylinder* M_{f_1, f_2} , which is the quotient space of

$$X \times [0, 1] \sqcup Y_0 \sqcup Y_1$$

by the equivalence relation:

$$(x, i) \sim f_i(x), i = 0, 1.$$

If $f : X \rightarrow Y, f_i : X \rightarrow Y_i, i = 0, 1$, are cellular maps of cell complexes, then the corresponding mapping cylinders and double mapping cylinders also have natural structures of cell complexes.

Morphisms of almost regular complexes.

DEFINITION 1.49. Let X and Y be almost regular cell complexes. A cellular map $f : X \rightarrow Y$ is said to be *almost regular* if for every n -cell σ in X either:

- (a) f *collapses* σ , i.e. $f(\sigma) \subset Y^{(n-1)}$, or
- (b) f maps the interior of σ homeomorphically onto the interior of an n -cell in Y .

An almost regular map is *regular* or *noncollapsing* if only (b) occurs.

For instance, a simplicial map of simplicial complexes is always almost regular, while a simplicial topological embedding of simplicial complexes is noncollapsing.

CHAPTER 2

Metric spaces

2.1. General metric spaces

A *metric space* is a set X endowed with a function $\text{dist} : X \times X \rightarrow \mathbb{R}$ satisfying the following properties:

- (M1) $\text{dist}(x, y) \geq 0$ for all $x, y \in X$; $\text{dist}(x, y) = 0$ if and only if $x = y$;
- (M2) (Symmetry) for all $x, y \in X$, $\text{dist}(y, x) = \text{dist}(x, y)$;
- (M3) (Triangle inequality) for all $x, y, z \in X$, $\text{dist}(x, z) \leq \text{dist}(x, y) + \text{dist}(y, z)$.

The function dist is called *metric* or *distance function*. Occasionally, we will relax the axiom (M1) and allow $\text{dist}(x, y) = 0$ even for distinct points $x, y \in X$; we will also allow dist to take infinite values, in which case we will interpret triangle inequalities following the usual calculus conventions ($a + \infty = \infty$ for every $a \in \mathbb{R} \cup \{\infty\}$, etc.). With these changes in the definition, we will refer to dist as a *pseudo-distance* or *pseudo-metric*.

Notation. We will use the notation d or dist to denote the metric on a metric space X . For $x \in X$ and $A \subset X$ we will use the notation $\text{dist}(x, A)$ for the *minimal distance* from x to A , i.e.

$$\text{dist}(x, A) = \inf\{d(x, a) : a \in A\}.$$

Similarly, given two subsets $A, B \subset X$, we define their *minimal distance*

$$\text{dist}(A, B) = \inf\{d(a, b) : a \in A, b \in B\}.$$

For subsets $A, B \subset X$ we let

$$\text{pdist}_{\text{Haus}}(A, B) = \max\left(\sup_{a \in A} \text{dist}(a, B), \sup_{b \in B} \text{dist}(b, A)\right)$$

denote the *Hausdorff (pseudo-) distance* between A and B in X . Two subsets of X are called *Hausdorff-close* if they are within finite Hausdorff distance from each other. See Section 2.4 for further details on this distance and its generalizations.

Given two maps $f_i : (X, \text{dist}_X) \rightarrow (Y, \text{dist}_Y)$, $i = 1, 2$, we define the *distance between these maps*

$$\text{dist}(f_1, f_2) := \sup_{x \in X} \text{dist}(f_1(x), f_2(x)) \in [0, \infty].$$

Let (X, dist) be a metric space. We will use the notation $\mathcal{N}_R(A)$ to denote the *open R -neighborhood* of a subset $A \subset X$, i.e. $\mathcal{N}_R(A) = \{x \in X : \text{dist}(x, A) < R\}$. In particular, if $A = \{a\}$ then $\mathcal{N}_R(A) = B(a, R)$ is the open R -ball centered at a .

We will use the notation $\overline{\mathcal{N}}_R(A)$, $\overline{B}(a, R)$ to denote the corresponding *closed neighborhoods* and *closed balls*, defined by non-strict inequalities.

We denote by $S(x, r)$ the *sphere with center x and radius r* , i.e. the set

$$\{y \in X : \text{dist}(y, x) = r\}.$$

We will use the notation AB to denote a geodesic segment connecting the point A to the point B in X : Note that such a segment may be non-unique, so our notation is slightly ambiguous. Similarly, we will use the notation $\triangle(A, B, C)$ or $T(A, B, C)$ for a geodesic triangle in X with the vertices A, B, C . The *perimeter* of a triangle is the sum of its side-lengths (lengths of its edges). Lastly, we will use the notation $\blacktriangle(A, B, C)$ for a solid triangle in a surface with the given vertices A, B and C . Precise definitions of geodesic segments and triangles will be given in section 2.2.

A metric space is said to satisfy the *ultrametric inequality* if

$$\text{dist}(x, z) \leq \max(\text{dist}(x, y), \text{dist}(y, z)), \forall x, y, z \in X.$$

We will see some examples of ultrametric spaces in section 2.9.

Every norm $\|\cdot\|$ on a vector space V defines a metric on V :

$$\text{dist}(u, v) = \|u - v\|.$$

The standard examples of norms on the n -dimensional real vector space V are:

$$\|v\|_p = \left(\sum_{i=1}^n |x_i|^p \right)^{1/p}, \quad 1 \leq p < \infty,$$

and

$$\|v\|_{max} = \|v\|_\infty = \max\{|x_1|, \dots, |x_n|\}.$$

In what follows, our default assumption, unless stated otherwise, is that \mathbb{R}^n is equipped with the Euclidean metric, defined by the ℓ_2 -norm $\|v\|_2$; we will also use the notation \mathbb{E}^n for the Euclidean n -space.

EXERCISE 2.1. Show that the Euclidean plane \mathbb{E}^2 satisfies the *parallelogram identity*: If A, B, C, D are vertices of a parallelogram P in \mathbb{E}^2 with the diagonals AC and BD , then

$$(2.1) \quad d^2(A, B) + d^2(B, C) + d^2(C, D) + d^2(D, A) = d^2(A, C) + d^2(B, D),$$

i.e. sum of squares of the lengths of the sides of P equals the sum of squares of the length of the diagonals of P .

If X, Y are metric spaces, the *product metric* on the direct product $X \times Y$ is defined by the formula

$$(2.2) \quad d^2((x_1, y_1), (x_2, y_2)) = d^2(x_1, x_2) + d^2(y_1, y_2).$$

We will need a *separation lemma*, which is standard (see for instance [Mun75, §32]), but we include a proof for the convenience of the reader.

LEMMA 2.2. *Every metric space X is perfectly normal.*

PROOF. Let $A, V \subset X$ be disjoint closed subsets. Both functions $\text{dist}_A, \text{dist}_V$, which assign to $x \in X$ its minimal distance to A and to V respectively, are clearly continuous. Therefore, the ratio

$$\sigma(x) := \frac{\text{dist}_A(x)}{\text{dist}_V(x)}, \quad \sigma : X \rightarrow [0, \infty],$$

is continuous as well. Let $\tau : [0, \infty] \rightarrow [0, 1]$ be a continuous monotone function such that $\tau(0) = 0, \tau(\infty) = 1$, e.g.

$$\tau(y) = \frac{2}{\pi} \arctan(y), \quad y \neq \infty, \quad \tau(\infty) := 1.$$

Then the composition $\rho := \tau \circ \sigma$ satisfies the required properties. \square

A metric space (X, dist) is called *proper* if for every $p \in X$ and $R > 0$ the closed ball $\overline{B}(p, R)$ is compact. In other words, the distance function $d_p(x) = d(p, x)$ is proper.

DEFINITION 2.3. Given a function $\phi : \mathbb{R}_+ \rightarrow \mathbb{N}$, a metric space X is called ϕ -*uniformly discrete* if each ball $\overline{B}(x, r) \subset X$ contains at most $\phi(r)$ points. A metric space is called *uniformly discrete* if it is ϕ -uniformly discrete for some function ϕ .

Note that every uniformly discrete metric space necessarily has discrete topology.

Given two metric spaces $(X, \text{dist}_X), (Y, \text{dist}_Y)$, a map $f : X \rightarrow Y$ is an *isometric embedding* if for every $x, x' \in X$

$$\text{dist}_Y(f(x), f(x')) = \text{dist}_X(x, x').$$

The image $f(X)$ of an isometric embedding is called an *isometric copy of X in Y* .

A surjective isometric embedding is called an *isometry*, and the metric spaces X and Y are called *isometric*. A surjective map $f : X \rightarrow Y$ is called a *similarity with factor λ* if for all $x, x' \in X$,

$$\text{dist}_Y(f(x), f(x')) = \lambda \text{dist}_X(x, x').$$

The group of isometries of a metric space X is denoted $\text{Isom}(X)$. A metric space is called *homogeneous* if $\text{Isom}(X)$ acts transitively on X , i.e. for every $x, y \in X$ there exists an isometry $f : X \rightarrow X$ such that $f(x) = y$.

2.2. Length metric spaces

Throughout this book, by a *path* in a topological space X we mean a continuous map $\mathbf{p} : [a, b] \rightarrow X$. A path is said to *join* (or *connect*) two points x, y if $\mathbf{p}(a) = x, \mathbf{p}(b) = y$. We will frequently conflate a path and its image.

Given a path \mathbf{p} in a metric space X , one defines the *length* of \mathbf{p} as follows. A partition

$$a = t_0 < t_1 < \dots < t_{n-1} < t_n = b$$

of the interval $[a, b]$ defines a finite collection of points $\mathbf{p}(t_0), \mathbf{p}(t_1), \dots, \mathbf{p}(t_{n-1}), \mathbf{p}(t_n)$ in the space X . The *length* of \mathbf{p} is then defined to be

$$(2.3) \quad \text{length}(\mathbf{p}) = \sup_{a=t_0 < t_1 < \dots < t_n=b} \sum_{i=0}^{n-1} \text{dist}(\mathbf{p}(t_i), \mathbf{p}(t_{i+1}))$$

where the supremum is taken over all possible partitions of $[a, b]$ and all integers n .

If the length of \mathbf{p} is finite then \mathbf{p} is called *rectifiable*, otherwise the path \mathbf{p} is called *non-rectifiable*.

EXERCISE 2.4. Consider a C^1 -smooth path in the Euclidean space $\mathbf{p} : [a, b] \rightarrow \mathbb{R}^n$, $\mathbf{p}(t) = (x_1(t), \dots, x_n(t))$. Prove that its length (defined above) is given by the familiar formula

$$\text{length}(\mathbf{p}) = \int_a^b \sqrt{[x'_1(t)]^2 + \dots + [x'_n(t)]^2} dt.$$

Similarly, if (M, g) is a connected Riemannian manifold and dist is the Riemannian distance function (see section 3.3), then the two notions of length, given by equations (3.1) and by (2.3), coincide for smooth paths.

EXERCISE 2.5. Prove that the graph of the function $f : [0, 1] \rightarrow \mathbb{R}$,

$$f(x) = \begin{cases} x \sin \frac{1}{x} & \text{if } 0 < x \leq 1, \\ 0 & \text{if } x = 0, \end{cases}$$

is a non-rectifiable path joining $(0, 0)$ and $(1, \sin(1))$.

Let (X, dist) be a metric space. We define a new metric dist_ℓ on X , known as the *induced intrinsic metric*: $\text{dist}_\ell(x, y)$ is the infimum of the lengths of all rectifiable paths joining x to y .

EXERCISE 2.6. Show that:

1. dist_ℓ is a metric on X with values in $[0, \infty]$.
2. $\text{dist} \leq \text{dist}_\ell$.

Suppose that $\mathbf{p} : [0, b] \rightarrow X$ is a path joining x to y and realizing the finite infimum in the definition of the distance $D = \text{dist}_\ell(x, y)$. We will (re)parameterize \mathbf{p} by its arc-length:

$$\mathbf{q}(s) = \mathbf{p}(t),$$

where

$$s = \text{length}(\mathbf{p}|_{[0, t]}).$$

The resulting path $\mathbf{q} : [0, D] \rightarrow (X, \text{dist}_\ell)$ is called a *geodesic segment* in (X, dist_ℓ) .

Note that in a path metric space, *a priori*, not every two points are connected by a geodesic. We extend the notion of geodesic to general metric spaces: A *geodesic* in a metric space (X, dist) is an isometric embedding \mathbf{g} of an interval in \mathbb{R} into X . Note that this notion is different from the one in Riemannian geometry, where geodesics are isometric embeddings only *locally*, and need not be arc-length parameterized. A geodesic is called a *geodesic ray* if it is defined on an interval $(-\infty, a]$ or $[a, +\infty)$, and it is called *bi-infinite* or *complete* if it is defined on \mathbb{R} . As with paths, we will frequently conflate geodesics and their images.

EXERCISE 2.7. Prove that for (X, dist_ℓ) the two notions of geodesics (for maps of finite intervals) agree.

DEFINITION 2.8. A metric space (X, dist) such that $\text{dist} = \text{dist}_\ell$ is called a *length (or path) metric space*.

DEFINITION 2.9. A metric space X is called *geodesic* if every two points in X are connected by a geodesic path. A subset A in a metric space X is called *convex* if for every two points $x, y \in A$ there exists a geodesic $\gamma : [0, D] \rightarrow A$ connecting x and y .

EXERCISE 2.10. Each geodesic metric space is locally path-connected.

A *geodesic triangle* $T = T(A, B, C)$ or $\Delta(A, B, C)$ with vertices A, B, C in a metric space X is a collection of geodesic segments AB, BC, CA in X . These segments are called *edges* of T . We would like to emphasize that *triangles* in this book are 1-dimensional objects; we will use the terminology *solid triangle* to denote the corresponding 2-dimensional object.

Later on, in Chapters 4 and 11 we will use *generalized triangles*, where some edges are geodesic rays or, even, complete geodesics. The corresponding vertices of the generalized triangles will be *points of the ideal boundary* of X .

- EXAMPLES 2.11. (1) \mathbb{R}^n with the Euclidean metric is a geodesic metric space.
 (2) $\mathbb{R}^n \setminus \{0\}$ with the Euclidean metric is a length metric space, but not a geodesic metric space.
 (3) The unit circle \mathbb{S}^1 with the metric inherited from the Euclidean metric of \mathbb{R}^2 (the chordal metric) is not a length metric space. The induced intrinsic metric on \mathbb{S}^1 is the one that measures distances as angles in radians, it is the distance function of the Riemannian metric induced by the embedding $\mathbb{S}^1 \rightarrow \mathbb{R}^2$.
 (4) The Riemannian distance function dist defined for a connected Riemannian manifold (M, g) (see section 3.3) is a path-metric. If this metric is complete, then the path-metric is geodesic.
 (5) Every connected graph equipped with the standard distance function (see section 2.3) is a geodesic metric space.

EXERCISE 2.12. If X, Y are geodesic metric spaces, so is $X \times Y$. If X, Y are path-metric spaces, so is $X \times Y$. Here $X \times Y$ is equipped with the product metric defined by the formula (2.2).

THEOREM 2.13 (Hopf–Rinow Theorem [Gro07]). *If a length metric space is complete and locally compact, then it is geodesic and proper.*

EXERCISE 2.14. Construct an example of a metric space X which is not a length metric space, so that X is complete, locally compact, but is not proper.

2.3. Graphs as length spaces

Let Γ be a connected graph. Recall that we are conflating Γ and its geometric realization; the notation $x \in \Gamma$ below will simply mean that x is a point of the geometric realization.

We introduce a path-metric dist on the geometric realization of Γ as follows. We declare every edge of Γ to be isometric to the unit interval in \mathbb{R} . **Then** the distance between any vertices of Γ is the combinatorial length of the shortest edge-path connecting these vertices. Of course, points of the interiors of edges of Γ are not connected by any edge-paths. Thus, we consider *fractional* edge-paths, where in addition to the edges of Γ we allow intervals contained in the edges. The length of such a fractional path is the sum of lengths of the intervals in the path. Then, for $x, y \in \Gamma$,

$$\text{dist}(x, y) = \inf_{\mathbf{p}} (\text{length}(\mathbf{p})),$$

where the infimum is taken over all fractional edge-paths \mathbf{p} in Γ connecting x to y . The metric dist is called the *standard* metric on Γ .

EXERCISE 2.15. a. Show that infimum is the same as minimum in this definition.

b. Show that every edge of Γ (treated as a unit interval) is isometrically embedded in (Γ, dist) .

c. Show that dist is a path-metric.

d. Show that dist is a complete metric.

The notion of a standard metric on a graph generalizes to the concept of a *metric graph*, which is a connected graph Γ equipped with a path-metric dist_ℓ . Such path-metric is, of course, uniquely determined by the lengths of edges of Γ with respect to the metric d .

EXAMPLE 2.16. Consider Γ which is the complete graph on 3 vertices (a triangle) and declare that two edges e_1, e_2 of Γ are unit intervals and the remaining edge e_3 of Γ has length 3. Let dist_ℓ be the corresponding path-metric on Γ . Then e_3 is not isometrically embedded in $(\Gamma, \text{dist}_\ell)$.

Diameters in graphs. Recall that the diameter of a metric space is the supremum of distances between its points. Suppose that Γ is a connected graph equipped with the standard metric. A subgraph Γ' of Γ is called a *diameter* of Γ , if Γ' is isomorphic to the n -interval I_n , where $n = \text{diam}(\Gamma)$ is the diameter of Γ . This should not cause a confusion since one *diameter* is a number, while the other *diameter* is a subgraph.

EXERCISE 2.17. Suppose that T is a tree of finite diameter n . Then:

1. Any two diameters of T have non-empty intersection.
2. The intersection C of all diameters of T is non-empty. The subtree C is the *core* of T .
3. Each connected component of $T \setminus C$ has diameter strictly less than n .

EXERCISE 2.18. Show that each connected graph of finite valence and infinite diameter contains an isometrically embedded copy of \mathbb{R}_+ .

LEMMA 2.19. *Suppose that $f : H \rightarrow T$ is a map of graphs, where H is the half-line and T is a tree, such that $\text{diam}(f(H))$ is finite. Then there exists a vertex $v \in T$ such that $f^{-1}(v)$ is unbounded.*

PROOF. The proof is by induction on $D = \text{diam}(f(H))$. If $D = 1$, there is nothing to prove. Suppose that D is at least 2. The image subgraph $f(H)$ is connected and, hence, is a subtree $A \subset T$. Let $C \subset A$ be the core of A as in Exercise 2.17. If there exists a vertex $a \in V(C)$ with infinite preimage $f^{-1}(a)$, we are done. Otherwise, there exists n such that

$$f([n, \infty)),$$

is disjoint from C . Since the subgraph $H' = [n, \infty)$ is isomorphic to the half-line H , we obtain a new map of graph $f|_{H'} : H' \rightarrow T$. The diameter of the image of this map is strictly less than D . Lemma follows from the induction hypothesis. \square

2.4. Hausdorff and Gromov–Hausdorff distances. Nets

The Hausdorff distance between two distinct spaces (for instance, between a space and a dense subspace in it) can be zero. The Hausdorff distance becomes a genuine distance only when restricted to certain classes of subsets, for instance,

to the class of compact subsets of a metric space. Still, for simplicity, we call it a *distance* or a *metric* in all cases.

The Hausdorff distance defines the topology of *Hausdorff-convergence* on the set \mathcal{K}_X of compact subsets of a metric space X . This topology extends to the set \mathcal{C}_X of closed subsets of X as follows. Given $\epsilon > 0$ and a compact $K \subset X$ we define the neighborhood $U_{\epsilon, K}$ of a closed subset $C \in \mathcal{C}_X$ to be

$$\{Z \in \mathcal{C}_X : \text{dist}_{Haus}(Z \cap K, C \cap K) < \epsilon\}.$$

This system of neighborhoods generates a topology on \mathcal{C}_X , called *Chabauty topology*. Thus, a sequence $C_i \in \mathcal{C}_X$ converges to a closed subset $C \in \mathcal{C}_X$ if and only if for every compact subset $K \subset X$,

$$\lim_{i \rightarrow \infty} (C_i \cap K) = C \cap K,$$

where the limit is in the topology of Hausdorff-convergence.

M. Gromov defined in [Gro81a, section 6] the *modified Hausdorff pseudo-distance* (also called the *Gromov-Hausdorff pseudo-distance*) on the class of proper metric spaces:

$$(2.4) \quad \text{pdist}_{GH}((X, d_X), (Y, d_Y)) = \inf_{(x, y) \in X \times Y} \inf\{\epsilon > 0 \mid \exists \text{ a pseudo-metric}$$

dist on $M = X \sqcup Y$, such that $\text{dist}(x, y) < \epsilon$, $\text{dist}|_X = d_X$, $\text{dist}|_Y = d_Y$ and

$$B(x, 1/\epsilon) \subset \mathcal{N}_\epsilon(Y), B(y, 1/\epsilon) \subset \mathcal{N}_\epsilon(X)\}.$$

For homogeneous metric spaces the modified Hausdorff pseudo-distance coincides with the pseudo-distance for the pointed metric spaces:

$$(2.5) \quad \text{dist}_H((X, d_X, x_0), (Y, d_Y, y_0)) = \inf\{\epsilon > 0 \mid \exists \text{ a pseudo-metric}$$

dist on $M = X \sqcup Y$ such that $\text{dist}(x_0, y_0) < \epsilon$, $\text{dist}|_X = d_X$, $\text{dist}|_Y = d_Y$,

$$B(x_0, 1/\epsilon) \subset \mathcal{N}_\epsilon(Y), B(y_0, 1/\epsilon) \subset \mathcal{N}_\epsilon(X)\}.$$

This pseudo-distance becomes a metric when restricted to the class of proper pointed metric spaces. Note that since we use pseudo-metrics in order to define d_{GH} and d_H , instead of considering pseudo-metrics on the disjoint union $X \sqcup Y$, we can as well consider pseudo-metrics on spaces Z such that X, Y embed isometrically in Z .

In order to simplify the terminology we shall refer to all three pseudo-distances as ‘distances’ or ‘metrics.’

One can associate to every metric space (X, dist) a discrete metric space at finite Hausdorff distance from X , as follows.

DEFINITION 2.20. An ϵ -separated subset A in X is a subset such that

$$\text{dist}(a_1, a_2) \geq \epsilon, \forall a_1, a_2 \in A, a_1 \neq a_2.$$

A subset S of a metric space X is said to be r -dense in X if the Hausdorff distance between S and X is at most r . In other words, for every $x \in X$, we have the inequality $\text{dist}(x, S) \leq r$.

DEFINITION 2.21. An ε -separated δ -net in a metric space X is a subset of X that is ε -separated and δ -dense.

An ε -separated net in X is a subset that is ε -separated and 2ε -dense.

When the constants ε and δ are not relevant we shall not mention them and simply speak of *separated nets*.

LEMMA 2.22. A maximal (with respect to inclusion) δ -separated set in X is a δ -separated net in X .

PROOF. Let N be a maximal δ -separated subset in X . For every $x \in X \setminus N$, the set $N \cup \{x\}$ is no longer δ -separated, by maximality of N . Hence there exists $y \in N$ such that $\text{dist}(x, y) < \delta$. \square

By Zorn's lemma a maximal δ -separated subset always exists. Thus, every metric space contains a δ -separated net, for any $\delta > 0$.

EXERCISE 2.23. Prove that if (X, dist) is compact then every separated net in X is finite; hence, every separated subset in X is finite.

DEFINITION 2.24 (Rips complex). Let (X, dist) be a metric space. For $R \geq 0$ we define a simplicial complex $\text{Rips}_R(X)$: Its vertices are points of X ; vertices x_0, x_1, \dots, x_n span a simplex if and only if for all i, j ,

$$\text{dist}(x_i, x_j) \leq R.$$

The simplicial complex $\text{Rips}_R(X)$ is called the *R-Rips complex of X*.

We will discuss Rips complexes in more detail in Section 9.2.1.

REMARK 2.25. The complex $\text{Rips}_r(X)$ was first introduced by Leopold Vietoris in [Vie27], who was primarily interested in the case of compact metric spaces X and small values of r . This complex was reinvented by Eliyahu Rips in 1980s with the primary goal of studying hyperbolic groups, where X is a hyperbolic group equipped with the word metric and r is large. Accordingly, the complex $\text{Rips}_r(X)$ is also known as the *Vietoris complex* and the *Vietoris-Rips complex*.

2.5. Lipschitz maps and Banach-Mazur distance

If one attempts to think of metrics spaces *categorically*, one wonders what is the right notion of *morphism* in metric geometry. It turns out that depending on the situation, one has to use different notions of morphisms. Lipschitz (especially 1-Lipschitz) and locally Lipschitz maps appear to be the most useful. However, as we will see throughout the book, other classes of maps are also important, especially quasisymmetric, quasisymmetric and uniformly proper maps.

2.5.1. Lipschitz and locally Lipschitz maps. A map $f : X \rightarrow Y$ between two metric spaces (X, dist_X) , (Y, dist_Y) is *L-Lipschitz*, where L is a positive number, if for all $x, x' \in X$

$$\text{dist}_Y(f(x), f(x')) \leq L \text{dist}_X(x, x').$$

A map which is L -Lipschitz for some $L > 0$ is called simply *Lipschitz*.

EXERCISE 2.26. 1. Show that every L -Lipschitz path $\mathbf{p} : [0, 1] \rightarrow X$ is rectifiable and $\text{length}(\mathbf{p}) \leq L$.

2. Show that a map $f : X \rightarrow Y$ is an isometry if and only if f is 1-Lipschitz and admits a 1-Lipschitz inverse.

For a Lipschitz function $f : X \rightarrow \mathbb{R}$ let $\text{Lip}(f)$ denote

$$(2.6) \quad \text{Lip}(f) := \inf\{L : f \text{ is } L\text{-Lipschitz}\}$$

EXERCISE 2.27. Suppose that f, g are Lipschitz functions on X . Let $\|f\|, \|g\|$ denote the sup-norms of f and g on X . Show that

1. “Sum rule”: $\text{Lip}(f + g) \leq \text{Lip}(f) + \text{Lip}(g)$.
2. “Product rule”: $\text{Lip}(fg) \leq \text{Lip}(f)\|g\| + \text{Lip}(g)\|f\|$.
3. “Ratio rule”:

$$\text{Lip}\left(\frac{f}{g}\right) \leq \frac{\text{Lip}(f)\|g\| + \text{Lip}(g)\|f\|}{\inf_{x \in X} g^2(x)}.$$

Note that in case when f is a smooth function on a Riemannian manifold (e.g., on \mathbb{R}^n), these formulae follow from the formulae for the derivatives of the sum, product and ratio of two functions.

The following is a fundamental theorem about Lipschitz maps between Euclidean spaces:

THEOREM 2.28 (Rademacher Theorem, see Theorem 3.1 in [Hei01]). *Let U be an open subset of \mathbb{R}^n and let $f : U \rightarrow \mathbb{R}^m$ be Lipschitz. Then f is differentiable at almost every point in U .*

A map $f : X \rightarrow Y$ is called *locally Lipschitz* if for every $x \in X$ there exists $\epsilon > 0$ so that the restriction $f|_{B(x, \epsilon)}$ is Lipschitz. We let $\text{Lip}_{\text{loc}}(X; Y)$ denote the space of locally Lipschitz maps $X \rightarrow Y$. We set $\text{Lip}_{\text{loc}}(X) := \text{Lip}_{\text{loc}}(X; \mathbb{R})$.

EXERCISE 2.29. Fix a point p in a metric space (X, dist) and define the function dist_p by $\text{dist}_p(x) := \text{dist}(x, p)$. Show that this function is 1-Lipschitz. Prove the same for the function $\text{dist}_A(x) = \text{dist}(x, A)$, where $A \subset X$ is a non-empty subset.

LEMMA 2.30 (Lipschitz bump-function). *Let $0 < R < \infty$. Then there exists a $\frac{1}{R}$ -Lipschitz function $\varphi = \varphi_{p, R}$ on X such that*

1. $\text{Supp}(\varphi) = \overline{B}(p, R)$.
2. $\varphi(p) = 1$.
3. $0 \leq \varphi \leq 1$ on X .

PROOF. We first define the function $\zeta : \mathbb{R}_+ \rightarrow [0, 1]$ which vanishes on the interval $[R, \infty)$, is linear on $[0, R]$ and equals 1 at 0. Then ζ is $\frac{1}{R}$ -Lipschitz. Now take $\varphi := \zeta \circ \text{dist}_p$. \square

LEMMA 2.31 (Lipschitz partition of unity). *Suppose that we are given a locally finite covering of a metric space X by a countable set of open R_i -balls $B_i := B(x_i, R_i)$, $i \in I \subset \mathbb{N}$. Then there exists a collection of Lipschitz functions η_i , $i \in I$, so that:*

1. $\sum_i \eta_i \equiv 1$.
2. $0 \leq \eta_i \leq 1$, $\forall i \in I$.
3. $\text{Supp}(\eta_i) \subset \overline{B}(x_i, R_i)$, $\forall i \in I$.

PROOF. For each i define the bump-function using Lemma 2.30:

$$\varphi_i := \varphi_{x_i, R_i}.$$

Then the function

$$\varphi := \sum_{i \in I} \varphi_i$$

is positive on X . Finally, define

$$\eta_i := \frac{\varphi_i}{\varphi}.$$

It is clear that the functions η_i satisfy all the required properties. \square

REMARK 2.32. Since the collection of balls $\{B_i\}$ is locally finite, it is clear that the function

$$L(x) := \sup_{i \in I, \eta_i(x) \neq 0} \text{Lip}(\eta_i)$$

is bounded on compact sets in X , however, in general, it is unbounded on X . We refer the reader to the equation (2.6) for the definition of $\text{Lip}(\eta_i)$.

From now on, we assume that X is a proper metric space.

PROPOSITION 2.33. $\text{Lip}_{\text{loc}}(X)$ is a dense subset in $C(X)$, the space of continuous functions $X \rightarrow \mathbb{R}$, equipped with the compact-open topology.

PROOF. Fix a base-point $o \in X$ and let $A_n, n \in \mathbb{N}$, denote the annulus

$$\{x \in X : n - 1 \leq \text{dist}(x, o) \leq n\}.$$

Let f be a continuous function on X . Pick $\epsilon > 0$. Our goal is to find a locally Lipschitz function g on X so that $|f(x) - g(x)| < \epsilon$ for all $x \in X$. Since f is uniformly continuous on compact sets, for each $n \in \mathbb{N}$ there exists $\delta = \delta(n, \epsilon)$ such that

$$\forall x, x' \in A_n, \quad \text{dist}(x, x') < \delta \Rightarrow |f(x) - f(x')| < \epsilon.$$

Therefore for each n we find a finite subset

$$X_n := \{x_{n,1}, \dots, x_{n,m_n}\} \subset A_n$$

so that for $r := \delta(n, \epsilon)/4, R := 2r$, the open balls $B_{n,j} := B(x_{n,j}, r)$ cover A_n . We reindex the set of points $\{x_{n,j}\}$ and the balls $B_{n,j}$ with a countable set I . Thus, we obtain an open locally finite covering of X by the balls $B_j, j \in I$. Let $\{\eta_j, j \in I\}$ denote the corresponding Lipschitz partition of unity. It is then clear that

$$g(x) := \sum_{i \in I} \eta_i(x) f(x_i)$$

is a locally Lipschitz function. For $x \in B_i$ let $J \subset I$ be such that

$$x \notin B(x_j, R_j), \quad \forall j \notin J.$$

Then $|f(x) - f(x_j)| < \epsilon$ for all $j \in J$. Therefore

$$|g(x) - f(x)| \leq \sum_{j \in J} \eta_j(x) |f(x_j) - f(x)| < \epsilon \sum_{j \in J} \eta_j(x) = \epsilon \sum_{i \in I} \eta_j(x) = \epsilon.$$

It follows that $|f(x) - g(x)| < \epsilon$ for all $x \in X$. \square

A relative version of Proposition 2.33 also holds:

PROPOSITION 2.34. Let $A \subset X$ be a closed subset contained in a subset U which is open in X . Then, for every $\epsilon > 0$ and every continuous function $f \in C(X)$ there exists a function $g \in C(X)$ so that:

1. g is locally Lipschitz on $X \setminus U$.
2. $\|f - g\| < \epsilon$.
3. $g|_A = f|_A$.

PROOF. For the closed set $V := X \setminus U$ pick a continuous function $\rho = \rho_{A,V}$ separating the sets A and V . Such a function exists, by Lemma 2.2. According to Proposition 2.33, there exists $h \in \text{Lip}_{\text{loc}}(X)$ such that $\|f - h\| < \epsilon$. Then take

$$g(x) := \rho(x)h(x) + (1 - \rho(x))f(x).$$

We leave it to the reader to verify that g satisfies all the requirements of the proposition. \square

2.5.2. Bi-Lipschitz maps. The Banach-Mazur distance. A map of metric spaces $f : X \rightarrow Y$ is *L-bi-Lipschitz*, for some constant $L \geq 1$, if it is a bijection and both f and f^{-1} are L -Lipschitz for some L ; equivalently, f is surjective and there exists a constant $L \geq 1$ such that for every $x, x' \in X$

$$\frac{1}{L} \text{dist}_X(x, x') \leq \text{dist}_Y(f(x), f(x')) \leq L \text{dist}_X(x, x').$$

A *bi-Lipschitz embedding* is defined by dropping surjectivity assumption.

EXAMPLE 2.35. Let (M, g) and (N, h) be two connected Riemannian manifolds (see section 3.3). Then a diffeomorphism $f : M \rightarrow N$ is L -bi-Lipschitz if and only if

$$L^{-1} \leq \sqrt{\frac{f^*h}{g}} \leq L.$$

In other words, for every tangent vector $v \in TM$,

$$L^{-1} \leq \frac{|df(v)|}{|v|} \leq L.$$

If there exists a bi-Lipschitz map $f : X \rightarrow Y$, the metric spaces (X, dist_X) and (Y, dist_Y) are called *bi-Lipschitz equivalent* or *bi-Lipschitz homeomorphic*. If dist_1 and dist_2 are two distances on the same metric space X such that the identity map $\text{id} : (X, \text{dist}_1) \rightarrow (X, \text{dist}_2)$ is bi-Lipschitz, then we say that dist_1 and dist_2 are *bi-Lipschitz equivalent*.

- EXAMPLES 2.36. (1) Any two metrics d_1, d_2 on \mathbb{R}^n defined by two norms on \mathbb{R}^n are bi-Lipschitz equivalent.
 (2) Any two left-invariant Riemannian metrics on a connected real Lie group define bi-Lipschitz equivalent distance functions.

EXAMPLE 2.37. If $T : V \rightarrow W$ is a continuous linear map between Banach spaces, then

$$\text{Lip}(T) = \|T\|,$$

the operator norm of T .

The *Banach-Mazur distance* $\text{dist}_{BM}(V, W)$ between two Banach spaces V and W is

$$\log \left(\inf_{T:V \rightarrow W} (\|T\| \cdot \|T^{-1}\|) \right),$$

where the infimum is taken over all bounded invertible linear maps $T : V \rightarrow W$ with bounded inverse. The reader can think of dist_{BM} as a Banach-space analogue (and precursor) of the Gromov-Hausdorff distance between metric spaces.

EXERCISE 2.38. Show that dist_{BM} is a metric on the set of n -dimensional Banach spaces.

THEOREM 2.39 (John's Theorem, see e.g. [MS86], Theorem 3.3). *For every pair of n -dimensional normed vector spaces V, W , $\text{dist}_{BM}(V, W) \leq \log(n)$.*

2.6. Hausdorff dimension

In this section we review the concept of *Hausdorff dimension* for metric spaces.

We let ω_n denote the volume of the unit Euclidean n -ball. The function ω_n extends to \mathbb{R}_+ by the formula

$$\omega_\alpha = \frac{\pi^{\alpha/2}}{\Gamma(1 + \alpha/2)},$$

where Γ is the Gamma-function.

Let K be a metric space and $\alpha > 0$. The α -Hausdorff measure $\mu_\alpha(K)$ is defined as

$$(2.7) \quad \mu_\alpha \liminf_{r \rightarrow 0} \sum_i r_i^\alpha,$$

where the infimum is taken over all countable coverings of K by balls $B(x_i, r_i)$, $r_i \leq r$. The motivation for this definition is that the volume of the Euclidean r -ball of dimension $n \in \mathbb{N}$ is $\omega_n r^n$; hence, the Lebesgue measure of a (measurable) subset of \mathbb{R}^n equals its n -Hausdorff measure. Euclidean spaces, of course, have integer dimension, the point of Hausdorff measure and dimension is to extend the definition to the non-integer case.

EXERCISE 2.40. Suppose that $f : X \rightarrow Y$ is an L -Lipschitz map between metric spaces. Show that

$$\mu_\alpha(f(X)) \leq L^\alpha \mu_\alpha(X).$$

The *Hausdorff dimension* of the metric space K is defined as:

$$\dim_H(K) := \inf\{\alpha : \mu_\alpha(K) = 0\}.$$

EXERCISE 2.41. Verify that the Euclidean space \mathbb{R}^n has Hausdorff dimension n .

We will need the following theorem:

THEOREM 2.42 (L. Sznirelman; see [HW41]). *The covering dimension $\dim(X)$ of a proper metric space X is at most the Hausdorff dimension $\dim_H(X)$.*

Let $A \subset X$ be a closed subset. Recall that $\mathbb{D}^n := \overline{B}(0, 1) \subset \mathbb{R}^n$ denotes the closed unit ball in \mathbb{R}^n . Define

$$C(X, A; \mathbb{B}^n) := \{f \in C(X, \mathbb{D}^n) : f(A) \subset \mathbb{S}^{n-1} = \partial \mathbb{D}^n\}.$$

An immediate consequence of Proposition 2.34 is the following.

COROLLARY 2.43. *For every function $f \in C(X, A; \mathbb{D}^n)$ and open set $U \subset X$ containing A , there exists a sequence of functions $g_i \in C(X, A; \mathbb{D}^n)$ such that for all $i \in \mathbb{N}$:*

1. $g_i|_A = f|_A$.
2. $g_i \in \text{Lip}_{loc}(X \setminus U; \mathbb{R}^n)$.

For a continuous map $f : X \rightarrow \mathbb{D}^n$ define $A = A_f$ as

$$A := f^{-1}(\mathbb{S}^{n-1}).$$

DEFINITION 2.44. The map f is *inessential* if it is homotopic rel. A to a map $f' : X \rightarrow \mathbb{S}^{n-1}$. An *essential map* is the one which is not inessential.

We will be using the following characterization of the covering dimension, due to P. S. Alexandrov:

THEOREM 2.45 (P. S. Alexandrov, see Theorem III.5 in [Nag83]). *A space X satisfies $\dim(X) < n$ if and only if every continuous map $f : X \rightarrow \mathbb{D}^n$ is inessential.*

We are now ready to prove Theorem 2.42.

Proof of Theorem 2.42. Suppose that $\dim_H(X) < n$. We will prove that $\dim(X) < n$ as well. We need to show that every continuous map $f : X \rightarrow \mathbb{D}^n$ is inessential. Let D denote the annulus

$$\{x \in \mathbb{R}^n : 1/2 < |x| \leq 1\}.$$

Set $A := f^{-1}(\mathbb{S}^{n-1})$ and $U := f^{-1}(D)$.

Take the sequence g_i given by Corollary 2.43. Since each g_i is homotopic to f rel. A , it suffices to show that some g_i is inessential. Since $f = \lim_{i \rightarrow \infty} g_i$, it follows that for all sufficiently large i ,

$$g_i(U) \cap B\left(0, \frac{1}{3}\right) = \emptyset.$$

We claim that the image of every such g_i misses a point in $B\left(0, \frac{1}{3}\right)$. Indeed, since $\dim_H(X) < n$, the n -dimensional Hausdorff measure of X is zero. However, each

$$g_i|_{X \setminus U}$$

is locally Lipschitz. Therefore $g_i(X \setminus U)$ has zero n -dimensional Hausdorff (and hence Lebesgue) measure, see Exercise 2.40. It follows that $g_i(X)$ misses a point y in $B\left(0, \frac{1}{3}\right)$. Composing g_i with the retraction $\mathbb{D}^n \setminus \{y\} \rightarrow \mathbb{S}^{n-1}$ we get a map $f' : X \rightarrow \mathbb{S}^{n-1}$ which is homotopic to f rel. A . Thus f is inessential and, therefore, $\dim(X) < n$. \square

2.7. Norms and valuations

In this and the following section we describe certain metric spaces of algebraic origin that will be used in the proof of the Tits alternative. We refer the reader to [Lan02, Chapter XII] for more details.

A *norm* or an *absolute value* on a ring R is a function $|\cdot|$ from R to \mathbb{R}_+ , which satisfies the following axioms:

1. $|x| = 0 \iff x = 0$.
2. $|xy| = |x| \cdot |y|$.
3. $|x + y| \leq |x| + |y|$.

An element $x \in R$ such that $|x| = 1$ is called a *unit*. A norm $|\cdot|$ is called *non-archimedean* if it satisfies the *ultrametric inequality*

$$|x + y| \leq \max(|x|, |y|).$$

According to Ostrowski's theorem, [Cas86], Theorem 1.1, if $(F, \|\cdot\|)$ is a normed field which is *not non-archimedean*, then there exists an isometric homeomorphic embedding

$$f : (F, \|\cdot\|) \hookrightarrow (\mathbb{C}, |\cdot|^\alpha),$$

where \mathbb{C} is equipped with the Euclidean norm given by a power of the absolute value of complex numbers, $|\cdot|^\alpha$, $\alpha > 0$.

Such norms $\|\cdot\|$ are called *archimedean*. We will be primarily interested in normed archimedean fields which are \mathbb{R} and \mathbb{C} with the usual norms given by the absolute value. In the case $F = \mathbb{Q}$, Ostrowski's theorem can be made even more precise: Every norm $\|\cdot\|$ on \mathbb{Q} arises as a power of the Euclidean norm or of a p -adic norm; see [Cas86], Theorem 2.1.

Below is an alternative approach to non-archimedean normed rings R . A function $\nu : R \rightarrow \mathbb{R} \cup \{\infty\}$ is called a *valuation* if it satisfies the following axioms:

1. $\nu(x) = \infty \iff x = 0$.
2. $\nu(xy) = \nu(x) + \nu(y)$.
3. $\nu(x + y) \geq \min(\nu(x), \nu(y))$.

Therefore, one converts a valuation to a non-archimedean norm by setting

$$|x| = c^{-\nu(x)}, x \neq 0, \quad |0| = 0,$$

where $c > 0$ is a fixed real number.

REMARK 2.46. More generally, one also considers valuations with values in arbitrary ordered abelian groups, but we will not need this.

A normed ring R is said to be *local* if it is locally compact as a metric space; a normed ring R is said to be *complete* if it is complete as a metric space. A norm on a field F is said to be *discrete* if the image Γ of $|\cdot| : F^\times = F \setminus \{0\} \rightarrow \mathbb{R}^\times$ is an infinite cyclic group. If a norm is discrete, then an element $\pi \in F$ such that $|\pi|$ is a generator of Γ satisfying $|\pi| < 1$, is called a *uniformizer* of F . If F is a field with valuation ν , then the subset

$$O_\nu = \{x \in F : \nu(x) \geq 0\}$$

is a subring in F , the *valuation ring* or the *ring of integers* in F .

EXERCISE 2.47. 1. Verify that every non-zero element of a field F with discrete norm has the form $\pi^k u$, where u is a unit.

2. Verify that every discrete norm is non-archimedean.

Below are the two main examples of fields with discrete norms:

1. The field \mathbb{Q}_p of p -adic numbers. Fix a prime number p . For each number $x = q/p^n \in \mathbb{Q}$ (where both numerator and denominator of q are not divisible by p) set $|x|_p := p^{-n}$. Then $|\cdot|_p$ is a non-archimedean norm on \mathbb{Q} , called the *p -adic norm*. The completion of \mathbb{Q} with respect to the p -adic norm is the field of p -adic numbers \mathbb{Q}_p . The ring of p -adic integers O_p intersects \mathbb{Q} along the subset consisting of (reduced) fractions $\frac{n}{m}$ where $m, n \in \mathbb{Z}$ and m is not divisible by p . Note that p is a uniformizer of \mathbb{Q}_p .

REMARK 2.48. We will not use the common notation \mathbb{Z}_p for O_p , in order to avoid the confusion with finite cyclic groups.

EXERCISE 2.49. Verify that O_p is open in \mathbb{Q}_p . Hint: Use the fact that $|x+y|_p \leq 1$ provided that $|x|_p \leq 1, |y|_p \leq 1$.

Recall that one can describe real numbers using infinite decimal sequences. There is a similar description of p -adic numbers using “base p arithmetic.” Namely,

we can identify p -adic numbers with semi-infinite Laurent series

$$\sum_{k=-n}^{\infty} a_k p^k,$$

where $n \in \mathbb{Z}$ and $a_k \in \{0, \dots, p-1\}$. Operations of addition and multiplication here are the usual operations with power series where we treat p as a formal variable, the only difference is that we still have to “carry to the right” as in the usual decimal arithmetic.

With this identification, $|x|_p = p^n$, where a_{-n} is the first non-zero coefficient in the power series. The corresponding valuation is $v(x) = -n$, $c = p$. In particular, the ring O_p is identified with the set of series

$$\sum_{k=0}^{\infty} a_k p^k.$$

REMARK 2.50. In other words, one can describe p -adic numbers as left-infinite sequences of (base p) digits

$$\cdots a_m a_{m-1} \cdots a_0 . a_{-1} \cdots a_{-n}$$

where $\forall i, a_i \in \{0, \dots, p-1\}$, and the algebraic operations require “carrying to the left” instead of carrying to the right.

EXERCISE 2.51. Show that in \mathbb{Q}_p ,

$$\sum_{k=0}^{\infty} p^k = \frac{1}{1-p}.$$

2. Let A be a field. Consider the ring $R = A[t, t^{-1}]$ of Laurent polynomials

$$f(t) = \sum_{k=n}^m a_k t^k.$$

Set $v(0) = \infty$ and for non-zero f let $v(f)$ be the least n so that $a_n \neq 0$. In other words, $v(f)$ is the order of vanishing of f at $0 \in R$.

EXERCISE 2.52. 1. Verify that v is a valuation on R . Define $|f| := e^{-v(f)}$.

2. Verify that the completion \widehat{R} of R with respect to the above norm is naturally isomorphic to the ring of semi-infinite formal Laurent series

$$f = \sum_{k=n}^{\infty} a_k t^k,$$

where $v(f)$ is the minimal n such that $a_n \neq 0$.

Let $A(t)$ be the field of rational functions in the variable t . We embed $A(t)$ in \widehat{R} using the rule

$$\frac{1}{1-at} = 1 + \sum_{n=1}^{\infty} a^n t^n.$$

If A is algebraically closed, every rational function is a product of a polynomial function and several functions of the form

$$\frac{1}{a_i - t},$$

so we obtain an embedding $A(t) \hookrightarrow \widehat{R}$ in this case. If A is not algebraically closed, proceed as follows. First, construct, as above, an embedding ι of $\bar{A}(t)$ to the completion of $\bar{A}[t, t^{-1}]$, where \bar{A} is the algebraic closure of A . Next, observe that this embedding is equivariant with respect to the Galois group $\text{Gal}(\bar{A}/A)$, where $\sigma \in \text{Gal}(\bar{A}/A)$ acts on Laurent series

$$f = \sum_{k=n}^{\infty} a_k t^k, a \in \bar{A},$$

by

$$f^\sigma = \sum_{k=n}^{\infty} a_k^\sigma t^k.$$

Therefore, $\iota(A(t)) \subset \widehat{R}$, $R = A[t, t^{-1}]$.

In any case, we obtain a norm on $A(t)$ by restricting the norm in \widehat{R} . Since $R \subset \iota A(t)$, it follows that \widehat{R} is the completion of $\iota A(t)$. In particular, \widehat{R} is a complete normed field.

EXERCISE 2.53. 1. Verify that \widehat{R} is local if and only if A is finite.

2. Show that t is a uniformizer of \widehat{R} .

3. At the first glance, it looks like \mathbb{Q}_p is the same as \widehat{R} for $A = \mathbb{Z}_p$, since elements of both are described using formal power series with coefficients in $\{0, \dots, p-1\}$. What is the difference between these fields?

The same construction works with Laurent polynomials of several variables. We let A be a field, $T = \{t_1, \dots, t_l\}$ a finite set of variables and consider first the ring of Laurent polynomials:

$$A[t_1^{\pm 1}, \dots, t_n^{\pm 1}]$$

in these variables. The degree of a monomial

$$a t_1^{k_1} \dots t_n^{k_n}$$

with $a \neq 0$ is defined as the sum

$$k_1 + \dots + k_n.$$

For a general Laurent polynomial p in the variables T , set $v(p) = d$ iff d is the lowest degree of all non-zero monomials in p . With this definition, one again gets a complete norm $|\cdot|$ on the field $A(t_1, \dots, t_n)$ of rational functions in the variables t_i , where

$$|p| = e^{-v(p)}.$$

EXERCISE 2.54. If A is finite, then the normed field $(A(t_1, \dots, t_n), |\cdot|)$ is local.

Similarly, we have the following lemma:

LEMMA 2.55. \mathbb{Q}_p is a local field.

PROOF. It suffices to show that the ring O_p of p -adic integers is compact. Since \mathbb{Q}_p is complete, we only need to show that O_p is closed and totally bounded, i.e. for every $\epsilon > 0$, O_p has a finite cover by closed ϵ -balls. The fact that O_p is closed follows from the fact that $|\cdot|_p : \mathbb{Q}_p \rightarrow \mathbb{R}$ is continuous and O_p is given by the inequality $O_p = \{x : |x|_p \leq 1\}$.

Let us check that O_p is totally bounded. For $\epsilon > 0$ pick $k \in \mathbb{N}$ such that $p^{-k} < \epsilon$. The ring $\mathbb{Z}/p^k\mathbb{Z}$ is finite, let $z_1, \dots, z_N \in \mathbb{Z} \setminus \{0\}$ (where $N = p^k$) denote representatives of the cosets in $\mathbb{Z}/p^k\mathbb{Z}$. We claim that the set of fractions

$$w_{ij} = \frac{z_i}{z_j}, \quad 1 \leq i, j \leq N,$$

forms a p^{-k} -net in $O_p \cap \mathbb{Q}$. Indeed, for a rational number $\frac{m}{n} \in O_p \cap \mathbb{Q}$, find $s, t \in \{z_1, \dots, z_N\}$ such that

$$s \equiv m, t \equiv n, \pmod{p^k}.$$

Then

$$\frac{m}{n} - \frac{s}{t} \in p^k O_p$$

and, hence,

$$\left| \frac{m}{n} - \frac{s}{t} \right|_p \leq p^{-k}.$$

Since $O_p \cap \mathbb{Q}$ is dense in O_p , it follows that

$$O_p \subset \bigcup_{i,j=1}^N \overline{B}(w_{ij}, \epsilon). \quad \square$$

For the next exercise, recall *Brouwer's Theorem* (see e.g. [HY88, Corollary 2-98]) that every compact metrizable totally disconnected and perfect topological space is homeomorphic to the Cantor set.

EXERCISE 2.56. Show that O_p is homeomorphic to the Cantor set. Hint: Verify that O_p is totally disconnected and perfect.

2.8. Norms on field extensions. Adeles

A proof of the following theorem can be found e.g., in [Lan02, Chapter XII.2, Proposition 2.5].

THEOREM 2.57. *Suppose that $(\mathbb{E}, |\cdot|)$ is a normed field and $\mathbb{E} \subset \mathbb{F}$ is a finite extension. Then the norm $|\cdot|$ extends to a norm $|\cdot|$ on \mathbb{F} and this extension is unique. If $(\mathbb{E}, |\cdot|)$ is a local field, then so is $(\mathbb{F}, |\cdot|)$.*

We note that the statement about local fields follows from the fact that if V is a finite-dimensional normed vector space over a local field, then V is locally compact.

Norms on number fields are used to define *rings of adeles* of these fields. We refer the reader to [Lan64, Chapter 6] for the detailed treatment of adeles.

We let $\text{Nor}(\mathbb{Q})$ denote the set of norms on \mathbb{Q} , $|\cdot| : \mathbb{Q} \rightarrow \mathbb{R}_+$, see Section 2.7. If \mathbb{F} is an algebraic number field (a finite algebraic extension of \mathbb{Q}), then we let $\text{Nor}(\mathbb{F})$ denote the set of norms on \mathbb{F} extending the ones on \mathbb{Q} . We will use the notation ν and $|\cdot|_\nu$ for the elements of $\text{Nor}(\mathbb{F})$ and O_ν for the corresponding rings of integers; we let ν_p denote the p -adic norm and its unique extension to \mathbb{F} . Note that for each $x \in \mathbb{Q}$, $x \in O_p$ (the ring of p -adic integers) for all but finitely many p 's, since x has only finitely many primes in its denominator.

For each ν we let \mathbb{F}_ν denote the completion of \mathbb{F} with respect to ν and set $N_\nu = [\mathbb{F}_\nu : \mathbb{Q}_\nu]$.

LEMMA 2.58 (Product formula). *For each $x \in \mathbb{F} \setminus \{0\}$ we have*

$$\prod_{\nu \in \text{Nor}(\mathbb{F})} (\nu(x))^{N_\nu} = 1.$$

PROOF. We will prove this in the case $\mathbb{F} = \mathbb{Q}$; the reader can find the proof for general number fields in [Lan64, Chapter 6]. If $x = p$ is prime, then $|p| = p$ for the archimedean norm, $\nu(p) = 1$ if $\nu \neq \nu_p$ is a non-archimedean norm and $\nu_p(p) = 1/p$. Thus, the product formula holds for prime numbers x . Since norms are multiplicative functions from \mathbb{Q}^\times to \mathbb{R}_+ , the product formula holds for arbitrary $x \neq 0$. \square

For a non-archimedean norm ν we let

$$O_\nu = \{x \in \mathbb{F}_\nu : |x|_\nu \leq 1\}$$

denote the ring of integers in \mathbb{F}_ν . Since ν is non-archimedean, O_ν is both closed and open in \mathbb{F}_ν .

DEFINITION 2.59. For a finitely generated algebraic number field \mathbb{F} , the ring of *adeles* is the *restricted product*

$$\mathbb{A}(\mathbb{F}) := \prod'_{\nu \in \text{Nor}(\mathbb{F})} \mathbb{F}_\nu,$$

i.e. the subset of the direct product

$$(2.8) \quad \prod_{\nu \in \text{Nor}(\mathbb{F})} \mathbb{F}_\nu$$

which consists of sequences whose projection to \mathbb{F}_ν belongs to O_ν for all but finitely many ν 's. The ring operations on $\mathbb{A}(\mathbb{F})$ are defined first on sequences in the infinite product which have only finitely many non-zero terms and then extends to the rest of $\mathbb{A}(\mathbb{F})$ by taking suitable limits.

Note that in the case $F = \mathbb{Q}$, the $\mathbb{A}(\mathbb{Q})$ is the restricted product

$$\mathbb{R} \times \prod'_{p \text{ is prime}} \mathbb{Q}_p.$$

Adelic topology. Open subsets in the adelic topology on $\mathbb{A}(\mathbb{F})$ are products of open sets of \mathbb{F}_ν for finitely many ν 's (including all archimedean ones) and of O_ν 's for the rest of ν 's. Then the ring operations are continuous with respect to this topology. Accordingly, we topologize the group $GL(n, \mathbb{A}(\mathbb{F}))$ using the product topology on $\mathbb{A}(\mathbb{F})^{n^2}$. With this topology, $GL(n, \mathbb{A}(\mathbb{F}))$ becomes a topological group. Tychonoff's theorem implies compactness of product sets of the form

$$\prod_{\nu \in \text{Nor}(\mathbb{F})} C_\nu,$$

where $C_\nu \subset \mathbb{F}_\nu$ is compact for each ν , which equals to O_ν for all but finitely many ν 's,

THEOREM 2.60 (See e.g. Chapter 6, Theorem 1 in [Lan64]). *The image $\iota(\mathbb{F})$ of the diagonal embedding $\mathbb{F} \hookrightarrow \mathbb{A}(\mathbb{F})$ is a discrete subset in $\mathbb{A}(\mathbb{F})$.*

PROOF. Since $\iota(\mathbb{F})$ is an additive subgroup of the topological group $\mathbb{A}(\mathbb{F})$, it suffices to verify that 0 is an isolated point of $\iota(\mathbb{F})$. Take the archimedean norms $\nu_1, \dots, \nu_m \in \text{Nor}(\mathbb{F})$ (there are only finitely many of them since the Galois group $\text{Gal}(\mathbb{F}/\mathbb{Q})$ is finite) and consider the open subset

$$U = \prod_{i=1}^m \{x \in \mathbb{F}_{\nu_i} : \nu_i(x) < 1/2\} \times \prod_{\mu \in \text{Nor}(\mathbb{F}) \setminus \{\nu_1, \dots, \nu_m\}} O_\mu$$

of $\mathbb{A}(\mathbb{F})$. Then for each $(x_\nu) \in U$,

$$\prod_{\nu \in \text{Nor}(\mathbb{F})} \nu(x_\nu) < 1/2 < 1.$$

Hence, by the product formula, the intersection of U with the image of \mathbb{F} in $\mathbb{A}(\mathbb{F})$ consists only of $\{0\}$. \square

In order to appreciate this theorem, note that $\mathbb{F} = \mathbb{Q}$ is dense in the completion of \mathbb{Q} with respect to every norm. We also observe that this theorem fails if we equip $\mathbb{A}(\mathbb{F})$ with the topology induced from the product topology on the product of all \mathbb{F}_ν 's.

COROLLARY 2.61. *The image of $GL(n, \mathbb{F})$ under the embedding $\iota : GL(n, \mathbb{F}) \rightarrow GL(n, \mathbb{A}(\mathbb{F}))$ (induced by the diagonal embedding $\mathbb{F} \rightarrow \mathbb{A}(\mathbb{F})$) is a discrete subgroup.*

Even though, the adelic topology on $\mathbb{A}(\mathbb{F})$ is strictly stronger than the product topology, we note, nevertheless, that for a finitely generated subgroup $L \leq \mathbb{F}$, the image of L under the diagonal embedding $\iota : \mathbb{F} \rightarrow \mathbb{A}(\mathbb{F})$ projects to O_ν for all but finitely many ν 's (since the generators of L have only finitely many denominators). Thus, the restriction of the adelic topology to $\iota(L)$ coincides with the restriction of the product topology. The same applies for finitely generated subgroups $\Gamma \leq GL(n, \mathbb{F})$, since such Γ is contained in $GL(n, \mathbb{F}')$ for a finitely generated subfield \mathbb{F}' of \mathbb{F} . We, thus, obtain:

COROLLARY 2.62. *Suppose that $\Gamma \leq GL(n, \mathbb{F})$ is a finitely generated subgroup which project to a relatively compact subgroup of $GL(n, \mathbb{F}_\nu)$ for every norm ν . Then Γ is finite.*

PROOF. Since Γ is finitely generated, the restriction of the product topology on

$$\prod_{\nu \in \text{Nor}(\mathbb{F})} GL(n, \mathbb{F}_\nu)$$

to $\iota(\Gamma)$ coincides with the adelic topology, since Γ projects to $GL(n, O_\nu)$ for all but finitely many ν 's. In the adelic topology, $\iota(\Gamma)$ is discrete, while, in the product topology, it is a closed subset of a set C , which is the product of compact subsets of the groups $GL(n, \mathbb{F}_\nu)$. Hence, by Tychonoff's Theorem, C is compact. Thus, $\iota(\Gamma)$ is a discrete compact topological space, which implies that $\iota(\Gamma)$ is finite. Since ι is injective, it follows that Γ is finite as well. \square

COROLLARY 2.63. *Suppose that $\alpha \in \bar{\mathbb{Q}}$ is an algebraic integer, i.e. a root of a monic polynomial $p(x)$ with integer coefficients. Then either α is a root of unity or $p(x)$ has a root β such that $|\beta| > 1$. In other words, there exists an element $\sigma \in \text{Gal}(\bar{\mathbb{Q}}/\mathbb{Q})$ which sends α to an algebraic number β with non-unit absolute value.*

PROOF. Let $\mathbb{F} = \mathbb{Q}(\alpha)$ and consider the cyclic subgroup $\Gamma < \mathbb{Q}^\times$ generated by α . Since $\alpha \in O_{\mathbb{F}}$, we conclude that α belongs to the ring of integers O_α for each non-archimedean norm ν of \mathbb{F} . Thus, Γ projects to the compact subgroup O_ν of $\mathbb{A}(\mathbb{F})$. For each archimedean norm ν , we have $|\alpha|_\nu = |\sigma(\alpha)|$, where $\sigma \in \text{Gal}(\mathbb{F}/\mathbb{Q})$. However, $\sigma(\alpha)$ is another root of $p(x)$. Therefore, either there exists a root β of $p(x)$ such that $|\beta| \neq 1$, or Γ projects to a compact subgroup of \mathbb{F}_ν for each ν , both archimedean and non-archimedean. In the latter case, by Corollary 2.62, the group Γ is finite. Hence, α is a root of unity. \square

Our next goal is to extend this corollary to the case of general finitely generated fields, including transcendental extensions of \mathbb{Q} , as well as fields of positive characteristic. The following theorem is Lemma 4.1 in [Tit72]:

THEOREM 2.64. *Let \mathbb{E} be a finitely generated field and suppose that $\alpha \in \mathbb{E}$ is not a root of unity. Then there exists an extension $(\mathbb{F}, |\cdot|)$ of \mathbb{E} , which is a local field with the norm $|\cdot|$, such that $|\alpha| \neq 1$.*

PROOF. Let $\mathbb{P} \subset \mathbb{E}$ denote the prime subfield of \mathbb{E} . Since \mathbb{E} is finitely generated over \mathbb{P} , there is a finite transcendence basis $T = \{t_1, \dots, t_n\}$ for \mathbb{E} over \mathbb{P} and \mathbb{E} is a finite extension of $\mathbb{P}(T) = \mathbb{P}(t_1, \dots, t_n)$ (cf. Chapter VI.1 of [Hun80]). Here $\mathbb{P}(t_1, \dots, t_n)$ is isomorphic to the field of rational functions with coefficients in \mathbb{P} and variables in T . We also let T' denote a (finite) transcendence basis of \mathbb{E} over $\mathbb{P}(\alpha)$.

There are two main cases to consider.

Case 1: \mathbb{E} has characteristic $p > 0$, equivalently, $\mathbb{P} \cong \mathbb{Z}_p$ for some p . If α were to be algebraic over \mathbb{P} , then $\mathbb{P}(\alpha)$ would be finite and, hence, $\alpha^i = \alpha^k$ for some $i \neq k$, implying that α is a root of unity. This is a contradiction. Therefore, α is transcendental over \mathbb{P} and, hence, we can assume that α is an element of T . Define the ring $A = \mathbb{P}[T]$ and let $I \subset A$ denote the ideal generated by T . As we explained in the previous section, there is a (unique) valuation v on $\mathbb{P}(T)$ (with the norm $|\cdot|$) such that

$$v(a) = k \iff a \in I^k \setminus I^{k+1}.$$

By the construction, $v(\alpha) = 1$ and, hence, $|\alpha| \neq 1$. The completion of $\mathbb{P}(T)$ with respect to the norm $|\cdot|$ is a local field, since \mathbb{P} is finite. We then extend the norm to a norm on \mathbb{E} ; the completion with respect to this norm is again a local field.

Case 2: \mathbb{E} has zero characteristic, equivalently, $\mathbb{P} \cong \mathbb{Q}$. Suppose, first, that α is a transcendental number. Then there exists an embedding

$$\mathbb{Q}(\alpha) \rightarrow \mathbb{C}$$

which sends α to a transcendental number whose absolute value is > 1 , e.g., $\alpha \mapsto \pi$. This embedding extends to an embedding $\mathbb{E} \rightarrow \mathbb{C}$, thereby finishing the proof.

Suppose, therefore, that α is an algebraic number, $\alpha \in \mathbb{Q}$. Let $p(x)$ be the minimal monic polynomial of α .

Subcase 2a. Assume first that p has integer coefficients. Then, since α is not a root of unity, by Corollary 2.63, one of the roots β of p has absolute value > 1 . Consider the Galois automorphism $\phi : \mathbb{Q}(\alpha) \rightarrow \mathbb{Q}(\alpha)$ sending α to β . We then extend ϕ to an embedding

$$\psi : \mathbb{E}(T' \cup \{\alpha\}) \rightarrow \mathbb{C},$$

by sending the elements of T to complex numbers which are algebraically independent over $\mathbb{Q}(\alpha)$. Lastly, since $\mathbb{E}(T' \cup \{\alpha\}) \subset \mathbb{E}$ is an algebraic extension and \mathbb{C} is algebraically closed, the embedding ψ extends to the required embedding $\mathbb{E} \rightarrow \mathbb{C}$.

Subcase 2b. Lastly, we consider the case when $p \in \mathbb{Q}[x]$ has a non-integer coefficient. We consider the infinite cyclic subgroup generated by α in $\mathbb{Q}(\alpha)^\times$ and the embedding

$$\langle \alpha \rangle \rightarrow \mathbb{Q}(\alpha) \rightarrow \mathbb{A}_\alpha,$$

where \mathbb{A}_α is the ring of adeles \mathbb{A}_α of the field $\mathbb{Q}(\alpha)$; here $\mathbb{Q}(\alpha) \rightarrow \mathbb{A}_\alpha$ is the diagonal embedding. Since the subgroup $\mathbb{Q}(\alpha) < \mathbb{A}_\alpha$ is discrete and $\langle \alpha \rangle$ is an infinite subgroup, Tychonoff compactness theorem implies that the projection of $\langle \alpha \rangle$ to at least one of the factors of \mathbb{A}_α is unbounded. If this factor were archimedean, we would obtain a Galois embedding $\mathbb{Q}(\alpha) \rightarrow \mathbb{C}$ sending α to $\beta \in \mathbb{C}$ whose absolute value is different from 1. This situation is already handled in the Subcase 2a. Suppose, therefore, that there is a prime number p such that $\langle \alpha \rangle$ is an unbounded subgroup of the p -adic completion of $\mathbb{Q}(\alpha)$, which means that $|\alpha|_p \neq 1$, where $|\cdot|_p$ is the extension of the p -adic norm to $\mathbb{Q}(\alpha)$. Next, extending the norm $|\cdot|_p$ from $\mathbb{Q}(\alpha)$ to \mathbb{E} and then taking the completion, we obtain an embedding to \mathbb{E} to a local field, α has non-unit norm. \square

2.9. Metrics on affine and projective spaces

In this section we will use complete normed fields to define metrics on affine and projective spaces. Consider the vector space $V = \mathbb{F}^n$ over a complete normed field \mathbb{F} , with the standard basis e_1, \dots, e_n . We equip V with the usual Euclidean/hermitian norm in the case \mathbb{F} is archimedean and with the max-norm

$$|(x_1, \dots, x_n)| = \max_i |x_i|$$

if \mathbb{F} is non-archimedean. We let $\langle \cdot, \cdot \rangle$ denote the standard inner/hermitian product on V in the archimedean case.

EXERCISE 2.65. Suppose that \mathbb{F} is non-archimedean. Show that the metric $|v - w|$ on V satisfies the ultrametric triangle inequality.

If \mathbb{F} is non-archimedean, define the group $K = GL(n, O)$, consisting of matrices A such that $A, A^{-1} \in Mat_n(O)$.

EXERCISE 2.66. If \mathbb{F} is a non-archimedean local field, show that the group K is compact with respect to the subspace topology induced from $Mat_n(\mathbb{F}) = \mathbb{F}^{n^2}$.

LEMMA 2.67. *The group K acts isometrically on V .*

PROOF. It suffices to show that elements $g \in K$ do not increase the norm on V . Let a_{ij} denote the matrix coefficients of g . Then, for a vector $v = \sum_i v_i e_i \in V$, the vector $w = g(v)$ has coordinates

$$w_j = \sum_i a_{ji} v_i.$$

Since $|a_{ij}| \leq 1$, the ultrametric inequality implies

$$|w| = \max_j |w_j|, \quad |w_j| \leq \max_i |a_{ji} v_i| \leq |v|.$$

Thus, $|g(v)| \leq |v|$. \square

If \mathbb{F} is archimedean, we let $K < GL(V)$ denote the orthogonal/hermitian subgroup preserving the inner/hermitian product on V . The following is a standard fact from the elementary linear algebra, which follows from the spectral theorem, see e.g. [Str06, §6.3]:

THEOREM 2.68 (Singular Value Decomposition Theorem). *If \mathbb{F} is archimedean, then every matrix $M \in \text{End}(V)$ admits a singular value decomposition*

$$M = UDV,$$

where $U, V \in K$ and D is a diagonal matrix with nonnegative entries arranged in the descending order. The diagonal entries of D are called the singular values of M .

We will also need a (slightly less well-known) analogue of the singular value decomposition in the case of non-archimedean normed fields, see e.g. [DF04, §12.2, Theorem 21]:

THEOREM 2.69 (Smith Normal Form Theorem). *Let \mathbb{F} be a field with discrete norm, and uniformizer π and ring of integers O . Then every matrix $M \in \text{Mat}_n(\mathbb{F})$ admits a Smith Normal Form decomposition*

$$M = LDU,$$

where D is diagonal with diagonal entries (d_1, \dots, d_n) , $d_i = \pi^{k_i}$, $i = 1, \dots, n$,

$$k_1 \geq k_2 \geq \dots \geq k_n,$$

and $L, U \in K = GL(n, O)$. The diagonal entries $d_i \in \mathbb{F}$ are called the invariant factors of M .

PROOF. First, note that permutation matrices belong to K ; the group K also contains upper and lower triangular matrices with coefficients in O , whose diagonal entries are units in \mathbb{F} . We then apply Gauss Elimination Algorithm to the matrix M . Note that the row operation of adding the z -multiple of the i -th row to the j -th row amounts to multiplication on the left with the lower-triangular elementary matrix $E_{ij}(z)$ with the ij -entry equal z . If $z \in O$, then $E_{ij} \in K$. Similarly, column operations amount to multiplication on the right by an upper-triangular elementary matrix. Observe also that dividing a row (column) by a unit in \mathbb{F} amounts to multiplying a matrix on left (right) by an appropriate diagonal matrix with unit entries on the diagonal.

We now describe row operations for the Gauss Elimination in detail (column operations will be similar). Consider (non-zero) i -th column of a matrix $A \in \text{End}(\mathbb{F}^n)$. We first multiply M on left and right by permutation matrices so that a_{ii} has the largest norm in the i -th column. By dividing rows on A by units in \mathbb{F} , we achieve that every entry in the i -th column is a power of π . Now, eliminating non-zero entries in the i -th column will require only row operations involving $\pi^{s_{ij}}$ -multiples of the i -th row, where $s_{ij} \geq 0$, i.e. $\pi^{s_{ij}} \in O$. Applying this form of Gauss Algorithm to M , we convert M to a diagonal matrix A , whose diagonal entries are powers of π and

$$A = L'MU', \quad L', M' \in GL(n, O).$$

Multiplying A on left and right by permutation matrices, we rearrange the diagonal entries to have weakly decreasing exponents. \square

Note that both singular value decomposition and Smith normal form decomposition both have the form:

$$M = UDV, \quad U, V \in K,$$

and D is diagonal. Such decomposition of the $Mat_n(\mathbb{F})$ is called the *Cartan decomposition*. To simplify the terminology, we will refer to the diagonal entries of D as *singular values* of M in both archimedean and non-archimedean cases.

EXERCISE 2.70. Deduce the Cartan decomposition in $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$, from the statement that given any Euclidean/hermitian bilinear form q on $V = \mathbb{F}^n$, there exists a basis orthogonal with respect to q and orthonormal with respect to the standard inner product

$$x_1\bar{y}_1 + \dots + x_n\bar{y}_n.$$

We now turn our discussion to projective spaces. The \mathbb{F} -projective space $P = \mathbb{F}P^{n-1}$ is the quotient of $\mathbb{F}^n \setminus \{0\}$ by the action of \mathbb{F}^\times via scalar multiplication.

NOTATION 2.71. Given a non-zero vector $v \in V$ let $[v]$ denote the projection of v to the projective space $P(V)$; similarly, for a subset $W \subset V$ we let $[W]$ denote the image of $W \setminus \{0\}$ under the canonical projection $V \rightarrow P(V)$. Given an invertible linear map $g : V \rightarrow V$, we will retain the notation g for the induced projective map $P(V) \rightarrow P(V)$.

Suppose now that \mathbb{F} is a normed field. Our next goal is to define the *chordal metric* on $P(V) = \mathbb{F}P^{n-1}$. In the case of an archimedean field \mathbb{F} , we define the Euclidean or hermitian norm on $V \wedge V$ by declaring basis vectors

$$e_i \wedge e_j, 1 \leq i < j \leq n$$

to be orthonormal. Then

$$|v \wedge w|^2 = |v|^2|w|^2 - \langle v, w \rangle \langle w, v \rangle = (|v| \cdot |w| \cdot |\sin(\varphi)|)^2$$

where $\varphi = \angle(v, w)$. In other words, $|v \wedge w|$ is the area of the parallelogram spanned by the vectors v and w .

In the case when \mathbb{F} is non-archimedean, we equip $V \wedge V$ with the max-norm so that

$$|v \wedge w| = \max_{i,j} |x_i y_j - x_j y_i|$$

where $v = (x_1, \dots, x_n)$, $w = (y_1, \dots, y_n)$.

DEFINITION 2.72. The *chordal metric* on $P(V)$ is defined by

$$d([v], [w]) = \frac{|v \wedge w|}{|v| \cdot |w|}.$$

In the non-archimedean case this definition is due to A. Néron [Nér64].

EXERCISE 2.73. 1. If \mathbb{F} is non-archimedean, show that the group $GL(n, O)$ preserves the chordal metric.

2. If $\mathbb{F} = \mathbb{R}$, show that the orthogonal group preserves the chordal metric.

3. If $\mathbb{F} = \mathbb{C}$, show that the unitary group preserves the chordal metric.

4. Show that each $g \in GL(n, \mathbb{K})$ is a Lipschitz homeomorphism with respect to the chordal metric.

It is clear that $d([v], [w]) = d([w], [v])$ and $d([v], [w]) = 0$ if and only if $[v] = [w]$. What is not so obvious is why d satisfies the triangle inequality. Note, however, that in the case of a non-archimedean field \mathbb{F} ,

$$d([v], [w]) \leq 1$$

for all $[v], [w] \in P = P(V)$. Indeed, pick unit vectors v, w representing $[v], [w]$; in particular, v_i, w_j belong to O for all i, j . Then the denominator in the definition of $d([v], [w])$ equals 1, while the numerator is ≤ 1 , since O is the ring of integers.

PROPOSITION 2.74. *If \mathbb{F} is non-archimedean, then d satisfies the triangle inequality.*

PROOF. We will verify the triangle inequality by giving an alternative description of the function d . We define *affine patches* on P to be the affine hyperplanes

$$A_j = \{x \in V : x_j = 1\} \subset V$$

together with the (injective) projections $A_j \rightarrow P$. Every affine patch is, of course, just a translate of \mathbb{F}^{n-1} , so that e_j is the translate of the origin. We then equip A_j with the restriction of the metric $|v - w|$ from V . Let $B_j \subset A_j$ denote the closed unit ball centered at e_j . In other words,

$$B_j = A_j \cap O^{n+1}.$$

We now set $d_j(x, y) = |x - y|$ if $x, y \in B_j$ and $d_j(x, y) = 1$ otherwise. It follows immediately from the ultrametric triangle inequality that d_j is a metric. Define for $[x], [y] \in P$ the function $\text{dist}([x], [y])$ by:

1. If there exists j so that $x, y \in B_j$ project to $[x], [y]$, then $\text{dist}([x], [y]) := d_j(x, y)$.
2. Otherwise, set $\text{dist}([x], [y]) = 1$.

If we knew that dist is well-defined (*a priori*, different indices j give different values of dist), it would be clear that dist satisfies the ultrametric triangle inequality. Proposition will, now, follow from:

LEMMA 2.75. $d([x], [y]) = \text{dist}([x], [y])$ for all points in P .

PROOF. The proof will break in two cases:

1. There exists k such that $[x], [y]$ lift to $x, y \in B_k$. To simplify the notation, we will assume that $k = n$. Since $x, y \in B_n$, $|x_i| \leq 1, |y_i| \leq 1$ for all i , and $x_n = y_n = 1$. In particular, $|x| = |y| = 1$. Hence, for every i ,

$$|x_i - y_i| = |x_i y_n - x_j y_n| \leq \max_j |x_i y_j - x_j y_i| \leq d([x], [y]),$$

which implies that

$$\text{dist}([x], [y]) \leq d([x], [y]).$$

We will now prove the opposite inequality:

$$\forall i, j \quad |x_i y_j - x_j y_i| \leq a := |x - y|.$$

There exist $z_i, z_j \in \mathbb{F}$ so that

$$y_i = x_i(1 + z_i), \quad y_j = x_j(1 + z_j),$$

where, if $x_i \neq 0, x_j \neq 0$,

$$z_i = \frac{y_i - x_i}{x_i}, \quad z_j = \frac{y_j - x_j}{x_j}.$$

We will consider the case $x_i x_j \neq 0$, leaving the exceptional cases to the reader.

Then

$$|z_i| \leq \frac{a}{|x_i|}, \quad |z_j| \leq \frac{a}{|x_j|}.$$

Computing $x_i y_j - x_j y_i$ using the new variables z_i, z_j , we obtain:

$$\begin{aligned} |x_i y_j - x_j y_i| &= |x_i x_j (1 + z_j) - x_i x_j (1 + z_i)| = |x_j x_j (z_j - z_i)| \leq \\ &|x_i x_j| \max(|z_i|, |z_j|) \leq |x_i x_j| \max\left(\frac{a}{|x_i|}, \frac{a}{|x_j|}\right) \leq a \max(|x_i|, |x_j|) \leq a, \end{aligned}$$

since $x_i, x_j \in O$.

2. Suppose that (1) does not happen. Since $d([x], [y]) \leq 1$ and $\text{dist}([x], [y]) = 1$ (in the second case), we just have to prove that

$$d([x], [y]) \geq 1.$$

Consider representatives x, y of points $[x], [y]$ and let i, j be the indices such that

$$|x_i| = |x|, \quad |y_j| = |y|.$$

Clearly, i, j are independent of the choices of the vectors x, y representing $[x], [y]$. Therefore, we choose x so that $x_i = 1$, which implies that $x_k \in O$ for all k . If $y_i = 0$ then

$$|x_i y_j - x_j y_i| = |y_j|$$

and

$$d([x], [y]) \geq \frac{\max_j |1 \cdot y_j|}{|y_j|} = 1.$$

Thus, we assume that $y_i \neq 0$. This allows us to choose $y \in A_i$ as well. Since (1) does not occur, $y \notin O^n$, which implies that $|y_j| > 1$. Now,

$$d([x], [y]) \geq \frac{|x_i y_j - x_j y_i|}{|x_i| \cdot |y_j|} = \frac{|y_j - x_j|}{|y_j|}.$$

Since $x_j \in O$ and $y_j \notin O$, the ultrametric inequality implies that $|y_j - x_j| = |y_j|$. Therefore,

$$\frac{|y_j - x_j|}{|y_j|} = \frac{|y_j|}{|y_j|} = 1$$

and $d([x], [y]) \geq 1$. This concludes the proof of lemma and proposition. \square

COROLLARY 2.76. *If \mathbb{K} is non-archimedean, then the metric d on P is locally isometric to the metric $|x - y|$ on the affine space \mathbb{F}^{n-1} .*

We now consider real and complex projective spaces. Choosing unit vectors u, v as representatives of points $[u], [v] \in P$, we get:

$$d([u], [v]) = \sin(\angle(u, v)),$$

where we normalize the angle to be in the interval $[0, \pi]$. Consider now three points $[u], [v], [w] \in P$; our goal is to verify the triangle inequality

$$d([u], [w]) \leq d([u], [v]) + d([v], [w]).$$

We choose unit vectors u, v, w representing these points so that

$$0 \leq \alpha = \angle(u, v) \leq \frac{\pi}{2}, \quad 0 \leq \beta = \angle(v, w) \leq \frac{\pi}{2}.$$

Then

$$\gamma = \angle(u, w) \leq \alpha + \beta$$

and the triangle inequality for the metric d is equivalent to the inequality

$$\sin(\gamma) \leq \sin(\alpha) + \sin(\beta).$$

We leave verification of the last inequality as an exercise to the reader. Thus, we obtain

THEOREM 2.77. *The chordal metric is a metric on P in both archimedean and non-archimedean cases.*

EXERCISE 2.78. Suppose that \mathbb{F} is a normed field (either non-archimedean or archimedean).

1. Verify that metric d determines the topology on P which is the quotient topology induced from $V \setminus \{0\}$.

2. Assuming that \mathbb{F} is local, verify that P is compact.

3. If the norm on \mathbb{F} is complete, show that the metric space (P, d) is complete.

4. If H is a hyperplane in $V = \mathbb{F}^n$, given as $\text{Ker } f$, where $f : V \rightarrow \mathbb{F}$ is a linear function, show that

$$\text{dist}([v], [H]) = \frac{|f(v)|}{\|v\| \|f\|}.$$

2.10. Quasiprojective transformations. Proximal transformations

In what follows, V is a finite-dimensional vector space of dimension n , over a local field \mathbb{F} . Each automorphism $g \in GL(V)$ of the vector space V projects to a projective transformation $g \in PGL(V)$, $g : P(V) \rightarrow P(V)$. Given g , we will always extend the norm from \mathbb{F} to the splitting field \mathbb{E} of the characteristic polynomial of g , in order to define norms of eigenvalues of g .

On the other hand, endomorphisms of V (i.e. linear maps $V \rightarrow V$) do not project, in general, to self-maps $P(V) \rightarrow P(V)$. Nevertheless, if $g \in \text{End}(V)$ is a linear transformation of rank $r > 0$ with kernel $\text{Ker}(g)$ and image $\text{Im}(g)$, then g determines a *quasiprojective transformation* g of $P(V)$, whose domain dom_g is the complement of $P(\text{Ker}(g))$ and whose image is $\text{Im}_g := P(\text{Im}(g))$. The number $r = \text{rank}(g)$ is called *the rank* of this quasiprojective transformation. The subspace $\text{Ker}_g := P(\text{Ker}(g))$ is the *kernel*, or the *indeterminacy set* of g . We let $\text{End}(P(V))$ denote the semigroup of quasiprojective transformations of $P(V)$. Rank 1 quasiprojective transformations are *quasiconstant maps*: Each quasiconstant map is undefined on a hyperplane in $P(V)$ and its image is a single point.

EXERCISE 2.79. For $h \in GL(V)$ and $g \in \text{End}(V)$ we have:

$$\text{Im}_{hg} = h(\text{Im}_g), \quad \text{Ker}_{hg} = \text{Ker}_g,$$

$$\text{Im}_{gh} = \text{Im}_g, \quad \text{Ker}_{gh} = h^{-1}(\text{Ker}_g).$$

The rank of a quasiprojective transformation can be detected locally:

EXERCISE 2.80. Suppose that $g \in \text{End}(P(V))$ and $U \subset \text{dom}_g \subset P(V)$ is a non-empty open subset. Then

$$\text{rank}(g) = \dim(g(U)) + 1.$$

We will topologize $\text{End}(P(V))$ using the operator norm topology on $\text{End}(V)$.

EXERCISE 2.81. 1. The function rank is lower semicontinuous on $End(P(V))$.

2. A sequence g_i converges to g in $End(P(V))$ if and only if it converges to g uniformly on compacts in dom_g . (In particular, each compact $C \subset dom_g$ is contained in dom_{g_i} for all but finitely many i 's.)

3. Suppose that $g \in End(V)$ is such that the dominant eigenvalue λ_1 of g satisfies $|\lambda_1| < 1$. Show that the sequence $g^k \in End(V)$ converges to $0 \in End(V)$. Hint: Show that for any norm on V and the corresponding norm on $End(V)$ we have

$$\|g^k\| \leq |\lambda_1|^k p(k),$$

where $p(k)$ is a polynomial in k of degree $\leq n - 1$.

THEOREM 2.82 (A convergence property). *The semigroup $End(P(V))$ is compact: Each sequence $g_i \in End(P(V))$ subconverges to a quasiprojective transformation.*

PROOF. We fix a basis in V . In the case when the field \mathbb{F} is archimedean we equip V with the inner product with respect to which the basis is orthonormal. In the case when \mathbb{F} is non-archimedean we equip V with the maximum-norm:

$$\|v\|_{max} = \max_{i=1, \dots, n} |v_i|, v = (v_1, \dots, v_n).$$

In either case, we equip V with the operator norm defined via $\|\cdot\|$ and let $K < GL(V)$ denote a maximal compact subgroup preserving the norm on V . We will use the Cartan decomposition $End(V) = K \cdot Diag(V) \cdot K$: Each $g \in End(V)$ has the form $g = k_g a_g k'_g$, where k_g, k'_g belong to the subgroup $K < GL(V)$ and a_g is a diagonal transformation whose diagonal entries are the singular values of g , see Section 2.9. Assuming $g \neq 0$, $a_g \neq 0$ as well and, by replacing g with its scalar multiple (which does not affect the corresponding quasiprojective endomorphism), we can assume that the dominant eigenvalue of a_g equals 1, i.e. $\|a_g\| = 1$. We apply this to elements of a sequence $g_i \in End(P(V))$ and obtain:

$$g_i = k_{g_i} a_{g_i} k'_{g_i}, \quad \|a_{g_i}\| = 1.$$

Since \mathbb{F} is a locally compact field, the sequences $k_{g_i}, a_{g_i}, k'_{g_i}$ subconverge in $End(V)$: The limits k, k' of convergent subsequences of k_{g_i}, k'_{g_i} belong to the group K , while a_{g_i} subconverges to an endomorphism a of the unit norm, in particular, this limit is different from 0. Thus, the sequence (g_i) subconverges to the non-zero endomorphism kak' . \square

LEMMA 2.83. *Suppose that (g_i) is a sequence in $GL(V)$ converging to a quasi-constant map \hat{g} . Then*

$$\lim_{i \rightarrow \infty} Lip(g_i) = 0$$

uniformly on compacts in $dom(\hat{g})$. In other words, for every compact

$$C \subset dom_{\hat{g}}$$

we have

$$\sup_{x, y \in C, x \neq y} \frac{d(g_i(x), g_i(y))}{d(x, y)} = 0$$

PROOF. In view of the Cartan decomposition, it suffices to analyze the case when each g_i is a diagonal matrix with the diagonal entries $\lambda_1 = 1, \lambda_2, \dots, \lambda_n$, satisfying

$$1 \geq |\lambda_{2,i}| \geq \dots \geq |\lambda_{n,i}|$$

such that

$$\lim_{i \rightarrow \infty} \lambda_{2,i} = 0.$$

In particular, the maps g_i preserve the affine hyperplane

$$A_1 = \{(1, x_2, \dots, x_n) : x_2, \dots, x_n \in \mathbb{F}\}$$

in V . We will identify A_1 with

$$\text{dom}_{\hat{g}} \subset P(V)$$

and, accordingly, lift C to a compact subset (again denoted by C) in A_1 .

We first consider the case of a non-archimedean field \mathbb{F} . We will use the action of g_i on A_1 in order to analyze the contraction properties of g_i . Since the sequence g_i restricted to C converges uniformly to e_1 , for all sufficiently large i , $g_i(C)$ is contained in the unit ball centered at e_1 . In view of Lemma 2.75, it is clear that the maps g_i do not increase distances between points in A_1 (measured in the metric d_1 on A_1). Furthermore, for all $x, y \in B(e_1, 1) \subset A_1$, we have

$$(2.9) \quad d_1(g_i(x), g_i(y)) \leq \lambda_{2,i} d_1(x, y)$$

and, hence, the Lipschitz constant of g_i converges to zero.

Consider now the archimedean case. As in the non-archimedean case, we let d_1 denote the restriction to A_1 of the metric defined via the maximum-norm on V . We leave it to the reader to check the inequalities:

$$D^2 d_1(x, y) \frac{|x - y|}{|x| |y|} \leq d([x], [y]) \leq n \|x - y\|_{max} = n d_1(x, y)$$

for all points $x, y \in A_1$ satisfying $\max(|x|, |y|) \leq D$. This shows that the map $(A_1, d_1) \rightarrow (P(V), d)$ is uniformly bilipschitz on each compact in A_1 . On the other hand, the map

$$g_i : (A_1, d_1) \rightarrow (A_1, d_1)$$

satisfies the inequality (2.9) for all $x, y \in A_1$. Lemma follows. \square

REMARK 2.84. It is useful to note here that while singular values depend on the choice of a basis in V , the limit quasiprojective transformation of the sequence (g_i) is, of course, independent of the basis. The same applies to the notion of proximality below.

The most important, for us, example of convergence to a quasiprojective transformation comes from iterations of a single invertible transformation: $g_i = g^i$, $i \in \mathbb{N}$. For $g \in \text{End}(V)$ we say that an eigenvalue λ_1 of g is *dominant* if it has algebraic multiplicity one and

$$|\lambda_1| > |\lambda_k|$$

for all eigenvalues λ_k of g different from λ_1 .

DEFINITION 2.85. An endomorphism g is called *proximal* if it has a dominant eigenvalue; an automorphism $g \in GL(V)$ is *very proximal* if both g and g^{-1} are proximal elements of $GL(V)$.

For a proximal endomorphism g we let $\tilde{A}_g \subset V$ denote the (one-dimensional) eigenspace corresponding to the dominant eigenvalue λ_1 and let $\tilde{E}_g \subset V$ denote the sum of the rest of the generalized eigenspaces of g . We project \tilde{A}_g and \tilde{E}_g , respectively, to a point A_g and a hyperplane E_g in the projective space $P(V)$. We will refer to A_g as the *attractive point* and E_g the *exceptional hyperplane* for the

action of a proximal projective transformation g on $\mathbb{P}(V)$. (The reason for the terminology will become clear from the next lemma.)

It is clear that proximality depends only on the projectivization of g .

We now work out limits of sequences $g^i \in \text{End}(P(V))$, when g is proximal. We already know that the sequence (g^i) of projective transformations subconverges to a quasiprojective transformation, the issue is to compute the rank, range and the kernel of the limit.

LEMMA 2.86. *If $g \in \text{End}(V)$ is a proximal endomorphism of $P(V)$, then each convergent subsequence in the sequence (g^k) of projective transformations converges to a rank 1 (quasiconstant) quasiprojective transformation \hat{g} . The image $\text{Im}_{\hat{g}}$ of \hat{g} equals $A_g := P(\tilde{A}_g)$ and the kernel $\text{Ker}_{\hat{g}}$ of \hat{g} equals $E_g := P(\tilde{E}_g)$.*

PROOF. We normalize g so that $\lambda_1 = 1$; hence, all eigenvalues of g restricted to \tilde{E}_g have absolute value < 1 . Clearly, the restriction of g to \tilde{A}_g is the identity, while, by Part 3 of Exercise 2.81, the restriction of the sequence g^i to \tilde{E}_g converges to the zero linear map. Lemma follows. \square

COROLLARY 2.87. *Given a proximal endomorphism $g \in \text{End}(P(V))$, for every $\varepsilon > 0$ there exists $N = N(\varepsilon)$ such that for every $i \geq N$, the projective transformation $g^i \in \text{End}(P(V))$ maps the complement of the ε -neighborhood of the hyperplane $E_g \subset P(V)$ inside the ball*

$$B(A_g, \varepsilon)$$

of radius ε and center A_g .

We will be using quasiprojective transformations and proximal elements of $GL(V)$ in the proof of the Tits' Alternative, Section 15.4.

2.11. Kernels and distance functions

A *kernel* on a set X is a symmetric map $\psi : X \times X \rightarrow \mathbb{R}_+$ such that $\psi(x, x) = 0$. (Symmetry of ψ means that $\psi(x, y) = \psi(y, x)$ for all x, y in X .) Fix $p \in X$ and define the associated *Gromov kernel*

$$k_p(x, y) := \frac{1}{2} (\psi(x, p) + \psi(p, y) - \psi(x, y)),$$

cf. section 11.3 for the definition of the *Gromov product* in metric spaces. Clearly,

$$\forall x \in X, \quad k_p(x, x) = \psi(x, p).$$

DEFINITION 2.88. 1. A kernel ψ is *positive semidefinite* if for every natural number n , every subset $\{x_1, \dots, x_n\} \subset X$ and every vector $\lambda \in \mathbb{R}^n$,

$$(2.10) \quad \sum_{i=1}^n \sum_{j=1}^n \lambda_i \lambda_j \psi(x_i, x_j) \geq 0.$$

2. A kernel ψ is *conditionally negative semidefinite* if for every $n \in \mathbb{N}$, every subset $\{x_1, \dots, x_n\} \subset X$ and every vector $\lambda \in \mathbb{R}^n$ with $\sum_{i=1}^n \lambda_i = 0$, the following holds:

$$(2.11) \quad \sum_{i=1}^n \sum_{j=1}^n \lambda_i \lambda_j \psi(x_i, x_j) \leq 0.$$

This is not a particularly transparent definition. A better way to think about this definition is in terms of the vector space $V = V(X)$ of consisting of functions with finite support $X \rightarrow \mathbb{R}$. Then each kernel ψ on X defines a symmetric bilinear form on V (denoted Ψ):

$$\Psi(f, g) = \sum_{x, y \in X} \psi(x, y) f(x) g(y).$$

With this notation, the left hand side of (2.10) becomes simply $\Psi(f, f)$, where

$$\lambda_i := f(x_i), \quad \text{Supp}(f) \subset \{x_1, \dots, x_n\} \subset X.$$

Thus, a kernel is positive semidefinite if and only if Ψ is a positive semidefinite bilinear form. Similarly, ψ is conditionally negative semidefinite if and only if the restriction of $-\Psi$ to the subspace V_0 consisting of functions with zero average, is a positive semidefinite bilinear form.

NOTATION 2.89. We will use the lower case letters to denote kernels and the corresponding upper case letters to denote the associated bilinear forms on V .

Below is yet another interpretation of the conditionally negative semidefinite kernels. For a subset $\{x_1, \dots, x_n\} \subset X$ define the symmetric matrix M with the entries

$$m_{ij} = -\psi(x_i, x_j), \quad 1 \leq i, j \leq n.$$

For $\lambda = (\lambda_1, \dots, \lambda_n)$, the left hand-side of the inequality (2.11) equals

$$q(\lambda) = \lambda^T M \lambda,$$

a symmetric bilinear form on \mathbb{R}^n . **Then** the condition (2.11) means that q is positive semi-definite on the hyperplane

$$\sum_{i=1}^n \lambda_i = 0$$

in \mathbb{R}^n . Suppose, for a moment, that this form is actually positive-definite. Since $\psi(x_i, x_j) \geq 0$, it follows that the form q on \mathbb{R}^n has signature $(n-1, 1)$. The standard basis vectors e_1, \dots, e_n in \mathbb{R}^n are null-vectors for q ; the condition $m_{ij} \leq 0$ amounts to the requirement that these vectors belong to the same, say, positive, light cone.

The following theorem gives yet another interpretation of conditionally negative semidefinite kernels in terms of embeddings in Hilbert spaces. It was first proven by J. Schoenberg in [Sch38] in the case of finite sets, but the same proof works for infinite sets as well.

THEOREM 2.90. *A kernel ψ on X is conditionally negative semidefinite if and only if there exists a map $F : X \rightarrow \mathcal{H}$ to a Hilbert space such that*

$$\psi(x, y) = \|F(x) - F(y)\|^2.$$

Here $\|\cdot\|$ denotes the norm on \mathcal{H} . Furthermore, if G is a group acting on X preserving the kernel ψ then the map F is equivariant with respect to a homomorphism $G \rightarrow \text{Isom}(\mathcal{H})$.

PROOF. 1. Suppose that the map F exists. Then, for every $p = x_0 \in X$, the associated Gromov kernel $k_p(x, y)$ equals

$$k_p(x, y) = \langle F(x), F(y) \rangle,$$

and, hence, for every finite subset $\{x_0, x_1, \dots, x_n\} \subset X$, the corresponding matrix with the entries $k_p(x_i, x_j)$ is the Gramm matrix of the set

$$\{y_i := F(x_i) - F(x_0) : i = 1, \dots, n\} \subset \mathcal{H}.$$

Hence, this matrix is positive semidefinite. Accordingly, Gromov kernel determines a positive semidefinite bilinear form on the vector space $V = V(X)$.

We will verify that ψ is conditionally negative semidefinite by considering subsets X_0 in X of the form $\{x_0, x_1, \dots, x_n\}$. (Since the point x_0 was arbitrary, this will suffice.)

Let $f : X_0 \rightarrow \mathbb{R}$ be such that

$$(2.12) \quad \sum_{i=0}^n f(x_i) = 0.$$

Thus,

$$f(x_0) := - \sum_{i=1}^n f(x_i).$$

Set $y_i := F(x_i)$, $i = 0, \dots, n$. Since the kernel K is positive semidefinite, we have

$$(2.13) \quad \sum_{i,j=1}^n (|y_0 - y_i|^2 + |y_0 - y_j|^2 - |y_i - y_j|^2) f(x_i) f(x_j) =$$

$$2 \sum_{i,j=1}^n k_p(x_i, x_j) f(x_i) f(x_j) \geq 0.$$

The left hand side of this equation equals

$$2 \left(\sum_{i=1}^n f(x_i) \right) \cdot \left(\sum_{j=1}^n |y_0 - y_j|^2 f(x_j) \right) -$$

$$\sum_{i,j=1}^n |y_i - y_j|^2 f(x_i) f(x_j).$$

Since $f(x_0) := - \sum_{i=1}^n f(x_i)$, we can rewrite this expression as

$$-f(x_0)^2 |y_0 - y_0|^2 - 2 \left(\sum_{j=1}^n |y_0 - y_j|^2 f(x_0) f(x_j) \right) - \sum_{i,j=1}^n |y_i - y_j|^2 f(x_i) f(x_j) =$$

$$\sum_{i,j=0}^n |y_i - y_j|^2 f(x_i) f(x_j) = \sum_{i,j=0}^n \psi(x_i, x_j) f(x_i) f(x_j).$$

Taking into account the inequality (2.13), we conclude that

$$(2.14) \quad \sum_{i,j=0}^n \psi(x_i, x_j) f(x_i) f(x_j) \leq 0.$$

In other words, the kernel ψ on X is conditionally negative semidefinite.

2. Suppose that ψ is conditionally negative semidefinite. Fix $p \in X$ and define the Gromov kernel

$$k_p(x, y) := \frac{1}{2} (\psi(x, p) + \psi(p, y) - \psi(x, y)).$$

The key to the proof is:

LEMMA 2.91. $k_p(x, y)$ is a positive semidefinite kernel on X .

PROOF. Consider a subset $X_0 = \{x_1, \dots, x_n\} \subset X$ and a function $f : X_0 \rightarrow \mathbb{R}$.

a. We first consider the case when $p \notin X_0$. Then we set $x_0 := p$ and extend the function f to p by

$$f(x_0) := - \sum_{i=1}^n f(x_i).$$

The resulting function $f : \{x_0, \dots, x_n\} \rightarrow \mathbb{R}$ satisfies (2.12) and, hence,

$$\sum_{i,j=0}^n \psi(x_i, x_j) f(x_i) f(x_j) \leq 0.$$

The same argument as in the first part of the proof of Theorem 2.90 (run in the reverse) then shows that

$$\sum_{i,j=1}^n k_p(x_i, x_j) f(x_i) f(x_j) \geq 0.$$

Thus, k_p is positive semidefinite on functions whose support is disjoint from $\{p\}$.

b. Suppose that $p \in X_0$, $f(p) = c \neq 0$. We define a new function $g(x) := f(x) - c\delta_p$. Here δ_p is the characteristic function of the subset $\{p\} \subset X$. Then $p \notin \text{Supp}(g)$ and, hence, by the Case (a),

$$K_p(g, g) \geq 0.$$

On the other hand,

$$K_p(f, f) = F(g, g) + 2cK(g, \delta_p) + c^2K(\delta_p, \delta_p) = F(g, g),$$

since the other two terms vanish (as $k_p(x, p) = 0$ for every $x \in X$). Thus, K_p is positive semidefinite. \square

Now, consider the vector space $V = V(X)$ equipped with the positive semidefinite bilinear form $\langle f, g \rangle = K(f, g)$. Define the Hilbert space \mathcal{H} as the metric completion of

$$V / \{f \in V : \langle f, f \rangle = 0\}.$$

Then we have a natural map $F : X \rightarrow \mathcal{H}$ which sends $x \in X$ to the projection of the δ -function δ_x (the indicator function $\mathbf{1}_x$); we obtain:

$$\langle F(x), F(y) \rangle = k_p(x, y).$$

Let us verify now that

$$(2.15) \quad \langle F(x) - F(y), F(x) - F(y) \rangle = \psi(x, y).$$

The left hand side of this expression equals

$$\langle F(x), F(x) \rangle + \langle F(y), F(y) \rangle - 2k_p(x, y) = \psi(x, p) + \psi(y, p) - 2k_p(x, y).$$

Then the equality (2.15) follows from the definition of the Gromov kernel k . The part of the theorem dealing with G -invariant kernels is clear from the construction. \square

Below we list several elementary properties of positive semidefinite and conditionally negative semidefinite kernels.

LEMMA 2.92. *Each kernel of the form $\psi(x, y) = f(x)f(y)$ is positive semidefinite.*

PROOF. This follows from the computation:

$$\sum_{i,j=1}^n \lambda_i \lambda_j \psi(x_i, x_j) = \left(\sum_{i=1}^n \lambda_i f(x_i) \right) \left(\sum_{j=1}^n \lambda_j f(x_j) \right) = \left(\sum_{i=1}^n \lambda_i f(x_i) \right)^2 \geq 0. \quad \square$$

Before proving the next lemma we will need the notion of *Hadamard product* of two matrices: If $A = (a_{ij})$ and $B = (b_{ij})$ are $n \times m$ matrices, then their *Hadamard product*, denoted by $A \circ B$ is the matrix with the entries $(a_{ij}b_{ij})$. The main property of the Hadamard product that we will need is known as the *Schur Product Theorem*:

THEOREM 2.93 (I. Schur). *If A, B are positive semidefinite $n \times n$ matrices, then their Hadamard product $A \circ B$ is again positive semidefinite.*

PROOF. A proof of Schur's Product Theorem reduces to two calculations: For each (row) vector $v \in \mathbb{R}^n$

$$v^T (A \circ B) v = \text{tr}(A \text{diag}(v) B \text{diag}(v))$$

(where $\text{diag}(v)$ is the diagonal matrix with the diagonal entries equal to $v_i, i = 1, \dots, n$). Then for the matrix $M = B^{1/2} \text{diag}(v) A^{1/2}$ (note that square roots exist since A and B are positive semidefinite) we have:

$$\begin{aligned} v^T (A \circ B) v &= \text{tr}(A^{1/2} A^{1/2} \text{diag}(v) B^{1/2} B^{1/2} \text{diag}(v)) = \\ &= \text{tr}(A^{1/2} \text{diag}(v) B^{1/2} B^{1/2} \text{diag}(v) A^{1/2}) = \text{tr}(M^T M) \geq 0. \quad \square \end{aligned}$$

LEMMA 2.94. *Sums and products of positive semidefinite kernels are again positive semidefinite. The set of positive semidefinite kernels is closed in the space of all kernels with respect to the topology of pointwise convergence.*

PROOF. The only statement which is not immediate from the definitions is that product of positive semidefinite kernels $\theta(x, y) = \varphi(x, y)\psi(x, y)$ is again positive semidefinite. In order to prove so it suffices to consider the case $x, y \in X = \{x_1, \dots, x_n\}$. Let $A = (a_{ij}), a_{ij} = \varphi(x_i, x_j)$, and $B = (b_{ij}), b_{ij} = \psi(x_i, x_j)$ denote the Gram matrices of the kernels φ and ψ . Then the product kernel θ is given by the matrix

$$C = A \circ B.$$

Since A and B are positive semidefinite, so is C and, hence, θ . \square

COROLLARY 2.95. *If a kernel $\psi(x, y)$ is positive semidefinite, so is the kernel $\exp(\psi(x, y))$.*

PROOF. This follows from the previous lemma since

$$\exp(t) = \sum_{n=0}^{\infty} \frac{t^n}{n!}. \quad \square$$

THEOREM 2.96 (J. Schoenberg, [Sch38]). *If $\psi(x, y)$ is a conditionally negative semidefinite kernel then for each $s > 0$ the function $\varphi(x, y) = \exp(-s\psi(x, y))$ is a positive semidefinite kernel.*

PROOF. If X is empty, there is nothing to prove, therefore, fix $p \in X$ and consider the kernel

$$k(x, y) = \psi(x, p) + \psi(y, p) - \psi(x, y)$$

(twice the Gromov kernel). This kernel is positive semidefinite according to Lemma 2.91. We have:

$$\exp(-s\psi(x, y)) = \exp(s k(x, y)) \exp(-s\psi(x, p)) \exp(-s\psi(y, p)).$$

The first term on the right hand side is positive semidefinite according to Corollary 2.95. The product of the other two terms is positive semidefinite by Lemma 2.92. Therefore, Lemma 2.94 implies that $\varphi(x, y)$ is positive semidefinite. \square

Note that Schoenberg uses Theorem 2.96 to prove in [Sch38] the following neat result: For every conditionally negative semidefinite kernel $\psi : X \times X \rightarrow \mathbb{R}_+$ and every $0 < \alpha \leq 1$, the power ψ^α is also a conditionally negative semidefinite kernel. In other words, if a metric space (X, dist) embeds isometrically into a Hilbert space, so does every metric space

$$(X, \text{dist}^\alpha), \quad 0 < \alpha \leq 1.$$

The main source of examples of conditionally negative semidefinite kernels comes from norms in L^p -spaces (the case $p = 2$ is covered by Theorem 2.90).

Before proceeding with the discussion on kernels, we wish to clarify our choices for L^p -spaces, with $p \in (0, 1)$.

REMARK 2.97. For a space $L^p(X, \mu)$ with $p \in (0, 1)$, $\|f\|_p = (\int |f|^p d\mu)^{\frac{1}{p}}$ no longer satisfies the usual triangular inequality, it only satisfies a similar inequality with a multiplicative factor added to the second term. On the other hand, $\|f\|_p^p$ is no longer a norm, but it does satisfy the triangular inequality, hence it defines a metric [KPR84].

Throughout the book, we consider L^p -spaces endowed with this metric, for $p \in (0, 1)$.

PROPOSITION 2.98 ([WW75], Theorem 4.10). *Let (Z, μ) be a measure space. Let $0 < p \leq 2$, and let $E = L^p(Z, \mu)$ be endowed with the norm $\|\cdot\|_p$. Then $\psi : E \times E \rightarrow \mathbb{R}$, $\psi(x, y) = \|x - y\|_p^p$ is a conditionally negative semidefinite kernel.*

On the other hand, according to Schoenberg's theorem 2.90, every conditionally negative semidefinite kernels comes from maps to Hilbert spaces. A corollary of this is a theorem first proven by Banach and Mazur:

COROLLARY 2.99. *For each $p \in (0, 2]$ there exists a linear isometric embedding of metric spaces*

$$(L^p(Z, \mu), d_p) \rightarrow \mathcal{H}$$

where \mathcal{H} is a Hilbert space and

$$d_p(f, g) = \left(\int_Z |f - g|^p d\mu \right)^{1/2}.$$

More generally:

THEOREM 2.100 (Theorems 1 and 7 in [BDCK66]). *Let $1 \leq p \leq q \leq 2$.*

- (1) *The normed space $(L^q(X, \mu), \|\cdot\|_q)$ can be embedded linearly and isometrically into*

$$(L^p(X', \mu'), \|\cdot\|_p)$$

for some measure space (X', μ') .

- (2) *If $L^p(X, \mu)$ has infinite dimension, then $(L^p(X, \mu), \|\cdot\|_p^\alpha)$ can be embedded isometrically into $(L^q(X', \mu'), \|\cdot\|_q)$ for some measured space (X', μ') , if and only if $0 < \alpha \leq \frac{p}{q}$.*

Differential geometry

In this book we will use some elementary Differential and Riemannian geometry, basics of which are reviewed in this chapter. All the manifolds that we consider are second countable.

3.1. Smooth manifolds

We expect the reader to know basics of differential topology, that can be found, for instance, in [GP10], [Hir76], [War83]. Below is only a brief review.

Unless stated otherwise, all maps between smooth manifolds, vector fields and differential forms are assumed to be infinitely differentiable.

We will use the notation $\Lambda^k(M)$ for the space of differential k -forms on M . Every vector field X on M defines the *contraction operator*

$$i_X : \Lambda^{\ell+1}(M) \rightarrow \Lambda^\ell(M), \quad i_X(\omega)(X_1, \dots, X_\ell) = \omega(X, X_1, \dots, X_\ell).$$

The *Lie derivative* along the vector field X is defined as

$$\begin{aligned} L_X : \Lambda^k(M) &\rightarrow \Lambda^k(M), \\ L_X(\omega) &= i_X d\omega + d(i_X \omega). \end{aligned}$$

For a smooth n -dimensional manifold M , a k -dimensional *submanifold* in M is a subset $N \subset M$ with the property that every point $p \in N$ is contained in the domain U of a chart $\varphi : U \rightarrow \mathbb{R}^n$ such that $\varphi(U \cap N) = \varphi(U) \cap \mathbb{R}^k$.

If $k = n$ then, by the inverse function theorem, N is an open subset in M ; in this case N is also called an *open submanifold* in M . (The same is true in the topological category, but the proof is harder and requires Brouwer's Invariance of Domain Theorem, see e.g. [Hat02], Theorem 2B.3.)

Suppose that $U \subset \mathbb{R}^n$ is an open subset. A *piecewise-smooth function* $f : U \rightarrow \mathbb{R}^m$ is a continuous function such that for every $x \in U$ there exists a neighborhood V of x in U , a diffeomorphism $\phi : V \rightarrow V' \subset \mathbb{R}^n$, a triangulation T of V' , so that the composition

$$f \circ \phi^{-1} : (V', T) \rightarrow \mathbb{R}^m$$

is smooth on each simplex. Note that composition $g \circ f$ is again piecewise-smooth, provided that g is smooth; however, composition of piecewise-smooth maps need not be piecewise-smooth.

One then defines *piecewise smooth k -dimensional submanifolds* N of a smooth manifold M . Such N is a topological submanifold which is locally the image of \mathbb{R}^k in \mathbb{R}^n under a piecewise-smooth homeomorphism $\mathbb{R}^k \rightarrow \mathbb{R}^n$. We refer the reader to [Thu97] for the detailed discussion of piecewise-smooth manifolds.

If $k = n - 1$ we also sometimes call a submanifold a (*piecewise smooth*) *hypersurface*.

Below we review two alternative ways of defining submanifolds. Consider a smooth map $f : M \rightarrow N$ of a m -dimensional manifold $M = M^m$ to an n -dimensional manifold $N = N^n$. The map $f : M \rightarrow N$ is called an *immersion* if for every $p \in M$, the linear map $df_p : T_p M \rightarrow T_{f(p)} N$ is injective. If, moreover, f defines a homeomorphism from M to $f(M)$ with the subspace topology, then f is called a *smooth embedding*.

EXERCISE 3.1. Construct an injective immersion $\mathbb{R} \rightarrow \mathbb{R}^2$ which is not a smooth embedding.

If N is a submanifold in M then the inclusion map $i : N \rightarrow M$ is a smooth embedding. This, in fact, provides an alternative definition for k -dimensional submanifolds: They are images of smooth embeddings with k -dimensional manifolds (see Corollary 3.4). Images of immersions provide a large class of subsets, called *immersed submanifolds*.

A smooth map $f : M^k \rightarrow N^n$ is called a *submersion* if for every $p \in M$, the linear map df_p is surjective. The following theorem can be found for instance, in [GP10], [Hir76], [War83].

THEOREM 3.2. (1) *If $f : M^m \rightarrow N^n$ is an immersion, then for every $p \in M$ and $q = f(p)$ there exists a chart $\varphi : U \rightarrow \mathbb{R}^m$ of M with $p \in U$, and a chart $\psi : V \rightarrow \mathbb{R}^n$ of N with $q \in V$ such that the composition*

$$\bar{f} = \psi \circ f \circ \varphi^{-1} : \varphi(U) \rightarrow \psi(V)$$

is of the form

$$\bar{f}(x_1, \dots, x_m) = (x_1, \dots, x_m, \underbrace{0, \dots, 0}_{n-m \text{ times}}).$$

(2) *If $f : M^m \rightarrow N^n$ is a submersion, then for every $p \in M$ and $q = f(p)$ there exists a chart $\varphi : U \rightarrow \mathbb{R}^m$ of M with $p \in U$, and a chart $\psi : V \rightarrow \mathbb{R}^n$ of N with $q \in V$ such that the composition*

$$\bar{f} = \psi \circ f \circ \varphi^{-1} : \varphi(U) \rightarrow \psi(V)$$

is of the form

$$\bar{f}(x_1, \dots, x_n, \dots, x_m) = (x_1, \dots, x_n).$$

(3) *The constant rank theorem is a combination of (1) and (2). Suppose that the derivative of $f : M^m \rightarrow N^n$ has constant rank r . Then then for every $p \in M$ and $q = f(p)$ there exists a chart $\varphi : U \rightarrow \mathbb{R}^m$ of M with $p \in U$, and a chart $\psi : V \rightarrow \mathbb{R}^n$ of N with $q \in V$ such that the composition*

$$\bar{f} = \psi \circ f \circ \varphi^{-1} : \varphi(U) \rightarrow \psi(V)$$

is of the form

$$\bar{f}(x_1, \dots, x_m) = (x_1, \dots, x_r, 0, \dots, 0).$$

In particular, $f(U)$ is a submanifold of dimension r in N .

EXERCISE 3.3. Prove Theorem 3.2. *Hint.* Use the Inverse Function Theorem and the Implicit Function Theorem from Vector Calculus.

COROLLARY 3.4. (1) *If $f : M^m \rightarrow N^n$ is a smooth embedding then $f(M^m)$ is a m -dimensional submanifold of N^n .*

- (2) If $f : M^m \rightarrow N^n$ is a submersion then for every $x \in N^n$ the fiber $f^{-1}(x)$ is a submanifold of codimension n .

EXERCISE 3.5. Every submersion $f : M \rightarrow N$ is an *open map*, i.e. the image of an open subset in M is an open subset in N .

Let $f : M^m \rightarrow N^n$ be a smooth map and $y \in N$ is a point such that for some $x \in f^{-1}(y)$, the map $df_x : T_x M \rightarrow T_y N, y = f(x)$, is not surjective. Then the point $y \in N$ is called a *singular value* of f . A point $y \in N$ which is not a singular value of f is called a *regular value* of f . Thus, for every regular value $y \in N$ of f , the preimage $f^{-1}(y)$ is either empty or a smooth submanifold of dimension $m - n$.

THEOREM 3.6 (Sard's theorem). *Almost every point $y \in N$ is a regular value of f .*

3.2. Smooth partition of unity

DEFINITION 3.7. Let M be a smooth manifold and $\mathcal{U} = \{B_i : i \in I\}$ a locally finite cover of M by open subsets diffeomorphic to Euclidean balls. A collection of smooth functions $\{\eta_i : i \in I\}$ on M is called a *smooth partition of unity* for the cover \mathcal{U} if the following conditions hold:

- (1) $\sum_i \eta_i \equiv 1$.
- (2) $0 \leq \eta_i \leq 1, \quad \forall i \in I$.
- (3) $\text{Supp}(\eta_i) \subset \overline{B_i}, \quad \forall i \in I$.

THEOREM 3.8. *Every open cover \mathcal{U} as above admits a smooth partition of unity.*

3.3. Riemannian metrics

A *Riemannian metric* (also known as the *metric tensor*) on a smooth n -dimensional manifold M , is a positive definite inner product $\langle \cdot, \cdot \rangle_p$ defined on the tangent spaces $T_p M$ of M ; this inner product is required to depend smoothly on the point $p \in M$. We will suppress the subscript p in this notation; we let $\|\cdot\|$ denote the norm on $T_p M$ determined by the Riemannian metric. The subspace of TM consisting of unit tangent vectors is a submanifold denoted UM and called the *unit tangent bundle*: UM is a smooth submanifold of TM and the restriction of the projection $TM \rightarrow M$ is a bundle, whose fibers are $n - 1$ -dimensional spheres.

The Riemannian metric is usually denoted $g = g_x = g(x), x \in M$ or ds^2 . We will use the notation dx^2 to denote the Euclidean Riemannian metric on \mathbb{R}^n :

$$dx^2 := dx_1^2 + \dots + dx_n^2.$$

Here and in what follows we use the convention that for tangent vectors u, v ,

$$dx_i dx_j(u, v) = u_i v_j$$

and dx_i^2 stands for $dx_i dx_i$. A Riemannian metric on an open subset $\Omega \subset \mathbb{R}^n$ is determined by its *Gramm matrix* $A_x, x \in \Omega$, where A_x is a positive-definite symmetric matrix depending smoothly on x :

$$\left\langle \frac{\partial}{\partial x_i}, \frac{\partial}{\partial x_j} \right\rangle_x = g_{ij}(x),$$

the ij -th entry of the matrix A_x .

A *Riemannian manifold* is a smooth manifold equipped with a Riemannian metric.

Two Riemannian metrics g, h on a manifold M are said to be *conformal* to each other, if $h_x = \lambda(x)g_x$, where $\lambda(x)$ is a smooth positive function on M , called the *conformal factor*. In matrix notation, we just multiply the matrix A_x of g_x by a scalar function. Such modification of Riemannian metrics does not change the angles between tangent vectors. A Riemannian metric g_x on a domain Ω in \mathbb{R}^n is called *conformally-Euclidean* if it is conformal to dx^2 , i.e. it is given by

$$\lambda(x)dx^2 = \lambda(x)(dx_1^2 + \dots + dx_n^2).$$

Thus, the square of the norm of a vector $v \in T_x\Omega$ with respect to g_x is given by

$$\lambda(x) \sum_{i=1}^n v_i^2.$$

Given an immersion $f : M^m \rightarrow N^n$ and a Riemannian metric g on N , one defines the *pull-back* Riemannian metric $f^*(g)$ by

$$\langle v, w \rangle_p = \langle df(v), df(w) \rangle_q, \quad p \in M, q = f(p) \in N,$$

where in the right-hand side we use the inner product defined by g and in the left-hand side the one defined by $f^*(g)$. It is useful to rewrite this definition in terms of symmetric matrices, when M, N are open subsets of \mathbb{R}^n . Let A_y be the matrix-function defining g . Then $f^*(g)$ is given by the matrix-function B_x , where

$$y = f(x), \quad B_x = (D_x f)^T A_y (D_x f)$$

and $D_x f$ is the Jacobian matrix of f at the point x .

Let us compute how pull-back works in “calculus terms” (this is useful for explicit computations of the pull-back metrics $f^*(g)$), when $g(y)$ is a Riemannian metric on an open subset U in \mathbb{R}^n . Suppose that

$$g(y) = \sum_{i,j} g_{ij}(y) dy_i dy_j$$

and $f = (f_1, \dots, f_n)$ is a diffeomorphism $V \subset \mathbb{R}^n \rightarrow U$. Then

$$f^*(g) = h,$$

$$h(x) = \sum_{i,j} g_{ij}(f(x)) df_i df_j.$$

Here for a function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$, e.g., $\phi(x) = f_i(x)$,

$$d\phi = \sum_{k=1}^n d_k \phi = \sum_{k=1}^n \frac{\partial \phi}{\partial x_k} dx_k,$$

and, thus,

$$df_i df_j = \sum_{k,l=1}^n \frac{\partial f_i}{\partial x_k} \frac{\partial f_j}{\partial x_l} dx_k dx_l.$$

A special case of the above is when N is a submanifold in a Riemannian manifold M . One can define a Riemannian metric on N either by using the inclusion map and the pull-back metric, or by considering, for every $p \in N$, the subspace $T_p N$ of $T_p M$, and restricting the inner product $\langle \cdot, \cdot \rangle_p$ to it. Both procedures define the same Riemannian metric on N .

Measurable Riemannian metrics. The same definition makes sense if the inner product depends only measurably on the point $p \in M$, equivalently, the

matrix-function A_x is only measurable. This generalization of Riemannian metrics will be used in our discussion of quasiconformal groups, Chapter 23, section 23.6.

Gradient and divergence. A Riemannian metric g on M defines isomorphisms between tangent and cotangent spaces of M :

$$T_p(M) \rightarrow T_p^*(M),$$

where each $v \in T_p(M)$ corresponds to the linear functional

$$v^* \in T_p^*(M), \quad v^*(w) = \langle v, w \rangle.$$

In particular, one defines the *gradient vector field* ∇u of a function $u : M \rightarrow \mathbb{R}$ by dualizing the 1-form du .

Suppose now that M is n -dimensional. For a vector field X on M , the *divergence* $\operatorname{div} X$ is a function on M , which, for every n -form ω , satisfies

$$\operatorname{div} X \omega = L_X(\omega),$$

where L_X is the Lie derivative along X .

In local coordinates, divergence and gradient are given by the formulae:

$$\operatorname{div} X = \sum_{i=1}^n \frac{1}{\sqrt{|g|}} \frac{\partial}{\partial x_i} \left(\sqrt{\det(g)} X^i \right),$$

and

$$(\nabla u)^i = \sum_{j=1}^n g^{ij} \frac{\partial u}{\partial x_j}.$$

Here

$$X = (X^1, \dots, X^n),$$

and $(g^{ij}) = (g_{ij})^{-1}$ is the inverse matrix of the metric tensor g .

Length and distance. Given a Riemannian metric on M , one defines the *length* of a path $\mathbf{p} : [a, b] \rightarrow M$ by

$$(3.1) \quad \operatorname{length}(\mathbf{p}) = \int_a^b \|\mathbf{p}'(t)\| dt.$$

By abusing the notation, we will frequently denote $\operatorname{length}(\mathbf{p})$ by $\operatorname{length}(\mathbf{p}([a, b]))$.

Then, provided that M is connected, one defines the Riemannian *distance function*

$$\operatorname{dist}(p, q) = \inf_{\mathbf{p}} \operatorname{length}(\mathbf{p}),$$

where the infimum is taken over all paths in M connecting p to q .

A smooth map $f : (M, g) \rightarrow (N, h)$ of Riemannian manifolds is called a *Riemannian isometry* if $f^*(h) = g$. In most cases, such maps do not preserve the Riemannian distances. This leads to a somewhat unfortunate terminological confusion, since the same name *isometry* is used to define maps between metric spaces which preserve the distance functions. Of course, if a Riemannian isometry $f : (M, g) \rightarrow (N, h)$ is also a diffeomorphism, then it preserves the Riemannian distance function and, hence, f is an isometry of Riemannian metric spaces.

A *Riemannian geodesic segment* is a path $\mathbf{p} : [a, b] \subset \mathbb{R} \rightarrow M$ which is a local length-minimizer, i.e.:

There exists $c \geq 0$ so that for all t_1, t_2 in J sufficiently close to each other,

$$\operatorname{dist}(\mathbf{p}(t_1), \mathbf{p}(t_2)) = \operatorname{length}(\mathbf{p}([t_1, t_2])) = c|t_1 - t_2|.$$

If $c = 1$, we say that \mathbf{p} has *unit speed*. Thus, a unit speed geodesic is a locally-distance preserving map from an interval to (M, g) . This definition extends to infinite geodesics in M , which are maps $\mathbf{p} : J \rightarrow M$, defined on intervals $J \subset M$, whose restrictions to each finite interval are finite geodesics. A Riemannian metric is said to be *complete* if every geodesic segment extends to a complete geodesic $\gamma : \mathbb{R} \rightarrow M$. According to the Hopf–Rinow theorem, a Riemannian metric on a connected manifold is complete if and only if the associated distance function is complete.

A smooth map $f : (M, g) \rightarrow (N, h)$ is called *totally-geodesic* if it maps geodesics in (M, g) to geodesics in (N, h) . If, in addition, $f^*(h) = g$, then such f is locally distance-preserving.

Injectivity and convexity radii. For every complete Riemannian manifold M and a point $p \in M$, there exists the *exponential map*

$$\exp_p : T_p M \rightarrow M,$$

which sends every vector $v \in T_p M$ to the point $\gamma_v(1)$, where $\gamma_v(t)$ is the unique geodesic in M with $\gamma(0) = p$ and $\gamma'(0) = v$. If $S(0, r) \subset T_p M$, $B(0, r) \subset T_p M$ are the round sphere and round ball of radius r , then

$$\exp(S(0, r)), \quad \exp(B(0, r)) \subset M$$

are the *geodesic r -sphere* and the *geodesic r -ball* in M centered at p .

The *injectivity radius* $\text{InjRad}(p)$ of M at the point $p \in M$ is the supremum of the numbers r so that $\exp_p|B(0, r)$ is a diffeomorphism to its image. The *radius of convexity* $\text{ConRad}(p)$ is the supremum of r 's so that $r \leq \text{InRad}(p)$ and $C = \exp_p(B(0, r))$ is a convex subset of M , i.e. every $x, y \in C$ are connected by a (distance-realizing) geodesic segment entirely contained in C . It is a basic fact of Riemannian geometry that for every $p \in M$,

$$\text{ConRad}(p) > 0,$$

see e.g. [dC92].

3.4. Riemannian volume

For every n -dimensional Riemannian manifold (M, g) one defines the *volume element* (or *volume density*) denoted dV (or dA if M is 2-dimensional). Given n vectors $v_1, \dots, v_n \in T_p M$, $dV(v_1 \wedge \dots \wedge v_n)$ is the volume of the parallelepiped in $T_p M$ spanned by these vectors. This volume is nothing but $\sqrt{|\det(G(v_1, \dots, v_n))|}$, where $G(v_1, \dots, v_n)$ is the Gram matrix with the entries $\langle v_i, v_j \rangle$. If $ds^2 = \rho^2(x)dx^2$, is a conformally-Euclidean metric on an open subset of \mathbb{R}^n , $\rho > 0$, then the volume density of ds^2 is given by

$$\rho^n(x)dx_1 \dots dx_n.$$

Thus, every Riemannian manifold has a canonical measure, given by the integral of its volume form

$$\text{mes}(E) = \int_A dV.$$

THEOREM 3.9 (Generalized Rademacher's theorem). *Let $f : M \rightarrow N$ be a Lipschitz map of Riemannian manifolds. Then f is differentiable almost everywhere.*

EXERCISE 3.10. Deduce Theorem 3.9 from Theorem 2.28 and the fact that M is second countable.

We now define *volumes of maps and submanifolds*. The simplest and the most familiar notion of volume of maps comes from the vector calculus. Let Ω be a bounded region in \mathbb{R}^n and $f : \Omega \rightarrow \mathbb{R}^n$ be a smooth map. Then the *geometric volume* of f is defined as

$$(3.2) \quad Vol(f) := \int_{\Omega} |J_f(x)| dx_1 \dots dx_n,$$

where J_f is the Jacobian determinant of f . Note that we are integrating here a non-negative quantity, hence, the geometric volume of a map is always non-negative. If f were 1-1 and $J_f(x) > 0$ for every x , then, of course,

$$Vol(f) = \int_{\Omega} J_f(x) dx_1 \dots dx_n = Vol(f(\Omega)).$$

More generally, if $f : \Omega \rightarrow \mathbb{R}^m$ (now, m need not be equal to n), then

$$Vol(f) = \int_{\Omega} \sqrt{\det(G_f)},$$

where G_f is the Gram matrix with the entries $\left\langle \frac{\partial f}{\partial x_i}, \frac{\partial f}{\partial x_j} \right\rangle$, where brackets denote the usual inner product in \mathbb{R}^m . In case f is an embedding, the reader will recognize in this formula the familiar expression for the volume of the submanifold $\Sigma = f(\Omega)$ in \mathbb{R}^m ,

$$Vol(f) = \int_{\Sigma} dS.$$

The Gram matrix above makes sense also for maps whose target is an m -dimensional Riemannian manifold (M, g) , with partial derivatives replaced with vectors $df(X_i)$ in M , where X_i are coordinate vector fields in Ω :

$$X_i = \frac{\partial}{\partial x_i}, i = 1, \dots, n.$$

Furthermore, one can take the domain of the map f to be an arbitrary smooth manifold N (possibly with boundary). The definition of volume still makes sense and is independent of the choice of local charts on N used to define the integral: This independence is a corollary of the change of variables formula in the integral in \mathbb{R}^n . More precisely, consider charts $\varphi_{\alpha} : U_{\alpha} \rightarrow V_{\alpha} \subset N$, so that $\{V_{\alpha}\}_{\alpha \in J}$ is a locally-finite open covering of N . Let $\{\eta_{\alpha}\}$ be a partition of unity on N corresponding to this cover. Then for $\zeta_{\alpha} = \eta_{\alpha} \circ \varphi_{\alpha}$, $f_{\alpha} = f \circ \varphi_{\alpha}$,

$$Vol(f) = \sum_{\alpha \in J} \int_{U_{\alpha}} \zeta_{\alpha} \sqrt{|\det(G_{f_{\alpha}})|} dx_1 \dots dx_n$$

In particular, if f is 1-1 and $\Sigma = f(N)$, then

$$Vol(f) = Vol(\Sigma).$$

Observe that the formula for $Vol(f)$ makes sense when $f : N \rightarrow M$ is merely Lipschitz, in view of Theorem 3.9.

Thus, one can define the volume of an immersed submanifold, as well as that of a piecewise smooth submanifold; in the latter case we subdivide a piecewise-smooth submanifold in a union of images of simplices under smooth maps.

By abuse of language, sometimes, when we consider an open submanifold N in M , so that boundary ∂N of N a submanifold of codimension 1, while we denote

the volume of N by $Vol(N)$, we shall call the volume of ∂N the *area*, and denote it by $Area(\partial N)$.

In Section 7.10.1 we will introduce combinatorial/simplicial/cellular analogues of the Riemannian volume of maps, for this reason, for a Lipschitz map $f : N \rightarrow M$ we will use the notation $Vol^{met}(f)$ for its Riemannian (metric volume).

EXERCISE 3.11. (1) Suppose that $f : \Omega \subset \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a smooth map so that $|d_x f(u)| \leq 1$ for every unit vector u and every $x \in \Omega$. Show that $|J_f(x)| \leq 1$ for every x and, in particular,

$$Vol(f(\Omega)) = \left| \int_{\Omega} J_f dx_1 \dots dx_n \right| \leq Vol(f) \leq Vol(\Omega).$$

Hint: Use the fact that under the linear map $A = d_x f$, the image of every r -ball is contained in an r -ball.

(2) Prove the same thing if the map f is merely 1-Lipschitz.

More general versions of the above exercises are the following.

EXERCISE 3.12. Let (M, g) and (N, h) be n -dimensional Riemannian manifolds.

(1) Let $f : M \rightarrow N$ be a smooth map such that for every $x \in M$, the norm of the linear map

$$df_x : (T_x M, \langle \cdot, \cdot \rangle_g) \rightarrow (T_x(x)N, \langle \cdot, \cdot \rangle_h)$$

is at most L .

Prove that $|J_f(x)| \leq L^n$ for every x and that for every open subset U of M

$$Vol(f(U)) \leq L^n Vol(U).$$

(2) Prove the same statement for an L -Lipschitz map $f : M \rightarrow N$.

A consequence of Theorem 3.2 is the following.

THEOREM 3.13. Consider a compact Riemannian manifold M^m , a submersion $f : M^m \rightarrow N^n$. For every $x \in N$ set $M_x := f^{-1}(x)$. Then, for every $p \in N$ and every $\epsilon > 0$ there exists an open neighborhood W of p such that for every $x \in W$,

$$1 - \epsilon \leq \frac{Vol(M_x)}{Vol(M_p)} \leq 1 + \epsilon.$$

PROOF. First note that, by compactness of M_p , for every neighborhood U of M_p there exists a neighborhood W of p such that $f^{-1}(W) \subset U$.

According to Theorem 3.2, (2), for every $x \in M_p$ there exists a chart of M , $\varphi_x : U_x \rightarrow \tilde{U}_x$, with U_x containing x , and a chart of N , $\psi_x : V_x \rightarrow \tilde{V}_x$ with V_x containing p , such that $\psi_x \circ f \circ \varphi_x^{-1}$ is a restriction of the projection to the first n coordinates. Without loss of generality we may assume that \tilde{U}_x is an open cube in \mathbb{R}^m . Therefore, \tilde{V}_x is also a cube in \mathbb{R}^n , and $\tilde{U}_x = \tilde{V}_x \times \tilde{Z}_x$, where \tilde{Z}_x is an open subset in \mathbb{R}^{m-n} .

Since M_p is compact, it can be covered by finitely many such domains of charts U_1, \dots, U_k . Let V_1, \dots, V_k be the corresponding domains of charts containing p . For the open neighborhood $U = \bigcup_{i=1}^k U_i$ of M_p consider an open neighborhood W of p , contained in $\bigcap_{i=1}^k V_i$, such that $f^{-1}(W) \subseteq U$.

For every $x \in W$, $M_x = \bigcup_{l=1}^k (U_l \cap M_x)$. Fix $l \in \{1, \dots, k\}$. Let $(g_{ij}(y))_{1 \leq i, j \leq n}$ be the matrix-valued function on \tilde{U}_l , defining the pull-back by φ_l of the Riemannian metric on M .

Since g_{ij} is continuous, there exists a neighborhood \tilde{W}_l of $\tilde{p} = \psi_l(p)$ such that for every $\tilde{x} \in \tilde{W}_l$ and for every $\tilde{t} \in \tilde{Z}_l$ we have,

$$(1 - \epsilon)^2 \leq \frac{\det [g_{ij}(\tilde{x}, \tilde{t})]_{n+1 \leq i, j \leq m}}{\det [g_{ij}(\tilde{p}, \tilde{t})]_{n+1 \leq i, j \leq m}} \leq (1 + \epsilon)^2.$$

Recall that the volumes of $M_x \cap U_l$ and of $M_p \cap U_l$ are obtained by integrating respectively

$$(\det [g_{ij}(\tilde{x}, \tilde{t})]_{n+1 \leq i, j \leq k})^{1/2}$$

and

$$(\det [g_{ij}(\tilde{p}, \tilde{t})]_{n+1 \leq i, j \leq k})^{1/2}$$

on Z_l . The volumes of M_x and M_p are obtained by combining this with a partition of unity.

It follows that for $x \in \bigcap_{i=1}^k \psi_i^{-1}(\tilde{W}_l)$,

$$1 - \epsilon \leq \frac{\text{Vol}(M_x)}{\text{Vol}(M_p)} \leq 1 + \epsilon.$$

□

Finally, we recall an important formula for volume computations:

THEOREM 3.14 (Coarea formula, see e.g. Theorem 6.3 in [Cha06] and 3.2.22 in [Fed69]). *Let $f : M \rightarrow (0, \infty)$ be a smooth function on a Riemannian manifold M . For almost every $t \in (0, \infty)$, the level set $\mathcal{H}_t := f^{-1}(t)$ is a smooth hypersurface in M ; let dA_t be the Riemannian area density induced on \mathcal{H}_t and dV be the Riemannian volume density of M . Then, for every function $g \in L^1(M)$,*

$$\int_M g |\nabla f| dV = \int_0^\infty dt \int_{\mathcal{H}_t} g dA_t.$$

3.5. Volume growth and isoperimetric functions. Cheeger constant

In this section we present several basic notions, initially introduced in Riemannian geometry and later adapted and used in group theory and in combinatorics. These notions and their *coarse analogues* will appear frequently in this book.

Volume growth. Given a Riemannian manifold M and a basepoint $x_0 \in M$, the (*volume*) *growth function* is defined as

$$\mathfrak{G}_{M, x_0}(r) := \text{Vol} B(x_0, r),$$

the volume of the metric ball of radius r and center at x_0 in M .

REMARKS 3.15. (1) For two different points x_0, y_0 , we have

$$\mathfrak{G}_{M, x_0}(r) \leq \mathfrak{G}_{M, y_0}(r + d), \quad \text{where } d = \text{dist}(x_0, y_0).$$

- (2) Suppose that the action of the isometry group of M is *cobounded* on M , i.e. there exists a constant κ such that the orbit of $B(x_0, \kappa)$ under the group $\text{Isom}(M)$, is the entire manifold M . (For instance, this is the case if M is a regular covering space of a compact Riemannian manifold.) Then, for every two basepoints x_0, y_0

$$\mathfrak{G}_{M, x_0}(r) \leq \mathfrak{G}_{M, y_0}(r + \kappa).$$

Thus, in this case, the growth rate of the function \mathfrak{G} does not depend on the choice of the basepoint.

We refer the reader to Section 8.7 for the detailed discussion of volume growth and its relation to *group growth*.

EXERCISE 3.16. Assume again that the action $\text{Isom}(M) \curvearrowright M$ is cobounded, the constant κ is as above, and that M is complete.

- (1) Prove that the growth function is *almost sub-multiplicative*, that is:

$$\mathfrak{G}_{M, x_0}((r + t)\kappa) \leq \mathfrak{G}_{M, x_0}(r\kappa)\mathfrak{G}_{M, x_0}((t + 1)\kappa).$$

- (2) Prove that the growth function of M is at most exponential, that is, there exists $a > 1$ such that

$$\mathfrak{G}_{M, x_0}(x) \leq a^x, \text{ for every } x \geq 0.$$

Isoperimetric inequalities and isoperimetric functions. Isoperimetric problems in geometry go back to the antiquity: (Dido's problem) *Which region of the given perimeter in the Euclidean plane \mathbb{R}^2 has the least area?* The answer "the round disk" is intuitively obvious, but, surprisingly, hard to prove. This classical problem explain the terminology *isoperimetric* below.

In general, isoperimetric problem in Riemannian geometry have the following *minimax* form.

Consider a complete connected n -dimensional Riemannian manifold M (which may or may not be closed). Fix a number $k \in \mathbb{N}$ and all consider closed k -dimensional submanifolds $Z \subset M$ (or maps $Z \rightarrow M$ of closed k -manifolds to M or, more generally, k -cycles in M). Assume, now, depending on the context, that each $Z \subset M$ bounds a $k + 1$ -dimensional submanifold, or a $k + 1$ -chain B , or that the map $Z \rightarrow M$ extends to a map $B \rightarrow M$, where B is a compact manifold with boundary equal to Z . The the latter case, one typically assumes that $Z = \mathbb{S}^k$ and B is the $k + 1$ -ball. To unify the notation, we will simply say that $Z = \partial B$, even in the case of maps $Z \rightarrow M$.

Next, among all these B 's (or their maps), one looks for the one of the least $k + 1$ -volume. (The minimum may not exist, in which case one takes the infimum.) This least volume is the *filling volume* of Z . Lastly, among all Z 's with $\text{Vol}_k(Z) \leq L$, one looks for the ones which have the largest filling volume (again, taking the supremum in general). This defines the *isoperimetric function* of M :

$$(3.3) \quad IP_{M, k}^{met} = IP_{M, k}(L) = \sup_{Z, \text{Vol}_k(Z) \leq L} \inf_{B, \partial B = Z} \text{Vol}_{k+1}(B).$$

REMARK 3.17. In Section 9.7 we introduce other isoperimetric functions of more *combinatorial* and *coarse geometric* nature. In order to distinguish the Riemannian isoperimetric functions for those, we will use the notation $IP_{M, k}^{met}$ when convenient.

In each setting (submanifolds, maps, cycles), we get a different isoperimetric function, of course.

We will be primarily interested in two cases (Z having codimension 1 and dimension 1 respectively):

1. $k = n - 1$, Z is a (smooth) closed hypersurface in M .
2. $k = 1$, $Z = \mathbb{S}^1$, where we consider Lipschitz maps $Z \rightarrow M$ and their extensions $B = \mathbb{D}^2 \rightarrow M$ (“filling disks”).

Perhaps surprisingly, asymptotic behavior of isoperimetric functions in these two cases goes long way towards determining the asymptotic geometry of M . Suppose, for instance, that M is a regular cover of a compact Riemannian manifold, with the group G of covering transformations. Then the dichotomy *linear/superlinear* for both isoperimetric functions serves as a major demarkation line in the world of finitely generated groups:

1. The condition $IP_{M,1}(L) \approx L$ (linear growth of the filling area) yields the class of *Gromov-hyperbolic groups* G . This linearity condition can be regarded as *asymptotically negative sectional curvature* of the manifold M (and the group G).
2. The condition $IP_{M^n, n-1}(L) \approx L$ yields the class of *nonamenable groups* G .

Here we are using the notation \approx introduced in the Definition 1.3.

The following Riemannian geometry theorem illustrates the power of this linear/nonlinear dichotomy:

THEOREM 3.18. *Suppose that M is a Riemannian manifold which is the universal cover of a compact Riemannian manifold. Then*

$$IP_{M,1}(L) \approx L \Rightarrow IP_{M,k}(L) \approx L,$$

for all $k \geq 2$. Here for $k \geq 2$ one can equally use either the homological filling or filling of maps of spheres by maps of disks. For the former, one needs to assume that $H_i(M) = 0, i \leq k$ and for the latter one requires that $\pi_i(M) = 0, i \leq k$.

As far as we know, this theorem does not have a “purely Riemannian” proof: One first verifies that the group $G = \pi_1(M)$ is Gromov-hyperbolic (Theorem 11.181), then proves that all such groups have linear isoperimetric functions of in all degrees [Lan00, Min01] and, then uses the approximate equality of isoperimetric functions of M and of G (cf. Theorem 9.75).

Below we discuss the “codimension 1” isoperimetric function in more detail. If M is connected and non-compact, each closed hypersurface in M bounds exactly one compact submanifold, which leads to

DEFINITION 3.19. Suppose that $F : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a function and M is a (connected) non-compact n -dimensional Riemannian manifold. Then M is said to satisfy the *isoperimetric inequality* of the form

$$Vol(\Omega) \leq F(Area(\partial\Omega)),$$

if this inequality holds for all open submanifolds $\Omega \subset M$ with compact closure and smooth boundary.

EXERCISE 3.20. The above definition is equivalent to the inequality

$$IP_{M, n-1}^{met}(L) \leq F(L)$$

for every $L > 0$. (Note: Hypersurfaces in the definition of $IP_{M, n-1}^{met}$ need not be connected.)

For instance, if M is the Euclidean plane, then

$$(3.4) \quad 4\pi \text{Ar}(c) \leq \ell^2(c),$$

for every loop c (with equality realized precisely in the case when c is a round circle). Thus,

$$IP_{\mathbb{R}^2}^{met}(\ell) = \frac{\ell^2}{4\pi}.$$

The Cheeger constant. As the main dichotomy in the case of codimension 1 isoperimetric inequality is linear/nonlinear, it makes sense to look at the ratio between areas of hypersurfaces in M and volumes of domains in M which they bound. If M is compact, connected and the hypersurface is connected, then there are exactly two such domains. In line with the definition of the isoperimetric function, we will be choosing the domain with the least volume. This motivates:

DEFINITION 3.21. The *Cheeger (isoperimetric) constant* $h(M)$ (or *isoperimetric ratio*) of M is the infimum of the ratios

$$\frac{\text{Area}(\partial\Omega)}{\min[\text{Vol}(\Omega), \text{Vol}(M \setminus \Omega)]},$$

where Ω varies over all open non-empty submanifolds with compact closure and smooth boundary.

In particular, if $h(M) \geq \kappa > 0$, then the following isoperimetric inequality holds in M :

$$\text{Vol}(\Omega) \leq \frac{1}{\kappa} \text{Area}(\partial\Omega).$$

Cheeger constant was defined by J. Cheeger for compact manifolds in [Che70]. Further details can be found for instance in P. Buser's book [Bus10]. Note that when M is a Riemannian manifold of infinite volume, one may replace the denominator in the ratio defining the Cheeger constant by $\text{Vol}(\Omega)$.

Assume now that M is the universal cover of a compact Riemannian manifold N . A natural question to ask is to what extent the growth function and the Cheeger constant of M depend on the choice of the Riemannian metric on N . The first question, in a way, was one of the origins of the Geometric Group Theory.

Vadim Efremovich [Efr53] noted that two growth functions corresponding to two different choices of metrics on N are asymptotically equal (see Definition 1.4) and, moreover, that their asymptotic equivalence class is determined by the fundamental group of N only. See Proposition 8.80 for a slightly more general statement.

A similar phenomenon occurs with the Cheeger constant: Positivity of $h(M)$ does not depend on the metric on N , it depends only on a certain property of $\pi_1(N)$, namely, the non-amenability, see Remark 18.15. This was proved much later by Robert Brooks [Bro81a, Bro82a]. Brooks' argument has a global analytic flavor, as it uses the connection established by Jeff Cheeger [Che70] between positivity of the isoperimetric constant and positivity of spectrum of the Laplace–Beltrami operator on M . This result was highly influential in global analysis on manifolds and harmonic analysis on graphs and manifolds.

3.6. Curvature

Instead of defining the Riemannian curvature tensor, we will only describe some properties of Riemannian curvature. First, if (M, g) is a 2-dimensional Riemannian manifold, one defines the *Gaussian curvature* of (M, g) , which is a smooth function $K : M \rightarrow \mathbb{R}$, whose values are denoted $K(p)$ and K_p .

More generally, for an n -dimensional Riemannian manifold (M, g) , one defines the *sectional curvature*, which is a function $\Lambda^2 TM \rightarrow \mathbb{R}$, denoted $K_p(u, v) = K_{p,g}(u, v)$:

$$K_p(u, v) = \frac{\langle R(u, v)u, v \rangle}{|u \wedge v|^2},$$

provided that $u, v \in T_p M$ are linearly independent. Here R is the Riemannian curvature tensor and $|u \wedge v|$ is the area of the parallelogram in $T_p M$ spanned by the vectors u, v . Sectional curvature depends only on the 2-plane P in $T_p M$ spanned by u and v . The curvature tensor $R(u, v)w$ does not change if we replace the metric g with a conformal metric $h = ag$, where $a > 0$ is a constant. Thus,

$$K_{p,h}(u, v) = a^{-1}K_{p,g}(u, v).$$

Totally geodesic Riemannian isometric immersions $f : (M, g) \rightarrow (N, h)$ preserve sectional curvature:

$$K_p(u, v) = K_q(df(u), df(v)), \quad q = f(p).$$

In particular, sectional curvature is invariant under Riemannian isometries of equidimensional Riemannian manifolds. In the case when M is 2-dimensional, $K_p(u, v) = K_p$, is the Gaussian curvature of M .

Gauss-Bonnet formula. Our next goal is to connect areas of triangles to curvature.

THEOREM 3.22 (Gauss-Bonnet formula). *Let (M, g) be a Riemannian surface with the Gaussian curvature $K(p), p \in M$ and the area form dA . Then for every 2-dimensional triangle $\blacktriangle \subset M$ with geodesic edges and vertex angles α, β, γ ,*

$$\int_{\blacktriangle} K(p)dA = (\alpha + \beta + \gamma) - \pi.$$

In particular, if $K(p)$ is constant equal κ , we get

$$-\kappa \text{Area}(\blacktriangle) = \pi - (\alpha + \beta + \gamma).$$

The quantity $\pi - (\alpha + \beta + \gamma)$ is called the *angle deficit* of the triangle \blacktriangle .

Curvature and volume. Below we describe the relation of uniform lower and upper bounds on the sectional curvature and the growth of volumes of balls, that will be used in the sequel. The references for these results are [BC01, Section 11.10], [CGT82], [Gro86], [G60] and [GHL04], Theorem 3.101, p. 140.

We will use the following notation: For $\kappa \in \mathbb{R}$, we let $A_\kappa(r)$ and $V_\kappa(r)$ denote the area of the sphere, respectively the volume of the ball of radius r , in the n -dimensional space of constant sectional curvature κ . We will also denote by $A(x, r)$ the area of the geodesic sphere of radius r and center x in a given Riemannian manifold M . Likewise, $V(x, r)$ will denote the volume of the geodesic ball centered at x and of radius r in M .

THEOREM 3.23 (Bishop–Gromov–Günther). *Let M be a complete n -dimensional Riemannian manifold.*

- (1) *Assume that the sectional curvature on M is at least a . Then, for every point $x \in M$:*
 - $A(x, r) \leq A_a(r)$ and $V(x, r) \leq V_a(r)$.
 - The functions $r \mapsto \frac{A(x, r)}{A_a(r)}$ and $r \mapsto \frac{V(x, r)}{V_a(r)}$ are non-increasing.
- (2) *Assume that the sectional curvature on M is at most b . Then, for every $x \in M$ with injectivity radius $\rho_x = \text{InjRad}_M(x)$:*
 - For all $r \in (0, \rho_x)$, we have $A(x, r) \geq A_b(r)$ and $V(x, r) \geq V_b(r)$.
 - The functions $r \mapsto \frac{A(x, r)}{A_b(r)}$ and $r \mapsto \frac{V(x, r)}{V_b(r)}$ are non-decreasing on the interval $(0, \rho_x)$.

The results (1) in the theorem above are also true if the Ricci curvature of M is at least $(n - 1)a$.

EXERCISE 3.24. Use this inequality to show that every n -dimensional Riemannian manifold M of nonnegative Ricci curvature has at most polynomial growth:

$$\mathfrak{G}_M(r) \lesssim r^n.$$

Theorem 3.23 follows from infinitesimal versions of the above inequalities (see Theorems 3.6 and 3.8 in [Cha06]). A consequence of the infinitesimal version of Theorem 3.23, (1), is the following theorem which will be useful in the proof of the quasiisometric invariance of positivity of the Cheeger constant:

THEOREM 3.25 (Buser’s inequality [Bus82], [Cha06], Theorem 6.8). *Let M be a complete n -dimensional manifold with sectional curvature at least a . Then there exists a positive constant λ depending on n, a and $r > 0$, such that the following holds. Given a hypersurface $\mathcal{H} \subset M$ and a ball $B(x, r) \subset M$ such that $B(x, r) \setminus \mathcal{H}$ is the union of two open subsets Ω_1, Ω_2 separated by \mathcal{H} , we have:*

$$\min [\text{Vol}(\Omega_1), \text{Vol}(\Omega_2)] \leq \lambda \text{Area} [\mathcal{H} \cap B(x, r)].$$

3.7. Riemannian manifolds of bounded geometry

DEFINITION 3.26. We say that a Riemannian manifold M has *bounded geometry* if it is connected, complete, has uniform upper and lower bounds for the sectional curvature:

$$a \leq K_p(u, v) \leq b$$

(for all $p \in M, u, v \in T_p(M)$) and a uniform lower bound for the injectivity radius:

$$\text{InjRad}(x) > \epsilon > 0.$$

Probably the correct terminology should be “uniformly locally bounded geometry”, but we prefer shortness to an accurate description. The numbers a, b, ϵ in this definition are called *geometric bounds* on M . For instance, every compact connected Riemannian manifold M has bounded geometry, every covering space of M (with pull-back Riemannian metric) also has bounded geometry. More generally, if M is connected, complete and the action of the isometry group on M is cobounded, then M has bounded geometry.

EXERCISE 3.27. Every non-compact manifold of bounded geometry has infinite volume.

REMARK 3.28. One frequently encounters weaker notions of bounded geometry for Riemannian manifold, e.g.:

1. There exists $L \geq 1$ and $R > 0$ such that every ball of radius R in M is L -bi-Lipschitz equivalent to the ball of radius R in \mathbb{R}^n . (This notion is used, for instance, by Gromov in [Gro93], §0.5.A₃).

2. The Ricci curvature of M has a uniform lower bound ([Cha06], [Cha01]).

For the purposes of this book, the restricted condition in Definition 3.26 suffices.

The following theorem connects Gromov's notion of bounded geometry with the one used in this book:

THEOREM 3.29 (See e.g. Theorem 1.14, [Att94]). *Let M be a Riemannian manifold of bounded geometry with geometric bounds a, b, ϵ . Then for every $x \in M$ and $0 < r < \epsilon/2$, the exponential map*

$$\exp_x : B(0, r) \rightarrow B(x, r) \subset M$$

is an L -bi-Lipschitz diffeomorphism, where $L = L(a, b, \epsilon)$.

This theorem also allows one to refine the notion of partition of unity in the context of Riemannian manifolds of bounded geometry:

LEMMA 3.30. *Let M be a Riemannian manifold of bounded geometry and let $\mathcal{U} = \{B_i = B(x_i, r_i) : i \in I\}$ be a locally finite covering of M by metric balls so that $\text{InjRad}_M(x_i) > 2r_i$ for every i and*

$$B\left(x_i, \frac{3}{4}r_i\right) \cap B\left(x_j, \frac{3}{4}r_j\right) = \emptyset, \quad \forall i \neq j.$$

Then \mathcal{U} admits a smooth partition of unity $\{\eta_i : i \in I\}$ which, in addition, satisfies the following properties:

1. $\eta_i \equiv 1$ on every ball $B(x_i, \frac{r_i}{2})$.
2. Every smooth functions η_i is L -Lipschitz for some L independent of i .

In what follows we keep the notation $V_\kappa(r)$ from Theorem 3.23 for the volume of a ball of radius r in the n -dimensional space of constant sectional curvature κ .

LEMMA 3.31. *Let M be complete n -dimensional Riemannian manifold with bounded geometry, let $a \leq b$ and $\rho > 0$ be such that the sectional curvature of M varies in the interval $[a, b]$ and that at every point of M the injectivity radius is larger than ρ . Then:*

- (1) *For every $\delta > 0$, every δ -separated set in M is ϕ -uniformly discrete, with $\phi(r) = \frac{V_a(r+\lambda)}{V_b(\lambda)}$, where λ is the minimum of $\frac{\delta}{2}$ and ρ .*
- (2) *For every $2\rho > \delta > 0$ and every maximal δ -separated set N in M , the multiplicity of the covering $\{B(x, \delta) \mid x \in N\}$ is at most $\frac{V_a(\frac{3\delta}{2})}{V_b(\frac{\delta}{2})}$.*

PROOF. (1) Let S be a δ -separated subset in M .

According to Theorem 3.23, for every point $x \in S$ and radius $r > 0$ we have:

$$V_a(r + \lambda) \geq \text{Vol}[B_M(x, r + \lambda)] \geq \text{card}[\overline{B}(x, r) \cap S] V_b(\lambda).$$

This inequality implies that

$$\text{card } [\overline{B}(x, r) \cap S] \leq \frac{V_a(r + \lambda)}{V_b(\lambda)},$$

whence, S with the induced metric is ϕ -uniformly discrete, with the required ϕ .

(2) Let F be a subset in N such that the intersection

$$\bigcap_{x \in F} B(x, \delta)$$

is non-empty. Let y be a point in this intersection. Then the ball $B(y, \frac{3\delta}{2})$ contains the disjoint union $\bigsqcup_{x \in F} B(x, \frac{\delta}{2})$, whence

$$V_a\left(\frac{3\delta}{2}\right) \geq \text{Vol}\left[B_M\left(y, \frac{3\delta}{2}\right)\right] \geq \text{card } [F] V_b\left(\frac{\delta}{2}\right).$$

□

3.8. Metric simplicial complexes of bounded geometry and systolic inequalities

In this section we describe a discretization of manifolds of bounded geometry via *metric simplicial complexes*. Another method of approximating of Riemannian manifolds by simplicial complexes will be described in Section 8.3, cf. Theorem 8.52.

Let X be a simplicial complex and d a path-metric on X . Then (X, d) is said to be a *metric simplicial complex* if the restriction of d to each simplex is isometric to a Euclidean simplex. The main example of a metric simplicial complex is a generalization of a graph with the standard metric described below.

Let X be a connected simplicial complex. As usual, we will often conflate X and its geometric realization. Metrize each k -simplex of X to be isometric to the standard k -simplex Δ^k in the Euclidean space:

$$\Delta^k = (\mathbb{R}_+)^{k+1} \cap \{x_0 + \dots + x_k = 1\}.$$

Thus, for each m -simplex Δ^m and its face Δ^k , the inclusion $\Delta^k \rightarrow \Delta^m$ is an isometric embedding. This allows us to define a path-metric on X so that each simplex is isometrically embedded in X , similarly to the definition of the standard metric on a graph and the Riemannian distance function. Namely, a *piecewise-linear (PL) path* \mathfrak{p} in X is a path $\mathfrak{p} : [a, b] \rightarrow X$, whose domain can be subdivided in finitely many intervals $[a_i, a_{i+1}]$ such that each restriction

$$\mathfrak{p}|_{[a_i, a_{i+1}]}$$

is a piecewise-linear path whose image is contained in a single simplex of X . Lengths of such paths are defined using the Euclidean metric on simplices of X . Then

$$d(x, y) = \inf_{\mathfrak{p}} \text{length}(\mathfrak{p}),$$

where the infimum is taken over all PL paths in X connecting x to y . The metric d is then a path-metric; we call this metric the *standard metric* on X .

EXERCISE 3.32. Verify that the standard metric is complete and that X is a geodesic metric space.

For Lipschitz maps $f : N \rightarrow X$ from smooth manifolds to simplicial complexes with at most countably many simplices, equipped with the standard piecewise-Euclidean structure we define the notion of volume as the integral

$$(3.5) \quad Vol^{met}(f) := \sum_{\sigma} \int_{f^{-1}(\text{int } \sigma)} \sqrt{G_f}$$

where the sum is taken over all simplices in X and the Gram matrix is defined for the map to the corresponding open simplex. The superscript *met* is used to distinguish this notion of volume from the combinatorial and coarse concepts in Chapter 9.7. The definition that we are using here is a special concept of volume of maps from smooth manifolds to metric spaces. We refer the reader to [Wen05] for the notion of volume of maps to general metric spaces. Given the metric volume of maps, we define metric isoperimetric functions $IP_{X,n}^{met}$ of the simplicial complex X exactly as in (3.3).

DEFINITION 3.33. A metric simplicial complex X has *bounded geometry* if it is connected and if there exist $L \geq 1$ and $N < \infty$ such that:

- every vertex of X is incident to at most N edges;
- the length of every edge is $\leq L$.
- The volume of every simplex is $\geq L^{-1}$.

In particular, the set of vertices of X with the induced metric is a uniformly discrete metric space.

Thus, a metric simplicial complex of bounded geometry is necessarily finite-dimensional.

- EXAMPLE 3.34.**
- If Y is a finite connected metric simplicial complex, then its universal cover (with the pull-back path metric) has bounded geometry.
 - A connected simplicial complex (with the standard metric) has bounded geometry if and only if there is a uniform bound on the valency of the vertices in its 1-skeleton.

Analogously to simplicial complexes of bounded geometry one defines *almost regular cell complexes* X of bounded geometry by requiring that:

- (a) Each cell c of X is contained in the image of at most D cells.
- (b) There are only finitely many combinatorial types of polyhedra which appear in the definition of attaching maps for cells in X .

Metric simplicial complexes of bounded geometry appear in the context of Riemannian manifolds with bounded geometry.

DEFINITION 3.35. Let M be a Riemannian manifold. A *bounded geometry triangulation* \mathcal{T} of M is a metric simplicial complex X of bounded geometry together with a bi-Lipschitz homeomorphism $\tau : X \rightarrow M$.

Every smooth manifold admits a triangulation (see [Cai61] for an especially simple proof); however, a general Riemannian manifold M will not have a uniform triangulation. An easy sufficient condition for uniformity of (any) triangulation of M is compactness of M . Lifting a finite triangulation \mathcal{T} of a compact Riemannian manifold M to its Riemannian covering $M' \rightarrow M$ results in a bounded geometry triangulation \mathcal{T}' of M' .

Proofs of the following theorem are outlined in [Att94, Theorem 1.14] and [ECH⁺92, Theorem 10.3.1]; a detailed proof in the case of hyperbolic manifolds can be found in [Bre09].

THEOREM 3.36. *Every Riemannian manifold of bounded geometry admits a bounded geometry triangulation. Furthermore, there exists a function $L = L(m, a, b, \epsilon)$ with the following property. Let M be an m -dimensional Riemannian manifold of bounded geometry with the geometric bounds a, b, ϵ . Then M admits a bounded geometry triangulation \mathcal{T} , which is a simplicial complex X equipped with the standard metric together with an L -bi-Lipschitz homeomorphism $\tau : X \rightarrow M$, such that geometric bounds on X depend only on m, a, b and ϵ .*

Given lack of a detailed proof, this theorem should be currently treated as a conjecture. Nevertheless, a *homotopy form* of this theorem is not all that hard:

THEOREM 3.37. *If (M, g) is a Riemannian manifold of bounded geometry then there exists a simplicial complex X of bounded geometry (with the standard metric) and a pair of L -Lipschitz maps*

$$f : M \rightarrow X, \quad \bar{f} : X \rightarrow M$$

which form a homotopy-equivalence between M and X . Furthermore, the homotopies

$$H : M \times [0, 1] \rightarrow M, \quad \bar{H} : X \times [0, 1] \rightarrow X$$

between $\bar{f} \circ f$ and id_M , and $f \circ \bar{f}$ and id_X respectively, are also L -Lipschitz maps. Here L and geometric bounds on X depend only on m, a, b and ϵ .

PROOF. Let $\epsilon > 0$ be the injectivity radius of M and $a \leq b$ the constants bounding the curvature. Pick

$$r < \min\left(\frac{\epsilon}{2}, \frac{1}{4}\pi|b|^{-1/2}\right).$$

Then r is smaller than the *convexity radius* of (M, g) , see e.g. [Pet16, p. 177].

For a probability measure μ whose support is contained in $B(x, r)$ define its center of mass $Center(\mu)$ as the unique point of minimum for the function

$$\varphi(z) := \int_{B(x, r)} d^2(z, y) d\mu(y),$$

cf. [BK81, §8].

Consider the cover of M by suitable open metric balls $B(x_i, r), i \in I$, where the centers x_i form an $\frac{r}{10}$ -separated $\frac{r}{2}$ -net in M . Let X denote the nerve of this cover; we identify the vertices of X with the centers x_i of the balls. Equip X with the standard metric.

For each simplex $\sigma = [x_1, \dots, x_n]$ in X we pick a point

$$y_\sigma \in B(x_1, r) \cap \dots \cap B(x_n, r).$$

The map f is defined via a suitable Lipschitz partition of unity $(\eta_i)_{i \in I}$ subordinate to the covering $\{B(x_i, r)\}_{i \in I}$ (cf. [Kap09, §6]), namely, we take the partition of unity as in Lemma 3.30. We identify X with a subcomplex in the infinite-dimensional simplex

$$\Delta^\infty \subset \bigoplus_{i \in I} \mathbb{R},$$

$$\Delta^\infty = \{(t_i)_{i \in I} \mid 0 \leq t_i, i \in I, \sum_{i \in I} t_i = 1\}.$$

Then f is defined as

$$y \mapsto (\eta_i(y))_{i \in I}.$$

In order to construct the map $\bar{f} : X \rightarrow M$ we first take the barycentric subdivision X' of X . The vertices v of X' are labelled by the simplices σ of X , $v = v_\sigma$. The map \bar{f} sends each vertex v_σ of X' to the point $y_\sigma \in M$. This map then extends to a map \bar{f} on each simplex τ of X' via Riemannian barycentric coordinates, see e.g. [BK81, §8]: If $x \in \tau$ has barycentric coordinates $(\lambda_1, \dots, \lambda_n)$ in τ and v_1, \dots, v_n are the vertices of τ then

$$\bar{f}(x) := \text{Center}\left(\sum_{i=1}^n \lambda_i \delta_{\bar{f}(v_i)}\right),$$

is where δ_y denotes the probability measure supported on the point $y \in M$.

In particular, if τ is a simplex of X' with the vertex $x_i = \nu(\tau)$ that is also a vertex of X , then for every vertex v of τ , $\bar{f}(v)$ belongs to $B(x_i, r)$ and, hence, by convexity of the latter, $\bar{f}(\tau)$ is contained in $B(x_i, r)$ as well.

In order to construct the homotopy

$$H : M \times [0, 1] \rightarrow M, H(y, 0) = \bar{f} \circ f, H(y, 1) = y, y \in M$$

we use the straight-line homotopy via minimizing geodesics in (M, g) (here we are again using the fact that the distance between y and $\bar{f} \circ f(y)$ is less than the convexity radius ϵ of (M, g)). The map H is Lipschitz because of uniform Lipschitz dependence of the minimal geodesic $y_1 y_2$ in M on the endpoints y_1, y_2 , provided that $d(y_1, y_2) < \epsilon$. The composition $h \circ \bar{h} : X \rightarrow X$ sends each simplex τ of X' into the star $St(x_i, X)$ of the vertex $x_i = \nu(\tau)$ in the complex X . For any two points $x, y \in St(x_i, X)$ we define the broken geodesic path $p_{x,y}(t)$ from x to y as the concatenation of the Euclidean geodesic segments

$$xx_i \star x_i y$$

parameterized by the unit interval $[0, 1]$ with the constant speed. Lastly, define the homotopy

$$\bar{H} : (x, t) \mapsto p_{x, h(x)}(t)$$

where $h(x) = f \circ \bar{f}(x)$, $x \in \tau \subset St(x_i, X)$. □

DEFINITION 3.38. We will refer to the complex together with the map, (X, f) , as a *Lipschitz simplicial model* of the Riemannian manifold (M, g) .

The main application of bounded geometry triangulations (or, Lipschitz simplicial models) in this book comes in the form of *systolic inequalities* which we describe below.

Let M be a Riemannian manifold. The k -*systole* $sys_k(M)$ of M is defined as the infimum of volumes of homologically non-trivial k -cycles in M . In the following proof, by abusing the terminology, we will conflate singular k -chains

$$S = \sum_{i=1}^N a_i \sigma_i$$

(where $a_i \in \mathbb{Z} \setminus \{0\}$ and σ_i 's are singular simplices) and their *support sets* in X , i.e. unions of images of the singular k -simplices σ_i .

THEOREM 3.39. *Every Riemannian manifold M of bounded geometry has positive k -systoles for all k .*

PROOF. Given M we take either a bounded geometry triangulation $\mathcal{T} = (X, \tau)$ of M (if it exists) or a Lipschitz simplicial model (X, f) . Since the map f is a bi-Lipschitz homotopy equivalence, it suffices to prove positivity of k -systoles for X . The key to the proof is the following *Deformation Theorem* of Federer and Fleming, which first appeared in their work on the Plateau Problem [FF60, §5]. Another proof of this fundamental fact can be found in Federer's book [Fed69, 4.2.9]; an especially readable proof is given in [ECH⁺92, Theorem 10.3.3]. Suppose that Δ^n is the standard n -simplex. For each interior point $x \in \Delta^n$ we define the *radial projection* $p_x : \Delta^n \setminus \{x\} \rightarrow \partial\Delta^n$. We will need:

THEOREM 3.40 (Deformation Theorem). *Suppose that S is a singular k -chain in Δ , $k < n$. Then for almost every point $x \in \Delta^n$, the k -volume of the chain $p_x(S)$ in $\partial\Delta$ does not exceed $CVol_k(S)$, where the constant C depends only on the dimension n of the simplex.*

We will refer to the projections p_x satisfying the conclusion of this theorem as *Federer-Flemming projections*.

Suppose now that S is a singular k -cycle in a bounded geometry D -dimensional simplicial complex X . In each n -simplex Δ^n in X whose dimension is greater than k , we apply a Federer–Flemming projection p_x to S . By combining these projections, we obtain a chain S_1 in $X^{(n-1)}$, which is homologous to S and whose volume is at most $CVol_k(S)$. After repeating the process at most $D - k$ times, we obtain a k -cycle S' in the k -skeleton $X^{(k)}$ (homologous to S); the volume of S' is at most $C^{D-k}Vol_k(S)$. Let V_k denote the volume of the standard k -simplex. If $Vol_k(S')$ is less than V_k , then S' cannot cover any k -simplex in X . Therefore, for each k -simplex Δ^k in X we apply a radial projection p_x to S' from any point x which does not belong to S' (at this stage, we no longer care about the volume of the image). The result is a k -cycle T in the $k - 1$ -skeleton $X^{(k-1)}$ of X , which is still homologous to S . However, $k - 1$ -dimensional simplicial complexes have zero k th homology groups, which means that T (and, hence, S) is homologically trivial. Therefore, assuming that $S \in Z_k(M)$ was homologically non-trivial, we obtain a lower bound on its volume:

$$Vol_k(S) \geq L^{-k} C^{k-D} V_k.$$

□

3.9. Harmonic functions

For the detailed discussion of the material in this section we refer the reader to [Li12] and [SY94].

Let M be a Riemannian manifold. Given a smooth function $f : M \rightarrow \mathbb{R}$, we define the *energy* of f as the integral

$$E(f) = \int_M |df|^2 dV = \int_M |\nabla f|^2 dV.$$

Here the gradient vector field ∇f is obtained by dualizing the differential 1-form df using the Riemannian metric on M . Note that energy is defined even if f only belongs to the Sobolev space $W_{loc}^{1,2}(M)$ of functions differentiable a.e. on M with locally square-integrable partial derivatives.

THEOREM 3.41 (Lower semicontinuity of the energy functional). *Let (f_i) be a sequence of functions in $W_{loc}^{1,2}(M)$ which converges (in $W_{loc}^{1,2}(M)$) to a function f . Then*

$$E(f) \leq \liminf_{i \rightarrow \infty} E(f_i).$$

DEFINITION 3.42. A function $h \in W_{loc}^{1,2}$ is called *harmonic* if it is *locally energy-minimizing*: For every point $p \in M$ and a small metric ball $B = B(p, r) \subset M$,

$$E(h|_B) \leq E(u), \quad \forall u : \bar{B} \rightarrow \mathbb{R}, u|_{\partial B} = h|_{\partial B}.$$

Equivalently, for every relatively compact open subset $\Omega \subset M$ with smooth boundary

$$E(h|_B) \leq E(u), \quad \forall u : \bar{\Omega} \rightarrow \mathbb{R}, u|_{\partial\Omega} = h|_{\partial\Omega}.$$

It turns out that harmonic functions h on M are automatically smooth and, moreover, satisfy the equation $\Delta h = 0$, where Δ is the *Laplace–Beltrami operator* on M :

$$\Delta u = \operatorname{div} \nabla u.$$

In local coordinates (assuming that M is n -dimensional)

$$\Delta u = \sum_{i,j=1}^n \frac{1}{\sqrt{|g|}} \frac{\partial}{\partial x_i} \left(g^{ij} \sqrt{|g|} \frac{\partial u}{\partial x_j} \right).$$

In terms of the Levi–Civita connection ∇ on M ,

$$\Delta(u) = \operatorname{Trace}(H(u)), \quad H(u)(X, Y) = \nabla_X \nabla_Y(u) - \nabla_{\nabla_X Y}(u),$$

where X, Y are vector fields on M . In local coordinates, setting

$$H_{ij} = H(u) \left(\frac{\partial}{\partial x_i}, \frac{\partial}{\partial x_j} \right),$$

we have

$$\operatorname{Trace}(H) = \sum_{i,j=1}^n g^{ij} H_{ij}.$$

If $M = \mathbb{R}^n$ with the flat metric, then Δ is the usual Laplace operator:

$$\Delta u = \sum_{i=1}^n \frac{\partial^2}{\partial x_i^2} u.$$

EXERCISE 3.43. Work out the formula for Δu in the case of a conformally-Euclidean metric g on an open subset of \mathbb{R}^n . Conclude that harmonicity with respect to g is equivalent to harmonicity with respect to the flat metric.

Lastly, if we use the normal (geodesic) coordinates on a Riemannian manifold then

$$\Delta u(p) = \sum_{i=1}^n \frac{\partial^2}{\partial x_i^2} u(0).$$

A function u on M is called *subharmonic* if

$$\Delta u \geq 0.$$

EXAMPLE 3.44. If $n = 1$ and $M = \mathbb{R}$ then a function is harmonic if and only if it is linear, and is subharmonic if and only if it is convex.

EXERCISE 3.45. Suppose that $h : M \rightarrow \mathbb{R}$ is a harmonic function and $f : \mathbb{R} \rightarrow \mathbb{R}$ is a smooth convex function. Then the composition $u = f \circ h$ is subharmonic. Hint: Verify that $\Delta u(p) \geq 0$ for every $p \in M$ using normal coordinates on M defined via the exponential map $\exp_p : T_p M \rightarrow M$. This reduces the problem to a Euclidean computation.

THEOREM 3.46 (Maximum Principle). *Suppose that M is connected, $\Omega \subset M$ is a relatively compact subset with smooth boundary and $h : M \rightarrow \mathbb{R}$ is a harmonic function. Then $h|_{\overline{\Omega}}$ attains maximum on the boundary of Ω and, moreover, if $h|_{\Omega}$ attains its maximum at a point of Ω , then h is constant.*

COROLLARY 3.47. *Let $h_i : M \rightarrow \mathbb{R}, i = 1, 2$ be two harmonic functions such that $h_1 \leq h_2$. Then either $h_1 = h_2$ or $h_1 < h_2$.*

PROOF. The difference $h = h_1 - h_2 \leq 0$ is also a harmonic function on M . Suppose that the subset $A = \{h_1(x) = h_2(x)\}$ is non-empty. Then for every relatively compact subset with smooth boundary $\Omega \subset M$ and $A \cap \Omega \neq \emptyset$, the maximum of $h|_{\Omega}$ is attained on

$$A \cap \Omega.$$

Therefore, $h|_{\Omega}$ is identically zero. Taking an exhaustion of M by subsets Ω as above, we conclude that h vanishes on the entire M . \square

THEOREM 3.48 (Li-Schoen's Mean Value Inequality for subharmonic functions). *Suppose that Ricci curvature of the Riemannian n -manifold M is bounded below by a constant r . Then there exists a function $C(n, r, R)$ such that for every nonnegative subharmonic function $u : M \rightarrow \mathbb{R}$, and normal ball $B(p, R)$, we have*

$$u^2(p) \leq C(n, r, R) \int_{B(p, R)} u^2 dV.$$

As a corollary, one obtains a similar mean value inequality for harmonic functions (without any positivity assumption):

COROLLARY 3.49. *Suppose that M, p and R satisfy the hypothesis of the previous theorem. Then for every harmonic function $h : M \rightarrow \mathbb{R}$ we have*

$$h^2(p) \leq \sqrt{C(n, r, R)} \int_{B(p, R)} h^2 dV.$$

PROOF. The composition $u = h^2$ of h with the convex function $x \mapsto x^2$ is subharmonic. Therefore,

$$u^2(p) \leq C(n, r, R) \int_{B(p, R)} u^2 dV.$$

Thus,

$$h^4(p) \leq C(n, r, R) \int_{B(p, R)} u^2 dV \leq C(n, r, R) \left(\int_{B(p, R)} u dV \right)^2,$$

which implies the inequality

$$h^2(p) \leq \sqrt{C(n, r, R)} \int_{B(p, R)} h^2 dV. \quad \square$$

THEOREM 3.50 (Yau’s gradient estimate). *Suppose that M^n is a complete n -dimensional Riemannian manifold with Ricci curvature $\geq a$. Then for every positive harmonic function h on M , every $x \in M$ with $\text{InjRad}(x) \geq \epsilon$,*

$$|\nabla h(x)| \leq C(\epsilon, n)h(x).$$

The following two theorems are a part of the so called *elliptic regularity theory* for solutions of second order elliptic PDEs, see e.g. [GT83].

THEOREM 3.51 (Derivative bounds). *For every harmonic function h on a manifold of bounded geometry, there exists $L(r)$ such that for every $x \in M$ and every harmonic function $h : M \rightarrow \mathbb{R}$, whose restriction to the ball $B(x, r)$ takes values in $[0, 1]$, we have*

$$|\nabla|\nabla h(x)|^2| \leq L(r),$$

as long as $\nabla h(x) \neq 0$.

Note that similar estimates hold for higher-order derivatives of harmonic functions; we will only need a bound on the second derivatives.

THEOREM 3.52 (Compactness Property). *Suppose that (f_i) is a sequence of harmonic functions on M so that there exists $p \in M$ for which the sequence $(f_i(p))$ is bounded. Then the family of functions (f_i) is precompact in $W_{loc}^{1,2}(M)$. Furthermore, every limit of a subsequence in (f_i) is a harmonic function.*

We will use these properties of harmonic functions in Chapter 21, in the proof of Stallings Theorem on ends of groups via harmonic functions. Since in the proof it suffices to work with 2-dimensional Riemannian manifolds (Riemann surfaces), the properties of harmonic functions we are using follow from more elementary properties of harmonic functions of one complex variable (real parts of holomorphic functions). For instance, the upper bounds on the first and second derivatives and Compactness Property follow from Cauchy’s integral formula; the maximum principle for harmonic functions follows from the maximum principle for holomorphic functions. Similarly, Corollary 3.49 follows from [Poisson’s Integral Formula].

3.10. Spectral interpretation of the Cheeger constant

Let M be a complete connected Riemannian manifold of infinite volume. Then the vector space $V = L^2(M) \cap C^\infty(M)$ contains no non-zero constant functions. We let Δ_M denote the restriction of the Laplace-Beltrami operator to the space V and let $\lambda_1(M)$ be the lowest eigenvalue of Δ_M . The number $\lambda_1(M)$ is also known as the *spectral gap* of the manifold M . The eigenvalue $\lambda_1(M)$ can be computed as

$$(3.6) \quad \inf \left\{ \frac{\int_M |\nabla f|^2}{\int_M f^2} \mid f : M \rightarrow \mathbb{R} \text{ is smooth, non-zero, with compact support} \right\},$$

see [CY75] or Chapter I of [SY94]. J. Cheeger proved in [Che70] that

$$\lambda_1(M) \geq \frac{1}{4}h^2(M),$$

where $h(M)$ is the Cheeger constant of M . Even though Cheeger’s original result was formulated for compact manifolds, his argument works for non-compact manifolds as well, see [SY94]. Cheeger’s inequality is complemented by the following

inequality due to P. Buser (see [Bus82], or [SY94]) which holds for all complete Riemannian manifolds whose Ricci curvature is bounded below by some $a \in \mathbb{R}$:

$$\lambda_1(M) \leq \alpha h(M) + \beta h^2(M),$$

for some $\alpha = \alpha(a), \beta = \beta(a)$. Combined, Cheeger and Buser inequalities imply

THEOREM 3.53. $h(M) = 0 \iff \lambda_1(M) = 0$.

3.11. Comparison geometry

In the setting of general metric spaces it is still possible to define a notion of (upper and lower bound for the) sectional curvature, which, moreover, coincide with the standard ones for Riemannian manifolds. This is done by comparing geodesic triangles in a metric space to geodesic triangles in a *model space of constant curvature*. In what follows, we only discuss the metric definition of upper bound for the sectional curvature, the lower bound case is similar (see e.g. [BBI01]) but will not be used in this book.

3.11.1. Alexandrov curvature and $CAT(\kappa)$ spaces. For a real number $\kappa \in \mathbb{R}$, we denote by X_κ the *model surface of constant curvature κ* . If $\kappa = 0$ then X_κ is the Euclidean plane. If $\kappa < 0$ then X_κ will be discussed in detail in Chapter 4, it is the *upper half-plane with the rescaled hyperbolic metric*:

$$X_\kappa = \left(\mathbf{U}^2, |\kappa|^{-1} \frac{dx^2 + dy^2}{y^2} \right).$$

If $\kappa > 0$ then X_κ is the 2-dimensional sphere $S\left(0, \frac{1}{\sqrt{\kappa}}\right)$ in \mathbb{R}^3 with the Riemannian metric induced from \mathbb{R}^3 .

Let X be a geodesic metric space, and let Δ be a geodesic triangle in X . Given $\kappa > 0$ we say that Δ is κ -compatible if its perimeter is at most $\frac{2\pi}{\sqrt{\kappa}}$. By default, every triangle is κ -compatible for $\kappa \leq 0$.

We will prove later on (see Section 4.11) the following:

LEMMA 3.54. *Let $\kappa \in \mathbb{R}$ and let $a \leq b \leq c$ be three numbers such that $c \leq a + b$ and $a + b + c < \frac{2\pi}{\sqrt{\kappa}}$ if $\kappa > 0$. Then there exists a geodesic triangle in X_κ with side-lengths a, b and c , and this triangle is unique up to congruence.*

Therefore, for every $\kappa \in \mathbb{R}$ and every κ -compatible triangle $\Delta = \Delta(A, B, C) \subset X$ with vertices $A, B, C \in X$ and lengths a, b, c of the opposite sides, there exists a triangle (unique, up to congruence)

$$\tilde{\Delta}(\tilde{A}, \tilde{B}, \tilde{C}) \subset X_\kappa$$

with the side-lengths a, b, c . The triangle $\tilde{\Delta}(\tilde{A}, \tilde{B}, \tilde{C})$ is called the κ -comparison triangle or a κ -Alexandrov triangle.

For every point P on, say, the side AB of Δ , we define the κ -comparison point $\tilde{P} \in \tilde{A}\tilde{B}$, such that

$$d(A, P) = d(\tilde{A}, \tilde{P}).$$

DEFINITION 3.55. We say that *the triangle Δ is $CAT(\kappa)$* if it is κ -compatible and for every pair of points P and Q on the triangle, their κ -comparison points \tilde{P}, \tilde{Q} satisfy

$$\text{dist}_{X_\kappa}(\tilde{P}, \tilde{Q}) \geq \text{dist}_X(P, Q) .$$

DEFINITION 3.56. (1) A $CAT(\kappa)$ -domain in X is an open convex set $U \subset X$, and such that all the geodesic triangles entirely contained in U are $CAT(\kappa)$.

(2) The space X has *Alexandrov curvature at most κ* if it is covered by $CAT(\kappa)$ -domains.

Note that a $CAT(\kappa)$ -domain U for $\kappa > 0$ must have diameter strictly less than $\frac{\pi}{\sqrt{\kappa}}$. Otherwise, one can construct geodesic triangles in U with two equal edges and the third reduced to a point, with perimeter $\geq \frac{2\pi}{\sqrt{\kappa}}$.

The point of Definition 3.56 is that it applies to non-Riemannian metric spaces where such notions as tangent vectors, Riemannian metric, curvature tensor cannot be defined, while one can still talk about curvature being bounded from above by κ .

PROPOSITION 3.57. *Let X be a Riemannian manifold. Its Alexandrov curvature is at most κ if and only if its sectional curvature in every point is $\leq \kappa$.*

PROOF. The “if” implication follows from the Rauch-Toponogov comparison theorem (see [dC92, Proposition 2.5]). For the “only if” implication we refer to [Rin61] or to [GHL04, Chapter III]. \square

DEFINITION 3.58. A metric space X is called a $CAT(\kappa)$ -space if the entire X is a $CAT(\kappa)$ -domain. We will use the definition only for $\kappa \leq 0$. A metric space X is said to be a $CAT(-\infty)$ -space if X is a $CAT(\kappa)$ -space for every κ .

Note that for the moment we do not assume X to be metrically complete. This is because there are naturally occurring incomplete $CAT(0)$ spaces, called *Euclidean buildings*, which, nevertheless, are *geodesically complete* (every geodesic segment is contained in a complete geodesic).

Clearly, every Hilbert space is $CAT(0)$.

EXERCISE 3.59. Let X be a simplicial tree with a path-metric d . Show that (X, d) is $CAT(-\infty)$.

This exercise leads to the following definition:

DEFINITION 3.60. A geodesic metric space X such that for every geodesic triangle in X with the sides xy, yz, zx , the side xy is contained in the union $yz \cup zx$, is called a *real tree*.

EXERCISE 3.61. 1. Show that a geodesic metric space X is a real tree if and only if X is $CAT(-\infty)$.

2. Consider the following metric space: Take the union of the x -axis in \mathbb{R}^2 and all vertical lines $\{x = q\}$, where q 's are rational numbers. Equip X with the path-metric induced from \mathbb{R}^2 . Show that X is a real tree.

We note that real trees are also called \mathbb{R} -trees or *metric trees* in the literature. A real tree is called *complete* if it is complete as a metric space. While the simplest examples of real trees are given by simplicial trees equipped with their standard path-metrics, we will see in Chapter 11 that other real trees also arise naturally in the Geometric Group Theory. We refer to Section 11.2 for further discussion of real trees.

EXERCISE 3.62. Let Γ be a connected metric graph with the path-metric. Show that Γ is a $CAT(1)$ if and only if Γ contains no circuits of length $< 2\pi$.

More interesting examples come from polygonal complexes. Their origins lie in two areas of mathematics, going back to 1940s and 1950s:

- The *small cancellation theory*, which is an area of the combinatorial group theory.
- Alexandrov's theory of spaces of curvature bounded from above.

Suppose that X is a connected almost regular 2-dimensional cell complex. We equip X with a path-metric where each 2-face is isometric to a constant curvature κ 2-dimensional polygon with unit edges. This defines structure of a metric graph on the link $Lk(v)$ of each vertex v of X , where each corner c of each 2-face F determines an edge of $Lk(v)$ whose length is the angle of F at c . We refer the reader to [BH99, Ba195] for proofs of the following theorem:

THEOREM 3.63. *The metric space X has Alexandrov curvature $\leq \kappa$ if and only if each connected component of the link $Lk(v)$ of each vertex v of X is a $CAT(1)$ space. (See [BH99], Theorem 5.20, Ch. II.5.)*

To make this theorem more concrete, we assume that each 2-dimensional face of X has n edges and for each vertex $v \in X$ the combinatorial length of the shortest circuit in the link $Lk(v)$ is at least m . Then Theorem 3.63 implies:

- COROLLARY 3.64.**
1. *Suppose that $\kappa = 0$, $n \geq 3$ and $m \geq 6$, or $n \geq 4$ and $m \geq 4$, or $n \geq 6$ and $m \geq 3$. Then X has Alexandrov curvature ≤ 0 .*
 2. *Suppose that $\kappa = -1$, $n \geq 3$ and $m \geq 7$, or $n \geq 4$ and $m \geq 5$, or $n \geq 6$ and $m \geq 4$, or $n \geq 7$ and $m \geq 3$. Then X has Alexandrov curvature ≤ -1 .*

Yes another class of examples of $CAT(0)$ spaces comes from *cube complexes*. The n -dimensional *cube* is the product of intervals

$$I^n = [0, 1]^n.$$

DEFINITION 3.65. An almost regular cell complex where each cell is isomorphic to a cube is called a *cube complex*.

We will always equip cube complexes with the *standard path-metric* where each n -dimensional face is isometric to the Euclidean cube I^n .

DEFINITION 3.66. A simplicial complex Y is called a *flag-complex* if whenever Y contains a 1-dimensional complex Z isomorphic to the 1-skeleton of the n -simplex Δ^n , the complex Y also contains a subcomplex W isomorphic to Δ^n such that W^1 equals Z .

THEOREM 3.67. *A simply-connected cube complex X is a $CAT(0)$ space if and only if the link of every vertex in X is a flag-complex.*

In the case of spaces of non-positive curvature one can connect local and global curvature bounds:

THEOREM 3.68 (Cartan-Hadamard Theorem). *If X is a simply connected complete metric space with Alexandrov curvature at most κ for some $\kappa \leq 0$, then X is a $CAT(\kappa)$ -space.*

We refer the reader to [Bal95] and [BH99] for proofs of this theorem, and a detailed discussion of $CAT(\kappa)$ -spaces, with $\kappa \leq 0$.

DEFINITION 3.69. Simply-connected complete Riemannian manifolds of sectional curvature ≤ 0 are called *Hadamard manifolds*. Thus, every Hadamard manifold is a $CAT(0)$ space.

An important property of $CAT(0)$ -spaces is *convexity of the distance function*. Suppose that X is a geodesic metric space. A function $F : X \times X \rightarrow \mathbb{R}$ is said to be *convex* if for every pair of geodesics $\alpha(s), \beta(s)$ in X (which are parameterized with constant, but not necessarily unit, speed), the function

$$f(s) = F(\alpha(s), \beta(s))$$

is a convex function of one variable. Thus, the distance function dist of X is convex, whenever for every pair of geodesics a_0a_1 and b_0b_1 in X , the points $a_s \in a_0a_1$ and $b_s \in b_0b_1$ such that $\text{dist}(a_0, a_s) = s \text{dist}(a_0, a_1)$ and $\text{dist}(b_0, b_s) = s \text{dist}(b_0, b_1)$, satisfy

$$(3.7) \quad \text{dist}(a_s, b_s) \leq (1-s)\text{dist}(a_0, b_0) + s \text{dist}(a_1, b_1).$$

Note that in the case of a normed vector space X , a function $f : X \times X \rightarrow \mathbb{R}$ is convex if and only if the epi-graph

$$\{(x, y, t) \in X^2 \times \mathbb{R} : f(x, y) \geq t\}$$

is convex.

PROPOSITION 3.70. *If a geodesic metric space X is $CAT(0)$ then the distance on X is convex.*

PROOF. Consider two geodesics a_0b_0 and a_1b_1 in X . On the geodesic a_0b_1 consider the point c_s such that $\text{dist}(a_0, c_s) = s \text{dist}(a_0, b_1)$. The fact that the triangle with edges a_0a_1 , a_0b_1 and a_1b_1 is $CAT(0)$ and the Thales theorem in \mathbb{R}^2 , imply that $\text{dist}(a_s, c_s) \leq s \text{dist}(a_1, b_1)$. The same argument applied to the triangle with edges a_0b_1 , a_0b_0 , b_0b_1 , implies that $\text{dist}(c_s, b_s) \leq (1-s)\text{dist}(a_0, b_0)$. The inequality (3.7) follows from

$$\text{dist}(a_s, b_s) \leq \text{dist}(a_s, c_s) + \text{dist}(c_s, b_s).$$

□

REMARK 3.71. The converse to this proposition is not true in general. Indeed, every strictly convex normed vector space has convex distance function but only Hilbert spaces (among normed vector spaces) are $CAT(0)$. See also [BH99], Example 1.18, page 169.

COROLLARY 3.72. *Every $CAT(0)$ -space X is uniquely geodesic, i.e. for any two points $p, q \in X$, the (arc-length parameterized) geodesic from p to q is unique.*

PROOF. It suffices to apply the inequality (3.7) to a geodesic bigon, that is, in the special case when $a_0 = b_0$ and $a_1 = b_1$. □




FIGURE 3.1. Argument for convexity of the distance function.

3.11.2. Cartan's fixed point theorem. Let X be a metric space and $A \subset X$ be a subset. Define the function

$$\rho(x) = \rho_A(x) = \sup_{a \in A} d^2(x, a).$$

PROPOSITION 3.73. *Let X be a complete CAT(0) space. Then for every bounded subset $A \subset X$, the function $\rho = \rho_A$ attains unique minimum in X .*

PROOF. Consider a sequence (x_n) in X such that

$$\lim_{n \rightarrow \infty} \rho(x_n) = r = \inf_{x \in X} \rho(x).$$

We claim that the sequence (x_n) is Cauchy. Given $\epsilon > 0$ let $x = x_i, x' = x_j$ be points in this sequence such that

$$r \leq \rho(x) < r + \epsilon, \quad r \leq \rho(x') < r + \epsilon.$$

Let p be the midpoint of $xx' \subset X$; hence, $r \leq \rho(p)$. Let $a \in A$ be such that

$$\rho(p) - \epsilon < d^2(p, a).$$

Consider the Euclidean comparison triangle $\tilde{T} = T(\tilde{x}, \tilde{x}', \tilde{a})$ for the triangle $T(x, x', a)$. In the Euclidean plane we have (by the parallelogram identity (2.1)):

$$d^2(\tilde{x}, \tilde{x}') + 4d^2(\tilde{a}, \tilde{p}) = 2(d^2(\tilde{a}, \tilde{x}) + d^2(\tilde{a}, \tilde{x}')).$$

Applying the comparison inequality for the triangles T and \tilde{T} , we obtain:

$$d(a, p) \leq d(\tilde{a}, \tilde{p}).$$

Thus:

$$d(x, x')^2 + 4(r - \epsilon) < d^2(x, x') + 4d^2(a, p) \leq 2(d^2(a, x) + d^2(a, x')) <$$

$$2(\rho(x) + \rho(x')) < 4r + 4\epsilon.$$

It follows that

$$d(x, x')^2 < 8\epsilon$$

and, therefore, the sequence (x_n) is Cauchy. By completeness of X , the function ρ attains minimum in X ; the same Cauchy argument implies that the point of minimum is unique. \square

As a corollary, we obtain a fixed-point theorem for isometric group actions on complete $CAT(0)$ spaces, which was first proven by E. Cartan in the context of Riemannian manifolds of nonpositive curvature and then extended by J. Tits to geodesic metric spaces with convex distance function:

THEOREM 3.74. *Let X be a complete $CAT(0)$ metric space and $G < \text{Isom}(X)$ be a subgroup which has bounded orbits: One (equivalently every) subset of the form*

$$G \cdot x = \{g(x) : g \in G\}$$

is bounded. Then G fixes a point in X .

PROOF. Let A denote a (bounded) orbit of G in X and let ρ_A be the corresponding function on X . Then, by the uniqueness of the minimum point m of ρ_A , the group G will fix m . \square

COROLLARY 3.75. *1. Every finite group action on a complete $CAT(0)$ space has a fixed point. For instance, every action of a finite group on a complete real tree or on a Hilbert space fixes a point.*

2. If G is a compact group acting isometrically and continuously on a Hilbert space \mathcal{H} , then G fixes a point in \mathcal{H} .

EXERCISE 3.76. Prove that this corollary holds for all real trees T , not necessarily complete ones. Hint: For a finite subset $F \subset T$ consider its *span* T_F , i.e. the union of all geodesic segments connecting points of F . Show that T_F is isometric to a complete metric tree and is G -invariant if F was. In fact, T_F is isometric to a finite metric simplicial complex (which, as a simplicial complex, is isomorphic to a finite simplicial tree).

DEFINITION 3.77. A group G is said to have the *Property FA* if for every isometric action $G \curvearrowright T$ on a complete real tree T , G fixes a point in T .

Thus, all finite groups have the Property FA.

3.11.3. Ideal boundary, horoballs and horospheres. In this section we discuss the notion of the *ideal boundary* of a metric space. This is a particularly useful concept when the metric space is $CAT(0)$, and it generalizes the concept introduced for non-positively curved simply connected Riemannian manifolds by P. Eberlein and B. O'Neill in [EO73, Section 1].

Let X be a geodesic metric space. Two geodesic rays ρ_1 and ρ_2 in X are called *asymptotic* if they are at finite Hausdorff distance; equivalently if the function $t \mapsto \text{dist}(\rho_1(t), \rho_2(t))$ is bounded on $[0, \infty)$.

Clearly, being asymptotic is an equivalence relation on the set of geodesic rays in X .

DEFINITION 3.78. The *ideal boundary* of a metric space X is the collection of equivalence classes of geodesic rays. It is usually denoted either by $\partial_\infty X$ or by $X(\infty)$.

An equivalence class $\xi \in \partial_\infty X$ is called an *ideal point* or *point at infinity* of X , and the fact that a geodesic ray ρ is contained in this class is sometimes expressed by the equality $\rho(\infty) = \xi$. When a geodesic ray ρ represents an equivalence class $\xi \in \partial_\infty X$, the ray ρ is said to be *asymptotic to* ξ .

The space of geodesic rays in X has a natural compact-open topology, or, equivalently, topology of uniform convergence on compacts (recall that we regard geodesic rays as maps from $[0, \infty)$ to X). Thus, we topologize $\partial_\infty X$ by giving it the quotient topology τ .

EXERCISE 3.79. Every isometry $g : X \rightarrow X$ induces a homeomorphism $g_\infty : \partial_\infty X \rightarrow \partial_\infty X$.

This exercise explains why we consider rays emanating from different points of X : Otherwise, most isometries of X would not act on $\partial_\infty X$.

Convention. From now on, in this section, we assume that X is a complete $CAT(0)$ metric space.

LEMMA 3.80. *If X is locally compact then for every point $x \in X$ and every point $\xi \in \partial_\infty X$ there exists a unique geodesic ray ρ with $\rho(0) = x$ and $\rho(\infty) = \xi$. We will also use the notation $x\xi$ for the ray ρ .*

PROOF. Let $r : [0, \infty) \rightarrow X$ be a geodesic ray with $r(\infty) = \xi$. For every $n \in \mathbb{N}$, according to Corollary 3.72, there exists a unique geodesic \mathfrak{g}_n joining x and $r(n)$. The convexity of the distance function implies that every \mathfrak{g}_n is at Hausdorff distance $\text{dist}(x, r(0))$ from the segment of r between $r(0)$ and $r(n)$.

By the Arzela-Ascoli Theorem, a subsequence \mathfrak{g}_{n_k} of geodesic segments converges in the compact-open topology to a geodesic ray ρ with $\rho(0) = x$. Moreover, ρ is at Hausdorff distance $\text{dist}(x, r(0))$ from r .

Assume that ρ_1 and ρ_2 are two asymptotic geodesic rays with $\rho_1(0) = \rho_2(0) = x$. Let M be such that $\text{dist}(\rho_1(t), \rho_2(t)) \leq M$, for every $t \geq 0$. Consider $t \in [0, \infty)$, and $\varepsilon > 0$ arbitrarily small. Convexity of the distance function implies that

$$\text{dist}(\rho_1(t), \rho_2(t)) \leq \varepsilon \text{dist}(\rho_1(t/\varepsilon), \rho_2(t/\varepsilon)) \leq \varepsilon M.$$

It follows that $\text{dist}(\rho_1(t), \rho_2(t)) = 0$ and, hence, $\rho_1 = \rho_2$. □

In particular, for a fixed point $p \in X$ one can identify the set $\overline{X} := X \sqcup \partial_\infty X$ with the set of geodesic segments and rays with initial point p . In what follows, we will equip \overline{X} with the topology induced from the compact-open topology on the space of these segments and rays.

- EXERCISE 3.81. (1) Prove that the embedding $X \hookrightarrow \overline{X}$ is a homeomorphism to its image.
- (2) Prove that the topology on \overline{X} is independent of the chosen basepoint p . In other words, given p and q two points in X , the map $[p, x] \mapsto [q, x]$ (with $x \in \overline{X}$) is a homeomorphism.
- (3) In the special case when X is a Hadamard manifold, show that for every point $p \in X$, the ideal boundary $\partial_\infty X$ is homeomorphic to the unit sphere S in the tangent space $T_p M$ via the map

$$v \in S \subset T_p M \rightarrow \exp_p(\mathbb{R}_+ v) \in \partial_\infty X.$$

An immediate consequence of the Arzela–Ascoli Theorem is that \overline{X} is compact, provided that X is locally compact.

Consider a geodesic ray $r : [0, \infty) \rightarrow X$, and an arbitrary point $x \in X$. The function $t \mapsto \text{dist}(x, r(t)) - t$ is decreasing (due to the triangle inequality) and bounded from below by $-\text{dist}(x, r(0))$. Therefore, there exists a limit

$$(3.8) \quad b_r(x) := \lim_{t \rightarrow \infty} [\text{dist}(x, r(t)) - t] .$$

DEFINITION 3.82. The function $b_r : X \rightarrow \mathbb{R}$ thus defined, is called the *Busemann function for the ray r* .

For a proof of the next result see e.g. [Bal95], Chapter 2, Proposition 2.5.

THEOREM 3.83. *If r_1 and r_2 are two asymptotic rays then $b_{r_1} - b_{r_2}$ is a constant function.*

In particular, it follows that the collections of sublevel sets and the level sets of a Busemann function do not depend on the ray r , but only on the point at infinity that r represents.

EXERCISE 3.84. Show that b_r is linear with slope -1 along the ray r . In particular,

$$\lim_{t \rightarrow \infty} b_r(t) = -\infty .$$

DEFINITION 3.85. A sublevel set of a Busemann function, $b_r^{-1}(-\infty, a]$ is called a (*closed*) *horoball with center $\xi = r(\infty)$* ; we denote such horoballs as $\overline{B}(\xi)$ or $\overline{B}(r)$. A level set $b_r^{-1}(a)$ of a Busemann function is called a *horosphere with center ξ* , it is denoted $\Sigma(\xi)$. In the case when X is 2-dimensional, horospheres are called *horocycles*. Lastly, an open sublevel set $b_r^{-1}(-\infty, a)$ is called an *open horoball with center $\xi = r(\infty)$* , and denoted $B(\xi)$ or $B(r)$.

Informally, one can think informally of horoballs $B(\xi)$ and horospheres $\Sigma(\xi)$ as metric balls and metric spheres of infinite radii in X , centered at ξ , whose radii are determined by the choice of the Busemann function b_r (which is determined only up to a constant) and by the choice of the value a of b_r .

LEMMA 3.86. *Let r be a geodesic ray and let B be the open horoball $b_r^{-1}(-\infty, 0)$. Then $B = \bigcup_{t \geq 0} B(r(t), t)$.*

PROOF. Indeed, if for a point x ,

$$b_r(x) = \lim_{t \rightarrow \infty} [\text{dist}(x, r(t)) - t] < 0,$$

then, for some sufficiently large t , $\text{dist}(x, r(t)) - t < 0$. Whence $x \in B(r(t), t)$.

Conversely, suppose that $x \in X$ is such that for some $s \geq 0$,

$$\text{dist}(x, r(s)) - s = \delta_s < 0.$$

Then, because the function $t \mapsto \text{dist}(x, r(t)) - t$ is decreasing, it follows that for every $t \geq s$,

$$\text{dist}(x, r(t)) - t \leq \delta_s .$$

Whence, $b_r(x) \leq \delta_s < 0$. □

LEMMA 3.87. *Let X be a CAT(0) space. Then every Busemann function on X is convex and 1-Lipschitz.*

PROOF. Recall that the distance function on any metric space is 1-Lipschitz. Since Busemann functions are limits of normalized distance functions, it follows that Busemann functions are 1-Lipschitz as well. (This part does not require the $CAT(0)$ assumption.) Similarly, since the distance function is convex, Busemann functions are also convex as limits of normalized distance functions. \square

Furthermore, if X is a Hadamard manifold, then every Busemann function b_r is smooth, with gradient of constant norm 1, see [BGS85].

LEMMA 3.88. *Assume that X is a complete $CAT(0)$ space. Then:*

- *Open and closed horoballs in X are convex sets.*
- *A closed horoball is the closure of an open horoball.*

PROOF. The first property follows immediately from the convexity of Busemann functions. Let $f = b_r$ be a Busemann function. Consider the closed horoball

$$\overline{B} = \{x : f(x) \leq t\}.$$

Since this horoball is a closed subset of X , it contains the closure of the open horoball

$$B = \{x : f(x) < t\}.$$

Suppose now that $f(x) = t$. Since $\lim_{s \rightarrow \infty} f(s) = -\infty$, there exists s such that $f(r(s)) < t$. Convexity of f implies that for $z = r(s)$,

$$f(y) < f(x) = t, \quad \forall y \in xz \setminus \{x\}.$$

Therefore, x belongs to the closure of the open horoball B , which implies that \overline{B} is the closure of B . \square

EXERCISE 3.89. 1. Suppose that X is the Euclidean space \mathbb{R}^n , r is the geodesic ray in X with $r(0) = 0$ and $r'(0) = u$, where u is a unit vector. Show that

$$b_r(x) = -x \cdot u.$$

In particular, closed (resp. open) horoballs in X are closed (resp. open) half-spaces, while horospheres are hyperplanes.

2. Construct an example of a proper $CAT(0)$ space and an open horoball $B \subset X$, $B \neq X$, so that B is not equal to the interior of the closed horoball \overline{B} . Can this happen in the case of Hadamard manifolds?

Hyperbolic Space

The real hyperbolic space is the oldest and easiest example of hyperbolic spaces, which will be discussed in detail in Chapter 11. The real hyperbolic space has its origin in the following classical question that has challenged the geometers for nearly 2000 years:

QUESTION. *Does Euclid’s fifth postulate follow from the rest of the axioms of Euclidean geometry? (The fifth postulate is equivalent to the statement that given a line L and a point P in the plane, there exists exactly one line through P parallel to L .)*

After a long history of unsuccessful attempts to establish a positive answer to this question, N.I. Lobachevski, J. Bolyai and C.F. Gauss independently (in the early 19th century) developed a theory of non-Euclidean geometry (which we now call “hyperbolic geometry”), where Euclid’s fifth postulate is replaced by the axiom:

“For every point P which does not belong to L , there are infinitely many lines through P parallel to L .”

Independence of the 5th postulate from the rest of the Euclidean axioms was proved by E. Beltrami in 1868, *via* a construction of a model of hyperbolic geometry. In this chapter we will use the unit ball and the upper half-space models of hyperbolic geometry, the latter of which is due to H. Poincaré.

Given the classical nature of the subject, there are many books about real hyperbolic spaces, for instance, [And05], [Bea83], [BP92], [Rat06], [Thu97]. Our treatment of hyperbolic spaces is not meant to be comprehensive, we only cover the material needed elsewhere in the book. The purpose of this chapter is threefold:

1. It motivates many ideas and constructions in more general *Gromov–hyperbolic spaces*, which appear in Chapter 11.
2. It provides the necessary geometric background for *lattices* in the isometry group $PO(n, 1)$ of hyperbolic n -space. This background will be needed in the proof of various rigidity theorems for such lattices, which are due to Mostow, Tukia and Schwartz (Chapters 23 and 24).
3. We will use some basic hyperbolic geometry as a technical tool in proofs of a purely group-theoretic theorem, Stallings’s theorem on ends of groups. Hyperbolic geometry appears in both proofs of this theorem given in the book, Chapters 20 and 21.

4.1. Moebius transformations

We will think of the sphere \mathbb{S}^n as the 1-point compactification of Euclidean n -space \mathbb{E}^n ,

$$\mathbb{S}^n = \widehat{\mathbb{E}^n} = \mathbb{E}^n \cup \{\infty\}$$

Accordingly, we will regard the 1-point compactification of a hyperplane in \mathbb{E}^n as a *round sphere* (of infinite radius) and the 1-point compactification of a line in \mathbb{E}^n as a *round circle*. Another way to justify this treatment of hyperplanes and lines is that hyperplanes in \mathbb{E}^n appear as Chabauty–limits of round spheres: Consider a sequence of round spheres $S(\mathbf{a}_i, R_i)$ in \mathbb{E}^n passing through the origin ($|\mathbf{a}_i| = R_i$) with the sequence R_i diverging to infinity.

EXERCISE 4.1. Every sequence of spheres as above subconverges to a linear hyperplane in \mathbb{R}^n .

The *inversion* in the radius r sphere $\Sigma = S(\mathbf{0}, r) = \{\mathbf{x} : |\mathbf{x}| = r\}$ is the map

$$J_\Sigma : \mathbf{x} \mapsto r^2 \frac{\mathbf{x}}{|\mathbf{x}|^2}, \quad J_\Sigma(\mathbf{0}) = \infty, \quad J_\Sigma(\infty) = \mathbf{0}.$$

One defines the inversion J_Σ in the sphere $\Sigma = S(\mathbf{a}, r) = \{\mathbf{x} : |\mathbf{x} - \mathbf{a}| = r\}$ by the formula

$$J_\Sigma = T_{\mathbf{a}} \circ J_{S(\mathbf{0}, r)} \circ T_{-\mathbf{a}}, \quad J_\Sigma(\mathbf{x}) = r^2 \frac{(\mathbf{x} - \mathbf{a})}{|\mathbf{x} - \mathbf{a}|^2} + \mathbf{a},$$

where $T_{\mathbf{a}}$ is the translation by the vector \mathbf{a} . Inversions map round spheres to round spheres and round circles to circles; inversions also preserve Euclidean angles. We will regard the reflection in a Euclidean hyperplane as an inversion (this inversion fixes the point ∞). This is justified by:

EXERCISE 4.2. Suppose that the sequence of spheres $\Sigma_i = S(\mathbf{a}_i, r_i)$ converges to a linear hyperplane Π in \mathbb{R}^n . Show that the sequence of inversions J_i in the spheres Σ_i converges uniformly on compact subsets in \mathbb{R}^n to the reflection in Π .

DEFINITION 4.3. A *Moebius transformation* of \mathbb{E}^n (or, more precisely, of \mathbb{S}^n) is a composition of finitely many inversions in \mathbb{E}^n . The group of all Moebius transformations of \mathbb{E}^n is denoted $Mob(\mathbb{E}^n)$ or $Mob(\mathbb{S}^n)$.

In particular, Moebius transformations preserve angles, send circles to circles and spheres to spheres.

For instance, every translation is a Moebius transformation, since it is the composition of two reflections in parallel hyperplanes. Every rotation in \mathbb{E}^n is the composition of at most n inversions (reflections), since every rotation in \mathbb{E}^2 is the composition of two reflections. Every dilation $\mathbf{x} \mapsto \lambda \mathbf{x}, \lambda > 0$ is the composition of two inversions in spheres centered at $\mathbf{0}$. Thus, the group of *Euclidean similarities*

$$Sym(\mathbb{E}^n) = \{g : g(\mathbf{x}) = \lambda A\mathbf{x} + \mathbf{b}, \lambda > 0, A \in O(n), \mathbf{b} \in \mathbb{R}^n\},$$

is a subgroup of $Mob(\mathbb{S}^n)$.

The *cross-ratio* of a quadruple of points in \mathbb{S}^n is defined as:

$$[\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w}] := \frac{|\mathbf{x} - \mathbf{y}| \cdot |\mathbf{z} - \mathbf{w}|}{|\mathbf{y} - \mathbf{z}| \cdot |\mathbf{w} - \mathbf{x}|}.$$

Here and in what follows we assume, by default, that $\mathbf{y} \neq \mathbf{z}, \mathbf{x} \neq \mathbf{w}$. In the formula for the cross-ratio we use the chordal distance on the sphere (defined via the standard embedding of \mathbb{S}^n in \mathbb{E}^{n+1}). Instead, we can identify, via the stereographic projection, \mathbb{S}^n with the extended Euclidean space $\widehat{\mathbb{E}^n} = \mathbb{E}^n \cup \{\infty\}$ and use Euclidean distances, provided that the points $\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w}$ are not equal to the point ∞ . Even if one of these points is ∞ , we can define the cross-ratio by declaring that the two infinities appearing in the fraction defining the cross-ratio cancel each other.

This cross-ratio turns out to be equal to the one defined via the chordal metric (since the stereographic projection is the restriction of a Moebius transformation, see Example 4.6).

THEOREM 4.4. 1. A map $g : \mathbb{S}^n \rightarrow \mathbb{S}^n$ is a Moebius transformation if and only if it preserves cross-ratios of quadruples of points in \mathbb{S}^n . 2. If a Moebius transformation g fixes the point ∞ in $\widehat{\mathbb{E}^n}$, then g is a Euclidean similarity.

We refer the reader to [Rat06, Theorems 4.3.1, 4.3.2] for a proof.

This theorem has an immediate corollary:

COROLLARY 4.5. The subgroup $Mob(\mathbb{S}^n)$ is closed in the topological group $Homeo(\mathbb{S}^n)$, equipped with the topology of pointwise convergence.

EXAMPLE 4.6. Let us construct a Moebius transformation σ sending the open unit ball $\mathbf{B}^n = B(0, 1) \subset \mathbb{E}^n$ to the upper half-space \mathbf{U}^n ,

$$\mathbf{U}^n = \{\mathbf{x} = (x_1, \dots, x_n) : x_n > 0\}.$$

We take σ to be the composition of translation $\mathbf{x} \mapsto \mathbf{x} + \mathbf{e}_n$, where $\mathbf{e}_n = (0, \dots, 0, 1)$, inversion J_Σ , where $\Sigma = \partial\mathbf{B}^n$, translation $\mathbf{x} \mapsto \mathbf{x} - \frac{1}{2}\mathbf{e}_n$ and, lastly, the dilation $\mathbf{x} \rightarrow 2\mathbf{x}$. The reader will notice that the restriction of σ to the boundary sphere Σ of \mathbf{B}^n is nothing but the stereographic projection with the pole at $-\mathbf{e}_n$.

Note that the map σ sends the origin $0 \in \mathbf{B}^n$ to the point $\mathbf{e}_n \in \mathbf{U}^n$.

Given a subset $A \subset \mathbb{S}^n$, we will use the notation $Mob(A)$ for the stabilizer of A in $Mob(\mathbb{S}^n)$.

EXERCISE 4.7. Each Moebius transformation $g \in Mob(\mathbf{B}^n)$ commutes with the inversion J in the boundary sphere of \mathbf{B}^n .

Low-dimensional Moebius transformations. Suppose now that $n = 2$. The group $SL(2, \mathbb{C})$ acts on the extended complex plane $\mathbb{S}^2 = \mathbb{C} \cup \{\infty\}$ by *linear-fractional transformations*:

$$(4.1) \quad \begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot z = \frac{az + b}{cz + d}.$$

Note that the matrix $-I$ is in the kernel of this action, thus, the action factors through the group $PSL(2, \mathbb{C}) = SL(2, \mathbb{C})/\pm I$. If we identify the complex-projective line \mathbb{CP}^1 with the sphere $\mathbb{S}^2 = \mathbb{C} \cup \infty$ via the map $[z : w] \mapsto z/w$, this action of $SL(2, \mathbb{C})$ on \mathbb{S}^2 is nothing but the action of $SL(2, \mathbb{C})$ on \mathbb{CP}^1 obtained via projection of the linear action of $SL(2, \mathbb{C})$ on $\mathbb{C}^2 \setminus 0$.

EXERCISE 4.8. Show the group $PSL(2, \mathbb{C})$ acts faithfully on \mathbb{S}^2 .

EXERCISE 4.9. Prove that the subgroup $SL(2, \mathbb{R}) \subset SL(2, \mathbb{C})$ preserves the upper half-plane $\mathbf{U}^2 = \{z : Im(z) > 0\}$. Moreover, $SL(2, \mathbb{R})$ is the stabilizer of \mathbf{U}^2 in $SL(2, \mathbb{C})$.

EXERCISE 4.10. Prove that each matrix in $SL(2, \mathbb{C})$ is either of the form

$$\begin{pmatrix} a & b \\ 0 & a^{-1} \end{pmatrix}$$

or it can be written as a product

$$\begin{pmatrix} a & b \\ 0 & a^{-1} \end{pmatrix} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}$$

Hint: If a matrix is not of the first type then it is a matrix

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

such that $c \neq 0$. Use this information and multiplications on the left and on the right by matrices

$$\begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}$$

to create zeroes on the diagonal in the matrix.

LEMMA 4.11. *$PSL(2, \mathbb{C})$ is the subgroup $Mob_+(\mathbb{S}^2)$ of Moebius transformations of \mathbb{S}^2 which preserve orientation.*

PROOF. 1. Every linear-fractional transformation is a composition of

$$j : z \mapsto z^{-1},$$

translations, dilations and rotations (see Exercise 4.10). Note that $j(z)$ is the composition of the complex conjugation with the inversion in the unit circle. Thus, $PSL(2, \mathbb{C}) \subset Mob_+(\mathbb{S}^2)$. Conversely, let $g \in Mob(\mathbb{S}^2)$ and $z_0 := g(\infty)$. Then $h = j \circ \tau \circ g$ fixes the point ∞ , where $\tau_0(z) = z - z_0$. Let $z_1 = h(0)$. Then composition f of h with the translation $\tau_1 : z \mapsto z - z_1$ has the property that $f(\infty) = \infty, f(0) = 0$. Thus, $f \in CO(2)$ and h preserves orientation. It follows that f has the form $f(z) = \lambda z$, for some $\lambda \in \mathbb{C} \setminus 0$. Since $f, \tau_0, \tau-1, j$ are linear-fractional transformation, it follows that g is also linear-fractional. \square

EXERCISE 4.12. Show that the group $Mob(\mathbb{S}^1)$ equals the group of real-linear fractional transformations

$$x \mapsto \frac{ax + b}{cx + d},$$

$ad - bc \neq 0, a, b, c, d \in \mathbb{R}$.

4.2. Real hyperbolic space

The easiest way to introduce the real-hyperbolic n -space \mathbb{H}^n is by using its models: Upper half-space, unit ball and the projectivization of the two-sheeted hyperboloid in the Lorentzian model. Different features of \mathbb{H}^n are best visible in different models.

Upper half-space model. We equip \mathbf{U}^n with the Riemannian metric

$$(4.2) \quad ds^2 = \frac{d\mathbf{x}^2}{x_n^2} = \frac{dx_1^2 + \dots + dx_n^2}{x_n^2}$$

The Riemannian manifold (\mathbf{U}^n, ds^2) is called the *n -dimensional hyperbolic space* and denoted \mathbb{H}^n . This space is also frequently called the *real-hyperbolic space*, in order to distinguish it from other spaces also called *hyperbolic* (e.g., complex-hyperbolic space, quaternionic-hyperbolic space, Gromov-hyperbolic space, etc.). We will use the terminology *hyperbolic space* for \mathbb{H}^n and add adjective *real* in case when other notions of hyperbolicity are involved in the discussion. In case $n = 2$,

we identify \mathbb{R}^2 with the complex plane, so that $\mathbf{U}^2 = \{z | \text{Im}(z) > 0\}$, $z = x + iy$, and

$$ds^2 = \frac{dx^2 + dy^2}{y^2}.$$

Note that the hyperbolic Riemannian metric ds^2 on \mathbf{U}^n is conformally-Euclidean, hence, hyperbolic angles are equal to the Euclidean angles. One computes hyperbolic volumes of solids in \mathbb{H}^n by the formula

$$\text{Vol}(\Omega) = \int_{\Omega} \frac{dx_1 \dots dx_n}{x_n^n}$$

Consider the projection to the x_n -axis in \mathbf{U}^n given by the formula

$$\pi : (x_1, \dots, x_n) \mapsto (0, \dots, 0, x_n).$$

EXERCISE 4.13. 1. Verify that $d_x \pi$ does not increase the length of tangent vectors $\mathbf{v} \in T_x \mathbb{H}^n$ for every $x \in \mathbb{H}^n$.

2. Verify that for a unit vector $\mathbf{v} \in T_x \mathbb{H}^n$, $\|d_x \pi(\mathbf{v})\| = 1$ if and only if \mathbf{v} is “vertical”, i.e. it has the form $(0, \dots, 0, v_n)$.

Here and in what follows, the norm $\|\cdot\|$ is the one with respect to the hyperbolic Riemannian metric on the tangent spaces to \mathbb{H}^n ; the notation $|\cdot|$ is reserved for the Euclidean norm.

EXERCISE 4.14. Suppose that $\mathbf{p} = a\mathbf{e}_n$, $\mathbf{q} = b\mathbf{e}_n$, where $0 < a < b$. Let α be the vertical path $\alpha(t) = (1-t)\mathbf{p} + t\mathbf{q}$, $t \in [0, 1]$ connecting \mathbf{p} to \mathbf{q} . Show that α is the shortest path (with respect to the hyperbolic metric) connecting \mathbf{p} to \mathbf{q} in \mathbb{H}^n . In particular, α is a hyperbolic geodesic and

$$d(\mathbf{p}, \mathbf{q}) = \log(b/a).$$

Hint: Use the previous exercise.

We note that the metric ds^2 on \mathbb{H}^n is clearly invariant under the “horizontal” Euclidean translations $\mathbf{x} \mapsto \mathbf{x} + \mathbf{v}$, where $\mathbf{v} = (v_1, \dots, v_{n-1}, 0)$ (since they preserve the Euclidean metric and the x_n -coordinate). Similarly, ds^2 is invariant under the dilations

$$h : \mathbf{x} \mapsto \lambda \mathbf{x}, \lambda > 0$$

since h scales both numerator and denominator in (4.2) by λ^2 . Lastly, ds^2 is invariant under Euclidean rotations which fix the x_n -axis (since they preserve the x_n -coordinate). Clearly, the group generated by such isometries of \mathbb{H}^n act transitively on \mathbb{H}^n , which means that \mathbb{H}^n is a *homogeneous Riemannian manifold*.

EXERCISE 4.15. Show that \mathbb{H}^n is a complete Riemannian manifold. You can either use homogeneity of \mathbb{H}^n or show directly that every Cauchy sequence in \mathbb{H}^n lies in a compact subset of \mathbb{H}^n .

EXERCISE 4.16. Show that the inversion $J = J_{\Sigma}$ in the unit sphere Σ centered at the origin, is an isometry of \mathbb{H}^n . The proof is an easy but (somewhat) tedious calculation, which is best done using *calculus* interpretation of the pull-back Riemannian metric.

EXERCISE 4.17. Show that every inversion preserving \mathbb{H}^n is an isometry of \mathbb{H}^n . To prove this, use compositions of the inversion J_{Σ} in the unit sphere with translations and dilations.

In order to see clearly other isometries of \mathbb{H}^n , it is useful to consider the *unit ball model* of the hyperbolic space.

Unit ball model. Consider the open unit Euclidean n -ball

$$\mathbf{B}^n := \{\mathbf{x} \in \mathbb{R}^n : |\mathbf{x}| < 1\}$$

in \mathbb{E}^n . We equip \mathbf{B}^n with the Riemannian metric

$$ds_{\mathbf{B}}^2 = 4 \frac{dx_1^2 + \dots + dx_n^2}{(1 - |\mathbf{x}|^2)^2}.$$

The Riemannian manifold (\mathbf{B}^n, ds^2) is called the *unit ball model* of the hyperbolic n -space. What is clear in this model is that the group $O(n)$ of orthogonal transformations of \mathbb{R}^n preserves $ds_{\mathbf{B}}^2$ (since its elements preserve $|x|$ and, hence, the denominator of $ds_{\mathbf{B}}^2$). The two models of the hyperbolic space are related by the Moebius transformation $\sigma : \mathbf{B}^n \rightarrow \mathbf{U}^n$ defined in the previous section.

EXERCISE 4.18. Show that $ds_{\mathbf{B}}^2 = \sigma^*(ds^2)$. The proof is again a straightforward calculation similar to the Exercise 4.16. Namely, first, pull-back ds^2 via the dilation $\mathbf{x} \rightarrow 2\mathbf{x}$, then apply pull-back via the translation $\mathbf{x} \mapsto \mathbf{x} - \frac{1}{2}\mathbf{e}_n$, etc. Thus, σ is an isometry of the Riemannian manifolds $(\mathbf{B}^n, ds_{\mathbf{B}}^2), (\mathbf{U}^n, ds^2)$.

LEMMA 4.19. *The group $O(n)$ is the stabilizer of $\mathbf{0}$ in the group of isometries of $(\mathbf{B}^n, ds_{\mathbf{B}}^2)$.*

PROOF. Note that if $g \in \text{Isom}(\mathbf{B}^n)$ fixes $\mathbf{0}$, then its derivative at the origin $dg_{\mathbf{0}}$ is an orthogonal transformation u . Thus, the derivative (at the origin) of the composition $h = u^{-1}g \in \text{Isom}(\mathbf{B}^n)$ is the identity. Therefore, for every geodesic γ in \mathbb{H}^n such that $\gamma(0) = 0, D_0h(\gamma'(0)) = \gamma'(0)$. Since each geodesic in a Riemannian manifold is uniquely determined by its initial point and initial velocity, we conclude that $h(\gamma(t)) = \gamma(t)$ for every t . By completeness of \mathbb{H}^n , for every $q \in \mathbf{B}^n$ there exists a geodesic γ connecting p to q . It follows that $h(q) = q$ and, therefore, $g = u \in O(n)$. \square

COROLLARY 4.20. *The stabilizer of the point $\mathbf{e}_n \in \mathbf{U}^n$ in the group $\text{Isom}(\mathbb{H}^n)$ is contained in the group of Moebius transformations.*

PROOF. Note that σ sends $\mathbf{0} \in \mathbf{B}^n$ to $\mathbf{e}_n \in \mathbf{U}^n$, and σ is Moebius. Thus, $\sigma : \mathbf{B}^n \rightarrow \mathbf{U}^n$ conjugates the stabilizer $O(n)$ of $\mathbf{0}$ in $\text{Isom}(\mathbf{B}^n, ds_{\mathbf{B}}^2)$ to the stabilizer $K = \sigma^{-1}O(n)\sigma$ of \mathbf{e}_n in $\text{Isom}(\mathbf{U}^n, ds^2)$. Since $O(n) \subset \text{Mob}(\mathbb{S}^n), \sigma \in \text{Mob}(\mathbb{S}^n)$, the claim follows. \square

COROLLARY 4.21. *a. $\text{Isom}(\mathbb{H}^n)$ equals the group $\text{Mob}(\mathbb{H}^n)$ of Moebius transformations of \mathbb{S}^n preserving \mathbb{H}^n .*

b. $\text{Isom}(\mathbb{H}^n)$ acts transitively on the unit tangent bundle $U\mathbb{H}^n$ of \mathbb{H}^n .

PROOF. a. Since the two models of \mathbb{H}^n differ by a Moebius transformation, it suffices to work with \mathbf{U}^n .

1. We already know that the $\text{Isom}(\mathbb{H}^n) \cap \text{Mob}(\mathbb{H}^n)$ contains a subgroup acting transitively on \mathbb{H}^n . We also know, that the stabilizer K of p in $\text{Isom}(\mathbb{H}^n)$ is contained in $\text{Mob}(\mathbb{H}^n)$. Thus, given $g \in \text{Isom}(\mathbb{H}^n)$ we first find

$$h \in \text{Mob}(\mathbb{H}^n) \cap \text{Isom}(\mathbb{H}^n)$$

such that $k = h \circ g(p) = p$. Since $k \in \text{Mob}(\mathbb{H}^n)$, we conclude that $\text{Isom}(\mathbb{H}^n) \leq \text{Mob}(\mathbb{H}^n)$.

2. We leave it to the reader to verify that the restriction homomorphism $\text{Mob}(\mathbb{H}^n) \rightarrow \text{Mob}(\mathbb{S}^{n-1})$ is injective. Every $g \in \text{Mob}(\mathbb{S}^{n-1})$ extends to a composition of inversions preserving \mathbb{H}^n . Thus, the above restriction map is a group isomorphism. We already know that inversions $J \in \text{Mob}(\mathbb{H}^n)$ are hyperbolic isometries. Thus, $\text{Mob}(\mathbb{H}^n) \subset \text{Isom}(\mathbb{H}^n)$.

b. Transitivity of the action of $\text{Isom}(\mathbb{H}^n)$ on $U\mathbb{H}^n$ follows from the fact that this group acts transitively on \mathbb{H}^n and that the stabilizer of p acts transitively on the set of unit vectors in $T_p\mathbb{H}^n$. \square

For the next lemma we recall that we treat straight lines as circles.

LEMMA 4.22. *Geodesics in \mathbb{H}^n are arcs of circles orthogonal to the boundary sphere of \mathbb{H}^n . Furthermore, for every such arc α in \mathbf{U}^n , there exists an isometry of \mathbb{H}^n which carries α to a segment of the x_n -axis.*

PROOF. It suffices to consider complete hyperbolic geodesics $\alpha : \mathbb{R} \rightarrow \mathbb{H}^n$. Since $\sigma : \mathbf{B}^n \rightarrow \mathbf{U}^n$ sends circles to circles and preserves angles, it again suffices to work with the upper half-space model. Let α be a hyperbolic geodesic in \mathbf{U}^n . Since $\text{Isom}(\mathbb{H}^n)$ acts transitively on $U\mathbb{H}^n$, there exists a hyperbolic isometry g such that the hyperbolic geodesic $\beta = g \circ \alpha$ satisfies: $\beta(0) = p = \mathbf{e}_n$ and the vector $\beta'(0)$ has the form $\mathbf{e}_n = (0, \dots, 0, 1)$. We already know that the curve

$$\gamma : t \mapsto e^t \mathbf{e}_n$$

is a hyperbolic geodesic, see Exercise 4.14. Furthermore, $\gamma'(0) = \mathbf{e}_n$ and $\gamma(0) = p$. Thus, $\beta = \gamma$ is a circle orthogonal to the boundary of \mathbb{H}^n . Since $\text{Isom}(\mathbb{H}^n) = \text{Mob}(\mathbb{H}^n)$ and Moebius transformations map circles to circles and preserve angles, lemma follows. \square

COROLLARY 4.23. *The space \mathbb{H}^n is uniquely geodesic, i.e. for every pair of points in \mathbb{H}^n there exists a unique unit speed geodesic segment connecting these points.*

PROOF. By the above lemma, it suffices to consider points p, q on the x_n -axis. But, according to Exercise 4.14, the vertical segment is the unique length-minimizing path between such p and q . \square

COROLLARY 4.24. *Let $H \subset \mathbb{H}^n$ be the intersection of \mathbb{H}^n with a round k -sphere orthogonal to the boundary of \mathbb{H}^n . Then H is a totally-geodesic subspace of \mathbb{H}^n , i.e. for every pair of points $p, q \in H$, the unique hyperbolic geodesic γ connecting p and q in \mathbb{H}^n , is contained in H . Furthermore, if $\iota : H \rightarrow \mathbb{H}^n$ is the embedding, then the Riemannian manifold $(H, \iota^* ds^2)$ is isometric to \mathbb{H}^k .*

PROOF. The first assertion follows from the description of geodesics in \mathbb{H}^n . To prove the second assertion, by applying an appropriate isometry of \mathbb{H}^n , it suffices to consider the case when H is contained in a coordinate k -dimensional subspace in \mathbb{R}^n :

$$H = \{(0, \dots, 0, x_{n-k+1}, \dots, x_n) : x_n > 0\}.$$

Then

$$\iota^* ds^2 = \frac{dx_{n-k+1}^2 + \dots + dx_n^2}{x_n^2}$$

is isometric to the hyperbolic metric on \mathbb{H}^k (by relabeling the coordinates). \square

We will refer to the submanifolds $H \subset \mathbb{H}^n$ as *hyperbolic subspaces*.

EXERCISE 4.25. Show that the hyperbolic plane violates the 5th Euclidean postulate: For every (geodesic) line $L \subset \mathbb{H}^2$ and every point $P \notin L$, there are infinitely many lines through P which are parallel to L (i.e. disjoint from L).

EXERCISE 4.26. Prove that:

- The unit sphere \mathbb{S}^{n-1} (with its standard topology) is the ideal boundary (in the sense of Definition 3.78) of the hyperbolic space \mathbb{H}^n in the unit ball model.
- The extended Euclidean space $\widehat{\mathbb{E}}^n = \mathbb{S}^n$ is the ideal boundary of the hyperbolic space \mathbb{H}^{n+1} in the upper half-space model.

Note that the Moebius transformation $\sigma : \mathbf{B}^n \rightarrow \mathbf{U}^n$ carries the ideal boundary of \mathbf{B}^n to the ideal boundary of \mathbf{U}^n . Observe also that all Moebius transformations which preserve \mathbb{H}^n in either model, induce Moebius transformations of the ideal boundary of \mathbb{H}^n .

Lorentzian model of \mathbb{H}^n . We refer the reader to [Rat06] and [Thu97] for the material below.

Consider the *Lorentzian space* $\mathbb{R}^{n,1}$, which is \mathbb{R}^{n+1} equipped with the indefinite nondegenerate quadratic form

$$Q(\mathbf{x}) = x_1^2 + \dots + x_n^2 - x_{n+1}^2,$$

which is the quadratic form of the inner product

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n x_i y_i - x_{n+1} y_{n+1}.$$

Let H denote the upper sheet of the 2-sheeted hyperboloid in $\mathbb{R}^{n,1}$:

$$Q(\mathbf{x}) = -1, \quad x_{n+1} > 0.$$

The restriction of Q to the tangent bundle of H is positive-definite and, hence, defines a Riemannian metric ds^2 on H . We identify the unit ball \mathbf{B}^n in \mathbb{R}^n with the ball

$$\{(x_1, \dots, x_n, 0) : x_1^2 + \dots + x_n^2 < 1\} \subset \mathbb{R}^{n+1}$$

via the inclusion $\mathbb{R}^n \hookrightarrow \mathbb{R}^{n+1}$,

$$(x_1, \dots, x_n) \mapsto (x_1, \dots, x_n, 0).$$

Let $\pi : H \rightarrow \mathbf{B}^n$ denote the radial projection from the point $-\mathbf{e}_{n+1}$:

$$\pi(\mathbf{x}) = t\mathbf{x} - (1-t)\mathbf{e}_{n+1}, \quad t = \frac{1}{x_{n+1} + 1}.$$

One then verifies that

$$\pi : (H, ds^2) \rightarrow \mathbb{H}^n = \left(\mathbf{B}^n, \frac{4d\mathbf{x}^2}{(1-|\mathbf{x}|^2)^2} \right)$$

is an isometry. Accordingly, intersections of H with k -dimensional linear subspaces of \mathbb{R}^{n+1} are k -dimensional hyperbolic subspaces of \mathbb{H}^n .

Instead of working with the upper sheet H of the hyperboloid $\{Q = -1\}$ it is sometimes convenient to work with the projectivization of this hyperboloid or, equivalently, of the open cone

$$\{Q(\mathbf{x}) < 0\}.$$

Then the stabilizer $O(n, 1)^+$ of H in $O(n, 1)$ is naturally isomorphic to the quotient $PO(n, 1) = O(n, 1)/\pm I$. The stabilizer of H in $O(n, 1)$ acts isometrically on H . Furthermore, this stabilizer is the entire isometry group of (H, ds^2) .

Thus, $\text{Isom}(\mathbb{H}^n) \cong PO(n, 1) \cong O(n, 1)^+ < O(n, 1)$; in particular, the Lie group $\text{Isom}(\mathbb{H}^n)$ is linear.

The distance function in \mathbb{H}^n in terms of the Lorentzian inner product is given by the formula:

$$(4.3) \quad \cosh d(\mathbf{x}, \mathbf{y}) = -\langle \mathbf{x}, \mathbf{y} \rangle,$$

which is a direct analogue of the familiar formula for the angular metric on the unit sphere in terms of the Euclidean inner product. In order to see this, it suffices to consider the 1-dimensional hyperbolic space \mathbb{H}^1 identified with the hyperbola $x_1^2 - x_2^2 = -1, x_2 > 0$, in $\mathbb{R}^{1,1}$. This hyperbola is parameterized as

$$\mathbf{x}(t) = (\sinh(t), \cosh(t)), \quad t \in \mathbb{R}.$$

It is immediate from the definition of the induced Riemannian metric on \mathbb{H}^1 that this is an isometric parameterization of \mathbb{H}^1 and, hence,

$$t = \text{dist}(\mathbf{e}_2, \mathbf{x}), \quad \mathbf{x} = \mathbf{x}(t).$$

Lastly,

$$\langle \mathbf{e}_2, \mathbf{x} \rangle = -\cosh(t).$$

The general case follows from transitivity of the isometry group of \mathbb{H}^n .

EXERCISE 4.27 (Rigidity of n -point configurations). Every n -tuple of points (p_1, \dots, p_n) in \mathbb{H}^n is uniquely determined, up to an isometry of \mathbb{H}^n , by their mutual distances

$$\text{dist}(p_i, p_j), \quad i < j.$$

In particular, a geodesic triangle in \mathbb{H}^n is uniquely determined (up to congruence) by its side-lengths. Hint: Use the distance formula (4.3) and the fact that a quadratic form is uniquely determined (up to an isometry) by its Gram matrix.

The Lorentzian model of \mathbb{H}^n is a luxury one has in studying real-hyperbolic spaces, as the unit ball and the upper half-space models work just fine. However, *linear models* become a necessity when dealing with other hyperbolic spaces, complex-hyperbolic and quaternionic ones (see Section 4.9), as the unit ball and upper half-spaces models for such spaces become much more awkward to use.

4.3. Classification of isometries

Every continuous map of a closed disk to itself has a fixed point. Since every isometry of \mathbb{H}^n (in the unit ball model) extends to a Moebius transformation of the closed ball \mathbb{D}^n , isometries of the hyperbolic space are classified by their fixed points in \mathbb{D}^n .

DEFINITION 4.28. An isometry $g \in \text{Isom}(\mathbb{H}^n)$ is *elliptic* if it fixes a point $x \in \mathbb{H}^n$.

Conjugating an elliptic isometry g (fixing $x \in \mathbb{H}^n$) by an isometry $h \in \text{Isom}(\mathbb{H}^n)$, sending x to the center of the ball \mathbf{B}^n , we obtain another elliptic isometry

$$f = hgh^{-1}$$

which fixes the center of \mathbf{B}^n . Since f commutes with the inversion J in the unit sphere \mathbb{S}^{n-1} , we obtain:

$$f(\infty) = JfJ(\infty) = Jf(0) = J(0) = \infty.$$

Therefore, in view of Theorem 4.4, we conclude that f has to be a Euclidean similarity fixing the origin and preserving the unit ball \mathbf{B}^n . Such f is necessarily an orthogonal transformation, an element of the orthogonal group $O(n)$. We obtain:

LEMMA 4.29. *An element $g \in \text{Isom}(\mathbb{H}^n) = \text{Mob}(\mathbf{B}^n)$ is elliptic if and only if g is conjugate in $\text{Isom}(\mathbb{H}^n)$ to an orthogonal transformation.*

Suppose that a Moebius transformation g of the boundary sphere \mathbb{S}^{n-1} fixes three distinct points $\xi_1, \xi_2, \xi_3 \in \mathbb{S}^{n-1}$. Let C denote the unique round circle through these three points. The circle C appears as the boundary circle of a unique hyperbolic plane $\mathbb{H}^2 \subset \mathbb{H}^n$. Since g fixes the points ξ_1, ξ_2, ξ_3 , it has to preserve C and, hence, \mathbb{H}^2 . Furthermore, g preserves the hyperbolic geodesic $\gamma \subset \mathbb{H}^2$ asymptotic to ξ_1, ξ_2 . There exists a unique horoball $B \subset \mathbb{H}^2$ centered at ξ_3 , whose boundary touches the geodesic γ ; we let $x \in \gamma$ denote this point of tangency. By combining these observations, we conclude that g fixes the point x and is, therefore, elliptic. Moreover, we also see that g fixes two linearly independent vectors $v_1, v_3 \in T_x\mathbb{H}^2$: These are the initial velocity vectors of the geodesic rays ρ_1, ρ_3 emanating from x and asymptotic to ξ_1, ξ_2 respectively. Therefore, g fixes x and acts as the identity map on the tangent plane $T_x\mathbb{H}^2$.

EXERCISE 4.30. Use these facts to conclude that the isometry g fixes the hyperbolic plane \mathbb{H}^2 and the circle C pointwise. Alternatively, argue that a linear-fractional transformation fixing three points in C is the identity map.

We, thus, obtain:

LEMMA 4.31. *Each isometry of \mathbb{H}^n fixing at least three points in the boundary sphere \mathbb{S}^{n-1} is elliptic and, moreover, fixes pointwise a hyperbolic plane in \mathbb{H}^n .*

Of course, elliptic isometries need not fix any points in \mathbb{S}^{n-1} , for instance, the antipodal map

$$\mathbf{x} \mapsto -\mathbf{x}, \quad \mathbf{x} \in \mathbf{B}^n$$

is an elliptic isometry which has unique fixed point in \mathbb{D}^n . Another example to keep in mind is that each rotation $g \in SO(3)$ is an elliptic isometry of $\mathbb{H}^3 = \mathbf{B}^3$, which has exactly two fixed points in the boundary sphere.

In view of Lemma 4.31, in order to classify non-elliptic isometries, we have to consider isometries with one or two fixed points in \mathbb{S}^{n-1} .

DEFINITION 4.32. An isometry g of \mathbb{H}^n is called *parabolic* if it has exactly one fixed point in the boundary sphere \mathbb{S}^{n-1} .

Note that such an isometry cannot be elliptic, since a fixed point $x \in \mathbb{H}^n$ together with a fixed point $\xi \in \mathbb{S}^{n-1}$ determine a unique geodesic $\gamma \subset \mathbb{H}^n$ through x asymptotic to ξ . Therefore, an isometry g fixing both x and ξ also fixes the entire geodesic γ and, hence, the second ideal boundary point $\hat{\xi} \in \mathbb{S}^{n-1}$ of γ .

It is now convenient to switch from the unit ball model to the upper half-space model \mathbf{U}^n ; we choose a Moebius transformation $h : \mathbf{B}^n \rightarrow \mathbf{U}^n$ which sends the

fixed point ξ of g to the point ∞ in \widehat{E}^n . Conjugating g via h , we obtain a parabolic isometry

$$f = ghg^{-1}$$

whose unique fixed point is ∞ . Such f has to act as a Euclidean similarity on \mathbb{E}^{n-1} which has no fixed points in \mathbb{E}^{n-1} .

EXERCISE 4.33. Suppose that $f \in \text{Sim}(\mathbb{E}^{n-1})$ has no fixed points in \mathbb{E}^{n-1} . Then f has the form

$$f(\mathbf{x}) = A\mathbf{x} + \mathbf{b},$$

with $A \in O(n-1)$.

A Euclidean isometry $f(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$ is called a *Euclidean skew motion* with the translational component \mathbf{b} if the vector \mathbf{b} is non-zero and is fixed by A . Note that we allow Euclidean translations as special cases as skew motions (with the identity orthogonal component A).

EXERCISE 4.34. 1. Suppose that $f(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$ is a Euclidean isometry without fixed points in \mathbb{E}^{n-1} . Then f is conjugate by a translation in \mathbb{R}^n to a Euclidean skew motion.

2. Conversely, Euclidean skew motions have no fixed points in \mathbb{E}^{n-1} .

To summarize:

LEMMA 4.35. *An isometry of \mathbb{H}^n is parabolic if and only if it is conjugate in $\text{Mob}(\mathbb{S}^n)$ to a Euclidean skew motion.*

The last class of isometries of \mathbb{H}^n consists of *hyperbolic isometries*. Each hyperbolic isometry g has exactly two fixed points $\xi, \hat{\xi}$ in the boundary sphere \mathbb{S}^{n-1} . In order to distinguish such isometries from elliptic isometries, consider the unique geodesic γ in \mathbb{H}^n asymptotic to the points $\xi, \hat{\xi}$. This geodesic has to be preserved by g . Therefore, g induces an isometry $\gamma \rightarrow \gamma$. The isometry group of \mathbb{R} consists of three types of elements:

1. The identity map.
2. Reflections, $R_a : x \mapsto a - x$, $a \in \mathbb{R}$.
3. Nontrivial translations $x \mapsto x + b$, $b \in \mathbb{R} \setminus \{0\}$.

It is clear that if g induces an isometry of type 1 or 2 of the geodesic γ , then g is necessarily elliptic. This leads to:

DEFINITION 4.36. An isometry $g \in \text{Isom}(\mathbb{H}^n)$ is *hyperbolic* if it preserves a geodesic $\gamma \subset \mathbb{H}^n$ and acts on this geodesic as a non-zero Euclidean translation $x \mapsto x + b$. The number b is called the *translation number* τ_g of g . The geodesic γ is called the *axis* of g .

EXERCISE 4.37. Show that each hyperbolic isometry has unique axis. Hint: Assuming that g has two distinct axes, consider the action of g on their ideal boundary points.

Note that g , of course, fixes the ideal points $\xi, \hat{\xi} \in \mathbb{S}^{n-1}$ of its γ . One can distinguish g from an elliptic isometry fixing these points by noting that g is hyperbolic if and only if its derivative at these points *is not an orthogonal transformation*.

EXERCISE 4.38. Prove this characterization of hyperbolic isometries in terms of their derivatives. Hint: First consider the case when g fixes 0 and ∞ , and consider the derivative at the origin. Then reduce the general case to this one.

As with the elliptic and parabolic isometries, we can conjugate each hyperbolic isometry g to a Euclidean similarity, by sending (via a Moebius transformation $h : \mathbf{B}^n \rightarrow \mathbf{U}^n$) the fixed points $\xi, \hat{\xi}$ to 0 and ∞ respectively. The conjugate Moebius transformation

$$f = hgh^{-1}$$

has the form

$$f(\mathbf{x}) = \lambda A\mathbf{x}, \quad A \in O(n-1), \quad \lambda > 0, \quad \lambda \neq 1.$$

The translation number τ_g equals

$$\tau_g = |\log(\lambda)|,$$

since

$$\text{dist}(\mathbf{e}_n, \lambda\mathbf{e}_n) = |\log(\lambda)|.$$

In the case when $n = 3$ and we can identify $\text{Isom}_+(\mathbb{H}^n)$ with the group $PSL(2, \mathbb{C})$, one can give a simple *numerical* characterization of elliptic, parabolic and hyperbolic isometries:

Suppose that g is an orientation-preserving Moebius transformation of \mathbb{C} , represented by the matrices $\pm A$,

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in SL(2, \mathbb{C}).$$

We assume that $A \neq \pm I$, i.e. g is not the identity map (in which case, g is, of course, elliptic).

- EXERCISE 4.39. 1. g is elliptic iff $\text{tr}(A) \in (-2, 2) \subset \mathbb{C}$.
 2. g is parabolic iff $\text{tr}(A) = \pm 2$.
 3. g is hyperbolic iff $\text{tr}(A) \notin [-2, 2]$.

Hint: Use the fact that each $g \in PSL(2, \mathbb{C})$ is conjugate to a Euclidean similarity.

Lastly, we note that the elliptic-parabolic-hyperbolic classification of isometries can be generalized in the context of $CAT(-1)$ spaces X . Instead of relying upon the (unavailable) fixed-point theorem for general continuous maps, one classifies isometries g of X according to their *translation numbers*:

$$\tau_g = \inf_{x \in X} d(x, gx).$$

- An isometry g is elliptic if $\tau_g = 0$ and the infimum in the definition of τ_g is realized in X .
- An isometry g is parabolic if $\tau_g = 0$ and the infimum in the definition of τ_g is not realized in X .
- An isometry g is hyperbolic if $\tau_g > 0$.

EXERCISE 4.40. Show that the classification of isometries of \mathbb{H}^n described in this section is equivalent to their classification via translation numbers.

4.4. Hyperbolic trigonometry

In this section we consider geometry of triangles in the hyperbolic plane. We refer to [Bea83, Rat06, Thu97] for the proofs of the hyperbolic trigonometric formulae introduced in this section. Recall that a (geodesic) triangle $T = T(A, B, C)$ as a 1-dimensional object. From the Euclidean viewpoint, a hyperbolic triangle T is a concatenation of circular arcs connecting points A, B, C in \mathbb{H}^2 , where the circles containing the arcs are orthogonal to the boundary of \mathbb{H}^2 . Besides such “conventional” triangles, it is useful to consider *generalized* hyperbolic triangles where some vertices are *ideal*, i.e. they belong to the (ideal) boundary circle of \mathbb{H}^2 . Such triangles are easiest to introduce by using the Euclidean interpretation of hyperbolic triangles: One simply allows some (or, even all) vertices A, B, C to be points on the boundary circle of \mathbb{H}^2 , the rest of the definition is exactly the same. *However, we no longer allow two vertices which belong to the boundary circle \mathbb{S}^1 to be the same.* More intrinsically, an triangle $T(A, B, C)$, where, say, B and C are in \mathbb{H}^2 and $A \in \mathbb{S}^1$ is the concatenation of the geodesic arc BC and geodesic rays CA and BA (although, the natural orientation of the latter is from A to B).

The vertices of T which happen to be points of the boundary circle \mathbb{S}^1 are called the *ideal vertices* of T . The *angle* of T at its ideal vertex A is just the Euclidean angle, which has to be zero, since both sides of T at A are orthogonal to the ideal boundary circle \mathbb{S}^1 .

In general, we will use the notation $\alpha = \angle_A(B, C)$ to denote the angle of T at A . From now on, a *hyperbolic triangle* means either a usual triangle or a triangle where some vertices are ideal. We still refer to such triangles as *triangles in \mathbb{H}^2* , even though, some of the vertices could lie on the ideal boundary, so, strictly speaking, an ideal hyperbolic triangle in \mathbb{H}^2 is not a subset of \mathbb{H}^2 . An *ideal* hyperbolic triangle, is a triangle where all the vertices are distinct ideal points in \mathbb{H}^2 . The same conventions will be used for hyperbolic triangles in \mathbb{H}^n .




FIGURE 4.1. Geometry of a general hyperbolic triangle.

1. General triangles. Consider hyperbolic triangles T in \mathbb{H}^2 with the side-lengths a, b, c and the opposite angles α, β, γ , see Figure 4.1.

a. Hyperbolic Sine Law:

$$(4.4) \quad \frac{\sinh(a)}{\sin(\alpha)} = \frac{\sinh(b)}{\sin(\beta)} = \frac{\sinh(c)}{\sin(\gamma)}.$$

b. Hyperbolic Cosine Law:

$$(4.5) \quad \cosh(c) = \cosh(a) \cosh(b) - \sinh(a) \sinh(b) \cos(\gamma)$$

c. Dual Hyperbolic Cosine Law:

$$(4.6) \quad \cos(\gamma) = -\cos(\alpha) \cos(\beta) + \sin(\alpha) \sin(\beta) \cosh(c)$$

2. Right triangles. Consider a right-angled hyperbolic triangle with the hypotenuse c , the other side-lengths a, b and the opposite angles α, β . Then hyperbolic cosine laws become:

$$(4.7) \quad \cosh(c) = \cosh(a) \cosh(b),$$

$$(4.8) \quad \cos(\alpha) = \sin(\beta) \cosh(a),$$

$$(4.9) \quad \cos(\alpha) = \frac{\tanh b}{\tanh c}$$

In particular,

$$(4.10) \quad \cos(\alpha) = \frac{\cosh(a) \sinh(b)}{\sinh(c)}.$$

3. First variation formula for right triangles. We now hold the side a fixed and vary the hypotenuse in the above right-angled triangle. By combining (4.7) and (4.5) we obtain the *First Variation Formula*:

$$(4.11) \quad c'(0) = \frac{\cosh(a) \sinh(b)}{\sinh(c)} b'(0) = \cos(\alpha) b'(0).$$

The equation $c'(0) = \cos(\alpha) b'(0)$ is a special case of the *First Variation Formula* in Riemannian geometry, which applies to general Riemannian manifolds.

As an application of the first variation formula, consider a hyperbolic triangle with vertices A, B, C , side-lengths a, b, c and the angles β, γ opposite to the sides b, c . Then

LEMMA 4.41. $a + b - c \geq ma$, where

$$m = \min\{|1 - \cos(\beta)|, |1 - \cos(\gamma)|\}.$$

PROOF. We let $g(t)$ denote the unit speed parameterizations of the segment BC , such that $g(0) = C, g(a) = B$. Let $c(t)$ denote the distance $\text{dist}(A, g(t))$ (such that $b = c(0), c = c(a)$) and let $\beta(t)$ denote the angle $\angle Ag(t)B$. We leave it to the reader to verify that

$$|1 - \cos(\beta(t))| \geq m.$$

Consider the function

$$f(t) = t + b - c(t), \quad f(0) = 0, \quad f(a) = a + b - c.$$

By the first variation formula,

$$c'(t) = \cos(\beta(t))$$

and, hence,

$$f'(t) = 1 - \cos(\beta(t)) \geq m$$

Thus,

$$a + b - c = f(a) \geq ma \quad \square$$

EXERCISE 4.42. [Monotonicity of the hyperbolic distance] Let $T_i, i = 1, 2$ be right hyperbolic triangles with vertices A_i, B_i, C_i (where A_i or B_i could be ideal vertices) so that $A = A_1 = A_2$, $A_1B_1 \subset A_2B_2$, $\alpha_1 = \alpha_2$ and $\gamma_1 = \gamma_2 = \pi/2$. See Figure 4.2. Then $a_1 \leq a_2$. Hint: Use (4.9).

In other words, if $\sigma(t), \tau(t)$ are hyperbolic geodesics with unit speed parameterizations, so that $\sigma(0) = \tau(0) = A \in \mathbb{H}^2$, then the distance $d(\sigma(t), \tau)$ from the point $\sigma(t)$ to the geodesic τ , is a monotonically increasing function of t .




FIGURE 4.2. Monotonicity of distance.

4.5. Triangles and curvature of \mathbb{H}^n

Given points $A, B, C \in \mathbb{H}^n$ we define the *hyperbolic triangle* $T = [A, B, C] = \Delta ABC$ with vertices A, B, C . We topologize the set $Tri(\mathbb{H}^n)$ of hyperbolic triangles T in \mathbb{H}^n by using topology on triples of vertices of T , i.e. a subspace topology in $(\bar{\mathbf{B}}^n)^3$.

EXERCISE 4.43. Angles of hyperbolic triangles are continuous functions on $Tri(\mathbb{H}^n)$.

EXERCISE 4.44. Every hyperbolic triangle T in \mathbb{H}^n is contained in (the compactification of) a 2-dimensional hyperbolic subspace $\mathbb{H}^2 \subset \mathbb{H}^n$. Hint: Consider a triangle $T = [A, B, C]$, where A, B belong to a common vertical line.

So far, we considered only geodesic hyperbolic triangles, we now introduce their 2-dimensional counterparts. First, let $T = T(A, B, C)$ be a generalized hyperbolic triangle in \mathbb{H}^2 . We will assume that T is *nondegenerate*, i.e. is not contained in a hyperbolic geodesic. Such triangle T cuts \mathbb{H}^2 into several connected components, exactly one of which is a convex region with the boundary equal to T itself. (For instance, if all vertices of T are points in \mathbb{H}^2 , then $\mathbb{H}^2 \setminus T$ consists of two components, while if T is an ideal triangle, then $\mathbb{H}^2 \setminus T$ is a disjoint union of four convex regions.) The closure of this region is called *solid* (generalized) hyperbolic triangle and denoted $\blacktriangle = \blacktriangle(A, B, C)$. If T is degenerate, we set $\blacktriangle := T$. More generally, if $T \subset \mathbb{H}^n$ is a hyperbolic triangle, then the *solid triangle* bounded by T is the solid triangle bounded by T in the hyperbolic plane $\mathbb{H}^2 \subset \mathbb{H}^n$ containing T . We will retain the notation \blacktriangle for solid triangles in \mathbb{H}^n .

EXERCISE 4.45. Let S be a hyperbolic triangle with the sides $\sigma_i, i = 1, 2, 3$. Then there exists an ideal hyperbolic triangle T in \mathbb{H}^2 with the sides $\tau_i, i = 1, 2, 3$, bounding solid triangle \blacktriangle , so that $S \subset \blacktriangle$ and σ_1 is contained in the side τ_1 of T . See Figure 4.3.




FIGURE 4.3. Triangles in the hyperbolic plane.

LEMMA 4.46. $\text{Isom}(\mathbb{H}^2)$ acts transitively on the set of ordered triples of pairwise distinct points in $\partial_\infty \mathbb{H}^2$.

PROOF. Let $a, b, c \in \mathbb{R} \cup \infty$ be distinct points. By applying an inversion we send a to ∞ , so we can assume $a = \infty$. By applying a translation in \mathbb{R} we get $b = 0$. Lastly, composing a map of the type $x \rightarrow \lambda x, \lambda \in \mathbb{R} \setminus 0$, we send c to 1.

The composition of the above maps is a Moebius transformation of \mathbb{S}^1 and, hence, equals to the restriction of an isometry of \mathbb{H}^2 . \square

COROLLARY 4.47. *All ideal hyperbolic triangles are congruent to each other.*

EXERCISE 4.48. Generalize the above corollary to: Every hyperbolic triangle is uniquely determined by its angles. Hint: Use hyperbolic trigonometry.

We will use the notation $T_{\alpha,\beta,\gamma}$ to denote unique (up to congruence) triangle with the angles α, β, γ .

EXERCISE 4.49. The group $Mob(\mathbb{S}^n)$ acts transitively on 3-point subsets of \mathbb{S}^n . (Hint: Use the fact that any triple of points in \mathbb{S}^n is contained in a round circle; then apply Lemma 4.46.)

LEMMA 4.50. *Suppose that $(x, y, z), (x', y', z'), (x_i, y_i, z_i), i \in \mathbb{N}$ are triples of distinct points in \mathbb{S}^n and*

$$(4.12) \quad \lim_{i \rightarrow \infty} (x_i, y_i, z_i) = (x', y', z').$$

Assume that $\gamma_i \in Mob(\mathbb{S}^n)$ are such that

$$\gamma_i(x, y, z) = (x_i, y_i, z_i).$$

Then the sequence (g_i) belongs to a compact subset of $Mob(\mathbb{S}^n)$.

PROOF. We let $T, T', T_i \subset \mathbb{H}^{n+1}$ denote the (unique) ideal triangles with the vertices $(x, y, z), (x', y', z'), (x_i, y_i, z_i)$ respectively. Then each g_i sends T to T_i and maps the center c of T to the center c_i of T_i . The limit (4.12) implies that

$$\lim_{i \rightarrow \infty} c_i = c',$$

where c' is the center of T' . The Arzela-Ascoli theorem now implies precompactness of the sequence (g_i) in $Isom(\mathbb{H}^{n+1})$ and, hence, in $Mob(\mathbb{S}^n)$. \square

We now return to the study of geometry of hyperbolic triangles.

Given a hyperbolic triangle T bounding a solid triangle \blacktriangle , the *area of T* is the area of \blacktriangle ,

$$Area(T) = \iint_{\blacktriangle} \frac{dx dy}{y^2}.$$

Area of a degenerate hyperbolic triangle is, of course, zero. Here is an example of the area calculation. Consider the triangle $T = T_{0,\alpha,\pi/2}$ (which has angles $\pi/2, 0, \alpha$). We can realize T as the triangle with the vertices $i, \infty, e^{i\alpha}$. Computing hyperbolic area of this triangle (and using the substitution $x = \cos(t), \alpha \leq t \leq \pi/2$), we obtain

$$Area(T) = \iint_{\blacktriangle} \frac{dx dy}{y^2} = \frac{\pi}{2} - \alpha.$$

For $T = T_{0,0,\alpha}$, we subdivide T in two right triangles congruent to $T_{0,\alpha/2,\pi/2}$ and, thus, obtain

$$(4.13) \quad Area(T_{0,0,\alpha}) = \pi - \alpha.$$

In particular, area of the ideal triangle equals π .

LEMMA 4.51. $Area(T_{\alpha,\beta,\gamma}) = \pi - (\alpha + \beta + \gamma)$.

PROOF. The proof given here is due to Gauss, it appears in the letter from Gauss to Bolyai, see [Gau73]. We realize $T = T_{\alpha,\beta,\gamma}$ as a part of the subdivision of an ideal triangle $T_{0,0,0}$ in four triangles, the rest of which are $T_{0,0,\alpha'}$, $T_{0,0,\beta'}$, $T_{0,0,\gamma'}$, where $\theta' = \pi - \theta$ is the complementary angle. See Figure 4.4. Using additivity of area and equation (4.13), we obtain the area formula for T . \square




FIGURE 4.4. Computation of area of the triangle T .

Curvature computation. Our next goal is to compute the sectional curvature of \mathbb{H}^n . Since $\text{Isom}(\mathbb{H}^n)$ acts transitively on pairs (p, P) , where $P \subset T_p\mathbb{H}^n$ is a 2-dimensional subspace, it follows that \mathbb{H}^n has *constant* sectional curvature κ (see Section 3.6). Since $\mathbb{H}^2 \subset \mathbb{H}^n$ is totally geodesic and isometrically embedded (in the sense of Riemannian geometry), κ is the same for \mathbb{H}^n as for \mathbb{H}^2 .

COROLLARY 4.52. *The Gaussian curvature κ of \mathbb{H}^2 equals -1 .*

PROOF. Instead of computing the curvature tensor (see e.g. [dC92] for the computation), we will use Gauss-Bonnet formula. Comparing the area computation given in Lemma 4.51 with Gauss-Bonnet formula (Theorem 3.22) we conclude that $\kappa = -1$. \square

Note that scaling properties of the sectional curvature (see Section 3.6) imply that the sectional curvature of

$$\left(\mathbf{U}^n, a \frac{dx^2}{x_n^2} \right)$$

equals $-a^{-1}$ for every $a > 0$.

4.6. Distance function on \mathbb{H}^n

We begin by defining the following quantities:

$$(4.14) \quad \text{dist}(z, w) = \text{arccosh} \left(1 + \frac{|z - w|^2}{2 \text{Im } z \text{Im } w} \right) \quad z, w \in \mathbf{U}^2$$

and, more generally,

$$(4.15) \quad \text{dist}(\mathbf{p}, \mathbf{q}) = \text{arccosh} \left(1 + \frac{|\mathbf{p} - \mathbf{q}|^2}{2p_n q_n} \right) \quad \mathbf{p}, \mathbf{q} \in \mathbf{U}^n$$

It is immediate that $\text{dist}(p, q) = \text{dist}(q, p)$ and that $\text{dist}(p, q) = 0$ if and only if $p = q$. However, it is, *a priori*, far from clear that dist satisfies the triangle inequality.

LEMMA 4.53. *dist is invariant under $\text{Isom}(\mathbb{H}^n) = \text{Mob}(\mathbf{U}^n)$.*

PROOF. First, it is clear that dist is invariant under the group $\text{Euc}(\mathbf{U}^n)$ of Euclidean isometries which preserve \mathbf{U}^n . Next, any two points in \mathbf{U}^n belong to a vertical half-plane in \mathbf{U}^n . Applying elements of $\text{Euc}(\mathbf{U}^n)$ to this half-plane, we can transform it to the coordinate half-plane $\mathbf{U}^2 \subset \mathbf{U}^n$. Thus, the problem reduces to the case $n = 2$ and orientation-preserving Moebius transformations of \mathbb{H}^2 . We leave it to the reader as an exercise to show that the map $z \mapsto -\frac{1}{z}$ (which is an element of $\text{PSL}(2, \mathbb{R})$) preserves the quantity

$$\frac{|z - w|^2}{\text{Im } z \text{ Im } w}$$

and, hence, the function dist . Now, the assertion follows from Exercise 4.10 and Lemma 4.11. \square

Recall that $d(p, q)$ denotes the hyperbolic distance between points $p, q \in \mathbf{U}^n$.

PROPOSITION 4.54. *$\text{dist}(p, q) = d(p, q)$ for all points $p, q \in \mathbb{H}^n$. In particular, the function dist is indeed a metric on \mathbb{H}^n .*

PROOF. As in the above lemma, it suffices to consider the case $n = 2$. We can also assume that $p \neq q$. First, suppose that $p = i$ and $q = ib$, $b > 1$. Then, by Exercise 4.14,

$$\text{dist}(p, q) = \int_1^b \frac{dt}{t} = \log(b), \quad \exp(d(p, q)) = b.$$

On the other hand, the formula (4.14) yields:

$$\text{dist}(p, q) = \text{arccosh} \left(1 + \frac{(b-1)^2}{2b} \right).$$

Hence,

$$\cosh(\text{dist}(p, q)) = \frac{e^{\text{dist}(p, q)} + e^{-\text{dist}(p, q)}}{2} = 1 + \frac{(b-1)^2}{2b}.$$

Now, the equality $\text{dist}(p, q) = d(p, q)$ follows from the identity

$$1 + \frac{(b-1)^2}{2b} = \frac{b + b^{-1}}{2}.$$

For general points p, q in \mathbb{H}^2 , by Lemma 4.22, there exists a hyperbolic isometry which sends p to i and q to a point of the form $ib, b \geq 1$. We already know that both hyperbolic distance d and the quantity dist are invariant under the action of $\text{Isom}(\mathbb{H}^2)$. Thus, the equality $d(p, q) = \text{dist}(p, q)$ follows from the special case of points on the y -axis. \square

EXERCISE 4.55. 1. Deduce from (4.14) that

$$\log \left(1 + \frac{|z - w|^2}{2 \operatorname{Im} z \operatorname{Im} w} \right) \leq d(z, w) \leq \log \left(1 + \frac{|z - w|^2}{2 \operatorname{Im} z \operatorname{Im} w} \right) + \log 2$$

for all points $z, w \in \mathbf{U}^2$.

2. Suppose that \hat{A}, \hat{B} are distinct points in \mathbb{S}^1 and A, B are points which belong to the geodesic in \mathbb{H}^2 connecting \hat{A} to \hat{B} . Show that

$$\operatorname{dist}(A, B) = \left| \log[A, B, \hat{B}, \hat{A}] \right|.$$

Hint: First do the computation when $\hat{A} = 0, \hat{B} = \infty$ in the upper half-plane model.

4.7. Hyperbolic balls and spheres

Pick a point $p \in \mathbb{H}^n$ and a positive real number R . Then the *hyperbolic sphere* of radius R centered at p is the set

$$S(p, R) = \{x \in \mathbb{H}^n : d(x, p) = R\}.$$

EXERCISE 4.56. 1. Prove that $S(\mathbf{e}_n, R) \subset \mathbb{H}^n = \mathbf{U}^n$ equals the Euclidean sphere of center $\cosh(R)\mathbf{e}_n$ and radius $\sinh(R)$. *Hint.* It follows immediately from the distance formula (4.14).

2. Suppose that $S \subset \mathbf{U}^n$ is the Euclidean sphere with Euclidean radius R and the center x so that $x_n = a$. Then $S = S(p, r)$, where the hyperbolic radius r equals

$$\frac{1}{2} (\log(a + R) - \log(a - R)).$$

Since the group of Euclidean similarities acts transitively on \mathbf{U}^n , it follows that every hyperbolic sphere is also a Euclidean sphere. A non-computational proof of this fact is as follows: Since the hyperbolic metric $ds_{\mathbf{B}}^2$ on \mathbf{B}^n is invariant under $O(n)$, it follows that hyperbolic spheres centered at 0 in \mathbf{B}^n are also Euclidean spheres. The general case follows from transitivity of $\operatorname{Isom}(\mathbb{H}^n)$ and the fact that isometries of \mathbb{H}^n are Moebius transformations, which, therefore, send Euclidean spheres to Euclidean spheres.

LEMMA 4.57. *If $B(x_1, R_1) \subset B(x_2, R_2)$ are hyperbolic balls, then $R_1 \leq R_2$.*

PROOF. It follows from the triangle inequality that the diameter of a metric ball $B(x, R)$ is the longest geodesic segment contained in $B(x, R)$. Therefore, let $\gamma \subset B(x_1, R_1)$ be a diameter. Then γ is contained in $B(x_2, R_2)$ and, hence, its length is $\leq 2R_2$. However, the length of γ is $2R_1$, therefore, $R_1 \leq R_2$. \square

EXERCISE 4.58. Show that this lemma fails for general metric spaces.

4.8. Horoballs and horospheres in \mathbb{H}^n

Horoballs and horospheres play prominent role in the theory of discrete groups of isometries of hyperbolic n -space, primarily due to the *thick-thin decomposition*, which we will discuss in detail in Section 12.6.3. Later on, in Chapter 24 we will deal with families of disjoint horoballs in \mathbb{H}^n , while proving Schwartz' theorem on quasi-isometric rigidity of non-uniform lattices.

Consider a geodesic ray $r = x\xi$ in $\mathbb{H}^n = \mathbf{B}^n$, connecting a point $x \in \mathbb{H}^n$ to a boundary point $\xi \in \mathbb{S}^{n-1}$. We let b_r denote the Busemann function on \mathbb{H}^n for the

ray r ($b_r(x) = 0$). By Lemma 3.86, the open horoball $B(\xi)$ defined by the inequality $b_r < 0$, equals the union of open balls

$$B(\xi) = \bigcup_{t \geq 0} B(r(t), t).$$

As we saw in Section 4.7, in particular Exercise 4.56, each ball $B(r(t), t)$ is a Euclidean ball centered in a point $r(T_t)$ with $T_t > t$. Therefore, this union of hyperbolic balls is the open Euclidean ball with boundary tangent to \mathbb{S}^{n-1} at ξ , and containing the point x . According to Lemma 3.88, the closed horoball and the horosphere defined by $b_r \leq 0$ and $b_r = 0$, respectively, are the closed Euclidean ball and its boundary sphere, both with the point ξ removed.

EXERCISE 4.59. The isometry group of \mathbb{H}^n acts transitively on the set of open horoballs in \mathbb{H}^n .

We conclude that the set of horoballs (closed or open) with center ξ is the same as the set of Euclidean balls in \mathbf{B}^n (closed or open) tangent to \mathbb{S}^{n-1} at ξ , with the point ξ removed.

Applying the map $\sigma : \mathbf{B}^n \rightarrow \mathbf{U}^n$ to horoballs and horospheres in \mathbf{B}^n , we obtain horoballs and horospheres in the upper-half space model \mathbf{U}^n of \mathbb{H}^n . Being a Moebius transformation, σ carries Euclidean spheres to Euclidean spheres (recall that a compactified Euclidean hyperplane is also regarded as a Euclidean sphere). Recall that $\sigma(-\mathbf{e}_n) = \infty$. Therefore, every horosphere in \mathbf{B}^n centered at $-\mathbf{e}_n$ is sent by σ to an $n - 1$ -dimensional Euclidean subspace E of \mathbf{U}^n whose compactification contains the point ∞ . Hence, E has to be a horizontal Euclidean subspace, i.e. a subspace of the form

$$\{\mathbf{x} \in \mathbf{U}^n : x_n = t\}$$

for some fixed $t > 0$. Restricting the metric ds^2 to such E we obtain the Euclidean metric rescaled by t^{-2} . Thus, the restriction of the Riemannian metric ds^2 to every horosphere is isometric to the Euclidean $n - 1$ space \mathbb{E}^{n-1} . When working with horoballs and horospheres we will frequently use their identification with Euclidean half-spaces and hyperplanes in \mathbf{U}^n .

On the other hand, the restriction of the hyperbolic distance function to a horosphere is very far from the Euclidean metric: It follows from Exercise 4.55 that as the distance D between points x, y in a fixed horosphere Σ tends to infinity, the distance $\text{dist}(x, y)$ in \mathbb{H}^n also tends to infinity, but logarithmically slower:

$$\text{dist}(x, y) \asymp \log(D).$$

Thus, horospheres in \mathbb{H}^n are *exponentially distorted*, see Section 8.9.

We next consider intersections of horoballs $B(\xi_1) \cap B(\xi_2)$. If $\xi_1 = \xi_2$ then this intersection is either $B(\xi_1)$ or $B(\xi_2)$, whichever of these horoballs is smaller. Suppose now that $\xi_1 \neq \xi_2$. The horoballs $B(\xi_1), B(\xi_2)$ are said to be *opposite* in this case. Using the upper half-space model, we find an isometry of \mathbb{H}^n sending ξ_2 to ∞ and $B(\xi_2)$ to $\{x_n > 1\}$. After applying this isometry, we can assume that $B(\xi_2) = \{x_n > 1\}$ and $B(\xi_1)$ is a Euclidean round ball. Then the intersection of the horoballs is clearly bounded and, furthermore, the intersection

$$B(\xi_1) \cap \{x_n = 1\}$$

is either empty or is a round Euclidean ball. This proves:

LEMMA 4.60. 1. The intersection of two horoballs with the same center is another horoball with the same center.

2. The intersection of two opposite horoballs is always bounded.

3. The intersection of a horoball with the horosphere Σ bounding an opposite horoball is either empty or is a metric ball with respect to the intrinsic (flat) Riemannian metric of Σ .

EXERCISE 4.61. Consider the upper half-space model for the hyperbolic space \mathbb{H}^n and the vertical geodesic ray r in \mathbb{H}^n :

$$r = \{(0, \dots, 0, x_n) : x_n \geq 1\}.$$

Show that the Busemann function b_r for the ray r is given by

$$b_r(x_1, \dots, x_n) = -\log(x_n).$$

Consider the boundary horosphere $\Sigma \subset \mathbb{H}^n$ of the horoball

$$B = \{(x_1, \dots, x_n) : x_n > 1\}.$$

Define the projection

$$\pi : B^c := \mathbb{H}^n \setminus B \rightarrow \Sigma, \quad \pi(x_1, \dots, x_{n-1}, x_n) = (x_1, \dots, x_{n-1}, 1).$$

EXERCISE 4.62. For $x \in B^c$, the norm (computed with respect to the hyperbolic metric) of the derivative $d\pi_x$ equals x_n . In particular, $\|d\pi_x\| \leq 1$ with equality iff $x \in \Sigma$.

We now switch to the Lorentzian model H of the hyperbolic n -space. In view of the projection $\pi : H \rightarrow \mathbf{B}^n = \mathbb{H}^n$, can identify the ideal boundary points $\xi \in \mathbb{S}^{n-1}$ of \mathbb{H}^n with lines in the *future light cone*

$$C^\uparrow = \{\mathbf{x} \in \mathbb{R}^{n+1} : x_{n+1} > 0, Q(\mathbf{x}) = 0\}.$$

EXERCISE 4.63. Given a point $\xi \in C^\uparrow$, show that the corresponding Busemann function b_ξ on \mathbb{H}^n (up to constant) equals

$$-\log(-\langle \mathbf{x}, \xi \rangle).$$

Accordingly, horospheres in \mathbb{H}^n are projections of intersections of affine hyperplanes $\{\langle \mathbf{x}, \xi \rangle = a\} \cap H$, where $a < 0$. Similarly, show that open horoballs are projections of the intersections

$$\{\langle \mathbf{x}, \xi \rangle > a\} \cap H, \quad a < 0.$$

4.9. \mathbb{H}^n as a symmetric space

A *symmetric space* is a complete simply connected Riemannian manifold X such that for every point p there exists a global isometry of X which is a geodesic symmetry σ_p with respect to p , that is, for every geodesic γ through p , $\sigma_p(\gamma(t)) = \gamma(-t)$, where $\gamma(0) = p$. We will discuss general symmetric spaces and their discrete groups of isometries in more details in Chapter 12; we also refer the reader to [BH99, II.10], [Ebe96] and [Hel01] for a detailed treatment.

Let us verify that each symmetric space X is a homogeneous Riemannian manifold. Indeed, given points $p, q \in X$, let m denote the midpoint of a geodesic connecting p to q . Then $\sigma_m(p) = q$. Thus, X can be naturally identified with the quotient G/K , where G is a Lie group (acting transitively and isometrically on X) and $K < G$ is a compact subgroup. In the case of the symmetric spaces of nonpositive curvature we are interested in, the group G is *semisimple* and K is its

maximal compact subgroup. Another important subgroup, in the non-positively curved case, is the *minimal parabolic subgroup* $B < G$, it is a minimal subgroup of G such that the quotient G/B is compact. Geometrically speaking, the quotient G/B is identified with the *Furstenberg boundary* of X .

In the case of negatively curved symmetric spaces, G/B is the ideal boundary $\partial_\infty X$ of X in the sense of Section 3.11.3. The solvable group B has a further decomposition as the semidirect product

$$N \rtimes (T \times K_B),$$

where the group T is abelian and the subgroup N is nilpotent, while $K_B < B$ is maximal compact. The subgroup T is a *maximal (split) torus* of G . Both groups play important role in geometry of symmetric spaces. The dimension of T is the *rank* of X (and of G). A symmetric space is negatively curved if and only if it has rank 1. In this situation, the group N acts simply-transitively on a horosphere in X . Accordingly, $\partial_\infty X$ can be identified with a one-point compactification of N . This algebraic description of $\partial_\infty X$ plays an important role in proofs of rigidity theorems for rank 1 symmetric spaces.

In this section we describe how the real-hyperbolic space fits into the general framework of symmetric spaces. We will also discuss briefly other negatively curved symmetric spaces, as it turns out that besides real-hyperbolic spaces \mathbb{H}^n , there are three other families of negatively curved symmetric spaces: $\mathbb{C}\mathbb{H}^n, n \geq 2$ (complex-hyperbolic spaces), $\mathbb{H}\mathbb{H}^n, n \geq 2$ (quaternionic hyperbolic spaces) and $\mathbb{O}\mathbb{H}^2$ (octonionic hyperbolic plane).

Generalities of negatively curved symmetric spaces $\mathbb{H}^n, \mathbb{C}\mathbb{H}^n, \mathbb{H}\mathbb{H}^n, \mathbb{O}\mathbb{H}^2$. All four classes symmetric spaces can be described via a “linear algebra” model, generalizing the Lorentzian model of \mathbb{H}^n , although things become quite complicated in the case of $\mathbb{O}\mathbb{H}^2$ due to lack of associativity.

In the first three cases, the symmetric space X appears as a projectivization of a certain open cone V_- in a vector space (or a module in the case of $\mathbb{H}\mathbb{H}^n$), equipped with a hermitian form $\langle \cdot, \cdot \rangle$. The distance function in X is given by the formula:

$$(4.16) \quad \cosh^2(\text{dist}(\mathbf{p}, \mathbf{q})) = \frac{\langle \mathbf{p}, \mathbf{q} \rangle \langle \mathbf{q}, \mathbf{p} \rangle}{\langle \mathbf{p}, \mathbf{p} \rangle \langle \mathbf{q}, \mathbf{q} \rangle},$$

where $\mathbf{p}, \mathbf{q} \in V_-$ represent points in X .

In the case of all negatively curved symmetric spaces, the maximal torus T is isomorphic to \mathbb{R}_+ , while the group N is *2-step nilpotent*. Accordingly, the Lie algebra of N splits (as a vector space) as a direct sum

$$\mathfrak{n} = \mathfrak{n}_1 \oplus \mathfrak{n}_2$$

and this decomposition is T -invariant (one of these Lie algebras is trivial in the real-hyperbolic case). The subalgebra \mathfrak{n}_2 is the Lie algebra of the center of N . Each element $t \in T$ acts on \mathfrak{n} with two distinct eigenvalues λ_1, λ_2 , which are evaluations on t of two homomorphisms $\lambda_1, \lambda_2 : T \rightarrow \mathbb{R}_+$, called *characters*.

Special features of rank 1 symmetric spaces. The rank one symmetric spaces X are also characterized among symmetric spaces by the property that any two segments of the same length are congruent in X , i.e. the subgroup $K < G$ (the stabilizer of a point $p \in X$) acts transitively on each R -sphere $S(p, R)$ centered at

p . Another distinguishing characteristic of negatively curved symmetric spaces X is that their horospheres are exponentially distorted in X (cf. Section 4.8), while for all other non-positively curved symmetric spaces, horospheres are quasi-isometrically embedded. Furthermore, two horoballs with distinct centers in negatively curved symmetric spaces have bounded intersection, while it is not the case for the rest of the symmetric spaces.

Real-hyperbolic spaces \mathbb{H}^n . We note that in the unit ball model of \mathbb{H}^n we clearly have the symmetry σ_p with respect to the origin $p = 0$, namely, $\sigma_0 : \mathbf{x} \mapsto -\mathbf{x}$. Since \mathbb{H}^n is homogeneous, it follows that it has a symmetry at every point. Thus, \mathbb{H}^n is a symmetric space.

EXERCISE 4.64. Prove that the linear-fractional transformation $\sigma_i \in PSL(2, \mathbb{R})$ defined by $\pm S_i$, where

$$S_i = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

fixes i and is a symmetry with respect to i .

We proved in Section 4.5 that \mathbb{H}^n has negative curvature -1 . In particular, it contains no totally-geodesic Euclidean subspaces of dimension ≥ 2 and, thus, \mathbb{H}^n has rank 1.

The isometry group of \mathbb{H}^n is $PO(n, 1)$, its maximal compact subgroup is $K \simeq O(n)$, its subgroup B is the semidirect product

$$\mathbb{R}^{n-1} \rtimes (\mathbb{R}_+ \times O(n-1)) = N \rtimes (T \times K_B).$$

In the upper half-space model, the group N consists of Euclidean translations in \mathbb{R}^{n-1} , while T consists of dilations $\mathbf{x} \mapsto t\mathbf{x}, t > 0$.

There are many properties which distinguish the real-hyperbolic space among other rank 1 symmetric spaces, for instance, the fact that the subgroup N is abelian, which, geometrically, reflects flatness of the intrinsic Riemannian metric of the horospheres in \mathbb{H}^n . Another example is the fact that only in the real-hyperbolic space triangles are uniquely determined by their side-lengths: This is false for other hyperbolic spaces.

Complex-hyperbolic spaces. Start with the complex vector space $V = \mathbb{C}^{n+1}$ equipped with the Hermitian form

$$\langle \mathbf{v}, \mathbf{w} \rangle = \sum_{k=1}^n v_k \bar{w}_k - v_{n+1} \bar{w}_{n+1}.$$

The group $U(n, 1)$ is the group of complex-linear automorphisms of \mathbb{C}^{n+1} preserving this bilinear form. Consider the negative cone

$$V_- = \{ \mathbf{v} : \langle \mathbf{v}, \mathbf{v} \rangle < 0 \} \subset \mathbb{C}^{n+1}.$$

Then the *complex-hyperbolic space* \mathbf{CH}^n is the projectivization of V_- . The group $G = PU(n, 1)$ acts naturally on $X = \mathbf{CH}^n$. One can describe the Riemannian metric on \mathbf{CH}^n as follows. Let $\mathbf{p} \in V_-$ be a vector such that $\langle \mathbf{p}, \mathbf{p} \rangle = 1$; the tangent space $T_{[\mathbf{p}]}X$ of X at the projection $[\mathbf{p}]$ of \mathbf{p} , is the projection of the orthogonal complement \mathbf{p}^\perp in \mathbb{C}^{n+1} . Let $\mathbf{v}, \mathbf{w} \in \mathbb{C}^{n+1}$ be vectors orthogonal to \mathbf{p} , i.e.

$$\langle \mathbf{p}, \mathbf{v} \rangle = \langle \mathbf{p}, \mathbf{w} \rangle = 0.$$

Then define

$$(\mathbf{v}, \mathbf{w})_{\mathbf{p}} := \operatorname{Re} \langle \mathbf{v}, \mathbf{w} \rangle.$$

This determines a $PU(n, 1)$ -invariant Riemannian metric on X . The corresponding distance function (4.16) is $PU(n, 1)$ -invariant. The geodesic symmetry fixing the point $[\mathbf{e}_{n+1}]$ is the projectivization of the diagonal matrix $\operatorname{Diag}(-1, \dots, -1, 1)$.

The maximal compact subgroup of $PU(n, 1)$ is $U(n)$, the nilpotent subgroup $N < B < G$ is the *Heisenberg group*, its Lie algebra splits as

$$\mathbb{C}^n \oplus \mathbb{R},$$

where one should think of \mathbb{R} as the set of imaginary complex numbers (the reason for this will become clear shortly).

An important special feature of complex-hyperbolic spaces is the fact that they are *Kähler manifolds*: The $PU(n, 1)$ -invariant complex structure on $\mathbb{C}\mathbb{H}^n$ is the restriction of the complex structure on the ambient complex-projective space. The corresponding *almost complex structure* on the tangent bundle of $\mathbb{C}\mathbb{H}^n$ is given by the multiplication by i :

$$J(\mathbf{v}) = i\mathbf{v}, \quad \mathbf{v} \in T_{\mathbf{p}}\mathbb{C}\mathbb{H}^n.$$

This complex structure is *hermitian*, i.e. J preserves the Riemannian metric on $\mathbb{C}\mathbb{H}^n$. Furthermore, J and $(\mathbf{v}, \mathbf{w})_{\mathbf{p}}$ together define a $PU(n, 1)$ -invariant *symplectic structure* on $\mathbb{C}\mathbb{H}^n$ (a closed nondegenerate 2-form), given by

$$\omega(\mathbf{v}, \mathbf{w}) = (\mathbf{v}, J\mathbf{w}).$$

This Kähler nature of $\mathbb{C}\mathbb{H}^n$ means that one can use tools of complex analysis and complex differential geometry in order to study complex-hyperbolic spaces and their quotients by discrete isometry groups.

As we noted earlier, geodesic triangles $T \subset \mathbb{C}\mathbb{H}^n$ are not uniquely determined by their side-lengths. The additional invariant which determines geodesic triangles is their *symplectic area*, which is defined as the integral

$$\int_S \omega$$

of the symplectic form ω on $\mathbb{C}\mathbb{H}^n$ over any surface $S \subset \mathbb{C}\mathbb{H}^n$ bounded by T . (The fact that the area is independent of the choice of S follows from Stokes Theorem, since the form ω is closed.)

Quaternionic-hyperbolic spaces. Consider the ring \mathbf{H} of *quaternions*; the elements of the quaternion ring have the form

$$q = x + iy + jz + kw, \quad x, y, z, w \in \mathbb{R}.$$

The quaternionic conjugation is given by

$$\bar{q} = x - iy - jz - kw$$

and

$$|q| = (q\bar{q})^{1/2} \in \mathbb{R}_+$$

is the quaternionic norm. A *unit quaternion* is a quaternion of the unit norm. Let V be a left $n + 1$ -dimensional free module over \mathbf{H} :

$$V = \{\mathbf{q} = (q_1, \dots, q_{n+1}) : q_m \in \mathbf{H}\}.$$

Consider the quaternionic-hermitian inner product of signature $(n, 1)$:

$$\langle \mathbf{p}, \mathbf{q} \rangle = \sum_{m=1}^n p_m \bar{q}_m - p_{n+1} \bar{q}_{n+1}.$$

Then the group $G = Sp(n, 1)$ is the group of automorphisms of the module V preserving this inner product. The quotient of V by the group of non-zero quaternions \mathbf{H}^\times (with respect to the left multiplication action) is the n -dimensional *quaternionic-projective space* PV . Analogously to the case of real and complex hyperbolic spaces, consider the negative cone

$$V_- = \{\mathbf{q} \in V : \langle \mathbf{q}, \mathbf{q} \rangle < 0\}.$$

The group G acts naturally on $PV_- \subset PV$ through the group $PSp(n, 1)$, the quotient of G by the subgroup of unit quaternions embedded in the subgroup of diagonal matrices in G ,

$$q \mapsto qI.$$

The space PV_- is called the n -dimensional *quaternionic-hyperbolic space* $\mathbf{H}\mathbb{H}^n$. As in the real and complex cases, the geodesic symmetry fixing the point $[\mathbf{e}_{n+1}]$ is the projectivization of the diagonal matrix $Diag(-1, \dots, -1, 1)$.

The maximal compact subgroup of G is $Sp(n)$, the Lie algebra of the nilpotent subgroup $N < B$ splits as a real vector space as

$$\mathfrak{n}_1 \oplus \mathfrak{n}_2 = \mathbb{H}^n \oplus Im(\mathbf{H}),$$

where $Im(\mathbf{H})$ is the 3-dimensional real vector space of imaginary quaternions.

The octonionic-hyperbolic plane. One defines the *octonionic-hyperbolic plane* $\mathbf{O}\mathbb{H}^2$ analogously to $\mathbf{H}\mathbb{H}^n$, only using the algebra \mathbf{O} of Cayley octonions instead of quaternions. An extra complication comes from the fact that the algebra \mathbf{O} is not associative, which means that one cannot talk about free \mathbf{O} -modules.

The space $\mathbf{O}\mathbb{H}^2$ has dimension 16. It is identified with the quotient G/K , where G is a real form of the exceptional Lie group F_4 and the maximal compact subgroup $K < G$ is isomorphic to $Spin(9)$, the 2-fold cover of the orthogonal group $SO(9)$. The Lie algebra of the nilpotent subgroup $N < B < G$ has dimension 15; it splits as a real vector space as

$$\mathfrak{n}_1 \oplus \mathfrak{n}_2 = \mathbf{O} \oplus Im(\mathbf{O}),$$

where $Im(\mathbf{O})$ is the 7-dimensional vector space consisting of imaginary octonions.

We refer to Mostow's book [Mos73] and Parker's survey [Par08] for a more detailed discussion of negatively curved symmetric spaces.

4.10. Inscribed radius and thinness of hyperbolic triangles

Suppose that T is a hyperbolic triangle in the hyperbolic plane \mathbb{H}^2 with the sides $\tau_i, i = 1, 2, 3$, so that T bounds the solid triangle \blacktriangle . For a point $x \in \blacktriangle$ define

$$\Delta_x(T) := \max_{i=1,2,3} d(x, \tau_i).$$

and

$$\Delta(T) := \inf_{x \in \blacktriangle} \Delta_x(T).$$

The goal of this section is to estimate $\Delta(T)$ from above. It is immediate that the infimum in the definition of $\Delta(T)$ is realized by a point $x_o \in \blacktriangle$ which is equidistant from all the three sides of T , i.e. by the intersection point of the angle bisectors.

Define the *inscribed radius* $\text{inrad}(T)$ of T is the supremum of radii of hyperbolic disks contained in \blacktriangle .

LEMMA 4.65. $\Delta(T) = \text{inrad}(T)$.

PROOF. Suppose that $D = B(X, R) \subset \blacktriangle$ is a hyperbolic disk. Unless D touches two sides of T , there exists a disk $D' = B(X', R') \subset \blacktriangle$ which contains D and, hence, has larger radius, see Lemma 4.57. Suppose, therefore, that $D \subset \blacktriangle$ touches two boundary edges of T , hence, the center X of D belongs to the bisector σ of the corner ABC of T . Unless D touches all three sides of T , we can move the center X of D along the bisector σ away from the vertex B so that the resulting disk $D' = B(X', R')$ still touches only the sides AB, BC of T . We claim that the (radius R' of D' is larger than the radius R of D . In order to prove this, consider hyperbolic triangles $[X, Y, B]$ and $[X', Y', B']$, where Y, Y' are the points of tangency between D, D' and the side BA . These right-angled triangles have the common angle $\angle XBY$ and satisfy

$$d(B, X) \leq d(B, X').$$

Thus, the inequality $R \leq R'$ follows from the Exercise 4.42. \square

Thus, we need to estimate the inradius of hyperbolic triangles from above. Recall that by Exercise 4.45, for every hyperbolic triangle S in \mathbb{H}^2 there exists an ideal hyperbolic triangle T , so that $S \subset \blacktriangle$, the solid triangle bounded by T . Clearly, $\text{inrad}(S) \leq \text{inrad}(T)$. Since all ideal hyperbolic triangles are congruent, it suffices to consider the ideal hyperbolic triangle T in \mathbf{U}^2 with the vertices $-1, 1, \infty$. The inscribed circle C in T has Euclidean center $(0, 2)$ and Euclidean radius 1. Therefore, by Exercise 4.56, its hyperbolic radius equals $\log(3)/2$. By combining these observations with Exercise 4.44, we obtain

PROPOSITION 4.66. *For every hyperbolic triangle T , $\Delta(T) = \text{inrad}(T) \leq \frac{\log(3)}{2}$. In particular, for every hyperbolic triangle in \mathbb{H}^n , there exists a point $p \in \mathbb{H}^n$ so that distance from p to all three sides of T is $\leq \frac{\log(3)}{2}$.*

Another way to measure thinness of a hyperbolic triangle T is to compute distance from points of one side of T to the union of the two other sides. Let T be a hyperbolic triangle with sides $\tau_j, j = 1, 2, 3$. Define

$$\delta(T) := \max_j \sup_{p \in \tau_j} d(p, \tau_{j+1} \cup \tau_{j+2}),$$

where indices of the sides of T are taken modulo 3. In other words, if $\delta = \delta(T)$ then each side of T is contained in the δ -neighborhood of the union of the other two sides.

PROPOSITION 4.67. *For every geodesic triangle S in \mathbb{H}^n , $\delta(S) \leq \text{arccosh}(\sqrt{2})$.*

PROOF. First of all, as above, it suffices to consider the case $n = 2$. Let $\sigma_j, j = 1, 2, 3$ denote the edges of S . We will estimate $d(p, \sigma_2 \cup \sigma_3)$ (from above) for $p \in \sigma_1$. We enlarge the hyperbolic triangle S to an ideal hyperbolic triangle T as in Figure 4.5. For every $p \in \sigma_1$, every geodesic segment g connecting p to a point of $\tau_2 \cup \tau_3$ has to cross $\sigma_2 \cup \sigma_3$. In particular,

$$d(p, \sigma_2 \cup \sigma_3) \leq d(p, \tau_2 \cup \tau_3).$$

Thus, it suffices to show that $\delta(T) \leq \text{arccosh}(\sqrt{2})$ for the ideal triangle T as above. We realize T as the triangle with the (ideal) vertices $A_1 = \infty, A_2 = -1, A_3 = 1$ in




FIGURE 4.5. Enlarging the hyperbolic triangle S .

$\partial_\infty \mathbb{H}^2$. We parameterize the sides $\tau_i = A_{j-1}A_{j+1}$, $j = 1, 2, 3$ modulo 3, according to their orientation. Then, by the Exercise 4.42, for every j ,

$$d(\tau_j(t), \tau_{j-1})$$

is monotonically increasing. Thus,

$$\sup_t d(\tau_1(t), \tau_2 \cup \tau_3)$$

is achieved at the point $p = \tau_1(t) = i = \sqrt{-1}$ and equals $d(p, q)$, where $q = -1 + \sqrt{2}i$. Then, using formula 4.15, we get $d(p, q) = \operatorname{arccosh}(\sqrt{2})$. Note that alternatively, one can get the formula for $d(p, q)$ from (4.8) by considering the right triangle $[p, q, -1]$ where the angle at p equals $\pi/4$. \square

As we will see in Section 11.1, the above propositions mean that *all hyperbolic triangles are uniformly thin*.

COROLLARY 4.68.

$$\sup_{T \in \operatorname{Tri}(\mathbb{H}^n)} \delta(T) = \operatorname{arccosh}(\sqrt{2}).$$

4.11. Existence-uniqueness theorem for triangles

Proof of Lemma 3.54. We will prove this result for the hyperbolic plane \mathbb{H}^2 , this will imply the lemma for all $\kappa < 0$ by rescaling the metric on \mathbb{H}^2 . We leave the cases $\kappa \geq 0$ to the reader as the proof is similar. The proof below is goes back to Euclid (in the case of \mathbb{E}^2). Let c denote the largest of the numbers a, b, c . Draw a geodesic $\gamma \subset \mathbb{H}^2$ through points x, y such that $d(x, y) = c$. Then

$$\gamma = \gamma_x \cup xy \cup \gamma_y,$$

where γ_x, γ_y are geodesic rays emanating from x and y respectively. Now, consider the hyperbolic circles $S(x, b)$ and $S(y, a)$ centered at x, y and having radii b, a respectively. Since $c \geq \max(a, b)$,

$$\gamma_x \cap S(y, a) \subset \{x\}, \quad \gamma_y \cap S(x, b) \subset \{y\},$$

while

$$S(x, b) \cap xy = p, \quad S(y, a) \cap xy = q.$$

By the triangle inequality on $c \leq a + b$, p separates q from y (and q separates x from p). Therefore, both the ball $B(x, b)$ and its complement contain points of the circle $S(y, a)$, which (by connectivity) implies that $S(x, b) \cap S(y, a) \neq \emptyset$. Therefore, the triangle with the side-lengths a, b, c exists. Uniqueness (up to congruence) of this triangle follows from Exercise 4.27; alternatively it can be derived from the hyperbolic cosine law. \square

Groups and their actions

This chapter covers some basic group-theoretic material as well as group actions on topological and metric spaces. We also briefly discuss Lie groups, group cohomology and its relation to the structural theory of groups. For detailed treatment of the basic group theory we refer to [Hal76] and [LS77].

Notation and terminology. With very few exceptions, in a group G we use the multiplication sign \cdot to denote its binary operation. We denote its identity element either by e or by 1 . We denote the inverse of an element $g \in G$ by g^{-1} . Note that for abelian groups the neutral element is usually denoted 0 , the inverse of x by $-x$ and the binary operation by $+$. We will use the notation

$$[x, y] = xyx^{-1}y^{-1}$$

for the *commutator* of elements x, y of a group G .

A surjective homomorphism is called an *epimorphism*, while an injective homomorphism is called a *monomorphism*. If two groups G and G' are isomorphic we write $G \simeq G'$. An isomorphism of groups $\varphi : G \rightarrow G'$ is also called an *automorphism*. In what follows, we denote by $\text{Aut}(G)$ the group of automorphisms of G .

We use the notation $H < G$ or $H \leq G$ to denote that H is a subgroup in G . Given a subgroup H in G :

- the *order* $|H|$ of H is its cardinality;
- the *index* of H in G , denoted $|G : H|$, is the common cardinality of the quotients G/H and $H \backslash G$.

The *order* of an element g in a group (G, \cdot) is the order of the subgroup $\langle g \rangle$ of G generated by g . In other words, the order of g is the minimal positive integer n such that $g^n = 1$. If no such integer exists, then g is said to be of *infinite order*. In this case, $\langle g \rangle$ is isomorphic to \mathbb{Z} .

For every positive integer m we denote by \mathbb{Z}_m the *cyclic group of order m* , $\mathbb{Z}/m\mathbb{Z}$. Given $x, y \in G$ we let x^y denote the conjugation of x by y , i.e. xyx^{-1} .

5.1. Subgroups

Given two subsets A, B in a group G we denote by AB the subset

$$\{ab : a \in A, b \in B\} \subset G.$$

Similarly, we will use the notation

$$A^{-1} = \{a^{-1} : a \in A\}.$$

A *normal subgroup* K in G is a subgroup such that for every $g \in G$, $gKg^{-1} = K$ (equivalently $gK = Kg$). We use the notation $K \triangleleft G$ to denote that K is a normal

subgroup in G . When H and K are subgroups of G and either H or K is a normal subgroup of G , the subset $HK \subset G$ becomes a subgroup of G .

A subgroup K of a group G is called *characteristic* if for every automorphism $\phi : G \rightarrow G$, $\phi(K) = K$. Note that every characteristic subgroup is normal (since conjugation is an automorphism). But not every normal subgroup is characteristic:

EXAMPLE 5.1. Let G be the group $(\mathbb{Z}^2, +)$. Since G is abelian, every subgroup is normal. But, for instance, the subgroup $\mathbb{Z} \times \{0\}$ is not invariant under the automorphism $\phi : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2$, $\phi(m, n) = (n, m)$.

DEFINITION 5.2. The *center* $Z(G)$ of a group G is defined as the subgroup consisting of elements $h \in G$ so that $[h, g] = 1$ for each $g \in G$.

It is easy to see that the center is a characteristic subgroup of G .

DEFINITION 5.3. A *subnormal descending series* in a group G is a series

$$G = N_0 \triangleright N_1 \triangleright \cdots \triangleright N_n \triangleright \cdots$$

such that N_{i+1} is a normal subgroup in N_i for every $i \geq 0$.

If all N_i 's are normal subgroups of G , then the series is called *normal*.

A subnormal series of a group is called a *refinement* of another subnormal series if the terms of the latter series all occur as terms in the former series.

The following is a basic result in group theory:

LEMMA 5.4. *If G is a group, $N \triangleleft G$, and $A \triangleleft B < G$, then BN/AN is isomorphic to $B/A(B \cap N)$.*

DEFINITION 5.5. Two subnormal series

$$G = A_0 \triangleright A_1 \triangleright \cdots \triangleright A_n = \{1\} \quad \text{and} \quad G = B_0 \triangleright B_1 \triangleright \cdots \triangleright B_m = \{1\}$$

are called *isomorphic* if $n = m$ and there exists a bijection between the sets of partial quotients $\{A_i/A_{i+1} \mid i = 1, \dots, n-1\}$ and $\{B_i/B_{i+1} \mid i = 1, \dots, n-1\}$ such that the corresponding quotients are isomorphic.

LEMMA 5.6. *Any two finite subnormal series*

$$G = H_0 \geq H_1 \geq \cdots \geq H_n = \{1\} \quad \text{and} \quad G = K_0 \geq K_1 \geq \cdots \geq K_m = \{1\}$$

possess isomorphic refinements.

PROOF. Define $H_{ij} = (K_j \cap H_i)H_{i+1}$. The following is a subnormal series

$$H_{i0} = H_i \geq H_{i1} \geq \cdots \geq H_{im} = H_{i+1}.$$

When inserting all these in the series of H_i one obtains the required refinement.

Likewise, define $K_{rs} = (H_s \cap K_r)K_{r+1}$ and by inserting the series

$$K_{r0} = K_r \geq K_{r1} \geq \cdots \geq K_{rn} = K_r$$

in the series of K_r , we define its refinement.

According to Lemma 5.4

$$H_{ij}/H_{ij+1} = (K_j \cap H_i)H_{i+1}/(K_{j+1} \cap H_i)H_{i+1} \simeq K_j \cap H_i/(K_{j+1} \cap H_i)(K_j \cap H_{i+1}).$$

Similarly, one proves that $K_{ji}/K_{ji+1} \simeq K_j \cap H_i/(K_{j+1} \cap H_i)(K_j \cap H_{i+1})$. \square

DEFINITION 5.7. A group G is a *torsion group* if all its elements have finite order.

A group G is said to be *without torsion* (or *torsion-free*) if all its non-trivial elements have infinite order.

Note that the subset $\text{Tor } G = \{g \in G \mid g \text{ of finite order}\}$ of the group G , sometimes called the *torsion* of G , is in general not a subgroup.

DEFINITION 5.8. A group G is said to have property $*$ *virtually* if some finite-index subgroup H of G has the property $*$.

For instance, a group is *virtually torsion-free* if it contains a torsion-free subgroup of finite index, a group is *virtually abelian* if it contains an abelian subgroup of finite index and a *virtually free group* is a group which contains a free subgroup of finite index.

REMARK 5.9. Note that this terminology widely used in group theory is not entirely consistent with the notion of *virtually isomorphic groups*, which involves not only taking finite-index subgroups but also quotients by finite normal subgroups.

The following properties of finite-index subgroups will be useful.

LEMMA 5.10. *If $N \triangleleft H$ and $H \triangleleft G$, N of finite index in H and H finitely generated, then N contains a finite-index subgroup K which is normal in G .*

PROOF. By hypothesis, the quotient group $F = H/N$ is finite. For an arbitrary $g \in G$ the conjugation by g is an automorphism of H , hence H/gNg^{-1} is isomorphic to F . A homomorphism $H \rightarrow F$ is completely determined by the images in F of elements of a finite generating set of H . Therefore there are finitely many such homomorphisms, and finitely many possible kernels of them. Thus, the set of subgroups gNg^{-1} , $g \in G$, forms a finite list N, N_1, \dots, N_k . The subgroup $K = \bigcap_{g \in G} gNg^{-1} = N \cap N_1 \cap \dots \cap N_k$ is normal in G and has finite index in N , since each of the subgroups N_1, \dots, N_k has finite index in H . \square

PROPOSITION 5.11. *Let G be a finitely generated group. Then:*

- (1) *For every $n \in \mathbb{N}$ there exist finitely many subgroups of index n in G .*
- (2) *Every finite-index subgroup H in G contains a subgroup K which is finite index and characteristic in G .*

PROOF. (1) Let $H \leq G$ be a subgroup of index n . We list the left cosets of H :

$$H = g_1 \cdot H, g_2 \cdot H, \dots, g_n \cdot H,$$

and label these cosets by the numbers $\{1, \dots, n\}$. The action by left multiplication of G on the set of left cosets of H defines a homomorphism $\phi : G \rightarrow S_n$ such that $\phi(G)$ acts transitively on $\{1, 2, \dots, n\}$ and H is the inverse image under ϕ of the stabilizer of 1 in S_n . Note that there are $(n-1)!$ ways of labeling the left cosets, each defining a different homomorphism with these properties.

Conversely, if $\phi : G \rightarrow S_n$ is such that $\phi(G)$ acts transitively on $\{1, 2, \dots, n\}$, then $G/\phi^{-1}(\text{Stab}(1))$ has cardinality n .

Since the group G is finitely generated, a homomorphism $\phi : G \rightarrow S_n$ is determined by the image of a generating finite set of G , hence there are finitely many

distinct such homomorphisms. The number of subgroups of index n in H is equal to the number η_n of homomorphisms $\phi : G \rightarrow S_n$ such that $\phi(G)$ acts transitively on $\{1, 2, \dots, n\}$, divided by $(n-1)!$.

(2) Let H be a subgroup of index n . For every automorphism $\varphi : G \rightarrow G$, $\varphi(H)$ is a subgroup of index n . According to (1) the set $\{\varphi(H) \mid \varphi \in \text{Aut}(G)\}$ is finite, equal $\{H, H_1, \dots, H_k\}$. It follows that

$$K = \bigcap_{\varphi \in \text{Aut}(G)} \varphi(H) = H \cap H_1 \cap \dots \cap H_k.$$

Then K is a characteristic subgroup of finite index in H hence in G . □

EXERCISE 5.12. Does the conclusion of Proposition 5.11 still hold for groups which are not finitely generated?

Let S be a subset in a group G , and let $H \leq G$ be a subgroup. The following are equivalent:

- (1) H is the smallest subgroup of G containing S ;
- (2) $H = \bigcap_{S \subset G_1 \leq G} G_1$;
- (3) $H = \{s_1 s_2 \cdots s_n : n \in \mathbb{N}, s_i \in S \text{ or } s_i^{-1} \in S \text{ for every } i \in \{1, 2, \dots, n\}\}$.

The subgroup H satisfying any of the above is denoted $H = \langle S \rangle$ and is said to be *generated by* S . The subset $S \subset H$ is called a *generating set* of H . The elements in S are called *generators* of H .

When S consists of a single element x , $\langle S \rangle$ is usually written as $\langle x \rangle$; it is the cyclic subgroup consisting of powers of x .

We say that a normal subgroup $K \triangleleft G$ is *normally generated* by a set $R \subset K$ if K is the smallest normal subgroup of G which contains R , i.e.

$$K = \bigcap_{R \subset N \triangleleft G} N.$$

We will use the notation

$$K = \langle\langle R \rangle\rangle$$

for this subgroup. The subgroup K is also called the *normal closure* or the *conjugate closure* of R in G . Other notations for K which appear in the literature are R^G and $\langle R \rangle^G$.

5.2. Virtual isomorphisms of groups and commensurators

In this section we consider a weakening of the notion of a group isomorphism to the one of a *virtual isomorphism*. This turns out to be the right *algebraic* concept in the context of Geometric Group Theory.

DEFINITION 5.13. (1) Two groups G_1 and G_2 are called *co-embeddable* if there exist injective group homomorphisms $G_1 \rightarrow G_2$ and $G_2 \rightarrow G_1$.

(2) The groups G_1 and G_2 are *commensurable* if there exist finite-index subgroups $H_i \leq G_i$, $i = 1, 2$, such that H_1 is isomorphic to H_2 .

An isomorphism $\varphi : H_1 \rightarrow H_2$ is called an *abstract commensurator* of G_1 and G_2 .

- (3) We say that two groups G_1 and G_2 are *virtually isomorphic* (abbreviated as VI) if there exist finite-index subgroups $H_i \subset G_i$ and finite normal subgroups $F_i \triangleleft H_i$, $i = 1, 2$, so that the quotients H_1/F_1 and H_2/F_2 are isomorphic.

An isomorphism $\varphi : H_1/F_1 \rightarrow H_2/F_2$ is called a *virtual isomorphism* of G_1 and G_2 . When $G_1 = G_2$, φ is called *virtual automorphism*.

EXAMPLE 5.14. All countable free groups are co-embeddable. However, a free group of infinite rank is not virtually isomorphic to a free group of infinite rank.

PROPOSITION 5.15. *All the relations in Definition 5.13 are equivalence relation between groups.*

PROOF. The fact that co-embeddability is an equivalence relation is immediate. It suffices to prove that virtual isomorphism is an equivalence relation. The only non-obvious property is transitivity. We need:

LEMMA 5.16. *Let F_1, F_2 be normal finite subgroups of a group G . Then their normal closure $F = \langle\langle F_1, F_2 \rangle\rangle$ (i.e. the smallest normal subgroup of G containing F_1 and F_2) is again finite.*

PROOF. Let $f_1 : G \rightarrow G_1 = G/F_1$, $f_2 : G_1 \rightarrow G_1/f_1(F_2)$ be the quotient maps. Since the kernel of each f_1, f_2 is finite, it follows that the kernel of $f = f_2 \circ f_1$ is finite as well. On the other hand, the kernel of f is clearly the subgroup F . \square

Suppose now that G_1 is VI to G_2 and G_2 is VI to G_3 . Then we have

$$F_i \triangleleft H_i < G_i, |G_i : H_i| < \infty, |F_i| < \infty, \quad i = 1, 2, 3,$$

and

$$F'_2 \triangleleft H'_2 < G_2, |G_2 : H'_2| < \infty, |F'_2| < \infty,$$

so that

$$H_1/F_1 \cong H_2/F_2, \quad H'_2/F'_2 \cong H_3/F_3.$$

The subgroup $H''_2 := H_2 \cap H'_2$ has finite index in G_2 . By the above lemma, the normal closure in H''_2

$$K_2 := \langle\langle F_2 \cap H''_2, F'_2 \cap H''_2 \rangle\rangle$$

is finite. We have quotient maps

$$f_i : H''_2 \rightarrow C_i = f_i(H''_2) \leq H_i/F_i, i = 1, 3,$$

with finite kernels and cokernels. The subgroups $E_i := f_i(K_2)$, are finite and normal in C_i , $i = 1, 3$. We let $H'_i, F'_i \subset H_i$ denote the preimages of C_i and E_i under the quotient maps $H_i \rightarrow H_i/F_i$, $i = 1, 3$. Then $|F'_i| < \infty, |G_i : H'_i| < \infty, i = 1, 3$. Lastly,

$$H'_i/F'_i \cong C_i/E_i \cong H''_2/K_2, i = 1, 3.$$

Therefore, G_1, G_3 are virtually isomorphic. \square

Given a group G , we define $VI(G)$ as the set of equivalence classes of virtual automorphisms of G with respect to the following equivalence relation. Two virtual automorphisms of G , $\varphi : H_1/F_1 \rightarrow H_2/F_2$ and $\psi : H'_1/F'_1 \rightarrow H'_2/F'_2$, are *equivalent* if for $i = 1, 2$, there exist \tilde{H}_i , a finite-index subgroup of $H_i \cap H'_i$, and \tilde{F}_i , a normal subgroup in \tilde{H}_i containing the intersections $\tilde{H}_i \cap F_i$ and $\tilde{H}_i \cap F'_i$, such that φ and ψ induce the same automorphism from \tilde{H}_1/\tilde{F}_1 to \tilde{H}_2/\tilde{F}_2 .

Lemma 5.16 implies that the composition induces a binary operation on $VI(G)$, and that $VI(G)$ with this operation becomes a group, called *the group of virtual automorphisms of G* .

Let $Comm(G)$ be the set of equivalence classes of virtual automorphisms of G with respect to the equivalence relation defined as above, with the normal subgroups F_i and F'_i trivial. As in the case of $VI(G)$, the set $Comm(G)$, endowed with the binary operation defined by the composition, becomes a group, called the *abstract commensurator of the group G* .

Similarly to the notion of virtually isomorphic groups and abstract commensurators of groups one defines *commensurable subgroups* and *commensurators* of subgroups:

DEFINITION 5.17. Two subgroups Γ_1, Γ_2 of a group G are called *commensurable* if their intersection has finite index in both Γ_1 and in Γ_2 . The *commensurator* of a subgroup $\Gamma < G$ is defined as the set of elements g in G such that the subgroups $\Gamma, g\Gamma g^{-1}$ are commensurable. The commensurator of a subgroup $\Gamma < G$ is denoted $Comm_G(\Gamma)$.

EXERCISE 5.18. Show that $Comm_G(\Gamma)$ is a subgroup of G .

EXERCISE 5.19. Show that for $G = SL(n, \mathbb{R})$ and $\Gamma = SL(n, \mathbb{Z})$, $Comm_G(\Gamma)$ contains $SL(n, \mathbb{Q})$.

5.3. Commutators and the commutator subgroup

Recall that the commutator of two elements x, y of a group G is defined as $[x, y] = xyx^{-1}y^{-1}$. Thus:

- two elements x, y commute, i.e. $xy = yx$, if and only if $[x, y] = 1$.
- $xy = [x, y]yx$.

Thus, the commutator $[x, y]$ ‘measures the degree of non-commutation’ of the elements x and y . In Lemma 13.30 we will prove some further properties of commutators.

Let H, K be two subgroups of G . We denote by $[H, K]$ the subgroup of G generated by all commutators $[h, k]$ with $h \in H, k \in K$.

DEFINITION 5.20. The *commutator subgroup* (or *derived subgroup*) of G is the subgroup $G' = [G, G]$. As above, we may say that the commutator subgroup G' of G ‘measures the degree of non-commutativity’ of the group G .

A group G is *abelian* if every two elements of G commute, i.e. $ab = ba$ for all $a, b \in G$.

EXERCISE 5.21. Suppose that S is a generating set of G . Then G is abelian if and only if $[a, b] = 1$ for all $a, b \in S$.

- PROPOSITION 5.22.
- (1) G' is a characteristic subgroup of G ;
 - (2) G is abelian if and only if $G' = \{1\}$;
 - (3) $G_{ab} = G/G'$ is an abelian group (called the abelianization of G);

- (4) if $\varphi : G \rightarrow A$ is a homomorphism to an abelian group A , then φ factors through the abelianization: Given the quotient map $p : G \rightarrow G_{ab}$, there exists a homomorphism $\bar{\varphi} : G_{ab} \rightarrow A$ such that $\varphi = \bar{\varphi} \circ p$.

PROOF. (1) The set $S = \{[x, y] \mid x, y \in G\}$ is a generating set of G' and for every automorphism $\psi : G \rightarrow G$, $\psi(S) = S$.

Part (2) follows from the equivalence $xy = yx \Leftrightarrow [x, y] = 1$, and (3) is an immediate consequence of (2).

Part (4) follows from the fact that $\varphi(S) = \{1\}$. □

Recall that the *finite dihedral group* of order $2n$, denoted by D_{2n} or $I_2(n)$, is the group of symmetries of the regular Euclidean n -gon, i.e. the group of isometries of the unit circle $S^1 \subset \mathbb{C}$ generated by the rotation $r(z) = e^{\frac{2\pi i}{n}} z$ and the reflection $s(z) = \bar{z}$. Likewise, the *infinite dihedral group* D_∞ is the group of isometries of \mathbb{Z} (with the metric induced from \mathbb{R}); the group D_∞ is generated by the translation $t(x) = x + 1$ and the symmetry $s(x) = -x$.

EXERCISE 5.23. Find the commutator subgroup and the abelianization for the finite dihedral group D_{2n} and for the infinite dihedral group D_∞ .

EXERCISE 5.24. Let S_n (the symmetric group on n symbols) be the group of permutations of the set $\{1, 2, \dots, n\}$, and $A_n < S_n$ be the alternating subgroup, consisting of even permutations.

- (1) Prove that for every $n \notin \{2, 4\}$ the group A_n is generated by the set of cycles of length 3.
- (2) Prove that if $n \geq 3$, then for every cycle σ of length 3 there exists $\rho \in S_n$ such that $\sigma^2 = \rho\sigma\rho^{-1}$.
- (3) Use (1) and (2) to find the commutator subgroup and the abelianization for A_n and for S_n .
- (4) Find the commutator subgroup and the abelianization for the group H of permutations of \mathbb{Z} defined in Example 7.8 in Chapter 7.

Note that it is not necessarily true that the commutator subgroup G' of G consists entirely of commutators $\{[x, y] : x, y \in G\}$. However, occasionally, every element of the derived subgroup is indeed a single commutator. For instance, every element of the alternating group $A_n < S_n$ is the commutator in S_n , see [Ore51].

This leads to an interesting invariant (of geometric flavor) called the *commutator norm* (or *commutator length*) $\ell_c(g)$ of $g \in G'$, which is the least number k so that g can be expressed as a product

$$g = [x_1, y_1] \cdots [x_k, y_k],$$

as well as the *stable commutator norm* of g :

$$\limsup_{n \rightarrow \infty} \frac{\ell_c(g^n)}{n}.$$

See [Bav91, Cal08, Cal09] for further details. For instance, if G is the free group on two generators (see Definition 7.19), then every non-trivial element of G' has stable commutator norm greater than 1.

5.4. Semidirect products and short exact sequences

Let $G_i, i \in I$, be a collection of groups. The *direct product* of these groups, denoted

$$G = \prod_{i \in I} G_i$$

is the Cartesian product of the sets G_i with the group operation given by

$$(a_i) \cdot (b_i) = (a_i b_i).$$

Note that each group G_i is the quotient of G by the (normal) subgroup

$$\prod_{j \in I \setminus \{i\}} G_j.$$

A group G is said to *split* as a direct product of its normal subgroups $N_i \triangleleft G, i = 1, \dots, k$, if one of the following equivalent statements holds:

- $G = N_1 \cdots N_k$ and

$$N_i \cap N_1 \cdots N_{i-1} \cdot N_{i+1} \cdots N_k = \{1\} \text{ for all } i;$$

- for every element g of G there exists a unique k -tuple

$$(n_1, \dots, n_k), n_i \in N_i, i = 1, \dots, k$$

such that $g = n_1 \cdots n_k$.

Then G is isomorphic to the direct product $N_1 \times \dots \times N_k$. Thus, finite direct products G can be defined either *extrinsically*, using groups N_i as quotients of G , or *intrinsically*, using normal subgroups N_i of G .

Similarly, one defines *semidirect products* of two groups, by taking the above *intrinsic* definition and relaxing the normality assumption:

DEFINITION 5.25. (1) (*with the ambient group as the given data*) A group G is said to *split as a semidirect product of two subgroups* N and H , which is denoted by $G = N \rtimes H$, if and only if N is a *normal subgroup* of G , H is a *subgroup* of G , and one of the following equivalent statements holds:

- $G = NH$ and $N \cap H = \{1\}$;
- $G = HN$ and $N \cap H = \{1\}$;
- for every element g of G there exists a unique $n \in N$ and $h \in H$ such that $g = nh$;
- for every element g of G there exists a unique $n \in N$ and $h \in H$ such that $g = hn$;
- there exists a *retraction* $G \rightarrow H$, i.e. a homomorphism which restricts to the identity on H , and whose kernel is N .

Observe that the map $\varphi : H \rightarrow \text{Aut}(N)$ defined by $\varphi(h)(n) = hnh^{-1}$, is a group homomorphism.

(2) (*with the quotient groups as the given data*) Given any two groups N and H (not necessarily subgroups of the same group) and a group homomorphism $\varphi : H \rightarrow \text{Aut}(N)$, one can define a new group $G = N \rtimes_{\varphi} H$ which is a semidirect product of a copy of N and a copy of H in the above sense, defined as follows. As a set, $N \rtimes_{\varphi} H$ is defined as the cartesian product $N \times H$. The binary operation $*$ on G is defined by

$$(n_1, h_1) * (n_2, h_2) = (n_1 \varphi(h_1)(n_2), h_1 h_2), \forall n_1, n_2 \in N \text{ and } h_1, h_2 \in H.$$

The group $G = N \rtimes_{\varphi} H$ is called the *semidirect product of N and H with respect to φ* .

REMARKS 5.26. (1) If a group G is the semidirect product of a normal subgroup N with a subgroup H in the sense of (1), then G is isomorphic to $N \rtimes_{\varphi} H$ defined as in (2), where

$$\varphi(h)(n) = hnh^{-1}.$$

(2) The group $N \rtimes_{\varphi} H$ defined in (2) is a semidirect product of the normal subgroup $N_1 = N \times \{1\}$ and the subgroup $H = \{1\} \times H$ in the sense of (1).

(3) If both N and H are normal subgroups in (1), then G is a direct product of N and H .

If φ is the trivial homomorphism, sending every element of H to the identity automorphism of N , then $N \rtimes_{\varphi} H$ is the direct product $N \times H$.

Here is yet another way to define semidirect products. An *exact sequence* is a sequence of groups and group homomorphisms

$$\dots G_{n-1} \xrightarrow{\varphi_{n-1}} G_n \xrightarrow{\varphi_n} G_{n+1} \dots$$

such that $\text{Im } \varphi_{n-1} = \text{Ker } \varphi_n$ for every n . A *short exact sequence* is an exact sequence of the form:

$$(5.1) \quad \{1\} \longrightarrow N \xrightarrow{\varphi} G \xrightarrow{\psi} H \longrightarrow \{1\}.$$

In other words, φ is an isomorphism from N to a normal subgroup $N' \triangleleft G$ and ψ descends to an isomorphism $G/N' \simeq H$.

DEFINITION 5.27. A short exact sequence *splits* if there exists a homomorphism $\sigma : H \rightarrow G$ (called a *section*) such that

$$\psi \circ \sigma = \text{Id}.$$

When the sequence splits we shall sometimes write it as

$$1 \rightarrow N \rightarrow G \xrightarrow{\cong} H \rightarrow 1.$$

Every split exact sequence determines a decomposition of G as the semidirect product $\varphi(N) \rtimes \sigma(H)$. Conversely, every semidirect product decomposition $G = N \rtimes H$ defines a split exact sequence, where φ is the identity embedding and $\psi : G \rightarrow H$ is the retraction.

EXAMPLES 5.28. (1) The dihedral group D_{2n} is isomorphic to $\mathbb{Z}_n \rtimes_{\varphi} \mathbb{Z}_2$, where $\varphi(1)(k) = n - k$.

(2) The infinite dihedral group D_{∞} is isomorphic to $\mathbb{Z} \rtimes_{\varphi} \mathbb{Z}_2$, where $\varphi(1)(k) = -k$.

(3) The permutation group S_n is the semidirect product of A_n and $\mathbb{Z}_2 = \{\text{Id}, (12)\}$.

(4) The group $(\text{Aff}(\mathbb{R}), \circ)$ of affine maps $f : \mathbb{R} \rightarrow \mathbb{R}, f(x) = ax + b$, with $a \in \mathbb{R}^*$ and $b \in \mathbb{R}$ is a semidirect product $\mathbb{R} \rtimes_{\varphi} \mathbb{R}^*$, where $\varphi(a)(x) = ax$.

PROPOSITION 5.29. (1) Every isometry ϕ of \mathbb{R}^n is of the form $\phi(\mathbf{x}) = A\mathbf{x} + \mathbf{b}$, where $\mathbf{b} \in \mathbb{R}^n$ and $A \in O(n)$.

- (2) The group $\text{Isom}(\mathbb{R}^n)$ splits as the semidirect product $\mathbb{R}^n \rtimes O(n)$, with the obvious action of the orthogonal group $O(n)$ on \mathbb{R}^n .

Sketch of proof of (1). For every vector $\mathbf{a} \in \mathbb{R}^n$ we denote by $T_{\mathbf{a}}$ the translation of vector \mathbf{a} , $\mathbf{x} \mapsto \mathbf{x} + \mathbf{a}$.

If $\phi(\mathbf{0}) = \mathbf{b}$, then the isometry $\psi = T_{-\mathbf{b}} \circ \phi$ fixes the origin $\mathbf{0}$. Thus, it suffices to prove that an isometry fixing the origin is an element of $O(n)$. Indeed:

- an isometry of \mathbb{R}^n preserves straight lines, because these are bi-infinite geodesics;
- an isometry is a homogeneous map, i.e. $\psi(\lambda \mathbf{v}) = \lambda \psi(\mathbf{v})$; this is due to the fact that (for $0 < \lambda \leq 1$) $\mathbf{w} = \lambda \mathbf{v}$ is the unique point in \mathbb{R}^n satisfying

$$d(\mathbf{0}, \mathbf{w}) + d(\mathbf{w}, \mathbf{v}) = d(\mathbf{0}, \mathbf{v}).$$

- an isometry map is an additive map, i.e. $\psi(\mathbf{a} + \mathbf{b}) = \psi(\mathbf{a}) + \psi(\mathbf{b})$ because an isometry preserves parallelograms.

Thus, ψ is a linear transformation of \mathbb{R}^n , $\psi(\mathbf{x}) = A\mathbf{x}$ for some matrix A . The orthogonality of the matrix A follows from the fact that the image of an orthonormal basis under ψ is again an orthonormal basis. \square

EXERCISE 5.30. 1. Prove the statement (2) of Proposition 5.29. Note that \mathbb{R}^n is identified with the group of translations of the n -dimensional affine space *via* the map $\mathbf{b} \mapsto T_{\mathbf{b}}$.

2. Suppose that G is a subgroup of $\text{Isom}(\mathbb{R}^n)$. Is it true that G is isomorphic to the semidirect product $T \rtimes Q$, where $T = G \cap \mathbb{R}^n$ and Q is the projection of G to $O(n)$?

In sections 5.9.5 and 5.9.6 we discuss semidirect products and short exact sequences in more detail.

5.5. Direct sums and wreath products

Let X be a non-empty set, and let $\mathcal{G} = \{G_x \mid x \in X\}$ be a collection of groups indexed by X . Consider the set of maps $\text{Map}_f(X, \mathcal{G})$ with finite support, i.e.

$$\text{Map}_f(X, \mathcal{G}) := \left\{ f : X \rightarrow \prod_{x \in X} G_x \mid f(x) \in G_x, f(x) \neq 1_{G_x} \right.$$

for only finitely many $x \in X \left. \right\}$.

DEFINITION 5.31. The *direct sum* $\bigoplus_{x \in X} G_x$ is defined as $\text{Map}_f(X, \mathcal{G})$, endowed with the pointwise multiplication of functions:

$$(f \cdot g)(x) = f(x) \cdot g(x), \forall x \in X.$$

Clearly, if A_x are abelian groups, then $\bigoplus_{x \in X} A_x$ is abelian.

When $G_x = G$ is the same group for all $x \in X$, the direct sum is the set of maps

$$\text{Map}_f(X, G) := \{f : X \rightarrow G \mid f(x) \neq 1_G \text{ for only finitely many } x \in X\},$$

and we denote it either by $\bigoplus_{x \in X} G$ or by $G^{\oplus X}$.

If, in this latter case, the set X is itself a group H , then there is a natural action of H on the direct sum, defined by

$$\varphi : H \rightarrow \text{Aut} \left(\bigoplus_{h \in H} G \right), \quad \varphi(h)f(x) = f(h^{-1}x), \quad \forall x \in H.$$

Thus, we define the semidirect product

$$(5.2) \quad \left(\bigoplus_{h \in H} G \right) \rtimes_{\varphi} H.$$

DEFINITION 5.32. The semidirect product (5.2) is called *the wreath product of G with H* , and it is denoted by $G \wr H$. The wreath product $G = \mathbb{Z}_2 \wr \mathbb{Z}$ is called the *lamplighter group*.

This useful construction is a source of many interesting examples in group theory, for instance, we will see in Section 14.5 how it is used to prove failure of QI rigidity of the class of virtually solvable groups.

5.6. Geometry of group actions

5.6.1. Group actions. Let G be a group or a semigroup and X be a set. An *action of G on X on the left* is a map

$$\mu : G \times X \rightarrow X, \quad \mu(g, a) = g(a),$$

so that

- (1) $\mu(1, x) = x$;
- (2) $\mu(g_1 g_2, x) = \mu(g_1, \mu(g_2, x))$ for all $g_1, g_2 \in G$ and $x \in X$.

REMARK 5.33. If G is a group, then the two properties above imply that

$$\mu(g, \mu(g^{-1}, x)) = x$$

for all $g \in G$ and $x \in X$.

An *action of G on X on the right* is a map

$$\mu : X \times G \rightarrow X, \quad \mu(a, g) = (a)g,$$

so that

- (1) $\mu(x, 1) = x$;
- (2) $\mu(x, g_1 g_2) = \mu(\mu(x, g_1), g_2)$ for all $g_1, g_2 \in G$ and $x \in X$.

Note that the difference between an action on the left and an action on the right is the order in which the elements of a product act.

If not specified, an action of a group G on a set X is always on the left, and it is often denoted $G \curvearrowright X$. Every left action amounts to a homomorphism from G to the group $\text{Bij}(X)$ of bijections of X . An action is called *effective* or *faithful* if this homomorphism is injective. Given an action $\mu : G \times X \rightarrow X$ we will use the notation $g(x)$ for $\mu(g, x)$.

If X is a metric space, an *isometric action* is an action so that $\mu(g, \cdot)$ is an isometry of X for each $g \in G$. In other words, an isometric action is a group homomorphism

$$G \rightarrow \text{Isom}(X).$$

A group action $G \curvearrowright X$ on a set X is called *free* if for every $x \in X$, the *stabilizer of x in G* ,

$$G_x = \{g \in G : g(x) = x\}$$

is $\{1\}$.

Given an action $\mu : G \curvearrowright X$, a map $f : X \rightarrow Y$ is called *G -invariant* if

$$f(\mu(g, x)) = f(x), \quad \forall g \in G, x \in X.$$

Given two actions $\mu : G \curvearrowright X$ and $\nu : G \curvearrowright Y$, a map $f : X \rightarrow Y$ is called *G -equivariant* if

$$f(\mu(g, x)) = \nu(g, f(x)), \quad \forall g \in G, x \in X.$$

In other words, for each $g \in G$ we have a commutative diagram,

$$\begin{array}{ccc} X & \xrightarrow{g} & X \\ f \downarrow & & \downarrow f \\ Y & \xrightarrow{g} & Y \end{array}$$

A *topological group* is a group G equipped with a structure of topological space, so that the group operations (multiplication and inversion) are continuous maps. If G is a group without a specified topology, we will always assume that G is *discrete*, i.e. it is given the discrete topology. When referring to homomorphisms or isomorphisms of topological groups, we will always mean continuous homomorphisms and homeomorphic isomorphisms.

The usual algebraic concepts have *local analogues* for topological groups. One says that a map $\phi : G_2 \rightarrow G_1$ is a *local embedding* of topological groups if it is continuous on its domain, which is an (open) neighborhood U of $1 \in G_2$, $\phi(1) = 1$ and

$$\phi(g_1 g_2) = \phi(g_1) \phi(g_2),$$

whenever all three elements $g_1, g_2, g_1 g_2$ belong to U .

Accordingly, topological groups G_1, G_2 are said to be *locally isomorphic* if G_1 locally embeds in G_2 and vice-versa.

A topological group G is called *locally compact*, if it admits a basis of topology at $1 \in G$ consisting of relatively compact neighborhoods. A topological group is σ -compact if it is the union of countably many compact subsets.

LEMMA 5.34. *Each open subgroup $H \leq G$ is also closed.*

PROOF. The complement $G \setminus H$ equals the union

$$\bigcup_{g \notin H} gH$$

of open subsets. Therefore, H is closed. \square

A topological group G is said to be *compactly generated* if there exists a compact subset $K \subset G$ generating G . Every compactly generated group is σ -compact. The converse is not true in general: For locally compact σ -compact spaces (even compactly generated groups) second countability may not hold. Nevertheless, for

every locally compact σ -compact group G there exists a compact normal subgroup N such that G/N is second countable [Com84, Theorem 3.7].

THEOREM 5.35 (See [Str74]). *Every locally compact second countable Hausdorff group has a proper left-invariant metric.*

THEOREM 5.36 (Cartan–Iwasawa–Mal’cev). *Every connected locally compact topological group contains a unique¹ maximal compact subgroup.*

We will use this theorem in the context of Lie groups discussed in the next section.

If G is a topological group and X is a topological space, a *continuous action* of G on X is a continuous map μ satisfying the above *action* axioms.

A continuous action $\mu : G \curvearrowright X$ is called *proper* if for every compact subsets $K_1, K_2 \subset X$, the set

$$G_{K_1, K_2} = \{g \in G : g(K_1) \cap K_2 \neq \emptyset\} \subset G$$

is compact. If G has discrete topology, a proper action is called *properly discontinuous* action, as G_{K_1, K_2} is finite.

EXERCISE 5.37. Suppose that X is locally compact, Hausdorff and $G \curvearrowright X$ is proper. Show that the quotient X/G is Hausdorff.

Recall that a topological space X is called *Baire* if it satisfies the Baire property: *Countable intersections of open dense subsets of X are dense in X .* According to Baire theorem, each complete metric space is Baire.

LEMMA 5.38 (R. Arens, [Are46]). *Suppose that $G \times X \rightarrow X$ is a continuous transitive action of a σ -compact group on a Hausdorff Baire space X . Then for each $x \in X$ the orbit map $G \rightarrow X, g \mapsto g(x)$ descends to a homeomorphism $\phi : G/G_x \rightarrow X$.*

PROOF. We let $p : G \rightarrow G/G_x$ denote the quotient map. The map ϕ is defined by $\phi(p(g)) = g(x)$. We leave it to the reader to verify that ϕ is a continuous bijection, equivariant with respect to the G -action on G/G_x and X . Since X is Hausdorff, for each compact subset $K \subset G$ the restriction of the map ϕ to $p(K)$ is a homeomorphism to its image, which is necessarily closed. Since G is σ -compact, there exists a countable collection $K_i, i \in I$, of compact subsets of G , whose union equals G . Since the orbit map $G \rightarrow X$ is surjective and X is Baire, there exists a compact subset $K \subset G$ whose image has non-empty interior in X . Therefore, the restriction of ϕ^{-1} to the interior U of $\phi(p(K))$ is continuous. Since ϕ is G -equivariant and the G -orbit of U is the entire X , we conclude that $\phi^{-1} : X \rightarrow G/G_x$ is continuous. \square

A topological action $G \curvearrowright X$ is called *cocompact* if there exists a compact $C \subset X$ so that

$$G \cdot C := \bigcup_{g \in G} gC = X.$$

EXERCISE 5.39. If the action $G \curvearrowright X$ is cocompact, then the quotient space X/G (equipped with the quotient topology) is compact.

The following is a converse to the above exercise:

¹up to conjugation

LEMMA 5.40. *Suppose that X is a locally compact space and the action $G \curvearrowright X$ is such that the quotient space X/G is compact. Then G acts cocompactly on X .*

PROOF. Let $p : X \rightarrow Y = X/G$ be the quotient. For every $x \in X$ choose a relatively compact (open) neighborhood $U_x \subset X$ of x . Then the collection

$$\{p(U_x)\}_{x \in X}$$

is an open cover of Y . Since Y is compact, this open cover has a finite subcover

$$\{p(U_{x_i}) : i = 1, \dots, n\}$$

The union

$$C := \bigcup_{i=1}^n \text{cl}(U_{x_i})$$

is compact in X and projects onto Y . Hence, $G \cdot C = X$. \square

In the context of non-proper metric spaces, the concept of a cocompact group action is replaced with the one of a *cobounded action*. An isometric action $G \curvearrowright X$ is called *cobounded* if there exists $D < \infty$ such that for some point $x \in X$,

$$\bigcup_{g \in G} g(B(x, D)) = X.$$

Equivalently, given any pair of points $x, y \in X$, there exists $g \in G$ such that $\text{dist}(g(x), y) \leq 2D$. Clearly, if X is proper, the action $G \curvearrowright X$ is cobounded if and only if it is cocompact. We call a metric space X *quasihomogeneous* if the action $\text{Isom}(X) \curvearrowright X$ is cobounded.

Similarly, we have to modify the notion of a properly discontinuous action. A continuous isometric action $G \curvearrowright X$ of a topological group G on a metric space is called *metrically proper* if for every bounded subset $B \subset X$, the set

$$G_{B,B} = \{g \in G : g(B) \cap B \neq \emptyset\}$$

is relatively compact in G .

Note that if X is a proper metric space then a continuous isometric G -action on X is proper if and only if it is metrically proper.

EXERCISE 5.41. 1. A continuous isometric action $G \curvearrowright X$ is metrically proper if and only if for some (equivalently, every) $x \in X$ the function $g \mapsto d(g(x), x)$ has relatively compact sublevel sets in G .

2. Suppose that G is a discrete group. Then an isometric action $G \curvearrowright X$ is metrically proper if and only if for every sequence (g_n) consisting of distinct elements of G and for some (equivalently, every) $x \in X$ we have

$$\lim_{n \rightarrow \infty} \text{dist}(g_n(x), x) = \infty.$$

EXAMPLE 5.42. Let G be an infinite discrete group equipped with the discrete metric, taking only the values 0 and 1. Then the action $G \curvearrowright G$ is properly discontinuous as a topological action, but is not metrically proper.

5.6.2. Linear actions. In this section, V will denote a finite-dimensional vector space over a field \mathbb{K} whose algebraic closure will be denoted $\bar{\mathbb{K}}$. We let $End(V)$ denote the algebra of (linear) endomorphisms of V and $GL(V)$ the group of invertible endomorphisms of V . Linear actions of groups G on V are called *representations* of G on V .

A group Γ which is isomorphic to a subgroup of $GL(V)$ for some V , is called a *matrix group* or a *linear group*.

LEMMA 5.43.

$$\tau : End(V) \times End(V) \rightarrow \mathbb{K}, \tau(A, B) = \text{tr}(AB^T)$$

is a nondegenerate bilinear form on $End(V)$, regarded as a vector space over \mathbb{K} .

PROOF. Representing A and B by their matrix entries $(a_{ij}), (b_{kl})$, we obtain:

$$\text{tr}(AB^T) = \sum_{i,j} a_{ij}b_{ij}.$$

Therefore, if for some i, j , $a_{ij} \neq 0$, we take B such that $b_{kl} = 0$ for all $(k, l) \neq (i, j)$ and $b_{ij} = 1$. Then $\text{tr}(AB^T) = a_{ij} \neq 0$. \square

If V is a vector space and $A \subset End(V)$ is a subsemigroup, then A is said to *act irreducibly* on V if V contains no proper subspaces $V' \subset V$ such that $aV' \subset V'$ for all $a \in A$. An action is said to be *absolutely irreducible* iff the corresponding action on the vector space $V \otimes_{\mathbb{K}} \bar{\mathbb{K}}$ is irreducible.

A proof of the following theorem can be found, for instance, in [Lan02, Chapter XVII, §3, Corollary 3.3]:

THEOREM 5.44 (Burnside's theorem). *If $A \subset End(V)$ is a subalgebra which acts absolutely irreducibly on a finite-dimensional vector space V , then $A = End(V)$. In particular, if $G \subset End(V)$ is a subsemigroup acting irreducibly, then G spans $End(V)$ as a vector space.*

LEMMA 5.45. *If a linear action of a group G on V is absolutely irreducible, then so is the action G on $W = \Lambda^k V$.*

PROOF. Since G spans $End(V)$, the action of G on W is absolutely reducible iff the action of $End(V)$ is. However, all exterior product representations of $End(V)$ are absolutely irreducible; this is a special case of irreducibility of Weyl modules, see e.g. [FH94, Theorem 6.3, Part 4]. \square

EXERCISE 5.46. Suppose that $\mathbb{K} \subset \mathbb{L}$ is a field extension and the linear action $G \curvearrowright V$ is absolutely irreducible. Show that the action of G on $V \otimes_{\mathbb{K}} \mathbb{L}$ (regarded as a vector space over \mathbb{L}) is also absolutely irreducible. Give example of an irreducible representation which is not absolutely irreducible.

5.6.3. Lie groups. References for this section are [FH94, Hel01, OV90, War83].

A *Lie algebra* is a vector space \mathfrak{g} over a field F , equipped with a binary operation $[\cdot, \cdot] : \mathfrak{g}^2 \rightarrow \mathfrak{g}$, called the *Lie bracket*, which satisfies the following axioms:

1. The Lie bracket is bilinear:

$$[\lambda x, y] = \lambda \mu [x, y], \quad [x + y, z] = [x, z] + [y, z]$$

for all $\lambda \in F$, $x, y, z \in \mathfrak{g}$.

2. The Lie bracket is anti-symmetric:

$$[x, y] = -[y, x].$$

3. The Lie bracket satisfies the *Jacobi identity*:

$$[x, [y, z]] + [y, [z, x]] + [z, [x, y]] = 0.$$

In this book we will consider only finite-dimensional real and complex Lie algebras, i.e. we will assume that $F = \mathbb{R}$ or $F = \mathbb{C}$ and \mathfrak{g} is finite-dimensional as a vector space.

EXAMPLE 5.47. Lie algebra \mathfrak{g} which is the vector space of $n \times n$ matrices $Mat_n(F)$ with coefficients in the field F , with the Lie bracket given by the commutator

$$[X, Y] = XY - YX.$$

An *ideal* in a Lie algebra \mathfrak{g} is a vector subspace $J \subset \mathfrak{g}$ such that for every $x \in \mathfrak{g}, y \in J$ we have:

$$[x, y] \in J.$$

For instance, the subspace J consisting of scalar multiples of the identity matrix $I \in Mat_n(F)$ is an ideal in $\mathfrak{g} = Mat_n(F)$. A Lie algebra \mathfrak{g} is called *simple* if it is not 1-dimensional and every ideal in \mathfrak{g} is either 0 or the entire \mathfrak{g} .

If $\mathfrak{g}_1, \dots, \mathfrak{g}_m$ are Lie algebras, their *direct sum* is the direct sum of the vector spaces

$$\mathfrak{g}_1 \oplus \dots \oplus \mathfrak{g}_m$$

with the Lie bracket, given by

$$[x_1 + \dots + x_m, y_1 + \dots + y_m] = \sum_{i=1}^m [x_i, y_i]$$

for $x_i, y_i \in \mathfrak{g}_i, i = 1, \dots, m$. A Lie algebra \mathfrak{g} is called *semisimple* if it is isomorphic to the direct sum of finitely many simple Lie algebras.

A *Lie group* is a group G which has the structure of a smooth manifold, so that the binary operation $G \times G \rightarrow G$ and inversion $g \mapsto g^{-1}, G \rightarrow G$ are smooth maps. Actually, every Lie group G can be made into a real analytic manifold with real analytic group operations. We will mostly use the notation e for the neutral element of G . We will assume that G is a real n -dimensional manifold, although we will sometimes also consider *complex Lie groups*. A *homomorphism* of Lie groups is a group homomorphism which is also a smooth map.

EXAMPLE 5.48. Groups $GL(n, \mathbb{R}), SL(n, \mathbb{R}), O(n), O(p, q)$ are (real) Lie groups. Every countable discrete group (a topological group with discrete topology) is a Lie group. (Recall that we require our manifolds to be second countable. If we were to drop this requirement, then any discrete group becomes a Lie group.)

Here $O(p, q)$ is the group of linear isometries of the quadratic form

$$x_1^2 + \dots + x_p^2 - x_{p+1}^2 - \dots - x_{p+q}^2$$

of the signature (p, q) . The most important, for us, case is the group $O(n, 1) \simeq O(1, n)$. The group $PO(n, 1) = O(n, 1) / \pm I$ is the group of isometries of the hyperbolic n -space.

EXERCISE 5.49. Show that the group $PO(n, 1)$ embeds in $O(n, 1)$ as the subgroup stabilizing the *future light cone*

$$x_1^2 + \dots + x_n^2 - x_{n+1}^2 > 0, \quad x_{n+1} > 0.$$

The tangent space $V = T_e G$ of a Lie group G at the identity element $e \in G$ has the structure of a Lie algebra, called the *Lie algebra* \mathfrak{g} of the group G .

EXAMPLE 5.50. 1. The Lie algebra $\mathfrak{sl}(n, \mathbb{C})$ of $SL(n, \mathbb{C})$ consists of trace-free $n \times n$ complex matrices. The Lie bracket operation on $\mathfrak{sl}(n, \mathbb{C})$ is given by

$$[A, B] = AB - BA.$$

2. The Lie algebra of the unitary subgroup $U(n) < GL(n, \mathbb{C})$ equals the space of skew-hermitian matrices

$$\mathfrak{u}(n) = \{A \in Mat_n(\mathbb{C}) : A = -A^*\}.$$

3. The Lie algebra of the orthogonal subgroup $O(n) < GL(n, \mathbb{R})$ equals the space of skew-symmetric matrices

$$\mathfrak{o}(n) = \{A \in Mat_n(\mathbb{R}) : A = -A^T\}.$$

EXERCISE 5.51. $\mathfrak{u}(n) \oplus i\mathfrak{u}(n) = Mat_n(\mathbb{C})$, is the Lie algebra of the group $GL(n, \mathbb{C})$.

Every Lie group G has a left-invariant Riemannian metric, i.e. a Riemannian metric invariant under the left multiplication

$$L_g : G \rightarrow G, \quad L_g(x) = gx$$

by elements of G . Indeed, pick a positive-definite inner product $\langle \cdot, \cdot \rangle_e$ on $T_e G$. The map $L_g : G \rightarrow G$ is a diffeomorphism and the action of G on itself *via* left multiplication is simply-transitive. We define the inner product $\langle \cdot, \cdot \rangle_g$ on $T_g G$ as the image of $\langle \cdot, \cdot \rangle_e$ under the derivative $DL_g : T_e G \rightarrow T_g G$. Similarly, if G is a compact Lie group, then it admits a bi-invariant Riemannian metric, i.e. a Riemannian metric invariant under both left and right multiplication. Namely, if $\langle \cdot, \cdot \rangle$ is a left-invariant Riemannian metric on G , define the right-invariant metric by the formula:

$$\langle u, v \rangle' := \int_G \langle DR_g(u), DR_g(v) \rangle dVol(g),$$

where $dVol$ is the volume form of the Riemannian metric $\langle \cdot, \cdot \rangle$ and R_g is the right multiplication by g :

$$R_g : G \rightarrow G, \quad R_g(x) = xg.$$

Every Lie group G acts on itself *via* inner automorphisms

$$\rho(g)(x) = Inn(g)(x) = gxg^{-1}.$$

This action is smooth and the identity element $e \in G$ is fixed by the entire group G . Therefore, G acts linearly on the tangent space $V = T_e G$ at the identity $e \in G$. The action of G on V is called *the adjoint representation of the group G* and denoted by Ad . Thus, one obtains a homomorphism

$$Ad : G \rightarrow GL(V).$$

LEMMA 5.52. *For every connected Lie group G the kernel of $Ad : G \rightarrow GL(V)$ is contained in the center of G .*

PROOF. There is a local diffeomorphism

$$\exp : V \rightarrow G$$

called *the exponential map* of the group G , sending $0 \in V$ to $e \in G$. In the case when $G = GL(n, \mathbb{R})$ this map is the ordinary matrix exponential map. The map \exp satisfies the identity

$$g \exp(v) g^{-1} = \exp(\text{Ad}(g)v), \quad \forall v \in V, g \in G.$$

Thus, if $\text{Ad}(g) = \text{Id}$, then g commutes with every element of G of the form $\exp(v)$, $v \in V$. The set of such elements is open in G . Now, if we are willing to use a real analytic structure on G , then it would immediately follow that g belongs to the center of G . Below is an alternative argument. Let $g \in \text{Ker}(\text{Ad})$. The centralizer $Z(g)$ of g in G is given by the equation

$$Z(g) = \{h \in G : [h, g] = 1\}.$$

Since the commutator is a continuous map, $Z(g)$ is a closed subgroup of G . Moreover, as we observed above, this subgroup has non-empty interior in G (containing e). Since $Z(g)$ acts transitively on itself by, say, left multiplication, $Z(g)$ is open in G . As G is connected, we conclude that $Z(g) = G$. Therefore the kernel of Ad is contained in the center of G . The opposite inclusion is immediate. \square

DEFINITION 5.53. A connected noncommutative Lie group G is called *simple* if G contains no closed connected proper normal subgroups.

Equivalently, a connected Lie group G is simple if its Lie algebra \mathfrak{g} is simple.

EXAMPLE 5.54. 1. The group $SL(2, \mathbb{R})$ is simple, but its center is isomorphic to \mathbb{Z}_2 . Thus, a simple Lie group need not be simple as an abstract group.

2. Examples of simple Lie groups are $SL(n, \mathbb{R})$, $O_0(p, q)$, $q \geq 2$, $q \geq 1$, unless $p = q = 2$, $Sp(n, \mathbb{R})$.

DEFINITION 5.55. A connected Lie group G is *semisimple* if its Lie algebra is semisimple. For instance, the Lie group $O_0(2, 2)$ is semisimple but not simple.

Below are several deep structural theorems about Lie groups:

THEOREM 5.56 (S. Lie). 1. For every finite-dimensional real Lie algebra \mathfrak{g} there exists a unique² simply-connected Lie group G whose Lie algebra is isomorphic to \mathfrak{g} .

2. Two Lie groups are locally isomorphic if and only if their Lie algebras are isomorphic.

THEOREM 5.57 (E. Cartan). Every closed subgroup H of a Lie group G has a structure of Lie group so that the inclusion $H \hookrightarrow G$ is an embedding of smooth manifolds.

The next theorem is a corollary of the *Peter-Weyl theorem*, see e.g., [OV90, Theorem 10, page 245]:

THEOREM 5.58. Every compact Lie group is linear, i.e. it embeds in $GL(V)$ for some finite-dimensional real vector space V .

²up to an isomorphism

While there are nonlinear (connected) Lie groups, e.g. the universal cover of $SL(2, \mathbb{R})$, each Lie group is *locally* linear.

THEOREM 5.59 (I. D. Ado, [Ado36]). *Every finite-dimensional Lie algebra \mathfrak{g} over a field F of characteristic zero (e.g. over the real numbers) admits a faithful finite-dimensional representation. In particular, every Lie group locally embeds in $GL(V)$ for some finite-dimensional real vector space V .*

We refer the reader to [FH94, Theorem E.4] for a proof. Note that if a Lie group G has discrete center, then the adjoint representation of G is a local embedding of G in $GL(\mathfrak{g})$, where \mathfrak{g} is the Lie algebra of G . The difficulty is in the case of groups with non-discrete center.

5.6.4. Haar measure and lattices.

DEFINITION 5.60. A (left) Haar measure on a locally compact Hausdorff topological group G is a countably additive, non-trivial measure μ on Borel subsets of G satisfying:

- (1) $\mu(gE) = \mu(E)$ for every $g \in G$ and every Borel subset $E \subset G$.
- (2) $\mu(K)$ is finite for every compact $K \subset G$.
- (3) Every Borel subset $E \subset G$ is *outer regular*:

$$\mu(E) = \inf\{\mu(U) : E \subset U, U \text{ is open in } G\}.$$

- (4) Every open set $E \subset G$ is *inner regular*:

$$\mu(E) = \sup\{\mu(K) : K \subset E, K \text{ is compact in } G\}.$$

Accordingly, topological groups G_1, G_2 are said to be *locally isomorphic* if G_1 locally embeds in G_2 and vice-versa.

By Haar's Theorem, see [Bou63], every locally compact Hausdorff topological group G admits a Haar measure and this measure is unique up to scaling. Similarly, one defines *right Haar measures*. In general, left and right Haar measures are not the same, but they are for some important classes of groups:

DEFINITION 5.61. A locally compact Hausdorff topological group G is *unimodular* if left and right Haar measures are constant multiples of each other.

Important examples of Haar measures come from Riemannian geometry. Let G be a Lie group. We equip G with a left-invariant Riemannian metric. The volume form of this metric defines a left Haar measure on G .

Let $\Gamma < G$ be a discrete subgroup of a locally compact Hausdorff topological group G . A *measurable fundamental set* in G for the left action of γ on G is a measurable subset of $D \subset G$ such that

$$\bigcup_{\gamma \in \Gamma} \gamma D = G, \quad \mu(\gamma D \cap D) = 0, \quad \forall \gamma \in \Gamma \setminus \{e\}.$$

LEMMA 5.62. *Every discrete subgroup $\Gamma < G$ admits a measurable fundamental set.*

PROOF. Since $\Gamma < G$ is discrete, there exists an open neighborhood V of $e \in G$ such that $\Gamma \cap V = \{e\}$. Since G is a topological group, there exists another open neighborhood U of $e \in G$, such that $UU^{-1} \subset V$. Then for $\gamma \in \Gamma$ we have

$$\gamma u = u', u \in U, u' \in U \Rightarrow \gamma = u'u^{-1} \in V \Rightarrow \gamma = e.$$

In other words, Γ -images of U are pairwise disjoint. Since G is a second countable, there exists a countable subset

$$E = \{g_i \in G : i \in \mathbb{N}\}$$

such that

$$G = \bigcup_i U g_i.$$

Clearly, each set

$$W_n := U g_n \setminus \bigcup_{i < n} \Gamma U g_i$$

is measurable, and so is their union

$$D = \bigcup_{n=1}^{\infty} W_n.$$

Let us verify that D is a measurable fundamental set. First, note that for every $x \in G$ there exists the least n such that $x \in U g_n$. Therefore,

$$G = \bigcup_{n=1}^{\infty} \left(U g_n \setminus \bigcup_{i < n} U g_i \right).$$

Next,

$$\begin{aligned} \Gamma \cdot D &= \bigcup_{n=1}^{\infty} \left(\Gamma U g_n \setminus \bigcup_{i < n} \Gamma U g_i \right) = \\ \Gamma \cdot \bigcup_{n=1}^{\infty} \left(U g_n \setminus \bigcup_{i < n} U g_i \right) &\supset \bigcup_{n=1}^{\infty} \left(U g_n \setminus \bigcup_{i < n} U g_i \right) = G. \end{aligned}$$

Therefore, $\Gamma \cdot D = G$. Next, suppose that

$$x \in \gamma D \cap D.$$

Then, for some n, m

$$x \in W_n \cap \gamma W_m.$$

If $m < n$, then

$$\gamma W_m \subset \Gamma \bigcup_{i < n} U g_i,$$

which is disjoint from W_n , a contradiction. Thus, $W_n \cap \gamma W_m = \emptyset$ for $m < n$ and all $\gamma \in \Gamma$. If $n < m$, then

$$W_n \cap \gamma W_m = \gamma^{-1} (\gamma W_n \cap W_m) = \emptyset.$$

Thus, $n = m$, which implies that

$$U g_n \cap \gamma U g_n \neq \emptyset \Rightarrow U \cap \gamma U \neq \emptyset \Rightarrow \gamma = e.$$

Therefore, for all $\gamma \in \Gamma \setminus \{e\}$, $\gamma D \cap D = \emptyset$. \square

If $\Gamma < G$ is a discrete subgroup, then the left Haar measure μ on G descends to a Borel measure $\bar{\mu}$ the quotient space $Q = \Gamma \backslash G$: If $A \subset G$ is a Borel subset such that the restriction of the projection $p : G \rightarrow Q$ is injective on A , then $\bar{\mu}(p(A)) = \mu(A)$. The measure $\bar{\mu}$ can be defined using a measurable fundamental domain D in G as:

$$\bar{\mu}(\Gamma A) := \mu(A \cap D).$$

Note that when G is unimodular, the measure $\bar{\mu}$ is invariant under the right action of G .

EXERCISE 5.63. Prove that $\bar{\mu}$ is independent of the choice of a measurable fundamental set D .

If G is a Lie group, then the measure $\bar{\mu}$ (up to a scalar multiple) can also be described by using the volume form of the projection to Q of a left-invariant Riemannian metric on G .

DEFINITION 5.64. Let G be a locally compact Hausdorff topological group and μ a left Haar measure on G . A *lattice* in G is a discrete subgroup $\Gamma < G$ so that the quotient $Q = \Gamma \backslash G$ has finite measure, $\bar{\mu}(Q) < \infty$. A lattice Γ is called *uniform* if the quotient Q is compact.

EXERCISE 5.65. If G is a Lie group acting transitively and faithfully on a Riemannian manifold X , then $\Gamma < G$ is a lattice if and only if the quotient space X/Γ has finite volume.

THEOREM 5.66. *A locally compact second countable Hausdorff group G is unimodular, provided that it contains a lattice.*

PROOF. For arbitrary $g \in G$ consider the push-forward $\nu = R_g(\mu)$ of the (left) Haar measure μ on G ; here R_g is the right multiplication by g :

$$\nu(E) = \mu(Eg^{-1}).$$

Then ν is also a left Haar measure on G . By the uniqueness of the Haar measure, $\nu = c\mu$ for some constant $c > 0$.

Let $\Gamma < G$ be a lattice and let $D \subset G$ be its measurable fundamental set. Then

$$0 < \mu(D) = \mu(\Gamma \backslash G) < \infty$$

since Γ is a lattice. For every $g \in G$, Dg is again a measurable fundamental set for Γ and, thus, $\mu(D) = \mu(Dg)$. Hence,

$$\mu(D) = \mu(Dg) = c\nu(D).$$

It follows that $c = 1$. Thus, μ is also a right Haar measure. \square

5.6.5. Geometric actions. Suppose now that X is a metric space. We will equip the group of isometries $\text{Isom}(X)$ of X with the *compact-open topology*, equivalent to the topology of uniform convergence on compact sets. A subgroup $G \leq \text{Isom}(X)$ is called *discrete* if it is discrete with respect to the subset topology.

EXERCISE 5.67. Suppose that X is proper. Show that the following are equivalent for a subgroup $G < \text{Isom}(X)$:

- G is discrete.
- The action $G \curvearrowright X$ is properly discontinuous.
- For every $x \in X$ and an infinite sequence of distinct elements $g_i \in G$,

$$\lim_{i \rightarrow \infty} d(x, g_i(x)) = \infty.$$

Hint: Use Arzela–Ascoli theorem.

DEFINITION 5.68. A *geometric action* of a group G on a metric space X is an isometric properly discontinuous cobounded action $G \curvearrowright X$.

For instance, if X is a homogeneous Riemannian manifold with the isometry group G and $\Gamma < G$ is a uniform lattice, then Γ acts geometrically on X . Note that every geometric action on a proper metric space is cocompact.

LEMMA 5.69. *Suppose that a group G acts geometrically on a proper metric space X . Then $G \backslash X$ has a metric defined by*

$$(5.3) \quad \text{dist}(\bar{a}, \bar{b}) = \inf\{\text{dist}(p, q) ; p \in Ga, q \in Gb\} = \inf\{\text{dist}(a, q) ; q \in Gb\},$$

where $\bar{a} = Ga$ and $\bar{b} = Gb$.

Moreover, this metric induces the quotient topology of $G \backslash X$.

PROOF. The infimum in (5.3) is attained, i.e. there exists $g \in G$ such that

$$\text{dist}(\bar{a}, \bar{b}) = \text{dist}(a, gb).$$

Indeed, take $g_0 \in G$ arbitrary, and let $R = \text{dist}(a, gb)$. Then

$$\text{dist}(\bar{a}, \bar{b}) = \inf\{\text{dist}(a, q) ; q \in Gb \cap \overline{B}(a, R)\}.$$

Now, for every $gb \in \overline{B}(a, R)$,

$$gg_0^{-1}\overline{B}(a, R) \cap \overline{B}(a, R) \neq \emptyset.$$

Since G acts properly discontinuously on X , this implies that the set $Gb \cap \overline{B}(a, R)$ is finite, hence the last infimum is over a finite set, and it is attained. We leave it to the reader to verify that dist is the Hausdorff distance between the orbits $G \cdot a$ and $G \cdot b$. Clearly the projection $X \rightarrow G \backslash X$ is a contraction. One can easily check that the topology induced by the metric dist on $G \backslash X$ coincides with the quotient topology. \square

5.7. Zariski topology and algebraic groups

The proof of the Tits Alternative relies in part on some basic results from theory of affine algebraic groups. We recall some terminology and results needed in the argument. For a more thorough presentation, see [Hum75] and [OV90].

The proof of the following general lemma is straightforward, and left as an exercise to the reader.

LEMMA 5.70. *For every commutative ring A the following two statements are equivalent:*

- (1) *every ideal in A is finitely generated;*
- (2) *the set of ideals satisfies the ascending chain condition (ACC), that is, every ascending chain of ideals*

$$I_1 \subseteq I_2 \subseteq \cdots \subseteq I_n \subseteq \cdots$$

stabilizes, i.e. there exists an integer N such that $I_n = I_N$ for every $n \geq N$.

DEFINITION 5.71. A commutative ring is called *noetherian*, if it satisfies one (hence, both) statements in Lemma 5.70.

Note that a field seen as a ring is always noetherian. Other examples of noetherian rings come from the following theorem:

THEOREM 5.72 (Hilbert's ideal basis theorem, see [DF04]). *If A is a noetherian ring then the ring of multivariable polynomials $A[X_1, \dots, X_n]$ is also noetherian.*

From now on, we fix a field \mathbb{K} .

DEFINITION 5.73. An *affine algebraic set* in \mathbb{K}^n is a subset Z in \mathbb{K}^n that is the solution set of a system of multivariable polynomial equations $p_j = 0$, $\forall j \in J$, with coefficients in \mathbb{K} :

$$Z = \{(x_1, \dots, x_n) \in \mathbb{K}^n ; p_j(x_1, \dots, x_n) = 0, j \in J\}.$$

We will frequently say “algebraic subset” or “affine variety” when referring to an affine algebraic set.

For instance, algebraic subsets in the affine line (the 1-dimensional vector space V over \mathbb{K}) are finite subsets and the entire of V , since every non-zero polynomial in one variable has at most finitely many zeroes.

There is a one-to-one map associating to every algebraic subset in \mathbb{K}^n an ideal in $\mathbb{K}[X_1, \dots, X_n]$:

$$Z \mapsto I_Z = \{p \in \mathbb{K}[X_1, \dots, X_n] ; p|_Z \equiv 0\}.$$

Note that I_Z is the kernel of the homomorphism $p \mapsto p|_Z$ from $\mathbb{K}[X_1, \dots, X_n]$ to the ring of functions on Z . Thus, the ring $\mathbb{K}[X_1, \dots, X_n]/I_Z$ may be seen as a ring of functions on Z ; this quotient ring is called the *coordinate ring of Z* or the *ring of polynomials on Z* , and denoted $\mathbb{K}[Z]$.

Theorem 5.72 and Lemma 5.70 imply the following.

LEMMA 5.74. (1) *The set of algebraic subsets of \mathbb{K}^n satisfies the descending chain condition (DCC): every descending chain of algebraic subsets*

$$Z_1 \supseteq Z_2 \supseteq \dots \supseteq Z_i \supseteq \dots$$

stabilizes, i.e. for some integer $N \geq 1$, $Z_i = Z_N$ for every $i \geq N$.

(2) *Every algebraic set is defined by finitely many equations.*

DEFINITION 5.75. A *morphism* between two affine varieties Y in \mathbb{K}^n and Z in \mathbb{K}^m is a map of the form $\varphi : Y \rightarrow Z$, $\varphi = (\varphi_1, \dots, \varphi_m)$, such that each φ_i is in $\mathbb{K}[Y]$, $i \in \{1, 2, \dots, m\}$.

Note that every morphism is induced by a morphism $\tilde{\varphi} : \mathbb{K}^n \rightarrow \mathbb{K}^m$, $\tilde{\varphi} = (\tilde{\varphi}_1, \dots, \tilde{\varphi}_m)$, with $\tilde{\varphi}_i : \mathbb{K}^n \rightarrow \mathbb{K}$ a polynomial function for every $i \in \{1, 2, \dots, m\}$.

An *isomorphism* between two affine varieties Y and Z is an invertible map $\varphi : Y \rightarrow Z$ such that both φ and φ^{-1} are morphisms. When $Y = Z$, an isomorphism is called an *automorphism*.

EXERCISE 5.76. 1. If $f : Y \rightarrow Z$ is a morphism of affine varieties and $W \subset Z$ is a subvariety, then $f^{-1}(W)$ is a subvariety in Y . In particular, every linear automorphism of $V = \mathbb{K}^n$ sends subvarieties to subvarieties and, hence, the notion of a subvariety is independent of the choice of a basis in V .

2. Show that the projection map $f : \mathbb{C}^2 \rightarrow \mathbb{C}$, $f(x, y) = x$, does not map subvarieties to subvarieties.

Let V be an n -dimensional vector space over a field \mathbb{K} . The *Zariski topology* on V is the topology having as closed sets all the algebraic subsets in V . It is clear that the intersection of algebraic subsets is again an algebraic subset. Let $Z = Z_1 \cup \dots \cup Z_\ell$ be a finite union of algebraic subsets, Z_i defined by the ideal I_{Z_i} . Then Z is defined by the ideal

$$I_Z = I_{Z_1} \cdot \dots \cdot I_{Z_\ell}$$

generated by the products

$$\prod_{i=1}^{\ell} p_i$$

of elements $p_i \in I_{Z_i}$.

The induced topology on a subvariety $Z \subseteq V$ is also called the Zariski topology. Note that this topology can also be defined directly using polynomial functions in $\mathbb{K}[Z]$. According to Exercise 5.76, morphisms between affine varieties are continuous with respect to the Zariski topologies.

The *Zariski closure* of a subset $E \subset V$ can also be defined by means of the set P_E of all polynomials which vanish on E , i.e. it coincides with

$$\{x \in V \mid p(x) = 0, \forall p \in P_E\}.$$

A subset $Y \subset Z$ in an affine variety is called *Zariski-dense* if its Zariski closure is the entire of Z .

Lemma 5.74, Part (1), implies that closed subsets in Zariski topology satisfy the descending chain condition (DCC).

DEFINITION 5.77. A topological space such that the closed sets satisfy the DCC is called *noetherian*.

LEMMA 5.78. *Every subspace of a noetherian topological space (with the subspace topology) is noetherian.*

PROOF. Let X be a space with topology \mathcal{T} such that (X, \mathcal{T}) is noetherian, and let Y be an arbitrary subset in X . Consider a descending chain of closed subsets in Y :

$$Z_1 \supseteq Z_2 \supseteq \dots \supseteq Z_n \supseteq \dots$$

Every Z_i is equal to $Y \cap C_i$ for some closed set C_i in X . We leave it to the reader to check that C_i can be taken equal to the closure $\overline{Z_i}$ of Z_i in X .

The descending chain of closed subsets in X ,

$$\overline{Z_1} \supseteq \overline{Z_2} \supseteq \dots \supseteq \overline{Z_n} \supseteq \dots$$

stabilizes, hence, so does the chain of the subsets Z_i . \square

PROPOSITION 5.79. *Every noetherian topological space X is compact.*

PROOF. Compactness of X is equivalent to the condition that for every family $\{Z_i : i \in I\}$ of closed subsets in X , if $\bigcap_{i \in I} Z_i = \emptyset$, then there exists a finite subset J of I such that $\bigcap_{j \in J} Z_j = \emptyset$. Assume that all finite intersections of a family as above are non-empty. Then we construct inductively a descending sequence of closed sets that never stabilizes. The initial step consists of picking an arbitrary set Z_{i_1} , with $i_1 \in I$. At the n th step we have a non-empty intersection $Z_{i_1} \cap Z_{i_2} \cap \dots \cap Z_{i_n}$; hence, there exists $Z_{i_{n+1}}$ with $i_{n+1} \in I$ such that $Z_{i_1} \cap Z_{i_2} \cap \dots \cap Z_{i_n} \cap Z_{i_{n+1}}$ is a non-empty proper closed subset of $Z_{i_1} \cap Z_{i_2} \cap \dots \cap Z_{i_n}$. \square

We now discuss a strong version of connectedness, relevant in the setting of noetherian spaces.

LEMMA 5.80. *For a topological space X the following properties are equivalent:*

- (1) *every open non-empty subset of X is dense in X ;*
- (2) *any two open non-empty subsets have non-empty intersection;*

(3) X cannot be written as a finite union of proper closed subsets.

We leave the proof of this lemma as an exercise to the reader.

DEFINITION 5.81. A topological space is called *irreducible* if it is non-empty and one of (hence all) the properties in Lemma 5.80 is (are) satisfied. A subset of a topological space is *irreducible* if, when endowed with the subset topology, it is an irreducible space.

EXERCISE 5.82. (1) Prove that \mathbb{K}^n with Zariski topology is irreducible, provided that the field \mathbb{K} is infinite.

(2) Prove that an algebraic variety Z is irreducible if and only if $\mathbb{K}[Z]$ does not contain zero divisors.

The following properties are straightforward and their proof is left as an exercise to the reader.

LEMMA 5.83. (1) *The image of an irreducible space under a continuous map is irreducible.*

(2) *The cartesian product of two irreducible spaces is an irreducible space, when endowed with the product topology.*

Note that the Zariski topology on $\mathbb{K}^{n+m} = \mathbb{K}^n \times \mathbb{K}^m$ is *not* the product topology (unless $nm = 0$ or \mathbb{K} is finite). Hence, irreducibility of products of irreducible varieties cannot be derived from Lemma 5.83.

LEMMA 5.84. *Let V_1, V_2 be finite-dimensional vector spaces over \mathbb{K} and let $Z_i \subset V_i, i = 1, 2$, be irreducible subvarieties. Then the product $Z := Z_1 \times Z_2 \subset V = V_1 \times V_2$ is an irreducible subvariety in the vector space V .*

PROOF. Let $Z = W_1 \cup W_2$ be a union of two proper subvarieties. For every $z \in Z_1$ the product $\{z\} \times Z_2$ is isomorphic to Z_2 (via the projection to the second factor) and, hence, is irreducible. On the other hand,

$$\{z\} \times Z_2 = ((\{z\} \times Z_2) \cap W_1) \cup ((\{z\} \times Z_2) \cap W_2)$$

is a union of two subvarieties. Thus, for every $z \in Z_1$, one of these subvarieties has to be the entire $\{z\} \times Z_2$. In other words, either $\{z\} \times Z_2 \subset W_1$ or $\{z\} \times Z_2 \subset W_2$. We then partition Z_1 in two subsets A_1, A_2 :

$$A_i = \{z \in Z_1 : \{z\} \times Z_2 \subset W_i\}, i = 1, 2.$$

Since each W_1, W_2 is a proper subvariety, both A_1, A_2 are proper subsets of Z_1 . We will now prove that both A_1, A_2 are subvarieties in Z_1 . We will consider the case of A_1 since the other case is obtained by relabeling. Let f_1, \dots, f_k denote generators of the ideal of W_1 . We will think of each f_i as a function of two variables $f = f(X_1, X_2)$, where X_k stands for the tuple of coordinates in $V_k, k = 1, 2$. Then

$$A_1 = \{z \in Z_2 : f_i(z, z_2) = 0, \forall z \in Z_1, i = 1, \dots, k\}.$$

However, for every fixed $z \in Z_1$, the function $f_i(z, \cdot)$ is a polynomial function $f_{i,z}$ on Z_2 . Therefore, A_1 is the solution set of the system of polynomial equations on Z_1 :

$$\{f_{i,z} = 0 : i = 1, \dots, k, z \in Z_1\}.$$

It follows that A_1 is a subvariety, which contradicts irreducibility of Z_2 . \square

LEMMA 5.85. Let (X, \mathcal{T}) be a topological space.

- (1) A subset Y of X is irreducible if and only if its closure \bar{Y} in X is irreducible.
- (2) If Y is irreducible and $Y \subseteq A \subseteq \bar{Y}$ then A is irreducible.
- (3) Every irreducible subset Y of X is contained in a maximal irreducible subset.
- (4) The maximal irreducible subsets of X are closed and they cover X .

PROOF. (1) For every open subset $U \subset X$, $U \cap Y \neq \emptyset$ if and only if $U \cap \bar{Y} \neq \emptyset$. This and Lemma 5.80, (2), imply the equivalence.

(2) Follows directly from (1).

(3) The family \mathcal{I}_Y of irreducible subsets containing Y has the property that every ascending chain has a maximal element, which is the union. It can be easily verified that the union is again irreducible, using Lemma 5.80, (2). It follows from Zorn's Lemma that \mathcal{I}_Y contains a maximal element.

(4) Singletons are irreducible and cover X . Now, the statement follows from (1) and (3), since \square

THEOREM 5.86. A noetherian topological space X is a union of finitely many distinct maximal irreducible subsets X_1, X_2, \dots, X_n , such that for every i , X_i is not contained in $\bigcup_{j \neq i} X_j$. Moreover, every maximal irreducible subset in X coincides with one of the subsets X_1, X_2, \dots, X_n . This decomposition of X is unique up to a renumbering of the X_i 's.

PROOF. Let \mathcal{F} be the collection of closed subsets of X that cannot be written as a finite union of maximal irreducible subsets. Assume that \mathcal{F} is non-empty. Since X is noetherian, \mathcal{F} satisfies the DCC, hence by Zorn's Lemma it contains a minimal element Y . As Y is not irreducible, it can be decomposed as $Y = Y_1 \cup Y_2$, where Y_i are closed and, by the minimality of Y , both Y_i decompose as finite unions of irreducible subsets (maximal in Y_i). According to Lemma 5.85, (3), Y itself can be written as union of finitely many maximal irreducible subsets, a contradiction. It follows that \mathcal{F} is empty.

If $X_i \subseteq \bigcup_{j \neq i} X_j$ then $X_i = \bigcup_{j \neq i} (X_j \cap X_i)$. As X_i is irreducible it follows that $X_i \subseteq X_j$ for some $j \neq i$, hence, by maximality, $X_i = X_j$, contradicting the fact that we took only distinct maximal irreducible subsets. A similar argument is used to prove that every maximal irreducible subset of X must coincide with one of the sets X_i .

Now assume that X can be also written as a union of distinct maximal irreducible subsets Y_1, Y_2, \dots, Y_m such that for every i , Y_i is not contained in $\bigcup_{j \neq i} Y_j$. For every $i \in \{1, 2, \dots, m\}$ there exists a unique $j_i \in \{1, 2, \dots, n\}$ such that $Y_i = X_{j_i}$. The map $i \mapsto j_i$ is injective, and if some $k \in \{1, 2, \dots, n\}$ is not in the image of this map then it follows that

$$X_k \subseteq \bigcup_{i=1}^m Y_i \subseteq \bigcup_{j \neq k} X_j,$$

a contradiction. \square

DEFINITION 5.87. The subsets X_i defined in Theorem 5.86 are called *the irreducible components* of X . In other words, irreducible components of X are maximal irreducible subsets of X .

Note that we can equip every Zariski-open subset U of a (finite-dimensional) vector space V with the Zariski topology, which is the subset topology with respect to the Zariski topology on V . Then U is also Noetherian. We will be using the Zariski topology primarily in the context of the group $GL(V)$, which we identify with the Zariski open subset of $V \otimes V^*$, the space of $n \times n$ matrices with non-zero determinant.

DEFINITION 5.88. An *algebraic subgroup* of $GL(V)$ is a Zariski-closed subgroup of $GL(V)$.

Given an algebraic subgroup G of $GL(V)$, the binary operation $G \times G \rightarrow G$, $(g, h) \mapsto gh$ is a morphism. The inversion map $g \mapsto g^{-1}$, as well as the left-multiplication and right-multiplication maps $g \mapsto ag$ and $g \mapsto ga$, by a fixed element $a \in G$, are automorphisms of the variety G .

EXAMPLE 5.89. (1) The subgroup $SL(V)$ of $GL(V)$ is algebraic, defined by the equation $\det(g) = 1$.

(2) The group $GL(n, \mathbb{K})$ can be identified with an algebraic subgroup of the group $SL(n+1, \mathbb{K})$ by mapping every matrix $A \in GL(n, \mathbb{K})$ to the matrix

$$\begin{pmatrix} A & 0 \\ 0 & \frac{1}{\det(A)} \end{pmatrix}.$$

Therefore, in what follows, it will not matter if we consider algebraic subgroups of $GL(n, \mathbb{K})$ or of $SL(n, \mathbb{K})$.

(3) The group $O(V)$ is an algebraic subgroup, as it is given by the system of equations $M^T M = \text{Id}_V$.

(4) More generally, given an arbitrary quadratic form q on V , its stabilizer $O(q)$ is obviously algebraic. A special instance of this is the *symplectic group* $Sp(2k, \mathbb{K})$, preserving the form with the following matrix (given with respect to the standard basis in $V = \mathbb{K}^{2n}$)

$$J = \begin{pmatrix} 0 & K \\ -K & 0 \end{pmatrix}, \text{ where } K = \begin{pmatrix} 0 & \dots & 1 \\ 0 & \ddots & 0 \\ 1 & \dots & 0 \end{pmatrix}.$$

LEMMA 5.90. If Γ is a subgroup of an algebraic group G , then its Zariski closure $\bar{\Gamma}$ in G is also a subgroup.

PROOF. We let $\mu : G \times G \rightarrow G, \lambda : G \rightarrow G$ denote the multiplication and inversion maps respectively. Both maps are continuous if we equip $G \times G, G$ with their respective Zariski topologies. Therefore, for each subset $E \subset G$,

$$\lambda(\bar{E}) \subset \overline{\lambda(E)},$$

which implies that

$$\lambda(\bar{\Gamma}) \subset \bar{\Gamma}.$$

Hence, $\bar{\Gamma}$ is closed under the inversion. Similarly, for each $g \in \Gamma$,

$$\mu(g, \bar{\Gamma}) \subset \bar{\Gamma}$$

and, thus, for each $h \in \bar{\Gamma}$,

$$\mu(\bar{\Gamma}, h) \subset \bar{\Gamma}.$$

It follows that $\mu(\bar{\Gamma}, \bar{\Gamma}) \subset \bar{\Gamma}$. □

When $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$, the vector space $V = \mathbb{K}^n$ also has the *standard* or *classical* topology, given by the suitable norm on V . We use the terminology *classical topology* for the induced topology on subsets of V . Classical topology, of course, is stronger than Zariski topology.

THEOREM 5.91 (See for instance Chapter 3, §2, in [OV90]). (1) *An algebraic subgroup of $GL(n, \mathbb{C})$ is irreducible in the Zariski topology if and only if it is connected in the classical topology.*

(2) *A connected (in classical topology) algebraic subgroup of $GL(n, \mathbb{R})$ is irreducible in the Zariski topology.*

We will not need this theorem; the following proposition will suffice for our purposes:

PROPOSITION 5.92. *Let G be an algebraic subgroup in $GL(V)$.*

(1) *Only one irreducible component of G contains the identity element. This component is called the identity component and is denoted by G_0 .*

(2) *The subset G_0 is a normal subgroup of finite index in G whose cosets are the irreducible components of G .*

PROOF. (1) Let X_1, \dots, X_k be irreducible components of G containing the identity. According to Lemma 5.84, the product set $X_1 \times \dots \times X_k$ is irreducible. Since the product map is a morphism, the subset $X_1 \cdots X_k \subset G$ is irreducible as well; hence by Lemma 5.85, (3), and by Theorem 5.86 this subset is contained in some X_j . The fact that every X_i with $i \in \{1, \dots, k\}$ is contained in $X_1 \cdots X_k$, hence in X_j , implies that $k = 1$.

(2) Since the inversion map $g \mapsto g^{-1}$ is an algebraic automorphism of G (but not a group automorphism, of course) it follows that G_0 is stable with respect to the inversion. Hence for every $g \in G_0$, gG_0 contains the identity element, and is an irreducible component. It follows that $gG_0 = G_0$. Likewise, for every $x \in G$, xG_0x^{-1} is an irreducible component containing the identity element, hence it equals G_0 . The cosets of G_0 (left or right) are images of G_0 under automorphisms, therefore also irreducible components. Thus, there can only be finitely many of them. □

REMARK 5.93. Proposition 5.92, (2), implies that for algebraic groups the irreducible components are disjoint. This is not true in general for algebraic varieties, consider, for instance, the subvariety $\{xy = 0\} \subset \mathbb{K}^2$.

We now relate Lie groups and algebraic groups. For $\mathbb{F} = \mathbb{R}$ or \mathbb{C} , each algebraic subgroup $G < GL(n, \mathbb{F})$ is necessarily closed in the classical topology, hence, is a (complex, resp. real) Lie subgroup of $GL(n, \mathbb{F})$ by Theorem 5.57. Below is a simpler argument which does not rely upon Cartan's theorem.

THEOREM 5.94. *Each algebraic subgroup $G < GL(n, \mathbb{F})$ is a Lie subgroup.*

PROOF. Since G is a subgroup, it is a (real or complex) submanifold in $GL(n, \mathbb{F})$ iff it contains a non-empty open subset which is a submanifold in $GL(n, \mathbb{F})$. The subgroup G is the zero-set of a polynomial map $p : GL(n, \mathbb{F}) \rightarrow \mathbb{F}^k$. Let r denote the maximum of ranks of the derivative dp of p on G , it is attained on an open non-empty subset U of G . Let V denote the subset of $GL(n, \mathbb{F})$ where dp has rank r . By the constant rank theorem (Theorem 3.2), V is a smooth submanifold of $GL(n, \mathbb{F})$. Applying the constant rank theorem to the restriction $p : V \rightarrow \mathbb{F}^k$, we conclude that $U = G \cap V$ is a smooth submanifold in V and, hence, in $GL(n, \mathbb{F})$. \square

5.8. Group actions on complexes

5.8.1. G -complexes. Let G be a (discrete) topological group and let X be a cell complex, defined via disjoint unions of balls U_n and attaching maps e^n , see section 1.7.2. We say that X is a G -complex, or that we have a *cellular action* $G \curvearrowright X$, if $G \times X \rightarrow X$ is a topological action and for every n we have a G -action $G \curvearrowright U_n$ such that the attaching map

$$e^n : \partial U_n \rightarrow X^{(n-1)}$$

is G -equivariant.

DEFINITION 5.95. A cellular action $G \curvearrowright X$ is said to be *without inversions* if whenever $g \in G$ preserves a cell s in X , it fixes this cell pointwise.

A topological action $G \curvearrowright X$ on the geometric realization of a simplicial complex is called *simplicial* if it sends simplices to simplices and is affine on each simplex. As with cellular actions, a simplicial complex equipped with a simplicial group action is called a *simplicial G -complex*.

Equivalently, one describes simplicial G -complexes as follows. Let $G \curvearrowright V(X)$ be a (set-theoretic) action of G on the vertex set of a simplicial complex X , which sends simplices to simplices. Then this action defines a simplicial action of G on X . (Use the canonical affine extension of $G \curvearrowright V(X)$ to the geometric realization of each simplex.)

The following is immediate from the definition of X'' , since barycentric subdivisions are canonical:

LEMMA 5.96. *Let X be an almost regular cell complex and let $G \curvearrowright X$ be an action without inversions. Then $G \curvearrowright X$ induces a simplicial action without inversions $G \curvearrowright X''$.*

EXERCISE 5.97. Given a cellular action $G \curvearrowright \tilde{X}$ on a cell complex \tilde{X} , there exists a simplicial complex \tilde{Z} , a simplicial action without inversions $G \curvearrowright \tilde{Z}$ and a G -equivariant homotopy-equivalence $\tilde{X} \rightarrow \tilde{Z}$. Moreover, if \tilde{X} is finite-dimensional, then \tilde{Z} can be also taken finite-dimensional. Hint: Follow the proof of the fact that every cell complex is homotopy-equivalent to a simplicial complex.

LEMMA 5.98. *Let X be a simplicial complex and $G \curvearrowright X$ be a free simplicial action. Then this action is properly discontinuous on X (in the weak topology).*

PROOF. Let K be a compact in X . Then K is contained in a finite union of simplices $\sigma_1, \dots, \sigma_k$ in X . Let $F \subset G$ be the subset consisting of elements $g \in G$ so that $gK \cap K \neq \emptyset$. Then, assuming that F is infinite, it contains distinct elements g, h such that $g(\sigma) = h(\sigma)$ for some $\sigma \in \{\sigma_1, \dots, \sigma_n\}$. Then $f := h^{-1}g(\sigma) = \sigma$. Since the action $G \curvearrowright X$ is linear on each simplex, f fixes a point in σ . This contradicts the assumption that the action of G on X is free. \square

5.8.2. Borel and Haefliger constructions. Every group G admits a *classifying space* $E(G)$, which is a contractible cell complex admitting a free cellular action $G \curvearrowright E(G)$. The space $E(G)$ is far from being unique, we will use the one obtained by the *Milnor's Construction*, see for instance [Hat02, Section 1.B]. A benefit of this construction is that $E(G)$ is a simplicial complex and the construction of $G \curvearrowright E(G)$ is canonical. Simplices in $E(G)$ are ordered tuples of elements of $g: [g_0, \dots, g_n]$ is an n -simplex, with the obvious inclusions of simplices. To verify the contractibility of $E = E(G)$, note that for every $i \geq k \geq 0$ the map

$$\pi_k(E^i) \rightarrow \pi_k(E^{i+1})$$

is trivial. Here and in what follows, E^i is the i -skeleton of $E = E(G)$.

The group G acts on $E(G)$ by the left multiplication

$$g \cdot [g_0, \dots, g_n] = [gg_0, \dots, gg_n].$$

Clearly, this action is free and, moreover, each simplex has trivial stabilizer. The action $G \curvearrowright E(G)$ has two obvious properties that we will be using:

- (1) If the group G is finite, then each skeleton E^i of $E(G)$ is a finite simplicial complex.
- (2) Every monomorphism $G_1 \hookrightarrow G_2$, induces a canonical equivariant embedding $E(G_1) \hookrightarrow E(G_2)$.

Suppose now that \tilde{X} is a cell complex and $G \curvearrowright \tilde{X}$ is a cellular action without inversions. The main goal of this section is to replace this action with a *free* cellular action $G \curvearrowright \hat{X}$ on a new cell complex \hat{X} such that there exists a G -equivariant homotopy-equivalence $\tilde{X} \rightarrow \hat{X}$. We will describe two constructions of complexes \hat{X} : the *Borel construction* and the *Haefliger construction*. The second will be the most useful to us; it first appeared in Haefliger's paper [Hae92].

First, we consider the product of \tilde{X} with the classifying space $E(G)$. The group G acts on $E(G) \times \tilde{X}$ diagonally. The product space $E(G) \times \tilde{X}$ equipped with the G -action is called the *Borel Construction*. We will use the notation B for the quotient

$$B = (E(G) \times \tilde{X})/G.$$

This space has a natural projection $p: B \rightarrow X = \tilde{X}/G$ coming from the coordinate projection $\tilde{p}: E(G) \times \tilde{X} \rightarrow \tilde{X}$. The product $E(G) \times \tilde{X}$ is a cell complex. The action of G on the product is free since G acts freely on the first factor. For every open cell σ in X , the fiber $p^{-1}(\sigma) \subset B$ is naturally homeomorphic to the quotient

$$(E(G) \times \tilde{\sigma})/G_{\tilde{\sigma}},$$

where $\tilde{\sigma}$ is a component of the preimage of σ in \tilde{X} and $G_{\tilde{\sigma}}$ is the stabilizer of $\tilde{\sigma}$ in G (the conjugacy class of this subgroup of G is independent of the choice of $\tilde{\sigma}$).

In view of Exercise 5.97, we will assume that \tilde{X} is a simplicial complex and $G \curvearrowright \tilde{X}$ is a simplicial action without inversions. The product $E(G) \times \tilde{X}$ then is a regular cell complex.

The second construction (due to Haefliger) is considerably more complicated. We summarize it in the following theorem:

THEOREM 5.99. *There exists a regular cell complex \hat{X} , a G -action without inversions on \hat{X} and a projection $\hat{q}: \hat{X} \rightarrow \tilde{X}$ such that:*

1. There exists a G -equivariant homotopy-equivalence $\widehat{h} : \widehat{X} \rightarrow E(G) \times \widetilde{X}$ such that

$$\tilde{p} \circ \widehat{h} = \widehat{q}.$$

2. For every open simplex $\tilde{\sigma}$ in the barycentric subdivision \widetilde{X}' of \widetilde{X} , $\widehat{q}^{-1}(\tilde{\sigma})$ is equivariantly isomorphic to the product $E(G_{\tilde{\sigma}}) \times \tilde{\sigma}$, where $G_{\tilde{\sigma}}$ is the stabilizer of $\tilde{\sigma}$ in G .

PROOF. We will construct \widehat{X} as a suitable covering space of a certain complex Y ; the map \widehat{h} will be a lift of a homotopy-equivalence $h : Y \rightarrow B$.

Set $X := \widetilde{X}/G$; it is an almost regular cell complex where cells are simplices. Let X' denote the barycentric subdivision of X . Then the vertices of X' are in bijective correspondence with the faces of X . The inclusion of faces of X induces a natural orientation of the edges of X' , where an edge $[u, v]$ is oriented whenever $\sigma \supset \tau$ with σ, τ the faces with the barycenters u, v respectively. Observe that the vertex set $V(X')$ of X' forms a natural poset where $u < v$ if and only if there exists an oriented edge $[u, v]$ in X' . Given a face σ of X' we let $\min(\sigma)$ and $\max(\sigma)$ denote, respectively, the minimal and the maximal vertices of σ . For every vertex $u \in X'$ corresponding to a simplex σ in X , we pick a simplex $\tilde{\sigma} \subset \widetilde{X}$ projecting to σ . Let $G_\sigma < G$ denote the stabilizer of $\tilde{\sigma}$; it equals the stabilizer G_v of a lift of v to $\tilde{\sigma}$. Then for each vertex v of X' we have a distinguished subcomplex $E(G_v) \subset E(G)$. Projecting $E(G_v)$ to B we obtain a subcomplex

$$X_v = E(G_v)/G_v \subset p^{-1}(v) \subset B = (E(G) \times \widetilde{X})/G,$$

with $\pi_1(X_v) \simeq G_v$. This subcomplex has a distinguished vertex x_v , the projection of the vertex of $E(G_v)$ defined by the neutral element 1 of the group G_v .

We fix the following data in the complex B . For each vertex v of X' we pick a base-point

$$x_v \in X_v \subset p^{-1}(v) \subset B.$$

We also pick a base-vertex v_0 of X' , set $x_0 := x_{v_0}$. Since $E(G) \times \widetilde{X} \rightarrow B$ is a covering map, we will identify the group G with a quotient of the fundamental group $\pi_1(B, x_0)$,

$$\pi_1(B, x_0) \rightarrow G.$$

We pick a maximal tree T in the 1-skeleton of the complex X' . We construct a (partial) section $s : T \rightarrow B$ of $p : B \rightarrow X$ so that for every vertex v of X' , $s(v) = x_v$ belongs to the complex X_v . In particular, for every vertex v of X' we obtain a path ξ_v in B connecting x_0 to x_v : This path equals the image under s of the geodesic path in T connecting o to v . Lastly, we extend the section s to every edge e of X' which is not in T . As the result, for every oriented edge ϵ of X' we obtain a path $\eta_\epsilon := s(\epsilon)$ in B . Furthermore, the system of paths ξ_v and η_ϵ , determines group isomorphisms

$$\pi_1(p^{-1}(v), x_v) \simeq G_v < G,$$

as well as isomorphisms of simplicial complexes $p^{-1}(v) \cong E(G)/G_v$, sending x_v to the projection of $[1] \in E(G)$.

The choice of the paths ξ and η yields the following elements of G , homomorphisms of groups and simplicial maps:

- i. For each oriented edge $\epsilon = [u, v]$ of X' we have an element $g_\epsilon \in G_v$ represented by the concatenation of paths

$$(5.4) \quad \xi_u \star \eta_\epsilon \star \xi_v^{-1}.$$

It is convenient to set

$$g_{\bar{\epsilon}} = g_{[v, u]} := g_{[u, v]}^{-1} = g_\epsilon^{-1}.$$

- ii. The conjugation by g_ϵ ,

$$(5.5) \quad \psi_\epsilon : g \mapsto g_\epsilon^{-1} g g_\epsilon$$

defines a monomorphism $\psi_\epsilon : G_u \rightarrow G_v$. More specifically, given a loop λ in X_u , based at x_u , and representing $g \in G_u$, the image of g in G_v is represented by the concatenation

$$s(\epsilon)^{-1} \star \lambda \star s(\epsilon).$$

- iii. For each oriented edge $\epsilon = [u, v]$ of X' we have a canonical simplicial embedding³

$$(5.6) \quad \Psi_\epsilon : (X_u, x_u) \rightarrow (X_v, x_v)$$

which induces the monomorphism $\psi_\epsilon : G_u \rightarrow G_v$.

- iv. Given an oriented 2-dimensional simplex $\tau = [v_1, v_2, v_3]$ in X' with edges $\alpha = [v_1, v_2], \beta = [v_2, v_3], \gamma = [v_1, v_3]$, there is no reason to expect g_γ to be the product of g_α and g_β . Thus, define the *monodromy*

$$(5.7) \quad g_\tau = g_\gamma^{-1} g_\alpha g_\beta = g_{[v_3, v_1]} g_{[v_1, v_2]} g_{[v_2, v_3]} \in G_{v_3}.$$

The element g_τ is represented by the concatenation of paths

$$\xi_{v_3} \star \eta_{\gamma^{-1}} \star \eta_\alpha \star \eta_\beta \star \xi_{v_3}^{-1}.$$

The reader can think of the map $\tau \mapsto g_\tau$, where τ runs over the two-dimensional faces of X' , as a *nonabelian cocycle* on the simplicial complex X' .

- v. We also define *monodromy maps*

$$(5.8) \quad \Psi_\tau : X_{v_1} \times [0, 1] \rightarrow X_{v_3},$$

such that:

- $\Psi_\tau(x, 0) = \Psi_\gamma(x)$, $x \in X_{v_1}$.
- $\Psi_\tau(x_{v_1}, t) = x_{v_3}$, $t \in \{0, 1\}$.
- The loop $\Psi_\tau(x_{v_1}, t)$, $t \in [0, 1]$ represents the element $g_\tau \in G_{v_3}$.

We will build a complex Y by attaching products $X_v \times I^n$ by induction on the dimension of the skeleta of X' . The attaching maps will be *guided* by certain maps θ of n -dimensional cubes I^n (n -fold products of the unit interval $I = [0, 1]$) to n -dimensional simplices in X .

We introduce some notation useful for constructing the maps θ and the attaching maps. For each $i = 1, \dots, n$ we have the following parallel facets of I^n :

$$I_i^- := I \times \dots \times I \times \{0\}_i \times I \times \dots \times I$$

³Recall that each group monomorphism $H_1 \rightarrow H_2$ defines a canonical simplicial embedding $E(H_1) \rightarrow E(H_2)$.

and

$$I_i^+ := I \times \dots \times I \times \{1\}_i \times I \times \dots \times I.$$

Here 0_i and 1_i means that 0 (resp. 1) appears in the i -th place. We let $\text{top}(I^n)$ denote the *top-facet* of I^n , namely, I_n^+ .

Furthermore, for $\sigma = [v_0, \dots, v_n]$ and $i = 1, \dots, n$, we let $\partial_i \sigma = [v_0, \dots, \hat{v}_i, \dots, v_n]$ denote the facet of σ obtained by skipping the vertex v_i .

We define maps

$$\theta_\sigma = \theta_{n,\sigma} : I^n \rightarrow \sigma = [v_0, \dots, v_n]$$

such that for $i = 1, \dots, n$

$$\theta_{n,\sigma} : I_i^- \rightarrow [v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_n],$$

and

$$\theta_{n,\sigma} : \text{top}(I^n) \rightarrow \{v_n\}.$$

We define these maps by induction on n . Namely, we start with the map

$$\{0\} = I^0 \rightarrow \{v_0\}.$$

Then, given a map

$$\theta_{k,\tau} : I^k \rightarrow \tau = [v_0, \dots, v_k],$$

we extend it to a map

$$\theta_{k+1,\rho} = I^{k+1} \rightarrow \rho = [v_0, \dots, v_{k+1}]$$

by sending the face $I^k \times \{1\}$ to v_{k+1} and then “coning off the map $\theta_{k,\tau}$, i.e. so that

$$\theta_{k+1,\rho}|_{I^k \times \{0\}} = \theta_{k,\tau}$$

and $\theta_{k+1,\rho}$ is linear on the vertical segments $r \times [0, 1]$, $r \in I^k$. From now on, we will suppress the subscript n for the maps $\theta_{n,\sigma}$ but will sometimes keep the subscript σ to indicate that θ_σ maps I^n to the simplex σ .

We now build a complex Y and a projection $q : Y \rightarrow X$.

For each vertex v of X' we have a pointed simplicial complex (X_v, x_v) isomorphic to $(E(G_v)/G_v, [1])$. Define

$$Y_0 := \coprod_{v \in V(X')} X_v, \quad q : Y \rightarrow X', \quad q(X_v) = \{v\}, \quad v \in V(X').$$

We then proceed inductively. Assume that the spaces Y_k , $k \leq n-1$, and the projections

$$q_k : Y_k \rightarrow X,$$

are constructed, as well as the maps

$$f_\sigma = f_{[v_0, \dots, v_k]} : X_{v_0} \times I^k \rightarrow Y_k, \quad \sigma = [v_0, \dots, v_k],$$

so that we have commutative diagrams

$$\begin{array}{ccc} X_{v_0} \times I^k & \xrightarrow{f_\sigma} & Y_k \\ \downarrow & & \downarrow q_k \\ I^k & \xrightarrow{\theta_\sigma} & \sigma \end{array}$$

We will construct a space Y_n and a projection $q_n : Y_n \rightarrow X'$. In order to construct Y_n , for every n -dimensional face $\sigma = [v_0, \dots, v_n]$ of X' we will need an *attaching map*

$$\partial f_\sigma : Y_{v_0} \times \partial I^n \rightarrow Y_{n-1},$$

whose image is contained in $q_{n-1}^{-1}(\theta_\sigma(\partial I^n))$. Then Y_n will be obtained as the quotient (by an equivalence relation \sim) of the disjoint union

$$Y_{n-1} \sqcup \coprod_{\sigma} X_{\min(\sigma)} \times I^n,$$

where the disjoint union is taken over all n -dimensional faces σ of X' and $\min(\sigma)$ is the minimal vertex of σ as defined earlier. The equivalence relation is given by

$$y \sim \partial f_\sigma(y), \quad y \in X_{\min(\sigma)} \times \partial I^n.$$

This will yield also maps

$$f_\sigma : X_{\min(\sigma)} \times I^n \rightarrow Y_n$$

extending the maps ∂f_σ .

For each n and each n -simplex $\sigma = [v_0, \dots, v_n]$ in X' , the map ∂f_σ will be defined on

$$X_{v_0} \times (\partial I^n - \text{top}(I^n))$$

by the same inductive formula independent of n . However, the definition of the map ∂f_σ on $\text{top}(I^n)$ will depend on n .

Definition of the attaching maps on $X_{v_0} \times (\partial I^n - \text{top}(I^n))$.

For each facet $I_i^-, i = 1, \dots, n$, we set

$$(5.9) \quad \begin{aligned} \partial f_\sigma(z, t_1, \dots, 0, \dots, t_n) &= f_{\partial_i \sigma}(z, t_1, \dots, t_{i-1}, t_{i+1}, \dots, t_n), \\ 0 \leq t_k \leq 1, k &= 1, \dots, i-1, i+1, \dots, n, z \in Y_{v_0}. \end{aligned}$$

In order to define the attaching map on the facet $I_i^+, i = 1, \dots, n-1$, we first note that, by the induction assumption, since $f_{[v_0, \dots, v_i]}$ is a lift of the map

$$\theta_{[v_0, \dots, v_i]} : I^i \rightarrow [v_0, \dots, v_i],$$

we have that

$$f_{[v_0, \dots, v_i]} : X_{v_0} \times \text{top}(I^i) \rightarrow X_{v_i}.$$

Furthermore, we also have an inductively defined map

$$f_{[v_i, \dots, v_n]} : X_{v_i} \times I^{n-i} \rightarrow Y_{n-i} \subset Y_{n-1}.$$

Therefore, we define the attaching map on $X_{v_0} \times I_i^+ = (X_{v_0} \times \text{top}(I^i)) \times I^{n-i}$, $i < n$, as the composition

$$(5.10) \quad \begin{array}{ccc} X_{v_0} \times I_i^+ & \xrightarrow{f_{[v_0, \dots, v_i]} \times \text{Id}} & X_{v_i} \times I^{n-i} \\ & \searrow & \downarrow f_{[v_i, \dots, v_n]} \\ & & Y_{n-i} \end{array}$$

We will introduce the attaching map on $X_{v_0} \times \text{top}(I^n)$ later; for now, let us check that the attaching maps introduced so far are well defined. For instance, consider the intersection of two facets

$$I_{ij}^{--} = I_i^- \cap I_j^-,$$

where $i < j < n$. We have

$$\begin{aligned} f_{\partial_i \sigma}(z, t_1, \dots, \hat{t}_i, \dots, 0_j, \dots, t_n) &= f_{\partial_i \sigma}(z, t_1, \dots, \hat{t}_i, \dots, \hat{t}_j, \dots, t_n) = \\ f_{\partial_j \sigma}(z, t_1, \dots, \hat{t}_i, \dots, \hat{t}_j, \dots, t_n) &= f_{\partial_j \sigma}(z, t_1, \dots, 0_i, \dots, \hat{t}_j, \dots, t_n). \end{aligned}$$

Therefore, ∂f_σ is well defined on I_{ij}^{--} .

Consider also the intersection of two facets

$$I_{ij}^{++} = I_i^+ \cap I_j^+, i < j.$$

We will compare two attaching maps on this face, which appear as restrictions of the attaching maps coming from the facets I_i^+ and I_j^+ . These restrictions are given by the formulae

$$(5.11) \quad f_{[v_i, \dots, v_n]}(f_{[v_0, \dots, v_i]}(z, t_1, \dots, 1_i), t_{i+1}, \dots, 1_j, \dots, t_n),$$

and

$$(5.12) \quad f_{[v_j, \dots, v_n]}(f_{[v_0, \dots, v_j]}(z, t_1, \dots, 1_i, \dots, 1_j), t_{j+1}, \dots, t_n),$$

respectively.

The first map, described in (5.11), is the composition of

$$f_{[v_0, \dots, v_i]} : X_{v_0} \times \text{top}(I^i) \rightarrow X_{v_i}$$

with the map $f_{[v_i, \dots, v_n]}$. Due to the inductive nature of the definition of the latter, it equals the composition of

$$f_{[v_i, \dots, v_j]} : X_{v_i} \times \text{top}(I^{j-i}) \rightarrow X_{v_j}$$

with the map $f_{[v_j, \dots, v_n]} : X_{v_j} \times I^{n-j} \rightarrow Y_{n-1}$.

The second map, described in (5.12), is the composition of the maps

$$f_{[v_0, \dots, v_j]} : X_{v_0} \times \text{top}(I^j) \rightarrow X_{v_j}, \quad t_i = 1,$$

and $f_{[v_j, \dots, v_n]} : X_{v_j} \times I^{n-j} \rightarrow Y_{n-1}$. The former map is again inductively defined as the composition

$$f_{[v_0, \dots, v_i]} : X_{v_0} \times \text{top}(I^i) \rightarrow X_{v_i}$$

with the map

$$f_{[v_i, \dots, v_j]} : X_{v_i} \times \text{top}(I^{j-i}) \rightarrow X_{v_j}.$$

From this description, it is immediate that the two restriction maps, defined by (5.11) and (5.12), are the same.

The proofs for the faces $I_i^+ \cap I_j^- = I_{ij}^{+-}$ and $I_i^- \cap I_j^+ = I_{ij}^{-+}$ ($i < j$) are similar and left to the reader.

Our next task is to define the attaching map on $X_{v_0} \times \text{top}(I^n)$. This map will be obtained by extending the already defined map on $X_{v_0} \times \partial(\text{top}(I^n))$. The definition depends on n , the case $n = 3$ it is the most complicated and we discuss this case last.

We will need the following technical lemma:

LEMMA 5.100. Let V, W be cell complexes with W aspherical and let m be an integer ≥ 2 . Let

$$\partial h : V \times \mathbb{S}^{m-1} \rightarrow W$$

be a cellular map such that for each $v \in V$ the map $\partial h : \{v\} \times \mathbb{S}^{m-1} \rightarrow W$ is null-homotopic. Then h extends to a continuous map $V \times \mathbb{D}^m \rightarrow W$.

PROOF. It suffices to consider the case when V is a simplicial complex. We construct the extension by induction on skeleta of V . By the null-homotopy assumption, the map

$$\partial h|_{V^0 \times \mathbb{S}^{m-1}}$$

extends to a map

$$h : V^0 \times \mathbb{D}^m \rightarrow W.$$

Consider $i \geq 1$ and suppose that the extension $h : V^{i-1} \times \mathbb{D}^m \rightarrow W$ is defined. For each i -dimensional simplex c in V , the product $c \times \mathbb{D}^m$ is a ball of dimension $i + m$ and the map h is already defined on the boundary sphere of this ball. Since this sphere has dimension $\geq m \geq 2$ and W is assumed to be aspherical, the map h extends to the ball $c \times \mathbb{D}^m$. \square

We now proceed with the construction of the attaching map on $X_{v_0} \times \text{top}(I^n)$.

1. For $n = 1$ we use the map $\Psi_{[v_0, v_1]} : X_{v_0} \rightarrow X_{v_1}$ for

$$\partial f_{[v_0, v_1]}|_{X_{v_0} \times \text{top}([0, 1])} = \partial f_{[v_0, v_1]}|_{X_{v_0} \times \{1\}}.$$

2. For $n = 2$ we use the monodromy map⁴

$$\Psi_{[v_0, v_1, v_2]} : X_{v_0} \times I \rightarrow X_{v_2}.$$

3. For $n \geq 4$ we let Σ denote the boundary of the top-face $\text{top}(I^n)$; Σ is, of course, a topological sphere of dimension $n - 2 \geq 2$. Notice that for each $x \in X_{v_0}$ the restriction map

$$\partial f_{[v_0, \dots, v_n]} : \{x\} \times \Sigma \rightarrow X_{v_n}$$

is null-homotopic since X_{v_n} is $K(G_{v_n}, 1)$. Therefore, according to Lemma 5.100, there exists a continuous extension of the map $X_{v_0} \times \Sigma \rightarrow X_{v_n}$ to a map

$$\partial f : X_{v_0} \times \text{top}(I^n) \rightarrow X_{v_n}.$$

4. The proof for $n = 3$ is similar to the one for $n \geq 4$, but we need to show that the maps

$$\partial f_{[v_0, \dots, v_3]} : \{x\} \times \Sigma \rightarrow X_{v_3}$$

are null-homotopic. This is not automatic since in this case $\Sigma \cong \mathbb{S}^1$ and X_{v_3} is not (in general) simply-connected. As X_{v_0} is connected, it suffices to show that the map

$$\partial f_{[v_0, \dots, v_3]} : \{x_{v_0}\} \times \Sigma \rightarrow X_{v_3}$$

is null-homotopic. In order to simplify the notation, we let ∂f denote $\partial f_{[v_0, \dots, v_3]}$, defined on the closure of $X_{v_0} \times (\text{cl}(I^3 - \text{top}(I^3)))$. The boundary of $\text{top}(I^3)$ consists of the segments

$$I_{23}^{-+}, I_{23}^{++}, I_{13}^{++}, I_{13}^{-+},$$

which we orient in the direction of the increase of their natural parameters, t_1 (for the first two segments) and t_2 (for the other two segments). The null-homotopy will follow from:

⁴given by equation (5.8)

LEMMA 5.101. *The loops based at $x_{v_3} \in X_{v_3}$,*

$$\lambda_1 = \partial f(I_{23}^{-+}) \star \partial f(I_{13}^{++}), \quad \lambda_2 = \partial f(I_{13}^{-+}) \star \partial f(I_{23}^{++})$$

are homotopic in X_{v_3} relative to the base-point.

PROOF. By construction (see equations (5.9) and (5.10)),

$$\partial f(I_{23}^{-+}) = \Psi_{[v_0, v_1, v_3]}(x_{v_0}, t), \quad 0 \leq t \leq 1,$$

$$\partial f(I_{13}^{++}) = \Psi_{[v_1, v_2, v_3]}(x_{v_1}, t), \quad 0 \leq t \leq 1,$$

$$\partial f(I_{13}^{-+}) = \Psi_{[v_0, v_2, v_3]}(x_{v_0}, t), \quad 0 \leq t \leq 1,$$

$$\partial f(I_{23}^{++}) = \Psi_{[v_2, v_3]}(\Psi_{[v_0, v_1, v_2]}(x_{v_0}, t)), \quad 0 \leq t \leq 1.$$

Hence, the based loop λ_1 represents the element

$$g_{[v_0, v_1, v_3]} g_{[v_1, v_2, v_3]} \in G_{v_3},$$

while the second loop, λ_2 , represents the element

$$g_{[v_0, v_2, v_3]} \psi_{[v_2, v_3]}(g_{[v_0, v_1, v_2]}) = g_{[v_0, v_2, v_3]} g_{[v_2, v_3]}^{-1} g_{[v_0, v_1, v_2]} g_{[v_2, v_3]}.$$

The first product, λ_1 , equals

$$g_{[v_3, v_0]} g_{[v_0, v_1]} g_{[v_1, v_3]} g_{[v_3, v_1]} g_{[v_1, v_2]} g_{[v_2, v_3]} = g_{[v_3, v_0]} g_{[v_0, v_1]} g_{[v_1, v_2]} g_{[v_2, v_3]},$$

and the second product λ_2 equals

$$g_{[v_3, v_0]} g_{[v_0, v_2]} g_{[v_2, v_3]}^{-1} g_{[v_2, v_0]} g_{[v_0, v_1]} g_{[v_1, v_2]} g_{[v_2, v_3]} = g_{[v_3, v_0]} g_{[v_0, v_1]} g_{[v_1, v_2]} g_{[v_2, v_3]}.$$

(See equations (5.7) and (5.4).) Lemma follows. \square

Thus, we obtain the required attaching map $\partial f_{[v_0, v_1, v_2, v_3]}$. This concludes the construction of the complex Y_n . We define the projection $q_n : Y_n \rightarrow X$ by using the maps

$$Y_\sigma \times I^n \rightarrow I^n \xrightarrow{\theta_\sigma} \sigma \subset X^n.$$

We let Y denote the direct limit of the complexes Y_n :

$$Y_0 \subset Y_1 \subset \dots \subset Y_n \subset \dots$$

The maps q_n then define a map $q : Y \rightarrow X$.

PROPOSITION 5.102. *There exists a homotopy-equivalence $h : Y \rightarrow B$ such that $p \circ h$ is homotopic to q , $p \circ h \simeq q$.*

PROOF. We will construct h by induction on the the dimension of the skeleta of X' . Recall that for each vertex v of X' , $p^{-1}(v)$ is naturally isomorphic (as a simplicial complex) to the quotient $E(G)/G_v$, inducing the isomorphism

$$\pi_1(p^{-1}(v), x_v) \rightarrow G_v.$$

The inclusions $G_v \hookrightarrow G$ induce G_v -equivariant simplicial embeddings

$$E(G_v) \rightarrow E(G)$$

which, therefore, project to homotopy-equivalences $h_v : X_v \rightarrow p^{-1}(v)$. This yields a map

$$h_0 : Y_0 \rightarrow p^{-1}(V(X'))$$

such that $p \circ h_0 \simeq q_0$.

For $n = 1$ we define a map $h_1 : Y_1 \rightarrow B$ by sending, for every oriented edge $\epsilon = [u, v]$ of X' , the product $X_u \times [0, 1]$ to $p^{-1}(e)$ so that

$$h_1|_{X_u \times \{u\}} = h_0|_{X_u \times \{u\}}$$

and the map

$$h_1|_{X_u \times \{v\}}$$

comes from the composition of natural simplicial maps

$$X_u = E(G_u)/G_u \xrightarrow{\Psi_\epsilon} X_v = E(G_v)/G_v \rightarrow E(G)/G_v.$$

The extension to the product $X_u \times (0, 1)$ is the projection of the straight-line homotopy in $E(G) \times [0, 1]$.

Suppose that $n \geq 1$ and a map $h_n : Y_n \rightarrow B$ satisfying $p \circ h_n \simeq q_n$ is constructed. Together with this map h_n we have a collection of maps

$$h_\sigma : X_{v_0} \times I^n \rightarrow X_{v_n} \times \sigma \subset p^{-1}(\sigma)$$

commuting with the maps $f_\sigma : X_{v_0} \times I^n \rightarrow Y_n$,

$$h_n \circ f_\sigma = h_\sigma.$$

Here $\sigma = [v_0, \dots, v_n]$ are n -dimensional simplices in X .

We now construct an extension h_{n+1} of the map h_n to Y_{n+1} . Consider an $(n+1)$ -dimensional simplex $\sigma = [v_0, \dots, v_{n+1}]$ in X' . We observe that the maps Ψ_ϵ (associated with the oriented edges of X' ; see (5.6)) yield natural embeddings

$$X_{v_0} \rightarrow X_{v_i} \rightarrow X_{v_{n+1}}, \quad 0 \leq i \leq n.$$

Therefore, the maps given by the induction hypothesis,

$$h_{\partial_i \sigma} : X_{\min(\partial_i \sigma)} \times I^n \rightarrow p^{-1}(\sigma),$$

yield maps

$$\partial h_\sigma : X_{v_0} \times \partial I^{n+1} \rightarrow X_{v_{n+1}} \times \partial \sigma.$$

Unless $n = 1$, since $X_{v_{n+1}}$ is aspherical, Lemma 5.100 yields an extension of ∂h_σ to

$$(5.13) \quad h_\sigma : X_{v_0} \times I^{n+1} \rightarrow X_{v_{n+1}} \times \sigma,$$

which projects to the map $\theta_\sigma : I^{n+1} \rightarrow \sigma$. When $n = 1$, one needs to verify that for each $x \in X_{v_0}$ the loop

$$\partial h_\sigma|_{\{x\} \times \partial I^2} \rightarrow X_{v_2} \times \partial \sigma$$

is null-homotopic. This follows from the equation

$$g_\tau^{-1} g_\gamma^{-1} g_\alpha g_\beta = 1,$$

see (5.7).

Since Y_{n+1} is obtained by attaching product spaces $X_{v_0} \times I^{n+1}$ to Y_n , the maps h_σ defined above, yield the required map

$$h_{n+1} : Y_{n+1} \rightarrow B.$$

The homotopies $p \circ h_n \simeq q$ extend to a homotopy $p \circ h_{n+1} \simeq q$ due to contractibility of simplices, cf. (5.13).

We next construct a homotopy-inverse $\bar{h} : B \rightarrow Y$. The construction is again by induction on the dimension of the skeleta $(X')^i$ of X' . For each vertex v of X' we define M_v , the *partial* star of v in X' , which consists of all the simplices in X' having v as their *maximal* vertex. For instance, if v was a vertex of X , then M_v

is the ordinary star of v in X' . In general, if v corresponds to a face c of X , then vertices of M_v correspond to the faces of X containing c . For each vertex $v \in X'$ we define the (partial) *star* $St(X_v, Y)$ of X_v in Y as the union

$$\bigcup_{\sigma \in M_v} f_\sigma(X_{\min(\sigma)} \times I^{\dim(\sigma)}) \subset Y,$$

where the union is taken over all simplices σ in M_v .

The maps

$$X_{v_0} \rightarrow X_{v_1} \rightarrow \dots \rightarrow X_{v_n}, \quad v_0 = \min(\sigma), v = v_n = \max(\sigma),$$

defined as compositions of edge-maps⁵ $\Psi_\epsilon, \epsilon = [v_i, v_{i+1}], i = 0, 1, \dots, n-1$, yield a deformation-retraction

$$St(X_w, Y) \rightarrow X_w.$$

In particular, each $St(X_w, Y)$ is homotopy-equivalent to X_w .

Observe furthermore, that for each vertex w of X' , the inclusion map

$$X_w \rightarrow p^{-1}(w) \cong E(G)/G_w$$

is a homotopy-equivalence whose homotopy-inverse

$$\bar{h}_w : p^{-1}(w) \rightarrow X_w$$

is a retraction to X_w . We then construct a homotopy-inverse map $\bar{h} : B \rightarrow Y$ by induction on the dimension of the skeleta of X' , starting with the maps $\bar{h}_w, w \in V(X')$. Assuming that \bar{h}_{n-1} is defined on $p^{-1}((X')^{n-1})$, we extend this map to $p^{-1}((X')^n)$, one n -dimensional simplex at a time, using Lemma 5.100 for the maps

$$\bar{h}_{n-1} : p^{-1}(\partial\sigma) \rightarrow St(X_v, Y), \quad v = \max(\sigma).$$

We leave it to the reader to verify that \bar{h} is a homotopy-inverse of h : This is again proven using Lemma 5.100. \square

We can now define the space \widehat{X} . In view of the homotopy-equivalence $Y \rightarrow B$, there exists a covering map

$$\widehat{X} \rightarrow Y$$

corresponding to the homomorphism of fundamental groups

$$\pi_1(Y) \xrightarrow{\cong} \pi_1(B) \rightarrow G,$$

where the latter is the homomorphism associated with the G -covering map

$$E(G) \times \tilde{X} \rightarrow B.$$

The homotopy-equivalence $h : Y \rightarrow B$ lifts to a G -equivariant homotopy-equivalence $\widehat{h} : \widehat{X} \rightarrow \tilde{X}$. This concludes the proof of Theorem 5.99. \square

We note that the complex Y has a natural filtration

$$F_0(Y) \subset F_1(Y) \subset \dots,$$

where $F_j(Y)$ is obtained by attaching, for every simplex σ in X' , not the entire X_v , for $v = \min(\sigma)$, but only $E^j(G_v)/G_v \times I^n$, where $n = \dim(\sigma)$ and $E^j(G_v)$ is the j -skeleton of $E(G_v)$. If each group $G_v, v \in V(X')$, is finite, then each $E^j(G_v)$ is also finite. In addition to this filtration, we also have the filtration

$$Y_0 \subset Y_1 \subset \dots$$

⁵see (5.6)

coming from the inductive construction of the complex Y .

One application of these observation is the following:

LEMMA 5.103. *Suppose that $G \curvearrowright \tilde{X}$ is a properly discontinuous cocompact action. Then the Haefliger model \hat{X} can be chosen so that the action $G \curvearrowright \hat{X}$ is cocompact on each skeleton.*

PROOF. Suppose that the cell complex \tilde{X} is n -dimensional, $n < \infty$. The simplicial complex \tilde{Z} in Exercise 5.97 can be chosen to have dimension n as well. Thus, compactness of \tilde{X}^i/G implies compactness of \tilde{Z}^i/G and, hence, compactness of the complexes $F_j(Y_i)$ defined above. In particular, every skeleton of Y is finite. \square

5.8.3. Groups of finite type. Consider a free group cellular action $G \curvearrowright X$. In the case when X is a simplicial complex or, more generally, an almost regular complex, Lemmas 5.96 and 5.98 imply that G acts properly discontinuously on X .

EXERCISE 5.104. Show that a free cellular group action on a cell complex is always properly discontinuous.

If G is a group admitting a cellular free and cocompact action on a graph Γ , then G is finitely generated, as, by the covering theory, $G \cong \pi_1(\Gamma/G)/p_*(\pi_1(\Gamma))$, where $p : \Gamma \rightarrow \Gamma/G$ is the covering map. Groups of *finite type* \mathbf{F}_n are higher-dimensional generalizations of this example.

DEFINITION 5.105. A group G is said to have *type* \mathbf{F}_n , $1 \leq n < \infty$, if it admits a cellular free and cocompact action on an $n - 1$ -connected n -dimensional cell complex \tilde{X} . A group G has type \mathbf{F}_∞ if it admits a cellular, free cocompact action on a contractible cell complex \tilde{X} , which is cocompact on each skeleton. A group G has *type* \mathbf{F} if there exists a finite $K(G, 1)$ complex Y , i.e. G acts cellularly, freely and cocompactly on a contractible finite-dimensional complex X .

In other words, G has type \mathbf{F}_n ($n < \infty$) if there exists a finite $n - 1$ -connected n -dimensional complex Y with $\pi_1(Y) \simeq G$. Similarly, G has type \mathbf{F}_∞ if and only if there exist an aspherical complex Y with $\pi_1(Y) \simeq G$ such that every skeleton of Y is finite.

EXAMPLE 5.106. Every finite group G has type \mathbf{F}_∞ .

PROOF. Use the action of G on its classifying space $E(G)$, see Section 5.8.2. \square Clearly,

$$\mathbf{F} \subset \mathbf{F}_\infty \subset \dots \mathbf{F}_n \subset \mathbf{F}_{n-1} \subset \dots \subset \mathbf{F}_1.$$

We refer the reader to [Geo08, Proposition 7.2.2] for the proof of the following theorem:

THEOREM 5.107. *A group G has type \mathbf{F}_∞ if and only if it has type \mathbf{F}_n for every n .*

EXAMPLE 5.108 (See [Bie76b] and [BB97]). Let F_2 be free group on two generators. Consider the group $G = F_2^n$ which is the n -fold direct product of F_2 . We equip G with the generating set

$$a_1, b_1, \dots, a_n, b_n,$$

where a_i, b_i are the free generators of the i -th direct factor of G . Define the homomorphism $\phi : G \rightarrow \mathbb{Z}$ which sends all generators a_i, b_i of G to the generator $1 \in \mathbb{Z}$. Let $K := \text{Ker}(\phi)$. Then K is of type \mathbf{F}_{n-1} but not of type \mathbf{F}_n .

In view of Lemma 5.103, we obtain:

COROLLARY 5.109. *A group G has type \mathbf{F}_n if and only if it admits a properly discontinuous cocompact cellular action on an $n - 1$ -connected n -dimensional cell complex \tilde{X} .*

PROOF. One direction is obvious. Suppose, therefore, that we have an action $G \curvearrowright \tilde{X}$ as in the statement of the corollary. If the action were free, it would follow immediately that G has type \mathbf{F}_n (cf. Definition 5.105). Consider now the general case. Since the action of G is properly discontinuous, the stabilizer of each cell is finite. We then apply the Haefliger construction (as in Lemma 5.103) to the action $G \curvearrowright \tilde{X}$ and obtain a free properly discontinuous action $G \curvearrowright \hat{X}$ cocompact on each n -dimensional skeleton \hat{X}^n , since for every finite group G_v each skeleton of $E(G_v)$ is a finite complex. Recall that the inclusion $\hat{X}^n \hookrightarrow \hat{X}$ induces monomorphisms of all homotopy groups π_j , $j \leq n - 1$. Since \tilde{X} is $n - 1$ -connected, the same holds for \hat{X} and, hence, for \hat{X}^n . Thus, G admits a free cocompact action on an $n - 1$ -connected complex \hat{X}^n . \square

5.9. Cohomology

The purpose of this section is to introduce cohomology of groups and to give explicit formulae for cocycles and coboundaries in small degrees. We refer the reader to [Bro82b, Chapter III, Section 1] for the more thorough discussion. We will also connect group cohomology to two group-theoretic constructions: *Semidirect products* and *coextensions*.

5.9.1. Group rings and modules. Suppose that R is a commutative ring with unit element 1. The *R -ring* RG of a group G is the set of formal sums $\sum_{g \in G} m_g g$, where m_g are elements of R which are equal to zero for all but finitely many values of g . The most important examples for us will be the integer group ring $\mathbb{Z}G$ and the rational group ring $\mathbb{Q}G$. So far, RG is just a set, but it becomes a ring once endowed with the two operations:

- addition:

$$\sum_{g \in G} m_g g + \sum_{g \in G} n_g g = \sum_{g \in G} (m_g + n_g) g$$

- multiplication defined by the convolution of maps to \mathbb{Z} , that is

$$\sum_{a \in G} m_a a + \sum_{b \in G} n_b b = \sum_{g \in G} \left(\sum_{ab=g} m_a n_b \right) g.$$

According to a Theorem of G. Higman [Hig40], every integer group ring is an integral domain. Both R and G embed as subsets of RG by identifying every $m \in \mathbb{Z}$ with $m1_G$ and every $g \in G$ with $1g$. Every group homomorphism $\varphi : G \rightarrow H$ induces a homomorphism between group rings, which by abuse of notation we shall denote also by φ . In particular, the trivial homomorphism $o : G \rightarrow \{1\}$ induces a retraction $o : \mathbb{Z}G \rightarrow R$, called the *augmentation*. If the homomorphism $\varphi : G \rightarrow H$

is an isomorphism, then so is the homomorphism between group rings. This implies that an action of a group G on another group H (by automorphisms) extends to an action of G on the group ring $\mathbb{Z}H$ (by automorphisms).

Let L be a ring and M be an abelian group. We say that M is a (left) L -module if we are given a map

$$(\ell, m) \mapsto \ell \cdot m, L \times M \rightarrow M,$$

which is additive in both variables and so that

$$(5.14) \quad (\ell_1 \star \ell_2) \cdot m = \ell_1 \cdot (\ell_2 \cdot m),$$

where \star denotes the multiplication operation in L .

Similarly, M is a *right* L -module if we are given an additive (in both variables) map

$$(m, \ell) \mapsto m \cdot \ell, M \times L \rightarrow M,$$

so that

$$(5.15) \quad m \cdot (\ell_1 \star \ell_2) = (m \cdot \ell_1) \cdot \ell_2.$$

Lastly, M is an L -bimodule if M has structure of both left and right L -module.

In the case when R is a field F (say, $R = \mathbb{Q}$), a left RG -module is an F -vector space endowed with a linear G -action. In the case $R = \mathbb{Z}$, we will refer to (left) $\mathbb{Z}G$ -modules simply as G -modules.

5.9.2. Group cohomology. Let G be a group and let M, N be left $\mathbb{Z}G$ -modules; then $\text{Hom}_G(M, N)$ denotes the \mathbb{Z} -submodule of G -invariant elements in the \mathbb{Z} -module $\text{Hom}(M, N)$, where G acts on homomorphisms (of abelian groups) $u : M \rightarrow N$ by the formula:

$$(gu)(m) = g \cdot u(g^{-1}m).$$

Suppose that C_* is a chain complex of abelian groups endowed with an action of G and A is a G -module, then $\text{Hom}_G(C_*, A) \subset \text{Hom}(C_*, A)$ is the chain subcomplex formed by the submodules $\text{Hom}_G(C_k, A)$ in $\text{Hom}(C_k, A)$. The *standard chain complex* $C_* = C_*(G)$ of G with coefficients in A is defined as follows:

$C_k(G) = \mathbb{Z} \otimes \prod_{i=0}^k G$, is the $\mathbb{Z}G$ -module freely generated by $(k+1)$ -tuples (g_0, \dots, g_k) of elements of G with the G -action given by

$$g \cdot (g_0, \dots, g_k) = (gg_0, \dots, gg_k).$$

The reader should think of each tuple as spanning a k -simplex. The boundary operator on this chain complex is the natural one:

$$\partial_k(g_0, \dots, g_k) = \sum_{i=0}^k (-1)^i (g_0, \dots, \hat{g}_i, \dots, g_k),$$

where \hat{g}_i means that we omit this entry in the $k+1$ -tuple. The dual cochain complex C^* is defined by:

$$C^k = \text{Hom}(C_k, A), \quad \delta_k(f)((g_0, \dots, g_{k+1})) = f(\partial_{k+1}(g_0, \dots, g_{k+1})), f \in C^k.$$

Thus, C_* and C^* are just the simplicial chain and cochain complexes of the simplicial complex defining the Milnor's classifying space EG of the group G (see Section 5.8.2), with which the reader is probably familiar with from a basic algebraic topology course.

Suppose for a moment that A is a trivial G -module. Then, for $BG = (EG)/G$, the simplicial cochain complex $C^*(BG, A)$ is naturally isomorphic to the subcomplex of G -invariant cochains in $C^*(G, A)$, i.e. the subcomplex

$$(C^*(G, A))^G = \text{Hom}_G(C_*, A).$$

If A is a non-trivial G -module, then $\text{Hom}_G(C_*, A)$ is still isomorphic to a certain natural cochain complex based on the simplicial complex $C_*(BG)$ (a cochain complex with twisted coefficients, or coefficients in a certain sheaf), but the definition is more involved and we will omit it.

DEFINITION 5.110. Define the subspaces of i -cocycles and i -coboundaries in $\text{Hom}_G(C_i, A)$ as

$$Z^i(G, A) := \text{Ker}(\delta_i), \quad B^i(G, A) := \text{Im}(\delta_{i-1}),$$

respectively. The cohomology groups of G with coefficients in the G -module A are defined as

$$H^*(G, A) := H_*(\text{Hom}_G(C_*, A)).$$

In other words,

$$H^i(G, A) = Z^i(G, A)/B^i(G, A).$$

In particular, if A is a trivial G -module, then $H^*(G, A) = H^*(BG, A)$.

DEFINITION 5.111. The (*integer*) *cohomological dimension* of a group G , is defined as

$$cd(G) = \sup\{q \in \mathbb{Z} : \exists A, \text{ a } \mathbb{Z}G\text{-module, such that } H^q(G, A) \neq 0\}.$$

Note that the definition of cohomological dimension we gave is, in fact, a theorem rather than the standard definition. We refer the reader to [Bro82b] for the usual definition of cohomological dimension in terms of projective resolutions.

EXAMPLE 5.112. 1. Suppose that G admits a $K(G, 1)$ CW complex X . Then $cd(G) \leq \dim(X)$.

2. If G is a non-trivial finite group, then $cd(G) = \infty$.

REMARK 5.113. 1. Analogously to the integer cohomological dimension, one defines the *rational* cohomological dimension $cd_{\mathbb{Q}}(G)$ as the supremum of degrees q such that $H^q(G, A) \neq 0$ for some $\mathbb{Q}G$ -module A , i.e. a vector space over \mathbb{Q} on which G acts linearly.⁶ One advantage $cd_{\mathbb{Q}}$ has over the integer cohomological dimension is that (unlike the latter) the former is an invariant under virtual isomorphisms of groups (see [Bro82b]) and, for finitely generated groups, is invariant under quasiisometries, see Theorem 9.64.

2. Similarly to the cohomological dimension of G one defines its *homological dimension* over a ring R , as the supremum of degrees q such that $H_q(G, A) \neq 0$ for some RG -module A .

3. One defines the *geometric dimension* of a group G as the least number k such there exists a k -dimensional cell complex X which is $K(G, 1)$. Thus, geometric dimension of G is always $\geq cd(G)$.

⁶Even more generally, given an arbitrary commutative ring R one defines the appropriate group cohomology using the group ring RG instead of $\mathbb{Z}G$, and the cohomological dimension $cd_R(G)$, see [Bro82b].

So far, all definitions looked very natural. Our next step is to reduce the number of variables in the definition of cochains by one using the fact that cochains in $\text{Hom}_G(C_k, A)$ are G -invariant. The drawback of this reduction, as we will see, will be lack of naturality, but the advantage will be new formulae for cohomology groups which are useful in some applications.

By G -invariance, for $f \in \text{Hom}_G(C_k, A)$ we have:

$$f(g_0, \dots, g_k) = g_0 \cdot f(1, g_0^{-1}g_1, \dots, g_0^{-1}g_k).$$

In other words, it suffices to restrict cochains to the set of $(k+1)$ -tuples where the first entry is $1 \in G$. Every such tuple has the form

$$(1, g_1, g_1g_2, \dots, g_1 \cdots g_k)$$

(we will see below why). The latter is commonly denoted

$$[g_1|g_2|\dots|g_k].$$

Note that, computing the value of the coboundary,

$$\delta_{k-1}f(1, g_1, g_1g_2, \dots, g_1 \cdots g_k) = \delta_{k-1}f([g_1|g_2|\dots|g_k])$$

we get

$$\begin{aligned} \delta_{k-1}f(1, g_1, g_1g_2, \dots, g_1 \cdots g_k) = \\ f(g_1, \dots, g_1 \cdots g_k) - f(1, g_1g_2, \dots, g_1 \cdots g_k) + f(1, g_1, g_1g_2g_3, \dots, g_1 \cdots g_k) - \dots = \\ g_1 \cdot f(1, g_2, \dots, g_2 \cdots g_k) - f([g_1g_2|g_3|\dots|g_k]) + f([g_1|g_2g_3|g_4|\dots|g_k]) - \dots = \\ g_1 \cdot f([g_2|\dots|g_k]) - f([g_1g_2|g_3|\dots|g_k]) + f([g_1|g_2g_3|g_4|\dots|g_k]) - \dots \end{aligned}$$

Thus,

$$\begin{aligned} \delta_{k-1}f([g_1|g_2|\dots|g_k]) = g_1 \cdot f([g_2|\dots|g_k]) - f([g_1g_2|g_3|\dots|g_k]) + \\ f([g_1|g_2g_3|g_4|\dots|g_k]) - \dots \end{aligned}$$

Then we let \bar{C}^k ($k \geq 1$) denote the abelian group of functions f sending k -tuples $[g_1|\dots|g_k]$ of elements of G to elements of A ; we equip these groups with the above coboundary homomorphisms δ_k . For $k = 0$, we have to use the empty symbol $[\]$, $f([\]) = a \in A$, so that such functions f are identified with elements of A . Thus, $\bar{C}_0 = A$ and the above formula for δ_0 reads as:

$$\delta_0 : a \mapsto c_a, \quad c_a([g]) = g \cdot a - a.$$

The resulting chain complex (\bar{C}_*, δ_*) is called the *inhomogeneous bar complex* of G with coefficients in A . We now compute the coboundary maps δ_k for this complex for small values of k :

- (1) $\delta_0 : a \mapsto f_a, \quad f_a([g]) = g \cdot a - a.$
- (2) $\delta_1(f)([g_1, g_2]) = g_1 \cdot f([g_2]) - f([g_1g_2]) + f([g_1]).$
- (3) $\delta_2(f)([g_1|g_2|g_3]) = g_1 \cdot f([g_2|g_3]) - f([g_1g_2|g_3]) + f([g_1|g_2g_3]) - f([g_1|g_2]).$

Therefore, spaces of coboundaries and cocycles for (\bar{C}_*, δ_*) in small degrees are (we now drop the bar notation for simplicity):

- (1) $B^1(G, A) = \{f_a : G \rightarrow A, \forall a \in A | f_a(g) = g \cdot a - a\}.$
- (2) $Z^1(G, A) = \{f : G \rightarrow A | f(g_1g_2) = f(g_1) + g_1 \cdot f(g_2)\}.$
- (3) $B^2(G, A) = \{h : G \times G \rightarrow A | \exists f : G \rightarrow A, h(g_1, g_2) = f(g_1) - f(g_1g_2) + g_1 \cdot f(g_2)\}.$

$$(4) \ Z^2(G, A) = \{f : G \times G \rightarrow A \mid g_1 \cdot f(g_2, g_3) - f(g_1, g_2) = f(g_1 g_2, g_3) - f(g_1, g_2 g_3)\}.$$

Let us look at the definition of $Z^1(G, A)$ more closely. In addition to the left action of G on A , we define a *trivial right action* of G on A : $a \cdot g = a$. Then a function $f : G \rightarrow A$ is a 1-cocycle if and only if

$$f(g_1 g_2) = f(g_1) \cdot g_2 + g_1 \cdot f(g_2).$$

The reader will immediately recognize here the Leibnitz formula for the derivative of the product. Hence, 1-cocycles $f \in Z^1(G, A)$ are called *derivations* of G with values in A . The 1-coboundaries are called *principal derivations* or *inner derivations*. If A is trivial as a left G -module, then, of course, all principal derivations are zero and derivations are just homomorphisms $G \rightarrow A$.

Nonabelian derivations. The notions of derivation and principal derivation can be extended to the case when the target group is nonabelian; we will use the notation N for the target group with the binary operation \star and $g \cdot n$ for the action of G on N by automorphisms, i.e.

$$g \cdot n = \varphi(g)(n), \quad \text{where } \varphi : G \rightarrow \text{Aut}(N) \text{ is a homomorphism.}$$

DEFINITION 5.114. A function $d : G \rightarrow N$ is called a *derivation* if

$$d(g_1 g_2) = d(g_1) \star g_1 \cdot d(g_2), \quad \forall g_1, g_2 \in G.$$

A derivation is called *principal* if it is of the form $d = d_n$, where

$$d_n(g) = n^{-1} \star (g \cdot n).$$

The space of derivations is denoted $\text{Der}(G, N)$ and the subspace of principal derivations is denoted $\text{Prin}(G, N)$ or, simply, $P(G, N)$.

EXERCISE 5.115. Verify that every principal derivation is indeed a derivation.

EXERCISE 5.116. Verify that every derivation d satisfies

- $d(1) = 1$;
- $d(g^{-1}) = g^{-1} \cdot [d(g)]^{-1}$.

We will use derivations in the context of free solvable groups in section 13.6. In section 5.9.5 we will discuss derivations in the context of semidirect products, while in section 5.9.6 we explain how second cohomology group $H^2(G, A)$ can be used to describe central coextensions.

Nonabelian cohomology. We would like to define the 1-st cohomology $H^1(G, N)$, where the group N is nonabelian and we have an action of G on N . The problem is that neither $\text{Der}(G, N)$ nor $\text{Prin}(G, N)$ is a group, so taking quotient $\text{Der}(G, N)/\text{Prin}(G, N)$ makes no sense. Nevertheless, we can think of the formula

$$f \mapsto f + d_a, a \in A,$$

in the abelian case (defining action of $\text{Prin}(G, A)$ on $\text{Der}(G, A)$) as the *left* action of the group A on $\text{Der}(G, A)$:

$$a(f) = f', \quad f'(g) = -a + f(g) + (g \cdot a).$$

The latter generalizes in the nonabelian case, as the group N acts to the left on $\text{Der}(G, N)$ by

$$n(f) = f', \quad f'(g) = n^{-1} \star f(g) \star (g \cdot n).$$

Then one defines $H^1(G, N)$ as the quotient

$$N \backslash \text{Der}(G, N).$$

EXAMPLE 5.117. 1. Suppose that G -action on N is trivial. Then $\text{Der}(G, N) = \text{Hom}(G, N)$ and N acts on homomorphisms $f : G \rightarrow N$ by postcomposition with inner automorphisms. Thus, $H^1(G, N)$ in this case is

$$N \backslash \text{Hom}(G, N),$$

the set of conjugacy classes of homomorphisms $G \rightarrow N$.

2. Suppose that $G \cong \mathbb{Z} = \langle 1 \rangle$ and the action φ of \mathbb{Z} on N is arbitrary. We have $\eta := \varphi(1) \in \text{Aut}(N)$. Then $H^1(G, N)$ is the set of *twisted conjugacy classes of elements of N* : Two elements $m_1, m_2 \in N$ are said to be in the same η -twisted conjugacy class if there exists $n \in N$ so that

$$m_2 = n^{-1} \star m_1 \star \eta(n).$$

Indeed, every derivation $d \in \text{Der}(\mathbb{Z}, N)$ is determined by the image $m = d(1) \in N$. Then two derivations d_i so that $m_i = d_i(1)$ ($i = 1, 2$) are in the same N -orbit if m_1, m_2 are in the same η -twisted conjugacy class.

5.9.3. Bounded cohomology of groups. *An isometric Banach $\mathbb{Z}G$ -module V is a Banach space equipped with an isometric action of the group G . Using $C_*(G)$, which is the bar-complex of G , one defines the *bounded cochain complex**

$$C_b^*(G, V) = \text{Hom}_{G,b}(C_*, V),$$

where $C_b^k(G, V)$ consists of G -equivariant bounded maps $G^{k+1} \rightarrow V$, with the usual coboundary operator. Accordingly, one defines the *bounded cohomology groups* of G with coefficients in V :

$$H_b^*(G, V) := H_*(C_b^*(G, V)).$$

Alternatively, one can use the subcomplex of bounded functions $\bar{C}_b^*(G, V)$ in the inhomogeneous bar-complex of the group G and obtain

$$H_b^k(G, V) \cong Z_b^k(G, V) / B_b^k(G, V),$$

where the spaces of cocycles and coboundaries on the right hand-side refer to the bounded elements of the group of homogeneous cocycles, respectively to the images by δ^{k-1} of bounded cochains.

The same definitions go through if instead of the entire V one uses a $\mathbb{Z}G$ -submodule $A \subset V$; then one defines the bounded cohomology groups $H_b^k(G, A)$ via maps $G^{k+1} \rightarrow A$.

We now consider the special case, when V (and, hence, A) is a trivial G -module. (The most important cases are, of course, $V = \mathbb{R}$ and $A = \mathbb{Z}$.) Then for a classifying space $Y = BG$ of G one defines the subcomplex $C_b^*(Y, A)$ of the cochain complex $C^*(Y, A)$. The homology of this subcomplex is the *bounded cohomology* $H_b^*(Y, A)$ of Y with coefficients in A .

EXERCISE 5.118. Verify that $H_b^*(Y, A) \cong H_b^*(G, A)$.

Note that the above isomorphism holds even if Y is not a $K(G, 1)$ but merely has G as its fundamental group, see [Bro81b] and [Gro82].

It is instructive to identify elements of $\bar{Z}_b^2(G, A)$, where A is a subgroup of \mathbb{R} , which appear as ordinary coboundaries: For $f \in C^1(G, A)$, i.e. $f : G \rightarrow A$,

$$\delta_1(f)([g_1, g_2]) = f(g_2) - f(g_1g_2) + f(g_1)$$

is a bounded 2-cocycle if and only if there exists a constant D so that for all $g_1, g_2 \in G$,

$$(5.16) \quad |f(g_1) + f(g_2) - f(g_1g_2)| \leq D.$$

In other words, such f is “almost a homomorphism $f : G \rightarrow A$ ”, with an error $\leq D$ in the definition of a homomorphism.

DEFINITION 5.119. A map $f : G \rightarrow \mathbb{R}$ is called a *quasimorphism* if it satisfies the inequality (5.16) for all $g_1, g_2 \in G$ and a fixed constant D .

Quasi-morphisms appear frequently in Geometric Group Theory; they were first used by R. Brooks in [Bro81b], who proved that, while for the free group F_n of rank $n \geq 2$, $H^2(F_n, \mathbb{R}) = 0$, nevertheless, the vector space $H_b^2(F_n, \mathbb{R})$ is infinite-dimensional. Namely, he constructed an infinite-dimensional space of equivalence classes of quasimorphisms $F_n \rightarrow \mathbb{R}$, where

$$f_1 \sim f_2 \iff \|f_1 - f_2\| < \infty.$$

Taking coboundaries of these quasimorphisms shows that $H_b^2(F_n, \mathbb{R})$ has infinite dimension.

Many interesting groups do not admit non-trivial homomorphisms of \mathbb{R} but admit unbounded quasimorphisms. For instance, a hyperbolic Coxeter group G does not admit non-trivial homomorphisms to \mathbb{R} . However, if G is a nonelementary hyperbolic group, it has infinite-dimensional space of equivalence classes of quasimorphisms, see [EF97a] for details. We refer the reader to Monod’s paper [Mon06] for a survey of applications of bounded cohomology of groups, as well as Calegari’s book [Cal09] for the in-depth discussion of quasimorphisms defined by the *commutator norm*.

We will encounter elements of the group $H_b^2(G, \mathbb{Z})$ in Section 11.19 when discussing central coextensions of hyperbolic groups, as we will be proving subjectivity of the homomorphism $H_b^2(G, \mathbb{Z}) \rightarrow H^2(G, \mathbb{Z})$.

Analogously to the bounded cohomology, one defines ℓ_p -cohomology and ℓ_p -homology groups, we refer the reader to [AG99, BP03, Pan95] for the detailed discussion.

5.9.4. Ring derivations. Our next goal is to extend the notion of derivation in the context of (noncommutative) rings. Typical rings that the reader should have in mind are *integer group rings*.

DEFINITION 5.120. Let M be an L -bimodule. A *derivation* (with respect to this bimodule structure) is a map $d : L \rightarrow M$ such that:

- (1) $d(\ell_1 + \ell_2) = d(\ell_1) + d(\ell_2)$,
- (2) $d(\ell_1 \star \ell_2) = d(\ell_1) \cdot \ell_2 + \ell_1 \cdot d(\ell_2)$.

The space of derivations is an abelian group, which will be denoted $Der(L, M)$.

Below is the key example of a bimodule that we will be using in the context of derivations. Let G, H be groups, $\varphi : G \rightarrow Bij(H)$ is an action of G on H by set-theoretic automorphisms. We let $L := \mathbb{Z}G, M := \mathbb{Z}H$ be the integer group

rings, where we regard the ring M as an abelian group and ignore its multiplicative structure.

Every action $\varphi : G \curvearrowright H$ determines the left L -module structure on M by:

$$\left(\sum_i a_i g_i\right) \cdot \left(\sum_j b_j h_j\right) := \sum_{i,j} a_i b_j g_i \cdot h_j, \quad a_i \in \mathbb{Z}, b_j \in \mathbb{Z},$$

where $g \cdot h = \varphi(g)(h)$ for $g \in G, h \in H$. We define the structure of a *right L -module* on M by:

$$(m, \ell) \mapsto m o(\ell) = o(\ell) m, \quad o(\ell) \in \mathbb{Z},$$

where $o : L \rightarrow \mathbb{Z}$ is the augmentation of $\mathbb{Z}G = L$.

Derivations with respect for the above group ring bimodules will be called *group ring derivations*.

EXERCISE 5.121. Verify the following properties of group ring derivations:

(P_1) $d(1_G) = 0$, whence $d(m) = 0$ for every $m \in \mathbb{Z}$.

(P_2) $d(g^{-1}) = -g^{-1} \cdot d(g)$.

(P_3) $d(g_1 \cdots g_m) = \sum_{i=1}^m (g_1 \cdots g_{i-1}) \cdot d(g_i) o(g_{i+1} \cdots g_m)$.

(P_4) Every derivation $d \in \text{Der}(\mathbb{Z}G, \mathbb{Z}H)$ is uniquely determined by its values $d(x)$ on the generators x of G .

Fox Calculus. We now consider the special case when $G = H = F_X$, is the free group on the generating set X . In this context, the theory of derivations was developed by R. H. Fox in [**Fox53**].

LEMMA 5.122. Every map $d : X \rightarrow M = \mathbb{Z}G$ extends to a group ring derivation $d \in \text{Der}(\mathbb{Z}G, M)$.

PROOF. We set

$$d(x^{-1}) = -x^{-1} \cdot d(x), \quad \forall x \in X$$

and $d(1) = 0$. We then extend d inductively to the free group G by

$$d(yu) = d(y) + y \cdot d(u),$$

where $y = x \in X$ or $y = x^{-1}$ and yu is a reduced word in the alphabet $X \cup X^{-1}$. Lastly, we extend d by additivity to the rest of the ring $L = \mathbb{Z}G$. In order to verify that d is a derivation, we need to check only that

$$d(uv) = d(u) + u \cdot d(v),$$

where $u, v \in F_X$. The verification is a straightforward induction on the length of the reduced word u and is left to the reader. \square

DEFINITION 5.123. To each generator $x_i \in X$ we associate a derivation ∂_i , called the *Fox derivative*, defined by $\partial_i x_j = \delta_{ij} \in \{0, 1\}$, which is regarded as the subset of $\mathbb{Z} \cdot 1_G \subset \mathbb{Z}G$. The maps ∂_i then extend to derivations $\partial_i \in \text{Der}(\mathbb{Z}F_X, \mathbb{Z}F_X)$ as in Lemma 5.122. In particular,

$$\partial_i(x_i^{-1}) = -x_i^{-1}.$$

Importance of the derivations ∂_i comes from:

PROPOSITION 5.124. Suppose that $G = F_r$ is the free group of rank $r < \infty$. Then every derivation $d \in \text{Der}(\mathbb{Z}G, \mathbb{Z}G)$ can be written as a sum

$$d = \sum_{i=1}^r k_i \partial_i, \quad \text{where } k_i = d(x_i) \in \mathbb{Z}.$$

Furthermore, $\text{Der}(\mathbb{Z}G, \mathbb{Z}G)$ is a free abelian group with the basis $\partial_i, i = 1, \dots, r$.

PROOF. The first assertion immediately follows from Exercise 5.121 (part (P_4)), and from the fact that both sides of the equation evaluated on x_j equal k_j . Thus, the derivations $\partial_i, i = 1, \dots, k$, generate $\text{Der}(\mathbb{Z}G, \mathbb{Z}G)$. Independence of these generators follows from the fact that $\partial_i x_j = \delta_{ij}$. \square

5.9.5. Derivations and split extensions. Components of homomorphisms to semidirect products.

DEFINITION 5.125. Let G and L be two groups and let N, H be subgroups in G .

- (1) Assume that $G = N \times H$. Every group homomorphism $F : L \rightarrow G$ splits as a product of two homomorphisms $F = (f_1, f_2)$, $f_1 : L \rightarrow N$ and $f_2 : L \rightarrow H$, called the *components* of F .
- (2) Assume now that G is a semidirect product $N \rtimes H$. Then every homomorphism $F : L \rightarrow G$ determines (and is determined by) a pair (d, f) , where
 - $f : L \rightarrow H$ is a homomorphism (the composition of F and the retraction $G \rightarrow H$);
 - a map $d = d_F : L \rightarrow N$, called *derivation* associated with F . The derivation d is determined by the formula

$$F(\ell) = d(\ell)f(\ell).$$

EXERCISE 5.126. Show that d is indeed a derivation in the sense of Section 5.9.2.

EXERCISE 5.127. Verify that for every derivation d and a homomorphism $f : L \rightarrow H$ there exists a homomorphism $F : L \rightarrow G$ with the components d, f .

Extensions and coextensions.

DEFINITION 5.128. Given a short exact sequence

$$\{1\} \longrightarrow N \longrightarrow G \longrightarrow H \longrightarrow \{1\},$$

we call the group G an *extension of N by H* or a *coextension of H by N* .⁷

Given two classes of groups \mathcal{A} and \mathcal{B} , the groups that can be obtained as extensions of N by H with $N \in \mathcal{A}$ and $H \in \mathcal{B}$, are called *\mathcal{A} -by- \mathcal{B} groups* (e.g. abelian-by-finite, nilpotent-by-free etc.).

⁷Our terminology is a bit nonstandard, as both constructions are called *extensions* in the literature. We settled on the coextension terminology following the paper [MN82] where it was used for semigroups.

Two extensions defined by the short exact sequences

$$\{1\} \longrightarrow N_i \xrightarrow{\varphi_i} G_i \xrightarrow{\psi_i} H_i \longrightarrow \{1\}$$

($i = 1, 2$) are *equivalent* if there exist isomorphisms

$$f_1 : N_1 \rightarrow N_2, \quad f_2 : G_1 \rightarrow G_2, \quad f_3 : H_1 \rightarrow H_2$$

that determine a commutative diagram:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & N_1 & \longrightarrow & G_1 & \longrightarrow & H_1 & \longrightarrow & 1 \\ & & \downarrow f_1 & & \downarrow f_2 & & \downarrow f_3 & & \\ 1 & \longrightarrow & N_2 & \longrightarrow & G_2 & \longrightarrow & H_2 & \longrightarrow & 1 \end{array}$$

We now use the notion of an isomorphism of exact sequences to reinterpret the notion of a split extension.

PROPOSITION 5.129. *Consider a short exact sequence*

$$(5.17) \quad 1 \rightarrow N \xrightarrow{\iota} G \xrightarrow{\pi} Q \rightarrow 1.$$

The following are equivalent:

- (1) *the sequence splits;*
- (2) *there exists a subgroup H in G such that the projection π restricted to H becomes an isomorphism.*
- (3) *the extension G is equivalent to an extension corresponding to a semidirect product $N \rtimes Q$;*
- (4) *there exists a subgroup H in G such that $G = N \rtimes H$.*

PROOF. It is clear that (2) \Rightarrow (1).

(1) \Rightarrow (2): Let $\sigma : Q \rightarrow \sigma(H) \subset G$ be a section. The equality $\pi \circ \sigma = \text{Id}_Q$ implies that π restricted to H is both surjective and injective.

The implication (3) \Rightarrow (4) is obvious.

(3) \Rightarrow (2): Assume that there exists H such that $\pi|_H$ is an isomorphism. The fact that it is surjective implies that $G = NH$. The fact that it is injective implies that $H \cap N = \{1\}$.

(2) \Rightarrow (3): Since π restricted to H is surjective, it follows that for every $g \in G$ there exists $h \in H$ such that $\pi(g) = \pi(h)$, hence $gh^{-1} \in \text{Ker } \pi = \text{Im } \iota$.

The intersection $\iota(N) \cap H$ is the preimage of 1 by π restricted to H , hence it must be $\{1\}$.

(4) \Rightarrow (2): The existence of the decomposition for every $g \in G$ implies that π restricted to H is surjective.

The uniqueness of the decomposition implies that $H \cap \text{Im } \iota = \{1\}$, whence π restricted to H is injective. \square

REMARK 5.130. Every sequence (5.17) where the group Q is free splits; see Lemma 7.23.

EXAMPLES 5.131. (1) For $n \geq 1$, the short exact sequence

$$1 \longrightarrow (2\mathbb{Z})^n \longrightarrow \mathbb{Z}^n \longrightarrow \mathbb{Z}_2^n \longrightarrow 1$$

does not split.

(2) Let F_n be a free group of rank $n \geq 2$ (see Definition 7.19) and let F'_n be its commutator subgroup (see Definition 5.20). Note that the abelianization of F_n as defined in Proposition 5.22, (3), is \mathbb{Z}^n . The short exact sequence

$$1 \longrightarrow F'_n \longrightarrow F_n \longrightarrow \mathbb{Z}^n \longrightarrow 1$$

does not split.

From now on, we restrict to the case of exact sequences

$$(5.18) \quad 1 \rightarrow A \xrightarrow{\iota} G \xrightarrow{\pi} Q \rightarrow 1,$$

where A is an abelian group. Recall that the set of derivations $Der(Q, A)$ has a natural structure of an abelian group.

REMARKS 5.132. (1) The short exact sequence (5.18) uniquely defines an action of Q on A . Indeed, G acts on A by conjugation and, since the kernel of this action contains A , it defines an action of Q on A . In what follows we shall denote this action by $(q, a) \mapsto q \cdot a$, and by φ the homomorphism $Q \rightarrow \text{Aut}(A)$ defined by this action.

(2) If the short exact sequence (5.18) splits, the group G is isomorphic to $A \rtimes_{\varphi} Q$.

Classification of splittings.

Below we discuss classification of all splittings of short exact sequences (5.18) which do split. We use the additive notation for the binary operation on A . We begin with few observations. From now on, we fix a section σ_0 and, hence, a semidirect product decomposition $G = A \rtimes Q$. Note that every splitting of a short exact sequence (5.18), is determined by a section $\sigma : Q \rightarrow G$. Furthermore, every section $\sigma : Q \rightarrow G$ is determined by its components (d_{σ}, π) with respect to the semidirect product decomposition given by σ_0 (see Remark 5.125). Since π is fixed, a section σ is uniquely determined by its derivation d_{σ} . Conversely, every derivation $d \in Der(Q, A)$ determines a section σ , so that $d = d_{\sigma}$. Thus, the set of sections of (5.18) is in bijective correspondence with the abelian group of derivations $Der(Q, A)$.

Our next goal is to discuss the equivalence relation between different sections (and derivations). We say that an automorphism $\alpha \in \text{Aut}(G)$ is a *shearing* (with respect to the semidirect product decomposition $G = A \rtimes Q$) if $\alpha(A) = A$, $\alpha|_A = \text{Id}$ and α projects to the identity on Q . Examples of shearing automorphisms are *principal shearing automorphisms*, which are given by conjugations by elements $a \in A$. It is clear that shearing automorphisms act on splittings of the short exact sequence (5.18).

EXERCISE 5.133. The group of shearing automorphisms of G is isomorphic to the abelian group $Der(Q, A)$: Every derivation $d \in Der(Q, A)$ determines a shearing automorphism $\alpha = \alpha_d$ of G by the formula

$$\alpha(a \star q) = (a + d(q)) \star q,$$

which gives the bijective correspondence.

In view of this exercise, the classification of splittings modulo shearing automorphisms yields a very boring answer: All sections are equivalent under the group of shearing transformations. A finer classification of splittings is given by the following definition. Two splittings σ_1, σ_2 are said to be *A-conjugate* if they differ by a principal shearing automorphism: There exists $a \in A$ such that

$$\sigma_2(q) = a\sigma_1(q)a^{-1}, \forall q \in Q.$$

If d_1, d_2 are the derivations corresponding to the sections σ_1, σ_2 , then

$$(d_2(q), q) = (a, 1)(d_1(q), q)(-a, 1) \Leftrightarrow d_2(q) = d_1(q) - [q \cdot a - a].$$

In other words, d_1, d_2 differ by the principal derivation corresponding to $a \in A$.

Thus, we proved the following

PROPOSITION 5.134. *A-conjugacy classes of splittings of the short exact sequence (5.18) are in bijective correspondence with the quotient*

$$\text{Der}(Q, A)/\text{Prin}(Q, A),$$

where $\text{Prin}(Q, A)$ is the subgroup of principal derivations.

Note that $\text{Der}(Q, A) \cong Z^1(Q, A)$, $\text{Prin}(Q, A) = B^1(Q, A)$ and the quotient $\text{Der}(Q, A)/\text{Prin}(Q, A)$ is $H^1(Q, A)$, the first cohomology group of Q with coefficients in the $\mathbb{Z}Q$ -module A .

Below is another application of $H^1(Q, A)$. Let L be a group and let

$$F : L \rightarrow G = A \rtimes Q$$

be a homomorphism. The group G , of course, acts on the homomorphisms F by postcomposition with inner automorphisms. Two homomorphisms are said to be *conjugate* if they belong to the same orbit of this G -action.

LEMMA 5.135. *1. A homomorphism $F : L \rightarrow G$ is conjugate to a homomorphism with the image in Q if and only if the derivation d_F of F is principal.*

2. Furthermore, suppose that $F_i : L \rightarrow G$ are homomorphisms with components $(d_i, \pi), i = 1, 2$. Then F_1 and F_2 are A -conjugate if and only if

$$[d_1] = [d_2] \in H^1(L, A).$$

PROOF. Let $g = qa \in G, a \in A, q \in Q$. If $(qa)F(\ell)(qa)^{-1} \in Q$, then $aF(\ell)a^{-1} \in Q$. Thus, for (1) it suffices to consider A -conjugation of homomorphisms $F : L \rightarrow G$. Hence, (2) \Rightarrow (1). To prove (2) we note that the composition of F with an inner automorphism defined by $a \in A$ has the derivation equal to $d_F - d_a$, where d_a is the principal derivation determined by a . \square

5.9.6. Central coextensions and second cohomology. We restrict ourselves to the case of central coextensions (a similar result holds for general extensions with abelian kernels, see e.g. [Bro82b]). In this case, A is trivial as a G -module and, hence, $H^*(G, A) \cong H^k(K(G, 1), A)$. This cohomology group can be also computed as $H^k(Y, A)$, where $G = \pi_1(Y)$ and Y is a $k + 1$ -connected cell complex.

Let G be a group and A an abelian group. A *central coextension* of G by A is a short exact sequence

$$1 \rightarrow A \xrightarrow{\iota} \tilde{G} \xrightarrow{r} G \rightarrow 1$$

where $\iota(A)$ is contained in the center of \tilde{G} . Choose a set-theoretic section

$$s : G \rightarrow \tilde{G}, s(1) = 1, r \circ s = \text{Id}.$$

Then the group \tilde{G} is identified (as a set) with the direct product $A \times G$. With this identification, the group operation on \tilde{G} has the form

$$(a, g) \cdot (b, h) = (a + b + f(g, h), gh),$$

where $f(1, 1) = 0 \in A$. Here the function $f : G \times G \rightarrow A$ measures the failure of s to be a homomorphism:

$$f(g, h) = s(g)s(h)(s(gh))^{-1}.$$

Not every function $f : G \times G \rightarrow A$ corresponds to a central extension:

EXERCISE 5.136. A function f gives rise to a central coextension if and only if it satisfies the *cocycle identity*:

$$f(g, h) + f(gh, k) = f(h, k) + f(g, hk).$$

In other words, the set of such functions is the abelian group of cocycles $Z^2(G, A)$, see Section 5.9.2. We will refer to f simply as a *cocycle*.

Two central coextensions are said to be *equivalent* if there exist an isomorphism τ making the following diagram commutative:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & A & \longrightarrow & \tilde{G}_1 & \longrightarrow & G & \longrightarrow & 1 \\ & & \downarrow \text{id} & & \downarrow \tau & & \downarrow \text{id} & & \\ 1 & \longrightarrow & A & \longrightarrow & \tilde{G}_2 & \longrightarrow & G & \longrightarrow & 1 \end{array}$$

EXERCISE 5.137. A coextension is trivial, meaning equivalent to the product $A \times G$, if and only if the central coextension splits.

We will use the notation $\mathbb{E}(G, A)$ to denote the set of equivalence classes of coextensions. In the language of cocycles, $r_1 \sim r_2$ if and only if

$$f_1 - f_2 = \delta c,$$

where $c : G \rightarrow A$, and

$$\delta c(g, h) = c(g) + c(h) - c(gh)$$

is the coboundary, $\delta c \in B^2(G, A)$. Recall that

$$H^2(G, A) = Z^2(G, A)/B^2(G, A)$$

is the 2-nd cohomology group of G with coefficients in A .

The set $\mathbb{E}(G, A)$ has natural structure of an abelian group, where the sum of two coextensions

$$A \rightarrow G_i \xrightarrow{r_i} G$$

is defined by

$$\begin{aligned} G_3 &= \{(g_1, g_2) \in G_1 \times G_2 \mid r_1(g_1) = r_2(g_2)\} \xrightarrow{r} G, \\ r(g_1, g_2) &= r_1(g_1) = r_2(g_2). \end{aligned}$$

The kernel of this coextension is the subgroup A embedded diagonally in $G_1 \times G_2$. In the language of cocycles $f : G \times G \rightarrow A$, the sum of coextensions corresponds to the sum of cocycles and the trivial element is represented by the cocycle $f = 0$.

To summarize:

THEOREM 5.138 (See Chapter IV in [Bro82b]). *There exists an isomorphism of abelian groups*

$$H^2(K(G, 1), A) \cong H^2(G, A) \rightarrow \mathbb{E}(G, A).$$

The conclusion, thus, is that a group G with non-trivial 2-nd cohomology group $H^2(G, A)$ admits non-trivial central coextensions with the kernel A . How does one construct groups with non-trivial $H^2(G, A)$? Suppose that G admits a finite 2-dimensional $K(G, 1)$ complex Y , such that $\chi(G) := \chi(Y) \geq 2$. Then for $A \cong \mathbb{Z}$, we have

$$\chi(G) = 1 - b_1(Y) + b_2(Y) \geq 2 \Rightarrow b_2(Y) > 0.$$

The universal coefficients theorem shows that, for such groups G , if A is an abelian group which admits an epimorphism to \mathbb{Z} , then $H^2(G, A) \neq 0$.

Pull-backs of central coextensions. We fix an abelian group A and consider behavior of the groups $\mathbb{E}(G, A)$ under group homomorphisms $f : G_1 \rightarrow G_2$.

LEMMA 5.139. *Every homomorphism $f : G_1 \rightarrow G_2$ induces a homomorphism*

$$f^* : \mathbb{E}(G_2, A) \rightarrow \mathbb{E}(G_1, A).$$

Moreover, f lifts to a homomorphism of the corresponding central extensions $\tilde{G}_1 \rightarrow \tilde{G}_2$.

PROOF. Given a central coextension e_2 :

$$0 \rightarrow A \rightarrow \tilde{G}_2 \xrightarrow{p_2} G_2 \rightarrow 1,$$

we define a group \tilde{G}_1 as the *fiber product*:

$$\tilde{G}_1 := \{(g_1, \tilde{g}_2) \in G_1 \times \tilde{G}_2 : f(g_1) = p_2(\tilde{g}_2)\}.$$

The reader will verify that \tilde{G}_1 is a subgroup of the direct product $G_1 \times \tilde{G}_2$. The subgroup $A < 1 \times \tilde{G}_2$ is contained in the center of the product group. The subgroup $\tilde{G}_1 < G_1 \times \tilde{G}_2$ admits two projections: The projection to the first factor, G_1 , which we denote p_1 and the projection to the second factor \tilde{G}_2 , which we denote \tilde{f} . Let us identify the kernel of the homomorphism p_1 :

$$p_1(g_1, \tilde{g}_2) = 1 \iff p_2(\tilde{g}_2) = 1 \iff \tilde{g}_2 \in A.$$

Therefore, the kernel of p_1 is naturally isomorphic to the group A . Hence, we obtain a central coextension $e_1 = f^*(e_2)$:

$$(5.19) \quad 0 \rightarrow A \rightarrow \tilde{G}_1 \rightarrow G_1 \rightarrow 1$$

and a homomorphism $\tilde{f} : \tilde{G}_1 \rightarrow \tilde{G}_2$, such that the following digram is commutative:

$$\begin{array}{ccc} \tilde{G}_1 & \xrightarrow{\tilde{f}} & \tilde{G}_2 \\ p_1 \downarrow & & \downarrow p_2 \\ G_1 & \xrightarrow{f} & G_1 \end{array} .$$

Thus, f indeed determines a natural map $f^* : \mathbb{E}(G_2, A) \rightarrow \mathbb{E}(G_1, A)$. We leave it to the reader to verify that f^* is a homomorphism. \square

- EXERCISE 5.140. 1. If f is surjective, so is \tilde{f} .
 2. If s_2 is a set-theoretic section of p_2 , then

$$s_1(g_1) = (g_1, s_2 f(g_1))$$

is a set-theoretic section of p_1 .

3. Use Part 2 to verify commutativity of the diagram:

$$\begin{array}{ccc} H^2(G_2, A) & \xrightarrow{H^2(f)} & H^2(G_2, A) \\ \downarrow & & \downarrow \\ \mathbb{E}(G_2, A) & \xrightarrow{f^*} & \mathbb{E}(G_2, A) \end{array} .$$

Here the vertical arrows are the isomorphisms given by Theorem 5.138 and $H^2(f)$ is the homomorphism of second cohomology groups induced by $f : G_1 \rightarrow G_2$. On the level of cocycles, the homomorphism $H^2(f)$ is given by

$$\begin{aligned} \omega_2 \in Z^2(G_2, A) &\mapsto \omega_1 \in Z^2(G_1, A), \\ \omega_1(x, y) &= \omega_2(f(x), f(y)). \end{aligned}$$

Let us also identify the kernel of the homomorphism \tilde{f} . Suppose that s_1, s_2 are the sections as in Part 2 of Exercise 5.140. Assume that the section s_2 is normalized: $s_2(1) = 1$. Then for each $k \in K = \text{Ker}(f)$, $s_1(k) = (k, 1)$, i.e. the restriction of s_1 to K is a homomorphism (even though, $s_1 : G_1 \rightarrow \tilde{G}_1$ is not, in general). Since \tilde{f} is the restriction of the projection to the second factor, we conclude that

$$\tilde{K} = \text{Ker}(\tilde{f}) = s_1(K).$$

In particular, kernels of f and \tilde{f} are isomorphic.

Suppose for a moment, that the central coextension (5.19) splits, i.e. there exists a homomorphism $s : G_1 \rightarrow \tilde{G}_1$ right-inverse to p_1 . Then the homomorphisms $s|_K$ and $s_1|_K$ differ by a homomorphism $\varphi : K \rightarrow A$:

$$s_1(k) = s(k)\varphi(k),$$

where we identify A with the subgroup $1 \times A < \tilde{G}_1$. Since the subgroup A is contained in the center of \tilde{G}_1 , we obtain:

$$\varphi(gkg^{-1}) = \varphi(k)$$

for all $k \in K, g \in G_1$. In other words, the action of G_1 by conjugation on K fixes the homomorphism φ .

We will be using central coextensions in Section 11.19 to give some interesting examples of quasiisometric groups and in Section 19.9 to give examples of failure of quasiisometric invariance of Property (T).

Median spaces and spaces with measured walls

Median spaces discussed in this chapter compose a large class of metric spaces containing, among others, metric trees, L^1 -spaces and 1-skeleta of $CAT(0)$ cube complexes equipped with the standard metric. Spaces with walls (more precisely, spaces with a measured wall structure) provide key examples of median spaces. In this section we establish basic properties of median spaces and spaces with walls which will be used later on, in chapter 19, in order to establish geometric criteria for Properties (T) and a-T-menability of groups in terms of their actions on median spaces and spaces with walls.

6.1. Median spaces

DEFINITION 6.1 (intervals and geodesic sequences). Let (X, pdist) be a pseudo-metric space. A point b is *between a and c* if $\text{pdist}(a, b) + \text{pdist}(b, c) = \text{pdist}(a, c)$. We denote by $I(a, c)$ the set of points that are between a and c , and we call $I(a, c)$ the *interval between a and c* . A (finite, discrete) *path* in (X, pdist) is a finite sequence of points (a_1, a_2, \dots, a_n) . A path is called a *geodesic sequence* if and only if

$$\text{pdist}(a_1, a_n) = \text{pdist}(a_1, a_2) + \text{pdist}(a_2, a_3) + \dots + \text{pdist}(a_{n-1}, a_n).$$

Thus, (a, b, c) is a geodesic sequence if and only if $b \in I(a, c)$.

DEFINITION 6.2 (median point). Let a, b, c be three points of a pseudo-metric space (X, dist) . We denote the intersection $I(a, b) \cap I(b, c) \cap I(a, c)$ by $M(a, b, c)$, and we call any point in $M(a, b, c)$ a *median point* for a, b, c .

We note that

$$I(a, b) = \{x \in X, x \in M(a, x, b)\}.$$

DEFINITION 6.3 (median spaces). A *median pseudo-metric space* is a pseudo-metric space in which for any three points x, y, z the set $M(x, y, z)$ is non-empty and of diameter zero (any two median points are at pseudo-distance 0). In particular a metric space is median if any three points x, y, z have one and only one median point, which we will denote by $m(x, y, z)$.

Note that a pseudo-metric space is median if and only if its metric quotient is median.

A subset $Y \subset X$ in a median space is a *median subspace* if for any three points x, y, z in Y , we have $M(x, y, z) \subset Y$. Note that Y is then median for the induced pseudo-metric. The intersection of median subspaces is a median subspace, thus any subset $Y \subset X$ is contained in a smallest median subspace, which we call the *median hull* of Y .

CONVENTION 6.4. Throughout the book, we call median metric spaces simply *median spaces*.

DEFINITION 6.5. We say that a metric space (X, dist) is *submedian* if it admits an isometric embedding into a median space.

A median space together with the ternary operation $(x, y, z) \mapsto m(x, y, z)$ is a particular instance of a *median algebra*. For literature on median algebras, we refer the reader to [Sho54a], [Sho54b], [Nie78], [Isb80], [BH83], [vdV93], [Bas01]. Geometrical approaches of median spaces can be found in [Rol16] and [Nic08].

In what follows we use some classical results in the theory of median algebras, leaving it as an exercise to the reader to reprove them in the metric context.

6.1.1. A review of median algebras. The notion of median algebra appeared as a common generalization of trees and *distributive lattices*, cf. Example 6.11. We recall here some basic definitions and properties related to median algebras. For proofs and further details we refer the reader to the books [vdV93], [Ver93], the surveys [BH83], [Isb80], and the papers [BK47], [Sho54a], [Sho54b] and [Rol16].

DEFINITION 6.6 (median algebra, the first definition). A *median algebra* is a set X endowed with a ternary operation $(a, b, c) \mapsto m(a, b, c)$ such that:

- (1) $m(a, a, b) = a$;
- (2) $m(a, b, c) = m(b, a, c) = m(b, c, a)$;
- (3) $m(m(a, b, c), d, e) = m(a, m(b, d, e), m(c, d, e))$.

The property (3) can be replaced by:

$$(3') \quad m(a, m(a, c, d), m(b, c, d)) = m(a, c, d).$$

The element $m(a, b, c)$ is the *median of the points a, b, c* . In a median algebra (X, m) , given any two points a, b the *interval with endpoints a, b* is the set

$$I(a, b) = \{x; x = m(a, b, x)\}.$$

This defines a map $I : X \times X \rightarrow \mathcal{P}(X)$. We say that points $x \in I(a, b)$ are *between a and b* .

A *homomorphism* of median algebras is a map $f : (X, m_X) \rightarrow (Y, m_Y)$ such that $m_Y(f(x), f(y), f(z)) = f(m_X(x, y, z))$. Equivalently, f is a homomorphism if and only if it preserves the betweenness relation. If moreover f is injective (bijective) then f is called *embedding* or *monomorphism* (respectively *isomorphism*) of median algebras.

The following are straightforward properties of median algebras, see e.g. [Sho54a] and [Rol16, §2].

LEMMA 6.7. *Let (X, m) be a median algebra. For $x, y, z \in X$ we have that*

- (1) $I(x, x) = \{x\}$;
- (2) $I(x, y) \cap I(x, z) = I(x, m(x, y, z))$;
- (3) $I(x, y) \cap I(x, z) \cap I(y, z) = \{m(x, y, z)\}$;
- (4) *if $a \in I(x, y)$ then for any t , $I(x, t) \cap I(y, t) \subseteq I(a, t)$ (equivalently $m(x, y, t) \in I(a, t)$);*
- (5) *if $x \in I(a, b)$ and $y \in I(x, b)$ then $x \in I(a, y)$.*

A sequence of points (a_1, a_2, \dots, a_n) is *geodesic* in the median algebra (X, m) if $a_i \in I(a_1, a_{i+1})$ for all $i = 2, \dots, n-1$. This is equivalent, by Lemma 6.7, part (5), to the condition that $a_{i+1} \in I(a_i, a_n)$ for all $i = 1, 2, \dots, n-2$.

LEMMA 6.8. *If (x, t, y) is a geodesic sequence, then:*

- (1) $I(x, t) \cup I(t, y) \subseteq I(x, y)$;
- (2) $I(x, t) \cap I(t, y) = \{t\}$.

According to [Sho54a], [Sho54b] there is an alternative definition of median algebras, using intervals:

DEFINITION 6.9 (median algebra, the second definition). A *median algebra* is a set X endowed with a map $I : X \times X \rightarrow \mathcal{P}(X)$ such that:

- (1) $I(x, x) = \{x\}$;
- (2) if $y \in I(x, z)$ then $I(x, y) \subset I(x, z)$;
- (3) for every x, y, z in X the intersection $I(x, y) \cap I(x, z) \cap I(y, z)$ has cardinality 1.

EXERCISE 6.10. Let (X, dist) be a median space. Then the metric intervals $I(x, y)$ satisfy the properties in Definition 6.9, and thus the metric median $(x, y, z) \mapsto m(x, y, z)$ defines a structure of median algebra on X .

EXAMPLE 6.11. For any set X , the power set $\mathcal{P}(X)$ is a median algebra when endowed with the Boolean median operation

$$(6.1) \quad m(A, B, C) = (A \cap B) \cup (A \cap C) \cup (B \cap C) = (A \cup B) \cap (A \cup C) \cap (B \cup C).$$

The median algebra $(\mathcal{P}(X), m)$ is called a *Boolean median algebra*.

EXERCISE 6.12. Show that in this example

$$(6.2) \quad I(A, B) = \{C ; A \cap B \subset C \subset A \cup B\}.$$

6.1.2. Convexity.

DEFINITION 6.13. A *convex subset* A in a median algebra is a subset such that for all $a, b \in A$, $I(a, b) \subset A$.

A subset h in a median space (X, m) is called a *convex half-space* if both h itself and the complementary set $h^c = X \setminus h$ are convex. The pair $\{h, h^c\}$ is called a *convex wall*. We denote by $\mathcal{D}_c(X)$ the set of convex half-spaces in X and by $\mathcal{W}_c(X)$ the set of convex walls in X . When there is no possibility of confusion we simply use the notations \mathcal{D}_c and \mathcal{W}_c .

EXERCISE 6.14. A subset A in a median algebra (X, m) is convex if and only if for every $x \in X$, and a, b in A , the element $m(a, x, b)$ is in A .

A convex wall $\{h, h^c\}$ is said to *separate* subsets $A, B \subset X$ if $A \subset h$ and $B \subset h^c$ or vice versa.

The above algebraic notion of convexity coincides with the metric notion of convexity introduced in Definition 6.23, in the case of the median algebra associated with a median metric space (see Exercise 6.10).

The following theorem shows abundance of convex walls in median algebras:

THEOREM 6.15. *Let X be a median algebra, and let A, B be convex non-empty disjoint subsets of X . Then there exists a convex wall separating A and B .*

A proof of Theorem 6.15 when A is a singleton can be found in [Nie78]; in its most general form it follows from [dV84, Theorem 2.5]. Other proofs can be found in [Bas01, §5.2] and in [Rol16, §2].

COROLLARY 6.16. *Given any two distinct points x, y in a median space (X, dist) there exists a convex wall $w = \{h, h^c\}$ with $x \in h, y \in h^c$.*

DEFINITION 6.17. Given a median algebra X , we define the map

$$\sigma : X \rightarrow \mathcal{P}(\mathcal{D}_c), \sigma(x) = \sigma_x = \{h \in \mathcal{D}_c ; x \in h\}.$$

EXERCISE 6.18. The map σ is a monomorphism of median algebras, where $\mathcal{P}(\mathcal{D}_c)$ is endowed with the median algebra structure described in Example 6.11. Hint: Use Theorem 6.15.

This shows that Boolean algebras, in a sense, are “universal” median algebras.

6.1.3. Examples of median metric spaces.

EXAMPLES 6.19. (1) In the real line \mathbb{R} , the metric intervals are precisely the closed order intervals, i.e. $I(x, y) = [x, y]$. The median function sends the triple (a, b, c) to the “middle” element of the set $\{a, b, c\}$, i.e.

$$m_{\mathbb{R}}(a, b, c) = a + b + c - [\max(a, b, c) + \min(a, b, c)].$$

(2) More generally, in \mathbb{R}^n with the ℓ_1 norm, the interval $I(\mathbf{x}, \mathbf{y})$ between two points $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ is the product of intervals

$$[x_1, y_1] \times \cdots \times [x_n, y_n].$$

The median point is given by

$$m(\mathbf{x}, \mathbf{y}, \mathbf{z}) = (m_{\mathbb{R}}(x_1, y_1, z_1), \dots, m_{\mathbb{R}}(x_n, y_n, z_n)).$$

(3) The ℓ_1 -product of two pseudo-metric spaces (X_1, pdist_1) and (X_2, pdist_2) is the set $X_1 \times X_2$, endowed with the pseudo-metric

$$\text{pdist}((x_1, x_2), (y_1, y_2)) = \text{pdist}_1(x_1, y_1) + \text{pdist}_2(x_2, y_2).$$

Intervals in the product are the Cartesian products of intervals, i.e.

$$I_{X_1 \times X_2}((x_1, x_2), (y_1, y_2)) = I_{X_1}(x_1, y_1) \times I_{X_2}(x_2, y_2).$$

The space $(X_1 \times X_2, \text{pdist})$ is median if and only if (X_1, pdist_1) and (X_2, pdist_2) are median (the components of a median point in $X_1 \times X_2$ are median points of the components).

(4) (trees) Every \mathbb{R} -tree is a median space, with intervals equal to closed geodesic segments.

(5) (motivating example: CAT(0) cube complexes) The 0-skeleton of a CAT(0) cube complex with the standard metric is a (discrete) median space. In fact, according to [Che00, Theorem 6.1] the converse is also true: the set of vertices of a simplicial graph is median if and only if the graph is the 1-skeleton of a CAT(0) cube complex.

(6) If (W, S) is a Coxeter system and dist_S is the word distance on W with respect to S then $(W, \text{dist}_S^{1/2})$ is submedian, [BJS88].

(7) Metric versions of Boolean median algebras. These examples are explained in Lemma 6.21

(8) (L^1 spaces) Every space $L^1(X, \mu)$ is median.

Given a measured space (X, \mathcal{B}, μ) we let $\mathcal{L}^1(X, \mu)$ denote the vector space of real-valued functions $f : X \rightarrow \mathbb{R}$ with finite integral

$$\|f\|_1 := \int_X |f(x)| d\mu(x).$$

Thus, $\mathcal{L}^1(X, \mu)$ is a pseudo-metric space with the pseudo-distance function given by $\text{dist}(f, g) = \|f - g\|_1$. The quotient metric space of $\mathcal{L}^1(X, \mu)$ is $L^1(X, \mu)$.

LEMMA 6.20. *For every measured space (X, \mathcal{B}, μ) , the metric space $L^1(X, \mu)$ is median.*

PROOF. It is enough to see that the vector space $\mathcal{L}^1(X, \mu)$ is a median pseudo-metric space. A function $p \in \mathcal{L}^1(X, \mu)$ belongs to $I(f, g)$ if and only if the set of points x such that $p(x)$ is not between $f(x)$ and $g(x)$ has measure 0. This is due to two observations:

- given $u, v \in \mathcal{L}^1(X, \mu)$, $\|u+v\|_1 = \|u\|_1 + \|v\|_1$ if and only if $uv \geq 0$ almost everywhere;
- hence, given $f, g, p \in \mathcal{L}^1(X, \mu)$, $\|f - g\|_1 = \|f - p\|_1 + \|p - g\|_1$ if and only if $p(x)$ is in the interval with endpoints $f(x), g(x)$ for almost every $x \in X$.

Define on $\mathcal{L}^1(X, \mu)$ the ternary operation $(f, g, h) \mapsto m(f, g, h)$ by

$$m(f, g, h)(x) = m_{\mathbb{R}}(f(x), g(x), h(x)).$$

The function

$$m = m(f, g, h) = f + g + h - \max(f, g, h) - \min(f, g, h)$$

is measurable because the sum, the maximum and the minimum of measurable functions is measurable. Since m is (a.e.) pointwise between f and g , $m \in I(f, g)$. Similarly, we have $m \in I(g, h)$ and $m \in I(f, h)$, and thus $m(f, g, h)$ is a median for the functions f, g, h . It follows that $M(f, g, h)$ is the set of functions that are almost everywhere equal to $m(f, g, h)$, and hence $\mathcal{L}^1(X, \mu)$ is a median pseudo-metric space. We conclude that $L^1(X, \mu)$ is median because it is the metric quotient of $\mathcal{L}^1(X, \mu)$. \square

A measure-theoretic version of the Example 6.11 yields a metric version of Boolean algebras, which, moreover, appear as subspaces of \mathcal{L}^1 -spaces. Let (X, \mathcal{B}, μ) denote a measured space. For each subset $A \subset X$, we define

$$\mathcal{B}_A = \{B \subseteq X \mid A \Delta B \in \mathcal{B}, \mu(A \Delta B) < +\infty\}.$$

Notice that we do not require the sets in \mathcal{B}_A to be measurable, only their symmetric difference with A should be. Denote as usual by χ_C the characteristic function of $C \subset X$. The map $\chi^A : \mathcal{B}_A \rightarrow \mathcal{L}^1(X, \mu)$ defined by $B \mapsto \chi_{A \Delta B}$ is injective. The range of χ^A consists of the set $\mathcal{S}^1(X, \mu)$ of all characteristic functions of measurable subsets with finite measure. Indeed the preimage of $\chi_{B'}$ (with $B' \in \mathcal{B}, \mu(B') < +\infty$) is the subset $B := A \Delta B'$. Observe that the \mathcal{L}^1 -pseudo-distance between two functions $\chi_{B'}$ and $\chi_{C'}$ in $\mathcal{S}^1(X, \mu)$ equals $\mu(B' \Delta C')$. Since we have

$$(A \Delta B) \Delta (A \Delta C) = B \Delta C,$$

it follows that for any two elements $B_1, B_2 \in \mathcal{B}_A$ the symmetric difference $B_1 \Delta B_2$ is measurable with finite measure, and the pull-back of the \mathcal{L}^1 -pseudo-distance via the bijection $\mathcal{B}_A \rightarrow \mathcal{S}^1(X, \mu)$ is the pseudo-metric $\text{pdist}_{\mathcal{B}}$ defined by $\text{pdist}_{\mathcal{B}}(B_1, B_2) =$

$\mu(B_1 \Delta B_2)$. The interval $I(B_1, B_2)$ is composed of subsets $C \in \mathcal{B}_A$ such that there exists C' satisfying $\mu(C' \Delta C) = 0$ (hence $C' \in \mathcal{B}_A$) and $B_1 \cap B_2 \subset C' \subset B_1 \cup B_2$.

LEMMA 6.21. $(\mathcal{B}_A, \text{pdist}_{\mathcal{B}})$ is a median pseudo-metric space, equivalently, the space $\mathcal{S}^1(X, \mu)$ is a median subspace of $\mathcal{L}^1(X, \mu)$, cf. Lemma 6.20.

PROOF. The claim follows from the explicit formula:

$$m(B_1, B_2, B_3) = (B_1 \cup B_2) \cap (B_1 \cup B_3) \cap (B_2 \cup B_3) = (B_1 \cap B_2) \cup (B_1 \cap B_3) \cup (B_2 \cap B_3).$$

We leave details to the reader. \square

Later on (Corollary 6.59) we will prove that every median space embeds isometrically as a median subspace of some space $L^1(X, \mu)$ (compare with the similar result in the context of median algebras appearing in Exercise 6.18).

REMARK 6.22. It is impossible, in general, to define for a given submedian space Y its “median completion,” that is, a median space \tilde{Y} containing an isometric copy of Y , and such that any isometric embedding of Y into a median space X extends to an isometric embedding $\tilde{Y} \rightarrow X$. This can be seen in the following example.

Let $E = \mathbb{R}^7$ endowed with the ℓ_1 norm, and let $\{\mathbf{e}_i ; i = 1, 2, \dots, 7\}$ be the canonical basis in E . Given $t \in [0, 1]$ let Y_t be the set composed of the four points A, B, C, D in E defined by $A = \frac{t}{2}(\mathbf{e}_1 + \mathbf{e}_2 + \mathbf{e}_3) + (1-t)\mathbf{e}_4$, $B = \frac{t}{2}(-\mathbf{e}_1 - \mathbf{e}_2 + \mathbf{e}_3) + (1-t)\mathbf{e}_5$, $C = \frac{t}{2}(\mathbf{e}_1 - \mathbf{e}_2 - \mathbf{e}_3) + (1-t)\mathbf{e}_6$, $D = \frac{t}{2}(-\mathbf{e}_1 + \mathbf{e}_2 - \mathbf{e}_3) + (1-t)\mathbf{e}_7$.

Any two distinct points in Y_t are at the ℓ_1 -distance 2. Thus, all the subsets $Y_t \subset E$ equipped with the ℓ_1 -distance are pairwise isometric. The median hull of Y_t in E is composed of Y_t itself and of the eight vertices of a cube of edge length t , namely,

$$\frac{t}{2}(\pm \mathbf{e}_1 \pm \mathbf{e}_2 \pm \mathbf{e}_3).$$

Thus, for two distinct values $t \neq t'$ the median hulls of Y_t and of $Y_{t'}$ are not isometric.

Furthermore, the median hull of Y_0 is a 5-point set $(Y_0 \cup \{\mathbf{0}\})$, while the median hull of $Y_t, t \neq 0$ consists of 12 points. Consequently, in general, the median hulls of two isometric submedian spaces may not even be isomorphic as median algebras.

6.1.4. Convexity and gate property in median spaces.

DEFINITION 6.23. Let (X, pdist) be a pseudo-metric space. A subset $Y \subset X$ is said to be *convex* if for all $a, b \in Y$, the set $I(a, b)$ is contained in Y . A subset $Y \subset X$ is *quasi-convex* if there exists $R < \infty$ such that for all $a, b \in Y$ the set $I(a, b)$ is contained in $\overline{\mathcal{N}}_R(Y)$. The *convex hull* of a subset $Y \subset X$ is the intersection of all convex subsets containing Y .

Note that any convex subspace of a median space is median but not vice versa, as for instance any subset E of cardinality two is a median subspace, while E might not be convex. The median hull of a subset is contained in the convex hull, and the example above shows the inclusion may be strict.

We now introduce a notion related to convexity in median spaces, which is commonly used in the theory of Tits buildings (see for example [Sch85]) and in graph theory ([Mul80], [vdV93]).

DEFINITION 6.24 (gate). Let (X, dist) be a metric space, let Y be a subset of X , and x a point in X .

We say that a point $p \in X$ is *between x and Y* if p is between x and every $y \in Y$. When a point $p \in Y$ is between x and Y , we say that p is a *gate between x and Y* . Note that there is always at most one gate p between x and Y , and that $\text{dist}(x, p) = \text{dist}(x, Y)$.

We say that a subset $Y \subset X$ is *gate-convex* if for every point $x \in X$ there exists a gate (in Y) between x and Y . We then denote by $\pi_Y(x)$ this gate, and call the map π_Y the *projection map onto Y* .

EXERCISE 6.25. π_Y restricts to the identity map on Y .

LEMMA 6.26 (gate-convex subsets). (1) *The projection map π_Y onto a gate-convex subset $Y \subset X$ is 1-Lipschitz.*

(2) *Any gate-convex subset is closed and convex.*

(3) *In a complete median space, any closed convex subset is gate-convex.*

In other words, for closed subsets of a complete median space, convexity is equivalent to gate-convexity.

PROOF. (1) Let x, x' be two points in a metric space X , and let p, p' be the respective gates between x, x' and a gate-convex subset Y . Since (x, p, p') and (x', p', p) are geodesic sequences, we have that

$$\begin{aligned} \text{dist}(x, p) + \text{dist}(p, p') &\leq \text{dist}(x, x') + \text{dist}(x', p') \\ \text{dist}(x', p') + \text{dist}(p', p) &\leq \text{dist}(x', x) + \text{dist}(x, p) \end{aligned}$$

By summing up the two inequalities, we conclude that $\text{dist}(p, p') \leq \text{dist}(x, x')$.

(2) Assume that Y is gate-convex and that (x, y, z) is a geodesic sequence with $x, z \in Y$. Let p be the gate between y and Y , so that (y, p, x) and (y, p, z) are geodesic sequences. Hence (x, p, y, p, z) is a geodesic sequence, which forces $y = p \in Y$.

Any point x in the closure of Y satisfies $\text{dist}(x, Y) = 0$. Thus, if p is the gate between x and Y we have $\text{dist}(x, p) = 0$, hence $x \in Y$. We conclude that Y is closed.

(3) Let Y be a closed convex subset of a complete median space X . For $x \in X$ choose a sequence $(y_k)_{k \geq 0}$ of points in Y such that $\text{dist}(y_k, x)$ tends to $\text{dist}(x, Y)$. First observe that $(y_k)_{k \geq 0}$ is a Cauchy sequence. Indeed, denote by $\epsilon_k = \text{dist}(y_k, x) - \text{dist}(Y, x)$, which clearly is a sequence of positive numbers converging to zero. Let $m_{k, \ell}$ be the median point of (x, y_k, y_ℓ) . Then

$$\text{dist}(x, y_k) + \text{dist}(x, y_\ell) = 2\text{dist}(x, m_{k, \ell}) + \text{dist}(y_k, y_\ell)$$

and so by convexity of Y we have

$$\text{dist}(x, y_k) + \text{dist}(x, y_\ell) \geq 2\text{dist}(x, Y) + \text{dist}(y_k, y_\ell).$$

It follows that $\text{dist}(y_k, y_\ell) \leq \epsilon_k + \epsilon_\ell$. Since X is complete, the sequence $(y_k)_{k \geq 0}$ has a limit p in X . Since Y is closed, the point p is in Y . Note that $\text{dist}(x, p) = \text{dist}(x, Y)$. It remains to check that p is between x and Y .

Let y be some point in Y , and let m be the median point of x, p, y . By convexity of Y we have $m \in Y$, so that $\text{dist}(x, m) \geq \text{dist}(x, Y)$. We also have $\text{dist}(x, p) =$

$\text{dist}(x, m) + \text{dist}(m, p)$. Since $\text{dist}(x, p) = \text{dist}(x, Y)$ we get $\text{dist}(m, p) = 0$ as desired. \square

We now prove that in a median space metric intervals are gate-convex.

LEMMA 6.27. *In a median metric space X any interval $I(a, b)$ is gate-convex, and the gate between an arbitrary point x and $I(a, b)$ is $m(x, a, b)$.*

PROOF. Consider an arbitrary point $x \in X$; let p be the median point $m(x, a, b)$ let and y be an arbitrary point in $I(a, b)$. We will show that (x, p, y) is a geodesic sequence.

We consider the median points $a' = m(x, a, y)$, $b' = m(x, b, y)$ and $p' = m(x, a', b')$. Note that $p' \in I(x, a') \subset I(x, a)$ and similarly $p' \in I(x, b)$. Since (a, y, b) , (a, a', y) and (y, b', b) are geodesic sequences, the sequence (a, a', y, b', b) is geodesic as well. Thus $I(a', b') \subset I(a, b)$, hence $p' \in I(a, b)$.

We proved that $p' \in I(x, a) \cap I(x, b) \cap I(a, b)$, which by the uniqueness of the median point, implies $p' = p$. It follows that $p \in I(x, a') \subset I(x, y)$. \square

We can now deduce that the median map is 1-Lipschitz, in each variable and on $X \times X \times X$ endowed with the ℓ_1 -metric.

COROLLARY 6.28. *Let X be a median space.*

- (1) *For any two fixed points $a, b \in X$ the interval $I(a, b)$ is closed and convex, and the map $x \mapsto m(x, a, b)$ is 1-Lipschitz.*
- (2) *The median map $m : X \times X \times X \rightarrow X$ is 1-Lipschitz (here $X \times X \times X$ is endowed with the ℓ_1 -product metric as defined in Example 6.19, (1)).*

PROOF. Combine Lemma 6.27 and Lemma 6.26, and use the fact that, given six points $a, b, c, a', b', c' \in X$, the distance between the median points $m(a, b, c)$ and $m(a', b', c')$ is at most

$$\text{dist}(m(a, b, c), m(a', b, c)) + \text{dist}(m(a', b, c), m(a', b', c)) + \text{dist}(m(a', b', c), m(a', b', c')).$$

\square

6.1.5. Rectangles and parallel pairs. In a median space X , the following notion of rectangle will allow us to treat median spaces as continuous versions of the 1-skeleta of $CAT(0)$ cube complexes.

DEFINITION 6.29. A *quadrilateral* in a metric space (X, dist) is a closed path (a, b, c, d, a) , which we denote by $[a, b, c, d]$. A quadrilateral $[a, b, c, d]$ is a *rectangle* if the four sequences (a, b, c) , (b, c, d) , (c, d, a) and (d, a, b) are geodesic.

REMARK 6.30. Suppose that X is a median metric space. Then:

- (1) By the triangle inequality, in a rectangle $[a, b, c, d]$ in X the following equalities hold: $\text{dist}(a, b) = \text{dist}(c, d)$, $\text{dist}(a, d) = \text{dist}(b, c)$ and $\text{dist}(a, c) = \text{dist}(b, d)$.
- (2) (rectangles in intervals) If $x, y \in I(a, b) \subset X$ then $[x, m(x, y, a), y, m(x, y, b)]$ is a rectangle.
- (3) (subdivision of rectangles) Let $[a, b, c, d] \subset X$ be a rectangle. Let $e \in I(a, d)$ and $f = m(e, b, c)$. Then $[a, b, f, e]$ and $[c, d, e, f]$ are rectangles.

DEFINITION 6.31. (parallelism on pairs) Two pairs (a, b) and (d, c) in X are *parallel* if $[a, b, c, d]$ is a rectangle.

The following property of median spaces is analogous to transitivity of parallelism for geodesics in $CAT(0)$ spaces:

PROPOSITION 6.32. *In a median space X the parallelism of pairs is an equivalence relation.*

PROOF. Suppose now that pairs (a, d) and (b, c) are parallel and the pairs (b, c) and (f, e) are parallel; in other words, we have two rectangles $[a, b, c, d]$ and $[b, c, e, f]$. We will show that (a, d) is parallel to (f, e) , i.e. the quadrilateral $[a, d, e, f]$ is also a rectangle.

We will prove that $f \in I(a, e)$; the rest of the proof of the rectangle properties is left as an exercise to the reader, as they are obtained by relabelling the points. Define the point

$$m = m(a, c, f) = I(a, c) \cap I(c, f) \cap I(f, a).$$

Regarding m as the gate between a and the interval $I(c, f)$, containing e , Lemma 6.27 implies that $m \in I(a, e)$, i.e. the triple (a, m, e) is geodesic. Similarly, regarding m as the gate between f and the interval $I(a, c)$ containing b , Lemma 6.27 implies that $m \in I(b, f)$, i.e. the triple (b, m, f) is geodesic.

Since the triples (b, m, f) and (b, f, e) are both geodesic, the quadruple (b, m, f, e) is geodesic and, hence, (m, f, e) is geodesic. Since (a, m, e) is geodesic and (m, f, e) is geodesic, we conclude that the quadruple (a, m, f, e) is geodesic. Hence, the triple (a, f, e) is geodesic, i.e. $f \in I(a, e)$ as required. \square

We now explain how to any 4-tuple of points one can associate a rectangle.

LEMMA 6.33. *Let $[x, a, y, b]$ be any quadrilateral in a median space. Then there exists a unique rectangle $[x', a', y', b']$ satisfying the following properties:*

- (1) *the following sequences are geodesic:*

$$(x, x', a', a), (a, a', y', y), (y, y', b', b), (b, b', x', x);$$

- (2) *(a, a', b', b) is a geodesic sequence;*
- (3) *(x, x', y') and (y, y', x') are geodesic sequences.*

PROOF. *Existence.* Let $x' = m(x, a, b)$ and $y' = m(y, a, b)$, and let $a' = m(a, x', y')$ and $b' = m(b, x', y')$ (see Figure 6.1). Then $[x', a', y', b']$ is a rectangle by Remark 6.30, part (3). Properties (1) and (2) follow immediately from the construction, property (3) follows from Lemma 6.27 applied to x and $y' \in I(a, b)$, respectively to y and $x' \in I(a, b)$.

Uniqueness. Let $[x', a', y', b']$ be a rectangle satisfying the three required properties. Properties (1), (2) and the fact that $[x', a', y', b']$ is a rectangle imply that $x' = m(x, a, b)$ and $y' = m(y, a, b)$. Again property (2) and the fact that $[x', a', y', b']$ is a rectangle imply that $a' = m(a, x', y')$ and $b' = m(b, x', y')$. \square




FIGURE 6.1. Central rectangle.

DEFINITION 6.34. We call the rectangle $[x', a', y', b']$ described in Lemma 6.33 the *central rectangle* associated with the quadrilateral $[x, a, y, b]$.

REMARK 6.35. Property (3) cannot be improved to “ (x, x', y', y) is a geodesic sequence”, as shown by the example of a unit cube in \mathbb{R}^3 , with a, b two opposite vertices of the lower horizontal face, and x, y the two opposite vertices of the upper horizontal face that are not above b or d (see Figure 6.2).

Note also that in general the central rectangle associated with $[x, a, y, b]$ is distinct from the central rectangle associated with $[a, x, b, y]$ (again see Figure 6.2).




FIGURE 6.2. Example of a central rectangle.

Property (3) in Lemma 6.33 can be slightly improved as follows.

LEMMA 6.36. *Let x, y, p, q be four points such that (x, p, q) and (p, q, y) are geodesic sequences. Then there exists a geodesic sequence (x, x', y', y) such that (x', y') and (p, q) are parallel.*

PROOF. Applying Lemma 6.33 to the quadrilateral $[p, q, y, x]$, we note that the resulting central rectangle $[p', q', y', x']$ satisfies $p' = p, q' = q$. □

6.1.6. Approximate geodesics and medians; completions of median spaces. In this section we prove that the median property is preserved under metric completion. We will need an auxiliary result stating that in a median space, approximate geodesics are close to geodesics, and approximate medians are close to medians. We begin by defining approximate geodesics and medians.

DEFINITION 6.37. Let (X, dist) be a metric space and let δ be a non-negative real number. We say that z is between x and y up to the error δ provided that

$$\text{dist}(x, z) + \text{dist}(z, y) \leq \text{dist}(x, y) + \delta.$$

We say that (a_1, a_2, \dots, a_n) is a δ -geodesic sequence if

$$\text{dist}(a_1, a_2) + \text{dist}(a_2, a_3) + \dots + \text{dist}(a_{n-1}, a_n) \leq \text{dist}(a_1, a_n) + \delta.$$

NOTATION 6.38. Let x, y be two points of X . We denote by $I_\delta(a, b)$ the set of points that are between a and b up to the error δ .

Let x, y, z be three points of X . We denote by $M_\delta(a, b, c)$ the intersection

$$I_{2\delta}(a, b) \cap I_{2\delta}(b, c) \cap I_{2\delta}(a, c).$$

In accordance with the previous notation, whenever $\delta = 0$, the subscript δ is dropped.

LEMMA 6.39. *Given $\delta, \delta' \geq 0$, for every $c \in I_\delta(a, b)$ the set $I_{\delta'}(a, c)$ is contained in $I_{\delta+\delta'}(a, b)$.*

DEFINITION 6.40. Let x, y, z be three points in a metric space. If $M_\delta(x, y, z)$ is non-empty then any point in it is called a δ -median point for x, y, z .

LEMMA 6.41. *Let (X, dist) be a median space, and a, b, c three arbitrary points in X .*

(i) *The set $I_{2\delta}(a, b)$ coincides with $\overline{N}_\delta(I(a, b))$.*

(ii) *The following sequence of inclusions holds:*

$$(6.3) \quad B(m(a, b, c), \delta) \subseteq M_\delta(a, b, c) \subseteq B(m(a, b, c), 3\delta).$$

PROOF. Statement (i) immediately follows from Lemma 6.27.

The first inclusion in (6.3) is obvious. We prove the second inclusion. Consider the median points $p_1 = m(p, a, b), p_2 = m(p, b, c), p_3 = m(p, a, c), q = m(p_1, b, c), r = m(q, a, c)$.

First we show that $r = m(a, b, c)$. Indeed $r \in I(a, c)$ by definition. We also have $r \in I(q, c)$, and since $q \in I(c, b)$ it follows that $r \in I(b, c)$. Finally we have $r \in I(a, q)$. Now $q \in I(p_1, b)$ and $p_1 \in I(a, b)$, so $q \in I(a, b)$. It follows that $r \in I(a, b)$.

It remains to estimate the distance between p and r . According to (i) and Lemma 6.27 the point p is within distance at most δ from p_1, p_2 and p_3 respectively.

By Corollary 6.28 we have $\text{dist}(p_2, q) \leq \text{dist}(p, p_1) \leq \delta$. Hence $\text{dist}(p, q) \leq 2\delta$. Applying Corollary 6.28 again we get $\text{dist}(p_3, r) \leq \text{dist}(p, q) \leq 2\delta$, consequently $\text{dist}(p, r) \leq 3\delta$. \square

The following result is also proved in [Ver93, Corollary II.3.5]. For the sake of completeness we give another proof here.

PROPOSITION 6.42. *The metric completion of a median space is a median space.*

PROOF. Let (X, dist) be a median space, and let $X \rightarrow \widehat{X}$ be the metric completion. For simplicity we denote the distance on \widehat{X} also by dist .

The median map $m : X \times X \times X \rightarrow X \subset \widehat{X}$ is 1-Lipschitz by Corollary 6.28. Thus it extends to a 1-Lipschitz map $\widehat{X} \times \widehat{X} \times \widehat{X} \rightarrow \widehat{X}$, also denoted by m .

Clearly for any three points a, b, c in \widehat{X} , the point $m(a, b, c)$ is median for a, b, c . We now prove that $m(a, b, c)$ is the unique median point for a, b, c . Let p be another median point for a, b, c . The points a, b, c are limits of sequences $(a_n), (b_n), (c_n)$ of points in X . Let m_n be the median point of a_n, b_n, c_n . Set $\delta_n = \text{dist}(a, a_n) + \text{dist}(b, b_n) + \text{dist}(c, c_n)$.

We show that p is a δ_n -median point for a_n, b_n, c_n . Indeed we have that $\text{dist}(a_n, p) + \text{dist}(p, b_n)$ is at most

$$\begin{aligned} \text{dist}(a_n, a) + \text{dist}(a, p) + \text{dist}(p, b) + \text{dist}(b, b_n) &= \text{dist}(a_n, a) + \text{dist}(a, b) + \text{dist}(b, b_n) \leq \\ &2\text{dist}(a, a_n) + \text{dist}(a_n, b_n) + 2\text{dist}(b, b_n) \leq \text{dist}(a_n, b_n) + 2\delta_n. \end{aligned}$$

The other inequalities are proved similarly.

The point p is also the limit of a sequence of points p_n in X , such that $\text{dist}(p, p_n) \leq \delta_n$. It follows that p_n is a $2\delta_n$ -median point for a_n, b_n, c_n . By Lemma 6.41 we then have that $\text{dist}(p_n, m_n) \leq 6\delta_n$. Since $\delta_n \rightarrow 0$ we get $p = m(a, b, c)$. \square

6.2. Spaces with measured walls

In this section we discuss *measured wall structures* on sets. Every such structure induces a pseudo-metric on the underlying set. The resulting class of metric spaces turns out to coincide with the class of submedian metric spaces (i.e. spaces which can be embedded isometrically in a median space). Examples of such submedian spaces that are not median include real hyperbolic spaces and complex hyperbolic spaces equipped with the square root of the Riemannian distance function.

6.2.1. Definition and basic properties. Following [HP98], a *wall* on a set X is a partition $X = h \sqcup h^c$ (where h is possibly empty or the entire X). A collection \mathcal{D} of subsets of X is called a *collection of half-spaces* if for every $h \in \mathcal{D}$ the complementary subset h^c is also in \mathcal{D} . Given a collection of half-spaces on X , defines a *collection of walls* on X , which is the collection $\mathcal{W}_{\mathcal{D}}$ of pairs $w = \{h, h^c\}$ with $h \in \mathcal{D}$. For a wall $w = \{h, h^c\}$ we call h and h^c the two *half-spaces bounding* w .

We say that a wall $w = \{h, h^c\}$ *separates* two disjoint subsets A, B in X if $A \subset h$ and $B \subset h^c$ or vice-versa. We denote by $\mathcal{W}(A|B)$ the set of walls separating

A and B . In particular, $\mathcal{W}(A|\emptyset)$ is the set of walls $w = \{h, h^c\}$ such that $A \subset h$ or $A \subset h^c$; hence $\mathcal{W}(\emptyset|\emptyset) = \mathcal{W}$.

When $A = \{x_1, \dots, x_n\}$ and $B = \{y_1, \dots, y_m\}$ we write

$$\mathcal{W}(A|B) = \mathcal{W}(x_1, \dots, x_n|y_1, \dots, y_m).$$

In particular, we use the notation $\mathcal{W}(x|y)$ instead of $\mathcal{W}(\{x\}|\{y\})$. We call any set of walls of the form $\mathcal{W}(x|y)$ a *wall-interval*. By convention, $\mathcal{W}(A|A) = \emptyset$ for every non-empty set A .

DEFINITION 6.43 (space with measured walls [CMV04]). A *space with measured walls* is a 4-uple $(X, \mathcal{W}, \mathcal{B}, \mu)$, where \mathcal{W} is a collection of walls on X , \mathcal{B} is a σ -algebra of subsets in \mathcal{W} and μ is a measure on \mathcal{B} , such that for every two points $x, y \in X$ the set of separating walls $\mathcal{W}(x|y)$ is in \mathcal{B} and it has finite measure. We denote by pdist_μ the pseudo-metric on X defined by $\text{pdist}_\mu(x, y) = \mu(\mathcal{W}(x|y))$, and we call it the *wall pseudo-metric*.

LEMMA 6.44. *The collection \mathcal{R} of disjoint unions $\bigsqcup_{i=1}^n \mathcal{W}(F_i|G_i)$, where $n \in \mathbb{N} \cup \{\infty\}$, and F_i, G_i are finite non-empty sets for every $i = 1, 2, \dots, n$, is a ring with respect to the boolean operations on the sets (complementation, intersection and union).*

PROOF. First observe that given finite sets F, F', G, G' :

- $\mathcal{W}(F|G) \cap \mathcal{W}(F'|G') = \mathcal{W}(F \cup F'|G \cup G') \sqcup \mathcal{W}(F \cup G'|G \cup F')$;
- $\mathcal{W}(F|G)^c = \bigsqcup_{S \sqcup T = F \cup G, \{S, T\} \neq \{F, G\}} \mathcal{W}(S|T)$.

From the above it follows that \mathcal{R} is closed with respect to the operation \setminus , and as it is also closed with respect to the intersection, and the union. \square

Theorem 1.11 and Lemma 6.44 imply the following.

PROPOSITION 6.45 (minimal data required for a structure of measured walls). *Let X be a set and let \mathcal{W} be a collection of walls on it. A structure of measured walls can be defined on (X, \mathcal{W}) if and only if on the ring \mathcal{R} composed of disjoint unions $\bigsqcup_{i=1}^n \mathcal{W}(F_i|G_i)$, where $n \in \mathbb{N} \cup \{\infty\}$, and $F_i, G_i, i = 1, 2, \dots, n$, are finite non-empty sets, there exists a premeasure μ such that for every $x, y \in X$, $\mu(\mathcal{W}(x|y))$ is finite.*

Let $(X, \mathcal{W}, \mathcal{B}, \mu)$ and $(X', \mathcal{W}', \mathcal{B}', \mu')$ be two spaces with measured walls, and let $\phi : X \rightarrow X'$ be a map.

DEFINITION 6.46. The map ϕ is a *homomorphism between spaces with measured walls* provided that:

- for any $w' = \{h', h'^c\} \in \mathcal{W}'$ we have $\{\phi^{-1}(h'), \phi^{-1}(h'^c)\} \in \mathcal{W}$ — this latter wall we denote by $\phi^*(w')$;
- the map $\phi^* : \mathcal{W}' \rightarrow \mathcal{W}$ is surjective and for every $B \in \mathcal{B}$, $(\phi^*)^{-1}(B) \in \mathcal{B}'$ and $\mu'((\phi^*)^{-1}(B)) = \mu(B)$.

EXERCISE 6.47. Every homomorphism ϕ as above preserves pseudo-distances.

Consider the set \mathcal{D} of half-spaces determined by \mathcal{W} , and the natural projection map $\mathfrak{p} : \mathcal{D} \rightarrow \mathcal{W}$, $h \mapsto \{h, h^c\}$. The preimages of the sets in \mathcal{B} define a σ -algebra on \mathcal{D} , which we denote by $\mathcal{B}^{\mathcal{D}}$; hence on \mathcal{D} we obtain a pull-back measure that we also denote by μ . This allows us to work either with \mathcal{D} or with \mathcal{W} . Notice that the σ -algebra $\mathcal{B}^{\mathcal{D}}$ does not separate points in \mathcal{D} , as sets in $\mathcal{B}^{\mathcal{D}}$ are unions of fibers of \mathfrak{p} .

DEFINITION 6.48. [CN05], [Nic04] A section \mathfrak{s} for \mathfrak{p} is called *admissible* if its image contains together with a half-space h all the half-spaces h' containing h .

In the sequel we identify an admissible section \mathfrak{s} with its image $\sigma = \mathfrak{s}(\mathcal{W})$; with this identification, an admissible section becomes a collection of half-spaces, σ , such that:

- for every wall $w = \{h, h^c\}$ either h or h^c is in σ , but never both;
- if $h \subset h'$ and $h \in \sigma$ then $h' \in \sigma$.

For any $x \in X$ we denote by \mathfrak{s}_x the section of \mathfrak{p} associating to each wall the half-space bounding it and containing x . Obviously this is an admissible section. We denote by σ_x its image, that is the set of half-spaces $h \in \mathcal{D}$ such that $x \in h$. Observe that σ_x is not necessarily in $\mathcal{B}^{\mathcal{D}}$. Note also that $\mathfrak{p}(\sigma_x \triangle \sigma_y) = \mathcal{W}(x|y)$.

Among standard examples of spaces with measured walls are the real hyperbolic spaces $\mathbb{H}_{\mathbb{R}}^n$. Here the half-spaces are closed or open geometric half-spaces, bounded by hyperbolic hyperplanes, isometric copies of $\mathbb{H}_{\mathbb{R}}^{n-1}$, so that each wall consists of one closed half-space and its (open) complement, as in Section 3 of [CMV04]. Recall that the full group of orientation-preserving isometries of $\mathbb{H}_{\mathbb{R}}^n$ is $SO_0(n, 1)$. The associated set of walls $\mathcal{W}_{\mathbb{H}_{\mathbb{R}}^n}$ is naturally identified with the homogeneous space $SO_0(n, 1)/SO_0(n-1, 1)$. The group $SO_0(n-1, 1)$ is unimodular, therefore there exists an $SO_0(n, 1)$ -invariant borelian measure $\mu_{\mathbb{H}_{\mathbb{R}}^n}$ on the set of walls [Nac65, Chapter 3, Corollary 4]. The set of walls separating two points is relatively compact and has finite measure. Thus

$$(\mathbb{H}_{\mathbb{R}}^n, \mathcal{W}_{\mathbb{H}_{\mathbb{R}}^n}, \mathcal{B}, \mu_{\mathbb{H}_{\mathbb{R}}^n})$$

is a space with measured walls. By Crofton's formula [Rob98, Proposition 2.1] the wall pseudo-metric on $\mathbb{H}_{\mathbb{R}}^n$ is a constant multiple of the usual hyperbolic distance function.

Another example is given by the vertex set $V(T)$ of a simplicial tree T . Every edge e of T defines a partition of $V(T)$ as follows. Let m denote the midpoint of e . Then $T - \{m\}$ consists of two components C^{\pm} and we let $V(T) = e^+ \sqcup e^-$, with e^{\pm} consisting of vertices contained in C^{\pm} . Thus, the set of edges of T defines a collection of walls \mathcal{W} . We equip \mathcal{W} with the counting measure. The reader will verify that this measure defines a structure of measured walls on $V(T)$.

More generally, suppose that T is a real tree which contains a dense subset $M \subset T$ of points such that for every $m \in M$ the complement $T - \{m\}$ consists of exactly two components, m^+, m^- . This again determines a wall structure \mathcal{D} on T defined via the collection of half-spaces $h = m^+, h^c = m^- \cup \{m\}$. For every geodesic arc $\alpha \subset T$ we let \mathcal{W}_{α} denote the set of walls defined by points in $M \cap \alpha$. The metric on α determines a measure on \mathcal{W}_{α} with the total mass of \mathcal{W}_{α} equal to

$$\sup_{p, q \in M \cap \alpha} d(p, q).$$

We define the σ -algebra \mathcal{B} in $2^{\mathcal{D}}$ generated by the subsets \mathcal{W}_{α} , with α 's the geodesic arcs in T . We leave it to the reader to verify that the measures on \mathcal{W}_{α} 's extend to a measure on \mathcal{B} , defining a measured walls structure on T . We refer the reader to [CMV04] for details.

The structure of measured walls behaves well with respect to pull-back.

LEMMA 6.49 (pull-back of a space with measured walls). *Let $(X, \mathcal{W}, \mathcal{B}, \mu)$ be a space with measured walls, let S be a set and $f : S \rightarrow X$ a map. There exists a pull-back structure of space with measured walls $(S, \mathcal{W}_S, \mathcal{B}_S, \mu_S)$ turning f into a homomorphism. Moreover:*

- (i) *if S is endowed with a pseudo-metric pdist and f is an isometry between (S, pdist) and (X, pdist_μ) , then the wall pseudo-metric pdist_{μ_S} coincides with the original pseudo-metric pdist ;*
- (ii) *if a group G acts on S by bijective transformations and on X by automorphisms of the space with measured walls, and if f is G -equivariant, then G acts on $(S, \mathcal{W}_S, \mathcal{B}_S, \mu_S)$ by automorphisms of the space with measured walls.*

PROOF. Define the set of walls \mathcal{W}_S on S as the set of walls $\{f^{-1}(h), f^{-1}(h^c)\}$, where $\{h, h^c\}$ is a wall in X . This yields a surjective map $f^* : \mathcal{W} \rightarrow \mathcal{W}_S$. We then consider the push-forward structure of measured space on \mathcal{W}_S . This defines a structure of a space with measured walls on S such that f is a homomorphism of spaces with measured walls.

(i) It is easily seen that for every $x, y \in S$, $(f^*)^{-1}(\mathcal{W}_S(x|y)) = \mathcal{W}(f(x), f(y))$, hence $\text{pdist}_{\mu_S}(x, y) = \text{pdist}_\mu(f(x), f(y)) = \text{pdist}(x, y)$.

(ii) If f is G -equivariant then the structure of a space with measured walls $(S, \mathcal{W}_S, \mathcal{B}_S, \mu_S)$ is G -invariant. \square

One of the main reasons for the interest in actions of groups on spaces with measured walls is given by the following result.

PROPOSITION 6.50 ([CMV04], [dCTV08]). *Let G be a group acting by automorphisms on a space with measured walls $(X, \mathcal{W}, \mathcal{B}, \mu)$. Let $p > 0$ and let π_p be the representation of G on $L^p(\mathcal{D}, \mu_{\mathcal{D}})$ defined by $\pi_p(g)f = f \circ g^{-1}$.*

Then for every $x \in X$, the map $b : G \rightarrow L^p(\mathcal{D}, \mu_{\mathcal{D}})$ defined by $b(g) = \chi_{\sigma_{gx}} - \chi_{\sigma_x}$ is a 1-cocycle in $Z^1(G, \pi_p)$. In other words, we have an action of G on $L^p(\mathcal{D}, \mu_{\mathcal{D}})$ by affine isometries defined by:

$$g \cdot f = \pi_p(g)f + b(g).$$

PROOF. We check the cocycle property for b . Indeed

$$b(gh) = \chi_{\sigma_{ghx}} - \chi_{\sigma_x} = \chi_{\sigma_{ghx}} - \chi_{\sigma_{gx}} + \chi_{\sigma_{gx}} - \chi_{\sigma_x} = \pi(g)b(h) + b(g).$$

\square

See Remark 2.97 for the definition of metric (and therefore the meaning of “an affine isometry”) on an L^p -space with $p \in (0, 1)$.

6.2.2. Relationship between median spaces and spaces with measured walls.

THEOREM 6.51. (1) *Any space X with measured walls embeds isometrically in a canonically associated median space $\mathcal{M}(X)$. Moreover, any homomorphism between two spaces with measured walls induces an isometry between the associated median spaces.*

- (2) Any median space (X, dist) has a canonical structure of a space with measured walls such that the wall metric coincides with the original metric. Moreover, any isometry between median spaces induces an isomorphism between the structures of measured walls.
- (3) Any median space (X, dist) embeds isometrically in $L^1(\mathcal{W}, \mu)$, for some measured space (\mathcal{W}, μ) .

We will prove this theorem in sections 6.2.3 and 6.2.4.

The fact that each median space embeds into an L^1 -space was known previously, although the embedding was not explicitly constructed, but obtained via a result of Assouad that a metric space embeds isometrically into an L^1 -space if and only if every finite subspace of it embeds (see section 10.9 of this book as well as [AD82], [Ass84], [Ass81], [Ver93]). That all finite median spaces can be embedded into ℓ^1 -spaces seems to be well known in graph theory; all proofs usually refer to finite median graphs only, but can be adapted to work for finite median spaces (see for instance [Mul80]). There exist even algorithms which isometrically embed a given median graph into an ℓ^1 -space; the same method yields algorithms in sub-quadratic time recognizing median graphs [HIK99]. The statement that all finite median spaces can be embedded into ℓ^1 was explicitly stated and proved for the first time in [Ver93, Theorem V.2.3].

It is moreover known that complete median normed spaces are linearly isometric to L^1 -spaces [Ver93, Theorem III.4.13].

We recall that there is no hope of defining a median space containing a space with measured walls and having the universality property with respect to embeddings into median spaces (see Remark 6.22). Nevertheless, the medianization $\mathcal{M}(X)$ of a space with measured walls X appearing in Theorem 6.51, Part (1), is canonically defined and is, in some sense, minimal. This is emphasized for instance by the fact that, under some extra assumptions, a space with measured walls X is at finite Hausdorff distance from $\mathcal{M}(X)$, see [CD17]. In particular, it is the case when X is the n -dimensional real hyperbolic space with the standard structure of space with measured walls.

6.2.3. Embedding a space with measured walls in a median space.

Let $(X, \mathcal{W}, \mathcal{B}, \mu)$ be a space with measured walls, and let x_0 be a base point in X .

Recall from Example 6.19, (7), that $\mathcal{B}_{\sigma_{x_0}}^{\mathcal{D}}$ denotes the collection of subsets $A \subset \mathcal{D}$ such that $A \triangle \sigma_{x_0} \in \mathcal{B}$ and $\mu(A \triangle \sigma_{x_0}) < +\infty$, and that endowed with the pseudo-metric $\text{pdist}_{\mu}(A, B) = \mu(A \triangle B)$ this collection becomes a median pseudo-metric space. The map

$$(6.4) \quad \chi^{x_0} : \mathcal{B}_{\sigma_{x_0}}^{\mathcal{D}} \rightarrow \mathcal{S}^1(\mathcal{D}, \mu), \quad \chi^{x_0}(A) = \chi_{A \triangle \sigma_{x_0}}$$

is an isometric embedding of $\mathcal{B}_{\sigma_{x_0}}^{\mathcal{D}}$ into the median subspace $\mathcal{S}^1(\mathcal{D}, \mu) \subset \mathcal{L}^1(\mathcal{D}, \mu)$, where $\mathcal{S}^1(\mathcal{D}, \mu) = \{\chi_B ; B \text{ measurable and } \mu(B) < +\infty\}$.

The formula $A \triangle \sigma_{x_1} = (A \triangle \sigma_{x_0}) \triangle (\sigma_{x_0} \triangle \sigma_{x_1})$ and the fact that $\sigma_{x_0} \triangle \sigma_{x_1}$ is measurable with finite measure shows that the median pseudo-metric spaces $\mathcal{B}_{\sigma_{x_0}}^{\mathcal{D}}$ and $\mathcal{B}_{\sigma_{x_1}}^{\mathcal{D}}$ are identical: we simply denote this space by $\mathcal{B}_X^{\mathcal{D}}$. In particular, $\sigma_x \in \mathcal{B}_X^{\mathcal{D}}$ for each $x \in X$.

For $x, y \in X$ we have $\text{pdist}_{\mu}(x, y) = \mu(\sigma_x \triangle \sigma_y)$, thus $x \mapsto \sigma_x$ is an isometric embedding of X into $(\mathcal{B}_X^{\mathcal{D}}, \text{pdist}_{\mu})$. Composing with the isometry $\chi^{x_0} : \mathcal{B}_X^{\mathcal{D}} \rightarrow$

$\mathcal{S}^1(\mathcal{D}, \mu)$, we get the following well-known result stating that a wall pseudo-distance is of type 1, in the terminology of [BDCK66, Troisième partie, §2]:

LEMMA 6.52. *Let $(X, \mathcal{W}, \mathcal{B}, \mu)$ be a space with measured walls, and $x_0 \in X$ a basepoint. Then the map $x \mapsto \chi_{\mathcal{W}(x|x_0)}$ is an isometry from X to $L^1(\mathcal{W}, \mu)$. Thus, if the wall pseudo-distance is a distance then (X, dist_μ) is isometric to a subset of $L^1(\mathcal{W}, \mu)$.*

We could probably define the median space associated to a space with measured walls $(X, \mathcal{W}, \mathcal{B}, \mu)$ to be the (closure of the) median hull of the isometric image of X inside $L^1(\mathcal{W}, \mu)$. We give here an alternative construction which is more intrinsic.

NOTATION 6.53. We denote by $\overline{\mathcal{M}}(X)$ the set of admissible sections, and by $\mathcal{M}(X)$ the intersection $\overline{\mathcal{M}}(X) \cap \mathcal{B}_X^{\mathcal{D}}$. Every section σ_x belongs to $\mathcal{M}(X)$, thus X isometrically embeds in $\mathcal{M}(X)$. We denote by $\iota : X \rightarrow \mathcal{M}(X)$ this isometric embedding.

PROPOSITION 6.54. *Let $(X, \mathcal{W}, \mathcal{B}, \mu)$ be a space with measured walls. Then:*

- (i) *The space $\mathcal{M}(X)$ defined as above is a median subspace of $\mathcal{B}_X^{\mathcal{D}}$.*
- (ii) *Any homomorphism $\phi : X \rightarrow X'$ between X and another space with measured walls $(X', \mathcal{W}', \mathcal{B}', \mu')$ induces an isometry $\mathcal{M}(X) \rightarrow \mathcal{M}(X')$.*
- (iii) *In particular, the group of automorphisms of $(X, \mathcal{W}, \mathcal{B}, \mu)$ acts by isometries on $\mathcal{M}(X)$.*

PROOF. (i) Given an arbitrary triple $(\sigma_1, \sigma_2, \sigma_3) \in \mathcal{M}(X)^3$, let us denote by $m(\sigma_1, \sigma_2, \sigma_3)$ the set of half-spaces h such that there exist at least two distinct indices $i, j \in \{1, 2, 3\}$ with $h \in \sigma_i, h \in \sigma_j$. In other words

$$m(\sigma_1, \sigma_2, \sigma_3) = (\sigma_1 \cap \sigma_2) \cup (\sigma_1 \cap \sigma_3) \cup (\sigma_2 \cap \sigma_3)$$

(see also Example 6.11).

Clearly $m = m(\sigma_1, \sigma_2, \sigma_3)$ belongs to $\overline{\mathcal{M}}(X)$. Fix a point x_0 in X and take $\chi_0 = \chi^{x_0}$, the function defined in (6.4). We want to show that

$$\chi_0(m) = m(\chi_0(\sigma_1), \chi_0(\sigma_2), \chi_0(\sigma_3)).$$

This will prove that $m \in \mathcal{B}_X^{\mathcal{D}}$ and that m is a median point of $\sigma_1, \sigma_2, \sigma_3$.

For our set-theoretical calculation it is convenient to treat characteristic functions as maps from \mathcal{D} to $\mathbb{Z}/2\mathbb{Z}$. We may then use the addition (mod. 2) and pointwise multiplication on these functions. We get

$$\chi_{A \cap B} = \chi_A \chi_B, \chi_{A \Delta B} = \chi_A + \chi_B, \chi_{A \cup B} = \chi_A + \chi_B + \chi_A \chi_B.$$

It follows easily that for any three subsets A, B, C we have

$$\chi_{(A \cap B) \cup (A \cap C) \cup (B \cap C)} = \chi_A \chi_B + \chi_A \chi_C + \chi_B \chi_C.$$

Thus

$$\chi_{[(A \cap B) \cup (A \cap C) \cup (B \cap C)]_{\Delta D}} = \chi_A \chi_B + \chi_A \chi_C + \chi_B \chi_C + \chi_D.$$

On the other hand

$$\begin{aligned} & \chi_{((A \Delta D) \cap (B \Delta D)) \cup ((A \Delta D) \cap (C \Delta D)) \cup ((B \Delta D) \cap (C \Delta D))} = \\ & (\chi_A + \chi_D)(\chi_B + \chi_D) + (\chi_A + \chi_D)(\chi_C + \chi_D) + (\chi_B + \chi_D)(\chi_C + \chi_D) = \\ & \chi_A \chi_B + \chi_A \chi_C + \chi_B \chi_C + 2\chi_A \chi_D + 2\chi_B \chi_D + 2\chi_C \chi_D + 3\chi_D = \\ & \chi_A \chi_B + \chi_A \chi_C + \chi_B \chi_C + \chi_D. \end{aligned}$$

We have thus checked that $[(A \cap B) \cup (A \cap C) \cup (B \cap C)] \Delta D$ coincides with

$$[(A \Delta D) \cap (B \Delta D)] \cup [(A \Delta D) \cap (C \Delta D)] \cup [(B \Delta D) \cap (C \Delta D)].$$

Applying this to $A = \sigma_1, B = \sigma_2, C = \sigma_3, D = \sigma_{x_0}$ yields the desired result.

(ii) Consider a homomorphism of spaces with measured walls $\phi : X \rightarrow X'$. It is easily seen that the surjective map $\phi^* : \mathcal{W}' \rightarrow \mathcal{W}$ induces a surjective map $\phi^* : \mathcal{D}' \rightarrow \mathcal{D}$ such that for every $B \in \mathcal{B}^{\mathcal{D}}, (\phi^*)^{-1}(B) \in \mathcal{B}^{\mathcal{D}'}$ and $\mu'((\phi^*)^{-1}(B)) = \mu(B)$.

Let σ denote any admissible section. Set

$$\phi_*(\sigma) = (\phi^*)^{-1}(\sigma) = \{h' \in \mathcal{D}' ; \phi^{-1}(h') \in \sigma\}.$$

Since ϕ is a homomorphism, $\phi_*(\sigma)$ is an admissible section of $(X', \mathcal{W}', \mathcal{B}', \mu')$. Note that $\phi_*(\sigma_x) = \sigma_{\phi(x)}$ and that $\phi_*(\sigma \Delta \sigma') = \phi_*(\sigma) \Delta \phi_*(\sigma')$. This implies that ϕ_* defines a map from $\mathcal{M}(X)$ to $\mathcal{M}(X')$. Moreover

$$\begin{aligned} \text{pdist}_{\mathcal{M}(X')}(\phi_*(\sigma), \phi_*(\sigma')) &= \mu'(\phi_*(\sigma) \Delta \phi_*(\sigma')) = \mu'(\phi_*(\sigma \Delta \sigma')) = \\ &= \mu'((\phi^*)^{-1}(\sigma \Delta \sigma')) = \mu(\sigma \Delta \sigma') = \text{pdist}_{\mathcal{M}(X)}(\sigma, \sigma'). \end{aligned}$$

Thus ϕ_* is an isometry.

The statement (iii) is an immediate consequence of (ii). \square

The results in Proposition 6.54 justify the following terminology.

DEFINITION 6.55. We call $\mathcal{M}(X)$ *the median space associated to $(X, \mathcal{W}, \mathcal{B}, \mu)$* .

REMARK 6.56. The space $\mathcal{M}(X)$ can be replaced by $\mathcal{M}_0(X)$, the metric completion of the median closure of X in $\mathcal{M}(X)$. The two spaces are, in general, different, but they become equal when the space $\mathcal{M}_0(X)$ is locally convex [Fio17].

The space $\mathcal{M}_0(X)$ has the advantage of being a complete geodesic metric median space when X is connected. This follows from the fact that $\mathcal{M}_0(X)$ is connected, and a result of Bowditch [Bow16], stating that a complete median space that is connected is geodesic. The connectedness of $\mathcal{M}_0(X)$ is due to the fact that the median map is 1-Lipschitz, and as the median completion of X in $\mathcal{M}(X)$ equals the increasing union of (connected) sets obtained by iterative applications of the median map to X , it is itself connected. The space $\mathcal{M}_0(X)$ is the metric completion of this latter median completion, hence it is itself connected.

The first part of Theorem 6.51 is now proved. The second part will be a corollary of Theorem 6.57 proven in the next section. In order to prove it we will need some preliminary results on the geometry of median spaces allowing to define measured walls in a consistent manner.

6.2.4. Median spaces have measured walls. The aim of this section is to prove the following.

THEOREM 6.57. *Let (X, dist) be a median space. Let \mathcal{W} be the set of convex walls, and let \mathcal{B} be the σ -algebra generated by the following subset of $\mathcal{P}(\mathcal{W})$:*

$$\mathcal{U} = \{\mathcal{W}(x|y) ; x, y \text{ points of } X\}.$$

Then there exists a measure μ on \mathcal{B} such that:

- (1) $\mu(\mathcal{W}(x|y)) = \text{dist}(x, y)$; consequently, the quadruple $(X, \mathcal{W}, \mathcal{B}, \mu)$ is a space with measured walls;
- (2) any isometry of (X, dist) is an automorphism of the space with measured walls $(X, \mathcal{W}, \mathcal{B}, \mu)$.

REMARK 6.58. In general, a measure μ on the σ -algebra \mathcal{B} is not uniquely defined by the condition (1) in Theorem 6.57. It follows from the Caratheodory's Theorem (theorem 1.11), the measure is uniquely determined if there exists, say, a sequence of points (x_n) in X such that $\mathcal{W} = \bigcup_{n,m} \mathcal{W}(x_n|x_m)$. This happens for instance if there exists a countable subset in X whose convex hull is the entire X . Uniqueness is also guaranteed when for some topology on \mathcal{W} the measure μ is borelian and \mathcal{W} is locally compact second countable.

Combining Theorem 6.57 above and Lemma 6.52 we get the following:

COROLLARY 6.59. *Let (X, dist) be a median space. Then X isometrically embeds in $L^1(\mathcal{W}, \mu)$, where (\mathcal{W}, μ) is as in Theorem 6.57. More precisely, given any $x_0 \in X$, the space X is isometric to*

$$\{\chi_{\mathcal{W}(x|x_0)}; x \in X\}$$

regarded as a subset of $L^1(\mathcal{W}, \mu)$ endowed with the induced metric.

COROLLARY 6.60. *A metric space (X, dist) is submedian in the sense of Definition 6.5 if and only if it admits a structure of a space with measured walls $(X, \mathcal{W}, \mathcal{B}, \mu)$ such that $\text{dist} = \text{dist}_\mu$. Moreover, all walls in \mathcal{W} may be assumed to be convex.*

PROOF. The direct part follows from Theorem 6.57 and Lemma 6.49. The converse part follows from Lemma 6.52. \square

REMARK 6.61. Corollary 6.60 for finite metric spaces was already known. More precisely, according to [Ass80] and [AD82] a finite metric space (X, dist) is isometrically ℓ^1 -embeddable if and only if

$$\text{dist} = \sum_{S \subseteq X} \lambda_S \delta_S,$$

where λ_S are non-negative real numbers, and $\delta_S(x, y) = 1$ if $x \neq y$ and $S \cap \{x, y\}$ has cardinality one, $\delta_S(x, y) = 0$ otherwise.

The strategy of the proof of Theorem 6.57 is to use Proposition 6.45. We first show that for any pair of finite non-empty sets F, G in X , $\mathcal{W}(F|G)$ is equal to $\mathcal{W}(a|b)$ for some pair of points a, b . In order to do this we need the following intermediate results.

LEMMA 6.62. *Let (x, y, z) be a geodesic sequence in (X, dist) . Then we have the following decomposition as a disjoint union:*

$$\mathcal{W}(x|z) = \mathcal{W}(x|y) \sqcup \mathcal{W}(y|z).$$

PROOF. First notice that by convexity of half-spaces, the intersection $\mathcal{W}(x|y) \cap \mathcal{W}(y|z)$ is empty. Then the inclusion $\mathcal{W}(x|z) \subseteq \mathcal{W}(x|y) \cup \mathcal{W}(y|z)$ is clear because if a half-space h contains x but does not contain z , then either h contains y (in which case the wall $\{h, h^c\}$ separates y from z) or h^c contains y (in which case the wall $\{h, h^c\}$ separates x from y). The inclusion $\mathcal{W}(x|y) \cup \mathcal{W}(y|z) \subseteq \mathcal{W}(x|z)$ holds because if h contains x and $y \notin h$, again, by convexity, we cannot have $z \in h$ and hence $\{h, h^c\}$ separates x from z . \square

As an immediate consequence we get the following:

COROLLARY 6.63. For any geodesic sequence (x_1, x_2, \dots, x_n) we have the following decomposition:

$$\mathcal{W}(x_1|x_n) = \mathcal{W}(x_1|x_2) \sqcup \dots \sqcup \mathcal{W}(x_{n-1}|x_n).$$

COROLLARY 6.64. If (x, y) and (x', y') are parallel pairs then

$$\mathcal{W}(x|y) = \mathcal{W}(x'|y') = \mathcal{W}(x, x'|y, y').$$

and

$$\mathcal{W}(x|y') = \mathcal{W}(x'|y) = \mathcal{W}(x|y) \sqcup \mathcal{W}(x|x').$$

LEMMA 6.65. Given three points x, y, z with the median point m , we have

$$\mathcal{W}(x|y, z) = \mathcal{W}(x|m).$$

PROOF. According to Lemma 6.62 we have that $\mathcal{W}(x|y) = \mathcal{W}(x|m) \sqcup \mathcal{W}(m|y)$ and that $\mathcal{W}(x|z) = \mathcal{W}(x|m) \sqcup \mathcal{W}(m|z)$. It follows that

$$\mathcal{W}(x|y, z) = \mathcal{W}(x|y) \cap \mathcal{W}(x|z) = \mathcal{W}(x|m) \sqcup (\mathcal{W}(m|y) \cap \mathcal{W}(m|z)).$$

But by convexity of the walls $\mathcal{W}(m|y) \cap \mathcal{W}(m|z) = \emptyset$, and we are done. \square

We will use intensively the following two operations:

DEFINITION 6.66 (projection and straightening). Let $(x, y), (a, b)$ be two pairs of points of a median space X .

The *projection of (x, y) with the target (a, b)* is the pair (x', y') defined by $x' = m(x, a, b), y' = m(y, a, b)$.

If furthermore $x, y \in I(a, b)$, we also consider the *straightening of the path (a, x, y, b)* , which by definition is the path (a, p, q, b) , where the pair (p, q) is defined by $p = m(a, x, y), q = m(b, x, y)$.

Observe that given two pairs of points $(x, y), (a, b)$, the central rectangle $[x', a', y', b']$ associated with $[x, a, y, b]$ (as in Definition 6.34) is obtained by first projecting (x, y) with the target (a, b) — this yields the pair (x', y') , and then straightening (a, x', y', b) — which yields the pair (a', b') . We now describe some properties of both procedures.

LEMMA 6.67. Let $(x, y), (a, b)$ be two pairs of points.

(1) Let (x', y') be the projection of (x, y) with target (a, b) . Then

$$\mathcal{W}(x'|y') = \mathcal{W}(x|y) \cap \mathcal{W}(a|b).$$

(2) Assume $x, y \in I(a, b)$, and let (p, q) be the projection of (a, b) with the target (x, y) . Then $[p, x, q, y]$ is a rectangle, $\mathcal{W}(p|q) = \mathcal{W}(x|y)$, and (a, p, q, b) is a geodesic sequence (thus (a, x, y, b) really has been straightened to a geodesic).

(3) Let $[x', a', y', b']$ be the central rectangle associated with $[x, a, y, b]$. Then

$$\mathcal{W}(x'|y') = \mathcal{W}(x|y) \cap \mathcal{W}(a|b), \mathcal{W}(x'|y') = \mathcal{W}(a'|b').$$

PROOF. Since the central rectangle is in fact obtained by composing the projecting and straightening operations, it is enough to prove the part 3 of the lemma.

The equality $\mathcal{W}(x'|y') = \mathcal{W}(a'|b')$ follows from Corollary 6.64.

By Lemma 6.65 we have $\mathcal{W}(x|x') = \mathcal{W}(x|a, b)$. In particular, $\mathcal{W}(x|x') \cap \mathcal{W}(a|b) = \emptyset$. Similarly,

$$\mathcal{W}(y|y') \cap \mathcal{W}(a|b) = \emptyset.$$

Consider now a half-space h such that $x \in h, y \notin h$ and $\{h, h^c\} \in \mathcal{W}(a|b)$. Since

$$\mathcal{W}(x|x') \cap \mathcal{W}(a|b) = \emptyset,$$

we deduce that $x' \in h$. Similarly we have $y' \in h^c$. We have thus proved that $\mathcal{W}(x|y) \cap \mathcal{W}(a|b) \subset \mathcal{W}(x'|y')$.

On the other hand, since $\mathcal{W}(x'|y') = \mathcal{W}(a'|b')$ and (a, a', b', b) is a geodesic, it follows that $\mathcal{W}(x'|y') \subset \mathcal{W}(a|b)$.

According to Lemma 6.36, (x', y') is parallel to a pair (x'', y'') such that the sequence (x, x'', y'', y) is geodesic. This and Corollary 6.64 imply that $\mathcal{W}(x'|y') \subset \mathcal{W}(x|y)$. \square

PROPOSITION 6.68. *Let F and G be two finite non-empty subsets in X . There exist two points $p, q \in X$ such that*

$$\mathcal{W}(F|G) = \mathcal{W}(p|q).$$

PROOF. We use an inductive argument over $n = \text{card } F + \text{card } G$. For $n = 2$ the result is obvious, while for $n = 3$ it is Lemma 6.65.

Assume that the statement holds for n and let F, G be such that $\text{card } F + \text{card } G = n + 1 \geq 3$. Without loss of generality we may assume that $\text{card } F \geq 2$. Then $F = F_1 \sqcup \{x\}$, and $\mathcal{W}(F|G) = \mathcal{W}(F_1|G) \cap \mathcal{W}(x|G)$. The inductive hypothesis implies that $\mathcal{W}(F_1|G) = \mathcal{W}(a|b)$ and $\mathcal{W}(x|G) = \mathcal{W}(c|d)$, for some points a, b, c, d . Hence $\mathcal{W}(F|G) = \mathcal{W}(a|b) \cap \mathcal{W}(c|d)$. We conclude by applying Lemma 6.67. \square

At this stage we have proven that the ring \mathcal{R} defined in Proposition 6.45 coincides with the set of disjoint unions $\bigsqcup_{i=1}^n \mathcal{W}(x_i|y_i)$. It remains to show that there is a premeasure $\mu : \mathcal{R} \rightarrow \mathbb{R}^+$ on the ring \mathcal{R} such that $\mu(\mathcal{W}(x|y)) = \text{dist}(x, y)$. We first define μ as an additive function.

LEMMA 6.69. *If $\mathcal{W}(x|y) = \mathcal{W}(a|b)$ then $\text{dist}(x, y) = \text{dist}(a, b)$.*

PROOF. First let (x', y') be the projection of (x, y) with target (a, b) . Then by Lemma 6.67(1) we have $\mathcal{W}(x'|y') = \mathcal{W}(x|y) \cap \mathcal{W}(a|b) = \mathcal{W}(a|b)$. By Corollary 6.28 the median map is 1-Lipschitz, thus $d(x', y') \leq d(x, y)$.

We now straighten (a, x', y', b) to (a, p, q, b) (thus (p, q) is the projection of (a, b) with target (x', y')). Then by Lemma 6.67(2) we have $\mathcal{W}(p|q) = \mathcal{W}(x'|y') = \mathcal{W}(a|b)$, and (a, p, q, b) is a geodesic sequence. By Corollary 6.63 we deduce $\mathcal{W}(a|p) = \mathcal{W}(q|b) = \emptyset$, and thus $a = p, q = b$. It follows that $d(a, b) = d(p, q)$, and thus by Corollary 6.64 we have $d(a, b) = d(x', y') \leq d(x, y)$. We conclude by symmetry. \square

PROPOSITION 6.70. *Assume that for two points a, b the set of walls $\mathcal{W}(a|b)$ decomposes as $\mathcal{W}(a|b) = \bigsqcup_{j=1}^n \mathcal{W}(x_j|y_j)$. Then there exists a geodesic sequence $(a_1 = a, a_2, \dots, a_{2^n} = b)$ and a partition $\{1, 2, \dots, 2^n - 1\} = I_1 \sqcup I_2 \sqcup \dots \sqcup I_n$ such that:*

- (1) *for each $j \in \{1, \dots, n\}$ the set I_j has 2^{j-1} elements and we have a decomposition of $\mathcal{W}(x_j|y_j) = \bigsqcup_{i \in I_j} \mathcal{W}(a_i|a_{i+1})$*
- (2) *for each $j \in \{1, \dots, n\}$ we have $\text{dist}(x_j, y_j) = \sum_{i \in I_j} \text{dist}(a_i, a_{i+1})$*

In particular, $\text{dist}(a, b) = \sum_j \text{dist}(x_j, y_j)$.

We easily deduce the following:

COROLLARY 6.71. *There is a unique additive function $\mu : \mathcal{R} \rightarrow \mathbb{R}^+$ such that $\mu(\mathcal{W}(x|y)) = \text{dist}(x, y)$.*

To prove the Proposition we need the following auxiliary result:

LEMMA 6.72. *In a median space (X, dist) , consider two geodesic sequences with common endpoints (a, p, q, b) and (a, p', q', b) , such that $\mathcal{W}(p|q) \cap \mathcal{W}(p'|q') = \emptyset$. Let (s, t) be the projection of (p', q') with target (a, p) . Similarly, let (u, v) be the projection of (p', q') with target (q, b) . Then $\text{dist}(p', q') = \text{dist}(s, t) + \text{dist}(u, v)$.*

PROOF. Consider two more points: $m = m(t, p', q')$, $n = m(u, p', q')$ (see Figure 6.3). Let us check that $[s, t, m, p']$ is a rectangle. By the construction (t, m, p') is a geodesic sequence. Since s, t are projection of p', q' onto the interval $I(a, p)$ we deduce that $(q', m, t, s), (p', s, t)$ are geodesic sequences. Since (x, p', q', y) is a geodesic sequence we see that (x, s, p', m, q', y) is geodesic.

We thus have $\text{dist}(p', m) = \text{dist}(s, t)$, and also $\mathcal{W}(p'|m) = \mathcal{W}(s|t)$ (by Corollary 6.64). Hence $\mathcal{W}(p'|m) = \mathcal{W}(a|p) \cap \mathcal{W}(p'|q')$ (by Lemma 6.67(1)). Similarly we get $\text{dist}(n, q') = \text{dist}(u, v)$, and $\mathcal{W}(n|q') = \mathcal{W}(q|b) \cap \mathcal{W}(p'|q')$.

We claim that $\mathcal{W}(m|q') = \mathcal{W}(q|b) \cap \mathcal{W}(p'|q')$. Indeed applying Lemma 6.62 several times we get

$$\mathcal{W}(p'|m) \sqcup \mathcal{W}(m|q') = \mathcal{W}(p'|q') \subset \mathcal{W}(a|b) = \mathcal{W}(a|p) \sqcup \mathcal{W}(p|q) \sqcup \mathcal{W}(q|b)$$

and the claim follows, since, by the assumption, $\mathcal{W}(p|q) \cap \mathcal{W}(p'|q') = \emptyset$ and we already have $\mathcal{W}(p'|m) = \mathcal{W}(a|p) \cap \mathcal{W}(p'|q')$.

We conclude that $\mathcal{W}(m|q') = \mathcal{W}(n|q')$. This implies the $\text{dist}(m, q') = \text{dist}(n, q') = \text{dist}(u, v)$, cf. Lemma 6.69. Since (p', m, q') is a geodesic sequence we get $\text{dist}(p', q') = \text{dist}(p', m) + \text{dist}(m, q') = \text{dist}(s, t) + \text{dist}(u, v)$. \square




FIGURE 6.3. The construction in Lemma 6.72.

PROOF OF PROPOSITION 6.70. We argue by induction on n . The case $n = 1$ follows from Lemma 6.69.

Now let us assume that $n > 1$ and that the lemma is holds for partitions of any wall-interval into $n - 1$ wall-subintervals. Notice first that, according to

Lemma 6.67, Part (1), and Lemma 6.69, after replacing (x_i, y_i) by its projection with target (a, b) , we can assume that the x_i 's and y_i 's belong to the interval $I(a, b)$.

We straighten (a, x_1, y_1, b) to (a, p_1, q_1, b) . Then by Lemma 6.67, Part (2), the sequence (a, p_1, q_1, b) is geodesic, and we have $\mathcal{W}(x_1|y_1) = \mathcal{W}(p_1|q_1)$.

By Lemma 6.62 we have $\mathcal{W}(a|b) = \mathcal{W}(a|p_1) \sqcup \mathcal{W}(p_1|q_1) \sqcup \mathcal{W}(q_1|b)$. It follows that $\mathcal{W}(a|p_1) \sqcup \mathcal{W}(q_1|b) = \sqcup_{i=2}^n \mathcal{W}(x_i|y_i)$.

We now straighten each path (a, x_i, y_i, b) to (a, p_i, q_i, b) (when $i > 1$). Again we have $\mathcal{W}(x_i|y_i) = \mathcal{W}(p_i|q_i)$ and moreover $\text{dist}(x_i, y_i) = \text{dist}(p_i, q_i)$ (since $[x_i, p_i, y_i, q_i]$ is a rectangle). Now let us project the points p_i and q_i onto $I(x, p_1)$ and $I(q_1, y)$. We set $s_i := m(p_i, x, p_1)$, $t_i := m(q_i, x, p_1)$, $u_i := m(p_i, q_1, y)$ and $v_i := m(q_i, q_1, y)$.

Applying again Lemma 6.67, Part (1), we see that $\mathcal{W}(p_i|q_i) \cap \mathcal{W}(a|p_1) = \mathcal{W}(s_i|t_i)$ and $\mathcal{W}(p_i|q_i) \cap \mathcal{W}(q_1|b) = \mathcal{W}(u_i|v_i)$. Thus $\mathcal{W}(p_i|q_i) = \mathcal{W}(s_i|t_i) \sqcup \mathcal{W}(u_i|v_i)$, and we get two decompositions:

$$\mathcal{W}(a|p_1) = \sqcup_{i=2}^n \mathcal{W}(s_i|t_i)$$

and

$$\mathcal{W}(q_1|b) = \sqcup_{i=2}^n \mathcal{W}(u_i|v_i).$$

We conclude the proof by applying the induction hypothesis to the two decompositions above, since Lemma 6.72 ensures that $\text{dist}(p_i, q_i) = \text{dist}(s_i, t_i) + \text{dist}(u_i, v_i)$. \square

The following proposition shows that the premeasure μ satisfies the property (M_1'') .

PROPOSITION 6.73. *Let (X, dist) be a median space, endowed with convex walls. If $(I_n)_{n \in \mathbb{N}}$ is a non-increasing sequence of finite disjoint unions of wall-intervals such that $\bigcap_n I_n = \emptyset$, then $I_k = \emptyset$ for k large enough.*

PROOF. In what follows we identify a half-space with its characteristic function. First note that the set of half-spaces bounding a convex wall (i.e. the set of convex subsets whose complement is convex as well) is a closed subset of $\{0, 1\}^X$. Then the set $\mathcal{D}(x|y)$ of half-spaces containing x but not y is a closed subset of the compact subset of $\{0, 1\}^X$ consisting in functions $f : X \rightarrow \{0, 1\}$ such that $f(x) = 1, f(y) = 0$. Therefore, $\mathcal{D}(x|y)$ is compact.

It is enough to argue when $I_0 = \mathcal{W}(x|y)$. Since $(I_n)_{n \in \mathbb{N}}$ is non-increasing for each n we have $I_n \subset \mathcal{W}(x|y)$. We then define H_n as the set of half-spaces h such that $\{h, h^c\} \in I_n$, and $x \in h$. It follows that $(H_n)_{n \in \mathbb{N}}$ is non-increasing, and has empty intersection. By projecting onto $I(x, y)$ we have

$$I_n = \prod_i \mathcal{W}(x_i|y_i)$$

for some points $x_i, y_i \in I(x, y)$ (Lemma 6.67(1)). We know that $\mathcal{W}(x_i|y_i) = \mathcal{W}(p_i|q_i)$ for $p_i = m(x, x_i, y_i)$, $q_i = m(y, x_i, y_i)$, and furthermore (x, p_i, q_i, y) is a geodesic sequence. Thus

$$H_n = \prod_i \mathcal{W}(p_i|q_i)$$

and H_n is compact. It follows that there exists k such that $H_k = \emptyset$, which implies that $I_k = \emptyset$. \square

We now have all the ingredients to finish the proof of Theorem 6.57.

PROOF OF THEOREM 6.57. That the premeasure μ is well-defined on \mathcal{R} is the content of Proposition 6.70. It obviously satisfies properties (M_0) and (M'_1) , while (M''_1) is proved in Proposition 6.73.

By the Carathéodory's Theorem, Theorem 1.11, μ^* restricted to \mathcal{A}^* is a measure extending μ , hence its restriction to \mathcal{B} is also a measure extending μ .

Obviously any isometry of (X, dist) defines a bijective transformation on \mathcal{W} preserving \mathcal{R} and the premeasure μ , hence the outer measure μ^* and \mathcal{A}^* , hence it defines an automorphism of the measured space $(\mathcal{W}, \mathcal{B}, \mu)$. \square

Finitely generated and finitely presented groups

7.1. Finitely generated groups

A group which has a finite generating set is called *finitely generated*.

DEFINITION 7.1. The *rank* of a finitely generated group G , denoted $\text{rank}(G)$, is the minimal number of generators of G .

REMARK 7.2. In French, the terminology for finitely generated groups is *groupe de type fini*. On the other hand, in English, being a *group of finite type* is a much stronger requirement than finite generation (typically, this means that the group has type \mathbf{F}_∞).

EXERCISE 7.3. Show that every finitely generated group is countable.

EXAMPLES 7.4. (1) The group $(\mathbb{Z}, +)$ is finitely generated by both $\{1\}$ and $\{-1\}$. Also, any set $\{p, q\}$ of coprime integers generates \mathbb{Z} .
 (2) The group $(\mathbb{Q}, +)$ is not finitely generated.

EXERCISE 7.5. Prove that the transposition (12) and the cycle $(12\dots n)$ generate the permutation group S_n .

REMARKS 7.6. (1) Every quotient \bar{G} of a finitely generated group G is finitely generated; we can take as generators of \bar{G} the images of the generators of G .

(2) If N is a normal subgroup of G , and both N and G/N are finitely generated, then G is finitely generated. Indeed, take a finite generating set $\{n_1, \dots, n_k\}$ for N , and a finite generating set $\{g_1N, \dots, g_mN\}$ for G/N . Then

$$\{g_i, n_j : 1 \leq i \leq m, 1 \leq j \leq k\}$$

is a finite generating set for G .

We will see in examples below that if N is a normal subgroup in a group G and G is finitely generated, it *does not* necessarily follow that N is finitely generated (not even if G is a semidirect product of N and G/N).

EXAMPLE 7.7. Let G be the wreath product $\mathbb{Z} \wr \mathbb{Z} \cong N \rtimes \mathbb{Z}$, where N is the (countably) infinite direct sum of copies of \mathbb{Z} . Then G is 2-generated (see Lemma 7.11). On the other hand, the subgroup N is not finitely generated.

EXAMPLE 7.8. Let H be the group of permutations of \mathbb{Z} generated by the transposition $t = (01)$ and the translation map $s(i) = i + 1$. Let H_i be the group of permutations of \mathbb{Z} supported on $[-i, i] = \{-i, -i + 1, \dots, 0, 1, \dots, i - 1, i\}$, and let

H_ω be the group of finitely supported permutations of \mathbb{Z} (i.e. the group of bijections $f : \mathbb{Z} \rightarrow \mathbb{Z}$ such that f is the identity outside a finite subset of \mathbb{Z}),

$$H_\omega = \bigcup_{i=0}^{\infty} H_i.$$

Then H_ω is a normal subgroup in H and $H/H_\omega \simeq \mathbb{Z}$, while H_ω is not finitely generated.

Indeed from the relation $s^k t s^{-k} = (k k + 1)$, $k \in \mathbb{Z}$, it immediately follows that H_ω is a subgroup in H . It is, likewise, easy to see that $s^k H_i s^{-k} \subset H_{i+k}$, whence $s^k H_\omega s^{-k} \subset H_\omega$ for every $k \in \mathbb{Z}$.

If g_1, \dots, g_k is a finite set generating H_ω , then there exists an $i \in \mathbb{N}$ so that all g_j 's are in H_i , hence $H_\omega = H_i$. On the other hand, clearly, H_i is a proper subgroup of H_ω .

EXERCISE 7.9. 1. Let F be a non-abelian free group (see Definition 7.19). Let $\varphi : F \rightarrow \mathbb{Z}$ be any non-trivial homomorphism. Prove that the kernel of φ is not finitely generated.

2. Let F be a free group of finite rank with free generators x_1, \dots, x_n ; set $G := F \times F$. Then G has the generating set

$$\{(x_i, 1), (1, x_j) : 1 \leq i, j \leq n\}.$$

Define homomorphism $\phi : G \rightarrow \mathbb{Z}$ sending every generator of G to 1 in \mathbb{Z} . Show that the kernel K of ϕ is finitely generated. Hint: Use the elements (x_i, x_j^{-1}) , $(x_i x_j^{-1}, 1)$, $(1, x_i x_j^{-1})$, $1 \leq i, j \leq n$, of the subgroup K .

We will see later that a *finite index* subgroup of a finitely generated group is always finitely generated (Lemma 7.85 or Theorem 8.37). The next lemma shows that extensions of finitely generated groups are again finitely generated:

LEMMA 7.10. *Suppose that we have a short exact sequence of groups*

$$1 \rightarrow G_1 \xrightarrow{i} G_2 \xrightarrow{\pi} G_3 \rightarrow 1$$

such that the groups G_1, G_3 are finitely generated. Then G_2 is also finitely generated.

PROOF. Let S_1, S_3 be finite generating sets of G_1, G_3 . For each $\bar{s} \in S_3$ pick $s \in \pi^{-1} \bar{s}$. We claim that

$$S_2 := i(S_1) \cup \{s | \bar{s} \in S_3\}$$

is a generating set of G_2 . Indeed, each $g \in G_2$ projects to $\pi(g)$, which is a product

$$\bar{s}_1^{\pm 1} \dots \bar{s}_k^{\pm 1}, \quad s_i \in S_3.$$

Therefore, by normality of $i(G_1)$ in G_2 , the element g itself has the form

$$h \cdot s_1^{\pm 1} \dots s_k^{\pm 1}, \quad h \in i(G_1).$$

Since h is a product of the elements $s \in S_1$ (and their inverses), the claim follows. \square

A similar proof applies to *wreath products*. Recall that the *wreath product* $A \wr C$ of groups A and C is the semidirect product

$$(\oplus_C A) \rtimes C$$

where C acts on the direct sum by precompositions: $f(x) \mapsto f(xc^{-1})$. Thus, elements of wreath products $A \wr C$ are pairs (f, c) , where $f : C \rightarrow A$ is a function with finite support and $c \in C$. The product structure on this set is given by the formula

$$(f_1(x), c_1) \cdot (f_2(x), c_2) = (f_1(xc_2^{-1})f_2(x), c_1c_2).$$

Here and below we use multiplicative notation when dealing with wreath products. For each $a \in A$ we define the function $\delta_a : C \rightarrow A$ is the function which sends $1 \in C$ to $a \in A$ and sends all other elements of C to $1 \in A$.

LEMMA 7.11. *If $a_i, i \in I, c_j, j \in J$ are generators of A and C , respectively, the elements $(1, c_j), j \in J$ and $(\delta_{a_i}, 1), i \in I$, generate $G_A := A \wr C$. In particular, if A and C are finitely generated, so is $A \wr C$.*

PROOF. It is enough to show that each $(f, 1) \in G_A$ is a product of the elements $(1, c_j), (\delta_{a_i}, 1)$. Since the maps δ_a generate

$$\bigoplus_{C} A,$$

it suffices to prove this claim for each δ_a . If

$$a = a_{i_1} \dots a_{i_k},$$

then, clearly,

$$(\delta_a, 1) = (\delta_{a_{i_1}}, 1) \dots (\delta_{a_{i_k}}, 1).$$

The conclusion of the Lemma follows. \square

Below we describe a finite generating set for the group $GL(n, \mathbb{Z})$. In the proof we use the *elementary matrices* $N_{i,j} = I_n + E_{i,j}$ ($i \neq j$); here I_n is the identity $n \times n$ matrix and the matrix $E_{i,j}$ has a unique non-zero entry 1 in the intersection of the i -th row and the j -th column.

PROPOSITION 7.12. *The group $GL(n, \mathbb{Z})$ is generated by*

$$s_1 = \begin{pmatrix} 0 & \dots & \dots & 0 & 1 \\ 1 & \ddots & & \vdots & 0 \\ 0 & \ddots & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \ddots & 0 \\ 0 & \dots & \dots & 0 & 1 \end{pmatrix}, s_2 = \begin{pmatrix} 0 & 1 & 0 & \dots & 0 & 0 \\ 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 & 0 \\ 0 & 0 & 0 & \dots & 0 & 1 \end{pmatrix},$$

$$s_3 = \begin{pmatrix} 1 & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 & 0 \\ 0 & 0 & 0 & \dots & 0 & 1 \end{pmatrix}, s_4 = \begin{pmatrix} -1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & 1 & 0 \\ 0 & 0 & 0 & \dots & 0 & 1 \end{pmatrix}.$$

PROOF. *Step 1.* The permutation group S_n acts (effectively) on \mathbb{Z}^n by permuting the basis vectors; we, thus, obtain a monomorphism $\varphi : S_n \rightarrow GL(n, \mathbb{Z})$, so that $\varphi(12\dots n) = s_1$, $\varphi(12) = s_2$. Consider now the corresponding action of S_n on $n \times n$ matrices. Multiplication of a matrix by s_1 on the left permutes rows cyclically, multiplication to the right does the same with columns. Multiplication

by s_2 on the left swaps the first two rows, multiplication to the right does the same with columns. Therefore, by multiplying an elementary matrix A by appropriate products of s_1, s_1^{-1} and s_2 on the left and on the right, we obtain the matrix s_3 . In view of Exercise 7.5, the permutation $(12 \dots n)$ and the transposition (12) generate the permutation group S_n . Thus, every elementary matrix N_{ij} is a product of s_1, s_1^{-1}, s_2 and s_3 .

Let d_j denote the diagonal matrix with the diagonal entries $(1, \dots, 1, -1, 1, \dots, 1)$, where -1 occurs in j -th place. Thus, $d_1 = s_4$. The same argument as above, shows that for every d_j and $s = (1j) \in S_n$, $sd_js = d_1$. Thus, all diagonal matrices d_j belong to the subgroup generated by s_1, s_2 and s_4 .

Step 2. Now, let g be an arbitrary element in $GL(n, \mathbb{Z})$. Let a_1, \dots, a_n be the entries of the first column of g . We will prove that there exists an element p in $\langle s_1, s_2, s_3, s_4 \rangle \subset GL(n, \mathbb{Z})$, such that pg has the entries $1, 0, \dots, 0$ in its first column. We argue by induction on $k = C_1(g) = |a_1| + \dots + |a_n|$. Note that $k \geq 1$. If $k = 1$, then (a_1, \dots, a_n) is a permutation of $(\pm 1, 0, \dots, 0)$; hence, it suffices to take p in $\langle s_1, s_2, s_4 \rangle$ permuting the rows so as to obtain $1, 0, \dots, 0$ in the first column.

Assume that the statement is true for all integers $1 \leq i < k$; we will prove it for k . After to permuting rows and multiplying by $d_1 = s_4$ and d_2 , we may assume that $a_1 > a_2 > 0$. Then $N_{1,2}d_2g$ has the following entries in the first column: $a_1 - a_2, -a_2, a_3, \dots, a_n$. Therefore, $C_1(N_{1,2}d_2g) < C_1(g)$. By the induction assumption, there exists an element p of $\langle s_1, s_2, s_3, s_4 \rangle$ such that $pN_{1,2}d_2g$ has the entries of its first column equal to $1, 0, \dots, 0$. This proves the claim.

Step 3. We leave it to the reader to check that for every pair of matrices $A, B \in GL(n-1, \mathbb{R})$ and row vectors $L = (l_1, \dots, l_{n-1})$ and $M = (m_1, \dots, m_{n-1})$

$$\begin{pmatrix} 1 & L \\ 0 & A \end{pmatrix} \cdot \begin{pmatrix} 1 & M \\ 0 & B \end{pmatrix} = \begin{pmatrix} 1 & M + LB \\ 0 & AB \end{pmatrix}.$$

Therefore, the set of matrices

$$\left\{ \begin{pmatrix} 1 & L \\ 0 & A \end{pmatrix} ; A \in GL(n-1, \mathbb{Z}), L \in \mathbb{Z}^{n-1} \right\}$$

is a subgroup of $GL(n, \mathbb{Z})$ isomorphic to $\mathbb{Z}^{n-1} \times GL(n-1, \mathbb{Z})$.

Using this, an induction on n and Step 2, one shows that there exists an element p in $\langle s_1, s_2, s_3, s_4 \rangle$ such that pg is upper triangular and with entries on the diagonal equal to 1. It therefore suffices to prove that every integer upper triangular matrix as above is in $\langle s_1, s_2, s_3, s_4 \rangle$. This can be done for instance by **multiplying such a matrix to the right with matrices of the form $d_1 N_{1i}^{a_{1i}} d_1$, until all the entries on the first row become zero, except the diagonal one which remains 1; then by multiplying with $d_2 N_{2i}^{a_{2i}} d_2$ to perform the same operation on the second row etc. In the end we obtain the identity matrix, and can therefore deduce that every integer upper triangular matrix with entries on the diagonal equal to 1 is a product of matrices $d_i, i \in \{1, 2, \dots, n\}$, and $N_{jk}, j, k \in \{1, 2, \dots, n\}$, and is therefore in $\langle s_1, s_2, s_3, s_4 \rangle$.** \square

EXERCISE 7.13. Let G be a finitely generated group and let S be an infinite set of generators of G . Show that there exists a finite subset F of S so that G is generated by F .

EXERCISE 7.14. An element g of the group G is a *non-generator* if for every generating set S of G , the complement $S \setminus \{g\}$ is still a generating set of G .

- (a) Prove that the set of non-generators forms a subgroup of G . This subgroup is called the *Frattini subgroup*.
- (b) Compute the Frattini subgroup of $(\mathbb{Z}, +)$.
- (c) Compute the Frattini subgroup of $(\mathbb{Z}^n, +)$. (*Hint:* You may use the fact that $\text{Aut}(\mathbb{Z}^n)$ is $GL(n, \mathbb{Z})$, and that the $GL(n, \mathbb{Z})$ -orbit of e_1 is the set of vectors (k_1, \dots, k_n) in \mathbb{Z}^n such that $\gcd(k_1, \dots, k_n) = 1$.)

DEFINITION 7.15. A group G is said to have bounded generation property (or is boundedly generated) if there exists a finite subset $\{t_1, \dots, t_m\} \subset G$ such that every $g \in G$ can be written as

$$g = t_1^{k_1} t_2^{k_2} \cdots t_m^{k_m},$$

where k_1, k_2, \dots, k_m are integers.

Clearly, all finitely generated abelian groups have the bounded generation property, and so are all finite groups. On the other hand, the nonabelian free groups, which we will introduce in the next section, obviously, do not have the bounded generation property. For other examples of boundedly generated groups see Proposition 13.74. We also note that Alexey Muranov [Mur05] constructed examples of infinite boundedly generated *simple* groups.

7.2. Free groups

Let X be a set. Its elements are called *letters* or *symbols*. We define the set of *inverse letters* (or *inverse symbols*) $X^{-1} = \{a^{-1} \mid a \in X\}$. We will think of $X \cup X^{-1}$ as an *alphabet*.

A *word* in $X \cup X^{-1}$ is a finite (possibly empty) string of letters in $X \cup X^{-1}$, i.e. an expression of the form

$$a_{i_1}^{\epsilon_1} a_{i_2}^{\epsilon_2} \cdots a_{i_k}^{\epsilon_k},$$

where $a_i \in X$, $\epsilon_i = \pm 1$; here $x^1 = x$ for every $x \in X$. We will use the notation 1 for the *empty word* (the one which has no letters).

CONVENTION 7.16. Sometimes, by abusing the terminology, we will refer to words in $X \cup X^{-1}$ as *words in X* .

Denote by X^* the set of words in the alphabet $X \cup X^{-1}$, where the empty word, denoted by 1, is included. For instance,

$$a_1 a_2 a_1^{-1} a_2 a_1 \in X^*.$$

The *length* of a word w is the number of letters in this word. The length of the empty word is 0.

A word $w \in X^*$ is *reduced* if it contains no pair of consecutive letters of the form aa^{-1} or $a^{-1}a$. The *reduction* of a word $w \in X^*$ is the deletion of all pairs of consecutive letters of the form aa^{-1} or $a^{-1}a$.

For instance, the words

$$1, a_2 a_1, a_1 a_2 a_1^{-1}$$

are reduced, while

$$a_2 a_1 a_1^{-1} a_3$$

is not reduced.

More generally, a word w is *cyclically reduced* if it is reduced and, in addition, the first and the last letters of w are not inverses of each other. Equivalently, conjugating w by an element of $X \cup X^{-1}$:

$$w' = awa^{-1}, \quad a \in X \cup X^{-1}$$

results in a word w' whose reduction has length \gg the length of w .

We define an equivalence relation on X^* by $w \sim w'$ if w can be obtained from w' by a finite sequence of reductions and their inverses, i.e. the relation \sim on X^* is generated by

$$ua_i a_i^{-1} v \sim uv, \quad ua_i^{-1} a_i v \sim uv$$

where $u, v \in X^*$.

PROPOSITION 7.17. *Any word $w \in X^*$ is equivalent to a unique reduced word.*

PROOF. *Existence.* We prove the statement by induction on the length of a word. For words of length 0 and 1 the statement is clearly true. Assume that it is true for words of length n and consider a word of length $n+1$, $w = a_1 \cdots a_n a_{n+1}$, where $a_i \in X \cup X^{-1}$. According to the induction hypothesis, there exists a reduced word $u = b_1 \cdots b_k$ with $b_j \in X \cup X^{-1}$ such that $a_2 \cdots a_{n+1} \sim u$. Then $w \sim a_1 u$. If $a_1 \neq b_1^{-1}$ then $a_1 u$ is reduced. If $a_1 = b_1^{-1}$ then $a_1 u \sim b_2 \cdots b_k$ and the latter word is reduced.

Uniqueness. Let $F(X)$ be the set of reduced words in $X \cup X^{-1}$. For every $a \in X \cup X^{-1}$ we define a map $L_a : F(X) \rightarrow F(X)$ by

$$L_a(b_1 \cdots b_k) = \begin{cases} ab_1 \cdots b_k & \text{if } a \neq b_1^{-1}, \\ b_2 \cdots b_k & \text{if } a = b_1^{-1}. \end{cases}$$

For every word $w = a_1 \cdots a_n$ define $L_w = L_{a_1} \circ \cdots \circ L_{a_n}$. For the empty word 1 define $L_1 = \text{id}$. It is easy to check that $L_a \circ L_{a^{-1}} = \text{id}$ for every $a \in X \cup X^{-1}$, and to deduce from it that $v \sim w$ implies $L_v = L_w$.

We prove by induction on the length that if w is reduced then $w = L_w(1)$. The statement clearly holds for w of length 0 and 1. Assume that it is true for reduced words of length n and let w be a reduced word of length $n+1$. Then $w = au$, where $a \in X \cup X^{-1}$ and u is a reduced word that does not begin with a^{-1} , i.e. such that $L_a(u) = au$. Then $L_w(1) = L_a \circ L_u(1) = L_a(u) = au = w$.

In order to prove uniqueness it suffices to prove that if $v \sim w$ and v, w are reduced then $v = w$. Since $v \sim w$ it follows that $L_v = L_w$, hence $L_v(1) = L_w(1)$, that is $v = w$. \square

EXERCISE 7.18. Give a geometric proof of this proposition using identification of $w \in X^*$ with the set of edge-paths \mathfrak{p}_w in a regular tree T of valence $2|X|$, which start at a fixed vertex v_0 . The reduced path \mathfrak{p}^* in T corresponding to the reduction w^* of w is the unique geodesic in T connecting v_0 to the terminal point of \mathfrak{p} . Uniqueness of w^* then translates to the fact that a tree contains no circuits.

Let $F(X)$ be the set of reduced words in $X \cup X^{-1}$. Proposition 7.17 implies that X^*/\sim can be identified with $F(X)$.

DEFINITION 7.19. The *free group over X* is the set $F(X)$ endowed with the product $*$ defined by: $w * w'$ is the unique reduced word equivalent to the word ww' . The unit is the empty word.

The cardinality of X is called the *rank* of the free group $F(X)$.

We note that, at the moment, we have two, a priori distinct, notions of rank for (finitely generated) free groups: One is the least number of generators and the second is the cardinality of the set X . We will see, however, that the two numbers are the same.

The set $F = F(X)$ with the product defined in Definition 7.19 is indeed a group. The inverse of a reduced word

$$w = a_{i_1}^{\epsilon_1} a_{i_2}^{\epsilon_2} \cdots a_{i_k}^{\epsilon_k}$$

is given by

$$w^{-1} = a_{i_k}^{-\epsilon_k} a_{i_{k-1}}^{-\epsilon_{k-1}} \cdots a_{i_1}^{-\epsilon_1}.$$

It is clear that the product ww^{-1} projects to the empty word 1 in F .

EXERCISE 7.20. A free group of rank at least 2 is not abelian. Thus, *free non-abelian* means ‘free of rank at least 2.’

The *free semigroup* $F^s(X)$ with the generating set X is defined in the fashion similar to $F(X)$, except that we only allow the words in the alphabet X (and not in X^{-1}), in particular the reduction is not needed.

PROPOSITION 7.21 (Universal property of free groups). *A map $\varphi : X \rightarrow G$ from the set X to a group G can be extended to a homomorphism $\Phi : F(X) \rightarrow G$ and this extension is unique.*

PROOF. *Existence.* The map φ can be extended to a map on $X \cup X^{-1}$ (which we denote also φ) by $\varphi(a^{-1}) = \varphi(a)^{-1}$.

For every reduced word $w = a_1 \cdots a_n$ in $F = F(X)$ define

$$\Phi(a_1 \cdots a_n) = \varphi(a_1) \cdots \varphi(a_n).$$

Set $\Phi(1_F) := 1_G$, the identity element of G . We leave it to the reader to check that Φ is a homomorphism.

Uniqueness. Let $\Psi : F(X) \rightarrow G$ be a homomorphism such that $\Psi(x) = \varphi(x)$ for every $x \in X$. Then for every reduced word $w = a_1 \cdots a_n$ in $F(X)$,

$$\Psi(w) = \Psi(a_1) \cdots \Psi(a_n) = \varphi(a_1) \cdots \varphi(a_n) = \Phi(w).$$

□

COROLLARY 7.22. *Every group is the quotient of a free group.*

PROOF. Apply Proposition 7.21 to the group G and a generating set X of G (e.g., $X = G$). □

LEMMA 7.23. *Every short exact sequence $1 \rightarrow N \rightarrow G \xrightarrow{r} F(X) \rightarrow 1$ splits. In particular, G contains a subgroup isomorphic to $F(X)$.*

PROOF. Indeed, for each $x \in X$ consider choose an element $t_x \in G$ projecting to x ; the map $x \mapsto t_x$ extends to a group homomorphism $s : F(X) \rightarrow G$. Composition $r \circ s$ is the identity homomorphism $F(X) \rightarrow F(X)$ (since it is the identity on the generating set X). Therefore, the homomorphism s is a splitting of the exact sequence. Since $r \circ s = \text{Id}$, s is a monomorphism. □

COROLLARY 7.24. *Every short exact sequence $1 \rightarrow N \rightarrow G \rightarrow \mathbb{Z} \rightarrow 1$ splits.*

7.3. Presentations of groups

Let G be a group and S a generating set of G . According to Proposition 7.21, the inclusion map $i : S \rightarrow G$ extends uniquely to an epimorphism $\pi_S : F(S) \rightarrow G$. The elements of $\text{Ker}(\pi_S)$ are called *relators* (or *relations*) of the group G with the generating set S .

N.B. In the above, by an abuse of language we used the symbol s to designate two different objects: s is a letter in $F(S)$, as well as an element in the group G .

If $R = \{r_i \mid i \in I\} \subset F(S)$ is such that $\text{Ker}(\pi_S)$ is normally generated by R (i.e. $\langle\langle R \rangle\rangle = \text{Ker}(\pi_S)$) then we say that the ordered pair (S, R) , usually denoted $\langle S \mid R \rangle$, is a *presentation of G* . The elements $r \in R$ are called *defining relators* (or *defining relations*) of the presentation $\langle S \mid R \rangle$.

A group G is said to be *finitely presented* if it admits a finite presentation, i.e. a presentation with finitely many generators and relators.

By abuse of language we also say that the generators $s \in S$ and the *relations* $r = 1, r \in R$, constitute a presentation of the group G . Sometimes we will write presentations in the form

$$\langle s_i, i \in I \mid r_j = 1, j \in J \rangle$$

where

$$S = \{x_i\}_{i \in I}, \quad R = \{r_j\}_{j \in J}.$$

If both S and R are finite, then the pair S, R is called a *finite presentation of G* . A group G is called *finitely presented* if it admits a finite presentation. Sometimes it is difficult, and even algorithmically impossible, to find a finite presentation of a finitely presented group, see [BW11].

Conversely, given an alphabet S and a set R of (reduced) words in the alphabet S , we can form the quotient

$$G := F(S) / \langle\langle R \rangle\rangle.$$

Then $\langle S \mid R \rangle$ is a presentation of G . By abusing notation, we will often write

$$G = \langle S \mid R \rangle,$$

if G is a group with the presentation $\langle S \mid R \rangle$. If w is a word in the generating set S , we will use $[w]$ to denote its projection to the group G . An alternative notation for the equality

$$[v] = [w]$$

is

$$v \equiv_G w.$$

Note that the significance of a presentation of a group is the following:

- every element in G can be written as a finite product $x_1 \cdots x_n$ with

$$x_i \in S \cup S^{-1} = \{s^{\pm 1} : s \in S\}$$

i.e. as a word in the alphabet $S \cup S^{-1}$;

- a word $w = x_1 \cdots x_n$ in the alphabet $S \cup S^{-1}$ is equal to the identity in G , $w \equiv_G 1$, if and only if in $F(S)$ the word w is the product of finitely many conjugates of words $r_i \in R$, i.e.

$$w = \prod_{i=1}^m r_i^{u_i}$$

for some $m \in \mathbb{N}$, $u_i \in F(S)$ and $r_i \in R$.

Below are few examples of group presentations:

EXAMPLES 7.25. (1) $\langle a_1, \dots, a_n \mid [a_i, a_j], 1 \leq i, j \leq n \rangle$ is a finite presentation of \mathbb{Z}^n ;

(2) $\langle x, y \mid x^n, y^2, yxyx \rangle$ is a presentation of the finite dihedral group D_{2n} ;

(3) $\langle x, y \mid x^2, y^3, [x, y] \rangle$ is a presentation of the cyclic group \mathbb{Z}_6 .

Let $\langle S \mid R \rangle$ be a presentation of a group G . Let H be a group and $\psi : S \rightarrow H$ be a map which “preserves the relators”, i.e. $\psi(r) = 1$ for every $r \in R$. Then:

LEMMA 7.26. *The map ψ extends to a group homomorphism $\psi : G \rightarrow H$.*

PROOF. By the universal property of free groups, the map ψ extends to a homomorphism $\tilde{\psi} : F(S) \rightarrow H$. We need to show that $\langle\langle R \rangle\rangle$ is contained in $\text{Ker}(\tilde{\psi})$. However, $\langle\langle R \rangle\rangle$ consists of products of elements of the form grg^{-1} , where $g \in F$, $r \in R$. Since $\tilde{\psi}(grg^{-1}) = 1$, the claim follows. \square

EXERCISE 7.27. The group $\bigoplus_{x \in X} \mathbb{Z}_2$ has the presentation

$$\langle x \in X \mid x^2, [x, y], \forall x, y \in X \rangle.$$

PROPOSITION 7.28 (Finite presentability is independent of the generating set). *Assume that a group G has finite presentation $\langle S \mid R \rangle$, and let $\langle X \mid T \rangle$ be an arbitrary presentation of G , such that X is finite. Then there exists a finite subset $T_0 \subset T$ such that $\langle X \mid T_0 \rangle$ is a presentation of G .*

PROOF. Every element $s \in S$ can be written as a word $a_s(X)$ in X . The map $i_{SX} : S \rightarrow F(X)$, $i_{SX}(s) = a_s(X)$ extends to a unique homomorphism $p : F(S) \rightarrow F(X)$. Moreover, since $\pi_X \circ i_{SX}$ is an inclusion map of S into G , and both π_S and $\pi_X \circ p$ are homomorphisms from $F(S)$ to G extending the inclusion map $S \rightarrow G$, by the uniqueness of the extension we have that

$$\pi_S = \pi_X \circ p.$$

This implies that $\text{Ker}(\pi_X)$ contains $p(r)$ for every $r \in R$.

Likewise, every $x \in X$ can be written as a word $b_x(S)$ in S , and this defines a map $i_{XS} : X \rightarrow F(S)$, $i_{XS}(x) = b_x(S)$, which extends to a homomorphism $q : F(X) \rightarrow F(S)$. A similar argument shows that $\pi_S \circ q = \pi_X$.

For every $x \in X$,

$$\pi_X(p(q(x))) = \pi_S(q(x)) = \pi_X(x).$$

This implies that for every $x \in X$, $x^{-1}p(q(x))$ is in $\text{Ker}(\pi_X)$. Let N be the normal subgroup of $F(X)$ normally generated by

$$\{p(r) \mid r \in R\} \cup \{x^{-1}p(q(x)) \mid x \in X\}.$$

We have that $N \leq \text{Ker}(\pi_X)$. Therefore, there is a natural projection

$$\text{proj} : F(X)/N \rightarrow F(X)/\text{Ker}(\pi_X).$$

Let $\bar{p} : F(S) \rightarrow F(X)/N$ be the homomorphism induced by p . Since $\bar{p}(r) = 1$ for all $r \in R$, it follows that $\bar{p}(\text{Ker} \pi_S) = 1$, hence, \bar{p} induces a homomorphism

$$\varphi : F(S)/\text{Ker}(\pi_S) \rightarrow F(X)/N.$$

We next observe that the homomorphism φ is onto. Indeed, $F(X)/N$ is generated by elements of the form $xN = p(q(x))N$, and the latter is the image under φ of $q(x)\text{Ker}(\pi_S)$.

Consider the homomorphism

$$\text{proj} \circ \varphi : F(S)/\text{Ker}(\pi_S) \rightarrow F(X)/\text{Ker}(\pi_X)$$

Both the domain and the target groups are isomorphic to G . Each element x of the generating set X is sent by the isomorphism $G \rightarrow F(S)/\text{Ker}(\pi_S)$ to $q(x)\text{Ker}(\pi_S)$. The same element x is sent by the isomorphism $G \rightarrow F(X)/\text{Ker}(\pi_X)$ to $x\text{Ker}(\pi_X)$. Note that

$$\text{proj} \circ \varphi(q(x)\text{Ker}(\pi_S)) = \text{proj}(xN) = x\text{Ker}(\pi_X).$$

This means that, modulo the two isomorphisms mentioned above, the map $\text{proj} \circ \varphi$ is id_G . This implies that φ is injective, hence, a bijection. Therefore, proj is also a bijection. This happens if and only if $N = \text{Ker}(\pi_X)$. In particular, $\text{Ker}(\pi_X)$ is normally generated by the finite set of relators

$$\mathfrak{R} = \{p(r) \mid r \in R\} \cup \{x^{-1}p(q(x)) \mid x \in X\}.$$

Since $\mathfrak{R} = \langle\langle T \rangle\rangle$, every relator $\rho \in \mathfrak{R}$ can be written as a product

$$\prod_{i \in I_\rho} t_i^{v_i}$$

with $v_i \in F(X)$, $t_i \in T$ and I_ρ finite. It follows that $\text{Ker}(\pi_X)$ is normally generated by the finite subset

$$T_0 = \bigcup_{\rho \in \mathfrak{R}} \{t_i \mid i \in I_\rho\}$$

of T . □

Proposition 7.28 can be reformulated as follows: If G is finitely presented, X is finite and

$$1 \rightarrow N \rightarrow F(X) \rightarrow G \rightarrow 1$$

is a short exact sequence, then N is normally generated by finitely many elements n_1, \dots, n_k . This can be generalized to an arbitrary short exact sequence:

LEMMA 7.29. *Consider a short exact sequence*

$$(7.1) \quad 1 \rightarrow N \rightarrow K \xrightarrow{\pi} G \rightarrow 1, \quad \text{with } K \text{ finitely generated.}$$

If G is finitely presented, then N is normally generated by finitely many elements $n_1, \dots, n_k \in N$.

PROOF. Let S be a finite generating set of K ; then $\bar{S} = \pi(S)$ is a finite generating set of G . Since G is finitely presented, by Proposition 7.28 there exist finitely many words r_1, \dots, r_k in S such that

$$\langle \bar{S} \mid r_1(\bar{S}), \dots, r_k(\bar{S}) \rangle$$

is a presentation of G .

Define $n_j = r_j(S)$, an element of N by the assumption.

Let n be an arbitrary element in N and $w(S)$ a word in S such that $n = w(S)$ in K . Then $w(\bar{S}) = \pi(n) = 1$, whence in $F(S)$ the word $w(S)$ is a product of finitely many conjugates of r_1, \dots, r_k . When projecting such a relation via $F(S) \rightarrow K$ we obtain that n is a product of finitely many conjugates of n_1, \dots, n_k . \square

PROPOSITION 7.30. *Suppose that N a normal subgroup of a group G . If both N and G/N are finitely presented then G is also finitely presented.*

PROOF. Let X be a finite generating set of N and let Y be a finite subset of G such that $\bar{Y} = \{yN \mid y \in Y\}$ is a generating set of G/N . Let $\langle X \mid r_1, \dots, r_k \rangle$ be a finite presentation of N and let $\langle \bar{Y} \mid \rho_1, \dots, \rho_m \rangle$ be a finite presentation of G/N . The group G is generated by $S = X \cup Y$ and this set of generators satisfies a list of relations of the following form:

$$(7.2) \quad r_i(X) = 1, 1 \leq i \leq k, \rho_j(Y) = u_j(X), 1 \leq j \leq m,$$

$$(7.3) \quad x^y = v_{xy}(X), x^{y^{-1}} = w_{xy}(X)$$

for some words u_j, v_{xy}, w_{xy} in S .

We claim that this is a complete set of defining relators of G .

All the relations above can be rewritten as $t(X, Y) = 1$ for a finite set T of words t in S . Let K be the normal subgroup of $F(S)$ normally generated by T .

The epimorphism $\pi_S : F(S) \rightarrow G$ defines an epimorphism $\varphi : F(S)/K \rightarrow G$. Let wK be an element in $\text{Ker}(\varphi)$, where w is a word in S . Due to the set of relations (7.3), there exist a word $w_1(X)$ in X and a word $w_2(Y)$ in Y , such that $wK = w_1(X)w_2(Y)K$.

Applying the projection $\pi : G \rightarrow G/N$, we see that $\pi(\varphi(wK)) = 1$, i.e. $\pi(\varphi(w_2(Y)K)) = 1$. This implies that $w_2(Y)$ is a product of finitely many conjugates of $\rho_i(Y)$, hence $w_2(Y)K$ is a product of finitely many conjugates of $u_j(X)K$, by the second set of relations in (7.2). This and the relations (7.3) imply that $w_1(X)w_2(Y)K = v(X)K$ for some word $v(X)$ in X . Then the image $\varphi(wK) = \varphi(v(X)K)$ is in N ; therefore, $v(X)$ is a product of finitely many conjugates of relators $r_i(X)$. This implies that $v(X)K = K$.

We have thus obtained that $\text{Ker}(\varphi)$ is trivial, hence φ is an isomorphism, equivalently that $K = \text{Ker}(\pi_S)$. This implies that $\text{Ker}(\pi_S)$ is normally generated by the finite set of relators listed in (7.2) and (7.3). \square

We continue with a list of finite presentations of some important groups:

EXAMPLES 7.31. (1) *Surface groups:*

$$\Pi_n = \langle a_1, b_1, \dots, a_n, b_n \mid [a_1, b_1] \cdots [a_n, b_n] \rangle,$$

is the fundamental group of the closed connected oriented surface of genus n , see e.g. [Hat02, Mas91].

- (2) *Right-angled Artin groups (RAAGs)*. Let \mathcal{G} be a finite graph with the vertex set $V = \{x_1, \dots, x_n\}$ and the edge set E consisting of the edges $\{[x_i, x_j]\}_{i,j}$. Define the *right-angled Artin group* by

$$A_{\mathcal{G}} := \langle V \mid [x_i, x_j], \text{ whenever } [x_i, x_j] \in E \rangle.$$

Here we commit a useful abuse of notation: In the first instance $[x_i, x_j]$ denotes the commutator and in the second instance it denotes the edge of \mathcal{G} connecting x_i to x_j .

EXERCISE 7.32. a. If \mathcal{G} contains no edges then $A_{\mathcal{G}}$ is a free group on n generators.

b. If \mathcal{G} is the complete graph on n vertices then

$$A_{\mathcal{G}} \cong \mathbb{Z}^n.$$

- (3) *Coxeter groups*. Let \mathcal{G} be a finite simple graph. Let V and E denote be the vertex and the edge set of \mathcal{G} respectively. Put a label $m(e) \in \mathbb{N} \setminus \{1\}$ on each edge $e = [x_i, x_j]$ of \mathcal{G} . Call the pair

$$\Gamma := (\mathcal{G}, m : E \rightarrow \mathbb{N} \setminus \{1\})$$

a *Coxeter graph*. Then Γ defines the *Coxeter group* C_{Γ} :

$$C_{\Gamma} := \left\langle x_i \in V \mid x_i^2, (x_i x_j)^{m(e)}, \text{ whenever there exists an edge } e = [x_i, x_j] \right\rangle.$$

See [Dav08] for the detailed discussion of Coxeter groups.

- (4) *Artin groups*. Let Γ be a Coxeter graph. Define

$$A_{\Gamma} := \left\langle x_i \in V \mid \underbrace{x_i x_j \cdots}_{m(e) \text{ terms}} = \left(\underbrace{x_j x_i \cdots}_{m(e) \text{ terms}} \right), \text{ whenever } e = [x_i, x_j] \in E \right\rangle.$$

Then A_{Γ} is a right-angled Artin group if and only if $m(e) = 2$ for every $e \in E$. In general, C_{Γ} is the quotient of A_{Γ} by the subgroup normally generated by the set

$$\{x_i^2 : x_i \in V\}.$$

- (5) *Shephard groups*: Let Γ be a Coxeter graph. Label vertices of Γ with natural numbers $n_x, x \in V(\Gamma)$. Now, take a group, a Shepherd group, S_{Γ} to be generated by vertices $x \in V$, subject to Artin relators and, in addition, relators

$$x^{n_x}, \quad x \in V.$$

Note that, in the case $n_x = 2$ for all $x \in V$, the group which we obtain is the Coxeter group C_{Γ} . Shephard groups (and von Dyck groups below) are *complex analogues* of Coxeter groups.

- (6) *Generalized von Dyck groups*: Let Γ be a labeled graph as in the previous example. Define a group D_{Γ} to be generated by vertices $x \in V$, subject to the relators

$$\begin{aligned} x^{n_x}, \quad x \in V; \\ (xy)^{m(e)}, e = [x, y] \in E. \end{aligned}$$

If Γ consists of a single edge, then D_{Γ} is called a *von Dyck group*. Every von Dyck group D_{Γ} is an index 2 subgroup in the Coxeter group C_{Δ} ,

where Δ is the triangle with edge-labels p, q, r , which are the vertex-edge labels of Γ .

(7) *Integer Heisenberg group:*

$$H_{2n+1}(\mathbb{Z}) := \langle x_1, \dots, x_n, y_1, \dots, y_n, z \mid [x_i, z] = 1, [y_j, z] = 1, [x_i, x_j] = 1, [y_i, y_j] = 1, [x_i, y_j] = z^{\delta_{ij}}, 1 \leq i, j \leq n \rangle.$$

(8) *Baumslag–Solitar groups:*

$$BS(m, n) = \langle a, b \mid ab^m a^{-1} = b^n \rangle,$$

where m, n are non-zero integers.

EXERCISE 7.33. Show that $H_{2n+1}(\mathbb{Z})$ is isomorphic to the group appearing in Example 13.36, (3).

The classes of groups described so far were defined combinatorially, in terms of their presentations. Below are several important classes of finitely presented groups which are defined *geometrically*:

- (1) $CAT(-1)$ groups: Groups G which act geometrically on $CAT(-1)$ metric spaces.
- (2) $CAT(0)$ groups: Groups G which act geometrically on $CAT(0)$ metric spaces.
- (3) Automatic groups: We refer the reader to [ECH⁺92] for the definition.
- (4) Hyperbolic and relatively hyperbolic groups, which will be defined in Chapter 11.
- (5) Semihyperbolic groups, see [JA95].

An important feature of finitely presented groups is provided by the following theorem, see e.g. [Hat02]:

THEOREM 7.34. *Every finitely generated group is the fundamental group of a smooth closed manifold of dimension 4.*

Laws in groups.

DEFINITION 7.35. An *identity* (or *law*) is a non-trivial reduced word

$$w = w(x_1, \dots, x_n)$$

in the letters x_1, \dots, x_n and their inverses. A group G is said to *satisfy the identity (law) $w = 1$* if this equality is satisfied in G whenever x_1, \dots, x_n are replaced by arbitrary elements in G . In other words, for the group

$$Q = \langle x_1, \dots, x_n \mid w \rangle,$$

the pull-back map

$$\text{Hom}(Q, G) \longrightarrow \text{Hom}(F_n, G)$$

is surjective.

EXAMPLES 7.36 (Groups satisfying a law). (1) Abelian groups. Here the law is

$$w(x_1, x_2) = x_1 x_2 x_1^{-1} x_2^{-1}.$$

(2) Solvable groups, see section 13.6, equation (13.10).

(3) Free Burnside groups. The *free Burnside group*

$$B(n, m) = \langle x_1, \dots, x_n \mid w^n \text{ for every word } w \text{ in } x_1^{\pm 1}, \dots, x_n^{\pm 1} \rangle.$$

It is known that these groups are infinite for sufficiently large m (see [Ady79], [Ol'91a], [Iva94], [Lys96], [DG] and references therein).

Note that free nonabelian groups (and, hence, groups containing them) do not satisfy any law.

7.4. The rank of a free group determines the group. Subgroups

PROPOSITION 7.37. *Two free groups $F(X)$ and $F(Y)$ are isomorphic if and only if X and Y have the same cardinality.*

PROOF. A bijection $\varphi : X \rightarrow Y$ extends to an isomorphism $\Phi : F(X) \rightarrow F(Y)$ by Proposition 7.21. Therefore, two free groups $F(X)$ and $F(Y)$ are isomorphic if X and Y have the same cardinality.

Conversely, let $\Phi : F(X) \rightarrow F(Y)$ be an isomorphism. Take $N := N(X) \leq F(X)$, the subgroup generated by the subset $\{g^2 : g \in F(X)\}$; clearly, N is normal in $F(X)$. Then $\Phi(N(X)) = N(Y)$ is the normal subgroup generated by $\{h^2 : h \in F(Y)\}$. It follows that Φ induces an isomorphism $\Psi : F(X)/N(X) \rightarrow F(Y)/N(Y)$.

LEMMA 7.38. *The quotient $\bar{F} := F/N$ is isomorphic to $A = \mathbb{Z}_2^{\oplus X}$, where $F = F(X)$.*

PROOF. Recall that A has the presentation

$$\langle x \in X \mid x^2, [x, y], \forall x, y \in X \rangle,$$

see Exercise 7.27. We now prove the assertion of the lemma. Let $\pi : F \rightarrow \bar{F}$ denote the quotient map. Since $\pi(g) = \pi(g^{-1})$ for all $g \in F$, we conclude that for all $g, h \in X$,

$$1 = \pi((hg)^2) = \pi([g, h]),$$

and, therefore, \bar{F} is abelian.

Consider the map $\eta : F \rightarrow A$ sending the generators of F to the obvious generators of A . Since A satisfies the law $a^2 = 1$ for all $a \in A$, it is clear that $\eta = \phi \circ \pi$, for some homomorphism $\phi : \bar{F} \rightarrow A$. We next construct the inverse ψ to ϕ . We define ψ on the generators $x \in X$ of A : $\psi(x) = \bar{x} = \pi(x)$. We need to show that ψ preserves the relators of A (as in Lemma 7.26): Since \bar{F} is abelian, $[\psi(x), \psi(y)] = 1$ for all $x, y \in X$. Moreover, $\psi(x)^2 = 1$ since \bar{F} also satisfies the law $g^2 = 1$. It is clear that ϕ, ψ are inverses to each other. \square

Thus, $F(X)/N(X)$ is isomorphic to $\mathbb{Z}_2^{\oplus X}$, while $F(Y)/N(Y)$ is isomorphic to $\mathbb{Z}_2^{\oplus Y}$. It follows that $\mathbb{Z}_2^{\oplus X} \cong \mathbb{Z}_2^{\oplus Y}$ as \mathbb{Z}_2 -vector spaces. Therefore, X and Y have the same cardinality, by uniqueness of the dimension of vector spaces. \square

REMARK 7.39. Proposition 7.37 implies that for every cardinal number n there exists, up to isomorphism, exactly one free group of rank n . We denote this group by F_n .

Recall that the *rank* of a finitely generated group G is the least number of generators of G . In other words,

$$\text{rank}(G) = \min\{r : \exists \text{ an epimorphism } F_r \rightarrow G\}.$$

COROLLARY 7.40. *For each finite n , the number n is the least cardinality of a generating set of F_n . In other words, $\text{rank}(F_n) = n$.*

PROOF. If this theorem fails, there exists a epimorphism

$$h : F(X) \rightarrow F(Y), \quad |X| = m < |Y| = n.$$

This epimorphism projects to an epimorphism of the abelian quotients

$$\bar{h} : A = F(X)/N(X) \rightarrow B = F(Y)/N(Y).$$

However, A and B are vector spaces over \mathbb{Z}_2 of dimensions m and n respectively. This contradicts the assumption that $m < n$. \square

THEOREM 7.41 (Nielsen–Schreier). *Any subgroup of a free group is a free group.*

This theorem will be proven in Corollary 7.79 using topological methods; see also [LS77, Proposition 2.11].

7.5. Free constructions: Amalgams of groups and graphs of groups

7.5.1. Amalgams. Amalgams (amalgamated free products and HNN extensions) allow one to build more complicated groups starting with a given pair of groups or a group and a pair of its subgroups which are isomorphic to each other.

Amalgamated free products. As a warm-up we first define the *free product* of groups $G_1 = \langle X_1 | R_1 \rangle, G_2 = \langle X_2 | R_2 \rangle$ by the presentation:

$$G_1 \star G_2 = \langle G_1, G_2 | \quad \rangle,$$

which is a shorthand for the presentation:

$$\langle X_1 \sqcup X_2 | R_1 \sqcup R_2 \rangle.$$

For instance, the free group of rank 2 is isomorphic to $\mathbb{Z} \star \mathbb{Z}$.

More generally, suppose that we are given subgroups $H_i \leq G_i$ ($i = 1, 2$) and an isomorphism

$$\phi : H_1 \rightarrow H_2.$$

Define the *amalgamated free product*

$$G_1 \star_{H_1 \cong H_2} G_2 = \langle G_1, G_2 | \phi(h)h^{-1}, h \in H_1 \rangle.$$

In other words, in addition to the relators in G_1, G_2 we identify $\phi(h)$ with h for each $h \in H_1$. A common shorthand for the amalgamated free product is

$$G_1 \star_H G_2,$$

where $H \cong H_1 \cong H_2$ (the embeddings of H into G_1 and G_2 are suppressed in this notation).

HNN extensions. This construction is named after G. Higman, B. Neumann and H. Neumann who first introduced it in [HNN49]. It is a variation on the amalgamated free product where $G_1 = G_2$. Namely, suppose that we are given

a group G , its subgroup H and a monomorphism $\phi : H \rightarrow G$. Then the HNN extension of G via ϕ is defined as

$$G \star_{H, \phi} = \langle G, t | t h t^{-1} = \phi(h), \forall h \in H \rangle.$$

A common shorthand for the HNN extension is

$$G \star_H$$

where the monomorphism ϕ is suppressed in this notation.

EXERCISE 7.42. Suppose that H is the trivial subgroup. Then

$$G \star_H \cong G \star \mathbb{Z}.$$

EXERCISE 7.43. Let $G = \langle S | R \rangle$, where R is a single relator which contains each letter $x \in X$ exactly twice (possibly as x^{-1}). Show that G is isomorphic to the free product of the fundamental group of a closed surface and a free group. Give an example where the free factor is non-trivial.

More generally, one defines *simultaneous HNN extension* of G along a collection of isomorphic subgroups: Suppose that we are given a collection of subgroups $H_j, j \in J$ of G and isomorphic embeddings $\phi_j : H_j \rightarrow G$. Then define the group

$$G \star_{\phi_j : H_j \rightarrow G, j \in J} = \langle G, t_j, j \in J | t_j h t_j^{-1} = \phi_j(h), \forall h \in H_j, j \in J \rangle.$$

7.5.2. Graphs of groups. In this section, graphs are no longer assumed to be simplicial, but are assumed to be connected. The notion of graphs of groups is a very useful generalization of both the amalgamated free product and the HNN extension.

Suppose that Γ is a graph. Assign to each vertex v of Γ a *vertex group* G_v ; assign to each edge e of Γ an *edge group* G_e . We orient each edge e so that its head is e_+ and the tail is e_- (this allows for the possibility that $e_+ = e_-$). Suppose, furthermore, that for each edge e we are given *monomorphisms*

$$\phi_{e_+} : G_e \rightarrow G_{e_+}, \phi_{e_-} : G_e \rightarrow G_{e_-}.$$

REMARK 7.44. More generally, one can allow non-injective homomorphisms

$$G_e \rightarrow G_{e_+}, \quad G_e \rightarrow G_{e_-},$$

but we will not consider them here, see [Mas91].

The oriented graph Γ together with the collection of vertex and edge groups and the monomorphisms $\phi_{e_{\pm}}$ is called a *graph of groups* \mathcal{G} based on the graph Γ .


Our next goal is to convert (connected) graphs of groups \mathcal{G} into groups. We first do this in the case when Γ is simply-connected, i.e. is a tree.

DEFINITION 7.45. Suppose that Γ is a tree. The *fundamental group* $\pi(\mathcal{G}) = \pi_1(\mathcal{G})$ of a graph of groups based on a tree Γ is a group G satisfying the following:


1. There is a collection of *compatible homomorphisms*

$$G_v \rightarrow G, G_e \rightarrow G, v \in V(\Gamma), e \in E(\Gamma),$$

i.e. that whenever $v = e_{\pm}$, we have the commutative diagram



2. The group G is *universal* with respect to the above property, i.e. given any group H and a collection of compatible homomorphisms $G_v \rightarrow H, G_e \rightarrow H$, there exists a unique homomorphism $G \rightarrow H$ such that we have commutative diagrams



for all $v \in V(\Gamma)$.

Note that the above definition easily implies that $G = \pi(\mathcal{G})$ is unique (up to an isomorphism). For the existence of $\pi(\mathcal{G})$ see [Ser80] and the discussion below. It is also a non-trivial (but not a very difficult to prove) fact that the homomorphisms $G_v \rightarrow G$ are injective.

Suppose now that Γ is connected but not simply-connected. We then let $T \subset \Gamma$ be a maximal subtree and $\mathcal{T} \subset \mathcal{G}$ be the corresponding subgraph of groups. Set $G_T := \pi(\mathcal{T})$. For each edge $e = [v, w] \in E(\Gamma)$ which is not in T , we have embeddings ψ_{e^\pm} obtained by composing $\phi_{e^\pm} : G_e \rightarrow G_v, G_w$ with embeddings $G_v, G_w \rightarrow G_T$. Thus, for each edge e which is not in T , we have two isomorphisms $G_e \rightarrow G_e^\pm < G_T$ and, accordingly, we obtain isomorphisms $G_e^- \rightarrow G_e^+$. Lastly, using these isomorphisms, define the simultaneous HNN extension G of G_T . Lastly, set $\pi(\mathcal{G}) = G$.

Whenever $G \cong \pi(\mathcal{G})$, we will say that \mathcal{G} determines a *graph of groups decomposition* of G . The decomposition \mathcal{G} is called *trivial* if there is a vertex v so that the natural homomorphism $G_v \rightarrow G$ is onto.

EXAMPLE 7.46. 1. Suppose that the graph Γ consists of a single edge e whose head e_+ is the vertex called 2 and the tail e_- is the vertex called 1. Assume that $\phi_{e_-}(G_e) = H_1 \leq G_1, \phi_{e_+}(G_e) = H_2 \leq G_2$. Then

$$\pi(\mathcal{G}) \cong G_1 \star_{H_1 \cong H_2} G_2.$$

2. Suppose that the graph Γ is a monogon, consisting of an edge e connecting the vertex called 1 to itself. Suppose, furthermore, $\phi_{e_-}(G_e) = H_1 \leq G_1, \phi_{e_+}(G_e) = H_2 \leq G_1$. Then

$$\pi(\mathcal{G}) \cong G_1 \star_{H_1 \cong H_2} .$$

Once this example is understood, one can show that for every graph of groups \mathcal{G} , the group $\pi_1(\mathcal{G})$ exists by describing this group in terms of generators and relators in the manner similar to the definition of the amalgamated free product and the HNN extension. In the next section we will see how to construct $\pi_1(\mathcal{G})$ using topology.

7.5.3. Converting graphs of groups into amalgams. Suppose that \mathcal{G} is a graph of groups and $G = \pi_1(\mathcal{G})$. Our goal is to convert \mathcal{G} into an amalgam decomposition of G . There are two cases to consider:

1. Suppose that the graph Γ underlying \mathcal{G} contains an oriented edge $e = [v_1, v_2]$ so that e separates Γ in the sense that the graph Γ' obtained from Γ by removing e (and keeping v_1, v_2) is a disjoint union of connected subgraphs $\Gamma_1 \sqcup \Gamma_2$, where $v_i \in V(\Gamma_i)$. Let \mathcal{G}_i denote the subgraph in the graph of groups \mathcal{G} , corresponding to $\Gamma_i, i = 1, 2$. Then set

$$G_i := \pi_1(\mathcal{G}_i), \quad i = 1, 2, \quad G_3 := G_e.$$

We have composition of embeddings $G_e \rightarrow G_{v_i} \rightarrow G_i \rightarrow G$. Then the universal property of $\pi_1(\mathcal{G}_i)$ and $\pi_1(\mathcal{G})$ implies that $G \cong G_1 \star_{G_3} G_2$: One simply verifies that G satisfies the universal property for the amalgam $G_1 \star_{G_3} G_2$.

2. Suppose that Γ contains an oriented edge $e = [v_1, v_2]$ such that e does not separate Γ . Let $\Gamma_1 := \Gamma'$, where Γ' is obtained from Γ by removing the edge e as in the Case 1. Set $G_1 := \pi_1(\mathcal{G}_1)$ as before. Then the embeddings

$$G_e \rightarrow G_{v_i}, i = 1, 2$$

induce embeddings $G_e \rightarrow G_i$ with the images H_1, H_2 respectively. Similarly to the Case 1, we obtain

$$G \cong G_1 \star_{G_e} = G_1 \star_{H_1 \cong H_2}$$

where the isomorphism $H_1 \rightarrow H_2$ is given by the composition

$$H_1 \rightarrow G_e \rightarrow H_2.$$

Clearly, \mathcal{G} is trivial if and only if the corresponding amalgam $G_1 \star_{G_3} G_2$ or $G_1 \star_{G_e}$ is trivial.

7.5.4. Topological interpretation of graphs of groups. Let \mathcal{G} be a graph of groups. Suppose that for all vertices and edges $v \in V(\Gamma)$ and $e \in E(\Gamma)$ we are given connected cell complexes M_v, M_e with the fundamental groups G_v, G_e respectively. For each edge $e = [v, w]$ assume that we are given a continuous map $f_{e\pm} : M_e \rightarrow M_{e\pm}$ which induces the monomorphism $\phi_{e\pm}$. This collection of spaces and maps is called a *graph of spaces*

$$\mathcal{G}_M := \{M_v, M_e, f_{e\pm} : M_e \rightarrow M_{e\pm} : v \in V(\Gamma), e \in E(\Gamma)\}.$$

In order to construct \mathcal{G}_M starting from \mathcal{G} , recall that each group G admits a cell complex $K(G, 1)$ whose fundamental group is G and whose universal cover is contractible, see Section 5.8.2. Given a group homomorphism $\phi : H \rightarrow G$, there exists a continuous map, unique up to homotopy,

$$f : K(H, 1) \rightarrow K(G, 1)$$

which induces the homomorphism ϕ . Then one can take $M_v := K(G_v, 1)$, $M_e := K(G_e, 1)$, etc.

To simplify the picture (although this is not the general case), the reader can think of each M_v as a manifold with several boundary components which are homeomorphic to M_{e_1}, M_{e_2}, \dots , where e_j are the edges having v as their *head* or *tail*. Then assume that the maps $f_{e\pm}$ are homeomorphisms onto the respective boundary components.

For each edge e we form the product $M_e \times [0, 1]$ and then form the double mapping cylinders for the maps $f_{e_{\pm}}$, i.e. identify points of $M_e \times \{0\}$ and $M_e \times \{1\}$ with their images under f_{e_-} and f_{e_+} respectively. Let M denote the resulting cell complex. It then follows from the Seifert–Van Kampen theorem [Mas91] that

THEOREM 7.47. *The group $\pi_1(M)$ is isomorphic to $\pi(\mathcal{G})$.*

This theorem allows one to think of the graphs of groups and their fundamental groups *topologically* rather than *algebraically*.

EXERCISE 7.48. Use the above interpretation to show that for each vertex $v \in V(\Gamma)$ the canonical homomorphism $G_v \rightarrow \pi(\mathcal{G})$ is injective.

EXAMPLE 7.49. The group $F(X)$ is isomorphic to $\pi_1(\bigvee_{x \in X} \mathbb{S}^1)$.

7.5.5. Constructing finite-index subgroups. In this section we use the topological interpretation of graphs of groups in order to construct finite-index subgroups. The main result (Theorem 7.51) will be used in the proof of quasi-isometric rigidity of virtually free groups in Chapter 20.

Let \mathcal{G} be a finite graph of groups. Suppose that we are given a *compatible* collection of finite index subgroups $G'_v < G_v, G'_e < G_e$ for each vertex and edge group of \mathcal{G} , i.e., a collection of subgroups such that whenever $v = e_{\pm}$, we have

$$G_v \cap \phi_{e_{\pm}}(G'_e) = G'_v \cap \phi_{e_{\pm}}(G_e).$$

We refer to this equality as the *compatibility condition*.

THEOREM 7.50. *For every compatible collection of finite-index subgroups as above, there exists a finite-index subgroup $G' < G$ such that*

$$G' \cap G_v = G'_v, \quad G' \cap G_e = G'_e$$

for every vertex v and edge e . Furthermore, $G' = \pi_1(\mathcal{G}')$, where \mathcal{G}' is another finite graph of groups, for which there exists a morphism of graphs of groups

$$p : \mathcal{G}' \rightarrow \mathcal{G}$$

inducing the inclusion $G' \hookrightarrow G$.

PROOF. This theorem is proven by John Hempel in [Hem87] (Theorem 2.2). Our proof mostly follows his arguments.

Let Γ denote the graph underlying \mathcal{G} . For each vertex group G_v (resp. edge group G_e) of \mathcal{G} we let X_v (resp. X_e) denote a classifying space of this group. Then, as in Section 7.5.4, we convert the graph of groups \mathcal{G} into a graph of spaces X , with vertex spaces X_v and edge spaces X_e . We will use the notation

$$f_{e_{\pm}} : X_e \rightarrow X_{e_{\pm}}$$

for the attaching maps inducing the monomorphisms $\phi_{e_{\pm}}$. It will be convenient to assume that distinct attaching maps have disjoint images.

We will construct the subgroup G' as the fundamental group of another graph of spaces X' which admits a finite cover $p : X' \rightarrow X$, such that $G' = p_*(\pi_1(X'))$. The group inclusions $G'_v \rightarrow G_v, G'_e \rightarrow G_e$ are induced by finite covers of spaces

$$X'_v \rightarrow X_v, \quad X'_e \rightarrow X_e.$$

We now assemble the spaces X'_v, X'_e into a finite connected graph of spaces X' . We let d_v, d_e denote the degrees of these covers, i.e.

$$d_v = |G_v : G'_v|.$$

Set

$$d = \prod_{v \in V(\Gamma)} d_v .$$

Now, for each $v \in V(\Gamma)$ we let \tilde{X}_v denote the disjoint union of d/d_v copies of X'_v .

We will use the notation X'_{v_i} for components of \tilde{X}_v .

Our next goal is to describe how to connect components X'_{v_i} to each other via copies of the double mapping cones for the maps $X'_e \rightarrow X'_v$. We then observe that by the definition of \tilde{X}_v and the compatibility assumption, for each edge e with $e_+ = v, e_- = w$, the number of components preimages of $f_{e_+}(X_e)$ in \tilde{X}_v equals the number of components of preimages of $f_{e_-}(X_e)$ in \tilde{X}_w . We therefore, match these subsets of \tilde{X}_v, \tilde{X}_w in pairs. For every such pair, we connect the corresponding vertices v_i, w_j by an edge e_{ij} . This defines a new graph $\tilde{\Gamma}'$ whose vertices are v_i 's and edges are e_{ij} 's, where v runs through the vertex set of Γ . The graph $\tilde{\Gamma}'$ is, a priori, disconnected, we pick a connected component Γ' of this graph.

We then construct a graph of spaces X' based on Γ' as follows. For each vertex v_i , we take, of course, X'_{v_i} as the associated vertex space. For every edge e_{ij} define

$$X'_{e_{ij}} = X'_e$$

where e_{ij} corresponds to the matching of preimages of $f_{e_{\pm}}(X_e)$. Accordingly, we let the map

$$f_{e_{ij+}} : X'_{e_{ij}} \rightarrow X'_{v_i}$$

be the lift of the attaching map $f_{e_+} : X_e \rightarrow X_v$. Note that these lifts exist by the compatibility assumption. We do the same for the vertex e_{ij-} . As the result, we obtain a connected graph of spaces.

We leave it to the reader to verify that the covering maps

$$X'_{v_i} \rightarrow X_v, \quad X'_{e_{ij}} \rightarrow X_e$$

assemble to a covering map $X' \rightarrow X$. This covering map is finite-to-one by the construction. It induces an embedding $G' = \pi_1(X') \rightarrow G = \pi_1(X)$. Again, by construction, this embedding satisfies the requirements of the theorem. \square

As an application we obtain:

THEOREM 7.51. *Let \mathcal{G} be a finite graph of finite groups. Then its fundamental group $G = \pi_1(\mathcal{G})$ is virtually free.*

PROOF. For each vertex group G_v of \mathcal{G} we let $G'_v < G_v$ be the trivial subgroup; we make the same choice for the edge groups. Let $G' < G$ denote the finite-index subgroup and the morphism

$$\mathcal{G}' \rightarrow \mathcal{G}$$

given by Theorem 7.50. By construction, \mathcal{G}' has trivial vertex groups. Hence, for the underlying graph Γ' of \mathcal{G}' we obtain

$$G' = \pi_1(\Gamma')$$

which is free. \square

7.5.6. Graphs of groups and group actions on trees. An *action* of a group G on a tree T is an action $G \curvearrowright T$ such that each element of G acts as an automorphism of T , i.e. such action is a homomorphism $G \rightarrow \text{Aut}(T)$. A tree T with the prescribed action $G \curvearrowright T$ is called a G -tree. An action $G \curvearrowright T$ is said to be *without inversions* if whenever $g \in G$ preserves an edge e of T , it fixes e pointwise. The action is called *bounded* (or *trivial*) if there is a vertex $v \in T$ fixed by the entire group G .

REMARK 7.52. Later on, in Chapter 11, we will encounter more complicated (non-simplicial) *real* trees and group actions on such trees.

Our next goal is to explain the relation between the graph of groups decompositions of G and actions of G on simplicial trees without inversions.

Suppose that $G \cong \pi(\mathcal{G})$ is a graph of groups decomposition of G . We associate with \mathcal{G} a graph of spaces $M = M_{\mathcal{G}}$ as in Section 7.5.4. Let X denote the universal cover of the corresponding cell complex M . Then X is the disjoint union of the copies of the universal covers $\tilde{M}_v, \tilde{M}_e \times (0, 1)$ of the complexes M_v and $M_e \times (0, 1)$. We will refer to this partitioning of X as the *tiling* of X . In other words, X has the structure of a graph of spaces, where each vertex/edge space is homeomorphic to $\tilde{M}_v, v \in V(\Gamma), \tilde{M}_e \times [0, 1], e \in E(\Gamma)$. Let T denote the graph corresponding to X : Each copy of \tilde{M}_v determines a vertex in T and each copy of $\tilde{M}_e \times [0, 1]$ determines an edge in T .

EXAMPLE 7.53. Suppose that Γ consists of two vertices 1 and 2 and the edge $[1, 2]$ connecting them, M_1 and M_2 are surfaces of genus 1 with a single boundary component each. Let M_e be the circle. We assume that the maps $f_{e_{\pm}}$ are homeomorphisms of this circle to the boundary circles of M_1, M_2 . Then M is a surface of genus 2. The graph T is sketched in Figure 7.1.

The Mayer–Vietoris theorem, applied to the above tiling of X , implies that $0 = H_1(X, \mathbb{Z}) \cong H_1(T, \mathbb{Z})$. Therefore, $T = T(\mathcal{G})$ is a tree. The group $G = \pi_1(M)$ acts on X by deck-transformations, preserving the tiling. Thus, we obtain the induced action $G \curvearrowright T$. If $g \in G$ preserves some $\tilde{M}_e \times (0, 1)$, then g comes from the fundamental group of M_e . Therefore, such g also preserves the orientation on the segment $[0, 1]$. Hence, the action $G \curvearrowright T$ is without inversions. Observe that the stabilizer of each \tilde{M}_v in G is conjugate in G to $\pi_1(M_v) = G_v$. Moreover, $T/G = \Gamma$.

EXAMPLE 7.54. Let $G = BS(n, m)$ be the Baumslag–Solitar group described in Example 7.31, (8). The group G clearly has the structure of a graph of groups since it is isomorphic to the HNN extension of \mathbb{Z} ,

$$\mathbb{Z} \star_{H_1 \cong H_2}$$

where the subgroups $H_1, H_2 \subset \mathbb{Z}$ have the indices n and m respectively. In order to construct the cell complex $K(G, 1)$, take the circle $\mathbb{S}^1 = M_v$, the cylinder $\mathbb{S}^1 \times [0, 1]$ and attach the ends to this cylinder to M_v by the maps of the degrees p and q respectively. Now, consider the associated G -tree T . Its vertices have valence $n+m$: Each vertex v has m incoming and n outgoing edges so that for each outgoing edge e we have $v = e_-$ and for each incoming edge we have $v = e_+$. The vertex stabilizer $G_v \cong \mathbb{Z}$ permutes (transitively) incoming and outgoing edges among each other. The stabilizer of each outgoing edge is the subgroup H_1 and the stabilizer of each incoming edge is the subgroup H_2 . Thus, the action of \mathbb{Z} on the set of incoming edges is *via* the group \mathbb{Z}/m and on the set of outgoing edges *via* the group \mathbb{Z}/n .




FIGURE 7.1. Universal cover of the genus 2 surface.

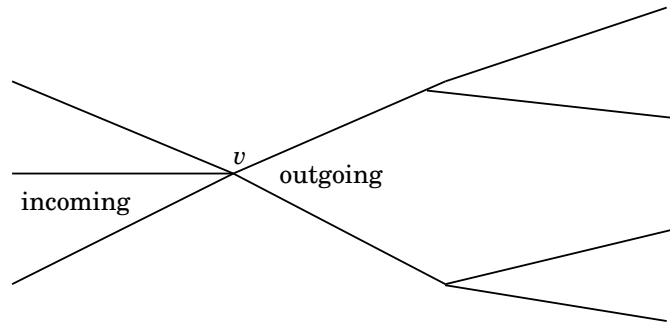


FIGURE 7.2. Tree for the group $BS(2, 3)$.

LEMMA 7.55. *The action $G \curvearrowright T$ is bounded if and only if the graph of groups decomposition of G is trivial.*

PROOF. Suppose that G fixes a vertex $\tilde{v} \in T$. Then $\pi_1(M_v) = G_v = G$, where $v \in \Gamma$ is the projection of \tilde{v} . Hence, the decomposition of G is trivial. Conversely, suppose that G_v maps onto G . Let $\tilde{v} \in T$ be the vertex which projects to v . Then $\pi_1(M_v)$ is the entire $\pi_1(M)$ and, hence, G preserves $\tilde{M}_{\tilde{v}}$. Therefore, the group G fixes \tilde{v} . \square

Conversely, each action of G on a simplicial tree T yields a realization of G as the fundamental group of a graph of groups \mathcal{G} , such that $T = T(\mathcal{G})$. Here is the construction of \mathcal{G} . Furthermore, an *unbounded* action leads to a *non-trivial* graph of groups.

If the action $G \curvearrowright T$ has inversions, we replace T with its barycentric subdivision T' . Then G acts on T' without inversions. If the action $G \curvearrowright T$ is unbounded, so is $G \curvearrowright T'$. Thus, from now on, we assume that G acts on T without inversions. **Then** the quotient T/G is a graph Γ : $V(\Gamma) = V(T)/G$ and $E(\Gamma) = E(T)/G$. For every vertex \tilde{v} and edge \tilde{e} of T we let $G_{\tilde{v}}$ and $G_{\tilde{e}}$ be their respective stabilizes in G . Clearly, whenever $\tilde{e} = [\tilde{v}, \tilde{w}]$, we get the embedding

$$G_{\tilde{e}} \rightarrow G_{\tilde{v}}.$$

If $g \in G$ maps oriented the edge $\tilde{e} = [\tilde{v}, \tilde{w}]$ to an oriented edge $\tilde{e}' = [\tilde{v}', \tilde{w}']$, we obtain isomorphisms

$$G_{\tilde{v}} \rightarrow G_{\tilde{v}'}, \quad G_{\tilde{w}} \rightarrow G_{\tilde{w}'}, \quad G_{\tilde{e}} \rightarrow G_{\tilde{e}'}$$

induced by conjugation *via* g and the following diagram is commutative:

$$\begin{array}{ccc} G_{\tilde{e}} & \longrightarrow & G_{\tilde{v}} \\ \downarrow & & \downarrow \\ G_{\tilde{e}'} & \longrightarrow & G_{\tilde{v}'} \end{array}$$

We set $G_v := G_{\tilde{v}}$, $G_e := G_{\tilde{e}}$, where v and e are the projections of \tilde{v} and edge \tilde{e} to Γ . For every edge e of Γ oriented as $e = [v, w]$, we define the monomorphism $G_e \rightarrow G_v$ as follows. By applying an appropriate element $g \in G$ as above, we can assume that $\tilde{e} = [\tilde{v}, \tilde{w}]$. We then define the embedding $G_e \rightarrow G_v$ to make the diagram

$$\begin{array}{ccc} G_{\tilde{e}} & \longrightarrow & G_{\tilde{v}} \\ \downarrow & & \downarrow \\ G_e & \longrightarrow & G_v \end{array}$$

commutative. The result is a graph of groups \mathcal{G} . We leave it to the reader to verify that the functor $(G \curvearrowright T) \rightarrow \mathcal{G}$ described above is the inverse of the functor $\mathcal{G} \rightarrow (G \curvearrowright T)$ for \mathcal{G} with $G = \pi_1(\mathcal{G})$. In particular, \mathcal{G} is trivial if and only if the action $G \curvearrowright T$ is bounded.

DEFINITION 7.56. $\mathcal{G} \rightarrow (G \curvearrowright T) \rightarrow \mathcal{G}$ is the *Bass–Serre correspondence* between realizations of groups as fundamental groups of graphs of groups and group actions on trees without inversions.

We refer the reader to [SW79] and [Ser80] for further details on the Bass–Serre correspondence. Below is a simple, yet non-obvious, example of application of this correspondence:

LEMMA 7.57. *Suppose that G is countable, but not finitely generated. Then G admits a non-trivial action on a simplicial tree.*

PROOF. Using countability of G , enumerate the elements of the group G and define an exhaustion of G by finitely generated subgroups:

$$G_1 \leq G_2 \leq G_3 \leq \dots$$

where $G_{n+1} = \langle G_n, g_{n+1} \rangle$. The inclusion homomorphisms

$$\iota_n : G_n \hookrightarrow G_{n+1}$$

determine an infinite graph of groups, where vertices are labeled $v_n, n \in \mathbb{N}$, the vertex groups are $G_{v_n} = G_n$, the edge groups are

$$G_{e_n} = G_n, \quad e_n = [v_n, v_{n+1}]$$

and the map $G_{e_n} \rightarrow G_{v_n}$ is the identity, while the map $G_{e_n} \rightarrow G_{v_{n+1}}$ is ι_n . We claim that the fundamental group $\pi_1(\mathcal{G})$ of this graph of groups is G itself. Indeed, we have natural inclusion homomorphisms

$$f_n : G_n \rightarrow G.$$

If H is a group and $h_n : G_n \rightarrow H$ are homomorphisms, such that

$$h_{n+1}|_{G_n} = h_n,$$

then h_n 's determine a homomorphism $h : G \rightarrow H$ by $h(g) = h_n(g)$ whenever $g \in G_n$. Uniqueness of h is also clear. Thus, G satisfies the universality property in the definition of $\pi_1(\mathcal{G})$ and, hence, $G \cong \pi_1(\mathcal{G})$.

Next, none of the vertex groups G_n maps onto G via the inclusion homomorphism ι_n . Therefore, the action $G \curvearrowright T$ of G on a simplicial tree, defined by the Bass-Serre correspondence, is non-trivial and without inversions. \square

7.6. Ping-pong lemma. Examples of free groups

The ping-pong lemma is a simple, yet powerful, tool for constructing free groups acting on sets. We will see in Chapter 15 how ping-pong is used for the proof of the *Tits Alternative*.

We begin with the *ping-lemma*, a version of the ping-pong lemma for semigroups:

LEMMA 7.58 (Ping-pong for semigroups). *Let X be a set, and let $g : X \rightarrow X$ and $h : X \rightarrow X$ be two injective maps. Suppose that $A \subset X$ is a non-empty subset such that $g(A), h(A)$ are disjoint subsets of A . Then g, h generate a free subsemigroup of rank 2 in the semigroup of self-maps $X \rightarrow X$. Moreover, for two distinct words w, w' in the generators g, h ,*

$$w(A) \cap w'(A) = \emptyset.$$

PROOF. Let w, w' be distinct non-empty words in the alphabet g, h . We claim that $w(A) \cap w'(A) = \emptyset$. We prove this by induction on the maximum of lengths $\ell(w), \ell(w')$ of w, w' . If both w, w' have unit length the claim is immediate. Suppose that the claim holds for all words w, w' such that $\max(\ell(w), \ell(w')) \leq n$. Let w, w' be distinct non-empty words in g, h such that $\ell(w) \leq \ell(w') = n + 1$. The words w, w' either have the same first letter (the prefix), or distinct prefixes. Suppose first that w, w' have the same prefix $x \in \{g, h\}$; then

$$w = xu \quad w' = xu', \quad y \neq y', \quad \max(\ell(u), \ell(u')) \leq n.$$

Then, by the induction hypothesis,

$$u(A) \cap u'(A) = \emptyset.$$

Injectivity of x implies that the sets $w(A) = xu(A)$ and $w'(A) = xu'(A)$ are also disjoint, as claimed. Suppose, next, that w, w' have distinct prefixes:

$$w = xu \quad w' = x'u', \quad \{x, x'\} = \{g, h\}.$$

Then $w(A) \subset x(A), w'(A) \subset x'(A)$ are disjoint and the claim follows. \square

EXERCISE 7.59. Suppose that $g \in \text{Bij}(X)$ is a bijection such that for some $A \subset X$,

$$g(A) \subsetneq A.$$

Then g has infinite order.

We next consider ping-pong for groups of bijections. The setup for the ping-pong lemma is a pair of bijections $g_1, g_2 \in \text{Bij}(X)$ (“ping-pong partners”) and a quadruple of non-empty subsets

$$B_i^\pm \subset X, \quad i = 1, 2.$$

Define

$$C_i^+ := B_i^+ \cup B_j^- \cup B_j^+, C_i^- := B_i^- \cup B_j^- \cup B_j^+ \quad \{i, j\} = \{1, 2\}.$$

As it was written before, with B the union of the four sets and $C_i^+ = B \setminus B_i^-$ it was wrong, for the argument to work we need C_i^+ to contain the whole of $B_j^- \cup B_j^+$. In particular, in the example below, the matrices with $r = 1$ would have satisfied the condition with $C_i^+ = B \setminus B_i^-$. We require that:

$$C_i^\pm \not\subset B_j^\pm \text{ and } C_i^\pm \not\subset B_j^\mp \text{ for all choices of } i, j \text{ and } +, -.$$

Typically, this is achieved by assuming that all the four sets B_1^\pm, B_2^\pm are pairwise disjoint and non-empty.

LEMMA 7.60 (Ping-pong, or table-tennis, lemma). *Let X, B_i^\pm, C_i^\pm be as above, and suppose that*

$$g_i^{\pm 1}(C_i^\pm) \subset B_i^\pm, \quad i = 1, 2.$$

Then the bijections g_1, g_2 generate a rank 2 free subgroup of $\text{Bij}(X)$.

PROOF. Let w be a non-empty reduced word in $\{g, g^{-1}, h, h^{-1}\}$. In order to prove that w corresponds to a non-identity element of $\text{Bij}(X)$, it suffices to check that $w(C_j^\pm) \subset B_i^\pm$ for some i, j and for some choice of $+$ or $-$. We claim that whenever w has the form

$$w = g_i^{\pm 1} u g_j^{\pm 1},$$

we have

$$w(C_j^\pm) \subset B_i^\pm.$$

This would immediately imply that w does not represent the identity map $X \rightarrow X$. The claim is proven by induction on the length $\ell(w)$ of w as in the proof of Lemma 7.58. The statement is clear if $\ell(w) = 1$. Suppose it holds for all words w' of length n , we will prove it for words w of length $n + 1$. Such w has the form

$$w = g_i^{\pm 1} w', \quad \ell(w') = n.$$

Since the prefix of w' cannot equal $g_i^{\mp 1}$ (as w is a reduced word), it follows from the induction hypothesis that (for some j and a choice of $+, -$)

$$w'(C_j^\pm) \subset C_i^\pm.$$

Since

$$g_i^{\pm 1} w'(C_j^\pm) \subset g_i^{\pm 1} (C_i^\pm) \subset B_i^\pm,$$

the claim follows. □

Lemma 7.60 extends to the case of free products of subgroups. The setup for this extension is a collection $\{G_i : i \in I\}$ of subgroups of $\text{Bij}(X)$, and of subsets $A_i \subset X$ ($i \in I$), whose union is denoted

$$A = \bigcup_{i \in I} A_i.$$

For each A_i define $A_i^c = A \setminus A_i$.

LEMMA 7.61 (The ping-pong lemma for free products). *Given the above data, suppose that:*

(1) *For each pair $i, j \in I$,*

$$A_i^c \not\subset A_j.$$

(2) *For each $i \in I$ and all $g \in G_i \setminus \{1\}$, we have the inclusion*

$$g(A_i^c) \subset A_i.$$

Then the natural homomorphism

$$\phi : \star_{i \in I} G_i \rightarrow \text{Bij}(X), \quad \phi|_{G_i} = \text{Id}_{G_i}, i \in I,$$

is a monomorphism.

PROOF. Consider a non-trivial word w in the alphabet

$$\bigcup_{i \in I} G_i,$$

where no two consecutive letters belong to the same G_k . Suppose that w has the prefix $g_i \in G_i \setminus \{1\}$ and the suffix $g_j \in G_j \setminus \{1\}$. We claim that

$$w(A_j^c) \subset A_i.$$

The proof is the induction on the length $\ell(w)$ of w . The claim is clear for $\ell(w) = 1$. Suppose that the claim holds for all words w' of the length n and let w be a word of the length $n + 1$. Then w has the form

$$w = g_i w', \quad \ell(w') = n,$$

where the suffix of w' is $g_j \in G_j$. Since the prefix of w' cannot equal to an element of G_i , it follows from the induction hypothesis that

$$w'(A_j^c) \subset A_i.$$

Hence, $w(A_j^c) \subset g_i(A_i^c) \subset A_i$. Since $A_j^c \not\subset A_i$, we conclude that $w(A_j^c) \neq A_j^c$ and, hence, $w \neq \text{Id}$. It follows that the homomorphism ϕ is injective. \square

In the following example we illustrate both forms of ping-pong.

EXAMPLE 7.62. For any real number $r \geq 2$ the matrices

$$g_1 = \begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix} \text{ and } g_2 = \begin{pmatrix} 1 & 0 \\ r & 1 \end{pmatrix}$$

generate a free subgroup of $SL(2, \mathbb{R})$.

First proof. The group $SL(2, \mathbb{R})$ acts (with the kernel $\pm I$) on the upper half plane $\mathbb{H}^2 = \{z \in \mathbb{C} \mid \Im(z) > 0\}$ by linear fractional transformations

$$z \mapsto \frac{az + b}{cz + d}.$$

Define quater-planes

$$B_1^+ = \{z \in \mathbb{H}^2 : \Re(z) > r/2, \quad B_1^- = \{z \in \mathbb{H}^2 : \Re(z) < -r/2\}$$

and open disks

$$B_2^+ := \{z \in \mathbb{H}^2 : |z - \frac{1}{r}| < \frac{1}{r}\}, \quad B_2^- := \{z \in \mathbb{H}^2 : |z + \frac{1}{r}| < \frac{1}{r}\}.$$

The reader will verify that $g_k, B_k^\pm, k = 1, 2$ satisfy the assumptions of Lemma 7.60. It follows that the group $\langle g_1, g_2 \rangle$ is free of rank 2.




FIGURE 7.3. Example of ping-pong.

Second proof. The group $SL(2, \mathbb{R})$ also acts linearly on \mathbb{R}^2 . Consider the infinite cyclic subgroups $G_k = \langle g_k \rangle, i = 1, 2$ of $SL(2, \mathbb{R})$. Define the following subsets of \mathbb{R}^2

$$A_1 = \left\{ \begin{pmatrix} x \\ y \end{pmatrix} : |x| > |y| \right\} \text{ and } A_2 = \left\{ \begin{pmatrix} x \\ y \end{pmatrix} : |x| < |y| \right\}.$$

Then for each $g \in G_1 \setminus \{1\}$, $g(A_2) \subset A_1$ and for each $g \in G_2 \setminus \{1\}$, $g(A_1) \subset A_2$. Therefore, the subgroup of $SL(2, \mathbb{R})$ generated by g_1, g_2 is free of rank 2 according to Lemma 7.61. \square

REMARK 7.63. The statement in the Example 7.62 no longer holds for $r = 1$. Indeed, in this case we have

$$g_1^{-1} g_2 g_1^{-1} = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

Thus, $(g_1^{-1} g_2 g_1^{-1})^2 = I$, and, hence, the group generated by g_1, g_2 is not free.

7.7. Free subgroups in $SU(2)$

As an application of ping-pong in $SL(2, \mathbb{R})$ and the formalism of algebraic groups, we will now give a “cheap” proof of the fact that the group $SU(2)$ contains a subgroup isomorphic to F_2 , the free group on two generators:

LEMMA 7.64. *The subset of monomorphisms $F_2 \rightarrow SU(2)$ is dense (with respect to the classical topology) in the variety $\text{Hom}(F_2, SU(2)) = SU(2) \times SU(2)$.*

PROOF. Consider the space $V = \text{Hom}(F_2, SL(2, \mathbb{C})) = SL(2, \mathbb{C}) \times SL(2, \mathbb{C})$; every element $w \in F_2$ defines a polynomial function

$$f_w : V \rightarrow SL(2, \mathbb{C}), \quad f_w(\rho) = \rho(w).$$

Since $SL(2, \mathbb{R}) \leq SL(2, \mathbb{C})$ contains a subgroup isomorphic to F_2 (see Example 7.62), it follows that for every $w \neq 1$, the function f_w takes values different from 1. In particular, the subset $E_w := f_w^{-1}(1)$ is a proper (complex) subvariety in V . Since $SL(2, \mathbb{C})$ is a connected complex manifold, the variety $SL(2, \mathbb{C})$ is irreducible; hence, V is irreducible as well. It follows that for every $w \neq 1$, E_w has empty interior (in the classical topology) in V . Suppose that for some $w \neq 1$, the intersection

$$E'_w := E_w \cap SU(2) \times SU(2)$$

contains a non-empty open subset U . In view of Exercise 5.51, $SU(2)$ is Zariski dense (over \mathbb{C}) in $SL(2, \mathbb{C})$; hence, U (and, thus, E_w) is Zariski dense in V . It then follows that $E_w = V$, which is false. Therefore, for every $w \neq 1$, the closed (in the classical topology) subset $E'_w \subset \text{Hom}(F_2, SU(2))$ has empty interior. Since F_2 is countable, by Baire category theorem, the union

$$E := \bigcup_{w \neq 1} E'_w$$

has empty interior in $\text{Hom}(F_2, SU(2))$. Since every $\rho \notin E$ is injective, lemma follows. \square

Since the group $SU(2)/\{\pm I\}$ is isomorphic to $SO(3)$, we also obtain:

COROLLARY 7.65. *The subset of monomorphisms $F_2 \rightarrow SO(3)$ is dense in the variety $\text{Hom}(F_2, SO(3))$.*

7.8. Ping-pong on projective spaces

We will frequently use the Ping-pong lemma in the case when X is a projective space. This application of the ping-pong argument is the key for the proof of the Tits Alternative.

Let V be an n -dimensional space over a local field \mathbb{K} , the reader should think of \mathbb{R} , \mathbb{C} , or \mathbb{Q}_p . We endow the projective space $P(V)$ with the metric d as in Section 2.9. We refer the reader to Section 2.10 for the notion of *proximity*, attractive points $A_g \in P(V)$ and exceptional hyperplanes $E_g \subset P(V)$ for proximal projective transformations.

DEFINITION 7.66. Two proximal elements $g, h \in GL(V)$ will be called *ping partners* if

$$A_g \notin E_h, \quad A_h \notin E_g.$$

Two very proximal elements $g, h \in GL(V)$ will be called *ping-pong partners* if all four pairs $(g, h), (g, h^{-1}), (g^{-1}, h)$ and (g^{-1}, h^{-1}) are ping-partners. In particular, the four points $A_g, A_{g^{-1}}, A_h, A_{h^{-1}}$ are all distinct.

For instance, if $n = 2$, then g, h are ping-pong partners if and only if they are both proximal and their four fixed points in the projective line $P(V)$ are pairwise distinct.

LEMMA 7.67. Assume that $g, h \in GL(n, \mathbb{K})$ are ping partners. Then there exists a positive integer N such that for all $m \geq N$, the powers g^m and h^m generate a rank two free subsemigroup of $GL(n, \mathbb{K})$. Similarly, if g, h are ping-pong partners, then there exists N such that for all $m \geq N$, g^m and h^m generate a rank two free subgroup of $GL(n, \mathbb{K})$.

PROOF. We prove the statement about ping-pong partners, since its proof will contain the proof in the case of ping-partners. Define

$$\varepsilon = \frac{1}{2} \min (\text{dist}(A_g, H(g) \cup E_h \cup E_{h^{-1}}), \text{dist}(A_{g^{-1}}, E_{g^{-1}} \cup E_h \cup E_{h^{-1}}), \\ \text{dist}(A_h, E_h \cup H(g) \cup E_{g^{-1}}), \text{dist}(A_{h^{-1}}, E_{h^{-1}} \cup H(g) \cup E_{g^{-1}})).$$

Since g, h are ping-pong partners, $\varepsilon > 0$. Next, by Corollary 2.87, there exists N such that for all $m \geq N$ we have:

1. $g^{\pm m} : P(V) \rightarrow P(V)$ maps the complement of the ε -neighborhood of $E_{g^{\pm 1}}$ inside the ball of radius ε and center $A_{g^{\pm 1}}$.
2. $h^{\pm m}$ maps the complement of the ε -neighborhood of $E_{h^{\pm 1}}$ inside the ball of radius ε and center $A_{h^{\pm 1}}$.

Set

$$A := B(A_g, \varepsilon) \sqcup B(A_{g^{-1}}, \varepsilon)$$

and

$$B := B(A_h, \varepsilon) \sqcup B(A_{h^{-1}}, \varepsilon).$$

Clearly,

$$g^{km}(A) \subseteq B$$

and

$$h^{km}(B) \subseteq A$$

for every $k \in \mathbb{Z} \setminus \{0\}$. Hence, by Lemma 7.61, regarded as projective transformations, g^m and h^m generate a free subgroup of rank 2 in $PGL(n, \mathbb{K})$. Therefore, the same holds for $g^m, h^m \in GL(n, \mathbb{K})$, see Lemma 7.23. \square

7.9. Cayley graphs

One of the central themes of Geometric Group Theory is treating groups as geometric objects. The oldest, and most common, way to ‘geometrize’ groups, by their *Cayley graphs*. Other ‘geometrizations’ of groups are given by simplicial complexes and Riemannian manifolds.

Every group may be turned into a geometric object (a graph) as follows. Given a group G and its generating set S , one defines the *Cayley graph* of G with respect to S . This is a directed graph $\text{Cayley}_{\text{dir}}(G, S)$ such that

- its set of vertices is G ;
- its set of oriented edges is (g, gs) , with $s \in S$.

Usually, the underlying non-oriented graph $\text{Cayley}(G, S)$ of $\text{Cayley}_{\text{dir}}(G, S)$, i.e. the graph such that:

- its set of vertices is G ;
- its set of edges consists of all pairs of elements in G , $\{g, h\}$, such that $h = gs$, with $s \in S$,

is also called the *Cayley graph of G with respect to S* .

We will also denote the notation \overline{gh} and $[g, h]$ for the edge $\{g, h\}$. In order to avoid the confusion with the notation for the commutator of the elements g and h we will always add the word *edge* in this situation.

EXERCISE 7.68. Show that the graph $\text{Cayley}(G, S)$ is connected.

One can attach a *color (label)* from S to each oriented edge in $\text{Cayley}_{\text{dir}}(G, S)$: the edge (g, gs) is labeled by s .

We endow the graph $\text{Cayley}(G, S)$ with the standard length metric (where every edge has unit length). The restriction of this metric to G is called *the word metric associated to S* and it is denoted by dist_S or d_S .

NOTATION 7.69. For an element $g \in G$ and a generating set S we denote $\text{dist}_S(1, g)$ by $|g|_S$, the *word norm* of g . With this notation, $\text{dist}_S(g, h) = |g^{-1}h|_S = |h^{-1}g|_S$.

CONVENTION 7.70. In this book, unless stated otherwise, all Cayley graphs are defined for finite generating sets S .

Much of the discussion in this section, though, remains valid for arbitrary generating sets, including infinite ones.

REMARK 7.71. 1. Every group acts on itself, on the left, by the left multiplication:

$$G \times G \rightarrow G, (g, h) \mapsto gh.$$

This action extends to any Cayley graph: if $[x, xs]$ is an edge of $\text{Cayley}(G, S)$ with the vertices x, xs , we extend g to the isometry

$$g : [x, xs] \rightarrow [gx, gxs]$$

between the unit intervals. Both actions $G \curvearrowright G$ and $G \curvearrowright \text{Cayley}(G, S)$ are by isometries. It is also clear that the action on G is free, while the action on $\text{Cayley}(G, S)$ is free if and only if none of the generators is of order two. Both actions are properly discontinuous and cocompact (provided that S is finite): The quotient $\text{Cayley}(G, S)/G$ is homeomorphic to the bouquet of n circles, where n is the cardinality of S , if no $s \in S$ satisfies $s^2 = 1$; while in the opposite case, for each generator of order two the corresponding circle must be replaced by an interval of length $\frac{1}{2}$ with one endpoint the basepoint of the bouquet.

2. The action of the group on itself by right multiplication defines maps

$$R_g : G \rightarrow G, R_g(h) = hg$$

that are, in general, not isometries with respect to a word metric, but are at finite distance from the identity map:

$$\text{dist}(\text{id}(h), R_g(h)) = |g|_S.$$

EXERCISE 7.72. Prove that the word metric on a group G associated to a generating set S may also be defined

(1) either as the unique maximal left-invariant metric on G such that

$$\text{dist}(1, s) = \text{dist}(1, s^{-1}) = 1, \forall s \in S;$$

(2) or by the following formula: $\text{dist}(g, h)$ is the length of the shortest word w in the alphabet $S \cup S^{-1}$ such that $w = g^{-1}h$ in G .

Below are two simple examples of Cayley graphs.

EXAMPLE 7.73. Consider the group \mathbb{Z}^2 with the set of generators

$$S = \{a = (1, 0), b = (0, 1)\}.$$

The Cayley graph $\text{Cayley}(G, S)$ is the square grid in the Euclidean plane: The vertices are points with integer coordinates, two vertices are connected by an edge if and only if either their first or their second coordinates differ by ± 1 . See Figure 7.4.




FIGURE 7.4. The Cayley graph of \mathbb{Z}^2 .

The Cayley graph of \mathbb{Z}^2 with respect to the generating set $\{(1, 0), (1, 1)\}$ has the same set of vertices as the above, but the vertical lines are replaced by diagonal lines.

EXAMPLE 7.74. Let G be the free group on two generators a, b . Take $S = \{a, b\}$. The Cayley graph $\text{Cayley}(G, S)$ is the 4-valent tree (there are four edges incident to each vertex). See Figure 7.5.

EXERCISE 7.75. 1. Show that every simplicial tree is contractible and, hence, simply-connected.

2. Show that, conversely, every simply-connected graph is a simplicial tree. (Hint: Verify that if a connected graph Γ is not a tree then $H_1(\Gamma) \neq 0$.)

THEOREM 7.76. *The fundamental group of every connected graph Γ is free.*

PROOF. By the Axiom of Choice (Zorn Lemma), Γ contains a maximal subtree $\Lambda \subset \Gamma$. Let Γ' denote the subdivision of Γ where every edge e in $\mathcal{E} = E(\Gamma) \setminus E(\Lambda)$ is subdivided in 3 sub-edges. For every such edge e let e' denote the middle 3rd. Now, add to Λ all the edges in $E(\Gamma')$ which are not of the form e' ($e \in \mathcal{E}$), and the vertices of such edges, of course, and let T' denote the resulting tree. Thus, we obtain a covering of Γ' by the simplicial tree T' and the subgraph $\Gamma_{\mathcal{E}}$ consisting of the pairwise disjoint edges e' ($e \in \mathcal{E}$), and the incident vertices. To this covering we can now apply Seifert—van Kampen Theorem and conclude (in view of the




FIGURE 7.5. The Cayley graph of the free group F_2 .

fact that T' is simply-connected) that $G = \pi_1(\Gamma)$ is free, with the free generators indexed by the set \mathcal{E} . \square

COROLLARY 7.77. 1. Every free group $F(X)$ is the fundamental group of the bouquet (wedge) B of $|X|$ circles. 2. The universal cover of B is a tree T , which is isomorphic to the Cayley graph of $F(X)$ with respect to the generating set X .

PROOF. 1. By Theorem 7.76, $G = \pi_1(B)$ is free; furthermore, the proof also shows that the generating set of G is identified with the set of edges of B . We now orient every edge of B using this identification. 2. The universal cover T of B is a simply-connected graph, hence, a tree. We lift the orientation of edges of B to orientation of edges of T . The group $F(X) = \pi_1(B)$ acts on T by covering transformations, hence, the action on the vertex $V(T)$ set of T is simply-transitive. Therefore, we obtain an identification of $V(T)$ with G . Let v be a vertex of T . By construction and the standard identification of $\pi_1(B)$ with covering transformations of T , every oriented edge e of B lifts to an oriented edge \tilde{e} of T of the form $[v, w]$. Conversely, every oriented edge $[v, w]$ of T projects to an oriented edge of B . Thus, we labeled all the oriented edges of T with generators of $F(X)$. Again, by the covering theory, if an oriented edge $[u, w]$ of T is labeled with a generator $x \in F(X)$, then x sends u to w . Thus, T is isomorphic to the Cayley graph of $F(X)$. \square

COROLLARY 7.78. A group G is free if and only if it can act freely by automorphisms on a simplicial tree T .

PROOF. By the covering theory, $G \cong \pi_1(\Gamma)$ where $\Gamma = T/G$. Now, by Theorem 7.76, $G = \pi_1(\Gamma)$ is free. See [Ser80] for another proof and the more general discussion of group actions on trees. \square

The concept of a simplicial tree generalizes to the one of a *real tree* (see Definition 3.60). There are non-free groups acting isometrically and freely on real trees, e.g., surface groups and free abelian groups. Rips proved that every finitely generated group acting freely and isometrically on a real tree is a free product of surface groups and free abelian groups, see e.g. [BF95, Kap01, CR13].

As an immediate application of Corollary 7.78 we obtain:

COROLLARY 7.79 (Nielsen–Schreier). *Every subgroup H of a free group F is itself free.*

PROOF. Realize the free group F as the fundamental group of a bouquet B of circles; the universal cover T of B is a simplicial tree. The subgroup $H \leq F$ also acts on T freely. Thus, H is free. \square

PROPOSITION 7.80. *The free group of rank 2 contains an isomorphic copy of F_m for every finite m and for $m = \aleph_0$. Moreover, for finite m , we can find a subgroup $F_m < F_2$ of finite index.*

PROOF. Let x, y denote the free generators of the group F_2 .

1. Define the epimorphism $\rho_m : F_2 \rightarrow \mathbb{Z}_m$ by sending x to 1 and y to 0. Then the kernel K_m of ρ_m has index m in F_2 . Then K_m is a finitely generated free group F . In order to compute the rank of F , it is convenient to argue topologically. Let R be a finite graph with the (free) fundamental group $\pi_1(R)$. Then $\chi(R) = 1 - b_1(R) = 1 - \text{rank}(\pi_1(R))$. Let R_2 be such a graph for F_2 , then $\chi(R_2) = 1 - 2 = -1$. Let $R \rightarrow R_2$ be the m -fold covering corresponding to the inclusion $K_m \hookrightarrow F_2$. Then $\chi(R) = m\chi(R_2) = -m$. Hence, $\text{rank}(K_m) = 1 - \chi(R) = 1 + m$. Thus, for every $n = 1 + m \geq 2$, we have a finite-index inclusion $F_n \hookrightarrow F_2$.

2. Let x, y be the two generators of F_2 . Let S be the subset consisting of all elements of F_2 of the form $x_k := y^k x y^{-k}$, for all $k \in \mathbb{N}$. We claim that the subgroup $\langle S \rangle$ generated by S is isomorphic to the free group of rank \aleph_0 .

Indeed, consider the set A_k of all reduced words with prefix $y^k x$. With the notation of Section 7.2, the transformation $L_{x_k} : F_2 \rightarrow F_2$ has the property that $L_{x_k}(A_j) \subset A_k$ for every $j \neq k$. Obviously, the sets A_k , $k \in \mathbb{N}$, are pairwise disjoint. This and Lemma 7.61, imply that $\{L_{x_k} : k \in \mathbb{N}\}$ generate a free subgroup in $\text{Bij}(F_2)$, hence so do $\{x_k : k \in \mathbb{N}\}$ in F_2 . \square

EXERCISE 7.81. Let G and H be finitely generated groups, with S and X respective finite generating sets. Consider the wreath product $G \wr H$ (see Definition 5.32), endowed with the finite generating set canonically associated to S and X described in Lemma 7.11. For every function $f : H \rightarrow G$ denote by $\text{supp } f$ the set of elements $h \in H$ such that $f(h) \neq \mathbf{1}_G$.

Let f and g be arbitrary functions from H to G with finite support, and h, k arbitrary elements in H . Prove that the word distance in $G \wr H$ from (f, h) to (g, k) with respect to the generating set mentioned above is

$$(7.4) \quad \text{dist}((f, h), (g, k)) = \sum_{x \in H} \text{dist}_S(f(x), g(x)) + \text{length}(\text{supp}(g^{-1}f); h, k),$$

where

$$\text{length}(\text{supp}(g^{-1}f); h, k)$$

is the length of the shortest path in $\text{Cayley}(H, X)$ starting in h , ending in k and whose image contains the set $\text{supp}(g^{-1}f)$.

Thus, we succeeded in assigning to every finitely generated group G a metric space $\text{Cayley}(G, S)$. The problem, however, is that this assignment

$$G \rightarrow \text{Cayley}(G, S)$$

is far from canonical: Different generating sets could yield completely different Cayley graphs. For instance, the trivial group has the presentations:

$$\langle \quad | \quad \rangle, \quad \langle a|a \rangle, \quad \langle a, b|ab, ab^2 \rangle, \dots,$$

which give rise to the non-isometric Cayley graphs:



FIGURE 7.6. Cayley graphs of the trivial group.

The same applies to the infinite cyclic group:




FIGURE 7.7. Cayley graphs of $\mathbb{Z} = \langle x \rangle$ and $\mathbb{Z} = \langle x, y | xy^{-1} \rangle$.

In the above examples we did not follow the convention that $S = S^{-1}$.

Note, however, that all Cayley graphs of the trivial group have finite diameter; the same, of course, applies to all finite groups. The Cayley graphs of \mathbb{Z} as above, although they are clearly non-isometric, are within finite distance from each other (when placed in the same Euclidean plane). Therefore, when seen from a (very) large distance (or by a person with a very poor vision), every Cayley graph of a finite group looks like a “fuzzy dot”; every Cayley graph of \mathbb{Z} looks like a “fuzzy line,” etc. Therefore, although non-isometric, they all “look alike”.

EXERCISE 7.82. (1) Prove that if S and \bar{S} are two finite generating sets of G , then the word metrics dist_S and $\text{dist}_{\bar{S}}$ on G are bi-Lipschitz equivalent, i.e. there exists $L > 0$ such that

$$(7.5) \quad \frac{1}{L} \text{dist}_S(g, g') \leq \text{dist}_{\bar{S}}(g, g') \leq L \text{dist}_S(g, g'), \forall g, g' \in G.$$

Hint: Verify the inequality (7.5) first for $g' = 1_G$ and $g \in S$; then verify the inequality for arbitrary $g \in G$ and $g' = 1_G$. Lastly, verify the inequality for all g, g' using left-invariance of word-metrics.

(2) Prove that an isomorphism between two finitely generated groups is a bi-Lipschitz map when the two groups are endowed with word metrics.

CONVENTION 7.83. From now on, unless otherwise stated, by a metric on a finitely generated group we mean a word metric coming from a finite generating set.

EXERCISE 7.84. Show that the Cayley graph of a finitely generated infinite group contains an isometric copy of \mathbb{R} , i.e. a bi-infinite geodesic. Hint: Apply Arzela-Ascoli theorem to a sequence of geodesic segments in the Cayley graph.

On the other hand, it is clear that no matter how poor one's vision is, the Cayley graphs of, say, $\{1\}$, \mathbb{Z} and \mathbb{Z}^2 all look different: They appear to have different “dimension” (0, 1 and 2 respectively).

Telling apart the Cayley graph Cayley_1 of \mathbb{Z}^2 from the Cayley graph Cayley_2 of the Coxeter group

$$\Delta := \Delta(4, 4, 4) := \langle a, b, c \mid a^2, b^2, c^2, (ab)^4, (bc)^4, (ca)^4 \rangle$$

seems more difficult: They both “appear” 2-dimensional. However, by looking at the larger pieces of Cayley_1 and Cayley_2 , the difference becomes more apparent: Within a given ball of radius R in Cayley_1 , there seems to be less vertices than in Cayley_2 . The former grows quadratically, while the latter grows exponentially fast as R goes to infinity.




FIGURE 7.8. Dual simplicial complex of the Cayley graph of Δ .

The goal of the rest of the book is to make sense of this “fuzzy math”.

In Section 8.1 we replace the notion of an *isometry* with the notion of a *quasi-isometry*, in order to capture what different Cayley graphs of the same group have in common.

LEMMA 7.85. *A finite index subgroup of a finitely generated group is finitely generated.*

PROOF. This lemma follows from Theorem 8.37 proven in the next chapter. We give here another proof, as the set of generators of the subgroup found here will be used in future applications.

Let G be a group and S a finite generating set of G , and let H be a finite-index subgroup in G . Then

$$G = H \sqcup \bigsqcup_{i=1}^k Hg_i$$

for some elements $g_i \in G$. Consider

$$R = \max_{1 \leq i \leq k} |g_i|_S.$$

Then $G = HB(1, R)$. We now prove that $X = H \cap B(1, 2R + 1)$ is a generating set of H .

Let h be an arbitrary element in H and let $g_0 = 1, g_1, \dots, g_n = h$ be the consecutive vertices on a geodesic in $\text{Cayley}(G, S)$ joining 1 and h . In particular, this implies that $\text{dist}_S(1, h) = n$.

For every $1 \leq i \leq n - 1$ there exist $h_i \in H$ such that $\text{dist}_S(g_i, h_i) \leq R$. Set $h_0 = 1$ and $h_n = h$. Then $\text{dist}_S(h_i, h_{i+1}) \leq 2R + 1$, hence $h_{i+1} = h_i x_i$ for some $x_i \in X$, for every $0 \leq i \leq n - 1$. It follows that $h = h_n = x_1 x_2 \cdots x_n$, whence X generates H and $|h|_X \leq |h|_S = n$. \square

Other geometric models of groups. Let G be a finitely generated group. Then G is the quotient group of a free group F_n . Therefore, if Y is any connected space whose fundamental group surjects to F_n , we obtain the homomorphism

$$\phi : \pi_1(Y) \rightarrow G.$$

Therefore, if Y is, say, locally simply-connected, we obtain the regular covering map

$$p : X \rightarrow Y$$

associated to the kernel of ϕ , such that the group of covering transformations of p is isomorphic to G . The group G acts properly discontinuously on X . We will be primarily interested in two cases:

1. Y is a compact CW-complex.
2. Y is a compact Riemannian manifold.

The structure of a CW-complex/Riemannian manifold, lifts from Y to X and the action of G preserves this structure: The action of G is cellular in the former case and is isometric in the latter case. If X is a simplicial complex and the action $G \curvearrowright X$ is simplicial, the *standard* metric dist on X is G -invariant and, hence, (X, dist) is a *simplicial geometric model* for the group G . If X is a Riemannian manifold, taking dist to be the Riemannian distance function we obtain a *Riemannian geometric model* for the group G .

In order to construct a CW-complex Y we can take, for instance, Y equal to the bouquet of n circles; the space X is, then, a Cayley graph of G . In order to get a Riemannian manifold Y , we can take Y to be a compact Riemannian surface of genus n , the epimorphism $\pi_1(Y) \rightarrow F_n$ is then given by

$$\phi : \left\langle a_1, b_1, \dots, a_n, b_n \mid \prod_{i=1}^n [a_i, b_i] \right\rangle \rightarrow \langle a_1, \dots, a_n \rangle,$$

$$\phi(a_i) = a_i, \quad \phi(b_i) = 1, i = 1, \dots, n.$$

In the case when G is finitely presented, one can do a bit better: Each finite presentation of G yields a finite *presentation complex* Y of G (see Definition 7.91), which is a finite CW-complex whose fundamental group is isomorphic to G . Hence, the universal cover X of Y is a simply-connected CW-complex and we obtain a cellular, free, properly discontinuous and cocompact action $G \curvearrowright X$. Since every compact CW-complex is homotopy-equivalent to a compact simplicial complex, we can also finite a simplicial complex X with the above properties.

Similarly, there exists a smooth closed m -manifold M ($m \geq 4$) whose fundamental group is isomorphic to G (see e.g. [Hat02]). Then we equip $Y = M$ with a Riemannian metric; lifting this metric to the universal cover $X \rightarrow Y$, we

obtain a simply-connected complete Riemannian manifold X and a free, properly discontinuous, isometric and cocompact action $G \curvearrowright X$.

Working with geometric models (simplicial or Riemannian) of groups G is a major theme and a key technical tool of Geometric Group Theory. We will use this tool throughout this book. As one example, we will use both simplicial and Riemannian geometric models in Chapters 20 and 21 in order to prove group-theoretic theorems by Stallings and Dunwoody.

On the other hand, when replacing a group with its geometric model, we are faced with the inevitable:

QUESTION 7.86. What do all these geometric models have in common?

We will discuss this question in detail in the next chapter.

7.10. Volumes of maps of cell complexes and Van Kampen diagrams

The goal of this section is to describe several notions of volumes of maps and to relate them to each other and to the word reductions in finitely presented groups. It turns out that most of these notions are equivalent, but, in few cases, there subtle differences.

7.10.1. Simplicial, cellular and combinatorial volumes of maps. Recall that in Section 3.4 we defined volumes of maps between Riemannian manifolds. Our next goal is to give simplicial/cellular/combinatorial analogues of Riemannian volumes.

DEFINITION 7.87. Let $f : Z \rightarrow X$ be a simplicial map of simplicial complexes. Then $Vol_n^{sim}(f)$, the *simplicial n -volume* of f , is defined as the number of n -dimensional simplices in Z which are mapped by f onto n -dimensional simplices.

The *combinatorial n -volume* $Vol_n^{com}(f)$ of f as the number of n -dimensional simplices in Z .

At the first glance the combinatorial volume appears to be a strange concept, as it is independent of the map f ; nevertheless, the definition turns out to be quite useful, see Section 9.7.

We next define a *cellular* analogue of the simplicial volume.

DEFINITION 7.88. Let X, Y be n -dimensional almost regular cell complexes. A cellular map $f : X \rightarrow Y$ is said to be *almost regular* if for every n -cell σ in X either:

- (a) f collapses σ , i.e. $f(\sigma) \subset Y^{(n-1)}$, or
 - (b) f maps the interior of σ homeomorphically to the interior of an n -cell in Y .
- An almost regular map is *regular* if only (b) occurs.

For instance, a simplicial map of simplicial complexes is almost regular, while a simplicial topological embedding of simplicial complexes is regular. The following definition first appeared in [AWP99]; we refer the reader to [BBFS09] and [ABD⁺13] for more geometric treatment.

DEFINITION 7.89. The *cellular n -volume* $Vol_n^{cell}(f)$ of an almost regular (cellular) map $f : Z \rightarrow X$ of almost regular cell complexes is the number of n -cells in Z which map homeomorphically onto n -cells in X .

The combinatorial/simplicial/cellular 1-volume is called *length* and the 2-volume is called *area*. They are denoted $length^{com}$ and $Area^{com}$, etc., respectively.

7.10.2. Topological interpretation of finite-presentability.

LEMMA 7.90. *A group G is isomorphic to the fundamental group of a finite cell complex Y if and only if G is finitely presented.*

PROOF. 1. Suppose that G has a finite presentation

$$\langle S|R \rangle = \langle x_1, \dots, x_n | r_1, \dots, r_m \rangle.$$

We construct a finite 2-dimensional cell-complex Y , as follows. The complex Y has unique vertex v . The 1-skeleton of Y is the n -rose, the bouquet of n circles $\gamma_1, \dots, \gamma_n$ with the common point v , the circles are labeled x_1, \dots, x_n . Observe that the free group $F(S)$ is isomorphic to $\pi_1(Y^{(1)}, v)$ where the isomorphism sends each x_i to the circle in $Y^{(1)}$ with the label x_i . Thus, every word w in X^* determines a based loop L_w in $Y^{(1)}$ with the base-point v . In particular, each relator r_i determines a loop $\alpha_i := L_{r_i}$. We then attach 2-cells $\sigma_1, \dots, \sigma_m$ to $Y^{(1)}$ using the maps $\alpha_i : \mathbb{S}^1 \rightarrow Y^{(1)}$ as the attaching maps. Let Y be the resulting cell complex. It is clear from the construction that the complex Y is *almost regular*.

We obtain a homomorphism $\phi : F(S) \rightarrow \pi_1(Y^{(1)}) \rightarrow \pi_1(Y)$. Since each r_i lies in the kernel of this homomorphism, ϕ descends to a homomorphism $\psi : G \rightarrow \pi_1(Y)$. It follows from the Seifert–van Kampen theorem (see [Hat02] or [Mas91]) that ψ is an isomorphism.

2. Suppose that Y is a finite complex with $G \cong \pi_1(Y)$. Pick a maximal subtree $T \subset Y^{(1)}$ and let X be the complex obtained by contracting T to a point. Since T is contractible, the resulting map $Y \rightarrow X$ (contracting T to a point $v \in X^{(0)}$) is a homotopy-equivalence. The 1-skeleton of X is an n -rose with the edges $\gamma_1, \dots, \gamma_n$ which we will label x_1, \dots, x_n . It now again follows from the Seifert–van Kampen theorem that X defines a finite presentation of G : The generators x_i are the loops γ_i and the relators are the attaching maps $\mathbb{S}^1 \rightarrow X^{(1)}$ of the 2-cells of X . \square

DEFINITION 7.91. The 2-dimensional complex Y constructed in the first part of the above proof is called the *presentation complex* of the presentation

$$\langle x_1, \dots, x_n | r_1, \dots, r_m \rangle.$$

7.10.3. Presentations of central coextensions. In this section we illustrate the concepts introduced earlier in the case of central coextensions. Let $f : F = F(S) \rightarrow G$ be an epimorphism. Consider a central coextension

$$0 \rightarrow A \rightarrow \tilde{G} = \tilde{G}_\omega \xrightarrow{p} G \rightarrow 1$$

associated with a cohomology class $\omega \in H^2(G, A)$. Our goal is to describe a presentation of \tilde{G} in terms of the presentation of G given by f . In Section 5.9.6, we discussed a *pull-back construction* for central coextensions. Applying this to the homomorphism f , we obtain a central coextension

$$0 \rightarrow A \rightarrow \tilde{F} \xrightarrow{q} F \rightarrow 1$$

and a commutative diagram of homomorphisms

$$\begin{array}{ccc} \tilde{F} & \xrightarrow{\tilde{f}} & \tilde{G} \\ \downarrow q & & \downarrow p \\ F & \xrightarrow{f} & G \end{array}.$$

The homomorphism \tilde{f} is surjective (Exercise 5.140, Part 1). Since F is free, its central coextension splits and there exists a homomorphism $s : F \rightarrow \tilde{F}$ right-inverse to q . Following our discussion in the end of Section 5.9.6, we pick a set-theoretic section s_1 of p , lift it to a set-theoretic section s_1 of q and observe that $\text{Ker}(\tilde{f}) = s_1(K) \cong K$, where $K = \text{Ker}(f)$, the normal closure of a set $R = \{R_i : i \in I\}$ of defining relators of G . Using the section s we define an isomorphism $\tilde{F} \cong F \times A$. With this identification, the restriction of the section s_1 to K is a homomorphism $\varphi : K \rightarrow A$ (invariant under conjugation by elements of F). By abusing the notation, we denote φ by φ_ω , even though, there are some choices involved in constructing φ from the central coextension. Then the group \tilde{G} is isomorphic to the quotient of $F \times A$ by the normal closure of the subset

$$\{R_i\varphi(R_i)^{-1} : i \in I\}.$$

In order to describe the corresponding presentation of \tilde{G} , we fix a presentation

$$\langle T | Q \rangle$$

of the group A , $T = \{t_j : j \in J\}$, $Q = \{Q_\ell : \ell \in L\}$. We then obtain the presentation

$$\langle S \sqcup T | Q, [x, t] = 1, x \in S, t \in T \rangle$$

of the group $F \times A$. Lastly, the presentation of the group \tilde{G} is:

$$\langle S \sqcup T | Q, [x, t] = 1, x \in S, t \in T, R_i = \varphi(R_i), i \in I \rangle.$$

EXAMPLE 7.92. Let G be the fundamental group of closed oriented surface Y of genus $n \geq 1$ with the standard presentation

$$\langle a_1, b_1, \dots, a_n, b_n | R = [a_1, b_1] \cdots [a_n, b_n] \rangle.$$

Since $Y = K(G, 1)$, $H^2(G) \cong H^2(Y) \cong \mathbb{Z}$. The space of F_{2n} -invariant homomorphisms

$$\langle \langle R \rangle \rangle \rightarrow A = \mathbb{Z}$$

is isomorphic to \mathbb{Z} (since every such homomorphism is determined by its restriction to R). Thus, central coextensions \tilde{G}_ω of G are indexed by integers $e \in \mathbb{Z}$:

$$\varphi : R \mapsto e \in \mathbb{Z}.$$

The group \tilde{G}_ω has the presentation

$$\langle a_1, b_1, \dots, a_n, b_n, t | [a_1, b_1] \cdots [a_n, b_n] = t, [a_i, t] = 1, [b_i, t] = 1, i = 1, \dots, n \rangle.$$

We next give (without a proof) two topological interpretations of the group \tilde{G} and its presentation. Suppose that the presentation complex Y of $\langle S | R \rangle$ is aspherical, i.e. $\pi_2(Y) = 0$. Then $H^2(G, A) \cong H^2(Y, A)$. We will use cellular cohomology in order to compute $H^2(Y, A)$. Since Y is 2-dimensional, $Z^2(Y, A) = C^2(Y, A)$. The generators of $C_2(Y)$ are 2-cells e_i , which are labelled by the relators $R_i \in R$. Let \mathcal{R} denote the free abelian group with the basis R . Then we have the isomorphism

$$\Psi : Z^2(Y, A) \cong \text{Hom}(\mathcal{R}, A),$$

which sends a cocycle $c : e_i \mapsto c(e_i) \in A$ to the homomorphism $\psi_c : R_i \mapsto c(e_i)$. Altering c by a coboundary, results in a new element of $\text{Hom}(\mathcal{R}, A)$. In other words, only the coset $\psi_c \Psi(B^2(Y, A))$ is determined by the cohomology class $[c] \in H^2(Y, A)$.

The class $[c]$ maps to a cohomology class $\omega \in H^2(G, A)$ under the isomorphism $H^2(Y, A) \cong H^2(G, A)$. The class ω , as we say before, determines (subject to some

ambiguity) a homomorphism φ_ω from the normal closure of R in F into A . This homomorphism is determined by its restriction to R (since φ_ω is F -invariant). Thus, both $[c] \in H^2(Y, A)$ and ω determine (equivalence classes) of homomorphisms $\psi_c, \phi_\omega : \mathcal{R} \rightarrow A$. One can verify that *these equivalence classes are the same*. Therefore, we obtain a somewhat more concrete description of the presentation of the group \tilde{G} : In addition to the relators of $F \times A$, we have the relators

$$R_i = \psi_c(R_i), \quad i \in I.$$

This is the first topological interpretation of the presentation of \tilde{G} .

The second topological interpretation requires *complex line bundles* $\xi : L \rightarrow Y$. Such line bundles are parameterized by the elements of $H^2(Y) = H^2(Y, \mathbb{Z})$. The cohomology class defining L is called the *first Chern class* $c_1(\xi)$ of the bundle ξ . We refer the reader to [Che95, pp. 33–34] for the details. Given a line bundle $L \rightarrow Y$, we define the space $L_o \subset L$ by removing the image of the zero section from L (then L_o is the total space of the \mathbb{C}^* -bundle associated with L). The fundamental group of \mathbb{C}^* is infinite cyclic and $\pi_1(Y)$ acts trivially on this group. Therefore, taking into the account the long exact sequence of homotopy groups of the fibration

$$\mathbb{C}^* \rightarrow L_o \rightarrow Y,$$

we obtain that $\pi_1(L_o)$ is isomorphic to a certain central coextension \tilde{G} of G :

$$0 \rightarrow \mathbb{Z} \rightarrow \tilde{G} \rightarrow G \rightarrow 1.$$

Then: *The cohomology class ω defining this central coextension maps to the first Chern class $c_1(\xi)$ under the isomorphism $H_1(G) \rightarrow H_1(Y)$.*

7.10.4. Dehn function and van Kampen diagrams. One of the oldest algorithmic problems in group theory is the *word problem*. This problem is *largely* controlled by the *Dehn function* of the group, which depends on the group presentation $\langle S|R \rangle$. In this section we define the Dehn function, van Kampen diagrams of finite presentations, and relate the latter to the word problem. We refer the reader to [LS77] for the more thorough treatment of this topic. The reader familiar with the treatment of van Kampen diagrams in [LS77] will notice that our definitions of diskoids and van Kampen diagrams are more general.

Suppose that w is a word in S , representing a trivial element of the group G with the presentation $\langle S|R \rangle$. How can we convince ourselves that, indeed, $w \equiv_G 1$? If R were empty, we could eliminate all the reductions in w , which will result in an empty word. In the case of non-empty R , we can try the same thing, namely, the reduction of w in $F = F(S)$. If the reduction results in an empty word, we are done; hence, we will work, in what follows, with non-empty reduced words. Thus, we will identify each w with a non-trivial element of the free group F . Any “proof” that such w is trivial in G would amount to finding a product decomposition of $w \in F$ of the form

$$(7.6) \quad w = \prod_{i=1}^k u_i r_i^{\pm 1} u_i^{-1},$$

where r_i 's are elements of R . Of course, it is to our advantage, to use as few defining relators r_i as we can, in order to get as short “proof” as possible. This leads us to

DEFINITION 7.93. The *algebraic area* of the (reduced) word w , such that $w \equiv_G 1$, is defined as the least number k of relators r_i used to describe w as a product

of conjugates of defining relators and their inverses. The algebraic area of w is denoted by $A(w)$.

The significance of this notion of area is that it captures the complexity of the word problem for the presentation $\langle S|R \rangle$ of the group G . In order to estimate “hardness” of the word problem, we then search for the words w of the largest area: The most reasonable way to do so, by analogy with the norms of linear operators, is to restrict to w 's of bounded word-length. This leads us to

DEFINITION 7.94 (Dehn function). The *Dehn function of the group G* (with respect to the finite presentation $\langle S|R \rangle$) is defined as

$$Dehn(\ell) := \max\{A(w) : |w| \leq \ell\}$$

where w 's are elements in S^* representing trivial words in G .

EXERCISE 7.95. Let $\langle S, |R \rangle, \langle S'|R' \rangle$ be finite presentation of the same group G . Show that the resulting Dehn functions $Dehn, Dehn'$ are approximately equivalent in the sense of the Definition 1.3.

In view of this property, we will frequently use the notation $Dehn_G$ for the Dehn function of G (with respect to some unspecified finite presentation of G).

Our next task is to describe a *geometric interpretation* of areas of words and Dehn functions. The classical tool for this task is *van Kampen diagrams*, which give topological interpretation to the product decompositions (7.6).

Van Kampen diagrams. Suppose that Y is the presentation complex of the presentation $\langle S|R \rangle$. Each nonempty reduced word w represents a certain closed edge-path c_w in $Y^{(1)}$ (here and in what follows, the base-point is the sole vertex y of Y): Each letter s in w corresponds to an oriented edge in $Y^{(1)}$ representing the corresponding generator of G or its inverse.

We will think of c_w as a regular map $\mathbb{S}^1 \rightarrow Y^{(1)}$, where \mathbb{S}^1 is the circle equipped with a certain fixed cell-complex structure as well as a base-vertex.

EXAMPLE 7.96. If $w \in S \cup S^{-1}$, then the almost regular cell complex structure on \mathbb{S}^1 will consist of a single vertex and a single edge. The map c_w is a topological embedding.

As the map c_w is null-homotopic, one can extend the map c_w to a cellular map $\mathbb{D}^2 \rightarrow Y$. We will see below that one can find an extension of c_w which is almost regular; we then will define the *combinatorial area* of c_w as the least combinatorial area of the resulting extension. The extension will collapse some 2-cells in \mathbb{D}^2 into the 1-skeleton of Y : These cells contribute nothing to the combinatorial area and we would like to get rid of them. *Van Kampen diagrams* are a convenient (and traditional) way to eliminate these dimension reductions.

DEFINITION 7.97. We say that a contractible finite planar almost regular cell complex $K \subset \mathbb{R}^2$ is a *diskoid* (a *tree of disks* or a *tree-graded disk*) if every edge of K is contained in the boundary of K in \mathbb{R}^2 .

In other words, K is obtained from a finite simplicial tree by replacing some vertices with (cellulated) 2-disks, which is why we think of K as a “tree of disks”. To simplify the picture, the reader can (at first) think of K as a single disk in \mathbb{R}^2 rather than a tree of disks. In what follows, we will assume that K is *non-trivial*: $K \neq \emptyset$ and K does not consist of a single vertex. (The case when K is a single

vertex would correspond to the case of the empty word w .) Note that the boundary ∂K of K in \mathbb{R}^2 is also an almost regular cell complex (a planar graph). However, the graph ∂K may have some valence 1 vertices, the *leaves* of ∂K : These leaves will not exist in the case of van Kampen diagrams of reduced words.




FIGURE 7.9. Example of a diskoid.

The complex K admits a *canonical enlargement* to a planar almost regular cell complex \widehat{K} homeomorphic to the disk \mathbb{D}^2 : The complement of K in \widehat{K} is homeomorphic to the annulus

$$(0, 1] \times \mathbb{S}^1.$$

The 2-cells in $\widehat{K} \setminus \text{int}(K)$, are rectangles. If e is an edge of ∂K which belongs to the closure of the interior of K in \mathbb{R}^2 , then e is the boundary edge of exactly one such rectangle, otherwise, e is the boundary edge of exactly two rectangles in $\widehat{K} \setminus \text{int}(K)$. Furthermore, every rectangular boundary face σ shares exactly one edge, called e_σ , with K . We refer to these 2-cells σ as the *boundary faces* of \widehat{K} . Thus, the number of boundary faces of \widehat{K} is at most twice the number of edges in W . We have the canonical retraction

$$\kappa : \widehat{K} \rightarrow K$$

sending each boundary face σ to the edge e_σ . See Figure 7.10.

Restricting κ to the boundary circle of the disk \widehat{K} , we obtain a regular cellular map $b : \mathbb{S}^1 \rightarrow \partial K$ tracing the boundary of K according to the orientation induced on the boundary arcs of K from the Euclidean plane. Here \mathbb{S}^1 is given the structure of a regular cell complex C coming from \widehat{K} . We will refer to b as the *boundary map* of K . For each boundary edge e of K not contained in the closure of the interior of K in \mathbb{R}^2 , the preimage $b^{-1}(e)$ contains exactly two edges in C .

We now describe a certain class of maps from diskoids to almost regular 2-dimensional cell complexes Y .

DEFINITION 7.98. A regular cellular map $h : K \rightarrow Y$ from a diskoid to an almost regular 2-dimensional cell complex Y is called a *van Kampen diagram in Y* . Suppose that $w \in S^*$ represents the identity in G and $c_w : \mathbb{S}^1 \rightarrow Y$ is the associated loop in the presentation complex Y of $\langle S|R \rangle$. If the composition $\partial h := h \circ b$ of h with the boundary map of K equals c_w , we will say that h is a *van Kampen diagram* of the word w .

It will be sometimes convenient to consider van Kampen diagrams not in presentation complexes but in their universal covers. This, of course, will make no difference as far as combinatorial areas and lengths are concerned.




FIGURE 7.10. Collapsing map κ .

It is customary to describe a van Kampen diagram by labeling *oriented edges* \bar{e} of K by the elements of $S \cup S^{-1}$ which correspond to the edges $h(\bar{e})$ of Y . (Recall that some boundary edges of the diskoid K have two opposite orientations defined by the boundary map b : They will be labelled by a generator and its inverse respectively.) For instance, for the standard presentation

$$\langle a, b | aba^{-1}b^{-1} \rangle,$$

of the group \mathbb{Z}^2 , a van Kampen diagram of the relator $[a, b]$ is described in the Figure 7.11.




FIGURE 7.11. A van Kampen diagram of the commutator $[a, b]$.

A van Kampen diagram of the relator $[a, b]^2$ for the same presentation is described in the Figure 7.12.




FIGURE 7.12. A van Kampen diagram of the relator $[a, b]^2$.

EXERCISE 7.99. Note that each van Kampen diagram $h : K \rightarrow Y$ extends to the canonical enlargement of K :

$$\widehat{h} = h \circ \kappa : \widehat{K} \rightarrow Y.$$

Then

$$\text{Area}(h) = \text{Area}^{\text{cell}}(h) = \text{Area}^{\text{cell}}(\widehat{h}).$$

LEMMA 7.100 (Van Kampen lemma). 1. For every word w in the alphabet $S \cup S^{-1}$, representing the identity element 1_G , there exists a van Kampen diagram $h : K \rightarrow Y$ such that the maps ∂h and $c = c_w$ are homotopic as maps $\mathbb{S}^1 \rightarrow Y^{(1)}$, rel. the base-vertex in \mathbb{S}^1 . Furthermore, $\text{Area}(h)$ equals $A(w)$.

2. If the word w is reduced, then there exists a van Kampen diagram h in Part 1 such that $\partial h = c_w$.

PROOF. 1. According to the product decomposition (7.6) of $w \in F(S)$, the circle \mathbb{S}^1 is subdivided into cyclically ordered and oriented cellular subarcs

$$\alpha_1^+ \cup \beta_1 \cup \alpha_1^- \cup \dots \cup \alpha_k^+ \cup \beta_k \cup \alpha_k^-,$$

so that:

- (1) The path $c|_{\alpha_i^+}$ represents the word u_i .
- (2) The path $c|_{\alpha_i^-}$ represents the word u_i^{-1} .
- (3) The path $c|_{\beta_i}$ represents the word $r_i^{\pm 1}$.

The orientation on the arcs α_i^\pm, β_i (induced from the standard orientation of unit circle), defines for each of these arcs the *head* and the *tail* vertex.

We then connect (some of) the vertices of C by chords in \mathbb{D}^2 as follows:

For each i , we connect the tail of α_i^+ to the head of α_i^- by the chord ϵ_i^+ and the head of α_i^+ to the tail of α_i^- by the chord ϵ_i^- .

The chords $\epsilon_i^\pm, \epsilon_j^\pm$ may cross only at the boundary circle \mathbb{S}^1 . See Figure 7.13. The chords ϵ_i^\pm , together with the original cell-complex structure C on \mathbb{S}^1 , define a regular cell complex structure \tilde{K} on \mathbb{D}^2 , where every vertex is in \mathbb{S}^1 . There are three types of 2-cells in \tilde{K} :

- 1 Cells A_i bounded by the “bigons” $\beta_i \cup \epsilon_i^-$.
- 2 Cells B_i bounded by “rectangles” $\alpha_i^+ \cup \epsilon_i^+ \cup \alpha_i^- \cup \epsilon_i^-$.
- 3 The rest, not having any edges in \mathbb{S}^1 .

Note that in type (2) we allow the degenerate case when α_i^\pm is a single vertex: Then the corresponding “rectangle” degenerates to a triangle. It can even become a bigon in case when the word u_i is empty. Similarly, there will be one case when a “bigon” is actually a monogon: $w = r_1$.

We now collapse each type (3) cell to a vertex and collapse each type (2) cell to an edge e_i (so that each α_i^\pm maps homeomorphically onto this edge while the chords ϵ_i^\pm map to the end-points of e_i). Note that α_i^\pm , with their orientation inherited from \mathbb{S}^1 , define two opposite orientations on e_i . The quotient complex K will be our diskoid. We also obtain the quotient *collapsing* map $\kappa : \tilde{K} \rightarrow K$. Because “rectangles” can degenerate to triangles, or “bigons”, the complex K is merely almost regular, not regular.

We define a map $h : K^{(1)} \rightarrow Y$ such that

$$h \circ \kappa|_{\alpha_i^\pm} = c_{u_i^\pm 1}$$

while

$$h \circ \kappa|_{\beta_i} = c_{r_i}.$$

Lastly, we extend h to the 2-cells $\kappa(A_i)$ in K : $h : \kappa(A_i) \rightarrow Y$ are the standard parameterizations of the 2-cells in Y corresponding to the defining relators r_i .

By the construction, h is a van Kampen diagram of w : The maps $h \circ \kappa$ and c_w are homotopic as based loops $\mathbb{S}^1 \rightarrow Y^{(1)}$. However, as maps they need not be the same, as the product decomposition (7.6) need not be a reduced word. The equality

$$Area^{sim}(h) = k$$

is immediate from the construction.

2. Suppose now that the word w is reduced. The boundary map ∂h reads off a word w' in S^* : The word w' is obtained by reading off the boundary labels defined via h . The word w' , by the construction, represents the same element of $F(S)$ as w . If w' were also reduced, we would be done. In general, however, w' is not reduced and, hence, we can find two adjacent boundary edges e_1, e_2 in ∂K , whose labels are inverses of each other. We then glue the edges e_1, e_2 together. The result is a new diskoid K_1 , where the projection of e_1, e_2 is no longer a boundary edge. The map h descends to a van Kampen diagram $h_1 : K_1 \rightarrow Y$. By repeating this procedure inductively we eliminate all boundary reductions and obtain a new van Kampen diagram $h' : K' \rightarrow Y$ with the required properties. \square

Lemma 7.100 shows that the algebraic area of w does not exceed the least combinatorial area of van Kampen diagrams of w .

EXERCISE 7.101. i. Show that $A(w)$ equals the least cellular area of all almost regular maps $f : \mathbb{D}^2 \rightarrow Y$ extending the cellular map c_w , where \mathbb{D}^2 is given the structure of an almost regular cell complex. Hint: Convert maps f into product decompositions of w as in (7.6).

ii. Combine Part (i) with the canonical extension \hat{h} of van Kampen diagrams $h : K \rightarrow Y$, to conclude that $A(w)$ equals

$$\min_{h:K \rightarrow Y} Area^{cell}(h),$$

where the minimum is taken over all Van Kampen diagrams of w in Y . In other words, the algebraic area of the words w (trivial in G) equals the cellular area of




FIGURE 7.13

the loops c_w in Y (the cellular area defined via van Kampen diagrams of w or, equivalently, almost regular extensions $\mathbb{D}^2 \rightarrow Y$ of c_w).

We will return to Dehn functions in Section 7.13 after discussing residual finiteness of groups.

7.11. Residual finiteness

Even though studying infinite groups is our primary focus, questions in group theory can be, sometimes, reduced to questions about finite groups. *Residual finiteness* is the concept that (sometimes) allows such a reduction.

DEFINITION 7.102. A group G is said to be *residually finite* if

$$\bigcap_{i \in I} G_i = \{1\},$$

where $\{G_i : i \in I\}$ is the set of all finite-index subgroups in G .

Clearly, subgroups of residually finite groups are also residually finite. In contrast, if G is an infinite simple group, then G cannot be residually-finite.

LEMMA 7.103. *A finitely generated group G is residually finite if and only if for every $g \in G \setminus \{1\}$, there exists a finite group Φ and a homomorphism $\varphi : G \rightarrow \Phi$, such that $\varphi(g) \neq 1$.*

PROOF. Suppose that G is residually finite. Then, for every $g \in G \setminus \{1\}$ there exists a finite-index subgroup $G_i \leq G$ so that $g \notin G_i$. It follows that G contains a normal subgroup of finite index $N_i \triangleleft G$, such that $N_i \leq G_i$. Clearly, $g \notin N_i$ and

$|G : N_i| < \infty$. Now, setting $\Phi := G/N_i$, we obtain the required homomorphism $\varphi : G \rightarrow \Phi$.

Conversely, suppose that for every $g \neq 1$ we have a homomorphism $\varphi_g : G \rightarrow \Phi_g$, where Φ_g is a finite group, so that $\varphi_g(g) \neq 1$. Setting $N_g := \text{Ker}(\varphi_g)$, we get

$$\bigcap_{g \in G} N_g = \{1\}.$$

The above intersection, of course, contains the intersection of all finite-index subgroups in G . \square

EXERCISE 7.104. Direct products of residually finite groups are again residually finite.

LEMMA 7.105. *If a group G contains a residually finite subgroup of finite index, then G itself is residually finite.*

PROOF. Let $H \leq G$ be a finite index residually finite subgroup. The intersection of all finite-index subgroups

$$(7.7) \quad \bigcap_{i \in I} H_i$$

of H is $\{1\}$. Since H has finite index in G and each $H_i \leq H$ as above has finite index in G , the intersection of all finite-index subgroups of G is contained in (7.7) and, hence, is trivial. \square

PROPOSITION 7.106. *A semidirect product of a finitely generated residually finite group with a (not necessarily finitely generated) residually finite group is also residually finite.*

PROOF. Let G be a group that splits as a semidirect product $H \rtimes Q$, where H and Q are residually finite, and H is moreover finitely generated. Let p denote the projection homomorphism $G \rightarrow Q$.

Consider $g \in G \setminus \{1\}$. If g does not belong to H , then $p(g) \neq 1$ and the residual finiteness of Q implies that there exists a homomorphism of Q to a finite group which sends $p(g)$ to a non-trivial element. By composing the homomorphisms, we obtain a homomorphism of G to a finite group which sends g to a non-trivial element.

Suppose, therefore, that g is in H . Let $F < H$ be a finite-index subgroup which does not contain g . Since H is finitely generated, Proposition 5.11, (2), implies that there exists a finite-index subgroup $A \leq F$ which is a characteristic subgroup of H . The subgroup $A \rtimes Q$ is a finite index subgroup in $G = H \rtimes Q$ that does not contain g . \square

REMARK 7.107. Proposition 7.106 cannot extend to short exact sequences that do not split, following the terminology of Definition 5.27. In other words, it is not true that if H and Q are residually finite, H is finitely generated, and there is a short exact sequence

$$1 \rightarrow H \rightarrow G \rightarrow Q \rightarrow 1,$$

then G is residually finite. Indeed, there exist coextensions

$$1 \rightarrow \mathbb{Z}_2 \rightarrow G \rightarrow Q \rightarrow 1$$

where Q is finitely generated residually finite, while G is not residually finite; see [Mil79].

COROLLARY 7.108. *Suppose that H is a finitely generated residually finite group and we have a cyclic extension of H , i.e. a group G which appears in a short exact sequence*

$$1 \rightarrow H \rightarrow G \xrightarrow{p} C \rightarrow 1,$$

where C is a cyclic group. Then G is also residually finite.

PROOF. When C is finite, the statement follows from Lemma 7.105. When C is infinite, that is $C \simeq \mathbb{Z}$, the short exact sequence splits and $G \simeq H \rtimes \mathbb{Z}$, by Corollary 7.24. The result now follows from Proposition 7.106. \square

A special case of this corollary is residual finiteness of groups virtually isomorphic to cyclic groups.

COROLLARY 7.109. *Each group G virtually isomorphic to \mathbb{Z} is residually finite and contains an infinite cyclic subgroup of finite index.*

PROOF. In view of Lemma 7.105, it suffices to show that if G is a finite coextension of the infinite cyclic group C ,

$$1 \rightarrow F \rightarrow G \xrightarrow{p} C \rightarrow 1$$

(where F is finite), then G contains an infinite cyclic subgroup of finite index. This is an immediate consequence of Corollary 7.108. \square

REMARK 7.110. The first part of Corollary 7.109 can be generalized, by replacing \mathbb{Z} with “a polycyclic group”; see Theorem 13.78.

EXAMPLE 7.111. The group $\Gamma = GL(n, \mathbb{Z})$ is residually finite. Indeed, we take subgroups $\Gamma(p) \leq \Gamma$, $\Gamma(p) = \text{Ker}(\varphi_p)$, where $\varphi_p : \Gamma \rightarrow GL(n, \mathbb{Z}_p)$ is the reduction modulo p . If $g \in \Gamma$ is a non-trivial element, we consider its non-zero off-diagonal entry $g_{ij} \neq 0$. Then $g_{ij} \not\equiv 0 \pmod{p}$, whenever $p > |g_{ij}|$. Thus, $\varphi_p(g) \neq 1$ and Γ is residually finite.

COROLLARY 7.112. *The free group F_2 of rank 2 is residually finite. Every free group of (at most) countable rank is residually finite.*

PROOF. As we saw in the Example 7.62 the group F_2 embeds in $SL(2, \mathbb{Z})$. Furthermore, every free group of (at most) countable rank embeds in F_2 (see Proposition 7.80). Now, the assertion follows from the Example 7.111. \square

We note that there are other proofs of residual finiteness of finitely generated free groups: Combinatorial (see [Hal49]), topological (see [Sta83]) and geometric (see [Sco78]).

EXERCISE 7.113. For an arbitrary cardinality r , the free group F_r of rank r is residually finite.

Less trivially,

THEOREM 7.114 (K. W. Gruenberg [Gru57]). *Free products of residually finite groups are again residually finite.*

The simple argument for $GL(n, \mathbb{Z})$ is a model for a proof of a harder theorem:

THEOREM 7.115 (A. I. Mal'cev [Mal40]). *Let Γ be a finitely generated subgroup of $GL(n, R)$, where R is a commutative ring with unity. Then Γ is residually finite.*

Mal'cev's theorem is complemented by the following result proven by A. Selberg and known as *Selberg's Lemma* [Sel60]:

THEOREM 7.116 (Selberg's Lemma). *Let Γ be a finitely generated subgroup of $GL(n, F)$, where F is a field of characteristic zero. Then Γ contains a torsion-free subgroup of finite index.*

Proofs of Mal'cev's and Selberg's theorems, will be given in the Appendix to this book, written by Bogdan Nica.

PROBLEM 7.117. It is known that all (finitely generated) Coxeter groups are linear; see e.g. [Bou02]. Is the same true for all Artin groups, Shephard groups, generalized von Dyck groups? (Note that even linearity of Artin Braid groups was unknown prior to [Big01].) Is it at least true that all these groups are residually finite?

Mal'cev's theorem implies that infinite finitely generated matrix groups cannot be simple. On the other hand, $PSL(2, \mathbb{Q})$ is a simple countable matrix group.

PROBLEM 7.118. Are there infinite simple discrete subgroups $\Gamma < SL(n, \mathbb{R})$, $n \gg 3$?

Here discreteness of Γ means that it is discrete with the subspace topology. One can prove that infinite discrete subgroups of $SO(n, 1)$, and, more generally, isometry groups of *rank 1 symmetric spaces*, cannot be simple: Given an infinite discrete subgroup $\Gamma < SO(n, 1)$, which does not preserve a line in \mathbb{R}^{n+1} , one shows (using a ping-pong argument) that there exists an infinite order element $g \in \Gamma$, such that the normal closure Λ of $\{g\}$ in Γ is a free subgroup of Γ . If $\Lambda = \Gamma$ then Γ is not simple (since non-trivial free groups are never simple); otherwise, Λ is a proper normal subgroup of Γ .

7.12. Hopfian and cohopfian properties

A group G is called *hopfian* if every epimorphism $G \rightarrow G$ is injective. Mal'cev prove in [Mal40] that every residually finite group is hopfian. On the other hand, many Baumslag-Solitar groups are not hopfian. Collins and Levin [CL83] gave a criterion for $BS(m, n)$ to be hopfian (for $|m| > 1, |n| > 1$): The numbers m and n should have the same set of prime divisors.

An example of a hopfian group with a nonhopfian subgroup of finite index is the Baumslag-Solitar group

$$BS(2, 4) = \langle a, b | ab^2a^{-1} = b^4 \rangle.$$

According to the criterion of Collins and Levin, this group is hopfian. Meskin in [Mes72] proved that $BS(2, 4)$ contains a nonhopfian subgroup of finite index.

We now turn to *cohopfian property*, which is *dual* to the hopfian property: Every injective endomorphism $f : G \rightarrow G$ is surjective. Of course, every finite group is cohopfian. Sela proved in [Sel97b] that every torsion-free 1-ended *hyperbolic group* is cohopfian. On the other hand, every free abelian group \mathbb{Z}^n is not cohopfian: The endomorphism $g \mapsto g^k$, $k > 1$, is injective but not surjective. However, there are finitely generated cohopfian nilpotent groups [Bel03]. Dekimpe and Deré [DD16] recently found a complete criterion for virtually nilpotent groups to be cohopfian, in particular, they proved that in this class of groups, cohopfian property is invariant under virtual isomorphisms.

EXERCISE 7.119. Each free group $F = F(X)$ is not cohopfian, provided that X is non-empty, of course.

We now give an example of a cohopfian virtually free group. Let

$$G_1 = \langle a, b, c \mid abc = 1, a^2 = 1, b^3 = 1, c^3 = 1 \rangle$$

be the alternating group A_4 . We leave it to the reader to check that the subgroup $C = \langle c \rangle$ of G_1 is *malnormal*: For each $g \in G_1$,

$$gCg^{-1} \cap C \neq \{1\} \iff g \in C.$$

(One way to verify this is to let A_4 act as a group of symmetries of the regular 3-dimensional tetrahedron.) Define the amalgam

$$G = G_1 \star_C G_2,$$

where G_2 is another copy of G_1 and let T be the associated Bass-Serre tree. We let $v_i \in V(T)$ be the vertex fixed by G_i , $i = 1, 2$, and let $e = [v_1, v_2] \in E(T)$ be the edge fixed by C . Malnormality of C in G_1 translates to the fact that C does not fix any edges of T besides e .

We claim that the group G is cohopfian. Suppose that $f : G \rightarrow G$ is an injective endomorphism. (In fact, it suffices to assume that the restrictions of f to G_1 and G_2 are injective and that $f(G_1) \neq f(G_2)$.) Since the groups G_1, G_2 are finite, their images $f(G_i)$ fix vertices in the tree T . After composing f with an automorphism of G , we can assume that $f(G_1) = G_1$ (i.e. $f(G_1)$ fixes v_1) and $f(C) = C$ (i.e. $f(C)$ fixes e). Since C fixes only the edge e of T , the group $f(G_2)$ has to fix the vertex v_2 of e and, hence, $f(G_2) = G_2$. Surjectivity of f follows.

On the other hand, being an amalgam of finite groups, the group G is commensurable to the free group F_2 , see Theorem 7.51.

7.13. Algorithmic problems in the combinatorial group theory

Presentations provide a ‘compact’ form of defining a group. They were introduced by Max Dehn in the early 20-th century. A central topic in combinatorial group theory is to derive algebraic information about a group from its presentation. Below is a list of problems formulated in this spirit, whose origin lies in the work of Max Dehn in the early 20th century.

Word Problem. Let $G = \langle S \mid R \rangle$ be a finitely presented group. Construct a Turing machine (or prove its non-existence) that, given a word w in the generating set X as its input, would determine if w represents the trivial element of G , i.e. if

$$w \in \langle\langle R \rangle\rangle.$$

Conjugacy Problem. Let $G = \langle S \mid R \rangle$ be a finitely presented group. Construct a Turing machine (or prove its non-existence) that, given a pair of word v, w in the generating set X , would determine if v and w represent conjugate elements of G , i.e. if there exists $g \in G$ so that

$$[w] = g^{-1}[v]g.$$

To simplify the language, we will state such problems below as: Given a finite presentation of G , determine if two elements of G are conjugate.

Simultaneous Conjugacy Problem. Given n -tuples pair of words

$$(v_1, \dots, v_n), \quad (w_1, \dots, w_n)$$

in the generating set X and a (finite) presentation $G = \langle S|R \rangle$, determine if there exists $g \in G$ so that

$$[w_i] = g^{-1}[v_i]g, i = 1, \dots, n.$$

Triviality Problem. Given a (finite) presentation $G = \langle S|R \rangle$ as an input, determine if G is trivial, i.e. equals $\{1\}$.

Isomorphism Problem. Given two (finite) presentations $G_i = \langle X_i|R_i \rangle, i = 1, 2$ as an input, determine if G_1 is isomorphic to G_2 .

Embedding Problem. Given two (finite) presentations $G_i = \langle X_i|R_i \rangle, i = 1, 2$ as an input, determine if G_1 is isomorphic to a subgroup of G_2 .

Membership Problem. Let G be a finitely presented group, $h_1, \dots, h_k \in G$ and H , the subgroup of G generated by the elements h_i . Given an element $g \in G$, determine if g belongs to H .

Note that a group with solvable conjugacy or membership problem, also has solvable word problem. It was discovered in the 1950-s in the work of P. S. Novikov, W. Boone and M. O. Rabin [Nov58, Boo57, Rab58] that all of the above problems are *algorithmically unsolvable*. For instance, in the case of the word problem, given a finite presentation $G = \langle S|R \rangle$, there is no algorithm whose input would be a (reduced) word w and the output YES is $w \equiv_G 1$ and NO if not. A. A. Fridman [Fri60] proved that certain groups have solvable word problem and unsolvable conjugacy problem. We will later see examples of groups with solvable word and conjugacy problems but unsolvable membership problem (Corollary 11.158). Furthermore, there are examples [BH05] of finitely presented groups with solvable conjugacy problem but unsolvable simultaneous conjugacy problem for every $n \geq 2$.

Nevertheless, the main message of the Geometric Group Theory is that *under various geometric assumptions on groups (and their subgroups), all of the above algorithmic problems are solvable*. Incidentally, the idea that geometry can help solving algorithmic problems also goes back to Max Dehn. Here are two simple examples of solvability of **the Word Problem**:

PROPOSITION 7.120. *Free groups of finite rank have solvable Word Problem.*

PROOF. Let F be a group freely generated by $X = \{x_1, \dots, x_n\}$. Given a word w in $X \sqcup X^{-1}$, we cancel recursively all possible pairs $x_i x_i^{-1}, x_i^{-1} x_i$ in w . Eventually, this results in a reduced word w' . If w' is non-empty, then w represents a non-trivial element of F , if w' is empty, then $w \equiv 1$ in F . \square

PROPOSITION 7.121. *Every finitely presented residually finite group has solvable word problem.*

PROOF. First, note that if Φ is a finite group, then it has solvable word problem (using the multiplication table in Φ we can “compute” every product of generators as an element of Φ and decide if this element is trivial or not). Given a residually finite group G with finite presentation $\langle S|R \rangle$ we will run two Turing machines T_1, T_2 simultaneously:

The machine T_1 will look for homomorphism $\varphi : G \rightarrow S_n$, where S_n is the symmetric group on n letters ($n \in \mathbb{N}$): The machine will try to send generators x_1, \dots, x_m of G to elements of S_m and then check if the images of the relators in G under this map are trivial or not. For every such homomorphism, T_1 will check if $\varphi(g) = 1$ or not. If T_1 finds φ so that $\varphi(g) \neq 1$, then $g \in G$ is non-trivial and the process stops.

The machine T_2 will list all the elements of the kernel N of the quotient homomorphism $F_m \rightarrow G$: It will multiply conjugates of the relators $r_j \in R$ by products of the generators $x_i \in X$ (and their inverses) and transforms the product to a reduced word. Every element of N is such a product, of course. We first write $g \in G$ as a reduced word w in generators x_i and their inverses. If T_2 finds that w equals one of the elements of N , then it stops and concludes that $g \equiv_G 1$.

The point of residual finiteness is that, eventually, one of the machines stops and we determine whether g is trivial or not. \square

The Dehn function $Dehn_G(n)$ of a group G (equipped with the finite presentation $\langle S|R \rangle$) quantifies (to some extent) the difficulty of solving the word problem in G :

THEOREM 7.122 (S. Gersten, [Ger93a]). *A group G has solvable word problem if and only if its Dehn function is recursive.*

We note that G has solvable word problem if and only if its Dehn function $Dehn(\ell)$ is merely bounded above by a recursive function $r(\ell)$. Indeed, given such a bound, one applies the machine T_2 from the proof of Proposition 7.121 to the word w of length ℓ . The number of van Kampen diagrams $h : W \rightarrow Y$ with reduced ∂h and of area $\leq r(\ell)$, is also bounded above by a recursive function of ℓ . Hence, the algorithm terminates in a finite amount of time either representing w as a product of conjugates of defining relators or verifying that such representation does not exist.

Coarse geometry

In this chapter we will *coarsify* familiar geometric concepts: In the context of the coarse geometry the exact geometric computations will not matter, what matters are the *asymptotics* of various geometric quantities. For instance, the exact computations of distances become irrelevant, as long as we have uniform linear bounds on the distances; accordingly, isometries will be coarsified to *quasiisometries*. In the process of coarsification, metric spaces will be frequently replaced with nets which *approximate* or *discretize* these metric spaces:

$$\text{coarsification} : (X, \text{dist}_X) \rightarrow \text{a net } N \subset X,$$

while maps between metric spaces will be replaced with maps between the respective nets. We will coarsify the notions of area, volume and isoperimetric inequalities: Various geometric quantities will be replaced with cardinalities of certain nets.

The drawback of the coarse geometry is that we will be missing the beauty of the precise formulae and sharp inequalities of the classical geometry: We will be unable to tell apart the Euclidean $n \times n$ square from the Euclidean disk of radius n . Accordingly, we will think of n^2 as their (coarse) areas. What we gain, however, is a theory particularly adapted to discrete groups, in which any choice among the different geometric models of the group yields equivalence associated geometric invariants. Another advantage of working in a setting where it is allowed to discretize is that algorithmic approaches become possible.

8.1. Quasi-isometry

In this section we define an important equivalence relation between metric spaces: The *quasiisometry*. It is this concept that will relate different geometric models of finitely generated groups which were introduced in the previous chapter. The quasiisometry of spaces has two equivalent definitions (both useful): One which is easy to visualize and the other which makes it easier to understand why it is an equivalence relation. We begin with the first definition, continue with the second and then prove their equivalence.

The notion of *quasiisometry*

$$f : (X, \text{dist}_X) \rightarrow (Y, \text{dist}_Y)$$

between two metric spaces appeared first in work of Mostow on *strong rigidity* of lattices in semisimple Lie groups, see e.g. [Mos73]. Mostow's notion (the one of a *pseudo-isometry*) was slightly more restrictive than the one we will be using:

$$L^{-1} \text{dist}_X(x, x') - A \leq \text{dist}_Y(f(x), f(x')) \leq L \text{dist}_X(x, x').$$

In particular, pseudo-isometries used by Mostow were continuous maps. Later on, it became clear that it makes sense to add an additive constant in this equation on the right hand side as well, and work with (typically) discontinuous maps.

DEFINITION 8.1. Two metric spaces (X, dist_X) and (Y, dist_Y) are called *quasi-isometric* if and only if there exist separated nets $A \subset X$ and $B \subset Y$, such that (A, dist_X) and (B, dist_Y) are bi-Lipschitz equivalent.

Thus, if we think of a separated net as a *discretization* of a metric space, then quasiisometric spaces are the ones which admit bi-Lipschitz discretizations.

- EXAMPLES 8.2. (1) A non-empty metric space of finite diameter is quasi-isometric to a point.
 (2) The space \mathbb{R}^n endowed with a norm is quasiisometric to \mathbb{Z}^n with the metric induced by that norm.

Historically, quasiisometry was introduced in order to formalize the relationship between some discrete metric spaces (most of the time, groups) and some “non-discrete” (or continuous) metric spaces like for instance Riemannian manifolds, etc. Examples of this is the relationship between finitely generated (or, finitely presented) groups and their *geometric models* introduced in Section 7.9.

When trying to prove that the quasiisometry relation is an equivalence relation, reflexivity and symmetry are straightforward, but, when attempting to prove transitivity, the following question naturally arises:

QUESTION 8.3 (M. Gromov, [Gro93], p. 23). Can a space contain two separated nets that are not bi-Lipschitz equivalent?

Gromov’s question was answered by

THEOREM 8.4 (D. Burago, B. Kleiner, [BK98]). *There exists a separated net N in \mathbb{R}^2 which is not bi-Lipschitz equivalent to \mathbb{Z}^2 .*

Along the same lines, Gromov asked whether two infinite finitely generated groups G and H that are quasiisometric are also bi-Lipschitz equivalent. The negative answer to this question was given by T. Dymarz [Dym10]. We discuss Gromov’s questions in more detail in Chapter 25.

Fortunately, there is a second equivalent way of defining quasiisometry of two metric spaces, based on loosening (coarsifying) the Lipschitz concept. The reader can think of the coarse Lipschitz notion defined below as a generalization of the traditional notion of continuity. Unlike the notion continuity, we will not care about behavior of maps on the small scale, as long as they behave “well” on the large scale.

DEFINITION 8.5. Let X, Y be metric spaces. A map $f : X \rightarrow Y$ is called (L, C) -coarse Lipschitz if

$$(8.1) \quad \text{dist}_Y(f(x), f(x')) \leq L \text{dist}_X(x, x') + C,$$

for all $x, x' \in X$. A map $f : X \rightarrow Y$ is called an (L, C) -quasiisometric embedding if

$$(8.2) \quad L^{-1} \text{dist}_X(x, x') - C \leq \text{dist}_Y(f(x), f(x')) \leq L \text{dist}_X(x, x') + C,$$

for all $x, x' \in X$. Note that a quasiisometric embedding does not have to be an embedding in the usual sense, however distant points have distinct images.

EXAMPLE 8.6. 1. The floor function $f : \mathbb{R} \rightarrow \mathbb{Z} \subset \mathbb{R}$, $f(x) = \lfloor x \rfloor$, is $(0, 1)$ -coarse Lipschitz. This function is a quasiisometric embedding $\mathbb{R} \rightarrow \mathbb{Z}$.

2. The function $f : \mathbb{R} \rightarrow \mathbb{R}$, $f(x) = x^2$, is not coarse Lipschitz.

LEMMA 8.7. *Suppose that G_1, G_2 are finitely generated groups equipped with word-metrics. Then every coarse Lipschitz map $f : G_1 \rightarrow G_2$ is K -Lipschitz for some K .*

PROOF. Let S_1, S_2 be finite generating sets of the groups G_1, G_2 . Suppose that $f : (G_1, \text{dist}_{S_1}) \rightarrow (G_2, \text{dist}_{S_2})$ is (L, C) -coarse Lipschitz. For every $s \in S_1$ and $g \in G_1$, we have

$$\text{dist}_{S_2}(f(g), f(sg)) \leq L + C.$$

Therefore, by the triangle inequalities, for all $g, h \in G_1$,

$$\text{dist}_{S_2}(f(g), f(h)) \leq (L + C)\text{dist}_{S_1}(g, h).$$

Hence, f is K -Lipschitz with $K = L + C$. \square

Nice thing about Lipschitz and coarse Lipschitz maps is that these classes of maps can be recognized *locally*.

LEMMA 8.8. *Consider a map $f : (X, \text{dist}_X) \rightarrow (Y, \text{dist}_Y)$ between metric spaces, where X is a geodesic metric space (but Y is not required to be geodesic). Suppose that r is a positive number such that for all $x, x' \in X$*

$$\text{dist}_X(x, x') \leq r \Rightarrow \text{dist}_Y(f(x), f(x')) \leq A.$$

Then f is $(\frac{A}{r}, A)$ -coarse Lipschitz.

PROOF. For points $x, x' \in X$ consider a geodesic $\gamma \subset X$ connecting x to x' . There exists a finite sequence

$$x_1 = x, x_2, \dots, x_n, x_{n+1} = x'$$

along the geodesic γ , such that

$$r(n-1) \leq D = \text{dist}_X(x, x') < rn, \quad \text{dist}_X(x_i, x_{i+1}) \leq r, i = 1, \dots, n.$$

Then, applying the triangle inequality and the fact that $\text{dist}_Y(f(x_i), f(x_{i+1})) \leq A$, we obtain:

$$\text{dist}_Y(f(x), f(x')) \leq nA \leq \frac{DA}{r} + A = \frac{A}{r}\text{dist}_X(x, x') + A. \quad \square$$

If X is a finite interval $[a, b]$ then an (L, C) -quasiisometric embedding $\mathfrak{q} : X \rightarrow Y$ is called an (L, C) -*quasigeodesic (segment)*. If $a = -\infty$ or $b = +\infty$ then \mathfrak{q} is called an (L, C) -*quasigeodesic ray*. If both $a = -\infty$ and $b = +\infty$, then \mathfrak{q} is called an (L, C) -*quasigeodesic line*. By abuse of terminology, the same names are used for the image of \mathfrak{q} .

In line with loosening the Lipschitz concept, we will also loosen the concept of the inverse map:

DEFINITION 8.9. Maps of metric spaces $f : X \rightarrow Y, \bar{f} : Y \rightarrow X$ are said to be C -coarse inverse to each other if

$$(8.3) \quad \text{dist}_X(\bar{f} \circ f, id_X) \leq C, \quad \text{dist}_Y(f \circ \bar{f}, id_Y) \leq C.$$

In particular, a 0-coarse inverse map is the inverse map in the usual sense. Lastly, we can define quasiisometries:

DEFINITION 8.10. A map $f : X \rightarrow Y$ between metric spaces is called a *quasi-isometry* if it is coarse Lipschitz and admits a coarse Lipschitz coarse inverse map. More precisely, f is an (L, C) -*quasiisometry* if f is (L, C) -coarse Lipschitz and there exists a (L, C) -coarse Lipschitz map $\bar{f} : Y \rightarrow X$ such that the maps f, \bar{f} are C -coarse inverse to each other.

Two metric spaces X, Y are *quasiisometric* if there exists a quasiisometry $X \rightarrow Y$.

A metric space X is called *quasigeodesic* if there exist constants (L, C) so that every pair of points in X can be connected by an (L, C) -quasigeodesic.

Most of the time, the *quasiisometry constants* L, C do not matter, hence, we shall use the words *quasiisometries*, *quasigeodesic* and *quasiisometric embeddings* without specifying the constants. We will frequently abbreviate *quasiisometry*, *quasiisometric* and *quasiisometrically* to *QI*.

- EXERCISE 8.11. (1) Prove that every quasiisometry $f : X \rightarrow Y$ is a quasiisometric embedding.
 (2) Prove that the coarse inverse of a quasiisometry is also a quasiisometry.
 (3) Prove that the composition of two quasiisometric embeddings is a quasiisometric embedding, and that the composition of two quasiisometries is a quasiisometry.
 (4) If $f, g : X \rightarrow Y$ are within finite distance from each other, i.e.

$$\text{dist}(f, g) < \infty,$$

and f is a quasiisometry, then g is also a quasiisometry.

- (5) Let $f_i : X \rightarrow X, i = 1, 2, 3$ be maps such that f_3 is (L_3, A_3) coarse Lipschitz and $\text{dist}(f_2, id_X) \leq A_2$. Then

$$\text{dist}(f_3 \circ f_1, f_3 \circ f_2 \circ f_1) \leq L_3 A_2 + A_3.$$

- (6) Prove that quasiisometry of metric spaces (defined as in Definition 8.5) is an equivalence relation.

EXERCISE 8.12. 1. Suppose that Y and Z are subsets of a metric space (X, dist) such that Z is contained in the r -neighborhood $\mathcal{N}_r(Y)$. Define the “nearest point projection” $\pi_Z : Y \rightarrow Z$, sending each $y \in Y$ to a point $z \in Z$ such that

$$\text{dist}_X(y, z) \leq r.$$

Show that π is a quasiisometric embedding.

2. Suppose that $f : X \rightarrow Y$ is a quasiisometric embedding such that $f(X)$ is r -dense in Y for some $r < \infty$. Show that f is a quasiisometry. *Hint:* Construct a coarse inverse \bar{f} to the map f by mapping a point $y \in Y$ to $x \in X$ such that

$$\text{dist}_Y(f(x), y) \leq r.$$

Maps $f : X \rightarrow Y$ such that $f(X)$ is r -dense in Y for some $r < \infty$, are *coarsely surjective*. Thus, we obtain:

COROLLARY 8.13. *A map $f : X \rightarrow Y$ is a quasiisometry if and only if f is a coarsely surjective quasiisometric embedding.*

EXAMPLE 8.14. The cylinder $X = \mathbb{S}^n \times \mathbb{R}$ with a product metric is quasiisometric to $Y = \mathbb{R}$; the quasiisometry is the projection to the second factor.

EXAMPLE 8.15. Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be an L -Lipschitz function. Then the map

$$f : \mathbb{R} \rightarrow \mathbb{R}^2, \quad f(x) = (x, h(x)),$$

is a QI embedding.

Indeed, f is $\sqrt{1 + L^2}$ -Lipschitz. On the other hand, clearly,

$$\text{dist}(x, y) \leq \text{dist}(f(x), f(y)),$$

for all $x, y \in \mathbb{R}$.

EXAMPLE 8.16. Let $\varphi : [1, \infty) \rightarrow \mathbb{R}_+$ be a differentiable function such that

$$\lim_{r \rightarrow \infty} \varphi(r) = \infty,$$

and there exists $C \in \mathbb{R}$ for which $|r\varphi'(r)| \leq C$ for all r . For instance, take $\varphi(r) = \log(r)$. Define the function $F : \mathbb{R}^2 \setminus B(0, 1) \rightarrow \mathbb{R}^2 \setminus B(0, 1)$ which, in the polar coordinates, takes the form

$$(r, \theta) \mapsto (r, \theta + \varphi(r)).$$

Hence F maps radial straight lines to spirals. Let us check that F is L -bi-Lipschitz for $L = \sqrt{1 + C^2}$. Indeed, the Euclidean metric in the polar coordinates takes the form

$$ds^2 = dr^2 + r^2 d\theta^2.$$

Then

$$F^*(ds^2) = ((r\varphi'(r))^2 + 1)dr^2 + r^2 d\theta^2$$

and the assertion follows. Extend F to the unit disk by the zero map. Therefore, $F : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, is a QI embedding. Since F is onto, it is a quasiisometry $\mathbb{R}^2 \rightarrow \mathbb{R}^2$.

PROPOSITION 8.17. *Two metric spaces (X, dist_X) and (Y, dist_Y) are quasiisometric in the sense of Definition 8.1 if and only if there exists a quasiisometry $f : X \rightarrow Y$.*

PROOF. Assume there exists an (L, C) -quasiisometry $f : X \rightarrow Y$. Let $\delta = L(C + 1)$ and let A be a δ -separated ε -net in X . Then $B = f(A)$ is a 1 -separated $(L\varepsilon + 2C)$ -net in Y . Moreover, for any $a, a' \in A$,

$$\text{dist}_Y(f(a), f(a')) \leq L\text{dist}_X(a, a') + C \leq \left(L + \frac{C}{\delta}\right) \text{dist}_X(a, a')$$

and

$$\text{dist}_Y(f(a), f(a')) \geq \frac{1}{L}\text{dist}_X(a, a') - C \geq \left(\frac{1}{L} - \frac{C}{\delta}\right) \text{dist}_X(a, a') =$$

$$\frac{1}{L(C + 1)} \text{dist}_X(a, a').$$

It follows that f , restricted to A and with target B , is a bi-Lipschitz map.

Conversely, assume that $A \subset X$ and $B \subset Y$ are two ε -separated δ -nets, and that there exists a surjective bi-Lipschitz map $g : A \rightarrow B$. We define a map $f : X \rightarrow Y$ as follows: For every $x \in X$ we choose a point $a_x \in A$ at distance at most δ from x and define $f(x) = g(a_x)$.

REMARK 8.18. The Axiom of Choice makes here yet another important appearance. We will discuss Axiom of Choice in more detail in Chapter 10. Nevertheless, when X is proper (for instance X is a finitely generated group with a word metric), there are finitely many possibilities for the point a_x . Hence, the Axiom of Choice is not required in this situation, as in the finite case it follows from the Zermelo–Fraenkel axioms.

Since $f(X) = g(A) = B$ it follows that Y is contained in the ε -tubular neighborhood of $f(X)$. For every $x, y \in X$,

$$\text{dist}_Y(f(x), f(y)) = \text{dist}_Y(g(a_x), g(a_y)) \leq L \text{dist}_X(a_x, a_y) \leq L(\text{dist}_X(x, y) + 2\varepsilon).$$

Also

$$\text{dist}_Y(f(x), f(y)) = \text{dist}_Y(g(a_x), g(a_y)) \geq \frac{1}{L} \text{dist}_X(a_x, a_y) \geq \frac{1}{L}(\text{dist}_X(x, y) - 2\varepsilon).$$

Now the proposition follows from Exercise 8.12. \square

Below is yet another variation on the definition of quasiisometry, based on relations.

First, some terminology: Given a relation $R \subset X \times Y$, for $x \in X$ let $R(x)$ denote $\{(x, y) \in X \times Y : (x, y) \in R\}$. Similarly, define $R(y)$ for $y \in Y$. Let π_X, π_Y denote the projections of $X \times Y$ to X and Y respectively.

DEFINITION 8.19. Let X and Y be metric spaces. A subset $R \subset X \times Y$ is called an (L, A) -quasiisometric relation if the following conditions hold:

1. Each $x \in X$ and each $y \in Y$ are within distance $\leq A$ from the projection of R to X and Y , respectively.
2. For all $x, x' \in \pi_X(R)$,

$$\text{dist}_{\text{Haus}}(\pi_Y(R(x)), \pi_Y(R(x'))) \leq L \text{dist}(x, x') + A.$$

3. Similarly, for all $y, y' \in \pi_Y(R)$,

$$\text{dist}_{\text{Haus}}(\pi_X(R(y)), \pi_X(R(y'))) \leq L \text{dist}(y, y') + A.$$

Observe that for any (L, A) -quasiisometric relation R , for all pair of points $x, x' \in X$, and $y \in R(x), y' \in R(x')$ we have

$$\frac{1}{L} \text{dist}(x, x') - \frac{A}{L} \leq \text{dist}(y, y') \leq L \text{dist}(x, x') + A.$$

The same inequality holds for all pairs of points $y, y' \in Y$, and $x \in R(y), x' \in R(y')$.

In particular, by using the Axiom of Choice as in the proof of Proposition 8.17, if R is an (L, A) -quasiisometric relation between non-empty metric spaces, then it induces an (L_1, A_1) -quasiisometry $X \rightarrow Y$. Conversely, every (L, A) -quasiisometry is an (L_2, A_2) -quasiisometric relation.

Quasi-isometry group of a space. Some quasiisometries $X \rightarrow X$ are more interesting than others. The *boring* quasiisometries are the ones which are within finite distance from the identity:

DEFINITION 8.20. Given a metric space (X, dist) we denote by $\mathcal{B}(X)$ the set of maps $f : X \rightarrow X$ (not necessarily bijections) which are *bounded perturbations of the identity*, i.e. maps such that

$$\text{dist}(f, \text{id}_X) = \sup_{x \in X} \text{dist}(f(x), x) < \infty.$$

In order to mod out the semigroup of quasiisometries $X \rightarrow X$ by $\mathcal{B}(X)$, one introduces a group $QI(X)$ defined below. Given a metric space (X, dist) , consider the set $QI(X)$ of equivalence classes $[f]$ of quasiisometries $f : X \rightarrow X$, where two quasiisometries f, g are equivalent if and only if

$$\text{dist}(f, g) < \infty.$$

In particular, the set of quasiisometries equivalent to id_X is $\mathcal{B}(X)$. Clearly, the composition is an associative binary operation on $QI(X)$.

EXERCISE 8.21. Show that the coarse inverse defines an inverse in $QI(X)$, and, hence, $QI(X)$ is a group.

DEFINITION 8.22. The group $(QI(X), \circ)$ is called the *group of quasiisometries of the metric space X* . When G is a finitely generated group, then $QI(G)$ will denote the group of quasiisometries of G equipped with the word metric.

Note that if S, S' are two finite generating sets of a group G then the identity map $(G, \text{dist}_S) \rightarrow (G, \text{dist}_{S'})$ is a quasiisometry, see Exercise 7.82.

EXERCISE 8.23. If $h : X \rightarrow X'$ is a quasiisometry of metric spaces, then the groups $QI(X), QI(X')$ are isomorphic; the isomorphism is given by the map

$$[f] \mapsto [h \circ f \circ \bar{h}],$$

where \bar{h} is a coarse inverse to h . Conclude that the group $QI(G)$ is independent of the generating set of G .

More importantly, we will see (Corollary 8.64) that every group quasiisometric to G admits a natural homomorphism to $QI(G)$.

Isometries and virtual isomorphisms. For every metric space X there is a natural homomorphism $q_X : \text{Isom}(X) \rightarrow QI(X)$, given by $f \mapsto [f]$. In general, this homomorphism is not injective. For instance, if $X = \mathbb{R}^n$, then the kernel of q_X is the full group of translations \mathbb{R}^n . Similarly, the entire group $G = \mathbb{Z}^n \times F$, where F is a finite group, maps trivially to $QI(G)$.

Suppose now that G is an arbitrary finitely generated group. Since G acts isometrically on (G, dist_S) (where S a finite generating set of G), we obtain a homomorphism $q_G : G \rightarrow QI(G)$. We will prove in Lemma 16.20 that the kernel K of this homomorphism is a subgroup such that for every $k \in K$ the G -centralizer of k has finite index in G . In particular, if $G = K$ then G is virtually abelian.

The group $VI(G)$ of virtual automorphisms of G defined in Section 5.2 also maps naturally to $QI(G)$. Indeed, suppose that an isomorphism

$$\phi : G_1/K_1 \rightarrow G_2/K_2$$

is a virtual automorphism of G ; here G_1, G_2 are finite-index subgroups of G and $K_i \triangleleft G_i$ are finite normal subgroups, $i = 1, 2$. Then ϕ is a quasiisometry; it lifts to a map

$$\psi : G_1 \rightarrow G_2, \quad \phi(gK_1) = \psi(g)K_2, g \in G_1.$$

We leave it to the reader to verify that ψ is also a quasiisometry. Since G_i 's are finite-index subgroups in G , they are nets in G . Therefore, as in the proof of Proposition 8.17, ψ extends to a quasiisometry

$$f = f_\psi : G \rightarrow G.$$

EXERCISE 8.24. Show that the map defined by $\phi \mapsto [f_\phi]$ is a homomorphism $VI(G) \rightarrow QI(G)$.

Coarse embeddings.

The notion of coarse embedding generalizes the concept of quasi-isometric embedding between two metric spaces. It has been introduced by Gromov in [Gro93, §7.E].

Let $\rho_\pm : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be two continuous functions such that $\rho_-(x) \leq \rho_+(x)$ for every $x \in \mathbb{R}_+$, and such that both functions have limit $+\infty$ at $+\infty$.

DEFINITION 8.25. A (ρ_-, ρ_+) -embedding of a metric space (X, dist_X) into a metric space (Y, dist_Y) is an embedding $\varphi : X \rightarrow Y$ such that

$$(8.4) \quad \rho_-(\text{dist}_X(x, y)) \leq \text{dist}_Y(\varphi(x), \varphi(y)) \leq \rho_+(\text{dist}_X(x, y)).$$

A *coarse embedding* is a (ρ_-, ρ_+) -embedding for some functions ρ_\pm .

Assume now that ρ_- is the inverse of ρ_+ (in particular, both are bijections). A (ρ_-, ρ_+) -transformation of a metric space (X, dist_X) is a bijection $\varphi : X \rightarrow X$ such that both φ and its inverse φ^{-1} satisfy the inequalities in (8.4). Under the same assumptions, if $\rho_+(x) = Lx$ for some $L \geq 1$, the corresponding transformation is an *L-bi-Lipschitz transformation*.

LEMMA 8.26. Let $f : X \rightarrow Y$ be a coarse embedding such that X is geodesic. Then one can take the function $\rho_+(x)$ equal to Lx on $[1, \infty)$, for some $L > 0$.

PROOF. Let ρ_\pm be such that f is a (ρ_-, ρ_+) -embedding, and let L be the supremum of ρ_+ over $[0, 1]$. For every two points a, b in X with $\text{dist}_X(a, b) \geq 1$, consider a finite sequence of consecutive points $x_0 = a, x_1, x_2, \dots, x_n, x_{n+1} = b$ on the same geodesic joining a and b , and such that $\text{dist}_X(x_i, x_{i+1}) = 1$ for $0 \leq i \leq n-1$, and $\text{dist}_X(x_n, x_{n+1}) < 1$. It follows that

$$\text{dist}_Y(f(x_i), f(x_{i+1})) \leq L, \forall 0 \leq i \leq n.$$

Therefore, by the triangular inequality

$$\text{dist}_Y(f(a), f(b)) \leq L(n+1) \leq 2L \text{dist}_X(a, b).$$

□

A coarse embedding of a geodesic metric space is always Lipschitz, therefore in this setting an equivalent notion is the following.

Uniformly proper maps. Let X, Y be topological spaces. Recall that a (continuous) map $f : X \rightarrow Y$ is called *proper* if the inverse image $f^{-1}(K)$ of each compact in Y is a compact in X . The next definition is a “coarsification” of the notion of a proper map:

DEFINITION 8.27. A map $f : X \rightarrow Y$ between proper metric spaces is called *uniformly proper* if f is coarse Lipschitz and there exists a function $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $\text{diam}(f^{-1}(B(y, R))) \leq \psi(R)$ for each $y \in Y, R \in \mathbb{R}_+$. Equivalently, there exists a proper continuous function $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that

$$\text{dist}(f(x), f(x')) \geq \eta(\text{dist}(x, x')).$$

The functions ζ and η are called *upper* and *lower distortion functions* of f respectively.

Clearly, every QI embedding is uniformly proper. Conversely, if f is uniformly proper with linear lower distortion function η , then f is a QI embedding. In Lemma 8.30 we will see how uniformly proper maps appear naturally in group theory.

EXERCISE 8.28. 1. Consider the arc-length parameterization $f : \mathbb{R} \rightarrow \mathbb{R}^2$ of the parabola $y = x^2$. Then f is uniformly proper but is not a QI embedding.

2. The following function is L -Lipschitz, proper, but not uniformly proper:

$$f(x) = (|x|, \arctan(x)), f : \mathbb{R} \rightarrow \mathbb{R}^2.$$

3. The function $\log : (0, \infty) \rightarrow \mathbb{R}$ is not uniformly proper.

4. Composition of uniformly proper maps is again uniformly proper.

5. If $f_1, f_2 : X \rightarrow Y$ are such that $\text{dist}(f_1, f_2) < \infty$ and f_1 is uniformly proper, then so is f_2 .

Even though, uniform properness is weaker than the requirement of a QI embedding, sometimes, the two notions coincide:

LEMMA 8.29. *Suppose that Y is a geodesic metric space, $f : X \rightarrow Y$ is a uniformly proper map whose image is r -dense in Y for some $r < \infty$. Then f is a quasiisometry.*

PROOF. We have to construct a coarse inverse to the map f . Given a point $y \in Y$ pick a point $\bar{f}(y) := x \in X$ such that $\text{dist}(f(x), y) \leq r$. Let us check that \bar{f} is coarse Lipschitz. Since Y is a geodesic metric space it suffices to verify that there is a constant A such that for all $y, y' \in Y$ with $\text{dist}(y, y') \leq 1$, one has:

$$\text{dist}(\bar{f}(y), \bar{f}(y')) \leq A.$$

Pick $t > 2r + 1$ which is in the image of the lower distortion function η . Then take $A \in \eta^{-1}(t)$. Hence, \bar{f} is also coarse Lipschitz. It is also clear that the maps f, \bar{f} are coarse inverse to each other. Hence, f is a quasiisometry. \square

LEMMA 8.30. *Suppose that G is a finitely generated group equipped with word metric and $G \curvearrowright X$ is a properly discontinuous isometric action on a metric space X . Then for every $o \in X$ the orbit map $f : G \rightarrow X$, $f(g) = g \cdot o$, is uniformly proper.*

PROOF. 1. Let S denote the finite generating set of G ; set

$$L = \max_{s \in S} (d_X(s(o), o)).$$

Then for every $g \in G$, $\sin S$, $d_S(gs, g) = 1$, while

$$d_X(gs(o), g(o)) = d_X(s(o), o) \leq L.$$

Therefore, by applying triangle inequalities, we conclude that f is L -Lipschitz.

2. Define the function

$$\eta(n) = \min\{d_X(go, o) : |g| = n\}.$$

Since the action $G \curvearrowright X$ is properly discontinuous,

$$\lim_{n \rightarrow \infty} \eta(n) = \infty.$$

We extend η linearly to unit intervals $[n, n + 1] \subset \mathbb{R}$ and retain the notation η for the extension. The extension $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous and proper. By definition of the function η , for every $g \in G$,

$$d_X(f(g), f(1)) = d_X(go, o) \geq \eta(d_S(g, 1)).$$

Since G acts on itself and on X isometrically, it follows that

$$d_X(f(g), f(h)) \geq \eta(d_S(g, h)), \quad \forall g, h \in G.$$

Thus, the map f is uniformly proper. \square

COROLLARY 8.31. *Let $H \leq G$ is a finitely generated subgroup of a finitely generated group G . Then the inclusion map $H \rightarrow G$ is uniformly proper, where we are using word metrics on G and H associated with their respective finite generating sets.*

We will discuss distortion of subgroups of finitely generated groups in more detail in Section 8.9.

Coarse convergence. So far, we coarsified geometric concepts. Below is a useful coarsification of an analytical concept.

DEFINITION 8.32. Suppose that (Y, d_Y) is a metric space and X is a set. A sequence of maps $f_i : X \rightarrow Y$ is said to *coarsely uniformly converge* to a map $f : X \rightarrow Y$ if there exists $R \in \mathbb{R}_+$ and $i_0 \in \mathbb{N}$ such that for all $i > i_0$ and all $x \in X$,

$$d_Y(f(x), f_i(x)) \leq R.$$

In other words, there exists i_0 such that for all $i > i_0$, $\text{dist}(f, f_i) < \infty$.

Note that the difference with the usual notion of uniform convergence is just one quantifier: $\forall R$ is replaced with $\exists R$.

Similarly, one defines coarse uniform convergence on compact subsets:

DEFINITION 8.33. Suppose that X is a topological. A sequence (f_i) of maps $X \rightarrow Y$ is said to *coarsely uniformly converge to a map $f : X \rightarrow Y$ on compact subsets*, if:

There exists a number $R < \infty$ so that for every compact $K \subset X$, there exists i_K so that for all $i > i_K$,

$$\forall x \in K, \quad d(f_i(x), f(x)) \leq R.$$

We will use the notation

$$\lim_{i \rightarrow \infty}^c f_i = f.$$

to denote the fact that the sequence (f_i) coarsely converges to f .

PROPOSITION 8.34 (Coarse Arzela–Ascoli theorem.). *Fix real numbers L, A and D and let X, Y be proper metric spaces so that X admits a separated R -net. Let $f_i : X \rightarrow Y$ be a sequence of (L_1, A_1) -Lipschitz maps, such that for some points $x_0 \in X, y_0 \in Y$ we have $d(f(x_0), y_0) \leq D$. Then there exists a subsequence (f_{i_k}) , and an (L_2, A_2) -Lipschitz map $f : X \rightarrow Y$, such that*

$$\lim_{k \rightarrow \infty}^c f_{i_k} = f.$$

Furthermore, if the maps f_i are (L_1, A_1) -quasiisometries, then f is also an (L_3, A_3) -quasiisometry.

PROOF. Let $N \subset X$ be a separated net. We can assume that $x_0 \in N$. Then the restrictions $f_i|_N$ are L' -Lipschitz maps and, by the usual Arzela–Ascoli theorem, the sequence $(f_i|_N)$ subconverges (uniformly on compact subsets) to an L' -Lipschitz map $f : N \rightarrow Y$. We extend f to X by the rule:

For $x \in X$ pick $x' \in N$ so that $d(x, x') \leq R$ and set $f(x) := f(x')$.

Then the map $f : X \rightarrow Y$ is (L_2, A_2) -Lipschitz. For a metric ball $B(x_0, r) \subset X$, $r \geq R$, there exists i_r so that for all $i \geq i_r$ and all $x \in N \cap B(x_0, r)$, we have $d(f_i(x), f(x)) \leq 1$. For arbitrary $x \in K$, we find $x' \in N \cap B(x_0, r + R)$ such that $d(x', x) \leq R$. Then

$$d(f_i(x), f(x)) \leq d(f_i(x'), f(x')) \leq L_1(R + 1) + A.$$

This proves coarse convergence. The statement about quasiisometries follows from the Exercise 8.11, part (4). \square

8.2. Group-theoretic examples of quasiisometries

We begin by noting that given a finitely generated group G endowed with a word metric, the set $\mathcal{B}(G)$ is particularly easy to describe. To begin with, it contains all the right translations $R_g : G \rightarrow G$, $R_g(x) = xg$ (see Remark 7.71).

LEMMA 8.35. *For a finitely generated group (G, dist_S) endowed with a word metric, the set of maps $\mathcal{B}(G)$ consists of piecewise right translations. That is, given a map $f \in \mathcal{B}(G)$ there exist finitely many elements h_1, \dots, h_n in G and a decomposition $G = T_1 \sqcup T_2 \sqcup \dots \sqcup T_n$ such that f restricted to T_i coincides with R_{h_i} .*

PROOF. Since $f \in \mathcal{B}(G)$, there exists a constant $R > 0$ such that for every $x \in G$, $\text{dist}(x, f(x)) \leq R$. This implies that $x^{-1}f(x) \in B(1, R)$. The ball $B(1, R)$ is a finite set. We enumerate its distinct elements $\{h_1, \dots, h_n\}$. Thus, for every $x \in G$ there exists h_i such that $f(x) = xh_i = R_{h_i}(x)$ for some $i \in \{1, 2, \dots, n\}$. We define $T_i = \{x \in G ; f(x) = R_{h_i}(x)\}$. If there exists $x \in T_i \cap T_j$ then $f(x) = xh_i = xh_j$, which implies $h_i = h_j$, a contradiction. \square

The main example of a quasiisometry, which partly justifies the interest in such maps, is given by Theorem 8.37, proved in the context of Riemannian manifolds first by A. Schwarz [Šva55] and, 13 years later, by J. Milnor [Mil68b]. At the time, both were motivated by relating volume growth in universal covers of compact Riemannian manifolds and growth of their fundamental groups. Note that sometimes, in the literature it is this theorem (stating the equivalence between the growth function of the fundamental group of a compact manifold and that of the universal cover of the manifold) that is referred to as the Milnor–Schwarz Theorem, and not Theorem 8.37 below.

In fact, it had been observed already by V.A. Efremovich in [Efr53] that two growth functions as above (i.e. of the volume of metric balls in the universal cover of a compact Riemannian manifold, and of the cardinality of balls in the fundamental group with a word metric) increase at the same rate.

REMARK 8.36 (What is in the name?). Schwarz is a German-Jewish name which was translated to Russian (presumably, at some point in the 19-th century) as Шварц. In the 1950-s, the AMS, in its infinite wisdom, decided to translate this name to English as Švarc. A. Schwarz himself eventually moved to the United States and is currently a colleague of the second author at University of California, Davis. See <http://www.math.ucdavis.edu/~schwarz/bion.pdf> for his mathematical autobiography. The transformation

$$\text{Schwarz} \rightarrow \text{Шварц} \rightarrow \check{\text{S}}\text{varc}$$

is a good example of a composition of a quasiisometry and its coarse inverse.

THEOREM 8.37 (Milnor–Schwarz). *Let (X, dist) be a proper geodesic metric space (which is equivalent, by Theorem 2.13, to X being a length metric space which is complete and locally compact) and let G be a group acting geometrically on X . Then:*

- (1) *The group G is finitely generated.*
- (2) *For any word metric dist_W on G and any point $x \in X$, the orbit map $G \rightarrow X$ given by $g \mapsto gx$ is a quasiisometry.*

PROOF. We denote the orbit of a point $y \in X$ by Gy . Given a subset A in X we denote by GA the union of all orbits Ga with $a \in A$.

Step 1: The generating set.

As every geometric action, the action $G \curvearrowright X$ is cobounded: There exists a closed ball \overline{B} of radius D such that $G\overline{B} = X$. Since X is proper, \overline{B} is compact. Define

$$S = \{s \in G : s\overline{B} \cap \overline{B} \neq \emptyset\}.$$

Note that S is finite because the action of G is proper, and that $1 \in S^{-1} = S$ by the definition of S . If $S = G$, then there is nothing to prove; we assume, therefore, that $G \neq S$.

Step 2: Outside of the generating set.

Now consider

$$2d := \inf\{\text{dist}(\overline{B}, g\overline{B}) : g \in G \setminus S\}.$$

Pick $g \in G \setminus S$; the distance $\text{dist}(\overline{B}, g\overline{B})$ is a positive constant R , by the definition of S . The subset $H \subset G$ consisting of elements $h \in G$ such that $\text{dist}(\overline{B}, h\overline{B}) \leq R$, is contained in the set

$$\{g \in G : g\overline{B}(x, D + R) \cap \overline{B}(x, D + R) \neq \emptyset\}$$

and, hence, the subset H is finite. Now,

$$\inf\{\text{dist}(\overline{B}, g\overline{B}) : g \in G \setminus S\} = \inf\{\text{dist}(\overline{B}, g\overline{B}) : g \in H \setminus S\}$$

and the latter infimum is over finitely many positive numbers. Therefore, there exists $h_0 \in H \setminus S$ such that $\text{dist}(\overline{B}, h_0\overline{B})$ realizes that infimum, which is, therefore, positive. By the definition, $\text{dist}(\overline{B}, g\overline{B}) < 2d$ implies that $g \in S$.

Step 3: G is finitely generated.

Consider a geodesic $[x, gx] \subset X$ and define

$$k = \left\lfloor \frac{\text{dist}(x, gx)}{d} \right\rfloor.$$

Then there exists a finite sequence of points on the geodesic $[x, gx]$,

$$y_0 = x, y_1, \dots, y_k, y_{k+1} = gx,$$

such that $\text{dist}(y_i, y_{i+1}) \leq d$ for every $i \in \{0, \dots, k\}$. For every $i \in \{1, \dots, k\}$ let $h_i \in G$ be such that $y_i \in h_i\overline{B}$. We take $h_0 = 1$ and $h_{k+1} = g$. As

$$\text{dist}(\overline{B}, h_i^{-1}h_{i+1}\overline{B}) = \text{dist}(h_i\overline{B}, h_{i+1}\overline{B}) \leq \text{dist}(y_i, y_{i+1}) \leq d,$$

it follows that $h_i^{-1}h_{i+1} = s_i \in S$, that is, $h_{i+1} = h_i s_i$. Then

$$(8.5) \quad g = h_{k+1} = s_0 s_1 \cdots s_k.$$

We have thus proved that G is generated by S , consequently, G is finitely generated.

Step 4: The quasiisometry.

Since all word metrics on G are bi-Lipschitz equivalent, it suffices to prove Part (2) for the word metric dist_S on G , where S is the finite generating set found as above for the chosen point x . The space X is contained in the $2D$ -tubular neighborhood of the image Gx of the orbit map $G \rightarrow X$. It, therefore, remains to prove that the orbit map is a quasiisometric embedding. By the equation (8.5),

$$|g|_S \leq k + 1 \leq \frac{1}{d} \text{dist}(x, gx) + 1.$$

On the other hand, by Lemma 8.30, the orbit map of an isometric properly discontinuous action of a finitely generated group, is L -Lipschitz for some L . Therefore,

$$d|g|_S - d \leq \text{dist}(x, gx) \leq L|g|_S,$$

equivalently,

$$d \cdot \text{dist}_S(1_G, g) - d \leq \text{dist}(x, gx) \leq L \text{dist}_S(1_G, g)$$

Since both the word metric dist_S and the metric dist on X are left-invariant with respect to the action of G , in the above inequality, 1_G can be replaced by any element $h \in G$. \square

EXERCISE 8.38. Verify that the orbit map in this proof is $2D$ -Lipschitz.

COROLLARY 8.39. *Given M a compact connected Riemannian manifold, let \widetilde{M} be its universal cover endowed with the pull-back Riemannian metric, so that the fundamental group $\pi_1(M)$ acts isometrically on \widetilde{M} .*

Then the group $\pi_1(M)$ is finitely generated, and the metric space \widetilde{M} is quasiisometric to $\pi_1(M)$ with some word metric.

Thus, the Milnor-Schwarz Theorem provides an answer to the question about the relation between different geometric models of a finitely generated group G : Different models are *quasiisometric* to each other and to the group G equipped with the word metric.

EXERCISE 8.40. Prove the Milnor-Schwarz Theorem replacing the assumption that X is a geodesic metric space by the hypothesis that X is a *quasigeodesic* metric space.

Our next goal is to prove several corollaries and generalizations of Theorem 8.37.

LEMMA 8.41. *Let (X, dist_i) , $i = 1, 2$, be proper geodesic metric spaces. Suppose that the action $G \curvearrowright X$ is geometric with respect to both metrics $\text{dist}_1, \text{dist}_2$. Then the identity map*

$$\text{Id} : (X, \text{dist}_1) \rightarrow (X, \text{dist}_2)$$

is a quasiisometry.

PROOF. The group G is finitely generated by Theorem 8.37; choose a word metric dist_G on G corresponding to any finite generating set. Pick a point $x_0 \in X$; then the orbit maps

$$f_i : (G, \text{dist}_G) \rightarrow (X, \text{dist}_i), \quad f_i(g) = g(x_0)$$

are quasiisometries, let \bar{f}_i denote their coarse inverses. Then the map

$$\text{Id} : (X, \text{dist}_1) \rightarrow (X, \text{dist}_2)$$

is within finite distance from the quasiisometry $f_2 \circ \bar{f}_1$. \square

COROLLARY 8.42. *Let $\text{dist}_1, \text{dist}_2$ be as in Lemma 8.41. Then any geodesic γ with respect to the metric dist_1 is a quasigeodesic with respect to the metric dist_2 .*

LEMMA 8.43. *Let $G \curvearrowright X$ be a geometric action on a proper geodesic metric space X . Suppose, in addition, that we have an isometric properly discontinuous action $G \curvearrowright X'$ on another metric space X' and a G -equivariant coarsely Lipschitz map $f : X \rightarrow X'$. Then f is uniformly proper.*

PROOF. Pick a point $p \in X$ and set $o := f(p)$. We equip G with a word metric corresponding to a finite generating set S of G ; then the orbit map $\phi : g \mapsto g(p), \phi : G \rightarrow X$ is a quasiisometry by Milnor–Schwarz theorem. We have the second orbit map $\psi : G \rightarrow X', \psi(g) = g(p)$. The map ψ is uniformly proper according to Lemma 8.30. We leave it to the reader to verify that

$$\text{dist}(f \circ \phi, \psi) < \infty.$$

Thus, the map $f \circ \phi$ is uniformly proper as well (see Exercise 8.28). Taking $\bar{\phi} : X \rightarrow G$, a coarse inverse to ϕ , we see that the composition

$$f \circ \phi \circ \bar{\phi}$$

is uniformly proper too. Since

$$\text{dist}(f \circ \phi \circ \bar{\phi}, f) < \infty,$$

we conclude that f is also uniformly proper. \square

Let $G \curvearrowright X, G \curvearrowright X'$ be isometric actions and let $f : X \rightarrow X'$ be a quasiisometric embedding. We say that f is G -*quasiequivariant* if for every $g \in G$

$$\text{dist}(g \circ f, f \circ g) \leq C,$$

where $C < \infty$ is independent of g .

EXERCISE 8.44. (1) Composition of quasiequivariant maps is again quasiequivariant.

(2) If $f : X \rightarrow X'$ is a quasiequivariant quasiisometry, then every coarse inverse $\bar{f} : X' \rightarrow X$ to the map f is also quasiequivariant.

LEMMA 8.45. *Suppose that X, X' are proper geodesic metric spaces, G is a group acting geometrically on X and X' respectively. Then there exists a G -quasiequivariant quasiisometry $f : X \rightarrow X'$.*

PROOF. Pick points $x \in X, x' \in X'$. According to Theorem 8.37, the orbit maps

$$G \rightarrow G \cdot x \hookrightarrow X, \quad G \rightarrow G \cdot x' \hookrightarrow X'$$

are quasiisometries. The statement now follows from the Exercise 8.44. \square

EXERCISE 8.46. Construct an example when in the setting of the lemma there is no quasiisometry $X \rightarrow X'$ which is G -equivariant in the traditional sense, i.e.

$$f \circ g = g \circ f$$

for all $g \in G$.

Below we discuss the relation between quasiisometry and virtual isomorphism, in view of Milnor–Schwarz Theorem.

COROLLARY 8.47. *Let G be a finitely generated group.*

- (1) *If G_1 is a finite index subgroup in G , then G_1 is also finitely generated; moreover the groups G and G_1 are quasiisometric.*
- (2) *Given a finite normal subgroup N in G , the groups G and G/N are quasiisometric.*
- (3) *Thus, two virtually isomorphic (VI) finitely generated groups are quasiisometric (QI).*

PROOF. (1) is a special case of Theorem 8.37, with $G_2 = G$ and X a Cayley graph of G .

(2) follows from Theorem 8.37 applied to the action of the group G on a Cayley graph of the group G/N .

(3) The last part is an immediate consequence of parts (1) and (2). \square

The next example shows that VI is not equivalent to QI.

EXAMPLE 8.48. Let A be a matrix diagonalizable over \mathbb{R} in $SL(2, \mathbb{Z})$ so that $A^2 \neq I$. Thus, the eigenvalues λ, λ^{-1} of A have the absolute value $\neq 1$. We will use the notation $Hyp(2, \mathbb{Z})$ for the set of such matrices. Define the action of \mathbb{Z} on \mathbb{Z}^2 so that the generator $1 \in \mathbb{Z}$ acts by the automorphism given by A . Let G_A denote the associated semidirect product $G_A := \mathbb{Z}^2 \rtimes_A \mathbb{Z}$. We leave it to the reader to verify that \mathbb{Z}^2 is a unique maximal normal abelian subgroup in G_A . By diagonalizing the matrix A , we see that the group G_A embeds as a discrete cocompact subgroup in the Lie group

$$Sol_3 = \mathbb{R}^2 \rtimes_D \mathbb{R}$$

where

$$D(t) = \begin{bmatrix} e^t & 0 \\ 0 & e^{-t} \end{bmatrix}, t \in \mathbb{R}.$$

In particular, G_A is torsion-free. The group Sol_3 has its left-invariant Riemannian metric; hence, G_A acts isometrically on Sol_3 , regarded as a metric space. Therefore, every group G_A as above is QI to Sol_3 . We now construct two groups G_A, G_B of the above type which are not VI to each other. Pick two matrices $A, B \in Hyp(2, \mathbb{Z})$ such that for every $n, m \in \mathbb{Z} \setminus \{0\}$, A^n is not conjugate to B^m . For instance, take

$$A = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}, B = \begin{bmatrix} 3 & 2 \\ 1 & 1 \end{bmatrix}.$$

(The above property of the powers of A and B follows by considering the eigenvalues of A and B and observing that the fields they generate are different quadratic extensions of \mathbb{Q} .) The group G_A is QI to G_B since they are both QI to Sol_3 . Let us check that G_A is not VI to G_B . First, since both G_A, G_B are torsion-free, it suffices to show that they are not commensurable, i.e. do not contain isomorphic finite-index subgroups. Let $H = H_A$ be a finite-index subgroup in G_A . Then H intersects the normal rank 2 free abelian subgroup of G_A along a rank 2 free abelian subgroup L_A . The image of H under the quotient homomorphism $G_A \rightarrow G_A/\mathbb{Z}^2 = \mathbb{Z}$ has to be an infinite cyclic subgroup, generated by some $n \in \mathbb{N}$. Therefore, H_A is isomorphic to $\mathbb{Z}^2 \rtimes_{A^n} \mathbb{Z}$. For the same reason, $H_B \cong \mathbb{Z}^2 \rtimes_{B^m} \mathbb{Z}$. Any isomorphism

$H_A \rightarrow H_B$ has to carry L_A isomorphically to L_B . However, this would imply that A^n is conjugate to B^m . Contradiction.

EXAMPLE 8.49. Another example where QI does not imply VI is as follows. Let S be a closed oriented surface of genus $n \geq 2$. Let $G_1 = \pi_1(S) \times \mathbb{Z}$. Let M be the total space of the unit tangent bundle $UT(S)$ of S . Then the fundamental group $G_2 = \pi_1(M)$ is a non-trivial central extension of $\pi_1(S)$:

$$1 \rightarrow \mathbb{Z} \rightarrow G_2 \rightarrow \pi_1(S) \rightarrow 1,$$

$$G_2 = \langle a_1, b_1, \dots, a_n, b_n, t[a_1, b_1] \cdots [a_n, b_n]t^{2n-2}, [a_i, t], [b_i, t], i = 1, \dots, n \rangle.$$

We leave it to the reader to check that passing to any finite-index subgroup in G_2 does not make it a trivial central extension of the fundamental group of a hyperbolic surface. On the other hand, since the group $\pi_1(S)$ is hyperbolic, the groups G_1 and G_2 are quasiisometric, see Section 11.19.

One more example of quasiisometry (which comes from a virtual isomorphism) is the following:

EXAMPLE 8.50. All non-abelian free groups of finite rank are quasiisometric to each other.

PROOF. We present two proofs: One is algebraic and the other is geometric.

1. **Algebraic proof.** We claim that all free groups $F_n, 2 \leq n < \infty$, are virtually isomorphic. By Proposition 7.80, for every $1 < m < \infty$, the group F_2 contains a finite-index subgroup isomorphic to F_m . Since virtual isomorphism is a transitive relation, which implies quasiisometry, the claim follows.

2. **Geometric proof.** The Cayley graph of F_n with respect to a set of n generators and their inverses is the regular simplicial tree of valence $2n$.

We claim that all regular simplicial trees of valence at least 3 (equipped with the standard metrics) are quasiisometric. Let \mathcal{T}_k denote the regular simplicial tree of valence k ; we will show that \mathcal{T}_3 is quasiisometric to \mathcal{T}_k for every $k \geq 4$.

We construct a countable collection \mathcal{C} of pairwise-disjoint embedded edge-paths c of length $k - 3$ in \mathcal{T}_k , such that every vertex in \mathcal{T}_k belongs to exactly one such path. See Figure 8.1, where the paths $c \in \mathcal{C}$ are drawn in tick lines. Let \mathcal{T} denote the tree obtained from \mathcal{T}_k by collapsing each path $c \in \mathcal{C}$ to a single vertex. We leave it to the reader to verify that the quotient tree \mathcal{T} is isomorphic to the valence k tree \mathcal{T}_k . The quotient map

$$q = q_k : \mathcal{T}_3 \rightarrow \mathcal{T}_k$$

is a morphism of trees. We also leave it to the reader to verify that the map q satisfies the inequality

$$\frac{1}{k-2} \text{dist}(x, y) - 1 \leq \text{dist}(q(x), q(y)) \leq \text{dist}(x, y)$$

for all vertices $x, y \in V(\mathcal{T}_3)$. Therefore, q is a surjective quasiisometric embedding and, hence, a quasiisometry. \square




FIGURE 8.1. All regular simplicial trees are quasiisometric.

8.3. A metric version of the Milnor–Schwarz Theorem

In the case of a Riemannian manifold, or more generally a metric space, without a geometric action of a group, one can still use a purely metric argument and create a discretization of the space, that is a simplicial graph quasiisometric to the space. We begin with a few simple observations.

LEMMA 8.51. *Let X and Y be two discrete metric spaces that are bi-Lipschitz equivalent. If X is uniformly discrete, then so is Y .*

PROOF. Assume $f : X \rightarrow Y$ is an L -bi-Lipschitz bijection, where $L \geq 1$, and assume that $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a function such that for every $r > 0$ every closed ball $\overline{B}(x, r)$ in X contains at most $\phi(r)$ points. Every closed ball $\overline{B}(y, R)$ in Y is in 1-to-1 correspondence with a subset of $B(f^{-1}(y), LR)$, whence it contains at most $\phi(LR)$ points. \square

Notation: Let A be a subset in a metric space. We denote by $\mathcal{G}_\kappa(A)$ the simplicial graph with set of vertices A and set of edges

$$\{(a_1, a_2) \mid a_1, a_2 \in A, 0 < \text{dist}(a_1, a_2) \leq \kappa\}.$$

In other words, $\mathcal{G}_\kappa(A)$ is the 1-skeleton of the Rips complex $\text{Rips}_\kappa(A)$.

As usual, we will equip each component of the graph $\mathcal{G}_\kappa(A)$ with the standard metric.

THEOREM 8.52. (1) *Let (X, dist) be a proper geodesic metric space (equivalently, a complete, locally compact length metric space, see Theorem 2.13). Let $N \subset X$ be an ε -separated δ -net, where $0 < \varepsilon < 2\delta < 1$ and let \mathcal{G} be the metric graph $\mathcal{G}_{\delta\varepsilon}(N)$. Then the graph \mathcal{G} is connected, and the metric space (X, dist) is quasiisometric to the graph \mathcal{G} . More precisely, for all $x, y \in N$ we have*

$$(8.6) \quad \frac{1}{8\delta} \text{dist}_X(x, y) \leq \text{dist}_{\mathcal{G}}(x, y) \leq \frac{3}{\varepsilon} \text{dist}_X(x, y).$$

- (2) *If, moreover, (X, dist) is either a complete Riemannian manifold of bounded geometry or a metric simplicial complex of bounded geometry, then \mathcal{G} is a graph of bounded geometry (see Definition 3.33).*

PROOF. (1) Our proof is modeled on the one of Milnor–Schwarz Theorem. Let x, y be two points in N . If $\text{dist}_X(x, y) \leq 8\delta$ then, by construction, $\text{dist}_{\mathcal{G}}(x, y) = 1$ and both inequalities in (8.6) hold. Let us suppose that $\text{dist}_X(x, y) > 8\delta$.

The distance $\text{dist}_{\mathcal{G}}(x, y)$ is the length s of an edge-path $e_1 e_2 \dots e_s$, where x is the tail vertex of e_1 and y is the head vertex of e_s . It follows that

$$\text{dist}_{\mathcal{G}}(x, y) = s \geq \frac{1}{8\delta} \text{dist}_X(x, y).$$

The distance $\text{dist}_X(x, y)$ is the length of a geodesic $\mathbf{c}: [0, \text{dist}_X(x, y)] \rightarrow X$. Let

$$t_0 = 0, t_1, t_2, \dots, t_m = \text{dist}_X(x, y)$$

be a sequence of numbers in $[0, \text{dist}_X(x, y)]$ such that $5\delta \leq t_{i+1} - t_i \leq 6\delta$, for every $i \in \{0, 1, \dots, m-1\}$.

Let $x_i = \mathbf{c}(t_i)$, $i \in \{0, 1, 2, \dots, m\}$. For every $i \in \{0, 1, 2, \dots, m\}$ there exists $w_i \in N$ such that $\text{dist}_X(x_i, w_i) \leq \delta$. We note that $w_0 = x, w_m = y$. The choice of t_i implies that

$$3\delta \leq \text{dist}_X(w_i, w_{i+1}) \leq 8\delta, \quad \text{for every } i \in \{0, \dots, m-1\}$$

In particular:

- w_i and w_{i+1} are the endpoints of an edge in \mathcal{G} , for every $i \in \{0, \dots, m-1\}$;
- $\text{dist}_X(x_i, x_{i+1}) \geq \text{dist}(w_i, w_{i+1}) - 2\delta \geq \text{dist}(w_i, w_{i+1}) - \frac{2}{3} \text{dist}(w_i, w_{i+1}) = \frac{1}{3} \text{dist}(w_i, w_{i+1})$.

We can then write

$$(8.7) \quad \text{dist}_X(x, y) = \sum_{i=0}^{m-1} \text{dist}_X(x_i, x_{i+1}) \geq \frac{1}{3} \sum_{i=0}^{m-1} \text{dist}(w_i, w_{i+1}) \geq \frac{\varepsilon}{3} m \geq \frac{\varepsilon}{3} \text{dist}_{\mathcal{G}}(x, y).$$

This inequality implies both connectivity of the graph \mathcal{G} and the required quasi-isometry estimates.

(2) According to the Example 3.34, the graph \mathcal{G} has bounded geometry if and only if its set of vertices with the induced simplicial distance is uniformly discrete. Lemma 8.51 implies that it suffices to show that the set of vertices of \mathcal{G} (i.e. the net N) with the metric induced from X is uniformly discrete.

When X is a Riemannian manifold, this follows from Lemma 3.31. When X is a simplicial complex this follows from the fact that the set of vertices of X is uniformly discrete. \square

Note that one can also discretize a Riemannian manifold M (i.e. of replace M by a quasiisometric simplicial complex) using Theorem 3.36, which implies:

THEOREM 8.53. *Every Riemannian manifold M of bounded geometry is quasi-isometric to a bounded geometry simplicial complex homeomorphic to M .*

8.4. Topological coupling

In this section we describe an alternative, and sometimes useful, *dynamical* criterion for quasiisometry between groups. This alternative definition then motivates the notion of *measure-equivalence* between groups. Neither notion will be used elsewhere in the book.

We begin with Gromov’s interpretation of quasiisometry between groups using the language of topological actions.

A *topological coupling* of topological groups G_1, G_2 , is a metrizable locally compact topological space X , together with two commuting cocompact properly discontinuous topological actions $\rho_i : G_i \rightarrow \text{Homeo}(X), i = 1, 2$. (The actions commute if and only if $\rho_1(g_1)\rho_2(g_2) = \rho_2(g_2)\rho_1(g_1)$ for all $g_i \in G_i, i = 1, 2$.) Note that the actions ρ_i are not required to be isometric. Below we will see some natural examples of topological couplings.

The following theorem was first proven by Gromov in [Gro93]; see also [dlH00, page 98].

THEOREM 8.54. *If G_1, G_2 are finitely generated groups, then G_1 is QI to G_2 if and only if there exists a topological coupling between these discrete groups.*

PROOF. 1. Suppose that G_1 is QI to G_2 . Then there exists an (L, A) -quasiisometry $\mathfrak{q} : G_1 \rightarrow G_2$. The map \mathfrak{q} is $(L + A)$ -Lipschitz (see Lemma 8.7). Consider the space X of (L, A) -quasiisometric maps $G_1 \rightarrow G_2$. We equip X with the topology of pointwise convergence. By the Arzela–Ascoli Theorem, X is locally compact.

The groups G_1, G_2 act on X as follows:

$$\rho_1(g_1)(f) := f \circ g_1^{-1}, \quad \rho_2(g_2)(f) := g_2 \circ f, \quad f \in X.$$

It is clear that these actions commute and are topological. For each (L, A) -quasiisometry $f \in X$, there exist $g_1 \in G_1, g_2 \in G_2$ such that

$$g_2 \circ f(1_{G_1}) = 1_{G_2}, f \circ g_1^{-1}(1_{G_2}) \in B(1, A) \subset G_2.$$

Therefore, by the Arzela–Ascoli Theorem, both actions (of G_1 and of G_2) are cocompact. We will check that ρ_2 is properly discontinuous as the case of ρ_1 is analogous. Let $K \subset X$ be a compact subset. Then there exists $R < \infty$ so that for every $f \in K$,

$$f(1_{G_1}) \in B(1, R).$$

If $g_2 \in G_2$ is such that $g_2 \circ f \in K$ for some $f \in K$, then

$$(8.8) \quad g_2(B(1_{G_2}, R)) \cap B(1_{G_2}, R) \neq \emptyset.$$

Since the action of G_2 on itself is free, it follows that the collection of $g_2 \in G_2$ satisfying (8.8) is finite. Hence, ρ_2 is properly discontinuous.

Lastly, the space X is metrizable, since it is locally compact, second countable and Hausdorff; more explicitly, one can define distance between functions as the Gromov–Hausdorff distance between their graphs. (Note that this metric is G_1 -invariant.)

2. Suppose that X is a topological coupling of G_1 and G_2 . If X were a geodesic metric space and the actions of G_1, G_2 were isometric, we would not need commutation of these action (as Milnor–Schwarz Theorem would apply). However, there are examples of QI groups which do not act geometrically on the same geodesic

metric space, see Theorem 8.37. Nevertheless, the construction of a quasiisometry below is pretty much the same as in the proof of the Milnor-Schwarz Theorem.

Since $G_i \curvearrowright X$ is cocompact, there exists a compact $K \subset X$ so that $G_i K = X$; pick a point $p \in K$. Then for each $g_i \in G_i$ there exists $\phi_i(g_i) \in G_{i+1}$ such that $g_i(p) \in \phi_i(g_i)(K)$; here and below i is taken modulo 2. We, thus, have the maps $\phi_i : G_i \rightarrow G_{i+1}$, $i = 1, 2$.

a. Let us check that these maps are Lipschitz. Let $s \in S_i$, a finite generating set of G_i , we will use the word metric on G_i with respect to S_i , $i = 1, 2$. Define C to be the union

$$\bigcup_{s \in S_1} s(K) \cup \bigcup_{s \in S_2} s(K).$$

Since ρ_i 's are properly discontinuous actions, the sets

$$G_i^C := \{h \in G_i : h(C) \cap C \neq \emptyset\}, i = 1, 2,$$

are finite. Therefore, the word-lengths of the elements of these sets are bounded by some $L < \infty$. Suppose now that $g_{i+1} = \phi_i(g_i)$, $s \in S_i$. Then

$$g_i(p) \in g_{i+1}(K), sg_i(p) \in g'_{i+1}(K),$$

for some $g'_{i+1} \in G_{i+1}$. Therefore,

$$sg_{i+1}(K) \cap g'_{i+1}(K) \neq \emptyset$$

and, hence,

$$g_{i+1}^{-1}g'_{i+1}(K) \cap s(K) \neq \emptyset.$$

(This is where we are using the fact that the actions of G_1 and G_2 on X commute.) Therefore, $g_{i+1}^{-1}g'_{i+1} \in G_{i+1}^C$, which implies that $d(g_{i+1}, g'_{i+1}) \leq L$. Consequently, ϕ_i is L -Lipschitz.

b. Set

$$\phi_i(g_i) = g_{i+1}, \phi_{i+1}(g_{i+1}) = g'_i.$$

Then $g_i(K) \cap g'_i(K) \neq \emptyset$ and, hence, $g_i^{-1}g'_i \in G_i^C$. Therefore, we conclude that

$$\text{dist}(\phi_{i+1} \circ \phi_i, \text{Id}_{G_i}) \leq L,$$

and, thus, the maps ϕ_1, ϕ_2 are coarse inverse to each other. Thus, $\phi_1 : G_1 \rightarrow G_2$ is a quasiisometry. \square

The more useful direction of this theorem is, of course, from QI to a topological coupling, see e.g. [Sha04, Sau06].

DEFINITION 8.55. Two groups G_1, G_2 are said to have a *common geometric model* if there exists a proper quasigeodesic metric space X such that G_1, G_2 both act geometrically on X .

In view of Theorem 8.37, if two groups have a common geometric model then they are quasiisometric. The following theorem shows that the converse is false:

THEOREM 8.56 (L. Mosher, M. Sageev, K. Whyte, [MSW03]). *Consider the groups*

$$G_1 := \mathbb{Z}_p * \mathbb{Z}_p, \quad G_2 := \mathbb{Z}_q * \mathbb{Z}_q,$$

where p, q are distinct odd primes. Then the groups G_1, G_2 are quasiisometric (since they are virtually isomorphic to the free group on two generators) but do not have a common geometric model.

This theorem, in particular, implies that in Theorem 8.54 one cannot assume that both group actions are isometric (for the same metric).

Measure–equivalence. The interpretation of quasiisometry of groups in terms of topological couplings was generalized by M. Gromov [Gro93] in the measure-theoretic context:

DEFINITION 8.57. A *measurable coupling* for two groups G_1, G_2 is a measure space (Ω, μ) such that G_1, G_2 admit commuting measure-preserving free actions on (Ω, μ) , which both admit a finite measure fundamental set in (Ω, μ) . Groups G_1, G_2 are called *measure-equivalent* if they admit a measurable coupling.

We refer the reader to [Gab05, Gab10] for further discussion of this fruitful concept.

8.5. Quasiactions

The notion of an *action* of a group on a space is frequently replaced, in the context of quasiisometries, by the one of a *quasiaction*. Recall that an *action* of a group G on a set X is a homomorphism $\phi : G \rightarrow \text{Aut}(X)$, where $\text{Aut}(X)$ is the group of bijections $X \rightarrow X$. Since quasiisometries are defined only up to “bounded error”, the concept of a homomorphism has to be modified when we use quasiisometries.

DEFINITION 8.58. Let G be a group and X be a metric space. An (L, A) -quasiaction of G on X is a map $\phi : G \rightarrow \text{Map}(X, X)$, such that:

- $\phi(g)$ is an (L, A) -quasiisometry of X for all $g \in G$.
- $d(\phi(1_G), \text{Id}_X) \leq A$.
- $d(\phi(g_1 g_2), \phi(g_1)\phi(g_2)) \leq A$ for all $g_1, g_2 \in G$.

By abusing the notation, we will denote quasiactions by $\phi : G \curvearrowright X$, even though, what we have is not an action.

The last two conditions can be informally summarized as: ϕ is “almost” a homomorphism with the error A .

Similarly, a *quasihomomorphism* from a group to another group equipped with a left-invariant metric is a map

$$\phi : G_1 \rightarrow (G_2, \text{dist})$$

which satisfies properties (2) and (3) of a quasiaction with respect to the metric dist on G_2 (the property (1) is automatic since G_1 will quasiact via isometries on G_2).

EXAMPLE 8.59. Suppose that G is a group and $\phi : G \rightarrow \mathbb{R} \subset \text{Isom}(\mathbb{R})$ is a function. Then ϕ , of course, satisfies (1), while properties (2) and (3) are equivalent to the single condition:

$$|\phi(g_1 g_2) - \phi(g_1) - \phi(g_2)| \leq A.$$

In other words, such maps ϕ are *quasimorphisms*, see Definition 5.119.

We refer the reader to [FK16] for the discussion of quasihomomorphisms with noncommutative targets.

We can also define proper discontinuity and cocompactness for quasiactions by analogy with isometric actions:

DEFINITION 8.60. Let $\phi : G \curvearrowright X$ be a quasiaction.

1. We say that ϕ is *properly discontinuous* if for every $x \in X, R \in \mathbb{R}_+$, the set

$$\{g \in G \mid d(x, \phi(g)(x)) \leq R\}$$

is finite. Note that if X proper and ϕ is an isometric action, this definition is equivalent to proper discontinuity of the action $\phi : G \curvearrowright X$.

2. We say that ϕ is *cobounded* if there exists $x \in X, R \in \mathbb{R}_+$ such that for every $x' \in X$ there exists $g \in G$, for which $d(x', \phi(g)(x)) \leq R$.

3. Lastly, we say that a quasiaction ϕ is *geometric* if it is both properly discontinuous and cobounded.

EXERCISE 8.61. Let $QI(X)$ denote the group of (equivalence classes of) quasiisometries $X \rightarrow X$. Show that every quasiaction $G \curvearrowright X$ determines a homomorphism $\widehat{\phi} : G \rightarrow QI(X)$ given by composing ϕ with the projection to $QI(X)$.

The *kernel* of the quasiaction $\phi : G \curvearrowright X$ is the kernel of the homomorphism $\widehat{\phi}$.

EXERCISE 8.62. Construct an example of a geometric quasiaction $G \curvearrowright \mathbb{R}$ whose kernel is the entire group G .

Below we explain how quasiactions appear in the context of QI rigidity problems. Suppose that G_1, G_2 are groups, $\psi_i : G_i \curvearrowright X_i$ are isometric actions; for instance, X_i could be G_i or its Cayley graph. Suppose that $f : X_1 \rightarrow X_2$ is a quasiisometry with coarse inverse \bar{f} . We then define a *conjugate* quasiaction $\phi = f^*(\psi_2)$ of G_2 on X_1 by

$$(8.9) \quad \phi(g) = \bar{f} \circ g \circ f.$$

More generally, we say that two quasiactions $\psi_i : G \curvearrowright X_i$ are (quasi) *conjugate* if there exists a quasiisometry $f : X_1 \rightarrow X_2$, such that ψ_1 and $f^*(\psi_2)$ project to the same homomorphism

$$G \rightarrow QI(X_1).$$

LEMMA 8.63. *Suppose that $\psi : G \curvearrowright X_2$ is a quasiaction, $f : X_1 \rightarrow X_2$ is a quasiisometry and $\phi = f^*(\psi)$ is defined by the formula (8.9). Then:*

1. $\phi = f^*(\psi)$ is a quasiaction of G on X_1 .

2. If ψ is *properly discontinuous* (respectively, *cobounded*, or *geometric*), then so is ϕ .

PROOF. 1. Suppose that f is an (L, A) -quasiisometry with coarse inverse \bar{f} . In view of Exercise 8.11, it is clear that ϕ satisfies Parts 1 and 2 of the definition of a quasiaction; we only have to verify Part (3):

$$\text{dist}(\phi(g_1g_2), \phi(g_1)\phi(g_2)) = \text{dist}(\bar{f}g_1g_2f, \bar{f}g_1f\bar{f}g_2f) \leq LA + A.$$

2. We will verify the statement about properly discontinuous quasiactions, since the proof for cobounded quasiactions is similar. Pick $x \in X, R \in \mathbb{R}_+$, and consider the subset

$$G_{x,R} = \{g \in G \mid d(x, \phi(g)(x)) \leq R\} \subset G.$$

By the definition, $\phi(g)(x) = \bar{f}gf(x)$. Thus, $d(x, \phi(g)(x)) \leq LR + 2A$. Hence, by proper discontinuity of the action $\psi : G \curvearrowright X_2$, the set $G_{x,R}$ is finite. \square

COROLLARY 8.64. *Let G_1 and G_2 be finitely generated quasiisometric groups and let $f : G_1 \rightarrow G_2$ be a quasiisometry. Then:*

1. *The quasiisometry f induces (by conjugating actions and quasiactions on G_2) an isomorphism $QI(f) : QI(G_2) \rightarrow QI(G_1)$ and a homomorphism $f_* : G_2 \rightarrow QI(G_1)$*

2. *The homomorphism f_* is quasiinjective: For every $K \geq 0$, the set of $g \in G_2$ such that $\text{dist}(f_*(g), \text{Id}_{G_1}) \leq K$, is finite.*

PROOF. The isomorphism $QI(f) : QI(G_2) \rightarrow QI(G_1)$ is defined by the formula (8.9). The inverse to this homomorphism is defined by switching the roles of f and \bar{f} . We leave it to the reader to verify that $QI(f)$ is an isomorphism. To define f_* we compose the homomorphism $G_2 \rightarrow QI(G_2)$ with $QI(f)$. Quasi-injectivity of f_* follows from the proper discontinuity of the action $G_2 \curvearrowright G_2$ by the left multiplication. \square

REMARK 8.65. For many groups $G = G_1$, if $h : G \rightarrow G$ is an (L, A) -quasiisometry which belongs to $\mathcal{B}(G)$, we also have $\text{dist}(f, \text{Id}_G) \leq D(L, A)$, where $D(L, A)$ depends only on L, A and (G, d_S) but not on f . For instance, this holds when G is a non-elementary hyperbolic group, see Lemma 11.112. This is also true for isometry groups of irreducible symmetric spaces and Euclidean buildings and many other spaces, see e.g. [KKL98]. In this situation, the kernel of f_* above is actually finite.

The following theorem is a weak converse to the construction of a conjugate quasiaction:

THEOREM 8.66 (B. Kleiner, B. Leeb, [KL09]). *Suppose that $\phi : G \curvearrowright X_1$ is a quasiaction. Then there exists a metric space X_2 , a quasiisometry $f : X_1 \rightarrow X_2$ and an isometric action $\psi : G \curvearrowright X_2$, such that f conjugates ψ to ϕ .*

Thus, every quasiaction is conjugate to an isometric action, but, *a priori*, on a different metric space. The main issue of the QI (quasiisometric) rigidity, discussed in the next section is:

Can one, under some conditions, take $X_2 = X_1$? More precisely: Given a quasiaction $G \curvearrowright X$ of a group G on a space $X = X_1$, can one find a conjugate isometric action $G \curvearrowright X$?

8.6. Quasi-isometric rigidity problems

So far, we succeeded in converting finitely generated groups into metric spaces, i.e. *treating groups as geometric objects*. All these spaces are quasiisometric to each other, but we would like to *reconstruct* (to the extent possible) the group G , as an *algebraic object*, from its geometric models (defined only up to a quasiisometry). In other words, we would like to know, to which extent the “geometrization map”

$geo : \text{Finitely generated groups} \rightarrow \text{metric spaces/quasiisometry}$

is injective?

Corollary 8.47 establishes a limitation on injectivity of *geo*: Virtually isomorphic groups are quasiisometric to each other. Therefore, the best we can hope for, is to recover a group from its (coarse) geometry up to virtual isomorphisms.

DEFINITION 8.67. 1. A (finitely generated) group G is called *QI rigid* if every group G' which is quasiisometric to G is, in fact, virtually isomorphic to G .

2. A group G is called *strongly QI rigid* if the natural map $Comm(G) \rightarrow QI(G)$ (from the group of virtual automorphisms of G to the group of self-quasiisometries of G) is surjective.

3. A subclass \mathcal{G} of the class of all (finitely generated) groups is called *QI rigid* if each group G which is quasiisometric to a member of \mathcal{G} , is virtually isomorphic to a member of \mathcal{G} .

4. A group G in a subclass \mathcal{G} (of all groups) is *QI rigid within \mathcal{G}* if any $G' \in \mathcal{G}$ which is quasiisometric to G , is virtually isomorphic to G .

In the purely geometric context, one can ask if a quasiisometry between metric spaces is within finite distance from an isometry of these spaces:

DEFINITION 8.68. 1. A metric space X is called *strongly QI rigid* if the natural map $Isom(X) \rightarrow QI(X)$ is surjective.

2. A more quantitative version of this property is the *uniform QI rigidity*: A space X is *uniformly QI rigid* if there exists $D(X, L, A) \in \mathbb{R}_+$ such that every (L, A) -quasiisometry $X \rightarrow X$ is within distance $\leq D(X, L, A)$ from an isometry $X \rightarrow X$.

3. More restrictively, one talks about QI rigidity *within a subclass \mathcal{M}* of the class of all metric spaces, by requiring that any two quasiisometric spaces in \mathcal{M} are, in fact, isometric.

A *QI rigidity theorem* is a theorem which establishes QI rigidity in the sense of any of the above definitions.

Most proofs of QI rigidity theorems proceed along the following route:

1. Suppose that the groups G_1, G_2 are quasiisometric. Find a “nice space” X_1 on which G_1 acts geometrically. Take a quasiisometry $f : X_1 \rightarrow X_2 = G_2$, where $\psi : G_2 \curvearrowright G_2$ is the action by the left multiplication.

2. Define the conjugate quasiaction $\phi = f^*(\psi)$ of G_2 on X_1 .

3. Show that the quasiaction ϕ has finite kernel (or, at least, identify the kernel, prove that it is, say, abelian).

4. Extend, if necessary, the quasiaction $G_2 \curvearrowright X_1$ to a quasiaction $\widehat{\phi}$ on a larger space \widehat{X}_1 .

5. Show that $\widehat{\phi}$ has the same projection to $QI(\widehat{X}_1)$ as a isometric action $\phi' : G_2 \curvearrowright \widehat{X}_1$ by verifying, for instance, that \widehat{X}_1 has very few quasiisometries, namely, every quasiisometry of X is within finite distance from an isometry. (Well, maybe not all quasiisometries of \widehat{X}_1 , but the ones which extend from X_1 .) Then conclude either that $G_2 \curvearrowright \widehat{X}_1$ is geometric, or, that the isometric actions of G_1, G_2 are commensurable, i.e. the images of G_1, G_2 in $Isom(\widehat{X}_2)$ have a common subgroup of finite index.

We will see how R. Schwartz’s proof of QI rigidity for non-uniform lattices follows this line of arguments: X_1 will be a truncated hyperbolic space and \widehat{X}_1 will be the hyperbolic space itself. The same is true for QI rigidity of higher rank non-uniform lattices (A. Eskin’s theorem [Esk98]). This is also true for uniform lattices in the isometry groups of non-positively curved symmetric spaces other than \mathbb{H}^n and $\mathbb{C}\mathbb{H}^n$ (P. Pansu, [Pan89], B. Kleiner and B. Leeb [KL98b]; A. Eskin and B. Farb [EF97b]), except one does not have to enlarge X_1 . Another example of such argument is the proof by M. Bourdon and H. Pajot [BP00] and X. Xie

[Xie06] of QI rigidity of groups acting geometrically on 2-dimensional hyperbolic buildings.

5'. Part 5 may fail if X has too many quasiisometries, e.g. if $X_1 = \mathbb{H}^n$ or $X_1 = \mathbb{C}\mathbb{H}^n$. Then, instead, one shows that every geometric quasiaction $G_2 \curvearrowright X_1$ is quasiconjugate to a geometric (isometric!) action. We will see such a proof in the case of the Sullivan–Tukia rigidity theorem for uniform lattices in $\text{Isom}(\mathbb{H}^n)$, $n \geq 3$. Similar arguments apply in the case of groups quasiisometric to the hyperbolic plane.

Not all quasiisometric rigidity theorems are proven in this fashion. An alternative route is to show QI rigidity of a certain algebraic property (P) is to show that it is equivalent to some geometric property (P'), which is QI invariant. Examples of such proofs are QI rigidity of the class of virtually nilpotent groups and of virtually free groups. The first property is equivalent, by Gromov's theorem, to the polynomial growth. The argument in the second case is less direct (see Theorem 20.45), but the key fact is that the geometric condition of having infinitely many ends is equivalent to the algebraic condition that a group splits (as a graph of groups) over a finite subgroup.

8.7. The growth function

Suppose that X is a discrete metric space (see Definition 2.3) and $x \in X$ is a base-point. We define the *growth function*

$$\mathfrak{G}_{X,x}(R) := \text{card } \bar{B}(x, R),$$

the cardinality of the closed R -ball centered at x . Similarly, given a connected simplicial complex X or a graph (equipped with the standard metric) and a vertex v as a base-point, the *growth function* of X is the growth function of its set of vertices with the base-point v .

We refer the reader to Notation 1.4 for the equivalence relation \asymp between functions used below.

LEMMA 8.69 (Equivalence class of growth is QI invariant.). *If (X, x_0) and (Y, y_0) are quasiisometric uniformly discrete pointed spaces, then $\mathfrak{G}_{X,x_0} \asymp \mathfrak{G}_{Y,y_0}$.*

PROOF. Let $f : X \rightarrow Y, \bar{f} : Y \rightarrow X$ be L -Lipschitz maps which are coarse inverse to each other (see Definition 8.5). We assume that f, \bar{f} satisfy

$$L^{-1}d(x, x') - A \leq d(f(x), f(x')), \quad L^{-1}d(y, y') - A \leq d(\bar{f}(y), \bar{f}(y')).$$

Let $D = \max(d(f(x_0), y_0), d(x_0, \bar{f}(y_0)))$. Then for each $R > 0$,

$$f(\bar{B}(x_0, R)) \subset \bar{B}(y_0, LR + D), \quad \bar{f}(\bar{B}(y_0, R)) \subset \bar{B}(x_0, LR + D),$$

while $f(x) = f(x')$ implies $d(x, x') \leq AL$. The same applies to the map \bar{f} . Since the spaces X and Y are uniformly discrete, both maps f, \bar{f} have multiplicity $\leq m$, where m is an upper bound for the cardinalities of closed LA -balls in X and Y . It follows that

$$\text{card } \bar{B}(x_0, R) \leq m \text{card } \bar{B}(y_0, LR + D)$$

and

$$\text{card } \bar{B}(y_0, R) \leq m \text{card } \bar{B}(x_0, LR + D). \quad \square$$

COROLLARY 8.70. $\mathfrak{G}_{X,x} \asymp \mathfrak{G}_{X,x'}$ for all $x, x' \in X$.

EXERCISE 8.71. Prove that the lemma and the corollary also hold for simplicial complexes and graphs of bounded geometry.

Henceforth we will suppress the choice of the base-point in the notation for the growth function.

EXERCISE 8.72. Show that for each (uniformly discrete) space X , $\mathfrak{G}_X(R) \preceq e^R$.

For a group G endowed with the word metric dist_S corresponding to a finite generating set S we sometimes will use the notation $\mathfrak{G}_S(R)$ for $\mathfrak{G}_G(R)$. Since G acts transitively on itself, this function does not depend on the choice of a base-point.

EXAMPLES 8.73. (1) If $G = \mathbb{Z}^k$ then $\mathfrak{G}_S \asymp x^k$ for every finite generating set S .

(2) If $G = F_k$ is the free group of finite rank $k \geq 2$ and S is the set of k generators then

$$\mathfrak{G}_S(n) = 1 + (q^n - 1) \frac{q+1}{q-1}, \quad q = 2k - 1.$$

EXERCISE 8.74. (1) Prove the two statements above.

(2) Conclude that \mathbb{Z}^m is quasiisometric to \mathbb{Z}^n if and only if $n = m$. (Cf. Lemma 8.69.)

(3) Compute the growth function for the group \mathbb{Z}^2 equipped with the generating set x, y , where $\{x, y\}$ is a basis of \mathbb{Z}^2 .

(4) Prove that for every $n \geq 2$ the group $SL(n, \mathbb{Z})$ has exponential growth.

PROPOSITION 8.75. (1) If S, S' are two finite generating sets of G then $\mathfrak{G}_S \asymp \mathfrak{G}_{S'}$. Thus one can speak about the growth function \mathfrak{G}_G of a group G , well defined up to the equivalence relation \asymp .

(2) If G is infinite, $\mathfrak{G}_S|_{\mathbb{N}}$ is strictly increasing.

(3) The growth function is sub-multiplicative:

$$\mathfrak{G}_S(r+t) \leq \mathfrak{G}_S(r)\mathfrak{G}_S(t).$$

(4) For each finitely generated group G , $\mathfrak{G}_G(r) \preceq 2^r$.

PROOF. (1) follows immediately from Lemma 8.69 and Milnor–Schwarz theorem.

(2) Consider two integers $n < m$. As G is infinite there exists $g \in G$ at distance $d \geq m$ from 1. The shortest path joining 1 and g in $\text{Cayley}(G, S)$ can be parameterized as an isometric embedding $p : [0, d] \rightarrow \text{Cayley}(G, S)$. The vertex $p(n+1)$ is an element of $\bar{B}(1, m) \setminus \bar{B}(1, n)$.

(3) follows immediately from the fact that

$$\bar{B}(1, n+m) \subseteq \bigcup_{y \in \bar{B}(1, n)} \bar{B}(y, m).$$

(4) follows from the existence of an epimorphism $\pi_S : F(S) \rightarrow G$, where S is a finite generating set of G . \square

The property (3) implies that the function $\ln \mathfrak{G}_S(n)$ is sub-additive, hence by the Fekete's Lemma, see e.g. [HP74, Theorem 7.6.1], there exists a (finite) limit

$$\lim_{n \rightarrow \infty} \frac{\ln \mathfrak{G}_S(n)}{n}.$$

Hence, we also get a finite limit

$$\gamma_S = \lim_{n \rightarrow \infty} \mathfrak{G}_S(n)^{\frac{1}{n}},$$

called *growth constant*. The property (2) implies that $\mathfrak{G}_S(n) \geq n$; whence, $\gamma_S \geq 1$.

DEFINITION 8.76. If $\gamma_S > 1$ then G is said to be of *exponential growth*. If $\gamma_S = 1$ then G is said to be of *sub-exponential growth*.

Note that by Proposition 8.75, (1), if there exists a finite generating set S such that $\gamma_S > 1$ then $\gamma_{S'} > 1$ for every other finite generating set S' . Likewise for equality to 1.

The notion of subexponential growth makes sense for (some classes of) general metric spaces.

DEFINITION 8.77. Let (X, dist) be a metric space for which the growth function is defined (e.g. a Riemannian manifold equipped with the Riemannian distance function, a discrete proper metric space, a locally finite simplicial complex). The space X is said to be of *sub-exponential growth* if for some basepoint $x_0 \in X$

$$\limsup_{n \rightarrow \infty} \frac{\ln \mathfrak{G}_{x_0, X}(n)}{n} = 0.$$

Since for every other basepoint y_0 , $\mathfrak{G}_{y_0, X}(n) \leq \mathfrak{G}_{x_0, X}(n + \text{dist}(x_0, y_0))$, it follows that the definition is independent of the choice of basepoint.

PROPOSITION 8.78. (a) *If H is a finitely generated subgroup in a finitely generated group G then $\mathfrak{G}_H \preceq \mathfrak{G}_G$.*

(b) *If H is a subgroup of finite index in G then $\mathfrak{G}_H \asymp \mathfrak{G}_G$.*

(c) *If N is a normal subgroup in G then $\mathfrak{G}_{G/N} \preceq \mathfrak{G}_G$*

(d) *If N is a finite normal subgroup in G then $\mathfrak{G}_{G/N} \asymp \mathfrak{G}_G$.*

PROOF. (a) If X is a finite generating set of H and S is a finite generating set of G containing X then $\text{Cayley}(H, X)$ is a subgraph of $\text{Cayley}(G, S)$ and $\text{dist}_X(1, h) \geq \text{dist}_S(1, h)$ for every $h \in H$. In particular the closed ball of radius r and center 1 in $\text{Cayley}(H, X)$ is contained in the closed ball of radius r and center 1 in $\text{Cayley}(G, S)$.

(b) and (d) are immediate corollaries of Lemma 8.69 and the Milnor–Schwarz theorem.

(c) Let S be a finite generating set in G , and let $\bar{S} = \{sN \mid s \in S, s \notin N\}$ be the corresponding finite generating set in G/N . The epimorphism $\pi : G \rightarrow G/N$ maps the ball of center 1 and radius r onto the ball of center 1 and radius r . □

Let G and H be two groups with finite generating sets S and X , respectively. A homomorphism $\varphi : G \rightarrow H$ is called *expanding* if there exist constants $\lambda > 1$ and $C \geq 0$ such that for every $g \in G$ with $|g|_S \geq C$

$$|\varphi(g)|_X \geq \lambda |g|_S.$$

Such homomorphisms generalize the notion of Euclidean similarities, which expand lengths of all vectors by a fixed constant.

EXERCISE 8.79. Let G be a group with a finite generating set S and $H \leq G$ a finite-index subgroup. We equip G with the word metric d_S and equip H with the metric which is the restriction of d_S . Assume that there exists an expanding homomorphism $\varphi : H \rightarrow G$ such that $\varphi(H)$ has finite index in G . Prove *Franks' Lemma*, that such group G has polynomial growth.

More importantly, one has the following generalization of Efremovich's theorem [Efr53]:

PROPOSITION 8.80 (Efremovich–Schwarz–Milnor). *Let M be a connected complete Riemannian manifold with bounded geometry. If M is quasiisometric to a graph \mathcal{G} with bounded geometry, then the growth function \mathfrak{G}_{M,x_0} and the growth function of \mathcal{G} with respect to an arbitrary vertex v , are equivalent in the sense of the equivalence relation \asymp .*

PROOF. The manifold M has bounded geometry, therefore its sectional curvature is at least a and at most b for some constants $a \leq b$; moreover, there exists a uniform lower bound $2\rho > 0$ on the injectivity radius of M at every point. Let n denote the dimension of M . We let $V(x, r)$ denote volume of r -ball centered at the point $x \in M$ and let $V_a(r)$ denote the volume of the r -ball in the complete simply-connected n -dimensional manifold of constant curvature a .

The fact that the sectional curvature is at least a implies, by Theorem 3.23, Part (1), that for every $r > 0$, $V(x, r) \leq V_a(r)$. Similarly, Theorem 3.23, Part (2), implies that the volume $V(x, \rho) \geq V_b(\rho)$.

Since M and \mathcal{G} are quasiisometric, by Definition 8.1 it follows that there exist $L \geq 1$, $C \geq 0$, two $2C$ -separated nets A in M and B in \mathcal{G} , respectively, and a L -bi-Lipschitz bijection $\mathfrak{q} : A \rightarrow B$. Without loss of generality we may assume that $C \geq \rho$; otherwise we choose a maximal 2ρ -separated subset A' of A and then restrict \mathfrak{q} to A' .

According to Remark 3.15, (2), we may assume without loss of generality that the base-point x_0 in M is contained in the net A , and that $\mathfrak{q}(x_0) = v$, the base vertex in \mathcal{G} .

For every $r > 0$ we have that

$$\begin{aligned} \mathfrak{G}_{M,x_0}(r) &\geq \text{card} [A \cap B_M(x_0, r - C)] V_b(\rho) \geq \text{card} \left[B \cap B_{\mathcal{G}} \left(1, \frac{r - C}{L} \right) \right] V_b(\rho) \\ &\geq \mathfrak{G}_{\mathcal{G}} \left(\frac{r - C}{L} \right) \frac{V_b(\rho)}{\mathfrak{G}_{\mathcal{G}}(2C)}. \end{aligned}$$

Conversely,

$$\begin{aligned} \mathfrak{G}_{M,x_0}(r) &\leq \text{card} [A \cap B_M(x_0, r + 2C)] V_a(2C) \leq \\ &\text{card} [B \cap B_{\mathcal{G}}(1, L(r + 2C))] V_a(2C) \leq \mathfrak{G}_{\mathcal{G}}(L(r + 2C)) V_a(2C). \end{aligned}$$

□

Thus, it follows from Theorem 7.34 that considering \asymp -equivalence classes of growth functions of universal covers of compact Riemannian manifolds is not different from considering equivalence classes of growth functions of finitely presented groups.

REMARK 8.81. Note that in view of Theorem 8.52, every connected Riemannian manifold of bounded geometry is quasiisometric to a graph of bounded geometry.

QUESTION 8.82. What is the set $Growth(groups)$ of the equivalence classes of growth functions of finitely generated groups?

QUESTION 8.83. What are the equivalence classes of growth functions for finitely presented groups?

This question is equivalent to

QUESTION 8.84. What is the set $Growth(manifolds)$ of equivalence classes of growth functions for universal covers of compact connected Riemannian manifolds?

Clearly, $Growth(manifolds) \subset Growth(groups)$. This inclusion is proper since R. Grigorchuk [Gri84a] proved that there exist uncountably many nonequivalent growth functions of finitely generated groups, while there are only countably many nonisomorphic finitely presented groups.

We will see later on that:

$$\{\exp(t), t^n, n \in \mathbb{N}\} \subset Growth(manifolds) \subset Growth(groups)$$

One can refine Question 8.84 by defining $Growth_n(manifolds)$ as the set of equivalence classes of growth functions of universal covers of n -dimensional compact connected Riemannian manifolds. Since every finitely presented group is the fundamental group of a closed smooth 4-dimensional manifold and growth function depends only on the fundamental group, we obtain:

$$Growth_4(manifolds) = Growth_n(manifolds), \quad \forall n \geq 4.$$

On the other hand:

THEOREM 8.85. $Growth_2(manifolds) = \{1, t^2, e^t\}$, $Growth_3(manifolds) = \{1, t, t^3, t^4, e^t\}$.

Below is an outline of the proof. Firstly, in view of classification of surfaces, for every closed connected oriented surface S we have:

- (1) If $\chi(S) = 2$ then $\pi_1(S) = \{1\}$ and growth function is trivial.
- (2) If $\chi(S) = 0$ then $\pi_1(S) = \mathbb{Z}^2$ and growth function is equivalent to t^2 .
- (3) If $\chi(S) < 0$ then $\pi_1(S)$ contains a free nonabelian subgroup, so growth function is exponential.

In the case of 3-dimensional manifolds, one has to appeal to Perelman's Geometrization Theorem. We refer to [Kap01] for the precise statement and definitions which appear below:

For every closed connected 3-dimensional manifold M one of the following holds:

- (1) M admits a Riemannian metric of constant positive curvature, in which case $\pi_1(M)$ is finite and has trivial growth.

- (2) M admits a Riemannian metric locally isometric to the product metric $\mathbb{S}^2 \times \mathbb{R}$. In this case growth function is linear.
- (3) M admits a flat Riemannian metric, so universal cover of M is isometric to \mathbb{R}^3 and growth function is t^3 .
- (4) M is homeomorphic to the quotient H_3/Γ , where H_3 is the 3-dimensional Heisenberg group and Γ is a uniform lattice in H_3 . In this case, in view of Exercise 14.4, growth function is t^4 .
- (5) The fundamental group of M is solvable but not virtually nilpotent, thus, by Wolf's Theorem (theorem 14.30), the growth function is exponential.
- (6) In all other cases, $\pi_1(M)$ contains free nonabelian subgroup; hence, its growth is exponential.

QUESTION 8.86 (J. Milnor [Mil68b]). Is it true that the growth of a finitely generated group is either polynomial (i.e. $\mathfrak{G}_S(t) \preceq t^d$ for some integer d) or exponential (i.e. $\gamma_S > 1$)?

R. Grigorchuk in [Gri83] (see also [Gri84a, Gri84b]) proved that Milnor's question has negative answer, by constructing finitely generated groups of *intermediate growth*, i.e. their growth is superpolynomial but subexponential. More precisely, Grigorchuk proved that for every sub-exponential function f there exists a group G_f of intermediate growth equipped with a finite generating set S_f whose growth function $\mathfrak{G}_{S_f}(n)$ is larger than $f(n)$ for infinitely many n . A. Erschler in [Ers04] adapted Grigorchuk's arguments to show that for every such function f , a direct product $G_f \times G_f$, equipped with the generating set $S = S_f \sqcup S_f$, has the growth function $\mathfrak{G}_S(n)$ satisfying $\mathfrak{G}_S(n) \geq f(n)$ for all but finitely many n .

The first explicit computations of growth functions (up to the equivalence relation \asymp) some groups of intermediate growth were done by L. Bartholdi and A. Erschler in [BE12]. For every $k \in \mathbb{N}$, they construct examples of torsion groups G_k and of torsion-free groups H_k such that their growth functions satisfy

$$\mathfrak{G}_{G_k}(x) \asymp \exp\left(x^{1-(1-\alpha)^k}\right),$$

and

$$\mathfrak{G}_{H_k}(x) \asymp \exp\left(\log x \cdot x^{1-(1-\alpha)^k}\right),$$

Here, α is the number satisfying $2^{3-\frac{3}{\alpha}} + 2^{2-\frac{2}{\alpha}} + 2^{1-\frac{1}{\alpha}} = 2$.

We note that all currently known groups of intermediate growth have growth larger than $2^{\sqrt{n}}$. Existence of finitely presented groups of intermediate growth is unknown. In particular the currently known examples of groups of intermediate growth do not answer Question 8.84.

8.8. Codimension one isoperimetric inequalities

One can define, in the setting of graphs, the following concepts, inspired by, and closely connected to, notions introduced in Riemannian geometry (see Definitions 3.19 and 3.21). Recall that for a subset $F \subset V$, F^c denotes its complement in V .

DEFINITION 8.87. An *isoperimetric inequality* in a graph \mathcal{G} of bounded geometry is an inequality satisfied by all finite subsets F of vertices, of the form

$$\text{card}(F) \leq f(F)g(\text{card } E(F, F^c)),$$

where f and g are real-valued functions, g defined on \mathbb{R}_+ .

DEFINITION 8.88. Let Γ be a graph of bounded geometry, with the vertex set V and edge set E . The *Cheeger constant* or the *Expansion Ratio* of the graph Γ is defined as

$$h(\Gamma) = \inf \left\{ \frac{|E(F, F^c)|}{|F|} : F \text{ is a finite non-empty subset of } V, |F| \leq \frac{|V|}{2} \right\}.$$

Here $E(F, F^c)$ is edge boundary for both F and F^c , i.e. the set of edges connecting F to F^c (see Definition 1.43). Thus, the condition $|F| \leq \frac{|V|}{2}$ insures that, in case V is finite, one picks the smallest of the two sets F and F^c in the definition of the Cheeger constant. Intuitively, finite graphs with small Cheeger constant can be separated by vertex sets which are relatively small comparing to the size of (the smallest component of) their complements. In contrast, graphs with large Cheeger constant are “hard to separate.”

EXERCISE 8.89. a. Let Γ be a single circuit with n vertices. Then $h(\Gamma) = \frac{2}{n}$.

b. Let $\Gamma = K_n$ be the complete graph on n vertices, i.e. Γ is the 1-dimensional skeleton of the $n - 1$ -dimensional simplex. Then

$$h(\Gamma) = \left\lfloor \frac{n}{2} \right\rfloor.$$

The inequalities in (1.4) imply that in every isoperimetric inequality, the edge-boundary can be replaced by the vertex boundary, if one replaces the function g by an asymptotically equal function (respectively the Cheeger constants by bi-Lipschitz equivalent values). Therefore, in what follows we choose freely whether to work with the edge-boundary or with the vertex-boundary, depending on which one is more convenient.

There exists an isoperimetric inequality satisfied in every Cayley graph of an infinite group.

PROPOSITION 8.90. *Let \mathcal{G} be the Cayley graph of a finitely generated infinite group. For every finite set F of vertices*

$$(8.10) \quad \text{card}(F) \leq [\text{diam}(F) + 1] \text{card}(\partial_V F).$$

PROOF. Assume that \mathcal{G} is the Cayley graph of an infinite group G with respect to a finite generating set S .

Let d be the diameter of F with respect to the word metric dist_S , and let g be an element in G such that $|g|_S = d + 1$. Let $g_0 = 1, g_1, g_2, \dots, g_d, g_{d+1} = g$ be the set of vertices on a geodesic joining 1 to g .

Given an arbitrary vertex $x \in F$, the element xg is at distance $d + 1$ from x ; therefore, by the definition of d it follows that $xg \in F^c$. In the finite sequence of vertices $x, xg_1, xg_2, \dots, xg_d, xg_{d+1} = xg$ consider the largest i such that $xg_i \in F$. Then $i < d + 1$ and $xg_{i+1} \in F^c$, whence $xg_{i+1} \in \partial_V F$, equivalently, $x \in [\partial_V F] g_{i+1}^{-1}$.

We have thus proved that $F \subseteq \bigcup_{i=1}^{d+1} [\partial_V F] g_i^{-1}$, which implies the inequality (8.10). \square

An argument similar in spirit, but more elaborate, allows to relate isoperimetric inequalities and growth functions:

PROPOSITION 8.91 (Coulhon-Saloff-Coste inequality). *Let \mathcal{G} be the Cayley graph of an infinite group G with respect to a finite generating set S , and let d be the cardinality of S . For every finite set F of vertices*

$$(8.11) \quad |F| \leq 2d k \text{card}(\partial_V F),$$

where k is the unique integer such that $\mathfrak{G}_S(k-1) \leq 2|F| < \mathfrak{G}_S(k)$.

PROOF. Our goal is to show that with the given choice of k , there exists an element $g \in B_S(1, k)$ such that for a certain fraction of the vertices x in F , the right-translates xg are in F^c . In what follows we omit the subscript S in our notation.

We consider the sum

$$\begin{aligned} \mathcal{S} &= \frac{1}{\mathfrak{G}(k)} \sum_{g \in B(1, k)} \text{card} \{x \in F \mid xg \in F^c\} = \frac{1}{\mathfrak{G}(k)} \sum_{g \in B(1, k)} \sum_{x \in F} \mathbf{1}_{F^c}(xg) = \\ &= \frac{1}{\mathfrak{G}(k)} \sum_{x \in F} \sum_{g \in B(1, k)} \mathbf{1}_{F^c}(xg) = \frac{1}{\mathfrak{G}(k)} \sum_{x \in F} \text{card} [B(x, k) \setminus F]. \end{aligned}$$

By the choice of k , the cardinality of each ball $B(x, k)$ is larger than $2|F|$, whence

$$\text{card} [B(x, k) \setminus F] \geq |F|.$$

The denominator $\mathfrak{G}(k) \leq d\mathfrak{G}(k-1) \leq 2d|F|$. We, therefore, find as a lower bound for the sum \mathcal{S} , the value

$$\frac{1}{2d|F|} \sum_{x \in F} |F| = \frac{|F|}{2d}.$$

It follows that

$$\frac{1}{\mathfrak{G}(k)} \sum_{g \in B(1, k)} \text{card} \{x \in F \mid xg \in F^c\} \geq \frac{|F|}{2d}.$$

The latter inequality implies that there exists $g \in B(1, k)$ such that

$$\text{card} \{x \in F \mid xg \in F^c\} \geq \frac{|F|}{2d}.$$

We now argue as in the proof of Proposition 8.90, and for the element $g \in B(1, k)$ thus found, we consider $g_0 = 1, g_1, g_2, \dots, g_{m-1}, g_m = g$ to be the set of vertices on a geodesic joining 1 to g , where $m \leq k$. The set $\{x \in F \mid xg \in F^c\}$ is contained in the union $\bigcup_{i=1}^m [\partial_V F] g_i^{-1}$; therefore, we obtain

$$\frac{|F|}{2d} \leq k |\partial_V F|.$$

□

REMARKS 8.92. Proposition 8.91 was initially proved in [VSCC92] for nilpotent groups using random walks. The proof reproduced above follows [CSC93].

COROLLARY 8.93. *Let G be an infinite finitely generated group and let F be an arbitrary set of elements in G .*

(1) *If $\mathfrak{G}_G \asymp x^n$ then*

$$|F| \leq K [\text{card} (\partial_V F)]^{\frac{n}{n-1}}.$$

(2) *If $\mathfrak{G}_G \asymp \exp(x)$ then*

$$\frac{|F|}{\ln(\text{card } F)} \leq K \text{card} (\partial_V F).$$

In both inequalities above, the boundary $\partial_V F$ is considered in the Cayley graph of G with respect to a finite generating set S , and K depends on S .

8.9. Distortion of a subgroup in a group

So far, we were primarily interested in quasiisometries and quasiisometric embeddings. In this section we will consider coarse Lipschitz maps which fail to be quasiisometric embeddings and our goal is to quantify failure of the quasiisometric embedding property. While our primary interest comes from finitely generated subgroups of finitely generated groups, we start with general definitions.

DEFINITION 8.94. Let $f : Y \rightarrow X$ be a coarse Lipschitz map. The *distortion* Δ_f of the map f is defined as

$$\Delta_f(t) = \sup\{\text{dist}_Y(y, y') : \text{dist}_X(f(y), f(y')) \leq t\}.$$

Note that the function Δ_f , in general, takes infinite values. It is clear from the definition, that f is uniformly proper if and only if Δ_f takes values in \mathbb{R} . It is also clear that f is a quasiisometric embedding if and only if Δ_f is bounded above by a linear function.

EXAMPLE 8.95. The function $f : \mathbb{R} \rightarrow \mathbb{R}$, $f(y) = \sqrt{y}$ has quadratic distortion.

An interesting special case of the distortion function is when X is a path-metric space, Y is a rectifiable connected subspace of X , equipped with the induced path-metric and f is the identity embedding:

$$\text{dist}(y, y') = \inf_{\gamma} \text{length}_X(\gamma)$$

where the infimum is taken over all paths in Y connecting y to y' , while the length of these paths is computed using the metric of X . In this setting, the distortion function of f is denoted Δ_X^Y . (Note that f itself is 1-Lipschitz.)

EXERCISE 8.96. 1. Compute the distortion of the parabola $y = x^2$ in the Euclidean plane.

2. Compute the distortion of the cubic curve $y = x^3$ in the Euclidean plane.

3. Composing $f : X \rightarrow Y$ with quasiisometries $X \rightarrow X'$ and $Y' \rightarrow Y$, preserves the \asymp equivalence class of the distortion function.

We now specialize to the group-theoretic setting. Suppose that G is a finitely generated group and $H \leq G$ is a finitely generated subgroup; we let S be a finite generating set of G and T be a finite generating set of H . We then have word-metrics dist_S on G and dist_T on H , and the identity embedding $f : H \rightarrow G$. In order to analyze the distortion of H in G (up to the \asymp equivalence relation), we are free to choose the generating sets S and T (see Part 3 of Exercise 8.96); in particular, we can assume that $T \subset S$ and, hence, $\text{Cayley}(H, T)$ is a subgraph of $\text{Cayley}(G, S)$. Since H acts transitively on itself via left multiplication, we obtain that

$$(8.12) \quad \Delta_G^H(n) = \max\{\text{dist}_T(1, h) \mid h \in H, \text{dist}_S(1, h) \leq n\}.$$

The subgroup H is called *undistorted* (in G) if $\Delta_G^H(n) \asymp n$, equivalently, the inclusion map $H \rightarrow G$ is a quasiisometric embedding.

In general, distortion functions for subgroups can be as bad as one can imagine, for instance, nonrecursive.

EXAMPLE 8.97. [Mikhailova's construction] Let Q be a finitely presented group with Dehn function $\delta(n)$. Let a_1, \dots, a_m be generators of Q and $\phi : F_m \rightarrow Q$ be

the epimorphism from the free group of rank m sending free generators of F_m to the elements a_i , $i = 1, \dots, m$. Consider the group $G = F_m \times F_m$ and its subgroup

$$H = \langle (g_1, g_2) \in G \mid \phi(g_1) = \phi(g_2) \rangle.$$

This construction of H is called *Mikhailova's construction*, it is a source of many pathological examples in group theory. The subgroup H is finitely generated and its distortion in G is $\asymp \delta(n)$. In particular, if Q has unsolvable word problem then its distortion in G is nonrecursive. We refer the reader to [OS01, Theorem 2] for further details.

Below are the basic properties of the distortion function:

PROPOSITION 8.98. (1) If \tilde{X} and \tilde{S} are finite generating sets of H and G , respectively, and $\tilde{\Delta}_G^H$ is the distortion function with respect to these generating sets, then $\tilde{\Delta}_G^H \asymp \Delta_G^H$. Thus up to the equivalence relation \asymp , the distortion function of the subgroup H in the group G is uniquely defined by H and G .

(2) For every finitely generated subgroup H in a finitely generated group G , $\Delta_G^H(n) \succeq n$.

(3) If H has finite index in G then $\Delta_G^H(n) \asymp n$.

(4) Let $K \triangleleft G$ is a finite normal subgroup and let $H \leq G$ be a finitely generated subgroup; set $\bar{G} := G/K$, $\bar{H} := H/K$. Then

$$\Delta_G^H \asymp \Delta_{\bar{G}}^{\bar{H}}.$$

(5) If $K \leq H \leq G$ then

$$\Delta_G^K \preceq \Delta_H^K \circ \Delta_G^H.$$

(6) Subgroups of finitely generated abelian groups are undistorted.

PROOF. (1) follows from Part 3 of Exercise 8.96.

(2) If we take finite generating sets S and T of G and H , respectively such that $T \subset S$, then the embedding $H \rightarrow G$ is 1-Lipschitz with respect to the resulting word metrics. Whence $\Delta_G^H(n) \geq n$.

(3) The statement follows immediately from the fact that the inclusion map $H \rightarrow G$ is a quasiisometry.

(4) This equivalence follows from the fact that the projections $G \rightarrow \bar{G}$ and $H \rightarrow \bar{H}$ are quasiisometries.

(5) Consider dist_K , dist_H and dist_G three word metrics, and an arbitrary element $k \in K$ such that $\text{dist}_G(1, k) \leq n$. Then $\text{dist}_H(1, k) \leq \Delta_G^H(n)$ whence

$$\text{dist}_K(1, k) \leq \Delta_H^K(\Delta_G^H(n)).$$

(6) By the classification theorem of finitely generated abelian groups (Theorem 13.7), every subgroup $H \leq G$ of an abelian group G is isomorphic to the direct product of a finite group and free abelian group. In particular, every finitely generated abelian group is virtually torsion-free. Therefore, by combining (3) and (5), it suffices to consider the case where G is torsion-free of rank n . Then G acts by translations geometrically on \mathbb{R}^n ; its rank m subgroup H also acts geometrically on a subspace $\mathbb{R}^m \subset \mathbb{R}^n$. Since \mathbb{R}^m is isometrically embedded in \mathbb{R}^n , it follows that the embedding $H \rightarrow G$ is quasiisometric. Hence, H is undistorted in G and $\Delta_G^H(n) \asymp n$. \square

CHAPTER 9

Coarse topology

So far, we succeeded in coarsifying Riemannian manifolds and groups, while treating metric spaces *up to quasiisometry*. The trouble is that, in a way, we succeeded all too well, and, seemingly, lost all the topological tools in the process. Indeed, quasiisometries lack continuity and uniformly discrete spaces have very boring (discrete) topology. The goal of this chapter is to describe tools of algebraic topology for studying quasiisometries and other concepts of the Geometric Group Theory. We will see how to define *coarse topological invariants* of metric spaces, which are robust enough to be stable under quasiisometries. The price we have to pay for this stability is that we will be forced to work not with simplicial/cell complexes and their (co)homology groups as it is done in algebraic topology, but with direct/inverse systems of such complexes and groups.

In this chapter we also introduce *metric cell complexes with bounded geometry*, which will provide a class of spaces for which application of algebraic topology (in the coarse setting) is possible.

Note that the coarse algebraic topology invariants defined and used in this chapter and in this book are quite basic (homology, coarse separation, Poincaré duality).

QUESTION 9.1. Are there any interesting coarse topology applications of other invariants of algebraic topology?

9.1. Ends

In this section we review the oldest coarse topological notion, the one of *ends* of a topological space. Even though we are primarily interested in coarse topology of metric spaces, we will also define ends in the more general, topological, setting. We refer the reader to [BH99] and [Geo08] for a more detailed treatment of ends of spaces.

9.1.1. The number of ends. We begin with the motivation. One of the simplest topological invariants of a space X is the number of its connected components or, more precisely, the cardinality of the set of its connected components. Alternatively, one can use the set $\pi_0(X)$ of path-connected components of X . Suppose, however, we are dealing with a connected (or path-connected) topological space. The next topological invariant one can try, is the number of connected components of complements to points or, more generally, finite subsets, of X . For instance, if one space can be disconnected by a point and the other cannot, then the two spaces are not homeomorphic. In the coarse setting (of metric spaces) a point is undistinguishable from a bounded subset. Therefore, one naturally looks for complementary components of bounded subsets, say, metric balls.

REMARK 9.2. In the topological setting, metric balls will be replaced with compact subsets. In order to maintain consistency between the two notions (metric and topological), we will later restrict to proper geodesic metric spaces on the metric side and locally compact, locally path-connected Hausdorff topological spaces on the topological side.

The trouble is that, say, a point might fail to disconnect a metric space, while a larger bounded (or compact) subset, might disconnect X . Moreover, some complementary components C of a bounded subset might be bounded themselves and, hence, such C “disappears” if we remove a larger bounded subset from X . Such bounded complementary components should be discarded, of course. This leads to the first, *numerical*, definition below, which suffices for many purposes. In what follows, for a subset B in X , B^c will denote the complement of B in X . For each closed subset $B \subset X$ we define the set $\pi_0^u(B^c) := \pi_0(U_B)$, where U_B is the union of unbounded path-connected components of B^c . (The letter u stays for *unbounded*). In the topological setting, being unbounded, of course, makes no sense. Thus, for a Hausdorff topological space X , we let U_B denote the union of path-components of $B^c = X \setminus B$ which are not relatively compact in X . We retain the notation $\pi_0^u(B^c)$ for the set $\pi_0(U_B)$.

From now on, let X be non-empty, locally compact, connected, locally path-connected, second countable Hausdorff topological space, e.g., a proper geodesic metric space.

DEFINITION 9.3. The *number of ends* of X is the supremum, taken over all compact subsets $K \subset X$, of cardinalities of $\pi_0^u(K^c)$. We will denote the number of ends of X by $\eta(X)$.

The reader has to be warned at this points that we, eventually, will define a certain set $\epsilon(X)$, called the *set of ends*, of the space X . The cardinality of this set equals $\eta(X)$ if either one of them is finite; in the infinite case, $\text{card}(\epsilon(X)) \geq \eta(X)$. In the group-theoretic setting, $\epsilon(X)$ will have the cardinality of continuum, once it is infinite. Nevertheless, what we will really care about, as far as groups are concerned, is finiteness or infiniteness of the number of ends. Thus, the distinction between $\eta(X)$ and $\text{card}(\epsilon(X))$ will not be that important.

According to our definition, X has zero number of ends iff X is compact; X has one end (is *one-ended*) iff X is non-compact and for each compact $K \subset X$, the complement K^c has exactly one unbounded component. The space X is *disconnected at infinity* iff X has at least two ends. The space X has infinitely many ends iff for every $n \in \mathbb{N}$, there exists a compact $K \subset X$ such that K^c has at least n unbounded complementary components.

It is clear that the number of ends is a topological invariant of X . Note also that for any compact subsets $K_1 \subset K_2 \subset X$ we have

$$\text{card}(\pi_0^u(K_2^c)) \geq \text{card}(\pi_0^u(K_1^c)).$$

In particular, in the definition of the number of ends of a proper geodesic metric space, we can take the supremum of cardinalities $\pi_0^u(B^c)$ over all metric balls in X ; equivalently, over all bounded subsets of X .

- EXERCISE 9.4. (1) The real line \mathbb{R} is 2-ended.
 (2) \mathbb{R}^n is one-ended for $n \geq 2$.

- (3) Suppose that X is a regular tree of finite valence $k \geq 3$. Then X has infinitely many ends.

The proof of the following lemma is a model for many arguments appearing in this chapter.

LEMMA 9.5. *The number of ends $\eta(X)$ is a quasiisometry invariant of X .*

PROOF. Let $f : X \rightarrow Y$ be an (L, A) -quasiisometry of (proper, geodesic) metric spaces. Suppose that $\eta(X) \geq n$, $n \in \mathbb{N}$. This means that there exists a metric ball $B = B(x, R) \subset X$ such that B^c consists of at least n unbounded components. The image of a bounded subset under quasiisometry is again bounded, while the image of an unbounded complementary component C is still unbounded. The trouble is that $f(C)$, of course, may fail to be connected and be contained in the complement of $f(B)$; moreover, images of distinct complementary components under f might be contained in the same complementary component of $f(B)$.

We will deal with these three problems one at a time. Consider another metric ball $B' = B(x, R')$, $R' \geq R$.

1. If $R' - R \geq t$, where $L^{-1}t - A > 0$, then for each component C' of $(B')^c$, its image $f(C')$ is disjoint from $f(B)$. (Thus, it suffices to take $R' > R + AL$.)
2. If $x_1, x_2 \in X$ are within distance ≤ 1 from each other, then

$$\text{dist}_Y(f(x_1), f(x_2)) \leq r := L + A.$$

Therefore, the r -neighborhood $\mathcal{N}_r(f(C'))$ of $f(C')$ in Y will be path-connected. In order for this neighborhood to be disjoint from $f(B)$, we need to increase R' a little bit: It suffices to take t such that $L^{-1}t - A > r$, i.e. $R' > R + L(A + r)$.

3. The last issue we have to address is slightly more difficult: So far, we only used the fact that f is a QI embedding. Considering the example of an isometric embedding of the line into the plane, we see what can go wrong without the assumption of coarse surjectivity of f . Suppose that C_1, C_2 are distinct unbounded components of B and $x_i \in C_i, i = 1, 2$, are points which are mapped to points $y_i = f(x_i)$ which are in the same complementary (path-connected) component of $\text{cl}(f(B))^c$. Pick a path p connecting y_1, y_2 and avoiding $\text{cl}(f(B))$. The composition of p with the coarse inverse \bar{f} to f , is not a path in X , so we have to coarsify p . We find a finite sequence $z_1 = y_1, z_2, \dots, z_n = y_2$ in the image of p , such that

$$\text{dist}_Y(z_i, z_{i+1}) \leq 1, \quad i = 1, \dots, n - 1.$$

Then points $w_i = \bar{f}(z_i) \in X$ satisfy

$$\text{dist}_X(w_i, w_{i+1}) \leq L + A, \quad i = 1, \dots, n - 1.$$

The points $x'_1 := w_1, x'_2 := w_n$ are within distance $\leq A$ from the points $w_0 := x_1, w_{n+1} := x_2$, respectively. Connecting the consecutive points $w_i, w_{i+1}, i = 0, \dots, n$, by geodesic segments in X results in a path q , connecting x_1 to x_2 . This is our replacement for the (likely discontinuous) path $\bar{f} \circ p$. We would like to ensure that the image of q is disjoint from B : This would result in a contradiction, as we assumed that $C_1 \neq C_2$ are distinct components of B^c . If the image $\text{Im}(p)$ of the path p lies outside of the ball $B(y, r')$, $y = f(x)$, then

$$\text{dist}_X(x, \text{Im}(q)) \geq R'' := L^{-1}r' - 3A - L.$$

We choose r' such that $R'' \geq R$. Therefore, if x_1, x_2 are sufficiently far away from x (and this is certainly possible to achieve since we assume that the sets C_1, C_2

are unbounded), then y_1, y_2 lie in distinct complementary components of $B(y, r')$. Thus, there exists a bounded subset $B' = B(y, r') \subset Y$ whose complement contains at least n unbounded components. We proved that $\eta(Y) \geq \eta(X)$.

Reversing the roles of X and Y , we conclude that $\eta(X) = \eta(Y)$. \square

In particular, we now can define the number of ends of finitely generated groups:

DEFINITION 9.6. Let G be a finitely generated group. Then the number of ends $\eta(G)$ is the number of ends of its Cayley graph.

In view of Lemma 9.5, the quantity $\eta(G)$ is well-defined, as the number of ends is independent of the generating set of G . Moreover, $\eta(G)$ is a quasiisometry invariant of G .

9.1.2. The space of ends. Our next goal is to define a set $\epsilon(X)$ of ends of a topological space X , such that $\text{card}(\epsilon(X)) = \eta(X)$ if either one is finite. We will also equip $\epsilon(X)$ with a topology, which we then use in order to compactify X by adding to it the set of ends. The idea is that the ends of X are encoded by *decreasing* families of complementary components of compact subsets of X . We refer the reader to Section 1.5 for the required background on inverse limits.

We again let X be a non-empty, locally compact, connected, locally path-connected, second countable Hausdorff topological space. In particular, X admits an exhaustion by a countable family $(B_n)_{n \in \mathbb{N}}$ of compact subsets as in Proposition 1.22. For instance, if X is a proper metric space (the case we are mostly interested in), we can take $B_n = \bar{B}(x, n)$, where $x \in X$ is a fixed point, $n \in \mathbb{N}$.

Define $\mathcal{K} = \mathcal{K}_X$, the poset of compact subsets of X with the partial order \leq given by the inclusion. It is clear that the poset \mathcal{K} is directed, as the union of two compact sets is again compact. For each $K \in \mathcal{K}$ we have the set $\pi_0(K^c)$ whose elements are connected (equivalently, path-connected) components of K^c . Whenever $K_1 \leq K_2$ are compact subsets of X , we have the associated map

$$f_{K_1, K_2} : \pi_0(K_2^c) \rightarrow \pi_0(K_1^c),$$

sending each component C_2 of K_2^c to the unique component C_1 of K_1^c such that C_2 is contained in C_1 .

EXERCISE 9.7. Verify that the resulting collection of maps f_{K_2, K_1} is an inverse system, i.e.

$$f_{K_1, K_2} \circ f_{K_2, K_3} = f_{K_1, K_3}, \quad f_{K, K} = \text{Id}.$$

We will use the notation $\pi_0(\mathcal{K}^c)$ for this inverse system.

DEFINITION 9.8. The *set of ends* of X , denoted $\epsilon(X)$, is the inverse limit of the inverse system $\pi_0(\mathcal{K}^c)$. We will equip $\epsilon(X)$ with the initial topology, where each $\pi_0(K^c)$ is equipped with the discrete topology.

EXERCISE 9.9. Show that the space $\epsilon(X)$ is totally disconnected and Hausdorff.

PROPOSITION 9.10. *For every compact $K \subset X$, the set $\pi_0^u(K^c)$ is finite.*

PROOF. We will assume that K is non-empty since the proof is clear otherwise. Since X admits an exhaustion by compact subsets, there exists a compact $K' \subset X$ whose interior contains K . We claim that only finitely many components of U_K have non-empty intersection with $X \setminus K'$. It suffices to exclude the case when U_K has countably infinitely many components $U_i, i \in \mathbb{N}$. Since X is path-connected,

for each component U_i , there exists a path connecting some $x \in K$ to $x_i \in U_i \setminus K'$. Let y_i be a point in this path which belongs to $\partial K'$. Since $\partial K'$ is compact, after passing to a subsequence, we can assume that

$$\lim_{i \rightarrow \infty} y_i = y \in \partial K'.$$

Then $V := X \setminus K$ is a neighborhood of y . Since X is locally path-connected, there exists a neighborhood W of y contained in V , such that for all $i \geq i_0$ points y, y_i are connected by a path contained in W . It follows that $U_i = U_{i_0}$ for all $i \geq i_0$. \square

In addition to the inverse system $(\pi_0(K^c))_{K \in \mathcal{K}}$, we also have, similarly defined, inverse systems

$$(\pi_0^u(K^c))_{K \in \mathcal{K}}$$

and

$$(\pi_0^u(B_n^c))_{n \in \mathbb{N}}$$

where we use the standard order on \mathbb{N} . Inclusion maps

$$\{B_n : n \in \mathbb{N}\} \hookrightarrow \mathcal{K}$$

and

$$\pi_0^u(K^c) \rightarrow \pi_0(K^c)$$

induce maps of inverse limits

$$\phi : \varprojlim \pi_0^u(B_n^c) \rightarrow \varprojlim \pi_0^u(K^c)$$

and

$$\psi : \varprojlim \pi_0^u(K^c) \rightarrow \varprojlim \pi_0(K^c) = \epsilon(X).$$

We again equip the inverse limits $\varprojlim \pi_0^u(B_n^c)$ and $\varprojlim \pi_0(K^c)$ with the initial topology.

Since each $\pi_0^u(K^c)$ is finite, the inverse limit

$$\varprojlim \pi_0^u(K^c)$$

is compact by Tychonoff's theorem, cf. Exercise 1.24.

In view of Proposition 9.10, each $K_1 \in \mathcal{K}$ is contained in the interior of $K_2 \in \mathcal{K}$, such that the image of the map $\pi_0(K_2^c) \rightarrow \pi_0(K_1^c)$ is contained in $\pi_0^u(K_1^c)$. Combined with the fact that (B_n) is cofinal in \mathcal{K} , Exercises 1.25 and 1.25 now imply that the maps ϕ and ψ are continuous bijections. Since the domain of each map is compact and the range is Hausdorff, it follows that the maps ϕ, ψ are homeomorphisms.

Therefore, we can identify elements of $\epsilon(X)$ with decreasing sequences, called *chains*,

$$C_1 \supset C_2 \supset \dots$$

of components of the sets B_i^c , $i \in \mathbb{N}$, defined with respect to a fixed exhaustion of X as above.

One way to think about ends of X according to the definition, is that an end of X is a map $e : \mathcal{K} \rightarrow 2^X$, which sends each compact $K \subset X$ to a component C of K^c , such that

$$K_1 \subset K_2 \Rightarrow e(K_2) \subset e(K_1).$$

The topology on $\epsilon(X)$ extends to a topology on $\bar{X} = X \cup \epsilon(X)$: The basis of topology at $e \in \epsilon(X)$ is the collection of subsets $B_{K,e} \subset \bar{X}$, $K \in \mathcal{K}$, where $B_{K,e} \cap X = e(K)$

and $B_{K,e} \cap \epsilon(X)$ consists of all maps $e' : \mathcal{K} \rightarrow 2^X$, such that $e'(K) = e(K)$. We will also refer to each set $e(K)$ as a *neighborhood of e in X* .

The topology on \bar{X} is, of course, the original one. It is then immediate that X is open and dense in \bar{X} .

We will say that a compact subset $K \subset X$ *separates* ends e, e' of X if e, e' belong to distinct components of $\bar{X} \setminus K$. Equivalently, there are unbounded components C, C' of K^c such that (C, e) is a neighborhood of e and (C', e') is a neighborhood of e' in $\epsilon(X)$.

EXERCISE 9.11. Every topological action $G \curvearrowright X$ extends to a topological action of G on \bar{X} .

REMARK 9.12. There is a terminological confusion here coming from the literature in differential geometry and geometric analysis, where X is a complete Riemannian manifold: An analyst would call each unbounded set C_i above, an end of X .

Here is yet another alternative description of the space $\epsilon(X)$. From each C_i we pick a point x_i . Then, for each chain (C_i) defining the end $e \in \epsilon(X)$, the sequence (x_i) , denoted x_\bullet , *represents* the end e . Given a sequence x_\bullet representing e , we connect each x_i to x_{i+1} by a path contained in C_i . The concatenation of these paths is a *ray* in X , i.e. a proper continuous map

$$r : \mathbb{R}_+ \rightarrow X, \quad r(i) = x_i.$$

Conversely, given a ray r in X , every sequence $t_i \in \mathbb{R}_+$ monotonically diverging to infinity, defines the sequence x_\bullet (with $x_i = r(t_i)$) which represents an end e of X . This end is independent of the choice of a sequence t_i .

Two rays r_1, r_2 represent the same end of X if and only if for every compact $K \subset X$ there exists T such that for all $t \geq T$ the points $r_1(t), r_2(t)$ lie in the same component of K^c . We refer the reader to [BH99] and [Geo08] for more detailed description of $\epsilon(X)$ and topology on \bar{X} using this interpretation of ends.

EXERCISE 9.13. 1. The space \bar{X} is Hausdorff.

2. If X is second countable, so is \bar{X} .

3. A sequence x_\bullet in X represents the end e if and only if it converges to e in the topology of \bar{X} .

4. If X is a metric space and $(x_i), (x'_i)$ are sequences within bounded distance from each other:

$$\sup_i \text{dist}(x_i, x'_i) < \infty$$

and (x_i) represents $e \in \epsilon(X)$ then (x'_i) also represents e .

An example of the space of ends is given by the Figure 9.1. The space X in this picture has five visibly different ends: $\epsilon_1, \dots, \epsilon_5$. We have $K_1 \subset K_2 \subset K_3$. The compact K_1 separates the ends ϵ_1, ϵ_2 . The next compact K_2 separates ϵ_3 from ϵ_4 . Finally, the compact K_3 separates ϵ_4 from ϵ_5 .

LEMMA 9.14. *The space $\bar{X} = X \cup \epsilon(X)$ is compact.*

PROOF. Let $\mathcal{V} = \{V_j\}_{j \in J}$ be an open cover of \bar{X} . Since $\epsilon(X)$ is compact, there is a finite subset

$$\{(K_i, e_i) : i = 1, \dots, n\} = \mathcal{V}_1 \subset \mathcal{V},$$




FIGURE 9.1. Ends of X .

which still covers $\epsilon(X)$. Here $K_i \in \mathcal{K}$ and $e_i \in \epsilon(X)$. Consider the open sets $C_i = e_i(K_i)$ (here we think of ends of X as maps $\mathcal{K} \rightarrow 2^X$). We claim that the closed set

$$A = X \setminus \bigcup_{i=1}^n C_i$$

is compact in X . If not, then there exists a sequence $x_k \in A$ which is not contained in any of the compact $K_i, i \in \mathbb{N}$. After passing to a subsequence in the sequence (x_k) , we get a decreasing sequence of complementary sets

$$C_{k_l} \subset X \setminus K_{k_l},$$

such that $x_{k_l} \in C_{k_l}$. This sequence of complementary sets defines an end $e \in \epsilon(X)$ not covered by any of the sets $(K_i, e_i), i = 1, \dots, n$, which is a contradiction.

Thus, $A \subset X$ is compact. Then there exists another finite subset $\mathcal{V}_2 \subset \mathcal{V}$, which covers A . Therefore,

$$\mathcal{V}_1 \cup \mathcal{V}_2$$

is a finite subcover of \bar{X} . □

COROLLARY 9.15. 1. \bar{X} is a compactification of X .

2. The space \bar{X} is normal.

EXERCISE 9.16. 1. $\eta(X)$ is finite if and only if $\epsilon(X)$ is finite.

2. If $\eta(X)$ is finite, then $\eta(X)$ equals the cardinality of $\epsilon(X)$.

Proofs of 1 and 2 amount to simply following the definitions of $\eta(X)$ and $\epsilon(X)$.

REMARK 9.17. One can think of the space of ends of X as its “ π_0 at infinity.” One can also define higher homotopy and (co)homology groups of X “at infinity”, by replacing $\pi_0(K^c)$ with suitable homotopy or (co)homology groups and then taking inverse (or direct) limit. See [Geo08].

PROPOSITION 9.18. *Every quasiisometry of proper geodesic metric spaces $X \rightarrow X'$ induces a homeomorphism $\epsilon(X) \rightarrow \epsilon(X')$.*

PROOF. The proof of this proposition follows the proof of Lemma 9.5 and below we only sketch the proof, leaving details to the reader.

Let $f : X \rightarrow X'$ be an (L, A) -quasiisometry. As in the proof of Lemma 9.5, we observe that there exists $r = r(L, A)$ such that for every connected subset $C \subset X$, the subset $\mathcal{N}_r(f(C)) \subset X'$ is also connected.

We define a map $\epsilon(f) : \epsilon(X) \rightarrow \epsilon(X')$ as follows. Suppose that $e \in \epsilon(X)$ is represented by a nested sequence (C_i) , where C_i is a component of K_i^c , $K_i = \bar{B}(x, i)$. Each connected subset $\mathcal{N}_r(f(C_{i'}))$ will be disjoint from $\bar{B}(x', i)$ ($x' = f(x)$), where $i \mapsto i'$ is a nonconstant linear function, depending only on L and A . Let D_i denote the (unbounded) component of $\bar{B}(x', i)$ containing $\mathcal{N}_r(f(C_{i'}))$. The sets D_i are nested: $D_{i+1} \subset D_i$, $i \in \mathbb{N}$. Therefore, the sequence (D_i) defines an end e' of X' , and we set $\epsilon(f)(e) := e'$. Proof of injectivity of the map $\epsilon(f)$ is the same as the 3rd part of the argument in Lemma 9.5.

In order to verify continuity of $\epsilon(f)$, let $D_i \subset Y$ be a neighborhood of $e' = \epsilon(f)(e)$, as above. Then, as we noted, $f(C_{i'}) \subset D_i$, where $i \rightarrow i'$ is a nonconstant linear function. Thus, if (C_j) is a chain representing e , then for all $j \geq i'$, $f(C_j) \subset D_i$. Therefore, the entire neighborhood of $e \in \epsilon(X)$ defined by the pair $(K_{i'}, e)$, is mapped by $\epsilon(f)$ into the neighborhood of e' , defined by the pair $(\bar{B}(x', i), e')$. Continuity of $\epsilon(f)$ follows.

In order to prove surjectivity of $\epsilon(f)$, take r such that $\mathcal{N}_r(f(X)) = X'$. Then, given a sequence (x'_i) in X' representing an end $e' \in \epsilon(X')$, find a sequence (x_i) in X such that

$$\text{dist}_Y(f(x_i), x'_i) \leq r.$$

Then the sequence (x_i) will converge to an end e of X and $\epsilon(f)(e) = e'$. \square

EXERCISE 9.19. 1. Show that every bounded perturbation of the identity $f : X \rightarrow X$ extends to the identity map $\epsilon(f) : \epsilon(X) \rightarrow \epsilon(X)$.

2. Suppose that $f : X \rightarrow Y, g : Y \rightarrow Z$ are quasiisometries. Show that

$$\epsilon(g \circ f) = \epsilon(g) \circ \epsilon(f).$$

3. Conclude that if $g : Y \rightarrow X$ is a coarse inverse to $f : X \rightarrow Y$, then $\epsilon(g)$ is the inverse of $\epsilon(f)$. This gives another proof, in Proposition 9.18, of the claim that $\epsilon(f)$ is invertible.

Proposition 9.18 immediately implies:

COROLLARY 9.20. *Quasi-isometric spaces have homeomorphic spaces of ends.*

EXERCISE 9.21. Suppose that X is a simplicial tree of finite valence, where all but finitely many vertices have the same valence $k \geq 3$. Then $\epsilon(X)$ is homeomorphic to the Cantor set.

Hint: In order to prove this, consider first the case when X is a binary rooted tree, i.e. a tree with one distinguished vertex v_0 (the root) of valence 2 and the rest of the vertices of the valence 3. Then consider the ternary Cantor set E . This set is obtained by intersecting closed subsets $A_i \subset [0, 1]$, $i \in \mathbb{N}$; each A_i is the disjoint union of 2^i closed intervals $J_{i,k}$. Similarly, for $i \in \mathbb{N}$, the complement to the closed ball $K_i = \bar{B}(v_0, i) \subset X$ consists of 2^i components $C_{i,k}$. Points in E are encoded by decreasing sequences of intervals $J_{i,k}$, while points in $\epsilon(X)$ are encoded

by chains $(C_{i,k})$. Now use bijections between the sets $\{J_{i,k} : k = 1, \dots, 2^i\}$ and $\{C_{i,k} : k = 1, \dots, 2^i\}$. For more general simplicial trees follow the geometric proof of the Example 8.50.

9.1.3. Ends of groups. Suppose that G is a finitely generated group. Then we define the space of ends $\epsilon(G)$ as the space of ends of its Cayley graph X . Corollary 9.20 shows that $\epsilon(G)$ is independent of the generating set. It follows from the Exercise 9.11 that the group G acts topologically on $\bar{X} = X \cup \epsilon(G)$. The same applies if instead of the Cayley graph we use as X a Riemannian manifold M , on which G acts isometrically, properly discontinuously and cocompactly.

A proof of the following theorem can be found for instance [BH99, Theorem 8.32]:

THEOREM 9.22 (Properties of $\epsilon(X)$). *1. Suppose that G is a finitely generated group. Then $\epsilon(G)$ consists of 0, 1, or 2 points, or is infinite. In the latter case, the topological space $\epsilon(G)$ is perfect. In particular, $\epsilon(G)$ is homeomorphic to the Cantor set.*

2. $\epsilon(G)$ is empty iff G is finite. $\epsilon(G)$ consists of 2-points if and only if G is virtually (infinite) cyclic. In particular, G is quasiisometric to \mathbb{Z} if and only if G is virtually isomorphic to \mathbb{Z} .

3. If G splits non-trivially over a finite subgroup then $|\epsilon(G)| > 1$.

Below we prove Part 2 of this theorem. Our proof (which we learned from Mladen Bestvina) is differential-geometric, in line with the arguments in Chapters 20 and 21. A combinatorial argument can be found in [BH99].

PROPOSITION 9.23. *Every 2-ended group G is virtually isomorphic to \mathbb{Z} and, hence, contains a finite-index subgroup isomorphic to \mathbb{Z} .*

PROOF. Let M be an n -dimensional oriented Riemannian manifold on which G acts isometrically, properly discontinuously and cocompactly, preserving orientation. We let $\omega \in \Lambda^n(M)$ denote the volume form of M . Since M is QI to G , the manifold M is also 2-ended. After passing to an index 2 subgroup of G , we can assume that G fixes the ends e_1, e_2 of M . Every compact connected hypersurface $S \subset M$ separating the ends of M has a canonical coorientation, such that the end e_1 lies to the left of S . Since M is oriented, we, therefore, obtain a canonical orientation on S . This orientation is preserved under the action of G . The oriented hypersurface S will be regarded below as a smooth singular cycle in M , an element of $Z_{n-1}(M)$. (This cycle is the image of the fundamental cycle of S under the map $Z_{n-1}(S) \rightarrow Z_{n-1}(M)$.) Accordingly, $-S$ is the hypersurface S with reversed orientation. We claim that for every $g \in G$, the oriented hypersurfaces $S, g(S)$, represent the same homology class in $H_{n-1}(M)$. Indeed, the hypersurface $g(S)$ still separates the ends of M . If $g(S) \cap S = \emptyset$, then, since M is 2-ended, there exists a compact submanifold $B \subset M$ whose (oriented) boundary equals $-S \cup g(S)$. Hence, $[S] = [g(S)] \in H_{n-1}(M)$. For arbitrary $g \in G$ we take $h \in G$ such that $h(S) \cap S = \emptyset$ and $h(S) \cap g(S)$ and obtain

$$[S] = [h(S)] = [g(S)].$$

We, thus, obtain a homomorphism

$$\phi : G \rightarrow \mathbb{R},$$

defined by

$$\phi(g) = \int_B \omega$$

where $B \in C_n(M)$ is a smooth singular chain such that

$$g(S) - S = B.$$

As we observed above, if $g(S) \cap S = \emptyset$, then B is realized by a submanifold in M , which implies that $\phi(g) \neq 0$ in this case. Since the action of G on M is properly discontinuous, the map $\phi : G \rightarrow \mathbb{R}$ is proper. In particular, its image is an infinite cyclic group and its kernel is finite. Therefore, the group G is virtually isomorphic to \mathbb{Z} . The existence of a finite-index subgroup of G isomorphic to \mathbb{Z} was proven in Corollary 7.109. \square

Part 3 of Theorem 9.22 has a deep and important converse:

THEOREM 9.24. *If $|\epsilon(G)| > 1$ then G splits non-trivially over a finite subgroup.*

This theorem is due to Stallings [Sta68] (in the torsion-free case) and Bergman [Ber68] for groups with torsion. A geometric proof could be found in Niblo's paper [Nib04] and a shorter, combinatorial, proof in Kron's paper [Krö10]. For finitely presented groups, there is an alternative combinatorial proof due to Dunwoody using minimal tracks, [Dun85]; a combinatorial version of this argument could be found in [DD89]. In Chapters 20 and 21 we prove Theorem 9.24 first for finitely presented, and then for all finitely generated groups. We will also prove QI rigidity of the class of virtually free groups.

An immediate corollary of Theorem 9.24 (and QI invariance of the number of ends) is

COROLLARY 9.25. *Suppose that a finitely generated group G splits non-trivially as $G_1 \star G_2$ and G' is a group quasiisometric to G . Then G' splits non-trivially as $G'_1 \star_F G'_2$ (amalgamated product) or as $G'_1 \star_F$ (HNN splitting), where F is a finite group.*

We conclude this section with a technical result which will be used in Section 21.3 for the proof of Stallings theorem via harmonic functions.

LEMMA 9.26. *Suppose that M is a complete connected n -dimensional Riemannian manifold and X is the corresponding metric space. Let $\chi : \epsilon(X) \rightarrow \{0, 1\}$ be the characteristic function of a clopen subset $A \subset \epsilon(X)$. Then χ admits a continuous extension $\varphi : \bar{X} \rightarrow [0, 1]$ which is smooth on M and $d\varphi|_M$ is compactly supported in M .*

PROOF. Since the space \bar{X} is normal, the disjoint closed sets A and $B = \epsilon(X) \setminus A$ admit disjoint open neighborhoods $U \subset \bar{X}$ and $V \subset \bar{X}$, respectively. We first extend χ to a function $\psi : U \cup V \rightarrow \{0, 1\}$, which is constant on U and on V . Next, by Tietze–Urysohn extension theorem (Theorem 1.16) the function $\psi : U \cup V \rightarrow \{0, 1\}$ admits a continuous extension $\zeta : \bar{X} \rightarrow \mathbb{R}$. After replacing ζ with $\zeta_0 = \max(\zeta, 0)$ and, afterwards, with $\zeta_1 = \min(\zeta, 1)$, we may assume that $\zeta : \bar{X} \rightarrow [0, 1]$. To get a smooth extension, consider a smooth partition of unity $\{\eta_i\}_{i \in I}$ corresponding to a locally finite open covering $\{U_i\}_{i \in I}$ of M via subsets diffeomorphic to the unit open ball $\mathbb{D} \subset \mathbb{R}^n$. We choose the functions η_i to have

unit integrals over U_i (with respect to the Lebesgue measure coming from \mathbb{D}). Using the diffeomorphisms $f_i : U_i \rightarrow \mathbf{B}$, we define convolutions

$$\zeta \star \eta_i(x) = \int_{\mathbb{D}} \zeta(f_i(y))\eta_i(f_i(x))dx.$$

The sum

$$\varphi = \sum_{i \in I} \zeta \star \eta_i$$

is the required extension. \square

9.2. Rips complexes and coarse homotopy theory

Connecting the dots. In the proof of Lemma 9.5, we saw an important principle of coarse topology: In order to recover a useful topological object from the image $f(C)$ of a set under a quasiisometry f , we first discretize C (replace C with a net $C' \subset C$) and then “connect the dots” in $f(C')$: Connect certain points (which are not too far from each other) in $f(C')$ by geodesic segments in the ambient space. How far the “connected dots” should be from each other is determined by geometry of the metric spaces involved and quasiisometric constants of f . The same principle will reappear in this section: “Connecting dots” will be replaced by taking a subcomplex of a suitable Rips complex Rips_R . The ambiguity in choosing the scale R (how far apart the “dots” can be) forces us to work with direct systems of Rips complexes and direct/inverse systems of the associated homotopy, homology and cohomology groups.

9.2.1. Rips complexes. Recall (Definition 2.24) that the R -Rips complex of a metric space X is the simplicial complex whose vertices are the points of X ; vertices x_1, \dots, x_n span a simplex if and only if

$$\text{dist}(x_i, x_j) \leq R, \forall i, j.$$

For each pair $0 \leq R_1 \leq R_2 < \infty$ we have a natural simplicial embedding

$$\iota_{R_1, R_2} : \text{Rips}_{R_1}(X) \rightarrow \text{Rips}_{R_2}(X),$$

such that

$$\iota_{R_1, R_3} = \iota_{R_2, R_3} \circ \iota_{R_1, R_2},$$

provided that $R_1 \leq R_2 \leq R_3$. Thus, the collection of Rips complexes of X forms a direct system $\text{Rips}_\bullet(X)$ of simplicial complexes indexed by positive real numbers.

Following the construction in Section 3.8, we metrize (connected) Rips complexes $\text{Rips}_R(X)$ using the *standard metric* on simplicial complexes. **Then** each embedding ι_{R_1, R_2} is isometric on every simplex and is 1-Lipschitz overall. Note that if X is uniformly discrete (see Definition 2.3), then for every R , the complex $\text{Rips}_R(X)$ is a simplicial complex of bounded geometry (Definition 3.33).

EXERCISE 9.27. 1. Suppose that $X = G$, a finitely generated group with a word metric. Show that for every R , the action of G on itself extends to a simplicial action of G on $\text{Rips}_R(G)$. Show that this action is geometric.

2. Show that a metric space X is quasigeodesic (see Section 8.1) if and only if for all sufficiently large R the Rips complex $\text{Rips}_R(X)$ is connected and the inclusion $X \rightarrow \text{Rips}_R(X)$ is a quasiisometry.

The following simple observation explains why Rips complexes are useful for analyzing quasiisometries:

LEMMA 9.28. *Let $f : X \rightarrow Y$ be an (L, A) -coarse Lipschitz map. Then f induces a simplicial map $\text{Rips}_R(X) \rightarrow \text{Rips}_{LR+A}(Y)$ for each $R \geq 0$. We retain the notation f for this simplicial map.*

PROOF. Consider an m -simplex σ in $\text{Rips}_R(X)$; the vertices of σ are distinct points $x_0, x_1, \dots, x_m \in X$ within distance $\leq R$ from each other. Since f is (L, A) -coarse Lipschitz, the points $f(x_0), \dots, f(x_m) \in Y$ are within distance $\leq LR + A$ from each other, hence, they span a simplex σ' of dimension $\leq m$ in $\text{Rips}_{LR+A}(Y)$. The map f sends vertices of σ to vertices of σ' . Thus, we have a simplicial map of simplicial complexes $\text{Rips}_R(X) \rightarrow \text{Rips}_{LR+A}(Y)$. \square

The idea behind the next definition is that the “coarse homotopy groups” of a metric space X are the homotopy groups of the Rips complexes $\text{Rips}_R(X)$ of X . Literally speaking, this does not make much sense since the above homotopy groups depend on R . To eliminate this dependence, we have to take into account the maps $\iota_{r,R}$.

DEFINITION 9.29. 1. A metric space X is *coarsely connected* if $\text{Rips}_r(X)$ is connected for some r . (Equivalently, $\text{Rips}_R(X)$ is connected for all sufficiently large R .)

2. A metric space X is *coarsely k -connected* if it is coarsely connected and for each r there exists $R \geq r$ such that the mapping $\text{Rips}_r(X) \rightarrow \text{Rips}_R(X)$ induces trivial maps of the homotopy groups

$$\pi_i(\text{Rips}_r(X), x) \rightarrow \pi_i(\text{Rips}_R(X), x),$$

for all $1 \leq i \leq k$ and $x \in X$.

In particular, X is *coarsely simply-connected* if it is coarsely 1-connected.

For instance, X is coarsely connected if there exists a number R such that each pair of points $x, y \in X$ can be connected by an R -chain of points $x_i \in X$, i.e. a finite sequence of points x_i , where $\text{dist}(x_i, x_{i+1}) \leq R$ for each i .

The definition of coarse k -connectedness is not quite satisfactory since it only deals with “vanishing” of coarse homotopy groups without actually defining these groups for a general metric space X . One way to deal with this issue is to consider *pro-groups*, which are direct systems

$$\pi_i(\text{Rips}_r(X)), r \in \mathbb{N},$$

of groups. Given such algebraic objects, one can define their *pro-homomorphisms*, *pro-monomorphisms*, etc., see [KK05] where this is done in the category of abelian groups (the homology groups). Alternatively, one can work with the direct limit of the homotopy groups.

REMARK 9.30. Arguing analogously to the proof of Lemma 9.54, one can show that if X is coarsely 1-connected then there exists R_0 such that the Rips complex $\text{Rips}_R(X)$ is 1-connected for all $R \geq R_0$. However, it is unclear to us if the same holds for coarsely k -connected metric spaces, $k \geq 2$.

9.2.2. Direct system of Rips complexes and coarse homotopy.

LEMMA 9.31. *Let X be a metric space. Then for $r, c < \infty$, each simplicial spherical cycle σ of diameter $\leq c$ in $\text{Rips}_r(X)$ bounds a singular disk of diameter $\leq r+c$ within $\text{Rips}_{r+c}(X)$. More precisely, every simplicial map of a triangulated $n-1$ sphere, $\sigma : \mathbb{S}^{n-1} \rightarrow \text{Rips}_r(X)$, extends to a simplicial map $\tau : \mathbb{D}^n \rightarrow \text{Rips}_{r+c}(X)$, where \mathbb{D}^n is a triangulated n -disk whose triangulation agrees with that one of \mathbb{S}^{n-1} .*

PROOF. Pick a vertex $x \in \text{Im}(\sigma)$. Then $\text{Rips}_{r+c}(X)$ contains the simplicial cone $C = \tau(\mathbb{D}^n)$ over $\text{Im}(\sigma)$ with vertex at x . Clearly, $\text{diam}(C) \leq r+c$. Coning off the map σ from the vertex x , defines an extension τ of σ to the n -disk, which we identify with the cone over \mathbb{S}^{n-1} . \square

Recall that the product of simplicial complexes $C \times [0, 1]$ admits a certain standard triangulation (determined by an ordering of vertices of X and the set $\{0, 1\}$). We will always equip this product simplicial complex with the standard metric.

PROPOSITION 9.32. *Let $f, g : X \rightarrow Y$ be maps within distance $\leq c$ from each other, which extend to simplicial maps*

$$f, g : \text{Rips}_{r_1}(X) \rightarrow \text{Rips}_{r_2}(Y).$$

Then for $r_3 = r_2 + c$, the maps

$$f, g : \text{Rips}_{r_1} \rightarrow \text{Rips}_{r_3}(Y)$$

are homotopic via a 1-Lipschitz homotopy $F : \text{Rips}_{r_1}(X) \times I \rightarrow \text{Rips}_{r_3}(Y)$. Furthermore, tracks of this homotopy have length $\leq (n+1)$, where $n = \dim(\text{Rips}_{r_1}(X))$.

PROOF. The map F of the zero-skeleton of $\text{Rips}_{r_1}(X) \times I$ is, of course, just $F(x, 0) = f(x), F(x, 1) = g(x)$. Let $\sigma \subset \text{Rips}_{r_1}(X) \times I$ be an i -simplex. Then

$$\text{diam}(F(V(\sigma))) \leq r_3 = r_2 + c,$$

where $V(\sigma)$ is the vertex set of σ . Therefore, F extends (linearly) from σ^0 to a (1-Lipschitz) map $F : \sigma \rightarrow \text{Rips}_{r_3}(Y)$ whose image is the simplex spanned by $F(\sigma^0)$.

To estimate the lengths of the tracks of the homotopy F , we note that for each $x \in \text{Rips}_{r_1}(X)$, the path $F(x, t)$ has length ≤ 1 since the interval $x \times I$ is covered by $\leq (n+1)$ simplices, each of which has unit diameter. \square

In view of the above lemma, we make the following definition:

DEFINITION 9.33. Maps $f, g : X \rightarrow Y$ are *coarsely homotopic* if for all r_1, r_2 , such that f and g extend to

$$f, g : \text{Rips}_{r_1}(X) \rightarrow \text{Rips}_{r_2}(Y),$$

there exist r_3 and r_4 so that the maps

$$f, g : \text{Rips}_{r_1}(X) \rightarrow \text{Rips}_{r_3}(Y)$$

are homotopic via a homotopy whose tracks have lengths $\leq r_4$.

We then say that a map $f : X \rightarrow Y$ determines a *coarse homotopy equivalence* (between the direct systems of Rips complexes of X, Y), if there exists a map $g : Y \rightarrow X$ such that the compositions $g \circ f, f \circ g$ are coarsely homotopic to the identity maps.

The next two corollaries, then, are immediate consequences of Proposition 9.32.

COROLLARY 9.34. *Let $f, g : X \rightarrow Y$ be L -Lipschitz maps within finite distance from each other. Then they are coarsely homotopic.*

COROLLARY 9.35. *If $f : X \rightarrow Y$ is a quasiisometry, then f induces a coarse homotopy-equivalence of the Rips complexes: $\text{Rips}_\bullet(X) \rightarrow \text{Rips}_\bullet(Y)$.*

The following corollary is a coarse analogue of the familiar fact that homotopy equivalence preserves connectivity properties of a space:

COROLLARY 9.36. *Coarse k -connectedness is a QI invariant.*

PROOF. Suppose that X' is a coarsely k -connected metric space and $f : X \rightarrow X'$ is an L -Lipschitz quasiisometry with L -Lipschitz coarse inverse $\bar{f} : X' \rightarrow X$. Let γ be a spherical i -cycle in $\text{Rips}_r(X)$, $0 \leq i \leq k$. Then we have the spherical i -cycle $f(\gamma) \subset \text{Rips}_{Lr}(X')$. Since X' is coarsely k -connected, there exists $r' \geq Lr$ such that $f(\gamma)$ bounds a singular $(i+1)$ -disk β within $\text{Rips}_{r'}(X')$. Consider now $\bar{f}(\beta) \subset \text{Rips}_{L^2r}(X)$. The boundary of this singular disk is a singular i -sphere $\bar{f}(\gamma)$. Since $\bar{f} \circ f$ is homotopic to Id within $\text{Rips}_{r''}(X)$, $r'' \geq L^2r$, there exists a singular cylinder σ in $\text{Rips}_{r''}(X)$ which cobounds γ and $\bar{f}(\gamma)$. Note that r'' does not depend on γ . By combining σ and $\bar{f}(\beta)$ we get a singular $(i+1)$ -disk in $\text{Rips}_{r''}(X)$ whose boundary is γ . Hence, X is coarsely k -connected. \square

9.3. Metric cell complexes

We now introduce a generalization of metric simplicial complexes, where the notion of bounded geometry does not imply finite-dimensionality. The objects that we will consider are called *metric cell complexes*, they are hybrids of metric spaces and CW complexes. The advantage of metric cell complexes over metric simplicial complexes is the same as of CW complexes over simplicial complexes in the traditional algebraic topology: CW complexes are more flexible.

A *metric cell complex* is a cell complex X together with a metric d defined on its 0-skeleton $X^{(0)}$. Note that if X is connected, its 1-skeleton $X^{(1)}$ is a graph, and, hence, can be equipped with the standard metric dist . The map $(X^{(0)}, d) \rightarrow (X^{(1)}, \text{dist})$, in general, need not be a quasiisometry. However, in the most interesting cases, coming from finitely generated groups, this map is actually an isometry. Therefore, we impose, from now on, the condition:

Axiom 1. The map $(X^{(0)}, d) \rightarrow (X^{(1)}, \text{dist})$ is a quasiisometry. Equivalently, X is a quasigeodesic metric space.

Even though this assumption could be avoided in what follows, restricting to complexes satisfying this axiom makes our discussion more intuitive.

Our first goal is to define, using the metric d , certain metric concepts on the entire complex X . We define inductively a map c , which sends cells in X to finite subsets of $X^{(0)}$ as follows. For a vertex $v \in X^{(0)}$ we set $c(v) = \{v\}$. Suppose that c is defined on $X^{(i)}$. For each closed $i+1$ -cell e , the *support* of e is the smallest subcomplex $\text{Supp}(e)$ of $X^{(i)}$, containing the image of the attaching map of e to $X^{(i)}$. We then set

$$c(\sigma) = c(\text{Supp}(e)).$$

For instance, for every 1-cell σ , $c(\sigma)$ consists of one or two vertices of X to which σ is attached.

REMARK 9.37. The reader familiar with the concepts of *controlled topology*, see e.g. [Ped95], will realize that the coarsely defined map $c : X \rightarrow X^{(0)}$ is a *control map* for X and $(X^{(0)}, d)$ is the *control space*. Metric cell complexes form a subclass of *metric chain complexes* defined in [KK05].

The *diameter* $\text{diam}(\sigma)$ of a cell σ in X is defined to be the diameter of $c(\sigma)$.

EXAMPLE 9.38. Take a connected simplicial complex X and restrict its standard metric to $X^{(0)}$. Then the diameter of a cell in X (as a simplicial complex) is the same as its diameter in the sense of metric cell complexes.

DEFINITION 9.39. A metric cell complex X is said to have *bounded geometry* if there exists a collections of increasing functions $\phi_k(r)$ and numbers $D_k < \infty$ such that the following axioms hold:

Axiom 2. For each ball $B(x, r) \subset X^{(0)}$, the set of k -cells σ such that $c(\sigma) \subset B(x, r)$, contains at most $\phi_k(r)$ cells.

Axiom 3. The diameter of each k -cell is at most $D_k = D_{k,X}$, $k \in \mathbb{N}$.

Axiom 4. $D_0 := \inf\{d(x, x') : x \neq x' \in X^{(0)}\} > 0$.

Note that we allow X to be infinite-dimensional. We will refer to the function $\phi_k(r)$ and the numbers D_k as *geometric bounds* on X , and set

$$(9.1) \quad D_X = \sup_{k>0} D_{k,X}.$$

The basic examples of metric cell complexes of bounded geometry are:

1. Simplicial complexes of bounded geometry.
2. Let M be a connected Riemannian manifold of bounded geometry and X is a simplicial complex defined in Theorem 3.36. Now, equip $X^{(0)}$ with the distance function d which is restriction of the Riemannian distance function on M to the vertex set of X .
3. $X^{(0)} := G$ is a finitely generated group with its word metric and X is the Cayley graph of G with the standard metric.
4. A covering space X of a connected finite cell complex Y . Equip $X^{(0)}$ with the restriction of the distance function dist on $X^{(1)}$.
5. Consider the spheres \mathbb{S}^n with the standard CW complex structure (single 0-cell and single n -cell). Then the cellular embeddings $\mathbb{S}^n \hookrightarrow \mathbb{S}^{n+1}$ give rise to an infinite-dimensional cell complex S^∞ . This complex has bounded geometry (since it has only one cell in every dimension). In view of this trivial example, the concept of metric cell complexes is more flexible than the one of simplicial complexes.

EXERCISE 9.40. 1. Suppose that X is a simplicial complex. Then the two notions of bounded geometry coincide for X . We will use this special class of metric cell complexes in Section 9.6.

2. If X is a metric cell complex of bounded geometry and $S \subset X$ is a connected subcomplex, then for every two vertices $u, v \in S$ there exists a chain $x_0 = u, x_1, \dots, x_m = v$, such that $d(x_i, x_{i+1}) \leq D_1$ for every i . In particular, if X is connected, the identity map $(X^{(0)}, d) \rightarrow (X^{(1)}, \text{dist})$ is D_1 -Lipschitz.

EXERCISE 9.41. Let X, Y be metric cell complexes. Then the product cell-complex $X \times Y$ is also a metric cell complex, where we equip the zero-skeleton

$X^{(0)} \times Y^{(0)}$ of $X \times Y$ with the product-metric. Furthermore, if X, Y have bounded geometry, then so does $X \times Y$.

We now continue defining metric concepts for metric cell complexes. The (coarse) R -ball $\mathbf{B}(x, R)$ centered at a vertex $x \in X^{(0)}$ is defined as the union of the cells σ in X such that $c(\sigma) \subset B(x, R)$.

We will say that the *diameter* $\text{diam}(S)$ of a subcomplex $S \subset X$ is the diameter of $c(S)$. Given a subcomplex $W \subset X$, we define the *closed R -neighborhood* $\bar{N}_R(W)$ of W in X to be the largest subcomplex $S \subset X$ such that for every $\sigma \in S$, there exists a vertex $\tau \in W$ such that $\text{dist}_{\text{Haus}}(c(v), c(w)) \leq R$. A cellular map $f : X \rightarrow Y$ between metric cell complexes is called *L -Lipschitz* if for every cell σ in X , we have $\text{diam}(f(\sigma)) \leq L$. In particular, the map

$$f : (X^{(0)}, d) \rightarrow (Y^{(0)}, d)$$

is $\frac{L}{D_0}$ -Lipschitz as a map of metric spaces.

EXERCISE 9.42. Suppose that $f_i : X_i \rightarrow X_{i+1}$ are L_i -Lipschitz for $i = 1, 2$. Show that $f_2 \circ f_1$ is L_3 -Lipschitz with

$$L_3 = L_2 \max_k (\phi_{X_2, k}(L_1)).$$

EXERCISE 9.43. Construct examples of a cellular map $f : X \rightarrow Y$ between metric graphs of bounded geometry, such that the restriction $f|_{X^{(0)}}$ is L -Lipschitz, but f is not L' -Lipschitz, for any $L' < \infty$.

The following definition is a version of the notion of uniformly proper maps of metric spaces in Definition 8.27. A map $f : X \rightarrow Y$ of metric cell complexes is called a *uniformly proper cellular map*, if f is cellular, L -Lipschitz for some $L < \infty$ and $f|_{X^{(0)}}$ is uniformly proper: There exists a proper function $\eta(R)$ such that

$$d(f(x), f(x')) \geq \eta(d(x, x')),$$

for all $x, x' \in X^{(0)}$. The function $\eta(R)$ is called the (*lower*) *distortion function* of f . We will frequently omit the adjective *cellular* when talking about uniformly proper maps of metric cell complexes.

For instance, suppose that H is a finitely generated group and $G \leq H$ is a finitely generated subgroup, whose generating set is contained in the one of H . Let X and Y denote the Cayley graphs of G and H , respectively. Then the inclusion map $X \rightarrow Y$ is a uniformly proper cellular map. As another example, suppose that G is the fundamental group of a finite cell complex X_1 , H is the fundamental group of a finite cell complex Y_1 and $f_1 : X_1 \rightarrow Y_1$ is a cellular map inducing the inclusion of fundamental groups $G \hookrightarrow H$. Let $f : X \rightarrow Y$ be a lift of f_1 to the universal covers X, Y of X_1, Y_1 , respectively. Then f is a uniformly proper cellular map.

We now relate metric cell complexes of bounded geometry to simplicial complexes of bounded geometry:

EXERCISE 9.44. Let X be a *finite-dimensional* metric cell complexes of bounded geometry. Then there exists a simplicial complex Y of bounded geometry and a cellular homotopy-equivalence $X \rightarrow Y$ which is a quasiisometry in the following sense: f and has homotopy-inverse \bar{f} so that:

1. Both f, \bar{f} are L -Lipschitz for some $L < \infty$.
2. $f \circ \bar{f}, \bar{f} \circ f$ are homotopic to the identity.

3. The maps $f : X^{(0)} \rightarrow Y^{(0)}$, $\bar{f} : Y^{(0)} \rightarrow X^{(0)}$ are coarse inverse to each other:

$$d(f \circ \bar{f}, \text{Id}) \leq A, \quad d(\bar{f} \circ f, \text{Id}) \leq A.$$

Hint: Apply the standard construction which converts a finite-dimensional CW-complex into a simplicial complex, see e.g. [Hat02].

Recall that quasiisometries are not necessarily continuous. In order to use algebraic topology, we, thus, have to approximate quasiisometries by cellular maps in the context of metric cell complexes, as it was done for Rips complexes (Lemma 9.28). In general, such approximation is, of course, impossible, since one complex in question can be, say, 0-dimensional and the other 1-dimensional. The *uniform contractibility* hypothesis allows one to resolve this issue.

Suppose that X and Y are cell complexes and $f : X \rightarrow Y$ is a cellular map. We will say that the map f is *k-null* if it induces zero map $\tilde{H}_0(X) \rightarrow \tilde{H}_0(Y)$ and trivial maps of all homotopy groups

$$\pi_i(X) \rightarrow \pi_i(Y), \quad 1 \leq i \leq k.$$

DEFINITION 9.45. A metric cell complex X is said to be *uniformly contractible* if there exists a continuous function $\psi(R)$ such that for every $x \in X^{(0)}$ the map

$$\mathbf{B}(x, R) \rightarrow \mathbf{B}(x, \psi(R))$$

is null-homotopic. Similarly, X is *uniformly k-connected* if there exists a function $\psi_k(R)$ such that for every $x \in X^{(0)}$ the map

$$(9.2) \quad \mathbf{B}(x, R) \hookrightarrow \mathbf{B}(x, \psi_k(R))$$

is *k-null*. We will refer to ψ, ψ_k as the *contractibility functions* of X . By the abuse of terminology, we will say that the inclusion of the balls (9.2) induces *trivial maps of homotopy groups* π_i , $i \leq k$. (The abuse comes from the fact that for $k = 0$ we use the reduced homology.)

The above definition implies, for instance, that the entire ball $\mathbf{B}(x, R)$ is contained in a single connected component of $\mathbf{B}(x, \psi_0(R))$, every loop in $\mathbf{B}(x, R)$ bounds a singular disk in $\mathbf{B}(x, \psi_1(R))$.

EXAMPLE 9.46. Suppose that X is a connected metric graph with the standard metric. Then X is uniformly 0-connected.

In general, even for simplicial complexes of bounded geometry, contractibility does not imply uniform contractibility. For instance, start with a triangulated 2-torus T^2 , let X be an infinite cyclic cover of T^2 . Of course, X is not contractible, but we attach a triangulated disk \mathbb{D}^2 to X along a simple homotopically non-trivial loop in $X^{(1)}$. The result is a contractible 2-dimensional simplicial complex Y which clearly has bounded geometry.

EXERCISE 9.47. Show that Y is not uniformly contractible.

We will see, nevertheless, in Lemma 9.51, that under certain assumptions (presence of a cocompact group action) contractibility implies uniform contractibility.

The reader uncomfortable with metric cell complexes in the proofs below, can think instead of Riemannian manifolds equipped with structures of CW-complexes, which appear as Riemannian cellular coverings of compact Riemannian manifolds,




FIGURE 9.2. Contractible but not uniformly contractible space.

which are given structures of finite CW complexes. Instead of the notions of diameter in metric cell complexes used in the book, the reader can think of the ordinary Riemannian diameters.

The following proposition is a metric analogue of the cellular approximation theorem:

PROPOSITION 9.48 (Lipschitz cellular approximation). *Suppose that X, Y are metric cell complexes, where X is finite-dimensional and has bounded geometry, Y is uniformly contractible, and $f : X^{(0)} \rightarrow Y^{(0)}$ is an L -Lipschitz map. Then f admits a (continuous) cellular extension $f : X \rightarrow Y$, which is an L' -Lipschitz map, where L' depends on L and geometric bounds on the complex X and the uniform contractibility function of Y . Furthermore, $f(X) \subset \tilde{N}_{L'}(f(X^{(0)}))$.*

PROOF. The proof of this proposition is a prototype of most of the proofs which appear in this and the following sections: It is a higher-dimensional version of the “collect the dots” process. The proof essentially amounts to a quantitative version of the proof of Whitehead’s theorem (see e.g. Theorem 4.5 in [Hat02]).

We extend f by induction on skeleta of X . We claim that (for certain constants $C_i, C'_{i+1}, i \geq 0$) we can construct a sequence of extensions $f_k : X^{(k)} \rightarrow Y^{(k)}$ such that:

1. $\text{diam}(f(\sigma)) \leq C_k$ for every k -cell $\sigma = \hat{e}(\mathbb{D}^k)$ in $X^{(k)}$.
2. $\text{diam}(f(\partial\tau)) \leq C'_{k+1}$, for every $(k+1)$ -cell τ in X .

Base of the induction. We already have $f = f_0 : X^{(0)} \rightarrow Y^{(0)}$ satisfying (1) with $C_0 = 0$. If x, x' belong to the boundary of a 1-cell τ in X then

$$\text{dist}(f(x), f(x')) \leq LD_1,$$

where $D_1 = D_{1,X}$ is the upper bound on the diameters of 1-cells in X . This establishes (2) in the base case.

Inductively, assume that $f = f_k$ was defined on $X^{(k)}$, so that (1) and (2) hold. Let $\sigma = \hat{e}(\mathbb{D}^{k+1})$ be a $(k+1)$ -cell in X . Note that

$$\text{diam}(f(\partial\sigma)) \leq C'_{k+1},$$

by the induction hypothesis. Then, using uniform contractibility of Y , we extend f to σ so that the diameter of the image of σ in Y is bounded above by C_{k+1} where $C_{k+1} = \psi(C'_k)$. Namely, the composition $f \circ e : \partial\mathbb{D}^{n+1} \rightarrow Y$ is null-homotopic and, hence, extends to a map $\mathbb{D}^{n+1} \rightarrow Y$ of controlled diameter. Without loss of generality (cf. Whitehead's cellular approximation theorem, [Hat02, Theorem 4.8]), we can assume that this extension h is cellular, i.e. its image is contained in $Y^{(n+1)}$. The extension of $f \circ e$ to \mathbb{D}^{n+1} determines the required extension of f to $\hat{e}(\mathbb{D}^{k+1})$:

$$f(x) := \tilde{f}(\hat{e}^{-1}(x)), x \in \sigma.$$

We thus obtain a cellular map $f : X^{(k+1)} \rightarrow Y^{(k+1)}$.

Let us verify that the new map $f : X^{(k+1)} \rightarrow Y^{(k+1)}$ satisfies (2).

Suppose that τ is a $(k+2)$ -cell in X . Then, since X has bounded geometry, $\text{diam}(\tau) \leq D_{k+2} = D_{k+2,X}$. In particular, $\partial\tau$ is connected and is contained in the union of at most $\phi(D_{k+2}, k+1)$ cells of dimension $k+1$. Therefore,

$$\text{diam}(f(\partial\tau)) \leq C_{k+1} \cdot \phi(D_{k+2}, k+1) =: C'_{k+2}.$$

This proves (2).

Since X is, say, n -dimensional, the induction terminates after n steps. The resulting map $f : X \rightarrow Y$ satisfies

$$L' := \text{diam}(f(\sigma)) \leq \max_{i=1,\dots,n} C_i,$$

for every cell σ in X . Therefore, $f : X \rightarrow Y$ is L' -Lipschitz. The second assertion of the proposition follows from the definition of C_i 's. \square

We note that Proposition 9.48 can be *relativized*:

LEMMA 9.49. *Suppose that X, Y are metric cell complexes, X is finite-dimensional and has bounded geometry, Y is uniformly contractible, and $Z \subset X$ is a subcomplex. Suppose that $f : Z \rightarrow Y$ is a continuous cellular map which extends to an L -Lipschitz map $f : X^{(0)} \rightarrow Y^{(0)}$. Then $f : Z \cup X^{(0)} \rightarrow Y$ admits a (continuous) cellular extension $g : X \rightarrow Y$, which is an L' -Lipschitz map, where L' depends on L , geometric bounds on X and contractibility function of Y .*

PROOF. The proof is the same induction on skeleta argument as in Proposition 9.48. \square

COROLLARY 9.50. *Suppose that X, Y are as above and $f_0, f_1 : X \rightarrow Y$ are L -Lipschitz cellular maps such that $\text{dist}(f_0, f_1) \leq C$ in the sense that*

$$d(f_0(x), f_1(x)) \leq C, \quad \forall x \in X^{(0)}.$$

Then there exists an L' -Lipschitz homotopy $f : X \times I \rightarrow Y$ between the maps f_0, f_1 .

PROOF. Consider the map $f_0 \cup f_1 : X \times \{0, 1\} \rightarrow Y$, where $X \times \{0, 1\}$ is a subcomplex in the metric cell complex $W := X \times I$ (see Exercise 9.41). Then the required extension $f : W \rightarrow Y$ of this map exists by Lemma 9.49. \square

9.4. Connectivity and coarse connectivity

Our next goal is to find a large supply of examples of metric spaces which are coarsely m -connected.

LEMMA 9.51. *If X is a finite-dimensional m -connected complex which admits a properly discontinuous, cocompact, cellular, isometric (on $X^{(0)}$) group action $G \curvearrowright X$, then X is uniformly m -connected.*

PROOF. Existence of the action $G \curvearrowright X$ implies that X is locally finite. Pick a base-vertex $x \in X$ and let $r < \infty$ be such that the G -orbit of $B(x, r) \subset X^{(0)}$ is the entire $X^{(0)}$. Therefore, if a subcomplex $C \subset X$ has diameter $\leq R/2$, there exists $g \in G$ such that $C' = g(C) \subset \mathbf{B}(x, r + R)$.

Since C is finite, its fundamental group $\pi_1(C')$ is finitely generated. Thus, simple connectivity of X implies that there exists a finite subcomplex $C'' \subset X$ such that each generator of $\pi_1(C')$ vanishes in $\pi_1(C'')$. Consider now $\pi_i(C')$, $2 \leq i \leq m$. Then, by Hurewicz theorem, the image of $\pi_i(C')$ in $\pi_i(X) \cong H_i(X)$, is contained in the image of $H_i(C')$ in $H_i(X)$. Since C' is a finite complex, we can choose C'' above such that the map $H_i(C') \rightarrow H_i(C'')$ is zero. To summarize, there exists a finite connected subcomplex C'' in X containing C' , such that all maps $\pi_i(C') \rightarrow \pi_i(C'')$ are trivial, $1 \leq i \leq m$.

Since C'' is a finite complex, there exists $R' < \infty$ such

$$C'' \subset \mathbf{B}(x, r + R + R').$$

Hence, the inclusion map

$$C' \rightarrow \mathbf{B}(x, r + R + R')$$

is m -null.

Set $\psi(k, r) = \rho = r + R'$. Therefore, taking into account action of G on X , we conclude that for each subcomplex $C \subset X$ of diameter $\leq R/2$, the inclusion map

$$C \rightarrow \mathcal{N}_\rho(C)$$

is m -null. □

Our next goal is to relate the notion of coarse n -connectivity from Section 9.2.2 to uniform n -connectivity for metric cell complexes.

THEOREM 9.52. *Suppose that X is a uniformly n -connected metric cell complex of bounded geometry. Then $Z := X^{(0)}$ is coarsely n -connected.*

PROOF. Let $\gamma : \mathbb{S}^k \rightarrow \text{Rips}_R(Z)$ be a spherical m -cycle in $\text{Rips}_R(Z)$, $0 \leq k \leq n$. Without loss of generality (using simplicial approximation) we can assume that γ is a simplicial cycle, i.e. the sphere \mathbb{S}^k is given a triangulation τ such that γ is a simplicial map.

LEMMA 9.53. *There exists a cellular map $\gamma' : (\mathbb{S}^k, \tau) \rightarrow X$ which agrees with γ on the vertex set of τ and such that $\text{diam}(\gamma'(\sigma)) \leq R'$, for each simplex $\sigma \in \tau$, where $R' \geq R$ depends only on R and contractibility functions $\psi_i(k, \cdot)$ of X , $i = 0, \dots, k$.*

PROOF. We construct γ' by induction on skeleta of (\mathbb{S}^k, τ) . The map is already defined on the 0-skeleton, namely, it is the map γ and images of all vertices of τ are within distance $\leq R$ from each other. This map extends to the 1-skeleton of τ : Given an edge $\sigma = [v, w]$ of τ , we extend γ to σ using a path of length $\leq \psi(1, R)$ in $X^{(1)}$ connecting the vertices $\gamma(v), \gamma(w) \in X^{(0)}$.

Suppose we constructed the required extension

$$\gamma' : \tau^{(i)} \rightarrow X$$

such that

$$\text{diam}(\gamma'(\sigma)) \leq R_i = R_i(R, \psi(k, \text{diam}(\sigma)))$$

for each i -simplex σ .

Let σ be an $i + 1$ -simplex in τ . We already have a map γ' defined on the boundary of σ and $\text{diam}(\gamma'(\partial\sigma)) \leq R_i$. Then, as in the proof of Proposition 9.48, using uniform contractibility of X , we extend γ' to the simplex σ , so that the new cellular map γ' map satisfies

$$\text{diam}(\gamma'(\sigma)) \leq \psi(i + 1, R_i).$$

This implies that the image $\gamma'(\sigma)$ is contained in $\mathbf{B}(\gamma(v), 2\psi(i + 1, R_i))$, where v is a vertex of σ . Thus,

$$\text{diam}(\gamma'(\tau^{i+1})) \leq R_{i+1} := R + \psi(i + 1, R_i).$$

Now, the lemma follows by induction. \square

Since X is k -connected, the map γ' extends to a cellular map $\gamma' : \mathbb{D}^{k+1} \rightarrow X^{(k+1)}$, where \mathbb{D}^{k+1} is a triangulated disk whose triangulation \mathcal{T} extends the triangulation τ of \mathbb{S}^k . Our next goal is to “push” γ' into a map $\gamma'' : \mathbb{D}^{k+1} \rightarrow \text{Rips}_{R'}(Z)$, relative to the restriction of γ to the vertex set of τ .

Let σ be a simplex in \mathcal{T} . A simplicial map is determined by images of vertices. By definition of the number R' , images of the vertices of σ under γ' are within distance $\leq R'$ from each other. Therefore, we have a canonical extension γ'' of $\gamma'|_{\sigma(0)}$ to a map $\sigma \rightarrow \text{Rips}_{R'}(Z)$. If σ_1 is a face of σ_2 , then $\gamma'' : \sigma_1 \rightarrow \text{Rips}_{R'}(Z)$ agrees with the restriction of $\gamma'' : \sigma_2 \rightarrow \text{Rips}_{R'}(Z)$, since maps are determined by their vertex values. We thus obtain a simplicial map

$$\gamma'' : \mathbb{D}^{k+1} \rightarrow \text{Rips}_{R'}(Z)$$

which agrees with γ on the boundary sphere. We conclude that the inclusion map $\text{Rips}_R(Z) \rightarrow \text{Rips}_{R'}(Z)$ is n -null, i.e. Z is coarsely n -connected. \square

So far we have seen, how to go from uniform k -connectivity of a metric cell complex X to coarse k -connectivity of its 0-skeleton. Our goal now is to go in the opposite direction: Convert a coarsely k -connected space to a uniformly k -connected metric cell complex.

LEMMA 9.54. *Let G be a finitely generated group with word metric. Then G is coarsely simply-connected if and only if $\text{Rips}_R(G)$ is simply-connected for all sufficiently large R .*

PROOF. One direction is clear, we only need to show that coarse simple connectivity of G implies that $\text{Rips}_R(G)$ is simply-connected for all sufficiently large R . Our argument is similar to the proof of Theorem 9.52. Note that the 1-skeleton of $\text{Rips}_1(G)$ is just a Cayley graph of G . Using coarse simple connectivity of G , we find $D \geq 1$ such that the map

$$\text{Rips}_1(G) \rightarrow \text{Rips}_D(G)$$

is 1-null (i.e. induces trivial map of fundamental groups). We claim that for all $R \geq D$ the Rips complex $\text{Rips}_R(G)$ is simply-connected. Let $\gamma \subset \text{Rips}_R(G)$ be a simplicial loop. For every edge $\gamma_i := [x_i, x_{i+1}]$ of γ we let $\gamma'_i \subset \text{Rips}_1(G)$ denote a geodesic path from x_i to x_{i+1} . The path γ'_i necessarily has length $\leq R$. Therefore, all the vertices of γ'_i are contained in the ball $B(x_i, R) \subset G$ and, hence, span a

simplex in $\text{Rips}_R(G)$. Thus, the paths γ_i, γ'_i are homotopic in $\text{Rips}_R(G)$, relative their end-points. Let γ' denote the loop in $\text{Rips}_1(G)$ which is the concatenation of the paths γ'_i . Then, by the above observation, γ' is freely homotopic to γ in $\text{Rips}_R(G)$. On the other hand, γ' is null-homotopic in $\text{Rips}_R(G)$ since the map

$$\pi_1(\text{Rips}_1(G)) \rightarrow \pi_1(\text{Rips}_R(G))$$

is trivial. We conclude that γ is null-homotopic in $\text{Rips}_R(G)$ as well. \square

COROLLARY 9.55. *Suppose that G is a finitely generated group with the word metric. Then G is finitely presented if and only if G is coarsely simply-connected. In particular, finite-presentability is a QI invariant.*

PROOF. Suppose that G is finitely presented and let Y be its finite presentation complex (see Definition 7.91). Then the universal cover X of Y is simply-connected. Hence, by Lemma 9.51, X is uniformly simply-connected and, by Theorem 9.52, the group G is coarsely simply-connected.

Conversely, suppose that G is coarsely simply-connected. By Lemma 9.54, the simplicial complex $\text{Rips}_R(G)$ is simply-connected for some R . The group G acts on $X := \text{Rips}_R(G)$ simplicially, properly discontinuously and cocompactly. Therefore, by Corollary 5.109, G admits a properly discontinuous, free cocompact action on another simply-connected cell complex Z . It follows that G is finitely presented. \square

We now proceed to $k \geq 2$. Recall (see Definition 5.105) that a group G has type \mathbf{F}_n ($n \leq \infty$) if it admits a free cellular action on a cell complex X such that for each $k \leq n$:

1. $X^{(k+1)}/G$ is compact.
2. $X^{(k+1)}$ is k -connected.

THEOREM 9.56 (M. Gromov, 1.C2 in [Gro93] and J. Alonso, [Alo94]). *Type \mathbf{F}_n is a QI invariant for each $n \leq \infty$.*

PROOF. Our argument is similar to the proof of Corollary 9.55, except we cannot rely on $n - 1$ -connectivity of Rips complexes $\text{Rips}_R(G)$ for large R . If G has type \mathbf{F}_n , then it admits a free cellular action $G \curvearrowright X$ on some $n - 1$ -connected cell complex X , so that the quotient of each skeleton (of dimension $\leq n$) is a finite complex. (In the case $n = \infty$, we require, of course, the entire X to be contractible and the quotient of each skeleton to be finite.) By combining Lemma 9.51 and Theorem 9.52, we see that the group G is coarsely $n - 1$ -connected. It remains to prove

PROPOSITION 9.57. *If G is a coarsely $n - 1$ -connected group, then G has type \mathbf{F}_n .*

PROOF. Note that we already proved this statement for $n = 2$: Coarsely simply-connected groups are finitely presented (Corollary 9.55). The proof below follows [KK05]. We break the argument in three parts: We first consider the case when G is torsion-free and $n < \infty$, then the case when G is still torsion-free but $n = \infty$ and, lastly, the general case.

Our goal is to build a complex X on which G acts as required by the definition of type \mathbf{F}_n . We construct this complex and the action by induction on skeleta $X^{(0)} \subset \dots \subset X^{(n-1)} \subset X^{(n)}$. Furthermore, we will inductively construct cellular G -equivariant maps

$$f_i : X_i = X^{(i)} \rightarrow Y_{R_i} = \text{Rips}_{R_i}(G)$$

$$\bar{f}_i : Y_{R_i}^{(i)} \rightarrow X_i, i = 0, \dots, n,$$

and (cellular G -equivariant) homotopies

$$H_i : X_{i-1} \times [0, 1] \rightarrow X_i$$

of

$$h_{i-1} := \bar{f}_{i-1} \circ f_{i-1} : X_{i-1} \rightarrow X_{i-1} \subset X_i$$

to the inclusion maps $X_{i-1} \hookrightarrow X_i$.

1. Torsion-free case, $n < \infty$. In this case the G -action on every Rips complex is free and cocompact. Our construction is by induction on i .

$i = 0$. We let $X_0 = G, R_0 = 0$ and let $f_0 = \bar{f}_0 : G \rightarrow G$ be the identity map.

$i = 1$. We let $R_1 = 1$ and let $X_1 = Y_{R_1}^{(1)}$ be the Cayley graph of G . Again, $f_1 = \bar{f}_1 = \text{Id}$, and, of course, $H_0(x, t) = x$.

$i = 2$. According to Lemma 9.54, there exists R_2 so that Y_R is simply-connected for all $R \geq R_2$. We then take $X_2 := Y_{R_2}^{(2)}$. Again, we let $f_2 = \bar{f}_2 = \text{Id}, H_1(x, t) = x$.

$i \Rightarrow i + 1$. Suppose now that $3 \leq i \leq n - 1, X_i, f_i, \bar{f}_i, H_i$ are constructed and R_i chosen; we will construct $X_{i+1}, f_{i+1}, \bar{f}_{i+1}$ and H_{i+1} .

In the arguments below we will be using unbased spherical cycles when dealing with homotopy groups of X_i : This is harmless since X_i is $i - 1$ -connected and we can identify homotopy and homology groups (in degree i) via Hurewicz theorem. Our first task is to extend the homotopy H_i from $X_{i-1} \times [0, 1]$ to $X_i \times [0, 1]$. This is impossible without increasing the dimension of X_i and this will be the first step of our construction.

LEMMA 9.58. *There exists a bounded geometry cell complex Z_{i+1} of dimension $i + 1$ whose i -skeleton is X_i , such that:*

1. *The G -action extends from X_i to a free cellular properly discontinuous cocompact action on Z_{i+1} .*

2. *The homotopy $H_i : X_{i-1} \times [0, 1] \rightarrow X_i$ extends to a G -equivariant homotopy $H_{i+1} : X_i \times [0, 1] \rightarrow Z_{i+1}$ between the map h_i and the inclusion map.*

PROOF. There are only finitely many i -cells in X_i modulo the G -action. It suffices to extend H_i to the finitely many cells $\hat{e}_\gamma : \mathbb{D}^i \rightarrow X_i$ in each G -orbit. Consider the $i + 1$ -ball $\mathbb{D}^i \times [0, 1]$. The homotopy H_i lifts to a homotopy $\hat{H}_i : \partial\mathbb{D}^i \times [0, 1] \rightarrow X_i$ between the map $h_{i-1} \circ e_\gamma$ and the attaching map e_γ ; furthermore, we are also given maps e_γ and $h_i \circ e_\gamma$ on $\mathbb{D}^i \times \{1\}$ and $\mathbb{D}^i \times \{1\}$ respectively. If we knew that the resulting map of the boundary sphere of $\mathbb{D}^i \times [0, 1]$

$$\epsilon_\gamma : \partial(\mathbb{D}^i \times [0, 1]) \rightarrow X_i$$

is null-homotopic, we would be able to construct the required extension \hat{H}_{i+1} . There is no reason, of course, for this null-homotopy (since X_i is only required to be $i - 1$ -connected and not i -connected). Therefore, we attach an $i + 1$ -cell to X_i along the map ϵ_γ .

Since the homotopy H_i was G -equivariant, we can attach these cells in G -equivariant fashion. The result is the G -complex Z_{i+1} . Proper discontinuity of the action of G on X_i ensures that Z_{i+1} has bounded geometry and the action $G \curvearrowright Z_{i+1}$ is properly discontinuous and cocompact. Freeness of the action of G follows from the fact that G is torsion-free. \square

The next step is to construct X_{i+1} by enlarging Z_{i+1} . Let $R' > R = R_i$ be such that the inclusion map

$$Y_R = \text{Rips}_R(G) \rightarrow Y_{R'} = \text{Rips}_{R'}(G)$$

is i -null. Since X_i is $i - 1$ -connected, the map

$$\bar{f}_i : Y_R^{(i)} \rightarrow X_i$$

extends to a cellular G -equivariant map

$$\tilde{f}_i : Y_{R'}^{(i)} \rightarrow X_i,$$

as in the proof of Proposition 9.48.

LEMMA 9.59. *There exists a finite set of spherical classes $[\sigma_\alpha], \alpha \in A'$, in $H_i(Z_{i+1})$, which generates $H_i(Z_{i+1})$ as a G -module.*

PROOF. We let $\{\Delta_\alpha : \alpha \in A\}$ denote the set of $i + 1$ -simplices in $Y_{R'}$. For each simplex Δ_α we let

$$\tau_\alpha : \partial\Delta_\alpha \rightarrow Y_{R'}$$

denote the inclusion map. We will identify the boundary of Δ_α with a triangulated sphere \mathbb{S}^i and think of the maps τ_α as spherical cycles in $Y_{R'}$. Since the map $H_i(Y_R) \rightarrow H_i(Y_{R'})$ is trivial, each $[\eta] \in H_i(Y_{R'})$ has the form

$$[\eta] = \sum_{\alpha \in A} z_\alpha [\partial\Delta_\alpha], \quad z_\alpha \in \mathbb{Z}.$$

In other words, in the group $H_i(Y_{R'}^{(i)})$ we have the equality

$$[\eta] = \sum_{\alpha \in A} z_\alpha [\tau_\alpha].$$

Since the action of G on Y_R is cocompact, there exists a finite subset $A' \subset A$, such that each cycle $\tau_\beta, \beta \in A$ belongs to the G -orbit of some $\tau_\alpha, \alpha \in A'$. In other words, the image M of $H_i(Y_{R'}^{(i)})$ in $H_i(Y_{R'})$ is a G -submodule of the finitely generated G -module M' with the generators

$$\{[\tau_\alpha] : \alpha \in A'\}.$$

Each $[\sigma] \in H_i(Z_{i+1})$ is represented by a (spherical) cycle σ in X_i and

$$(\bar{f}_i \circ f_i)_*([\sigma]) = [\sigma]$$

in $H_i(Z_{i+1})$ because of the homotopy H_{i+1} . Therefore, $[\sigma]$ belongs to the finitely generated G -module $(\tilde{f}_i)_*(M')$ whose generators are represented by spherical cycles

$$\sigma_\alpha := \tilde{f}_i(\tau_\alpha), \alpha \in A'.$$

We conclude that the G -module $H_i(Z_{i+1})$ is generated by the finite set $[\sigma_\alpha], \alpha \in A'$. \square

We now use the maps $e_\alpha = g \circ \sigma_\alpha, \alpha \in A'$, as attaching maps for $i + 1$ -cells, and let X_{i+1} denote the cell complex obtained by (equivariantly) attaching cells to Z_{i+1} along these maps. Recall, for a future reference, that $\hat{e}_\alpha : \mathbb{D}^{i+1} \rightarrow X_{i+1}$ denotes the $i + 1$ -cell defined via the attaching map e_α . The G -action extends from Z_{i+1} to a free cocompact properly discontinuous action on X_{i+1} . By the construction X_{i+1} is i -connected, since we killed $\pi_i(Z_{i+1})$ by attaching $i + 1$ -cells along its generators.

We next construct the maps f_{i+1} and \bar{f}_{i+1} . To construct the map $f_{i+1} : X_{i+1} \rightarrow Y_{R'}$, for each $\alpha \in A', g \in G$, we extend the map $f_i \circ g\sigma_\alpha : \mathbb{S}^i \rightarrow X_i$ to \mathbb{D}^{i+1} G -equivariantly using vanishing of the map

$$\pi_i(Y_R) \rightarrow \pi_i(Y_{R'}).$$

The construction of \bar{f}_{i+1} is similar: We already have an equivariant map $\tilde{f}_i : Y_{R'}^{(i)} \rightarrow X_i$. We extend this map to each $i + 1$ -simplex $g\Delta_\alpha \cong \mathbb{D}^{i+1}, \alpha \in A'$, using the map $g \circ \hat{e}_\alpha : \mathbb{D}^{i+1} \rightarrow X^{(i+1)}$. This concludes the proof in the case when G is torsion-free and n is finite. Note that at the last step of the construction, we only get a homotopy H_n between h_{n-1} and Id: As we noted above, there is no reason for the map h_n to be homotopic to the identity.

2. $n = \infty$. The inductive construction described in the proof, runs indefinitely. We obtain an increasing sequence of $i - 1$ -connected i -dimensional G -complexes X_i . Let X be the union

$$\bigcup_{i \geq 0} X_i$$

equipped with the weak topology. Since each X_i is $i - 1$ -connected, the complex X is contractible. The group G acts cellularly and freely on X , since it acts this way on each i -skeleton. The quotients X_i/G are finite for every $i \in \mathbb{N}$ and the action of G on each X_i is free and properly discontinuous. This concludes the proof in the case of torsion-free groups G .

3. General Case. We now explain what to do in the case when G is not torsion-free. The main problem is that a group G with torsion will not act freely on its Rips complexes. Thus, while equivariant maps f_i still exist, we would be unable to construct equivariant maps $\bar{f}_i : \text{Rips}_R(G) \rightarrow X_i$. Furthermore, it could happen that, for large R , the complex Y_R is contractible: This is clearly true if G is finite, it also holds for all Gromov-hyperbolic groups. If we were to have f_i and \bar{f}_i as before, we would be able to conclude that X_i is contractible for large i , while a group with torsion cannot act freely on a contractible cell complex.

We, therefore, have to modify the construction. For each R we let W_R denote the barycentric subdivision of $Y_R^{(i)} = \text{Rips}_R(G)^{(i)}$. Then G acts on W_R without inversions (see Definition 5.95). Let \widehat{W}_R denote the regular cell complex obtained by applying the *Haefliger construction* to W_R , see Section 5.8. The complex \widehat{W}_R is infinite-dimensional if G has torsion, but this does not cause trouble since at each step of induction we work only with finite skeleta. The action $G \curvearrowright W_R$ lifts to a free (properly discontinuous) action $G \curvearrowright \widehat{W}_R$ which is cocompact on each skeleton. We then can apply the arguments from the torsion-free case to the complexes \widehat{W}_R instead of $\text{Rips}_R(G)$. The key is that, since the action of G on \widehat{W}_R is free, the construction of the equivariant maps $\bar{f}_i : Y_{R_i}^{(i)} \rightarrow X^{(i)}$ goes through. Note also that in the first steps of the induction we used the fact that Y_R is simply-connected for sufficiently large R in order to construct $X^{(2)}$. Since the projection $\widehat{W}_R \rightarrow W_R$ is a homotopy-equivalence, the 2-skeleton of \widehat{W}_R is simply-connected for the same values of R . □

This finishes the proof of Theorem 9.56 as well. □

COROLLARY 9.60. *The condition of having the type \mathbf{F}_n , $1 \leq n \leq \infty$, is a VI invariant.*

The condition \mathbf{F}_n has *cohomological analogues*, for instance, the condition \mathbf{FP}_n , see [Bro82b]. The arguments used in this section apply in the context of \mathbf{FP}_n -groups as well, see Alonso's paper [Aea91] as well as Proposition 11.4 in [KK05]. The main difference is that instead of metric cell complexes, one works with metric chain complexes and instead of k -connectedness of the system of Rips complexes, one uses acyclicity over commutative rings R :

THEOREM 9.61 (M. Kapovich, B. Kleiner, [KK05]). *Let R be a commutative ring with a unit. Then the property of being \mathbf{FP}_n over R is QI invariant.*

QUESTION 9.62. 1. Is the homological dimension (over \mathbb{Q}) of a group a quasi-isometric invariant?

2. Suppose that G has geometric dimension $n < \infty$. Is there a bounded geometry uniformly contractible n -dimensional metric cell complex with free G -action $G \curvearrowright X$?

3. Is the geometric dimension a quasiisometric invariant for torsion-free groups?

4. Is the property of having the type \mathbf{F} invariant under quasiisometries?

QUESTION 9.63. Suppose that G_1, G_2 are finitely generated torsion-free quasi-isometric groups. Is it true that $cd(G_1) = cd(G_2)$?

According to R. Sauer, [Sau06], the problem reduces to showing that G_1 has finite cohomological dimension if and only if G_2 does.

Note that cohomological dimension is (mostly) known to equal geometric dimension, except there could be groups satisfying

$$2 = cd(G) \leq gd(G) \leq 3,$$

see [Bro82b]. On the other hand,

$$cd(G) \leq hd(G) \leq cd(G) + 1,$$

see [Bie76a]. Here cd stands for cohomological dimension, gd is the geometric dimension and hd is the homological dimension. QI invariance of cohomological dimension (over \mathbb{Q}) was proven by R. Sauer:

THEOREM 9.64 (R. Sauer [Sau06]). *The cohomological dimension $cd_{\mathbb{Q}}$ of a group (over \mathbb{Q}) is a QI invariant. Moreover, if G_1, G_2 are groups and $f : G_1 \rightarrow G_2$ is a quasiisometric embedding, then $cd_{\mathbb{Q}}(G_1) \leq cd_{\mathbb{Q}}(G_2)$.*

Note that partial results on QI invariance of cohomological dimension were proven earlier by P. Pansu [Pan83] (for virtually nilpotent groups), S. Gersten, [Ger93b] (for groups of type FP_n) and Y. Shalom, [Sha04] (for amenable groups).

9.5. Retractions

The goal of this section is to give a non-equivariant version of the construction of the retractions ρ_i from the proof of Proposition 9.57 in the previous section.

Suppose that X, Y are uniformly contractible finite-dimensional metric cell complexes of bounded geometry. Consider a uniformly proper map $f : X \rightarrow Y$. Our goal is to define a *coarse left-inverse* to f , a *retraction* ρ which maps an r -neighborhood of $V := f(X)$ back to X .

LEMMA 9.65. *Under the above assumptions, there exist numbers L, L', A , a function $R = R(r)$ which depend only on the distortion function of f and on the geometry of X and Y , such that:*

1. For every $r \in \mathbb{N}$ there exists a cellular L -Lipschitz map $\rho = \rho_r : \mathcal{N}_r(V) \rightarrow X$ so that $\text{dist}(\rho \circ f, \text{Id}_X) \leq A$. Here and below we equip $W^{(0)}$ with the restriction of the path-metric on the metric graph $W^{(1)}$ in order to satisfy Axiom 1 of metric cell complexes.

2. $\rho \circ f$ is homotopic to the identity by an L' -Lipschitz cellular homotopy.

3. The composition $h = f \circ \rho : \mathcal{N}_r(V) \rightarrow V \subset \mathcal{N}_R(V)$ is homotopic to the identity embedding $\text{Id} : V \rightarrow \mathcal{N}_R(V)$.

4. If $r_1 \leq r_2$, then $\rho_{r_2}|_{\mathcal{N}_{r_1}(V)} = \rho_{r_1}$.

PROOF. Let $D_0 = 0, D_1, D_2, \dots$ denote the geometric bounds on Y and

$$\max_{k>0} D_k = D < \infty.$$

Since f is uniformly proper, there exists a proper monotonic function $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that

$$\eta(d(x, x')) \leq d(f(x), f(x')), \forall x, x' \in X^{(0)}.$$

Let A_0, A_1 denote real numbers for which

$$\eta(t) > 0, \quad \forall t > A_0,$$

$$\eta(t) > 2r + D_1, \quad \forall t > A_1.$$

Recall that the neighborhood $W := \bar{\mathcal{N}}_r(V)$ is a subcomplex of Y . For each vertex $y \in W^{(0)}$ we pick a vertex $\rho(y) := x \in X^{(0)}$ such that the distance $\text{dist}(y, f(x))$ is the smallest possible. If there are several such points x , we pick one of them arbitrarily. The fact that f is uniformly proper, ensures that

$$\text{dist}(\rho \circ f, \text{Id}_{X^{(0)}}) \leq A := A_0.$$

Indeed, if $\rho(f(x)) = x'$, then $f(x) = f(x')$; if $d(x, x') > A_0$, then

$$0 < \eta(d(x, x')) \leq d(f(x), f(x')),$$

contradicting that $f(x) = f(x')$. Thus, by our choice of the metric on $W^{(0)}$ coming from $W^{(1)}$, we conclude that ρ is A_1 -Lipschitz.

Next, observe also that for each 1-cell σ in W , $\text{diam}(\rho(\partial\sigma)) \leq A_1$. Indeed, if $\partial\sigma = \{y_1, y_2\}$, then $d(y_1, y_2) \leq D_1$, by the definition of a metric cell complex. For $y'_i := f(x_i)$, $d(y_i, y'_i) \leq r$. Thus, $d(y'_1, y'_2) \leq 2r + D_1$ and $d(x_1, x_2) \leq A_1$, by the definition of A_1 . Now, existence of L -Lipschitz extension $\rho : W \rightarrow X$ follows from Proposition 9.48. This proves (1).

Part (2) follows from Corollary 9.50. To prove Part (3), observe that $h = f \circ \rho : \bar{\mathcal{N}}_r(V) \rightarrow V$ is L' -Lipschitz (see Exercise 9.42), $\text{dist}(h, \text{Id}) \leq r$. Now, (3) follows from Corollary 9.50 since Y is also uniformly contractible.

Lastly, in order to guarantee (4), we can construct the retractions ρ_r by induction on the values of r and using the Extension Lemma 9.49. \square

COROLLARY 9.66. *There exists a function $\alpha(r) \geq r$, such that for every r the map $h = f \circ \rho : \mathcal{N}_r(V) \rightarrow \mathcal{N}_{\alpha(r)}(V)$ is properly homotopic to the identity, where $V = f(X)$.*

We will think of this lemma and its corollary as a proper homotopy-equivalence between X and the direct system of metric cell complexes $\bar{\mathcal{N}}_R(V)$, $R \geq 1$. Recall that the usual proper homotopy-equivalence induces isomorphisms of compactly supported cohomology groups. In our case we get an ‘‘approximate isomorphism’’ of $H_c^*(X)$ to the inverse system of compactly supported cohomology groups $H_c^*(\bar{\mathcal{N}}_R(V))$:

COROLLARY 9.67. 1. The induced maps $\rho_R^* : H_c^*(X) \rightarrow H_c^*(\overline{\mathcal{N}}_R(V))$ are injective.

2. The induced maps ρ_R^* are approximately surjective in the sense that the subgroup $\text{coker}(\rho_{\alpha(R)}^*)$ maps to zero under the map induced by the restriction map

$$\text{rest}_R : H_c^*(\overline{\mathcal{N}}_{\alpha(R)}(V)) \rightarrow H_c^*(\overline{\mathcal{N}}_R(V)).$$

PROOF. 1. Follows from the fact that $\rho \circ f$ is properly homotopic to the identity and, hence, induces the identity map of $H_c^*(X)$, which means that f^* is the right-inverse to ρ_R^* .

2. By Corollary 9.66 the restriction map rest_R equals the map $\rho_R^* \circ f^*$. Therefore, the cohomology group $H_c^*(\overline{\mathcal{N}}_{\alpha(R)}(V))$ maps via rest_R to the image of ρ_R^* . The second claim follows. \square

9.6. Poincaré duality and coarse separation

In this section we discuss coarse implications of Poincaré duality in the context of triangulated manifolds. For a more general version of Poincaré duality, we refer the reader to [Roe03]; this concept was coarsified in [KK05], where coarse Poincaré duality was introduced and used in the context of metric cell complexes. We will be working with metric cell complexes which are simplicial complexes, the main reason being that Poincaré duality has cleaner statement in this case.

Let X be a connected simplicial complex of bounded geometry, which is a triangulation of a (possibly non-compact) n -dimensional manifold without boundary. Suppose that $W \subset X$ is a subcomplex, which is a triangulated manifold (possibly with boundary). We will use the notation W' to denote its first barycentric subdivision. We then have the Poincaré duality isomorphisms (see e.g. [Dol80, 7.12])

$$P_k : H_c^k(W) \rightarrow H_{n-k}(W, \partial W) = H_{n-k}(X, X \setminus W),$$

see e.g. [Dol80, 7.13]. Here, H_c^* are the cohomology groups with compact support. The Poincaré duality isomorphisms are *natural* in the sense that they commute with proper embeddings of manifolds and manifold pairs. Furthermore, the isomorphisms P_k move cocycles by *uniformly bounded amount*: Suppose that $\zeta \in Z_c^k(W)$ is a simplicial cocycle supported on a compact subcomplex $K \subset W$. Then the corresponding relative cycle $P_k(\zeta) \in Z_{n-k}(W, \partial W)$ is represented by a simplicial chain in W' , where each simplex has non-empty intersection with K .

EXERCISE 9.68. If $W \subsetneq X$ is a proper subcomplex, then $H_c^n(W) = 0$.

We will also have to use the Poincaré duality in the context of subcomplexes $V \subset X$ which are not submanifolds with boundary. Such V , nevertheless, admits a (closed) regular neighborhood $W = \mathcal{N}(V)$, which is a submanifold with boundary. The neighborhood W is homotopy-equivalent to V .

In this section we will present two applications of Poincaré duality to the coarse topology of X .

Coarse surjectivity

THEOREM 9.69. Let X, Y be uniformly contractible simplicial complexes of bounded geometry homeomorphic to \mathbb{R}^n . Then every uniformly cellular proper map $f : X \rightarrow Y$ is surjective.

PROOF. Assume to the contrary, i.e. $V = f(X) \neq Y$ is a proper subcomplex. Thus, $H_c^n(V) = 0$ by Exercise 9.68. Let $\rho : V \rightarrow X$ be a retraction constructed in Lemma 9.65. By Lemma 9.65, the composition $h = \rho \circ f : X \rightarrow X$ is properly homotopic to the identity. Thus, this map has to induce an isomorphism $H_c^*(X) \rightarrow H_c^*(X)$. However, $H_c^n(X) \cong \mathbb{Z}$ since X is homeomorphic to \mathbb{R}^n , while $H_c^n(V) = 0$. Contradiction. \square

COROLLARY 9.70. *Let X, Y be as above and let $f : X^{(0)} \rightarrow Y^{(0)}$ be a quasiisometric embedding. Then f is a quasiisometry.*

PROOF. Combine Proposition 9.48 with Theorem 9.69. \square

Coarse separation.

Suppose that X is a simplicial complex and $W \subset X$ is a subcomplex. Consider $\mathcal{N}_R(W)$, the open metric R -neighborhoods of W in X , and their complements C_R in X .

For a component $C \subset C_R$ define the *inradius*, $\text{Inrad}(C)$, of C to be the supremum of radii of balls $\mathbf{B}(x, R)$ in X contained in C . A component C is called *shallow* if $\text{Inrad}(C)$ is finite and *deep* if $\text{Inrad}(C) = \infty$.

EXAMPLE 9.71. Suppose that W is compact. Then deep complementary components of C_R are components of infinite diameter. These are the components which appear as neighborhoods of ends of X .

A subcomplex W is said to *coarsely separate* X if there is R such that $\mathcal{N}_R(W)$ has at least two distinct deep complementary components.

EXAMPLE 9.72. A simple properly embedded curve Γ in \mathbb{R}^2 need not coarsely separate \mathbb{R}^2 (see Figure 9.3). A straight line in \mathbb{R}^2 coarsely separates \mathbb{R}^2 .



FIGURE 9.3. A separating curve which does not coarsely separate the plane.

The following theorem is a coarse analogue of the Jordan separation theorem which states that for the image of an arbitrary proper embedding $f : \mathbb{R}^{n-1} \rightarrow \mathbb{R}^n$ separates \mathbb{R}^n into exactly two components. This topological theorem follows immediately from the Jordan separation theorem for spheres, since we can take the one-point compactifications of $A = f(\mathbb{R}^{n-1})$ and \mathbb{R}^n . Properness of f ensures that the compactification of A is homeomorphic to \mathbb{S}^{n-1} . The proof of the coarse Jordan separation theorem follows the same arguments and the proof of the topological separation theorem (via Poincaré duality).

THEOREM 9.73 (Coarse Jordan separation). *Suppose that X and Y are uniformly contractible simplicial complexes of bounded geometry, homeomorphic to \mathbb{R}^{n-1} and \mathbb{R}^n , respectively. Then for each uniformly proper simplicial map $f : X \rightarrow Y$, the image $V = f(X)$ coarsely separates Y . Moreover, for all sufficiently large R , $Y \setminus \mathcal{N}_R(V)$ has exactly two deep components.*

PROOF. Actually, our proof will use the assumption on the topology of X only weakly: To get coarse separation it suffices to assume that $H_c^{n-1}(X) \neq 0$.

Recall that in Section 9.5 we constructed a system of retractions

$$\rho_R : \mathcal{N}_R(V) \rightarrow X, \quad R \in \mathbb{N}$$

and proper homotopy-equivalences $f \circ \rho \equiv \text{Id}$, for which

$$\rho_R \circ f|_{\mathcal{N}_R(V)} \equiv \text{Id} : \mathcal{N}_R(V) \rightarrow \mathcal{N}_{\alpha(R)}(V).$$

Furthermore, we have the *restriction maps*

$$\text{rest}_{R_1, R_2} : H_c^*(\tilde{\mathcal{N}}_{R_2}(V)) \rightarrow H_c^*(\tilde{\mathcal{N}}_{R_1}(V)), \quad R_1 \leq R_2.$$

These maps satisfy

$$\text{rest}_{R_1, R_2} \circ \rho_{R_2}^* = \rho_{R_1}^*$$

by Part 4 of Lemma 9.65. We also have the projection maps

$$\text{proj}_{R_1, R_2} : H_*(Y, Y - \tilde{\mathcal{N}}_{R_2}(V)) \rightarrow H_*(Y, Y - \tilde{\mathcal{N}}_{R_1}(V)), \quad R_1 \leq R_2,$$

induced by inclusion maps of pairs $(Y, Y - \tilde{\mathcal{N}}_{R_2}(V)) \hookrightarrow (Y, Y - \tilde{\mathcal{N}}_{R_1}(V))$. The Poincaré duality in Y also gives us a system of isomorphisms

$$P : H_c^*(\tilde{\mathcal{N}}_R(V)) \cong H_{n-*}(Y, Y - \mathcal{N}_R(V)).$$

By naturality of the Poincaré duality, we have a commutative diagram:

$$\begin{array}{ccc} H_c^*(\tilde{\mathcal{N}}_{R_2}(V)) & \xrightarrow{P} & H_{n-*}(Y, C_{R_2}) \\ \text{rest}_{R_1, R_2} \downarrow & & \downarrow \text{proj}_{R_1, R_2} \\ H_c^*(\tilde{\mathcal{N}}_{R_1}(V)) & \xrightarrow{P} & H_{n-*}(Y, C_{R_1}) \end{array}$$

where P 's are the Poincaré duality isomorphisms.

Let ω be a generator of $H_c^{n-1}(X) \cong \mathbb{R}$. Given $R > 0$ consider the pull-back $\omega_R := \rho_R^*(\omega)$ and the relative cycle $\sigma_R = P(\omega_R)$. Then $\omega_r = \text{rest}_{r, R}(\omega_R)$ and

$$\sigma_r = \text{proj}_{r, R}(\sigma_R) \in H_1(Y, C_r),$$

for all $r < R$, see Figure 9.4. Observe that for every r , ω_r is non-zero, since $f^* \circ \rho^* = \text{Id}$ on the compactly supported cohomology of X . Hence, every σ_r is non-zero as well.

Contractibility of Y and the long exact sequence of the homology groups of the pair (Y, C_r) implies that

$$H_1(Y, C_r) \cong \tilde{H}_0(C_r).$$

We let τ_r denote the image of σ_r under this isomorphism. The class τ_r is represented by a 0-cycle, the boundary of the chain representing σ_r . Running the Poincaré duality in the reverse and using the fact that ω is a generator of $H_c^{n-1}(X)$, we see that τ_r is represented by the difference $y'_r - y''_r$, where $y'_r, y''_r \in C_r$. Nontriviality




FIGURE 9.4. Coarse separation.

of τ_r means that y_r', y_r'' belong to distinct components C_r', C_r'' of C_r . Furthermore, since for $r < R$,

$$proj_{r,R}(\sigma_R) = \sigma_r,$$

it follows that

$$C_R' \subset C_r', \quad C_R'' \subset C_r''.$$

Because this can be done for arbitrarily large r, R , we conclude that both components C_r', C_r'' are deep. The same argument run in the reverse implies that there are exactly two deep complementary components. \square

We refer to [FS96] and [KK05] for further discussion and generalization of coarse separation and coarse Poincaré/Alexander duality.

9.7. Metric filling functions

This is a technical section. Here we define *coarse* notions of loops, filling disks, isoperimetric functions and minimal filling area in the setting of geodesic metric spaces, following [Bow91] and [Gro93]. We then relate them to the notions of volume, area and Dehn functions defined earlier, in sections 3.4, 7.10.1, 7.10.4. We further show that growth rates of the functions thus defined are stable under quasiisometry.

Throughout this section, (X, dist) will be a coarsely connected metric space. Thus, there exists a constant ρ , which we fix once and for all, such that the Rips complex $\text{Rips}_R(X)$ is connected for all $R \geq \rho$. Given any pair of points $x, y \in X$, consider the shortest edge-path

$$p_{xy} = [x = q_{xy}(0), q_{xy}(1)] \cup \dots \cup [q_{xy}(n-1), q_{xy}(n) = y],$$

in $\text{Rips}_\rho(X)$ connecting x to y . The map $q_{xy} : \{0, \dots, n\} \rightarrow X$ is the corresponding *vertex-path* connecting x to y . We let $\ell(p_{x,y}) = n$ denote the combinatorial length p_{xy} .

Since X is a quasigeodesic metric space, the map

$$X \rightarrow \text{Rips}_\rho(X)$$

is a quasiisometry. In particular,

$$\frac{1}{\rho} \text{dist}_X(x, y) \leq \text{length}(p_{x,y}) \leq k \text{dist}_X(x, y) + a$$

for some uniform constants k and a .

9.7.1. Coarse isoperimetric functions and coarse filling radius. Our first goal is to discretize/coarsify the usual notions of Lipschitz maps of the unit circle and the unit disk to X . We fix a number $\delta > 0$, the *scale of coarsification*. Our definitions follow the ones in [Gro93, Chapter 5] and [Bow91].

For a triangulation \mathcal{T} of the circle \mathbb{S}^1 , we let $V(\mathcal{T})$ and $E(\mathcal{T})$ denote the vertex and edge sets of \mathcal{T} . Similarly, if \mathcal{D} is a triangulation of the disk \mathbb{D}^2 , extending \mathcal{T} , then $V(\mathcal{D}), E(\mathcal{D}), F(\mathcal{D})$ will denote the sets of vertices, edges and 2-dimensional faces of \mathcal{D} .

A (*coarse*) δ -loop in X is a pair consisting of a triangulated circle $(\mathbb{S}^1, \mathcal{T})$ and a map

$$\mathbf{c} : V(\mathcal{T}) \rightarrow X$$

such that for every edge $e = [u, w] \in E(\mathcal{T})$,

$$(9.3) \quad \text{dist}_X(\mathbf{c}(u), \mathbf{c}(w)) \leq \delta.$$

If X is geodesic, one can define a *geodesic extension* of \mathbf{c} to a Lipschitz map of the entire circle, sending each edge of the triangulation to a geodesic segment connecting images of its end-points. In view of non-uniqueness of geodesics in X , the geodesic extension is not unique, nevertheless, by abusing the notation, we will denote it by $\tilde{\mathbf{c}}$.

We let $\Omega_\delta(X)$ denote the space of δ -loops in X . We then define the δ -length function

$$\ell = \ell_\delta : \Omega_\delta(X) \rightarrow \mathbb{N}.$$

The δ -length $\ell_\delta(\mathbf{c})$ of the δ -loop \mathbf{c} is the number of edges in the triangulation \mathcal{T} .

Similarly, a (*coarse*) δ -disk is a map

$$\mathfrak{d} : V(\mathcal{D}) \rightarrow X$$

satisfying the inequality (9.3) for every edge of \mathcal{D} . The δ -disk \mathfrak{d} is a (*coarse*) *filling disk* of the coarse loop \mathbf{c} , if the triangulations \mathcal{T} and \mathcal{D} agree on the boundary circle and

$$\mathbf{c} = \mathfrak{d}|_{V(\mathcal{T})}.$$

Thus, a filling disk is a discretization of the notion of a Lipschitz-continuous extension $\mathbb{D}^2 \rightarrow X$ of a Lipschitz-continuous map $\mathbb{S}^1 \rightarrow X$.

Let τ be a 2-face of \mathcal{D} , with the vertex set $V(\tau)$. The restrictions

$$\mathfrak{d}|_{V(\tau)}, \tau \in F(\mathcal{D}),$$

are called *bricks* of the coarse filling disk \mathfrak{d} .

The combinatorial area of a coarse filling disk (see Section 7.10.1) is the number of 2-simplices in the triangulation \mathcal{D} , i.e. the number of bricks.

DEFINITION 9.74. The δ -filling area of the coarse loop \mathbf{c} is defined to be the minimum of combinatorial areas $\text{Area}^{\text{com}}(\mathfrak{d})$ of δ -filling disks \mathfrak{d} of \mathbf{c} . We will use both notation $\text{Ar}_\delta(\mathbf{c})$ and $\text{P}(\mathbf{c}, \delta)$ for the δ -filling area.

To motivate the definition, suppose for a moment that X is a Riemannian manifold of bounded geometry. Then every brick in a δ -filling disk can be filled in with a smooth triangle of the area $\leq C\delta^2$, where C is a uniform constant. Therefore, the filling area, in this case, approximates (as δ tends to zero) δ^2 times the least area of a singular disk in X bounded by the loop $\tilde{\mathbf{c}}$.

We, likewise, define the δ -filling radius function as

$$r_\delta : \Omega_\delta(X) \rightarrow \mathbb{R}_+,$$

$$r_\delta(\mathbf{c}) = \inf \left\{ \max_{x \in V(\mathcal{D})} \text{dist}_X(\mathfrak{d}(x), \mathbf{c}(V(\mathcal{T}))) : \mathfrak{d} \text{ is a } \delta\text{-filling disk of the loop } \mathbf{c} \right\}.$$

Again, in the case when X is a Riemannian manifold of bounded geometry, the function r_δ approximates (as δ tends to zero) the radius of the least radius singular 2-disk bounding the geodesic extension of the loop \mathbf{c} .

Both functions Ar_δ and r_δ depend on the parameter δ , and may take infinite values. In order to obtain real-valued functions, we add the hypothesis that X is coarsely simply connected, i.e. there exists $\mu \geq \rho > 0$, such that for all $\delta \geq \mu$, every δ -loop in X admits a δ -filling disk. Recall (Corollary 9.36) that quasiisometries preserve coarse simple connectivity.

Suppose that X is μ -simply connected and $\delta \geq \mu$. We define the δ -coarse (one-dimensional) isoperimetric function

$$\text{Ar}_\delta = \text{Ar}_{\delta, X} = IP_{\delta, X, 1}^{\text{coarse}} : \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$$

by

$$\text{Ar}_\delta(\ell) := \sup\{\text{Ar}_\delta(\mathbf{c}) : \mathbf{c} \in \Omega_\delta(X), \ell_\delta(\mathbf{c}) \leq \ell\},$$

i.e. the maximal δ -area needed to fill in a coarse loop of δ -length at most ℓ . When δ is fixed, we will also refer to the function Ar_δ as the *coarse isoperimetric function* or the *coarse filling area function*.

The function Ar_δ is a coarsification of the classical isoperimetric functions from the Riemannian geometry $IP_M = IP_{M, 1}$ defined via maps of 2-disks into a Riemannian manifold M , see Section 3.5.

The following theorem relates the coarse isoperimetric functions to the Riemannian ones:

THEOREM 9.75. (Cf. [BT02].) *If M is a simply-connected Riemannian manifold of bounded geometry, and $X = V(\mathcal{G})$, where \mathcal{G} is a graph approximating M as in Section 8.3, then for all $\delta > 0$,*

$$IP_M(\ell) \approx \text{Ar}_{\delta, X}(\ell).$$

Likewise, using the radius function we define the δ -filling radius function as

$$r = r_{\delta, X} : \mathbb{R}_+ \rightarrow \mathbb{R}_+, \quad r_\delta(\ell) = \sup\{r_\delta(\mathbf{c}) : \mathbf{c} \in \Omega_\delta(X), \ell(\mathbf{c}) \leq \ell\}.$$

Again, we regard r_δ as a *coarse filling radius function*.

EXAMPLE 9.76. In order to get a better feel for the δ -filling area function, let us estimate Ar_δ (from below) in the case $X = \mathbb{R}^2$.

Suppose that \mathbf{c} is a δ -loop in \mathbb{R}^2 with a δ -filling disk $\mathfrak{d} : V(\mathcal{D}) \rightarrow \mathbb{R}^2$. We let

$$g : \mathbb{D}^2 \rightarrow \mathbb{R}^2$$

denote the unique extension of \mathfrak{d} , which is affine on every simplex in \mathcal{D} . Then, in view of Heron's formula for the triangle area, for every 2-simplex τ in \mathcal{D} we obtain the inequality:

$$\text{Area}(g(\tau)) \leq \frac{\sqrt{3}}{4} \delta^2$$

Therefore, summing up over all 2-simplices τ in \mathcal{D} , we obtain an upper bound on the area of the polygon in \mathbb{R}^2 , which is the image of g :

$$(9.4) \quad \text{Area}(g(\mathbb{D}^2)) \leq \sum_{\tau} \delta^2 \frac{\sqrt{3}}{4} \leq \delta^2 \frac{\sqrt{3}}{4} \text{Ar}_{\delta}(\mathfrak{c}).$$

9.7.2. Quasi-isometric invariance of coarse filling functions. We first consider the behavior of the isoperimetric functions and the filling radius under the change of the parameter δ :

LEMMA 9.77. *Suppose that X is a ρ -coarsely connected metric space, i.e. $\text{Rips}_{\rho}(X)$ is connected. Assume also that Ar_{δ_1} takes only finite values. Then there exists $K = K(\delta_1, \delta_2, k, a)$ such that for all $\delta_2 \geq \delta_1 \geq \rho$,*

$$\text{Ar}_{\delta_1}(\ell) \leq \text{Ar}_{\delta_2}(\ell) \leq \text{Ar}_{\delta_1}(K) \text{Ar}_{\delta_2}(\ell)$$

and

$$r_{\delta_1}(\ell) \leq r_{\delta_2}(\ell) \leq r_{\delta_2}(\delta_1) r_{\delta_1}(\ell).$$

In particular, if both Ar_{δ_1} , Ar_{δ_2} are real-valued functions, then

$$\text{Ar}_{\delta_1} \asymp \text{Ar}_{\delta_2} \quad \text{and} \quad r_{\delta_1} \asymp r_{\delta_2}.$$

PROOF. We will consider only the isoperimetric function as the proof for the filling radius function is nearly the same. The inequality

$$\text{Ar}_{\delta_1} \leq \text{Ar}_{\delta_2}$$

is immediate from the definition, as each δ_1 -filling disk for a δ_1 -loop \mathfrak{c} is also a δ_2 -filling disk for the same loop. Consider a coarse loop $\mathfrak{c} \in \Omega_{\delta_1}(X)$. Since $\delta_1 \leq \delta_2$, we can treat \mathfrak{c} as a δ_2 -loop in X .

Let \mathfrak{d} be a δ_2 -filling disk of \mathfrak{c} . Our goal is to replace the triangulation \mathcal{D} (this is a triangulation of the 2-disk associated with the map \mathfrak{d}) with its subdivision $\tilde{\mathcal{D}}$ and extend the map \mathfrak{d} to a δ_1 -filling disk

$$\tilde{\mathfrak{d}} : V(\tilde{\mathcal{D}}) \rightarrow X.$$

Let τ be one of the 2-simplices of \mathcal{D} and

$$\sigma : V(\tau) = \{u_1, u_2, u_3\} \rightarrow X, \sigma(u_j) = x_j, \quad j = 1, 2, 3,$$

be the corresponding brick of \mathfrak{d} . Since $\delta_1 \geq \rho$, we obtain vertex-paths

$$q_j = q_{x_j x_{j+1}} : \{0, \dots, N_j + 1\} \rightarrow X,$$

connecting x_j to x_{j+1} ($j \in \{1, 2, 3\}$), defined in the beginning of this section. (Here and in what follows, the we compute $j + 1$ modulo 3.) Thus,

$$N_j \leq \lambda = \lceil k\delta_2 + a \rceil,$$

and for all $i = 0, \dots, N_j$,

$$\text{dist}_X(q_j(i), q_j(i + 1)) \leq \delta_1,$$

$j = 1, 2, 3$. We then subdivide each edge of the simplex τ into at most λ new edges and using the maps q_1, q_2, q_3 we define a map

$$\mathbf{c}_\tau : V(\mathcal{T}_\tau) \rightarrow X,$$

where \mathcal{T}_τ is the resulting triangulation of the boundary of τ .




FIGURE 9.5

The new map \mathbf{c}_τ agrees with \mathbf{c} on the vertices u_1, u_2, u_3 and sends vertices on the edge $[u_j, u_{j+1}]$ to the points $q_j(i)$, $i = 0, \dots, N_j$, $j = 1, 2, 3$. By the construction, the map \mathbf{c}_τ is a δ_1 -loop in X . Furthermore,

$$\ell_{\delta_1}(\mathbf{c}_\tau) \leq \delta_1(N_1 + N_2 + N_3) \leq K := 3\delta_1(k\delta_2 + a + 1)$$

Therefore, we can fill in this coarse loop with a δ_1 -disk of δ_1 -area at most

$$Ar_{\delta_1}(\mathbf{c}_\tau) \leq Ar_{\delta_1}(K).$$

By repeating this filling for each brick in \mathcal{D} , we construct a δ_1 -filling disk $\tilde{\mathfrak{d}}$ for the δ_1 -loop \mathbf{c} , such that

$$Area(\tilde{\mathfrak{d}}) \leq Ar_{\delta_1}(K)Ar_{\delta_2}(\ell_{\delta_1}(\mathbf{c})).$$

Therefore,

$$Ar_{\delta_1}(\ell) \leq Ar_{\delta_1}(K)Ar_{\delta_2}(\ell),$$

for every ℓ . □

We can now prove quasiisometric invariance of the filling area and filling radius functions:

THEOREM 9.78. *Suppose that X_1, X_2 are quasiisometric ρ -path connected metric spaces with real-valued coarse isoperimetric functions. Then their coarse isoperimetric functions and, respectively their filling radii, functions, are approximately equivalent in the sense of Definition 1.3.*

PROOF. We again consider only the coarse isoperimetric function and leave the case of the filling radius function as an exercise to the reader. Our proof is parallel to the one of Corollary 9.36. Let $f : X_1 \rightarrow X_2$ be an (L, A) -quasiisometry with coarse inverse $\bar{f} : X_2 \rightarrow X_1$. Consider a δ_1 -loop $\mathbf{c}_1 \in \Omega_{\delta_1}(X_1)$ of the length ℓ .

The composition $\mathbf{c}_2 = f \circ \mathbf{c}_1$ is a δ_2 -loop in X_2 , where

$$\delta_2 = L\delta_1 + A.$$

Since $\delta_2 \geq \delta_1$ and $Ar_{\delta_2}(L\delta_1 + A) < \infty$, the coarse loop \mathbf{c}_2 admits a δ_2 -filling disk

$$\mathfrak{d}_2 : V(\mathcal{D}_2) \rightarrow X_2,$$

where \mathcal{D}_2 is a triangulation of \mathbb{D}^2 .

Next, apply the coarse inverse map \bar{f} to the coarse disk \mathfrak{d}_2 : The composition

$$\mathfrak{d}_3 := \bar{f} \circ \mathfrak{d}_2$$

is a δ_3 -filling disk for the coarse loop

$$\mathbf{c}_3 = \bar{f} \circ \mathbf{c}_2,$$

where

$$\delta_3 = L\delta_2 + A.$$

The δ_3 -length of \mathbf{c}_3 is the same as the one of \mathbf{c}_1 and

$$Ar_{\delta_3}(\mathbf{c}_3) = Ar_{\delta_1}(\mathbf{c}_1),$$

since we did not change the triangulation of the unit circle and the unit disk. Of course, the coarse loop \mathbf{c}_3 is not the same as \mathbf{c}_1 , but they are within distance $\leq A$ from each other:

$$\text{dist}_{X_1}(\mathbf{c}_1(v), \mathbf{c}_3(v)) \leq A,$$

for every vertex v of the triangulation of the circle \mathbb{S}^1 . Observe now that A does not exceed δ_3 . Therefore, we can add to the coarse disk \mathfrak{d}_3 a ‘‘coarse annulus’’ $\mathfrak{a} : \mathcal{A} \rightarrow X$, as in the Figure 9.6. This requires adding ℓ vertices, 3ℓ edges and 2ℓ faces to the original triangulation \mathcal{D} of the disk \mathbb{D}^2 . We let $\tilde{\mathcal{D}}$ denote the new triangulation of the 2-disk. We let \mathfrak{d}_4 denote the extension of the map \mathfrak{d}_3 via the map \mathbf{c}_1 of the boundary vertices.

The result is a δ_3 -coarse disk

$$\mathfrak{d}_4 : V(\tilde{\mathcal{D}}) \rightarrow X$$

extending the map \mathbf{c}_1 ; the combinatorial area of this disk is

$$2\ell + \text{Area}(\mathfrak{d}_2) \leq 2\ell + Ar_{\delta_2, X_2}(\ell)$$

This proves that

$$Ar_{\delta_3, X_1}(\ell) \leq 2\ell + Ar_{\delta_2, X_2}(\ell).$$

Taking into account Lemma 9.77, we conclude that

$$Ar_{\delta, X_1} \lesssim Ar_{\delta, X_2}$$

for any $\delta \geq \mu$. Therefore, the spaces X_1, X_2 have approximately equivalent coarse isoperimetric functions:

$$Ar_{X_1} \approx Ar_{X_2}. \quad \square$$

An immediate corollary of this theorem is that the approximate growth rates of the filling area and filling radius, are quasiisometry invariants of the metric space X . The order of the filling function of a metric space X is also called *the filling order of X* . If the coarse isoperimetric function $Ar(\ell)$ of a metric space X satisfies $Ar(\ell) \prec \ell$ or ℓ^2 or e^ℓ , it is said that the space X *satisfies a linear, quadratic or exponential isoperimetric inequality*.




FIGURE 9.6. A coarse annulus

Filling area/radius in the Rips complex. Suppose that X is a μ -simply connected metric space and $\delta \geq \mu$. Instead of filling coarse loops in X by δ -disks, one can fill in polygonal loops in $P = \text{Rips}_\delta(X)$ by simplicial maps of triangulated disks. Let \mathfrak{c} be a δ -loop in X . Then we have a triangulation of the circle \mathbb{S}^1 such that $\text{diam}(\mathfrak{c}(\partial e)) \leq \delta$ for every edge e of the triangulation. Thus, we define an edge-loop $\mathfrak{c}_\delta = \tilde{\mathfrak{c}}$ in P by connecting points $\mathfrak{c}(\partial e)$ by the edges in P (provided that these points are distinct, of course). We will think of \mathfrak{c}_δ as a simplicial map $\mathbb{S}^1 \rightarrow P$ (this map may send some edges to vertices). Then

$$\text{length}^{\text{com}}(\mathfrak{c}_\delta) = \ell(\mathfrak{c}).$$

It is clear that for $\delta > 0$ the map

$$\begin{aligned} \{\delta\text{-loops in } X \text{ of length } \leq \ell\} &\rightarrow \{\text{edge-loops in } P \text{ of length } \leq \ell\} \\ \mathfrak{c} &\mapsto \mathfrak{c}_\delta \end{aligned}$$

is surjective. Furthermore, every δ -disk \mathcal{D} which fills in \mathfrak{c} , yields a simplicial map $\mathcal{D}_\delta : \mathbb{D}^2 \rightarrow P$ which is an extension of \mathfrak{c}_δ . The combinatorial area is preserved under this construction:

$$\text{Area}(\mathcal{D}_\delta) = \text{Area}(\mathcal{D}).$$

We leave it to the reader to verify that the above procedure yields *all* simplicial maps $\mathbb{D}^2 \rightarrow P$ extending \mathfrak{c}_δ and we obtain

$$\text{Area}(\mathfrak{c}_\delta) = \text{Ar}_\delta(\mathfrak{c}).$$

Summarizing all this, we obtain

$$A_{\text{Rips}_\delta(X)}(\ell) = \text{Ar}_\delta(\ell),$$

where the left hand side is defined analogously to the function Ar , only using simplicial maps to the Rips complex instead of δ -maps to X itself. The same argument applies to the filling radius and we obtain:

OBSERVATION 9.79. Studying the coarse filling area and filling radius functions in X (up to the equivalence relation \approx) is equivalent to studying the simplicial filling area and filling radius functions in $Rips_\delta(X)$.

Lastly, we relate the filling area function to the Dehn function:

THEOREM 9.80. *For every finitely presented group G , the coarse isoperimetric function and the Dehn function are also approximately equivalent.*

PROOF. Let G be a finitely presented group with the finite presentation $\langle S|R \rangle$, and equipped with the word metric dist_S . We let μ denote the length of the longest relator in R and let $Dehn_G$ denote the Dehn function associated with the presentation $\langle S|R \rangle$.

EXERCISE 9.81. For every finitely presented group G , the metric space (G, dist_S) is μ -simply connected.

We will prove the approximate inequality

$$Ar_{\mu,G} \lesssim Dehn_G$$

and leave the opposite inequality as an exercise to the reader.

We let Y denote the presentation complex of $\langle S|R \rangle$ and let \tilde{Y} denote its universal cover: The vertex set of the complex \tilde{Y} is the group G , the 1-skeleton of \tilde{Y} is the Cayley graph $\text{Cayley}(G, S)$ of G (with respect to the generating set S).

Let \mathfrak{c} be a coarse loop in G , an element of $\Omega_1(G)$; this coarse loop defines an (almost) regular cellular map $c = \tilde{\mathfrak{c}} : \mathbb{S}^1 \rightarrow \text{Cayley}(G, S)$, from the triangulated unit circle. Our goal is to estimate above the coarse filling area of \mathfrak{c} via van Kampen diagrams extending the map c .

The loop c projects to a map $c_w : \mathbb{S}^1 \rightarrow Y^{(1)}$ corresponding to some, possibly nonreduced, word w in S . (See Section 7.10.4.) Replacing w with its free reduction w' will change very little:

$$A(w') = A(w)$$

and the lift c' of the loop $c_{w'}$ will satisfy

$$\text{Ar}_\mu(\mathfrak{c}) \leq \text{Ar}_\mu(\mathfrak{c}') + \ell(\mathfrak{c}).$$

Here \mathfrak{c}' is the restriction of c' to the vertex set of the triangulation of \mathbb{S}^1 . Therefore, in what follows, we will assume that w is reduced.

Every van Kampen diagram $h : K \rightarrow Y$ of the word w lifts to a map

$$f : K \rightarrow \tilde{Y},$$

whose boundary value $\partial f : \mathbb{S}^1 \rightarrow \tilde{Y}$ is a lift of c_w . The van Kampen diagram f extends to an almost regular map $g = \hat{f} : \tilde{K} \rightarrow \tilde{Y}$ as in Section 7.10.4 and

$$\text{Area}^{com}(g) = \text{Area}(h).$$

By the construction, \hat{f} sends cells of \tilde{K} to cells of $\text{Cayley}(G, S)$; the only problem is that \tilde{K} is not a simplicial complex. However, the second barycentric subdivision of \tilde{K} is a triangulation \mathcal{D} of the disk \mathbb{D}^2 . The total number of faces of \mathcal{D} is at most $12\mu \text{Area}(h)$. In order to extend g to the vertices of \mathcal{D} , for each vertex $v \in V(\mathcal{D})$,

we let σ_v be the smallest cell of \widehat{K} containing v . Lastly, let $g(v)$ be an arbitrarily chosen vertex in $g(\sigma_v)$. Thus, we define the new map $\mathfrak{d} : V(\mathcal{D}) \rightarrow G = V(\tilde{Y})$, equal to the restriction of g to $V(\mathcal{D})$. The map \mathfrak{d} is a μ -disk in G extending the coarse loop \mathfrak{c} . We obtain

$$\text{Ar}_\mu(\mathfrak{c}) \leq 12\mu \text{Area}(h) \leq 12\mu \text{Dehn}_G(\ell) \leq 12\mu \text{Dehn}_G(\ell(\mathfrak{c})).$$

where ℓ is the word-length of w . The approximate inequality

$$\text{Ar}_\mu \lesssim \text{Dehn}_G$$

follows. □

The filling radius function is not as commonly used in Geometric Group Theory as the coarse isoperimetric function and the Dehn function. As we noted earlier, Gersten proved that solvability of the word problem for G is equivalent to the recursivity of its Dehn function. In the same paper [Ger93a] Gersten also proved:

PROPOSITION 9.82. *For a finitely presented group G the following are equivalent.*

1. G has solvable word problem.
2. The filling radius function of G is recursive.

9.7.3. Higher Dehn functions. The Dehn functions $\text{Dehn}(\ell)$ generalize to “higher Dehn functions” (“higher isoperimetric functions”) Dehn_n for groups G of type \mathbf{F}_n , with $\text{Dehn} = \text{Dehn}_1$. All the definitions amount to a coarsification of the Riemannian isoperimetric function $IP_{M,n}$ responsible for the least volume extension of maps of n -spheres to maps of $n + 1$ -balls, see Section 3.5.

Below we list four higher isoperimetric functions. They are all defined in a similar fashion as a max-min of a certain geometric quantity; the differences come from different notions of “volume” used for these functions. We purposefully restrict ourselves to the discussion of functions which measure complexity of extending maps of spheres to maps of balls, there are other isoperimetric functions dealing with “filling in” cycles and currents, see [ABD⁺13, Wen05].

9.7.3.A. *Combinatorial, simplicial and cellular isoperimetric functions.* Recall that in Definition 7.87 we introduced the notions of *simplicial* and *combinatorial* volumes of maps between simplicial complexes.

The *simplicial* (resp. *combinatorial*) *filling volume* of a simplicial map $f : Z \rightarrow X$ of a triangulated n -dimensional sphere Z into X is defined as

$$\text{FillVol}^{\text{sim}}(f) := \inf_{\hat{f}:W \rightarrow X} \text{Vol}_{n+1}^{\text{sim}}(\hat{f}),$$

and

$$\text{FillVol}^{\text{com}}(f) := \inf_{\hat{f}:W \rightarrow X} \text{Vol}_{n+1}^{\text{com}}(\hat{f}),$$

where in both definitions the infimum is taken over all extensions \hat{f} of f , such that $f : W \rightarrow X$ is a simplicial map of a triangulated $n + 1$ -dimensional ball W whose boundary is Z .

DEFINITION 9.83. The *simplicial* (resp. *combinatorial*) filling functions of a simplicial complex X are

$$\begin{aligned} IP_{X,n}^{\text{sim}} : A &\mapsto \sup\{\text{FillVol}^{\text{sim}}(f) \mid \text{Vol}_n^{\text{sim}}(f) \leq A\}, \\ IP_{X,n}^{\text{com}} : A &\mapsto \sup\{\text{FillVol}^{\text{com}}(f) \mid \text{Vol}_n^{\text{com}}(f) \leq A\}, \end{aligned}$$

where the supremum in both definitions is taken over all simplicial maps $f : Z \rightarrow X$ of triangulated n -dimensional spheres.

One also defines isoperimetric functions (higher Dehn functions) $\delta_{X,n}$ for bounded geometry almost regular cell complexes X using maps of spheres (and balls) equipped with structure of CW complexes and then defining the cellular n -volume by counting numbers of n -cells in the domain which are mapped homeomorphically onto n -cells in the target, see Definition 7.89.

With this definition, we define the *cellular filling volume* $FillVol^{cell}(f)$ of almost regular maps $f : Z \rightarrow X$ of the n -sphere (equipped with the structure of an almost regular cell complex Z) and the *cellular n th order isoperimetric function*

$$\delta_{X,n} = IP_{X,n}^{cell}$$

of X by repeating verbatim the simplicial definition.

EXERCISE 9.84. Let X be an almost regular bounded geometry cell complex and let X'' be the simplicial complex equal to the second barycentric subdivision of X . Then

$$IP_{X,n}^{cell} \approx IP_{X'',n}^{sim}$$

for all $n \geq 1$.

THEOREM 9.85. (See [AWP99, Ril03].) Suppose that X, Y are quasisisometric uniformly n -connected bounded geometry simplicial complexes with finite isoperimetric functions $\delta_{X,k}, \delta_{Y,k}$, $0 \leq k \leq n$. Then

$$\delta_{X,n} \approx \delta_{Y,n}.$$

DEFINITION 9.86. Let G be a group of type \mathbf{F}_n , $1 \leq n < \infty$. The n -th order Dehn function, $Dehn_{G,n}$ is defined as the function $\delta_{X,n}$ for some cell n -connected complex X on which G acts properly discontinuously and cocompactly.

In view of Theorem 9.85, the asymptotic equivalence class of $Dehn_{G,n}$ is independent of the choice of X .

The following result was proven by P. Papasoglou:

THEOREM 9.87 (P. Papasoglou, [Pap00]). The second Dehn function $Dehn_2$ of a group of type \mathbf{F}_3 is bounded above by a recursive function.

This theorem represents a striking contrast with the fact that there are finitely presented groups with unsolvable word problem and, hence, Dehn function which is not bounded above by any recursive function.

Here is the idea of the proof of Theorem 9.87. Let Y be a finite cell complex with $\pi_1(Y) \cong G$ and $\pi_2(Y) = 0$. Consider cellular maps $s : \mathbb{S}^2 \rightarrow Y$. Every such map s is null-homotopic and for every ℓ , there are only finitely many such maps with $Vol_2^{sim}(s) \leq \ell$. The key then is to design an algorithm which, for each s , finds some extension $f : \mathbb{D}^3 \rightarrow Y$: This algorithm uses the algorithmic recognition of 3-dimensional balls (and, hence, fails for $Dehn_n$, $n \geq 3$). One then computes the simplicial volume $Vol_3^{sim}(f)$ and takes the maximum $\Delta(\ell)$, over all maps s . The resulting function $\Delta(\ell)$ gives the required recursive upper bound:

$$Dehn_2 \lesssim \Delta(\ell).$$

We observe that one gets only an upper bound on $Dehn_2$ since the filling maps f above might not be optimal. Observe also that this proof also fails for the ordinary

Dehn function since the presentation complex Y is (usually) not simply connected and the recognizing which loops in Y are null-homotopic is algorithmically impossible.

9.7.3.B. *Coarse isoperimetric functions.* The *coarse higher isoperimetric functions* generalize the coarse 1-dimensional function defined earlier; our discussion is inspired by [Gro93, Chapter 5]. The definitions that we are about to give are modeled on the definition of the coarse 1-dimensional isoperimetric function given in Section 9.7.1.

Let (X, dist) be a metric space and $\mu > 0$ is a positive number, the measure of “coarseness” of maps into X .

DEFINITION 9.88. Let Z be a simplicial complex. A μ -coarse map $Z \rightarrow X$ is a map $f : Z^0 \rightarrow X$ defined on the vertex set of Z such that for every edge $[u, v]$ in Z , $\text{dist}(f(u), f(v)) \leq \mu$.

If Z is a subcomplex in W then the *restriction* of a μ -coarse map $f : W \rightarrow X$ to Z is the restriction g of $f : W^0 \rightarrow X$ to the vertex set of Z , such that g is μ -coarse. The map f is then said to be the μ -coarse *extension* of the map g .

DEFINITION 9.89. The n -dimensional volume $\text{Vol}_{\mu, n}^{\text{coarse}}(f)$ of a μ -coarse map $f : Z \rightarrow X$ is the number of n -dimensional simplices in Z .

DEFINITION 9.90. Let X be a metric space and $\mu > 0$. Suppose that Z is a triangulated n -dimensional sphere and $f : Z \rightarrow X$ is a μ -coarse map. The *coarse filling volume* $\text{FillVol}_{\mu}^{\text{coarse}}(f)$ of f is

$$\min_{\hat{f}} \text{Vol}_{\mu, n+1}^{\text{coarse}}(\hat{f})$$

where the minimum is taken over all μ -coarse extensions \hat{f} of f , where $\hat{f} : W \rightarrow X$ is a μ -coarse map of a triangulated $n + 1$ -dimensional ball whose boundary is Z . (The triangulation W of \mathbb{B}^{n+1} is not fixed, of course, but is required to coincide with the given triangulation Z of \mathbb{S}^n .)

Lastly:

DEFINITION 9.91.

$$IP_{\mu, X, n}^{\text{coarse}} : A \mapsto \sup\{\text{FillVol}_{\mu}^{\text{coarse}}(f) : \text{Vol}_{\mu, n}^{\text{coarse}}(f) \leq A\}$$

is the μ -coarse n -dimensional isoperimetric function of X , where the supremum is taken over all μ -coarse maps $f : Z \rightarrow X$ of triangulated n -dimensional spheres.

The reader will notice that for $n = 1$, we obtain

$$\text{Ar}_X(\ell) = IP_{\mu, X, 1}^{\text{coarse}}(\ell).$$

EXERCISE 9.92. For every metric space X and $\mu > 0$ we have

$$IP_{\mu, X, n}^{\text{coarse}} = IP_{\text{Rips}_{\mu}(X), n}^{\text{com}}.$$

Note that both IP^{coarse} and $\text{FillVol}^{\text{coarse}}$ are allowed to take infinite values. We say that a function $\mathbb{R}_+ \rightarrow [0, \infty]$ is *finite* if it takes values in $[0, \infty)$.

EXERCISE 9.93. Suppose that $IP_{\rho, X, k}^{\text{coarse}}$ is a finite function for $0 \leq k \leq n$. Then for all $\mu \geq \rho$ the functions

$$IP_{\mu, X, k}^{\text{coarse}}, 0 \leq k \leq n$$

are also finite.

The proofs of Lemma 9.77 and Theorem 9.78 go through almost verbatim for the higher coarse isoperimetric functions and we obtain:

THEOREM 9.94. 1. Suppose that $\rho > 0$ is such that the functions $IP_{\rho, X, k}^{coarse}$ are finite for $0 \leq k \leq n - 1$. Then for all $\mu \geq \rho, \nu \geq \rho$ we have

$$IP_{\mu, X, n}^{coarse} \approx IP_{\nu, X, n}^{coarse}.$$

2. Suppose that X, Y are quasiisometric metric spaces such that (for some $\rho > 0$) $IP_{\rho, X, k}^{coarse}, IP_{\rho, Y, k}^{coarse}$, $0 \leq k \leq n - 1$, are finite functions. Then for all sufficiently large μ , we have

$$IP_{\mu, X, n}^{coarse} \approx IP_{\mu, Y, n}^{coarse}.$$

EXERCISE 9.95. Prove Theorem 9.85 using the arguments of the proof of Theorem 9.78.

9.7.3.C. *Relation between different isoperimetric functions.* As we noted above, we have equivalence of the higher order isoperimetric functions

$$\delta_{X, n} \approx IP_{X, n}^{cell} \approx IP_{X, n}^{sim}.$$

In the case $n = 1$, Theorem 9.80 also yields

$$IP_{G, 1}^{coarse} \approx Dehn_G$$

for finitely presented groups G .

EXERCISE 9.96. Extend the proof of Theorem 9.80 to complexes X of bounded geometry and prove

$$IP_{X, 1}^{com} \approx IP_{X, 1}^{coarse} \approx IP_{X, 1}^{sim}.$$

However, it appears to be unknown if the last approximate equivalence also holds for $n \geq 2$ even for universal covers of finite simplicial complexes.

We next turn to the metric isoperimetric functions introduced in (3.3) and (3.5):

THEOREM 9.97. Suppose that (M, g) is a Riemannian manifold of bounded geometry, $\tau : X \rightarrow M$ is a bounded geometry triangulation of M . Then

$$IP_{M, n}^{met} \approx IP_{X, n}^{com}$$

for all $n \geq 1$.

This theorem is proven by Groft in [Gro09] in the case when (M, g) covers a compact Riemannian manifold (M', g') and the triangulation is lifted from a triangulation of (M', g') . However, examining Groft's proof one sees that the existence (M', g') is used only to ensure the existence of a bounded geometry triangulation of (M, g) .

THEOREM 9.98. If X is a simplicial complex of bounded geometry then

$$IP_{X, n}^{met} \approx IP_{X, n}^{com}$$

for all $n \geq 1$.

This theorem (as Theorem 9.97) is essentially contained in Groft's paper [Gro09], who proved this result for covering spaces of compact Lipschitz neighborhood retracts, but his proof, as in the manifold case, uses the covering space assumption only to guarantee uniform control on Lipschitz constants of local retractions.

Lastly, we relate metric isoperimetric functions of Riemannian manifolds of bounded geometry and their Lipschitz simplicial models.

COROLLARY 9.99. *Let (M, g) be a Riemannian manifold of bounded geometry and take its Lipschitz simplicial model X as in Theorem 3.37. Then*

$$IP_{M,n}^{met} \approx IP_{X,n}^{met}, \quad n \geq 1.$$

PROOF. The proof is similar to that of Theorem 9.78. We will use the notation from Theorem 3.37. Let $\sigma : \mathbb{S}^n \rightarrow M$ be a Lipschitz map of the (n -dimensional metric) volume A . Then the volume of $f \circ \sigma : \mathbb{S}^n \rightarrow X$ is at most $L^n A$. Hence, there exists a Lipschitz map $\varphi : \mathbb{B}^{n+1} \rightarrow X$ of volume $\leq V = IP_{X,n}(L^n A)$. The volume of $\psi = \tilde{f} \circ \varphi$ is at most $L^{n+1} V$. Consider the map

$$\eta : \mathbb{S}^n \times [0, 1] \rightarrow M, \eta(x, t) = H(\psi(x), t).$$

This map interpolates between σ and ψ , and its volume is at most $L^{n+1} A$. Now, we attach the spherical shell $\mathbb{S}^n \times [0, 1]$ to the unit ball so that $\mathbb{S}^n \times 1$ is attached to the boundary sphere $\mathbb{S}^n = \partial \mathbb{B}^{n+1}$. The result is again a ball \mathbb{D}^{n+1} and we define the Lipschitz map

$$\hat{\sigma} : \mathbb{D}^{n+1} \rightarrow M$$

to be equal to η on the spherical shell and ψ on \mathbb{B}^n , so that $\hat{\sigma}$ restricts to the map σ on $\partial \mathbb{D}^{n+1}$. The volume of $\hat{\sigma}$ is at most $L^{n+1}(A + V)$. It follows that

$$IP_{M,n}^{met} \lesssim IP_{X,n}^{met}.$$

The opposite inequality is proven in a similar fashion using the homotopy \tilde{H} instead of H . \square

COROLLARY 9.100. *If (M, g) is a Riemannian manifold of bounded geometry and (X, f) is its Lipschitz simplicial model, then*

$$IP_{M,n}^{met} \approx IP_{X,n}^{sim}, \quad n \geq 1.$$

9.7.4. Coarse Besikovitch inequality. In this section we will prove coarse analogues of the classical *Besikovitch inequality* (see e.g. [BZ88]).

Let $Q \subset \mathbb{R}^2$ denote the unit square; then the topological circle $C = \partial Q$ has the natural structure of a simplicial complex with the consecutive edges e_1, \dots, e_4 . Subdividing the edges of Q further, we obtain a triangulated topological circle $(\mathbb{S}^1, \mathcal{T})$, which will be used in the proposition below. We will regard the sides e_i ($i = 1, 2, 3, 4$) of Q as subcomplexes of \mathcal{T} . A *topological quadrilateral* in a topological space X is a continuous map $f : C \rightarrow X$. Similarly, one defines a *topological triangle* in X as a continuous map form $f : T \rightarrow X$, where $T \subset \mathbb{R}^2$ is a triangle with the edges e_1, e_2, e_3 (recall that triangles are always treated as 1-dimensional objects). Again, we will regard the edges of T as subcomplexes of a fixed triangulation \mathcal{T} of T , refining the original simplicial structure.

Given a topological quadrilateral $f : \partial Q \rightarrow X$ in a metric space X , we define its *separation* $\text{sep}(f)$ as the pair (d_1, d_2) , where

$$d_i = \text{dist}(f(e_i), f(e_{i+2})), \quad i = 1, 2,$$

where dist is the minimal distance between subsets of X (see Section 2.1). For instance, suppose that $f : Q \rightarrow \mathbb{R}^2$ is an affine map, whose image is the parallelogram P , with the side-lengths s_1 (the length of $f(e_1)$) and s_2 (the length of $f(e_2)$) and the angle α between $f(e_1), f(e_2)$. Then the separation of $f|_{\partial Q}$ equals

$$(d_1, d_2) = (s_2 \sin(\alpha), s_1 \sin(\alpha)).$$

It is then immediate that

$$\text{Area}(P) \leq d_1 d_2.$$

Besikovitch proved that the same inequality holds for topological quadrilaterals in arbitrary metric spaces, where Area of a topological quadrilateral is understood as the least area of 2-disks that its bounds in X .

The notion of *minsize* for topological triangles defined below is an analogue of separation for topological quadrilaterals.

DEFINITION 9.101. The *minimal size* (*minsize*) of a topological triangle $f : T \rightarrow X$ is defined as

$$\text{minsize}(f) = \inf\{\text{diam}\{f(y_1), f(y_2), f(y_3)\} ; y_i \in e_i, i = 1, 2, 3\}.$$

Next, we coarsify the notions of topological triangles and quadrilaterals, their minsize and separation. In what follows, we fix X , a μ -coarsely simply connected metric space, $\delta \geq \mu$ and (C, \mathcal{T}) , a triangulated topological circle (a subdivided quadrilateral or a triangle). A δ -loop $\mathbf{c} : V(\mathcal{T}) \rightarrow X$ will be regarded as a *coarse quadrilateral*, resp. a *coarse triangle* in X .

DEFINITION 9.102. The *separation* of a coarse quadrilateral \mathbf{c} is defined as the pair (d_1, d_2) , where

$$d_i = \text{dist}(\mathbf{c}(V(e_i)), \mathbf{c}(V(e_{i+2}))), \quad i = 1, 2.$$

The *minsize* of a coarse triangle \mathbf{c} is defined as

$$\text{minsize}(\mathbf{c}) = \min\{\text{diam}\{\mathbf{c}(y_1), \mathbf{c}(y_2), \mathbf{c}(y_3)\} ; y_i \in V(e_i), i = 1, 2, 3\}.$$

PROPOSITION 9.103 (The coarse Besikovitch inequality). *With the notation as above, for each coarse quadrilateral $\mathbf{c} \in \Omega_\delta(X)$ we have*

$$\text{Ar}_\delta(\mathbf{c}) \geq \frac{2}{\sqrt{3}\delta^2} d_1 d_2.$$

PROOF. Our proof follows closely the proof of the classical Besikovitch inequality. Consider the plane \mathbb{R}^2 , whose points will be denoted (s, t) . Define the map $\beta : X \rightarrow \mathbb{R}^2$,

$$\beta(x) = (\text{dist}(x, \mathbf{c}(V(e_1))), \text{dist}(x, \mathbf{c}(V(e_2)))).$$

Since each component of β is a 1-Lipschitz map, the map β itself is $\sqrt{2}$ -Lipschitz. Define the composition

$$\beta \circ \mathbf{c} : V(\mathcal{T}) \rightarrow \mathbb{R}^2$$

and its geodesic extension f . Then the image $f(e_1) \subset \mathbb{R}^2$ is a vertical segment connecting the origin to a point $(0, t_1)$, with $t_1 \geq d_2$, while $f(e_2)$ is a horizontal segment connecting the origin to a point $(s_2, 0)$, with $s_2 \geq d_1$.

Similarly, the image $f(e_3)$ is a path to the right of the vertical line $\{s = d_1\}$ and $f(e_4)$ is another path above the horizontal line $t = d_2$. Thus, the rectangle $R \subset \mathbb{R}^2$ with the vertices

$$(0, 0), (d_1, 0), (d_1, d_2), (0, d_2),$$

is separated from the infinity by the curve $\beta \circ \tilde{\mathfrak{c}}(\mathbb{S}^1)$ (see Figure 9.7).

In particular, the image of any continuous extension g of the map f to the entire square Q , contains the rectangle R . Thus,

$$\text{Area}(g(Q)) \geq \text{Area}(R) = d_1 d_2.$$

By taking into the account the fact that the map β is $\sqrt{2}$ -Lipschitz, and the inequality (9.4), we obtain

$$d_1 d_2 \leq \epsilon^2 \frac{\sqrt{3}}{4} \text{Ar}_\epsilon(\beta \circ \mathfrak{c}),$$

where $\epsilon = \sqrt{2}\delta$.

Consider a δ -filling disk \mathfrak{d} of the δ -loop \mathfrak{c} and let g be the extension of the map $\beta \circ \mathfrak{d}$, defined as in Example 9.76. The simplicial area of $\beta \circ \mathfrak{d}$ is, of course, exactly the same as the one of the map \mathfrak{d} . Putting this all together:

$$\text{sArea}(\mathfrak{d}) \geq \text{Ar}_\delta(\mathfrak{c}) \geq \frac{2}{\sqrt{3}\delta^2} d_1 d_2,$$

as required. □




FIGURE 9.7. The map β .

Besikovitch's inequality generalizes from maps of squares to maps of triangles: This generalization has interesting applications to δ -hyperbolic spaces which will be discussed in Section 11.22.1.

PROPOSITION 9.104 (Minsize inequality). *Let X be a μ -simply connected metric space and let $\delta \geq \mu$. Then each coarse topological triangle $\mathfrak{c} : V(\mathcal{T}) \rightarrow X$, $\mathfrak{c} \in \Omega_\delta(X)$, satisfies the minsized inequality*

$$\text{Ar}_\delta(\mathfrak{c}) \geq \frac{1}{2\sqrt{3}\delta^2} [\text{minsized}(\mathfrak{c})]^2.$$

PROOF. The proof is analogous to the one for coarse quadrilaterals. Define the $\sqrt{2}$ -Lipschitz map $\beta : X \rightarrow \mathbb{R}^2$,

$$\beta(x) = (\beta_1, \beta_2) = (\text{dist}(x, \mathfrak{c}(V(e_1))), \text{dist}(x, \mathfrak{c}(V(e_2)))) ,$$

the composition $\beta \circ \mathfrak{c} : V(\mathcal{T}) \rightarrow \mathbb{R}$ and the geodesic extension $f = (f_1, f_2) : T \rightarrow \mathbb{R}^2$ of the latter.

As in the proof of the coarse Besikovitch inequality for quadrilaterals, f sends the edges e_1, e_2 to coordinate segments, while the restriction of f to e_3 satisfies:

$$\max(f_1(x), f_2(x)) \geq \frac{m}{2}, \forall x \in e_3,$$

where $m = \text{minsize}(\mathbf{c})$. Therefore, the image of any continuous extension g of f contains the square with the vertices

$$(0, 0), \left(\frac{m}{2}, 0\right), \left(\frac{m}{2}, \frac{m}{2}\right), \left(0, \frac{m}{2}\right).$$

Arguing as in the case of coarse quadrilaterals, we obtain the estimate

$$\frac{m^2}{4} \leq \text{Area}(g)$$

and, hence,

$$\text{Ar}_\delta(\mathbf{c}) \geq \frac{2}{\sqrt{3}\delta^2} \frac{m^2}{4} = \frac{1}{2\sqrt{3}\delta^2} m^2.$$

□

Ultralimits of Metric Spaces

Let $(X_i)_{i \in I}$ be an indexed family of metric spaces. The goal of this chapter is to describe the asymptotic behavior of the family (X_i) by studying limits of indexed families of finite subsets $Y_i \subset X_i$. Ultrafilters are an efficient technical device for simultaneously taking limits of all such families of subspaces and putting them together to form one object, namely an *ultralimit* of (X_i) . The price to pay for this efficiency is that our discussion will have to rely upon a version of the Axiom of Choice.

10.1. The Axiom of Choice and its weaker versions

We first recall that the *Zermelo–Fraenkel axioms (ZF)* form a list of axioms which are the basis of axiomatic set theory in its standard form, see for instance [Kun80], [HJ99], [Jec03].

The Axiom of Choice (AC) can be seen as a rule of building sets from other sets. It was first formulated by Ernesto Zermelo in [Zer04]. According to work of Kurt Gödel and Paul Cohen, the Axiom of Choice is logically independent of the Zermelo–Fraenkel axioms (i.e. neither it nor its negation can be proven in ZF).

Given a non-empty collection \mathcal{S} of non-empty sets, a *choice function* defined on \mathcal{S} is a function $f : \mathcal{S} \rightarrow \cup_{A \in \mathcal{S}} A$, such that for every set A in \mathcal{S} , $f(A)$ is an element of A . In other words, a choice function on \mathcal{S} is an element of the Cartesian product $\prod_{A \in \mathcal{S}} A$.

Axiom of choice: *On any non-empty collection of non-empty sets there exists a choice function. Equivalently, the Cartesian product of a non-empty family of non-empty sets is non-empty:*

$$\mathcal{S} \neq \emptyset \quad \& \quad \forall A \in \mathcal{S}, A \neq \emptyset \Rightarrow \prod_{A \in \mathcal{S}} A \neq \emptyset.$$

We will use the abbreviation ZFC for the Zermelo–Fraenkel axioms plus the Axiom of Choice.

REMARK 10.1. If $\mathcal{S} = \{A\}$ then the existence of f follows from the fact that A is non-empty. If \mathcal{S} is finite or countable, the existence of a choice function can be proved by induction. Thus, if the collection \mathcal{S} is finite or countable then the existence of a choice function follows from ZF.

REMARK 10.2. Assuming ZF, the Axiom of Choice is equivalent to each of the following statements (see [HJ99] and [RR85] for a much longer list):

- (1) *Zorn’s lemma:* Suppose that S is a partially ordered set where every totally ordered subset has an upper bound. Then S has a maximal element.

- (2) Every vector space has a basis.
- (3) Every ideal in a unital ring is contained in a maximal ideal.
- (4) If A is a subset in a topological space X and B is a subset in a topological space Y , the closure of $A \times B$ in $X \times Y$ is equal to the product of the closure of A in X with the closure of B in Y .
- (5) *Tychonoff's theorem*: If $(X_i)_{i \in I}$ is a collection of non-empty compact topological spaces, then $\prod_{i \in I} X_i$ is compact.

REMARK 10.3. The following statements require the Axiom of Choice (i.e. are unprovable in ZF, but hold in ZFC), see [HJ99, RR85]:

- (1) Every union of countably many countable sets is countable.
- (2) The Nielsen–Schreier theorem: Every subgroup of a free group is free (Theorem 7.41), to ensure the existence of a maximal subtree. (Note that the Axiom of Choice is needed only for free groups of uncountable rank.)

Note that for finitely generated free groups (and we are mostly interested in these) the Nielsen–Schreier theorem does not require the Axiom of Choice.

In ZF, we have the following irreversible sequence of implications:

Axiom of choice \Rightarrow Ultrafilter lemma \Rightarrow Hahn–Banach extension theorem.

The first implication is easy (see Lemma 10.18), it was proved to be irreversible in [Hal64]. Proof of the second implication can be found in [LRN51], [Lux62], [Lux67], [Lux69], while proofs of its irreversibility is can be found in [Pin72] and [Pin74].

Thus, the Hahn–Banach extension theorem (see below) can be seen as the analyst's Axiom of Choice, in a weaker form.

THEOREM 10.4 (Hahn–Banach Theorem, see e.g. [Roy68]). *Let V be a real vector space, U a subspace of V , and $\varphi : U \rightarrow \mathbb{R}$ a linear function. Let $p : V \rightarrow \mathbb{R}$ be a map with the following properties:*

$$p(\lambda x) = \lambda p(x) \text{ and } p(x + y) \leq p(x) + p(y), \forall x, y \in V, \lambda \in [0, +\infty),$$

such that $\varphi(x) \leq p(x)$ for every $x \in U$. Then there exists a linear extension of φ , $\bar{\varphi} : V \rightarrow \mathbb{R}$ such that $\bar{\varphi}(x) \leq p(x)$ for every $x \in V$.

In order to state the Ultrafilter Lemma (which we will use to prove existence of ultrafilters and, hence, ultralimits and asymptotic cones), we first define *filters*. We refer the reader to [Bou65, §I.6.4] for the basic properties of filters and ultrafilters, and to [Kei10] for an in depth survey, including ultraproducts.

DEFINITION 10.5. A *filter* \mathcal{F} on a set I is a collection of subsets of I satisfying the following conditions:

- (F₁) $\emptyset \notin \mathcal{F}$.
- (F₂) If $A, B \in \mathcal{F}$ then $A \cap B \in \mathcal{F}$.
- (F₃) If $A \in \mathcal{F}$, $A \subseteq B \subseteq I$, then $B \in \mathcal{F}$.

EXERCISE 10.6. Given an infinite set I , prove that the collection of all complements of finite sets is a filter on I . This filter is called *the Fréchet filter* (or the *cofinite filter*); it is used to define the *cofinite topology* on a topological space.

DEFINITION 10.7. Subsets $A \subset I$ which belong to a filter \mathcal{F} are called *\mathcal{F} -large*. We say that a property (P) holds for *\mathcal{F} -all i* if (P) is satisfied for all i in some \mathcal{F} -large set.

DEFINITION 10.8. A *base of a filter* on a set I is a subset \mathcal{B} of the power set 2^I of I , which satisfies the properties:

- (B_1) If $B_i \in \mathcal{B}$, $i = 1, 2$, then $B_1 \cap B_2$ contains an element of \mathcal{B} ;
- (B_2) $\emptyset \notin \mathcal{B}$ and \mathcal{B} is not empty.

As an example, consider a point x in a topological space X . We let \mathcal{F}_x denote the system of neighborhoods of x , i.e. all subsets $N \subset X$ which contain x together with some open neighborhood of x . Then \mathcal{F}_x is a filter. A neighborhood basis of x is an example of a base of a filter. This topological intuition is somewhat useful when thinking about filters on \mathbb{N} (more precisely, *non-principal ultrafilters* defined below): Such a filter can be regarded as a system of *punctured neighborhoods* of ∞ in a topology on $\mathbb{N} \cup \{\infty\}$.

EXERCISE 10.9. If \mathcal{B} is a base of a filter, then the set $\langle \mathcal{B} \rangle$ of subsets of I containing some $B \in \mathcal{B}$ is a filter.

We will say that $\langle \mathcal{B} \rangle$ is the filter *generated* by \mathcal{B} . Thus, one can generate filters using bases in the same fashion one generates a topology using its neighborhood bases.

Given a set I , we let $Filter(I) \subset 2^{2^I}$ denote the set of all filters on I . In particular, $Filter(I)$ has a natural partially order given by the inclusion. If $\mathcal{F}_\alpha, \alpha \in A$, is a (non-empty) collection of filters on I , then the union

$$\mathcal{B} = \bigcup_{\alpha \in A} \mathcal{F}_\alpha$$

is not (in general) a filter on I , but it is a base of a filter. Therefore, $\langle \mathcal{B} \rangle$ is a filter on I .

EXERCISE 10.10. Use this construction to show that every totally ordered subset A of $Filter(I)$ has an upper bound in $Filter(I)$.

REMARK 10.11. The set $Filter(I)$ has the same cardinality as 2^{2^I} , see [Pos37].

DEFINITION 10.12. An *ultrafilter* on a set I is a filter \mathcal{U} on I which is a maximal element in the ordered set $Filter(I)$. Equivalently, an ultrafilter can be defined (see [Bou65, §I.6.4]) as a collection of subsets of I satisfying the conditions (F_1), (F_2), (F_3) defining a filter and the additional condition:

$$(F_4) \quad \text{For every } A \subseteq I, \text{ either } A \in \mathcal{U} \text{ or } A^c = I \setminus A \in \mathcal{U}.$$

One direction in this equivalence is clear: If \mathcal{F} satisfies the axioms (F_1)—(F_4), then \mathcal{F} has to be a maximal filter, since every strictly larger filter would have as its members a subset $A \subset I$ as well as A^c , but $A \cap A^c = \emptyset$, contradicting the axioms (F_1) and (F_2).

EXERCISE 10.13. Given a set I , take a point $x \in I$ and consider the collection \mathcal{U}_x of subsets of I containing x . Prove that \mathcal{U}_x is an ultrafilter on I .

EXERCISE 10.14. Given the set \mathbb{Z} of integers, prove, using Zorn's lemma, that there exists an ultrafilter containing all the non-trivial subgroups of \mathbb{Z} . Such an ultrafilter is called *profinite ultrafilter*. Hint: In the double power set 2^{2^I} consider the partially ordered subset S consisting of all filters which contain all the non-trivial subgroups of \mathbb{Z} .

DEFINITION 10.15. An ultrafilter as in Exercise 10.13 is called a *principal (or atomic) ultrafilter*. A filter that cannot be defined in such a way is called a *non-principal (or free) ultrafilter*.

PROPOSITION 10.16. *An ultrafilter on an infinite set I is non-principal if and only if it contains the Fréchet filter.*

PROOF. We will prove the equivalence between the negations of the two statements. A principal ultrafilter \mathcal{U}_x on I defined by a point x contains $\{x\}$; hence, by (F_4) , it does not contain $I \setminus \{x\}$ which is an element of the Fréchet filter.

Let now \mathcal{U} be an ultrafilter that does not contain the Fréchet filter. This and the axiom (F_4) imply that \mathcal{U} contains a finite subset F of I . If

$$F \cap \bigcap_{A \in \mathcal{U}} A = \emptyset,$$

then there exist $A_1, \dots, A_n \in \mathcal{U}$ such that

$$F \cap A_1 \cap \dots \cap A_n = \emptyset.$$

This and the property (F_2) contradict the property (F_1) .

It follows that

$$F \cap \bigcap_{A \in \mathcal{U}} A = F' \neq \emptyset,$$

in particular, given an element $x \in F'$, \mathcal{U} is contained in the principal ultrafilter \mathcal{U}_x . The maximality of \mathcal{U} implies that $\mathcal{U} = \mathcal{U}_x$. \square

EXERCISE 10.17. (1) Let J be an infinite subset of I . Prove (using Zorn's lemma) that there exists a non-principal ultrafilter \mathcal{U} such that $J \in \mathcal{U}$.

(2) Let $J_1 \supset J_2 \supset J_3 \supset \dots \supset J_m \supset \dots$ be an infinite sequence of infinite subsets of I . Prove that there exists a non-principal ultrafilter containing all J_m , $\forall m \in \mathbb{N}$, as its elements.

LEMMA 10.18 (The Ultrafilter Lemma). *Every filter on a set I is a subset of some ultrafilter on I .*

PROOF. Let \mathcal{F} be the Fréchet filter of I . By Zorn's lemma (cf. Exercise 10.10), there exists a maximal filter \mathcal{U} on I containing \mathcal{F} . By maximality, \mathcal{U} is an ultrafilter; \mathcal{U} is non-principal by Proposition 10.16. \square

In ZF, the Axiom of Choice is equivalent to Zorn's lemma, and the latter, as we just saw, implies the Ultrafilter Lemma.

Here is an alternative way to define ultrafilters:

DEFINITION 10.19. *An ultrafilter on a set I is a finitely additive measure ω on the set I , such that ω takes only the values 0 and 1, and such that $\omega(I) = 1$.*

We would like to stress that each subset of I is supposed to be measurable with respect to ω , in contrast to the measures (like the Lebesgue measure) that one usually encounters in analysis.

In order to verify equivalence two definitions, consider a measure ω , as Definition 10.19,. This measure defines a subset $\mathcal{U} \subset 2^I$,

$$(10.1) \quad J \in \mathcal{U} \iff \omega(J) = 1, \quad J \notin \mathcal{U} \iff \omega(J) = 0.$$

Conversely, given an ultrafilter \mathcal{U} , we define a measure ω on 2^I by the equations (10.1). We leave it to the reader to check that the filter axioms exactly match the finitely additive measure axioms.

Note that for an atomic ultrafilter \mathcal{U}_x defined in Example 10.13, the corresponding measure is the (atomic) Dirac measure δ_x .

DEFINITION 10.20. A *non-principal ultrafilter on a set I* is a finitely additive measure $\omega : 2^I \rightarrow \{0, 1\}$ such that $\omega(I) = 1$ and $\omega(F) = 0$ for every finite subset F of I .

EXERCISE 10.21. Prove the equivalence between Definitions 10.15 and 10.20.

Thus, in what follows, we will use the terminology *ultrafilter* for both maximal filters on I and finitely additive measures on I as above.

REMARK 10.22. Suppose that ω is an ultrafilter on I . Then:

- (1) If $\omega(A_1 \sqcup \cdots \sqcup A_n) = 1$, then there exists $i_0 \in \{1, 2, \dots, n\}$ such that $\omega(A_{i_0}) = 1$ and $\omega(A_j) = 0$ for every $j \neq i_0$.
- (2) If $\omega(A) = 1$ and $\omega(B) = 1$ then $\omega(A \cap B) = 1$.

NOTATION 10.23. Let $(A_i)_{i \in I}$ and $(B_i)_{i \in I}$ be two families of sets or numbers indexed by I , and let \mathcal{R} be a relation which holds for A_i and B_i , for every $i \in I$. We then write $A_i \mathcal{R}_\omega B_i$ if and only if $A_i \mathcal{R} B_i$ ω -almost surely, that is

$$\omega(\{i \in I \mid A_i \mathcal{R} B_i\}) = 1.$$

Examples of such \mathcal{R}_ω 's are: $=_\omega$, $<_\omega$, \subset_ω . For instance, suppose that $(x_n), (y_n)$ are sequences of real numbers and ω is an ultrafilter on \mathbb{N} , such that for ω -all $n \in \mathbb{N}$, $x_n < y_n$. Then we will say that

$$x_n <_\omega y_n.$$

Below we explain how existence of non-principal ultrafilters implies the Hahn–Banach in the following special case: V is the real vector space of bounded sequences of real numbers $\mathbf{x} = (x_n)$, $U \subset V$ is the subspace of convergent sequences of real numbers, p is the sup-norm

$$\|\mathbf{x}\|_\infty = \sup_{n \in \mathbb{N}} |x_n|$$

and $\varphi : U \rightarrow \mathbb{R}$ is the limit function, i.e.

$$\varphi(\mathbf{x}) = \lim_{n \rightarrow \infty} x_n.$$

Instead of the sup-norm, we can as well take

$$p : \mathbf{x} \mapsto \limsup x_n.$$

In other words, we will show how, using a non-principal ultrafilter, one can extend the notion of limit from convergent sequences to bounded sequences. The main tool in this extension is the concept of an *ultralimit*, which we will frequently use in the book.

DEFINITION 10.24. [Ultralimit of a function] Given a function $f : I \rightarrow Y$ (where Y is a topological space) define the ω -limit

$$\omega\text{-}\lim_i f(i)$$

to be a point $y \in Y$ such that for every neighborhood U of y , the pre-image $f^{-1}U$ belongs to ω . The point y is called the *ultralimit* of the function f .

Note that, in general, an ultralimit need not be unique. However, it is unique in the case when Y is Hausdorff:

LEMMA 10.25. 1. *If Y is compact, then every function $f : I \rightarrow Y$ has an ultralimit.*

2. *If Y is Hausdorff, then every function $f : I \rightarrow Y$ has at most one ultralimit.*

PROOF. 1. To prove existence of a limit, assume that there is no point $y \in Y$ satisfying the definition of the ultralimit. Then each point $z \in Y$ possesses a neighborhood U_z such that $f^{-1}U_z \notin \omega$. By compactness, we can cover Y with finitely many of these neighborhoods U_{z_i} , $i = 1, \dots, n$. Therefore,

$$I = \bigcup_{i=1}^n f^{-1}(U_{z_i})$$

and, thus,

$$\emptyset = \bigcap_{i=1}^n (I \setminus f^{-1}(U_{z_i})) \in \omega.$$

This contradicts the definition of a filter.

2. The proof of uniqueness of ultralimits is the same as for uniqueness of ordinary limits in Hausdorff spaces. Suppose that $f : I \rightarrow Y$ has two ultralimits $y_1 \neq y_2$. Since Y is Hausdorff, the points y_1, y_2 have disjoint neighborhoods U_1, U_2 . By the assumption, both sets $f^{-1}(U_1), f^{-1}(U_2)$ are ω -large. However, their intersection is empty since $U_1 \cap U_2$ is empty. This contradicts Axiom (F_2) of filters. (Note that in this part of the proof we did not use the assumption that ω is an ultrafilter, only that it is a filter.) \square

EXAMPLE 10.26. Suppose that $I = \mathbb{N}$ and $x_i = (-1)^i$. Then ω -lim x_i is either -1 or 1 , depending on whether the set of odd or even numbers belongs to ω .

Note that the ω -limit satisfies the “usual “calculus properties,” e.g., linearity:

$$\omega\text{-lim}(\lambda f + \mu g) = \lambda \omega\text{-lim } f + \mu \omega\text{-lim } g$$

for all bounded functions $f, g : I \rightarrow \mathbb{R}$. (Boundedness is needed to ensure existence of ultralimits.) Now, we can prove Hahn–Banach theorem for the space of convergent sequences U , the space of all bounded sequences V and the functional $\varphi := \lim : U \rightarrow \mathbb{R}$. We take

$$\bar{\varphi}((x_i)) = \omega\text{-lim } x_i.$$

Lemma 10.25 implies that every bounded function $f : I \rightarrow \mathbb{R}$ has an ultralimit. In the case when the ordinary limit $\lim_{i \rightarrow \infty} x_i$ exists, it equals the ultralimit $\omega\text{-lim } x_i$. We leave it to the reader to check the inequality

$$\omega\text{-lim } x_i \leq p((x_i))$$

for both $p((x_i)) = \sup_i |x_i|$ and $p(x_i) = \limsup x_i$. This proves Hahn–Banach theorem (in the special case).

EXERCISE 10.27. Show that the ω -limit of a function $f : I \rightarrow Y$ is an accumulation point of the subset $f(I) \subset Y$.

Conversely, if y is an accumulation point of $\{f(i)\}_{i \in I}$, then there is a non-principal ultrafilter ω with $\omega\text{-lim } f = y$, namely an ultrafilter containing the filter \mathcal{F} on I , which is the preimage of the neighborhood basis of y under f .

Thus, an ultrafilter is a device which selects accumulation points for subsets A in compact Hausdorff spaces Y , in a coherent manner.

Note that when the ultrafilter is principal, that is $\omega = \delta_{i_0}$ for some $i_0 \in I$, and Y is Hausdorff, the δ_{i_0} -limit of a function $f : I \rightarrow Y$ is simply the element $f(i_0)$, which is not very interesting. Thus, when considering ω -limits we shall always choose the ultrafilter ω to be non-principal.

REMARK 10.28. Recall that when we have a countable collection of sequences

$$\mathbf{x}^{(k)} = \left(x_n^{(k)} \right)_{n \in \mathbb{N}}, k \in \mathbb{N}, x_n^{(k)} \in X,$$

where X is a compact space, we can select a subset of indices $I \subset \mathbb{N}$, such that for every $k \in \mathbb{N}$ the subsequence $\left(x_i^{(k)} \right)_{i \in I}$ converges. This is achieved by the diagonal procedure. The ω -limit allows, in some sense, to do the same for an uncountable collection of (uncountable) sets. Thus, an ultralimit can be seen as an uncountable version of the diagonal procedure.

Note also that for applications in Geometric Group Theory, most of the time, one considers only countable index sets I . Thus, in principle, one can avoid using ultrafilters at the expense of getting complicated proofs involving passage to multiple subsequences.

Using ultralimits of maps we will later define ultralimits of sequences of metric spaces; in particular, given metric space (X, dist) , we will define an “image of (X, dist) seen from infinitely far away” (an asymptotic cone of (X, dist)). Ultralimits and asymptotic cones will be among key technical tools used in this book.

10.2. Ultrafilters and the Stone-Čech compactification

Let X be a Hausdorff topological space. The *Stone-Čech compactification* of X is a pair consisting of a compact Hausdorff topological space βX and a continuous map $X \rightarrow \beta X$ which satisfies the following universal property:

For every continuous map $f : X \rightarrow Y$, where Y is a compact Hausdorff space, there exists a unique continuous map $g : \beta X \rightarrow Y$, such that the following diagram commutes:

$$\begin{array}{ccc} X & \xrightarrow{\quad} & \beta X \\ & \searrow & \downarrow g \\ & & Y \end{array}$$

This universal property implies uniqueness of the Stone-Čech compactification in the sense that for any two such compactifications $c_1 : X \rightarrow X', c_2 : X \rightarrow X''$, there exists a homeomorphism $h : X' \rightarrow X''$ such that $c_2 = h \circ c_1$.

EXERCISE 10.29. Show that $X \rightarrow \beta X$ is injective and its image is dense in βX .

In view of this exercise, we will regard X as a subset of βX , so that βX is a compactification of X .

We will now explain how to construct βX using ultrafilters, provided that X has discrete topology, e.g., $X = \mathbb{N}$. We declare βX to be the set of *all ultrafilters* on X . **Then** βX is a subset $Filter(X)$, which, in turn, is a subset of the power set

$$2^{2^X}.$$

We equip 2^X and, hence, 2^{2^X} , with the product topology and the subset $\beta X \subset 2^{2^X}$ with the subspace topology.

EXERCISE 10.30. Show that the subset $\beta X \subset 2^{2^X}$ is closed. Thus, by the Tychonoff's theorem, βX is also compact. Since X is Hausdorff, so is 2^X and, hence, 2^{2^X} .

Every $x \in X$ determines the principal ultrafilter δ_x ; thus, we obtain an embedding $X \hookrightarrow \beta X$, $x \mapsto \delta_x$. This embedding is continuous since X has discrete topology. Therefore, from now, on we will regard X as a subset of βX .

EXERCISE 10.31. Let $\omega \in \beta X$ be a non-principal ultrafilter. Show that for every neighborhood U of ω in βX , the intersection $X \cap U$ is an ω -large set. Conversely, for every ω -large set $A \subset X$, there exists a neighborhood U of ω in βX such that $A = U \cap X$. In particular, X is dense in βX .

We will now verify the universal property of βX . Let $f : X \rightarrow Y$ be a continuous map to a compact Hausdorff space. For every $\omega \in \beta X \setminus X$ we set

$$g(\omega) := \omega\text{-lim } f.$$

By the definition of the ultralimit of a map, for every point $y \in Y$ and its neighborhood V in Y , the preimage $A = f^{-1}(V)$ is ω -large. Therefore, by Exercise 10.31, there exists a neighborhood U of ω in βX such that $A = U \cap X$. This proves that the map g is continuous. Hence, g is the required continuous extension of f . Uniqueness of g follows from the fact that X is dense in βX .

10.3. Elements of nonstandard algebra

Our discussion of nonstandard algebra mostly follows [Gol98], [dDW84].

Given an ultrafilter ω on I and a collection of sets $X_i, i \in I$, define the *ultraproduct*

$$\prod_{i \in I} X_i / \omega$$

to be the collection of equivalence classes of maps

$$f : I \rightarrow \bigcup_{i \in I} X_i$$

with $f(i) \in X_i$ for every $i \in I$, with respect to the equivalence relation $f \sim g$ defined by the property that $f(i) = g(i)$ for ω -all i . Thus, an ultraproduct is a certain quotient of the ordinary product of the sets X_i .

The equivalence class of a map f in the ultraproduct is denoted by f^ω . When the map is given by an indexed family of values $(x_i)_{i \in I}$, where $x_i = f(i)$, we will also use the notation $(x_i)^\omega$ for the equivalence class.

When $X_i = X$ for all $i \in I$, the ultraproduct is called the *ultrapower* of X and denoted by X^ω . Every subset A of X can be embedded into X^ω by

$$a \mapsto \hat{a} := (a)^\omega.$$

We let \widehat{A} denote the image of A in X^ω .

Note that any algebraic structure on X (group, ring, order, order, etc.) defines the same structure on X^ω , e.g., if G is a group then G^ω is a group, etc. When $X = \mathbb{K}$ is either \mathbb{N}, \mathbb{Z} or \mathbb{R} , the ultrapower \mathbb{K}^ω is sometimes called the *nonstandard extension of \mathbb{K}* , and the elements in $\mathbb{K}^\omega \setminus \mathbb{K}$ are called *nonstandard elements*. If X is totally ordered then X^ω is totally ordered as well: $f^\omega \leq g^\omega$ (for $f, g \in X^\omega$) if and only if $f(i) \leq_\omega g(i)$, with the Notation 10.23. Since ω is an ultrafilter, it follows that \leq_ω is a total order: This is where ultraproducts are superior to the ordinary products, since the ordinary product of totally ordered sets is (in general) only partially ordered.

In particular, we define the ordered semigroup \mathbb{N}^ω (the nonstandard natural numbers) and the ordered field \mathbb{R}^ω (the nonstandard real numbers).

DEFINITION 10.32. An element $R \in \mathbb{R}^\omega$ is called *infinitely large* if given any $r \in \mathbb{R} \subset \mathbb{R}^\omega$, one has $R \geq \widehat{r}$. Note that given any $R \in \mathbb{R}^\omega$, there exists $n \in \mathbb{N}^\omega$ such that $n > R$.

EXERCISE 10.33. Prove that $R = (R_i)^\omega \in \mathbb{R}^\omega$ is infinitely large if and only if $\omega\text{-lim}_i R_i = +\infty$.

DEFINITION 10.34 (Internal subsets). A subset $W^\omega \subset X^\omega$ is called *internal* if “membership in W can be determined by coordinate-wise computation”, i.e. if for each $i \in I$ there is a subset $W_i \subset X$ such that for $f \in X^I$

$$f^\omega \in W^\omega \iff f(i) \in_\omega W_i.$$

(Recall that the latter means that $f(i) \in W_i$ for ω -all i .) The sets W_i are called *coordinates* of W . We will write $W^\omega = (W_i)^\omega$.

- LEMMA 10.35. (1) *If an internal subset A^ω is defined by a family of subsets of bounded cardinality $A_i = \{a_i^1, \dots, a_i^k\}$, then $A^\omega = \{a_\omega^1, \dots, a_\omega^k\}$, where $a_\omega^j = (a_i^j)^\omega$.*
(2) *In particular, if an internal subset A^ω is defined by a constant family of finite subsets $A_i = A \subseteq X$ then $A^\omega = \widehat{A}$.*
(3) *Every finite subset in X^ω is internal.*

PROOF. (1) Let $x = (x_i)^\omega \in A^\omega$. The set of indices decomposes as $I = I_1 \sqcup \dots \sqcup I_k$, where $I_j = \{i \in I; x_i = a_i^j\}$. Then there exists $j \in \{1, \dots, k\}$ such that $\omega(I_j) = 1$, that is $x_i =_\omega a_i^j$, and $x = a_\omega^j$.

(2) is an immediate consequence of (1).

(3) Let U be a subset in X^ω of cardinality k , and let x_1, \dots, x_k be its elements. Each element x_r is of the form $(x_i^r)^\omega$ and ω -almost surely $x_i^r \neq x_i^s$ when $r \neq s$. Therefore ω -almost surely the set $A_i = \{x_i^1, \dots, x_i^k\}$ has cardinality k . It follows that $A^\omega = (A_i)^\omega$ has cardinality k , according to (1), and it contains U . Therefore $U = A^\omega$. \square

LEMMA 10.36. *If A is an infinite subset in X , then \widehat{A} is not internal.*

PROOF. Assume $\widehat{A} = (B_i)^\omega$ for a family $(B_i)_{i \in I}$ of subsets. For every $a \in A$, $\hat{a} \in (B_i)^\omega$, i.e.

$$(10.2) \quad a \in B_i \quad \omega - \text{almost surely.}$$

Take an infinite sequence $a_1, a_2, \dots, a_k, \dots$ of distinct elements in A . Consider the nested sequence of sets

$$I_k = \{i \in I \mid \{a_1, a_2, \dots, a_k\} \subseteq B_i\}.$$

From (10.2) and Remark 10.22, (2), it follows that $\omega(I_k) = 1$ for every k .

The intersection $J = \bigcap_{n \geq 1} I_n$ has ω -measure either 0 or 1. Assume first that $\omega(J) = 0$. Since

$$I_1 = \bigsqcup_{k=1}^{\infty} (I_k \setminus I_{k+1}) \sqcup J,$$

it follows that the set

$$J' = \bigsqcup_{k=1}^{\infty} (I_k \setminus I_{k+1})$$

has $\omega(J') = 1$.

Define the indexed family (x_i) such that $x_i = a_k$ for every $i \in I_k \setminus I_{k+1}$. By the definition, $x_i \in B_i$ for every $i \in J'$. Thus

$$(x_i)^\omega \in (B_i)^\omega = \widehat{A},$$

which implies that $x_i = a$ ω -a.s. for some $a \in A$.

Let $E = \{i \in I \mid x_i = a\}$, $\omega(E) = 1$. Remark 10.22, (2), implies that $E \cap J' \neq \emptyset$, hence, for some $k \in \mathbb{N}$,

$$E \cap (I_k \setminus I_{k+1}) \neq \emptyset.$$

For $i \in E \cap (I_k \setminus I_{k+1})$ we have $x_i = a = a_k$.

The fact that $\omega(I_{k+1}) = 1$ implies that $E \cap I_{k+1} \cap J' \neq \emptyset$. Hence, for some $j \geq k+1$,

$$E \cap (I_j \setminus I_{j+1}) \neq \emptyset.$$

For an index i in $E \cap (I_j \setminus I_{j+1})$ we have the equality $x_i = a = a_j$. But as $j > k$, $a_j \neq a_k$, and, thus, we obtain a contradiction.

Assume now that $\omega(J) = 1$. Suppose that this occurs for every sequence (a_k) of distinct elements in A . It follows that ω -almost surely $A \subseteq B_i$. \square

DEFINITION 10.37 (internal maps). A map $f^\omega : X^\omega \rightarrow Y^\omega$ is *internal* if there exists an indexed family of maps $f_i : X_i \rightarrow Y_i$, $i \in I$, such that $f^\omega(x^\omega) = (f_i(x_i))^\omega$.

Note that the range of an internal map is an internal set.

For instance, given a collection of metric spaces (X_i, dist_i) , one defines a metric dist^ω on X^ω as the internal function $\text{dist}^\omega : X^\omega \times X^\omega \rightarrow \mathbb{R}^\omega$ given by the collection of functions (dist_i) , that is $\text{dist}^\omega : X^\omega \times X^\omega \rightarrow \mathbb{R}^\omega$,

$$(10.3) \quad \text{dist}^\omega((x_i)^\omega, (y_i)^\omega) = (\text{dist}_i(x_i, y_i))^\omega.$$

The main problem is that dist^ω does not take values in \mathbb{R} but in \mathbb{R}^ω .

Let (II) be a property of a structure on a set X that can be expressed using elements, subsets, \in , \subset , \subseteq , $=$ and the logical quantifiers \exists, \forall, \wedge (and), \vee (or), \neg (not) and \Rightarrow (implies).

The *non-standard interpretation* $(\Pi)^\omega$ of (Π) is the statement obtained by replacing “ $x \in X$ ” with “ $x^\omega \in X^\omega$ ”, and “ A subset of X ” with “ A^ω internal subset of X^ω ”.

THEOREM 10.38 (Łoś’ Theorem, see e.g. [BS69], [Kei76], Chapter 1, [dDW84], p.361). *A property (Π) is true in X if and only if its non-standard interpretation $(\Pi)^\omega$ is true in X^ω .*

We will use the following special cases of this theorem when proving Gromov’s theorem on groups of polynomial growth:

- COROLLARY 10.39.**
- (1) *Every non-empty internal subset in \mathbb{R}^ω that is bounded from above (below) has a supremum (infimum).*
 - (2) *Every non-empty internal subset in \mathbb{N}^ω that is bounded from above (below) has a maximal (minimal) element.*

COROLLARY 10.40 (non-standard induction). *If a non-empty internal subset A^ω in \mathbb{N}^ω satisfies the properties:*

- $\hat{1} \in A^\omega$;
- for every $n^\omega \in A^\omega$, $n^\omega + 1 \in A^\omega$,

then $A^\omega = \mathbb{N}^\omega$.

- EXERCISE 10.41.**
- (1) Give a direct proof of Corollary 10.39, (1), for \mathbb{R}^ω .
 - (2) Deduce Corollary 10.39 from Theorem 10.38.
 - (3) Deduce Corollary 10.40 from Corollary 10.39.

Suppose we are given $a_n \in \mathbb{R}^\omega$, where $n \in \mathbb{N}^\omega$. Using the nonstandard induction principle one can define the nonstandard products:

$$a_1 \cdots a_n, n \in \mathbb{N}^\omega,$$

using the internal function $f : \mathbb{N}^\omega \rightarrow \mathbb{R}^\omega$, given by $f(1) = a_1$, $f(n+1) = f(n)a_{n+1}$.

Various properties of groups can be characterized in terms of ultrapowers, as explained below and in Chapter 18, Section 18.8.

Ultrapowers and laws in groups.

Suppose that G satisfies a law $w(x_1, \dots, x_n) = 1$. Then the ordinary product

$$G^I = \prod_{i \in I} G$$

also satisfies this law: For every function $f \in G^I$ and all $i \in I$,

$$w(f_1, \dots, f_n)(i) = w(f_1(i), \dots, f_n(i)) = 1.$$

Therefore, being a quotient of G^I , the group G^ω satisfies the law $w(x_1, \dots, x_n) = 1$ as well.

Moreover:

LEMMA 10.42 (See Lemma 6.15 in [DS05b]). *A group G satisfies a law if and only if for one (equivalently, every) non-principal ultrafilter ω on \mathbb{N} , the ultrapower G^ω does not contain free non-abelian subgroups.*

PROOF. For the direct implication note that if G satisfies a law, then G^ω also satisfies the same law. Since a free nonabelian group cannot satisfy a law, the claim follows.

For the converse implication, let ω be an ultrafilter on \mathbb{N} , and assume that G does not satisfy any law. Enumerate all the reduced words u_1, u_2, \dots in two variables x, y and define the sequence of iterated left-commutators:

$$v_1 = u_1, v_2 = [u_1, u_2], v_3 = [v_2, u_3], v_4 = [v_3, u_4], \dots,$$

$$v_n = [[u_1, u_2], \dots, u_{n-1}], u_n], \dots$$

We will think of v_n 's as words in x, y .

Since G does not satisfy any law, for every n there exists a pair (x_n, y_n) of elements in G such that $v_n(x_n, y_n) \neq 1$ in G . Consider the corresponding elements $x = (x_n)^\omega, y = (y_n)^\omega$ in the ultrapower G^ω . We claim that the subgroup $F < G^\omega$ generated by x and y is free. Suppose that the subgroup F satisfies a reduced relation. That relation is given by a reduced word u_i for some $i \in \mathbb{N}$. Hence, $u_i(x_n, y_n) = 1$ ω -almost surely. In particular, since ω is a non-principal ultrafilter, for some $n > i$, $u_i(x_n, y_n) = 1$. But then $v_n(x_n, y_n) = 1$ since u_i appears in the iterated commutator v_n , contradicting the choice of x_n, y_n . \square

10.4. Ultralimits of families of metric spaces

Let $(X_i)_{i \in I}$ be a family of metric spaces parameterized by an infinite set I .

CONVENTION 10.43. From now on, all ultrafilters are non-principal, and we will omit mentioning this property henceforth.

For an ultrafilter ω on I we define the ultralimit

$$X_\omega = \omega\text{-}\lim_i X_i$$

as follows. Let $\prod_i X_i$ be the product of the sets X_i , i.e. it is the set of indexed families of points $(x_i)_{i \in I}$ with $x_i \in X_i$. Define the distance between two points $(x_i), (y_i) \in \prod_i X_i$ by

$$\text{dist}_\omega((x_i), (y_i)) := \omega\text{-}\lim(i \mapsto \text{dist}_{X_i}(x_i, y_i)),$$

where we take the ultralimit of the function $i \mapsto \text{dist}_{X_i}(x_i, y_i)$ with values in the compact set $[0, \infty]$. The function dist_ω is a pseudo-distance on $\prod_i X_i$ with values in $[0, \infty]$. Set

$$(X_\omega, \text{dist}_\omega) := \left(\prod_i X_i, \text{dist}_i \right) / \sim$$

where we identify points with zero dist_ω -distance. In the case when $X_i = Y$, for all i , the ultralimit $(X_\omega, \text{dist}_\omega)$ is called a *constant ultralimit*.

The reader will notice similarities between this construction and the Cauchy–Bourbaki completion of a metric space. The difference is that we allow distinct metric spaces instead of a single space and, even if $X_i = Y$ for all i , we do not restrict to indexed families of points (x_i) which are Cauchy. The price we have to pay for this is that, at the moment, dist_ω is merely a pseudo-metric, as it takes infinite values (unless the spaces X_i have uniformly bounded diameter).

Given an indexed family of points $(x_i)_{i \in I}$ with $x_i \in X_i$ we denote the equivalence class corresponding to it either by x_ω or by $\omega\text{-}\lim x_i$.

EXERCISE 10.44. If $(X_\omega, \text{dist}_\omega)$ is a constant ultralimit of a sequence of compact metric spaces $X_i = Y$, then X_ω is isometric to Y for all ultrafilters ω .

If the spaces X_i do not have uniformly bounded diameter, then the ultralimit X_ω decomposes into (in general, uncountably many) components consisting of points at mutually finite distance. In order to pick one of these components, we introduce a family of base-points e_i in X_i . The pair (X_i, e_i) is called a *based* or *pointed* metric space. The indexed family (e_i) defines a base-point $e = e_\omega$ in X_ω and we set

$$X_{\omega, \mathbf{e}} := \{x_\omega \in X_\omega \mid \text{dist}_\omega(x_\omega, e_\omega) < \infty\}.$$

We define the *based ultralimit* as

$$\omega\text{-}\lim_i (X_i, e_i) := (X_{\omega, \mathbf{e}}, e_\omega).$$

By abusing the notation, we will frequently drop \mathbf{e} in the notation $X_{\omega, \mathbf{e}}$ when the choice of the base-point is clear. Given a family of subsets $A_i \subset X_i$ we let A_ω denote the subset of $X_{\omega, \mathbf{e}}$ represented by indexed families $(a_i)_{i \in I}$, $a_i \in A_i$.

EXERCISE 10.45. Let $X = \mathbb{R}^n$ with the Euclidean metric. Then for every sequence $e_i \in X$, $\omega\text{-}\lim(X, e_i) \cong (\mathbb{R}^n, 0)$.

The following theorem relates Gromov-Hausdorff convergence and ultralimits:

THEOREM 10.46 (M. Kapovich and B. Leeb, [KL95]). *Suppose that*

$$(X_i, \text{dist}_{X_i}, x_i)_{i \in \mathbb{N}}$$

is a sequence of proper metric spaces Gromov-Hausdorff converging to a pointed proper metric space (X, dist_X, x) . Then for all ultrafilters ω there exists an isometry between $\omega\text{-}\lim(X_i, \text{dist}_{X_i}, x_i)$ and (X, dist_X, x) sending $x_\omega = \omega\text{-}\lim x_n$ to x .

PROOF. In view of the properness assumption (and the Arzela-Ascoli theorem), it suffices to show that for each $r > 0$, the Gromov-Hausdorff limit of the sequence pointed of closed metric balls $(\bar{B}(x_i, r), x_i)$ in X_i is isometric to $\omega\text{-}\lim(\bar{B}(x_i, r), \text{dist}_i, x_i)$, where dist_i is the restriction of the distance function dist_{X_i} to $\bar{B}(x_i, r)$. Therefore, the problem reduces to the case when $X_i, i \in \mathbb{N}$, and Y are all compact. We realize Gromov-Hausdorff convergence as Hausdorff convergence in a compact metric space Y , i.e. embed each X_i and X isometrically into Y via isometric maps

$$f_i : X_i \rightarrow X'_i := f_i(X_i) \subset Y,$$

such that the Hausdorff limit of the sequence (X'_i) is $X' \cong X$:

$$\lim_{\text{Haus}} X'_i = X'.$$

Then the sequence of isometric embeddings there is an isometric embedding

$$f_\omega : X_\omega \rightarrow \omega\text{-}\lim Y = Y.$$

Since ω is non-principal, the ω -limit is independent of any finite collections of X_i 's and we get:

$$f_\omega(X_\omega) \subset \bigcap_{i_0 \in I} \overline{\bigcup_{i \geq i_0} X'_i} = X'.$$

On the other hand, $X \subset f_\omega(X_\omega)$ since $f_\omega((x_i)) = x$ whenever $\lim_{i \in I} x_i = x \in Y$. Hence, $X' = f_\omega(X_\omega)$. \square

EXAMPLE 10.47. Suppose that X_i is the sequence of spheres of radius $R_i \rightarrow \infty$ in \mathbb{E}^n with the induced path-metric. Then

$$X_\omega = \omega\text{-lim}(X_i, x_i) \cong \mathbb{E}^{n-1}$$

for any choice of base-points $x_i \in X_i$. Indeed, for each fixed r , define the sequence of subsets closed r -balls $Y_i(x_i, r) \subset X_i$. Then, since the sequence R_i diverges to infinity, the sequence of spaces $Y_i(x_i, r)$ Gromov-Hausdorff converges to the closed r -ball $\bar{B}(o, r) \subset \mathbb{E}^{n-1}$. Therefore, by the above lemma, for each r there is an isometry $h_r : (\bar{B}(o, r), o) \rightarrow \omega\text{-lim}(Y_i(x_i, r), x_i)$. It is clear that

$$X_\omega = \bigcup_{r>0} Y_\omega(r),$$

where each $Y_\omega(r)$ is isometric to $\omega\text{-lim}(Y_i(x_i, r), x_i)$. Composing the isometries h_r and $\omega\text{-lim}(Y_i(x_i, r), x_i) \rightarrow (Y_\omega(r), x_r)$ and taking an ultralimit as $r \rightarrow \infty$, we obtain the required isometry $\mathbb{E}^{n-1} \rightarrow X_\omega$.

LEMMA 10.48 (Functoriality of ultralimits). 1. Let $(X_i, e_i), (X'_i, e'_i), i \in I$, be families of pointed metric spaces with ultralimits X_ω, X'_ω , respectively. Let $f_i : (X_i, e_i) \rightarrow (X'_i, e'_i)$ be isometric embeddings such that

$$\omega\text{-lim dist}(f(e_i), e'_i) < \infty,$$

i.e.

$$\text{dist}(f(e_i), e'_i) \leq \text{Const, for } \omega\text{-all } i,$$

Then the maps f_i yield an isometric embedding of the ultralimits $f_\omega : X_\omega \rightarrow X'_\omega$.

2. If each f_i is an isometry, then so is f_ω .

3. $\Phi_\omega : (f_i) \mapsto f_\omega$ preserves compositions:

$$\Phi_\omega : (g_i \circ f_i) = \Phi_\omega((g_i)) \circ \Phi_\omega((f_i)).$$

PROOF. We define f_ω as

$$f_\omega((x_i)) = (f_i(x'_i)).$$

By the definition of distances in X_ω and X'_ω ,

$$d(f_\omega(x_\omega), f_\omega(y_\omega)) = \omega\text{-lim } d(f_i(x_i), f_i(y_i)) = \omega\text{-lim } d(x_i, y_i) = d(x_\omega, y_\omega)$$

for any pair of points $x_\omega, y_\omega \in X_\omega$. If each f_i is surjective, then, clearly, f_ω is surjective as well. The composition property is clear as well. \square

EXERCISE 10.49. Show that injectivity of each f_i does not imply injectivity of f_ω .

The map f_ω defined in this lemma is called the *ultralimit* of the sequence of maps (f_i) . An important example illustrating this lemma is the case when each X_i is an interval in \mathbb{R} and, hence, each f_i is a geodesic in Y_i . Then the ultralimit $f_\omega : J_\omega \rightarrow X_\omega$ is a geodesic in X_ω (here J_ω is an interval in \mathbb{R}).

DEFINITION 10.50. Geodesics $f_\omega : J_\omega \rightarrow X_\omega$ are called *limit geodesics* in X_ω .

In general, X_ω contains geodesics which are not limit geodesics. In the extreme case, Y_i may contain only constant geodesics, while Y_ω is a geodesic metric space (containing more than one point). For instance, let $X = \mathbb{Q}$ with the metric induced from \mathbb{R} . Of course, \mathbb{Q} contains no nonconstant geodesics, but

$$\omega\text{-lim}(X, 0) \cong (\mathbb{R}, 0),$$

see Exercise 10.62.

LEMMA 10.51. *Each ultralimit (X_ω, e_ω) of a sequence of pointed geodesic metric spaces (X_i, e_i) is again a geodesic metric space.*

PROOF. Let $x_\omega = (x_i), y_\omega = (y_i)$ be points in X_ω . Let $\gamma_i : [0, T_i] \rightarrow X_i$ be geodesics connecting x_i to y_i . Clearly,

$$\omega\text{-lim } T_i = T = d(x_\omega, y_\omega) = T < \infty.$$

We define the ultralimit γ_ω of the maps γ_i . Then $\gamma_\omega : [0, T] \rightarrow X_\omega$ is a geodesic connecting x_ω to y_ω . \square

EXERCISE 10.52. Let X be a path-metric space. Then every constant ultralimit of X is a geodesic metric space.

We now return to the discussion of basic properties of ultralimits.

LEMMA 10.53. *Let (X_i, e_i) be pointed $CAT(\kappa_i)$ metric spaces, $\kappa_i \leq 0$, and $\kappa = \omega\text{-lim } \kappa_i$. Then the ultralimit (X_ω, e_ω) of the sequence (X_i, e_i) is again a pointed $CAT(\kappa)$ space.*

PROOF. It is clear that comparison inequalities for triangles in X_i yield comparison inequalities for limit triangles in X_ω . It remains to show that X_ω is a uniquely geodesic metric space, in which case every geodesic segment in X_ω is a limit geodesic. Suppose that $m_\omega \in X_\omega$ is a point such that

$$d(x_\omega, z_\omega) + d(z_\omega, y_\omega) = d(x_\omega, y_\omega),$$

equivalently, z_ω belongs to *some geodesic* connecting x_ω to y_ω .

Thus, if $z_i \in X_i$ is a sequence representing z_ω , then

$$0 \leq d(x_i, z_i) + d(z_i, y_i) = d(x_i, y_i) \leq \eta_i, \quad \omega\text{-lim } \eta_i = 0.$$

Let us assume that $s_i = d(x_i, z_i) \leq d(z_i, y_i)$ and consider the point $q_i \in x_i y_i$ within distance s_i from x_i . Compare the triangle $T_i = T(x_i, y_i, z_i)$ with the Euclidean triangle using the comparison points $p_i = z_i$ and q_i . In the Euclidean comparison triangle \tilde{T}_i , we have

$$\omega\text{-lim } d(\tilde{z}_i, \tilde{q}_i) = 0$$

(since the constant ultralimit of the sequence of Euclidean planes is the Euclidean plane and, hence, is uniquely geodesic). Since, by the $CAT(0)$ -comparison inequality,

$$d(z_i, q_i) \leq d(\tilde{z}_i, \tilde{q}_i),$$

we conclude that $(q_i) = z_\omega$ in the space X_ω . Thus, z_ω lies on the limit geodesic connecting x_ω and y_ω . \square

EXERCISE 10.54. Show that every ultralimit of any sequence of median spaces is again median. Hint: Follow the proof of Proposition 6.42.

LEMMA 10.55 (Ultralimits preserve direct products of metric spaces). *Suppose that $X_i = X'_i \times X''_i, i \in I$, is an indexed family of direct products of metric spaces, i.e. the metrics on X_i are given by the Pythagorean formula (2.2). Then for every ω and a family of base-points $e_i \in X_i, e_i = e'_i \times e''_i$ we have an isometry*

$$X_\omega = \omega\text{-lim}(X_i, e_i) \cong \omega\text{-lim}(X'_i, e'_i) \times \omega\text{-lim}(X''_i, e''_i).$$

PROOF. By the definition of an ultralimit, as a set, X_ω splits naturally as a direct product of two ultralimits $X'_\omega = \omega\text{-lim}(X'_i, e'_i)$ and $X''_\omega = \omega\text{-lim}(X''_i, e''_i)$. Consider points $x_\omega = x'_\omega \times x''_\omega$ and $y_\omega = y'_\omega \times y''_\omega$ in X_ω , where $x'_\omega = (x'_i), y'_\omega = (y'_i), x''_\omega = (x''_i), y''_\omega = (y''_i)$. The distance between these points is given by

$$\begin{aligned} \text{dist}^2(x_\omega, y_\omega) &= \omega\text{-lim dist}^2_{X_i}(x_i, y_i) = \\ &= \omega\text{-lim} \left(\text{dist}^2_{X'_i}(x'_i, y'_i) + \text{dist}^2_{X''_i}(x''_i, y''_i) \right) = \\ &= \omega\text{-lim dist}^2_{X'_i}(x'_i, y'_i) + \omega\text{-lim dist}^2_{X''_i}(x''_i, y''_i) = \text{dist}^2_{X'_\omega}(x'_\omega, y'_\omega) + \text{dist}^2_{X''_\omega}(x''_\omega, y''_\omega). \end{aligned}$$

Lemma follows. \square

10.5. Completeness of ultralimits and incompleteness of ultrafilters

So far, our discussion of ultralimits did not depend on the nature of the set I and the ultrafilter ω (as long as the latter was non-principal). In this section we discuss the question of completeness of ultralimits of families of metric spaces. It turns out that the answer depends on the ultrafilter.

DEFINITION 10.56. An ultrafilter ω is called *countably complete* if it is closed under countable intersections.

Each principal ultrafilter is obviously countably complete. (In fact, a principal ultrafilter is closed under arbitrary intersections.) On the other hand, as we will see soon, any non-principal ultrafilter on a countable set is countably incomplete, and, hence, for the purposes of Geometric Group Theory, countably complete ultrafilters are irrelevant. Existence of countably complete non-principal ultrafilters is unprovable in ZFC, we refer the refer to [Kei10] for details and references.

Below is a characterization of complete ultrafilters that we will need.

LEMMA 10.57. *The following are equivalent for an ultrafilter ω on a set I :*

1. ω is countably incomplete.
2. There exists a map $\nu : I \rightarrow \mathbb{N}$, which sends ω to a non-principal ultrafilter, $\nu(\omega)$, i.e. for each finite subset $S \subset \mathbb{N}$ the preimage of S under ν is ω -large.

PROOF. 1. Suppose that ω is countably incomplete and, hence, there exists a sequence (J_n) of ω -large subsets of I with intersection J not in ω . By taking finite intersections of the sets J_n , we can assume that the sequence J_n is strictly decreasing:

$$J_1 \supset J_2 \supset J_3 \supset \dots$$

We define the following function $\nu : I \rightarrow \mathbb{N} \cup \{\infty\}$:

For each $i \in I$ we let $\nu(i)$ denote the supremum of the set

$$\{n : i \in J_n\}.$$

If this set is empty, we, of course, have $\nu(i) = 1$; if this set is unbounded, $\nu(i) = \infty$. Clearly, $\nu^{-1}(\infty) = J$ is not ω -large. If there exists a finite subset $[1, n] \subset \mathbb{N}$ such that $K = \nu^{-1}([1, n])$ is ω -large, then K is disjoint from J_{n+1} , which is a contradiction. Hence, $\nu(\omega)$ is a non-principal ultrafilter.

2. Suppose there exists a map $\nu : I \rightarrow \mathbb{N}$ which sends ω to a non-principal ultrafilter. Then for each interval $[n, \infty) \subset \mathbb{N}$, the preimage $\nu^{-1}([n, \infty))$ is ω -large. The intersection of these preimages is empty and, hence, ω is countably incomplete. \square

COROLLARY 10.58. *Any non-principal ultrafilter on a countable set is countably incomplete.*

Countably complete ultrafilters behave essentially like principal ultrafilters, as far as convergence in metric spaces is concerned:

LEMMA 10.59. *Suppose that X is a first countable Hausdorff topological space and $(x_i)_{i \in I}$ is an indexed family in X which ω -converges to some $x \in X$. Then the family $(x_i)_{i \in I}$ is ω -constant:*

$$\omega(\{i \in I : x_i = x\}) = 1.$$

PROOF. Consider a countable basis of topology U_n at the point x . Then for each n ,

$$\omega(\{i : x_i \in U_n\}) = 1.$$

Since the intersection J of the sets $\{i : x_i \in U_n\}$ is still ω -large and X is Hausdorff, we conclude that for ω -all i 's, $x_i = x$. \square

COROLLARY 10.60. *Suppose that ω is countably complete. Then for each family $(X_i, e_i)_{i \in I}$ of pointed metric spaces, the ultralimit $\omega\text{-lim}(X_i, e_i)$ is isometric to the pointed ultraproduct*

$$\left(\prod_{i \in I} (X_i, e_i) \right) / \omega.$$

In particular, the constant ultralimit $(X, e)_{i \in I}$ of an incomplete metric space X is still incomplete.

Thus, countably complete ultrafilters lead to incomplete ultralimits. We now turn to countably incomplete ultrafilters. The reader interested only in Geometric Group Theory applications, can safely assume here that the index set I is countable.

LEMMA 10.61. *Let $(Y_i)_{i \in I}$ be a family of metric spaces, and for every i let X_i be a dense subset in Y_i . Then for every countably incomplete ultrafilter ω on I , the natural isometric embedding of the ultralimit $\omega\text{-lim}_i X_i$ into the ultralimit $Y_\omega = \omega\text{-lim}_i Y_i$ is surjective. In particular, this holds when $Y_i = \widehat{X}_i$, the metric completion of X_i .*

PROOF. We first give a proof in the case when I is countable, since it is based on a *diagonal subsequence* argument probably familiar to the reader. We will identify I with the set of the natural numbers \mathbb{N} and consider a point $y_\omega \in Y_\omega$ and the corresponding indexed family $(y_i)_{i \in \mathbb{N}}$. By density of X_i in Y_i , for each i there exists a sequence $(x_{in})_{n \in \mathbb{N}}$ in X_i whose limit is y_i . For each i we choose n_i such that for $x_i := x_{in_i}$,

$$\text{dist}_{Y_i}(x_i, y_i) < \frac{1}{i}.$$

It follows that

$$\omega\text{-lim dist}_{Y_i}(x_i, y_i) = 0$$

and, hence, $(x_i) = (y_i)$ in Y_ω .

Suppose now that ω is a general countably incomplete ultrafilter and $\nu : I \rightarrow \mathbb{N}$ is a mapping which sends ω to a non-principal ultrafilter on \mathbb{N} . For each i choose $x_i \in X_i$ such that

$$\text{dist}_{Y_i}(x_i, y_i) < \frac{1}{\nu(i)}.$$

(Here we are using the AC!) Since $\nu(\omega)$ is a non-principal ultrafilter on \mathbb{N} ,

$$\omega\text{-lim} \frac{1}{\nu(i)} = 0,$$

hence,

$$\omega\text{-lim} \text{dist}_{Y_i}(x_i, y_i) = 0$$

as well. We again conclude that $(x_i) = (y_i)$ in Y_ω . \square

EXERCISE 10.62. Let Y be a proper metric space, take a subset $X \subset Y$ equipped with the restriction metric. Then for each countably incomplete ultrafilter ω , the constant ultralimit $\omega\text{-lim}(X, e)$ is naturally isometric to (\bar{X}, e) , where \bar{X} is the closure of X in Y . Hint: Use Exercise 10.44.

In the next proposition, we make no assumptions about completeness of the ultrafilter ω :

PROPOSITION 10.63. *Every based ultralimit $\omega\text{-lim}_i(X_i, e_i)$ of a family of complete metric spaces is a complete metric space.*

PROOF. We will prove that every Cauchy sequence $(x^{(k)})$ in $X_{\omega, e}$ contains a convergent subsequence, this will imply that $(x^{(k)})$ converges as well. We select a subsequence (which we again denote $(x^{(k)})$) such that

$$\text{dist}_\omega(x^{(k)}, x^{(k+1)}) < \frac{1}{2^k}.$$

Equivalently,

$$\omega\text{-lim}_i \text{dist}_i(x_i^{(k)}, x_i^{(k+1)}) < \frac{1}{2^k},$$

which implies that

$$\text{dist}_i(x_i^{(k)}, x_i^{(k+1)}) < \frac{1}{2^k}, \quad \omega - \text{a.s.},$$

i.e. for every k the following set is ω -large:

$$I_k = \left\{ i \in I; \text{dist}_i(x_i^{(k)}, x_i^{(k+1)}) < \frac{1}{2^k} \right\}.$$

We can assume that $I_{k+1} \subseteq I_k$, otherwise we replace I_{k+1} with $I_{k+1} \cap I_k$. Thus, we obtain a nested sequence of subsets I_k in I :

$$I_1 \supset I_2 \supset I_3 \supset \dots$$

Case 1. Assume first that the intersection $J := \bigcap_{k \geq 1} I_k$ of these subsets is also ω -large. (This will be always the case if ω is countably complete.)

For every $i \in J$ the sequence $(x_i^{(k)})$ is Cauchy, therefore, since the space X_i is complete, this sequence converges to some $y_i \in X_i$. The inequalities

$$\text{dist}_i(x_i^{(k)}, x_i^{(k+1)}) < \frac{1}{2^k}, \quad k \in \mathbb{N},$$

imply that for every $m > k$,

$$\text{dist}_i(x_i^{(k)}, x_i^{(m)}) < \frac{1}{2^{k-1}}.$$

The latter gives, by taking the limit as $m \rightarrow \infty$, that

$$\text{dist}_i(x_i^{(k)}, y_i) \leq \frac{1}{2^{k-1}}.$$

Hence,

$$\text{dist}_\omega \left(x^{(k)}, y_\omega \right) \leq \frac{1}{2^{k-1}},$$

for $y_\omega = \omega\text{-lim } y_i$. We have, thus, obtained a limit y_ω of the sequence $(x^{(k)})$.

Case 2. Assume now that $\omega(J) = 0$. Since for every $k \geq 1$ we have that

$$I_k = J \sqcup \bigsqcup_{j=k}^{\infty} (I_j \setminus I_{j+1})$$

and $\omega(I_k) = 1$, it follows that

$$\omega \left(\bigsqcup_{j=k}^{\infty} (I_j \setminus I_{j+1}) \right) = 1.$$

We define subsets

$$J_k := \bigsqcup_{j=k}^{\infty} (I_j \setminus I_{j+1}) \subset I.$$

We claim that the limit point of the sequence $(x^{(k)})$ is $y_\omega = (y_i) \in Y_\omega$, where $y_i = x_i^{(k)}$ whenever $i \in I_k \setminus I_{k+1}$. This defines y_i for all $i \in J_1$. We extend this definition to the rest of I arbitrarily: Values taken on ω -small sets of indices $i \in I$ do not matter.

For every

$$i \in J_k = \bigsqcup_{j=k}^{\infty} (I_j \setminus I_{j+1})$$

there exists $j \geq k$ such that $i \in I_j \setminus I_{j+1}$. By the definition, $y_i = x_i^{(j)}$.

Since

$$i \in I_j \subseteq I_{j-1} \subseteq \cdots \subseteq I_{k+1} \subseteq I_k$$

we may write

$$\begin{aligned} \text{dist}_i \left(x_i^{(k)}, y_i \right) &\leq \text{dist}_i \left(x_i^{(k)}, x_i^{(k+1)} \right) + \cdots + \text{dist}_i \left(x_i^{(j-1)}, x_i^{(j)} \right) \leq \\ &\frac{1}{2^k} + \frac{1}{2^{k+1}} + \cdots + \frac{1}{2^{j-1}} \leq \frac{1}{2^k} \frac{1}{1 - \frac{1}{2}} = \frac{1}{2^{k-1}}. \end{aligned}$$

Thus, we have

$$\text{dist}_\omega \left(x^{(k)}, y_\omega \right) \leq \frac{1}{2^{k-1}},$$

which implies that the sequence $x^{(k)}$ indeed converges to y_ω . \square

COROLLARY 10.64. *Suppose that ω is a countably incomplete ultrafilter (e.g., I is countable). Then:*

1. *The ultralimit $\omega\text{-lim}(X_i, e_i)$ of any family of based metric spaces $(X_i, e_i)_{i \in I}$ is complete.*
2. *For each family $A_i \subset X_i$ of subsets, the ultralimit $A_\omega \subset X_\omega$ is a closed subset.*

PROOF. This is a combination of Lemma 10.61 with Proposition 10.63. \square

Thus, (countably) incomplete ultrafilters, lead to metric completeness of ultralimits.

EXAMPLE 10.65. Let $(\mathcal{H}_i)_{i \in \mathbb{N}}$ be a sequence of Hilbert spaces and let $Y_i = S(0, R_i) \subset \mathcal{H}_i$ be metric spheres of radii R_i diverging to infinity. Then for each non-principal ultrafilter ω and choice of base-points $y_i \in Y_i$, the ultralimit $\omega\text{-lim}(Y_i, y_i)$ is isometric to a Hilbert space. Indeed, in view of the Example 10.47, if we fix n and let Σ_i denote the intersection of Y_i with an n -dimensional subspace in \mathcal{H}_i containing y_i , then $\omega\text{-lim}(\Sigma_i, y_i) \cong \mathbb{E}^{n-1}$. It follows that Y_ω is a complete (see Proposition 10.63) geodesic metric space such that:

1. Each finite subset of Y_ω is isometric to a subset of a Euclidean space.
2. For each geodesic segment $xy \subset Y_\omega$ there exists a complete geodesic in Y_ω containing xy .

Combining the first property with Theorem 2.90, we conclude that there exists an isometric embedding $\phi : Y_\omega \rightarrow \mathcal{H}$, where \mathcal{H} is a Hilbert space. Without loss of generality, we may assume that $\phi(y_\omega) = 0$. Since Y_ω is a geodesic metric space and \mathcal{H} is uniquely geodesic with geodesics given by line segments, the image $\phi(Y_\omega)$ is a convex subset of \mathcal{H} . Furthermore, the second property mentioned above implies with each point $y \neq 0$, the subset $\phi(Y_\omega)$ contains the line $\mathbb{R}y \subset \mathcal{H}$. It follows that $\phi(Y_\omega)$ is a linear subspace V in \mathcal{H} . Completeness of V follows from that of Y_ω .

10.6. Asymptotic cones of metric spaces

The concept of an asymptotic cone was first introduced in the Geometric Group Theory by van den Dries and Wilkie in [dDW84], although its version for groups of polynomial growth was already used by Gromov in [Gro81a], who used Gromov-Hausdorff convergence as a tool. Asymptotic cones (and ultralimits) for general metric spaces were defined by Gromov in [Gro93]. The idea is to construct, for a metric space (X, dist) , its “image” seen from “infinitely far.” More precisely, one defines the notion of a *limit* of a sequence of metric spaces $(X, \varepsilon \text{dist})$, $\varepsilon > 0$, as $\varepsilon \rightarrow 0$.

Let (X, d_X) be a metric space and ω be a non-principal ultrafilter on I . For each positive real number λ we define the new metric space $\lambda X = (X, \lambda d_X)$ by rescaling the metric d_X . Suppose that we are given a family $\boldsymbol{\lambda} = (\lambda_i)_{i \in I}$ of positive real numbers indexed by I such that $\omega\text{-lim} \lambda_i = 0$ and a family $\mathbf{e} = (e_i)_{i \in I}$ of base-points $e_i \in X$ indexed by I . Given this data, the *asymptotic cone* $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ of X is defined as the based ultralimit of rescaled copies of X :

$$\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda}) := \omega\text{-lim}_i (\lambda_i \cdot X, e_i).$$

For a family of points $(x_i)_{i \in I}$ in X , the corresponding subset in the asymptotic cone $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$, which is either a one-point set, or the empty set if $\omega\text{-lim} \lambda_i \text{dist}(x_i, e_i) = \infty$, is denoted by $\omega\text{-lim} x_i$.

The family $\boldsymbol{\lambda} = (\lambda_i)_{i \in I}$ is called the *scaling family*. When either the scaling family or the family of base-points are irrelevant, they are omitted from the notation.

Thus, to each metric space X we attach a collection of metric spaces $\text{Cones}(X)$ consisting of all asymptotic cones $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ of X , that is of all the “images of X seen from infinitely far.” The first questions to ask are: How large is the collection $\text{Cones}(X)$ for specific metric spaces X , and what features of X are inherited by the metric spaces in $\text{Cones}(X)$.

An asymptotic cone of a finitely generated group is the asymptotic cone of this group regarded as a metric space, where we use the word metric defined by the

given finite generating set. As we will see below, the *bi-Lipschitz homeomorphism class* of such asymptotic cone is independent of the generating set and the choice of base-points, but does depend on the ultrafilter ω and the scaling family λ .

We begin by noting that the choice of base-points is irrelevant for spaces that are quasihomogeneous:

EXERCISE 10.66. [See also Proposition 10.72.] When the space X is quasihomogeneous, all cones defined by the same fixed ultrafilter ω and sequence of scaling constants λ , are isometric.

Another simple observation is:

REMARK 10.67. Let α be a positive real number. The map

$$I_\alpha : \text{Cone}_\omega(X, \mathbf{e}, \lambda) \rightarrow \text{Cone}_\omega(X, \mathbf{e}, \alpha\lambda), \quad I_\alpha(\omega\text{-lim } x_i) = \omega\text{-lim } x_i$$

is a similarity: It multiplies all the distances by the factor α . Thus, for a fixed metric space X , the collection of asymptotic cones $\text{Cones}(X)$ is stable with respect to rescaling of the metric on X .

In particular, since the Euclidean space \mathbb{E}^n is proper, homogeneous and *self-similar* (\mathbb{E}^n is isometric to $\alpha\mathbb{E}^n$ for each $\alpha > 0$), it follows that

$$\text{Cone}_\omega \mathbb{E}^n \cong \mathbb{E}^n.$$

The same applies to all finite-dimensional normed vector spaces $(V, \|\cdot\|)$:

$$\text{Cone}_\omega(V, \|\cdot\|) \cong (V, \|\cdot\|).$$

Lemmata 10.55 and 10.51 imply that asymptotic cones preserve direct product decompositions of metric spaces and geodesic metric spaces:

COROLLARY 10.68. (1) $\text{Cone}_\omega(X \times Y) = \text{Cone}_\omega(X) \times \text{Cone}_\omega(Y)$.
 (2) *The asymptotic cone of a geodesic space is a geodesic space.*

DEFINITION 10.69. Given a family $(A_i)_{i \in I}$ of subsets of (X, dist) , we denote either by $\omega\text{-lim } A_i$ or by A_ω the subset of $\text{Cone}_\omega(X, \mathbf{e}, \lambda)$ that consists of all the elements $\omega\text{-lim } x_i$ such that $x_i \in A_i$ ω -almost surely. We call $\omega\text{-lim } A_i$ the *limit set* of the family $(A_i)_{i \in I}$.

Note that if $\omega\text{-lim } \lambda_i \text{dist}(e_i, A_i) = \infty$ then the set $\omega\text{-lim } A_i$ is empty.

PROPOSITION 10.70 (Van den Dries and Wilkie, Proposition 4.2 in [dDW84]). *If ω is countably incomplete (e.g., the index set I is countable) then:*

- (1) *Any asymptotic cone (with respect to ω) of a metric space is complete.*
- (2) *For each family $A_i \subset X_i$, the limit set $\omega\text{-lim } A_i$ is a closed subset of $\text{Cone}_\omega(X, \mathbf{e}, \lambda)$.*

PROOF. This is an immediate consequence of Proposition 10.63. □

In Definition 10.50 we introduced the notion of *limit geodesics* in the ultralimit of a sequence of metric spaces. Let $\gamma_i : [a_i, b_i] \rightarrow X$ be a family of geodesics with the limit geodesic γ_ω in $\text{Cone}_\omega(X, \mathbf{e}, \lambda)$.

EXERCISE 10.71. Show that the image of γ_ω is the limit set of the sequence of images of the geodesics γ_i .

We saw earlier that geodesics in the ultralimit may fail to be limit geodesics. However, in our example, we took a sequence of metric spaces which were not geodesic. It turns out that, in general, there exist geodesics in $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ that are not limit geodesic, even when X is the Cayley graph of a finitely generated group with a word metric. An example of this can be found in [Dru09].

Suppose that X is a metric space and $G \subset \text{Isom}(X)$ is a subgroup. Given a non-principal ultrafilter ω consider the ultraproduct $G^\omega = \prod_{i \in I} G/\omega$. For a family of positive real numbers $\boldsymbol{\lambda} = (\lambda_i)_{i \in I}$ such that $\omega\text{-lim } \lambda_i = 0$ and a family of base-points $\mathbf{e} = (e_i)$ in X , let $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ be the corresponding asymptotic cone. In view of Lemma 10.48, the group G^ω acts isometrically on the ultralimit

$$U := \omega\text{-lim}(\lambda_i \cdot X).$$

Let $G_{\mathbf{e}}^\omega \subset G^\omega$ denote the stabilizer in G^ω of the component $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda}) \subset U$. In other words,

$$G_{\mathbf{e}}^\omega = \{(g_i)^\omega \in G^\omega : \omega\text{-lim } \lambda_i \text{dist}(g_i(e_i), e_i) < \infty\}.$$

There is a natural homomorphism $G_{\mathbf{e}}^\omega \rightarrow \text{Isom}(\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda}))$. Observe also that if (e_i) is a bounded family in X then the group G has a diagonal embedding into $G_{\mathbf{e}}^\omega$.

PROPOSITION 10.72. *Suppose that $G \subset \text{Isom}(X)$ and the action $G \curvearrowright X$ is cobounded. Then for every asymptotic cone $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ the action $G_{\mathbf{e}}^\omega \curvearrowright \text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ is transitive. In particular, $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ is a homogeneous metric space.*

PROOF. Let $D < \infty$ be such that $G \cdot x$ is a D -net in X . Given two indexed families $(x_i), (y_i)$ of points in X , there exists an indexed family (g_i) of elements of G such that

$$\text{dist}(g_i(x_i), y_i) \leq 2D.$$

Therefore, if $g_\omega := (g_i)^\omega \in G^\omega$, then $g_\omega(\omega\text{-lim}_i x_i) = \omega\text{-lim}_i y_i$. Hence the action

$$G^\omega \curvearrowright U = \omega\text{-lim}_i(\lambda_i \cdot X)$$

is transitive. It follows that the action $G_{\mathbf{e}}^\omega \curvearrowright \text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ is transitive as well. \square

EXERCISE 10.73. 1. Construct an example of a metric space X , a bounded sequence (e_i) and an asymptotic cone $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ so that for the isometry group $G = \text{Isom}(X)$ the action $G_{\mathbf{e}}^\omega \curvearrowright \text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ is not effective, i.e. the homomorphism

$$G_{\mathbf{e}}^\omega \rightarrow \text{Isom}(\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda}))$$

has non-trivial kernel. Construct an example when the kernel of the above homomorphism contains the entire group G embedded diagonally in $G_{\mathbf{e}}^\omega$.

2. Show that $\text{Ker}(G \rightarrow \text{Isom}(X))$ is contained in $\text{Ker}(G \rightarrow \text{Isom}(X_\omega))$.

Suppose that X admits a cocompact discrete action of a subgroup $G < \text{Isom}(X)$. The problem of how large the class of spaces $\text{Cones}(X)$ can be, that is the problem of the dependence of the topological/metric type of $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ on the ultrafilter ω and the scaling sequence $\boldsymbol{\lambda}$, is, in general, quite hard. In some special cases, it is related to the Continuum Hypothesis (the hypothesis stating that there

is no cardinal number between \aleph_0 and 2^{\aleph_0}). Consider, for concreteness, the group $SL(n, \mathbb{R})$, $n \geq 3$, equipped with a fixed left-invariant metric.

Kramer, Shelah, Tent and Thomas have shown in [KSTT05] that:

- (1) If the Continuum Hypothesis (CH) is not true then the group $SL(n, \mathbb{R})$, $n \geq 3$, has $2^{2^{\aleph_0}}$ non-homeomorphic asymptotic cones.
- (2) If the CH is true then all asymptotic cones of $SL(n, \mathbb{R})$, $n \geq 3$, are isometric. Moreover, under the same assumption, a finitely generated group (with a fixed finite generating set) has at most continuum of non-isometric asymptotic cones.

Moreover, according to Theorem 1.4 in Blake Thornton's PhD thesis [Tho02]:

THEOREM 10.74. *Assuming CH, if X is a non-positively curved symmetric space then all asymptotic cones of X are isometric to each other, i.e. up to isometry asymptotic cones are independent of the scaling sequence and the choice of a nonprincipal ultrafilter.*

The case of $SL(2, \mathbb{R})$ was settled independently of the CH by A. Dyubina-Erschler and I. Polterovich (see Theorem 11.174).

Chronologically, the first non-trivial example of metric space X such that the set $\text{Cones}(X)$ contains very few elements (up to bilipschitz homeomorphisms) is that of virtually nilpotent groups, and is due to P. Pansu, see Theorem 16.28.

C. Druţu and M. Sapir constructed in [DS05b] an example of two-generated and recursively presented (but not finitely presented) group with continuum of pairwise non-homeomorphic asymptotic cones. The construction is independent of the Continuum Hypothesis. The example can be adapted so that at least one asymptotic cone is a real tree.

Note that if a finitely presented group G has *one* asymptotic cone which is a real tree, then the group is hyperbolic and hence *every* asymptotic cone of G is a real tree, see Theorem 11.170.

Historical remarks. The first instance (that we are aware of) where asymptotic cones of metric spaces were defined is the 1966 paper [BDCK66], where this is done in the context of normed vector spaces. Their definition, though, works for all metric spaces.

On the other hand, Gromov introduced the modified Hausdorff distance (see Section 8.1 for a definition) and the corresponding limits of sequences of pointed metric spaces in his work on groups of polynomial growth [Gro81a]. This approach is no longer appropriate in the case of more general metric spaces, as we will explain below.

Firstly, the modified Hausdorff distance does not distinguish between a space its dense subset, therefore in order to have a well defined limit one has to require *a priori* for the limit be complete.

Secondly, if a pointed sequence of proper geodesic metric spaces $(X_n, \text{dist}_n, x_n)$ converges to a complete geodesic metric space (X, dist, x) in the modified Hausdorff distance, then the limit space X is proper. Indeed given a ball $B(x, R)$ in X , for every ϵ there exists an n such that $B(x, R)$ is at Hausdorff distance at most ϵ from the ball $B(x_n, R)$ in X_n . From this and the fact that all spaces X_n are proper it follows that for every sequence (y_n) in $B(x, R)$ and every ϵ there exists a subsequence of (y_n) of diameter $\leq \epsilon$. A diagonal argument and completeness of

X allow to conclude that (y_n) has a convergent subsequence, and therefore that $B(x, R)$ is compact.

In view of Theorem 10.46, for a proper geodesic metric space (X, dist) , the existence of a sequence of pointed metric spaces of the form $(X, \lambda_n \text{dist}, e_n)$ convergent in the modified Hausdorff metric, implies the existence of proper asymptotic cones. On the other hand, if X is, for instance, the hyperbolic plane or a non-elementary hyperbolic group, no asymptotic cone of X is proper, see Theorem 11.174. Therefore, in such a case, the sequence $(X, \frac{1}{n} \text{dist})$ has no subsequence convergent with respect to the modified Hausdorff metric.

10.7. Ultralimits of asymptotic cones are asymptotic cones

In this section we show that ultralimits of asymptotic cones are asymptotic cones, following [DS05b]. To this end, we first describe a construction of ultrafilters on Cartesian products that generalizes the standard notion of product of ultrafilters, as defined in [She78, Definition 3.2 in Chapter VI]. In what follows, we view ultrafilters as in Definition 10.19. Throughout the section, ω will denote an ultrafilter on a set I and $\mu = (\mu_i)_{i \in I}$ a family, indexed by I , of ultrafilters on a set J .

DEFINITION 10.75. We define a new ultrafilter $\omega\mu$ on $I \times J$ such that for every subset A in $I \times J$, $\omega\mu(A)$ is equal to the ω -measure of the set of all $i \in I$ such that $\mu_i(A \cap (\{i\} \times J)) = 1$.

LEMMA 10.76. $\omega\mu$ is an ultrafilter over $I \times J$.

PROOF. It suffices to prove that $\omega\mu$ is finitely additive and that it takes the zero value on finite sets.

We first prove that $\omega\mu$ is finitely additive, using the fact that ω and μ_i are finitely additive. Let A and B be two disjoint subsets of $I \times J$. Fix an arbitrary $i \in I$. The sets $A \cap (\{i\} \times J)$ and $B \cap (\{i\} \times J)$ are disjoint, hence

$$\mu_i((A \cup B) \cap (\{i\} \times J)) = \mu_i(A \cap (\{i\} \times J)) + \mu_i(B \cap (\{i\} \times J)).$$

The finite additivity of ω implies that

$$\omega\mu(A \sqcup B) = \omega\mu(A) + \omega\mu(B).$$

Also, given a finite subset A of $I \times J$, $\omega\mu(A) = 0$. Indeed, since the set of i 's for which $\mu_i(A \cap (\{i\} \times J)) = 1$ is empty, $\omega\mu(A) = 0$ by definition. \square

LEMMA 10.77 (double ultralimit of real numbers). For every doubly indexed family of real numbers α_{ij} , $i \in I, j \in J$ we have that

$$(10.4) \quad \omega\mu\text{-lim } \alpha_{ij} = \omega\text{-lim } (\mu_i\text{-lim}_j \alpha_{ij}),$$

where the second limit on the right hand side is taken with respect to $j \in J$.

PROOF. Let a be the limit $\omega\mu\text{-lim } \alpha_{ij}$. For every neighborhood U of a ,

$$\begin{aligned} \omega\mu \{(i, j) \mid \alpha_{ij} \in U\} &= 1 \Leftrightarrow \\ \omega \{i \in I \mid \mu_i \{j \mid \alpha_{ij} \in U\} = 1\} &= 1. \end{aligned}$$

This implies that

$$\omega \{i \in I \mid \mu_i\text{-lim}_j \alpha_{ij} \in \overline{U}\} = 1,$$

which, in turn, implies that

$$\omega\text{-lim}(\mu_i\text{-lim}_j \alpha_{ij}) \in \overline{U}.$$

This holds for every neighborhood U of $a \in \mathbb{R} \cup \{\pm\infty\}$. Therefore, we conclude that

$$\omega\text{-lim}(\mu_i\text{-lim} \alpha_{ij}) = a.$$

□

Lemma 10.77 implies a similar result for ultralimits of spaces.

PROPOSITION 10.78 (double ultralimit of spaces). *Let $(X_{ij}, \text{dist}_{ij})$ be a doubly indexed sequence of metric spaces, $(i, j) \in I \times J$, and let $e = (e_{ij})$ be a doubly indexed sequence of points $e_{ij} \in X_{ij}$. We denote by e_i the sequence $(e_{ij})_{j \in J}$.*

Then the map

$$(10.5) \quad \omega\mu\text{-lim}(x_{ij}) \mapsto \omega\text{-lim}(\mu_i\text{-lim} x_{ij}),$$

is an isometry from

$$\omega\mu\text{-lim}(X_{ij}, e_{ij})$$

onto

$$\omega\text{-lim}(\mu_i\text{-lim}(X_{ij}, e_{ij}), e'_i)$$

where, $e'_i = \mu_i\text{-lim} e_{ij}$.

COROLLARY 10.79 (ultralimits of asymptotic cones are asymptotic cones). *Let X be a metric space. Consider double indexed families of points $e = (e_{ij})_{(i,j) \in I \times J}$ in X and of positive real numbers $\lambda = (\lambda_{ij})_{(i,j) \in I \times J}$ such that*

$$\mu_i\text{-lim}_j \lambda_{ij} = 0$$

for every $i \in I$. Let $\text{Cone}_{\mu_i}(X, (e_{ij}), (\lambda_{ij}))$ be the corresponding asymptotic cone of X . The map

$$(10.6) \quad \omega\mu\text{-lim}(x_{ij}) \mapsto \omega\text{-lim}(\mu_i\text{-lim}(x_{ij})),$$

is an isometry from $\text{Cone}_{\omega\mu}(X, e, \lambda)$ onto

$$\omega\text{-lim}(\text{Cone}_{\mu_i}(X, (e_{ij}), (\lambda_{ij})), \mu_i\text{-lim} e_{ij}).$$

PROOF. The statement follows from Proposition 10.78. The only thing to be proved here is that

$$\omega\mu\text{-lim} \lambda_{ij} = 0$$

Let $\varepsilon > 0$. For every $i \in I$ we have that

$$\mu_i\text{-lim} \lambda_{ij} = 0,$$

whence,

$$\mu_i \{j \in I \mid \lambda_{ij} < \varepsilon\} = 1.$$

It follows that

$$\{i \in I \mid \mu_i \{j \in I \mid \lambda_{ij} < \varepsilon\} = 1\} = I,$$

therefore, the ω -measure of this set is 1. We conclude that

$$\omega\mu \{(i, j) \in I \times J \mid \lambda_{ij} < \varepsilon\} = 1. \quad \square$$

COROLLARY 10.80. *Let X be a metric space. The collection of all asymptotic cones of X is stable with respect to rescaling, ultralimits and taking asymptotic cones.*

PROOF. It is an immediate consequence of Corollary 10.79 and Remark 10.67. \square

COROLLARY 10.81. *Let X, Y be metric spaces such that all asymptotic cones of X are isometric to Y . Then all asymptotic cones of Y are isometric to Y .*

This, in particular, implies that the following are *examples of metric spaces isometric to all their asymptotic cones*.

- EXAMPLES 10.82. (1) The 2^{\aleph_0} -universal real tree T_C , according to Theorem 11.174.
 (2) A non-discrete Euclidean building that is the asymptotic cone of $SL(n, \mathbb{R})$, $n \geq 3$, under the Continuum Hypothesis, according to [KSTT05] and [KL98b].
 (3) A graded nilpotent Lie group with a Carnot–Caratheodory metric, according to Theorem 16.28 of P. Pansu.

10.8. Asymptotic cones and quasiisometries

The following simple lemma shows why asymptotic cones are useful in studying quasiisometries, since they become bi-Lipschitz maps of asymptotic cones, and the latter maps are much easier to handle. It is a direct generalization of Lemma 10.48 on functoriality of ultralimits with respect to isometries.

LEMMA 10.83. *Let $(X, e_i), (X', e'_i)$ be pointed metric spaces, and let $\lambda = (\lambda_i)$ be a scaling family. Define the asymptotic cones*

$$X_\omega = \text{Cone}_\omega(X, (e_i), \lambda), \quad X'_\omega = \text{Cone}_\omega(X', (e'_i), \lambda).$$

Then the following holds for every family of (L, A) -coarse Lipschitz maps $f_i : X \rightarrow X'$, satisfying

$$\omega\text{-}\lim d(f_i(e_i), e'_i) < \infty :$$

1. *The ultralimit $f_\omega : X_\omega \rightarrow X'_\omega$ of the family (f_i) ,*

$$f_\omega((x_i)) := (f_i(x_i)),$$

is L -Lipschitz.

2. *If f_i is an (L, A) -quasiisometric embedding, then f_ω is an L -bi-Lipschitz embedding.*

3. *The correspondence $\Phi_\omega : (f_i) \mapsto f_\omega$ is functorial:*

$$\Phi_\omega : (g_i \circ f_i) \mapsto g_\omega \circ f_\omega.$$

4. *If $X = X'$ and f_i 's have uniformly bounded displacement, i.e. for ω -all i ,*

$$\text{dist}(f_i(x), x) \leq A, \quad \forall x \in X,$$

then $f_\omega = \text{Id}_X$.

5. *If each f_i is an (L, A) -quasiisometry, then f_ω is an L -bi-Lipschitz homeomorphism.*

PROOF. 1. We have the inequalities:

$$\frac{1}{\lambda_i} \text{dist}(f_i(x_i), f_i(y_i)) \leq L \frac{1}{\lambda_i} \text{dist}(x_i, y_i) + \frac{A}{\lambda_i}.$$

Passing to the ω -limit, we obtain

$$\text{dist}_\omega(f_\omega(x_\omega), f_\omega(y_\omega)) \leq L \text{dist}_\omega(x_\omega, y_\omega).$$

where $x_\omega = (x_i), y_\omega = (y_i)$. Thus, f_ω is L -Lipschitz.

2. In this case we also have the inequalities

$$L^{-1} \frac{1}{\lambda_i} \text{dist}(x_i, y_i) - \frac{A}{\lambda_i} \leq \frac{1}{\lambda_i} \text{dist}(f_i(x_i), f_i(y_i)),$$

which, after passing to the ultralimit, become

$$L^{-1} \text{dist}_\omega(x_\omega, y_\omega) \leq \text{dist}_\omega(f_\omega(x_\omega), f_\omega(y_\omega)).$$

Thus, f_ω is an L -bi-Lipschitz embedding.

Parts 3 and 4 are clear. Part 5 follows from 1, 3 and 4. \square

One may ask if a converse to this lemma is true, for instance: Does the existence of a (coarse Lipschitz) map between metric spaces that induces bi-Lipschitz maps between asymptotic cones imply quasiisometry of the original metric spaces? We say that two spaces are *asymptotically bi-Lipschitz* if the latter holds. (This notion is introduced in [dC11].) See Remark 16.29 for an example of asymptotically bi-Lipschitz spaces which are not quasiisometric to each other.

Here is an example of application of asymptotic cones to the study of quasiisometries.

LEMMA 10.84. *Suppose that $X = \mathbb{E}^n$ or \mathbb{R}_+ and $f : X \rightarrow X$ is an (L, A) -quasiisometric embedding. Then f is a quasiisometry, furthermore, $\mathcal{N}_C(f(X)) = X$, for some $C = C(L, A)$.*

PROOF. We will give a proof in the case of \mathbb{E}^n as the other case is analogous. Suppose that the assertion is false, i.e. there is a sequence of (L, A) -quasiisometric embeddings $f_i : \mathbb{E}^n \rightarrow \mathbb{E}^n$, sequence of real numbers r_i diverging to infinity and points $y_i \in \mathbb{E}^n$ such that $\text{dist}(y_i, f(\mathbb{E}^n)) = r_i$. Let $x_i \in \mathbb{E}^n$ be a point such that $\text{dist}(f(x_i), y_i) \leq r_i + 1$. Using x_i, y_i as base-points on the domain and range for f_i , rescale the metrics on the domain and the range by $\lambda_i = \frac{1}{r_i}$ and take the corresponding ultralimits. In the limit we get a bi-Lipschitz embedding

$$f_\omega : \mathbb{E}^n \rightarrow \mathbb{E}^n,$$

whose image misses the point $y_\omega \in \mathbb{E}^n$. However each bi-Lipschitz embedding of Euclidean spaces is necessarily proper, therefore, by the invariance of domain theorem, the image of f_ω is both closed and open. Contradiction. \square

REMARK 10.85. Alternatively, one can prove the above lemma (without using ultralimits) as follows: Approximate f by a continuous mapping g . Then, since g is proper, it has to be onto.

COROLLARY 10.86. *\mathbb{E}^n is quasiisometric to \mathbb{E}^m if and only if $n = m$.*

On the other hand, one cannot use ultralimits (at least directly) to prove that hyperbolic spaces of different dimensions are not quasiisometric to each other: All their ultralimits are isometric to the same universal real tree.

10.9. Assouad-type theorems

In this section we prove Assouad's Theorem (and some of its generalizations) stating that the problem of embedding a space into one of the spaces from a specific collection \mathcal{C} can be decided by looking at finite subspaces, provided that both the

type of embedding considered and the collection \mathcal{C} are stable with respect to ultralimits. The arguments in this section were inspired by arguments in [BDCK66, Troisième partie, §2, pp. 252].

DEFINITION 10.87. *Consider a class \mathcal{C} of metric spaces. We say that \mathcal{C} is stable with respect to ultralimits if for every set of indices I , every nonprincipal ultrafilter ω on I , every collection $(X_i, \text{dist}_i)_{i \in I}$ of metric spaces in \mathcal{C} and every sets of base-points $(e_i)_{i \in I}$ with $e_i \in X_i$, the ultralimit $\omega\text{-lim}(X_i, e_i, \text{dist}_i)$ is isometric to a metric space in \mathcal{C} .*

We say that \mathcal{C} is stable with respect to rescaled ultralimits if for every choice of I , ω , $(X_i, \text{dist}_i)_{i \in I}$ and $(e_i)_{i \in I}$ as above, and, moreover, every indexed set of positive real numbers $(\lambda_i)_{i \in I}$, the ultralimit of rescaled spaces $\omega\text{-lim}(X_i, e_i, \lambda_i \text{dist}_i)$ is isometric to a metric space in \mathcal{C} .

Note that in this definition we are not making any assumptions about the limits $\omega\text{-lim } \lambda_i$; in particular, they are allowed to be zero and ∞ .

EXAMPLE 10.88. The class of $CAT(0)$ spaces is stable with respect to rescaled ultralimits.

Since in a normed vector space V the scaling $x \mapsto \lambda x$, $\lambda \in \mathbb{R}_+$, scales the metric by λ , the metric space $(V, \lambda \text{dist})$ is isometric to (V, dist) , where $\text{dist}(u, v) = \|u - v\|$. Therefore, taking rescaled ultralimits of normed spaces is the same as taking their ultralimits. Since ultralimits of families of complete metric spaces are complete, we conclude that the class of Banach spaces is stable under ultralimits. In section 19.1 we will prove that the classes of Hilbert spaces and of abstract L^p -spaces are stable under ultralimits.

In what follows we consider two continuous functions $\rho_{\pm} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ as in Definition 8.25 of coarse embeddings, i.e. such that $\rho_-(x) \leq \rho_+(x)$ for every $x \in \mathbb{R}_+$, and such that both functions have limit $+\infty$ at $+\infty$.

THEOREM 10.89. *Let \mathcal{C} be a collection of metric spaces stable with respect to ultralimits.*

A metric space (X, dist) has a (ρ_-, ρ_+) -embedding into some space Y in \mathcal{C} if and only if every finite subset of X has a (ρ_-, ρ_+) -embedding into some space in \mathcal{C} .

PROOF. The direct implication is obvious, we will prove the converse. Let (X, dist) be a metric space such that for every finite subset F in X endowed with the induced metric, there exists a (ρ_-, ρ_+) -embedding $\varphi_F : F \rightarrow Y_F$, where (Y_F, dist_F) is a metric space in \mathcal{C} . We fix a base-point e in X . In every finite subset of X we fix a base-point e_F , such that $e_F = e$ whenever $e \in F$, and we denote $\varphi_F(e_F)$ by y_F .

Let I be the collection of all finite subsets of X . Let \mathcal{B} be the collection of subsets of I of the form $I_F = \{F' \in I \mid F \subseteq F'\}$, where F is a fixed element of I . Then \mathcal{B} is the base of a filter. Indeed:

1. $I_{F_1} \cap I_{F_2} = I_{F_1 \cup F_2}$.
2. For every F , I_F contains F and, hence, is non-empty.
3. $I = I_{\emptyset} \in \mathcal{B}$.

Therefore, it follows from Exercise 10.9 and the Ultrafilter Lemma 10.18 that there exists an ultrafilter ω on I such that for every finite subset of $F \subset X$, $\omega(I_F) = 1$.

Consider the ultralimits $X_\omega = \omega\text{-lim}(X, e, \text{dist})$ and $Y_\omega = \omega\text{-lim}(Y_F, y_F, \text{dist}_F)$. By hypothesis, the space $(Y_\omega, y_\omega, \text{dist}_\omega)$ belongs to the class \mathcal{C} .

We have the diagonal isometric embedding $\iota : X \rightarrow X_\omega$, $\iota(x) = x_\omega$. Set $X_\omega^0 := \iota(X)$. We define a map

$$\varphi_\omega : X_\omega^0 \rightarrow Y_\omega$$

by $\varphi_\omega(x_\omega) := \omega\text{-lim } z_F$, where $z_F = \varphi_F(x)$ whenever $x \in F$, and $z_F = y_F$ when $x \notin F$.

Let us check that φ_ω is a (ρ_-, ρ_+) -embedding. Consider two points x_ω, x'_ω in X_ω^0 . Recall that $\omega(I_{\{x, x'\}}) = 1$ by the definition of ω . Therefore, if $\varphi_\omega(x_\omega) = \omega\text{-lim } z_F$ and $\varphi_\omega(x'_\omega) = \omega\text{-lim } z'_F$, then ω -almost surely $z_F = \varphi_F(x)$ and $z'_F = \varphi_F(x')$. Hence, ω -almost surely

$$\rho_-(\text{dist}(x, x')) \leq \text{dist}_F(z_F, z'_F) \leq \rho_+(\text{dist}(x, x')).$$

By passing to the ultralimit we obtain

$$\rho_-(\text{dist}(x, x')) \leq \text{dist}_\omega(\varphi_\omega(x_\omega), \varphi_\omega(x'_\omega)) \leq \rho_+(\text{dist}(x, x')).$$

□

The following result first appeared in [BDCK66, Troisième partie, §2, pp. 252].

COROLLARY 10.90. *Let p be a real number in $[1, \infty)$. A metric space (X, dist) has a (ρ_-, ρ_+) -embedding into an L^p -space if and only if every finite subset of X has such a (ρ_-, ρ_+) -embedding.*

COROLLARY 10.91 (Assouad’s Theorem [WW75], Corollary 5.6). *Let p be a real number in $[1, \infty)$. A metric space (X, dist) has an isometric embedding into an L^p -space if and only if every finite subset of X has such an isometric embedding.*

Note that the same statement holds if one replaces “isometry” by “ (L, C) -quasi-isometry”, with fixed $L \geq 1$ and $C \geq 0$.

It may now be the right place to explain the importance of coarse embeddings. This increased when G. Yu, following a suggestion of Gromov [Gro93], proved in [Yu00] that every discrete metric space which embeds coarsely into a Hilbert space satisfies the Coarse Baum-Connes Conjecture. In particular, if the considered space is a finitely generated group with a word metric, its coarse embeddability into a Hilbert space implies the Novikov Conjecture. This latter result has been later extended to groups with a coarse embedding into special kinds of Banach spaces by Kasparov and Yu [KY12].

When introducing the notion of coarse embedding, Gromov asked if every separable metric space coarsely embeds in a Hilbert space [Gro93, p.218]. Following a counter-example to this initial question due to Dranishnikov, Gong, Lafforgue and Yu [DGLY02], M. Gromov noted that an obstruction to non-embeddability might be the presence of graphs with expanding properties [Gro00]. Later on, he constructed an infinite finitely generated group with a Cayley graph in which a family of expanders are coarsely embedded [Gro03] (see also [AD]). Gromov’s construction has been improved recently by D. Osajda, who constructed finitely generated groups with a family of expanders isometrically embedded into one of their Cayley graphs [Osa14].

Gromov–hyperbolic spaces and groups

The goal of this chapter is to review basic properties of δ -hyperbolic spaces and word-hyperbolic groups, which are far-reaching generalizations of the real-hyperbolic space \mathbb{H}^n and of groups acting geometrically on \mathbb{H}^n . The advantage of δ -hyperbolicity is that it can be defined in the context of arbitrary metric spaces which need not even be geodesic. These spaces were introduced in the seminal essay by Mikhael Gromov on hyperbolic groups [Gro87], although ideas of *combinatorial curvature* and (in retrospect) hyperbolic properties of finitely generated groups are much older. These ideas go back to work of Max Dehn (Dehn algorithm for the word problem in hyperbolic surface groups), Martin Greendlinger (small cancellation theory), Alexandr Ol’shanskiĭ (who used what we now call *relative hyperbolicity* in order to construct finitely generated groups with exotic properties) and many others.

11.1. Hyperbolicity according to Rips

We begin our discussion of δ -hyperbolic spaces with the notion of hyperbolicity in the context of geodesic metric spaces, which (according to Gromov) is due to Ilya (Eliyahu) Rips. This definitions will be then applied to Cayley graphs of groups, leading to the concept of *hyperbolic groups* discussed later in this chapter. Rips notion of hyperbolicity is based on the thinness properties of hyperbolic triangles which are established in Section 4.10.

Let (X, d) be a geodesic metric space. As in Section 4.5, a geodesic triangle T in X is a concatenation of three geodesic segments τ_1, τ_2, τ_3 connecting the points A_1, A_2, A_3 (vertices of T) in the natural cyclic order. Unlike the real-hyperbolic space, we no longer have uniqueness of geodesics, thus T is not (in general) determined by its vertices. We define a measure of the thinness of T similar to the one in Section 4.10 of Chapter 4.

DEFINITION 11.1. The *thinness radius* of the geodesic triangle T is the number

$$\delta(T) := \max_{j=1,2,3} \left(\sup_{p \in \tau_j} d(p, \tau_{j+1} \cup \tau_{j+2}) \right),$$

A triangle T is called *δ -thin* if $\delta(T) \leq \delta$.

DEFINITION 11.2 (Rips’ definition of hyperbolicity). A geodesic hyperbolic space X is called *δ -hyperbolic* (in the sense of Rips) if every geodesic triangle T in X is δ -thin. The infimum of all δ ’s for which X is δ -hyperbolic is called the *hyperbolicity constant* of X .

A space X which is δ -hyperbolic for some $\delta < \infty$ is called *Rips-hyperbolic*. In what follows, we will refer to δ -hyperbolic spaces in the sense of Rips simply as being δ -hyperbolic.

Below are several simple but important geometric features of δ -hyperbolic spaces.

First of all, note that general Rips-hyperbolic metric spaces X are by no means uniquely geodesic; the notation xy used in what follows will mean that xy is *some* geodesic connecting x to y . The next lemma shows that geodesics in X between the given pair of points are “almost unique” and justifies, to some extent, the abuse of notation that we are committing.

LEMMA 11.3 (Thin bigon property). *If X is δ -hyperbolic, then all geodesics xy, zx with $d(y, z) \leq D$ are at Hausdorff distance at most $D + \delta$ from each other. In particular, if α, β are geodesics connecting points $x, y \in X$, then $\text{dist}_{\text{Haus}}(\alpha, \beta) \leq \delta$.*

PROOF. Every point p on xy is, either at distance at most δ from xz , or at distance at most δ from yz ; in the latter case p is at distance at most $D + \delta$ from xz . \square

Lemma 11.4 below is the *fellow-traveling property of hyperbolic geodesics*, which sharpens the conclusion of Lemma 11.3.

LEMMA 11.4 (Fellow-traveling property). *Let $\alpha(t), \beta(t)$ be geodesics in a δ -hyperbolic space X , such that $\alpha(0) = \beta(0) = o$ and $d(\alpha(t_0), \beta(t_0)) \leq D$ for some $t_0 \geq 0$. Then for all $t \in [0, t_0]$,*

$$d(\alpha(t), \beta(t)) \leq 2(D + \delta).$$

PROOF. By the previous lemma, for every $t \in [0, t_0]$ there exists $s \in [0, t_0]$ so that

$$d(\beta(t), \alpha(s)) \leq c = \delta + D.$$

Applying the triangle inequality, we see that

$$|t - s| \leq c,$$

hence, $d(\alpha(t), \beta(t)) \leq 2c = 2(\delta + D)$. \square

The notion of thin triangles generalizes naturally to the concept of *thin polygons*. A *geodesic n -gon* P in a metric space X is a concatenation of geodesic segments $\sigma_i, i = 1, \dots, n$, connecting points $p_i, i = 1, \dots, n$, in the natural cyclic order. The points p_i are called the *vertices* of the polygon P and the geodesics σ_i are called the *sides* or the *edges* of P . A polygon P is called η -thin if every side of P is contained in the η -neighborhood of the union of the other sides.

EXERCISE 11.5. Suppose that X is a δ -hyperbolic metric space. Show that every n -gon in X is $\delta(n - 2)$ -thin. Hint: Triangulate an n -gon P by $n - 3$ diagonals emanating from a single vertex. Now, use δ -thinness of triangles in X inductively.

We next improve the estimate provided by this exercise.

LEMMA 11.6 (thin polygons). *If X is δ -hyperbolic then, for $n \geq 2$, every geodesic $n + 1$ -gon in X is η_n -thin with*

$$\eta_n = \delta \lceil \log_2 n \rceil.$$

PROOF. First of all, it suffices to consider the case $n = 2^k$, for otherwise we add to the polygon edges of zero length until the number of sides reaches $2^k + 1$. We prove the estimate on thinness of $(2^k + 1)$ -gons by induction on k . For $k = 1$ the statement amounts to the δ -thinness of triangles. Suppose that k is at least 2 and

that the thinness estimate holds for all $(2^{k-1} + 1)$ -gons. Take a geodesic $n + 1$ -gon P with the sides $\tau_i = p_i p_{i+1}$, $i = 0, \dots, n - 1$, $\tau = p_n p_0$. and consider the edge $\tau = \tau_n$ of P .

We subdivide P into three pieces by introducing the diagonals $p_0 p_m$ and $p_m p_n$, where $m = 2^{k-1}$. These pieces are two $2^m + 1$ -gons and one triangle:

$$P' = p_0 p_1 \dots p_m, \quad P'' = p_m p_{m+1} \dots p_n, \quad T = p_0 p_m p_n.$$

By the induction hypothesis, the polygons P', P'' are $\delta(k-1)$ -thin, while the triangle T is δ -thin. Therefore, τ is contained in the $\delta(k-1+1) = \delta k$ -neighborhood of the union of the other sides of P . \square

We now give some examples of Rips-hyperbolic metric spaces.

EXAMPLE 11.7. (1) Proposition 4.66 implies that \mathbb{H}^n is δ -hyperbolic for $\delta = \arccos(\sqrt{2})$.

- (2) Suppose that (X, d) is δ -hyperbolic and $a > 0$. Then the metric space $(X, a \cdot d)$ is $a\delta$ -hyperbolic. Indeed, distances in $(X, a \cdot d)$ are obtained from distances in (X, d) by multiplication by a . Therefore, the same is true for distances between the edges of geodesic triangles.
- (3) Let X_κ is the model surface of curvature $\kappa < 0$ as in Section 3.11.1. Then X_κ is δ -hyperbolic for

$$\delta_\kappa = |\kappa|^{-1/2} \arccos(\sqrt{2}).$$

Indeed, the Riemannian metric on X_κ is obtained by multiplying the Riemannian metric on \mathbb{H}^2 by $|\kappa|^{-1}$. This has the effect of multiplying all distances in \mathbb{H}^2 by $|\kappa|^{-1/2}$. Hence, if d is the distance function on \mathbb{H}^2 then $|\kappa|^{-1/2}d$ is the distance function on X_κ .

- (4) Suppose that X is a $CAT(\kappa)$ -space where $\kappa < 0$, see Section 3.11.1. Then X is δ_κ -hyperbolic. Indeed, all triangles in X are thinner than triangles in X_κ . Therefore, given a geodesic triangle T with the edges $\tau_i, i = 1, 2, 3$, and a point $p_1 \in \tau_1$ we take the comparison triangle $\tilde{T} \subset X_\kappa$ and the comparison point $\tilde{p}_1 \in \tilde{\tau}_1 \subset \tilde{T}$. Since \tilde{T} is δ_κ -thin, there exists a point $\tilde{p}_i \in \tilde{\tau}_i, i = 2$ or $i = 3$, such that $d(\tilde{p}_1, \tilde{p}_i) \leq \delta_\kappa$. Let $p_i \in \tau_i$ be the comparison point of \tilde{p}_i . By the comparison inequality,

$$d(p_1, p_i) \leq d(\tilde{p}_1, \tilde{p}_i) \leq \delta_\kappa,$$

and, hence, T is δ_κ -thin. In particular, if X is a simply-connected complete Riemannian manifold of sectional curvature $\leq \kappa < 0$, then X is δ_κ -hyperbolic.

- (5) Let X be a simplicial tree, and d be a path-metric on X . Then, by the Exercise 3.59, X is $CAT(-\infty)$. Thus, by (4), X is δ_κ -hyperbolic for every $\delta_\kappa = |\kappa|^{-1/2} \arccos(\sqrt{2})$. Since

$$\inf_{\kappa} \delta_\kappa = 0,$$

it follows that X is 0-hyperbolic. Of course, this fact one can easily see directly by observing that every triangle in X is a tripod.

- (6) Every geodesic metric space of diameter $\leq \delta < \infty$ is δ -hyperbolic.

EXERCISE 11.8. Let X be the circle of radius R in \mathbb{R}^2 with the induced path-metric d . Thus, (X, d) has diameter πR . Show that X is $\pi R/2$ -hyperbolic and is not δ -hyperbolic for any $\delta < \pi R/2$.

Not every geodesic metric space is hyperbolic:

EXAMPLE 11.9. For instance, let us verify that the Euclidean plane \mathbb{E}^2 is not δ -hyperbolic for any δ . Pick a nondegenerate triangle $T \subset \mathbb{R}^2$. Then $\delta(T) = k > 0$ for some k . Therefore, if we scale T by a positive constant c , then $\delta(cT) = ck$. Sending $c \rightarrow \infty$, shows that \mathbb{E}^2 is not δ -hyperbolic for any $\delta > 0$. More generally, if a metric space X contains an isometrically embedded copy of \mathbb{E}^2 , then X is not hyperbolic.

Here is an example of a metric space which is not hyperbolic, but does not contain a quasiisometrically embedded copy of \mathbb{E}^2 either. Consider the wedge X of countably many circles C_i each given with a path-metric of the overall length $2\pi i$, $i \in \mathbb{N}$. We equip X with the path-metric such that each C_i is isometrically embedded. Exercise 11.8 shows that X is not hyperbolic.

EXERCISE 11.10. Show that X contains no quasiisometrically embedded copy of \mathbb{E}^2 . Hint: Use coarse topology.

More interesting examples of non-hyperbolic spaces containing no quasi-isometrically embedded copies of \mathbb{E}^2 are given by various solvable groups, e.g. the solvable Lie group Sol_3 and the Cayley graph of the Baumslag–Solitar group $BS(n, 1)$, see [Bur99].

Below we describe briefly another measure of thinness of triangles which can be used as an alternative definition of Rips–hyperbolicity. It is also related to the minimal size of triangles, described in Definition 9.101, consequently it is related to the filling area of the triangle *via* a Besikovitch–type inequality as described in Proposition 9.104.

DEFINITION 11.11. For a geodesic triangle $T \subset X$ with the sides τ_1, τ_2, τ_3 , define the *inradius* of T to be

$$\Delta(T) := \inf_{x \in X} \max_{i=1,2,3} d(x, \tau_i).$$

In the case of the real-hyperbolic plane, as we saw in Lemma 4.65, this definition coincides with the radius of the largest circle inscribed in T . Clearly, $\Delta(T) \leq \delta(T)$ and

$$\Delta(T) \leq \text{minsize}(T)$$

We next show that

$$(11.1) \quad \text{minsize}(T) \leq 2\delta(T).$$

Indeed, suppose that $T = T(x_1, x_2, x_3)$ and the side τ_i of T connects x_i to x_{i+1} (i is taken modulo 3). Let a denote the length of τ_1 .

Define the continuous function

$$f(t) = d(\tau_1(t), \tau_2) - d(\tau_1(t), \tau_3),$$

it takes the value $d(x_1, x_2) = a$ at $t = 0$ and the value $-a$ at $t = a$. Therefore, by the intermediate value theorem, there exists $t_1 \in [0, a]$ such that

$$d(\tau_1(t_1), \tau_2) = d(\tau_1(t_1), \tau_3) \leq \delta(T).$$

Taking $p_1 = \tau_1(t_1)$ and $p_i \in \tau_i, i = 2, 3$ to be the points closest to p_1 , we get

$$d(p_1, p_2) \leq \delta, d(p_1, p_3) \leq \delta(T),$$

hence,

$$\text{minsize}(T) \leq 2\delta(T).$$

Hyperbolicity and combings. One might criticize Rips' definition of hyperbolicity by observing that it is difficult to verify that the given geodesic metric space is hyperbolic, as Rips' definition requires one to identify geodesic segments in the space. The notion of *thin bicombing* below can be used to circumvent this problem; for instance, it was used successfully to verify hyperbolicity of the *curve complex* by Bowditch in [Bow06b].

Let Γ be a connected graph with the standard metric. A *combing* of Γ is a map c which associates to every pair of vertices u, v in Γ an edge-path p_{uv} in Γ connecting u to v . A combing c is called a *bicombing* if p_{uv} equals p_{vu} run in the reverse. A combing c is said to be *consistent* if for every vertices u, v in Γ and every integer subinterval $[t, s]$ in the domain of p_{uv} , the restriction of p_{uv} to $[t, s]$ equals $p_{u',v'}$, where $u' = p_{uv}(t), v' = p_{uv}(s)$. A combing is called *proper* if there is a constant C such that whenever $d(u, v) \leq 1$, we also have

$$\text{length}(p_{uv}) \leq C.$$

A bicombing is called *thin* if it is consistent, proper and there exists a constant δ such that for every triple of vertices u, v, w in Γ there exists a vertex x within distance $\leq \delta$ from the images of all three paths

$$p_{uv}, \quad p_{vw}, \quad p_{wu}.$$

More generally, one defines the notion of a thin bicombing for a general metric space X by assuming that the paths p_{uv} connecting points of X are 1-Lipschitz and repeating the rest of the definition. It is now immediate that every Rips-hyperbolic metric space admits a thin bicombing, namely, the one given by geodesics. Conversely:

THEOREM 11.12. (*U. Hamenstädt, [Ham07, Proposition 3.5]*) *If a geodesic metric space X (or a connected graph with the standard metric) admits a thin bicombing, then X is Rips-hyperbolic. Furthermore, the paths p_{uv} are $(L, 0)$ -quasigeodesic for some L .*

11.2. Geometry and topology of real trees

In this section we consider in more detail a special class of hyperbolic spaces, the *real trees*. In view of Definition 3.60, a geodesic metric space is a real tree if and only if it is 0-hyperbolic.

LEMMA 11.13. *If X is a real tree then any two points in X are connected by a unique topological arc in X .*

PROOF. Let $D = d(x, y)$. Consider a topological arc, i.e. a continuous injective map $\alpha : [0, 1] \rightarrow X, x = \alpha(0), y = \alpha(1)$. Let $\alpha^* = xy, \alpha^* : [0, D] \rightarrow X$ be a geodesic connecting x to y . (This geodesic is unique by 0-hyperbolicity of X .) We claim that the image of α contains the image of α^* . Indeed, we can approximate α by piecewise-geodesic (nonembedded!) arcs

$$\alpha_n = x_0x_1 \cup \dots \cup x_{n-1}x_n, \quad x_0 = x, \quad x_n = y.$$

Since the $n + 1$ -gon P in X , which is the concatenation of α_n with yx is 0-thin, $\alpha^* \subset \alpha_n$, cf. Lemma 11.6. Therefore, the image of α also contains the image

of α^* . Consider the continuous map $(\alpha^*)^{-1} \circ \alpha : [0, D] \rightarrow [0, D]$. Applying the intermediate value theorem to this function, we see that the images of α and α^* are equal. \square

EXERCISE 11.14. Prove the converse to this lemma: If X is a path-metric space where any two points are connected by a unique topological arc, then X is isometric to a real tree. In particular, if X is a path-metric space homeomorphic to a tree, then X is isometric to a tree.

We refer the reader to [Bow91] for further discussion of characterizations of metric trees.

DEFINITION 11.15. Let T be a real tree and p be a point in T . The *space of directions* at p , denoted Σ_p , is defined as the space of germs of geodesics in T emanating from p , i.e. the quotient $\Sigma_p := \mathfrak{R}_p / \sim$, where

$$\mathfrak{R}_p = \{r : [0, a) \rightarrow T \mid a > 0, r \text{ is isometric}, r(0) = p\}$$

and

$$r_1 \sim r_2 \iff \exists \varepsilon > 0 \text{ such that } r_1|_{[0, \varepsilon)} \equiv r_2|_{[0, \varepsilon)}.$$

By Lemma 11.13, for every topological arc $c : [a, b] \rightarrow T$ in a tree, the image of c coincides with the geodesic segment $c(a)c(b)$. It follows that we may also define Σ_p as the space of germs of topological arcs instead of geodesic arcs.

DEFINITION 11.16. Define the *valence* $val(p)$ of a point p in a real tree T to be the cardinality of the set Σ_p . A *branch-point* of T is a point p of valence ≥ 3 . The *valence* of T is the supremum of valences of points in T .

EXERCISE 11.17. Show that $val(p)$ equals the number of connected components of $T \setminus \{p\}$.

DEFINITION 11.18. A real tree T is called α -*universal* if every real tree with valence at most α can be isometrically embedded into T .

We refer the reader to [MNO92] for a study of universal trees. In particular, the following holds:

THEOREM 11.19 ([MNO92]). *For every cardinal number $\alpha > 2$ there exists an α -universal tree, and it is unique up to isometry.*

11.3. Gromov hyperbolicity

One drawback of the Rips definition of hyperbolicity is that it uses geodesics. Below is an alternative definition of hyperbolicity, due to Gromov, where one needs to verify certain inequalities only for quadruples of points in a metric space (which need not be geodesic). Gromov's definition is less intuitive than the one of Rips, but, as we will see, it is more suitable in certain situations.

Let (X, dist) be a metric space (*which is no longer required to be geodesic*). Pick a base-point $p \in X$. For each $x \in X$ set $|x|_p := \text{dist}(x, p)$ and define the *Gromov product*

$$(x, y)_p := \frac{1}{2} (|x|_p + |y|_p - \text{dist}(x, y)).$$

Note that the triangle inequality immediately implies that $(x, y)_p \geq 0$ for all x, y, p ; the Gromov product measures how far the triangle inequality for the points x, y, p is from being an equality.

REMARK 11.20. The Gromov product is a generalization of the inner product in vector spaces with p serving as the origin. For instance, suppose that $X = \mathbb{R}^n$ with the usual inner product, $p = 0$ and $|v|_p := \|v\|$ for $v \in \mathbb{R}^n$. Then

$$\frac{1}{2} (|x|_p^2 + |y|_p^2 - \|x - y\|^2) = x \cdot y.$$

EXERCISE 11.21. Suppose that X is a metric tree. Then $(x, y)_p$ is the distance $\text{dist}(p, \gamma)$ from p to the geodesic segment $\gamma = xy$.

For general metric spaces general, a direct calculation using triangle inequalities shows that all points $p, x, y, z \in X$ satisfy the inequality

$$(p, x)_z + (p, y)_z \leq |z|_p - (x, y)_p$$

with the equality

$$(11.2) \quad (p, x)_z + (p, y)_z = |z|_p - (x, y)_p.$$

if and only $d(x, z) + d(z, y) = d(x, y)$. Thus, for every $z \in \gamma = xy$,

$$(x, y)_p = d(z, p) - (p, x)_z - (p, y)_z \leq d(z, p).$$

In particular, $(x, y)_p \leq \text{dist}(p, \gamma)$.

LEMMA 11.22. Suppose that X is δ -hyperbolic in the sense of Rips. Then the Gromov product in X is “comparable” to $\text{dist}(p, xy)$: For every $x, y, p \in X$ and geodesic xy ,

$$(x, y)_p \leq \text{dist}(p, xy) \leq (x, y)_p + 2\delta.$$

PROOF. The inequality $(x, y)_p \leq \text{dist}(p, xy)$ was proved above; thus, we have to establish the other inequality. Note that since the triangle $T(p, x, y)$ is δ -thin, for each point $z \in xy$ we have

$$\min\{(x, p)_z, (y, p)_z\} \leq \min\{\text{dist}(z, px), \text{dist}(z, py)\} \leq \delta.$$

By continuity of the distance function, there exists a point $z \in xy$ such that $(x, p)_z, (y, p)_z \leq \delta$. By applying the equality (11.2) we get:

$$|z|_p - (x, y)_p = (p, x)_z + (p, y)_z \leq 2\delta.$$

Since $|z|_p \leq \text{dist}(p, xy)$, we conclude that $\text{dist}(p, xy) \leq (x, y)_p + 2\delta$. □

For each pointed metric space (X, p) we define its *Gromov-hyperbolicity constant* $\delta_p = \delta_p(X) \in [0, \infty]$ as

$$\delta_p := \sup\{\min((x, z)_p, (y, z)_p) - (x, y)_p\},$$

where the supremum is taken over all triples of points $x, y, z \in X$.

EXERCISE 11.23. If $\delta_p \leq \delta$ then $\delta_q \leq 2\delta$ for all $q \in X$.

DEFINITION 11.24. A metric space X is said to be δ -hyperbolic in the sense of Gromov, if $\delta_p \leq \delta < \infty$ for all $p \in X$. In other words, for every quadruple $x, y, z, p \in X$, we have

$$(x, y)_p \geq \min((x, z)_p, (y, z)_p) - \delta.$$

EXERCISE 11.25. The real line with the usual metric is 0-hyperbolic in the sense of Gromov.

EXERCISE 11.26. Each δ -hyperbolic space in the sense of Gromov satisfies

$$(x, u)_p \geq \min\{(x, y)_p, (x, z)_p, (z, u)_p\} - 2\delta$$

for all $x, y, z, u, p \in X$.

Computing Gromov–hyperbolicity constant for a given metric space is, typically, not an easy task. We will see that all real trees are 0-hyperbolic in Gromov’s sense. It was recently proven by Nica and Spakula [NŠ16] that the Gromov–hyperbolicity constant for the hyperbolic plane \mathbb{H}^2 is $\log(2)$.

We next compare the two notions of hyperbolicity introduced so far.

LEMMA 11.27. *If a metric space X is δ -hyperbolic in the sense of Rips, then it is 3δ -hyperbolic in the sense of Gromov.*

PROOF. Consider points $x, y, z, p \in X$ and the geodesic triangle $T(x, y, z) \subset X$ with vertices x, y, z . Let $m \in xy$ be the point nearest to p . Then, since the triangle $T(x, y, z)$ is δ -thin, there exists a point $n \in xz \cup yz$ such that $\text{dist}(n, m) \leq \delta$. Assume that $n \in yz$. Then, by Lemma 11.22,

$$(y, z)_p \leq \text{dist}(p, yz) \leq \text{dist}(p, xy) + \delta.$$

On the other hand, by the same Lemma 11.22,

$$\text{dist}(p, xy) \leq (x, y)_p - 2\delta.$$

Combining these two inequalities, we obtain

$$(y, z)_p \leq (x, y)_p - 3\delta.$$

Therefore,

$$(x, y)_p \geq \min\{(x, z)_p, (y, z)_p\} - 3\delta.$$

□

We now prove a “converse” to this lemma:

LEMMA 11.28. *If X is a geodesic metric space which is δ -hyperbolic in the sense Gromov, then X is 2δ -hyperbolic in the sense of Rips.*

PROOF. 1. We first show that in such space geodesics connecting any pair of points are “almost” unique, i.e. if α is a geodesic connecting x to y and p is a point in X such that

$$\text{dist}(x, p) + \text{dist}(p, y) \leq \text{dist}(x, y) + 2\delta$$

then $\text{dist}(p, \alpha) \leq 2\delta$. We suppose that $\text{dist}(p, x) \leq \text{dist}(p, y)$. If $\text{dist}(p, x) \geq \text{dist}(x, y)$ then $\text{dist}(x, y) \leq 2\delta$ and thus

$$\min(\text{dist}(p, x), \text{dist}(p, y)) \leq 2\delta,$$

and we are done.

Therefore, assume that $\text{dist}(p, x) < \text{dist}(x, y)$ and let $z \in \alpha$ be such that $\text{dist}(z, y) = \text{dist}(p, y)$. Since X is δ -hyperbolic in the sense Gromov,

$$(x, y)_p \geq \min\{(x, z)_p, (y, z)_p\} - \delta.$$

Thus we can assume that $(x, y)_p \geq (x, z)_p$. Then

$$\text{dist}(y, p) - \text{dist}(x, y) \geq \text{dist}(z, p) - \text{dist}(x, z) - 2\delta \iff$$

$$\text{dist}(z, p) \leq 2\delta.$$

We conclude that $\text{dist}(p, \alpha) \leq 2\delta$.

2. Consider now a geodesic triangle $T(x, y, p) \subset X$ and let $z \in xy$. Our goal is to show that z belongs to $\mathcal{N}_{4\delta}(px \cup py)$. We have:

$$(x, y)_p \geq \min((x, z)_p, (y, z)_p) - \delta.$$

Assume that $(x, y)_p \geq (x, z)_p - \delta$. Set $\alpha := py$. We will show that $z \in \mathcal{N}_{2\delta}(\alpha)$.

By combining $\text{dist}(x, z) + \text{dist}(y, z) = \text{dist}(x, y)$ and $(x, y)_p \geq (x, z)_p - \delta$, we obtain

$$\text{dist}(y, p) \geq \text{dist}(y, z) + \text{dist}(z, p) - 2\delta.$$

Therefore, by Part 1, $z \in \mathcal{N}_{2\delta}(\alpha)$ and hence the triangle $T(x, y, z)$ is 2δ -thin. \square

COROLLARY 11.29 (M. Gromov, [Gro87], section 6.3C.). *For geodesic metric spaces, Gromov-hyperbolicity is equivalent to Rips-hyperbolicity.*

Another corollary of the Lemmata 11.27 and 11.28 is:

COROLLARY 11.30. *A geodesic metric space is a real tree if and only if it is 0-hyperbolic in the sense of Gromov.*

This corollary has a “converse” (see e.g. [Dre84] or [GdlH90, Ch. 2, Proposition 6]):

THEOREM 11.31. *Every 0-hyperbolic metric space (in the sense of Gromov) admits an isometric embedding into a tree.*

Furthermore:

THEOREM 11.32 (M. Bonk, O. Schramm [BS00]). *Every δ -hyperbolic metric space (in the sense of Gromov) admits an isometric embedding into a geodesic metric space which is also δ -hyperbolic.*

QUESTION 11.33. Does there exist a \aleph -quasiuniversal δ -hyperbolic space, i.e. a Gromov-hyperbolic metric space X such that every δ -hyperbolic metric space Y of cardinality $\leq \aleph$, admits an (L, A) quasiisometric embedding into X , with L and A depending only on δ ?

A partial positive answer to this question is provided by a universality theorem of Bonk and Schramm [BS00], see Theorem 11.218.

We next consider behavior of hyperbolicity under quasiisometries.

EXERCISE 11.34. Gromov-hyperbolicity is invariant under $(1, A)$ -quasiisometries.

EXERCISE 11.35. Let X be a metric space and $N \subset X$ be an R -net. Show that the embedding $N \hookrightarrow X$ is an $(1, R)$ -quasiisometry. Thus, X is Gromov-hyperbolic if and only if N is Gromov-hyperbolic. In particular, a group (G, d_S) with word metric d_S is Gromov-hyperbolic if and only if the Cayley graph $\Gamma_{G,S}$ of G is Rips-hyperbolic.

The drawback is that for general nongeodesic metric spaces, Gromov-hyperbolicity fails to be QI invariant:

EXAMPLE 11.36 (Gromov-hyperbolicity is not QI invariant). This example is taken from [Väi05]. Consider the graph X of the function $y = |x|$, where the metric on X is the restriction of the metric on \mathbb{R}^2 . (This is not a path-metric!) Then the map $f : \mathbb{R} \rightarrow X, f(x) = (x, |x|)$ is a quasiisometry:

$$|x - x'| \leq d(f(x), f(x')) \leq \sqrt{2}|x - x'|.$$

Let $p = (0, 0)$ be the base-point in X and for $t > 0$ we let $x := (2t, 2t)$, $y := (-2t, 2t)$ and $z := (t, t)$. The reader will verify that

$$\min((x, z)_p, (y, z)_p) - (x, y)_p = t \left(\frac{7\sqrt{2}}{2} - 3 \right) > t.$$

Therefore, the quantity $\min((x, z)_p, (y, z)_p) - (x, y)_p$ is unbounded from above as $t \rightarrow \infty$ and hence X is not δ -hyperbolic for any $\delta < \infty$. In particular, X is QI to a Gromov-hyperbolic space \mathbb{R} , but is not Gromov-hyperbolic itself. We will see, as a corollary of Morse Lemma (Corollary 11.43), that in the context of geodesic spaces, hyperbolicity is a QI invariant.

11.4. Ultralimits and stability of geodesics in Rips-hyperbolic spaces

In this section we will see that every hyperbolic geodesic metric spaces X *asymptotically* resembles a tree. This property will be used to prove *Morse Lemma*, which establishes that quasigeodesics in δ -hyperbolic spaces are uniformly close to geodesics.

LEMMA 11.37. *Let $(X_i)_{i \in \mathbb{N}}$ be a sequence of geodesic δ_i -hyperbolic spaces with δ_i tending to 0. Then for every non-principal ultrafilter ω each component of the ultralimit X_ω is a metric tree.*

PROOF. First, according to Lemma 10.51, ultralimit of geodesic metric spaces is again a geodesic metric space. Thus, in view of Lemma 11.28, it suffices to verify that X_ω is 0-hyperbolic in the sense of Gromov (since it will be 0-hyperbolic in the sense of Rips and, hence, a metric tree). This is one of the few cases where Gromov-hyperbolicity is superior to Rips-hyperbolicity: It suffices to check 0-hyperbolicity condition only for quadruples of points.

We know that for every quadruple x_i, y_i, z_i, p_i in X_i ,

$$(x_i, y_i)_{p_i} \geq \min((x_i, z_i)_{p_i}, (y_i, z_i)_{p_i}) - \delta_i.$$

By taking the ultralimit of this inequality, we obtain (for every quadruple of points $x_\omega, y_\omega, z_\omega, p_\omega$ in X_ω):

$$(x_\omega, y_\omega)_{p_\omega} \geq \min((x_\omega, z_\omega)_{p_\omega}, (y_\omega, z_\omega)_{p_\omega}),$$

since $\omega\text{-lim } \delta_i = 0$. Thus, X_ω is 0-hyperbolic. \square

COROLLARY 11.38. *Every geodesic in the tree X_ω is a limit geodesic.*

PROOF. 1. Suppose first that $x_\omega y_\omega$ is a geodesic segment in X_ω , $x_\omega = (x_i), y_\omega = (y_i)$. Then the ultralimit of geodesic segments $x_i y_i \subset X_i$ is a geodesic segment connecting x_ω to y_ω . Since each component of X_ω is 0-hyperbolic, it is uniquely geodesic, i.e. there exists a unique geodesic segment connecting x_ω to y_ω .

2. We consider the case of biinfinite geodesics in X_ω and leave the proof for geodesic rays to the reader. Let $l_\omega \subset X_\omega$ be a biinfinite geodesic parameterized by the isometric embedding $\gamma_\omega : \mathbb{R} \rightarrow X_\omega$. Take the points $x_{\omega, n} := \gamma_\omega(n), y_{\omega, n} = \gamma_\omega(-n)$. For each n the finite geodesic segment $x_{\omega, n} y_{\omega, n}$ is the ultralimit of geodesic segments $x_{i, n} y_{i, n} \subset X_i$. Then the entire l_ω is the ultralimit of the sequence of geodesic segments $x_{i, n} y_{i, n}$. \square

EXERCISE 11.39. Find a flaw in the following “proof” of Lemma 11.37: Since X_i is δ_i -hyperbolic, it follows that every geodesic triangle T_i in X_i is δ_i -thin. Suppose that $\omega\text{-lim } d(x_i, e_i) < \infty$, $\omega\text{-lim } d(p_i, e_i) < \infty$. Taking the limit in the definition of thinness of triangles, we conclude that the ultralimit of triangles $T_\omega = \omega\text{-lim } T_i \subset X_\pm$ is 0-thin. Therefore, every geodesic triangle in X_ω is 0-thin.

The following fundamental theorem in the theory of hyperbolic spaces is called *Morse Lemma* or *stability of hyperbolic geodesics*.

THEOREM 11.40 (Morse Lemma). *There exists a function $D = D(L, A, \delta)$, such that the following holds. If X be a δ -hyperbolic geodesic space, then for every (L, A) -quasigeodesic $f : [a, b] \rightarrow X$, the Hausdorff distance between the image of f and a geodesic segment $f(a)f(b) \subset X$ is at most D .*

PROOF. Set $c = d(f(a), f(b))$. Given a quasigeodesic f and $f^* : [0, c] \rightarrow X$ parameterizing the geodesic $f(a)f(b)$, we define two numbers:

$$D_f = \sup_{t \in [a, b]} d(f(t), \text{Im}(f^*))$$

and

$$D_f^* = \sup_{t \in [0, c]} d(f^*(t), \text{Im}(f)).$$

Then

$$\text{dist}_{\text{Haus}}(\text{Im}(f), \text{Im}(f^*)) = \max(D_f, D_f^*).$$

We will prove that D_f is uniformly bounded in terms of L, A, δ , since the proof for D_f^* is completely analogous.

Suppose that the quantities D_f are not uniformly bounded, that is, exists a sequence of (L, A) -quasigeodesics $f_n : [-n, n] \rightarrow X_n$ in δ -hyperbolic geodesic metric spaces X_n , such that

$$\lim_{n \rightarrow \infty} D_n = \infty,$$

where $D_n = D_{f_n}$. Pick points $t_n \in [-n, n]$ such that for $\gamma_n^* = f_n^*([-n, n])$ and $x_n := f_n(t_n)$, we have:

$$|\text{dist}(x_n, \gamma_n^*) - D_n| \leq 1.$$

In other words, the points x_n “almost” realize the maximal distance between the points of $f_n([-n, n])$ and the geodesic γ_n^* .

Define the sequence λ of scaling factors

$$\lambda_n = \frac{1}{D_n}.$$

As in Lemma 10.83, we consider two sequences of pointed metric spaces

$$(\lambda_n X_n, x_n), \quad (\lambda_n [-n, n], t_n).$$

Note that the ultralimit $\omega\text{-lim } \frac{n}{D_n}$ could be infinite, however, it cannot be zero. Let

$$(X_\omega, x_\omega) = \omega\text{-lim } (\lambda_n X_n, x_n)$$

and

$$(Y, y) := \omega\text{-lim } (\lambda_n [-n, n], t_n).$$

The metric space Y is either a nondegenerate segment in \mathbb{R} or a closed geodesic ray in \mathbb{R} or the whole real line. Note that the distance from the points of the image of f_n to γ_n^* in the rescaled metric space $\lambda_n X_n$ is at most $1 + \lambda_n$. Each map

$$f_n : Y_n \rightarrow \lambda_n X_n$$

is an $(L, A\lambda_n)$ -quasigeodesic. Therefore, the ultralimit

$$f_\omega = \omega\text{-lim } f_n : (Y, y) \rightarrow (X_\omega, x_\omega)$$

is an L -bi-Lipschitz map (cf. Lemma 10.83). In particular this map is a continuous embedding and the image of f_ω is a geodesic γ in X_ω , see Lemma 11.13.

On the other hand, the sequence of geodesic segments $\gamma_n^* \subset \lambda_n X_n$ also ω -converges to a geodesic $\gamma^* \subset X_\omega$, this geodesic is either a finite geodesic segment or a geodesic ray or a complete geodesic. In any case, by our choice of the points x_n , γ is contained in the 1-neighborhood of the geodesic γ^* and, at the same time, $\gamma \neq \gamma^*$ since $x_\omega \in \gamma \setminus \gamma^*$.

The last step of the proof is to get a contradiction with the fact that X_ω is a real tree. If γ^* is a finite geodesic, it connects the end-points of the geodesic γ , thereby contradicting the fact that each metric tree is uniquely geodesic. Suppose that γ^* is a complete geodesic. We then pick two points $y_\omega, z_\omega \in \gamma$ such that x_ω is the midpoint of the geodesic segment $y_\omega z_\omega$, while the distance between y_ω and z_ω is sufficiently high, say, larger than 4. We next find points $y'_\omega \in \gamma^*, z'_\omega \in \gamma^*$ within distance ≤ 1 from y_ω and z_ω respectively. Since

$$d(y_\omega, x_\omega) = d(z_\omega, x_\omega) \geq 2,$$

the point x_ω does not belong to union

$$y_\omega y'_\omega \cup z_\omega z'_\omega.$$

At the same time, x_ω does not belong to the geodesic segment $y'_\omega z'_\omega$ since the latter is contained in the geodesic γ^* . Therefore, the side $y_\omega z_\omega$ of the geodesic quadrilateral Q with the vertices

$$y_\omega, z_\omega, z'_\omega, y'_\omega$$

is not contained in the union of the three other sides. This contradicts the fact that Q is 0-thin. We proof in the remaining case, when γ is a geodesic ray is similar and is left to the reader. \square

HISTORICAL REMARK 11.41. The first version of this theorem was proven by Morse in [Mor24] in the following setting. Consider a compact surface S equipped with two Riemannian metrics g_1, g_2 of negative curvature. Now, lift, the metrics g_1, g_2 to the universal cover of S . Then each geodesic with respect to the lift \tilde{g}_1 of g_1 is a (uniform) quasigeodesic with respect to the lift \tilde{g}_2 of g_2 . Morse proved that all geodesics with respect to \tilde{g}_1 are uniformly close to the geodesics with respect to \tilde{g}_2 , as long as their end-points are the same. Later on, Busemann, [Bus65], proved a version of this lemma in the case of \mathbb{H}^n , where metrics in question were not necessarily Riemannian. A version in terms of quasigeodesics is due to Mostow [Mos73], in the context of negatively curved symmetric spaces, although his proof is general. The first proof for general δ -hyperbolic geodesic metric spaces is due to Gromov, see [Gro87, 7.2.A]. Of course, neither Morse, nor Busemann, nor Mostow, nor Gromov used ultralimits: Their proofs were based on an analysis of nearest-point projections to geodesics. We will give an effective proof of the Morse Lemma in Section 11.10.

REMARK 11.42. Stability of geodesics fails in the Euclidean plane \mathbb{E}^2 and, hence, for general $CAT(0)$ spaces. Nevertheless, some versions of the Morse Lemma remain true for non-hyperbolic $CAT(0)$ spaces, see [Sul14] and [KLP14].

COROLLARY 11.43 (QI invariance of hyperbolicity). *Suppose that X, X' are quasi-isometric geodesic metric spaces and X' is hyperbolic. Then X is also hyperbolic.*

PROOF. Suppose that X' is δ' -hyperbolic and $f : X \rightarrow X'$ is an (L, A) -quasiisometry and $f' : X' \rightarrow X$ is its coarse inverse. Pick a geodesic triangle $T \subset X$. Its image under f is a quasigeodesic triangle S in X' whose sides are (L, A) -quasigeodesic. Therefore each of the quasigeodesic sides σ_i of S is within distance $\leq D = D(L, A, \delta')$ from a geodesic σ_i^* connecting the end-points of this side. See Figure 11.1. The geodesic triangle S^* formed by the segments $\sigma_1^*, \sigma_2^*, \sigma_3^*$ is δ' -thin. Therefore, the quasigeodesic triangle $f'(S^*) \subset X$ is $\epsilon := L\delta' + A$ -thin, i.e. each quasigeodesic $\tau'_i := f'(\sigma_i^*)$ is within distance $\leq \epsilon$ from the union τ'_{i-1}, τ'_{i+1} . However,

$$\text{dist}_{\text{Haus}}(\tau_i, \tau'_i) \leq LD + 2A.$$

Putting this all together, we conclude that the triangle T is δ -thin with

$$\delta = 2(LD + 2A) + \epsilon = 2(LD + 2A) + L\delta' + A. \quad \square$$




FIGURE 11.1. Quasiisometric image of a geodesic triangle.

Observe that in Morse Lemma, we are not claiming, of course, that the distance $d(f(t), f^*(t))$ is uniformly bounded, only that for every t there exist s and s^* such that

$$d(f(t), f^*(s)) \leq D,$$

and

$$d(f^*(t), f(s^*)) \leq D.$$

Here $s = s(t), s^* = s^*(t)$. However, applying triangle inequalities we get for $B = A + D$ the following estimates:

$$(11.3) \quad L^{-1}t - B \leq s \leq Lt + B$$

and

$$(11.4) \quad L^{-1}(t - B) \leq s^* \leq L(t + B)$$

Lastly, we note that Proposition 11.167 proven later on, which characterizes hyperbolic geodesic metric spaces as the ones for which every asymptotic cone is a tree, provides an alternative proof of QI invariance of hyperbolicity: If two quasiisometric metric spaces have bi-Lipschitz homeomorphic asymptotic cones and one cone is a tree, then the other cone is also a tree.

11.5. Local geodesics in hyperbolic spaces

A map $p : I \rightarrow X$ of an interval $I \subset \mathbb{R}$ into a metric space X , is called a k -local geodesic if the restriction of p to each length k subinterval $I' \subset I$ is an isometric embedding. The notion of local geodesics is in line with the concept of geodesics in Riemannian geometry: (Unit speed) Riemannian geodesics are not required to be isometric embeddings, but *locally* they always are. If M is a Riemannian manifold with injectivity radius $\geq \epsilon > 0$, then every unit speed Riemannian geodesic in M is an ϵ -local geodesic in the metric sense.

EXERCISE 11.44. Suppose that X is a real tree and k is a positive number. Then each k -local geodesic in X is a geodesic.

A coarse version of this exercise works for general hyperbolic spaces as well, it is due to Gromov [Gro87, 7.2.B], see also [BH99, Ch. III.H, Theorem 1.13] and [CDP90, Ch. 3, Theorem 1.4] for a more general version:

THEOREM 11.45 (Local geodesics are uniform quasigeodesics). *Suppose that X is a δ -hyperbolic geodesic metric space in the sense of Rips, $\delta > 0$. Then for $k = 6\delta$, each k -local geodesic in X is a $(3, 4\delta)$ -quasigeodesic.*

PROOF. For each pair of points $x, y \in X$ we partition X in two half-spaces

$$\mathcal{D}(x, y) = \{z \in X : d(x, z) \leq d(z, y)\}, \quad \mathcal{D}(y, x) = \{z \in X : d(y, z) \leq d(z, x)\}.$$

The intersection of these half-spaces is the bisector $Bis(x, y)$ of the pair (x, y) , consisting of all points equidistant from x and y .




FIGURE 11.2. Bisector and half-spaces.

The key to the proof is the following:

LEMMA 11.46. *Consider three points $x_0, x_1, x_2 \in X$ such that*

$$\epsilon = 3\delta = d(x_0, x_1) = d(x_1, x_2) = \frac{1}{2}d(x_0, x_2).$$

Then:

1. $\text{dist}(x_0, \mathcal{D}(x_1, x_0)) \geq \delta$, $\text{dist}(x_2, \mathcal{D}(x_1, x_2)) \geq \delta$.
2. The half-space $\mathcal{D}(x_1, x_0)$ contains $\mathcal{D}(x_2, x_1)$ and, moreover,

$$\text{dist}(\mathcal{D}(x_0, x_1), \mathcal{D}(x_2, x_1)) \geq \delta/2.$$

PROOF. 1. Let y be a point in the bisector $Bis(x_0, x_1)$ nearest to x_0 and let $m \in x_0x_1$ be the midpoint of a geodesic segment x_0x_1 connecting x_0 to x_1 . Since the geodesic triangle $\Delta(x_0, y, x_1)$ is δ -thin, the distance from m to one of the two sides x_0y, yx_1 of this triangle does not exceed δ . We will assume that this side is x_0y , since the other case is obtained by relabeling.




FIGURE 11.3. Nested half-spaces.

Let $z \in x_0y$ be the point closest to m . Then, by the triangle inequality,

$$d(x_0, y) \geq d(x_0, z) \geq d(x_0, m) - \delta = \frac{\epsilon}{2} - \delta.$$

Since $\epsilon = 3\delta$, we obtain

$$\text{dist}(x, Bis(x_0, x_1)) = d(x_0, y) \geq \delta/2.$$

2. Take points $y_1 \in \mathcal{D}(x_0, x_1)$, $y_2 \in \mathcal{D}(x_2, x_1)$, i.e.

$$D_0 = d(x_0, y_1) \leq D_1 = d(y_1, x_1), \quad D_3 = d(x_2, y_2) \leq D_2 = d(y_2, x_1).$$

Our goal is to estimate the distance

$$\eta = d(y_1, y_2)$$

from below, this will provide a lower bound on the distance between the half-spaces.

We let u denote the midpoint of a geodesic segment y_1y_2 . Then

$$d(x_0, y_1) \leq D_0 + \frac{\eta}{2}, \quad d(x_2, y_2) \leq D_3 + \frac{\eta}{2}.$$

We also have

$$D = d(u, x_1) \geq \max(D_1 - \frac{\eta}{2}, D_2 - \frac{\eta}{2}).$$

Since the triangle $\Delta(x_0, u, x_2)$ is δ -thin, and, by the hypothesis of Lemma, x_1 lies on the geodesic segment x_0x_2 , the distance from x_1 to one of the sides x_0u, x_2u of this triangle is at most δ . We will assume that this side is x_0u . Let $v \in x_0u$ denote the point closest to x_1 ; this point divides the segment x_0u into two subsegments vu, vx_0 of the lengths D'_1, D''_1 respectively. We have

$$D'_1 \geq D - \delta \geq D_1 - \delta - \frac{\eta}{2} \geq D_0 - \delta - \frac{\eta}{2},$$




FIGURE 11.4. Estimating distance between half-spaces.

which implies

$$D_1'' = d(x_0, u) - D_1' \leq (D_0 + \frac{\eta}{2}) - (D_0 - \delta - \frac{\eta}{2}) = \eta + \delta.$$

Combining this with the triangle inequality for $\Delta(x_0, v, x_1)$, we obtain:

$$\epsilon \leq D_1'' + \delta \leq \eta + 2\delta,$$

implying

$$\eta \geq \epsilon - 2\delta = 3\delta - 2\delta = \delta. \quad \square$$

We now can prove the theorem. Let q be a 6δ -local geodesic in X . We first consider the special case when q has length $n\epsilon$, $n \in \mathbb{N}$ (where, as before, $\epsilon = 3\delta$) and estimate from below the distance between the end-points of q in terms of the length of q . We subdivide q into n subsegments

$$x_0x_1, x_1x_2, \dots, x_{n-1}x_n$$

of length ϵ . Since q is a $k = 2\epsilon$ -local geodesic, the unions

$$x_{i-1}x_i \cup x_ix_{i+1}, \quad i = 1, \dots, n-1$$

are geodesic segments in X . Therefore, by Lemma 11.46,

$$\text{dist}(\mathcal{D}(x_{i-1}, x_i), \mathcal{D}(x_{i+1}, x_i)) \geq \delta$$

for each i . Furthermore, the distances from x_0 to $\mathcal{D}(x_2, x_1)$ and from x_n to $\mathcal{D}(x_{n-1}, x_n)$ are at least $\delta/2$. Thus, every path q' connecting x_0 to x_n has length at least

$$\frac{\delta}{2} + (n-1)\delta + \frac{\delta}{2} = n\delta,$$

which is the length of q divided by 3.

Consider now the general case. Suppose that p is a k -local geodesic in X , x_0, x are points on p such that the length ℓ of p between them equals $n\epsilon + \sigma$, $0 < \sigma < \epsilon$. We represent the portion of p between x_0, x as the concatenation of two sub-paths: q (of length $n\epsilon$, connecting x_0 to x_n) and r (of length σ , connecting x_n to x). Then (by the special case considered above)

$$3d(x_0, x_n) \geq \text{length}(q) \geq \ell - \epsilon,$$




FIGURE 11.5. Estimating the distance between the end-points of a local geodesic.

$$3d(x_0, x) > 3d(x_0, x_n) - 3\epsilon \geq \ell - 4\epsilon,$$

$$\frac{1}{3}\ell - \frac{4}{3}\epsilon = \frac{1}{3}\ell - 4\delta < d(x_0, x).$$

Hence, p is a $(3, 4\delta)$ -quasigeodesic in X . \square

11.6. Quasiconvexity in hyperbolic spaces

The usual notion of convexity is not particularly useful in the context of hyperbolic geodesic metric spaces, it is replaced with the one of *quasiconvexity*.

DEFINITION 11.47. Let X be a geodesic metric space and $Y \subset X$. Then the *quasiconvex hull* $H(Y)$ of Y in X is the union of all geodesics $y_1y_2 \subset X$, with the end-points y_1, y_2 contained in Y .

Accordingly, a subset $Y \subset X$ is called *R -quasiconvex* if $H(Y) \subset \mathcal{N}_R(Y)$. A subset Y is called *quasiconvex* if it is quasiconvex for some $R < \infty$.

Let X be a δ -hyperbolic geodesic metric space.

The thin triangle property immediately implies that subsets of a δ -hyperbolic geodesic metric space satisfy:

1. Every metric ball $B(x, R)$ in X is δ -quasiconvex.
2. Suppose that $Y_i \subset X$ is R_i -quasiconvex, $i = 1, 2$, and $Y_1 \cap Y_2 \neq \emptyset$. Then $Y_1 \cup Y_2$ is $R_1 + R_2 + \delta$ -quasiconvex.

Thus, quasiconvex subsets behave somewhat differently from the convex ones, since the union of convex sets (with non-empty intersection) need not be convex.

An example of a non-quasiconvex subset is a horosphere in \mathbb{H}^n : Its quasiconvex hull is the horoball bounded by this horosphere.

EXERCISE 11.48. The quasiconvex hull of any subset $Y \subset X$ of a δ -hyperbolic geodesic metric space, is 2δ -quasiconvex in X . Hint: Use the fact that quadrilaterals in X are 2δ -thin.

Thus, quasiconvex hulls are quasiconvex.

The following results connect quasiconvexity and quasiisometry for subsets of Gromov-hyperbolic geodesic metric spaces.

THEOREM 11.49. *Let X, Y be geodesic metric spaces, such that X is δ -hyperbolic geodesic metric space. Then for every quasiisometric embedding $f : Y \rightarrow X$, the image $f(Y)$ is quasiconvex in X .*

PROOF. Let $y_1, y_2 \in Y$ and $\alpha = y_1 y_2 \subset Y$ be a geodesic connecting y_1 to y_2 . Since f is an (L, A) quasiisometric embedding, $\beta = f(\alpha)$ is an (L, A) quasigeodesic in X . By the Morse Lemma,

$$\text{dist}_{\text{Haus}}(\beta, \beta^*) \leq R = D(L, A, \delta),$$

where β^* is any geodesic in X connecting $x_1 = f(y_1)$ to $x_2 = f(y_2)$. Therefore, $\beta^* \subset \mathcal{N}_R(f(Y))$, and $f(Y)$ is R -quasiconvex. \square

The map $f : Y \rightarrow f(Y)$ is a quasiisometry, where we use the restriction of the metric from X to define a metric on $f(Y)$. Of course, $f(Y)$ is not a geodesic metric space, but it is quasiconvex; thus, applying the same arguments as in the proof of Theorem 11.43, we conclude that Y is also hyperbolic.

Conversely, let $Y \subset X$ be a coarsely connected subset, i.e. there exists a constant $c < \infty$ such that the complex $\text{Rips}_C(Y)$ is connected for all $C \geq c$, where we again use the restriction of the metric d from X to Y to define the Rips complex. Then we define a path-metric $d_{Y,C}$ on Y by looking at infima of lengths of edge-paths in $\text{Rips}_C(Y)$ connecting points of Y . The following is a converse to Theorem 11.49:

THEOREM 11.50. *Suppose that $Y \subset X$ is coarsely connected and Y is quasiconvex in X . Then the identity map $f : (Y, d_{Y,C}) \rightarrow (X, \text{dist}_X)$ is a quasiisometric embedding for all $C \geq 2c + 1$.*

PROOF. Let C be such that $H(Y) \subset \mathcal{N}_C(Y)$. First, if $d_Y(y, y') \leq C$ then $\text{dist}_X(y, y') \leq C$ as well. Hence, f is coarsely Lipschitz. Let $y, y' \in Y$ and γ is a geodesic in X of length L connecting y, y' . Subdivide γ into $n = \lfloor L \rfloor$ of unit subsegments and a subsegment of the length $L - n$:

$$z_0 z_1, \dots, z_{n-1} z_n, \quad z_n, z_{n+1},$$

where $z_0 = y, z_{n+1} = y'$. Since each z_i belongs to $\mathcal{N}_c(Y)$, there exist points $y_i \in Y$ such that $\text{dist}_X(y_i, z_i) \leq c$, where we take $y_0 = z_0, y_{n+1} = z_{n+1}$. Then

$$\text{dist}_X(z_i, z_{i+1}) \leq 2c + 1 \leq C$$

and, hence, z_i, z_{i+1} are connected by an edge (of length C) in $\text{Rips}_C(Y)$. Now it is clear that

$$d_{Y,C}(y, y') \leq C(n + 1) \leq C \text{dist}_X(y, y') + C. \quad \square$$

REMARK 11.51. It is proven in [Bow94] that in the context of subsets of negatively pinched complete simply-connected Riemannian manifolds X , quasiconvex hulls $Hull(Y)$ are essentially the same as convex hulls $H(Y)$:

There exists a function $L = L(C)$ such that for every C -quasiconvex subset $Y \subset X$,

$$H(Y) \subset Hull(Y) \subset \mathcal{N}_{L(C)}(Y).$$

11.7. Nearest-point projections

In general, nearest-point projections to geodesics in δ -hyperbolic geodesic spaces are not well defined. The following lemma shows, nevertheless, that they are *coarsely-well defined*:

Let γ be a geodesic in δ -hyperbolic geodesic space X . For a point $x \in X$ let $p = \pi_\gamma(x)$ be a point nearest to x .

LEMMA 11.52. *Let $p' \in \gamma$ be such that $d(x, p') < d(x, p) + R$. Then*

$$d(p, p') \leq 2(R + 2\delta).$$

In particular, if $p, p' \in \gamma$ are both nearest to x then

$$d(p, p') \leq 4\delta.$$

PROOF. Consider the geodesics α, α' connecting x to p and p' respectively. Let $q' \in \alpha'$ be the point within distance $\delta + R$ from p' (this point exists unless $d(x, p) < \delta + R$ in which case $d(p, p') \leq 2(\delta + R)$ by the triangle inequality). Since the triangle $\Delta(x, p, p')$ is δ -thin, there exists a point

$$q \in xp \cup pp' \subset xp \cup \gamma$$

within distance δ from q . If $q \in \gamma$, we obtain a contradiction with the fact that the point p is nearest to x on γ (the point q will be closer). Thus, $q \in xp$. By the triangle inequality

$$d(x, p') - (R + \delta) = d(x, q') \leq d(x, q) + \delta \leq d(x, p) - d(q, p) + \delta.$$

Thus,

$$d(q, p) \leq d(x, p) - d(x, p') + R + 2\delta \leq R + 2\delta.$$

Since $d(p', q) \leq R + 2\delta$, we obtain $d(p', p) \leq 2(R + 2\delta)$. □

This lemma can be strengthened, we now show that the nearest-point projection to a quasigeodesic subspace in a hyperbolic space is coarse Lipschitz:

LEMMA 11.53. *Let $X' \subset X$ be an R -quasiconvex subset. Then the nearest-point projection $\pi = \pi_{X'} : X \rightarrow X'$ is $(2, 2R + 9\delta)$ -coarse Lipschitz.*

PROOF. Suppose that $x, y \in X$ such that $d(x, y) = D$. Let $x' = \pi(x), y' = \pi(y)$. Consider the quadrilateral formed by geodesic segments $xy \cup [y, y'], [y', x'] \cup [x', x]$. Since this quadrilateral is 2δ -thin, there exists a point $q \in x'y'$ which is within distance $\leq 2\delta$ from $x'x \cup xy$ and $xy \cup yy'$.

Case 1. We first assume that there are points $x'' \in x'x, y'' \in yy'$ such that

$$d(q, x'') \leq 2\delta, d(q, y'') \leq 2\delta.$$

Let $q' \in X'$ be a point within distance $\leq R$ from q . By considering the paths

$$xx'' \cup x''q \cup qq', \quad yy'' \cup y''q \cup qq'$$

and using the fact that $x' = \pi(x), y' = \pi(y)$, we conclude that

$$d(x', x'') \leq R + 2\delta, \quad d(y', y'') \leq R + 2\delta.$$

Therefore,

$$d(x', y') \leq 2R + 9\delta.$$

Case 2. Suppose that there exists a point $q'' \in xy$ such that $d(q, q'') \leq 2\delta$. Setting $D_1 = d(x, q''), D_2 = d(y, q'')$, we obtain

$$d(x, x') \leq d(x, q') \leq D_1 + R + 2\delta, d(y, y') \leq d(y, q') \leq D_2 + R + 2\delta$$

which implies that

$$d(x', y') \leq 2D + 2R + 4\delta.$$

In either case, $d(x', y') \leq 2d(x, y) + 2R + 9\delta$. □




FIGURE 11.6. Projection to a quasiconvex subset.

11.8. Geometry of triangles in Rips–hyperbolic spaces

In the case of real-hyperbolic space we relied upon hyperbolic trigonometry in order to study geodesic triangles. Trigonometry no longer makes sense in the context of Rips–hyperbolic spaces X , so instead one compares geodesic triangles in X to geodesic triangles in real trees, i.e. to tripods, in the manner similar to the comparison theorems for $CAT(\kappa)$ -spaces. In this section we describe comparison maps to tripods, called *collapsing maps*. We will see that such maps are $(1, 14\delta)$ -quasiisometries. We will use the collapsing maps in order to get a detailed information about geometry of triangles in X .

A *tripod* \tilde{T} is a metric graph, which, as a graph, is isomorphic to the 3-pod, see Example 1.34. We will use the notation o for the center of the tripod. By abusing the notation, we will regard a tripod \tilde{T} as a geodesic triangle whose vertices are the extreme points (leaves) \tilde{x}_i of \tilde{T} ; hence, we will use the notation $\mathcal{T} = \tilde{T} = T(\tilde{x}_1, \tilde{x}_2, \tilde{x}_3)$. Accordingly, the *side-lengths* of a tripod are lengths of the sides of the corresponding triangle.

REMARK 11.54. Using the symbol \sim in the notation for a tripod is motivated by the comparison geometry, as we will compare geodesic triangles in δ -hyperbolic spaces with the tripods \tilde{T} : This is analogous to comparing geodesic triangles in metric spaces to geodesic triangles in constant curvature spaces, see Definition 3.56.




FIGURE 11.7. Collapsing map of a triangle to a tripod.

EXERCISE 11.55. For any three numbers $a_i \in \mathbb{R}_+, i = 1, 2, 3$ satisfying the triangle inequalities $a_i \leq a_j + a_k$ ($\{1, 2, 3\} = \{i, j, k\}$), there exists a unique (up to isometry) tripod $\tilde{T} = \mathcal{T}_{a_1, a_2, a_3}$ with the side-lengths a_1, a_2, a_3 .

Now, given a geodesic triangle $T = T(x_1, x_2, x_3)$ with side-lengths $a_i, i = 1, 2, 3$ in a metric space X , there exists a unique (possibly up to postcomposition with an isometry $\tilde{T} \rightarrow \tilde{T}$) map κ to the *comparison* tripod \tilde{T} ,

$$\kappa : T \rightarrow \tilde{T} = \mathcal{T}_{a_1, a_2, a_3},$$

which restricts to an isometry on every edge of T : The map κ sends the vertices x_i of T to the leaves \tilde{x}_i of the tripod \tilde{T} . The map κ is called the *collapsing map* for T . We say that the points $x, y \in T$ are *dual* to each other if $\kappa(x) = \kappa(y)$.

EXERCISE 11.56. The collapsing map κ is 1-Lipschitz and preserves the Gromov-products $(x_i, x_j)_{x_k}$.

Then

$$(x_i, x_j)_{x_k} = d(\tilde{x}_k, [\tilde{x}_i, \tilde{x}_j]) = d(\tilde{x}_k, o).$$

By taking the preimage of $o \in \tilde{T}$ under the maps $\kappa|_{(x_i x_j)}$ we obtain points

$$x_{ij} \in x_i x_j$$

called the *central points* of the triangle T :

$$d(x_i, x_{ij}) = (x_j, x_k)_{x_i}.$$

LEMMA 11.57 (Approximation of triangles by tripods). *Assume that a geodesic metric space X is δ -hyperbolic in the sense of Rips, and consider an arbitrary geodesic triangle $T = \Delta(x_1, x_2, x_3)$ with the central points $x_{ij} \in x_i x_j$. Then for every $\{i, j, k\} = \{1, 2, 3\}$ we have:*

1. $d(x_{ij}, x_{jk}) \leq 6\delta$.
2. $d_{\text{Haus}}(x_j x_{ji}, x_j x_{kj}) \leq 7\delta$.
3. *Distances between dual points in T are $\leq 14\delta$. In detail: Suppose that $\alpha_{ji}, \alpha_{jk} : [0, t_j] \rightarrow X$ ($t_j = d(x_j, x_{ij}) = d(x_j, x_{jk})$) are unit speed parameterizations of geodesic segments $x_j x_{ji}, x_j x_{jk}$. Then*

$$d(\alpha_{ji}(t), \alpha_{jk}(t)) \leq 14\delta$$

for all $t \in [0, t_j]$.

PROOF. The geodesic $x_i x_j$ is covered by the closed subsets $\overline{N}_\delta(x_i x_k)$ and $\overline{N}_\delta(x_j x_k)$, hence by connectedness there exists a point p on $x_i x_j$ at distance at most δ from both $x_i x_k$ and $x_j x_k$. Let $p' \in x_i x_k$ and $p'' \in x_j x_k$ be points at distance at most δ from p . The inequality

$$(x_j, x_k)_{x_i} = \frac{1}{2} [d(x_i, p) + d(p, x_j) + d(x_i, p') + d(p', x_k) - d(x_j, p'') - d(p'', x_k)]$$

combined with the triangle inequality implies that

$$|(x_j, x_k)_{x_i} - d(x_i, p)| \leq 2\delta,$$

and, hence $d(x_{ij}, p) \leq 2\delta$. Then $d(x_{ik}, p') \leq 3\delta$, whence $d(x_{ij}, x_{ik}) \leq 6\delta$. It remains to apply Lemma 11.3 to obtain 2 and Lemma 11.4 to obtain 3. \square

We thus obtain

PROPOSITION 11.58. κ is a $(1, 14\delta)$ -quasiisometry.

PROOF. The map κ is a surjective 1-Lipschitz map. On the other hand, Part 3 of the above lemma implies that

$$d(x, y) - 14\delta \leq d(\kappa(x), \kappa(y))$$

for all $x, y \in T$. \square

Proposition 11.58 allows one to reduce (up to a uniformly bounded error) study of geodesic triangles in δ -hyperbolic spaces to study of tripods. For instance suppose that $m_{ij} \in x_i x_j$ are points such that

$$d(m_{ij}, m_{jk}) \leq r$$

for all i, j, k . We already know that this property holds for the central points x_{ij} of T (with $r = 6\delta$). Next result shows that points m_{ij} have to be uniformly close to the central points:

COROLLARY 11.59. Under the above assumptions, $d(m_{ij}, x_{ij}) \leq r + 14\delta$.

PROOF. Since κ is 1-Lipschitz,

$$d(\kappa(m_{ik}), \kappa(m_{jk})) \leq r$$

for all i, j, k . By definition of the map κ , all three points $\kappa(m_{ij})$ cannot lie in the same leg of the tripod \tilde{T} , except when one of them is the center o of the tripod. Therefore, $d(\kappa(m_{ij}), o) \leq r$ for all i, j . Since κ is $(1, 14\delta)$ -quasiisometry,

$$d(m_{ij}, x_{ij}) \leq d(\kappa(m_{ik}), \kappa(m_{jk})) + 14\delta \leq r + 14\delta.$$

DEFINITION 11.60. We say that a point $p \in X$ is an R -centroid of a triangle $T \subset X$ if distances from p to all three sides of T are $\leq R$.

COROLLARY 11.61. Every two R -centroids of T are within distance at most $\phi(R) = 4R + 28\delta$ from each other.

PROOF. Given an R -centroid p , let $m_{ij} \in x_i x_j$ be the nearest points to p . Then

$$d(m_{ij}, m_{jk}) \leq 2R$$

for all i, j, k . By previous corollary,

$$d(m_{ij}, x_{ij}) \leq 2R + 14\delta.$$

Thus, triangle inequalities imply that every two centroids are within distance at most $2(2R + 14\delta)$ from each other. \square

Let $p_3 \in \gamma_{12} = x_1x_2$ be a point closest to x_3 . Taking $R = 2\delta$ and combining Lemma 11.22 with Lemma 11.52, we obtain:

COROLLARY 11.62. $d(p_3, x_{12}) \leq 2(2\delta + 2\delta) = 6\delta$.

We now can define a continuous coarse inverse $\bar{\kappa}$ to κ as follows: We map the geodesic segment $\tilde{x}_1\tilde{x}_2 \subset \tilde{T}$ isometrically to a geodesic x_1x_2 . We send $o\tilde{x}_3$ onto a geodesic $x_{12}x_3$ by an affine map. Since

$$d(x_{12}, x_{32}) \leq 6\delta$$

and

$$d(x_3, x_{32}) = d(\tilde{x}_3, 0),$$

we conclude that the map $\bar{\kappa}$ is $(1, 6\delta)$ -Lipschitz.

EXERCISE 11.63.

$$d(\bar{\kappa} \circ \kappa, \text{Id}) \leq 32\delta.$$

11.9. Divergence of geodesics in hyperbolic metric spaces

Another important feature of hyperbolic spaces is the *exponential divergence* of its geodesic rays. This can be deduced from the thinness of polygons described in Lemma 11.6, as shown below. Our arguments are inspired by those in [Pap03].

LEMMA 11.64. *Let X be a geodesic metric space, δ -hyperbolic in the sense of Rips. If xy is a geodesic of length $2r$ and m is its midpoint, then every path joining x, y outside the ball $B(m, r)$ has length at least $2^{\frac{r-1}{\delta}}$.*

PROOF. Consider such a path \mathbf{p} , of length ℓ . We divide this path first into two arcs of length $\frac{\ell}{2}$, then into four arcs of length $\frac{\ell}{4}$ etc., until we obtain $n = 2^k$ arcs of length $\frac{\ell}{2^k} \leq 2$. The minimal k satisfying this condition equals $\lceil \log_2 \ell \rceil$. Let $x_0 = x, x_1, \dots, x_n = y$ be the consecutive subdivision points on \mathbf{p} obtained after this procedure. Lemma 11.6 applied to a geodesic polygon $x_0x_1 \dots x_n$ implies that m is contained in the $k\delta$ -neighborhood of the broken geodesic

$$\mathbf{q} = \bigcup_{i=0}^{n-1} x_i x_{i+1}.$$

Let $p \in \mathbf{q}$ be the point closest to m . Since $d(x_i, m) \geq r$ for each i and $d(x_i, p) \leq 1$ for some i , we conclude that

$$r \leq k\delta + 1$$

and, hence,

$$r - 1 \leq \delta \log_2(\ell), \quad 2^{\frac{r-1}{\delta}} \leq \ell. \quad \square$$

The content of the next two lemmas can be described by saying that *geodesic rays in hyperbolic spaces diverge (at least) exponentially fast*.

LEMMA 11.65. *Let X be a geodesic metric space, δ -hyperbolic in the sense of Rips, and let x and y be two points on the sphere $S(o, R)$ such that $\text{dist}(x, y) = 2r$. Then every path joining x and y outside $\bar{B}(o, R)$ has length at least $\psi(r) = 2^{\frac{r-1}{\delta}-3} - 12\delta$.*

PROOF. Let $m \in xy$ be the midpoint. Since $d(o, x) = d(o, y)$, it follows that m is also one of the centroids of the triangle $T(x, y, o)$ in the sense of Section 11.8. Then, by using Lemma 11.57 (Part 1), we see that $d(m, o) \leq (R-r) + 6\delta$. Therefore, the closed ball $\overline{B}(m, r - 6\delta)$ is contained in $\overline{B}(o, R)$. Let \mathfrak{p} be a path joining x and y outside $\overline{B}(o, R)$, and let xx' and $y'y'$ be subsegments of xy of length 6δ . Lemma 11.64 implies that the path $x'x \cup \mathfrak{p} \cup yy'$ has length at least

$$2^{\frac{r-6\delta-1}{\delta}}$$

whence \mathfrak{p} has length at least

$$2^{\frac{r-1}{\delta}-3} - 12\delta.$$

□

LEMMA 11.66. *Let X be a δ -hyperbolic in the sense of Rips, and let x and y be two points on the sphere $S(o, r_1 + r_2)$ such that there exist two geodesics xo and yo intersecting the sphere $S(o, r_1)$ in two points x', y' at distance larger than 14δ . Then every path joining x and y outside $B(o, r_1 + r_2)$ has length at least $\psi(r_2 - 15\delta) = 2^{\frac{r_2-1}{\delta}-18} - 12\delta$.*

PROOF. Let m be the midpoint m of xy ; since $T(x, y, o)$ is isosceles, m is one of the centroids of this triangle. Since $d(x', y') > 14\delta$, they cannot be dual points on $\Delta(x, y, o)$ in the sense of Section 11.8. Let $x'', y'' \in xy$ be dual to x', y' . Thus (by Lemma 11.57 (Part 3)),

$$d(o, x'') \leq r_1 + 14\delta, d(o, y'') \leq r_1 + 14\delta.$$

Furthermore, by the definition of dual points, since m is a centroid of $\Delta(x, y, o)$, m belongs to the segment $x''y'' \subset xy$. Thus, by quasiconvexity of metric balls, see Section 11.6,

$$d(m, o) \leq r_1 + 14\delta + \delta = r_1 + 15\delta.$$

By the triangle inequality,

$$r_1 + r_2 = d(x, o) \leq r + d(m, o) \leq r + r_1 + 15\delta, \quad r_2 - 15\delta \leq r.$$

Since the function ψ in Lemma 11.65 is increasing,

$$\psi(r_2 - 15\delta) \leq \psi(r).$$

Combining this with Lemma 11.65 (where we take $R = r_1 + r_2$), we obtain the required inequality. □

COROLLARY 11.67. *Suppose that $\rho, \rho' \in \text{Ray}_p(X)$ are inequivalent rays. Then for every sequence t_n diverging to ∞ ,*

$$\lim_{i \rightarrow \infty} d(\rho(t_i), \rho'(t_i)) = \infty.$$

PROOF. Suppose to the contrary, there exists a divergent sequence t_i such that $d(\rho(t_i), \rho'(t_i)) \leq D$. Then, by Lemma 11.4, for every $t \leq t_i$,

$$d(\rho(t), \rho'(t)) \leq 2(D + \delta).$$

Since $\lim_{i \rightarrow \infty} t_i = \infty$, it follows that $\rho \sim \rho'$. A contradiction. □

We now promote the conclusion of Lemmas 11.64 and 11.66 to the notion of *divergence* of geodesics and spaces. In both definitions, X is a 1-ended geodesic metric space (which need not be hyperbolic).

DEFINITION 11.68 (Divergence of a geodesic). Let $\gamma : \mathbb{R} \rightarrow X$ be a biinfinite geodesic in X . This geodesic is said to have *divergence* $\geq \zeta(r)$ if for the points $x = \gamma(-r), y = \gamma(r)$, the infimum of lengths of paths \mathfrak{p} connecting x to y outside of the ball $B(\gamma(0), r)$ is $\geq \zeta(r)$.

The assumption that X is 1-ended in this definition is needed to ensure that the paths \mathfrak{p} connecting x to y exist.

EXERCISE 11.69. Suppose that X is the Cayley graph of a finitely generated group. Show that each geodesic in X has at most exponential divergence.

In view of this definition, lemma 11.64 says that in every δ -hyperbolic geodesic metric space, every biinfinite geodesic has divergence $\geq 2^{\frac{r-1}{2\delta}}$: Geodesics in Rips-hyperbolic spaces have at least exponential divergence. Similarly to the definition of divergence of geodesics, one defines divergence of quasigeodesics.

We refer the reader to [Ger94, KL98a, Mac13, DR09, BC12, AK11, Sul14] for a more detailed treatment of divergence of geodesics in metric spaces.

Lemma 11.66 suggests a notion of divergence of a space based on *uniform divergence of pairs of geodesic rays* rather than of geodesics.

DEFINITION 11.70 (Uniform divergence of a space). A continuous function $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}$ is called a *uniform divergence function* for X if for every point $o \in X$ and geodesic segments $\alpha = ox, \beta = oy$ in X , for all $r, R \in \mathbb{R}_+$ satisfying

$$R + r \leq \min(d(o, x), d(o, y)), \quad d(\alpha(r), \beta(r)) \geq \eta(0)$$

and every path \mathfrak{p} in $X \setminus B(o, r + R)$ connecting $\alpha(R + r)$ to $\beta(R + r)$, we have

$$\text{length}(\mathfrak{p}) \geq \eta(R).$$

For instance, in view of Lemma 11.66, every hyperbolic geodesic metric space has an exponential uniform divergence function.

THEOREM 11.71 (P. Papasoglu, [Pap95c]). *If X is a geodesic metric space with proper uniform divergence function, then X is hyperbolic.*

11.10. Morse Lemma revisited

In this section we use Lemma 11.64 to give another proof of the Morse Lemma, this time with an explicit bound on the distance between quasigeodesic and geodesic paths. We note that a more refined (and sharp) estimate in the Morse Lemma is established by V. Schur in [Sch13].

THEOREM 11.72. *For every (L, A) -quasigeodesic $\mathfrak{q} : [0, T] \rightarrow X$ in a δ -hyperbolic geodesic space X , the image of \mathfrak{q} is within Hausdorff distance $\leq D = D(L, A, \delta)$ from every geodesic $xy \subset X$ connecting the endpoints of \mathfrak{q} . The function D can be estimated from above as*

$$D \leq L(A + 1 + 2R_*)(L + A),$$

where

$$R_* = R_*(L, A, \delta) \leq \frac{L + 2A}{6} + 2\delta \log_2(7L(L + A)) + 2\delta \log_2(\delta)$$

provided that δ is at least 2.

PROOF. Let $\mathfrak{q} : [0, T] \rightarrow X$ be an (L, A) -quasigeodesic path in X . We let $x = \mathfrak{q}(0), y = \mathfrak{q}(T)$. Then the set

$$N = \mathfrak{q}([0, T] \cap \mathbb{Z}) \cup \{y\}$$

is an $(L + A)/2$ -net in the image of \mathfrak{q} . Note that for $i, i + 1 \in [0, T] \cap \mathbb{Z}$,

$$d(x_i, x_{i+1}) \leq (L + A)$$

and

$$d(x_n, y) \leq L + A,$$

where $n = \lfloor T \rfloor$.

Take a geodesic segment xy connecting the endpoints x, y of \mathfrak{q} . We let $m \in xy$ denote the point within the largest distance (denoted R) from the net N . We parameterize the segment xy via an isometric map $\gamma : [a, b] \rightarrow X$ such that $\gamma(a) = x, \gamma(b) = y, \gamma(0) = m$. We first consider the *generic* case when $a \leq -2R, 2R \leq b$. We mark four points

$$z = \gamma(-2R), x' = \gamma(-R), y' = \gamma(R), w = \gamma(2R)$$

in the segment xy and let $x_i = \mathfrak{q}(i), x_j = \mathfrak{q}(j)$ denote the points in the net N closest to z and w respectively. Due to our choice of m to be the point in xy farthest from N , we obtain the bound

$$\max(d(z, x_i), d(w, x_j)) \leq R.$$

We will consider the case $i \leq j$ for convenience of the notation and leave the case $i \geq j$ to the reader. Then we have a broken geodesic β connecting x_i and x_j :

$$\beta = x_i x_{i+1} \cdots x_{j-1} x_j.$$

Since \mathfrak{q} is an (L, A) -quasigeodesic, the length $\ell(\beta)$ of the path β is at most

$$(L + A)(j - i) \leq (L + A)L(d(x_i, x_j) + A).$$

Since $d(x_i, x_j) \leq 6R + (L + A)$, we obtain the bound

$$\ell(\beta) \leq (L + A)L(6R + (L + A) + A) = (L + A)L(6R + L + 2A).$$

Since none of the points $x_k, k = i, \dots, j$ belongs to the open ball $B(m, R)$, we conclude that the piecewise-geodesic concatenation

$$\alpha = x'z \star zx_i \star \beta \star x_jw \star wy'$$

is disjoint from the open ball $B(m, R)$. Therefore, in view of Lemma 11.64, the length of α is at least

$$2^{\frac{R-1}{\delta}}$$

which is a superlinear function of R . On the other hand, α has length at most

$$4R + \ell(\beta) \leq 4R + (L + A)L(6R + L + 2A)$$

which is a linear function of R . Therefore, the inequality

$$(11.5) \quad 2^{\frac{R-1}{\delta}} \leq 4R + (L + A)L(6R + L + 2A)$$

forces $R \leq R_*$ for some $R_* = R_*(L, A, \delta)$. Below we will estimate R_* (from above) explicitly.




FIGURE 11.8. The path α .

EXERCISE 11.73. For $\delta \geq 2$, $c_1 > 0$, $c_2 > 0$, if

$$2^{\frac{R-1}{\delta}} \leq c_1 R + c_2$$

then

$$R \leq \frac{c_2}{c_1} + 2\delta \log_2(c_1) + 2\delta \log_2(\delta)$$

Therefore, we obtain the estimate:

$$\begin{aligned} R_* &\leq \frac{L(L+A)(L+2A)}{4+6L(L+A)} + 2\delta \log_2(4+6L(L+A)) + 2\delta \log_2(\delta) \leq \\ &\frac{L+2A}{6} + 2\delta \log_2(7L(L+A)) + 2\delta \log_2(\delta). \end{aligned}$$

We next consider the *nongeneric* cases, i.e. when $d(m, y) < 2R$ or $d(x, m) < 2R$. There are several subcases to analyze, we will deal with the case

$$d(x, m) \geq 2R, \quad R \leq d(m, y) < 2R$$

and leave the other two possibilities to the reader since they are similar. We define the points x' , z and x_i as before, but now use the point y to play the role of x_j . The broken path β above will be replaced with the broken path

$$\beta := x_i x_{i+1} \cdots x_n y$$

and we will use the concatenation $x'z \star z x_i \star \beta \star y y'$. With this modification, the same inequality (11.5) still holds and we again conclude that $R \leq R_*$, where R_* is the same function as above.

So far, we proved that the geodesic segment xy is contained in the R_* -neighborhood of the image of \mathfrak{q} . More precisely, we proved that for each $t \in [a, b]$ there exists $s = s(t) \in \{0, 1, \dots, n, T\}$ such that $d(\mathfrak{q}(s), \gamma(t)) \leq R_*$. This choice of $s(t)$ need not be unique, but we will use $s(a) = 0$, $s(b) = T$ since the respective distances in X will be zero in this situation. It is convenient at this point to reparameterize the geodesic xy so that $a = 0$. The function $t \mapsto s$, of course, is not continuous, but it is coarse Lipschitz:

$$|s(t) - s(t+1)| \leq L(A+1+2R_*)$$

for all $t, t+1 \in [0, b]$. We will replace this function with a piecewise-linear function as follows. For every $t \in [0, b] \cap \mathbb{Z}$, we set $f(t) := s(t)$. We extend this function linearly over each unit interval contained in $[0, b]$ and having the form $[i, i+1]$, $i \in \mathbb{Z}$ or $[[b], b]$. By abusing the terminology, we will refer to these intervals as *integer intervals*. The resulting function f is continuous on the interval $[0, b]$ and maps it

onto the interval $[0, T]$. Moreover, since q is (L, A) -quasigeodesic, the function f maps each integer interval to an interval of the length at most $C = L(A + 1 + 2R_*)$.

Now, for $s \in [0, T]$ find t such that $f(t) = s$. Then t belongs to one of the integer subintervals $[i, i + 1] \subset [0, b]$ (or the interval $[i, b]$, $i = \lfloor b \rfloor$). We have $d(\gamma(i), q(s(i))) \leq R_*$ and, furthermore,

$$d(q(s), q(s(i))) \leq C(L + A).$$

Therefore, $q(s)$ is within distance $\leq C(L + A) = L(A + 1 + 2R_*)(L + A)$ from the geodesic xy . Since

$$R_* \leq D := L(A + 1 + 2R_*)(L + A)$$

(as $L \geq 1$), we conclude that the Hausdorff distance between the geodesic xy and the quasigeodesic q is at most D . \square

11.11. Ideal boundaries

We consider the general notion of the *ideal boundary* defined in Section 3.11.3, in the special case when X is geodesic, δ -hyperbolic and locally compact (equivalently, proper). We start by proving an analogue of Lemma 11.74 in the context of hyperbolic spaces.

LEMMA 11.74. *Suppose that X is geodesic, δ -hyperbolic and locally compact (equivalently, proper). Then for each $x \in X$ and $\xi \in \partial_\infty X$, there exists a geodesic ray ρ with the initial point x , asymptotic to ξ .*

PROOF. Let ρ' be a geodesic ray asymptotic to ξ , with the initial point x_0 . Consider a sequence of geodesic segments $\gamma_n : [0, D_n] \rightarrow X$, connecting p to $x_n = \rho'(n)$, where $D_n = d(x, \rho'(n))$. The δ -hyperbolicity of X implies that the image of γ_n is at Hausdorff distance at most $\delta + \text{dist}(x, x_0)$ from x_0x_n , where x_0x_n is the initial subsegment of ρ' . Combining properness of X with the Arzela-Ascoli theorem, we see that the maps γ_n subconverge to a geodesic ray ρ , $\rho(0) = p$. Clearly, the image of ρ is at Hausdorff distance at most $\delta + \text{dist}(x, x_0)$ from the image of ρ . In particular, ρ is asymptotic to ξ . \square

In view of Lemma 11.74, in order to understand $\partial_\infty X$ it suffices to restrict to the set $\text{Ray}_x(X)$ of geodesic rays in X emanating from $x \in X$. The important difference between this lemma and the one for $CAT(0)$ spaces (Lemma 11.74) is that the ray ρ now may not be unique. Nevertheless we will continue to use the notation $x\xi$, which now means that $x\xi$ is *one* of the geodesic rays with the initial point x and asymptotic to ξ . This abuse of notation is harmless in view of the following lemma which generalizes the *fellow-travelling property* for geodesic segments in hyperbolic spaces.

LEMMA 11.75 (Asymptotic rays are uniformly close). *Let ρ_1, ρ_2 be asymptotic geodesic rays in X with the common initial point $\rho_1(0) = \rho_2(0) = x$. Then for each t ,*

$$d(\rho_1(t), \rho_2(t)) \leq 2\delta.$$

PROOF. Suppose that the rays ρ_1, ρ_2 are within distance $\leq C$ from each other. Take T much larger than t . Then (since the rays are asymptotic) there exists $S \in \mathbb{R}_+$ such that

$$d(\rho_1(T), \rho_2(S)) \leq C.$$

By δ -thinness of the triangle $\Delta(p, \rho_1(T), \rho_2(S))$, the point $\rho_1(t)$ is within distance $\leq \delta$ from a point either on $p\rho_2(S)$ or on $\rho_1(T)\rho_2(S)$. Since the length of $\rho_1(T)\rho_2(S)$ is $\leq C$ and T is much larger than t , it follows that there exists t' such that

$$\text{dist}(\rho_1(t), \rho_2(t')) \leq \delta.$$

By the triangle inequality, $|t - t'| \leq \delta$ and, hence, $\text{dist}(\rho_1(t), \rho_2(t)) \leq 2\delta$. \square

Our next goal is to extend the topology τ defined on $\partial_\infty X$ (i.e. the quotient topology of the compact-open topology on the set of all rays in X , see Section 3.11.3) to a topology on the union $\bar{X} = X \cup \partial_\infty X$. There are several natural ways to do so, all resulting in the same topology, which is a compactification of X by its ideal boundary $\partial_\infty X$.

Shadow topology $\mathcal{T}_{x,k}$ on \bar{X} . Our next goal is to topologize \bar{X} and to describe some basic properties of this topology. We fix a point $x \in X$ and a number $k \geq 3\delta$. For each ideal boundary point $\xi \in \partial_\infty X$ we fix a geodesic ray $\rho = x\xi$ asymptotic to ξ . We define the topology $\mathcal{T}_{x,k}$ on \bar{X} by declaring that its basis at points $z \in X$ consists of open metric balls $B(z, r)$, $r > 0$, and defining basic neighborhoods $U_y = U_{x,y,k}(\xi)$ at points $\xi \in \partial_\infty X$ as

$$U_y = \{z \in \bar{X} : \forall xz, xz \cap B(y, k) \neq \emptyset\},$$

where $y = \rho(t)$, $t \geq 0$.




FIGURE 11.9. Shadow topology.

Note that the requirement in this definition is that *each* geodesic segment or a ray from x to z intersects the open ball $B(y, r)$. We need to check that the collection of basic sets we defined is indeed a basis of topology. It follows from Lemma 11.75 that ξ belongs to $U_y(\xi)$ for each $y = \rho(t)$. Furthermore, δ -hyperbolicity of X implies that for every $t \geq 0$

$$(11.6) \quad U_{\rho(t')} \subset U_{\rho(t)},$$

provided that t' is at least $t + k + \delta$. Therefore,

$$U_{y_3} \subset U_{y_1} \cap U_{y_2}, y_i = \rho(t_i), t_3 = \max(t_1, t_2) + k + \delta.$$

We next have to verify that each basic set is open, more precisely, each point $u \in U_y$ is contained in a basic set $U_z \subset U_y$. Suppose that $u \in U_y \cap X$ and $u_n \in X$ is a sequence converging to u . Assume for a moment that for each n there exists a

geodesic segment xu_n disjoint from the open ball $B(y, k)$. Then a subsequential limit of xu_n would connect x to u and also avoid the ball $B(y, k)$. (Here we are using properness of X .) However, this would imply that $u \notin U_y$, which is a contradiction.

The proof for boundary points is similar. Suppose that $\xi \in U_y(\eta) \cap \partial_\infty X$; let $\rho = x\xi$ be a ray connecting x to ξ . We again assume that there is a sequence t_n diverging to infinity, points $u_n \in B(\rho(t_n), k)$ and geodesic segments $\gamma_n = xu_n$ which all avoid the open ball $B(y, r)$. By δ -thinness of geodesic triangles, each segment γ_n is contained in the $k + \delta$ -neighborhood of the ray ρ . Therefore, after passing to a subsequence, we obtain a limit ray γ (of the segments γ_n), which is asymptotic to ξ and which also avoids the ball $B(y, k)$. However, this contradicts the assumption that ξ belongs to U_y .

To summarize, we now have a topology $\mathcal{T}_{x,k}$ on the set \bar{X} .

LEMMA 11.76. *With respect to the topology $\mathcal{T}_{x,k}$:*

1. X is an open and dense subset of \bar{X} .
2. \bar{X} is first countable.
3. \bar{X} is Hausdorff.

PROOF. 1. Openness of X is clear. Density of X follows from the fact that for each $\xi \in \partial_\infty X$ the sequence $(\rho(n))_{n \in \mathbb{N}}$ converges to ξ , where $\rho = x\xi$.

2. The first countability is clear at the points $x \in X$; at the points $\xi \in \partial_\infty X$, the first countability follows from the fact that the sets

$$U_{\rho(n)}, n \in \mathbb{N},$$

form a basis of topology at ξ , see (11.6).

3. We will check that any two distinct points $\xi_1, \xi_2 \in \partial_\infty X$ have disjoint neighborhoods and leave the other cases as an exercise to the reader. Since the geodesic rays $x\xi_i$, $i = 1, 2$, diverge, there exist points $y_i \in x\xi_i$ such that

$$\text{dist}(y_i, x\xi_{3-i}) > k + \delta, \quad i = 1, 2.$$

We claim that the basic neighborhoods U_{y_1}, U_{y_2} are disjoint. Otherwise, there exists $z \in xu \cap B(y_2, k)$ and a geodesic xz which intersects $B(y, k)$. Since the triangle $\Delta(x, y_2, z)$ is δ -thin, it follows that the point $xu \cap B(y, k)$ is within distance $\leq \delta$ from the subsegment xy_2 of $x\xi_2$. However, this contradicts the assumption that the minimal distance from y_1 to $x\xi_2$ is $> k + \delta$. Interchanging the roles of the points y_1, y_2 , we conclude that $U_{y_1} \cap U_{y_2} = \emptyset$. \square

Consider the set $\text{Geo}_x(X)$ consisting of geodesics in X (finite or half-infinite) emanating from x . In order to ensure that all maps in $\text{Geo}_x(X)$ have the same domain, we extend each geodesic segment $\gamma : [0, T] \rightarrow X$ by the constant map to the half-line $[T, \infty)$. We equip $\text{Geo}_x(X)$ with the compact-open topology (equivalently, the topology of uniform convergence on compacts). The space $\text{Geo}_x(X)$ is compact by the Arzela–Ascoli theorem. Since X is Hausdorff, so is $\text{Geo}_x(X)$. There is a natural quotient map $\epsilon : \text{Geo}_x(X) \rightarrow \bar{X}$ which sends a finite geodesic or a geodesic ray emanating from x to its terminal point in \bar{X} :

$$\epsilon : xy \mapsto y, \quad y \in \bar{X}.$$

LEMMA 11.77. *The map $\epsilon : \text{Geo}_x(X) \rightarrow \bar{X}$ is continuous.*

PROOF. The statement is clear for finite geodesic segments. Let $\xi \in \partial_\infty X$ be an ideal boundary point and let γ_n denote a sequence of geodesic segments/rays,

$\gamma_n = xx_n$, $x_n \in \bar{X}$, such that

$$\lim_{n \rightarrow \infty} \gamma_n = \gamma,$$

where γ is a ray asymptotic to ξ . We claim that the sequence $\epsilon(ga_n) = x_n$ converges to ξ in the topology $\mathcal{T}_{x,k}$. Pick a point y on the geodesic ray $x\xi$ (which could be different from γ). Since the rays γ and ρ are within distance $\leq 2\delta$ from each other, and the convergence $\gamma_n \rightarrow \gamma$ is uniform on compacts, for every there exists N such that for all $n > N$, the intersection

$$\gamma_n \cap B(y, 3\delta) \subset \gamma_n \cap B(y, k)$$

is non-empty. Hence, x_n belongs to U_y . \square

COROLLARY 11.78. 1. ϵ is a closed map. In particular, $\epsilon : Geo_x(X) \rightarrow (\bar{X}, \mathcal{T}_{x,k})$ is a quotient map.

2. $(\bar{X}, \mathcal{T}_{x,k})$ is a compact topological space.
3. $(\bar{X}, \mathcal{T}_{x,k})$ is a compactification of X .

PROOF. 1. The statement follows from the fact that $Geo_x(X)$ is compact and $(\bar{X}, \mathcal{T}_{x,k})$ is Hausdorff.

2. Continuous image of a compact topological space is again compact.

3. This part follows from openness and density of X in \bar{X} combined with compactness of \bar{X} . \square

COROLLARY 11.79. 1. For all $k_1, k_2 \geq 3\delta$, the topologies $\mathcal{T}_{x,k_1}, \mathcal{T}_{x,k_2}$ are equal.

2. The topology $\mathcal{T}_{x,k}$ is independent of the choice of geodesic rays $x\xi$ used to define $\mathcal{T}_{x,k}$

PROOF. For all different choices of k 's and the rays, the topologies are the quotient topologies of $Geo_x(X)$ with respect to the same quotient map ϵ . \square

The topology on \bar{X} is independent of the choice of a base-point x :

LEMMA 11.80. For all $x_1, x_2 \in X$, $\mathcal{T}_{x_1,k} = \mathcal{T}_{x_2,k}$.

PROOF. The equality of two topologies is clear at the points of X . Consider, therefore, an ideal boundary point $\xi \in \partial_\infty X$. We use a geodesic segment x_1x_2 and geodesic rays $x_i\xi$ to form a generalized geodesic triangle $\Delta(x_1, x_2, \xi)$ in X . Since this triangle is 2δ -thin, we pick points $y_i \in x_i\xi$ within distance $\leq 2\delta$ from each other. Then each $u \in B(y_1, k)$ is contained in the ball $B(y_2, k + \delta)$. Therefore, each basic neighborhood $U_{y_2,k}$ of ξ in the topology $\mathcal{T}_{x_2,k}$ is contained in the basic neighborhood $U_{y_2,k+\delta}$ of ξ in the topology $\mathcal{T}_{x_1,k+\delta}$. Hence, the topology $\mathcal{T}_{x_2,k}$ is finer than $\mathcal{T}_{x_1,k+\delta} = \mathcal{T}_{x_1,k}$. Switching the roles of x_1 and x_2 , we conclude that $\mathcal{T}_{x_1,k} = \mathcal{T}_{x_2,k}$. \square

In view of these basic results, from now on, we will omit the subscripts in the notation for the topology on \bar{X} . Using the identification of \bar{X} with the quotient space of $Geo_x(X)$, we can also give an alternative description of converging sequences in \bar{X} .

LEMMA 11.81. For a sequence (x_n) in \bar{X} and a point $\xi \in \partial_\infty X$, the following are equivalent:

1. $\lim_{n \rightarrow \infty} x_n = \xi$.
2. Every convergent subsequence in xx_n converges to a ray asymptotic to ξ .

PROOF. (1) \Rightarrow (2). Suppose first that $\lim_{n \rightarrow \infty} x_n = x$. Let $\gamma_m = xx_{n_m}$ be a convergent sequence of segments/rays and let γ be their limiting ray in $Geo_x(X)$. Pick an arbitrary point $y = \rho(t)$ on the ray $\rho = x\xi$. Since $\lim_{n \rightarrow \infty} x_n = \xi$, the intersections $\gamma_n \cap B(y, k) \neq \emptyset$ for all sufficiently large n . Hence, each subsegment $\gamma_n([0, T])$ is contained in the $k + \delta$ -neighborhood of the ray $x\xi$. Since this holds for all T , we conclude that the limit ray γ is also contained in the $k + \delta$ -neighborhood of $x\xi$. It follows that the ray γ is asymptotic to ξ .

(2) \Rightarrow (1). After passing to subsequences, we can assume that the sequence xx_n converges to a ray $\rho = x\xi$ and that the sequence (x_n) converges to an ideal boundary point η . Continuity of the map ϵ now implies that $\epsilon(\rho) = \xi$. \square

We owe the following remark to Bernhard Leeb:

REMARK 11.82. Even if a sequence (x_n) converges, this does not imply that there exists a convergent sequence of geodesic rays xx_n .

We compute two examples of compactifications and ideal boundaries of hyperbolic spaces.

1. Suppose that $X = \mathbb{H}^n$ is the real-hyperbolic space. We claim that \bar{X} is naturally homeomorphic to the closed ball \mathbb{D}^n , where we use the unit ball model of \mathbb{H}^n . Let o denote the center of the unit ball \mathbf{B}^n . The map $\epsilon : Geo_o(X) \rightarrow \mathbb{D}^n$ is a bijection. The fact that the map ϵ restricts to a homeomorphism $X \rightarrow \mathbb{H}^n$ is clear. Bicontinuity of this map at the points of $\partial_\infty X$ follows from the fact that a sequence $\gamma_n \in Geo_o(X)$ converges to a geodesic ray ρ iff

$$\lim_{n \rightarrow \infty} \gamma'_n(0) = \rho'(0).$$

2. Suppose that X is a simplicial tree of finite constant valence $val(X) \geq 3$, equipped with the standard metric. Fix a vertex $p \in X$. Since X is a $CAT(-\infty)$ -space, the map $\epsilon : Geo_p(X) \rightarrow \bar{X}$ is a bijection, hence, a homeomorphism (in view of the quotient topology on \bar{X}).

We that $\partial_\infty X$ is homeomorphic to the Cantor set. Since we know that $\partial_\infty X$ is compact and Hausdorff, it suffices to verify that $\partial_\infty X$ is totally disconnected and contains no isolated points. Let $\rho \in Ray_p(X)$ be a ray. For each n pick a ray $\rho_n \in Ray_p(X)$ which coincides with ρ on $[0, n]$, but $\rho_n(t) \neq \rho(t)$ for all $t > n$ (this is where we use the assumption that $val(X) \geq 3$). It is then clear that

$$\lim_{n \rightarrow \infty} \rho_n = \rho$$

uniformly on compacts. Hence, $\partial_\infty X$ has no isolated points. Recall that for $k = \frac{1}{2}$, we have open sets $U_{n,k}(\rho)$ forming a basis of neighborhoods of ρ . We also note that each $U_{n,k}(\rho)$ is also closed, since (for a tree X as in our example) it is also given by

$$\{\rho' : \rho(t) = \rho'(t), t \in [0, n]\}.$$

Therefore, $\partial_\infty X$ is totally-disconnected as for any pair of distinct points $\rho, \rho' \in Ray_p(X)$, they have open, closed and disjoint neighborhoods $U_{n,k}(\rho), U_{n,k}(\rho')$. Thus, $\partial_\infty X$ is compact, Hausdorff, perfect, consists of at least 2 points and is totally-disconnected. Therefore, $\partial_\infty X$ is homeomorphic to the Cantor set.

We now return to the discussion of ideal boundaries of arbitrary proper geodesic hyperbolic spaces.

LEMMA 11.83 (The visibility property). *Let X be a proper geodesic Gromov–hyperbolic space. Then for each pair of distinct points $\xi, \eta \in \partial_\infty X$ there exists a geodesic γ in X which is asymptotic to both ξ and η .*

PROOF. Consider geodesic rays ρ, ρ' emanating from the same point $p \in X$ and asymptotic to ξ, η respectively. Since $\xi \neq \eta$ (Corollary 11.67), for each $R < \infty$ the set

$$K(R) := \{x \in X : \text{dist}(x, \rho) \leq R, \text{dist}(x, \rho') \leq R\}$$

is compact. Consider the sequences $x_n := \rho(n), x'_n := \rho'(n)$ on ρ, ρ' respectively. Since the triangles $T(p, x_n, x'_n)$ are δ -thin, each segment $\gamma_n := x_n x'_n$ contains a point within distance $\leq \delta$ from both $p x_n, p x'_n$, i.e. $\gamma_n \cap K(\delta) \neq \emptyset$. Therefore, by the Arzela–Ascoli theorem, the sequence of geodesic segments γ_n subconverges to a complete geodesic γ in X . Since

$$\gamma \subset \mathcal{N}_\delta(\rho \cup \rho'),$$

it follows that γ is asymptotic to ξ and to η . \square

EXERCISE 11.84. Suppose that X is δ -hyperbolic. Show that there are no complete geodesics γ in X such that

$$\lim_{n \rightarrow \infty} \gamma(-n) = \lim_{n \rightarrow \infty} \gamma(n).$$

Hint: Use the fact that geodesic bigons in X are δ -thin.

EXERCISE 11.85 (Ideal bigons are 2δ -thin). Suppose that $\alpha, \beta : \mathbb{R} \rightarrow X$ are geodesics in X which are both asymptotic to points $\xi, \eta \in \partial_\infty X$. Then

$$\text{dist}_{\text{Haus}}(\alpha, \beta) \leq 2\delta.$$

Hint: For $n \in \mathbb{N}$ define

$$z_n, w_n \in \beta(\mathbb{R})$$

to be the nearest points to $x_n = \alpha(n), y_n = \alpha(-n)$. Let $x_n y_n, z_n w_n$ be the subsegments of α, β between x_n, y_n and y_n, z_n respectively. Now use the fact that the quadrilateral

$$x_n y_n \cup y_n w_n \cup w_n z_n \cup z_n x_n$$

is 2δ -thin.

Triangles in \bar{X} . We now generalize (geodesic) triangles in X to triangles with some vertices in $\partial_\infty X$, similarly to the definitions made in Section 4.4. Namely a (generalized) triangle in \bar{X} is a concatenation of geodesics connecting three points A, B, C in \bar{X} ; geodesics are now allowed to be finite, half-infinite and infinite. The points A, B, C are called the vertices of the triangle. As in the case of \mathbb{H}^n , we do not allow two ideal vertices of a triangle T to coincide. By abusing the terminology, we will again refer to such generalized triangles as *hyperbolic triangles*. As with hyperbolic geodesics, we continue to use the notation $T(A, B, C)$ and $\Delta(A, B, C)$ for geodesic triangles with the vertices A, B and C , even though geodesics connecting the vertices are not unique.

An ideal triangle is a triangle where all three vertices are in $\partial_\infty X$. We topologize the set $\text{Tri}(X)$ of hyperbolic triangles in X by the compact-open topology on the set of their geodesic edges. Given a hyperbolic triangle $T = T(A, B, C)$ in X , we find a sequence of finite triangles $T_i \subset X$ whose vertices converge to the respective vertices of T . Passing to a subsequence if necessary and taking the limit of the sides

of the triangles T_i , we obtain geodesics connecting the vertices A, B, C of T . The resulting triangle T' , of course, need not be equal to T (since geodesics connecting points in \bar{X} are not, in general, not be unique), however, in view of Exercise 11.85, sides of T' are within distance $\leq 2\delta$ from the respective sides of T . We will say that the sequence of triangles T_i *coarsely converges* to the triangle T (cf. Definition 8.33).

EXERCISE 11.86. Every (generalized) hyperbolic triangle T in X is 5δ -thin. In particular,

$$\text{minsize}(T) \leq 4\delta.$$

Hint: Use a sequence of finite triangles coarsely converging to T and the fact that finite triangles are δ -thin.

Centroids of triangles with ideal vertices. We now return to the discussion of proper geometric metric spaces X which are δ -hyperbolic in the sense of Rips. Exercise 11.86 allows one to define *centroids* of triangles T in \bar{X} . As in Definition 11.60 we say that a point $p \in X$ is an R -centroid of T if p is within distance $\leq R$ from all three sides of T . Furthermore, we will say that p is a *centroid* of T if

$$d(p, \tau_i) \leq 5\delta, i = 1, 2, 3.$$

LEMMA 11.87. *The distance between any two R -centroids of a triangle T is at most*

$$r(R, \delta) = 4R + 32\delta.$$

PROOF. Let p, q be R -centroids of T . We coarsely approximate T by a sequence of finite triangles $T_i \subset X$. Then for every $\epsilon > 0$, for all sufficiently large i , the points p, q are $R + 2\delta + \epsilon$ -centroids of T_i . Therefore, by Corollary 11.61 applied to the triangles T_i ,

$$d(p, q) \leq \phi(R + 2\delta + \epsilon) = 4(R + 2\delta + \epsilon) + 28\delta = 4R + 32\delta + 2\epsilon$$

Since this holds for every $\epsilon > 0$, we conclude that $d(p, q) \leq 4R + 32\delta$. \square

NOTATION 11.88. Given a topological space Z , we let $\text{Trip}(Z)$ denote the set of ordered triples of pairwise distinct elements of Z , equipped with the subspace topology induced from Z^3 .

We define the correspondence

$$\text{center} : \text{Trip}(\partial_\infty X) \rightarrow X$$

which sends every triple of distinct points in $\partial_\infty X$ first to the set of ideal triangles T that they span and then to the set of centroids of these ideal triangles. Lemma 11.87 implies:

COROLLARY 11.89. *For every $\xi \in \text{Trip}(\partial_\infty X)$,*

$$\text{diam}(\text{center}(\xi)) \leq r(7\delta, \delta) = 60\delta.$$

EXERCISE 11.90. Suppose that γ_n are geodesics in X asymptotic to points $\zeta_n, \eta_n \in \partial_\infty X$ and such that

$$\lim_{n \rightarrow \infty} \zeta_n = \zeta, \quad \lim_{n \rightarrow \infty} \eta_n = \eta, \quad \eta \neq \zeta.$$

Show that the sequence (γ_n) subconverges to a geodesic asymptotic to both ξ and η .

Use this exercise to conclude:

EXERCISE 11.91. If $K \subset \text{Trip}(\partial_\infty X)$ is a compact subset, then $\text{center}(K)$ is a bounded subset of X .

Conversely, prove:

EXERCISE 11.92. Let $B \subset X$ be a bounded subset and $K \subset \text{Trip}(\partial_\infty X)$ be a subset such that $\text{center}(K) \subset B$. Show that K is relatively compact in $\text{Trip}(\partial_\infty X)$. Hint: For every $\xi \in K$, every ideal edge of a triangle spanned by ξ intersects the 5δ -neighborhood of B . Now, use the Arzela-Ascoli theorem.

Loosely speaking, the two exercises show that the correspondence center is coarsely continuous (the image of a compact is bounded) and coarsely proper (the preimage of a bounded subset is relatively compact).

Cone topology. Suppose that X is a proper hyperbolic geodesic metric space. Later on, it will be convenient to use another topology on \bar{X} , called *the cone (or, radial) topology*. This topology is not equivalent to the topology $\text{mathcal{T}}$: With few exceptions, \bar{X} is non-compact with respect to this topology (even if $X = \mathbb{H}^n, n \geq 2$).

DEFINITION 11.93. Fix a base point $p \in X$. We use the metric topology on X and will say that a sequence $x_i \in X$ *conically converges* to a point $\xi \in \partial_\infty X$ if there is a constant R such that $x_i \in \mathcal{N}_R(p\xi)$ and

$$\lim_{i \rightarrow \infty} d(p, x_i) = \infty.$$

A subset $C \subset \bar{X}$ is closed in the conical topology if its intersection with X is closed in the metric topology of X and $C \cap \partial_\infty X$ contains conical limits of sequences in $C \cap X$. We will refer to the resulting topology as the *cone topology* on \bar{X} .

EXERCISE 11.94. If a sequence (x_i) converges to $\xi \in \partial_\infty X$ in the cone topology, then it also converges to ξ in the topology τ on \bar{X} .

As an example, consider $X = \mathbb{H}^n$ in the upper half-space model, $\xi = \mathbf{0} \in \mathbb{R}^{n-1}$, and let L be any vertical hyperbolic geodesic asymptotic to ξ . Then a sequence $x_i \in X$ converges ξ in the cone topology if and only if all the points x_i belong to some Euclidean cone with the axis L , while the Euclidean distance from x_i to $\mathbf{0}$ tends to zero. See Figure 11.10. This explains the name *cone topology*.

EXERCISE 11.95. Suppose that a sequence (x_i) converges to a point $\xi \in \partial_\infty \mathbb{H}^n$ along a horosphere centered at ξ . Show that the sequence (x_i) contains no convergent subsequences in the cone topology on \bar{X} .

11.12. Gromov bordification of Gromov–hyperbolic spaces

The definition of \bar{X} and its topology, used in the previous section, worked fine for geodesic hyperbolic metric spaces. Gromov extended this definition to the case when X is an arbitrary (non-empty) δ -hyperbolic metric space.

Pick a base-point $p \in X$. A sequence (x_n) in X is said to *converge at infinity* if

$$\lim_{(m,n) \rightarrow \infty} (x_m, x_n)_p = \infty.$$




FIGURE 11.10. Convergence in the cone topology.

In particular, for such a sequence,

$$\lim_{n \rightarrow \infty} d(p, x_n) = \infty.$$

Define the relation \sim on sequences converging at infinity by

$$(x_n) \sim (y_n) \iff \lim_{n \rightarrow \infty} (x_n, y_n)_p = \infty.$$

EXERCISE 11.96. 1. Show that \sim is an equivalence relation using Definition 11.24.

2. Show that each sequence (x_n) converging at infinity is equivalent to each subsequence in (x_n) .

3. Show that if $(x_n), (y_n)$ are inequivalent sequences converging at infinity, then

$$\sup_{m,n} (x_m, y_n)_p < \infty.$$

4. Show that for two sequences $(x_m), (y_n)$,

$$\lim_{(m,n) \rightarrow \infty} (x_m, y_n)_p = \infty$$

if and only if

$$\lim_{(m,n) \rightarrow \infty} (x_m, y_n)_q = \infty,$$

for all $q \in X$.

5. Suppose that $\mathbb{N} \rightarrow X, n \mapsto x_n$ is an isometric embedding. Show that the sequence (x_n) converges at infinity.

DEFINITION 11.97. The *Gromov boundary* $\partial_{Gromov} X$ of X consists of equivalence classes of sequences converging at infinity. Given a sequence (x_n) converging at infinity, we will use the notation $[x_n]$ for the equivalence class of this sequences. The union $X \cup \partial_{Gromov} X$ is the *Gromov bordification* of X .

It will be convenient to extend the notation $[x_n]$ for sequences (x_n) which converge in X ; we set

$$[x_n] = \lim_{n \rightarrow \infty} x_n \in X$$

for such sequences.

The Gromov product extends to $X \cup \partial_{\text{Gromov}} X$ by taking limits of Gromov products in X :

1.

$$(\xi, \eta)_p = \inf \liminf_{(m,n) \rightarrow \infty} (x_m, y_n)_p,$$

where the infimum is taken over all sequences $(x_m), (y_n)$ representing ξ and η respectively.

2.

$$(\xi, y)_p = \inf \liminf_{n \rightarrow \infty} (x_n, y)_p,$$

where the sequence (x_n) represents ξ . The infimum in this definition is again taken over all sequences (x_m) representing ξ .

REMARK 11.98. Taking the infimum and \liminf in this definition is, by no means, the only choice. However, all four possible choices in the definition of $(\xi, \eta)_p$ and $(\xi, y)_p$ differ by $\leq 2\delta$, see [Väi05].

EXERCISE 11.99. For points $x, y \in X \cup \partial_{\text{Gromov}} X$ we have

$$(x, y)_p = \inf \{ \liminf_{i \rightarrow \infty} (x_i, y_i)_p \}$$

where the infimum is taken over all sequences $(x_i), (y_i)$ in X such that $x = [x_i], y = [y_i]$.

LEMMA 11.100. For all points $p \in X, x, y, z \in X \cup \partial_{\text{Gromov}} X$ we have the inequality

$$(x, y)_p \geq \min\{(y, z)_p, (z, x)_p\} + \delta.$$

PROOF. For each point x, y, z we consider sequences $(x_i), (y_i), (z_i)$ in X such that $x = [x_i], y = [y_i], z = [z_i]$. We assume that $(x_i), (y_i)$ are chosen so that

$$(x, y)_p = \lim_{i \rightarrow \infty} (x_i, y_i)_p = (x, y)_p.$$

Then for each i we have

$$(x_i, y_i)_p \geq \min\{(y_i, z_i)_p, (z_i, x_i)_p\} + \delta.$$

Then

$$\begin{aligned} (x, y)_p &= \lim_{i \rightarrow \infty} (x_i, y_i)_p \geq \liminf_{i \rightarrow \infty} \min\{(y_i, z_i)_p, (z_i, x_i)_p\} + \delta \geq \\ &\quad \min\{\liminf_{i \rightarrow \infty} (y_i, z_i)_p, \liminf_{i \rightarrow \infty} (z_i, x_i)_p\} + \delta \end{aligned}$$

Now the claim follows from the Exercise 11.99. \square

The Gromov topology on

$$\bar{X} = X \cup \partial_{\text{Gromov}} X$$

is the metric topology on X , while a basis of topology at $\xi \in \partial_{\text{Gromov}} X$ consists of the sets

$$\mathcal{U}_{\xi, R} = \{x \in \bar{X} : (\xi, x)_p > R\}.$$

Thus, a sequence (x_n) in \bar{X} converges to $\xi \in \partial_{\text{Gromov}} X$ if and only if

$$\lim_{n \rightarrow \infty} (x_n, \xi)_p = \infty.$$

LEMMA 11.101. *A sequence (x_n) converges to $\xi \in \partial_{Gromov} X$ if and only if (x_n) converges at infinity and $[x_n] = \xi$.*

PROOF. 1. Suppose that (x_n) converges to $\xi \in \partial_{Gromov} X$. Then

$$\lim_{m \rightarrow \infty} (x_m, \xi)_p = \lim_{m \rightarrow \infty} (x_m, \xi)_p = \infty.$$

By Lemma 11.100, we have

$$(x_m, x_n)_p \geq \min\{(x_m, \xi)_p, (x_n, \xi)_p\} - \delta \rightarrow \infty, \text{ as } n, m \rightarrow \infty.$$

Hence, the sequence (x_n) converges at infinity. Let (y_n) be a sequence in X representing ξ . We claim that $(x_n) \sim (y_n)$. Indeed,

$$(x_n, y_n) \geq \min\{(x_n, \xi)_p, (y_n, \xi)_p\} - \delta \rightarrow \infty_{n \rightarrow \infty}.$$

Thus, $[x_n] = \xi$.

2. Suppose that $[x_n] = \xi$. For each n we pick a sequence $(y_{nm})_{m \in \mathbb{N}}$ representing ξ such that

$$(x_n, \xi)_p = \lim_{m \rightarrow \infty} (x_n, y_{nm})_p.$$

Taking a diagonal subsequence $(y_{n_k, m_k})_{k \in \mathbb{N}}$ which represents ξ , we obtain

$$\lim_{k \rightarrow \infty} (x_{n_k}, \xi)_p = \lim_{k \rightarrow \infty} (x_{n_k}, y_{n_k, m_k})_p = \infty.$$

Since

$$\lim_{k \rightarrow \infty} (x_{n_k}, x_k)_p = \infty,$$

Lemma 11.100 applied to the points x_{n_k}, x_n and ξ implies that

$$\lim_{k \rightarrow \infty} (x_k, \xi)_p = \infty. \quad \square$$

Suppose now that X is a geodesic metric space which is a δ_1 -hyperbolic (in the sense of Rips) and δ_2 -hyperbolic (in Gromov's sense). We define a map

$$h : X \cup \partial_\infty X \rightarrow X \cup \partial_{Gromov} X$$

which is the identity on X and sends $\xi = [\rho]$ in $\partial_\infty X$ to the equivalence class of the sequence $(\rho(n))$.

EXERCISE 11.102. The map h is well-defined, i.e.:

1. If $\rho : \mathbb{R}_+ \rightarrow X$ is a geodesics ray then the sequence $(\rho(n))$ converges at infinity.

2. If two rays ρ_1, ρ_2 are asymptotic then $[\rho_1(n)] = [\rho_2(n)]$.

LEMMA 11.103. *If X is a proper geodesic metric space then the map h is a bijection.*

PROOF. 1. Injectivity of h . Suppose that ρ_1, ρ_2 are rays in X emanating from $p \in X$, such that $[\rho_1(n)] = [\rho_2(n)]$. Set $x_n = \rho_1(n), y_n = \rho_2(n)$ and for each n choose a geodesic $x_n y_n$ in X . For each n let $z_n \in x_n y_n$ be a point within distance $\leq \delta_1$ from both sides $p x_n, p y_n$ of the geodesic triangle $T(p x_n y_n)$. Let $x'_n \in p x_n, y'_n \in p y_n$ be points within distance $\leq \delta_1$ from z_n . Since $[x_n] = [y_n]$,

$$\lim_{n \rightarrow \infty} d(p, z_n) = \infty,$$

which implies that

$$\lim_{n \rightarrow \infty} d(p, x'_n) = \lim_{n \rightarrow \infty} d(p, y'_n) = \infty.$$

Since $d(x'_n, y'_n) \leq 2\delta_1$, in view of δ_1 -thinness of the triangle $T(px'_n y'_n)$, we have that the Hausdorff distance between the geodesic segments px'_n, py'_n is $\leq 2\delta_1$. We conclude that the rays ρ_1, ρ_2 are Hausdorff-close to each other. Hence, the map h is injective.

2. Surjectivity of h . Let (x_n) be a sequence in X converging at infinity. Since X is proper, the sequence of geodesic segments px_n contains a subsequence px_{n_k} which converges (uniformly on compacts) to a geodesic ray $p\xi$ in X . For each k let $x'_{n_k} \in px_{n_k}$ denote a point within distance $\leq \delta$ from $p\xi$, and such that the distance $d(p, x'_{n_k})$ is maximal among all such points. Then

$$[x_n] = [x_{n_k}] = [x'_{n_k}].$$

Let $y_{n_k} \in p\xi$ denote a point within distance $\leq \delta_1$ from x'_{n_k} . Then

$$\lim_{n \rightarrow \infty} d(p, x'_{n_k}) = \infty$$

and $[x'_{n_k}] = [y'_{n_k}]$. It follows that h sends the equivalence class of the ray $p\xi$ to the equivalence class of the sequences $[x_n]$. \square

THEOREM 11.104. *The map h is a homeomorphism.*

PROOF. We will verify continuity of h and h^{-1} at each point $\xi \in \partial_\infty X$. We fix $k \geq 10\delta$ and consider the topology $\mathcal{T}_{p,k}$ of shadow-convergence in \bar{X} .

1. Pick a ray $p\xi = \rho(\mathbb{R}_+)$ asymptotic to ξ . Suppose that (x_n) is a sequence in \bar{X} which shadow-converges to $\xi \in \partial_\infty X$. Then there exists a sequence $t_n \in \mathbb{R}_+$ diverging to infinity, such that for $y_n = \rho(t_n)$, the segment (or a ray) px_n intersects the ball $B(y_n, k)$ at a point x'_n . Since

$$\begin{aligned} [x'_n] &= [y_n] = \xi \\ \lim_{n \rightarrow \infty} (x'_n, \xi) &= \infty, \end{aligned}$$

see Lemma 11.101. On the other hand,

$$(x_n, x'_n)_p = d(p, x'_n) \rightarrow \infty.$$

The inequality

$$(x_n, \xi)_p \geq \min\{(x_n, x'_n)_p, (x'_n, \xi)_p\} + \delta_2$$

now implies that

$$\lim_{n \rightarrow \infty} (x_n, \xi)_p = \infty.$$

Therefore, the map h is continuous.

2. Suppose that $x_n \in X$ is a sequence converging at infinity, $[x_n] = \eta \in \partial_{Gromov} X$. We let $x\xi$ denote a geodesic ray in X with $\xi = h^{-1}(\eta)$. We will show that (x_n) shadow-converges to ξ . For each n pick a geodesic ray $x_n\xi$. Since

$$(x_n, \xi)_p \rightarrow \infty,$$

the minimal distance from p to $x_n\xi$ diverges to ∞ . Let $z_n \in x_n\xi$ be a point within distance $\leq 4\delta_1$ from

$$x_n p \cup p\xi.$$

Let $y_n \in p\xi$, $u_n \in px_n$ be points within distance $\leq 4\delta_1$ from z_n . Since

$$\lim_{n \rightarrow \infty} d(p, z_n) = \lim_{n \rightarrow \infty} d(p, y_n) = \infty$$

and $u_n \in B(y_n, 8\delta_1)$, we conclude that the sequence (x_n) shadow-converges to ξ . The proof in the case when x_n is a sequence in $\partial_{Gromov} X$ converging to η is similar

and is left to the reader. (Alternatively, continuity of h^{-1} follows from Lemma 1.18.) \square

In view of this theorem, we will be identifying the visual and Gromov ideal boundaries of X .

11.13. Boundary extension of quasiisometries of hyperbolic spaces

The goal of this section is to explain how quasiisometries of Rips-hyperbolic spaces extend to their ideal boundaries.

11.13.1. Extended Morse Lemma. We first extend the Morse lemma to the case of quasigeodesic rays and complete geodesics.

LEMMA 11.105 (Extended Morse Lemma). *Suppose that X is a proper δ -hyperbolic geodesic space. Let ρ be an (L, A) -quasigeodesic ray or a complete (L, A) -quasigeodesic. Then there is ρ^* , which is either a geodesic ray or a complete geodesic in X , such that the Hausdorff distance between the images of ρ and ρ^* is at most $D(L, A, \delta)$. Here D is the function which appears in the Morse lemma.*

Moreover, there are two functions $s = s(t), s^ = s^*(t)$ such that:*

$$(11.7) \quad L^{-1}t - B \leq s \leq Lt + B,$$

$$(11.8) \quad L^{-1}(t - B) \leq s^* \leq L(t + B),$$

and for every t ,

$$d(\rho(t), \rho^*(s)) \leq D, \quad d(\rho^*(t), \rho(s^*)) \leq D.$$

Here $B = A + D$.

PROOF. We will consider only the case of quasigeodesic rays $\rho : [0, \infty) \rightarrow X$ as the other case is similar. For each i we define the finite quasigeodesic

$$\rho_i := \rho|_{[0, i]}$$

and the geodesic segment $\rho_i^* = px_i$, connecting the points $p = \rho(0), x_i = \rho(i)$. According to the Morse lemma,

$$\text{dist}_{\text{Haus}}(\rho_i, \rho_i^*) \leq D(L, A, \delta).$$

By properness of X , the sequence of geodesic segments ρ_i^* subconverges to a complete geodesic ray ρ^* . It is clear that

$$\text{dist}_{\text{Haus}}(\rho, \rho^*) \leq D(L, A, \delta).$$

The estimates (11.7) and (11.8) follow from the inequalities (11.3) and (11.4) in the case of finite geodesic segments. \square

COROLLARY 11.106. *If ρ is a quasigeodesic ray as in the above lemma, there exists a point $\xi \in \partial_\infty X$ such that $\lim_{t \rightarrow \infty} \rho(t) = \xi$.*

PROOF. According to Lemma 11.105, the quasigeodesic ray ρ is close to a geodesic ray $\rho^* = p\xi$. Since $d(\rho(t), p\xi) \leq D$ for all t , it follows that

$$\lim_{t \rightarrow \infty} \rho(t) = \xi. \quad \square$$

We will refer to the point η as $\rho(\infty)$. Note that if ρ' is another quasigeodesic ray Hausdorff-close to ρ , then $\rho(\infty) = \rho'(\infty)$.

Below is another useful application of the Extended Morse Lemma. Given a geodesic γ in X we let $\pi_\gamma : X \rightarrow \gamma$ denote a nearest-point projection.

PROPOSITION 11.107 (Quasiisometries commute with projections). *There exists $C = C(L, A, \delta)$ such that the following holds. Let X, X' be proper δ -hyperbolic geodesic metric spaces and let $f : X \rightarrow X'$ be an (L, A) -quasiisometry. Let α be a (finite or infinite) geodesic in X , and let $\beta \subset X'$ be a geodesic which is $D(L, A, \delta)$ -close to $f(\alpha)$. Then the map f almost commutes with the nearest-point projections π_α, π_β :*

$$d(f\pi_\alpha(x), \pi_\beta f(x)) \leq C, \quad \forall x \in X.$$

PROOF. For each (finite or infinite) geodesic $\gamma \subset X$ consider the triangle $\Delta = \Delta_{x,\gamma}$ where one side of Δ is γ and x is a vertex: The other two sides of Δ are geodesics connecting x to the (finite or ideal) end-points of γ . We use the same definition for triangles in X' .

Let $c = \text{center}(\Delta) \in \gamma$ denote a centroid of Δ : The distance from c to each side of Δ is $\leq 6\delta$. By Corollary 11.62,

$$d(c, \pi_\gamma(x)) \leq 21\delta$$

for all $x \in X$. Consider now a geodesic $\alpha \subset X$, a point $x \in X$ and its image $y = f(x)$ in X' . We let β be a geodesic in X' within distance $\leq D(L, A, \delta)$ from $f(\alpha)$.




FIGURE 11.11. Quasiisometries almost commute with projections.

Applying f to the centroid $c_{x,\alpha} = c(\Delta_{x,\alpha})$, we obtain a point $a \in X'$ whose distance to each side of the quasigeodesic triangle $f(\Delta_{x,\alpha})$ is $\leq 2\delta L + A$. Hence, the distance from a to each side of the geodesic triangle $\Delta_{y,\beta}$ is at most $R := 2\delta L + A + D(L, A, \delta)$. Hence, a is an R -centroid of $\Delta_{y,\beta}$. By Lemma 11.87, it follows that the distance from a to the centroid $c_{y,\beta} = \text{center}(\Delta_{y,\beta})$ is at most $8R + 32\delta$. Since $d(\pi_\beta(y), c(\Delta_{y,\beta})) \leq 21\delta$, we obtain:

$$d(f(\pi_\alpha(x)), \pi_\beta f(x)) \leq C := 21\delta + 8R + 27\delta + 21\delta L + A. \quad \square$$

11.13.2. The extension theorem. We are now ready to prove the main theorem of this section, which is a fundamental fact of the theory of hyperbolic spaces:

THEOREM 11.108 (Extension Theorem). *Suppose that $f : X \rightarrow X'$ is a quasi-isometry between two Rips-hyperbolic proper metric spaces. Then f admits a homeomorphic extension $f_\infty : \partial_\infty X \rightarrow \partial_\infty X'$. This extension is such that the map $\bar{f} = f \cup f_\infty$ is continuous at each point $\eta \in \partial_\infty X$. The extension satisfies the following functoriality properties:*

1. For every pair of quasiisometries $f_i : X_i \rightarrow X_{i+1}, i = 1, 2$, we have

$$(f_2 \circ f_1)_\infty = (f_2)_\infty \circ (f_1)_\infty.$$

2. For every pair of quasiisometries $f_1, f_2 : X \rightarrow X'$ satisfying $\text{dist}(f_1, f_2) < \infty$, we have

$$(f_2)_\infty = (f_1)_\infty.$$

PROOF. First, we construct the extension f_∞ .

Given $\xi \in \partial_\infty X$, we pick a sequence (x_n) in X representing ξ . We claim that the sequence $(y_n), y_n = f(x_n)$, converges at infinity in Y . Indeed, according to Lemma 11.22, for any pair of indices m, n , we have

$$\begin{aligned} (x_m, x_n)_p &\leq \text{dist}(p, x_m x_n) \leq L \text{dist}(f(p), f(x_m x_n)) + LA \leq \\ &LD \text{dist}(f(p), y_m y_n) + LA \leq LD((y_m, y_n)_{f(p)} + 2\delta) + LA \end{aligned}$$

where $D = D(L, A, \delta)$ is the constant from the Morse Lemma. Therefore,

$$\lim_{m, n \rightarrow \infty} (y_m, y_n)_{f(p)} = \infty.$$

The same argument shows that if $[x_n] = [x'_n] = \xi$ then $[f(x_n)] = [f(x'_n)]$. Therefore, we set $f_\infty(\xi) := [y_n]$. We next show that for each $\xi \in \partial_\infty X$ the restriction of \bar{f} to $X \cup \{\xi\}$ is continuous at ξ . Since \bar{X} and \bar{Y} are first countable, it suffices to verify sequential continuity. We have that $x_n \in X$ converges to ξ if and only if $[x_n] = \xi$. Since $\xi' = f_\infty(\xi) = [f(x_n)]$, it follows that the sequence $(f(x_n))$ converges to ξ' . Therefore, the restriction of \bar{f} to $X \cup \{\xi\}$ is continuous at ξ . Since \bar{X} is compact and Hausdorff, it is regular. Part 2 of Lemma 1.18 implies continuity of $\bar{f} : \bar{X} \rightarrow \bar{Y}$ at each $\xi \in \partial_\infty X$.

We next check the functoriality properties (1) and (2) of the extension maps. Suppose that $\xi = [x_n] \in \partial_\infty X$ (where (x_n) is a sequence in X converging at infinity), $\eta = (f_1)_\infty(\xi)$. Then

$$\eta = [f_1(x_n)], \quad (f_2)_\infty(\eta) = [f_2 \circ f_1(x_n)] = (f_2 \circ f_1)_\infty(\xi).$$

This implies Property 1. To verify Property 2, note that for each $\xi = [x_n]$ is also represented by the sequence $(y_n), y_n = f(x_n)$, since the distances $d(x_n, y_n)$ are uniformly bounded.

Lastly, we verify that f_∞ is a homeomorphism. Let g be a coarse inverse of $f : X \rightarrow X'$; this coarse inverse also has a continuous extension

$$g_\infty : \partial_\infty X' \rightarrow \partial_\infty X.$$

Since $\text{dist}(g \circ f, \text{Id}_X) < \infty$ and $\text{dist}(f \circ g, \text{Id}_{X'}) < \infty$ by the functoriality properties, we obtain:

$$\begin{aligned} \text{Id}_{X'} &= (f \circ g)_\infty = f_\infty \circ g_\infty, \\ \text{Id}_X &= (g \circ f)_\infty = g_\infty \circ f_\infty. \end{aligned}$$

Hence, g_∞ is the continuous inverse of f_∞ , and

$$f_\infty : \partial_\infty X \rightarrow \partial_\infty X'$$

is a homeomorphism. \square

EXERCISE 11.109. Suppose that f is merely a QI embedding $X \rightarrow X'$. Show that the continuous extension f_∞ given by $[x_n] \mapsto [f(x_n)]$ is injective.

HISTORICAL REMARK 11.110. The above extension theorem was first proven by Efremovich and Tikhomirova in [ET64] for the real-hyperbolic space and, soon afterwards, reproved by Mostow [Mos73]. We will see in Chapter 22 that the homeomorphisms f_∞ are *quasimoebius*, in particular, they enjoy certain regularity properties which are critical for proving QI rigidity theorems in the context of hyperbolic groups and spaces.

The next lemma is a simple but useful corollary of Theorem 11.108 (the functoriality part):

COROLLARY 11.111. *Suppose that f, g, h are quasiisometries of X such that $\text{dist}(h, g \circ f) < \infty$. Then*

$$h_\infty = g_\infty \circ f_\infty.$$

In particular, if $f : X \rightarrow X'$ is a quasiisometry quasiequivariant with respect to isometric group actions $G \curvearrowright X, G \curvearrowright X'$, then f_∞ is also G -equivariant.

We thus obtained a functor from quasiisometries between Rips-hyperbolic spaces to homeomorphisms between their boundaries.

The following lemma is a “converse” to the second functoriality property in Theorem 11.108:

LEMMA 11.112. *Let X and Y be proper geodesic δ -hyperbolic spaces. In addition we assume that there exists $R < \infty$ such that every $x \in X$ is an R -centroid of an ideal triangle T_x in X . Then quasiisometries $f, f' : X \rightarrow Y$ with equal extension maps $f_\infty = f'_\infty$ are uniformly close to each other. More precisely, there exists $D(L, A, R, \delta)$ such that each pair of (L, A) -quasiisometries $f, f' : X \rightarrow Y$ with $f_\infty = f'_\infty$, satisfies:*

$$\text{dist}(f, f') \leq D(L, A, R, \delta).$$

PROOF. By Lemma 11.105, for each $x \in X$ the points $y = f(x), y' = f'(x)$ are C -centroids of an ideal geodesic triangle $S \subset Y$ whose ideal vertices are the images of the ideal vertices of T_x under f_∞ . Here $C = LR + A + D(L, A, \delta)$. Lemma 11.87 implies that C -centroids of ideal triangles are uniformly close to each other:

$$d(y, y') \leq r(C, \delta).$$

We conclude that

$$d(f(x), f'(x)) \leq D(L, A, R, \delta) = 2(LR + A) + r(C, \delta). \quad \square$$

Suppose that X is hyperbolic and $\partial_\infty X$ contains at least 3 points. Then X has at least one ideal triangle and, hence, at least one centroid of an ideal triangle. If, in addition, X is quasihomogeneous, then, for some $R < \infty$, every $x \in X$ is an R -centroid of an ideal triangles in X . Thus, the above lemma applies to the real-hyperbolic space, more generally, all negatively curved symmetric space, and, as we will sees soon, all non-elementary hyperbolic groups.

EXERCISE 11.113. Suppose that X is a complete simply-connected Riemannian manifold of dimension > 1 and sectional curvature $\leq a < 0$. Show that every point $x \in X$ is an R -centroid of an ideal triangle, for some uniform R .

EXAMPLE 11.114. The line $X = \mathbb{R}$ is 0-hyperbolic, its ideal boundary consists of two points. Take a translation $f : X \rightarrow X$, $f(x) = x + a$. Then f_∞ is the identity map of $\{-\infty, \infty\}$ but there is no bound on the distance from f to the identity.

Here is an important corollary of Theorem 11.108 and Lemma 11.112:

COROLLARY 11.115. *Let X be a Rips-hyperbolic space. Then the map $f \mapsto f_\infty$, sending quasiisometries of X to homeomorphisms of $\partial_\infty X$, descends to a homomorphism $QI(X) \rightarrow \text{Homeo}(X)$. Furthermore, under the hypothesis of Lemma 11.112, this homomorphism is injective.*

In Section 22.5 we will identify the image of this homomorphism in the case of the real-hyperbolic space \mathbb{H}^n , it will be the subgroup of $\text{Homeo}(\mathbb{S}^{n-1})$ consisting of *quasimobius* homeomorphisms.

11.13.3. Boundary extension and quasiactions. In view of Corollary 11.115, we have

COROLLARY 11.116. *Suppose that X is a Rips-hyperbolic space. Then every quasiaction ϕ of a group G on X extends (by $g \mapsto \phi(g)_\infty$) to an action ϕ_∞ of G on $\partial_\infty X$ by homeomorphisms.*

LEMMA 11.117. *Suppose that X satisfies the hypothesis of Lemma 11.112 and $G \curvearrowright X$ is a properly discontinuous quasiaction. Then the kernel for the associated boundary action ϕ_∞ is finite.*

PROOF. The kernel K of ϕ_∞ consists of the elements $g \in G$ such that the distance from $\phi(g)$ to the identity is finite. Since $\phi(g)$ is an (L, A) -quasiisometry of X , it follows from Lemma 11.112, that

$$\text{dist}(\phi(g), \text{Id}_X) \leq D(L, A, R, \delta).$$

As ϕ was properly discontinuous, the subgroup K is finite. \square

11.13.4. Conical limit points of quasiactions. Suppose that ϕ is a quasi-action of a group G on a Rips-hyperbolic space X . A point $\xi \in \partial_\infty X$ is called a *conical limit point* for the quasiaction ϕ if there exists a sequence $g_i \in G$ such that $\phi(g_i)(x)$ converges to ξ in the conical topology. In other words, for some (equivalently every) geodesic ray $\gamma \subset X$ asymptotic to ξ , and some (equivalently every) point $x \in X$, there exists a constant $R < \infty$ such that:

$$\lim_{i \rightarrow \infty} \phi(g_i)(x) = \xi$$

and

$$d(\phi(g_i)(x), \gamma) \leq R, \quad \text{for all } i.$$

LEMMA 11.118. *Suppose that $\psi : G \curvearrowright X$ is a cobounded quasiaction. Then every point of the ideal boundary $\partial_\infty X$ is a conical limit point for ψ .*

PROOF. Let $\xi \in \partial_\infty X$ and let $x_i \in X$ be a sequence converging to ξ in the conical topology (e.g., we can take $x_i = \gamma(i)$, where γ is a geodesic ray in X asymptotic to ξ). Fix a point $x \in X$ and R such that for every $x' \in X$ there exists $g \in G$ satisfying

$$d(x', \phi(g)(x)) \leq R.$$

Then, by coboundedness of the quasiaction ψ , there exists a sequence $g_i \in G$ for which

$$d(x_i, \phi(g_i)(x)) \leq R.$$

It follows that ξ is a conical limit point of the quasiaction ψ . □

COROLLARY 11.119. *Suppose that G is a finitely generated group, $f : X \rightarrow G$ is a quasisisometry and $G \curvearrowright G$ is the isometric action by the left multiplications. Let $\psi : G \curvearrowright X$ be the quasiaction, obtained by conjugating $G \curvearrowright G$ via f . Then every point of $\partial_\infty X$ is a conical limit point for the quasiaction ψ .*

PROOF. The action $G \curvearrowright G$ by the left multiplications is cobounded, hence, the conjugate quasiaction $\psi : G \curvearrowright X$ is also cobounded. □

If ϕ_∞ is a topological action of a group G on $\partial_\infty X$, obtained by the extension of a quasiaction ϕ of G on X , then *conical limit points* of the action $G \curvearrowright \partial_\infty X$ are defined as the conical limit points for the quasiaction $G \curvearrowright X$.

11.14. Hyperbolic groups

We now come to the *raison d'être* for δ -hyperbolic spaces, namely, *hyperbolic groups*.

DEFINITION 11.120. A finitely generated group G is called *Gromov-hyperbolic* or *word-hyperbolic*, or simply *hyperbolic*, if one of its Cayley graphs is hyperbolic. A hyperbolic group is called *elementary* if it is virtually cyclic. A hyperbolic group is called non-elementary otherwise.

Immediate examples of hyperbolic groups are:

- EXAMPLE 11.121. 1. Trivially, finite groups are hyperbolic.
 2. Every finitely generated free group is hyperbolic: Taking the Cayley graph corresponding to a free generating set, we obtain a simplicial tree, which is 0-hyperbolic.

We will see more examples of hyperbolic groups below.

Many examples of hyperbolic groups can be constructed *via the small cancellation* theory, see e.g. [GdlH90, GS90, IS98]. For instance, let G be a 1-relator group with the presentation

$$\langle x_1, \dots, x_n \mid w^m \rangle,$$

where $m \geq 2$ and w is a cyclically reduced word in the generators x_i . Then G is hyperbolic. (This was proven by B. B. Newman in [New68, Theorem 3] before the notion of hyperbolic groups was introduced; Newman proved that for such groups G the *Dehn's algorithm* applies, which is equivalent to hyperbolicity, see Section 11.16.)

Below is a combinatorial characterization of hyperbolic groups among Coxeter groups. Let Γ be a finite Coxeter graph and $G = C_\Gamma$ the corresponding Coxeter group. A *parabolic subgroup* of Γ is the Coxeter subgroup defined by a full subgraph

Λ of Γ . It is clear that every parabolic subgroup of G admits a natural homomorphism to G , sending the generators $g_v, v \in V(\Lambda)$, to the generators g_v of G . As it turns out that such homomorphisms are always injective, see e.g. [Hum97], page 113.

THEOREM 11.122 (G. Moussong [Mou88]). *A Coxeter group G is Gromov–hyperbolic if and only if the following condition holds:*

No parabolic subgroup of G is virtually isomorphic to the direct product of two infinite groups.

In particular, a Coxeter group is hyperbolic if and only if it contains no free abelian subgroups of rank 2.

PROBLEM 11.123. Is there a similar characterization of Gromov–hyperbolic groups among Shephard groups and generalized von Dyck groups?

Another outstanding open problem of the same flavor is:

PROBLEM 11.124 (flat closing problem). Suppose that G is a $CAT(0)$ group. Is it true that G contains a subgroup isomorphic to \mathbb{Z}^2 ?

This problem is open even for fundamental groups of closed Riemannian manifolds of nonpositive curvature and for 2-dimensional $CAT(0)$ groups.

Since changing generating sets does not affect the quasiisometry type of the Cayley graph and Rips–hyperbolicity is invariant under quasiisometries (Corollary 11.43), we conclude that a group G is hyperbolic if and only if all its Cayley graphs are hyperbolic. Furthermore, if groups G, G' are quasiisometric, then G is hyperbolic if and only if G' is hyperbolic. In particular, if G, G' are virtually isomorphic, then G is hyperbolic if and only if G' is hyperbolic. For instance, all virtually free groups are hyperbolic.

In view of the Milnor–Schwarz Theorem:

OBSERVATION 11.125. If G is a group acting geometrically on a Rips–hyperbolic metric space, then G is also hyperbolic.

DEFINITION 11.126. A group G is called $CAT(\kappa)$ if it admits a geometric action on a $CAT(\kappa)$ space.

Thus, every $CAT(-1)$ group is hyperbolic. In particular, fundamental groups of compact Riemannian manifolds of negative curvature are hyperbolic. If S is a compact connected surface then $\pi_1(S)$ is hyperbolic if and only if S is neither the torus nor the Klein bottle.

The following is an outstanding open problem in Geometric Group Theory:

PROBLEM 11.127. Construct a hyperbolic group G which is not a $CAT(-1)$ group.

Here are some examples of non-hyperbolic groups:

1. \mathbb{Z}^n is not hyperbolic for every $n \geq 2$. Indeed, \mathbb{Z}^n is QI to \mathbb{R}^n and \mathbb{R}^n is not hyperbolic (see Example 11.9).
2. A deeper fact is that hyperbolic groups cannot contain subgroups isomorphic to \mathbb{Z}^2 .
3. More generally, if G contains a solvable subgroup S , then G is not hyperbolic unless S is virtually cyclic.

4. Even more generally, for every subgroup S of a hyperbolic group G , the group S is either elementary hyperbolic or contains a nonabelian free subgroup. In particular, every amenable subgroup of a hyperbolic group is virtually cyclic.

5. Furthermore, if $C \triangleleft G$ is a cyclic normal subgroup of a hyperbolic group, then either C is finite, or G/C is finite.

We refer the reader to [BH99] for the proofs of 2, 3, 4 and 5.

REMARK 11.128. There are hyperbolic groups which contain non-hyperbolic finitely generated subgroups, see Theorem 11.156. A subgroup $H \leq G$ of a hyperbolic group G is called *quasiconvex* if it is a quasiconvex subset of a Cayley graph of G . If $H \leq G$ is a quasiconvex subgroup, then, according to Theorem 11.50, H is quasiisometrically embedded in G and, hence, is hyperbolic itself.

Examples of quasiconvex subgroups are given by finite subgroups (which is clear) and (less obviously) infinite cyclic subgroups. Let G be a hyperbolic group with a word metric d . Define the *translation length* of $g \in G$ as

$$\|g\| := \lim_{n \rightarrow \infty} \frac{d(g^n, e)}{n}.$$

It is clear that $\|g\| = 0$ if g has finite order. On the other hand, every cyclic subgroup $\langle g \rangle \subset G$ is quasiconvex and $\|g\| > 0$ for every g of infinite order, see Chapter III.Γ, Propositions 3.10, 3.15 of [BH99].

Obstructions to hyperbolicity.

If a finitely generated group G satisfies one of the following, then G is not hyperbolic:

- (1) G contains an amenable subgroup which is not virtually cyclic.
- (2) G contains an infinite cyclic subgroup which is not quasiisometrically embedded, i.e. an infinite order element g such that $\|g\| = 0$.
- (3) G has infinite cohomological dimension over \mathbb{Q} .
- (4) G does not contain a free nonabelian subgroup, and G is not two-ended.
- (5) G does not have the type \mathbf{F}_∞ .
- (6) G contains infinitely many conjugacy classes of finite subgroups.
- (7) G has unbounded torsion.
- (8) G does not admit a uniformly proper map to a Hilbert space.
- (9) G is not hopfian.

Proofs of 1–7 can be found in [BH99], while 8 and 9 are proven by Z. Sela in [Sel92] and [Sel99].

Strangely, all known examples of groups of the type \mathbf{F}_3 contain either \mathbb{Z}^2 or a solvable Baumslag–Solitar subgroup $BS(p, 1)$, cf. [Bra99]. In some cases, e.g., fundamental groups of compact 3-dimensional manifolds or free-by-cyclic groups (see [Bri00]), absence of such subgroups implies hyperbolicity. In the case of 3-dimensional manifolds, this result is a corollary of Perelman’s Geometrization Theorem (it suffices to rule out free abelian subgroups of rank 2 in this case). The following is a well-known open problem:

PROBLEM 11.129. Construct an example of a non-hyperbolic group of the type \mathbf{F}_∞ which contains no Baumslag–Solitar subgroups.

11.15. Ideal boundaries of hyperbolic groups

We define the *ideal boundary* $\partial_\infty G$ of a hyperbolic group G as the ideal boundary of some (every) Cayley graph of G : It follows from Theorem 11.108, that boundaries of different Cayley graphs are equivariantly homeomorphic. Here are two simple examples of ideal boundaries of hyperbolic groups.

Since $\partial_\infty \mathbb{H}^n = \mathbb{S}^{n-1}$, we conclude that for the fundamental groups G of closed hyperbolic n -manifolds, $\partial_\infty G$ is homeomorphic to the sphere \mathbb{S}^{n-1} . The same applies to the fundamental groups of compact negatively curved n -dimensional Riemannian manifolds. Similarly, if $G = F_n$ is the free group of rank $n \geq 2$, then free generating set S of G yields the Cayley graph $X = \Gamma_{G,S}$, which is a simplicial tree of constant valence > 2 . Therefore, as we saw in Section 11.11, $\partial_\infty X$ is homeomorphic to the Cantor set. Thus, $\partial_\infty F_n$ is the Cantor set as well.

LEMMA 11.130. *Let G be a hyperbolic group and $Z = \partial_\infty G$. Then Z consists of 0, 2 or continuum of points; in the latter case Z is perfect. In the first two cases, G is elementary, otherwise G is non-elementary. In the latter case, the kernel of the action $G \curvearrowright Z$ is the unique maximal finite normal subgroup of G .*

PROOF. Let X be a Cayley graph of G . If G is finite, then X is bounded and, hence $Z = \emptyset$. Thus, we assume that G is infinite. By Exercise 7.84, X contains a complete geodesic γ , thus, Z has at least two distinct points, the limit points of γ . If $\text{dist}_{\text{Haus}}(\gamma, X) < \infty$, X is quasiisometric to \mathbb{R} and, hence, G is 2-ended. Therefore, G is virtually cyclic by Part 3 of Theorem 9.22.

We assume, therefore, that $\text{dist}_{\text{Haus}}(\gamma, X) = \infty$, while $|Z| = 2$. Then there exists a sequence of vertices $x_n \in X$ satisfying $\lim \text{dist}(x_n, \gamma) = \infty$. Let $y_n \in \gamma$ be a nearest vertex to x_n and $g_n \in G$ be such that $g_n(y_n) = e \in G$. Then applying g_n to the union of geodesics

$$x_n y_n \cup \gamma$$

and taking the limit as $n \rightarrow \infty$, we obtain a complete geodesic $\beta \subset X$ (the limit of a subsequence $g_n(\gamma)$) and a geodesic ray ρ meeting β at e , such that for every $x \in \rho$, e is a nearest point on γ to x . Therefore, $\rho(\infty)$ is a point different from $\gamma(\pm\infty)$, and Z contains at least three distinct points, a contradiction.

Assume, now that $|Z| \geq 3$. The orbit $G \cdot e = G$ is a 1-net in X and, we are, therefore, in the situation of Lemma 11.112. Let K denote the kernel of the action $G \curvearrowright Z$. Then every $k \in K$ moves every point in X by $\leq D(1, 0, 1, \delta)$, where D is the function defined in Lemma 11.112. It follows that K is a finite group. Since G is infinite, Z is also infinite. Suppose that $F \triangleleft G$ is a finite normal subgroup. Since the quotient map $G \rightarrow \bar{G} = G/F$ is a quasiisometry, it induces a G -equivariant homeomorphism $\partial_\infty G \rightarrow \partial_\infty \bar{G}$. Since F acts trivially on $\partial_\infty \bar{G}$, it acts trivially on $\partial_\infty G$. It follows that K is the unique maximal finite normal subgroup of G .

Let $\xi \in Z$ and let ρ be a ray in X asymptotic to ξ . **Then** there exists a sequence $g_n \in G$ for which $g_n(e) = x_n \in \rho$. Let $\gamma \subset X$ be a complete geodesic asymptotic to points η, ζ different from ξ . We leave it to the reader to verify that either

$$\lim_n g_n(\eta) = \xi,$$

or

$$\lim_n g_n(\zeta) = \xi.$$

Since Z is infinite, we can choose ζ, η such that their images under the given sequence g_n are not all equal to ξ . Thus, ξ is an accumulation point of Z and, hence,

Z is a perfect topological space. Since Z is second countable, infinite, compact and Hausdorff, it follows that Z has the cardinality of continuum. \square

We next describe some dynamical properties of the actions of hyperbolic groups on their ideal boundaries.

DEFINITION 11.131. Let $G < \text{Homeo}(Z)$ be a group of homeomorphisms of a compact Hausdorff space Z . The group G is said to be a *convergence group* if G acts properly discontinuously on $\text{Trip}(Z)$, where $\text{Trip}(Z)$ is the set of triples of distinct elements of Z . A convergence group G is said to be a *uniform* if $\text{Trip}(Z)/G$ is compact.

THEOREM 11.132 (P. Tukia, [Tuk94]). *Suppose that X is a proper δ -hyperbolic geodesic metric space with the ideal boundary $Z = \partial_\infty X$ consisting of at least three points. Let $G \curvearrowright X$ be an isometric action and $G \curvearrowright Z$ be the corresponding topological action. Then the action $G \curvearrowright X$ is geometric if and only if $G \curvearrowright Z$ is a uniform convergence action.*

PROOF. Recall that we have a correspondence $center : \text{Trip}(Z) \rightarrow X$ sending each triple of distinct points in Z to the set of centroids of the corresponding ideal triangles. Furthermore, by Corollary 11.89, for every $\xi \in \text{Trip}(Z)$,

$$\text{diam}(center(\xi)) \leq 60\delta.$$

Clearly, the correspondence $center$ is G -equivariant. Moreover, the image of every compact K in $\text{Trip}(Z)$ under $center$ is bounded (see Exercise 11.91).

Assume now that the action $G \curvearrowright X$ is geometric. Given a compact subset $K \subset \text{Trip}(Z)$, suppose for a moment that the set

$$G_K := \{g \in G \mid gK \cap K \neq \emptyset\}$$

is infinite. Then there exists a sequence $\xi_n \in K$ and an infinite sequence $g_n \in G$, $g_0 = e \in G$, $g_n(\xi_n) \in K$ for all $n > 0$. The diameter of the set

$$E = \left(\bigcup_n center(g_n(\xi_n)) \right) \subset X$$

is bounded and each g_n sends some $p_n \in E$ to an element of E . This, however, contradicts proper discontinuity of the action of G on X . Thus, the action $G \curvearrowright \text{Trip}(Z)$ is properly discontinuous.

Similarly, since the action $G \curvearrowright X$ is cobounded, the G -orbit of some metric ball $B(p, R)$ covers the entire X . Thus, using equivariance of $center$, for every $\xi \in \text{Trip}(Z)$, there exists $g \in G$ such that

$$center(g\xi) \subset B = B(x, R + 60\delta).$$

Since $center^{-1}(B)$ is relatively compact in $\text{Trip}(Z)$ (see Exercise 11.92), we conclude that G acts cocompactly on $\text{Trip}(Z)$. We conclude that $G < \text{Homeo}(Z)$ is a uniform convergence group.

The proof of the converse is essentially the same argument run in the reverse. Let $K \subset \text{Trip}(Z)$ be a compact whose G -orbit is the entire $\text{Trip}(Z)$. Then the set $center(K)$, which is the union of sets of centroids of points $\xi' \in K$, is a bounded subset $B \subset X$. By equivariance of the correspondence $center$, it follows that the G -orbit of B is the entire X . Hence, the action $G \curvearrowright X$ is cobounded. The argument for proper discontinuity of the action $G \curvearrowright \text{Trip}(Z)$ is similar, we use the fact that

the preimage of a sufficiently large metric ball $B \subset X$ under the correspondence *center* is non-empty and relatively compact in $\text{Trip}(Z)$. Then proper discontinuity of the action $G \curvearrowright X$ follows from proper discontinuity of $G \curvearrowright \text{Trip}(Z)$. \square

COROLLARY 11.133. *Suppose that G is a nonelementary hyperbolic group. Then the image \bar{G} of G in $\text{Homeo}(\partial_\infty G)$ is a uniform convergence group.*

The converse to Theorem 11.132 is a deep theorem of B. Bowditch [**Bow98c**]:

THEOREM 11.134. *Let Z be a perfect compact metrizable space consisting of more than one point. Suppose that $G < \text{Homeo}(Z)$ is a uniform convergence group. Then G is hyperbolic and, moreover, there exists an equivariant homeomorphism $Z \rightarrow \partial_\infty G$.*

Note that in the proof of Part 1 of Theorem 11.132 we did not really need the property that the action of G on itself was isometric, a geometric quasi-action (see Definition 8.60) suffices:

THEOREM 11.135. *Suppose that X is a δ -hyperbolic proper geodesic metric space. Assume that there exists R such that every point in X is an R -centroid of an ideal triangle in X . Let $\phi : G \curvearrowright X$ be a geometric quasi-action. Then the extension $\phi_\infty : G \rightarrow \text{Homeo}(Z)$, $Z = \partial_\infty X$, of the quasi-action ϕ to a topological action of G on Z , is a uniform convergence action.*

PROOF. The proof of this result closely follows the proof of Theorem 11.132; the only difference is that ideal triangles $T \subset X$ are not mapped to ideal triangles by quasi-isometries $\phi(g)$, $g \in G$. However, ideal quasigeodesic triangles $\phi(g)(T)$ are uniformly close to ideal triangles which suffices for the proof. \square

The next theorem relates the spaces of ends $\epsilon(X)$ and ideal boundaries of hyperbolic spaces:

THEOREM 11.136. *Suppose that X is a Rips-hyperbolic proper metric space. Then there exists a continuous surjection*

$$\eta : \partial_\infty X \rightarrow \epsilon(X)$$

such that the preimages $\eta^{-1}(\xi)$ are connected components of $\partial_\infty X$. Moreover, the map η is equivariant with respect to the isometry group of X .

We refer the reader to [**GdlH90**, Chapter 7, Proposition 17] for a proof.

COROLLARY 11.137. *An infinite hyperbolic group is one-ended if and only if $\partial_\infty G$ is connected.*

Below is a brief review of topological properties of boundaries of hyperbolic groups.

DEFINITION 11.138. A point z in a connected topological space Z is called a *cut-point* if $Z - \{z\}$ is not connected. A 2-point subset $Y = \{z_1, z_2\}$ in Z is called a *cut-pair* if $Z - Y$ is not connected.

THEOREM 11.139. [**B. Bowditch**, [**Bow98b**]] *If G is a one-ended hyperbolic group then $\partial_\infty G$ has no cut-points. If $\partial_\infty G$ contains a cut-pair then either G splits as the fundamental group of a graph of groups with two-ended edge groups or $\partial_\infty G$ is homeomorphic to \mathbb{S}^1 .*

Note that if $\partial_\infty G$ is homeomorphic to \mathbb{S}^1 then G acts properly discontinuously, isometrically and cocompactly on \mathbb{H}^2 ; this theorem is due to Tukia, Gabai, Casson and Jungreis, see Section 23.7 and references therein.

A combination of the first part of Bowditch's theorem with the earlier work of M. Bestvina and G. Mess [BM91] yields:

THEOREM 11.140 (M. Bestvina, G. Mess; B. Bowditch). *If a group G is hyperbolic and is one-ended then $\partial_\infty G$ is locally connected.*

Theorems 11.139 and 11.140 allow one to analyze, to large extent, one-dimensional boundaries of hyperbolic groups:

THEOREM 11.141 (M. Kapovich, B. Kleiner, [KK00]). *If G is a hyperbolic one-ended group and $\partial_\infty G$ is one-dimensional then one the the following holds:*

1. $\partial_\infty G$ contains a cut-pair.
2. $\partial_\infty G$ is homeomorphic to the Sierpinsky carpet.
3. $\partial_\infty G$ is homeomorphic to the Menger curve.

While many hyperbolic groups have the Menger curve as their ideal boundary, conjecturally, the class of hyperbolic groups whose boundary is the Sierpinsky carpet is more limited;

CONJECTURE 11.142 (M. Kapovich, B. Kleiner, [KK00]). *If G is a hyperbolic group and $\partial_\infty G$ is homeomorphic to the Sierpinsky carpet, then G acts isometrically, properly discontinuously and cocompactly on a non-empty closed convex subset of \mathbb{H}^3 .*

It is proven in [KK00] that this conjecture would follow from the following conjecture made by J. Cannon:

CONJECTURE 11.143 (Cannon's Conjecture). *If G is a hyperbolic group with $\partial_\infty G$ homeomorphic to \mathbb{S}^2 , then G acts isometrically, properly discontinuously and cocompactly on \mathbb{H}^3 .*

This conjecture was verified by M. Bourdon and B. Kleiner, [BK13], in the case of Coxeter groups.

One can further ask what happens when $\partial_\infty G$ is a topological sphere. A Hausdorff second countable topological space M is called an k -dimensional (co)homology manifold if for every $x \in M$,

$$\check{H}^*(M - \{x\}; \mathbb{Z}) \cong H^*(\mathbb{R}^k - \{0\}; \mathbb{Z}),$$

where \check{H}^* denotes the Čech cohomology. A topological space M is called an (integral) homology k -dimensional sphere if

$$\check{H}^*(M; \mathbb{Z}) \cong H^*(\mathbb{S}^k; \mathbb{Z}).$$

CONJECTURE 11.144 (C.T.C. Wall). *If G is a hyperbolic group and $\partial_\infty G$ is a (co)homology manifold which is, moreover, a (co)homology $n - 1$ -sphere then G acts properly discontinuously and cocompactly on a contractible topological n -dimensional manifold.*

A partial confirmation to this conjecture comes from the following result:

THEOREM 11.145 (A. Bartels, W. Lueck, S. Weinberger, [BLW10]). *Wall's conjecture holds for $n \geq 5$ provided that $\partial_\infty G$ is homeomorphic to \mathbb{S}^{n-1} .*

We refer the reader to [KB02] for the more detailed discussion of ideal boundaries of hyperbolic groups.

11.16. Linear isoperimetric inequality and Dehn algorithm for hyperbolic groups

Let G be a hyperbolic group, we suppose that Γ is a δ -hyperbolic Cayley graph of G . We will assume that $\delta \geq 2$ is a natural number. Recall that a *loop* in Γ is required to be a closed edge-path. Since the group G acts transitively on the vertices of Γ , the number of G -orbits of loops of length $\leq 12\delta$ in Γ is bounded. We attach a 2-cell along every such loop. Let X denote the resulting cell complex; the action of G on Γ extends naturally to a cellular action on X . Recall that for a loop γ in Γ , $\ell(\gamma)$ denotes the length of γ and $A(\gamma)$ the least combinatorial area of a disk in X bounding γ , see Section 7.10.

Our goal is to show that X is simply-connected and satisfies a linear isoperimetric inequality. We will prove a somewhat stronger statement. Namely, suppose that X is a connected two-dimensional cell complex whose 1-skeleton $X^{(1)}$ (equipped with the standard metric) is δ -hyperbolic (with δ a natural number) and such that for every loop γ of length $\leq 12\delta$ in X , $A(\gamma) \leq K < \infty$. The following theorem was first proven by Gromov in Section 2.3 of [Gro87]:

THEOREM 11.146 (Hyperbolicity implies linear isoperimetric inequality). *Under the above assumptions, for every loop $\gamma \subset X$,*

$$(11.9) \quad A(\gamma) \leq K\ell(\gamma).$$

Since the argument in the proof of the theorem is by induction on the length of γ , the following proposition is the key. In the proposition, by saying that a loop γ based at a vertex $v \in X^{(1)}$ is a product of two loops γ_1, γ_2 in $X^{(1)}$, we mean that γ_1, γ_2 are also based at v and that γ represents the same element of $\pi_1(X^{(1)}, v)$ as the (concatenation) product $\gamma_1 \star \gamma_2$. Furthermore, d is the standard metric on the graph $X^{(1)}$.

PROPOSITION 11.147. *Every loop γ in $X^{(1)}$ of length larger than 12δ is a product of two loops, one of the length $\leq 12\delta$ and another one of the length $< \ell(\gamma)$.*

PROOF. We assume that γ is parameterized by its arc-length, and that $\ell(\gamma) = n$.

Case 1. Assume that there exists a vertex $u = \gamma(t)$ such that the vertex $v = \gamma(t + 6\delta)$ satisfies $d(u, v) < 6\delta$. After reparameterizing γ , we may assume that $t = 0$. Let p denote a geodesic vu in $X^{(1)}$ and $-p$ the same geodesic run in the reverse. Then γ is the product of the loops

$$\gamma_1 = \gamma([0, 6\delta]) \star p$$

and

$$\gamma_2 = (-p) \star \gamma([6\delta, n]).$$

(Here $u = \gamma(0)$ serves as a base-point.) Since $\ell(p) < \ell(\gamma([0, 6\delta]))$, we have $\ell(\gamma_1) \leq 12\delta$ and $\ell(\gamma_2) < \ell(\gamma_1)$. Thus, the statement of the proposition holds in the Case 1.

Case 2. Assume now that for every integer t , $d(\gamma(t), \gamma(t + 6\delta)) = 6\delta$, where $t + 6\delta$ is considered modulo n . In other words, every subarc of γ of length 6δ is a geodesic segment in $X^{(1)}$.




FIGURE 11.12. Case 1.

Set $v_0 = \gamma(0)$ and let $v = \gamma(t)$ denote a vertex whose distance $d(v, v_0)$ to v_0 is the largest possible, in particular it is at least 6δ .




FIGURE 11.13. Case 2.

Define the vertices $v_{\pm} = \gamma(t \pm 3\delta)$ on γ and consider the geodesic triangle $T = v_0 v_- v_+$ with the edge $v_- v_+$ equal to the geodesic subarc of γ between these vertices. Since the triangle T is δ -thin, the point $v \in v_- v_+$ is within distance $\leq \delta$ either from the side $v_0 v_-$ or from $v_0 v_+$. After reparameterizing γ in the reverse direction if necessary, we may assume that there exists a vertex $u \in v_0 v_+$ within distance $\leq \delta$ from v . Set

$$r = d(v_0, u), \quad s = d(u, v_+).$$

Then, by the triangle inequalities, $d(v_0, v) \leq r + \delta$, while $s \geq 3\delta - \delta = 2\delta$. Therefore,

$$d(v_0, v) = r + s \geq r + 2\delta > r + \delta \geq d(v_0, v).$$

This contradicts our choice of v as the point in $X^{(0)}$ on γ with the largest distance to v_0 . We, thus, conclude that the Case 2 cannot occur. \square

Proof of Theorem 11.146. The proof of the inequality (11.9) is by induction on the length of γ .

1. If $\ell(\gamma) \leq 12\delta$ then $A(\gamma) \leq K \leq K\ell(\gamma)$.

2. Suppose that the inequality holds for $\ell(\gamma) \leq n$, $n \geq 12\delta$. If $\ell(\gamma) = n + 1$, then γ is the product of loops γ_1, γ_2 as in Proposition 11.147: $\ell(\gamma_2) < \ell(\gamma)$, $\ell(\gamma_1) \leq 12\delta$. Then, inductively,

$$A(\gamma_2) \leq K\ell(\gamma_2), \quad A(\gamma_1) \leq K,$$

and, thus,

$$A(\gamma) \leq A(\gamma_2) + A(\gamma_1) \leq K\ell(\gamma_2) + K \leq K\ell(\gamma). \quad \square$$

Below are two corollaries of Proposition 11.147, which was the key to the proof of the linear isoperimetric inequality.

COROLLARY 11.148 (M. Gromov, [Gro87]). *Every hyperbolic group is finitely presented .*

PROOF. Proposition 11.147 means that every loop in the Cayley graph of Γ is a product of loops of length $\leq 12\delta$. Attaching 2-cells to Γ along the G -images of these loops we obtain a simply-connected complex X on which G acts geometrically. Thus, G is finitely presented . \square

COROLLARY 11.149 (M. Gromov, [Gro87], section 6.8N). *Let X be a coarsely connected Rips-hyperbolic metric space. Then X satisfies the linear isoperimetric inequality:*

$$Ar_\mu(\mathbf{c}) \leq K\ell(\mathbf{c})$$

for all sufficiently large μ and for appropriate $K = K(\mu)$.

PROOF. Quasiisometry invariance of coarse isoperimetric functions implies that it suffices to prove the assertion for Γ , the 1-skeleton of a connected R -Rips complex $Rips_R(X)$ of X . By Proposition 11.147, every loop γ in Γ is a product of $\leq \ell(\gamma)$ loops of length $\leq 12\delta$, where Γ is δ -hyperbolic in the sense of Rips. Therefore, for any $\mu \geq 12\delta$, we get

$$Ar_\mu(\gamma) \leq \ell(\gamma). \quad \square$$

Dehn algorithm. A (finite) presentation $\langle X|R \rangle$ is called *Dehn* if for every non-trivial word w representing $1 \in G$, the word w contains more than half of a defining relator. A word w is called *Dehn-reduced* if it contains no more than half of any relator. Given a word w , we can inductively reduce the length of w by replacing subwords u in w with u' such that $u'u^{-1}$ is a relator and $|u'| < |u|$. This, of course, does not change the element g of G represented by w . As the length of w decreases on each step, eventually, we get a Dehn-reduced word v representing $g \in G$. Since the presentation $\langle X|R \rangle$ is Dehn, either $v = 1$ (in which case $g = 1$) or $v \neq 1$ in which case $g \neq 1$. This algorithm is, probably, the simplest way to solve the word problem in groups. It is also, historically, the oldest: Max Dehn introduced it in order to solve the word problem for hyperbolic surface groups.

Geometrically, Dehn reduction represents a based homotopy of the path in X represented by the word w (the base-point is $1 \in G$). Similarly, one defines the *cyclic* Dehn reduction, where the reduction is applied to (unbased) loops represented by w and the *cyclically Dehn* presentation: If w is a null-homotopic loop in X then this loop contains a subarc which is more than half of a relator. Again, if G admits a cyclically Dehn presentation then the word problem in G is solvable.

LEMMA 11.150. *Each δ -hyperbolic group G admits a finite (cyclically) Dehn presentation.*

PROOF. Start with an arbitrary finite presentation of G . Then add to the list of relators all the words of length $\leq 12\delta$ representing the identity in G . Since the set of such words is finite, we obtain a new finite presentation of the group G . The fact that the new presentation is (cyclically) Dehn is just the induction step of the proof of Proposition 11.146. \square

Note, however, that the construction of a (cyclically) Dehn presentation requires solvability of the word problem for G (or, rather, for the words of the length

$\leq 12\delta$) and, hence, is not *a priori* algorithmic. Nevertheless, we will see below that a Dehn presentation for δ -hyperbolic groups (with known δ) is algorithmically computable.

The converse of Proposition 11.146 is true as well, i.e. if a finitely presented group satisfies a linear isoperimetric inequality then it is hyperbolic. We shall discuss this in Section 11.22.

11.17. The small cancellation theory

As we noted earlier, one of the origins of the theory of hyperbolic groups is the *small cancellation theory*. In this section we briefly discuss one class of *small cancellation conditions*, namely, $C'(r)$. We refer the reader to the books by Lyndon and Schupp [LS77], Ol'shanskii [Ol'91a], and the appendix by Strebel to [GdlH90], for details.

Consider a presentation $P = \langle X | R \rangle$. We define a new set of relators R^* by first symmetrizing R (adding the relator R_k^{-1} for each relator $R_k \in R$) and then cyclically conjugating each relator by generators $x_i \in X$. The new set of relators R^* is symmetric ($R = R^{-1}$) and is invariant under conjugation via generators $x_i \in X$ and their inverses.

DEFINITION 11.151. The presentation $P^* = \langle X | R^* \rangle$ is *reduced* if each relator in R^* is reduced and no relator is repeated.

A (non-empty) word w in $X \cup X^{-1}$ is called a *piece* with respect to the presentation P^* if w appears as a common prefix in two distinct elements R_i, R_j of R^* . The *relative length* of a piece w is

$$\ell_{rel}(w) = \max_{R_i} \frac{|w|}{|R_i|},$$

where the maximum is taken over all relators $R_i \in R^*$ containing w as a prefix.

DEFINITION 11.152. For $\lambda > 0$, a presentation P is said to satisfy the *small cancellation condition* $C'(\lambda)$ if $\ell_{rel}(w) \leq \lambda$ for each piece w .

THEOREM 11.153. *For each presentation P satisfying the condition $C'(1/7)$, the presentation P^* is cyclically Dehn.*

COROLLARY 11.154. *If a group G admits a finite presentation satisfying the condition $C'(1/7)$, then G is hyperbolic.*

THEOREM 11.155 (I. Chiswell, D. Collins and J. Huebschmann, [CCH81]; S. Gersten, [Ger87]). *The presentation complex of a presentation P satisfying the condition $C'(1/6)$ is aspherical provided that no relator is a proper power.*

11.18. The Rips construction

The goal of this section is to describe the *Rips construction*, which associates a hyperbolic group to an arbitrary finite presentation of an arbitrary group.

THEOREM 11.156 (The Rips Construction, E. Rips [Rip82]). *Let Q be a group with a finite presentation $\langle A | R \rangle$. Then this presentations gives rise to a short exact sequence*

$$1 \rightarrow K \rightarrow G \rightarrow Q \rightarrow 1$$

where G is hyperbolic and K is finitely generated. Furthermore, the group K in this construction is finitely-presentable if and only if Q is finite.

PROOF. Let $A = \{a_1, \dots, a_m\}$, $R = \{R_1, \dots, R_n\}$. For $i = 1, \dots, m, j = 1, 2$, pick even natural numbers $r_i < s_i, p_{ij} < q_{ij}, u_{ij} < v_{ij}$, such that all the intervals

$$[r_i, s_i], \quad [p_{ij}, q_{ij}] \quad [u_{ij}, v_{ij}], i = 1, \dots, m, j = 1, 2$$

are pairwise disjoint and all the numbers $r_i, s_i, p_{ij}, q_{ij}, u_{ij}, v_{ij}$ are at least 10 times larger than the lengths of the words in R . Define the group G by the presentation P , where the generators are $a_1, \dots, a_m, b_1, b_2$, and the relators are:

$$(11.10) \quad R_i b_1 b_2^{r_i} b_1 b_2^{r_i+1} \dots b_1 b_2^{s_i}, \quad i = 1, \dots, n.$$

$$(11.11) \quad a_i^{-1} b_j a_i b_1 b_2^{u_{ij}} b_1 b_2^{u_{ij}+1} \dots b_1 b_2^{v_{ij}}, \quad i = 1, \dots, m, j = 1, 2.$$

$$(11.12) \quad a_i b_j a_i^{-1} b_1 b_2^{p_{ij}} b_1 b_2^{p_{ij}+1} \dots b_1 b_2^{q_{ij}}, \quad i = 1, \dots, m, j = 1, 2.$$

Now, define the map

$$\tilde{\phi}(a_i) = a_i, i = 1, \dots, m, \quad \phi(b_j) = 1, \quad j = 1, 2.$$

The map $\tilde{\phi}$ extends to an epimorphism $F_{m+2} \rightarrow Q$ which sends all the relators R_k to $1 \in Q$; therefore, it descends to an epimorphism $\phi : G \rightarrow Q$. We claim that the kernel K of ϕ is generated by b_1, b_2 . First, the kernel, of course, contains b_1, b_2 . The subgroup generated by b_1, b_2 is clearly normal in G , because of the relators (11.11) and (11.12). Thus, indeed, b_1, b_2 generate K .

EXERCISE 11.157. The presentation P satisfies the small cancellation condition $C'(1/7)$. Hint: Show that the product of generators b_1, b_2 appearing at the end of each relator cannot get cancelled when we multiply conjugates of the relators in P and their inverses.

In particular, the group G is hyperbolic. In view of Theorem 11.155, the presentation complex of the presentation P is aspherical. Therefore, G has cohomological dimension ≤ 2 .

Lastly, we will verify that K cannot be finitely-presentable, unless Q is finite. R. Bieri proved in [Bie76b, Theorem B] that if G is a group of cohomological dimension ≤ 2 and $H \triangleleft G$ is a finitely-presentable normal subgroup of infinite index, then H is free.

Suppose that the subgroup K is free. Then the rank of K is at most 2 since K is 2-generated. The elements $a_1, a_2 \in G$ act on K as automorphisms (by conjugation). However, considering the action of a_1, a_2 on the abelianization of K , we see that because p_{ij}, q_{ij} are even, the images of the generators b_1, b_2 cannot generate the abelianization of K . A similar argument shows that K cannot be cyclic; therefore, K is trivial and, hence, $b_1 \equiv b_2 \equiv 1$ in G . However, this clearly contradicts the fact that the presentation (11.10) — (11.12) is a Dehn presentation. \square

This argument is a typical example of the *small cancellation theory*. Rips in his paper [Rip82], did not use the language of hyperbolic groups (which did not yet exist!), but the language of the small cancellation theory.

The Rips construction shows that there are hyperbolic groups which contain non-hyperbolic finitely generated subgroups. Furthermore,

COROLLARY 11.158. *Some hyperbolic groups have unsolvable membership problem.*

PROOF. Indeed, start with a finitely presented group Q with unsolvable word problem and apply the Rips construction to Q . Then $g \in G$ belongs to the normal subgroup $K \triangleleft G$ if and only if g maps to the identity in Q . Since Q has unsolvable word problem, the problem of membership of g in K is unsolvable as well. \square

On the other hand, the membership problem is solvable for quasiconvex subgroups, see Theorem 11.214.

11.19. Central coextensions of hyperbolic groups and quasiisometries

We now consider a central coextension

$$(11.13) \quad 1 \rightarrow A \rightarrow \tilde{G} \xrightarrow{r} G \rightarrow 1$$

with A a finitely generated abelian group and G hyperbolic. The main result of this section is:

THEOREM 11.159 (W. Neumann, L. Reeves, [NR97a]). *The group \tilde{G} is QI to $A \times G$.*

PROOF. In the case when $A \cong \mathbb{Z}$, the first published proof belongs to S. Gersten [Ger92], although, it appears that D.B.A. Epstein and G. Mess also knew this result. Our proof follows the one in [NR97a].

First of all, since an epimorphism with finite kernel is a quasiisometry, it suffices to consider the case when A is free abelian of finite rank. Our main goal is to construct a Lipschitz section (which is not a homomorphism!) $s : G \rightarrow \tilde{G}$ of the sequence (11.13). We first consider the case when A is infinite cyclic. Each fiber $r^{-1}(g), g \in G$, admits a canonical bijection to \mathbb{Z} :

$$ga \mapsto a \in A.$$

This defines a natural order \leq on $r^{-1}(g)$. We let ι denote the embedding

$$\mathbb{Z} \simeq A \hookrightarrow \tilde{G}.$$

Fix \mathcal{X} , a symmetric generating set of \tilde{G} ; we will use the same name for its image under r . We let $\langle \mathcal{X} | \mathcal{R} \rangle$ be a finite presentation of G . We will use the notation $|w|$ for the word length with respect to this generating set, $w \in \mathcal{X}^*$, where \mathcal{X}^* is the set of all words in \mathcal{X} , as in Section 7.2. Lastly, let \tilde{w} and \bar{w} denote the elements of \tilde{G} and G respectively, represented by $w \in \mathcal{X}^*$.

LEMMA 11.160. *There is $C \in \mathbb{N}$ such that for every $g \in G$ the subsets*

$$\{\tilde{w}\iota(-C|w|) : w \in \mathcal{X}^*, \bar{w} = g\} \subset r^{-1}(g)$$

are bounded from above with respect to the order \leq .

PROOF. We will use the fact that G satisfies the linear isoperimetric inequality

$$\text{Area}(\alpha) \leq K|\alpha|$$

for every $\alpha \in \mathcal{X}^*$ representing the identity in G . We will assume that $K \in \mathbb{N}$. For each $R \in \mathcal{X}^*$ such that $R^{\pm 1}$ is a defining relator for G , the word R represents some $\tilde{R} \in A$. Therefore, since G is finitely presented, we define a natural number T for which

$$\iota(T) = \max\{\tilde{R} : R^{\pm 1} \text{ is a defining relator of } G\}.$$

We claim that for each $u \in \mathcal{X}^*$ representing the identity in G ,

$$(11.14) \quad \iota(TArea(u)) \geq \tilde{u} \in A.$$

Since general relators u of G are products of words of the form hRh^{-1} , $R \in \mathcal{R}$, (where $Area(u)$ is at most the number of these terms in the product) it suffices to verify that for $w = h^{-1}Rh$,

$$\tilde{w} \leq \iota(T),$$

where R is a defining relator of G and $h \in \mathcal{X}^*$. The latter inequality follows from the fact that the multiplication by \bar{h} (resp. \bar{h}^{-1}) determines an order isomorphism (resp. its inverse) between $r^{-1}(1)$ and $r^{-1}(\bar{h})$.

Set $C := TK$. We are now ready to prove lemma. Let w, v be in \mathcal{X}^* representing the same element $g \in G$. Set $u := v^{-1}$. Then $q = wu$ represents the identity and, hence, by (11.14),

$$\tilde{q} = \tilde{w}\tilde{u} \leq \iota(C|q|) = \iota(C|w|) + \iota(C|u|).$$

We now switch to the addition notation for $A \simeq \mathbb{Z}$. **Then**

$$w - v \leq \iota(C|w|) + \iota(C|v|),$$

and

$$w - \iota(C|w|) \leq v + \iota(C|v|).$$

Therefore, taking v to be a fixed word representing g , we conclude that all the differences $w - \iota(C|w|)$ are bounded from above. \square

In view of this lemma, we define a section $s : G \rightarrow \tilde{G}$

$$s(g) := \max\{\tilde{w}\iota(-C|w|) : w \in \mathcal{X}^*, \bar{w} = g\}$$

of the exact sequence (11.13). The unique word $w = w_g$ realizing the maximum in the definition of s is called *maximizing*. The section s , of course, need not be a group homomorphism. We will see, nevertheless, that it is not far from being one. Define the cocycle

$$\sigma(g_1, g_2) := s(g_1)s(g_2) - s(g_1g_2),$$

where the difference is taking place in $r^{-1}(g_1g_2)$. The next lemma does not use hyperbolicity of G , only the definition of s .

LEMMA 11.161. *The set $\sigma(G, \mathcal{X})$ is finite.*

PROOF. Let $x \in \mathcal{X}$, $g \in G$. We have to estimate the difference

$$s(g)x - s(gx).$$

Let w_1 and w_2 denote maximizing words for g and gx respectively. Note that the word w_1x also represents gx . Therefore, by the definition of s ,

$$\tilde{w}_1\tilde{x}\iota(-C(|w_1| + 1)) \leq \tilde{w}_2\iota(-C|w_2|).$$

Hence, there exists $a \in A$, $a \geq 0$, satisfying

$$\tilde{w}_1\iota(-C|w_1|)\tilde{x}\iota(-C)a = \tilde{w}_2\iota(-C|w_2|)$$

and, thus

$$(11.15) \quad s(g)\tilde{x}\iota(-C)a = s(gx).$$

Since w_2x^{-1} represents g , we similarly obtain

$$(11.16) \quad s(gx)\tilde{x}^{-1}\iota(-C)b = s(g), \quad b \geq 0, b \in A.$$

By combining equations (11.15) and (11.16), and switching to the additive notation for the group operation in A we get

$$a + b = \iota(2C).$$

Since $a \geq 0, b \geq 0$, we conclude that $-\iota(C) \leq a - \iota(C) \leq \iota(C)$. Therefore, (11.15) implies that

$$|s(g)x - s(gx)| \leq C.$$

Since the finite interval $[-\iota(C), \iota(C)]$ in A is a finite set, lemma follows. \square

REMARK 11.162. Actually, more is true: The image of $\sigma : G \times G \rightarrow A$ is a finite set; in other words, the map $s : G \rightarrow \tilde{G}$ is a *quasihomomorphism* and the extension class of the central coextension (11.13) is a bounded cohomology class.

Moreover, all (degree $d \geq 2$) cohomology classes of hyperbolic groups are *bounded*: The natural homomorphism

$$H_b^d(G, A) \rightarrow H^d(G, A)$$

is surjective, see Section 5.9.3 for the definition of the bounded cohomology groups H_b^* . However, the proof is more difficult; we refer the reader to [Min01] for the details.

Letting L denote the maximum of the word lengths (with respect to the generating set \mathcal{X}) of the elements in the sets $\sigma(G, \mathcal{X}), \sigma(\mathcal{X}, G)$, we conclude (in view of Lemma 11.161) that the map $s : G \rightarrow \tilde{G}$ is $(L + 1)$ -Lipschitz. Given the section $s : G \rightarrow \tilde{G}$, we define the projection $\phi = \phi_s : \tilde{G} \rightarrow A$ by

$$(11.17) \quad \phi(\tilde{g}) = \tilde{g} - s \circ r(\tilde{g}).$$

It is immediate that ϕ is Lipschitz since s is Lipschitz.

We now extend this construction to the case of central coextensions with free abelian kernel of finite rank. Let $A = \prod_{i=1}^n A_i, A_i \cong \mathbb{Z}$. Consider the central coextension (11.13). The homomorphisms $A \rightarrow A_i$ induce quotient maps $\eta_i : \tilde{G} \rightarrow \tilde{G}_i$ with the kernels $\prod_{j \neq i} A_j$. Each \tilde{G}_i , in turn, is a central coextension

$$(11.18) \quad 1 \rightarrow A_i \rightarrow \tilde{G}_i \xrightarrow{r_i} G \rightarrow 1.$$

Assuming that each central coextension (11.18) has a Lipschitz section s_i , we obtain the corresponding Lipschitz projection $\phi_i : \tilde{G}_i \rightarrow A_i$ given by (11.17). This yields a Lipschitz projection

$$\Phi : \tilde{G} \rightarrow A, \Phi = (\phi_1 \circ \eta_1, \dots, \phi_n \circ \eta_n).$$

We now set

$$s(r(\tilde{g})) := \tilde{g} - \Phi(\tilde{g}).$$

It is straightforward to verify that s is well defined and is Lipschitz, provided that each s_i is. We thus obtain:

COROLLARY 11.163. *Given a finitely generated free abelian group A and a hyperbolic group G , each central coextension (11.13) admits a Lipschitz section $s : G \rightarrow \tilde{G}$ and a Lipschitz projection $\Phi : \tilde{G} \rightarrow A$ given by*

$$\Phi(\tilde{g}) = \tilde{g} - s(r(\tilde{g})).$$

Using this corollary, we define the map

$$h : G \times A \rightarrow \tilde{G}, \quad h(g, a) = s(g) + \iota(a)$$

and its inverse

$$h^{-1} : \tilde{G} \rightarrow G \times A, \quad \hat{h}(\tilde{g}) = (r(\tilde{g}), \Phi(\tilde{g})).$$

Since homomorphisms are 1-Lipschitz while the maps r and Φ are Lipschitz, we conclude that h is a bi-Lipschitz quasiisometry. \square

REMARK 11.164. The same proof goes through in the case of an arbitrary finitely generated group G and a central coextension (11.13) given by a *bounded 2-nd cohomology class*, cf. [Ger92].

EXAMPLE 11.165. Let $G = \mathbb{Z}^2$, $A = \mathbb{Z}$. Since $H^2(G, \mathbb{Z}) = H^2(T^2, \mathbb{Z}) \cong \mathbb{Z}$, the group G admits non-trivial central coextensions with the kernel A , for instance, the integer Heisenberg group H_3 . The group \tilde{G} for such a coextension is nilpotent but not virtually abelian. Hence, by Pansu's theorem (Theorem 16.26), \tilde{G} is not quasiisometric to $G \times A = \mathbb{Z}^3$.

One can ask if Theorem 11.159 generalizes to other normal coextensions of hyperbolic groups G . We note that Theorem 11.159 does not extend, say, to the case where A is a non-elementary hyperbolic group and the action $G \curvearrowright A$ is trivial. The reason is the *quasiisometric rigidity* for products of certain types of groups proven in [KKL98]. A special case of this theorem says that if G_1, \dots, G_n are non-elementary hyperbolic groups, then quasiisometries of the product $G = G_1 \times \dots \times G_n$ quasipreserve the product structure:

THEOREM 11.166. *Let $\pi_j : G \rightarrow G_j, j = 1, \dots, n$ be natural projections. Then for each (L, A) -quasiisometry $f : G \rightarrow G$, there is $C = C(G, L, A) < \infty$, such that, up to composing with a permutation of quasiisometric factors G_k , the map f is within distance $\leq C$ from a product map $f_1 \times \dots \times f_n$, where each $f_i : G_i \rightarrow G_i$ is a quasiisometry and C depends only on δ, n, L and A .*

11.20. Characterization of hyperbolicity using asymptotic cones

The goal of this section is to strengthen the relation between hyperbolicity of geodesic metric spaces and 0-hyperbolicity of their asymptotic cones.

PROPOSITION 11.167 (§2.A, [Gro93]). *Let (X, dist) be a geodesic metric space. Assume that either of the following two conditions holds:*

- (a) *There exists a non-principal ultrafilter ω such that for all sequences $\mathbf{e} = (e_n)_{n \in \mathbb{N}}$ of base-points $e_n \in X$ and $\boldsymbol{\lambda} = (\lambda_n)_{n \in \mathbb{N}}$ of scaling constants with $\omega\text{-lim } \lambda_n = 0$, the asymptotic cone $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ is a real tree.*
- (b) *For every non-principal ultrafilter ω and every sequence $\mathbf{e} = (e_n)_{n \in \mathbb{N}}$ of base-points, the asymptotic cone $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ is a real tree, where $\boldsymbol{\lambda} = (n^{-1})$.*

Then (X, dist) is hyperbolic.

The proof of Proposition 11.167 relies on the following lemma, whose proof follows closely the proof of the Morse Lemma (Theorem 11.40).

LEMMA 11.168. Assume that a geodesic metric space (X, dist) satisfies either the property (a) or the property (b) in Proposition 11.167. Then there exists $M > 0$ such that for every geodesic triangle $\Delta(x, y, z) \subset X$ with $\text{dist}(y, z) \geq 1$, the two edges with the endpoint x are at Hausdorff distance at most $M \text{dist}(y, z)$.

PROOF. Suppose to the contrary that there exist sequences of triples of points $x_n, y_n, z_n \in X$, such that $\text{dist}(y_n, z_n) \geq 1$ and

$$\text{dist}_{\text{Haus}}(x_n y_n, x_n z_n) = M_n \text{dist}(y_n, z_n),$$

such that $M_n \rightarrow \infty$. Let a_n be a point on $x_n y_n$ such that

$$\delta_n := \text{dist}(a_n, x_n z_n) = \text{dist}_{\text{Haus}}(x_n y_n, x_n z_n).$$

Since $\delta_n \geq M_n$, it follows that $\delta_n \rightarrow \infty$.

(a) Assume that the condition **(a)** holds. Consider the sequence of base-points $\mathbf{a} = (a_n)_{n \in \mathbb{N}}$ and the sequence of scaling constants $\boldsymbol{\lambda} = (1/\delta_n)_{n \in \mathbb{N}}$. In the asymptotic cone $\text{Cone}_\omega(X, \mathbf{a}, \boldsymbol{\lambda})$, the limits of $x_n y_n$ and $x_n z_n$ are at Hausdorff distance 1.

The triangle inequalities imply that the limits

$$\omega\text{-lim} \frac{\text{dist}(y_n, a_n)}{\delta_n} \quad \text{and} \quad \omega\text{-lim} \frac{\text{dist}(z_n, a_n)}{\delta_n}$$

are either both finite or both infinite. It follows that the limits of $x_n y_n$ and $x_n z_n$ are either two distinct geodesics joining the points $x_\omega = (x_n)$ and the point $y_\omega = (y_n) = z_\omega = (z_n)$, or two distinct asymptotic rays with common origin, or two distinct geodesics asymptotic on both sides. As in the proof of Theorem 11.40, all these cases are impossible in a real tree.

(b) Suppose that the condition **(b)** holds. Define an infinite subset

$$\mathcal{S} = \{ \lfloor \delta_n \rfloor : n \in \mathbb{N} \} \subset \mathbb{N}.$$

By Exercise 10.17, there exists a nonprincipal ultrafilter ω on \mathbb{N} such that $\omega(\mathcal{S}) = 1$. We define sequences $(x'_m), (y'_m), (z'_m)$ and (a'_m) in X , as follows. For every $m \in \mathcal{S}$ we choose an $n \in \mathbb{N}$ with $\lfloor \delta_n \rfloor = m$ and set

$$(x'_m, y'_m, z'_m, a'_m) = (x_n, y_n, z_n, a_n).$$

For m not in \mathcal{S} we make an arbitrary choice of the quadruple (x'_m, y'_m, z'_m, a'_m) . Lastly, define the scaling sequence $\boldsymbol{\lambda} = (m^{-1})$.

We now repeat the arguments in the part (a) of the proof for the asymptotic cone $\text{Cone}_\omega(X, \mathbf{a}', \boldsymbol{\lambda})$ and the limit geodesics $\omega\text{-lim} x'_m y'_m$ and $\omega\text{-lim} x'_m z'_m$. \square

PROOF OF PROPOSITION 11.167. Suppose that the geodesic space X is not hyperbolic. For every geodesic triangle Δ in X and a point $a \in \Delta$ we define the quantity $d_\Delta(a)$, which is the minimal distance from a to the union of the two opposite sides of Δ . Since X is assumed to be non-hyperbolic, for every $n \in \mathbb{N}$ there exists a geodesic triangle

$$\Delta_n = \Delta(x_n, y_n, z_n)$$

(with the sides $x_n y_n, y_n z_n, z_n x_n$) and points $a_n \in x_n y_n, b_n \in y_n z_n$, such that

$$d_n := d_{\Delta_n}(a_n) = \text{dist}(a_n, b_n) \geq n.$$

Here we choose $a_n \in x_n y_n$ to maximize the function d_{Δ_n} . We also pick a point $c_n \in x_n z_n$ which realizes the distance

$$\delta_n := \text{dist}(a_n, x_n z_n) \geq d_n.$$




FIGURE 11.14. Fat triangles.

(a) Suppose that the condition **(a)** is satisfied. We use the sequence of base-points $\mathbf{a} = (a_n)$ and scaling factors $\boldsymbol{\lambda} = (1/d_n)$ to define the asymptotic cone

$$\mathbf{K} = \text{Cone}_\omega(X, \mathbf{a}, \boldsymbol{\lambda}).$$

We next analyze the ultralimit of the sequence of geodesic triangles Δ_n .

There are two cases to consider:

A) $\omega\text{-lim } \frac{\delta_n}{d_n} < +\infty$.

By Lemma 11.168, we have that

$$\text{dist}_{\text{Haus}}(a_n x_n, c_n x_n) \leq M \cdot \delta_n.$$

Therefore the limits of $a_n x_n$ and $c_n x_n$ are either two geodesic segments with a common endpoint or two asymptotic rays. The same is true of the pairs of segments $a_n y_n, b_n y_n$ and $b_n z_n, c_n z_n$, respectively. It follows that the limit $\omega\text{-lim } \Delta_n$ is a geodesic triangle Δ with vertices $x, y, z \in \mathbf{K} \cup \partial_\infty \mathbf{K}$. The point $a = \omega\text{-lim } a_n \in xy$ is such that $\text{dist}(a, xz \cup yz) \geq 1$, which implies that Δ is not a tripod. This contradicts the fact that \mathbf{K} is a real tree.

B) $\omega\text{-lim } \frac{\delta_n}{d_n} = +\infty$.

This also implies that

$$\omega\text{-lim } \frac{\text{dist}(a_n, x_n)}{d_n} = +\infty \text{ and } \omega\text{-lim } \frac{\text{dist}(a_n, z_n)}{d_n} = +\infty.$$

By Lemma 11.168, we have

$$\text{dist}_{\text{Haus}}(a_n y_n, b_n y_n) \leq M \cdot d_n.$$

Thus, the respective limits of the sequences of segments $x_n y_n$ and $y_n z_n$ are either two rays with the common origin $y = \omega\text{-lim } y_n$ or two complete geodesics asymptotic in one direction. We denote them xy and yz , respectively, with $y \in$

$\mathbf{K} \cup \partial_\infty \mathbf{K}, x, z \in \partial_\infty \mathbf{K}$. The limit of $\{x_n, z_n\}$ in this case is empty (it is “out of sight”).

The choice of a_n implies that any point of $b_n z_n$ must be at a distance at most d_n from $x_n y_n \cup x_n z_n$. Therefore, all points on the ray bz are at the distance at most 1 from xy . It follows that xy and yz are either asymptotic rays emanating from y or complete geodesics asymptotic in both directions and at the Hausdorff distance 1. We again obtain a contradiction with the fact that \mathbf{K} is a real tree.

We conclude that the condition in **(a)** implies that X is δ -hyperbolic, for some $\delta > 0$.

(b) Suppose the condition **(b)** holds. Define $\mathcal{S} = \{\lfloor d_n \rfloor : n \in \mathbb{N}\}$, and let ω be a non-principal ultrafilter on \mathbb{N} such that $\omega(\mathcal{S}) = 1$ (see Exercise 10.17). We consider a sequence (Δ'_m) of geodesic triangles and a sequence (a'_m) of points on these triangles with the property that whenever $m \in \mathcal{S}$, $\Delta'_m = \Delta_n$ and $a'_m = a_n$, for some n such that $\lfloor d_n \rfloor = m$.

In the asymptotic cone $\text{Cone}_\omega(X, \mathbf{a}', (m^{-1}))$, with $\mathbf{a}' = (a'_m)$ we consider the limit of the sequence of triangles (Δ'_m) . We then argue as in the case when the condition **(a)** holds and, similarly, obtain a contradiction with the fact that the cone is a real tree. It follows that the condition **(b)** also implies hyperbolicity of X . \square

REMARK 11.169. An immediate consequence of Proposition 11.167 is an alternative proof of quasiisometric invariance of Rips-hyperbolicity among geodesic metric spaces: A quasiisometry between two spaces induces a bi-Lipschitz map between asymptotic cones, and a metric space bi-Lipschitz equivalent to a real tree is a real tree.

As a special case, consider Proposition 11.167 in the context of hyperbolic groups: A finitely generated group G is hyperbolic if and only if every asymptotic cone of G is a real tree. A finitely generated group G is called *lacunary-hyperbolic* if at least one asymptotic cone of G is a tree. Theory of such groups is developed in [OOS09], where many examples of non-hyperbolic lacunary hyperbolic groups are constructed. Thus, having one tree as an asymptotic cone is not enough to guarantee hyperbolicity of a finitely generated group. On the other hand:

THEOREM 11.170 (M. Kapovich, B. Kleiner [OOS09]). *Suppose that G is a finitely-presented group. Then G is hyperbolic if and only if one asymptotic cone of G is a tree.*

PROOF. Below we present a proof of this theorem which we owe to Thomas Delzant. We will need the following “local-to-global” characterization of hyperbolic spaces, which is a variation on Gromov’s “local-to-global” criterion established in [Gro87]:

THEOREM 11.171 (B. Bowditch, [Bow91], Theorem 8.1.2). *For every δ there exists δ' , such that for every m there exists R for which the following holds. If Y be an m -locally simply-connected R -locally δ -hyperbolic geodesic metric space, then Y is δ' -hyperbolic.*

Here, a space Y is *R -locally δ -hyperbolic* if every R -ball in Y with the path-metric induced from Y is δ -hyperbolic. Instead of defining m -locally simply-connected spaces, we note that every simply-connected simplicial complex equipped with the

standard metric, satisfies this condition for every $m > 0$. We refer to [Bow91, Section 8.1] for the precise definition. We will be applying this theorem in the case when $\delta = 1$, $m = 1$ and let δ' and R denote the resulting constants.

We now proceed with the proof suggested to us by Thomas Delzant. Suppose that G is a finitely presented group, one of whose asymptotic cones is a real tree. Let X be a simply-connected simplicial complex on which G acts freely, simplicially and cocompactly. We equip X with the standard path-metric dist . Then (X, dist) is quasiisometric to G . Suppose that ω is an ultrafilter on \mathbb{N} , (λ_n) is a scaling sequence converging to zero, and X_ω is the asymptotic cone of X with respect to this sequence, such that X_ω is isometric to a tree. Consider the sequence of metric spaces $X_n = (X, \lambda_n \text{dist})$. Then, since X_ω is a tree, by taking a diagonal sequence, there exists a pair of sequences r_n, δ_n with

$$\omega\text{-lim } r_n = \infty, \quad \omega\text{-lim } \delta_n = 0,$$

such that for ω -all n , every r_n -ball in X_n is δ_n -hyperbolic. In particular, for ω -all n , every R -ball in X_n is 1-hyperbolic. Therefore, by Theorem 11.171, the space X_n is δ' -hyperbolic for ω -all n . Since X_n is a rescaled copy of X , it follows that X (and, hence, G) is hyperbolic as well. \square

We now continue discussion of properties of trees which appear as asymptotic cones of hyperbolic spaces.

THEOREM 11.172. *Let X be a geodesic hyperbolic space which admits a geometric action of a group G . Then all the asymptotic cones of X are real trees where every point is a branch-point with valence equal the cardinality of $\partial_\infty X$.*

PROOF. STEP 1. By Theorem 8.37, the group G is finitely generated and hyperbolic and every Cayley graph Γ of G is quasiisometric to X . It follows that there exists a bi-Lipschitz bijection between the asymptotic cones

$$\Phi : \text{Cone}_\omega(G, \mathbf{e}, \boldsymbol{\lambda}) \rightarrow \text{Cone}_\omega(X, \mathbf{x}, \boldsymbol{\lambda}),$$

where x is a base-point in X , and \mathbf{e}, \mathbf{x} denote the constant sequences equal to $e \in G$ (the neutral element in G), and respectively to $x \in X$. Moreover, $\Phi(\mathbf{e}_\omega) = \mathbf{x}_\omega$. The map Φ thus determines a bijection between the space of directions $\Sigma_{\mathbf{e}_\omega}$ in the cone of Γ and the space of directions $\Sigma_{\mathbf{x}_\omega}$ in the cone of X . It suffices, therefore, to compute the cardinality of $\Sigma_{\mathbf{e}_\omega}$. For simplicity, in what follows, we denote the asymptotic cone $\text{Cone}_\omega(G, \mathbf{e}, \boldsymbol{\lambda})$ by G_ω .

STEP 2. We now construct an injective map from $\partial_\infty G$ to the space of directions at \mathbf{e}_ω in the asymptotic cone G_ω . Each point $\xi \in \partial_\infty G$ determines a collection of rays $e\xi$ in G within distance $\leq 2\delta$ from each other. The ultralimits of all these rays determine the same geodesic ray in G_ω . Taking the direction of this ray at the origin, we obtain a map

$$\text{Log} : \partial_\infty G \rightarrow \Sigma_{\mathbf{e}_\omega}.$$

We need to verify injectivity of this map. To this end, consider two geodesic rays $\rho_i : [0, \infty) \rightarrow \Gamma$, $\rho_i(0) = 1 \in G$, asymptotic to distinct points $\xi_i \in \partial_\infty G$, $i = 1, 2$. The ultralimits ρ_i^ω of these geodesic rays are geodesic rays in G_ω emanating from the point \mathbf{e}_ω . Proving that $\text{Log}(\xi_1) \neq \text{Log}(\xi_2)$ amounts to showing that for all $s > 0, t > 0$,

$$\text{dist}(\rho_1^\omega(s), \rho_2^\omega(t)) = s + t.$$

Since $\xi_1 \neq \xi_2$, for all positive values of s and t , the sequence of Gromov-products

$$\left(\rho_1 \left(\frac{s}{\lambda_n} \right), \rho_2 \left(\frac{t}{\lambda_n} \right) \right)_e$$

ω -converges to $(\xi_1, \xi_2)_e \in \mathbb{R}$. Therefore,

$$\begin{aligned} \text{dist}(\rho_1^\omega(s), \rho_2^\omega(t)) &= \omega\text{-lim } \lambda_n \text{dist} \left(\rho_1 \left(\frac{t}{\lambda_n} \right), \rho_2 \left(\frac{s}{\lambda_n} \right) \right) = \\ &= \omega\text{-lim} \left[t + s - 2\lambda_n \left(\rho_1 \left(\frac{t}{\lambda_n} \right), \rho_2 \left(\frac{s}{\lambda_n} \right) \right)_e \right] = t + s. \end{aligned}$$

Thus, $\text{Log}(\xi_1) \neq \text{Log}(\xi_2)$.

STEP 3. We argue that every direction of Γ_ω at e_ω is determined by a sequence of geodesic rays emanating from e in Γ . The argument below was suggested to us by Panos Papasoglu.

Elements of Σ_{e_ω} are represented by nondegenerate geodesic segments $e_\omega g_\omega$, where $g_\omega \in G_\omega$ is represented by a sequence (g_n) in G with $|g_n| \asymp \lambda_n^{-1}$ as $n \rightarrow \infty$.

We will need:

LEMMA 11.173 (Geodesic segments are uniformly close to geodesic rays). *Let Γ be a δ -hyperbolic, in the sense of Rips, Cayley graph of a group G . Then there exists a constant M such that each geodesic segment $s \subset \Gamma$ is contained in the M -neighborhood of a geodesic ray $v\xi$ in Γ .*

PROOF. We will consider the case when the group G is infinite, otherwise, there is nothing to prove.

Without loss of generality, we may assume that δ is a natural number. Furthermore, taking into account the isometric G -action on Γ , we may assume that the geodesic segment s is represented by an edge-path in Γ starting at the vertex $1 \in G = V(\Gamma)$. Let X denote the generating set of G used to define the Cayley graph Γ . Then vertex-paths in Γ starting at 1, can be described as finite or semi-infinite words in the alphabet $X \cup X^{-1}$. By abusing the terminology, we will use the same notation for paths in Γ as for the corresponding words. We will refer to a word as *geodesic* if it represents a geodesic path in Γ . Consider the set T of words in this alphabet, which define geodesics of length $k = 6\delta$ in Γ . Then there exists $R = R(k)$ such that each finite geodesic word p of length $\geq R$ contains at least two *disjoint subwords* equal to w for some $w \in T$, i.e. p has the form

$$w_0 w w_1 w w_2$$

where w_i 's are subwords of p . Given such partition of a finite geodesic path p , we define an infinite path q :

$$w_0 w w_1 w w_1 w w_1 w \dots$$

alternating the subwords w and w_1 infinitely many times. As w has length k , the path q is k -local geodesic, since each length k subword u in q appears as a subword in p , and p is geodesic.

Consider now a finite geodesic word p of length $\geq R$ and break p as the product of subwords:

$$p = p_1 p_2,$$

where p_2 has length R . Then, as above, partition p_2 as the product $w_0 w w_1 w w_2$ and define an infinite word q using this partition. Lastly, take the infinite word

$$q' = p_1 q.$$

Since w_2 has length $\leq R$, the path p is contained in the R -neighborhood of q' . By the construction, q' is a k -local geodesic. Taking into account Theorem 11.45, we conclude that q' is an $(3, 4\delta)$ -quasigeodesic ray in Γ . By the Extended Morse Lemma, q' is $D = D(3, 4\delta)$ -Hausdorff close to a geodesic ray 1ξ in Γ . Therefore, for $M = R + D$, the original path p is contained in the M -neighborhood of 1ξ . \square

We conclude (using this Lemma) that every direction of Γ_ω in e_ω is the germ of a limit ray. We then have a surjective map from the set of sequences in $\partial_\infty G$ to Σ_{e_ω} :

$$\{(\xi_n)_{n \in \mathbb{N}} : \xi_n \in \partial_\infty \Gamma\} = (\partial_\infty \Gamma)^\mathbb{N} \rightarrow \Sigma_{e_\omega}.$$

Steps 2 and 3 imply that for a non-elementary hyperbolic group, the cardinality of $\Sigma_{[e_\omega]}$ is the same as of $\partial_\infty G$, i.e. continuum. If G is an elementary hyperbolic group, then its asymptotic cone is a line and, theorem holds in this case as well. \square

A. Dyubina–Erschler and I. Polterovich ([DP01], [DP98]) proved a stronger result than Theorem 11.172:

THEOREM 11.174 ([DP01], [DP98]). *Let \mathcal{A} be the 2^{\aleph_0} -universal tree, as defined in Theorem 11.19. Then:*

- (a) *Every asymptotic cone of a non-elementary hyperbolic group is isometric to \mathcal{A} .*
- (b) *Every asymptotic cone of a complete, simply connected Riemannian manifold with strictly sectional curvature (i.e., curvature $\leq -\kappa < 0$), is isometric to \mathcal{A} .*

A consequence of Theorem 11.174 is that asymptotic cones of non-elementary hyperbolic groups and of complete, simply connected Riemannian manifolds of strictly negative sectional curvature cannot be distinguished from each other.

11.21. Size of loops

In this section we show that the characterization of hyperbolicity using asymptotic cones allows one to define hyperbolicity of a space in terms of *size of its loops*. Throughout this section, X denotes a geodesic metric space.

11.21.1. The minsize. One quantity that measures the size of geodesic triangles is the *minimal size* introduced in Definition 9.101 for topological triangles, which, of course, apply to geodesic triangles in X . This leads to the following definition:

DEFINITION 11.175. The *minimal size function* (the *minsize function*),

$$\text{minsize} = \text{minsize}_X : \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{\infty\},$$

is defined by

$$\text{minsize}(\ell) = \sup\{\text{minsize}(\Delta) : \Delta \text{ a geodesic triangle of perimeter } \leq \ell\}.$$

Note that according to (11.1), for each δ -hyperbolic (in the sense of Rips) metric space X , the function minsize is bounded above by 2δ . We will see below that the “converse” is also true, i.e. when the function minsize is bounded, the space X is hyperbolic. Moreover, M. Gromov proved [Gro87, §6] that sublinear growth of minsize is enough to conclude that a space is hyperbolic. With the characterization of hyperbolicity using asymptotic cones, the proof of this result is straightforward:

PROPOSITION 11.176. *A geodesic metric space X is hyperbolic if and only if $\text{minsize}(\ell) = o(\ell)$.*

PROOF. As noted above, one implication immediately follows from Lemma 11.57. Conversely, assume that $\text{minsize}(\ell) = o(\ell)$. We begin by proving that in each asymptotic cone of X , every finite geodesic is a limit geodesic, in the sense of Definition 10.50. More precisely:

LEMMA 11.177. *Let $\gamma = a_\omega b_\omega$ be a geodesic segment in the asymptotic cone $X_\omega = \text{Cone}_\omega(X, e, \lambda)$, where the points a_ω, b_ω are represented by the sequences $(a_i), (b_i)$ respectively. Then for every geodesic $a_i b_i \subset X$ connecting a_i to b_i ,*

$$\omega\text{-lim } a_i b_i = \gamma.$$

PROOF. Let $c_\omega = (c_i)$ be a point on γ . Consider an arbitrary geodesic triangle $\Delta_i \subset X$ with vertices a_i, b_i, c_i and the perimeter ℓ_i . Since

$$2 d(a_\omega, b_\omega) = \omega\text{-lim } \lambda_i \ell_i < \infty$$

and $\text{minsize}(\Delta_i) = o(\ell_i)$, we get

$$\omega\text{-lim } \lambda_i \text{minsize}(\Delta_i) = 0.$$

Taking the points x_i, y_i, z_i on the sides of Δ_i realizing the minsize of Δ_i , we conclude:

$$\omega\text{-lim } \lambda_i \text{diam}(x_i, y_i, z_i) = 0.$$

In particular, the sequences $(x_i), (y_i), (z_i)$ represent the same point $x_\omega \in X_\omega$. Then

$$\text{dist}(a_\omega, b_\omega) \leq \text{dist}(a_\omega, x_\omega) + \text{dist}(x_\omega, b_\omega) \leq$$

$$\text{dist}(a_\omega, x_\omega) + \text{dist}(x_\omega, b_\omega) + 2\text{dist}(x_\omega, c_\omega) = \text{dist}(a_\omega, c_\omega) + \text{dist}(c_\omega, b_\omega).$$

The first and the last term in the above sequence of inequalities are equal, hence all inequalities become equalities, in particular $c_\omega = x_\omega$. Thus c_ω belongs to the ultralimit $\omega\text{-lim } a_i b_i$ and lemma follows. \square

If one asymptotic cone $X_\omega = \text{Cone}_\omega(X, e, \lambda)$ is not a real tree, then it contains a geodesic triangle Δ_ω which is not a tripod. Without loss of generality we may assume that the geodesic triangle is a simple loop in X_ω . By the above lemma, the geodesic triangle Δ_ω is the ultralimit of a sequence of geodesic triangles (Δ_i) , with perimeters of the order $O\left(\frac{1}{\lambda_i}\right)$. The fact that $\text{minsize}(\Delta_i) = o\left(\frac{1}{\lambda_i}\right)$ implies that the three edges of Δ have a common point, a contradiction. \square

We note that Gromov in [Gro87, Proposition 6.6.F] proved a stronger version of Proposition 11.176:

THEOREM 11.178. *There exists a universal constant $\varepsilon_0 > 0$ such that if in a geodesic metric space X all geodesic triangles with length $\geq L_0$, for some L_0 , have*

$$\text{minsize}(\Delta) \leq \varepsilon_0 \cdot \text{perimeter}(\Delta),$$

then X is hyperbolic.

11.21.2. The constriction. Another way of measuring the size of loops in a space X is through their *constriction* function. We define the constriction function only for simple loops in X primarily for the notational convenience, the definition and the results generalize without difficulty if one considers non-simple loops.

We fix a constant $\lambda \in (0, \frac{1}{2})$. For a Lipschitz loop $c : \mathbb{S}^1 \rightarrow X$ of length ℓ , we define the λ -constriction of the loop c as $\text{constr}_\lambda(c)$, which is the infimum of $d(x, y)$, where the infimum is taken over all points x, y separating $c(\mathbb{S}^1)$ into two arcs of length at least $\lambda\ell$. Thus, the higher constriction means less distortion of c in X . The λ -constriction function, $\text{constr}_\lambda : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, of a metric space X is defined as

$$\text{constr}_\lambda(\ell) = \sup\{\text{constr}_\lambda(c) : c \text{ is a Lipschitz simple loop in } X \text{ of length } \leq \ell\}.$$

Note that when $\lambda \leq \mu$, $\text{constr}_\lambda \leq \text{constr}_\mu$, and $\text{constr}_\lambda(\ell) \leq \ell$.

PROPOSITION 11.179 ([Dru01], Proposition 3.5). *For geodesic metric spaces X the following are equivalent:*

- (1) X is δ -hyperbolic in the sense of Rips, for some $\delta > 0$;
- (2) there exists $\lambda \in (0, \frac{1}{4}]$ such that $\text{constr}_\lambda(\ell) = o(\ell)$;
- (3) for all $\lambda \in (0, \frac{1}{4}]$ and $\ell > 1$,

$$\text{constr}_\lambda(\ell) \leq 4\delta(\log_2(\ell + 12\delta) + 3) + 2.$$

REMARK 11.180. One cannot obtain a better order than $O(\log \ell)$ for the constriction function in hyperbolic spaces. This can be seen by considering metric circles of length ℓ lying on a horosphere in \mathbb{H}^3 .

PROOF. Our main tool, as before, are asymptotic cones of X .

We begin by arguing that (2) implies (1). In what follows we define *limit triangles* in an asymptotic cone $X_\omega = \text{Cone}_\omega(X, e, \lambda)$, to be the triangles in X_ω whose edges are limit geodesics. Note that such triangles *a priori* need not be themselves limits of sequences of geodesic triangles in X : They are merely limits of sequences of geodesic hexagons.

First note that (2) implies that every limit triangle in every asymptotic cone $\text{Cone}_\omega(X, e, \lambda)$ is a tripod. Indeed, if one assumes that one limit triangle is not a tripod, without loss of generality one can assume that this triangle forms a simple loop in X_ω . This triangle is the limit of a sequence of geodesic hexagons (H_i) , with three edges of lengths of the order $O(\frac{1}{\lambda_i})$, alternating with three edges of lengths of the order $o(\frac{1}{\lambda_i})$. (We leave it to the reader to verify that such hexagons may be chosen to be simple.) Since $\text{constr}_\lambda(H_i) = o(\frac{1}{\lambda_i})$ we obtain that $\omega\text{-lim } H_i$ is not simple, a contradiction.

It remains to show that every geodesic segment in every asymptotic cone of X is a limit geodesic. The proof is similar to that of Lemma 11.177.

Let $\gamma = a_\omega b_\omega$ be a geodesic in a cone $\text{Cone}_\omega(X, e, \lambda)$, where $a_\omega = (a_i)$ and $b_\omega = (b_i)$. We let $c_\omega = (c_i)$ be an arbitrary point in γ . We already know that every limit geodesic triangle $\Delta(a_\omega, b_\omega, c_\omega) \subset X_\omega$ is a tripod. If c_ω does not coincide with the center of this tripod, then

$$\text{dist}(a_\omega, c_\omega) + \text{dist}(c_\omega, b_\omega) > \text{dist}(a_\omega, b_\omega),$$

a contradiction. Thus, c_ω belongs to $\omega\text{-lim } a_i b_i$ and, hence, $\gamma = \omega\text{-lim } a_i b_i$.

We thus proved that every geodesic triangle in every asymptotic cone of X is a tripod, hence every asymptotic cone is a real tree. It follows that X is hyperbolic.

Clearly, (3) implies (2). We will prove that (1) implies (3). By monotonicity of the constriction function (as a function of λ), it suffices to prove (3) for $\lambda = \frac{1}{4}$. We denote $\text{constr}_{\frac{1}{4}}(c)$ simply by constr .

Consider an arbitrary closed Lipschitz curve $c : \mathbb{S}^1 \rightarrow X$ of length ℓ . We orient the circle and will use the notation α_{pq} to denote the oriented arc of the image of c connecting p to q . Let x, y, z be three points on $c(\mathbb{S}^1)$ connected by the arcs $\alpha_{xy}, \alpha_{yz}, \alpha_{zx}$ in $c(\mathbb{S}^1)$, such that the first two arcs have length $\frac{\ell}{4}$. Let $t \in \alpha_{zx}$ be the point minimizing the distance to y in X . Clearly,

$$R := \text{dist}(y, t) \geq \text{constr}$$

and for each point $s \in \alpha_{zx}$, $d(s, y) \geq R$. The point t splits the arc α_{zx} into two subarcs α_{zt}, α_{tx} . Without loss of generality, we can assume that the length of α_{tx} is $\geq \frac{\ell}{4}$. Let $\alpha_{xx'}$ be the maximal subarc of α_{xy} disjoint from $B(y, r)$ (we allow $x = x'$). We set

$$r := \frac{d(x', t)}{2}.$$

Since α_{tx} has length $\geq \frac{\ell}{4}$, we obtain

$$2r \geq \text{constr}.$$




FIGURE 11.15. Constriction.

The arc $\alpha_{tx'}$ connects the points t, x' of the metric sphere $S(y, R)$ outside of the open ball $B(y, R)$. Therefore, according to Lemma 11.65,

$$\ell \geq \ell(\alpha_{tx'}) \geq 2^{\frac{r-1}{2\delta}-3} - 12\delta$$

and, thus,

$$\text{constr} \leq 2r \leq 4\delta (\log_2(\ell + 12\delta) + 3) + 2.$$

The inequality in (3) follows. \square

11.22. Filling invariants of hyperbolic spaces

Recall that for every μ -simply connected geodesic metric space X we defined (see Section 9.7) the *filling area function* (or, *isoperimetric function*) $Ar(\ell) = Ar_{\mu,X}(\ell)$, which computes the least upper bound on the areas of disks bounding loops of lengths $\leq \ell$ in X . We also defined the *filling radius function* $r_{\mu,X}(\ell)$ which computes the least upper bounds on radii of such disks. The goal of this section is to relate both invariants to hyperbolicity of the space X .

11.22.1. Filling area. Recall also that hyperbolicity of X implies linearity of $Ar(\ell)$, see Corollary 11.149. In this section we will prove the converse. Moreover, we will prove that there is a gap between the quadratic filling order and the linear isoperimetric order. Namely, as soon as the isoperimetric inequality is less than quadratic, it has to be linear and the space has to be hyperbolic:

THEOREM 11.181 (Subquadratic filling, §2.3, §6.8, [Gro87]). *If a coarsely simply-connected geodesic metric space X has the isoperimetric function $Ar(\ell) = o(\ell^2)$, then X is hyperbolic.*

Note that there is a second gap for the possible filling orders of groups:

THEOREM 11.182 ([Ol'91b], [Bat99]). *If a group G admits a finite presentation which has the Dehn function $Dehn(\ell) = o(\ell)$, then G is either free or finite.*

Proofs of Theorem 11.181 can be found in [Ol'91b], [Pap95b], [Bow95a] and [Dru01]. B. Bowditch makes use of only two properties of the area function in his proof: The *quadrangle* (or *Besikovitch*) inequality (see Proposition 9.103) and a certain *theta-property*. In fact, as we will see below, only the quadrangle inequality or its triangle counterpart, the *minsize inequality* (see Proposition 9.104) are needed. Also, we will see that it suffices to have subquadratic isoperimetric function for geodesic triangles.

Proof of Theorem 11.181. Let X be a μ -simply-connected geodesic metric space with the isoperimetric function Ar_X and the minsize function $\text{minsize}_X : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, see Definition 11.175. According to Proposition 9.104, for every $\delta \geq \mu$,

$$[\text{minsize}_X(\ell)]^2 \leq 2\sqrt{3}\mu^2 Ar_X(\ell),$$

whence $Ar_X(\ell) = o(\ell^2)$ implies that $\text{minsize}_X(\ell) = o(\ell)$. Using Proposition 11.176, we conclude that X is hyperbolic. \square

The strongest known version of the converse to Corollary 11.149 is:

THEOREM 11.183 (Strong subquadratic filling theorem, see §2.3, §6.8 of [Gro87], and also [Ol'91b], [Pap96]). *Let X be a μ -simply connected geodesic metric space. If there exist sufficiently large N and L , and $\epsilon > 0$ sufficiently small, such that every loop \mathbf{c} in X with $N \leq Ar_\delta(\mathbf{c}) \leq LN$ satisfies*

$$Ar_\delta(\mathbf{c}) \leq \epsilon[\text{length}(\mathbf{c})]^2,$$

then the space X is hyperbolic.

It seems impossible to prove this theorem using asymptotic cones.

In Theorem 11.183 it suffices to consider only geodesic triangles Δ instead of all loops, and to replace the condition $N \leq Ar_\delta(\Delta) \leq LN$ by $\text{length}(\Delta) \geq N$. This

follows immediately from Theorem 11.178 and the minsize inequality in Proposition 9.104.

M. Coornaert, T. Delzant and A. Papadopoulos have shown that if X is a complete simply connected Riemannian manifold which is *reasonable* (see [CDP90, Chapter 6, §1] for a definition of this notion; for instance if X admits a geometric group action, then X is reasonable) then the constant ϵ in the previous theorem only has to be smaller than $\frac{1}{16\pi}$, see [CDP90, Chapter 6, Theorem 2.1].

In terms of the multiplicative constant, a sharp inequality was proved by S. Wenger.

THEOREM 11.184 (S. Wenger [Wen08]). *Let X be a geodesic metric space. Assume that there exists $\epsilon > 0$ and $\ell_0 > 0$ such that every Lipschitz loop \mathbf{c} of length $\text{length}(\mathbf{c})$ at least ℓ_0 in X bounds a Lipschitz disk $\mathfrak{d} : \mathbb{D}^2 \rightarrow X$ with*

$$\text{Area}(\mathfrak{d}) \leq \frac{1 - \epsilon}{4\pi} \text{length}(\mathbf{c})^2.$$

Then X is hyperbolic.

In the Euclidean space one has the classical isoperimetric inequality

$$\text{Area}(\mathfrak{d}) \leq \frac{1}{4\pi} \text{length}(\mathbf{c})^2,$$

with equality if and only if \mathbf{c} is a circle and \mathfrak{d} a planar disk.

Note that the quantity $\text{Area}(\mathfrak{d})$ appearing in Theorem 11.184 is a generalization of the notion of the geometric area used in this book. If the Lipschitz map $\phi : \mathbb{D}^2 \rightarrow X$ is injective almost everywhere, then $\text{Area}(\phi)$ is the 2-dimensional Hausdorff measure of its image. In the case of a Lipschitz map to a Riemannian manifold, $\text{Area}(\phi)$ is the *area of a map* defined in Section 3.4. When the target is a general geodesic metric space, $\text{Area}(\phi)$ is obtained by suitably interpreting the *Jacobian* $J_x(\phi)$ in the integral formula

$$\text{Area}(\phi) = \int_{\mathbb{D}^2} |J_x \phi(x)|.$$

11.22.2. Filling radius. Another application of the results of Section 11.21 is a characterization of hyperbolic spaces in terms of their filling radii.

PROPOSITION 11.185 ([Gro87], §6, [Dru01], §3). *For a geodesic μ -simply connected metric space X the following statements are equivalent:*

- (1) *X is hyperbolic;*
- (2) *the filling radius of X has sublinear growth: $r_X(\ell) = o(\ell)$;*
- (3) *the filling radius of X grows at most logarithmically: $r_X(\ell) = O(\log \ell)$.*

PROOF. In what follows, we let $\text{Ar} = \text{Ar}_\mu$ denote the μ -filling area function in the sense of Section 9.7, defined for loops in the space X .

We first prove that (1) \Rightarrow (3). According to the linear isoperimetric inequality for hyperbolic spaces (see Corollary 11.149), there exists a constant K depending only on X such that

$$(11.19) \quad \text{Ar}(\mathbf{c}) \leq K \ell_X(\mathbf{c}).$$

Here $\text{Ar}(\mathbf{c})$ is the μ -area of a least-area μ -disk $\mathfrak{d} : \mathcal{D}^{(0)} \rightarrow X$ bounding \mathbf{c} . Recall also that the *combinatorial length* and *area* of a simplicial complex is the number of 1-simplices and 2-simplices respectively in this complex. Thus, for a loop \mathbf{c} as above, we have

$$\ell_X(\mathbf{c}) \leq \mu \text{ length}(\mathcal{C}),$$

where \mathcal{C} is the triangulation of the circle \mathbb{S}^1 so that vertices of any edge are mapped by \mathbf{c} to points within distance $\leq \mu$ in X .

Consider now a loop $\mathbf{c} : \mathbb{S}^1 \rightarrow X$ of metric length ℓ and a least area μ -disk $\mathfrak{d} : \mathcal{D}^{(0)} \rightarrow X$ filling \mathbf{c} ; thus, $\text{Ar}(\mathbf{c}) \leq K\ell$.

Let $v \in \mathcal{D}^{(0)}$ be a vertex such that its image $a = \mathfrak{d}(v)$ is at maximal distance r from $\mathbf{c}(\mathbb{S}^1)$. For every $1 \leq j \leq k$, with

$$k = \lfloor \frac{r}{\mu} \rfloor$$

we define a subcomplex \mathcal{D}_j of \mathcal{D} : \mathcal{D}_j is the maximal connected subcomplex in \mathcal{D} containing v , so that every vertex in \mathcal{D}_j could be connected to v by a *gallery* (in the sense of Section 1.7.1) of 2-dimensional simplices σ in \mathcal{D} such that

$$\mathfrak{d}(\sigma^{(0)}) \subset \overline{B}(a, j\mu).$$

For instance, \mathcal{D}_1 contains the star of v in \mathcal{D} . Let Ar_j be the number of 2-simplices in \mathcal{D}_j .

For each $j \leq k - 1$ the geometric realization \mathcal{D}_j of the subcomplex \mathcal{D}_j is homeomorphic to a 2-dimensional disk with several disks removed from the interior. (As usual, we will conflate a simplicial complex and its geometric realization.) Therefore, the boundary $\partial\mathcal{D}_j$ of \mathcal{D}_j in \mathbb{D}^2 is a union of several disjoint topological circles, while all the edges of \mathcal{D}_j are interior edges for \mathcal{D} . We denote by s_j the outermost circle in $\partial\mathcal{D}_j$, i.e. s_j bounds a triangulated disk $\mathcal{D}'_j \subset \mathcal{D}$, such that $\mathcal{D}_j \subset \mathcal{D}'_j$. Let $\text{length}(\partial\mathcal{D}_j)$ and $\text{length}(s_j)$ denote the number of edges of $\partial\mathcal{D}_j$ and of s_j respectively.

By the definition, every edge of \mathcal{D}_j is an interior edge of \mathcal{D}_{j+1} and belongs to a 2-simplex of \mathcal{D}_{j+1} . Note also that if σ is a 2-simplex in \mathcal{D} and two edges of σ belong to \mathcal{D}_j , then σ belongs to \mathcal{D}_j as well. Therefore,

$$\text{Ar}_{j+1} \geq \text{Ar}_j + \frac{1}{3} \text{length}(\partial\mathcal{D}_j) \geq \text{Ar}_j + \frac{1}{3} \text{length}(s_j).$$

Since \mathfrak{d} is a least area filling disk for \mathbf{c} it follows that each disk $\mathfrak{d}|_{\mathcal{D}'_j}$ is a least area disk bounding the loop $\mathfrak{d}|_{s_j}$. In particular, by the isoperimetric inequality in X ,

$$\text{Ar}_j = \text{Area}(\mathcal{D}_j) \leq \text{Area}(\mathcal{D}'_j) \leq K\ell_X(\mathfrak{d}(s_j)) \leq K\mu \text{ length}(s_j)$$

We have thus obtained that

$$\text{Ar}_{j+1} \geq \left(1 + \frac{1}{3\mu K}\right) \text{Ar}_j.$$

It follows that

$$K\ell \geq \text{Ar}(\mathfrak{d}) \geq \left(1 + \frac{1}{3\mu K}\right)^k$$

whence,

$$r \leq \mu(k+1) \leq \mu \left(\frac{\ln \ell + \ln K}{\ln \left(1 + \frac{1}{3\mu K}\right)} + 1 \right).$$

Clearly (3) \Rightarrow (2). It remains to prove the implication (2) \Rightarrow (1).

We first show that (2) implies that in every asymptotic cone $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$, all geodesic triangles that are limits of geodesic triangles in X (i.e. $\boldsymbol{\Delta} = \omega\text{-lim } \Delta_i$) are tripods. We assume that $\boldsymbol{\Delta}$ is not a point. Every geodesic triangle Δ_i can be regarded as a loop $\mathbf{c}_i : \mathbb{S}^1 \rightarrow \Delta_i$, and can be filled by a μ -disk $\mathfrak{d}_i : \mathcal{D}^{(1)} \rightarrow X$ of the filling radius $r_i = r(\mathfrak{d}_i) = o(\text{length}(\Delta_i))$. In particular, $\omega\text{-lim}_i \lambda_i r_i = 0$.

Let $x_i y_i, y_i z_i$ and $z_i x_i$ be the three geodesic edges of Δ_i , and let $\bar{x}_i, \bar{y}_i, \bar{z}_i$ be the three points on \mathbb{S}^1 corresponding to the three vertices x_i, y_i, z_i . Consider a path $\bar{\mathbf{p}}_i$ in the 1-skeleton of \mathcal{D} with endpoints \bar{y}_i and \bar{z}_i such that $\bar{\mathbf{p}}_i$ together with the arc of \mathbb{S}^1 with endpoints \bar{y}_i, \bar{z}_i encloses a maximal number of triangles with \mathfrak{d}_i -images in the r_i -neighborhood of $y_i z_i$. Every edge of $\bar{\mathbf{p}}_i$ that is not in \mathbb{S}^1 is contained in a 2-simplex whose third vertex has \mathfrak{d}_i -image in the r_i -neighborhood of $y_i x_i \cup x_i z_i$. The edges in $\bar{\mathbf{p}}_i$ that are in \mathbb{S}^1 are either between \bar{x}_i, \bar{y}_i or between \bar{x}_i, \bar{z}_i .

Thus, $\bar{\mathbf{p}}_i$ has \mathfrak{d}_i -image \mathbf{p}_i in the $(r_i + \mu)$ -neighborhood of $y_i x_i \cup x_i z_i$. See Figure 11.16.




FIGURE 11.16. The path $\bar{\mathbf{p}}_i$ and its image \mathbf{p}_i .

Consider an arbitrary vertex \bar{u} on \mathbb{S}^1 between \bar{y}_i, \bar{z}_i and its image $u \in y_i z_i$. We have that $\mathbf{p}_i \subset \bar{\mathcal{N}}_{r_i + \mu}(y_i u) \cup \bar{\mathcal{N}}_{r_i + \mu}(u z_i)$, where $y_i u$ and $u z_i$ are sub-geodesics of $y_i z_i$.

By connectedness, there exists a point $u' \in \mathbf{p}_i$ at distance at most $r_i + \mu$ from a point $u_1 \in y_i u$, and from a point $u_2 \in u z_i$. As the three points u_1, u, u_2 are aligned on a geodesic and $\text{dist}(u_1, u_2) \leq 2(r_i + \mu)$ it follows that, say, $\text{dist}(u_1, u) \leq r_i + \mu$, whence $\text{dist}(u, u') \leq 3(r_i + \mu)$. Since the point \bar{u} was arbitrary, we have thus proved that $y_i z_i$ is in $\bar{\mathcal{N}}_{3r_i + 3\mu}(\mathbf{p}_i)$, therefore, it is in $\bar{\mathcal{N}}_{4r_i + 4\mu}(y_i x_i \cup x_i z_i)$. This implies that in $\boldsymbol{\Delta}$ one edge is contained in the union of the other two. The same argument done for each edge implies that $\boldsymbol{\Delta}$ is a tripod.

From this, one can deduce that every triangle in the cone is a tripod. In order to do this it suffices to show that every geodesic in the cone is a limit geodesic. Consider a geodesic in $\text{Cone}_\omega(X, \mathbf{e}, \boldsymbol{\lambda})$ with the endpoints $x_\omega = (x_i)$ and $y_\omega = (y_i)$ and an arbitrary point $z_\omega = (z_i)$ on this geodesic. Geodesic triangles $\Delta_i = \Delta(x_i, y_i, z_i)$ yield a tripod $\Delta_\omega = \Delta(x_\omega, y_\omega, z_\omega)$ in the asymptotic cone, but since,

$$\text{dist}(x_\omega, z_\omega) + \text{dist}(z_\omega, y_\omega) = \text{dist}(x_\omega, y_\omega),$$

it follows that the tripod must be degenerate. Thus $z_\omega \in \omega\text{-lim } x_i y_i$. \square

REMARK 11.186. 1. One can show that in Part (3) of the proposition, given a loop $\mathbf{c} : \mathbb{S}^1 \rightarrow X$ of length ℓ , a filling disk \mathfrak{d} minimizing the area has the filling radius $r(\mathfrak{d}) = O(\log \ell)$.

2. The logarithmic order in (3) cannot be improved, as shown by the example of the horizontal circle in the upper half-space model of \mathbb{H}^3 .

3. In view of this proposition (as in the case of the filling area) there is a gap between the linear order of the filling radius and the logarithmic one.

Analogously to the filling area, for the radius too there is a stronger version of the implication *sublinear filling radius* \implies *hyperbolicity*, similar to Theorem 11.183.

PROPOSITION 11.187 (M. Gromov; P. Papasoglu [Pap98]). *Let Γ be a finitely presented group. If there exists $\ell_0 > 0$ such that*

$$r(\ell) \leq \frac{\ell}{73}, \quad \forall \ell \geq \ell_0,$$

then the group Γ is hyperbolic.

QUESTION 11.188 ([Pap98]). Find a proof of Proposition 11.187 with the constant $\frac{1}{73}$ replaced by $\frac{1}{8}$.

As the proof of Proposition 11.187 relies on the bigon criterion for hyperbolicity ([Pap95c]; see also Section 11.24), and there is now a metric version of this criterion (see [CN07] and the work of Pomroy quoted within), it is natural to ask also the following

QUESTION 11.189. Find a version of Proposition 11.187 for general metric spaces, with a constant that can be made effective (and sharp) for complete simply connected Riemannian manifolds.

11.22.3. Orders of Dehn functions of non-hyperbolic groups and higher Dehn functions. As we saw earlier, there is a gap between linear and quadratic orders for Dehn functions of groups. It is natural to ask what happens to the growth orders. For each $k \in \mathbb{N}$ define the subset $A_k \subset \mathbb{R}$ consisting of numbers α such that n^α is the order of the k -th Dehn function of a group G of type \mathbf{F}_{k+1} . Since there are only countably many finitely presented functions, each set A_k is countable. M. Sapir, J. Birget and E. Rips in [SBR02] gave a detailed description of the set A_1 , which, in particular, implies that the intervals $(0, 1)$ and $(1, 2)$ are the only gaps in A_1 :

THEOREM 11.190 (M. Sapir, J. Birget, E. Rips). *The closure of A_1 contains the half-line $[2, \infty)$. Moreover, let M be a not necessarily deterministic Turing machine with time function $T(n)$ for which $T(n)^4$ is superadditive. Then there exists a finitely presented group $G(M)$ with Dehn function equivalent to $T(n)^4$.*

In particular, this theorem shows that subpolynomial Dehn functions need not be of the type n^α ; according to the following theorem, this can happen even for nilpotent groups:

THEOREM 11.191 (S. Wenger, [Wen11]). *There exists a nilpotent group G with the order of Dehn(G) strictly larger n^2 and at most $n^2 \log(n)$.*

Dehn functions can be very fast-growing:

THEOREM 11.192 (W. Dison and T. Riley, [DR13]). *For each k , there exists a finitely presented group Hydra_k (a “hydra group”), whose Dehn function is equivalent to the k -th Ackermann function,*

$$\text{Hydra}_k = \langle a_1, \dots, a_k, p, t \mid t^{-1}a_it = a_ia_{i+1}, i > 1, t^{-1}a_1t = a_1, [p, a_it] = 1, \forall i \rangle$$

Recall that Ackermann functions are defined by the recursive formula:

$$\begin{aligned} A_1(n) &:= 2n \\ A_{k+1}(n) &:= A_k^{(n)}(1), \end{aligned}$$

where (n) means the n -fold composition.

Surprisingly, there are no gaps in the orders of higher Dehn functions:

THEOREM 11.193 (N. Brady, M. Forester, [BF10]). *For each $k \geq 2$, the closure of A_k contains the half-line $[1, \infty)$.*

We recall that one can define isoperimetric functions using homological fillings rather than the homotopical fillings (as in the definition of the Dehn functions), see (3.3).

THEOREM 11.194 (A. Abrams, N. Brady, P. Dani and R. Young, [ABDY13]). *There are finitely presented groups for which the (classical) Dehn function is not equivalent to the homological Dehn function.*

11.23. Asymptotic cones, actions on trees and isometric actions on hyperbolic spaces

Let G be a finitely generated group with the generating set g_1, \dots, g_m ; let X be a metric space. Given a homomorphism $\rho : G \rightarrow \text{Isom}(X)$, we define the following function:

$$(11.20) \quad d_\rho(x) := \max_k d(\rho(g_k)(x), x)$$

and set

$$d_\rho := \inf_{x \in X} d_\rho(x).$$

The function $d_\rho(x)$ does not necessarily attain its infimum in X , hence, we choose $x_\rho \in X$ to be a point such that

$$d_\rho(x) - d_\rho \leq 1.$$

Such points x_ρ are called *min-max* points of ρ for obvious reasons. The set of min-max points could be unbounded, but, as we will see, this does not matter. Thus, high value of d_ρ means that *all* points of X move a lot by at least one of the generators of $\rho(G)$, while small value of d_ρ means that at least one point in X is moved only a bit by all the generators of G .

EXAMPLE 11.195. 1. Let $X = \mathbb{H}^n$, $G = \langle g \rangle$ be infinite cyclic group, where $\rho(g) \in \text{Isom}(X)$ is a hyperbolic translation along a geodesic $L \subset X$ with the translation number $t > 1$, e.g., $\rho(g)(\mathbf{x}) = e^t \mathbf{x}$ in the upper half-space model. Then $d_\rho = t$ and we can take $x_\rho \in L$, since L is the set of point of minima for $d_\rho(x)$.

2. Suppose that $X = \mathbb{H}^n = \mathbf{U}^n$ and G are the same as above, but $\rho(g)$ is a parabolic isometry, e.g. $\rho(g)(\mathbf{x}) = \mathbf{x} + \mathbf{u}$, where $\mathbf{u} \in \mathbb{R}^{n-1}$ is a unit vector. Then d_ρ does not attain its infimum, $d_\rho = 0$ and we can take as x_ρ any point $\mathbf{x} \in \mathbf{U}^n$ with $x_n \geq 1$.

3. Suppose that X is the same, but G is no longer required to be cyclic. Assume that $\rho(G)$ fixes a unique point $x_o \in X$. Then $d_\rho = 0$ and the set of min-max points is contained in a metric ball centered at x_o . The radius of this ball could be estimated from above independently of G and ρ . (The latter is non-trivial.)

Suppose $\sigma \in \text{Isom}(X)$ and we replace the original representation ρ with the conjugate representation

$$\rho' = \rho^\sigma : g \mapsto \sigma\rho(g)\sigma^{-1}, g \in G.$$

EXERCISE 11.196. Verify that $d_\rho = d_{\rho'}$ and that as $x_{\rho'}$ one can take $\sigma(x_\rho)$.

Thus, conjugating ρ by an isometry, does not change the geometry of the action, but moves min-max points in a predictable manner.

The set $\text{Hom}(G, \text{Isom}(X))$ embeds in $(\text{Isom}(X))^m$ since every ρ is determined by the m -tuple

$$(\rho(g_1), \dots, \rho(g_m)).$$

As usual, we equip the group $\text{Isom}(X)$ with the topology of uniform convergence on compacts and the set $\text{Hom}(G, \text{Isom}(X))$ with the subspace topology.

EXERCISE 11.197. Show that the topology on $\text{Hom}(G, \text{Isom}(X))$ is independent of the finite generating set. Hint: Embed $\text{Hom}(G, \text{Isom}(X))$ in the product of countably many copies of $\text{Isom}(X)$ (indexed by the elements of G) and relate the topology on $\text{Hom}(G, \text{Isom}(X))$ to the product topology on the infinite product.

Suppose now that the metric space X is proper. Pick a base-point $o \in X$. Then the Arzela-Ascoli theorem implies that for every D the subset

$$\text{Hom}(G, \text{Isom}(X))_{o,D} = \{\rho : G \rightarrow \text{Isom}(X) \mid d_\rho(o) \leq D\}$$

is compact. We next consider the quotient

$$\text{Rep}(G, \text{Isom}(X)) = \text{Hom}(G, \text{Isom}(X)) / \text{Isom}(X),$$

where $\text{Isom}(X)$ acts on $\text{Hom}(G, \text{Isom}(X))$ by conjugation $\rho \mapsto \rho^\sigma$. We equip $\text{Rep}(G, \text{Isom}(X))$ with the quotient topology. In general, this topology is not Hausdorff:

EXAMPLE 11.198. Let $G = \langle g \rangle$ is infinite cyclic, $X = \mathbb{H}^n$. Show that the trivial representation $\rho_0 : G \rightarrow 1 \in \text{Isom}(X)$ and the representation ρ_1 , where $\rho_1(g)$ acts as a parabolic translation, project to points $[\rho_i]$ in $\text{Rep}(G, \text{Isom}(X))$, such that every neighborhood of $[\rho_0]$ contains $[\rho_1]$. Hence, $\text{Rep}(G, \text{Isom}(X))$ is not even T_1 in this example.

EXERCISE 11.199. Let X be a graph (not necessarily locally-finite) with the standard metric and consider the subset $\text{Hom}_f(G, \text{Isom}(X))$ consisting of representations ρ which give rise to the free actions $G/\text{Ker}(\rho) \curvearrowright X$. Then

$$\text{Rep}_f(G, \text{Isom}(X)) = \text{Hom}_f(G, \text{Isom}(X)) / \text{Isom}(X)$$

is Hausdorff.

We will be primarily interested in compactness rather than Hausdorff properties of $\text{Rep}(G, \text{Isom}(X))$. Define

$$\text{Hom}_D(G, \text{Isom}(X)) = \{\rho : G \rightarrow \text{Isom}(X) \mid d_\rho \leq D\}.$$

Similarly, for a subgroup $H \subset \text{Isom}(X)$, define

$$\text{Hom}_D(G, H) = \text{Hom}_D(G, \text{Isom}(X)) \cap \text{Hom}(G, H).$$

LEMMA 11.200. *Suppose that $H \subset \text{Isom}(X)$ is a closed subgroup whose action on X is cobounded. Then for every $D \in \mathbb{R}_+$, the quotient $\text{Rep}_D(G, H) = \text{Hom}_D(G, H)/H$ is compact.*

PROOF. Let $o \in X, R < \infty$ be such that the orbit of $\bar{B}(o, R)$ under the H -action is the entire space X . For every $\rho \in \text{Hom}(G, H)$ we pick $\sigma \in H$ such that some min-max point x_ρ of ρ satisfies:

$$\sigma(x_\rho) \in \bar{B}(o, R).$$

Then, using **a** conjugation by such σ 's, for each equivalence class $[\rho] \in \text{Rep}_D(G, H)$ we choose a representative ρ with $x_\rho \in \bar{B}(o, R)$. It follows that for every such ρ

$$\rho \in \text{Hom}(G, H) \cap \text{Hom}(G, \text{Isom}(X))_{o, D'}, \quad D' = D + 2R.$$

This set is compact and, hence, its projection $\text{Rep}_D(G, H)$ is also compact. \square

In view of this lemma, even if X is not proper, we say that a sequence of representations $\rho_i : G \rightarrow \text{Isom}(X)$ *diverges* if

$$\lim_{i \rightarrow \infty} d_{\rho_i} = \infty.$$

DEFINITION 11.201. We say that an isometric action of a group on a real tree T is *unbounded* if the group does not fix a point in T .

PROPOSITION 11.202 (M. Bestvina [**Bes88**] and F. Paulin, [**Pau88**]). *Suppose that (ρ_i) is a diverging sequence of representations $\rho_i : G \rightarrow H < \text{Isom}(X)$, where X is a Rips-hyperbolic metric space. Then G admits an unbounded isometric action on a real tree.*

PROOF. Let $p_i = x_{\rho_i}$ be min-max points of ρ_i 's. Take $\lambda_i := (d_{\rho_i})^{-1}$ and consider the corresponding asymptotic cone $\mathbf{X} = \text{Cone}_\omega(X, \mathbf{p}, \lambda)$ of the space X ; here $\mathbf{p} = (p_i)$. According to Lemma 11.37, the metric space \mathbf{X} in this asymptotic cone is a real tree. Furthermore, the sequence of group actions ρ_i converges to an isometric action $\rho_\omega : G \curvearrowright \mathbf{X}$, defined by:

$$\rho_\omega(g)(x_\omega) = (\rho_i(x_i)),$$

the key here is that *all generators $\rho_i(g_k)$ of $\rho_i(G)$ move the base-point $p_i \in \lambda_i X$ by $\leq \lambda_i(d_{\rho_i} + 1)$* . The ultralimit of the latter quantity is equal to 1. Furthermore, for ω -all i one of the generators, say $g = g_k$, satisfies

$$|d_{\rho_i} - d(\rho_i(g)(p_i), p_i)| \leq 1$$

in X . Thus, the element $\rho_\omega(g)$ will move the point $\mathbf{p} \in \mathbf{X}$ exactly by 1. Because p_i was a min-max point of ρ_i , it follows that

$$d_{\rho_\omega} = 1.$$

In particular, the isometric action $\rho_\omega : G \curvearrowright \mathbf{X}$ has no fixed point, i.e. is unbounded. \square

One of the important applications of this proposition is:

THEOREM 11.203 (F. Paulin, [**Pau91a**]). *Suppose that G is a finitely generated group with the property FA and H is a hyperbolic group. Then, up to conjugation in H , there are only finitely many homomorphisms $G \rightarrow H$.*

PROOF. Let X be a Cayley graph of H , then $H < \text{Isom}(X)$, X is proper and Rips–hyperbolic. By the above proposition, if $\text{Hom}(G, H)/H$ is non-compact, then G has an unbounded isometric action on a real tree. This contradicts the assumption that G has the property FA. Suppose, therefore, that $\text{Hom}(G, H)/H$ is compact. If this quotient is infinite, pick a sequence $\rho_i \in \text{Hom}(G, H)$ of pairwise non-conjugate representations. Without loss of generality, by replacing ρ_i 's by their conjugates, we can assume that min-max points p_i of ρ_i 's are in $\overline{B}(e, 1)$. Therefore, after passing to a subsequence if necessary, the sequence of representations ρ_i converges. However, the action of H on itself is free, hence, for every generator g of G , the sequence $\rho_i(g)$ is eventually constant. It follows that the entire sequence (ρ_i) consists of only finitely many distinct representations. Contradiction. Thus, $\text{Hom}(G, H)/H$ is finite. \square

This theorem is one of many results bounding the number of homomorphisms from a group to a hyperbolic group. Having the Property FA is a very strong restriction on the group, thus, typically, one improves the Proposition 11.202 by making stronger assumptions on representations $G \rightarrow H$ and, accordingly, stronger conclusions about the action of G on the tree, for instance:

THEOREM 11.204. *Suppose that G is a group, H is a hyperbolic group, X is a Cayley graph of H and $\rho_i : G \rightarrow H, i \in \mathbb{N}$, is a sequence of faithful representations as in Proposition 11.202. Then the action $G \curvearrowright T$ of G on a real tree as in Proposition 11.202, is small, i.e. the stabilizer of every non-trivial geodesic segment is virtually cyclic.*

Given this theorem, whose proof can be found e.g. in [Pau91b], one then (typically) uses the *Rips Theory*, which converts small actions (satisfying some mild restrictions which will hold in the case of groups G which embed in hyperbolic groups) $G \curvearrowright T$, into graph-of groups decompositions of G with virtually cyclic edge groups. We refer the reader to [BF95, RS94, Kap01] for the details. As the result, one obtains:

THEOREM 11.205 (E. Rips, Z. Sela, [RS94]). *Suppose that G does not split over a virtually cyclic subgroup. Then for every hyperbolic group H , $\text{Hom}_{inj}(G, H)/H$ is finite, where Hom_{inj} consists of injective homomorphisms. In particular, if G is itself hyperbolic, then $\text{Out}(G) = \text{Aut}(G)/\text{Inn}(G)$ is finite.*

Some interesting and important groups G , like surface groups, do split over virtually cyclic subgroups. In this case, one cannot in general expect $\text{Hom}_{inj}(G, H)/H$ to be finite. However, it turns out that the only reason for the lack of finiteness is the fact that one can precompose homomorphisms $G \rightarrow H$ with automorphisms of G itself:

THEOREM 11.206 (E. Rips, Z. Sela, [RS94]). *Suppose that G is a 1-ended finitely generated group. Then for every hyperbolic group H , the set*

$$\text{Aut}(G) \backslash \text{Hom}_{inj}(G, H) / H$$

is finite. Here $\text{Aut}(G)$ acts on $\text{Hom}(G, H)$ by precomposition.

11.24. Summary of equivalent definitions of hyperbolicity

Below we give a list of equivalent definitions of hyperbolicity for a (finitely generated) group G (some of these definitions we saw earlier).

- (1) Some/every Cayley graph of G is Rips-hyperbolic, i.e. has δ -thin triangles.
- (2) G is Gromov-hyperbolic, i.e. when equipped with the word metric for some (every) finite generating set, the group G satisfies Gromov's inequality for the Gromov-product.
- (3) Some/every Cayley graph of G has δ -thin bigons for some $\delta < \infty$, P. Papasoglu [Pap95c].
- (4) G admits a Dehn-presentation.
- (5) G is finitely presented and satisfies a linear isoperimetric inequality.
- (6) G is finitely presented and satisfies subquadratic isoperimetric inequality.
- (7) G is finitely presented and its isoperimetric function satisfies

$$IP(\ell) \leq \frac{1 - \epsilon}{4\pi} \ell^2$$

for some $\epsilon > 0$ and all sufficiently large ℓ .

- (8) G is finitely presented and has sublinear filling radius.
- (9) Every asymptotic cone of G is a real tree, M. Gromov [Gro93], see also Proposition 11.167.
- (10) G is finitely presented and one asymptotic cone of G is a real tree, M. Kapovich, B. Kleiner [KK09], see also Theorem 11.170.
- (11) The minsize function of one (every) Cayley graph of G is sublinear, see Proposition 11.176.
- (12) For some $\lambda \in (0, \frac{1}{4}]$ the λ -constriction function constr_λ of G is sublinear, C. Drutu [Dru01], see also Proposition 11.179.
- (13) Some/every Cayley graph of G has proper uniform divergence, P. Papasoglu [Pap95c].
- (14) Some/every Cayley graph of G has exponential uniform divergence, P. Papasoglu [Pap95c].
- (15) Some/every Cayley graph of G admits a thin bicombing, B. Bowditch and U. Hamenstädt, [Ham07, Proposition 3.5], see Theorem 11.12.
- (16) G is finitely presented and the canonical map between bounded and ordinary cohomology groups with coefficients in Banach $\mathbb{Z}G$ -modules,

$$H_b^*(G, V) \rightarrow H^*(G, V)$$

is surjective, I. Mineyev [Min01].

- (17) G is finitely presented and

$$\ell_1 H_1(G, \mathbb{R}) = \overline{\ell_1 H_2(G, \mathbb{R})} = 0,$$

D. Allcock and S. Gersten [AG99].

- (18) Either G is virtually cyclic or G acts topologically on a compact perfect metrizable space Z of infinite cardinality, such that the induced action $G \curvearrowright \text{Trip}(Z)$ is properly discontinuous and cocompact, B. Bowditch [Bow98c], see also Theorem 11.134.

11.25. Further properties of hyperbolic groups

We conclude this chapter with a list of properties of hyperbolic groups not discussed earlier in the chapter:

1. Hyperbolic groups are ubiquitous:

THEOREM 11.207 (See e.g. [Del96]). *Let G be a non-elementary δ -hyperbolic group. Then there exists N , such that for every collection $g_1, \dots, g_k \in G$ of elements with translation lengths $\|g_1\| = \dots = \|g_k\| \geq 1000\delta$, such that each pair of subgroups $\langle g_i, g_j \rangle, i \neq j$ is nonelementary, the following holds:*

- i. The subgroup generated by the elements g_i^N and all their conjugates is free.*
- ii. The quotient group $G / \langle\langle g_1^n, \dots, g_k^n \rangle\rangle$ is again non-elementary hyperbolic for all sufficiently large n . In particular, infinite hyperbolic groups are never simple.*

Thus, by starting with, say, a nonabelian free group $F_n = G$, and adding to its presentation one relator of the form w^n at a time (where n 's are large), one obtains non-elementary hyperbolic groups.

“Most” groups are hyperbolic:

THEOREM 11.208 (A. Ol'shanskiĭ [Ol'92]). *Fix $k \in \mathbb{N}, k \geq 2$ and let*

$$A = \{a^{\pm 1}, a^{\pm 2}, \dots, a_k^{\pm 1}\}$$

be an alphabet. Fix $i \in \mathbb{N}$ and let (n_1, \dots, n_i) be a sequence of natural numbers. Let $N = N(k, i, n_1, \dots, n_i)$ be the number of group presentations

$$\langle a_1, \dots, a_k | r_1, \dots, r_i \rangle$$

such that r_1, \dots, r_i are reduced words in the alphabet A such that the length of r_j is $n_j, j = 1, 2, \dots, i$. If N_h is the number of presentations as above which define hyperbolic groups and if $n = \min\{n_1, \dots, n_i\}$, then

$$\lim_{n \rightarrow \infty} \frac{N_h}{N} = 1,$$

and convergence is exponentially fast.

The model of randomness which appears in this theorem is, by no means, unique; below are two other models (among many others). We refer the reader to [Gro03], [Ghy04], [Oll04], [KS08] for further discussion of random groups.

- i. Fix the number $k \geq 2$ and consider the set $B(n)$ of presentations

$$\langle x_1, \dots, x_k | R_1, \dots, R_l \rangle,$$

where the total length of the words R_1, \dots, R_l is $\leq n$. Then a class \mathcal{C} of k -generated groups is said to consist of random groups if

$$\lim_{n \rightarrow \infty} \frac{\text{card}(B(n) \cap \mathcal{C})}{\text{card} B(n)} = 1.$$

- ii. Here is another notion of randomness: Fix the number l of relators, assume that all relators have the same length n ; this defines a class of presentations $S(k, l, n)$. Then require

$$\lim_{n \rightarrow \infty} \frac{\text{card}(S(k, l, n) \cap \mathcal{C})}{\text{card} S(k, l, n)} = 1.$$

See [KS08] for a comparison of various notions of randomness for groups. In all existing models of randomness, once random groups are infinite, they are hyperbolic with Menger curve as the ideal boundary, see e.g. [DGP11].

- 2. Hyperbolic groups have quotients with “exotic” properties:

THEOREM 11.209 (A. Ol’shanskiĭ, [Ol’91c]). *Every non-elementary torsion-free hyperbolic group admits a quotient which is an infinite torsion group, where every non-trivial element has the same order.*

THEOREM 11.210 (A. Ol’shanskiĭ, [Ol’95], T. Delzant [Del96]). *Every non-elementary hyperbolic group G is SQ-universal, i.e. every countable group embeds in a quotient of G .*

Theorems 11.211, 11.212, 11.213 below first appeared in Gromov’s paper [Gro87]; other proofs could be found for instance in [Aea91], [BH99], [ECH⁺92], [ECH⁺92], [GdlH90].

3. Hyperbolic groups have finite type:

THEOREM 11.211. *Let G be δ -hyperbolic. Then there exists $D_0 = D_0(\delta)$ such that for all $D \geq D_0$ the Rips complex $\text{Rips}_D(G)$ is contractible. In particular, G has type \mathbf{F}_∞ .*

4. Hyperbolic groups have controlled torsion:

THEOREM 11.212. *Let G be hyperbolic. Then G contains only finitely many conjugacy classes of finite subgroups.*

5. Hyperbolic groups have solvable algorithmic problems:

THEOREM 11.213. *Every δ -hyperbolic group has solvable word and conjugacy problems.*

Furthermore:

THEOREM 11.214 (I. Kapovich, [Kap96]). *The membership problem is decidable for quasiconvex subgroups of hyperbolic groups: Let G be hyperbolic and $H < G$ be a quasiconvex subgroup of a δ -hyperbolic group. Then the problem of membership in H is decidable.*

The isomorphism problem is decidable:

THEOREM 11.215 (Z. Sela, [Sel95]; F. Dahmani and V. Guirardel [DG11]). *There is an algorithm whose input is a pair (P_1, P_2) of finite presentations of δ -hyperbolic groups G_1, G_2 , and the output is YES if G_1, G_2 are isomorphic and NO if they are not.*

Note that Sela proved this theorem only for “rigid” torsion-free 1-ended hyperbolic groups (rigidity here means that the group does not split as a graph of groups with cyclic edge groups). This result was extended to all hyperbolic groups by Dahmani and Guirardel.

6. Hyperbolic groups are hopfian:

THEOREM 11.216 (Z. Sela, [Sel99]). *For every hyperbolic group G and every epimorphism $\phi : G \rightarrow G$, $\text{Ker}(\phi) = 1$.*

Note that every finitely generated residually finite group is hopfian, but the converse, in general, is false. An outstanding open problem is to determine if all hyperbolic groups are residually finite (it is widely expected that the answer is negative). Every finitely generated linear group is residually finite, but there are nonlinear hyperbolic groups, see [Kap05]. It is very likely that some (or even all) of the nonlinear hyperbolic groups described in [Kap05] are not residually finite.

7. Hyperbolic groups tend to be co-hopfian:

THEOREM 11.217 (Z. Sela, [Sel97a]). *For every 1-ended hyperbolic group G , every monomorphism $\phi : G \rightarrow G$ is surjective, i.e. such G is co-hopfian.*

8. All hyperbolic groups admit QI embeddings in some hyperbolic space \mathbb{H}^n :

THEOREM 11.218 (M. Bonk, O. Schramm [BS00]). *For every hyperbolic group G there exists n , such that G admits a quasiisometric embedding in \mathbb{H}^n .*

9. (M. Gromov; [Gro87]. See also [BH99] and [GdlH90]). Hyperbolic groups have type \mathbf{F}_∞ . Moreover, there exists $D_0 = D_0(\delta)$ such that for every δ -hyperbolic group G and all $D \geq D_0$, the Rips complex $\text{Rips}_D(G)$ is contractible.

11.26. Relatively hyperbolic spaces and groups

Relatively hyperbolic groups were introduced by M. Gromov in the same paper [Gro87] as hyperbolic groups. While hyperbolic groups are modeled on uniform lattices in negatively curved symmetric spaces, relatively hyperbolic groups are modeled on non-uniform lattices in negatively curved spaces and, more generally, fundamental groups of complete Riemannian manifolds of finite volume and curvature $\leq -a^2 < 0$. A good picture is that of truncated hyperbolic spaces defined in Chapter 12 (see Figure 12.1). These are metric spaces hyperbolic relative to the boundary horospheres. In general, one considers a geodesic metric space X and a collection \mathcal{A} of subsets of it (called *peripheral subsets* when the relative hyperbolicity conditions are fulfilled).

The metric definition of relative hyperbolicity consists of three conditions, the main one being very similar to the condition of thin triangles for hyperbolic spaces.

DEFINITION 11.219. We say that X is $(*)$ -relatively hyperbolic with respect to \mathcal{A} if for every $C \geq 0$ there exist two constants σ and δ such for every triangle $T \subset X$ with $(1, C)$ -quasi-geodesic edges, either there exists a point at distance at most σ from each of the sides of T , or there exists a subset $A \in \mathcal{A}$ such that its σ -neighborhood $\mathcal{N}_\sigma(A)$ intersects each of the sides of the triangle; moreover, for every vertex of the triangle, the two edges issuing from it enter $\mathcal{N}_\sigma(A)$ in two points at distance at most δ away from each other.

Clearly $(*)$ -relative hyperbolicity is a rather weak condition. For instance, every geodesic hyperbolic space is $(*)$ -hyperbolic relative to every family of subsets covering it.

DEFINITION 11.220. A space X is *hyperbolic relative* to \mathcal{A} if it is $(*)$ -hyperbolic relative to \mathcal{A} , and moreover, the following properties are satisfied:

- (α_1) For every $r > 0$, the r -neighborhoods of any two distinct subsets in \mathcal{A} intersect in a set of diameter at most $D = D(r)$.
- (α_2) Every geodesic of length ℓ with endpoints at distance at most $\frac{\ell}{3}$ from a set $A \in \mathcal{A}$, intersects the M -tubular neighborhood of A , with some uniform constant M .

DEFINITION 11.221. A finitely generated group G is *hyperbolic relative to* a finite set of subgroups H_1, \dots, H_n if, endowed with a word metric, G is hyperbolic in the sense of Definition 11.220 relative to the collection \mathcal{A} of left cosets

$$\{gH_i : i \in \{1, 2, \dots, n\}, g \in G\}.$$




FIGURE 11.17. The second case of Definition 11.219.

It follows from the definition that all H_i 's are finitely generated, since the three metric conditions imply that the peripheral subsets are quasi-convex. The groups H_i are called the *peripheral subgroups* of the relatively hyperbolic structure on G .

THEOREM 11.222 (C. Druţu, D. Osin, M. Sapir, [DS05b],[Osi06],[Dru09]). *Relative hyperbolicity in the sense of Definition 11.221 is equivalent to (strong) relative hyperbolicity as defined in [Gro87].*

Other characterizations of relative hyperbolicity can be found in the papers [Bow12], [Far98], [Dah03b], [DS05b], [Osi06]. Here and in what follows, by *relative hyperbolicity* we always mean *strong* relative hyperbolicity in the sense of Definition 11.220; we will always assume that *every H_i has infinite index in G* .

In the list of properties in Definition 11.220, one cannot drop the property (α_1) , as shown by the examples of groups in [OOS09] and in [BDM09, §7.1].

Many properties similar to those of hyperbolic groups are proved in the relatively hyperbolic case, in particular a Morse lemma, a characterization in terms of asymptotic cones [DS05b], of relative linear filling [Osi06], and of action on the boundary as a convergence group [Yam04].

Hyperbolic groups are clearly relatively hyperbolic with the peripheral subgroup $\{1\}$. Other examples of relatively hyperbolic groups include:

- (1) G is hyperbolic and each H_i is quasiconvex and *almost malnormal* in G (see [Far98]). Almost malnormality of a subgroup $H \leq G$ means that for every $g \in G \setminus H$,

$$|gHg^{-1} \cap H| < \infty.$$

- (2) G is the fundamental group of a finite graph of groups with finite edge groups; then G is hyperbolic relative to the vertex groups, see [Bow12].
- (3) Fundamental groups of complete finite volume manifolds of pinched negative curvature; the peripheral subgroups are the fundamental groups of their cusps ([Bow12], [Far98]).

- (4) Fully residually free groups, also known as limit groups of Sela; they have as peripheral subgroups a finite list of maximal abelian non-cyclic subgroups [Dah03a].

Similarly to hyperbolic groups, relatively hyperbolic groups are used to construct examples of infinite finitely generated groups with exotic properties. Denis Osin used in [Osi10] direct limits of relatively hyperbolic groups to construct torsion-free two-generated groups with exactly two conjugacy classes (i.e. all elements $\neq 1$ are conjugate to each other).

Study of relatively hyperbolic groups is a very active and rapidly developing area of Geometric Group Theory. We refer the reader to [Bow12, DG08, Dru09, BDM09, DS05b, DS05a, DS07, Ger09, Ger12, GP13, MR08, Osi06, Osi10, Yam04] for further reading.

Lattices in Lie groups

In Section 5.6.4 we defined lattices in general locally compact groups. In this chapter we consider lattices in Lie groups. While our main motivation comes from lattices in the Lie group $PO(n, 1)$, the isometry group of the hyperbolic n -space, most of our discussion here is general. Lattices in Lie groups (as well as in “ p -adic Lie groups”) play a prominent role in Geometric Group Theory for several reasons:

1. Rigidity theorems for lattices proven by Mostow and Margulis played a key role in the development of basic concepts and tools of geometric group theory.
2. Lattices act on homogeneous spaces, which provide nice *geometric models* for studying the coarse geometry of lattices themselves.
3. Having a nice geometric model helps to prove QI rigidity theorems for lattices: Somehow, the rigid geometric nature of homogeneous (primarily symmetric) spaces, translates into QI rigidity of lattices.

While the exact nature of QI rigidity for lattices in general connected Lie groups is still unclear, QI rigidity of lattices in semisimple Lie groups is now well-understood; see Chapter 25. We refer the reader to Gelander’s survey [Gel14] for a detailed review of properties of lattices in Lie groups.

12.1. Semisimple Lie groups and their symmetric spaces

Consider a Lie group G with a compact subgroup $K < G$. In view of the uniqueness (up to scaling) of the Haar measure μ on G , we can define μ as follows. Pick arbitrarily a positive definite bilinear form on the tangent space $T_e G$, where $e \in G$ is the identity element. Then, using the fact that G acts on itself smoothly and simply-transitively by the left multiplication, we spread this bilinear form from $T_e G$ to the rest of the tangent bundle TG . The result is a left-invariant Riemannian metric $h = \langle \cdot, \cdot \rangle$ on G and, hence, a G -invariant volume form. This volume form yields, by integration, the measure μ . This basic construction has an important modification. We let the compact subgroup $K < G$ act on G by the right multiplication. Compactness of K allows us to average the metric h :

$$Av(h) = \frac{1}{\text{mes}(K)} \int_K R_k^*(h) dk.$$

Here the integration is with respect to the Haar measure on K and R_k is the right multiplication:

$$R_k(g) = gk.$$

The metric $Av(h)$ is then both left-invariant with respect to the action of G and right-invariant with respect to the action of K . This *left-right* invariant Riemannian metric on G descends to a G -invariant Riemannian metric on the manifold $X = G/K$ and we obtain a homogeneous Riemannian manifold X . Conversely, as

explained in Section 5.6.4, one can lift the measure (defined via an invariant volume form) from X to a Haar measure on G .

The homomorphism

$$\rho : G \rightarrow \text{Isom}(X)$$

defined by the isometric action of G on X , is not, in general, injective, as G might have a normal subgroup contained in K . For instance, the action of the group $G = SL(2, \mathbb{R})$ on the hyperbolic plane $\mathbb{H}^2 \cong G/K$, $K = SO(2)$, has non-trivial kernel, equal to the center $\{\pm I\}$ of G . Furthermore, the image of the group G in $\text{Isom}(X)$ can have infinite index. Nevertheless, in view of the transitivity of the action of G on X and the compactness of K , both the kernel $\text{Ker}(\rho)$ of ρ and its *cokernel*, the quotient $\text{Isom}(X)/\rho(G)$, are compact.

EXAMPLE 12.1. Consider the group $SU(2)$ with a *biinvariant Riemannian metric*, i.e. a metric invariant under both left and right multiplication. The group $SU(2) \times SU(2)$ maps to $\text{Isom}(X)$, $X = SU(2)$, where the first factor acts via the left multiplication, while the second factor acts by the right multiplication. The kernel of the homomorphism $SU(2) \times SU(2) \rightarrow \text{Isom}(X)$ equals $\mathbb{Z}_2 = \langle (-I, -I) \rangle$, where $-I$ is the negative of the identity matrix in $SU(2)$. We leave it to the reader to verify that the Riemannian manifold X is isometric to a round 3-dimensional sphere (of some radius).

Some Lie groups, and their lattices, are better-behaved and more interesting than others. In Section 5.6.3 we defined the class of *semisimple Lie groups*. We now impose further conditions on the groups in this class:

CONVENTION 12.2. In order to simplify the terminology, from now on, when referring to a Lie group G as *semisimple*, we will always assume that G is linear (i.e. admits a monomorphism $G \rightarrow GL(N, \mathbb{R})$ for some N), has finitely many components and does not contain non-trivial compact connected normal subgroups.

The latter assumption guarantees that the kernel of $\rho : G \rightarrow \text{Isom}(X)$ is finite. Every semisimple group G contains a unique, up to conjugation, maximal compact subgroup K . Moreover, the quotient homogeneous space $X = G/K$ (with the projection of any left-right invariant metric on G) is a non-positively curved simply-connected complete Riemannian manifold. The G -invariant metric on X is essentially unique; for instance, if the group G is simple, then this metric is unique up to a multiplicative constant factor. The manifolds X obtained in this way are *symmetric spaces of non-compact type*. One interesting feature of such spaces is that $\rho(G)$ has finite index in $\text{Isom}(X)$. In other words, the homomorphism ρ has both finite kernel and cokernel.

The *rank of the symmetric space* X , $\text{rank}(X)$, is defined as the dimension of a *maximal flat* in the associated symmetric space $X = G/K$, i.e. the maximal r such that there exists an isometric embedding of the Euclidean r -space into X :

$$\mathbb{E}^r \rightarrow (X, \text{dist}_X),$$

where dist_X is the Riemannian distance function on X . For instance, X has rank one if and only if X is negatively curved. The *real rank*, $\text{rank}(G)$, of the group G can be defined geometrically as the rank of the symmetric space $X = G/K$. (There is an alternative algebraic definition of the rank of G , which we will not give here.)

DEFINITION 12.3. Suppose that X is a non-positively curved simply-connected symmetric space of non-compact type, i.e. $\text{Isom}(X)$ is a semisimple Lie group.

Then X admits a *de Rham decomposition*, i.e. an isometric splitting as the direct product

$$X = \prod_{i=1}^n X_i$$

of non-flat *irreducible symmetric spaces* X_i , i.e. spaces which themselves do not split as non-trivial products.

The de Rham splitting is preserved by the group $\text{Isom}(X)$, except that some factors might be permuted by isometries. Furthermore, the splitting is unique (up to permuting the indices) and it reflects the algebraic decomposition of the Lie algebra of G as the direct sum of simple subalgebras

$$\mathfrak{g} = \bigoplus_{i=1}^n \mathfrak{g}_i.$$

Namely, after rearranging the indices, for each i , the Lie algebra \mathfrak{g}_i is isomorphic to the Lie algebra of $\text{Isom}(X_i)$.

12.2. Lattices

Suppose now that G is a semisimple Lie group, $K < G$ is a compact subgroup, $X = G/K$ is the associated symmetric space, as in the previous section. Every subgroup Γ of G also acts isometrically on X and we will frequently identify Γ with its image, $\rho(\Gamma)$, in $\text{Isom}(X)$ (since ρ has finite kernel, this abuse of terminology is *mostly harmless*). Compactness of K (and, hence, finiteness of its Haar measure) has three immediate, but important, consequences:

LEMMA 12.4. 1. A subgroup $\Gamma < G$ is discrete if and only if Γ is discrete as a subgroup of $\text{Isom}(X)$, if and only if Γ acts properly discontinuously on X .

2. A discrete subgroup $\Gamma < G$ is a lattice if and only if the quotient $M = \Gamma \backslash X$ has finite volume.

3. A lattice $\Gamma < G$ is uniform (see Section 5.6.4) if and only if the quotient $\Gamma \backslash X$ is compact.

We have to warn the reader at this point that the quotient space M is, in general, not a manifold but an *orbifold*, since the group Γ can fail to act freely on X . If Γ acts freely on X , then we can compute the volume of M via the projection of the Riemannian metric from X . Otherwise, one can define the volume of M , say, by fundamental sets as in Section 5.6.4. It is a non-trivial fact that each lattice in G is finitely generated: This result is clear for uniform lattices in Lie groups with finitely many components, but it is difficult in general. Furthermore, T. Gelander [Gel11] established the following relation between the minimal number of generators of lattices and volumes of their quotient spaces:

THEOREM 12.5. $\text{rank}(\Gamma) \leq C \cdot \text{Vol}(\Gamma \backslash X)$, where C is a constant depending only on X .

Here $\text{rank}(\Gamma)$, the *rank* of Γ , is the minimum of cardinalities of its generating sets.

Since we are assuming that the Lie group G is linear, each lattice in G is virtually **torsion-free** (according to Selberg's Lemma, Theorem 7.116). For torsion-free lattices $\Gamma < G$, the projection $X \rightarrow M = \Gamma \backslash X$ is a covering map and the

Riemannian metric projects from X to a Riemannian metric on the smooth manifold M .

The quotients $M = \Gamma \backslash X$ of symmetric spaces X by discrete subgroups $\Gamma < G$ are called *locally-symmetric spaces*. Quotients defined by torsion-free subgroups Γ are Riemannian manifolds locally isometric to X .

EXERCISE 12.6. Recall that two subgroups Γ_1, Γ_2 of a group G are called *commensurable* if $\Gamma_1 \cap \Gamma_2$ has finite index in Γ_1, Γ_2 . Show that if two subgroups Γ_1, Γ_2 in a locally compact group G are commensurable, then Γ_1 is a lattice if and only if Γ_2 is a lattice. Furthermore, show that the lattice Γ_1 is uniform if and only if Γ_2 is uniform.

Let

$$X \cong \prod_{i=1}^n X_i$$

be the de Rham decomposition of the symmetric space X . A lattice Γ in the isometry group $\text{Isom}(X)$ is called *irreducible* if it is not commensurable to a lattice $\Gamma' < \text{Isom}(X)$ of the product form:

$$\Gamma' = \Gamma'_1 \times \Gamma'_2, \quad \Gamma'_1 < \text{Isom}(Y_1), \Gamma'_2 < \text{Isom}(Y_2),$$

where $X = Y_1 \times Y_2$ is a non-trivial product decomposition.

12.3. Examples of lattices

The most basic examples of lattices are given by the linear groups $SL(n, \mathbb{Z})$ with integer entries and their finite-index subgroups. Proving that $\Gamma = SL(n, \mathbb{Z})$ is a lattice in $G = SL(n, \mathbb{R})$ is not easy: While discreteness is clear, finiteness of volume of the quotient is not obvious. Proving finiteness requires constructing a certain subset S , called a *Siegel set* in G , which contains a fundamental domain of Γ . Then one verifies that S indeed has finite volume, from which it is immediate that $\text{Vol}(\Gamma \backslash G) < \infty$.

Arithmetic groups generalize this example and provide a rich and interesting source of lattices in all semisimple Lie groups.

DEFINITION 12.7. An *arithmetic subgroup* in a semisimple Lie group G is a subgroup of G commensurable to a subgroup of the form

$$\Gamma := \phi^{-1}(GL(N, \mathbb{Z})),$$

for a (continuous) homomorphism $\phi : G \rightarrow GL(N, \mathbb{R})$ with compact kernel.

As in the case of $SL(n, \mathbb{Z})$, it is clear that every arithmetic subgroup is discrete in G . It is a much deeper theorem that every arithmetic subgroup is a lattice in a Lie subgroup $H \leq G$, see e.g. [Mar91, Rag72]. We refer the reader to [Bor63] and [Rag72] for proofs of the following theorem:

THEOREM 12.8 (A. Borel). *Every semisimple Lie group G contains both uniform and non-uniform arithmetic lattices. Furthermore, G contains infinitely many commensurability classes of arithmetic lattices, both uniform and non-uniform.*

In other words, arithmetic lattices are ubiquitous. Arithmetic lattices are also interesting, since they provide connections between various fields of mathematics

(geometry, topology, analysis, ergodic theory) and number theory: Many number-theoretic results and conjectures can be stated (and proven!) in the form of properties of various lattices and their quotient spaces.

Fuchsian groups. We will refer to lattices in $PO(2, 1) = \text{Isom}(\mathbb{H}^2)$ as *Fuchsian groups*. Apart from the groups $SL(n, \mathbb{Z})$, these are the most studied lattices, whose investigation goes back to the second half of the nineteenth century. For instance, every finitely generated free group and the fundamental group of every closed connected surface of negative Euler characteristic, is isomorphic a Fuchsian group. If S is a compact oriented surface of genus $p > 1$, and Π_p is its fundamental group, then the space of conjugacy classes of isomorphisms from Π_p to lattices in $PSL(2, \mathbb{R}) < PO(2, 1)$ is a manifold of dimension $6p - 6$.

EXAMPLE 12.9. Consider the group $G = PO(2, 1)$ and a non-uniform lattice $\Gamma < G$. After passing to a finite-index subgroup in Γ , we may assume that Γ is torsion-free. Then the quotient \mathbb{H}^2/Γ is a non-compact surface with fundamental group Γ . Therefore, Γ is a free group of finite rank.

EXERCISE 12.10. Show that the groups Γ in the above example cannot be cyclic.

Bianchi groups. We now describe a very concrete class of non-uniform arithmetic lattices in the isometry group of the hyperbolic 3-space, called *Bianchi groups*. Let $D \in \mathbb{Z}$ be a *square-free negative integer*, i.e. an integer which is not divisible by the square of a prime number. Consider the *imaginary quadratic field*

$$\mathbb{Q}(\sqrt{D}) = \{a + \sqrt{D}b : a, b \in \mathbb{Q}\}$$

in \mathbb{C} . Set

$$\begin{aligned} \omega &:= \sqrt{D}, \text{ if } D \equiv 2, 3 \pmod{4} \\ \omega &:= \frac{1 + \sqrt{D}}{2}, \text{ if } D \equiv 1 \pmod{4} \end{aligned}$$

The ring of integers of $\mathbb{Q}(\sqrt{D})$ is

$$O_D = \{a + \omega b : a, b \in \mathbb{Z}\}.$$

For instance, if $D = -1$, then O_D is the ring of *Gaussian integers*

$$\{a + ib : a, b \in \mathbb{Z}\}.$$

A *Bianchi group* is a subgroup of the form

$$SL(2, O_D) < SL(2, \mathbb{C}),$$

for some D . Since the ring O_D is discrete in \mathbb{C} , it is immediate that every Bianchi subgroup is discrete in $SL(2, \mathbb{C})$. By abusing the terminology, one also refers to the group $PSL(2, O_D)$ as a Bianchi subgroup of $PSL(2, \mathbb{C})$.

Bianchi groups Γ are arithmetic lattices in $SL(2, \mathbb{C})$; in particular, the quotients \mathbb{H}^3/Γ have finite volume. Furthermore, every non-uniform arithmetic lattice in $SL(2, \mathbb{C})$ is commensurable to a Bianchi group. We refer the reader to [MR03] for the detailed discussion of these and other facts about Bianchi groups.

12.4. Rigidity and superrigidity

The example of Fuchsian groups shows that lattices in $PO(2,1)$ are highly flexible: They typically admit a continuum of non-conjugate representations (as lattices) into $PO(2,1)$. The theory of lattices in general semisimple Lie groups took off in the late 1950s, when it was discovered that, with the exception of the case of $PO(2,1)$, these lattices actually tend to be quite rigid. This development culminated in the fundamental rigidity theorems due to Mostow and Margulis which we recall below. Of course, in order to get rigidity, in addition to Fuchsian groups, one has to exclude their products in the products of $PO(2,1)$'s. This explains the *irreducibility* assumptions in rigidity theorems. In order to keep the statements simple, we first formulate the rigidity results in the context of *simple* Lie groups and, after that, for semisimple groups.

The proof of the next theorem can be found in Mostow's book [Mos73]

THEOREM 12.11 (G. D. Mostow, Strong Rigidity Theorem). *1. Let G_1, G_2 be connected linear non-compact simple Lie group with trivial centers, not isomorphic to $PSL(2, \mathbb{R})$. Then for any two lattices $\Gamma_1 < G_1, \Gamma_2 < G_2$, every isomorphism*

$$\phi : \Gamma_1 \rightarrow \Gamma_2$$

extends to an isomorphism $G_1 \rightarrow G_2$. Geometrically speaking, the isomorphism ϕ is induced by a similarity $f : X_1 \rightarrow X_2$ of the associated symmetric spaces $X_i = G_i/K_i$. (The mapping f becomes an isometry after one replaces the metric on X_2 by its appropriate scalar multiple.)

2. Assume that the groups G_i are connected semisimple, without non-trivial normal compact subgroups. Assume also that both G_1 and G_2 are not isomorphic to $PSL(2, \mathbb{R})$. Then for any two irreducible lattices $\Gamma_1 < G_1, \Gamma_2 < G_2$, every isomorphism

$$\phi : \Gamma_1 \rightarrow \Gamma_2$$

extends to an isomorphism $G_1 \rightarrow G_2$.

Mostow originally proved his theorem only for lattices in $G = PO(n,1)$, the isometry group of the hyperbolic n -space, $n \geq 3$. In Section 24.3 we will give another proof of Mostow's theorem for $PO(n,1)$.

In the case when the spaces X_i have rank 1 (i.e. are negatively curved), Mostow's proof is along the same lines as for the real-hyperbolic space: He first constructs an equivariant quasiisometry between symmetric spaces, then extends this quasiisometry to the ideal boundaries, establishes that the extension is quasi-conformal and then proves that this quasiconformal extension is, in fact, conformal. It is the last step where the assumption that X is not isometric to the hyperbolic plane is used. Mostow used ergodic theory arguments in the last step of his proof; we will be using the *zooming argument*, which seems to have its origin in Gromov's paper [Gro81b].

In the case of symmetric spaces of rank ≥ 2 Mostow's proof starts in the same fashion, but then he uses the theory of Tits buildings at infinity of X_i 's instead of quasiconformal analysis.

Since Mostow's pioneering work, other, very different, proofs of his rigidity theorem have emerged. For instance, for $X = \mathbb{H}^n$ there are very different proofs due to Gromov and Thurston [BP92] (based on bounded cohomology) and due to

Besson, Gourtois and Gallot [BCG96, BCG98]. The latter proof is differential-geometric in nature and avoids analyzing the boundary maps. In the complex-hyperbolic setting, Siu [Siu80] gave a proof using *Kähler geometry*; his proof is based on harmonic maps between complex-hyperbolic manifolds.

In his theorem Mostow assumes an isomorphism between two lattices. Margulis' Superrigidity Theorem below goes one step further: Margulis considers *arbitrary* homomorphisms from lattices $\Gamma < G$ into the group $GL(N, \mathbb{R})$, allowing, for instance, images to be non-discrete. Of course, there is a price to be paid for this level of generality on the side of homomorphisms, one has to restrict the class of Lie groups G . Namely, for every n , there exist arithmetic lattices Γ (both uniform and non-uniform) in the groups $PO(n, 1)$ and $PU(n, 1)$, such that Γ has infinite abelianization, i.e. there exists an epimorphism $\Gamma \rightarrow \mathbb{Z}$ (see [Mil76], [Kaz75]).

THEOREM 12.12 (G. Margulis, Archimedean Superrigidity Theorem). *1. Suppose that G is a simple connected (linear) Lie group and $\Gamma < G$ is a lattice. Assume, moreover, that G has rank at least two. Then for every homomorphism*

$$\phi : \Gamma \rightarrow GL(N, \mathbb{R}),$$

either the image $\phi(\Gamma)$ is relatively compact, or there exists a finite-index subgroup $\Gamma' < \Gamma$ such that the restriction $\phi|_{\Gamma'}$ extends to a homomorphism $G \rightarrow GL(N, \mathbb{R})$.

2. The same conclusion holds if the group G is semisimple, of rank ≥ 2 and $\Gamma < G$ is an irreducible lattice.

Margulis also proved a *non-archimedean superrigidity theorem*, which deals with representations of lattices Γ as above into the groups $GL(N, \mathbb{Q}_p)$, where the conclusion is exactly the same as before. Instead of trying to formulate the non-archimedean superrigidity theorem in full generality, we will only state a special case:

THEOREM 12.13. *Suppose that Γ is an irreducible lattice in a semisimple Lie group G of rank ≥ 2 . Then each action of Γ on a simplicial tree has a fixed point.*

The full non-archimedean superrigidity theorem states the existence of a fixed point for isometric actions of irreducible lattices on higher-dimensional generalizations of trees, which are called *Euclidean buildings*. As an application of these remarkable rigidity theorems, Margulis proved:

THEOREM 12.14 (G. Margulis, Arithmeticity Theorem). *Every lattice Γ satisfying the hypotheses of Theorem 12.12 is arithmetic.*

We refer the reader to Margulis' book [Mar91] for the proofs. The Margulis Arithmeticity Theorem was extended to lattices in the groups $\text{Isom}(\mathbb{H}^n)$ ($n \geq 2$) and $\text{Isom}(\mathbb{O}\mathbb{H}^2)$ by K. Corlette [Cor92] and by M. Gromov and R. Schoen [GS92]. The combination of these arithmeticity theorems yields:

THEOREM 12.15. *Suppose that $\Gamma < \text{Isom}(X)$ is an irreducible lattice, where X is a non-positively curved symmetric space not isometric (up to rescaling) to a real-hyperbolic space \mathbb{H}^n and complex-hyperbolic space $\mathbb{C}\mathbb{H}^n$. Then Γ is arithmetic.*

It follows from the work of M. Gromov and I. Piatetsky-Shapiro [GPS88] that for each $n \geq 3$, the group $PO(n, 1) = \text{Isom}(\mathbb{H}^n)$ contains infinitely many VI classes of both uniform and non-uniform non-arithmetic lattices. On the other hand, only finitely many VI classes of non-arithmetic lattices are known in $PU(2, 1)$ and

$PU(3,1)$ (the groups of biholomorphic isometries of the complex-hyperbolic plane and complex-hyperbolic 3-space). No non-arithmetic lattices are currently known in the groups $PU(n,1)$, $n \geq 4$.

12.5. Commensurators of lattices

Recall (see Section 5.2) that the commensurator of a subgroup Γ in a group G is the subgroup $Comm_G(\Gamma) < G$ consisting of elements $g \in G$ such that the groups $g\Gamma g^{-1}$ and Γ are commensurable, i.e.

$$|\Gamma : g\Gamma g^{-1} \cap \Gamma| < \infty$$

and

$$|g\Gamma g^{-1} : g\Gamma g^{-1} \cap \Gamma| < \infty.$$

Below we consider commensurators in the case Γ is a lattice in a Lie group G .

EXERCISE 12.16. Let $\Gamma := SL(2, O_D) < G := SL(2, \mathbb{C})$ be a Bianchi group.

1. Show that $Comm_G(\Gamma) < SL(2, \mathbb{Q}(\omega))$. In particular, $Comm_G(\Gamma)$ is dense in G .

2. Show that the set of fixed points of parabolic elements in Γ (in the upper half-space model of \mathbb{H}^3) is

$$\mathbb{Q}(\omega) \cup \{\infty\}.$$

3. Show that $Comm_G(\Gamma) = SL(2, \mathbb{Q}(\omega))$.

G. Margulis proved (see [Mar91], Chapter IX, Theorem B and Lemma 2.7; see also [Zim84], Theorem 6.2.5) that a lattice in a semisimple Lie group G is arithmetic if and only if its commensurator is dense in G .

12.6. Lattices in $PO(n,1)$

We now turn to the case of lattices in the isometry group $PO(n,1)$ of the real-hyperbolic n -space \mathbb{H}^n . The material discussed here will be used in Chapters 23, 24 in the proofs of QI rigidity theorems for lattices.

12.6.1. Zariski density. The next lemma is a basic result about discrete subgroups of $PO(n,1)$.

LEMMA 12.17. *Suppose that α is a hyperbolic isometry of \mathbb{H}^n and β is a parabolic isometry, which have a common fixed point ξ in the boundary sphere \mathbb{S}^{n-1} of \mathbb{H}^n . Then the subgroup $\Gamma < PO(n,1)$ generated by α and β is not discrete.*

PROOF. We will identify \mathbb{S}^{n-1} with $\mathbb{R}^{n-1} \cup \{\infty\}$ so that ξ corresponds to the point ∞ and the second fixed point of α corresponds to $0 \in \mathbb{R}^{n-1}$. Then α is a similarity

$$\mathbf{x} \mapsto \lambda A\mathbf{x}, A \in O(n-1), \lambda \neq 0, |\lambda| \neq 1,$$

and β is a (skew) translation

$$\mathbf{x} \mapsto B\mathbf{x} + \mathbf{v}, B \in O(n-1), \mathbf{v} \neq 0.$$

Here and below the \mathbf{x} 's are vectors in \mathbb{R}^{n-1} . Suppose first that $|\lambda| < 1$. Then consider the following sequence of conjugates of β in Γ :

$$\beta_k = \alpha^k \beta \alpha^{-k} : \mathbf{x} \mapsto C_k \mathbf{x} + \lambda^k A^k \mathbf{v},$$

where $C_k \in O(n-1)$ and, of course, $A^k \in O(n-1)$ as well. Since the group $O(n-1)$ is compact, the sequence C_k subconverges to some $C \in O(n-1)$. At the same time,

$$\lim_{k \rightarrow \infty} \lambda^k A^k \mathbf{v} = 0.$$

Therefore, a subsequential limit of the sequence of Moebius transformations (β_k) is the orthogonal transformation

$$\mathbf{x} \mapsto C\mathbf{x}.$$

It follows that Γ is a non-discrete subgroup of $Mob(\mathbb{S}^{n-1})$.

The case $|\lambda| > 1$ is similar: Instead of the sequence β_k above, one uses the sequence $\alpha^{-k}\beta\alpha^k$. \square

COROLLARY 12.18. *Suppose that $\alpha_1, \alpha_2 \in PO(n,1)$ are hyperbolic elements which generate a discrete subgroup of $PO(n,1)$ and have a common fixed point ξ in \mathbb{S}^{n-1} . Then α_1, α_2 share both fixed points; in particular, the subgroup they generate has an invariant geodesic in \mathbb{H}^n , asymptotic to their common fixed points.*

PROOF. Suppose that the fixed points of α_1 , respectively, α_2 that are different from ξ are distinct. The reader will verify that the commutator $\beta = [\alpha_1, \alpha_2]$ is a parabolic element of $PO(n,1)$. Since β clearly fixes ξ , we get a contradiction with Lemma 12.17. \square

THEOREM 12.19. *Suppose that n is at least 2. Then:*

1. *No lattice $\Gamma < PO(n,1)$ can have a proper invariant hyperbolic subspace $H \subset \mathbb{H}^n$.*
2. *No lattice $\Gamma < PO(n,1)$ can have a fixed point in the boundary sphere \mathbb{S}^{n-1} .*

PROOF. 1. Suppose that such a subspace H exists. We let $\pi : \mathbb{H}^n \rightarrow H$ be the nearest-point projection. Since the subspace H is Γ -invariant, the projection π is Γ -equivariant. The mapping π extends continuously to a Γ -equivariant mapping

$$\pi : Y := \mathbb{H}^n \cup (\mathbb{S}^{n-1} \setminus \partial_\infty H) \rightarrow H.$$

Proper discontinuity of the action of Γ on H implies proper discontinuity of the action of Γ on Y . Therefore, there exists a point $y \in Y \cap \mathbb{S}^{n-1}$ with trivial Γ -stabilizer and its neighborhood U in Y such that

$$\gamma U \cap U = \emptyset$$

for all $\gamma \in \Gamma \setminus \{1\}$. The neighborhood U contains metric balls $B(x, R) \subset U \cap \mathbb{H}^n$ of arbitrarily large radius R and, hence, arbitrarily large volume. Since the projection

$$\mathbb{H}^n \rightarrow M = \Gamma \backslash \mathbb{H}^n$$

is injective on the balls $B(x, R)$, we conclude that the space M has infinite volume, a contradiction.

2. The argument for fixed points at infinity is similar. Suppose that $\Gamma < PO(n,1)$ is a discrete subgroup fixing a point $\xi \in \mathbb{S}^{n-1}$. In view of Lemma 12.17, either all elements of Γ are parabolic and elliptic, or all its elements are hyperbolic and elliptic. We consider the former case and leave the latter to the reader as an exercise. Since Γ consists only of parabolic and elliptic elements, it preserves each horoball $B \subset \mathbb{H}^{n-1}$ centered at the point ξ . We now repeat the argument from part 1 using the nearest-point projection to B instead of the nearest-point projection to an invariant hyperbolic subspace. \square

COROLLARY 12.20. *If $\Gamma < PO(n, 1)$ is a lattice, it cannot have a finite orbit in \mathbb{S}^{n-1} .*

COROLLARY 12.21. *If $\Gamma < PO(n, 1)$ is a lattice, then Γ cannot contain non-trivial finite normal subgroups.*

PROOF. Suppose that $\Phi \triangleleft \Gamma$ is a finite normal subgroup. According to Corollary 3.75, the fixed-point set F of Φ in \mathbb{H}^n is non-empty. The set F is the intersection of fixed-point sets of the elements of Γ ; the latter are hyperbolic subspaces of \mathbb{H}^n . Therefore, F itself is a hyperbolic subspace in \mathbb{H}^n . Since the subgroup Φ is normal in Γ , the set F has to be invariant under Γ . By Theorem 12.19 a lattice in $PO(n, 1)$ cannot have a proper invariant hyperbolic subspace; it follows that $F = \mathbb{H}^n$, i.e. the subgroup Φ is trivial. \square

Note that every connected subgroup in $PO(n, 1)$ either has index 2 (i.e. is the subgroup $PO_o(n, 1)$ of orientation-preserving isometries of \mathbb{H}^n) or has a fixed point in \mathbb{S}^{n-1} , or has a proper invariant hyperbolic subspace in \mathbb{H}^n ; see [Gre62]. Therefore, we conclude that a lattice in $PO(n, 1)$ cannot be contained in a connected subgroup of $PO(n, 1)$, other than in $PO_o(n, 1)$.

The properties of lattices in $PO(n, 1)$ established above are elementary manifestations of a harder theorem, due to A. Borel [Bor60]:

THEOREM 12.22 (Borel Density Theorem). *Suppose that G is an algebraic Lie group. Then every lattice $\Gamma < G$ is Zariski dense.*

COROLLARY 12.23. *Every lattice Γ in a semisimple algebraic group G has finite center.*

PROOF. If Γ has infinite center, so does the Zariski closure $\bar{\Gamma} < G$. The Borel Density Theorem implies that G has infinite center. Since the center of an algebraic group is an algebraic subgroup, it follows that the center of G has positive dimension. By passing to the Lie algebra \mathfrak{g} of G , we conclude that the center of \mathfrak{g} is also non-trivial. This, however, contradicts the assumption that the group G is semisimple. \square

12.6.2. Parabolic elements and non-compactness. Consider the upper half-space model of the hyperbolic space \mathbb{H}^n . Recall that (open) horoballs in \mathbb{H}^n with center at the point $\infty \in \partial_\infty \mathbb{H}^n$ are Euclidean half-spaces of the form

$$B_t = \{(x_1, \dots, x_n) : x_n > t\}, \quad t > 0.$$

Accordingly, horospheres centered at the point ∞ are boundaries of horoballs:

$$\Sigma_t = \{(x_1, \dots, x_n) : x_n = t\}, \quad t > 0.$$

Define the projection $\Pi : \mathbb{R}_+^n \rightarrow \mathbb{R}^{n-1}$,

$$\Pi : (x_1, \dots, x_{n-1}, x_n) \mapsto (x_1, \dots, x_{n-1}).$$

For each $\mathbf{x} \in \mathbb{R}^{n-1}$ we set $\mathbf{x}(t) := \Pi^{-1}(\mathbf{x}) \cap \Sigma_t$:

$$\mathbf{x}(t) = (x_1, \dots, x_{n-1}, t),$$

where $\mathbf{x} = (x_1, \dots, x_{n-1})$.

LEMMA 12.24. *Suppose that $\Gamma < PO(n, 1)$ is a discrete subgroup containing a parabolic element γ . Then Γ cannot be a uniform lattice in $PO(n, 1)$.*

PROOF. Suppose to the contrary, that Γ is a uniform lattice. Without loss of generality, by conjugating Γ by an element of $PO(n, 1)$, we may assume that the unique fixed point of the parabolic element $\gamma \in \Gamma$ is the point $\infty \in \mathbb{S}^{n-1}$. Therefore, γ acts as a Euclidean isometry on \mathbb{R}_+^n , which, after conjugating further, has the form:

$$\gamma : \mathbf{x} \mapsto A\mathbf{x} + \mathbf{v}, \quad A \in O(n-1), \quad \mathbf{v} \in \mathbb{R}^{n-1} \setminus \{0\}, \quad A\mathbf{v} = \mathbf{v}.$$

The isometry γ preserves the Euclidean straight line $L \subset \mathbb{R}^{n-1}$ spanned by the vector \mathbf{v} . Furthermore, the restriction of γ to L is the translation $\mathbf{x} \mapsto \mathbf{x} + \mathbf{v}$. Then, for each $t > 0$ and $\mathbf{x} \in L$, by integrating the hyperbolic length element along the Euclidean line segment c_t connecting $\mathbf{x}(t)$ and $\gamma\mathbf{x}(t)$, we obtain:

$$\text{length}(c_t) = \frac{|\mathbf{v}|}{t}.$$

Therefore,

$$\lim_{t \rightarrow \infty} \text{dist}(\mathbf{x}(t), \gamma\mathbf{x}(t)) = 0.$$

Since the action $\Gamma \curvearrowright \mathbb{H}^n$ is cocompact, for each t exists $\alpha_t \in \Gamma$ such that

$$\text{dist}(\alpha_t(\mathbf{x}(t)), p) \leq R,$$

where $p \in \mathbb{H}^n$ is a base-point and R is a constant. Then the conjugate element of Γ

$$\gamma_t = \alpha_t \gamma \alpha_t^{-1}$$

moves the point $\alpha_t(\mathbf{x}(t)) \in B(p, R)$ by a distance not exceeding $\frac{|\mathbf{v}|}{t}$. In view of compactness of the ball $K = \bar{B}(p, R)$, there exists a sequence t_i diverging to infinity, such that the sequence

$$q_i = \alpha_{t_i}(\mathbf{x}(t_i))$$

converges to some $q \in \mathbb{H}^n$. Thus

$$\lim_{i \rightarrow \infty} \text{dist}(q, \gamma_{t_i}(q)) = 0.$$

Since the elements $\gamma_{t_i} \in \Gamma$ are not elliptic, we obtain a contradiction with the discreteness of the group Γ . \square

12.6.3. Thick-thin decomposition. The idea of the *thick-thin decomposition* of locally symmetric spaces $M = X/\Gamma$ is that such M splits naturally into a *thin part* M_{thin} , which has a reasonably simple topological structure, and a *thick part* M_{thick} which, typically, has a complicated topology, but whose geometry is bounded. When Γ is a lattice, the thick part of M turns out to be compact.

For simplicity, we will state and use the thick-thin decomposition only for lattices $\Gamma < PO(n, 1)$, even though a version of it also holds for general discrete subgroups of $PO(n, 1)$ and other semisimple Lie groups. Also, for simplicity in this section we consider all our quotient spaces to the right.

THEOREM 12.25 (Thick-thin decomposition). *Suppose that Γ is a non-uniform lattice in $PO(n, 1)$. Then:*

1. *There exists an (infinite) collection C of open horoballs $C := \{B_j, j \in J\}$, with pairwise disjoint closures, such that*

$$\Omega := \mathbb{H}^n \setminus \bigcup_{j \in J} B_j$$

is Γ -invariant and $M_c := \Omega/\Gamma$ is compact.

2. *Every parabolic element $\gamma \in \Gamma$ preserves (exactly) one of the horoballs B_j .*

The proof of this theorem is based on a mild generalization of the Zassenhaus Theorem due to Kazhdan and Margulis; see e.g. [BP92], [Kap01], [Rat06], [Thu97].

We note that the stabilizer Γ_j of each horoball B_j in this theorem cannot contain hyperbolic elements (since they do not preserve horoballs); therefore, Γ_j consists only of parabolic and elliptic elements. In view of the compactness of M_c , the quotient $T_j := \Sigma_j/\Gamma_j$ of each horosphere $\Sigma_j \subset \mathbb{H}^n$ bounding B_j , is compact. On the other hand, since Γ_j preserves horospheres with the same center as Σ_j , we have

$$\bar{B}_j/\Gamma_j \cong T_j \times [0, \infty),$$

where \cong above means ‘homeomorphic’.

The quotient M_c is called the *thick part*, M_{thick} , of $M = \mathbb{H}^n/\Gamma$, while its (non-compact) complement $M \setminus M_c$ is called the *thin part* of M . If Γ is torsion-free, then it acts freely on \mathbb{H}^n and M has a natural structure of a hyperbolic manifold of finite volume. If Γ is not torsion-free, then M is a *hyperbolic orbifold*. In view of the above observations, M is compact if and only if $M_{thin} = \emptyset$, equivalently, $C = \emptyset$,




FIGURE 12.1. Truncated hyperbolic space and thick-thin decomposition.

The set Ω is called a *truncated hyperbolic space*. The boundary horospheres Σ_j of Ω are called *peripheral horospheres*.

LEMMA 12.26. *The truncated hyperbolic space Ω is contractible.*

PROOF. Since all closed horoballs \bar{B}_j and all horospheres Σ_j are simply-connected, Seifert – van Kampen Theorem implies that $\pi_1(\mathbb{H}^n)$ is isomorphic to the group $\pi_1(\Omega)$. Hence, Ω is simply-connected. Vanishing of all homology groups $H_k(\Omega)$, $k \geq 2$, follows from the Mayer – Vietoris sequence. Therefore, the Hurewicz Theorem implies that Ω is contractible. \square

COROLLARY 12.27. *The group Γ is virtually torsion-free and has type \mathbf{F}_∞ .*

PROOF. The group Γ acts properly discontinuously and cocompactly on the simply-connected space Ω . Therefore, Γ is finitely generated. Since the group $PO(n, 1)$ is linear and Γ is finitely generated, the group Γ is virtually torsion-free (by Selberg’s Lemma). Let $\Gamma' < \Gamma$ be a finite index torsion-free subgroup. This subgroup acts smoothly, properly discontinuously, freely and cocompactly on the smooth manifold with boundary Ω . Therefore, the smooth quotient manifold Ω/Γ'

admits a finite triangulation. Contractibility of Ω now implies that Γ' has the type **F**. Since type \mathbf{F}_∞ is a virtual isomorphism invariant (Corollary 9.60), the group Γ also has type \mathbf{F}_∞ .

COROLLARY 12.28. *A lattice $\Gamma < PO(n, 1)$ is uniform if and only if it does not contain parabolic elements. Moreover, if Γ is non-uniform, it contains a parabolic subgroup isomorphic to \mathbb{Z}^{n-1} .*

PROOF. Since Γ acts cocompactly on Ω , it follows that each subgroup $\Gamma_j < \Gamma$ acts cocompactly on the corresponding horosphere Σ_j , which is isometric to \mathbb{R}^{n-1} . Therefore, Γ_j is isomorphic to a uniform lattice in $\text{Isom}(\mathbb{E}^{n-1})$. Bieberbach proved (see e.g. [Rat06, Theorem 7.5.2]) that each lattice in $\text{Isom}(\mathbb{E}^{n-1})$ contains a finite-index subgroup isomorphic to \mathbb{Z}^{n-1} . All non-trivial elements of this subgroup of Γ_j have to be parabolic since they preserve the horosphere Σ_j . \square

The next theorem is a sharpening of this corollary. We refer the reader to Section 11.13.4 for the definition of a conical limit point of a discrete group action on a hyperbolic space.

THEOREM 12.29. *If Γ is a lattice, then every point $\xi \in \partial_\infty \mathbb{H}^n$ is either a conical limit point or a parabolic fixed point.*

PROOF. Let ρ be a geodesic ray in \mathbb{H}^n asymptotic to ξ . This ray projects to a ray $\bar{\rho}$ in $M = \mathbb{H}^n/\Gamma$. Two things may occur:

Case 1: There exists $T \geq 0$ and a component $M_j = B_j/\Gamma_j$ of the thin part of M , such that for all $t \geq T$, $\bar{\rho}(t)$ belongs to M_j . Then the ray $\rho([T, \infty))$ is entirely contained in a Γ -translate B of the horoball B_j . However, if a horoball contains a geodesic ray, then this ray is asymptotic to the center of the horoball. It follows that the point ξ (to which ρ is asymptotic) is fixed by the subgroup $\Gamma_j < \Gamma$, which, as we know, contains parabolic elements.

Case 2: There exists a sequence $t_i \in \mathbb{R}_+$ diverging to ∞ such that for each i , $\bar{\rho}(t_i)$ belongs to the thick part M_c of M . Since M_c is compact, there exists a compact set $C \subset \mathbb{H}^n$ and a sequence of elements $\gamma_i \in \Gamma$ such that $\rho(t_i) \in \gamma_i(C)$. Arguing as in the proof of Lemma 11.118, we conclude that ξ is a conical limit point of Γ . (In Lemma 11.118 we assumed that the action of Γ on a Gromov-hyperbolic space is cobounded, but, in fact, all that we needed was a geodesic ρ , a compact set C and sequences t_i, γ_i as above.) \square

REMARK 12.30. The above theorem holds for all negatively curved symmetric spaces X and its converse holds as well (cf. [Bow95b]): A discrete subgroup $\Gamma < \text{Isom}(X)$ is a lattice if and only if every point of $\partial_\infty X$ is either a conical limit point of Γ or is a *bounded* parabolic fixed point.

12.7. Central coextensions

Recall that central coextensions

$$1 \rightarrow \mathbb{Z} \rightarrow \tilde{\Gamma} \rightarrow \Gamma \rightarrow 1$$

are classified by elements of the cohomology group $H^2(\Gamma, \mathbb{Z})$. In this section we describe some classes of lattices which admit non-trivial central coextensions.

For each subgroup $\Gamma' < \Gamma$ we have the *restriction homomorphism*:

$$H^2(\Gamma, \mathbb{Q}) \rightarrow H^2(\Gamma', \mathbb{Q})$$

defined by restricting cocycles to the subgroup Γ' . In general, this homomorphism may have large kernel (e.g., if we take $\Gamma' = \{1\}$). However, for finite-index subgroups $\Gamma' < \Gamma$ the behavior of cohomology classes is more predictable:

LEMMA 12.31. *Let $\Gamma' < \Gamma$ be a finite-index subgroup. Then the restriction homomorphism $H^2(\Gamma, \mathbb{Q}) \rightarrow H^2(\Gamma', \mathbb{Q})$ is injective.*

A proof of this lemma can be found, for instance, in [Bro82b, Chapter III, Proposition 10.4], it is an application of the *transfer argument*, which allows one to push cochains from Γ' to Γ by *averaging them* (this is where finite index and rational coefficients are used).

In particular, if a central coextension of Γ is given by a cohomology class which has non-zero projection to $H^2(\Gamma, \mathbb{Q})$, then this central extension remains non-trivial over every finite-index subgroups $\Gamma' < \Gamma$.

We will need a class of lattices with non-trivial second Betti numbers, i.e. non-vanishing $H^2(\Gamma, \mathbb{Q})$.

THEOREM 12.32. *1. For every $n \geq 2$, the group $PO(n, 1)$ contains uniform lattices Γ with non-vanishing $H^2(\Gamma, \mathbb{Q})$.*

2. If Γ is a torsion-free uniform lattice in $PU(n, 1)$, then $H^2(\Gamma, \mathbb{Q}) \neq 0$.

3. Every torsion-free uniform lattice $\Gamma < SO_o(n, 2)$ has non-zero $H^2(\Gamma, \mathbb{Q})$. Here $SO_o(n, 2)$ is the identity component of the Lie group $SO(n, 2)$.

PROOF. 1. For every $n \geq 2$, the group $PO(n, 1)$ contains torsion-free uniform lattices Γ with non-vanishing $H^2(\Gamma, \mathbb{Q}) \cong H^2(\mathbb{H}^n/\Gamma, \mathbb{Q})$; see [MR81].

2. A multiple of the Kähler form on the complex-hyperbolic space projects to the quotient manifold $M = \mathbf{C}\mathbb{H}^n/\Gamma$ and defines a non-zero element of $H^2(\Gamma, \mathbb{Q})$.

3. The same argument with the Kähler class applies in this case as well; see Toledo's appendix to [Ger92]. The non-zero cohomology class comes from the first Chern class in $H^2(M, \mathbb{Z})$, where M is the locally-symmetric space of Γ , $M = \Gamma \backslash SO(n, 2)/[SO(2) \times SO(n)]$. \square

COROLLARY 12.33. *Every lattice Γ in Theorem 12.32 admits a central coextension*

$$1 \rightarrow \mathbb{Z} \rightarrow \tilde{\Gamma} \rightarrow \Gamma \rightarrow 1$$

which satisfies the following properties:

a. The coextension does not split over any finite-index subgroup $\Gamma' < \Gamma$. In particular, $\tilde{\Gamma}$ is not virtually isomorphic to a product group $\Gamma' \times \mathbb{Z}$.

b. The group $\tilde{\Gamma}$ is quasiisometric to the product $\Gamma \times \mathbb{Z}$.

PROOF. (a) A multiple of a non-zero cohomology class $H^2(\Gamma, \mathbb{Q})$ is a non-trivial integral cohomology class. The latter defines a non-trivial central coextension

$$(12.1) \quad 1 \rightarrow A = \mathbb{Z} \rightarrow \tilde{\Gamma} \rightarrow \Gamma \rightarrow 1.$$

As noted above, this central coextension does not split over any finite-index subgroup $\Gamma' < \Gamma$. In order to see that the group $\tilde{\Gamma}$ is not isomorphic to the direct product $\Gamma \times \mathbb{Z}$, we note that Γ has trivial center (see Corollary 12.23). Therefore, any isomorphism

$$\tilde{\phi} : \tilde{\Gamma} \rightarrow \Gamma_1 \times A_1, \quad A_1 \cong \mathbb{Z}$$

would send A isomorphically to A_1 . Therefore, $\tilde{\phi}$ would project to an isomorphism $\phi : \Gamma \rightarrow \Gamma_1$. This would imply non-triviality of the central coextension (12.1),

resulting in a contradiction. The same argument applies to finite-index subgroups of Γ . Part (a) follows.

(b) The symmetric spaces \mathbb{H}^n and $\mathbf{C}\mathbb{H}^n$ associated with the Lie groups $PO(n, 1)$ and $PU(n, 1)$ are negatively curved. Hence, the uniform lattices $\Gamma < PO(n, 1), \Gamma < PU(n, 1)$ are Gromov–hyperbolic. Theorem 11.159 implies that each central coextension $\tilde{\Gamma}$ of such a lattice is quasiisometric to the direct product $\Gamma \times \mathbb{Z}$.

The argument in the case of lattices in $SO(n, 2)$ is less direct, since the associated symmetric space $X = SO(n, 2)/[SO(2) \times SO(n)]$ is not Gromov–hyperbolic. However, for each uniform torsion-free lattice $\Gamma < SO(n, 2)$, the non-trivial class in part 3 of Theorem 12.32 is *bounded*: It lies in the image of the natural homomorphism

$$H_b^2(\Gamma, \mathbb{Q}) \rightarrow H^2(\Gamma, \mathbb{Q}).$$

See Toledo’s appendix to [Ger92]. Therefore, the central coextension $\tilde{\Gamma}$ defined by the class ω is quasiisometric to the product $\Gamma \times \mathbb{Z}$; see Remark 11.164. \square

EXERCISE 12.34. Prove a generalization of Corollary 12.33 to central coextensions with the kernel \mathbb{Z}^k , $k \geq 2$.

Solvable groups

This chapter covers basic properties of general solvable groups and some special classes of solvable groups: Abelian, nilpotent and polycyclic groups. These properties will be used in proofs of theorems about growth of groups. Much of this material is algebraic rather than geometric, we decided to keep it in the book for the sake of completeness. Solvable and polycyclic groups appear naturally in the framework of *poly- X -groups*, where X is a certain class of groups: A group G is said to be poly- X if it admits a **subnormal descending series**:

$$G = G_0 \triangleright G_1 \triangleright \dots \triangleright G_k \triangleright G_{k+1} = \{1\},$$

such that each successive quotient G_i/G_{i+1} belongs to the class X . Solvable groups will be **obtained** by taking X to be the class of abelian groups, while polycyclic groups will use the class of cyclic groups (a further refinement of the definition uses X consisting of infinite cyclic groups, all isomorphic to each other, of course). As an aside, we note that there are other interesting classes of poly- X groups which we will not be discussing in the book, like *poly-free* groups, important examples of which are given by the *pure braid groups*.

Notation. For abelian groups G we will frequently use the notation mg or $m \cdot g$ for the m -fold sum

$$\underbrace{g + \dots + g}_{m \text{ times}}$$

with $m \in \mathbb{N}$. This extends to $m \in \mathbb{Z}$ by declaring that

$$0 \cdot g = 0 \in G$$

and that

$$-(m \cdot g) = (-m) \cdot g.$$

13.1. Free abelian groups

DEFINITION 13.1. A group G is called *free abelian* on a generating set S if it is isomorphic to the direct sum

$$\bigoplus_{s \in S} \mathbb{Z}.$$

The minimal cardinality of S is called *the rank of G* and denoted $\text{rank}(G)$, the set S is called a *basis* of G .

Of course, if $|S| = n$, $G \cong \mathbb{Z}^n$. Given an abelian group G , we define its subgroup

$$2G = \{2x : x \in G\}.$$

Clearly, this subgroup is *characteristic* in G , i.e. is invariant under all automorphisms of G . Then, for the free abelian group $G = \bigoplus_{s \in S} \mathbb{Z}$, the quotient $G/2G$ is isomorphic to

$$\bigoplus_{s \in S} \mathbb{Z}_2,$$

which has natural structure of a vector space over \mathbb{Z}_2 with basis S . Since any two bases of a vector space have the same cardinality, it follows that two bases of a free abelian group have the same cardinality, equal to $\text{rank}(G)$.

EXERCISE 13.2. Every free abelian group is torsion-free.

Below is a characterization of free abelian groups by a *universality property*:

THEOREM 13.3. *Let G be an abelian group and X is a subset of G . The group G is free abelian with basis X if and only if it satisfies the following universality property: For every abelian group A , every map $f : X \rightarrow A$ extends to a unique homomorphism $f : G \rightarrow A$.*

PROOF. Suppose that G is free abelian with the basis X . Every element $g \in G$ is uniquely represented as a sum

$$g = \sum_{x \in X} m_x \cdot x, m_x \in \mathbb{Z}$$

with only finitely many non-zero terms. Then we extend f to G by

$$f(g) = \sum_{x \in X} m_x \cdot f(x).$$

It is clear that this extension is unique.

Conversely, assume that $(G_1, X_1), (G_2, X_2)$ satisfy the universality property and $f : X_1 \rightarrow X_2$ is a bijection. Then f and $f^{-1} = \bar{f} : X_2 \rightarrow X_1$ admit homomorphic extensions $F : G_1 \rightarrow G_2, \bar{F} : G_2 \rightarrow G_1$ respectively. The compositions $\bar{F} \circ F, F \circ \bar{F}$ are homomorphisms $\phi : G_1 \rightarrow G_1, \psi : G_2 \rightarrow G_2$, respectively. These homomorphisms extend the identity maps $X_2 \rightarrow X_2, X_1 \rightarrow X_1$. By the uniqueness part of the universality property, it follows that ϕ and ψ are the identity maps. Therefore, the homomorphism $F : G_1 \rightarrow G_2$ is an isomorphism. Applying this to $G_1 = G, X_1 = X$ and G_2 equal to the free abelian group with the basis $X_2 = X_1 = X$, we conclude that G is free abelian with the basis X . \square

COROLLARY 13.4. *Let $0 \rightarrow A \rightarrow B \xrightarrow{r} C \rightarrow 0$ be a short exact sequence of abelian groups, where C is free abelian. Then this sequence splits and $B \cong A \oplus C$.*

PROOF. Let $c_i, i \in I$, denote a basis of C . Then, since r is surjective, for every c_i there exists $b_i \in B$ such that $r(b_i) = c_i$. By the universal property of free abelian groups, the map $s : c_i \rightarrow b_i$ extends to a homomorphism $s : C \rightarrow B$ such that $r \circ s = \text{Id}$. \square

EXERCISE 13.5. Show that a group G is free abelian with the basis S if and only if G admits the presentation

$$\langle S \mid [s, s'] = 1, \forall s, s' \in S \rangle.$$

The following theorem is the abelian analogue of the Nielsen–Schreier theorem (Theorem 7.41), although, we are unaware of a topological or geometric proof:

THEOREM 13.6. 1. *Subgroups of free abelian groups are again free abelian.*
 2. *If $G < F$ is a subgroup of a free abelian group F , then $\text{rank}(G) \leq \text{rank}(F)$.*

PROOF. Let X be a basis of a free abelian group $F = A_X$. For each subset Y of X let A_Y be the free group with the basis Y , thus A_Y embeds naturally as a free abelian subgroup A_Y in F . We fix a subgroup $G \leq F$ once and for all; for each $Y \subset X$ we let G_Y denote the intersection $G \cap A_Y$.

Define the set S consisting of triples (G_Y, B, ϕ) , where Y ranges over the set of all subsets of X such that G_Y is free with a basis of cardinality at most the cardinality of X ; the sets B are bases of such G_Y , and ϕ is an injective map $\phi : B \rightarrow X$.

The set S is non-empty, as we can take $Y = \emptyset$.

We define a partial order \leq on S by:

$$(G_Y, B, \phi) \leq (G_Z, C, \psi) \iff Y \subset Z, B \subset C, \quad \phi = \psi|_B.$$

Suppose that L is a chain in the above order indexed by an ordered set M :

$$\{(G_{Y_m}, B_m, \phi_m), m \in M\}, (G_{Y_m}, B_m, \phi_m) \leq (G_{Y_n}, B_n, \phi_n) \iff m \leq n.$$

Then the union

$$\bigcup_{m \in M} G_{Y_m}$$

is again a subgroup in F and the set

$$C = \bigcup_{m \in M} B_m$$

is a basis in the above group. Furthermore, the maps ϕ_m determine an embedding $\psi : C \hookrightarrow X$. Thus,

$$\left(\bigcup_{m \in M} G_{Y_m}, C, \psi \right) \in S.$$

Therefore, by Zorn's Lemma, there exists a maximal element (G_Y, B, ϕ) of S . If $Y = X$ then $G_Y = G$ and we are done. Suppose that there exists $x \in X \setminus Y$. Set $Z := Y \cup \{x\}$. We will show that G_Z is still free abelian with a basis C containing B and ϕ extends to an embedding $\psi : Z \rightarrow X$. If $G_Z = G_Y$, we take $C = B$, $\psi = \phi$. Otherwise, assume that $G_Z/G_Y \neq 0$. The quotient A_Z/A_Y is isomorphic to \mathbb{Z} and is generated by the image \bar{x} of x . The image of G_Z in this quotient is isomorphic to G_Z/G_Y and is generated by some $n \cdot \bar{x}$, $n \in \mathbb{Z} \setminus 0$. Let $g \in G_Z$ be an element which maps to $n \cdot \bar{x}$. The mapping $G_Z/G_Y \rightarrow \langle g \rangle$ splits the sequence

$$0 \rightarrow G_Y \rightarrow G_Z \rightarrow G_Z/G_Y = \mathbb{Z} \rightarrow 0$$

and, hence,

$$G_Z \cong G_Y \oplus \langle g \rangle.$$

This means that $C := B \cup \{g\}$ is a basis of G_Z ; we extend ϕ to C by $\psi(g) = x$. Thus, $(G_Z, C, \psi) \in S$. This contradicts maximality of (G_Y, B, ϕ) .

We conclude that G is free abelian and its basis embeds in a basis of F . \square

13.2. Classification of finitely generated abelian groups

THEOREM 13.7. *Every finitely generated abelian group A is isomorphic to a finite direct sum of cyclic groups.*

PROOF. The proof below is taken from [Mil12]. The proof is induction on the number of generators of A .

If A is 1-generated, the assertion is clear. Assume that the assertion holds for abelian groups with $\leq n - 1$ generators and suppose that A is an abelian group generated by n elements. Consider all ordered generating sets (a_1, \dots, a_n) of A . Among such generating sets choose one, $S = (a_1, \dots, a_n)$, such that the order of a_1 (denoted $|a_1|$) is the least possible. We claim that

$$A \cong \langle a_1 \rangle \oplus A' = \langle a_1 \rangle \oplus \langle a_2, \dots, a_n \rangle.$$

(This claim will imply the assertion since, inductively, A' splits as a direct sum of cyclic groups.) Indeed, if A is not the direct sum as above, then we have a non-trivial relation

$$(13.1) \quad \sum_{i=1}^n r_i a_i = 0, r_i \in \mathbb{Z}, r_1 a_1 \neq 0.$$

Without loss of generality, $0 < r_1 < |a_1|$ and $r_i \geq 0, i = 1, \dots, n$ (otherwise, we replace a_i 's with $-a_i$ whenever $r_i < 0$). Furthermore, let $d = \gcd(r_1, \dots, r_n)$ be the greatest common divisor of the numbers $r_i, i = 1, \dots, n$. Set $q_i := \frac{r_i}{d}$.

LEMMA 13.8. *Suppose that a_1, \dots, a_n are generators of A and $q_1, \dots, q_n \in \mathbb{Z}_+$ are such that $\gcd(q_1, \dots, q_n) = 1$. Then there exists a new generating set b_1, \dots, b_n of A such that*

$$b_1 = \sum_{i=1}^n q_i a_i.$$

PROOF. The proof of this lemma is a form of the Euclid's algorithm for computation of gcd. Note that $q := q_1 + \dots + q_n \geq 1$. The proof of lemma is induction on q . If $q = 1$ then $b_1 \in \{a_1, \dots, a_n\}$ and lemma follows. Suppose the assertion holds for all $q < m$, we will prove the claim for $q = m > 1$. After rearranging the indices, we can assume that $q_1 \geq q_2 > 0$.

Clearly, the set $\{a_1, a_1 + a_2, a_3, \dots, a_n\}$ generates A . Furthermore,

$$\gcd(q_1 - q_2, q_2, q_3, \dots, q_n) = 1$$

and

$$q' := (q_1 - q_2) + q_2 + q_3 + \dots + q_n < m$$

Thus, by the induction hypothesis, there exists a generating set b'_1, \dots, b'_n of A , where

$$b'_1 = (q_1 - q_2)a_1 + q_2(a_1 + a_2) + q_3 a_3 + \dots + q_n a_n.$$

However, $b_1 = b'_1$. Lemma follows. \square

In view of this lemma, we get a new generating set b_1, \dots, b_n of A such that

$$b_1 = \sum_{i=1}^n \frac{r_i}{d} a_i.$$

The equation (13.1) implies that $db_1 = 0$ and $d \leq r_1 < |a_1|$. Thus, the ordered generating set (b_1, \dots, b_n) of A has the property that $|b_1| < |a_1|$, contradicting our choice of S . Theorem follows. \square

For a prime p , an abelian group A is called a p -group if every element $a \in A$ has the order which is a power of p . Clearly, each subgroup and each quotient of a p -group is again a p -group.

EXERCISE 13.9. A finite abelian group A is a p -group if and only if $|A| = p^\ell$ for some ℓ .

Given an abelian group A , we let $A(p)$ denote the subset of A consisting of elements whose order is a power of p . Since the sum of two elements of the orders p^k, p^m has the order p^n , where $n = \max(k, m)$, the subset $A(p)$ is a subgroup of A . A group T is said to be a *torsion group* if every element of T has finite order. For every abelian group G , the set $\text{Tor}(G)$ of finite-order elements is a subgroup T of G , called the *torsion subgroup* $T \leq G$. This subgroup of G is characteristic.

EXERCISE 13.10. Every finitely generated abelian torsion group is finite.

THEOREM 13.11 (classification of abelian groups). *Suppose that A is a finitely generated abelian group. Then there exist an integer $r \geq 0$, and k -tuples of prime numbers (p_1, \dots, p_k) and natural numbers (m_1, \dots, m_k) , for which*

$$(13.2) \quad A \simeq \mathbb{Z}^r \times \mathbb{Z}_{p_1}^{m_1} \times \cdots \times \mathbb{Z}_{p_k}^{m_k}.$$

Here $p_1 \leq p_2 \leq \dots \leq p_k$, and whenever $p_i = p_{i+1}$, we have $m_i \geq m_{i+1}$. Furthermore, the number r , and the k -tuples (p_1, \dots, p_k) and (m_1, \dots, m_k) are uniquely determined by A .

PROOF. By Theorem 13.7, A is isomorphic to the direct product of finitely many cyclic groups

$$C_1 \times \cdots \times C_r \times C_{r+1} \times \cdots \times C_n,$$

where C_i is infinite cyclic for $i \leq r$ and finite cyclic for $i > r$.

EXERCISE 13.12. (Chinese remainder theorem) $\mathbb{Z}_s \times \mathbb{Z}_t \cong \mathbb{Z}_{st}$ if and only if the numbers s, t are coprime.

In view of this exercise, we can split every finite cyclic group C_i as a direct product of cyclic groups whose orders are prime powers. This proves existence of the decomposition (13.2).

We now consider the uniqueness part of the theorem. We first note that

$$\text{Tor}(A) = C_{r+1} \times \cdots \times C_n,$$

which implies that

$$C_1 \times \cdots \times C_r \simeq \mathbb{Z}^r \simeq A/\text{Tor}(A).$$

Since the subgroup $\text{Tor}(A)$ is characteristic in A , it follows that the number r is uniquely determined by A .

Thus, in order to prove uniqueness of p_i 's and m_i 's it suffices to assume that A is finite. Since the primes p_i are the prime divisors of the order of A , the uniqueness question reduces to the case when $|A| = p^\ell$, i.e. when $A = A(p)$ is an abelian p -group. Suppose that A is an abelian p -group and

$$A \cong \mathbb{Z}_{p^{m_1}} \times \cdots \times \mathbb{Z}_{p^{m_k}}, \quad m_1 \geq \dots \geq m_k.$$

Set $m = m_1$ and let $m_1 = m_2 = \dots = m_d > m_{d+1}$. Clearly, the number p^m is the largest order of an element of A . The subgroup A_m of A generated by elements of

this order is clearly characteristic and equals the d -fold direct product of copies of \mathbb{Z}_p^m ,

$$\mathbb{Z}_p^{m_1} \times \cdots \times \mathbb{Z}_p^{m_d}$$

in the above factorization of A . Hence, the number m_k and the number d depend only on the group A . We then divide A by A_m and proceed by induction. \square

EXERCISE 13.13. The number r equals the rank of a maximal free abelian subgroup of A .

We will refer to the number r as the *free rank* of the abelian group A , in order to distinguish it from the notion of rank in Definition 7.1. Theorem 13.7 implies that each finitely generated abelian group is isomorphic to a direct sum of finitely many cyclic groups C_i , which are unique up to an isomorphism.

DEFINITION 13.14. Generators of cyclic subgroups C_i such that

$$A = \bigoplus_{i=1}^s C_i$$

will be called *standard generators* of A . (These generators, of course, are not uniquely determined by A .)

Below are several immediate corollaries of Theorem 13.7.

COROLLARY 13.15. *Each finite abelian group A is isomorphic to the direct product of abelian p -groups:*

$$A \simeq A(p_1) \times \cdots \times A(p_k),$$

where p_1, \dots, p_k are the prime divisors of $|A|$.

COROLLARY 13.16. *Every finitely generated abelian group G is polycyclic, i.e. G possesses a finite descending series*

$$(13.3) \quad G = N_0 \supseteq N_1 \supseteq \cdots \supseteq N_n \supseteq N_{n+1} = \{1\},$$

such that every quotient N_i/N_{i+1} is cyclic.

COROLLARY 13.17. *Every finitely generated abelian group A is isomorphic to the direct product $F \times \text{Tor}(A)$, where F is a free abelian group.*

COROLLARY 13.18. *A finitely generated abelian group is free abelian if and only if it is torsion-free.*

EXERCISE 13.19. 1. Show that the torsion-free abelian group \mathbb{Q} is not a free abelian group.

2. Show that the image of the free abelian group F in A is not a characteristic subgroup of A (unless $A \simeq F$ or $A = \text{Tor}(A)$).

COROLLARY 13.20. *Let G be an abelian group generated by n elements. Then every subgroup H of G is finitely generated (with $\leq n$ generators).*

PROOF. Theorem 13.3 implies that there exists an epimorphism $\phi : \mathbb{Z}^n \rightarrow A$. Let $A := \phi^{-1}(H)$. Then, by Theorem 13.6, the subgroup A is free of rank $m \leq n$. Therefore, H is also m -generated. \square

REMARK 13.21. Groups where every subgroup is finitely generated are called *noetherian*. We will see that all polycyclic groups has this property; we will discuss noetherian groups in more detail in Section 13.8.

EXERCISE 13.22. Construct an example of a finitely generated abelian group G and a subgroup $H \leq G$, such that there is no direct product decomposition $G = F \times \text{Tor}(G)$ for which $H = (F \cap H) \times (\text{Tor}(G) \cap H)$. Hint: Take $G = \mathbb{Z} \times \mathbb{Z}_2$ and H infinite cyclic.

EXERCISE 13.23. Let F be a free abelian group of rank n and $B = \{x_1, \dots, x_n\}$ be a generating set of F . Then B is a basis of F . Conclude that n equals the minimal cardinality of all generating sets of F . Thus, the notion of *rank* for (finitely generated) free abelian groups agrees with the notion of rank introduced in the beginning of Section 7.1.

The classification of finitely generated abelian groups allows one find a simple geometric model for such groups:

LEMMA 13.24. *Every finitely generated abelian group G of free rank n admits a geometric (in the sense of Definition 5.68) action on the Euclidean space \mathbb{E}^n , such that every element of G acts as a translation. In particular, G is quasiisometric to \mathbb{E}^n .*

PROOF. Let $G = \mathbb{Z}^n \times \text{Tor}(G)$. We let $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ denote a basis of \mathbb{Z}^n , and let \mathbb{R}^n be the vector space with the basis $\mathbf{e}_1, \dots, \mathbf{e}_n$. We equip \mathbb{R}^n with the standard Euclidean metric where the basis $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ is orthonormal and let \mathbb{E}^n be the corresponding Euclidean n -space. Then every $g = \sum_{i=1}^n a_i \mathbf{e}_i \in \mathbb{Z}^n$ acts on \mathbb{E}^n as the translation by the vector $\mathbf{a} = (a_1, \dots, a_n)$. This action of \mathbb{Z}^n extends to G by declaring that every $g \in \text{Tor}(G)$ acts on \mathbb{E}^n trivially. We leave it to the reader to check that this action is geometric and the quotient \mathbb{E}^n/G is the n -torus T^n . \square

13.3. Automorphisms of \mathbb{Z}^n

THEOREM 13.25. *The group of automorphisms of \mathbb{Z}^n is isomorphic to $GL(n, \mathbb{Z})$.*

PROOF. Consider the basis $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ of \mathbb{Z}^n , where

$$\mathbf{e}_i = (\underbrace{0, \dots, 0}_{i-1 \text{ times}}, 1, \underbrace{0, \dots, 0}_{n-i \text{ times}}).$$

Let $\phi : \mathbb{Z}^n \rightarrow \mathbb{Z}^n$ be an automorphism. Set

$$(13.4) \quad \phi(\mathbf{e}_i) = \sum_{j=1}^n m_{ij} \mathbf{e}_j.$$

We, thus, obtain a map $\mu : \phi \mapsto M_\phi = (m_{ij})$, where M_ϕ is a matrix with integer entries. We leave it to the reader to check that $\mu(\phi \circ \psi) = M_\phi M_\psi$. It follows that $\mu(\phi) \in GL(n, \mathbb{Z})$ for every $\phi \in \text{Aut}(\mathbb{Z}^n)$.

Given a matrix $M \in GL(n, \mathbb{Z})$, we define an endomorphism

$$\phi : \mathbb{Z}^n \rightarrow \mathbb{Z}^n,$$

using the equation (13.4). Since the map $\nu : M \mapsto \phi$ respects the composition, it follows that $\nu : GL(n, \mathbb{Z}) \rightarrow \text{Aut}(\mathbb{Z}^n)$ is a homomorphism and $\mu = \nu^{-1}$. \square

Below we establish several properties of automorphisms of free abelian groups that are interesting by themselves and will also be useful in Chapter 14, in the proof of the Milnor–Wolf Theorem.

LEMMA 13.26. Let $\mathbf{v} = (v_1, \dots, v_n) \in G = \mathbb{Z}^n$ be a vector with $\gcd(v_1, \dots, v_n) = 1$. Then $H = G/\langle \mathbf{v} \rangle$ is free abelian of rank $n - 1$. Moreover, there exists a basis $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{n-1}, \mathbf{v}\}$ of G such that $\{\mathbf{y}_1 + \langle \mathbf{v} \rangle, \dots, \mathbf{y}_{n-1} + \langle \mathbf{v} \rangle\}$ is a basis of H .

PROOF. First, let us show that the group H is free abelian; since this group is finitely generated, it suffices to verify that it is torsion-free. We will use the notation $x \mapsto \bar{x}$ for the quotient map $G \rightarrow H$.

Let $u \in G$ be such that $\bar{u} \in H$ has finite order k . Then $ku \in \langle \mathbf{v} \rangle$, i.e. $ku = m\mathbf{v}$ for some $m \in \mathbb{Z}$. Since $\gcd(v_1, \dots, v_n) = 1$, it follows that $k|m$ and, hence, $u \in \langle \mathbf{v} \rangle$, $\bar{u} = \bar{1}$.

Thus, $H = \mathbb{Z}^n/\langle \mathbf{v} \rangle$ is torsion-free, and, hence, it is free abelian of finite rank m . Next, the homomorphism $G \rightarrow H$ extends to a surjective linear map $\mathbb{R}^n \rightarrow \mathbb{R}^m$, whose kernel is the line spanned by v . Therefore, $m = n - 1$.

Let $\{\bar{x}_1, \dots, \bar{x}_{n-1}\}$ be a basis on H . The map

$$\bar{x}_i \mapsto x_i, i = 1, \dots, n - 1,$$

extends to a group monomorphism $H \rightarrow G$; thus, the set $\{x_1, \dots, x_{n-1}, v\}$ generates \mathbb{Z}^n . It follows that $\{x_1, \dots, x_{n-1}, v\}$ is a basis of G . \square

LEMMA 13.27. If a matrix M in $GL(n, \mathbb{Z})$ has all eigenvalues equal to 1 then there exists a finite ascending series of subgroups

$$\{1\} = \Lambda_0 \leq \Lambda_1 \leq \dots \leq \Lambda_{n-1} \leq \Lambda_n = \mathbb{Z}^n$$

such that $\Lambda_i \simeq \mathbb{Z}^i$, $\Lambda_{i+1}/\Lambda_i \simeq \mathbb{Z}$ for all $i \geq 0$, $M(\Lambda_i) = \Lambda_i$ and M acts on Λ_{i+1}/Λ_i as the identity.

PROOF. Since M has eigenvalue 1, there exists a vector $v = (v_1, \dots, v_n) \in \mathbb{Z}^n$ such that $\gcd(v_1, \dots, v_n) = 1$ and $Mv = v$. Then M induces an automorphism of $H = \mathbb{Z}^n/\langle v \rangle \simeq \mathbb{Z}^{n-1}$ and the matrix \bar{M} of this automorphism has only 1 as an eigenvalue. This follows immediately when writing the matrix of the automorphism M with respect to a basis $\{x_1, x_2, \dots, x_{n-1}, v\}$ of \mathbb{Z}^n as in Lemma 13.26 and looking at the characteristic polynomial. Now, lemma follows by induction on n . \square

The following lemma is a special case of a classical result of L. Kronecker; see [Kro57] or Proposition 1.2.1 in [GdlHJ89]. Our proof follows Kronecker's original argument.

LEMMA 13.28. Let $M \in GL(n, \mathbb{Z})$ be a matrix such that each eigenvalue of M has absolute value 1. Then all the eigenvalues of M are roots of unity.

PROOF. Recall that for each $n \times n$ matrix A with the eigenvalues μ_1, \dots, μ_n (here and below, we repeat the eigenvalues if necessary, according to their multiplicities) the characteristic polynomial $p_A(t)$ equals

$$\sum_{i=0}^n a_{n-i} t^i,$$

where, by Vieta's formulae,

$$a_i = \det(A)(-1)^n \sigma_i(\mu_1, \dots, \mu_n),$$

and σ_i is the i th elementary symmetric polynomial:

$$\sigma_i(x_1, \dots, x_n) = \sum_{1 \leq j_1 < \dots < j_i \leq n} x_{j_1} \dots x_{j_i}.$$

We now return to the integer square matrix M as in lemma and let $\lambda_1, \dots, \lambda_n$ denote its eigenvalues. Consider the sequence of matrices $M^k, k \in \mathbb{N}$. The eigenvalues of M^k are $\lambda_1^k, \dots, \lambda_n^k$, which, by the assumption, all have the absolute value 1. Therefore, the coefficients of the characteristic polynomials $p_k(t) := p_{M^k}(t)$ of M^k are uniformly bounded, independently on k . Since the matrices M^k belong to $GL(n, \mathbb{Z})$, there are only finitely many distinct characteristic polynomials of the matrices M^k . Hence, there exists an infinite sequence $k_1 < k_2 < k_3 < \dots$, such that

$$p_{k_1}(t) = p_{k_2}(t) = p_{k_3}(t) = \dots$$

It follows that there are distinct members of this sequence, $q < r$, such that

$$\lambda_1^q = \lambda_1^r, \dots, \lambda_n^q = \lambda_n^r.$$

Hence, for each $i = 1, \dots, n$

$$\lambda_i^{r-q} = 1,$$

which means that each eigenvalue of M is a root of unity. \square

LEMMA 13.29. *If a matrix M in $GL(n, \mathbb{Z})$ has one eigenvalue λ of absolute value at least 2 then there exists a vector $\mathbf{v} \in \mathbb{Z}^n$ such that the following map is injective:*

$$(13.5) \quad \begin{aligned} \Phi : \bigoplus_{n \in \mathbb{Z}_+} \mathbb{Z}_2 &\longrightarrow \mathbb{Z}^n \\ \Phi : (s_n)_n &\mapsto s_0 \mathbf{v} + s_1 M \mathbf{v} + \dots + s_n M^n \mathbf{v} + \dots \end{aligned}$$

PROOF. The matrix M defines an automorphism $\varphi : \mathbb{Z}^n \rightarrow \mathbb{Z}^n$, $\varphi(\mathbf{v}) = M \mathbf{v}$. The dual map φ^* has the matrix M^T in the dual canonical basis. Therefore, it also has the eigenvalue λ and, hence, there exists a linear form $f : \mathbb{C}^n \rightarrow \mathbb{C}$ such that $\varphi^*(f) = f \circ \varphi = \lambda f$.

Take $\mathbf{v} \in \mathbb{Z}^n \setminus \text{Ker } f$. Assume that the map Φ is not injective. It follows that there exist some $(t_n)_n$, $t_n \in \{-1, 0, 1\}$, such that

$$t_0 \mathbf{v} + t_1 M \mathbf{v} + \dots + t_n M^n \mathbf{v} + \dots = 0.$$

Let N be the largest integer such that $t_N \neq 0$. Then

$$M^N \mathbf{v} = r_0 \mathbf{v} + r_1 M \mathbf{v} + \dots + r_{N-1} M^{N-1} \mathbf{v}$$

where $r_i \in \{-1, 0, 1\}$. By applying f to the equality we obtain

$$(r_0 + r_1 \lambda + \dots + r_{N-1} \lambda^{N-1}) f(\mathbf{v}) = \lambda^N f(\mathbf{v}),$$

whence

$$|\lambda|^N \leq \sum_{i=1}^{N-1} |\lambda|^i = \frac{|\lambda|^N - 1}{|\lambda| - 1} \leq |\lambda|^N - 1,$$

a contradiction. \square

13.4. Nilpotent groups

Recall that $[x, y] = xyx^{-1}y^{-1}$ is the commutator of the elements x, y in a group G and that $x^g := gxg^{-1}$ is the g -conjugate of x in G . We begin the discussion of nilpotent groups with some useful commutator identities:

LEMMA 13.30. *Let (G, \cdot) be a group and x, y, z elements in G . The following identities hold:*

- (1) $[x, y]^{-1} = [y, x]$;
- (2) $[x^{-1}, y] = [x^{-1}, [y, x]] [y, x]$;
- (3) $[x, yz] = [x, y] [y, [x, z]] [x, z]$;
- (4) $[xy, z] = [x, [y, z]] [y, z] [x, z]$.
- (5) $[x, y]^g = [x^g, y^g]$.

PROOF. (1) and (2) are immediate, (4) follows from (3) and (1). It remains to prove (3). Since $[y, [x, z]] [x, z] = y[x, z]y^{-1}$ we have that

$$[x, y] [y, [x, z]] [x, z] = xyx^{-1}[x, z]y^{-1} = xyzx^{-1}z^{-1}y^{-1} = [x, yz].$$

We leave the last identity as an exercise to the reader. \square

NOTATION 13.31. For every x_1, \dots, x_n in a group G we denote by $[x_1, \dots, x_n]$ the n -fold left-commutator

$$[[[x_1, x_2], \dots, x_{n-1}], x_n].$$

We declare that the 1-fold left commutator $[x]$ is simply x .

EXERCISE 13.32. $[x_1, \dots, x_n]^g = [x_1^g, \dots, x_n^g]$.

Recall that for subsets A, B in a group G , $[A, B]$ denotes the subgroup of G generated by all the commutators $[a, b]$, $a \in A, b \in B$. In what follows we also use:

NOTATION 13.33. Given n subgroups H_1, H_2, \dots, H_n in a group G we denote by $[H_1, \dots, H_n]$ the subgroup $[\dots [H_1, H_2], \dots, H_n] \leq G$.

We define the lower central series of a group G ,

$$C^1G \supseteq C^2G \supseteq \dots \supseteq C^mG \supseteq \dots,$$

inductively by:

$$C^1G = G, \quad C^{n+1}G = [C^nG, G].$$

In particular, each C^kG is a characteristic subgroup of G . We will see later on (Proposition 13.62) that

$$[C^iG, C^kG] \leq C^{i+k}G.$$

Note that $C^2G = [G, G] = G'$ is the commutator subgroup, or the derived subgroup, of G .

- EXERCISE 13.34. 1. The subgroup $C^kG \leq G$ is normal in G .
 2. $C^{m+1}G = [G, C^mG]$.

DEFINITION 13.35. A group G is called k -step nilpotent if $C^{k+1}G = \{1\}$. The minimal k for which G is k -step nilpotent is called the (nilpotency) class of G .

- EXAMPLES 13.36. (1) Every non-trivial abelian group is nilpotent of class 1.
 (2) The group $\mathcal{U}_n(\mathbb{K})$ of upper triangular $n \times n$ matrices with 1 on the diagonal and entries in a ring \mathbb{K} , is nilpotent of class $n - 1$ (see Exercise 13.38).

(3) The *Heisenberg group*

$$H_{2n+1}(\mathbb{K}) = \left\{ \begin{pmatrix} 1 & x_1 & x_2 & \dots & \dots & x_n & z \\ 0 & 1 & 0 & \dots & \dots & 0 & y_n \\ 0 & 0 & 1 & \dots & \dots & 0 & y_{n-1} \\ \vdots & \vdots & \ddots & \ddots & & \vdots & \vdots \\ 0 & 0 & \dots & \dots & 1 & 0 & y_2 \\ 0 & 0 & \dots & \dots & 0 & 1 & y_1 \\ 0 & 0 & \dots & \dots & \dots & 0 & 1 \end{pmatrix} ; x_1, \dots, x_n, y_1, \dots, y_n, z \in \mathbb{K} \right\}$$

is nilpotent of class 2.

Taking $\mathbb{K} = \mathbb{Z}$, we obtain the *integer Heisenberg group*

$$H_{2n+1}(\mathbb{Z}).$$

The group $H_{2n+1}(\mathbb{Z})$ is finitely generated; we can take as generators the elementary matrices $N_{ij} = I + E_{ij}$ with

$$(i, j) \in \{(1, 2), \dots, (1, n+1), (2, n), \dots, (n+1, n)\}.$$

All the groups $H_{2n+1}(\mathbb{K})$ are nilpotent of class 2. Indeed $C^2 H_{2n+1}(\mathbb{K})$ is the subgroup $x_i = y_i = 0, i = 1, \dots, n$.

(4) We will see later (Proposition 14.1) that semidirect products $\mathbb{Z}^n \rtimes_A \mathbb{Z}$ with the matrix $A \in \mathcal{U}_n(\mathbb{Z})$, are nilpotent.

EXERCISE 13.37. Which of the permutation groups S_n are nilpotent? Which of these groups are solvable?

EXERCISE 13.38. The goal of this exercise is to prove that the group $\mathcal{U}_n(\mathbb{K})$ is nilpotent of class $n - 1$.

Let $\mathcal{U}_{n,k}(\mathbb{K})$ be the subset of $\mathcal{U}_n(\mathbb{K})$ formed by matrices (a_{ij}) such that $a_{ij} = \delta_{ij}$ for $j < i + k$. Note that $\mathcal{U}_{n,1}(\mathbb{K}) = \mathcal{U}_n(\mathbb{K})$.

(1) Prove that for every $k \geq 1$ the map

$$\begin{aligned} \varphi_k : \mathcal{U}_{n,k}(\mathbb{K}) &\rightarrow (\mathbb{K}^{n-k}, +) \\ A = (a_{i,j}) &\mapsto (a_{1,k+1}, a_{2,k+2}, \dots, a_{n-k,n}) \end{aligned}$$

is a homomorphism. Deduce that $(\mathcal{U}_{n,k}(\mathbb{K}))' \subset \mathcal{U}_{n,k+1}(\mathbb{K})$ and that $\mathcal{U}_{n,k+1}(\mathbb{K}) \triangleleft \mathcal{U}_{n,k}(\mathbb{K})$ for every $k \geq 1$.

(2) Let E_{ij} be the matrix with all entries 0 except the (i, j) -entry, which is equal to 1. Consider the triangular matrix $T_{ij}(a) = I + aE_{ij}$.

Deduce from (1), using induction, that $\mathcal{U}_{n,k}$ is generated by the set

$$\{T_{ij}(a) \mid j \geq i + k, a \in \mathbb{K}\}.$$

(3) Prove that for every three distinct numbers i, j, k in $\{1, 2, \dots, n\}$

$$[T_{ij}(a), T_{jk}(b)] = T_{ik}(ab), [T_{ij}(a), T_{ki}(b)] = T_{kj}(-ab),$$

and that for all quadruples of distinct numbers i, j, k, ℓ ,

$$[T_{ij}(a), T_{k\ell}(b)] = I.$$

(4) Prove that $C^k \mathcal{U}_n(\mathbb{K}) \leq \mathcal{U}_{n,k+1}(\mathbb{K})$ for every $k \geq 0$. Deduce that $\mathcal{U}_n(\mathbb{K})$ is nilpotent.

Remark. All the arguments above work also when all matrices have integer entries. In this case (2) implies that $\mathcal{U}_n(\mathbb{Z})$ is generated by $\{T_{ij}(1) \mid j \geq i + 1\}$.

EXERCISE 13.39. The group $\mathcal{U}_n(\mathbb{K})$ is torsion-free provided that \mathbb{K} has zero characteristic.

A combination of deep theorems by Mal'cev and Ado shows that each finitely generated torsion-free nilpotent group embeds in \mathcal{U}_n , for some n :

THEOREM 13.40 (A. I. Mal'cev [Mal49b]). *Every finitely generated torsion-free nilpotent group Γ of class k embeds as a uniform lattice in a simply-connected nilpotent Lie group N of class k . Furthermore, the group N and the embedding $\Gamma \rightarrow N$ are unique up to an isomorphism.*

THEOREM 13.41 (Ado–Engel theorem). *Every simply-connected nilpotent Lie group N embeds into $\mathcal{U}_n(\mathbb{R})$ for some n .*

REMARK 13.42. We are attributing this theorem to Ado and Engel, but, as usual, the history is more complicated. Ado (see Theorem 5.59) proved linearity of finite-dimensional real Lie algebras; in the case of nilpotent Lie algebras \mathfrak{n} , the faithful linear representation

$$r : \mathfrak{n} \rightarrow \text{End}(\mathbb{R}^n) = \mathfrak{gl}_n(\mathbb{R})$$

sends each element of \mathfrak{n} to a nilpotent linear transformation, i.e. a linear endomorphism A such that $A^k = 0$ for some k . Much earlier, Engel sketched a proof, details of which were written by his student, Umlauf in his PhD thesis [Uml91], that any subalgebra of $\mathfrak{gl}_n(\mathbb{R})$ consisting entirely of nilpotent endomorphisms is conjugate to a subalgebra of the algebra of upper-triangular matrices with zeroes on the diagonal, we refer to [FH94, Theorem 9.9] for a modern proof, cf. Theorem 14.43. Thus, we can assume that $r(\mathfrak{n})$ is contained in the algebra u_n of such matrices. Now, if N is a simply-connected Lie group with the Lie algebra \mathfrak{n} , then, via exponentiation, r induces a representation $\rho : N \rightarrow \mathcal{U}_n(\mathbb{R})$, which has to be faithful, since the exponential map $\exp : u_n \rightarrow \mathcal{U}_n(\mathbb{R})$ is bijective. Note that this proof, in particular, implies that N is contractible, since its exponential map has to be a homeomorphism as well.

Since simply-connected nilpotent Lie groups are contractible, it follows that each finitely generated torsion-free nilpotent group Γ has type \mathbf{F} , i.e. admits a finite $K(\Gamma, 1)$, namely, N/Γ . In particular, the cohomological dimension $cd(\Gamma)$ of Γ is at most $n = \dim(N)$. Since N/Γ is a closed orientable manifold, $H^n(\Gamma) \cong H^n(N/\Gamma) \cong \mathbb{Z}$. Therefore,

$$cd(\Gamma) = \dim(N).$$

COROLLARY 13.43. *Each finitely generated torsion-free nilpotent group is residually finite.*

We will see later on, Theorem 13.77, that all polycyclic groups are residually finite, which shows residual finiteness of all finitely generated nilpotent groups.

We now proceed with establishing some basic properties of lower central series and nilpotent groups.

LEMMA 13.44. *If S is a generating set of a group G (not necessarily nilpotent), then for every k the subgroup $C^k G$ is generated by the k -fold left commutators in S and their inverses, together with $C^{k+1} G$.*

PROOF. We prove the assertion by induction on k . For $k = 1$ the statement is clear, since 1-fold commutators of elements of S are just elements of S . Assume that the assertion holds for some $k \geq 1$ and consider $C^{k+1}G$.

By definition, $C^{k+1}G$ is generated by all commutators $[c_k, g]$ with $c_k \in C^kG$ and $g \in G$. The induction hypothesis and normality of $C^{k+1}G$ in G imply that $c_k = \ell_1^{\pm 1} \cdots \ell_m^{\pm 1}x$, where $m \in \mathbb{N}$, ℓ_i are k -fold left commutators in S and $x \in C^{k+1}G$.

According to Lemma 13.30, (4),

$$[c_k, g] = [\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}x, g] = [\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}, [x, g]][x, g][\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}, g].$$

The first two factors are in $C^{k+2}G$, so it remains to deal with the third.

We write $g = s_1 \cdots s_r$, where $s_i \in S$, and we prove that $[\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}, s_1 \cdots s_r]$ is a product of $(k+1)$ -fold left commutators in S and their inverses, and of elements in $C^{k+2}G$; our proof is another induction, this time on $m+r \geq 2$.

For the case $m+r = 2$ it suffices to note that $[\ell^{-1}, s] = [\ell^{-1}, [s, \ell]][s, \ell]$. The first factor is in $C^{k+2}G$, the second is the inverse of a $(k+1)$ -fold left commutator.

Assume that the statement is true for $m+r = n \geq 2$. We now prove it for $m+r = n+1$.

Suppose that $m \geq 2$. We apply Lemma 13.30, (4), and obtain that

$$[\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}, s_1 \cdots s_r] = [\ell_1^{\pm 1} \cdots \ell_{m-1}^{\pm 1}, [\ell_m^{\pm 1}, g]][\ell_m^{\pm 1}, s_1 \cdots s_r][\ell_1^{\pm 1} \cdots \ell_{m-1}^{\pm 1}, s_1 \cdots s_r].$$

The first factor is in $C^{k+2}G$, and for the second and the third the induction hypothesis applies.

Likewise, if $r \geq 2$ then we apply Part 3 of Lemma 13.30, and write

$$\begin{aligned} & [\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}, s_1 \cdots s_r] = \\ & [\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}, s_1 \cdots s_{r-1}][s_1 \cdots s_{r-1}, [\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}, s_r]][\ell_1^{\pm 1} \cdots \ell_m^{\pm 1}, s_r]. \quad \square \end{aligned}$$

COROLLARY 13.45. *If G is nilpotent, then C^nG is generated by k -fold left commutators in S and their inverses, where $k \geq n$. In particular, if G is finitely generated, so is each group C^nG .*

PROOF. Suppose that $C^{m+1}G = \{1\}$. Then C^mG is generated by the m -fold left commutators in S and their inverses. By applying the reverse induction in n , each C^nG is generated by the set of all k -fold left commutators of elements of S and their inverses, $k \geq n$. \square

Thus, if G is finitely generated, each quotient $C^iG/C^{i+1}G$ is a finitely generated abelian group and, hence, we define two important invariants of finitely generated nilpotent groups:

DEFINITION 13.46. Let G be a finitely generated nilpotent group of class k . Let m_i denote the free rank of the abelian group $C^iG/C^{i+1}G$; define the *Hirsch length* of G

$$h(G) = \sum_{i=1}^k m_i$$

and the *homogeneous dimension* of G ,

$$d(G) = \sum_{i=1}^k im_i.$$

In the next chapter we will give a geometric interpretation of the number $d(G)$. For now, we note that for torsion-free nilpotent groups G , the Hirsch length $h(G)$ equals the dimension of the simply-connected nilpotent group N into which G embeds as a uniform lattice, [Mal49b]; hence, $h(G)$ is the cohomological dimension of G in this case.

DEFINITION 13.47. Given natural numbers k and m , the k -step m -generated free nilpotent group is the quotient $N_{m,k}$ of the free group of rank m , F_m , by the normal subgroup $C^{k+1}F_m$.

We will refer to the images of the free generators of F_m as *free nilpotent generators* of $N_{m,k}$.

Note that the free abelian group of rank m is the 1-step m -generated free nilpotent group.

A consequence of Proposition 7.21 is the following.

PROPOSITION 13.48 (Universality property of free nilpotent groups). *For every k -step nilpotent group G equipped with a generating set $X = \{x_1, \dots, x_m\}$, there exists an epimorphism $\psi : N_{m,k} \rightarrow G$ sending free nilpotent generators s_1, \dots, s_m of $N_{m,k}$ to the generators x_1, \dots, x_m respectively. In particular, every k -step m -generated nilpotent group is a quotient of $N_{m,k}$.*

PROOF. Take a generating set X of a k -step nilpotent group G , such that X has cardinality m . The homomorphism $\phi : F(S) = F_m \rightarrow G$, sending $s_i \mapsto x_i, i = 1, \dots, m$, defined in Proposition 7.21 contains $C^{k+1}F(S)$ in its kernel. Therefore, ϕ projects to an epimorphism $\psi : N_{m,k} \rightarrow G$ as required. \square

So far, we were describing nilpotent groups “from the top-down”, starting from the group G and then looking at the chain of decreasing subgroups. It is also useful to have a “bottom-up” description of nilpotent groups, which we present below.

Recall that the center of a group H is denoted $Z(H)$. Given a group G , consider the sequence of normal subgroups $Z_i(G) \triangleleft G$ defined inductively by:

- $Z_0(G) = \{1\}$.
- If $Z_i(G) \triangleleft G$ is defined and $\pi_i : G \rightarrow G/Z_i(G)$ is the quotient map, then

$$Z_{i+1}(G) = \pi_i^{-1}(Z(G/Z_i(G))).$$

Note that $Z_{i+1}(G)$ is normal in G , as the preimage of a normal subgroup of a quotient of G . In particular,

$$Z_{i+1}(G)/Z_i(G) \cong Z(G/Z_i(G)).$$

PROPOSITION 13.49. *The group G is k -step nilpotent if and only if $Z_k(G) = G$.*

PROOF. Assume that G is nilpotent of class k . We prove by induction on $i \geq 0$ that $C^{k+1-i}G \leq Z_i(G)$. For $i = 0$ we have equality. Assume that

$$C^{k+1-i}G \leq Z_i(G).$$

For every $g \in C^{k-i}G$ and every $x \in G$, $[g, x] \in C^{k+1-i}G \leq Z_i(G)$, whence $gZ_i(G)$ is in the center of $G/Z_i(G)$, i.e. $g \in Z_{i+1}(G)$. Thence, the inclusion follows by induction. For $i = k$ the inclusion becomes $C^1G = G \leq Z_k(G)$, hence, $Z_k(G) = G$.

Conversely, assume that there exists k such that $Z_k(G) = G$. We prove by induction on $j \geq 1$ that $C^jG \leq Z_{k+1-j}(G)$. For $j = 1$ the two are equal. Assume

that the inclusion is true for j . The subgroup $C^{j+1}G$ is generated by commutators $[c, g]$ with $c \in C^jG$ and $g \in G$. Since $c \in C^jG \leq Z_{k+1-j}(G)$, by the definition of $Z_{k+1-j}(G)$, the element c commutes with g modulo $Z_{k-j}(G)$, equivalently $[c, g] \in Z_{k-j}(G)$. This implies that $[c, g] \in Z_{k-j}(G)$. It follows that $C^{j+1}G \leq Z_{k-j}(G)$.

For $j = k+1$ this gives $C^{k+1}G \leq Z_0(G) = \{1\}$, hence G is k -step nilpotent. \square

DEFINITION 13.50. The ascending series

$$Z_0(G) = \{1\} \triangleleft Z_1(G) \triangleleft \dots \triangleleft Z_i(G) \triangleleft Z_{i+1}(G) \triangleleft \dots$$

of normal subgroups of G is called the *upper central series* of G .

In view of Proposition 13.49, a group G is nilpotent if and only if its upper central series is finite, and its nilpotency class is the minimal k such that $Z_k(G) = G$.

EXERCISE 13.51. Any central coextension of a nilpotent group is again nilpotent.

REMARK 13.52. Yet another equivalent definition of a nilpotent group is to require that the group admits a finite normal series

$$\{1\} = \Gamma_0 \triangleleft \dots \triangleleft \Gamma_i \triangleleft \Gamma_{i+1} \triangleleft \dots \triangleleft \Gamma_{n-1} \triangleleft \Gamma_n = G,$$

such that $\Gamma_{i+1}/\Gamma_i \leq Z(G/\Gamma_i)$, or, equivalently, $[G, \Gamma_{i+1}] \leq \Gamma_i$. In particular, the quotients Γ_{i+1}/Γ_i are abelian for each i . We will need only the fact that the existence of such a normal series implies that G is n -step nilpotent. Indeed, the condition $\Gamma_{i+1}/\Gamma_i \leq Z(G/\Gamma_i)$ implies that $\Gamma_i \leq Z_i(G)$ for every i . In particular, $G = Z_n(G)$. Now, the assertion follows from Proposition 13.49. We refer to [Hal76, Theorem 10.2.2] for further details.

The following example shows that the difference between lower and upper central series of groups can be quite substantial, in particular, $C^{k+1-i}G \leq Z_i(G)$ could be of infinite index:

EXAMPLE 13.53. We start with the integer Heisenberg group H ; it is 2-step nilpotent, $C^2H = H' = Z(H) \cong \mathbb{Z}$. Next, take $G = H \times \mathbb{Z}$. Then G is still 2-step nilpotent, but now $C^2G = C^2H \cong \mathbb{Z}$, while $Z(G) \cong \mathbb{Z}^2$.

EXERCISE 13.54. Construct an example of a 2-step nilpotent group G with torsion-free center, such that G/C^2G is not torsion-free.

The following useful lemma is a converse to Corollary 13.45:

LEMMA 13.55. Let S be a generating set of a group G . Suppose that all $N+1$ -fold commutators $[s_1, \dots, s_{N+1}]$ of elements of S are trivial. Then G is N -step nilpotent.

PROOF. Let G_n be the subgroup of Γ generated by the n -fold commutators $y_n = [s_1, \dots, s_n]$ of generators $s_i \in S$ of the group G . For every generator x of G and every generator y_n of G_n we have:

$$[y_n, x] = y_n x y_n^{-1} x^{-1} \in G_{n+1} \leq G_n.$$

Since $y_n \in G_n$, it follows that $x y_n^{-1} x^{-1} \in G_n$ which implies that G_n is a normal subgroup of G .

We claim that for every n , G_{n-1}/G_n embeds (under the map induced by inclusion $G_{n-1} \hookrightarrow G$) in the center of G/G_n . To simplify the notation, we will

regard G_{n-1}/G_n as a subgroup of G/G_n . The proof of this statement is the reverse induction on n .

The subgroup G_{N+1} is trivial, hence it is contained in the center of G . Suppose that the assertion holds for $n = k + 1$, we will now prove it for $n = k$. To show that G_{k-1}/G_k is in the center of G/G_k it is enough to verify that for all elements \bar{z} and \bar{w} of generating sets of G_{n-1}/G_n and G/G_n , respectively, the commutator $[\bar{z}, \bar{w}]$ is trivial.

The group G is generated by the set S , the group G_{n-1} is generated by the $n-1$ -fold commutators y_{n-1} of elements $x \in S$. Thus, the groups G_{n-1}/G_n and G/G_n are generated by the projections \bar{x}, \bar{y}_{n-1} of the elements x, y_{n-1} . By definition of G_n we have: $[y_{n-1}, x] \in G_n$, thus, dividing by G_n , we obtain $[\bar{y}_{n-1}, \bar{x}] = 1$. Thus, $G_{n-1}/G_n \leq Z(G/G_n)$ for every n and Lemma follows from Remark 13.52. \square

- LEMMA 13.56. (1) *Every subgroup of a nilpotent group is nilpotent.*
 (2) *If G is nilpotent and $N \triangleleft G$ then G/N is nilpotent.*
 (3) *The direct product of a family of nilpotent groups is again nilpotent.*

PROOF. (1) Let H be a subgroup in a nilpotent group G . Then $C^i H \leq C^i G$. Hence, if G is k -step nilpotent then $C^{k+1} H = \{1\}$.

(2) If $\pi : G \rightarrow G/N$ is the quotient map, $\pi(C^i G) = C^i(G/N)$.

(3) The assertion follows from the equality

$$C^j\left(\prod_{i \in I} G_i\right) = \prod_{i \in I} C^j G_i. \quad \square$$

THEOREM 13.57. *Every subgroup of a finitely generated nilpotent group is finitely generated, i.e. finitely generated nilpotent groups are noetherian.*

PROOF. We argue by induction on the class of nilpotency k . For $k = 1$ the group is abelian and the statement is already proven in Corollary 13.20. Assume that the assertion holds for k , let G be a nilpotent group of class $k+1$ and let $H \leq G$ be a subgroup. By the induction hypothesis $H_1 = H \cap C^2 G$ and $H_2 = H/(H \cap C^2 G)$ are both finitely generated. Thus, H fits in the short exact sequence

$$1 \rightarrow H_1 \rightarrow H \xrightarrow{\pi} H_2 \rightarrow 1,$$

where H_1, H_2 are finitely generated. Lemma 7.10 then shows that H is also finitely generated. \square

Our next goal is to prove some structural results for nilpotent groups. We begin the “calculus of commutators.”

LEMMA 13.58. *If A, B, C are normal subgroups in a group G , then the subgroup $[A, B, C] \leq G$ is generated by the commutators $[a, b, c]$ with $a \in A, b \in B, c \in C$.*

PROOF. By the definition, $[A, B, C]$ is generated by the commutators $[k, c]$ with $k \in [A, B]$ and $c \in C$. The element k is a product $t_1 \cdots t_n$, where each t_i is equal either to a commutator $[a, b]$ or to a commutator $[b, a]$, $a \in A, b \in B$.

We prove, by the induction on n , that $[k, c]$ is a product of finitely many commutators $[a, b, c]$ and their inverses. For $n = 1$ we only need to consider the case $[t^{-1}, c]$, where $t = [a, b]$. By Lemma 13.30, (2),

$$[t^{-1}, c] = [c, t]^{t^{-1}} = [c^{t^{-1}}, t] = [c', t] = [a, b, c']^{-1}.$$

In the second equality above we applied the identity $\phi([x, y]) = [\phi(x), \phi(y)]$ for the inner automorphism $\phi(x) = x^{t^{-1}}$.

Assume that the statement is true when k is a product of n commutators t_i and consider $k = k_1 t$, where t is equal to either a commutator $[a, b]$ or a commutator $[b, a]$, and k_1 is a product of n such commutators. According to Lemma 13.30, (4),

$$[k_1 t, c] = [t, c]^{k_1} [k_1, c].$$

Both factors are products of finitely many commutators $[a, b, c]$ and their inverses, by the induction hypothesis and the fact that A, B, C are normal subgroups and, thus, are invariant under conjugation. \square

EXERCISE 13.59. Prove the same result for $[H_1, \dots, H_n]$, where all H_i are normal subgroups of G .

LEMMA 13.60 (The Hall identity). *Given a group G and three arbitrary elements x, y, z in G , the following identity holds:*

$$(13.6) \quad [x^{-1}, y, z]^x [z^{-1}, x, y]^z [y^{-1}, z, x]^y = 1.$$

PROOF. The factor $[x^{-1}, y, z]^x$ equals $yx y^{-1} z y x^{-1} y^{-1} x z^{-1} x^{-1}$. The other two factors can be obtained by proper cyclic permutation and a direct calculation shows that all the terms cancel and the product is 1. \square

COROLLARY 13.61. *Assume that A, B, C are normal subgroups in G . Then*

$$(13.7) \quad [A, B, C] \leq [B, C, A][C, A, B].$$

The next proposition shows that the lower central series of G is *graded* with respect to commutators:

PROPOSITION 13.62. *Let $C^k G$ be the k -th group in the lower central series of a group G . Then for every $i, j \geq 1$*

$$(13.8) \quad [C^i G, C^j G] \leq C^{i+j} G.$$

PROOF. We prove by induction on $i \geq 1$ that for every $j \geq 1$, the inclusion (13.8) holds.

For $i = 1$ this follows from the definition of $C^k G$. Assume that the statement is true for i . Consider $j \geq 1$ arbitrary.

$$[C^{i+1} G, C^j G] = [C^i G, G, C^j G] \leq [G, C^j G, C^i G][C^j G, C^i G, G] \leq$$

$$[C^{j+1} G, C^i G][C^{j+i} G, G] = [C^i G, C^{j+1} G][C^{j+i} G, G] \leq C^{j+i+1} G,$$

since $[C^i G, C^{j+1} G] \leq C^{j+i+1} G$ by the induction hypothesis. \square

We now prove that, as for abelian groups, all elements of finite order in a finitely generated nilpotent group form a finite subgroup. We will need the following lemma:

LEMMA 13.63. *Let G be a nilpotent group of class k . For every $x \in G$ the subgroup H generated by x and $C^2 G$ is a normal subgroup, which is nilpotent of class $\leq k - 1$.*

PROOF. By normality of $C^2 G$ in G , the subgroup H can be described as

$$H = \{x^m c \mid m \in \mathbb{Z}, c \in C^2 G\}.$$

For every $g \in G$, and $h \in H$, $h = x^m c$, $ghg^{-1} = x^m [x^{-m}, g] g c g^{-1}$, and, since the last two factors are in $C^2 G$, the whole product is in H . Hence, H is normal in G .

We now prove that $C^2 H \leq C^3 G$, which will imply that H is of class $\leq k - 1$ and, thereby conclude the proof of lemma.

Let h, h' be two elements in H , $h = x^m c_1$, $h' = x^n c_2$ with $c_i \in C^2 G$. Then, according to Lemma 13.30, (3),

$$[h, h'] = [h, x^n c_2] = [h, x^n] [x^n, [h, c_2]] [h, c_2].$$

The last term is in $C^3 G$, hence the middle term is in $C^4 G$.

For $[h, x^n] = [x^m c_1, x^n]$ we apply Lemma 13.30, (4), and obtain

$$[h, h'] = [x^m, [c_1, x^n]] [c_1, x^n].$$

Since the last term is in $C^3 G$ and the first in $C^4 G$, lemma follows. \square

THEOREM 13.64. *Let G be a nilpotent group. The set of all finite order elements forms a characteristic subgroup of G , called the torsion subgroup of G and denoted by $\text{Tor } G$.*

PROOF. We argue by induction on the class of nilpotency k of G . For $k = 1$ the G group is abelian and the assertion is clear. Assume that the statement is true for all nilpotent groups of class $\leq k$, and consider a $(k + 1)$ -step nilpotent group G .

It suffices to prove that for two arbitrary elements a, b of finite order in G , the product ab is likewise of finite order. The subgroup $B = \langle b, C^2 G \rangle$ is nilpotent of class $\leq k$, according to Lemma 13.63. By the induction hypothesis, the set of finite order elements of B is a subgroup $\text{Tor } B \leq B$, which is necessarily characteristic in B . Since B is normal in G it follows that $\text{Tor } B$ is normal in G .

Assume that a is of order m . Then

$$(ab)^m = aba^{-1} a^2 ba^{-2} a^3 b \dots a^{-m+1} a^m ba^{-m},$$

and right-hand side is a product of conjugates of b , hence it is in $\text{Tor } B$. We conclude that $(ab)^m$ is of finite order. \square

PROPOSITION 13.65. *A finitely generated nilpotent torsion group is finite.*

PROOF. We again argue by induction on the nilpotency class n of the group G . For $n = 1$ we apply Exercise 13.10. Assume that the property holds for all nilpotent groups of class at most n and consider G , a finitely generated torsion group that is $(n + 1)$ -step nilpotent. Then $C^2 G$ and $G/C^2 G$ are finite, by the induction hypothesis, whence G is finite as well. \square

COROLLARY 13.66. *Let G be a finitely generated nilpotent group. Then the torsion subgroup $\text{Tor } G$ is finite.*

Brought the statement below here (it was after the proof of Proposition 13.70), and renamed it Proposition -it was called Corollary, though unclear of what.

PROPOSITION 13.67. *If G is nilpotent then $\bar{G} := G/\text{Tor } G$ is torsion-free.*

PROOF. Each element $\bar{x} \in \bar{G}$, is the image of $x = ty \in G$ under the quotient map $\pi : G \rightarrow \bar{G}$, where $t \in \text{Tor } G$. Then $1 = (\bar{x})^k$ would imply that

$$1 = (\bar{x})^k = \pi(y^k),$$

$y^k \in \text{Tor } G$ and, hence, $y \in \text{Tor } G$. It follows that $\bar{x} = 1$. \square

COROLLARY 13.68. *If G is nilpotent then the Hirsch length $h(G)$ equals $cd_{\mathbb{Q}}(G)$, the cohomological dimension of G over \mathbb{Q} .*

PROOF. Let $K := \text{Tor}(G)$, $\Lambda := G/\text{Tor}(G)$ and M be a $\mathbb{Q}G$ -module. We also let M^K denote the submodule of K -invariants in M . Since K is finite, we obtain $H^i(K, M) = 0$ for all $i > 0$. Then the Lyndon-Hochschild-Serre spectral sequence for group cohomology, yields isomorphisms

$$H^i(\Lambda, M^K) \cong H^i(G, M), \quad i > 0,$$

cf. Theorem 2 in [HS53]. It follows that $cd_{\mathbb{Q}}(G) \leq h = h(\Lambda) = cd_{\mathbb{Q}}(\Lambda)$: the vanishing of the cohomology of Λ in degrees $> h$ implies the vanishing of the cohomology of G in the same degrees. To see the converse, consider $M = \mathbb{Q}$, the trivial $\mathbb{Q}G$ -module (and trivial $\mathbb{Q}\Lambda$ -module); in this case, of course, $M = M^K$. Since

$$\mathbb{Q} \cong H^h(\Lambda, \mathbb{Q}),$$

it follows that

$$H^h(\Lambda, \mathbb{Q}) = H^h(G, \mathbb{Q}) \neq 0.$$

Therefore, $cd_{\mathbb{Q}}(G) = h = h(G)$. □

EXERCISE 13.69. Let D_{∞} be the infinite dihedral group.

- (1) Give an example of two elements a, b of finite order in D_{∞} such that their product ab is of infinite order.
- (2) Is D_{∞} a nilpotent group?
- (3) Are any of the finite dihedral groups D_{2n} nilpotent?

Changed the statement below into a proposition, stated it as a 3-fold equivalence, changed the proof accordingly and fixed the many mistakes in it.

PROPOSITION 13.70 (A. I. Mal'cev, [Mal49a]). *Let G be a nilpotent group. The following are equivalent:*

- (a) $Z(G)$ is torsion free;
- (b) Each quotient $Z_{i+1}(G)/Z_i(G)$ is torsion-free;
- (c) G is torsion-free.

PROOF. Clearly (c) \Rightarrow (a).

(a) \Rightarrow (b). We argue by induction on the nilpotency class n of G . The assertion is clear for $n = 1$; assume it holds for all nilpotent groups of class $< n$. We first prove that the group $Z_2(G)/Z_1(G)$ is torsion-free.

We will show that for each non-trivial element $\bar{x} \in Z_2(G)/Z_1(G)$, there exists a homomorphism $\varphi \in \text{Hom}(Z_2(G)/Z_1(G), Z_1(G))$ such that $\varphi(\bar{x}) \neq 1$. Since $Z_1(G)$ is torsion-free this would imply that $Z_2(G)/Z_1(G)$ is torsion-free as well. Let $x \in Z_2(G)$ be the element which projects to $\bar{x} \in Z_2(G)/Z_1(G)$. Thus $x \notin Z_1(G)$, therefore there exists an element $g \in G$ such that $[g, x] \in Z_1(G) \setminus \{1\}$. Define the map $\tilde{\varphi} : Z_2(G) \rightarrow Z_1(G)$ by:

$$\tilde{\varphi}(y) := [y, g],$$

where $g \in G$ is the element above (such that $[g, x] \neq 1$). Clearly, $\tilde{\varphi}(x) \neq 1$; since $Z_1(G)$ is the center of G , the map $\tilde{\varphi}$ descends to a map $\varphi : Z_2(G)/Z_1(G) \rightarrow Z_1(G)$. It follows from Part 3 of Lemma 13.30 that $\tilde{\varphi}$ is a homomorphism. Hence, φ is a homomorphism as well. Since $Z_1(G)$ is torsion-free, it follows that $Z_2(G)/Z_1(G)$ is torsion-free too. Now, we replace G by the group $\bar{G} = G/Z_1(G)$.

Since $Z_2(G)/Z_1(G)$ is torsion-free, the group \bar{G} has torsion-free center. Hence, by the induction hypothesis, $Z_i(\bar{G})/Z_{i-1}(\bar{G})$ is torsion-free for every $i \geq 1$. However,

$$Z_i(\bar{G})/Z_{i-1}(\bar{G}) \cong Z_{i+1}(G)/Z_i(G),$$

for every $i \geq 1$. Thus, every group $Z_{i+1}(G)/Z_i(G)$ is torsion-free, proving (b).

(b) \Rightarrow (c). In view of (b), for each i , $m \neq 0$ and each $x \in Z_i(G) \setminus Z_{i-1}(G)$ we have that $x^m \notin Z_{i-1}(G)$. Thus $x^m \neq 1$. Therefore, G is torsion-free. \square

Moved the remark below from the Polycyclic groups section (it was after Proposition 13.83) to here.

REMARK 13.71. Proposition 13.70 does not imply that for torsion-free nilpotent groups the quotients $C^i G/C^{i+1}G$ are torsion-free. This is, in general, false. Indeed, given an integer $p \geq 2$, consider the following subgroup G of the integer Heisenberg group $H_3(\mathbb{Z})$:

$$G = \left\{ \begin{pmatrix} 1 & k & n \\ 0 & 1 & pm \\ 0 & 0 & 1 \end{pmatrix} ; k, m, n \in \mathbb{Z} \right\}.$$

Since $H_3(\mathbb{Z})$ is poly- C_∞ , so is G . On the other hand, the commutator subgroup in G is:

$$G' = \left\{ \begin{pmatrix} 1 & 0 & pm \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} ; n \in \mathbb{Z} \right\}.$$

The quotient G/G' is isomorphic to $\mathbb{Z}^2 \times \mathbb{Z}_p$.

13.5. Polycyclic groups

DEFINITION 13.72. A group G is *polycyclic* if it admits a subnormal descending series

$$(13.9) \quad G = N_0 \triangleright N_1 \triangleright \dots \triangleright N_n \triangleright N_{n+1} = \{1\}$$

such that N_i/N_{i+1} is cyclic for all $i \geq 0$.

A series as in (13.9) is called a *cyclic series*, and its *length* is the number of non-trivial groups in this sequence, this number is $\leq n + 1$ in (13.9). The *length* $\ell(G)$ of a polycyclic group is the least length of a cyclic series of G .

If, moreover, N_i/N_{i+1} is infinite cyclic for all $i \geq 0$, then the group G is called *poly- C_∞* and the series is called a *C_∞ -series*.

We declare the trivial group to be poly- C_∞ as well.

REMARK 13.73. If G is poly- C_∞ then Corollary 7.24 implies that $N_i \simeq N_{i+1} \rtimes \mathbb{Z}$ for every $i \geq 0$; thus, the group G is obtained from $N_n \simeq \mathbb{Z}$ by *successive semidirect products with \mathbb{Z}* .

For general polycyclic groups G the above is no longer true, for instance, G could be a finite group. However, the above property is *almost true* for G : *Every polycyclic group contains a normal subgroup of finite index which is poly- C_∞* (see Proposition 13.81).

PROPOSITION 13.74. (1) *A polycyclic group has the bounded generation property in the sense of Definition 7.15. More precisely, let G be a group with a cyclic series (13.9) of length n and let t_i be such that $t_i N_{i+1}$ is a*

generator of N_i/N_{i+1} . Then every $g \in G$ can be written as $g = t_1^{k_1} \cdots t_n^{k_n}$, with k_1, \dots, k_n in \mathbb{Z} .

- (2) A polycyclic torsion group is finite.
- (3) Any subgroup of a polycyclic group is polycyclic, and, hence, finitely generated.
- (4) If N is a normal subgroup in a polycyclic group G , then G/N is polycyclic.
- (5) If $N \triangleleft G$ and both N and G/N are polycyclic then G is polycyclic.
- (6) Properties (3) and (5) hold with ‘polycyclic’ replaced by ‘poly- C_∞ ’, but not (4).

PROOF. Part (1) follows by an easy induction on n .

Part (2) follows immediately from (1).

(3). Let H be a subgroup in G . Given a cyclic series for G as above, the intersections $H \cap N_i$ define a cyclic series for H .

(4). The proof is by induction on the length $\ell(G) = n$. For $n = 1$, G is cyclic and any quotient of G is also cyclic.

Assume that the statement is true for all $k \leq n$, and consider a group G with $\ell(G) = n + 1$. Let N_1 be the first term distinct from G in this cyclic series. By the induction hypothesis, $N_1/(N_1 \cap N) \simeq N_1N/N$ is polycyclic. The subgroup N_1N/N is normal in G/N and $(G/N)/(N_1N/N) \simeq G/N_1N$ is cyclic, as it is a quotient of G/N_1 . It follows that G/N is polycyclic.

(5) Consider the cyclic series

$$G/N = Q_0 \geq Q_1 \geq \cdots \geq Q_n = \{\bar{1}\}$$

and

$$N = N_0 \geq N_1 \geq \cdots \geq N_k = \{1\}.$$

Given the quotient map $\pi : G \rightarrow G/N$ and $H_i := \pi^{-1}(Q_i)$, the following is a cyclic series for G :

$$G \geq H_1 \geq \cdots \geq H_n = N = N_0 \geq N_1 \geq \cdots \geq N_k = \{1\}.$$

(6) The proofs of properties (3) and (5) with ‘polycyclic’ replaced by ‘poly- C_∞ ’ are identical. A counter-example for (4) with ‘polycyclic’ replaced by ‘poly- C_∞ ’ is $G = \mathbb{Z}$, $N = 2\mathbb{Z}$. \square

REMARKS 13.75. (1) If G is polycyclic then, in general, the subset $\text{Tor } G \subset G$ of finite order elements in G is neither a subgroup **nor a finite set**.

Consider for instance the infinite dihedral group D_∞ . This group can be realized as the group of isometries of \mathbb{R} generated by the symmetry $s : \mathbb{R} \rightarrow \mathbb{R}$, $s(x) = -x$ and the translation $t : \mathbb{R} \rightarrow \mathbb{R}$, $t(x) = x + 1$, and as noted before (see Section 5.3) $D_\infty = \langle t \rangle \rtimes \langle s \rangle$. Therefore D_∞ is polycyclic by Proposition 13.74, (5), but $\text{Tor } D_\infty$ is the union of a left coset and the trivial subgroup:

$$\text{Tor } G = s \langle t \rangle \cup \{1\}.$$

(2) Every polycyclic group is virtually torsion-free (see Proposition 13.81).

PROPOSITION 13.76. *Every finitely generated nilpotent group is polycyclic.*

PROOF. This may be proved using Proposition 13.74, Part (5), and an induction on the nilpotency class or directly, by constructing a series as in (13.9) as follows: Consider the finite descending series with terms $C^k G$. For every $k \geq 1$, $C^k G / C^{k+1} G$ is finitely generated abelian (see Corollary 13.45). According to the classification of finitely generated abelian groups, there exists a finite subnormal descending series

$$C^k G = A_0 \geq A_1 \geq \cdots \geq A_n \geq A_{n+1} = C^{k+1} G$$

such that every quotient A_i / A_{i+1} is cyclic. By inserting all these finite descending series into the one defined by $C^k G$'s, we obtain a finite subnormal cyclic series for G . \square

THEOREM 13.77 (K. A. Hirsch, [Hir38]). *All polycyclic groups are residually finite.*

We will prove a **slightly** stronger statement, generalizing Corollary 7.109.

THEOREM 13.78. *If a group G is virtually isomorphic to a polycyclic group, then G is residually finite.*

PROOF. Since for a subgroup $G_0 \leq G$ of finite index, G_0 is residually finite if and only if G is, the problem reduces to the following: If $F \triangleleft G$ is a finite subgroup and $G_1 \cong G/F$ is polycyclic, then G is residually finite. The proof is by induction on the cyclic length $\ell(G_1)$ of G_1 . For $\ell(G_1) = 1$ the statement follows from Corollary 7.108.

Assuming that the claim holds for all groups G_1 of cyclic length n , consider the case when $\ell(G_1) = n + 1$. Then G_1 contains a normal subgroup G_2 such that $C = G_1/G_2$ is cyclic. Let H denote the preimage of G_2 under the quotient homomorphism $G \rightarrow G_1$. We, thus, have the short exact sequence

$$1 \rightarrow H \rightarrow G \rightarrow C \rightarrow 1.$$

By the induction hypothesis, the group H is residually finite. The subgroup G_2 of G_1 is finitely generated according to Proposition 13.74. Hence, H is finitely generated as well. Corollary 7.108 implies that G is residually finite. This concludes the proof of the theorem. \square

An edifying example of a polycyclic group is the following.

PROPOSITION 13.79. *Let $m, n \geq 1$ be two integers, and let $\varphi : \mathbb{Z}^n \rightarrow \text{Aut}(\mathbb{Z}^m)$ be a homomorphism.*

The semidirect product $G = \mathbb{Z}^m \rtimes_{\varphi} \mathbb{Z}^n$ is a poly- C_{∞} group.

PROOF. The quotient G/\mathbb{Z}^m is isomorphic to \mathbb{Z}^n . Therefore by Proposition 13.74, (6), the group G is poly- C_{∞} . \square

EXERCISE 13.80. Let $\mathcal{T}_n(\mathbb{K})$ be the group of invertible upper-triangular $n \times n$ matrices with entries in a field \mathbb{K} .

- (1) Prove that $\mathcal{T}_n(\mathbb{K})$ is a semidirect product of its nilpotent subgroup $\mathcal{U}_n(\mathbb{K})$ introduced in Exercise 13.38, and the subgroup of diagonal matrices.

- (2) Prove that, if \mathbb{K} has zero characteristic, the subgroup of $\mathcal{T}_n(\mathbb{K})$ generated by $I + E_{12}$ and by the diagonal matrix with $(-1, 1, \dots, 1)$ on the diagonal is isomorphic to the infinite dihedral group D_∞ . Deduce that $\mathcal{T}_n(\mathbb{K})$ is not nilpotent.

PROPOSITION 13.81. *A polycyclic group G contains a normal subgroup of finite index which is poly- C_∞ .*

PROOF. We argue by induction on the length $\ell(G) = n$. For $n = 1$ the group G is cyclic and the statement obviously true. Assume that the assertion is true for n and consider a polycyclic group G having a cyclic series (13.9).

The induction hypothesis implies that N_1 contains a normal subgroup S of finite index which is poly- C_∞ . Lemma 5.10 shows that S has a finite-index subgroup S_1 which is normal in G . Proposition 13.74, Part (6), implies that S_1 is poly- C_∞ as well.

If G/N_1 is finite then S_1 has finite index in G .

Assume that G/N_1 is infinite cyclic. Then the group $K = G/S_1$ contains the finite normal subgroup $F = N_1/S_1$ such that K/F is isomorphic to \mathbb{Z} . Corollary 7.24 implies that K is a semidirect product of F and an infinite cyclic subgroup $\langle x \rangle$. The conjugation by x defines an automorphism of F and since $\text{Aut}(F)$ is finite, there exists r such that the conjugation by x^r is the identity on F . Hence $F \langle x^r \rangle$ is a finite-index subgroup in K and it is a direct product of F and $\langle x^r \rangle$. We conclude that $\langle x^r \rangle$ is a finite index normal subgroup of K . We have that $\langle x^r \rangle = G_1/S_1$, where G_1 is a finite index normal subgroup in G , and G_1 is poly- C_∞ since S_1 is poly- C_∞ . \square

COROLLARY 13.82. (a) *A poly- C_∞ group is torsion-free.*

(b) *A polycyclic group is virtually torsion-free.*

PROOF. In view of Proposition 13.81, it suffices to prove (a). Consider a poly- C_∞ group G . We argue by induction on the cyclic length $\ell(G) = n$. For $n = 1$, the group G is infinite cyclic and the statement obviously holds. Assume that the statement is true for all groups of cyclic length at most n and consider a group G with $\ell(G) = n + 1$ and the cyclic series (13.9). Let g be an element of finite order in G . Then its image in the infinite cyclic quotient G/N_1 is the identity, hence $g \in N_1$. The induction hypothesis implies that $g = 1$. \square

PROPOSITION 13.83. *Let G be a finitely generated nilpotent group. The following are equivalent:*

- (1) *G is poly- C_∞ ;*
- (2) *G is torsion-free;*
- (3) *the center of G is torsion-free.*

PROOF. Implication (1) \Rightarrow (2) is Corollary 13.82, (a), while the implication (2) \Rightarrow (3) is obvious. The implication (3) \Rightarrow (1) follows from Proposition 13.70. \square

PROPOSITION 13.84. *Every polycyclic group is finitely presented.*

PROOF. The proof is an easy induction on the minimal length of a cyclic series, combined with Proposition 7.30. \square

One parameter measuring the complexity of the “poly- C_∞ part” of any polycyclic group is the *Hirsch number* (generalizing the Hirsch length for nilpotent groups), defined as follows:

PROPOSITION 13.85. *The number of infinite factors in a cyclic series of a polycyclic group G is the same for all series. This number is called the Hirsch number (or Hirsch length) of G .*

PROOF. The proof will follow from the following observation on cyclic series:

LEMMA 13.86. *Any refinement of a cyclic series is also cyclic. Moreover, the number of quotients isomorphic to \mathbb{Z} is the same for both series.*

PROOF. Consider a cyclic series

$$H_0 = G \geq H_1 \geq \dots \geq H_n = \{1\}.$$

A refinement of this series is composed of the following sub-series

$$H_i = R_k \geq R_{k+1} \geq \dots \geq R_{k+m} = H_{i+1}.$$

Each quotient R_j/R_{j+1} embeds naturally as a subgroup in H_i/R_{j+1} , and the latter is a quotient of the cyclic group H_i/H_{i+1} ; hence all quotients are cyclic. If H_i/H_{i+1} is finite then all quotients R_j/R_{j+1} are finite.

Assume now that $H_i/H_{i+1} \simeq \mathbb{Z}$. We prove by induction on $m \geq 1$ that exactly one among the quotients R_j/R_{j+1} is isomorphic to \mathbb{Z} , and the other quotients are finite. For $m = 1$ the statement is clear. Assume that it is true for m and consider the case of $m + 1$.

If H_i/R_{k+m} is finite then all R_j/R_{j+1} with $j \leq k + m - 1$ are finite. On the other, under this assumption, R_{k+m}/R_{k+m+1} cannot be finite, otherwise H_i/H_{i+1} would be finite.

Assume that $H_i/R_{k+m} \simeq \mathbb{Z}$. The induction hypothesis implies that exactly one quotient R_j/R_{j+1} with $j \leq k + m - 1$ is isomorphic to \mathbb{Z} and the others are finite. The quotient R_{k+m}/R_{k+m+1} is a subgroup of $H_i/R_{k+m} \simeq \mathbb{Z}$ such that the quotient by this subgroup is also isomorphic to \mathbb{Z} . This can only happen when R_{k+m}/R_{k+m+1} is trivial. \square

Proposition 13.85 now follows from Lemmas 13.86 and 5.6. \square

EXERCISE 13.87. Show that for each finitely generated nilpotent group the Hirsch number equals the *Hirsch length* $h(G)$, defined earlier.

In view of this exercise, the Hirsch number for a polycyclic group G will be again denoted $h(G)$.

A natural question to ask is the following.

QUESTION 13.88. Since poly- C_∞ groups are constructed by successive semidirect products with \mathbb{Z} , is there a way to detect during this construction whether the group is nilpotent or not?

The answer to this question will be given in Section 14.3 and it has some interesting relation to the growth of groups.

13.6. Solvable groups: Definition and basic properties

Recall that G' denotes the *derived subgroup* $[G, G]$ of the group G . Given a group G , we define its *iterated commutator subgroups* $G^{(k)}$ inductively by:

$$G^{(0)} = G, G^{(1)} = G', \dots, G^{(k+1)} = \left(G^{(k)}\right)', \dots$$

The descending series

$$G \supseteq G' \supseteq \dots \supseteq G^{(k)} \supseteq G^{(k+1)} \supseteq \dots$$

is called the *derived series* of the group G .

Note that all subgroups $G^{(k)}$ are *characteristic* in G .

DEFINITION 13.89. A group G is *solvable* if there exists k such that $G^{(k)} = \{1\}$. The minimal k such that $G^{(k)} = \{1\}$ is called the *derived length* of G and the group G itself is called *k -step solvable*. A solvable group of derived length at most two is called *metabelian*.

We will use the notation $\ell_{\text{der}}(G)$ for the derived length.

In particular, every solvable group G of derived length **at most** k satisfies the law:

$$(13.10) \quad \llbracket x_1, \dots, x_{2^k} \rrbracket = 1, \forall x_1, \dots, x_{2^k} \in G.$$

Here and in what follows,

$$(13.11) \quad \llbracket x_1, \dots, x_{2^k} \rrbracket := [\llbracket x_1, \dots, x_{2^{k-1}} \rrbracket, \llbracket x_{2^{k-1}+1}, \dots, x_{2^k} \rrbracket]$$

and $\llbracket x_1, x_2 \rrbracket = [x_1, x_2]$.

EXERCISE 13.90. 1. Find the values $n \in \mathbb{N}$ for which the symmetric group S_n is solvable. Show that the groups S_3, S_4 are nilpotent.

2. Show that if a group G satisfies the law (13.10), then it is solvable of **derived** length $\leq k$.

PROPOSITION 13.91. (1) *If N is a normal subgroup in G and both N and G/N are solvable, then G is solvable. If the derived lengths of G/N and N are at most d, d' respectively, then the derived length of G is at most $d + d'$. In other words, the derived length is subadditive:*

$$\ell_{\text{der}}(G) \leq \ell_{\text{der}}(N) + \ell_{\text{der}}(G/N).$$

(2) *Every subgroup H of a solvable group G is solvable and*

$$\ell_{\text{der}}(H) \leq \ell_{\text{der}}(G).$$

(3) *If G is solvable and $N \triangleleft G$, then G/N is solvable and*

$$\ell_{\text{der}}(G/N) \leq \ell_{\text{der}}(G).$$

Note that the statement (1) is not true when ‘solvable’ is replaced by ‘nilpotent’, consider, for instance, the infinite dihedral group D_∞ .

PROOF. (1) We are assuming that G/N is solvable of derived length d and N is solvable of derived length d' . Since $(G/N)^{(d)} = \{1\}$ it follows that $G^{(d)} \leq N$. Then, as $G^{(d+i)} \leq N^{(i)}$, we obtain $G^{(d+d')} = \{1\}$.

(2) Note that for every subgroup H of a group G , $H' \leq G'$. Thus, by induction,

$$H^{(i)} \leq G^{(i)}.$$

If G is solvable of derived length k then $G^{(k)} = \{1\}$; thus $H^{(k)} = \{1\}$ as well and, hence, H is also solvable.

(3) Consider the quotient map $\pi : G \rightarrow G/N$. It is immediate that $\pi(G^{(i)}) = (G/N)^{(i)}$, in particular if G is solvable then G/N is solvable. \square

For the next exercise, we will need the following definition: A finite sequence of vector subspaces

$$V_0 \subsetneq V_1 \subsetneq \cdots \subsetneq V_k$$

in a vector space V is called a *flag* in V . If the number of the subspaces in such a sequence is maximal possible (equal $\dim(V) + 1$), the flag is called *full* or *complete*. In other words, $\dim(V_i) = i$ for all members of this sequence.

EXERCISE 13.92. (1) Prove that the subgroup $\mathcal{T}_n(\mathbb{K})$ of upper-triangular matrices in $GL(n, \mathbb{K})$, where \mathbb{K} is a field, is solvable. [Hint: you may use Exercise 13.38.]

(2) Use Part (1) to show that for a finite-dimensional vector space V , the subgroup G of $GL(V)$ consisting of elements g preserving a complete flag in V (i.e. $gV_i = V_i$, for every $g \in G$ and every i) is solvable.

(3) Let V be a \mathbb{K} -vector space of dimension n , and let

$$V_0 = 0 \subset V_1 \subset \cdots \subset V_{k-1} \subset V_k = V$$

be a flag, not necessarily complete. Let G be a subgroup of $GL(V)$ preserving this flag. For every $i \in \{1, 2, \dots, k-1\}$ let ρ_i be the projection $G \rightarrow GL(V_{i+1}/V_i)$. Prove that if every $\rho_i(G)$ is solvable, then G is also solvable.

EXERCISE 13.93. 1. Let \mathbb{F}_k denote the field with k elements. Use the 1-dimensional vector subspaces in \mathbb{F}_k^2 to construct a homomorphism $GL(2, \mathbb{F}_k) \rightarrow S_n$ for an appropriate n .

2. Prove that $GL(2, \mathbb{F}_2)$ and $GL(2, \mathbb{F}_3)$ are solvable.

THEOREM 13.94 (Direct limits of virtually solvable groups). *Suppose that we have a direct system $G_i, i \in I$, of virtually solvable groups satisfying the following:*

1. *The derived length of each solvable subgroup in G_i is at most d for all i .*
2. *Each G_i contains a normal solvable subgroup H_i of index $\leq c$.*

Then the direct limit G of this system is again virtually solvable and contains a normal solvable subgroup H of index $\leq c$ and derived length at most d .

PROOF. We start the proof with several simple observations. If a group M is virtually solvable and $N_1, N_2 \triangleleft M$ are normal solvable subgroups of finite index, then the subgroup $N \leq F$ generated by N_1, N_2 is again solvable and normal in M . Therefore, without loss of generality, we may assume that each H_i is maximal among all normal solvable subgroups of finite index in G_i . By the hypothesis, the derived length of each H_i does not exceed d . Clearly, we retain the property that $|G_i : H_i| \leq c$ for all i .

Suppose now that $m \in I$ is such that the index $n = |G_m : H_m| \leq c$ is maximal among all indices $|G_i : H_i|, i \in I$. Define $J = \{j \in I : m \leq j\}$. Then for each homomorphism $f_{mi} : G_m \rightarrow G_i$ of the direct system, we have:

- (1) $H_{mi} := f_{mi}^{-1}(H_i)$ is a normal solvable subgroup of finite index in G_m .
Hence, $H_{im} \leq H_m$.
- (2)

$$n = |G_m : H_m| \leq |G_m : H_{im}| = |\langle H_i, f_{mi}(G_m) \rangle : H_i| \leq |G_i : H_i| \leq n.$$

In particular, $|G_i : H_i| = n$ and $f_{mi}(H_m) \leq H_i$. Applying this to any pair $i, j \in J, i \leq j$, we obtain:

$$f_{ij}(H_i) \leq H_j.$$

Hence, for $i, j \in J$, the restrictions of the homomorphisms f_{ij} to H_i define a direct system of solvable groups of derived length $\leq d$. Taking the direct limit $\lim_{j \in J} H_j$, we obtain a solvable subgroup $H \leq G$ (see Exercise 1.29). The reader will verify (using Exercises 1.30 and 13.90) that:

- (1) The subgroup H is normal in G .
- (2) $|G : H| = n$.
- (3) The derived length of H is at most d .

This concludes the proof. □

13.7. Free solvable groups and Magnus embedding

As in the case of nilpotent groups, there exist universal objects in the class of solvable groups that we now describe.

DEFINITION 13.95. Given two integers $k, m \geq 1$, the *free solvable group of derived length k with m generators* is the quotient of the free group F_m by the normal subgroup $F_m^{(k)}$.

When $k = 2$ we call the corresponding group *free metabelian group with m generators*.

NOTATION 13.96. In what follows we use the notation $S_{m,k}$ for the free solvable group of derived length k and with m generators. Note that $S_{m,1}$ is \mathbb{Z}^m .

PROPOSITION 13.97 (Universal property of free solvable groups). *Every solvable group with m generators and of derived length k , is a quotient of $S_{m,k}$.*

PROOF. Let G be a solvable group of derived length k and let X be a generating set of G of cardinality m . The map defined in Proposition 7.21 contains $F(X)^{(k)}$ in its kernel, therefore it defines an epimorphism from the free solvable group $S_{m,k}$ to G . □

Our next goal is to define the *Magnus embedding* of the free solvable group $S_{r,k+1}$ into the wreath product $\mathbb{Z}^r \wr S_{r,k}$. Since $\mathbb{Z}^r \wr S_{r,k}$ is a semidirect product, Remark 5.125, (2), implies that in order to define a homomorphism

$$S_{r,k+1} \rightarrow \mathbb{Z}^r \wr S_{r,k}$$

one has to specify a homomorphism $\pi : S_{r,k+1} \rightarrow S_{r,k}$ and a derivation

$$d \in \text{Der}(S_{r,k+1}, \bigoplus_{S_{r,k}} \mathbb{Z}^r).$$

Here we will use the following action of $S_{r,k+1}$ on $\bigoplus_{S_{r,k}} \mathbb{Z}^r$: We compose π with the action of $S_{r,k}$ on itself *via* left multiplication.

To simplify the notation, we let $F = F_r$ denote the free group on r generators x_1, \dots, x_r . First, since $F/F^{(m)} = S_{r,m}$ for every m , and $F^{(k+1)} \leq F^{(k)}$, we have a natural quotient homomorphism

$$\pi : S_{r,k+1} \rightarrow S_{r,k}.$$

We now proceed to construct the derivation d . We will use definitions and results of Section 5.9.4. Note that $\bigoplus_{S_{r,k}} \mathbb{Z}^r$ is isomorphic (as a free abelian group) to

$$M_1 \oplus \dots \oplus M_r$$

where for every i , $M_i = M = \mathbb{Z}S_{r,k}$, the group algebra of $S_{r,k}$. Since $S_{r,k}$ is the quotient of $F = F_r$, every derivation $\partial \in \text{Der}(\mathbb{Z}F, \mathbb{Z}F)$ projects to a derivation (denoted $\widehat{\partial}$) in $\text{Der}(\mathbb{Z}F, \mathbb{Z}S_{r,k})$. Thus, derivations $\partial_i \in \text{Der}(\mathbb{Z}F, \mathbb{Z}F)$ introduced in Section 5.9.4, projects to derivations $\widehat{\partial}_i \in \text{Der}(\mathbb{Z}F, M)$. Furthermore, every derivation $\widehat{\partial}_i \in \text{Der}(\mathbb{Z}F, M)$ extends to a derivation $d_i : \mathbb{Z}F \rightarrow \bigoplus_{S_{r,k}} \mathbb{Z}^r$ by

$$d_i : w \mapsto (0, \dots, \widehat{\partial}_i(w), \dots, 0)$$

where we place $\widehat{\partial}_i(w)$ in the i -th slot. Since a sum of derivations is again a derivation, we obtain a derivation

$$d = (\widehat{\partial}_1, \dots, \widehat{\partial}_r) = d_1 + \dots + d_r \in \text{Der}(\mathbb{Z}F, \bigoplus_{S_{r,k}} \mathbb{Z}^r).$$

For simplicity, in what follows, we denote $F^{(k)}$ by N and, accordingly, $F^{(k+1)}$ by N' . Thus, $S_{r,k} = F/N$ and $S_{r,k+1} = F/N'$.

LEMMA 13.98. *The derivation d projects to a derivation*

$$\bar{d} \in \text{Der}(\mathbb{Z}S_{r,k+1}, \bigoplus_{S_{r,k}} \mathbb{Z}^r).$$

PROOF. Let us check that N' is in the kernel of d . Indeed, given a commutator $[x, y]$ with x, y in N , property (P_3) in Exercise 5.121 implies that (by computing in $\mathbb{Z}F$)

$$\partial_i[x, y] = (1 - yx^{-1})\partial_i x + x(1 - yx^{-1}y^{-1})\partial_i y.$$

Since both $x, y \in N$ project to 1 in $S_{r,k}$, they act trivially on $M = \mathbb{Z}S_{r,k}$, it follows that

$$(1 - yx^{-1}) \cdot \xi = 0 \text{ and } x(1 - yx^{-1}y^{-1}) \cdot \eta = 0, \quad \forall \xi, \eta \in M.$$

Hence, $d_i([x, y]) = 0$ for every i and, thus, $d([x, y]) = 0$. Therefore, $d(N') = 0$ since the group N' is generated by commutators $[x, y], x, y \in N$. For arbitrary $g \in F, h \in N'$, we have

$$d(gn) = d(g) + g \cdot d(n) = d(g).$$

Thus, the derivation d projects to a derivation $\bar{d} \in \text{Der}((\mathbb{Z}S_{r,k+1}, \bigoplus_{S_{r,k}} \mathbb{Z}^r)$,

$$\bar{d}(gN') = d(g). \quad \square$$

Thus, according to Remark 5.125, the pair (d, π) determines a homomorphism

$$\mathfrak{M} : S_{r,k+1} \rightarrow \mathbb{Z}^r \wr S_{r,k}.$$

THEOREM 13.99 (W. Magnus [Mag39]). *The homomorphism \mathfrak{M} is injective; \mathfrak{M} is called the Magnus embedding.*

We refer to [Fox53, Section (4.9)] for the proof of injectivity of \mathfrak{M} . Remarkably, the Magnus embedding also has nice geometric features. The following theorem was proven independently by A. Sale [Sal12, Sal15] and S. Vassileva [Vas12]:

THEOREM 13.100 (A. Sale, S. Vassileva). *The Magnus embedding is a quasi-isometric embedding.*

Clearly, the Magnus embedding is a useful tool for studying free solvable groups by induction on the derived length.

13.8. Solvable versus polycyclic

PROPOSITION 13.101. *Every polycyclic group G is solvable.*

PROOF. This follows immediately by an induction argument on the cyclic length of G and Part (1) of Proposition 13.91. \square

DEFINITION 13.102. A group is said to be *noetherian*, or *to satisfy the maximal condition* if for every increasing sequence of subgroups

$$(13.12) \quad H_1 \leq H_2 \leq \dots \leq H_n \leq \dots$$

there exists N such that $H_n = H_N$ for every $n \geq N$.

PROPOSITION 13.103. *A group G is noetherian if and only if every subgroup of G is finitely generated.*

PROOF. Assume that G is a Noetherian group, and let $H \leq G$ be a subgroup which is not finitely generated. Pick $h_1 \in H \setminus \{1\}$ and let $H_1 = \langle h_1 \rangle$. Inductively, assume that

$$H_1 < H_2 < \dots < H_n$$

is a strictly increasing sequence of finitely generated subgroups of H , pick $h_{n+1} \in H \setminus H_n$, and set $H_{n+1} = \langle H_n, h_{n+1} \rangle$. We thus have a strictly increasing infinite sequence of subgroups of G , contradicting the assumption that G is Noetherian.

Conversely, assume that all subgroups of G are finitely generated, and consider an increasing sequence of subgroups as in (13.12). Then $H = \bigcup_{n \geq 1} H_n$ is a subgroup, hence generated by a finite set S . There exists N such that $S \subseteq H_N$, hence $H_N = H = H_n$ for every $n \geq N$. \square

PROPOSITION 13.104. *A solvable group is polycyclic if and only if it is noetherian.*

PROOF. The ‘only if’ part follows immediately from Parts (1) and (3) of Proposition 13.74. Let G be a noetherian solvable group. We prove by induction on the derived length k that G is polycyclic.

For $k = 1$ the group is abelian, and since, by hypothesis, G is finitely generated, it is polycyclic.

Assume that the statement is true for k and consider a solvable group G of derived length $k + 1$. The commutator subgroup $G' \leq G$ is also Noetherian and solvable of derived length k . Hence, by the induction hypothesis, G' is polycyclic. The abelianization $G_{ab} = G/G'$ is finitely generated (because G is, by hypothesis), hence it is polycyclic. It follows that G is polycyclic by Proposition 13.74 (5). \square

By Proposition 13.101 every nilpotent group is solvable. A natural question to ask is to find a relationship between nilpotency class and derived length.

PROPOSITION 13.105. (1) For every group G and every $i \geq 0$,

$$G^{(i)} \leq C^{2^i} G.$$

(2) If G is a k -step nilpotent group then its derived length is at most

$$[\log_2 k] + 1.$$

PROOF. (1) The statement is obviously true for $i = 0$. Assume that it is true for i . Then

$$G^{(i+1)} = [G^{(i)}, G^{(i)}] \leq [C^{2^i} G, C^{2^i} G] \leq C^{2^{i+1}} G.$$

In the last inclusion we applied Proposition 13.62.

(2) follows immediately from (1). \square

REMARK 13.106. The derived length can be much smaller than the nilpotency class: the dihedral subgroup D_{2n} with $n = 2^k$ is k -step nilpotent and metabelian.

EXERCISE 13.107. If G_1 is noetherian and G_2 is virtually isomorphic to G_1 , then G_2 is also noetherian.

REMARK 13.108. There are noetherian groups that are not virtually polycyclic, e.g. *Tarski monsters*: finitely generated groups such that every proper subgroup is cyclic, see [OI'91a].

An instructive example of solvable group is the *lamplighter group*. This group is the wreath product $G = \mathbb{Z}_2 \wr \mathbb{Z}$ in the sense of Definition 5.32.

EXERCISE 13.109. Prove that if K, H are solvable groups then $K \wr H$ is solvable. In particular, the lamplighter group G is solvable (even metabelian).

In view of Lemma 7.11, since wreath products of finitely generated groups are finitely generated as well, the lamplighter group is finitely generated. On the other hand:

- (1) Not all subgroups in the lamplighter group G are finitely generated: the subgroup $\bigoplus_{n \in \mathbb{Z}} \mathbb{Z}_2$ of G is not finitely generated.
- (2) The lamplighter group G is not virtually torsion-free: For any finite-index subgroup $H \leq G$, $H \cap \bigoplus_{n \in \mathbb{Z}} \mathbb{Z}_2$ has finite index in $\bigoplus_{n \in \mathbb{Z}} \mathbb{Z}_2$; in particular this intersection is infinite and contains elements of order 2.

Both (1) and (2) imply that the lamplighter group is *not polycyclic*.

(3) The commutator subgroup G' of the lamplighter group G coincides with the following subgroup of $\bigoplus_{n \in \mathbb{Z}} \mathbb{Z}_2$:

$$(13.13) \quad C = \{f : \mathbb{Z} \rightarrow \mathbb{Z}_2 \mid \text{Supp}(f) \text{ has even cardinality}\},$$

where $\text{Supp}(f) = \{n \in \mathbb{Z} \mid f(n) = 1\}$.

[NB. The notation here is additive, the identity element is 0.]

In particular, G' is not finitely generated and the group G is metabelian (since G' abelian).

We prove (3). First of all, C is clearly a subgroup. Note also that

$$(f, m)^{-1} = (-\varphi(-m)f, -m),$$

where φ is the action of \mathbb{Z} on the space of functions $f : \mathbb{Z} \rightarrow \mathbb{Z}_2$ via *shift*: For $m \in \mathbb{Z}$,

$$\varphi(m) : f(x) \mapsto f(x + m).$$

If we think of functions f as biinfinite sequences, then $\varphi(m)$ acts on a sequence via shifting all the indices by m . A straightforward calculation gives

$$[(f, m), (g, n)] = (f - g - \varphi(n)f + \varphi(m)g, 0).$$

Now, observe that either $\text{Supp}(f)$ and $\text{Supp}(\varphi(n)f)$ are disjoint, in which case $\text{Supp}(f - \varphi(n)f)$ has cardinality twice the cardinality of $\text{Supp} f$, or they overlap on a set of cardinality k ; in the latter case, $\text{Supp}(f - \varphi(n)f)$ has cardinality twice the cardinality of $\text{Supp} f$ minus $2k$. The same holds for $\text{Supp}(-g + \varphi(m)g)$. Since C is a subgroup,

$$(f - g - \varphi(n)f + \varphi(m)g) = (f - \varphi(n)f - (g - \varphi(m)g)) \in C.$$

This shows that $G' \leq C$.

Consider the opposite inclusion. The subgroup C is generated by functions $f : \mathbb{Z} \rightarrow \mathbb{Z}_2$, $\text{Supp} f = \{a, b\}$, where a, b are distinct integers; thus, it suffices to show that $(f, 0) \in G'$. Let $\delta_a : \mathbb{Z} \rightarrow \mathbb{Z}_2$, $\text{Supp} \delta_a = \{a\}$. Then

$$[(\delta_a, 0), (0, b - a)] = (\delta_a - \varphi(b - a)\delta_a, 0) = (f, 0)$$

which implies that $(f, 0) \in G'$.

We conclude this section by noting that, unlike polycyclic groups, solvable groups may not be finitely presented. An example of such a group is the wreath product $\mathbb{Z} \wr \mathbb{Z}$ [Bie79]. We refer to the same paper for a survey on finitely presented solvable groups. Nevertheless, a solvable group may be finitely presented without being polycyclic; for instance the Baumslag–Solitar group

$$G = BS(1, p) = \langle a, b \mid aba^{-1} = b^p \rangle$$

is metabelian but not polycyclic (for $|p| \geq 2$). The derived subgroup G' of G is isomorphic to the additive group of p -adic rational numbers, i.e. rational numbers whose denominators are powers of p . In particular, G' is not finitely generated. Hence, in view of Proposition 13.74, G is not polycyclic.

EXERCISE 13.110. Show that the group $G = BS(1, p)$ is metabelian.

Geometric aspects of solvable groups

In this chapter we discuss several geometric aspects of solvable groups:

- Distortion of subgroups in nilpotent groups.
- Growth of solvable groups: We will compute growth rates of nilpotent groups and prove Milnor–Wolf theorem that a solvable group has polynomial growth if and only if it is virtually nilpotent.
- Erschler’s examples establishing failure of quasiisometry invariance of the class of (virtually) solvable groups.
- Theorems of Zassenhaus and Jordan dealing with discrete subgroups of Lie groups. Jordan’s theorem shows that finite subgroups of $GL(n, \mathbb{R})$ contain abelian subgroups of *uniformly bounded index*.

Many of these results will play important role in the proof of Gromov’s theorem on groups of polynomial growth.

14.1. Wolf’s Theorem for semidirect products $\mathbb{Z}^n \rtimes \mathbb{Z}$

In this section we explain how to provide an affirmative answer to Question 8.86 in the case of semidirect products $\mathbb{Z}^n \rtimes \mathbb{Z}$. This easy example helps to understand the general case of polycyclic groups and the general Wolf’s Theorem.

Note that the semidirect product is defined by a homomorphism $\varphi : \mathbb{Z} \rightarrow \text{Aut}(\mathbb{Z}^n) = GL(n, \mathbb{Z})$, and the latter is determined by $\theta = \varphi(1)$, which is represented by a matrix $M \in GL(n, \mathbb{Z})$. Therefore the same semidirect product is also denoted $\mathbb{Z}^n \rtimes_{\theta} \mathbb{Z} = \mathbb{Z}^n \rtimes_M \mathbb{Z}$.

PROPOSITION 14.1. *A semidirect product $G = \mathbb{Z}^n \rtimes_M \mathbb{Z}$ is*

- (1) *either virtually nilpotent (when M has all eigenvalues of absolute value 1);*
- (2) *or of exponential growth (when M has at least one eigenvalue of absolute value $\neq 1$).*

REMARKS 14.2. (1) The group $G = \mathbb{Z}^n \rtimes_M \mathbb{Z}$ is nilpotent if M has all eigenvalues equal to 1 (see Case (1) of the proof of the proposition).

- (2) The same is not in general true if M has all eigenvalues of absolute value 1. The group $G = \mathbb{Z} \rtimes_M \mathbb{Z}$ with $M = (-1)$ is a counter-example: It admits a quotient which is the infinite dihedral group and the latter is not nilpotent. In this example, the group $G = \mathbb{Z} \rtimes_M \mathbb{Z}$ is polycyclic, virtually nilpotent but not nilpotent. In particular, the statement (1) in Proposition 14.1 cannot be improved to ‘ $G = \mathbb{Z}^n \rtimes_M \mathbb{Z}$ is nilpotent’.

PROOF. Note that $\mathbb{Z}^n \rtimes_{\theta N} \mathbb{Z}$ is a subgroup of finite index in $G = \mathbb{Z}^n \rtimes_{\theta} \mathbb{Z}$ (corresponding to the replacement of the second factor \mathbb{Z} by $N\mathbb{Z}$). Thus, we may

replace M by some power of M , and replace G with a finite-index subgroup. We will retain the notation G and M for the finite-index subgroup and the power of M . **Then** the matrix $M \in GL(n, \mathbb{Z})$ will have no non-trivial roots of unity as eigenvalues. In view of Lemma 13.28, this means that for every eigenvalue $\lambda \neq 1$ of M , $|\lambda| \neq 1$.

We have two cases to consider.

(1) We prove the statement by induction on n . For $n = 0$ there is nothing to prove; we assume, therefore, that the statement holds for $n - 1$. The matrix M has only eigenvalues equal 1. Lemma 13.27 then implies that there exists a finite series

$$\{1\} = H_n \leq H_{n-1} \leq \dots \leq H_1 \leq A = H_0 = \mathbb{Z}^n$$

such that $H_i \simeq \mathbb{Z}^{n-i}$, each quotient H_i/H_{i+1} is cyclic, the automorphism θ preserves each H_i and induces the identity automorphism on H_i/H_{i+1} . Thus, θ acts via the identity on H_{n-1} . In particular, the subgroup H_{n-1} is central in G ; the automorphism θ projects to an automorphism $\bar{\theta} : \bar{A} \rightarrow \bar{A}$, $\bar{A} = A/H_{n-1}$. The automorphism $\bar{\theta}$ preserves the central series

$$\{1\} = \bar{H}_{n-1} \leq \dots \leq \bar{H}_1 \cong \mathbb{Z}^{n-1},$$

(where $\bar{H}_i = H_i/H_{n-1}$) and induces trivial automorphism of each quotient

$$\bar{H}_i/\bar{H}_{i+1} \cong H_i/H_{i+1}.$$

By the induction hypothesis, the group

$$\bar{G} = \bar{A} \rtimes_{\bar{\theta}} \mathbb{Z} \cong G/H_{n-1},$$

is nilpotent. Since central coextensions of nilpotent groups are again nilpotent (Exercise 13.51), we conclude that the group G is nilpotent as well.

(2) Assume that M has an eigenvalue with absolute value strictly greater than 1. After replacing θ with its power θ^N if necessary, we may assume that the matrix M has an eigenvalue with absolute value at least 2.

Lemma 13.29 applied to M implies that there exists an element $v \in \mathbb{Z}^n$ such that distinct elements $s = (s_k) \in \bigoplus_{k \geq 0} \mathbb{Z}_2$ define distinct vectors

$$s_0 v + s_1 M v + \dots + s_n M^n v + \dots$$

in \mathbb{Z}^n . With the multiplicative notation for the binary operation in G , the above vectors correspond to distinct elements

$$g_s = v^{s_0} (t v t^{-1})^{s_1} \dots (t^k v t^{-k})^{s_k} \dots \in G.$$

Now, consider the set Σ_K of sequences $s = (s_k)$ for which $s_k = 0, \forall k \geq K + 1$. Then the map

$$\Sigma_K \rightarrow G, \quad s \mapsto g_s$$

is injective and its image consists of 2^{K+1} distinct elements g_s . Assume that the generating set of G contains the elements t and v . With respect to this generating set, the word-length $|g_s|$ is at most $3K + 1$ for every $s \in \Sigma_K$. Thus, for every K we obtain 2^{K+1} distinct elements of G of length at most $3K + 1$, whence G has exponential growth. \square

REMARK 14.3. What remains to be proven is that the two cases in Proposition 14.1 are mutually exclusive, i.e. that a nilpotent group cannot have exponential growth. We shall prove in Section 14.2 that nilpotent (hence virtually nilpotent) groups have in fact polynomial growth. In the next section we compute growth for the integer Heisenberg group $H_3(\mathbb{Z})$.

In order to analyze growth functions of solvable groups, we first have to discuss *distortion* (see Section 8.9) of subgroups of solvable groups. This will be done in sections 14.1.2 and 14.1.3.

14.1.1. Geometry of $H_3(\mathbb{Z})$. In this section we discuss in detail the geometric concepts introduced so far in the case of the *integer Heisenberg group*

$$G = H_3(\mathbb{Z}) = \left\{ U_{klm} = \begin{pmatrix} 1 & k & m \\ 0 & 1 & l \\ 0 & 0 & 1 \end{pmatrix} ; k, l, m \in \mathbb{Z} \right\}.$$

When convenient, we will also use the notation $U_{k,l,m}$ instead of U_{klm} .

EXERCISE 14.4. (1) Show that the elements $x = U_{100}$, $y = U_{010}$, $z = U_{001}$ generate G and satisfy the relation

$$[x, y] = z.$$

(2) Prove that $U_{klm} = x^k y^l z^{m-kl}$ for every $k, l, m \in \mathbb{Z}$. This in particular shows that every element of G can be written as $x^k y^l z^m$ with $k, l, m \in \mathbb{Z}$, and that this decomposition is unique for every element (since it is entirely determined by its matrix entries).

(3) Prove that $[x^k, y^l] = z^{kl}$.

We let S denote the generating set $\{x, y\}$ of G , and let $|g|$ denote the distance $\text{dist}_S(1, g)$, $g \in G$.

EXERCISE 14.5. Use Part 3 of Exercise 14.4 to show that

$$(14.1) \quad |x^k y^l z^m| \leq |k| + |l| + 4\sqrt{|m| + 1}.$$

LEMMA 14.6. *The group G has the presentation*

$$\langle X, Y, Z \mid [X, Y] = Z, [X, Z] = [Y, Z] = 1 \rangle.$$

PROOF. The group $H = \langle X, Y, Z \mid [X, Y] = Z, [X, Z] = [Y, Z] = 1 \rangle$ has a homomorphism ϕ to G defined by

$$\phi(X) = x, \quad \phi(Y) = y, \quad \phi(Z) = z.$$

We verify that ϕ is an isomorphism. To this end, consider the commutative diagram

$$\begin{array}{ccccccccc} 1 & \longrightarrow & \langle Z \rangle & \longrightarrow & H & \xrightarrow{ab} & \mathbb{Z}^2 & \longrightarrow & 1 \\ & & \downarrow \eta & & \downarrow \phi & & \downarrow \psi & & \downarrow \\ 1 & \longrightarrow & \langle z \rangle & \longrightarrow & G & \xrightarrow{ab} & \mathbb{Z}^2 & \longrightarrow & 1 \end{array}$$

where the homomorphisms $ab : H \rightarrow H_{ab} \simeq \mathbb{Z}^2$ and $ab : G \rightarrow G_{ab} \simeq \mathbb{Z}^2$ are the abelianization homomorphisms. We leave it to the reader to check that η and ψ are isomorphisms and to conclude from this that ϕ is an isomorphism as well. \square

By abusing the notation, we will continue to use the letters x, y for the images of the generators $x, y \in S$ under the abelianization homomorphism $G \rightarrow G^{ab}$. We will identify the Cayley graph of G^{ab} (with respect to the generating set $\{x, y\}$) with the coordinate grid in the plane \mathbb{R}^2 . We will use the coordinates X, Y in the plane so that the generators x, y correspond to the vectors $(1, 0), (0, 1)$ respectively.

Each word $w = w(x, y)$, representing the identity element of G^{ab} , defines a piecewise-linear oriented loop L_w in the plane, with edges of the unit length, where every edge is parallel to one of the coordinate axes; the loop L_w starts and ends at the origin. This loop, treated as 1-cycle in \mathbb{R}^2 , bounds a 2-chain D in \mathbb{R}^2 and we define the *signed area* $a(w) = a(L_w)$ as the integral

$$\int_D dXdY.$$

This integral, is, of course, independent of the choice of D . For instance,

$$a(x^p y^q x^{-p} y^{-q}) = pq.$$

EXERCISE 14.7. Show that

$$4\sqrt{a(w)} \leq \text{length}(L_w).$$

In the following lemma we describe a procedure of converting a word $w(x, y, z)$ into a *normal form* $x^k y^l z^m$; this is the simplest case of the similar process used in Section 14.1.3 for the proof of Proposition 14.20. Note that the redundant generating set $\{x, y, z\}$ for the group G , will be called a *closed lcs generating set* in Section 14.1.3.

LEMMA 14.8. *If $w = w(x, y)$ represents the element $z^m \in G$, then*

$$m = a(w).$$

PROOF. We will convert w to its normal form z^m by inductively moving all the letters $x^{\pm 1}$ to the left and the letters $z^{\pm 1}$ to the right. The induction is on the number of *inversions* in the word w , i.e. occurrences of letters $y^{\pm 1}$ to the left of the letters $x^{\pm 1}$.

If w has the form

$$y^q x^p u(x, y),$$

we convert it to the word

$$x^p y^q z^{-pq} u(x, y)$$

and then to

$$x^p y^q u(x, y) z^{-pq}$$

using the fact that z is a central element of G . If $L_w = \partial D$ and w' denotes $x^p y^q u(x, y)$, then $L_{w'}$ is bounded by the sum of two chains:

$$D + Q,$$

where Q is represented by the oriented rectangle bounding the loot L_c ,

$$c = x^p y^q x^{-p} y^{-q}.$$

In particular,

$$pq = a(Q)$$

and, hence,

$$a(w) = a(w') - pq.$$

The word w' has less inversions than w and, hence, inductively, we obtain:

$$w' =_G z^{a'}, a' = a(w').$$

Since

$$z^m =_G w =_G w' z^{-pq}$$

we obtain

$$m = a' - pq = a(w').$$

The case when the word w has the form

$$x^p y^q x^r v(x, y)$$

is similar (we commute the elements y^q, x^r instead) and is left to the reader. \square

Since $\text{length}(w) \geq 4\sqrt{|m|}$, we conclude that

$$4\sqrt{|m|} \leq |z^m| \leq 4\sqrt{|m| + 1},$$

and, hence,

$$4\sqrt{|m|} \leq |z^m| \leq 8\sqrt{|m|},$$

In other words, the central subgroup $\langle z \rangle$ in G has quadratic distortion.

We now consider words $w = w(x, y)$ representing arbitrary elements g of G . Suppose that g projects to $x^p y^q \in G^{ab}$. Then the word

$$w' = w y^{-q} x^{-p}$$

represents the identity element of G^{ab} and, hence (by the lemma),

$$w' =_G z^a, \quad a = a(w').$$

It follows that

$$g = x^p y^q z^a.$$

We obtain:

COROLLARY 14.9. *If $w(x, y)$ represents $x^k y^l z^m \in G$, then:*

1. *w represents the product $x^k y^l \in G^{ab}$ and, more importantly,*
2. *$m = a(w')$, where $w' = w y^{-l} x^{-k}$.*
- 3.

$$|w| \geq \max(|k| + |l|, \sqrt{|m|}) \geq \frac{1}{2}(|k| + |l| + \sqrt{|m|}).$$

COROLLARY 14.10. 1. *For each $n \in \mathbb{N}$ the map*

$$\varphi_n : x \mapsto x^n, \quad y \mapsto y^n$$

defines an endomorphism of G , such that

$$(14.2) \quad \varphi_n(U_{k,l,m}) = U_{nk,nl,n^2m}$$

and $|G : \varphi_n(G)| = n^4$.

2. *The endomorphism φ_n is expanding for each $n > 8$, see Section 8.7 for the definition.*

PROOF. 1. We will check the equation (14.2) and leave the rest to the reader. The endomorphism φ_n descends to the endomorphism of the abelianization

$$\bar{\varphi}_n : \langle x \rangle \oplus \langle y \rangle \mapsto \langle x^n \rangle \oplus \langle y^n \rangle,$$

which implies that φ_n sends $U_{k,l,m}$ to $U_{nk,nl,p}$ for some p . It remains to compute p . The endomorphism $\bar{\varphi}_n$ extends to the dilation $\mathbf{v} \mapsto n\mathbf{v}$ of \mathbb{R}^2 , which scales all length by n and all the areas by n^2 . Therefore, φ_n sends each word $w = w(x, y)$ to a word of the signed area

$$n^2 a(w).$$

Now, the claim follows from Corollary 14.9.

2. To verify the expansion property, note that

$$|U_{k,l,m}| \leq |k| + |l| + 4\sqrt{|m| + 1}$$

while

$$|\varphi_n(U_{k,l,m})| \geq (|nk| + |nl| + n\sqrt{|m|})/2.$$

Hence, for each $n > 8$, there exists $c > 1$ such that for all $g \in G \setminus \{1\}$,

$$|\varphi_n(g)| > c|g|,$$

which means that φ_n is expanding. □

We next compute the growth function of the group G :

LEMMA 14.11. $\mathfrak{G}_G(n) \asymp n^4$.

PROOF. We first note that the box

$$B_n := \{U_{klm} : -n \leq k, l \leq n, -n^2 + 1 \leq m \leq n^2 - 1\}$$

contains at least $4n^2(n^2 - 1)$ elements and each $U_{klm} \in B_n$ satisfies

$$|U_{klm}| \leq 2n + 4\sqrt{n^2} = 6n.$$

Thus, for all $n > 2$, the ball $B(1, 6n) \subset G$ contains at least $3n^4$ elements and, hence the growth function of G satisfies

$$n^4 \preceq \mathfrak{G}_G(n).$$

We next estimate the growth of G from above. The image of the ball $B(1, n) \subset G$ under the abelianization homomorphism $f : G \rightarrow G^{ab}$ equals the ball $B(1, n) \subset \langle x \rangle \oplus \langle y \rangle$. The latter has $4n^2 + 1$ elements. Since each $U_{klm} \in B(1, n)$ satisfies

$$n \geq \sqrt{|m|},$$

it follows that

$$f^{-1}(x^k y^l) \cap B(1, n)$$

contains at most $2n^2 + 1$ elements. Thus, the ball $B(1, n) \subset G$ has cardinality at most

$$(4n^2 + 1)(2n^2 + 1)$$

and $\mathfrak{G}_G(n) \lesssim n^4$. □

14.1.2. Distortion of subgroups of solvable groups.

LEMMA 14.12. *Let $G = \mathbb{Z}^m \rtimes_M \mathbb{Z}$, where $M \in GL(m, \mathbb{Z})$.*

If M has an eigenvalue with absolute value different from 1 then

$$(14.3) \quad \Delta_G^{\mathbb{Z}^m}(n) \asymp e^n.$$

PROOF. Note that (14.3) is equivalent to the existence of constants $b \geq a > 1$ and $c_i > 0$, $i = 1, 2$, such that for every $n \in \mathbb{N}$,

$$(14.4) \quad c_1 a^n \leq \Delta_G^{\mathbb{Z}^m}(n) \leq c_2 b^n.$$

Lower bound. There exists N such that M^N has an eigenvalue with absolute value at least 2. According to Proposition 8.98, we may replace in our arguments the group G by the finite-index subgroup $\mathbb{Z}^m \rtimes (N\mathbb{Z})$. Thus, without loss of generality, we may assume that M has an eigenvalue with absolute value at least 2.

Lemma 13.29 implies that there exists a vector $v \in \mathbb{Z}^m$ such that the map

$$\begin{aligned} \mathbb{Z}_2^{k+1} &\rightarrow \mathbb{Z}^m \\ s = (s_n) &\mapsto s_0 v + s_1 M v + \dots + s_k M^k v \end{aligned}$$

is injective. If we denote by t the generator of the factor \mathbb{Z} and we use the multiplicative notation for the operation in the group G , then the element

$$w_s = s_0 v + s_1 M v + \dots + s_k M^k v \in \mathbb{Z}^m$$

can be rewritten as

$$w_s = v^{s_0} (t v t^{-1})^{s_1} \dots (t^k v t^{-k})^{s_k}.$$

Thus we obtain 2^{k+1} elements of \mathbb{Z}^m of the form w_s , and if we assume that t and v are in the generating set defining the metric, the length of all these elements is at most $3k + 1$.

In the subgroup \mathbb{Z}^m we consider the generating set $X = \{e_i \mid 1 \leq i \leq m\}$, where e_i is the i -th element in the canonical basis. Then for every $w \in \mathbb{Z}^m$, $|w|_X = |w_1| + \dots + |w_m|$, i.e. $|w|_X = \|w\|_1$, where $\|\cdot\|_1$ denotes the ℓ_1 -norm on \mathbb{R}^m .

Define the number

$$r = \max\{\|w_s\|_1 : s = (s_n) \in \mathbb{Z}_2^{k+1}\}.$$

The ball in $(\mathbb{Z}^m, \|\cdot\|_1)$ with center 0 and radius r contains all the products w_s , i.e. 2^{k+1} elements, whence $r^m \geq 2^{k+1}$, and $r \geq a_1^k$, where $a_1 = 2^{\frac{1}{m}}$.

We have thus obtained that $\Delta_G^{\mathbb{Z}^m}(3k + 1) \geq a_1^k$, whence $\Delta_G^{\mathbb{Z}^m}(n) \geq a^n$, where $a = a_1^{\frac{1}{3}}$.

Upper bound. Consider the generating set $X = \{e_i \mid 1 \leq i \leq m\}$ in \mathbb{Z}^m and the generating set $S = X \cup \{t\}$ in G . Let w be an element of \mathbb{Z}^m such that $|w|_S \leq n$. It follows that

$$(14.5) \quad w = t^{k_0} v_1 t^{k_1} v_2 \dots t^{k_{\ell-1}} v_{\ell} t^{k_{\ell}},$$

where $k_j \in \mathbb{Z}$, k_0 and k_{ℓ} possibly equal to 0 but all the other exponents of t are non-zero, $v_j \in \mathbb{Z}^m$, and

$$\sum_{j=0}^{\ell} |k_j| + \sum_{j=1}^{\ell} \|v_j\|_1 \leq n.$$

We may rewrite (14.5) as

$$(14.6) \quad w = (t^{k_0}v_1t^{-k_0}) (t^{k_0+k_1}v_2t^{-k_0-k_1}) \dots (t^{k_0+\dots+k_{\ell-1}}v_{\ell}t^{-k_0-\dots-k_{\ell-1}}) t^{k_0+\dots+k_{\ell-1}+k_{\ell}}.$$

The uniqueness of the decomposition of every element in G as wt^q with $w \in \mathbb{Z}^m$ and $q \in \mathbb{Z}$, implies that $k_0 + \dots + k_{\ell-1} + k_{\ell} = 0$. With this correction, the decomposition in (14.6), rewritten with the additive notation and using the fact that $t^kvt^{-k} = M^k v$ for every $v \in \mathbb{Z}^m$, is as follows

$$w = M^{k_0}v_1 + M^{k_0+k_1}v_2 + \dots + M^{k_0+\dots+k_{\ell-1}}v_{\ell}.$$

Let α_+ be the maximum among absolute values of the eigenvalues of M , α_- be the maximum of absolute values of eigenvalue of M^{-1} ; set $\alpha = \max(\alpha_+, \alpha_-)$.

In $GL(m, \mathbb{C})$ the matrix M can be written as $PDUP^{-1}$, where D is diagonal, U is upper triangular with entries 1 on the diagonal and $DU = UD$ (the multiplicative Jordan decomposition of M).

Then $M^k = PD^kU^kP^{-1}$, and $\|M^k\| \leq \lambda \|D^k\| \|U^k\| \leq \lambda' \alpha^{|k|} k^m \leq \mu \beta^{|k|}$, for an arbitrary $\beta > \alpha$ and all sufficiently large values of k . Therefore,

$$\begin{aligned} \|w\|_1 &\leq \|M^{k_0}\| \|v_1\|_1 + \|M^{k_0+k_1}\| \|v_2\|_1 + \dots + \|M^{k_0+\dots+k_{\ell-1}}\| \|v_{\ell}\|_1 \leq \\ &\beta^{|k_0|} \|v_1\|_1 + \beta^{|k_0|+|k_1|} \|v_2\|_1 + \dots + \beta^{|k_0|+\dots+|k_{\ell-1}|} \|v_{\ell}\|_1 \leq \beta^n n \leq \beta^{2n}. \end{aligned}$$

We thus conclude that $\Delta_G^{\mathbb{Z}^m}(n) \leq \beta^{2n}$. \square

EXAMPLE 14.13. Let $G := \langle a, b : aba^{-1} = b^p \rangle$, $p \geq 2$. Then the subgroup $H = \langle b \rangle$ is exponentially distorted in G .

PROOF. To establish the lower exponential bound note that:

$$g_n := a^n b a^{-n} = b^{p^n},$$

hence $d_G(1, g_n) = 2n + 1$, $d_H(1, g_n) = p^n$, hence

$$\Delta_G^H(R) \geq p^{\lfloor (R-1)/2 \rfloor}.$$

We will leave the upper exponential bound as an exercise. \square

EXERCISE 14.14. Consider the group

$$G = \left\{ \begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix} ; a = 2^n, b = \frac{m}{2^k}, n, m, k \in \mathbb{Z} \right\}.$$

Note that G has a finite generating set consisting of matrices $d = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$ and $u = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$.

- (1) Prove that the group G has exponential growth.
- (2) Prove that the cyclic subgroup generated by u has exponential distortion.

14.1.3. Distortion of subgroups in nilpotent groups. The goal of this section is to estimate the distortion function Δ_G^H of subgroups $H \leq G$ of nilpotent groups. These estimates will be used in the proof of the Bass–Guivarc’h Theorem (Theorem 14.26).

LEMMA 14.15. *Let G be a finitely generated nilpotent group of class N and let $C^N G$ be the last non-trivial term in its lower central series. If S is a finite set of generators for G and g is an arbitrary element in $C^N G$ then there exists a constant $\lambda = \lambda(S, g)$ such that*

$$|g^n|_S \leq \lambda n^{\frac{1}{N}} \text{ for every } n \in \mathbb{N}.$$

PROOF. We argue by induction on N . The statement is clearly true for $N = 1$. Assume that it is true for N and consider G , an $(N + 1)$ -step nilpotent group.

Note that $C^N G$ is central in G , in particular it is abelian. The subgroup $C^N G$ has a finite set of generators of the form $[s, c]$, with $s \in S$ and $c \in C^{N-1} G$ (e.g., we can take as generators of $C^N G$ the inverses of N -fold left commutators of generators of G , see Lemma 13.44). Since $C^N G$ is abelian, it suffices to prove the statement of lemma for g equal to one of these generators $[s, c]$.

The formulae (3) and (4) in Lemma 13.30, imply that for every $x, x' \in G$ and $y, y' \in C^{N-1} G$ we have

$$(14.7) \quad [x, yy'] = [x, y][x, y'] \text{ and } [xx', y] = [x, y][x', y].$$

Here we used the fact that $C^N G$ is central in G to deduce that $[y, [x, y']] = 1$ and $[x, [x', y]] = 1$, and to swap $[x, y]$ and $[x', y]$.

In particular

$$(14.8) \quad [x, y^a] = [x, y]^a \text{ and } [x^b, y] = [x, y]^b.$$

Given n , we let q denote the smallest integer such that $q > n^{\frac{1}{N}}$. Note that our goal is to show that $|[s, c]^n|_S$ is bounded by λq for a suitable choice of λ .

There exist two positive integers a, b such that $n = aq^{N-1} + b$ and $0 \leq b < q^{N-1}$; moreover, $n < q^N$ implies that $a < q$. The formulas in (14.8) then imply that

$$[s, c]^n = \left[s^a, c^{q^{N-1}} \right] [s, c^b].$$

The induction hypothesis applied to the group $G/C^N G$ (where the finite generating set of the quotient is the image of S), and to the element $c \in C^{N-1} G$, implies that $c^{q^{N-1}} = k_1 z_1$ and $c^b = k_2 z_2$, where $|k_i|_S \leq \mu q$, for a constant $\mu = \mu(S, c)$, and $z_i \in C^N G$, for $i = 1, 2$.

The formulas (14.7) imply that for every $x \in G$, $[x, k_i z_i] = [x, k_i]$. Therefore

$$[s, c]^n = [s^a, k_1] [s, k_2],$$

whence $[s, c]^n$ has S -length at most

$$2(a + \mu q) + 2(1 + \mu q) \leq 4(1 + \mu)q \leq 8(1 + \mu)n^{\frac{1}{N}}.$$

Thus, we can take $\lambda = 8(1 + \mu)$. \square

COROLLARY 14.16. *Let G be a finitely generated nilpotent group of class N and let $H := C^N G$ be the last non-trivial term in its lower central series. Then:*

- (1) *The restriction of the distance function from G to H satisfies*

$$\text{dist}_G(1, g) \leq \text{dist}_H(1, g)^{\frac{1}{N}}, \quad g \in H.$$

(2) If H is infinite then its distortion function satisfies $\Delta_G^H(n) \succeq n^N$.

PROOF. The group $H = C^N G$ is abelian, hence isomorphic to $\mathbb{Z}^m \times F$ for some $m \in \mathbb{N}$ and a finite abelian group F . Let $\{t_1, \dots, t_m\}$ be a basis for \mathbb{Z}^m and let τ_1, \dots, τ_q be the respective generators of the cyclic factors of F . We consider the word metric in H corresponding to the generating set $\{t_1, \dots, t_m, \tau_1, \dots, \tau_q\}$. Take the shortest word w in this generating set representing g ,

$$g = t_1^{\alpha_1} \cdots t_m^{\alpha_m} \tau_1^{\beta_1} \cdots \tau_q^{\beta_q}.$$

Then

$$(14.9) \quad \text{dist}_H(1, g) = \sum_{i=1}^m |\alpha_i| + \sum_{j=1}^q |\beta_j|.$$

Let D denote the diameter of the finite group F with respect to dist_S and let

$$\lambda := \max_i \lambda(S, t_i),$$

where $\lambda(S, t_i)$ is as in Lemma 14.15. Then:

$$(14.10) \quad |g|_S \leq \sum_{i=1}^m |t_i^{\alpha_i}|_S + |\tau_1^{\beta_1} \cdots \tau_q^{\beta_q}|_S \leq \lambda \sum_{i=1}^m |\alpha_i|^{\frac{1}{N}} + D.$$

Now, the statement (1) follows from (14.9) and (14.10). The statement (2) is an immediate consequence of (1). \square

Our next goal is to prove the inequalities opposite to those in Corollary 14.16.

LEMMA 14.17. *Let X be a finite generating set for a nilpotent group G . There exists a finite generating set \widehat{X} containing X such that:*

- (i) \widehat{X} has the following two properties:
 1. For every $x, y \in \widehat{X}$, $[x, y] \in \widehat{X}$.
 2. Whenever $x \in \widehat{X} \cap C^i G$ projects to an element of finite order d in $C^i G / C^{i+1} G$, we have $x^d \in \widehat{X}$.
- (ii) Every set Y containing X and satisfying properties 1 and 2 contains \widehat{X} .

PROOF. (i) We define inductively on $i \geq 1$ finite subsets $T_i \subset C^i G$. Let $T_1 = X$. The set T_2 will be composed of all the commutators of elements of T_1 , and of all the powers x^d of elements $x \in T_1$ contained in some $C^k G \setminus C^{k+1} G$ and projecting to an element of order $d < \infty$ in $C^k G / C^{k+1} G$. Clearly $T_2 \subset C^2 G$.

Assume that we have defined T_1, \dots, T_i . We choose T_{i+1} to consist of all the commutators $[x, y]$ and $[y, x]$ with $x \in T_i$ and

$$y \in \bigcup_{r=1}^i T_r,$$

and of all the powers x^d of elements $x \in T_i$ contained in some $C^k G \setminus C^{k+1} G$ and projecting to an element of finite order d in $C^k G / C^{k+1} G$. The inclusion $T_i \subset C^i G$ implies that $T_{i+1} \subset C^{i+1} G$.

We then take

$$\widehat{X} := \bigcup_{i \geq 1} T_i.$$

Given $x, y \in \widehat{X}$ there exists $i \leq j$ such that $x \in T_i$ and $y \in T_j$ (or the other way around), which implies that $[x, y] \in T_{j+1}$.

Every element $z \in \widehat{X} \cap C^k G$ projecting onto a non-trivial element of finite order d in $C^k G / C^{k+1} G$ must be contained in some T_i , therefore $z^d \in T_{i+1}$.

(ii) Let Y be a set containing X and satisfying properties 1 and 2. Then Y must contain T_2 and an easy induction on $i \geq 1$ will show that Y must contain each T_i , and hence \widehat{X} . \square

DEFINITION 14.18. Let G be a finitely generated nilpotent group. We call a finite generating set S of G an *lcs-generating set* (where lcs stands for the ‘lower central series’) if for every $i \geq 1$, the subset $F^i(S) := S \cap C^i G$ generates $C^i G$. For such a generating set, we denote by $F_i(S)$ the complement $F^i(S) \setminus F^{i+1}(S)$. We say that an lcs-generating set T of G is *closed* if $T = \widehat{T}$, where \widehat{T} is defined as in Lemma 14.17.

Note that for any generating set X , the set \widehat{X} is a closed lcs-generating set, according to Lemma 13.44. Observe also that the projection of an lcs-generating set to each quotient $G/C^i G$ is again an lcs-generating set.

DEFINITION 14.19. If G is a finitely generated nilpotent group and S is an *lcs-generating set* of G , then for any word w in $S \cup S^{-1}$ we define its length $|w|_S$ as usual and its *m-length* $|w|_m$ as the number of occurrences of letters from $F_m(S) \cup (F_m(S))^{-1}$ in the word w . The *lcs-length* of a word w is the finite sequence $(|w|_1, \dots, |w|_N)$, where N is the class of G . An element g in G is said to have *lcs-length at most* (r_1, \dots, r_N) ,

$$lcs_S(g) \leq (r_1, \dots, r_N),$$

if g can be expressed as a word in $S \cup S^{-1}$ of lcs-length (m_1, \dots, m_N) , with $m_i \leq r_i$ for all $i, 1 \leq i \leq N$.

We are now ready to prove the following:

PROPOSITION 14.20. *Let G be a finitely generated nilpotent group of class N and let $H := C^k G$, $k \leq N$. Then:*

(1) *For $g \in H$, the distance function satisfies*

$$\text{dist}_H(1, g) \preceq \text{dist}_G(1, g)^N.$$

(2) *The distortion function $\Delta_G^H(n)$ of H in G satisfies*

$$\Delta_G^H(n) \preceq n^N.$$

(3) *Moreover, when $H = C^N G$,*

$$\Delta_G^H(n) \asymp n^N.$$

The statement (2) is an immediate consequence of (1), while (3) follows from the relation $\text{dist}_G(1, g)^N \preceq \text{dist}_H(1, g)$ proven for $H = C^N G$ in Corollary 14.16, (1).

In what follows we prove $\text{dist}_H(1, g) \preceq \text{dist}_G(1, g)^N$, for g an arbitrary element in a subgroup $H = C^k G$. The main step in the proof is the following lemma.

LEMMA 14.21. Let G be a finitely generated nilpotent group of class N with a closed lcs generating set S . Then there exists a sequence of closed lcs-generating sets $S^{(k)}$ of the group G , $k = 1, \dots, N$, such that the following holds:

For every pair of numbers $\lambda \geq 1$, $r \geq 1$ and every element $g \in C^k G$ with $lcs_S(g) \leq (\lambda r, \lambda r^2, \dots, \lambda r^N)$, we have:

(I)

$$lcs_{S^{(k)}}(g) \leq \left(\underbrace{0, \dots, 0}_{k-1}, \lambda_k r^k, \lambda_k r^{k+1}, \dots, \lambda_k r^N \right),$$

where λ_k depends only on S and on λ ;

(II) furthermore, for $k = N$, the standard word length satisfies

$$|g|_{S^{(N)}} \leq \lambda_N r^N.$$

PROOF. We construct the generating sets $S^{(k)}$, $k = 1, \dots, N$ by induction on k . For $k = 1$ we simply take $S^{(1)} = \widehat{S}$.

We describe only the induction step $1 \rightarrow 2$, since the general induction $n \rightarrow n+1$ is identical (replacing $G = C^1 G$ with $C^n G$). We also verify the inequality

$$lcs_{S^{(2)}}(g) \leq (0, \lambda_2 r^2, \lambda_2 r^3, \dots, \lambda_2 r^N).$$

Part (II) will be an immediate corollary of Part (I).

Step 1: Construction of the generating sets. Our goal is to define (given a generating set $S^{(1)}$ and a number $\lambda \geq 1$) a new closed lcs-generating set $S^{(2)}$ for G and a constant λ_2 such that, whenever $g \in C^2 G$ satisfies

$$lcs_{S^{(1)}}(g) \leq (\lambda r, \lambda r^2, \dots, \lambda r^N), \text{ for some } r \geq 1,$$

we also have

$$(14.11) \quad lcs_{S^{(2)}}(g) \leq (0, \lambda_2 r^2, \dots, \lambda_2 r^N).$$

Consider a list of elements in G ,

$$A = \{t_1, \dots, t_m\},$$

that projects to a *standard* generating set of the abelianization $G_{ab} = G/G' = G/C^2 G$ (see Definition 13.14).

We take first a new generating set constructed as follows: $Z = A \cup F^2(S^{(1)})$, and then its closure $S^{(2)} = \widehat{Z}$. In $S^{(2)}$ we have thus replaced $F_1(S^{(1)})$ with A , and added to $F^2(S^{(1)})$ powers $t_i^{d_i}$ of the elements t_i in A that project to standard generators of finite order d_i in G_{ab} , and commutators with elements from A or with powers $t_i^{d_i}$ as described above. For an element g with

$$lcs_{S^{(1)}}(g) \leq (r_1, \dots, r_N),$$

we can deduce that

$$lcs_{S^{(2)}}(g) \leq (\rho_1, \dots, \rho_N),$$

where $\rho_1 \leq \eta r_1$, and $\rho_i \leq r_i + \eta r_1$, with η a constant depending on $S^{(1)}$ and on A .

In particular, an inequality of the form

$$lcs_{S^{(1)}}(g) \leq (\lambda r, \dots, \lambda r^N), \text{ for some } r \geq 1$$

will be preserved for $S^{(2)}$, with λ increased but independent of r .

Let $g \in C^2G$ be an element satisfying an inequality as above, for $S^{(1)}$, hence with an increased λ for $S^{(2)}$. Let w_0 be a word in $S^{(2)}$ representing g and such that $|w_0|_i \leq \lambda r^i$. Our next goal is to modify w_0 and create a different word representing g that would justify the inequality (14.11).

For simplicity, we assume that all the letters t_1, \dots, t_m appear with at least one positive or negative power in w_0 , otherwise we would have to replace in the argument below this sequence of letters by the finite subsequence t_{q_1}, \dots, t_{q_u} of the letters that actually appear in w_0 , with either positive or negative power, with the rest of the proof carried almost *verbatim*.

Let ℓ_i denote the number of times the letters $t_i^{\pm 1}$ appear in w_0 and let ℓ denote $|w_0|_1$. We have that $\ell_1 + \dots + \ell_m = \ell$.

By induction on $0 \leq j \leq \ell_1$ we construct a word w_j such that

- (i) $w_j = t_1^{\alpha_j} u_j$ with $|\alpha_j| \leq j$;
- (ii) in u_j , $t_1^{\pm 1}$ occur $\ell_1 - j$ times;
- (iii) $|w_{j+1}|_1 \leq |w_j|_1$, $|w_{j+1}|_2 \leq |w_j|_2 + 1$ and for every $3 \leq k \leq N$, $|w_{j+1}|_k \leq |w_j|_k + |w_j|_{k-1}$.

Let us explain how w_{j+1} is obtained from w_j (the initial step, for $j = 0$, is similar). We consider the left-most occurrence of a letter $t = t_1^{\pm 1}$ in the word u_j , and move it to the left using the formulas

$$xt = tx [x^{-1}, t^{-1}].$$

Whenever we see the product $t_1 t_1^{-1}$ or $t_1^{-1} t_1$, we cancel these generators. If t_1 projects to a generator of order d_1 in G_{ab} , and, at some point of the process, we see a d_1 -fold product

$$\underbrace{t_1 \dots t_1}_{d_1 \text{ times}},$$

we replace this product with the generator $t_1^{d_1} \in F^2(S^{(2)})$.

Once we finish the induction process with $j = \ell_1$, we continue with another induction for $\ell_1 \leq j \leq \ell_1 + \ell_2$, and move the letters $t_2^{\pm 1}$ as far left as possible, but not beyond the power t_1 constructed earlier. We then proceed inductively for $t_2^{\pm 1}, \dots, t_m^{\pm 1}$, and obtain thus a sequence of words w_j with $j \leq \ell_1 + \ell_2 + \dots + \ell_m = \ell$.

In the end of the induction process, we convert w_0 to a word w_ℓ of the form

$$t_1^{\epsilon_1} \dots t_m^{\epsilon_m} w',$$

where w' is a word in the alphabet $F^2(S^{(2)}) \cup [F^2(S^{(2)})]^{-1}$ and, whenever t_j projects to an element of finite order d_j in G_{ab} , we have

$$0 \leq \epsilon_j < d_j.$$

Since the set $F_1(S^{(2)})$ projects to a standard generating set of G_{ab} and $g \in C^2G$, it follows that

$$t_1^{\epsilon_1} \dots t_m^{\epsilon_m} = 1$$

in G ; thus, the element $g \in G$ is represented by the word w' .

Step 2: Estimating the lcs-length of the word w' .

Using (iii) we obtain that for every $i = 1, \dots, N$, and $j = 1, \dots, \ell$,

$$(14.12) \quad |w_j|_i \leq \sum_{s=0}^a \binom{j}{s} |w_0|_{i-s}$$

where $a = \min(i-1, j)$. The induction step follows from the formula

$$\binom{j}{s} + \binom{j}{s-1} = \binom{j+1}{s}.$$

Since $|w_0|_{i-s} \leq \lambda r^{i-s}$ and

$$\binom{j}{s} \leq j^s \leq \ell^s \leq \lambda^s r^s,$$

we obtain:

$$|w_j|_i \leq \sum_{s=0}^a \ell^s \lambda r^{i-s} \leq \sum_{s=0}^a \lambda^s r^s \lambda r^{i-s} \leq \lambda \left(\sum_{s=0}^a \lambda^s \right) r^i \leq (a+1) \lambda^{a+1} r^i.$$

In particular, for $2 \leq i \leq N$,

$$|w'|_i \leq |w_\ell|_i \leq (a+1) \lambda^{a+1} r^i \leq \lambda_2 r^i,$$

where

$$\lambda_2 = N \lambda^N.$$

As we observed before, the construction of the generating sets $S^{(i)}$ is similar, we note that for $g \in C^{i+1}G$ the inequality

$$lcs_{S^{(i)}}(g) \leq \underbrace{(0, \dots, 0)}_{i-1}, \lambda_i r^i, \lambda_i r^{i+1}, \dots, \lambda_i r^N$$

implies

$$lcs_{S^{(i+1)}}(g) \leq \underbrace{(0, \dots, 0)}_i, \lambda_{i+1} r^{i+1}, \lambda_{i+1} r^{i+2}, \dots, \lambda_{i+1} r^N.$$

This finishes the proof of **(I)**. In particular, for $g \in C^N G$, we have

$$|g|_{S^{(N)}} \leq \lambda_N r^N,$$

where λ_N is independent of r and g . □

Now, we can conclude the proof of Proposition 14.20. We start with an lcs-generating set S of G . Let g be an arbitrary element in $C^k G$, with $1 \leq k \leq N$, and let $|g|_S = r$. For $S^{(1)} = \widehat{S}$, we can write

$$lcs_{S^{(1)}}(g) \leq (r, r, \dots, r) \leq (r, r^2, \dots, r^N).$$

Applying Lemma 14.21, we obtain a new lcs generating set $T := S^{(k)}$ of G such that

$$lcs_T(g) \leq \underbrace{(0, \dots, 0)}_{k-1}, \lambda_k r^k, \dots, \lambda_k r^N.$$

It follows that

$$|g|_{F^k(T)} \leq \lambda_k N r^N \leq \lambda_k N |g|_S^N.$$

□

Proposition 14.20 generalizes to all subgroups:

PROPOSITION 14.22. *Let G be a finitely generated nilpotent group of class N . Then, for every subgroup H in G ,*

$$\Delta_G^H(n) \leq n^\gamma,$$

where $\gamma = N!$.

PROOF. We will first prove a general lemma about distortion of subgroups. Suppose that we have a commutative diagram of finitely generated groups where the horizontal sequences are short exact and the vertical arrows are injective:

$$\begin{array}{ccccccccc}
1 & \longrightarrow & H' & \longrightarrow & H & \xrightarrow{\psi} & H'' & \longrightarrow & 1 \\
& & \uparrow f' & & \uparrow f & & \uparrow f'' & & \\
1 & \longrightarrow & G' & \longrightarrow & G & \xrightarrow{\phi} & G'' & \longrightarrow & 1
\end{array}$$

We will, henceforth, identify H', H, H'' with subgroups of G', G, G'' respectively. We let X, Y denote finite symmetric generating sets of G, H respectively, such that $f(Y) \subset X$, $X' = X \cap G'$, $X'' = \phi(X)$, $Y' = \psi(Y)$, $Y'' = Y \cap H'$ are generating sets of the respective groups.

LEMMA 14.23.

$$\Delta_G^H(m) \leq 2\Delta_{G'}^{H'} \circ \Delta_G^{G'}(2\Delta_{G''}^{H''}(m)).$$

PROOF. We take $h \in H$ such that $|f(h)|_X = m$. Then

$$|\phi(f(h))|_{X''} \leq m.$$

The distortion bound on the inclusion $H'' \hookrightarrow G''$ yields the existence of a word w'' in the alphabet Y'' whose length is

$$|w''|_{Y''} \leq \Delta_{G''}^{H''}(m),$$

such that w'' represents the projection $\psi(h)$. We lift w'' to a word \tilde{w}'' in Y and let $\tilde{h}'' \in H$ be the element represented by \tilde{w}'' . The product

$$h' := h \cdot (\tilde{h}'')^{-1},$$

belongs to the subgroup H' . We would like to bound the norm $|h'|_{Y'}$ from above using the distortion function

$$\Delta_{G'}^{H'}.$$

To this end, consider the word $v'' = f''(w'')$ (in the alphabet X'') and its lift \tilde{v}'' (in the alphabet X); the latter word represents the element $f(\tilde{h}'') \in G$. Clearly, v'' has length $\leq \Delta_{G''}^{H''}(m)$. We have

$$f'(h') = f(h) \cdot f(\tilde{h}'')^{-1},$$

which implies that

$$|f'(h')|_X \leq |f(h)|_X + |\tilde{v}''| \leq m + \Delta_{G''}^{H''}(m).$$

Therefore,

$$|f'(h')|_{X'} \leq \Delta_{G'}^{G'}(m + \Delta_{G''}^{H''}(m)),$$

and, hence,

$$|h'|_{Y'} \leq \Delta_{G'}^{H'}\left(\Delta_{G'}^{G'}(m + \Delta_{G''}^{H''}(m))\right).$$

By putting the inequalities together, we obtain

$$|h|_Y \leq |w''|_{Y''} + |h'|_{Y'} \leq \Delta_{G''}^{H''}(m) + \Delta_{G'}^{H'}\left(\Delta_{G'}^{G'}(m + \Delta_{G''}^{H''}(m))\right).$$

Since each distortion function Δ_A^B satisfies

$$\Delta_A^B(k) \geq k,$$

we obtain:

$$|h|_Y \leq 2\Delta_{G'}^{H'} \left(\Delta_G^{G'} (2\Delta_{G''}^{H''}(m)) \right),$$

as required. \square

With this lemma in mind, we will now prove the proposition by induction on the nilpotency class N of G . The statement is obviously true for $N = 1$ since subgroups of abelian groups are undistorted. Assume that proposition holds for N and consider a subgroup H in a group G of nilpotency class $N+1$. Set $G' = C^{N+1}G$, $H' := H \cap G'$ and let $G'' = G/G'$, $H'' = H/H'$, the projection of H into G'' . We have (for suitable finite generating sets):

$$\Delta_{G''}^{H''}(m) \leq C_1 m^{N!},$$

for some constant C_1 , by the induction hypothesis. We also have

$$\Delta_{G'}^{H'}(s) = s,$$

since subgroups of abelian groups are undistorted, and

$$\Delta_G^{G'}(t) \leq C_2 \cdot t^{N+1}$$

(by Proposition 14.20). Now, we can write

$$\Delta_G^H(m) \leq 2C_2 (2C_1 m^{N!})^{N+1} = 2^{N+2} C_2 C_1^{N+1} m^{(N+1)!}. \quad \square$$

In fact, a stronger statement holds:

THEOREM 14.24. *For every infinite subgroup H in a finitely generated nilpotent group G there exists a rational positive number α such that*

$$\Delta_G^H(n) \asymp n^\alpha.$$

Theorem 14.24 was originally proven by M. Gromov in [Gro93] (see also [Var99]); later on, an explicit computation of the possible exponents α was established by D. Osin in [Osi01]. More precisely, given an element of infinite order h in a nilpotent group G , its *weight* in G , $\nu_G(h)$, is defined as the maximal i such that $\langle h \rangle \cap C^i G \neq \{1\}$. The exponent α in Theorem 14.24 is the maximum of the fractions

$$\frac{\nu_G(h)}{\nu_H(h)}$$

over all elements $h \in H$ of infinite order.

Another interesting consequence of Lemma 14.21 is a control on the exponents of the bounded generation property for nilpotent groups.

PROPOSITION 14.25 (Controlled bounded generation for nilpotent groups). *Let G be a finitely generated nilpotent group of class N . For each i we let*

$$S_i = \{t_{i1}, \dots, t_{iq_i}\} \subset C^i G$$

be a subset projecting bijectively to a standard generating set of the abelian group $C^i G / C^{i+1} G$. Define

$$S = \bigcup_{i=1}^N S_i$$

(thus, S is an lcs-generating set of G). Then every element $g \in G$ can be written as a product

$$g = \prod_{i=1}^N t_{i1}^{m_{i1}} \cdots t_{iq_i}^{m_{iq_i}},$$

so that:

1. $m_{ij} \in \mathbb{Z}$ and $0 \leq m_{ij} < d_{ij}$, if $d_{ij} < \infty$ is the order of the projection of t_{ij} to $C^i G / C^{i+1} G$.

2. $|m_{i1}| + \dots + |m_{iq_i}| \leq C |g|_S^i$, for every $i \in \{1, \dots, k\}$, where C is a constant depending only on G and on S .

PROOF. We argue by induction on the class N . For $N = 1$ the group is abelian and the statement is obvious. Assume that the statement is true for the class $N - 1$ and let G be a nilpotent group of class $N \geq 2$. Let S_i and S be as in the statement of the proposition, and let g be an arbitrary element in G . The induction hypothesis implies that $g = pc$, where $c \in C^{N-1} G$ and

$$p = \prod_{i=1}^{N-1} t_{i1}^{m_{i1}} \cdots t_{iq_i}^{m_{iq_i}},$$

where $m_{ij} \in \mathbb{Z}$ are such that $0 \leq m_{ij} < d_{ij}$ (if the order d_{ij} is finite) and

$$(14.13) \quad |m_{i1}| + \dots + |m_{iq_i}| \leq C |g|_S^i,$$

for every $i \in \{1, \dots, k\}$, where C is a constant depending only on G and S .

Then, by the inequalities (14.13), the element $c = p^{-1}g$ in $C^{N-1} G$ has lcs-length with respect to S at most $(\lambda r, \lambda r^2, \dots, \lambda r^N)$, where $r = |g|_S$ and $\lambda = C + 1$. Without loss of generality, we may assume that S is replaced by its closure \widehat{S} , since taking the closure only decreases the lcs-length.

Lemma 14.21 then implies that there exists a new generating set T of G , (determined by S), such that

$$lcs_T(c) \leq (0, \dots, 0, \mu r^N),$$

where μ only depends on T . By the construction of this generating set, $T \cap C^{N-1} G$ is a generating set of the abelian group $C^{N-1} G$. Thus

$$c = t_{N1}^{m_{N1}} \cdots t_{Nq_N}^{m_{Nq_N}},$$

where

$$|m_{N1}| + \dots + |m_{Nq_N}| \leq \eta |c|_{T \cap C^{N-1} G} \leq \eta \mu r^N,$$

where η depends only on S and on T . Now, the assertion follows by combining the product decompositions of p and c . \square

14.2. Polynomial growth of nilpotent groups

Let G be a finitely generated nilpotent group and $d = d(G)$ is the *homogeneous dimension* of G , see Definition 13.46.

THEOREM 14.26 (Bass–Guivarc’h Theorem). *The growth function of G satisfies*

$$(14.14) \quad \mathfrak{G}_G(n) \asymp n^d.$$

PROOF. In the proof below, m_i is the free rank of $C^i G / C^{i+1} G$; the constants λ_i depend only on the generating set of the group G . We will use the notation $B_G(1, r)$ to denote the r -ball in the group G centered at $1 \in G$, with respect to the word metric defined by a suitable finite generating set of G .

We argue by induction on the class k of nilpotency of G . For $k = 1, d = m_1$, the group G is abelian of free rank d and the statement is obvious.

Assume that the statement holds for $k - 1$ and consider G of nilpotency class $k \geq 2$; let $H = C^k G$ be the last non-trivial subgroup in the lower central series of G . Let $d_1 := d - k m_k$ be the dimension of G/H .

If H is finite then $m_k = 0$ and we apply the induction hypothesis for G/H ; since G and G/H have equivalent growth functions the result follows.

We now assume that H is infinite, i.e. $m_k \geq 1$. Fix a finite generating set S of G and use its projection as the generating set of G/H .

Lower bound. By our choice of generating sets, the ball $B_G(1, r)$ maps onto the ball $B_{G/H}(1, r)$ under the projection $G \rightarrow G/H$. The induction hypothesis applied to G/H implies that there exists λ_1 for which

$$N = \text{card}(B_{G/H}(1, n)) \geq \lambda_1 n^{d_1}.$$

Let $\{g_1, \dots, g_N\} \subset B_G(1, n)$ denote the preimage of $B_{G/H}(1, n)$. Since the abelian group H has growth function t^{m_k} , and, according to Part (1) of Corollary 14.16, (for some μ independent of n)

$$B_G(1, n) \cap H \supset B_H(1, \mu n^k),$$

we conclude that

$$\text{card}(B_G(1, n) \cap H) \geq \lambda_2 n^{k m_k},$$

for some λ_2 independent of n .

Therefore, the ball $B_G(1, 2n)$ contains the set

$$\bigcup_{i=1}^N g_i (B_G(1, n) \cap H)$$

of cardinality at least

$$N \lambda_2 n^{k m_k} \geq \lambda_1 \lambda_2 n^{d_1 + k m_k} = \lambda_3 n^d = \lambda_3 2^{-d} (2n)^d.$$

Thus, for every even $t = 2n$,

$$\mathfrak{G}_G(t) \geq \lambda_4 t^d.$$

Since the growth function is increasing, the inequality above is easily extended to odd t , with a suitable λ_4 .

Upper bound. The proof is analogous to the lower bound. Recall that the image of $B_G(1, n)$ in G/H is the ball $B_{G/H}(1, n)$. By the inductive hypothesis there exist at most $\lambda_5 n^{d_1}$ elements

$$\bar{g}_1, \dots, \bar{g}_N \in B_{G/H}(1, n),$$

which are projections of elements $g_i \in B_G(1, n)$, $i = 1, \dots, N$. Each element $g \in B_G(1, n)$ can be written as

$$g = g_i h \in g_i H \cap B_G(1, n)$$

for some $i \in \{1, \dots, N\}$. We have

$$|h|_S \leq |g|_S + |g_i|_S \leq 2n.$$

By Proposition 14.20 there are at most $\lambda_6 n^{km_k}$ elements $h \in H$ satisfying this inequality. It then follows that there are at most $\lambda_5 \lambda_6 n^{d_1 + km_k} = \lambda_7 n^d$ distinct elements $g \in B_G(1, n)$. \square

14.3. Wolf's Theorem

In this section we classify virtually polycyclic groups according to their growth.

NOTATION 14.27. If G is a group, a semidirect product $G \rtimes_{\Phi} \mathbb{Z}$ is defined by a homomorphism $\Phi : \mathbb{Z} \rightarrow \text{Aut}(G)$. The latter homomorphism is entirely determined by $\Phi(1) = \varphi$. Following the notation in Section 14.1, we set

$$S = G \rtimes_{\varphi} \mathbb{Z} := G \rtimes_{\Phi} \mathbb{Z}$$

PROPOSITION 14.28. *Let G be a finitely generated nilpotent group and let $\varphi \in \text{Aut}(G)$. Then the polycyclic group $P = G \rtimes_{\varphi} \mathbb{Z}$ is*

- (1) *either virtually nilpotent;*
- (2) *or has exponential growth.*

REMARK 14.29. The statement (1) in Proposition 14.28 cannot be improved to ' P is nilpotent', see Remark 14.2, Part (2).

PROOF. We note that replacing φ with a power will replace P with a finite-index subgroup, and, hence, will not affect virtual nilpotency of P and its growth rate. The automorphism φ preserves the lower central series of G ; let θ_i denote the restriction of φ to $C^i G$, $i \geq 1$. Then θ_i projects to an automorphism φ_i of the finitely generated abelian group $B_i := C^i G / C^{i+1} G$. The automorphism φ_i induces an automorphism ψ_i of $\text{Tor } B_i$ and an automorphism $\bar{\varphi}_i$ of $B_i / \text{Tor } B_i \simeq \mathbb{Z}^{m_i}$. Each choice of a basis for $B_i / \text{Tor } B_i$ associates to the automorphism $\bar{\varphi}_i$ a matrix M_i in $GL(m_i, \mathbb{Z})$. All the conditions below are independent of the choice of a basis, therefore in what follows we assume that an arbitrary fixed basis is chosen in each $B_i / \text{Tor } B_i$. As in the proof of Proposition 14.1, we have two cases to consider:

(1) All matrices M_i only have eigenvalues of absolute value 1; hence, by Lemma 13.28, all the eigenvalues are roots of unity. Then there exists N such that the matrices of the automorphisms $\bar{\varphi}_i^N$ have only eigenvalues equal to 1 and the induced automorphisms of finite abelian groups

$$\psi_i : \text{Tor } B_i \rightarrow \text{Tor } B_i$$

are all equal to the identity. Without loss of generality we may therefore assume that the matrices M_i of all the φ_i 's have all eigenvalues equal to 1, and that all the ψ_i are the identity automorphisms.

Lemma 13.27 applied to each $\bar{\varphi}_i$ and to each $\psi_i = \text{id}_{\text{Tor } B_i}$, imply that the lower central series of G is a sub-series of a cyclic series

$$\{1\} = H_n \leq H_{n-1} \leq \dots \leq H_1 \leq H_0 = G,$$

where each H_i / H_{i+1} is cyclic, φ preserves each H_i and induces the identity map on H_i / H_{i+1} . We denote by t the generator of the semidirect factor \mathbb{Z} in the decomposition $P = G \rtimes \mathbb{Z}$. By the definition of the semidirect product, for every $g \in G$, $tgt^{-1} = \varphi(g)$. The fact that φ acts as the identity on each H_i / H_{i+1} implies that $t^k(hH_{i+1})t^{-k} = hH_{i+1}$ for every h in H_i ; equivalently

$$(14.15) \quad [t^k, h] \in H_{i+1}$$

for every such h .

Since G contains the kernel $C^2P = [P, P]$ of the abelianization map $G \rightarrow \mathbb{Z}$, it follows that $C^2P \leq G$. We claim that for every $i \geq 0$, $[P, H_i] \subseteq H_{i+1}$. Indeed, consider an arbitrary commutator $[h, s]$, $h \in H_i, s \in P$. Since s has the form $s = gt^k$, with $g \in G$ and $k \in \mathbb{Z}$, we obtain:

$$[h, s] = [h, gt^k] = [h, g][g, [h, t^k]][h, t^k],$$

in view of the commutator identity (3) in Lemma 13.30.

According to (14.15), $[h, t^k] \in H_{i+1}$. Also, since the lower central series of G is a subseries of (H_i) , there exists $r \geq 1$ such that $C^rG \geq H_i \geq H_{i+1} \geq C^{r+1}G$. Then, $h \in H_i \leq C^rG$ and

$$[h, g] \in C^{r+1}G \leq H_{i+1}.$$

Likewise, as $[h, t^k] \in H_{i+1} \leq C^rG$, the commutator

$$[g, [h, t^k]] \in C^{r+1}G \leq H_{i+1}.$$

By putting it all together, we conclude that $[h, s] \in H_{i+1}$ and, hence, $[P, H_i] \subseteq H_{i+1}$.

An easy induction now shows that $C^{i+2}P \leq H_i$ for every $i \geq 1$; in particular, $C^{n+2}P \leq H_n = \{1\}$. Therefore, P is virtually nilpotent.

(2) Assume that at least one matrix M_i has an eigenvalue with absolute value strictly greater than 1, in particular, $m_i \geq 2$. The group P contains the subgroup

$$P_i := C^iG \rtimes_{\theta_i} \mathbb{Z}.$$

Furthermore, the subgroup $C^{i+1}G$ is normal in P_i and

$$P_i/C^{i+1}G \simeq B_i \rtimes_{\varphi_i} \mathbb{Z},$$

where $B_i = C^iG/C^{i+1}G$. Lastly,

$$(B_i \rtimes_{\varphi_i} \mathbb{Z})/\text{Tor } B_i \cong \mathbb{Z}^{m_i} \rtimes_{M_i} \mathbb{Z}.$$

According to Proposition 14.1, the group $\mathbb{Z}^{m_i} \rtimes_{M_i} \mathbb{Z}$ has exponential growth. Therefore, in view of Proposition 8.78, parts (a) and (c), the groups $B_i \rtimes_{\varphi_i} \mathbb{Z}$, $P_i/C^{i+1}G$, P_i , and, hence, P , all have exponential growth. Thus, in the case (2), S has exponential growth. \square

We use Proposition 14.28 combined with Proposition 5.11 on subgroups of finite index in finitely generated groups to prove Wolf's Theorem [Wol68]:

THEOREM 14.30 (Wolf's Theorem). *A polycyclic group is either virtually nilpotent or has exponential growth.*

PROOF. According to Proposition 13.81, it suffices to prove the statement for poly- C_∞ groups. Let G be a poly- C_∞ group, and consider a finite subnormal descending series

$$G = N_0 \geq N_1 \geq \dots \geq N_n \geq N_{n+1} = \{1\}$$

such that $N_i/N_{i+1} \simeq \mathbb{Z}$ for every $i \geq 0$. We argue by induction on n . For $n = 0$ the group G is infinite cyclic and the statement is obvious. Assume that the assertion of theorem holds for n and consider the case of $n + 1$. By the induction hypothesis, the subgroup $N_1 \leq G$ is either virtually nilpotent or has exponential growth. In the second case the group G itself has exponential growth.

Assume that N_1 is virtually nilpotent. Corollary 7.24 implies that G decomposes as a semidirect product $N_1 \rtimes_{\theta} \mathbb{Z}$, corresponding to a homomorphism $\Psi : \mathbb{Z} \rightarrow \text{Aut}(N_1)$, $\theta = \Psi(1)$.

By hypothesis, N_1 contains a nilpotent subgroup H of finite index. According to Proposition 5.11, Part (2), we may moreover assume that H is characteristic in N_1 . In particular H is invariant under the automorphisms ψ . We retain the notation θ for the restriction $\theta|_H$. Therefore, $H \rtimes_{\theta} \mathbb{Z}$ is a normal subgroup of G . Moreover, $H \rtimes_{\theta} \mathbb{Z}$ has finite index in G , since $G/(H \rtimes_{\theta} \mathbb{Z})$ is the quotient of the finite group N_1/H .

By Proposition 14.28, $H \rtimes_{\theta} \mathbb{Z}$ is either virtually nilpotent or of exponential growth. Therefore, the same alternative holds for $N_1 \rtimes_{\theta} \mathbb{Z} = G$. \square

14.4. Milnor's theorem

THEOREM 14.31 (J. Milnor, [Mil68a]). *A finitely generated solvable group is either polycyclic or has exponential growth.*

We begin the proof by establishing a property of conjugates implied by sub-exponential growth:

LEMMA 14.32. *If a finitely generated group G has sub-exponential growth then for all $\beta_1, \dots, \beta_m, g \in G$, the set of conjugates*

$$\{g^k \beta_i g^{-k} \mid k \in \mathbb{Z}, i = 1, \dots, m\}$$

generates a finitely generated subgroup $N \leq G$.

PROOF. It suffices to prove lemma for $m = 1$, since N is generated by the subgroups

$$N_i = \langle g^k \beta_i g^{-k} \mid k \in \mathbb{Z} \rangle, \quad i = 1, \dots, m.$$

NOTATION 14.33. We set $\alpha := \beta_1$ and let α_k denote $g^k \alpha g^{-k}$ for $k \in \mathbb{Z}$. In the proof we will be identifying \mathbb{Z}_2 with the subset $\{0, 1\}$ of \mathbb{Z} .

The goal is to prove that finitely many elements in the set $\{\alpha_k \mid k \in \mathbb{Z}\}$ generate the subgroup N .

Consider the map

$$\begin{aligned} \mu &= \mu_m : \prod_{i=0}^m \mathbb{Z}_2 \rightarrow G \\ \mu : (s_i) &\mapsto g\alpha^{s_0} g\alpha^{s_1} \cdots g\alpha^{s_m}. \end{aligned}$$

EXERCISE 14.34. Verify that

$$g\alpha^{s_0} g\alpha^{s_1} \cdots g\alpha^{s_m} = \alpha_1^{s_0} \alpha_2^{s_1} \cdots \alpha_{m+1}^{s_m} g^{m+1}.$$

If for every $m \in \mathbb{N}$ the map μ is injective then for each sequence (s_i) we have 2^{m+1} products as above, and if $g, g\alpha$ are in the set of generators of G , all these products are in $B_G(1, m+1)$. This contradicts the hypothesis that G has sub-exponential growth. It follows that there exists some m and two distinct sequences $(s_i), (t_i)$ in $\prod_{i=0}^m \mathbb{Z}_2$ such that

$$(14.16) \quad g\alpha^{s_0} g\alpha^{s_1} \cdots g\alpha^{s_m} = g\alpha^{t_0} g\alpha^{t_1} \cdots g\alpha^{t_m}.$$

Assume that m is minimal with this property. This, in particular, implies that $s_0 \neq t_0$ and $s_m \neq t_m$. In view of the exercise, the equality (14.16) becomes

$$\alpha_1^{s_0} \alpha_2^{s_1} \cdots \alpha_{m+1}^{s_m} = \alpha_1^{t_0} \alpha_2^{t_1} \cdots \alpha_{m+1}^{t_m}.$$

Since $s_m \neq t_m$ and $s_m, t_m \in \{0, 1\}$, it follows that $s_m - t_m = \pm 1$. Then

$$(14.17) \quad \alpha_{m+1}^{\pm 1} = \alpha_m^{-s_m-1} \cdots \alpha_2^{-s_1} \alpha_1^{t_0-s_0} \alpha_2^{t_1} \cdots \alpha_m^{t_{m-1}}.$$

If in (14.17) we conjugate by g , we obtain that

$$\alpha_{m+2}^{\pm 1} = \alpha_{m+1}^{-s_m-1} \cdots \alpha_3^{-s_1} \alpha_2^{t_0-s_0} \alpha_3^{t_1} \cdots \alpha_{m+1}^{t_{m-1}}.$$

This and (14.17) imply that α_{m+2} is a product of powers of $\alpha_1, \dots, \alpha_m$. Then, by induction, every α_n with $n \in \mathbb{N}$ is a product of powers of $\alpha_1, \dots, \alpha_m$, and the same is true for α_n with $n \in \mathbb{Z}$ by replacing g with g^{-1} . Therefore, every generator α_n of N belongs to the subgroup of N generated by the elements $\alpha_1, \dots, \alpha_m$ and the elements $\alpha_1, \dots, \alpha_m$ generate N . \square

EXERCISE 14.35. Use Lemma 14.32 to prove that the finitely generated group H described in Example 7.8 has exponential growth.

We now are ready to prove Theorem 14.31; our proof by induction on the derived length d of G . For $d = 1$ the group G is finitely generated abelian and the statement is immediate. Assume that the alternative holds for finitely generated solvable groups of derived length $\leq d$ and consider G of derived length $d + 1$. Then $H = G/G^{(d)}$ is finitely generated solvable of derived length d . By the induction hypothesis, either H has exponential growth or H is polycyclic. If H has exponential growth then G has exponential growth too (see statement (c) in Proposition 8.78).

Assume therefore that H is polycyclic. In particular, H is finitely presented by Proposition 13.84. Theorem 14.31 will follow from:

LEMMA 14.36. *Consider a short exact sequence*

$$(14.18) \quad 1 \rightarrow A \rightarrow G \xrightarrow{\pi} H \rightarrow 1, \quad \text{with } A \text{ abelian and } G \text{ finitely generated.}$$

If H is polycyclic then G is either polycyclic or has exponential growth.

PROOF. We assume that G has sub-exponential growth and will prove that G is polycyclic. The group G is polycyclic iff A is finitely generated. Since H is polycyclic, it has the bounded generation property (see Proposition 13.74); hence, there exist finitely many elements h_1, \dots, h_q in H such that every element $h \in H$ can be written as

$$h = h_1^{m_1} h_2^{m_2} \cdots h_q^{m_q}, \quad \text{with } m_1, m_2, \dots, m_q \in \mathbb{Z}.$$

Choose $g_i \in G$ such that $\pi(g_i) = h_i$ for every $i \in \{1, 2, \dots, q\}$. Then every element $g \in G$ can be written as

$$(14.19) \quad g = g_1^{m_1} g_2^{m_2} \cdots g_q^{m_q} a, \quad \text{with } m_1, m_2, \dots, m_q \in \mathbb{Z} \text{ and } a \in A.$$

Since H is finitely presented, by Lemma 7.29 there exist finitely many elements a_1, \dots, a_k in A such that every element in A is a product of G -conjugates of a_1, \dots, a_k . According to (14.19), all the conjugates of a_j are of the form

$$(14.20) \quad g_1^{m_1} g_2^{m_2} \cdots g_q^{m_q} a_j (g_1^{m_1} g_2^{m_2} \cdots g_q^{m_q})^{-1}.$$

By Lemma 14.32, the subgroup A_q generated by the conjugates $g_q^m a_j g_q^{-m}$ with $m \in \mathbb{Z}$ and $j \in \{1, \dots, k\}$ is finitely generated. Let S_q be its finite generating set.

Then the conjugates $g_{q-1}^n g_q^m a_j g_q^{-m} g_{q-1}^{-n}$ with $m, n \in \mathbb{Z}$ and $j \in \{1, \dots, k\}$ are in the subgroup A_{q-1} of A generated by $g_{q-1}^n s g_{q-1}^{-n}$ with $n \in \mathbb{Z}$ and $s \in S_q$. Again Lemma 14.32 implies that the subgroup A_{q-1} is finitely generated. Continuing inductively, we conclude that the group A generated by all the conjugates in (14.20), is finitely generated. Hence, G is polycyclic. \square

This also concludes the proof of Milnor's theorem, Theorem 14.31. \square
By combining theorems of Milnor and Wolf we obtain:

THEOREM 14.37. *Every finitely generated solvable group either is virtually nilpotent or it has exponential growth.*

This was strengthened by J. Rosenblatt as follows:

THEOREM 14.38 (J. Rosenblatt, [Ros74]). *Every finitely generated solvable group either is virtually nilpotent or it contains a free non-abelian subsemigroup.*

Another application of Lemma 14.32 is the following proposition which will be used in the proof of Gromov's theorem on groups of polynomial growth:

PROPOSITION 14.39. *Suppose that G is a finitely generated group of sub-exponential growth, which fits in a short exact sequence*

$$1 \rightarrow K \rightarrow G \xrightarrow{\pi} \mathbb{Z} \rightarrow 1.$$

Then K is finitely generated. Moreover, if $\mathfrak{G}_G(R) \preceq R^d$ then $\mathfrak{G}_K(R) \preceq R^{d-1}$.

PROOF. Let $\gamma \in G$ be an element which projects to the generator 1 of \mathbb{Z} . Let $\{f_1, \dots, f_k\}$ denote a set of generators of G . Then for each i there exists $s_i \in \mathbb{Z}$ such for $g_i := f_i \gamma^{s_i}$, $i = 1, \dots, k$, we have

$$\pi(g_i) = 0 \in \mathbb{Z}.$$

Clearly, the set $\{g_1, \dots, g_k, \gamma\}$ generates G . Without loss of generality we may assume that each generator g_i is non-trivial. Define

$$S := \{\gamma_{m,i} := \gamma^m g_i \gamma^{-m} ; m \in \mathbb{Z}, i = 1, \dots, k\}.$$

We claim that the (infinite) set S generates K . Indeed, clearly, $S \subset K$. Every $g \in K$ can be written as a word $w = w(g_1, \dots, g_k, \gamma)$ in the letters $g_1^{\pm 1}, \dots, g_k^{\pm 1}, \gamma^{\pm 1}$. We then move all entries of powers of γ in the word w to the end of w by using the identities

$$\gamma^m g_i = \gamma_{m,i} \gamma^m.$$

As the result, we obtain a word $w' = u \gamma^a$ in the alphabet $S \cup S^{-1} \cup \{\gamma, \gamma^{-1}\}$, where u contains only the letters in $S \cup S^{-1}$ and $a \in \mathbb{Z}$. Since g projects to $0 \in \mathbb{Z}$, $a = 0$. Claim follows.

Lemma 14.32 implies that there exists $M(i)$ such that the subgroup K is generated by the finite set

$$\{\gamma_{l,i} ; |l| \leq M(i), i = 1, \dots, k\}.$$

This proves the first assertion of the Proposition.

Now let us prove the second assertion which estimates the growth function of K . Take a finite generating set Y of the subgroup K and set $X := Y \cup \{\gamma\}$, where γ is as above. Then X is a generating set of G . Given $n \in \mathbb{N}$ let $N := \mathfrak{G}_Y(n)$,

where \mathfrak{G}_Y is the growth function of K with respect to the generating set Y . Thus, there exists a subset

$$H := \{h_1, \dots, h_N\} \subset K$$

where $|h_i|_Y \leq n$ and $h_i \neq h_j$ for all $i \neq j$. We obtain a set T of $(2n+1) \cdot N$ pairwise distinct elements

$$h_i \gamma^j, \quad -n \leq j \leq n, \quad i = 1, \dots, N.$$

It is clear that $|h_i \gamma^j|_X \leq 2n$ for each $h_i \gamma^j \in T$. Therefore,

$$n\mathfrak{G}_Y(n) \leq (2n+1)\mathfrak{G}_Y(n) = (2n+1)N \leq \mathfrak{G}_X(2n) \leq C(2n)^d = 2^d C \cdot n^d,$$

for some constant C depending only on X . It follows that

$$\mathfrak{G}_Y(n) \leq 2^d C \cdot n^{d-1} \preceq n^{d-1}. \quad \square$$

14.5. Failure of QI rigidity for solvable groups

THEOREM 14.40 (A. Dyubina–Erschler, [Dyu00]). *The class of finitely generated (virtually) solvable groups is not QI rigid: There are groups quasiisometric to solvable groups, but not virtually solvable.*

PROOF. The groups that will be used in the proof are wreath products $G_A := A \wr C$ of finitely generated groups. Given (finite) generating sets $a_i, i \in I, c_j, j \in J$ of A and C , respectively, we will use the finite set of generators of G_A introduced in Lemma 7.11 and the corresponding Cayley graphs. In what follows we use the multiplicative notation when dealing with wreath products.

Suppose now that A, B are finite groups of the same order, where A is abelian, say, cyclic, and B is a non-solvable group. For instance, we can take A to be the group \mathbb{Z}_{60} and B is the alternating group A_5 (which is a simple nonabelian group of order 60). We declare each non-trivial element of these groups to be a generator. Let C be a finitely generated infinite abelian group, say, \mathbb{Z} , and consider the wreath products $G_A := A \wr C, G_B := B \wr C$. Let $\varphi : A \rightarrow B$ be a bijection sending 1 to 1. Extend this bijection to a map

$$\Phi : G_A \rightarrow G_B, \quad \Phi(f, c) = (\varphi \circ f, c).$$

LEMMA 14.41. Φ extends to an isometry of Cayley graphs.

PROOF. First, the inverse map Φ^{-1} is given by $\Phi(f, c) = (\varphi^{-1} \circ f, c)$. We now check that Φ preserves edges of the Cayley graphs. The group G_A has two types of generators: $(1, c_j)$ and $(\delta_a, 1)$, where $c_j \in X$, a finite generating set of C and $a \in A$ are all non-trivial elements of A . The same holds for the group G_B .

1. Consider the edges connecting $(f(x), c)$ to $(f(xc_j^{-1}), cc_j)$ in $\text{Cayley}(G_A)$. Applying Φ to the vertices of such edges we obtain

$$(\varphi \circ f(x), c), \quad (\varphi \circ f(xc_j^{-1}), cc_j).$$

Clearly, they are again within unit distance in $\text{Cayley}(G_B)$, since they differ by $(1, c_j)$.

2. Consider the edges connecting $(f(x), c), (f(x)\delta_a(x), c)$ in $\text{Cayley}(G_A)$. Applying Φ to the vertices we obtain

$$(\varphi \circ f(x), c), \quad (\varphi \circ f(x)\delta_b(x), c),$$

where $b = \varphi(a)$. Again, we obtain vertices which differ by $(\delta_b, 1)$, so they are within unit distance in $\text{Cayley}(G_B)$ as well. \square

LEMMA 14.42. *The group G_B is not virtually solvable.*

PROOF. Let $\psi : G_B \rightarrow F$ be a homomorphism to a finite group. The restriction of ψ to the subgroup $B_c < \oplus_C B$ consisting of maps $f : C \rightarrow B$ with support $\{c\}$, is determined by a homomorphism $\psi_c : B \rightarrow F$. There are only finitely many such homomorphisms, while C is an infinite group. Thus, we find $c_1 \neq c_2 \in C$ such that

$$\psi_{c_1} = \psi_{c_2} = \eta.$$

The kernel of the restriction $\psi|_{B_{c_1} \oplus B_{c_2}}$ consists of pairs

$$(b_1, b_2) \in B_{c_1} \oplus B_{c_2} = B \times B, \quad \eta(b_1) = \eta(b_2)^{-1},$$

and contains the subgroup

$$\{(b, b^{-1}), b \in B\}.$$

The latter this subgroup is isomorphic to B (via the projection to the first factor in the product $B \times B$). Thus, kernel of ψ contains a subgroup isomorphic to B and, therefore, is not solvable. We conclude that G is not virtually solvable. \square

The combination of these two lemmas implies the theorem. \square

14.6. Virtually nilpotent subgroups of $GL(n)$

In this section we collect various properties about virtually nilpotent subgroups of $GL(n, \mathbb{K})$ for arbitrary fields \mathbb{K} , the main focus of which is to show that under certain conditions, a subgroup of $GL(n, \mathbb{K})$ is nilpotent or virtually nilpotent. These results will be used in the proof of the Tits Alternative in the next chapter. In what follows, V will denote a finite-dimensional vector space over a field \mathbb{K} ; we will also use the notation W for a finite-dimensional vector space over the algebraic closure $\bar{\mathbb{K}}$ of \mathbb{K} . We let $End(V)$ denote the algebra of (linear) endomorphisms of V and $GL(V)$ the group of invertible endomorphisms of V .

In what follows we will repeatedly use Burnside's Theorem on irreducible linear actions of semigroups (Theorem 5.44) and Lemma 5.43. Kolchin's Theorem below and its generalizations are useful applications of Burnside's Theorem.

Let V be an n -dimensional vector space over a field \mathbb{K} . Recall that an endomorphism $h \in End(V)$ of V is *nilpotent* if $h^k = 0$ for some $k > 0$. Equivalently, in some basis, h can be written as an upper triangular matrix with zeroes on the diagonal. Automorphisms of V of the form $I + h$, with h nilpotent, are called *unipotent*. Here and in what follows, I is the identity map $V \rightarrow V$. An automorphism g of V is called *quasiunipotent* if all the eigenvalues of g are roots of unity in $\bar{\mathbb{K}}$. Equivalently, g is quasiunipotent if g^k is unipotent for some $k > 0$. A subgroup $G < GL(V)$ is *unipotent* (respectively, *quasiunipotent*) if every element of G is unipotent (respectively, quasiunipotent).

THEOREM 14.43 (Kolchin's theorem). *Suppose that $\mathbb{K} = \bar{\mathbb{K}}$ and $G < GL(V)$ is a unipotent subgroup. Then G is conjugate to a subgroup of the group of invertible upper-triangular matrices $\mathcal{T}_n(\mathbb{K})$.*

PROOF. The proof is by induction on the dimension n of V . The claim is clear for $n = 1$, hence, we assume that $n > 1$. The statement of the theorem amounts to the claim that G preserves a full flag

$$0 \subset V_1 \subset \dots \subset V_{n-1} \subset V,$$

where $i = \dim(V_i)$ for each i . Indeed, given such a flag, we will inductively pick basis elements $\mathbf{e}_i \in V_i$ such that $\{\mathbf{e}_1, \dots, \mathbf{e}_i\}$ is a basis in V_i . With respect to this basis, each subgroup of $GL(V)$ preserving the flag will be contained in $\mathcal{T}_n(\mathbb{K})$.

Suppose first that the action of G on V is reducible, that is G preserves a proper subspace $V' \subset V$. Then we obtain two induced actions of G on V' (by restriction) and on $V'' = V/V'$ (by projection). Since both actions preserve full flags in V', V'' (by the induction hypothesis), the combination of these flags yields a full G -invariant flag in V .

Therefore, we will assume that the action of G on V is irreducible. For each $g \in G$ the endomorphism $g' = g - I$ is nilpotent, hence, has zero trace. Therefore, for all $x \in G$, we have

$$\text{tr}(g'x) = \text{tr}(gx - x) = \text{tr}(I) - \text{tr}(I) = 0.$$

Since, by Burnside's theorem, G spans $\text{End}(V)$, we conclude that for each $x \in \text{End}(V)$ and each $g \in G$,

$$\text{tr}(g'x) = 0.$$

Using the fact that τ is a nondegenerate pairing on $\text{End}(V)$, we conclude that $g' = 0$ for all $g \in G$, i.e. $G = \{1\}$. \square

The following theorem is a minor variation on Kolchin's theorem.

PROPOSITION 14.44. *Suppose that $\mathbb{K} = \bar{\mathbb{K}}$, $G < GL(V)$ is quasiunipotent and, moreover, there exists an upper bound α on the orders of all eigenvalues of elements $g \in G$. Then G contains a finite index subgroup conjugate into the group of upper triangular matrices $\mathcal{T}_n(\mathbb{K})$. The index depends only on V and on α .*

PROOF. The proof follows closely the proof of Kolchin's Theorem. As in Kolchin's Theorem, the proof is by induction on the dimension of V and it suffices to consider the case of subgroups acting irreducibly on V . For each $g \in G$ define a linear map

$$T_g : \text{End}(V) \rightarrow \mathbb{K}, \quad T_g(x) = \text{tr}(gx).$$

Since $\tau, \tau(A, B) = \text{tr}(AB)$, is a nondegenerate pairing on $\text{End}(V)$ (see Lemma 5.43), for $g_1 \neq g_2 \in G$, we get $T_{g_1} \neq T_{g_2}$. As we assumed that the orders of the eigenvalues of elements of G are uniformly bounded, the set of traces of the elements of G is finite. Therefore, for each $g \in G$, the set

$$\{T_g(x) : x \in G\}$$

is a certain finite set $C \subset \mathbb{K}$, independent of g . By Burnside's Theorem, G spans the algebra $\text{End}(V)$, which implies that for each $g \in G$, the map T_g is determined by its restriction to G . Thus, the set

$$\{T_g : \text{End}(V) \rightarrow \mathbb{K} | g \in G\}$$

is finite. We, therefore, conclude that the group G is finite. \square

Suppose that $G < \mathcal{T}_n(\mathbb{K})$ is quasiunipotent with an upper bound on the orders of the eigenvalues. Then there exists $k > 0$ such that g^k is unipotent for each $g \in G$. Therefore, G contains a finite index subgroup G_1 contained in $\mathcal{U}_n(\mathbb{K})$. Since (see Example 13.36) the group $\mathcal{U}_n(\mathbb{K})$ is nilpotent, we obtain:

COROLLARY 14.45. *Suppose that $G < GL(V)$ is quasiunipotent and, moreover, there exists an upper bound α on the orders of all the eigenvalues of elements $g \in G$. Then G is virtually nilpotent. Moreover, the index of the nilpotent subgroup in G depends only on V and α .*

Restriction of scalars. Let the field \mathbb{F} be a finite extension of a field \mathbb{E} ; in other words, \mathbb{F} is a k -dimensional vector space over \mathbb{E} , where $k < \infty$. Thus, we obtain an isomorphism of abelian groups

$$\mathbb{F} \rightarrow \mathbb{E}^k.$$

Accordingly, we obtain a monomorphism

$$GL(n, \mathbb{F}) \hookrightarrow GL(nk, \mathbb{E}).$$

This construction, embedding $GL(n, \mathbb{F})$ into $G \hookrightarrow GL(nk, \mathbb{E})$ is called the *restriction of scalars*.

Restricting to finitely generated subgroups $G < GL(V)$ allows one to eliminate the upper bound assumption on the orders of eigenvalues. Since G is finitely generated, there exists a finitely generated field \mathbb{K} , such that $G < GL(n, \mathbb{K})$.

PROPOSITION 14.46. *Suppose that \mathbb{K} is a finitely generated field and $V = \mathbb{K}^n$. Then each quasiunipotent subgroup $G < GL(V)$ contains a finite index subgroup G' conjugate to a subgroup of $U_n(\mathbb{K})$. Furthermore, G contains such a nilpotent subgroup $G' < G$ of index $\leq q = q(V)$, which is independent of G .*

PROOF. Let $\mathbb{P} \subset \mathbb{K}$ be the prime field. Since \mathbb{K} is finitely generated, we can find inclusions

$$\mathbb{P} \subset \mathbb{P}(T) \subset \mathbb{K},$$

where $\mathbb{P} \subset \mathbb{P}(T)$ is a purely transcendental extension with finite basis T , and $\mathbb{P}(T) \subset \mathbb{K}$ is a finite algebraic extension. By applying the *restriction of scalars* procedure, we re-embed the subgroup $G < GL(n, \mathbb{K})$ into the group $GL(nd, \mathbb{P}(T))$, where $d = |\mathbb{K} : \mathbb{P}(T)|$. We leave it to the reader to verify that the image of the new embedding is still quasiunipotent. We claim that the set of orders of eigenvalues of elements $g \in G < GL(nd, \mathbb{P}(T))$ is finite.

Case 1: \mathbb{K} has zero characteristic, i.e. $\mathbb{P} \cong \mathbb{Q}$. Let $\chi_g(x)$ denote the characteristic polynomial of an element $g \in G < GL(nd, \mathbb{P}(T))$; its roots are roots of unity. Since the coefficients of $\chi_g(x)$ are symmetric polynomials of the roots of unity, it follows that all the coefficients of χ_g are algebraic integers. However, the extension $\mathbb{Q} \subset \mathbb{Q}(T)$ is purely transcendental, therefore, all the coefficients of χ_g actually belong to \mathbb{Z} . As in the proof of Lemma 13.28, the set of coefficients of $\chi_g, g \in G$, is bounded, therefore, it is finite. This means that the orders of eigenvalues of the elements of G are uniformly bounded.

Case 2: \mathbb{K} has characteristic p , i.e. $\mathbb{P} = \mathbb{Z}_p$ for some p . The argument is even simpler than the one in the case of zero characteristic. The coefficients of the characteristic polynomials of $g \in G < GL(nd, \mathbb{P}(T))$ all belong to \mathbb{P} . However, the field \mathbb{P} is finite, therefore, the set of eigenvalues of the elements $g \in G$ is finite.

Now, the first claim of the proposition follows from Proposition 14.44. To verify the second claim we notice that the bound that we obtained on the orders of the eigenvalues of $g \in G$ depends only on n and \mathbb{K} , and we again use Proposition 14.44. \square

14.7. Discreteness and nilpotence in Lie groups

The goal of this section is to prove the theorems of Zassenhaus and Jordan. These theorems deal, respectively, with discrete subgroups Γ of Lie groups G (with finitely many components). The Theorem of Zassenhaus shows that, appropriately defined, “small elements” of Γ generate a nilpotent subgroup of G . Jordan’s theorem establishes that finite subgroups of G are “almost abelian”: Every finite group Γ contains an abelian subgroup, whose index in Γ is uniformly bounded. Historically, Jordan’s theorem was proven first and Zassenhaus proved his theorem afterwards. We will provide proofs in the reverse order, and we will be using Zassenhaus’ results in order to prove Jordan’s theorem.

14.7.1. Some useful linear algebra. We begin by discussing some basic linear algebra to be used in the proof of Jordan’s theorem.

Suppose that V is a real inner product vector space (e.g., a Hilbert space, but we do not insist on the completeness of the norm), with the inner product denoted by $\mathbf{x} \cdot \mathbf{y}$ and the norm denoted by $\|x\|$. We will endow the complexification $V^{\mathbb{C}}$ of V with the inner product

$$(\mathbf{x} + i\mathbf{y}) \cdot (\mathbf{u} + i\mathbf{v}) = \mathbf{x} \cdot \mathbf{u} + \mathbf{y} \cdot \mathbf{v}.$$

Recall that the *operator norm* of a bounded linear transformation $A \in \text{End}(V)$ is

$$\|A\| := \sup_{\|\mathbf{x}\|=1} \|A\mathbf{x}\|.$$

Then A extends naturally to a complex-linear transformation of $V^{\mathbb{C}}$ whose operator norm is $\leq \sqrt{2}\|A\|$. In what follows we will use the notation $\nu(A) = \|A - I\|$ for the distance from A to the identity in $\text{End}(V)$.

The set of automorphisms A of V preserving the inner product, i.e. $A\mathbf{x} \cdot A\mathbf{y} = \mathbf{x} \cdot \mathbf{y}$, is denoted by $O(V)$, and its elements are called *orthogonal transformations*.

LEMMA 14.47. *Suppose that $A \in O(V)$ and $\nu(A) < \sqrt{2}$. Then A is a rotation with rotation angles $< \pi/2$, i.e. for every non-zero vector $\mathbf{x} \in V$,*

$$A\mathbf{x} \cdot \mathbf{x} > 0.$$

PROOF. By the assumption,

$$\|A\mathbf{x} - \mathbf{x}\| < \sqrt{2}$$

for all unit vectors $\mathbf{x} \in V$. Denoting by \mathbf{y} the difference vector $A\mathbf{x} - \mathbf{x}$, we obtain:

$$2 > \mathbf{y} \cdot \mathbf{y} = (A\mathbf{x} - \mathbf{x}) \cdot (A\mathbf{x} - \mathbf{x}) = 2 - 2(A\mathbf{x} \cdot \mathbf{x}).$$

Hence, $A\mathbf{x} \cdot \mathbf{x} > 0$. □

COROLLARY 14.48. *The same conclusion holds for the complexification of A .*

PROOF. Let $\mathbf{v} = \mathbf{x} + i\mathbf{y} \in V^{\mathbb{C}}$. Then

$$A\mathbf{v} \cdot \mathbf{v} = (A\mathbf{x} + iA\mathbf{y}) \cdot (\mathbf{x} + i\mathbf{y}) = A\mathbf{x} \cdot \mathbf{x} + A\mathbf{y} \cdot \mathbf{y} > 0$$

by the above lemma. □

LEMMA 14.49. *Suppose that $A, B \in O(V)$ and $\nu(B) < \sqrt{2}$. Then*

$$[A, BAB^{-1}] = 1 \iff [A, B] = 1.$$

PROOF. Since one implication is clear, we assume that $[A, BAB^{-1}] = 1$. Let λ_j 's be the (complex) eigenvalues of A . Then the complexification $V^{\mathbb{C}}$ splits as an A -invariant orthogonal sum

$$\bigoplus_j V^{\lambda_j},$$

where on each $V_j = V^{\lambda_j} = V_A^{\lambda_j}$ the orthogonal transformation A acts *via* multiplication by λ_j . Here we assume that for $j \neq k$, $\lambda_j \neq \lambda_k$. We refer to this orthogonal decomposition of $V^{\mathbb{C}}$ as \mathcal{F}_A . Then, clearly,

$$B(\mathcal{F}_A) = \mathcal{F}_{BAB^{-1}}$$

for any two orthogonal transformations $A, B \in O(V)$. Since A commutes with BAB^{-1} , A has to preserve the decomposition $\mathcal{F}_{BAB^{-1}}$ and, moreover, has to send each $W_j := V_{BAB^{-1}}^{\lambda_j} = B(V^{\lambda_j})$ to itself. What are the eigenvalues for this action of A on W_j ? They are λ_k 's for which V^{λ_k} has non-trivial intersection with W_j . However, if $\lambda_j \neq \lambda_k$ then V^{λ_j} is orthogonal to V^{λ_k} and, hence, by Corollary 14.48, B cannot send a (non-zero) vector from one space to the other. Therefore, in this case, $W_j \cap V_k = 0$. This leaves us with only one choice of the eigenvalue for the restriction $A|_{W_j}$, namely λ_j . (Since the restriction has to have some eigenvalues!) Thus, $W_j \subset V_j$. However, B sends V_j to W_j injectively, so $W_j = V_j$ and we conclude that $B(V_j) = V_j$. Since A acts on V_j *via* multiplication by λ_j , it follows that $B|_{V_j}$ commutes with $A|_{V_j}$. This holds for all j , hence, $[A, B] = 1$. \square

14.7.2. Zassenhaus neighborhoods. We now define ‘smallness’ in a Lie group: “Small” elements will be those which belong to a Zassenhaus neighborhood defined below.

DEFINITION 14.50. Let G be a topological group. A *Zassenhaus neighborhood* in G is an (open) neighborhood of the identity in G , denoted U or U_G , which satisfies the following:

1. The commutator map sends $U \times U$ to U .
2. There exists a continuous function $\sigma : U \rightarrow \mathbb{R}$ such that $0 = \sigma(1)$ is the minimal value of σ and

$$\sigma([A, B]) < \min(\sigma(A), \sigma(B))$$

for all $A \neq 1, B \neq 1$ in U .

We will refer to σ as a *Zassenhaus function*.

Note that if $H < G$ is a topological subgroup and U_G is a Zassenhaus neighborhood of G then $U_H := U_G \cap H$ is a Zassenhaus neighborhood of H .

We will see that every Lie group has Zassenhaus neighborhoods. We start with some examples. Let V be a Hilbert space (the reader can think of finite-dimensional V since this is the only case that we will need). We let $End(V)$ denote the semigroup of bounded linear operators in V and let $GL(V) \subset End(V)$ be the general linear group of V , i.e. the group of invertible operators A such that both A and A^{-1} are bounded. We again equip $End(V)$ with the operator norm.

LEMMA 14.51. *Let $G = O(V)$ be the orthogonal group of V . Then the set U given by the inequality $\nu(A) < 1/4$ is a Zassenhaus neighborhood in G .*

PROOF. We will use the function $\sigma = \nu$ as the Zassenhaus function. We will show that for all $A, B \in U \setminus 1$,

$$\nu([A, B]) < \min(\nu(A), \nu(B)),$$

which will also imply that $[\cdot, \cdot] : U \times U \rightarrow U$.

First, observe that the multiplication by orthogonal transformations in G preserves the operator norm on $End(V)$. Applying this twice to operators A, B such that $\nu(A) \leq \nu(B)$, we obtain:

$$\begin{aligned} \nu([A, B]) &= \|AB - BA\| = \|(A - B)(A - I) - (A - I)(A - B)\| \leq \\ &\|(A - B)(A - I)\| + \|(A - I)(A - B)\| \leq \\ &2\nu(A)\|A - B\| = 2\nu(A)(\nu(A) + \nu(B)). \end{aligned}$$

Since $\nu(A) \leq \nu(B) < 1/4$, we obtain

$$\nu([A, B]) < 2\nu(A) \left(\frac{1}{4} + \frac{1}{4} \right) = \nu(A). \quad \square$$

LEMMA 14.52. *Let $G = GL(V)$ be the general linear group of V , i.e. group of invertible operators A such that both A and A^{-1} are bounded. We set $\sigma(A) := \max(\nu(A), \nu(A^{-1}))$ for $A \in G$. Then the set U given by the inequality $\sigma(A) < 1/8$ is a Zassenhaus neighborhood in G with the Zassenhaus function σ .*

PROOF. Our proof follows the same lines as in the orthogonal case. We will show that

$$\|ABA^{-1}B^{-1} - I\| < \min(\sigma(A), \sigma(B)).$$

The inequality

$$\|(ABA^{-1}B^{-1})^{-1} - I\| < \min(\sigma(A), \sigma(B))$$

will follow by interchanging A and B . We again assume that $\sigma(A) \leq \sigma(B)$. Observe that $\|CD\| \leq \|C\| \cdot \|D\|$ for all $C, D \in End(V)$. Applying this twice, we get:

$$\|ABA^{-1}B^{-1} - I\| \leq \|B^{-1}\| \|ABA^{-1} - B\| \leq \|A^{-1}\| \|B^{-1}\| \|AB - BA\|.$$

If $\sigma(C) < c$ then $\|C^{-1}\| < 1 + c$ for every $C \in G$. Thus,

$$\|ABA^{-1}B^{-1} - I\| < (1 + c)^2 \|AB - BA\|$$

provided that $\sigma(A) < c, \sigma(B) < c$. The rest of the proof is the same as in the orthogonal case:

$$\|AB - BA\| \leq 2\sigma(A)(\sigma(A) + \sigma(B)) \leq 4\sigma(B)\sigma(A).$$

Putting it all together:

$$\|ABA^{-1}B^{-1} - I\| < 4c(1 + c)^2 \sigma(A).$$

Since for $c = 1/8$, $4c(1 + c)^2 = \frac{1}{2} \left(\frac{9}{8}\right)^2 < 1$, we conclude that

$$\|ABA^{-1}B^{-1} - I\| < \sigma(A).$$

Thus,

$$\sigma([A, B]) < \min(\sigma(A), \sigma(B))$$

for all $A, B \in U$. □

REMARK 14.53. The above proofs, at first glance, look like a trickery. What is really happening in the proof? Consider $G = GL(n, \mathbb{R})$. The idea behind the proof is that the commutator map has zero 1-st derivative at the point $(1, 1) \in G \times G$ (which one can easily see by using the Taylor expansion $A^{-1} = I - a + a^2 \dots$ for a matrix of the form $A = I + a$ where a has small norm). Thus, by the basic calculus, $[A, B]$ will be “closer” to $I \in G$ than $A = I + a$ and $B = I + b$ if a, b are sufficiently small. The above proofs provide explicit estimates for this argument.

We will say that a topological group G admits a basis of Zassenhaus neighborhoods if $1 \in G$ admits a basis of topology consisting of Zassenhaus neighborhoods.

COROLLARY 14.54. *Suppose that G is a linear Lie group. Then $1 \in G$ admits a basis of Zassenhaus neighborhoods.*

PROOF. First, suppose that $G = GL(V)$. Then the sets $U_t = \sigma^{-1}(t), t \in (0, \frac{1}{8})$ are Zassenhaus neighborhoods and their intersection is $1 \in G$. If G is a Lie group which admits a continuous closed embedding $\phi : G \rightarrow GL(V)$, the sets $\phi^{-1}(U_t)$ will serve as a Zassenhaus basis. \square

Note that being a subgroup of $GL(n, \mathbb{R})$ is not really necessary for this corollary since the conclusion is *local* at the identity in G : If a topological group G_1 admits a basis of Zassenhaus neighborhoods and G_2 is a locally compact group which *locally embeds* (see Section 5.6.1) in G_1 , then G_2 also admits a basis of Zassenhaus neighborhoods. In view of Ado’s theorem (Theorem 5.59), every Lie group locally embeds in $GL(V)$ for some finite-dimensional real vector space V .

COROLLARY 14.55. *Every Lie group admits a basis of Zassenhaus neighborhoods.*

Why are Zassenhaus neighborhoods useful? We assume now that G is a locally compact group which admits a basis of Zassenhaus neighborhoods and fix a Zassenhaus neighborhood $\Omega \subset G$ whose closure is compact and is contained in another Zassenhaus neighborhood $U \subset G$. Define inductively subsets $\Omega^{(i)}$ as $\Omega^{(i+1)} = [\Omega, \Omega^{(i)}], \Omega^{(0)} := \Omega$. Since Ω is Zassenhaus,

$$\Omega^{(i+1)} \subset \Omega^{(i)}$$

for all i .

LEMMA 14.56. $E := \bigcap_i \overline{\Omega^{(i)}} = \{1\}$.

PROOF. Clearly, E is compact and $[\Omega, E] = E$. Suppose that $E \neq \{1\}$ and take $f \in E$ with maximal $\sigma(f) > 0$, where σ is the function in the definition of a Zassenhaus neighborhood. **Then** $f = [g, h]$ for some $g \in \Omega, h \in E$ and, hence,

$$\sigma(f) < \min(\sigma(g), \sigma(h)) \leq \sigma(h).$$

Contradiction. \square

THEOREM 14.57 (The Zassenhaus theorem). *Suppose that G is a locally compact group which admits a relatively compact Zassenhaus neighborhood Ω . Assume that $\Gamma < G$ is a discrete subgroup generated by the subset $X := \Gamma \cap \Omega$. Then Γ is nilpotent. In particular, this property holds for all Lie groups.*

PROOF. In view of Lemma 13.44, it suffices to show that there exists n such that all n -fold iterated commutators of elements of X are trivial. By the definition of

$\Omega^{(i)}$, all i -fold iterated commutators of X are contained in $\Omega^{(i)}$. Since Γ is discrete and Ω is relatively compact, we can have only finitely many distinct non-trivial iterated commutators of elements of X . Since

$$\bigcap_i \overline{\Omega^{(i)}} = \{1\},$$

there exists n such that $\Omega^{(n)}$ is disjoint from all these non-trivial commutators. Thus, all n -fold iterated commutators of the elements of X are trivial. Hence, by Lemma 13.55, the group Γ is nilpotent. \square

In section 14.7.3 we will see how the Zassenhaus Theorem can be strengthened in the context of *finite subgroups* of Lie groups (Jordan's theorem).

Below is an application of the Zassenhaus Theorem to orthogonal groups. We equip a finite-dimensional real vector space V with a Euclidean structure and let $O(V)$ denote the group of orthogonal transformations. Recall that for $A \in O(V)$, $\nu(A)$ is the operator norm $\|A - I\|$.

Let U denote a Zassenhaus neighborhood of the identity in $O(V)$ such that

$$U \subset \{A : \nu(A) < \sqrt{2}\}.$$

For instance, in view of Lemma 14.51, we can take U given by the inequality $\nu(A) < 1/4$.

LEMMA 14.58. *Every nilpotent subgroup $G < O(V)$ generated by a subset $U' \subset U$, is abelian.*

PROOF. Consider the lower central series of G :

$$C^1G = G, \dots, G^i = [G, C^{i-1}G], i = 1, \dots, n,$$

where $C^{n+1}G = \{1\}$ and $C^nG \neq \{1\}$. We need to show that $n = 1$. Assume that $n \geq 2$ and consider an $(n + 1)$ -fold iterated commutator of generators of G ; this iterated commutator has the form:

$$[[B, A], A] \in C^{n+1}G = 1$$

where $A \in U'$ and $B \in C^nG$ is an n -fold iterated commutator of elements of U' . Thus, A commutes with $BAB^{-1}A^{-1}$. Since A commutes with A^{-1} , we then conclude that A commutes with BAB^{-1} . By the definition of a Zassenhaus neighborhood, all the iterated commutators of elements in U are also in U , hence B satisfies the inequality $\|B - I\| < \sqrt{2}$.

Lemma 14.49 implies that A commutes with B , hence every n -fold iterated commutator of generators of G is trivial. Thus, $C^nG = \{1\}$ by Lemma 13.44. Contradiction. \square

14.7.3. Jordan's theorem. Notice that even the group $SO(2)$ contains finite subgroups of arbitrarily high order, but these subgroups, of course, are all abelian. Considering the group $O(2)$ we find some non-abelian subgroups of arbitrarily high order, but they all contain abelian subgroups of index 2. Jordan's theorem below shows that a similar statement holds for finite subgroups of other Lie groups as well:

THEOREM 14.59 (Jordan's theorem). *Let L be a Lie group with finitely many connected components. Then there exists a number $q = q(L)$ such that each finite subgroup G in L contains an abelian subgroup of index $\leq q$.*

PROOF. If L is not connected, we replace L with L_0 , the identity component of L and G with $G_0 := G \cap L_0$. Since $|G : G_0| \leq |L : L_0|$, it suffices to prove theorem only for subgroups of connected Lie groups. Thus, we assume in what follows that L is connected.

1. We first consider the most interesting case, when the Lie group L is $K = O(V)$, the orthogonal group of a finite-dimensional Euclidean vector space V .

Let Ω denote a relatively compact Zassenhaus neighborhood of $O(V)$ given by the inequality

$$\{A : \nu(A) < 1/4\}.$$

The finite subgroup $G \subset K$ is clearly discrete, therefore the subgroup $H := \langle G \cap \Omega \rangle$ is nilpotent by the Zassenhaus Theorem. By Lemma 14.58, every nilpotent subgroup generated by elements of Ω is abelian. It, therefore, follows that H is abelian.

It remains to estimate the index $|G : H|$. Let $U \subset \Omega$ be a neighborhood of 1 in $K = O(V)$ such that $U \cdot U^{-1} \subset \Omega$ (i.e. products of pairs of elements xy^{-1} , $x, y \in U$, belong to Ω). Let q denote $\text{Vol}(K)/\text{Vol}(U)$, where Vol is the volume induced by a bi-invariant Riemannian metric on K .

LEMMA 14.60. $|G : H| \leq q$.

PROOF. Let $x_1, \dots, x_{q+1} \in G$ be distinct coset representatives for G/H . Then

$$\sum_{i=1}^{q+1} \text{Vol}(x_i U) = (q+1)\text{Vol}(U) > \text{Vol}(K).$$

Hence there are $i \neq j$ such that $x_i U \cap x_j U \neq \emptyset$. It follows that $x_j^{-1} x_i \in U U^{-1} \subset \Omega$ and, hence, $x_j^{-1} x_i \in H$. Contradiction. \square

This proves Jordan's theorem for subgroups of $O(V)$.

2. Consider now the case when either G or L has trivial center. Consider the adjoint representation $\rho : L \rightarrow GL(V)$, where V is the Lie algebra of L . This representation need not be faithful, but the kernel C of this representation is contained in the center of L , see Lemma 5.52. Hence, the kernel C of the homomorphism

$$\rho : G \rightarrow \bar{G} \leq GL(V)$$

is contained in the center of G . Under our assumptions, C is the trivial group and, hence, $G \cong \bar{G}$.

Next, we construct a G -invariant Euclidean metric on V : Start with an arbitrary positive-definite quadratic form μ_0 on V and then set

$$\mu := \sum_{g \in G} g^*(\mu_0).$$

It is clear that the quadratic form μ is again positive definite and invariant under \bar{G} . With respect to this quadratic form, $\rho(G) \leq O(V)$. Thus, by Part 1, there exists an abelian subgroup $A := H < G$ of index $\leq q(O(V))$, where q depends only on L .

3. We now consider the general case where we can no longer use elementary arguments. First, by Cartan–Iwasawa–Mal'cev theorem (Theorem 5.36), every connected locally compact contains unique, up to conjugation, maximal compact subgroup. We find such subgroup K in L . By Cartan's theorem (Theorem 5.57), every closed subgroup of a Lie group is again a Lie group. Hence, K is also a Lie group.

Since $G < L$ is finite, it is compact, and, thus, is conjugate to a subgroup of K . Next, every compact Lie group is linear, Theorem 5.58. Thus, we can assume that K is contained in $GL(V)$ for some finite-dimensional vector space V .

Note that the choice of V depends only on K , hence, on L . Now, we proceed as in part 2. Note that the requirement that either G or L has trivial center was used in part 2 only to ensure that G has a faithful representation in some $GL(V)$. This proves Jordan's theorem. \square

REMARK 14.61. The above proof does not provide an explicit bound on $q(L)$. Boris Weisfeiler [Wei84] proved that for $n > 63$, $q(GL(n, \mathbb{C})) \leq (n + 2)!$, which is nearly optimal since $GL(n, \mathbb{C})$ contains the permutation group S_n which has the order $n!$. Weisfeiler obtained this result in 1984, shortly before he, tragically, disappeared in Chile in 1985. On August 21 of 2012 a Chilean judge ordered the arrest of eight retired police and military officers in connection with the kidnapping and disappearance of Weisfeiler. According to the court filings, the suspects were to have been prosecuted for "aggravated kidnapping" and "complicity" in the disappearance of a U.S. citizen between January 3–5, 1985; the filings did not mention where Weisfeiler might have been taken after his detention or what may have happened to him afterwards. The case was closed in 2016, after the judge ruled the disappearance a common crime, for which the statute of limitations had passed, and not a violation of human rights. This ruling is currently being appealed.

14.8. Virtually solvable subgroups of $GL(n, \mathbb{C})$

In this section we prove an analogue of Jordan's Theorem for virtually solvable subgroups of $GL(n, \mathbb{C})$. This material will be needed only for the proof of Tits' Alternative for infinitely generated subgroups of $GL(n, \mathbb{C})$; the reader not interested in infinitely generated groups can skip it.

Throughout this section \mathbb{K} will be an algebraically closed field and $V \simeq \mathbb{K}^n$, an n -dimensional vector space over \mathbb{K} . At some point, we will restrict to $\mathbb{K} = \mathbb{C}$, but for most of the discussion, this is not needed.

Let G be a subgroup of $GL(V)$. We will think of V as a G -module, or, more precisely, as a $\mathbb{K}G$ -module. In particular, we will talk about G -submodules and quotient modules: The former are G -invariant subspaces W of V , the latter are quotients V/W , where W is a G -submodule. Recall that the G -module V is reducible if there exists a proper G -submodule $W \subset V$. We say that a subgroup $G < GL(V)$ is *triangular* (or the G -module V is *triangular*) if there exists a *complete flag*

$$\mathcal{F} = (0 \subset V_1 \subset \dots \subset V_n = V)$$

of G -submodules in V , where $\dim(V_i) = i$ for each i . Of course, reducibility makes sense only for modules of dimension > 1 ; however, by abusing the terminology, we will regard modules of dimension ≤ 1 as reducible by default.

DEFINITION 14.62. A subgroup $B < GL(V)$ is called a *Borel subgroup* if B is the stabilizer of a complete flag \mathcal{F} in V .

In other words, Borel subgroups of $GL(V)$ are maximal triangular subgroups of $GL(V)$.

LEMMA 14.63. *Suppose that G is an abstract group such that every $\mathbb{K}G$ -module of dimension $\leq n$ is reducible. Then every n -dimensional $\mathbb{K}G$ -module V is triangular.*

PROOF. Since $G \curvearrowright V$ is reducible, there exists a proper submodule $W \subset V$. Thus $\dim(W) < n$ and $\dim(V/W) < n$. The assertion follows by induction on the dimension. \square

LEMMA 14.64. *Suppose that $G < GL(V)$ is triangular, where $\dim(V) = n$. Then the fixed-point set $\text{Fix}(G)$ of the action of G on the projective space $P(V)$ is non-empty and consists of a disjoint union of projective subspaces $P(W_\ell)$, $\ell = 1, \dots, k$, so that the subspaces $W_i \subset V$ are linearly independent, i.e.:*

$$\text{Span}(\{W_1, \dots, W_k\}) = \bigoplus_{\ell=1}^k W_\ell.$$

In particular, $k \leq \dim(V)$.

PROOF. Since G is triangular, there exists a complete flag of G -submodules in V ,

$$0 \subset V_1 \subset \dots \subset V_n = V.$$

In what follows, we fix a basis $\{e_1, \dots, e_n\}$ in V so that $V_i = \text{Span}(\{e_1, \dots, e_i\})$ and identify each $g \in GL(V)$ with its matrix in this basis. We let $a_{ij}(g)$ denote the i, j matrix coefficient of g .

Since G is triangular, the maps $\chi_i : g \rightarrow a_{ii}(g)$, $i = 1, \dots, n$, are homomorphisms $G \rightarrow \mathbb{K}^\times$. Such homomorphisms are called (*multiplicative*) *characters* of G and the (multiplicative) group of characters of G is denoted by $X(G)$. We let $J \subset \{1, \dots, n\}$ be the set of indices j such that

$$a_{ij}(g) = a_{ji}(g) = 0, \forall g \in G, \forall i \neq j.$$

We split the set J into disjoint subsets J_1, \dots, J_m , $m \leq n$, which are preimages of single characters $\chi \in X(G)$ under the map

$$j \in J \mapsto \chi_j \in X(G).$$

Set $W_\ell := \text{Span}(\{e_i, i \in J_\ell\})$. It is clear that G fixes each $P(W_\ell)$ pointwise since each $g \in G$ acts on W_ℓ via the scalar multiplication by $\chi_\ell(g)$. We leave it to the reader to check that

$$\bigcup_{\ell=1}^m P(W_\ell)$$

is the entire fixed-point set $\text{Fix}(G)$. \square

Let V be an n -dimensional vector space over \mathbb{K} . Recall that the field \mathbb{K} is algebraically closed. In what follows, the topology on subgroups of $GL(V)$ is always the Zariski topology, in particular, connectedness always means Zariski-connectedness.

THEOREM 14.65 (A. Borel). *Every connected solvable subgroup $G < GL(V)$ is triangular.*

PROOF. In view of Lemma 14.63, it suffices to prove that the G -module V is reducible. The proof is an induction on the derived length d of G .

We first recall a few facts about eigenvalues of elements of $GL(V)$. Let $Z_{GL(V)}$ denote the center of $GL(V)$, i.e. the group of automorphisms of the form $\mu \cdot \text{Id}_V$, $\mu \in \mathbb{K}^\times$.

Let $g \in GL(V) \setminus Z_{GL(V)}$. Then g has linearly independent eigenspaces $E_{\lambda_j}(g)$, $j = 1, \dots, k$, labeled by the corresponding eigenvalues λ_j , $1 \leq j \leq k$, where $2 \leq k \leq n$. We let $\mathcal{E}(g)$ denote the set of (unlabeled) eigenspaces

$$\{E_{\lambda_j}(g), j = 1, \dots, k\}.$$

Let B_g denote the abelian subgroup of $GL(V)$ generated by g and the center $Z_{GL(V)}$. Then for every $g' \in B_g$, $\mathcal{E}(g') = \mathcal{E}(g)$ (with the different eigenvalues, of course). If $N(B_g)$ denotes the normalizer of B_g in G , then $N(B_g)$ preserves the set $\mathcal{E}(g)$, however, elements of $N(B_g)$ can permute the elements of $\mathcal{E}(g)$. (Note that $N(B_g)$ is, in general, larger than $N(\langle g \rangle)$, the normalizer of $\langle g \rangle$ in G .) Since $\mathcal{E}(g)$ has cardinality $\leq n$, there is a subgroup $N^o = N^o(B_g) < N(B_g)$ of index $\leq n!$ that fixes the set $\mathcal{E}(g)$ elementwise, i.e. every $h \in N^o$ will preserve each $E_{\lambda}(g)$, where λ 's are the eigenvalues of g . Of course, h need not act trivially on $E_{\lambda}(g)$. Since $g \notin Z_G$, this means that there exists a proper N^o -invariant subspace $E_{\lambda}(g) \subset V$.

We next prove several lemmas needed for the proof of Borel's Theorem.

LEMMA 14.66. *Let A be an abelian subgroup of $GL(V)$. Then the A -module V is reducible.*

PROOF. If $A \leq Z_{GL(V)}$, there is nothing to prove. Assume, therefore, that A contains an element $g \notin Z_{GL(V)}$. Since $A \leq N(B_g)$, it follows that A preserves the collection of subspaces $\mathcal{E}(g)$. Since A is abelian, it cannot permute these subspaces. Therefore, A preserves the proper subspace $E_{\lambda_1}(g) \subset V$ and hence $A \curvearrowright V$ is reducible. \square

LEMMA 14.67. *Suppose that $G < GL(V)$ is a connected metabelian group, such that $G' = [G, G] \leq Z_{GL(V)}$. Then the G -module V is reducible.*

PROOF. The proof is analogous to the proof of the previous lemma. If G is contained in the center of $GL(V)$, there is nothing to prove. Pick, therefore, some $g \in G \setminus Z_{GL(V)}$. Since the image of G in $PGL(V)$ is abelian, the group G is contained in $N(B_g)$. Since G is connected, it cannot permute the elements of $\mathcal{E}(g)$. Hence G preserves each $E_{\lambda_i}(g)$. Every subspace $E_{\lambda_i}(g)$ is proper, thus the G -module V is reducible. \square

Similarly, we have:

LEMMA 14.68. *Let $G < GL(V)$ be a metabelian group whose projection to $PGL(V)$ is abelian. Then G contains a reducible subgroup of index $\leq n!$.*

PROOF. We argue as in the proof of the previous lemma, except G may permute the elements of $\mathcal{E}(g)$. However, it will contain an index $\leq n!$ subgroup which preserves each $E_{\lambda_j}(g)$ and the assertion follows. \square

We can now prove Theorem 14.65. Lemma 14.66 proves the theorem for abelian groups, i.e. solvable groups of derived length 1. Suppose the assertion holds for all connected groups of derived length $< d$ and let $G < GL(V)$ be a connected solvable group of derived length d . Then $G' = [G, G]$ has derived length $< d$. Thus, by the induction hypothesis, G' is triangular. By Lemma 14.64, $\text{Fix}(G') \subset P(V)$ is

a non-empty disjoint union of independent projective subspaces $P(W_i), i = 1, \dots, \ell$. Since G' is normal in G , $\text{Fix}(G')$ is invariant under G . Since G is connected, it cannot interchange the components $P(W_i)$ of $\text{Fix}(G)$. Therefore, it has to preserve each $P(W_i)$. If one of the $P(W_i)$'s is a proper projective subspace in $P(V)$, then W_i is G -invariant and hence the G -module V is reducible. Therefore, we assume that $\ell = 1$ and $W_1 = V$, i.e. G' acts trivially on $P(V)$. This means that $G' < Z_{GL(V)}$ is abelian and hence G is 2-step nilpotent. Now, the assertion follows from Lemma 14.67. This concludes the proof of Theorem 14.65. \square

REMARK 14.69. Theorem 14.65 is false for disconnected solvable subgroups of $GL(V)$. Take $n = 2$, let A be the group of diagonal matrices in $SL(2, \mathbb{C})$ and let

$$s = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}.$$

Then s normalizes A and $s^2 \in A$. We let G be the group generated by A and s ; it is isomorphic to the semidirect product of A and \mathbb{Z}_2 . In particular, G is solvable of derived length 2. On the other hand, it is clear that the G -module \mathbb{C}^2 is irreducible.

The following proposition is a converse to Theorem 14.65:

PROPOSITION 14.70. *For $V = \mathbb{K}^n$, each Borel subgroup $B < GL(V)$ is solvable of derived length n . In particular, the derived length of every connected solvable subgroup of $GL_n(\mathbb{K})$ is at most n .*

PROOF. The proof is induction on n . The assertion is clear for $n = 1$ as $GL_1(\mathbb{K}) \simeq \mathbb{K}^\times$ is abelian. Suppose it holds for $n' = n - 1$, we will prove it for n . Let

$$0 = V_0 \subset V_1 \subset \dots \subset V_n = V$$

be a complete flag invariant under B . Let $B^{(i)} := [B^{(i-1)}, B^{(i-1)}], B^{(0)} = B$ be the derived series of B .

Set $W := V/V_1$, let B_W be the image of B in $GL(W)$. The kernel K of the homomorphism $B \rightarrow B_W$ is isomorphic to \mathbb{K}^\times . The group B_W preserves the complete flag

$$0 = W_0 := V_1/V_1 \subset W_1 := V_2/V_1 \subset \dots \subset W = V/V_1.$$

Therefore, by the induction assumption, B_W has derived length $n - 1$. Thus $B^{(n)} := [B^{(n-1)}, B^{(n-1)}] \subset K \simeq \mathbb{K}^\times$. Since \mathbb{K}^\times is commutative, $[B^{(n)}, B^{(n)}] = 0$, i.e. B has derived length n . \square

COROLLARY 14.71. *A connected subgroup of $GL(V)$ is solvable if and only if it is a subgroup of a Borel subgroup $B < G$.*

We now specialize to $\mathbb{K} = \mathbb{C}$. Actually, most of the proof will go through for any algebraically closed field, except for a place where we invoke Jordan's Theorem.

THEOREM 14.72. *There exist functions $\nu(n)$ and $\delta(n)$ such that every virtually solvable subgroup $\Gamma < GL(n, \mathbb{C})$ contains a solvable subgroup Λ of index $\leq \nu(n)$ and of derived length $\leq \delta(n)$.*

PROOF. Let d denote the derived length of solvable subgroup of finite index in Γ . Let $\bar{\Gamma}$ denote the Zariski closure of Γ in $GL(V)$. Then $\bar{\Gamma}$ has only finitely many (Zariski) connected components (see Theorem 5.86).

LEMMA 14.73. *The group $\bar{\Gamma}$ contains a finite-index subgroup which is solvable of derived length at most d .*

PROOF. We will use k -fold iterated commutators

$$\llbracket g_1, \dots, g_{2^k} \rrbracket,$$

as defined in the equation (13.11). Let $\Gamma^o < \Gamma$ denote a solvable subgroup of derived length d and finite index m in Γ ; thus

$$\Gamma = \gamma_1 \Gamma^o \cup \dots \cup \gamma_m \Gamma^o.$$

The group Γ^o satisfies polynomial equations of the form $\llbracket g_1, \dots, g_{2^d} \rrbracket = 1$. Therefore, Γ satisfies the following polynomial equations in the variables g_j :

$$\gamma_i \llbracket g_1, \dots, g_{2^d} \rrbracket = 1, i = 1, \dots, m.$$

Hence, the Zariski closure $\bar{\Gamma}$ of Γ satisfies the same set of polynomial equations. It follows that $\bar{\Gamma}$ contains a subgroup of index m which is solvable of derived length at most d . \square

Let G denote the (Zariski) connected component of the identity of $\bar{\Gamma}$; in particular, G is a normal subgroup of $\bar{\Gamma}$.

LEMMA 14.74. *The group G is solvable of derived length at most n .*

PROOF. Let $H \triangleleft G$ be the maximal solvable subgroup of derived length d of finite index. Thus, as above, H is given by imposing polynomial equations of the form $\llbracket g_1, \dots, g_{2^d} \rrbracket = 1$ on tuples of elements of G . In particular, H is Zariski closed. Since H has finite index in G , it is also open. On the other hand, G is connected, therefore $G = H$, i.e. G is solvable and has derived length $\leq n$ by Proposition 14.70. \square

It is clear that $\Gamma \cap G$ is a finite index subgroup of Γ , and $|\Gamma : \Gamma \cap G| \leq |\bar{\Gamma} : G|$. Unfortunately, the index $|\bar{\Gamma} : G|$ could be arbitrarily large. We will see, however, that we can enlarge G to a (possibly disconnected) subgroup $\hat{G} \leq \bar{\Gamma}$ which is still solvable but has a uniform upper bound on the index $|\bar{\Gamma} : \hat{G}|$ and a uniform bound on the derived length.

We will prove the existence of a bound on the index and the derived length by a dimension induction. The base case where $n = 1$ is clear, so we assume that for each $n' < n$ and each virtually solvable subgroup $\Gamma' < GL_{n'}(\mathbb{C})$ there exists a solvable group \hat{G}'

$$G' \leq \hat{G}' \leq \bar{\Gamma}'$$

as required, with a uniform bound $\nu(n')$ on the index $|\bar{\Gamma}' : \hat{G}'|$ and so that the derived length of \hat{G}' is at most $\delta(n') \leq \delta(n - 1)$.

Let $\mathcal{W} := \{W_1, \dots, W_\ell\}$ denote the maximal collection of (independent) subspaces in V so that G fixes each projective space $P(W_i)$ pointwise (see Theorem 14.65 and Lemma 14.64). In particular, $\ell \leq n$. Since G is normal in $\bar{\Gamma}$, the collection \mathcal{W} is invariant under $\bar{\Gamma}$. Let $K \leq \bar{\Gamma}$ denote the kernel of the action of $\bar{\Gamma}$ on the set \mathcal{W} . Clearly, $G \leq K$ and $|\bar{\Gamma} : K| \leq \ell! \leq n!$. We will, therefore, study the pair of groups $G \leq K$.

REMARK 14.75. Note that we just proved that every virtually solvable subgroup $\Gamma < GL(n, \mathbb{C})$ contains a reducible subgroup of index $\leq n!c(n)$, where $c(n) := q(PGL(n, \mathbb{C}))$ is the function from Jordan's Theorem, Theorem 14.59. Indeed, if $\ell > 1$, the subgroup $K \cap \Gamma$ (of index $\leq n!$) preserves a proper subspace $V_1 \in \mathcal{V}$. If $\ell = 1$, then G is contained in $Z_{GL(V)}$ and hence Γ projects to a finite subgroup $\Phi < PGL(V)$. After replacing Φ with an abelian subgroup A of index $\leq q(PGL(V))$ (using Jordan's Theorem 14.59), we obtain a metabelian group $\tilde{A} < \Gamma$ whose center is contained in $Z_{GL(V)}$. Now the assertion follows from Lemma 14.68.

The group K preserves each W_i and, by construction, the group G acts trivially on each $P(W_i)$. Therefore, the image Q_i of K/G in $PGL(W_i)$ is finite. (The finite group K/G need not act faithfully on $P(W_i)$.) By Jordan's Theorem, the group Q_i contains an abelian subgroup of index $\leq c(\dim(W_i)) \leq c(n)$. Hence, K contains a subgroup $N \triangleleft K$ of index at most

$$\prod_{i=1}^{\ell} c(\dim(W_i)) \leq c(n)^n$$

which acts as an abelian group on

$$\prod_{i=1}^{\ell} P(W_i).$$

We again note that $G \leq N$. The image of the restriction homomorphism $\phi : N \rightarrow GL(U)$,

$$U := W_1 \oplus \dots \oplus W_{\ell}$$

is therefore a *metabelian* group M .

We also have the homomorphism $\psi : N \rightarrow GL(W)$, $W = V/U$ with image N_W . The group N_W contains the connected solvable subgroup $G_W := \psi(G)$ of finite index. To identify the intersection $\text{Ker}(\phi) \cap \text{Ker}(\psi)$ we observe that $V \simeq U \oplus W$ and, *via* this identification, the group N acts by matrices of block-triangular form:

$$\begin{bmatrix} x & y \\ 0 & z \end{bmatrix}$$

where $x \in M$, $z \in N_W$. Then the kernel of the homomorphism $\phi \times \psi : N \rightarrow M \times N_W$ consists of matrices of the upper-triangular form

$$\begin{bmatrix} I & y \\ 0 & I \end{bmatrix}.$$

Thus, by Proposition 14.70, $L = \text{Ker}(\phi \times \psi)$ is solvable of derived length $\leq n$.

By the induction hypothesis, there exists a solvable group \widehat{G}_W of derived length $\leq \delta(n-1)$, such that

$$G_W \leq \widehat{G}_W \leq N_W$$

and $|N_W : \widehat{G}_W| \leq \nu(n-1)$. Therefore, for $\widehat{G} := (\phi \times \psi)^{-1}(M \times \widehat{G}_W)$, we obtain a commutative diagram

$$\begin{array}{ccccccc} 1 & \rightarrow & L & \longrightarrow & N & \xrightarrow{\phi \times \psi} & M \times N_W & \longrightarrow & 1 \\ & & \parallel & & \uparrow \iota' & & \uparrow \iota & & \\ 1 & \rightarrow & L & \longrightarrow & \widehat{G} & \xrightarrow{\phi \times \psi} & M \times \widehat{G}_W & \longrightarrow & 1 \end{array}$$

where ι is the inclusion of an index $i \leq \nu(n-1)$ subgroup and, hence, ι' is also an inclusion of an index i subgroup. Furthermore, L is solvable of derived length $\leq n$, $M \times \widehat{G}_W$ is solvable of derived length $\leq \max(2, \delta(n-1))$.

Putting it all together, we obtain the inequality

$$|\overline{\Gamma} : \widehat{G}| \leq \nu(n) := \nu(n-1)n!(c(n))^n,$$

where \widehat{G} is solvable of derived length $\leq \delta(n) := \max(2, \delta(n-1)) + n$. Intersecting \widehat{G} with Γ we obtain $\Lambda < \Gamma$ of index at most $\nu(n)$ and derived length $\leq \delta(n)$. Theorem 14.72 follows. \square

EXERCISE 14.76. Prove that Lemma 14.74 also holds if we replace Zariski closure with the closure with respect to the standard topology, when $\mathbb{K} = \mathbb{C}$.

In Theorem 14.72 we limited ourselves to the case of the field $\mathbb{K} = \mathbb{C}$. The following results treat the case of general fields (we owe the reference to Yves de Cornulier):

THEOREM 14.77 (A. I. Mal'cev, [Mal51]; see also Theorem 3.21 in [Rob72]). *Let G be a solvable subgroup of $GL(n, \mathbb{K})$, where \mathbb{K} is an algebraically closed field. Then G contains a normal triangular subgroup with finite index not exceeding*

$$\prod_{i=1}^n (i!) (i^2 f(i^2))^i,$$

where $f(i)$ is the maximum number of automorphisms of a group of order $\leq i$.

A corollary of Mal'cev's theorem is a result proven earlier by H. Zassenhaus:

THEOREM 14.78 (H. Zassenhaus [Zas38]; see also Theorem 3.23 in [Rob72]). *For every field \mathbb{F} , the derived length of a solvable subgroup of $GL(n, \mathbb{F})$ is at most $\theta(n)$, where $\theta(n)$ is independent of \mathbb{F} .*

The Tits Alternative

In this chapter we will prove

THEOREM 15.1 (The Tits Alternative, [Tit72]). *Let L be a Lie group with finitely many connected components and let $\Gamma < L$ be a finitely generated subgroup. Then either Γ is virtually solvable or Γ contains a free nonabelian subgroup.*

The main corollary of the Tits Alternative that we will use (in the proof of Gromov's theorem on groups of polynomial growth) is:

COROLLARY 15.2. *Suppose that Γ is a finitely generated subgroup of $GL(n, \mathbb{R})$. Then Γ is either virtually nilpotent or has exponential growth. In particular, Γ has either polynomial or exponential growth.*

PROOF. By the Tits Alternative, either Γ contains a nonabelian free subgroup (and hence Γ has exponential growth) or Γ is virtually solvable. For virtually solvable groups the assertion follows from Theorem 14.37. \square

Note that this corollary requires only a weaker version of the Tits Alternative, namely, existence of free subsemigroups, whose proof is slightly easier than that of the full Tits Alternative.

In view of Corollary 15.2, Milnor's Conjecture (Conjecture 8.86), has affirmative answer for finitely generated subgroups of Lie groups L with finite $\pi_0(L)$. Since the Lie group L in Theorem 15.1 has only finitely many connected components, the intersection Γ_0 of Γ with the identity component $L_0 \leq L$ has finite index in Γ . Therefore, it suffices to consider the case of connected Lie groups G . Consider the adjoint representation $Ad : G \rightarrow GL(\mathfrak{g})$, where \mathfrak{g} is the Lie algebra of G . The kernel of this homomorphism is a central subgroup of Γ (see Lemma 5.52). Since central coextension of a virtually solvable group is again virtually solvable and Γ contains a free nonabelian subgroup if and only if $Ad(\Gamma)$ does, we reduce the problem to the case of subgroups of $GL(n, \mathbb{R})$. As it turns out, even if we are interested in matrix groups with real coefficients, one still has to consider matrix groups over other fields, and we are lead to proving

THEOREM 15.3 (J. Tits). *The Tits alternative holds for finitely generated subgroups $\Gamma < GL(n, \mathbb{K})$, where \mathbb{K} is an arbitrary field. Either Γ is virtually solvable or contains a free nonabelian subgroup.*

Since Γ is finitely generated, we can assume that \mathbb{K} is finitely generated as well: Otherwise, we replace it with the subfield generated (over the prime field of \mathbb{K}) by the matrix entries of the generators of Γ . Moreover, it clearly suffices to consider the case of infinite fields \mathbb{K} , which we will assume from now on.

REMARK 15.4. In Section 15.6, we prove an analogue of Theorem 15.3 in the case of subgroups $\Gamma < GL(n, \mathbb{F})$ which are not required to be finitely generated, but the field \mathbb{F} has zero characteristic.

15.1. Outline of the proof

In this section we give an outline of the proof of the Tits alternative, breaking it into a sequence of theorems, which will be proven later on in this chapter. In what follows, V will denote a finite-dimensional vector space over a field \mathbb{K} whose algebraic closure will be denoted $\bar{\mathbb{K}}$. We let $End(V)$ denote the algebra of (linear) endomorphisms of V and $GL(V)$ the group of invertible endomorphisms of V .

We present a step-by-step outline of the proof of the Tits alternative for finitely generated matrix groups, i.e. finitely generated subgroups $\Gamma < GL(n, \mathbb{K})$, where \mathbb{K} 's is an arbitrary field. The proof is by induction on n . The statement is clear for $n = 1$, hence, we assume that it holds for all $m < n$.

Step 1: Reduction to the case of virtually absolutely irreducible subgroups.

Suppose that Γ contains a finite-index subgroup Γ_1 which preserves a proper subspace of $\bar{\mathbb{K}}^n$, i.e. the action of Γ on $\bar{\mathbb{K}}^n$ is not irreducible. We will identify the proper subspace as above with $\bar{\mathbb{K}}^m$, $1 \leq m < n$.

Then Γ_1 projects to a subgroup $\Gamma_2 < GL(n - m, \bar{\mathbb{K}})$ which, by the induction assumption, is either virtually solvable or contains a nonabelian free subgroup. In the latter case, Γ contains a nonabelian free subgroup as well. In the former case, we consider the kernel $\Gamma_3 \triangleleft \Gamma_1$ of the homomorphism $\Gamma_1 \rightarrow \Gamma_2$. The subgroup Γ_3 is contained in the semidirect product

$$N \rtimes GL(m, \bar{\mathbb{K}})$$

where N is the nilpotent subgroup of $GL(n, \bar{\mathbb{K}})$ consisting of the block-triangular matrices of the form

$$\begin{bmatrix} I_{n-m} & \star \\ 0 & I_m \end{bmatrix}.$$

This subgroup is the kernel of the restriction homomorphism r from the stabilizer of $\bar{\mathbb{K}}^m$ in $GL(n, \bar{\mathbb{K}})$ to the subspace $\bar{\mathbb{K}}^m$. If $r(\Gamma_3)$ contains a free nonabelian subgroup, then so does Γ . Otherwise, by the induction assumption, $r(\Gamma_3)$ is virtually solvable. Since $\Gamma_4 = \Gamma_3 \cap N$ is solvable, we obtain two short exact sequences

$$1 \rightarrow \Gamma_4 \rightarrow \Gamma_3 \rightarrow \Gamma_3/\Gamma_4 \rightarrow 1,$$

$$1 \rightarrow \Gamma_3 \rightarrow \Gamma_1 \rightarrow \Gamma_2 \rightarrow 1.$$

The projection $\Gamma_5 \cong \Gamma_3/\Gamma_4$ of Γ_3 to $GL(m, \bar{\mathbb{K}})$ is virtually solvable by the induction hypothesis; virtual solvability of Γ_4 then implies virtual solvability of Γ_3 . Since Γ_2 is virtually solvable, so is Γ_1 , see Proposition 13.91.

Therefore, it suffices to consider subgroups Γ of $GL(n, \mathbb{K})$ for various fields \mathbb{K} such that each finite-index subgroup of Γ acts absolutely irreducibly on \mathbb{K}^n .

Step 2: Reduction to the case of subgroups of $SL(n, \mathbb{K})$.

Given a finitely generated subgroup $\Gamma < GL(n, \mathbb{K})$ we define the homomorphism

$$\phi : GL(n, \mathbb{K}) \rightarrow SL(n, \mathbb{K}), \quad \phi(g) = g \det(g)^{-1}.$$

The kernel of this homomorphism is abelian, contained in the center of $GL(n, \mathbb{K})$. Therefore, Γ is virtually solvable if and only if $\phi(\Gamma)$ is; similarly, Γ contains a free

nonabelian subgroup if and only if $\phi(\Gamma)$ does. We leave it as an exercise to the reader to prove:

EXERCISE 15.5. Every finite-index subgroup of Γ acts absolutely irreducibly on \mathbb{K}^n if and only if the same is true for $\phi(\Gamma)$.

Therefore, it suffices to analyze the case of finitely generated subgroups of $SL(n, \mathbb{K})$ for various fields \mathbb{K} .

Step 3: Reduction to the case of subgroups which contain *non-distal elements*.

Recall that an element $a \in \mathbb{K}$ of a normed field \mathbb{K} is a *unit* if a has unit norm. An element $g \in GL(n, \mathbb{K})$ is called *distal* if all the eigenvalues of g are units.

Consider a finitely generated virtually absolutely irreducible subgroup $\Gamma < SL(n, \mathbb{F})$, where \mathbb{F} is a finitely generated field. If every eigenvalue of each element of Γ is a root of unity, then Γ is virtually nilpotent by Proposition 14.46. Suppose, therefore, that λ is an eigenvalue of some $\gamma \in \Gamma$, which is not a root of unity. Let \mathbb{E} denote the extension of \mathbb{F} , which is the splitting field of the characteristic polynomial of γ . By Theorem 2.64, there exists an embedding $\mathbb{E} \hookrightarrow (\mathbb{K}, |\cdot|)$ into a normed local field, such that the image of λ in \mathbb{K} is not a unit. Hence, the image of γ under the embedding $\Gamma \hookrightarrow SL(n, \mathbb{K})$, is non-distal.

This reduces the problem to the case of local fields \mathbb{K} and finitely generated subgroups $\Gamma < SL(n, \mathbb{K})$, acting absolutely irreducibly on \mathbb{K}^n , such that at least one element $\gamma \in \Gamma$ is non-distal (cf. Exercise 5.46). The claim is that such Γ contains a free nonabelian subgroup.

Step 4: Finding very proximal elements. Suppose that $\Gamma < SL(n, \mathbb{K})$ is a subgroup satisfying the conclusion of Step 3 and $\gamma \in \Gamma$ is a non-distal element. Not all the norms of the eigenvalues of γ are the same, since, their product equals 1.

We let $\lambda_1, \dots, \lambda_n \in \mathbb{K}$ denote the eigenvalues of γ , ordered so that

$$|\lambda_1| = |\lambda_2| = \dots = |\lambda_k| > |\lambda_{k+1}| \geq \dots \geq |\lambda_n|$$

We now consider the action of $SL(n, \mathbb{K})$ on the exterior product

$$W = \Lambda^k V, \quad V = \mathbb{K}^n.$$

According to Lemma 5.45, the action of G on $W = \Lambda^k V$ is again absolutely irreducible.

By the construction, the action of γ on W is *proximal*: The dominant eigenvalue of γ is λ_1^k , which has multiplicity one (see Section 2.10 for the definition of dominance). If one is interested only in finding free subsemigroups of Γ , proximal elements suffice for constructing semi ping-pong pairs in Γ , see Section 15.3. However, finding free subgroups requires more work:

THEOREM 15.6. Γ contains an element β whose action on W is very proximal β , i.e. both β and β^{-1} are proximal elements of $GL(W)$.

Our arguments in this part of the proof are inspired by the metric geometry arguments in the papers by Breuillard and Gelander, [BG03, BG08b], where they proved an *effective version* of the Tits alternative.

Step 5: Finding ping-pong partners. Suppose that $\Gamma < SL(n, \mathbb{K}) = SL(V)$ is absolutely irreducible and Γ contains an element γ whose action on $W = \Lambda^k V$ is very proximal.

THEOREM 15.7. *The subgroup $\Gamma < GL(W)$ contains a pair of ping-pong partners α, β in the sense of Definition 7.66.*

Now that $\Gamma < GL(W)$ contains a pair of ping-pong partners α, β , there exists $m > 0$ such that the subgroup of Γ generated by α^m, β^m is free of rank 2, see Lemma 7.67. This step will conclude the proof of Theorem 15.3.

Clearly, the proof of Theorem 15.3 is now reduced to Theorems 15.6 and 15.7.

15.2. Separating sets

DEFINITION 15.8. A subset $F \subset PGL(V)$ is called *m-separating* if for every choice of points $p_1, \dots, p_m \in P = P(V)$ and hyperplanes $H_1, \dots, H_m \subset P$, there exists $f \in F$ such that

$$f^{\pm 1}(p_i) \notin H_j, \forall i, j = 1, \dots, m.$$

PROPOSITION 15.9. *Let $\Gamma < GL(V)$ be a subgroup acting irreducibly on V with the property that its Zariski closure is Zariski-irreducible. For every m , Γ contains a finite *m-separating* subset F .*

PROOF. Let $\bar{\Gamma}$ be the Zariski closure of Γ in $GL(V)$. Let P^\vee denote the space of hyperplanes in P (i.e. the projective space of the dual of V). For each $g \in \bar{\Gamma}$ let $M_g \subset P^m \times (P^\vee)^m$ denote the collection of $2m$ -tuples

$$(p_1, \dots, p_m, H_1, \dots, H_m)$$

such that

$$g(p_i) \in H_j \text{ or } g^{-1}(p_i) \in H_j$$

for some $i, j = 1, \dots, m$.

LEMMA 15.10. *If Γ is as in Proposition 15.9 then*

$$\bigcap_{g \in \Gamma} M_g = \emptyset.$$

PROOF. Suppose to the contrary that the intersection is non-empty. Then there exists a $2m$ -tuple $(p_1, \dots, p_m, H_1, \dots, H_m)$ such that for every $g \in \Gamma$,

$$(15.1) \quad \exists i, j \text{ such that } g(p_i) \in H_j \text{ or } g^{-1}(p_i) \in H_j.$$

The set of elements $g \in GL(V)$ such that (15.1) holds for the given $2m$ -tuple is Zariski-closed, and G is the Zariski closure of Γ , hence all $g \in G$ also satisfy (15.1).

Let $G_{p_i, H_j}^\pm \subset \bar{\Gamma}$ denote the subset set consisting of those $g \in \bar{\Gamma}$ for which

$$g^{\pm 1}(p_i) \in H_j.$$

Clearly, these subsets are Zariski-closed and cover the group $\bar{\Gamma}$. Since $\bar{\Gamma}$ Zariski-irreducible, it follows that one of these sets, say G_{p_i, H_j}^+ , is the entire of $\bar{\Gamma}$. Therefore, for every $g \in G$, $g(p_i) \in H_j$. Thus, projectivization of the vector subspace L spanned by the $\bar{\Gamma}$ -orbit (of lines) $\bar{\Gamma} \cdot p_i$ is contained in H_j . The subspace L is proper and G -invariant. This contradicts the hypothesis that $\bar{\Gamma}$ acts irreducibly on V . \square

We now finish the proof of Proposition 15.9. Let M_g^c denote the complement of M_g in $P^m \times (P^\vee)^m$. This set is Zariski open. By Lemma 15.10, the sets M_g^c

$(g \in \Gamma)$ cover the space $P^m \times (P^\vee)^m$. By Lemma 5.74, there exists a finite subset $F \subset \Gamma$ so that

$$\bigcap_{f \in F} M_f = \bigcap_{g \in \Gamma} M_g = \emptyset.$$

In other words,

$$\bigcup_{f \in F} M_f^c = P^m \times (P^\vee)^m.$$

This subset F is the one whose existence is claimed by Proposition 15.9. \square

15.3. Proof of existence of free subsemigroup

In this section we prove a weaker form of the Tits Alternative. The reader who is only interested in the proof of Gromov's theorem on groups of polynomial growth can read this proof instead of the one given in Section 15.4, since it is technically much simpler. We refer the reader to Section 2.10 for the definitions of subspaces $E_\alpha, A_\alpha \subset P(V)$ associated with linear transformations $\alpha \in GL(V)$ and to Definition 7.66 for the notion of *ping-partners*.

THEOREM 15.11. *Let $\Gamma < GL(n, \mathbb{K})$, $n > 1$, be a subgroup which acts virtually irreducibly on $V = \mathbb{K}^n$ and contains a proximal element α . Then Γ contains a free subsemigroup of rank 2.*

PROOF. Let $\bar{\Gamma}$ denote the Zariski closure of Γ in $GL(n, \mathbb{K})$. If $\bar{\Gamma}$ is Zariski-reducible, we replace Γ with the finite-index subgroup $\Gamma_0 < \Gamma$, the intersection of Γ with the identity component of $\bar{\Gamma}$, cf. Proposition 5.92. Thus, we will assume that $\bar{\Gamma}$ is Zariski-irreducible. We claim that Γ contains an element γ such that the elements $\alpha, \beta = \gamma\alpha\gamma^{-1}$ of the group Γ are *ping-partners* (Definition 7.66). Indeed, since Γ contains a 2-separating subset F (see Proposition 15.9), there exists $\gamma \in F$ such that

$$(15.2) \quad \{\gamma(A_\alpha), \gamma^{-1}(A_\alpha)\} \cap E_\alpha = \emptyset.$$

Being a conjugate of α , the element $\beta = \gamma\alpha\gamma^{-1}$ is also proximal and

$$A_\beta = \gamma(A_\alpha), \quad E_\beta = \gamma(E_\alpha).$$

Then (15.2) implies that

$$A_\alpha \notin E_\beta, \quad A_\beta \notin E_\alpha,$$

which means that α, β are ping-partners. Therefore, by Lemma 7.67, there exists $m > 0$ such that the subsemigroup of Γ generated by α^m, β^m is free of rank two. \square

15.4. Existence of very proximal elements: Proof of Theorem 15.6

In what follows, V is an n -dimensional vector space over an (infinite) local field \mathbb{K} , $n = \dim(V) > 1$. We fix a basis e_1, \dots, e_n in V . Then the norm $|\cdot|$ on \mathbb{K} determines the Euclidean norms $\|\cdot\|$ on V .

NOTATION 15.12. We let $P(V)$ denote the projective space of V . When there is no possibility of confusion we do not mention the vector space anymore, and simply denote the projective space by P .

In this section we show how to find very proximal elements in a subgroup $\Gamma < GL(V)$, assuming that Γ contains a proximal element g . We will find such very proximal elements in the form

$$f'g^i f g^{-i},$$

where f, f' belong to some finite 4-separating subset of Γ . While this appears to be a linear algebra problem, we will use geometric and topological arguments instead. To this end, in the following section we will present some *geometric conditions* for proximality, based on the contraction principle. Below, for each non-zero vector $v \in V$ we let $[v] \in P(V)$ denote its projection to the projective space; we will use the metric on $P(V)$ and metric balls $B([v], R) \subset P(V)$, see Section 2.9.

15.4.1. Proximality criteria.

PROPOSITION 15.13 (Proximality criterion-I). *Suppose that $g \in GL(V)$ and u is an eigenvector of g such that for some $R > 0$ the restriction of g to $B([u], R)$ is $\frac{1}{2}$ -Lipschitz. Then g is proximal and $[u] = A_g$.*

PROOF. Suppose that g is not proximal and let $v \in V$ be an eigenvector of g linearly independent of u , such that $|\lambda_v| \geq |\lambda_u|$, where λ_u, λ_v are the eigenvalues for the eigenvectors u, v respectively. The linear transformation g preserves the 2-dimensional subspace $W = \text{Span}(u, v) \subset V$. The inequality $|\lambda_v| \geq |\lambda_u|$ implies that the sequence $g^i, i \in \mathbb{N}$ is either bounded in $GL(W)$ or its projection to $PGL(W)$ converges to $[v]$ uniformly on compacts in $P(W) \setminus \{[u]\}$, cf. Theorem 2.82. However, since the restriction of g to $B([u], R)$ is $\frac{1}{2}$ -Lipschitz, $g(B([u], R)) \subset B([u], R)$ and $B([u], R) \neq \{[u]\}$ since the field \mathbb{K} is infinite and local. Contradiction. \square

PROPOSITION 15.14 (Proximality criterion-II). *Suppose that $g_i \in GL(V)$ is a sequence such that for some $a \in P$ and $R > 0$, the sequence $g_i|_{\bar{B}(a, R)}$ converges uniformly to a . Then for all but finitely many values of i the elements g_i are proximal. Moreover,*

$$\lim_{i \rightarrow \infty} A_{g_i} = a.$$

PROOF. Since the sequence g_i converges to a on the ball $B(a, R)$, it follows that the sequence of projective transformations g_i converges to the quasiconstant map with the image a on the projective space $P(V)$, see Theorem 2.82 and Exercise 2.80. It follows from Lemma 2.83 that

$$\lim_{i \rightarrow \infty} \text{Lip}(g_i|_{B(a, R)}) = 0.$$

Pick a positive number $\epsilon < R$. After discarding finitely many members of the sequence g_i , we can assume that the restrictions of g_i to $\bar{B}(a, R)$ are $\frac{1}{2}$ -Lipschitz and that

$$g_i(\bar{B}(a, R)) \subset B(a, \epsilon).$$

Therefore, for each $g = g_i$ and $m \geq 1$, we have

$$\text{diam}(g^m(\bar{B}(a, R))) \leq 2^{-m}.$$

Since the normed field \mathbb{K} is complete, the sequence of iterates g^m converges on $\bar{B}(a, R)$ to a point $[u]$ contained in $B(a, \epsilon)$, which, therefore, has to be fixed by g . By Proposition 15.13, u is an eigenvector of g with dominant eigenvalue. Hence, g is proximal with $A_g \in B(a, \epsilon)$. This also shows that

$$\lim_{i \rightarrow \infty} A_{g_i} = a. \quad \square$$

PROPOSITION 15.15 (Proximality criterion-III). *Suppose that $g_i \in GL(V)$ is a sequence such that for some $c \in P$ and $R > 0$, the sequence $g_i|_{\bar{B}(c,R)}$ converges uniformly to a point $a \in P(V)$. Assume also that the sequence (g_i) converges to some quasiprojective transformation $\hat{g} \in End(V)$, whose kernel $\text{Ker}_{\hat{g}} \subset P(V)$ does not contain the point a . Then all but finitely many members of the sequence g_i are proximal and*

$$\lim_{i \rightarrow \infty} A_{g_i} = a.$$

PROOF. The proof is similar to that of Proposition 15.14. The sequence (g_i) converges to \hat{g} uniformly on compacts in $P(V) \setminus \text{Ker}_{\hat{g}}$. By Exercise 2.80, \hat{g} is a quasiconstant, since it maps the entire ball $B(c, R)$ to the point a . Since $\text{Ker}_{\hat{g}}$ does not contain a , the sequence of restrictions of g_i 's to a small ball $B(a, \epsilon)$ converges uniformly to a . Now, the claim follows from Proposition 15.14. \square

15.4.2. Constructing very proximal elements. We now prove one of the two main results of this section:

THEOREM 15.16. *Let $F \subset GL(V)$ be a 4-separating subset and let $g \in GL(V)$ be a proximal element. Then there exist $f, f' \in F$ such that for infinitely many i 's, the elements*

$$g_i = f' g^i f g^{-i}$$

are very proximal.

PROOF. We break the proof into two propositions, first ensuring proximality of g_i 's and the second verifying proximality of g_i^{-1} .

PROPOSITION 15.17. *Let $g \in GL(V)$ be a proximal element and $F \subset GL(V)$ be a 2-separating subset. Then for each infinite subset $I \subset \mathbb{N}$ there exist $f, f' \in F$ and an infinite subset $J \subset I$, such that the products $g_i = f' g^i f g^{-i}$, $i \in J$, satisfy:*

1. *Each g_i is proximal.*
- 2.

$$\lim_{i \rightarrow \infty, i \in J} A_{g_i} = f'(A_g).$$

PROOF. Since g is proximal, the sequence $(g^i)_{i \in \mathbb{N}}$ converges to a quasiconstant map

$$\hat{g} : P \setminus E_g \rightarrow A_g.$$

We first consider the sequence g^{-i} , $i \in I$. By the convergence property (Theorem 2.82), this sequence subconverges to some non-invertible quasiprojective transformation $\check{g} \in End(P(V))$. We let $J \subset I$ denote the subset such that

$$\lim_{i \rightarrow \infty, i \in J} g^{-i} = \check{g}.$$

We pick a small ball $\bar{B}(x, \epsilon)$ disjoint from $\text{Ker}(\check{g})$, its image under \check{g} is contained in a small ball $\bar{B}(\check{g}(x), L\epsilon) \subset \text{Im}(\check{g})$, where L is the Lipschitz constant of \check{g} . Since the set F is 2-separating, there exists $f \in F$ such that

$$f(\check{g}(x)) \notin E_g.$$

By using a sufficiently small ϵ , we can assume that $f(\bar{B}(\check{g}(x), L\epsilon))$ is disjoint from E_g as well. We let

$$E_{\hat{g} \circ f \circ \check{g}} = \text{Ker}_{\hat{g} \circ f \circ \check{g}}$$

denote the indeterminacy set of the quasiconstant map $\hat{g} \circ f \circ \check{g}$, whose image is A_g . To be more precise, this set is the hyperplane in P such that the suitable subsequence of

$$g^i \circ f \circ g^{-i}$$

converges to the constant A_g away from this hyperplane. (The limit is indeed a quasiconstant map since its restriction to the ball $B(x, \epsilon)$ is a constant map, see Exercise 2.80.) Using again the fact that F is a 2-separating subset, we find $f' \in F$ such that

$$f'(A_g) \notin E_{\hat{g} \circ f \circ \check{g}}.$$

Thus, the composition

$$g_i := f' \circ g^i \circ f \circ g^{-i}$$

converges to the constant $f'(A_g)$ uniformly on compacts away from the hyperplane

$$E_{\hat{g} \circ f \circ \check{g}},$$

which is disjoint from $f'(A_g)$. Therefore, according to Proposition 15.15, for all but finitely many $i \in J$, the composition g_i is proximal and

$$\lim_{i \rightarrow \infty, i \in J} A_{g_i} = f'(A_g). \quad \square$$

Our goal, of course, is to find *very proximal elements*, not just proximal ones. We will see now that this can be achieved by using compositions of the same type as in Proposition 15.17, but with a slightly more careful choice of the separating elements f, f' .

PROPOSITION 15.18. *Let $g \in GL(V)$ be a proximal element and $F \subset GL(V)$ be a 4-separating subset. Then for each infinite subset $I \subset \mathbb{N}$ there exist $f, f' \in F$ and an infinite subset $J \subset I$, such that the transformations $g_i = f g^i f g^{-i}$, $i \in J$, satisfy:*

1. g_i, g_i^{-1} are both proximal.
- 2.

$$\lim_{i \rightarrow \infty, i \in J} A_{g_i} = A_g, \quad \lim_{i \rightarrow \infty, i \in J} A_{g_i^{-1}} = A_g.$$

PROOF. We follow the proof of Proposition 15.17, except that we now impose more restrictions on the elements f, f' (using the 4-separation property).

1. We take $f \in F$ such that, in addition to the earlier requirement, we have:

$$f^{-1}(\check{g}(x)) \notin E_g.$$

2. Similarly, we take f' which, additionally, satisfies

$$f'(x) \notin E_{\hat{g} f \check{g}}.$$

Then, taking into account the fact that

$$g_i^{-1} = g^i \circ f^{-1} \circ g^{-i} \circ (f')^{-1},$$

and arguing as in the proof of Proposition 15.17, we obtain that the sequence g_i^{-1} converges on a small ball around $f'^{-1}(x)$ to the constant A_g . Hence, we conclude that (for infinitely many values of i), not only g_i is proximal, but also g_i^{-1} is proximal and

$$\lim_{i \rightarrow \infty} A_{g_i^{-1}} = A_g. \quad \square$$

This also concludes the proof of Theorem 15.16. □

15.5. Finding ping-pong partners: Proof of Theorem 15.7

Let $F \subset GL(V)$ be a 2-separating subset and suppose that $g \in GL(V)$ is a very proximal element.

LEMMA 15.19. *There exists $f \in F$ such that the elements $g, h = fgf^{-1}$ are a ping-pong partners.*

PROOF. For each $f \in GL(V)$, setting $h = fgf^{-1}$, we get:

$$A_h = f(A_g), \quad E_h = f(E_g), \quad A_{h^{-1}} = f(A_{g^{-1}}), \quad E_{h^{-1}} = f(E_{g^{-1}}).$$

In order to ensure that g, h are ping-pong partners, we need to find f such that

$$A_{h^{\pm 1}} \not\subset E_g \cup E_{g^{-1}}, \quad \text{equivalently} \quad f(\{A_g, A_{g^{-1}}\}) \cap (E_g \cup E_{g^{-1}}) = \emptyset,$$

and

$$A_{g^{\pm 1}} \not\subset E_h \cup E_{h^{-1}}, \quad \text{equivalently} \quad f^{-1}(\{A_g, A_{g^{-1}}\}) \cap (E_g \cup E_{g^{-1}}) = \emptyset.$$

The existence of such an f is immediate from the assumption that the subset F is 2-separating. \square

COROLLARY 15.20. *Suppose that $\Gamma < GL(V)$ acts irreducibly on V , the Zariski closure of Γ is an irreducible subgroup of $GL(V)$ and that Γ contains a proximal element. Then Γ contains a pair of ping-pong partners α, β .*

PROOF. This is a combination of Theorem 15.16 and Lemma 15.19. \square

THEOREM 15.21. *Let $\Gamma \leq GL(V)$ be a finitely generated subgroup, acting virtually irreducibly on V and containing a proximal element. Then Γ contains a free non-abelian subgroup.*

PROOF. It remains to summarize what we already have done. After passing to a finite-index subgroup in Γ , we can assume that the Zariski closure of Γ is a Zariski-irreducible subgroup of $GL(V)$. Then, by 15.9, Γ contains a 4-separating subset $F \subset \Gamma$. Using this subset, in Section 15.4.2, we converted a proximal element of Γ into a very proximal element and then (Lemma 15.19) into a pair of ping-pong partners $\alpha, \beta \in \Gamma$. Lastly, by Lemma 7.67, for some $m > 0$, the group generated by α^m, β^m is free of rank 2. \square

This also concludes the proof of the Tits Alternative (Theorem 15.3) for finitely generated subgroups of $GL(n, \mathbb{K})$, where \mathbb{K} is an arbitrary field, as well as the Tits Alternative for subgroups of Lie groups (Theorem 15.1).

As a consequence of the proof of the Tits Alternative we obtain:

THEOREM 15.22. *Let Γ be a finitely generated group that does not contain a free non-abelian subgroup. Then:*

- (1) *If Γ is a subgroup of a real algebraic group G then the Zariski closure $\bar{\Gamma}$ of Γ in G is virtually solvable.*
- (2) *If Γ is a subgroup of a Lie group G with finitely many connected components, then the closure $\hat{\Gamma}$ of Γ in the Lie group G (with respect to the standard topology) is virtually solvable.*

Furthermore, $\bar{\Gamma}$ and $\hat{\Gamma}$ contain a solvable subgroup S of derived length at most $\delta = \delta(G)$ and the index at most $\nu = \nu(G)$.

PROOF. We will prove the first statement as the proof of the second statement is completely analogous (cf. Exercise 14.76). First of all, after replacing G with its finite-index subgroup, we may assume that the group G is irreducible (with respect to the Zariski topology). According to the Tits Alternative, the group Γ is virtually solvable. Furthermore, the adjoint representation $\rho : G \rightarrow GL(\mathfrak{g})$ has abelian kernel. As in the proof of Theorem 14.72, we conclude that the group $G_1 := \overline{\rho(\Gamma)}$ is virtually solvable. According to the same theorem, γ_1 contains a solvable subgroup S_1 of index $\leq \nu(n)$ and derived length $\leq \delta(n)$, where n is the dimension of G . The preimage $\rho^{-1}(G_1) \leq G$ contains the algebraic closure $\overline{\Gamma}$ and $|\rho^{-1}(G_1) : \rho^{-1}(S_1)| \leq \nu(n)$, while $\rho^{-1}(S_1)$ has derived length $\leq \delta(n) + 1$. The same estimates hold for the intersections $S := \rho^{-1}(S_1) \cap \overline{\Gamma}$. Theorem follows. \square

15.6. The Tits Alternative without finite generation assumption

THEOREM 15.23 (The Tits Alternative without finite generation assumption). *Let \mathbb{K} be a field of zero characteristic and Γ be a subgroup of $GL(n, \mathbb{K})$. Then either Γ is virtually solvable or Γ contains a free nonabelian subgroup.*

PROOF. We will need the following elementary lemma:

LEMMA 15.24. *Every countable field \mathbb{L} of zero characteristic embeds into \mathbb{C} .*

PROOF. Since \mathbb{L} has characteristic zero, its prime subfield \mathbb{P} is isomorphic to \mathbb{Q} . Then \mathbb{L} is an extension of the form

$$\mathbb{P} \subset \mathbb{E} \subset \mathbb{L},$$

where $\mathbb{P} \subset \mathbb{E}$ is a purely transcendental extension and $\mathbb{E} \subset \mathbb{L}$ is an algebraic extension (see [Hun80, Chapter VI.1]). Since \mathbb{L} is countable, \mathbb{E}/\mathbb{P} has countable dimension and, therefore,

$$\mathbb{E} = \mathbb{P}(u_1, \dots, u_m)$$

or

$$\mathbb{E} = \mathbb{P}(u_1, \dots, u_m, \dots).$$

Sending variables u_j to algebraically independent transcendental numbers $z_j \in \mathbb{C}$, we then obtain an embedding $\mathbb{E} \hookrightarrow \mathbb{C}$. Since \mathbb{C} is algebraically closed, the algebraic closure of \mathbb{E} embeds in \mathbb{C} . Therefore, \mathbb{L} also embeds in \mathbb{C} . \square

We now return to the proof of the theorem. The group Γ is the direct limit of the direct system of its finitely generated subgroups Γ_i . Let $\mathbb{K}_i \subset \mathbb{K}$ denote the subfield generated by the matrix entries of the generators of Γ_i . Then $\Gamma_i \leq GL(n, \mathbb{K}_i)$. Since \mathbb{K} (and, hence, every \mathbb{K}_i) has zero characteristic, the field \mathbb{K}_i embeds in \mathbb{C} (see Lemma 15.24).

If one of the groups Γ_i contains a free nonabelian subgroup, then so does Γ . Assume, therefore, that this does not happen. Then, in view of the Tits Alternative (for finitely generated matrix groups), each Γ_i is virtually solvable. For $\nu = \nu(GL(n, \mathbb{C}))$ and $\delta = \delta(GL(n, \mathbb{C}))$, for every i there exists a subgroup $\Lambda_i \leq \Gamma_i$ of index $\leq \delta$, such that Λ_i has derived length $\leq \delta$ (see Theorem 14.72). In view of Theorem 13.94, the group Γ is also virtually solvable. \square

COROLLARY 15.25. *$SU(2)$ contains a subgroup isomorphic to F_2 .*

PROOF. The group $SU(2)$ is connected, therefore, it has no proper finite-index subgroups. The group $SU(2)$ cannot be solvable, for instance, because it contains A_5 , the alternating group on 5 symbols, which is a nonabelian finite simple group. Alternatively, if $SU(2) < SL(2, \mathbb{C})$ were solvable, it would preserve a proper subspace in \mathbb{C}^2 according to Theorem 14.65, which is not the case. Now, the claim follows from Theorem 15.23.

15.7. Groups satisfying the Tits Alternative

One says that a group G *satisfies the Tits' Alternative* if it is either virtually solvable or contains a free nonabelian subgroups.

Classes of groups (besides those covered by Theorems 15.1 and 15.23) satisfying the Tits' Alternative are:

- (1) Subgroups of Gromov hyperbolic groups (see [Gro87, §8.2.F], [GdlH90, Chapter 8]).
- (2) Subgroups of the mapping class group, see [Iva92].
- (3) Subgroups of $Out(F_n)$, see [BFH00, BFH05, BFH04].
- (4) Fundamental groups of 3-dimensional manifolds.
- (5) Fundamental groups of compact manifolds of nonpositive curvature, see [Bal95].
- (6) Groups acting isometrically properly discontinuously and cocompactly on two-dimensional $CAT(0)$ complexes [BŠ99, BB95].
- (7) Subgroups of cube complex groups (Sageev-Wise, [SW05]): If G acts properly on a finite-dimensional cube complex and has a bound on order of finite subgroups, then each subgroup of G either contains F_2 or is virtually abelian.
- (8) Certain classes of $CAT(0)$ groups, see [Xie06], [AM15].

Gromov's Theorem

The main objective of this chapter is to prove the converse of Bass–Guivarc'h Theorem 14.26. We refer the reader to Section 8.7 for the definition of the growth function \mathfrak{G} for finitely generated groups.

THEOREM 16.1 (M. Gromov, [Gro81a]). *If Γ is a finitely generated group of polynomial growth then Γ is virtually nilpotent.*

We will actually prove a slightly stronger version (Theorem 16.3 below) of Theorem 16.1, which is due to van der Dries and Wilkie [dDW84] (our proof mainly follows [dDW84]).

DEFINITION 16.2. A finitely generated group Γ has *weakly polynomial growth* of degree $\leq a$ if there exists a sequence of positive numbers R_n diverging to infinity and a pair of numbers C and a , for which

$$\mathfrak{G}(R_n) \leq CR_n^a, \forall n \in \mathbb{N}.$$

THEOREM 16.3. *If Γ has weakly polynomial growth then it is virtually nilpotent.*

Gromov's proof of polynomial growth theorem relies heavily upon the work of Gleason, Yamabe, Montgomery and Zippin on topological transformation groups. Therefore in the following section we review some of the results in the theory of topological transformation groups.

16.1. Topological transformation groups

Throughout this section we will consider only *metrizable* topological spaces X . We will topologize the group of homeomorphisms $\text{Homeo}(X)$ *via* the compact-open topology and, thus, obtain a continuous action $\text{Homeo}(X) \times X \rightarrow X$.

The following definition (Property A of topological groups introduced by Montgomery and Zippin) should not be confused with the Property A in Geometric Group Theory, introduced by G. Yu in [Yu00].

DEFINITION 16.4. [Property A, section 6.2 of [MZ74]] Suppose that H is a separable, locally compact topological group. Then H is said to satisfy Property A if for each neighborhood V of the identity $e \in H$ there exists a compact subgroup $K \subset H$ such that $K \subset V$ and H/K (equipped with the quotient topology) is a Lie group.

In other words, the group H can be *approximated by* the Lie groups H/K . Here is an example to keep in mind. Let H be the additive group \mathbb{Q}_p of p -adic numbers. The sets

$$H_{i,p} := \{x \in \mathbb{Q}_p : |x|_p \leq p^{-i}\}, i \in \mathbb{N},$$

are open and form a basis of topology at $0 \in \mathbb{Q}_p$. For instance, for $i = 0$, $H_{0,p}$ is the group of p -adic integers O_p . Now, the fact that the p -adic norm $|x|_p$ is non-archimedean implies that $H_{i,p}$ is a subgroup of H . Furthermore, this subgroup is closed and, therefore, compact, see Lemma 2.55. The quotient $H/H_{i,p}$ has discrete quotient topology since $H_{i,p}$ is open in H . Hence, $G_{i,p} = H/H_{i,p}$ is a Lie group. In particular, H has Property A.

THEOREM 16.5 (H. Yamabe, [Yam53]; see also [MZ74], Chapter IV). *Each separable locally compact group H contains an open subgroup $\hat{H} \leq H$ such that \hat{H} satisfies Property A.*

The following theorem is proven in [MZ74], section 6.3, Corollary on page 243.

THEOREM 16.6 (D. Montgomery and L. Zippin). *Suppose that X is a topological space which is connected, locally connected, finite-dimensional and locally compact. Suppose that H is a separable locally compact group satisfying Property A, $H \times X \rightarrow X$ is a topological action which is effective and transitive. Then H is a Lie group.*

We are mainly interested in the following corollary for metric spaces.

THEOREM 16.7. *Let X be a metric space which is proper, connected, locally connected and finite-dimensional. Let H be a closed subgroup in $\text{Homeo}(X)$ with the compact-open topology, such that $H \curvearrowright X$ is transitive. If there exists $L \in \mathbb{R}$ such that each $h \in H$ is L -Lipschitz, then the group H is a Lie group with finitely many connected components.*

PROOF. It is clear that $H \times X \rightarrow X$ is a continuous effective action. It follows from the Arzela-Ascoli theorem that H is locally compact.

- LEMMA 16.8.** 1. *The group H is separable.*
 2. *For each open subgroup $U \subset H$, the quotient H/U is countable.*

PROOF. 1. Pick a point $x \in X$. Given $r \in \mathbb{R}_+$, consider the subset

$$H_r = \{h \in H : \text{dist}(x, h(x)) \leq r\}.$$

By the Arzela-Ascoli theorem, each H_r is a compact set. Therefore

$$H = \bigcup_{r \in \mathbb{N}} H_r$$

is a countable union of compact subsets. Thus, it suffices to prove separability of each H_r . Recall that $\bar{B}(x, R)$ denotes the closed R -ball in X centered at the point x . For each $R \in \mathbb{R}_+$ define the map

$$\phi_R : H \rightarrow C_L(\bar{B}(x, R), X)$$

given by the restriction $h \mapsto h|_{\bar{B}(x, R)}$. Here $C_L(\bar{B}(x, R), X)$ is the space of L -Lipschitz maps from $\bar{B}(x, R)$ to X . Observe that $C_L(\bar{B}(x, R), X)$ is metrizable via

$$\text{dist}(f, g) = \max_{y \in \bar{B}(x, R)} \text{dist}(f(y), g(y)).$$

Thus, the image of H_r in each $C_L(\bar{B}(x, R), X)$ is a compact metrizable space. We claim now that each $\phi_R(H_r)$ is separable. Indeed, for each $i \in \mathbb{N}$ take $\mathcal{E}_i \subset \phi_R(H_r)$ to be a $\frac{1}{i}$ -net. The union

$$\bigcup_{i \in \mathbb{N}} \mathcal{E}_i$$

is a dense countable subset of $\phi_R(H_r)$. On the other hand, the group H (as a topological space) is homeomorphic to the *inverse limit*

$$\varprojlim_{R \in \mathbb{N}} \phi_R(H),$$

see Section 1.5.

Let $E_i \subset \phi_i(H_r)$ be a dense countable subset. For each element $h_i \in E_i$ consider a sequence $(g_j) = \tilde{h}_i$ in the above inverse limit such that $g_i = h_i$. Let $\tilde{h}_i \in H$ be the element corresponding to this sequence (g_j) . It is clear now that

$$\bigcup_{i \in \mathbb{N}} \{\tilde{h}_i \in H ; h_i \in E_i\}$$

is a dense countable subset of H_r .

2. Let $I \subset H$ be a dense countable set. The subsets

$$hU, h \in H$$

are open subsets of H such that $hU = gU$ or $hU \cap gU = \emptyset$ for all $g, h \in H$. The countable subset I intersects every $hU, h \in H$. Therefore, the above collection of open subsets of H consists of countably many elements. \square

Thus, we now know that the topological group H is locally compact and separable. Therefore, by Theorem 16.5, H contains an open subgroup $\hat{H} \leq H$ satisfying Property A.

LEMMA 16.9. *For every $x \in X$, the orbit $Y := \hat{H}x \subset X$ is open in X .*

PROOF. As we noted earlier, the group H is σ -compact. Therefore, by the Arens Theorem (Lemma 5.38), the orbit map $h \rightarrow h(x)$ projects a homeomorphism $H/H_x \rightarrow X$, where H_x is the stabilizer of x in H . Since the quotient map $H \rightarrow H/H_x$ is clearly open, openness of \hat{H} implies that Y is open in X . \square

We now can conclude the proof of Theorem 16.7. Let $Z \subset Y$ be the connected component of x in $Y := \hat{H}x$ as above. The stabilizer $F \subset \hat{H}$ of Z is both closed and open in \hat{H} . Therefore, F again has the Property A and the assumptions of Theorem 16.6 are satisfied by the action $F \curvearrowright Z$. It follows that F is a Lie group. Since $F \subset H$ is an open subgroup, the group H is a Lie group as well. Let K be the stabilizer of x in H . The subgroup K is a compact Lie group and, therefore, has only finitely many connected components. Since the action $H \curvearrowright X$ is transitive, X is homeomorphic to H/K , see Lemma 5.38. Connectedness of X now implies that H has only finitely many connected components. \square

16.2. Regular Growth Theorem

We now proceed to construct, for a group Γ of weakly polynomial growth, a representation $\rho : \Gamma \rightarrow \text{Isom}(X)$, where X is a homogeneous metric space as in Theorem 16.7.

The first naive attempt would be to take X to be a Cayley graph $\text{Cayley}(\Gamma, S)$ of Γ . But in that case $\text{Isom}(X)$ does not act transitively on X . If we replace the Cayley graph with its set of vertices, then we achieve homogeneity but loose connectedness. The ingenious idea of Gromov is to take X to be a limit of rescaled Cayley graphs $(\text{Cayley}(\Gamma, S), \lambda_n \text{dist})$, where λ_n is a sequence of positive numbers converging to 0. Gromov originally used Gromov-Hausdorff convergence to define the limit; we

will take X to be an asymptotic cone of $\text{Cayley}(\Gamma, S)$ instead; equivalently X is an asymptotic cone of Γ with the word metric. Such an asymptotic cone inherits both the homogeneity from Γ (see Proposition 10.72) and the property of being geodesic from $\text{Cayley}(\Gamma, S)$ (see (2) in Proposition 10.68). In particular it is connected and locally connected. The asymptotic cone X is also complete, by Proposition 10.70. These properties and the Hopf-Rinow Theorem 2.13 imply that in order to prove properness of X it suffices to prove local compactness.

To sum up, if we wish to apply Theorem 16.7 to an asymptotic cone, it remains to use the hypothesis of polynomial growth to find an asymptotic cone that is locally compact and finite dimensional. In what follows we explain how to choose a scaling sequence λ such that $X_\omega = \text{Cone}_\omega(\Gamma, \mathbf{1}, \lambda)$ has both properties.

REMARK 16.10. We note that once it is proven that $\text{Isom}(X_\omega)$ is a Lie group, one still has to address the issue that the homomorphisms $\Gamma \rightarrow \text{Isom}(X_\omega)$ arising as ultralimits of sequences of isometric actions of Γ on its Cayley graph (with rescaled metric), may have finite images.

A metric space X is called *p-doubling* if each R -ball in X can be covered by p balls of radius $R/2$. One way to show that a metric space X is doubling is to estimate the *packing number* of R -balls in X . The *packing number* $p(\bar{B})$ of a ball $\bar{B} = \bar{B}(x, R) \subset X$ is the supremum of cardinalities of $R/2$ -separated subsets \mathcal{N} of \bar{B} . If \mathcal{N} is a maximal subset as above, then

$$\forall x \in \bar{B} \exists y \in \mathcal{N} \text{ such that } d(x, y) \leq R/2.$$

(This condition is slightly stronger than the one of being an $R/2$ -net.) In other words, the collection of closed balls $\{\bar{B}(x, R/2) : x \in \mathcal{N}\}$ is a covering of B . Thus, there exist a covering of B by $p(B)$ balls of radius $R/2$. If $p(\bar{B}(x, R)) \leq p$ for every x and R , then X has *packing number* $\leq p$; such X is necessarily *p-doubling*. The reader should compare this (trivial) statement with the statement of the Regular Growth Theorem below.

EXERCISE 16.11. Show that doubling implies polynomial growth for uniformly discrete spaces.

Note that, being scale-invariant and invariant under ultralimits, the doubling property passes to asymptotic cones. The following lemma, although logically unnecessary for the proof of Gromov's theorem, motivates its arguments.

LEMMA 16.12. *If X is p-doubling then the Hausdorff dimension of X is at most $\log_2(p)$.*

PROOF. Consider a metric ball $\bar{B} = \bar{B}(o, R) \subset X$. We first cover \bar{B} by balls $\bar{B}(x_i, R/2)$, $i = 1, \dots, p$. We then cover each of the new balls by balls of radius $R/4$ and proceed inductively. On the n -th step of induction we have a covering of \bar{B} by p^n balls of radius $2^{-n}R$. The corresponding finite sum of in the definition of the α -Hausdorff measure of \bar{B} (see (2.7)) then equals

$$\sum_{i=1}^{p^n} 2^{-n\alpha} R^\alpha = R^\alpha \left(\frac{p}{2^\alpha}\right)^n.$$

This quantity converges to 0 as $n \rightarrow \infty$ provided that $p < 2^\alpha$, i.e. $\alpha > \log_2(p)$. Thus, $\mu_\alpha(B) = 0$ for every metric ball in X . Representing X as a countable union of concentric metric balls, we conclude that $\mu_\alpha(X) = 0$ for every $\alpha > \log_2(p)$. \square

Thus, every asymptotic cone of a doubling metric space has finite Hausdorff and, hence, finite covering, dimension, see Theorem 2.42.

Although there are spaces of polynomial growth which are not doubling, the Regular Growth Theorem below shows that groups of polynomial growth exhibit doubling-like behavior, which suffices for proving that suitable asymptotic cones are finite-dimensional.

Our discussion below follows the paper of L. Van den Dries and A. Wilkie, [dDW84], Gromov's original statement and proof of the Regular Growth Theorem were different (although, some key arguments were quite similar).

THEOREM 16.13 (Regular growth theorem). *Let Γ be a finitely generated group. Assume that there exists a sequence (R_n) such that $R = (R_n)^\omega$ is an infinitely large number in the ultrapower \mathbb{R}_+^ω and that the growth function satisfies:*

$$(16.1) \quad \mathfrak{G}_\Gamma(R_n) = |B_\Gamma(1, R_n)| \leq CR_n^a, \text{ for } \omega\text{-all } n \in \mathbb{N},$$

where $C > 0$ and $a \in \mathbb{N}$ are constants independent of n . Let $\epsilon > 0$.

Then there exists $\eta \in [\log R, R] \subset \mathbb{R}_+^\omega$ such that the ball $B(1, \frac{\eta}{4})$ in the ultrapower Γ^ω endowed with the nonstandard metric defined in (10.3) satisfies the following:

For every $i \in \mathbb{N}$, $i \geq 4$, all the $\frac{\eta}{i}$ -separated subsets in the ball $\bar{B}(1, \frac{\eta}{4})$ have cardinality at most $i^{a+\epsilon}$.

In particular, taking $i = 8$, we see that every $\frac{\eta}{4}$ -ball in Γ^ω has packing number $\leq 8^{a+\epsilon}$ (with respect to the nonstandard metric).

We refer the reader to Definition 10.32 and Exercise 10.33 for the discussion of infinitely large numbers. The difference between the assertion of this theorem and the statement that Γ^ω has finite packing number is that *we are not estimating packing numbers of all metric balls, but only of metric balls of certain radii*.

PROOF. Suppose to the contrary that for every $\eta \in [\log R, R] \subset \mathbb{R}_+^\omega$ there exists $i \in \mathbb{N}$, $i \geq 4$, such that the ball $B(1, \frac{\eta}{4})$ contains more than $i^{a+\epsilon}$ points that are $\frac{\eta}{i}$ -separated.

Then we define the function

$$\iota : [\log R, R] \rightarrow \mathbb{N}^\omega, \quad \iota(\eta) \text{ is the smallest } i \in \mathbb{N} \text{ for which the above holds.}$$

The range of ι is \mathbb{N} , which is identified with $\widehat{\mathbb{N}} \subset \mathbb{N}^\omega$.

It is easy to check that ι is an internal map defined by the sequence of maps:

$\iota_n : [\log R_n, R_n] \rightarrow \mathbb{N}$, $\iota_n(r) =$ the minimal $i \in \mathbb{N}$, $i \geq 4$, such that $B_\Gamma(1, \frac{r}{4})$ contains more than $i^{a+\epsilon}$ points that are $\frac{r}{i}$ -separated.

The image of ι is therefore internal, and contained in $\widehat{\mathbb{N}} \subset \mathbb{N}^\omega$. According to Lemma 10.36, the image of ι has to be finite. Thus, there exists $K \in \mathbb{N}$ such that

$$\iota(\eta) \in [4, K], \quad \forall \eta \in [\log R, R].$$

This implies that for every $\eta \in [\log R, R]$ there exists $i = \iota(\eta) \in \{4, \dots, K\}$ such that the ball $B(1, \frac{\eta}{2})$ contains more than $i^{a+\epsilon}$ disjoint balls of radii $\frac{\eta}{2i}$.

In particular, taking $\eta = R$, we see that there exists $i_1 = \iota(R) \in \{4, \dots, K\}$ such that the ball $B(1, \frac{R}{4}) \subset \Gamma^\omega$ contains at least $i_1^{a+\epsilon}$ disjoint balls

$$B\left(x_1(1), \frac{R}{2i_1}\right), B\left(x_2(1), \frac{R}{2i_1}\right), \dots, B\left(x_{t_1}(1), \frac{R}{2i_1}\right) \text{ with } t_1 \geq i_1^{a+\epsilon}.$$

Since Γ^ω is a group which acts on itself isometrically and transitively, all the balls in this list are isometric to $B\left(1, \frac{R}{2i_1}\right)$.

EXERCISE 16.14. For every natural number k and every infinitely large number R ,

$$k \log(R) < R.$$

Thus, $\frac{R}{i_1} \in [\log R, R]$; hence there exists $i_2 = \iota\left(\frac{R}{i_1}\right)$ such that the ball $B\left(1, \frac{R}{4i_1}\right)$ contains at least $i_2^{a+\epsilon}$ disjoint balls of radii $\frac{R}{2i_1i_2}$.

It follows that $B\left(1, \frac{R}{2}\right)$ contains a family of disjoint balls

$$B\left(x_1(2), \frac{R}{2i_1i_2}\right), B\left(x_2(2), \frac{R}{2i_1i_2}\right), \dots, B\left(x_{t_2}(2), \frac{R}{2i_1i_2}\right) \text{ with } t_2 \geq i_1^{a+\epsilon}i_2^{a+\epsilon}.$$

We continue *via* the nonstandard induction. Consider $u \in \mathbb{N}^\omega$ such that $B\left(1, \frac{R}{2}\right)$ contains a family of disjoint balls

$$B\left(x_1(u), \frac{R}{2i_1i_2 \cdots i_u}\right), B\left(x_2(u), \frac{R}{2i_1i_2 \cdots i_u}\right), \dots, B\left(x_{t_u}(u), \frac{R}{2i_1i_2 \cdots i_u}\right),$$

with $t_u \geq (i_1i_2 \cdots i_u)^{a+\epsilon}$.

We construct the next generation of points

$$x_1(u+1), \dots, x_{t_{u+1}}(u+1)$$

by considering, within each ball

$$B\left(x_j(u), \frac{R}{2i_1i_2 \cdots i_u}\right)$$

the centers of $i_{u+1}^{a+\epsilon}$ disjoint balls of radii

$$\frac{R}{2i_1i_2 \cdots i_u i_{u+1}}, \quad 1 \leq j \leq t_{u+1}.$$

Here and below

$$i_{u+1} = \iota\left(\frac{R}{i_1i_2 \cdots i_u}\right),$$

where the product $i_1 \cdots i_{u+1}$ is defined *via* the nonstandard induction as in the end of Section 10.3.

Thus, we obtain injective maps sending $[1, t_{u+1}] \subset \mathbb{N}^\omega$ to $B(1, R/2)$, $j \mapsto x_j(u+1)$.

This induction process continues as long as $R/(i_1 \cdots i_u) \geq \log R$. Recall that $i_j \geq 2$, hence

$$\frac{R}{i_1 \cdots i_u} \leq 2^{-u} R.$$

Therefore, if $u > \log R - \log \log R$ then

$$\frac{R}{i_1 \cdots i_u} < \log R.$$

Thus, there exists $u \in \mathbb{N}^\omega$ such that

$$\frac{R}{i_1i_2 \cdots i_{u+1}} < \log R \leq \frac{R}{i_1i_2 \cdots i_u} \leq \frac{KR}{i_1i_2 \cdots i_{u+1}} \Leftrightarrow$$

$$(16.2) \quad \frac{R}{\log R} < i_1 i_2 \cdots i_{u+1} \leq \frac{KR}{\log R}.$$

Let's "count" the "number" (nonstandard of course!) of *centers*, points $x_i(k)$, we constructed between the first step of the induction and the u -th step of the induction:

We get $i_1^{a+\epsilon} i_2^{a+\epsilon} \cdots i_{u+1}^{a+\epsilon}$ points, i.e. we obtain a bijection from the interval

$$[1, i_1^{a+\epsilon} i_2^{a+\epsilon} \cdots i_{u+1}^{a+\epsilon}] \subset \mathbb{N}^\omega$$

to the set of centers $x_i(k)$. On the other hand, from (16.2) we get:

$$\left(\frac{R}{\log R} \right)^{a+\epsilon} \leq (i_1 i_2 \cdots i_{u+1})^{a+\epsilon}.$$

What does this inequality actually mean? Recall that R and u are represented by sequences of real and natural numbers R_n, u_n , respectively. The above inequality thus implies that for ω -all $n \in \mathbb{N}$, we have:

$$\left(\frac{R_n}{\log R_n} \right)^{a+\epsilon} \leq |B(1, R_n)| \leq CR_n^a.$$

Accordingly,

$$R_n^\epsilon \leq C(\log(R_n))^{a+\epsilon},$$

for ω -all $n \in \mathbb{N}$. This contradicts the assumption that R is infinitely large, cf. Exercise 16.14. \square

16.3. Consequences of the Regular Growth Theorem

PROPOSITION 16.15. *Let Γ be a finitely generated group for which there exists an infinitely large number $R = (R_n)^\omega$ in the ultrapower \mathbb{R}_+^ω such that the growth function satisfies (16.1). Fix real numbers a and $\epsilon > 0$ as in Theorem 16.13 and let $\eta = (\eta_n)$ be a sequence provided by the conclusion of Regular Growth Theorem 16.13; let $\lambda = (\lambda_n)$ with $\lambda_n = \frac{1}{\eta_n}$.*

Then the asymptotic cone $X_\omega = \text{Cone}_\omega(\Gamma; 1, \lambda)$ is

- (a) *locally compact;*
- (b) *has Hausdorff dimension at most $a + \epsilon$. In particular, in view of Theorem 2.42, X_ω has finite covering dimension.*

PROOF. (a) Since X_ω is homogeneous, it suffices to prove that the closed ball $C = \bar{B}(1, \frac{1}{4}) \subset X_\omega$ is compact. Since C is complete, it suffices to show that it is *totally bounded*, i.e. for every $\delta > 0$ there exists a finite cover of C by δ -balls (see [Nag85]).

Let dist denote the word metric on Γ . By Theorem 16.13, the closed ball $\bar{B}(1, \frac{1}{4}) \subset (\Gamma, \lambda_n \text{dist})$ satisfies the property that for every integer $i \geq 4$, every $\frac{1}{i}$ -separated subset $E \subset \bar{B}(1, \frac{1}{4})$ contains at most $i^{a+\epsilon}$ points. The same assertion clearly holds for the ultralimit X_ω . Therefore, we pick some $i \in \mathbb{N}$ such that $\frac{1}{i} < \delta$ and choose (by Zorn's lemma) a maximal $\frac{1}{i}$ -separated subset $E \subset C$. Then, by maximality (see Lemma 2.22),

$$C \subset \bigcup_{x \in E} \bar{B}(x, \frac{1}{i}) \subset \bigcup_{x \in E} \bar{B}(x, \delta).$$

We, thus, have a finite cover of C by δ -balls and, therefore, C is compact.

(b) We first verify that the Hausdorff dimension of the ball $\bar{B}(1, 1/4)$ is at most $a + \epsilon$. Pick $\alpha > a + \epsilon$. For each i consider a maximal $\frac{1}{i}$ -separated subset $x_{1\omega}, x_{2\omega}, \dots, x_{t\omega}$ in $B(1, 1/4)$, with $t \leq i^{a+\epsilon}$.

Then $\bar{B}(1, 1/4)$ is covered by the balls

$$\bar{B}(x_{j\omega}, 1/i), j = 1, \dots, t.$$

We get:

$$\sum_{j=1}^t (1/i)^\alpha \leq i^{a+\epsilon}/i^\alpha = i^{a+\epsilon-\alpha}.$$

Since $\alpha > a + \epsilon$, $\lim_{i \rightarrow \infty} i^{a+\epsilon-\alpha} = 0$. Hence $\mu_\alpha(B(1, 1/4)) = 0$.

Thus by homogeneity of X_ω , $\dim_{Haus}(\bar{B}(x, 1/4)) \leq a + \epsilon$ for each $x \in X_\omega$.

By (a) and Theorem 2.13 X_ω is proper, hence it is covered by countably many balls $\bar{B}(x_n, 1/4)$, $n \in \mathbb{N}$. For every $\alpha > a + \epsilon$, additivity of μ_α implies that

$$\mu_\alpha(X_\omega) \leq \sum_{n=1}^{\infty} \mu_\alpha(\bar{B}(x_n, 1/4)) = 0.$$

Therefore $\dim_{Haus}(X_\omega) \leq a + \epsilon$. □

16.4. Weakly polynomial growth

Here we prove several elementary properties of groups of weakly polynomial growth (Definition 16.2) that will be used in the next section.

LEMMA 16.16. *If Γ has weakly polynomial growth then for every normal subgroup $N \triangleleft \Gamma$, the quotient Γ/N also has weakly polynomial growth.*

PROOF. We equip $Q = \Gamma/N$ with the generating set which is the image of the finite generating set of Γ . Then $B_Q(1, R)$ is the image of $B_\Gamma(1, R)$. Hence,

$$|B_Q(1, R)| \leq |B_\Gamma(1, R)|$$

and, therefore, Q also has weakly polynomial growth. □

LEMMA 16.17. *If Γ has exponential growth then it cannot have weakly polynomial growth.*

PROOF. Since Γ has exponential growth,

$$\lim_{r \rightarrow \infty} \frac{1}{r} \log(\mathfrak{G}(r)) > 0.$$

Suppose that Γ has weakly polynomial growth. This means that growth function of Γ satisfies

$$\mathfrak{G}(R_n) = |B_\Gamma(1, R_n)| \leq CR_n^a$$

for a certain sequence (R_n) diverging to infinity and constants C and a . Hence,

$$\frac{1}{R_n} \log(\mathfrak{G}(R_n)) \leq \frac{\log(C)}{R_n} + \frac{a}{R_n} \log(R_n).$$

However,

$$\lim_{R \rightarrow \infty} \left(\frac{\log(C)}{R} + \frac{a}{R} \log(R) \right) = 0.$$

Contradiction. □

LEMMA 16.18. *Let Γ be a finitely generated subgroup of a Lie group G with finitely many components. If Γ has weakly polynomial growth then Γ is virtually nilpotent.*

PROOF. According to Tits' alternative, either Γ contains a free nonabelian subgroup or is virtually solvable. In the former case, Γ cannot have weakly polynomial growth (see Lemma 16.17). Thus Γ is virtually solvable. Similarly, by applying Theorem 14.37, since Γ has weakly polynomial growth, Γ has to be is virtually nilpotent. \square

16.5. Displacement function

In this section we discuss certain metric properties of action of a finitely generated group Γ on itself by left translations. These properties will be used to prove Gromov's theorem. We fix a finite generating set S of Γ , the Cayley graph $\text{Cayley}(\Gamma, S)$ and the corresponding word metric on Γ .

We define certain *displacement functions* Δ for the action $\Gamma \curvearrowright \Gamma$ by left multiplication. For every $\gamma \in \Gamma$, $x \in \text{Cayley}(\Gamma, S)$ and $r \geq 0$ we define the function measuring the maximal displacement by γ on the ball $\bar{B}(x, r) \subset \text{Cayley}(\Gamma, S)$:

$$\Delta(\gamma, x, r) = \max\{\text{dist}(y, \gamma y) : y \in \bar{B}(x, r)\}.$$

When $x = 1$ we use the notation $\Delta(\gamma, r)$ for the displacement function.

For a subset of $F \subset \Gamma$, define

$$\Delta(F, x, r) = \sup_{\gamma \in F} \Delta(\gamma, x, r).$$

Likewise, we write $\Delta(F, r)$ when $x = 1$.

Clearly, for every $g \in \Gamma$,

$$\Delta(F, g, r) = \Delta(g^{-1}Fg, r).$$

LEMMA 16.19. *Fix $r > 0$ and a finite subset F in Γ . Then the function $\text{Cayley}(\Gamma, S) \rightarrow \mathbb{R}$, $x \mapsto \Delta(F, x, r)$ is 2-Lipschitz.*

PROOF. Let x, y be two points in $\text{Cayley}(\Gamma, S)$. Let p be an arbitrary point in $B(x, r) \subset \text{Cayley}(\Gamma, S)$. A geodesic in $\text{Cayley}(\Gamma, S)$ connecting p to y has length at most $r + \text{dist}(x, y)$, hence it contains a point $q \in B(y, r)$ with $\text{dist}(p, q) \leq \text{dist}(x, y)$. For an arbitrary $\gamma \in F$,

$$\text{dist}(p, \gamma p) \leq \text{dist}(q, \gamma q) + 2\text{dist}(x, y) \leq \Delta(F, y, r) + 2\text{dist}(x, y).$$

It follows that $\Delta(F, x, r) \leq \Delta(F, y, r) + 2\text{dist}(x, y)$. The inequality $\Delta(F, y, r) \leq \Delta(F, x, r) + 2\text{dist}(x, y)$ is proved similarly. \square

LEMMA 16.20. *Suppose that $\Delta(S, r)$ is bounded as a function of r . Then Γ is virtually abelian.*

PROOF. Suppose that $\text{dist}(sx, x) \leq C$ for all $x \in \Gamma$ and $s \in S$. Then

$$\text{dist}(x^{-1}sx, 1) \leq C,$$

and, therefore, the conjugacy class of s in Γ has cardinality $\leq \mathfrak{G}_\Gamma(C) = N$. We claim that the centralizer $Z_\Gamma(s)$ of s in Γ has finite index in Γ : Indeed, if $x_0, \dots, x_N \in \Gamma$ then there are i, k , $0 \leq i \neq k \leq N$, such that

$$x_i^{-1}sx_i = x_k^{-1}sx_k \Rightarrow [x_kx_i^{-1}, s] = 1 \Rightarrow x_kx_i^{-1} \in Z_\Gamma(s).$$

Thus, the intersection

$$A := \bigcap_{s \in S} Z_\Gamma(s)$$

has finite index in Γ . Therefore, A is an abelian subgroup of finite index in Γ . \square

Note that there are virtually abelian groups Γ with unbounded displacement functions $\Delta(S, r)$, the simplest example is $\mathbb{Z}_2 \star \mathbb{Z}_2$.

EXERCISE 16.21. Show that the displacement function $\Delta(S, r)$ of Γ is bounded as a function of r if and only if Γ contains a finite normal subgroup K such that Γ/K is free abelian.

16.6. Proof of Gromov's theorem

In this section we prove Theorem 16.3 and, hence, Theorem 16.1 as well. Let Γ be a group satisfying the assumptions of Theorem 16.3 and $a, \epsilon, R \in \mathbb{R}^*, \eta \in \mathbb{R}^*$ be the quantities appearing in Theorem 16.3. In what follows we fix a finite generating set S of Γ and the corresponding Cayley graph $\text{Cayley}(\Gamma, S)$.

Suppose that Γ has weakly polynomial growth with respect to a sequence (R_n) diverging to infinity. Take the diverging sequence (η_n) given by the Regular Growth Theorem applied to the group Γ . Set $\lambda = (\lambda_n)$ with $\lambda_n = \frac{1}{\eta_n}$. Construct the asymptotic cone $X_\omega = \text{Cone}_\omega(\Gamma; 1, \lambda)$ of the Cayley graph of Γ *via* rescaling by the sequence λ_n and considering the constant sequence e of base-points equal to the identity in Γ . By Proposition 16.15, the metric space X_ω is connected, locally connected, finite-dimensional and proper.

According to Proposition 10.72, we have a homomorphism

$$\alpha : \Gamma_e^\omega \rightarrow L := \text{Isom}(X_\omega)$$

such that $\alpha(\Gamma_e^\omega)$ acts on X_ω transitively. We also get a homomorphism

$$\rho : \Gamma \rightarrow L, \rho = \iota \circ \alpha,$$

where $\iota : \Gamma \hookrightarrow \Gamma_e^\omega$ is the diagonal embedding $\iota(\gamma) = (\gamma)^\omega$. Since the isometric action $L \curvearrowright X_\omega$ is effective and transitive, according to Theorem 16.7, the group L is a Lie group with finitely many components.

REMARK 16.22. Observe that, in view of local compactness of X_ω , the point-stabilizer L_y for $y \in X_\omega$ is a compact subgroup in L . Therefore X_ω (homeomorphic to L/L_x , see Lemma 5.38) can be given an L -invariant Riemannian metric ds^2 . Hence, since X_ω is connected, by using the exponential map with respect to ds^2 we see that if $g \in L$ fixes an open ball in X_ω pointwise, then $g = id$.

The subgroup $\rho(\Gamma) \leq L$ has weakly polynomial growth because Γ has weakly polynomial growth (see Lemma 16.16). By Lemma 16.18, $\rho(\Gamma)$ is virtually nilpotent.

The main problem is that ρ may have large kernel. Indeed, if Γ is abelian then the homomorphism ρ is actually trivial. An induction argument on the degree d of weakly polynomial growth allows to get around this problem and prove Gromov's Theorem. In the induction step, we shall use ρ to construct an epimorphism $\Gamma \rightarrow \mathbb{Z}$, and then apply Proposition 14.39.

If $d = 0$, then $\mathfrak{G}_\Gamma(R_n)$ is bounded. Since the growth function is monotonic, it follows that Γ is finite and there is nothing to prove.

Suppose that each group Γ of weakly polynomial growth of degree $\leq d - 1$ is virtually nilpotent. Let Γ be a group of weakly polynomial growth of degree $\leq d$, i.e.

$$\mathfrak{G}_\Gamma(R_n) \leq C_\Gamma R_n^d,$$

for some sequence (R_n) diverging to infinity. There are two cases to consider:

(a) The image of the homomorphism ρ above is infinite. Then there exists a finite-index subgroup $\Gamma_1 \leq \Gamma$ such that $\rho(\Gamma_1)$ is a torsion-free infinite nilpotent group. The latter has infinite abelianization, hence, we get an epimorphism $\phi : \Gamma_1 \rightarrow \mathbb{Z}$. If $K = \text{Ker}(\phi)$ is not finitely generated, then Γ_1 has exponential growth (see Proposition 14.39), which is a contradiction. Therefore, K is finitely generated. Repeating the arguments in the proof of Proposition 14.39 verbatim we see that K has weakly polynomial growth of degree $\leq d - 1$. It follows that, by the induction hypothesis, K is a virtually nilpotent group. Therefore, Γ_1 is solvable. Applying Lemma 16.17, we conclude that Γ_1 (and, hence, Γ) is virtually nilpotent.

(b) $\rho(\Gamma)$ is finite. First we note that we can reduce to the case when $\rho(\Gamma) = \{1\}$. Indeed, consider the subgroup of finite index $\Gamma' := \text{Ker}(\rho) \leq \Gamma$. For every $\gamma \in \Gamma'$, we have that

$$(16.3) \quad \text{dist}_\Gamma(x_n, \gamma x_n) = o(\eta_n),$$

for every sequence $(x_n) \in \Gamma^\mathbb{N}$ with $\text{dist}_\Gamma(1, x_n) = O(\eta_n)$. Since the inclusion map $\Gamma' \hookrightarrow \Gamma$ is a quasiisometry, the estimate (16.3) holds for sequences (x_n) in Γ' and $\text{dist}_{\Gamma'}$. Thus, Γ' acts trivially on its own asymptotic cone $\text{Cone}_\omega(\Gamma'; 1, \lambda)$, and it clearly suffices to prove that Γ' is virtually nilpotent.

Hence, from now on we assume that $\rho(\Gamma) = \{1\}$. The next exercise clarifies the metric significance of this condition.

EXERCISE 16.23. Let Δ denote the displacement function for the action of Γ on itself *via* left multiplication introduced in Section 16.5. Show that the condition $\text{Ker } \rho = \Gamma$ is equivalent to the property that

$$(16.4) \quad \omega\text{-lim} \frac{\Delta(S, R\eta_n)}{\eta_n} = 0, \text{ for every } R > 0.$$

In other words, all generators of Γ act on Γ with *sublinear* (with respect to $(\eta_n)^\omega$) displacement.

Let $q = q(L)$ denote the constant given by Jordan's theorem applied to the group L . Consider the intersection Γ' of all the subgroups in Γ of index at most q , and let S' denote a finite set generating Γ' . We keep the notation S for a finite generating set of the group Γ .

If the function $\Delta(S', r)$ were bounded then Γ' (and, hence, Γ) would be virtually abelian (Lemma 16.20), which would conclude the proof. Thus, we assume that $\Delta(S', r)$ diverges to infinity as $r \rightarrow \infty$.

LEMMA 16.24. *For every $\theta \in (0, 1]$ there exists a sequence (x_n) in Γ such that*

$$(16.5) \quad \omega\text{-lim} \frac{\Delta(x_n^{-1} S' x_n, \eta_n)}{\eta_n} = \theta.$$

PROOF. By (16.4), for ω -all $n \in \mathbb{N}$ we have $\Delta(S', \eta_n) \leq \theta \eta_n / 2$. Thus, there exists a subset $I \subset \mathbb{N}$ of ω -measure 1 such that for all $n \in I$, there exists a $p_n \in \Gamma$ such that $\Delta(S', p_n, \eta_n) \leq \eta_n / 2$. Fix $n \in I$. Since the function $\Delta(S', r)$ diverges to infinity, there exists $q_n \in \Gamma$ such that

$$\Delta(S', q_n, \eta_n) \geq \max_{s \in S'} \text{dist}(q_n, sq_n) > 2\eta_n.$$

The Cayley graph $\text{Cayley}(\Gamma, S)$ is connected and the function $\text{Cayley}(\Gamma, S) \rightarrow \mathbb{R}$, $p \mapsto \Delta(S', p, \eta_n)$ is continuous by Lemma 16.19. Hence, for ω -all n , there exists a $y_n \in \text{Cayley}(\Gamma, S)$ such that

$$\Delta(S', y_n, \eta_n) = \theta \eta_n.$$

The point y_n is not necessarily in the vertex set of the Cayley graph $\text{Cayley}(\Gamma, S)$. Pick a point $x_n \in \Gamma$ within distance $\frac{1}{2}$ from y_n . Again by Lemma 16.19

$$|\Delta(S', x_n, \eta_n) - \theta \eta_n| \leq 1.$$

It follows that $|\Delta(x_n^{-1} S' x_n, \eta_n) - \theta \eta_n| \leq 1$ and, therefore,

$$\omega\text{-lim} \frac{\Delta(x_n^{-1} S' x_n, \eta_n)}{\eta_n} = \theta. \quad \square$$

For every $0 < \epsilon \leq 1$ we consider a sequence (x_n) as in Lemma 16.24 and define the homomorphism

$$\rho_\theta : \Gamma \rightarrow \Gamma^\omega, \rho_\theta(g) = (x_n^{-1} g x_n)^\omega \in \Gamma^\omega.$$

Since $\Delta(x_n^{-1} S' x_n, \eta_n) = O(\theta \eta_n)$, the image of ρ_ϵ is contained in L . Clearly, the image of ρ_θ is non-trivial. If for some $\theta > 0$, $\rho_\theta(\Gamma)$ is infinite, then we are done as in (a). Hence we assume that $\rho_\theta(\Gamma)$ is finite for all $\theta \in (0, 1]$.

Next, we reduce the problem to the case when all the groups $\rho_\theta(\Gamma)$ are finite abelian. For each θ consider the preimage Γ_θ in Γ of the abelian subgroup in $\rho_\theta(\Gamma)$ which is given by Jordan's theorem applied to L . The index of Γ_θ in Γ is at most q . This implies that the group Γ' described before Lemma 16.24 is contained in Γ_θ for every $\theta > 0$. It follows that $\rho_\theta(\Gamma')$ is finite abelian for every $\theta \in (0, 1]$. Since L is a Lie group, there exists U_δ , a neighborhood of $1 \in L$, which contains no non-trivial finite subgroups. Without loss of generality, we may assume that U_δ has the form

$$U_\delta = \{u \in L : \Delta(u, 1_\omega, 1) < \delta\},$$

for some $\delta > 0$. Thus, for each natural number M and $i \leq M$, we have

$$\Delta(u^i, 1_\omega, 1) < M\delta.$$

By our choice of x_n , for every generator $s \in S'$ (of the group Γ'),

$$\Delta(\rho_\theta(s), 1_\omega, 1) \leq \theta$$

and for one of the generators the inequality is an equality. Assume there exists an $M \in \mathbb{N}$ such that the order $|\rho_\theta(\Gamma')|$ is at most M for all $\theta \in (0, 1]$. Therefore, for every $g \in \Gamma'$,

$$\Delta(\rho_\theta(g), 1_\omega, 1) \leq M\theta.$$

Choose θ such that $M\theta < \delta$. Since U_δ contains no non-trivial subgroups, we conclude that $\rho_\theta(\Gamma') = \{1\}$, which is a contradiction. Therefore,

$$\limsup_{\theta \rightarrow 0} |\rho_\theta(\Gamma')| = \infty.$$

This means that Γ' admits epimorphisms to finite abelian groups of arbitrarily large order. All such homomorphisms have to factor through the abelianization $(\Gamma')_{ab}$ of the group Γ' , therefore, the group $(\Gamma')_{ab}$ is infinite. Since $(\Gamma')_{ab}$ is finitely generated we conclude that Γ' admits an epimorphism to \mathbb{Z} . We apply Proposition 14.39 and the induction hypothesis, and conclude that Γ' is virtually nilpotent. Thus, Γ is virtually nilpotent too, and we are done. This concludes the proof of Theorem 16.3 and, hence, of Theorem 16.1. \square

16.7. Quasi-isometric rigidity of nilpotent and abelian groups

Gromov's theorem has several spectacular corollaries, proving that certain algebraic properties of groups can be recovered from the coarse geometric information.

THEOREM 16.25 (M. Gromov). *Suppose that Γ_1, Γ_2 are quasiisometric finitely generated groups and Γ_1 is virtually nilpotent. Then Γ_2 is virtually nilpotent.*

PROOF. Being virtually nilpotent, Γ_1 has polynomial growth of degree d (Theorem 14.26). Since growth is invariant under quasiisometry, Γ_2 also has polynomial growth of degree d . By Theorem 16.1, Γ_2 is virtually nilpotent. \square

Note that an alternative proof of this theorem (which does not use Gromov's theorem) was given by Y. Shalom [Sha04].

THEOREM 16.26 (P. Pansu). *Suppose that Γ_1, Γ_2 are quasiisometric finitely generated groups and Γ_1 is virtually abelian. Then Γ_2 is virtually isomorphic to Γ_1 .*

PROOF. Without loss of generality, we may assume that Γ_1 is abelian. Let d denote the rank of Γ_1 . Then $\mathfrak{G}_{\Gamma_1}(t) \asymp t^d$. Furthermore, d is the rational cohomological dimension of Γ_1 . Then, by quasiisometry invariance of growth, Γ_2 also has growth $\asymp t^d$. As we just saw above, Γ_2 is virtually nilpotent. Let $\Gamma_3 \leq \Gamma_2$ denote a nilpotent subgroup of finite index in Γ_2 . Let $\Gamma := \Gamma_3 / \text{Tor } \Gamma_3$. By Bass–Guivarc'h Theorem (Theorem 14.26),

$$d = d(\Gamma) = \sum_{i=1}^k i m_i,$$

where m_i is the rank of $C^i\Gamma/C^{i+1}\Gamma$. Recall that the rational cohomological dimension is a quasiisometry invariant, see Theorem 9.64.

Therefore,

$$d = cd(\Gamma) = \sum_{i=1}^k m_i,$$

and

$$\sum_{i=1}^k i m_i = \sum_{i=1}^k m_i.$$

The latter implies that $k = 1$, i.e. Γ is abelian. Virtual isomorphism of the groups Γ_1 and Γ (and, hence, of Γ_2 as well) follows from the equality of their ranks. \square

16.8. Further developments

The following version of Gromov’s theorem was proved by M. Sapir [Sap15], using the recent work of E. Hrushovsky [Hru12] on *approximate groups*. Sapir’s theorem answered affirmatively a question posed by van den Dries and Wilkie in [dDW84].

THEOREM 16.27. *If Γ is a finitely generated group such that some asymptotic cone of Γ is locally compact, then Γ is virtually nilpotent.*

Note that a weaker version of this theorem was proven earlier by F. Point [Poi95], who was assuming, in addition, that the asymptotic cone has finite Minkowski dimension.

Below we review some properties of asymptotic cones of nilpotent groups.

Let (Γ, dist) be a finitely generated nilpotent group endowed with a word metric, let $\text{Tor}(\Gamma)$ be the torsion subgroup of Γ and let H be the torsion-free nilpotent group $\Gamma/\text{Tor}(\Gamma)$. Recall that according to Mal’cev’s Theorem 13.40, the nilpotent group H is isomorphic to a uniform lattice in a connected nilpotent Lie group N .

With every k -step nilpotent Lie group N with Lie algebra \mathfrak{n} one associates the *associated graded Lie algebra* $\bar{\mathfrak{n}}$ obtained as the direct sum

$$\bigoplus_{i=1}^k \mathfrak{c}^i \mathfrak{n} / \mathfrak{c}^{i+1} \mathfrak{n},$$

where $\mathfrak{c}^i \mathfrak{n}$ is the Lie algebra of $C^i N$. Every finite-dimensional Lie algebra is the Lie algebra of a connected Lie group; thus, consider the connected nilpotent Lie group \bar{N} with the Lie algebra $\bar{\mathfrak{n}}$. The group \bar{N} is called the *associated graded Lie group* of the group Γ and of the Lie group N . We refer to Pansu’s paper [Pan83] for the definition of the Carnot-Caratheodory metric appearing in the following theorem:

- THEOREM 16.28** (P. Pansu, [Pan83]). **(a)** *All the asymptotic cones of the finitely generated nilpotent group Γ are bilipschitz homeomorphic to the graded Lie group \bar{N} endowed with a Carnot-Caratheodory metric dist_{CC} .*
- (b)** *For every sequence $\varepsilon_j > 0$ converging to 0 and every word metric dist on Γ , the sequence of metric spaces $(\Gamma, \varepsilon_j \cdot \text{dist})$ converges in the modified Hausdorff metric to $(\bar{N}, \text{dist}_{CC})$.*
- (c)** *The sub-bundle in \bar{N} defining the Carnot-Caratheodory metric is independent of the word metric on Γ , only the norm on this subbundle depends on the word metric.*
- (d)** *The dimension of \bar{N} equals the rational cohomological dimension of Γ , which, in turn, equals*

$$cd_{\mathbb{Q}}(\Gamma) = \sum_{i=1}^k m_i,$$

where m_i is the rank of the abelian quotient $C^i \Gamma / C^{i+1} \Gamma$.

- e)** *The Hausdorff dimension of $(\bar{N}, \text{dist}_{CC})$ equals to the degree of polynomial growth of Γ , that is, to*

$$d(\Gamma) = \sum_{i=1}^k i m_i$$

Note that, according to Theorem 10.46, (a) implies (b) in Pansu's theorem. We further note that \overline{N} , treated as a Lie group, is also a quasiisometry invariant of Γ , see [Pan89].

REMARK 16.29. One says that two metric spaces are *asymptotically bi-Lipschitz* if their asymptotic cones are bi-Lipschitz homeomorphic. Pansu's theorem above has lead to examples of asymptotically bi-Lipschitz nilpotent groups, which are not quasiisometric. Indeed, by Pansu's theorem, every two finitely generated nilpotent groups $\Gamma_i, i = 1, 2$, with isomorphic associated graded Lie groups \overline{N}_i , are asymptotically bi-Lipschitz. Y. Benoist constructed two nilpotent groups Γ_1, Γ_2 with isomorphic associated graded Lie groups $\overline{N}_1, \overline{N}_2$, but distinct *virtual Betti numbers*. Y. Shalom proved that for of finitely generated nilpotent groups virtual Betti numbers are quasiisometry invariant. Therefore, Benoist's groups Γ_1, Γ_2 are asymptotically bi-Lipschitz but not quasiisometric. We refer to Shalom's paper [Sha04, p. 151-152] for the details.

The Banach–Tarski paradox

In this chapter we discuss the *Banach–Tarski Paradox*, which relies upon existence of free nonabelian subgroups in orthogonal groups $O(n)$, $n \geq 3$ and also connects to the notion of amenability, which will be discussed in detail in the next chapter.

17.1. Paradoxical decompositions

The Banach–Tarski Paradox deals with decompositions of subsets in the Euclidean space into congruent pieces and rearranging them via isometries. We begin with discussing the concepts involved in this process for general sets and group actions. In what follows, X is a set and $G \leq \text{Bij}(X)$ is a group of bijections. From the geometric viewpoint, the most interesting case is that of $X = \mathbb{E}^n$ (the Euclidean n -space) and G a subgroup of the group of isometries of \mathbb{E}^n . Another useful example which we discuss in the next chapter is when $X = G$ is a group and G acts on itself by left multiplication.

DEFINITION 17.1. Two subsets A, B in X are *G -congruent* if there exists $g \in G$ such that $g(A) = B$. The restriction $g|_A$ is called a *G -congruence* from A to B .

DEFINITION 17.2. A bijection $\phi : A \rightarrow B$ between two subsets of X is called a *piecewise G -congruence* if the subsets A, B admit partitions into non-empty parts,

$$A = A_1 \sqcup \dots \sqcup A_k, \quad B = B_1 \sqcup \dots \sqcup B_k$$

such that the restrictions

$$(17.1) \quad \phi_i = \phi|_{A_i} : A_i \rightarrow B_i, \quad i = 1, \dots, k,$$

are G -congruences. Accordingly, two subsets A, B are called *piecewise G -congruent* if there exists a piecewise G -congruence $A \rightarrow B$.

EXERCISE 17.3. Prove that piecewise G -congruence is an equivalence relation.

DEFINITION 17.4. A subset $E \subset X$ is *G -paradoxical* if it is non-empty and admits a partition $E = E' \sqcup E''$ such that E' and E'' are both piecewise G -congruent to E . In detail: There exist partitions

$$E' = A'_1 \sqcup \dots \sqcup A'_k, \quad E'' = A''_1 \sqcup \dots \sqcup A''_l$$

and bijections $\phi' : E' \rightarrow E, \phi'' : E'' \rightarrow E$ which restrict to congruences

$$\phi'_i : E'_i \rightarrow \phi'_i(E'_i) \subset E, \quad \phi''_j : E''_j \rightarrow \phi''_j(E''_j) \subset E,$$

$i = 1, \dots, k, j = 1, \dots, l$. The subsets $E'_i, E''_j, 1 \leq i \leq k, 1 \leq j \leq l$, are called *pieces* of the *G -paradoxical decomposition*

$$E'_1 \sqcup \dots \sqcup E'_k \sqcup E''_1 \sqcup \dots \sqcup E''_l$$

of the subset $E \subset X$.

EXERCISE 17.5. If $A, B \subset X$ are piecewise G -congruent and A is G -paradoxical, then so is B .

A group action $G \curvearrowright X$ is called *paradoxical* if the set X is paradoxical with respect to this action. A group action $G \curvearrowright X$ is called *weakly paradoxical* if there exists a G -paradoxical subset $E \subset X$. Thus, every paradoxical action is also weakly paradoxical.

In the context of groups, considering the G -action on itself via left multiplication

$$L : G \times G \rightarrow G, \quad L(g, x) = L_g(x) = gx,$$

we arrive to the following definition:

DEFINITION 17.6. A group G is *paradoxical* (resp. *weakly paradoxical*) if the action $L : G \times G \rightarrow G$ is paradoxical (resp. weakly paradoxical).

Next, we prove several facts about piecewise congruences and paradoxical decompositions.

LEMMA 17.7. *Suppose that $S \subset G$ is a paradoxical subset, $G \curvearrowright X$ is an action such that for some $x \in X$ the orbit map*

$$f : G \rightarrow X, \quad g \mapsto g(x)$$

restricts to an injective map on S . Then X is G -paradoxical.

PROOF. Let

$$S = S' \sqcup S'', \quad S' = S'_1 \sqcup \dots \sqcup S'_k, \quad S'' = S''_1 \sqcup \dots \sqcup S''_l$$

be a G -paradoxical decomposition of S with the piecewise-congruences $\phi' : S' \rightarrow S, \phi'' : S'' \rightarrow S$, where

$$\phi'_i|_{S'_i} = g_i|_{S'_i}, \quad \phi''_j|_{S''_j} = g_j|_{S''_j}, \quad 1 \leq i \leq k, \quad 1 \leq j \leq l.$$

Define the partitioned subset $E = E' \sqcup E'' \subset X$ as $E = f(S), E' = f(S'), E'' = f(S'')$. Furthermore, define bijections

$$\psi' : E' \rightarrow E, \quad \psi'' : E'' \rightarrow E$$

by

$$\psi'(f(s)) = f(\phi'(s)), \quad \psi''(f(s)) = f(\phi''(s)).$$

It follows that the restriction of ψ' to $f(S'_i)$ is given by g'_i and the restriction of ψ'' to $f(S''_j)$ is given by g''_j , $1 \leq i \leq k$, $1 \leq j \leq l$. Therefore, ψ', ψ'' are piecewise G -congruences. \square

LEMMA 17.8. *Suppose that $H \leq G$ is an infinite cyclic subgroup preserving a subset $A \subset X$. Suppose that $E \subset A$ is such that $hE \cap E = \emptyset$ for all $h \in H \setminus \{1\}$. Then A is piecewise H -congruent to $A \setminus E$. In particular, if H acts freely on its orbit $Hx \subset A$, then A is piecewise H -congruent to $A \setminus \{x\}$.*

PROOF. Let g be a generator of H . Define the partition $A = A_1 \sqcup A_2$,

$$A_1 = \bigcup_{n \in \mathbb{Z}_+} g^n E, \quad A_2 := A \setminus A_1.$$

Now consider the map $\phi : A \rightarrow A \setminus E$ which is the identity on A_2 and is $g|_{A_1}$ on A_1 . Then ϕ is a piecewise H -congruence. \square

COROLLARY 17.9. *Let $G = SO(n)$. Then for each $n \geq 2$, the unit sphere $\mathbb{S}^{n-1} \subset \mathbb{E}^n$ is piecewise G -congruent to $\mathbb{S}^{n-1} \setminus \{p\}$, where $p \in \mathbb{S}^{n-1}$ is any point.*

PROOF. As usual, we identify \mathbb{E}^n with the vector space \mathbb{R}^n equipped with the standard Euclidean metric. Without loss of generality, we may assume that p belongs to $\mathbb{R}^2 \cap \mathbb{S}^{n-1}$. Let $h \in O(2)$ (the orthogonal group of \mathbb{R}^2) be an infinite order rotation. Then no power $h^k, k \neq 0$, fixes p . Extending h by the identity to the orthogonal complement of \mathbb{R}^2 in \mathbb{R}^n , we obtain an isometry $g \in SO(n)$ such that no power $g^k, k \neq 0$, fixes p . Now claim follows from Lemma 17.8. \square

The next lemma shows how to “double” paradoxical subsets:

LEMMA 17.10. *Suppose that $A \subset X$ is a G -paradoxical subset. Then A is piecewise G -congruent to any subset $B \subset X$ of the form*

$$B = B_1 \sqcup \dots \sqcup B_k,$$

where each B_i is G -congruent to A .

PROOF. It suffices to consider $k = 2$ as the general case follows by induction. Let $A = A_1 \sqcup A_2$ be a G -paradoxical decomposition and $\phi_i : A_1 \rightarrow A$ are piecewise G -congruences. Then composing ϕ_i with a G -congruence $\psi_i : A \rightarrow B_i$ ($i = 1, 2$), we obtain the required piecewise G -congruence $A \rightarrow B$. \square

REMARK 17.11. We note that instead of taking finite partitions, one can also take countable partitions; this leads to the notion of *countable G -congruence* (and countably paradoxical decompositions), but we will not discuss it in the book as its relation to the Geometric Group Theory is only tangential. We refer the reader to [Wag85] for the details.

We now specialize to the case $X = \mathbb{E}^n$, which is the n -dimensional Euclidean space and G the group of isometries of X . Building upon earlier work of Vitali [Vit05] and F. Hausdorff [Hau14], S. Banach and A. Tarski proved in [BT24] the following:

THEOREM 17.12 (Banach-Tarski paradox). *For $n \geq 3$, any two bounded subsets with non-empty interior in \mathbb{E}^n are piecewise-congruent.*

A corollary of this theorem is much better known:

COROLLARY 17.13. *Let n be at least 3 and let G denote the group $\text{Isom}(\mathbb{E}^n)$ of isometries of the Euclidean n -space.*

- (1) *Every closed ball in \mathbb{E}^n is G -paradoxical.*
- (2) *For $m \in \mathbb{N}$, every closed ball in \mathbb{E}^n is piecewise G -congruent to the disjoint union of m isometric copies of this ball in \mathbb{E}^n (one can “double” the ball).*
- (3) *Any two round n -balls in \mathbb{E}^n are piecewise G -congruent.*

We note that Part 2 of the corollary follows from Part 1 and Lemma 17.10.

REMARK 17.14. The Banach-Tarski paradox implies that there are no finitely-additive measures defined on *all* subsets of the Euclidean space of dimension at least 3 which are invariant with respect to isometries and take positive value on the unit cube. In particular, the congruent pieces A_i, B_i are not *Lebesgue measurable*.

REMARK 17.15 (Banach-Tarski paradox and the Axiom of Choice). The Banach-Tarski paradox is neither provable nor disprovable with Zermelo-Fraenkel axioms (ZF) only: It is impossible to prove that the unit ball in \mathbb{E}^3 is paradoxical in ZF, it is also impossible to prove it is not paradoxical. An extra axiom is needed, e.g., the Axiom of Choice (AC). In fact, work of M. Foreman & F. Wehrung [FW91] and J. Pawlikowski [Paw91] shows that the Banach-Tarski paradox can be proved assuming ZF and the Hahn-Banach theorem (which is a weaker axiom than AC, see Section 10.1).

In this book we will prove only Parts 1 (and, hence, Part 2) of Corollary 17.13; we refer the reader to [Wag85] for a proof of Theorem 17.12. We only note here that Theorem 17.12 is derived from the doubling of a ball (Part 2 of Corollary 17.13) by using the Banach-Bernstein-Schroeder theorem (see [Wag85]).

REMARK 17.16. Inspired by the Hausdorff's argument, R. M. Robinson, answering a question of von Neumann, proved in [Rob47] that five is the minimal number of pieces in a paradoxical decomposition of the unit 3-dimensional ball. See Section 18.7 for a discussion on the minimal number of pieces in a paradoxical decomposition.

REMARK 17.17. It turns out that the sphere \mathbb{S}^{n-1} can be partitioned into 2^{\aleph_0} pieces, so that each piece is piecewise congruent to \mathbb{S}^{n-1} , see [Wag85].

17.2. Step 1: A paradoxical decomposition of the free group F_2

Let F_2 be the free group of rank 2 with generators a, b . Given u , a reduced word in a, b, a^{-1}, b^{-1} , we denote by \mathcal{W}_u the set of reduced words in a, b, a^{-1}, b^{-1} with the prefix u . Every $x \in F_2$ defines a map $L_x : F_2 \rightarrow F_2$, $L_x(y) = xy$ (left translation by x).

Then

$$(17.2) \quad F_2 = \{1\} \sqcup \mathcal{W}_a \sqcup \mathcal{W}_{a^{-1}} \sqcup \mathcal{W}_b \sqcup \mathcal{W}_{b^{-1}}$$

but also $F_2 = L_a \mathcal{W}_{a^{-1}} \sqcup \mathcal{W}_a$, and $F_2 = L_b \mathcal{W}_{b^{-1}} \sqcup \mathcal{W}_b$. We slightly modify the above partition in order to include $\{1\}$ into one of the other four subsets.

Consider the following modifications of \mathcal{W}_a and $\mathcal{W}_{a^{-1}}$:

$$\mathcal{W}'_a = \mathcal{W}_a \setminus \{a^n ; n \in \mathbb{Z}\} \text{ and } \mathcal{W}'_{a^{-1}} = \mathcal{W}_{a^{-1}} \cup \{a^n ; n \in \mathbb{Z}\}.$$

Then

$$(17.3) \quad F_2 = (\mathcal{W}'_{a^{-1}} \sqcup \mathcal{W}'_a) \sqcup (\mathcal{W}_{b^{-1}} \sqcup \mathcal{W}_b)$$

and

$$F_2 = L_a \mathcal{W}'_{a^{-1}} \sqcup \mathcal{W}'_a = L_b \mathcal{W}_{b^{-1}} \sqcup \mathcal{W}_b.$$

Therefore, (17.3) is a G -paradoxical decomposition (with four pieces) of the group F_2 with $G \leq \text{Bij}(F_n)$ the group F_2 acting on itself by left multiplication, i.e. G is the image of $L : F_2 \rightarrow \text{Bij}(F_2)$, $L(u) = L_u$.

We thus proved:

LEMMA 17.18. *The free group F_2 is paradoxical.*

EXERCISE 17.19. Prove that every free group $F_n, n \geq 2$, is paradoxical.

LEMMA 17.20. *Suppose that X is a non-empty set and $\rho : F_2 \times X \rightarrow X$ is a free action of F_2 . Then X is F_2 -paradoxical, with a paradoxical decomposition consisting of four pieces.*

PROOF. According to the axiom of choice there exists a subset $D \subset X$ which intersects every F_2 -orbit in X exactly once. For subsets $R \subset F_2$ and $S \subset X$ we set

$$R \cdot S := \{\rho(g)(x) : g \in R, x \in S\}.$$

We now partition E as:

$$E'_1 = \mathcal{W}'_{a^{-1}} \cdot D, \quad E'_2 = \mathcal{W}'_a \cdot D, \quad E''_1 := \mathcal{W}_{b^{-1}} \cdot D, \quad E''_2 := \mathcal{W}_b \cdot D,$$

where

$$F_2 = (\mathcal{W}'_{a^{-1}} \sqcup \mathcal{W}'_a) \sqcup (\mathcal{W}_{b^{-1}} \sqcup \mathcal{W}_b)$$

is the F_2 -paradoxical decomposition of F_2 defined on on the Step 1. Then we have piecewise F_2 -congruences

$$\begin{aligned} \phi' : E'_1 \sqcup E'_2 &\rightarrow E, & \phi'|_{E'_1} &= \rho(a)|_{E'_1}, & \phi'|_{E'_2} &= \text{Id}, \\ \phi'' : E''_1 \sqcup E''_2 &\rightarrow E, & \phi''|_{E''_1} &= \rho(b)|_{E''_1}, & \phi''|_{E''_2} &= \text{Id}. \end{aligned}$$

□

CONVENTION 17.21. For the rest of the chapter, $G = \text{Isom}(\mathbb{E}^n)$ and for simplicity of the notation we will refer to G -congruences simply as congruences and to G -paradoxical decompositions simply as paradoxical decompositions.

17.3. Step 2: The Hausdorff paradox

In this section we prove the Hausdorff Paradox (and its generalization in higher dimensions), a historic precursor to the Banach–Tarski theorem. Recall that \mathbb{S}^2 is the unit sphere in \mathbb{E}^3 centered at the origin.

THEOREM 17.22 (Hausdorff Paradox). *There exists a countable subset $C \subset \mathbb{S}^2$ such that $\mathbb{S}^2 \setminus C$ is paradoxical.*

PROOF. Recall that there exists a isomorphism $\rho : F_2 \rightarrow H \leq SO(3)$: This can be viewed as a corollary of the Tits Alternative (Corollary 15.25), it was also proven in more directly Corollary 7.65. Let $C \subset \mathbb{S}^2$ denote the set of fixed points of elements of $H \setminus \{1\}$, this set is clearly H -invariant. The action of H on $E = \mathbb{S}^2 \setminus C$ is free and, hence, Lemma 17.20 implies that E is H -paradoxical. Since H acts isometrically on \mathbb{E}^3 , theorem follows. □

We next extend Hausdorff Paradox in higher dimensions. Given a representation $\eta : SU(2) \rightarrow GL(n, \mathbb{R})$ we let Fix_η denote the union of linear subspaces in \mathbb{R}^n fixed by all non-trivial elements of $SU(2)$ and Fix'_η denote the union of subspaces in \mathbb{R}^n fixed by all noncentral elements of $SU(2)$.

LEMMA 17.23. *For each $n \geq 3$ there exists a representation $\eta : SU(2) \rightarrow SO(n)$, such that:*

1. $\text{Fix}_\eta = \{0\}$, if n is divisible by 4.
2. Fix_η is a line or a plane, if n is congruent to 1 or 2 mod 4.
3. Fix'_η is a countable union of real lines, if n is congruent to 3 mod 4.

PROOF. 1. If $n = 4k$, we let

$$\eta_{4k} : SU(2) \rightarrow (SU(2))^k = \underbrace{SU(2) \times \dots \times SU(2)}_{k \text{ times}}$$

denote the diagonal embedding. Viewing $(SU(2))^k$ as a subgroup of $SU(2k) \subset SO(n)$ we obtain a representation $\eta : SU(2) \rightarrow SO(n)$ with $\text{Fix}_\eta = \{0\}$.

2. If $n = 4k + 1$ or $n = 4k + 2$, we identify $SO(4k)$ with a subgroup of $SO(n)$ preserving a $4k$ -dimensional subspace V and fixing pointwise its orthogonal complement V^\perp . Then for the representation

$$\eta : SU(2) \xrightarrow{\eta_{4k}} SO(4k) < SO(n).$$

we have $\text{Fix}_\eta = V^\perp$.

3. Lastly, if $n = 4k + 3$, we use the product representation

$$\eta_{4k} \times \zeta : SU(2) \rightarrow SO(4k) \times SO(3) < SO(4k + 3),$$

where $\zeta : SU(2) \rightarrow SO(3)$ is the universal cover (whose kernel is the center of $SU(2)$). The group $SO(4k)$ fixes a 3-dimensional subspace $W \subset \mathbb{R}^n$ and Fix'_η is a countable union of lines in W . \square

THEOREM 17.24. *For each $n \geq 3$, there exists a proper subset $C \subset \mathbb{S}^{n-1}$ whose complement $\mathbb{S}^{n-1} \setminus C$ is paradoxical. Furthermore, C is either empty (if n divisible by 4), or is countable (if n is odd) or is a single great circle (if n is even, not divisible by 4).*

PROOF. Given a monomorphism $\iota : F_2 \rightarrow SU(2)$ we let $\rho : F_2 \rightarrow SO(n), n \geq 4$, denote the composition of ι with the representation $\eta : SU(2) \rightarrow SO(n)$ constructed in Lemma 17.23. Note that ρ is a monomorphism since kernel of η can only contain central elements of $SU(2)$, hence, only the identity element of $\iota(F_2)$. Define the subset

$$C \subset \mathbb{S}^{n-1}$$

as the (countable) union of fixed-point set of non-trivial elements of $\rho(F_2)$. If n is divisible by 4, the subset C is empty. Now, the group $\rho(F_2)$ acts freely on $E := \mathbb{S}^{n-1} \setminus C$ theorem follows from Lemma 17.20. \square

17.4. Step 3: Spheres of dimension ≥ 2 are paradoxical

LEMMA 17.25. *For each $n \geq 2$, every subset $E \subset \mathbb{S}^{n-1}$ with countable complement is piecewise-congruent to \mathbb{S}^{n-1} .*

PROOF. Let C denote the complement of E in \mathbb{S}^{n-1} . We claim that there exists a codimension 2 subspace $F \subset \mathbb{R}^n$, which is disjoint from C . Indeed, for each $c \in C$ the set $L_c \subset \text{Lin}(\mathbb{R}^n, \mathbb{R}^2)$ of linear maps $\lambda : \mathbb{R}^n \rightarrow \mathbb{R}^2$ with $c \in \text{Ker}(\lambda)$, is nowhere dense in $\text{Lin}(\mathbb{R}^n, \mathbb{R}^2)$. Therefore, by the Baire Theorem, there exists a linear map $\lambda : \mathbb{R}^n \rightarrow \mathbb{R}^2$ whose kernel is disjoint from C . Then F is the subspace we needed. We identify $SO(2)$ with the subgroup of $SO(n)$ fixing F pointwise. Then for any two elements $c_1, c_2 \in C$ there exists at most one $g \in SO(2)$ such that $g(c_1) = c_2$. Since $SO(2)$ is uncountable, we conclude that there exists an element $g \in SO(2) < SO(n)$ such that

$$g^k(C) \cap C = \emptyset, \quad \forall k \in \mathbb{Z} \setminus \{0\}.$$

Now, the assertion follows from Lemma 17.8. \square

EXERCISE 17.26. Suppose that $F \subset \mathbb{R}^n$ is a 2-dimensional subspace and $n \geq 2$. Then there exists $g \in SO(n)$ such that for all $k \in \mathbb{Z} \setminus \{0\}$, $g^k F \cap F = \{0\}$.

It now follows from Lemma 17.8 that each subset $E \subset \mathbb{S}^{n-1}$ whose complement is a great circle, is piecewise-congruent to \mathbb{S}^{n-1} . According to Theorem 17.24, for each $n \geq 3$ there exists a paradoxical subset $E \subset \mathbb{S}^{n-1}$ whose complement is either countable or is a great circle. Since E is piecewise-congruent to \mathbb{S}^{n-1} , we obtain:

THEOREM 17.27. *The sphere \mathbb{S}^{n-1} is $SO(n)$ -paradoxical for all $n \geq 3$.*

17.5. Step 4: Euclidean unit balls are paradoxical

According to Theorem 17.27, for each $n \geq 3$, there exists a partition

$$S = \mathbb{S}^{n-1} = E' \sqcup E''$$

and piecewise $SO(n)$ -congruences $\phi' : E' \rightarrow S$ and $\phi'' : E'' \rightarrow S$. We define *radial extensions*

$$\hat{E}' = \{\lambda x : \lambda \in (0, 1], x \in E'\}, \quad \hat{E}'' = \{\lambda x : \lambda \in (0, 1], x \in E''\}$$

of the subsets E', E'' . Accordingly, we let $\hat{\phi}', \hat{\phi}''$ be the radial extensions of the piecewise $SO(n)$ -congruences ϕ', ϕ'' . Both maps $\hat{\phi}', \hat{\phi}''$ are piecewise $SO(n)$ -congruences

$$\hat{\phi}' : \hat{E}' \rightarrow \mathbb{B}^n \setminus \{0\}, \quad \hat{\phi}'' : \hat{E}'' \rightarrow \mathbb{B}^n \setminus \{0\}.$$

Therefore, the *punctured ball* $\mathbb{B}^n \setminus \{0\}$ is $SO(n)$ -paradoxical.

LEMMA 17.28. *The punctured unit ball $\mathbb{B}^n \setminus \{0\}$ is piecewise congruent to \mathbb{B}^n .*

Let $\Sigma \subset \mathbb{B}^n$ be a round sphere containing the origin 0. According to Lemma 17.25, there exists a piecewise congruence $\psi : \Sigma \setminus \{0\} \rightarrow \Sigma$. We then define a piecewise congruence

$$\phi : \mathbb{B}^n \setminus \{0\} \rightarrow \mathbb{B}^n$$

as the identity on $\mathbb{B}^n \setminus \Sigma$ and ψ on $\Sigma \setminus \{0\}$. □

Therefore, since the punctured ball $\mathbb{B}^n \setminus \{0\}$ is paradoxical, so is the ball \mathbb{B}^n . This concludes the proof of Corollary 17.13, Parts 1 and 2, for $n \leq 3$. □

Amenability and paradoxical decomposition.

In this chapter we discuss in detail two important concepts behind the Banach-Tarski paradox: Amenability and paradoxical decompositions. Although both properties were first introduced for groups (of isometries), it turns out that amenability can be defined in purely metric terms, in the context of graphs of bounded geometry. We shall begin by discussing amenability for graphs, then we will turn to the case of groups, and after that, to the opposite property of being *paradoxical*.

CONVENTION 18.1. Throughout the chapter, all the graphs are assumed to be non-empty. In the case of a non-connected graph \mathcal{G} , we declare the distance between vertices in different components of \mathcal{G} to be infinite.

18.1. Amenable graphs

We refer the reader to Definition 1.43 for definitions of various boundaries of subgraphs of a graph.

DEFINITION 18.2. A (non-empty) graph \mathcal{G} is called *amenable* if there exists a sequence Φ_n of finite subsets of V such that

$$(18.1) \quad \lim_{n \rightarrow \infty} \frac{|\partial^V \Phi_n|}{|\Phi_n|} = 0.$$

Such sequence Φ_n is called a *Følner sequence for the graph \mathcal{G}* .

Note that if \mathcal{G} has finite valence C , then the vertex boundaries $\partial_V \Phi_n$ and the exterior vertex boundaries $\partial^V \Phi_n$ satisfy

$$\begin{aligned} |\partial_V \Phi_n| &\leq |E(\Phi_n, F_n^c)| \leq C |\partial^V \Phi_n| \\ |\partial^V \Phi_n| &\leq |E(\Phi_n, F_n^c)| \leq C |\partial_V \Phi_n|. \end{aligned}$$

Therefore in this case, (18.1) is equivalent to

$$\lim_{n \rightarrow \infty} \frac{|\partial_V \Phi_n|}{|\Phi_n|} = 0$$

and

$$\lim_{n \rightarrow \infty} \frac{|E(\Phi_n, \Phi_n^c)|}{|\Phi_n|} = 0.$$

It is immediate from the definition that every finite graph is amenable (take $\Phi_n = V$).

An infinite connected graph is amenable if and only if its Cheeger constant, as described in Definition 8.88, is zero. This equivalent characterization of amenability does not extend to finite graphs, which all have positive Cheeger constants.

We describe in what follows various metric properties equivalent to non-amenability. Our arguments are adapted from [dlHGCS99]. The only tool that will be needed is Hall–Rado Marriage Theorem from graph theory, stated below.

Let $Bip(Y, Z; E)$ denote the bipartite graph with vertex set V split as $V = Y \sqcup Z$, and the edge-set E . Given two integers $k, l \geq 1$, a *perfect (k, l) -matching* of $Bip(Y, Z; E)$ is a subset $M \subset E$ such that each vertex in Y is the endpoint of exactly k edges in M , while each vertex in Z is the endpoint of exactly l edges in M .

THEOREM 18.3 (Hall-Rado [Bol79], §III.2). *Let $Bip(Y, Z; E)$ be a locally finite bipartite graph and let $k \geq 1$ be an integer such that:*

- *For every finite subset $A \subset Y$, its exterior vertex-boundary $\partial^V A$ contains at least $k|A|$ elements.*
- *For every finite subset B in Z , its exterior vertex-boundary contains at least $|B|$ elements.*

Then $Bip(Y, Z; E)$ has a perfect $(k, 1)$ -matching.

Given a discrete metric space (X, dist) , two (not necessarily disjoint) subsets Y, Z in X , and a real number $C \geq 0$, one defines a bipartite graph $Bip_C(Y, Z)$, with the vertex set $Y \sqcup Z$, where two vertices $y \in Y$ and $z \in Z$ are connected by an edge in $Bip_C(Y, Z)$ if and only if $\text{dist}(y, z) \leq C$. (The reader will recognize here a version of the Rips complex of a metric space.) We will use this construction in the case when $Y = Z = X$, then the vertex set of $Bip(X, X)$ will consist of two copies of the set X .

In what follows, given a subset $\Phi \subset V$ of the vertex set of a graph \mathcal{G} , we will use the notation $\overline{\mathcal{N}}_C(\Phi)$ and $\mathcal{N}_C(\Phi)$ to denote the “closed” and “open” C -neighborhoods of Φ in V :

$$\overline{\mathcal{N}}_C(\Phi) = \{v \in V : \text{dist}(v, \Phi) \leq C\}, \quad \mathcal{N}_C(\Phi) = \{v \in V : \text{dist}(v, \Phi) < C\}.$$

THEOREM 18.4. *Let \mathcal{G} be a connected graph of bounded geometry, with vertex set V and edge set E , endowed, as usual, with the standard metric. The following conditions are equivalent:*

- (a) \mathcal{G} is non-amenable.
- (b) \mathcal{G} satisfies the following expansion condition: *There exists a constant $C > 0$ such that for every finite non-empty subset $\Phi \subset V$, the set $\overline{\mathcal{N}}_C(\Phi) \subset V$ contains at least twice as many vertices as Φ .*
- (b') *For some (equivalently, every) $\beta > 1$ there exists $C > 0$ such that $\overline{\mathcal{N}}_C(\Phi) \cap V$ has cardinality at least β times the cardinality of Φ .*
- (c) *There exists a constant $C > 0$ such that the graph $Bip_C(V, V)$ has a perfect $(2, 1)$ -matching.*
- (d) *There exists a map $f \in \mathcal{B}(V)$ (see Definition 8.20) such that for every $v \in V$ the preimage $f^{-1}(v)$ contains exactly two elements.*
- (d') (Gromov’s condition) *there exists a map $f \in \mathcal{B}(V)$ such that for every $v \in V$ the pre-image $f^{-1}(v)$ contains at least two elements.*

PROOF. Let $m \geq 1$ denote the valence of the graph \mathcal{G} .

(b) \iff (b'). Observe that for every $\alpha > 1, C > 0$,

$$\forall \Phi, |\overline{\mathcal{N}}_C(\Phi)| \geq \alpha|\Phi| \implies \forall k \in \mathbb{N}, |\overline{\mathcal{N}}_{kC}(\Phi)| \geq \alpha^k|\Phi|.$$

Therefore, (b) \iff (b').

(a) \implies (b'). The graph \mathcal{G} is non-amenable if and only if its Cheeger constant is positive. In other words, there exists $\eta > 0$ such that for every finite set of vertices F , $|E(F, F^c)| \geq \eta|\Phi|$. This implies that $\overline{\mathcal{N}}_1(\Phi)$ contains at least $(1 + \frac{\eta}{m})|\Phi|$ vertices.

(b) \implies (c). Let C be the constant as in the expansion property. We form the bipartite graph $Bip_C(Y, Z)$, where Y, Z are two copies of V . Clearly, the graph $Bip_C(Y, Z)$ is locally finite. For any finite subset A in V , since $|\overline{\mathcal{N}}_C(A) \cap V| \geq 2|A|$, it follows that the edge-boundary of A in $Bip_C(Y, Z)$ has at least $2|A|$ elements, where we embed A in either one of the copies of V in $Bip_C(Y, Z)$. It follows by Theorem 18.3 that $Bip_C(Y, Z)$ has a perfect $(2, 1)$ -matching.

(c) \implies (d). The matching in (c) defines a map $f : Z = V \rightarrow Y = V$, so that $\text{dist}_{\mathcal{G}}(z, f(z)) \leq C$. Hence, $f \in \mathcal{B}(V)$ and $|f^{-1}(y)| = 2$ for every $y \in V$.

The implication (d) \implies (d') is obvious. We show that (d') \implies (b). According to (d'), there exists a constant $M > 0$ and a map $f : V \rightarrow V$ such that for every $x \in V$, $\text{dist}(x, f(x)) \leq M$, and $|f^{-1}(y)| \geq 2$ for every $y \in V$. For every finite non-empty set $F \subset V$, $f^{-1}(\Phi)$ is contained in $\mathcal{N}_M(\Phi)$ and it has at least twice as many elements. Thus, (b) is satisfied.

Thus, we proved that the properties (b) through (d') are equivalent.

It remains to be shown that (b) \implies (a). By hypothesis, there exists a constant C such that for every finite non-empty subset $\Phi \subset V$, $|\overline{\mathcal{N}}_C(\Phi) \cap V| \geq 2|\Phi|$. Without loss of generality, we may assume that C is a positive integer. Recall that $\partial_V \Phi$ is the vertex-boundary of the subset $\Phi \subset V$. Since $\overline{\mathcal{N}}_C(\Phi) = \Phi \cup \mathcal{N}_C(\partial_V \Phi)$, it follows that $|\mathcal{N}_C(\partial_V \Phi) \setminus \Phi| \geq |\Phi|$.

Recall that the graph \mathcal{G} has finite valence $m \geq 1$. Therefore,

$$|\overline{\mathcal{N}}_C(\partial_V \Phi)| \leq m^C |\partial_V \Phi|.$$

We have, thus, obtained that for every finite non-empty set $\Phi \subset V$,

$$|E(\Phi, \Phi^c)| \geq |\partial_V \Phi| \geq \frac{1}{m^C} |\mathcal{N}_C(\partial_V \Phi)| \geq \frac{1}{m^C} |\Phi|.$$

Therefore, the Cheeger constant of \mathcal{G} is at least $\frac{1}{m^C} > 0$, and the graph is non-amenable. \square

EXERCISE 18.5. Show that a sequence $\Phi_n \subset V$ is Følner if and only if for every $C \in \mathbb{R}_+$

$$\lim_{n \rightarrow \infty} \frac{|\mathcal{N}_C(\Phi_n)|}{|\Phi_n|} = 1.$$

Some graphs with bounded geometry admit Følner sequences which consist of metric balls. A proof of the following property (in the context of Cayley graphs) first appeared in [AVS57].

PROPOSITION 18.6. *A graph \mathcal{G} of bounded geometry and sub-exponential growth (in the sense of Definition 8.77) is amenable and has the property that for every basepoint $v_0 \in V$ (where V is the vertex set of \mathcal{G}) there exists a Følner sequence consisting of metric balls with center v_0 .*

PROOF. Let v_0 be an arbitrary vertex in \mathcal{G} . We equip the vertex set V of \mathcal{G} with the restriction of the standard metric on \mathcal{G} and set

$$\mathfrak{G}_{v_0, V}(n) = |\bar{B}(v_0, n)|,$$

here and in what follows $\bar{B}(x, n)$ is the ball of center x and radius x in V . Our goal is to show that for every $\varepsilon > 0$ there exists a radius R_ε such that $\partial_V \bar{B}(v_0, R_\varepsilon)$ has cardinality at most $\varepsilon |\bar{B}(v_0, R_\varepsilon)|$.

We argue by contradiction and assume that there exists $\varepsilon > 0$ such that for every integer $R > 0$,

$$|\partial_V \bar{B}(v_0, R)| \geq \varepsilon |\bar{B}(v_0, R)|.$$

(Since \mathcal{G} has bounded geometry, considering vertex–boundary is equivalent to considering the edge–boundary.) This inequality implies that

$$|\bar{B}(v_0, R+1)| \geq (1+\varepsilon)|\bar{B}(v_0, R)|.$$

Applying the latter inequality inductively we obtain

$$\forall n \in \mathbb{N}, \quad |\bar{B}(v_0, n)| \geq (1+\varepsilon)^n,$$

whence

$$\limsup_{n \rightarrow \infty} \frac{\ln \mathfrak{G}_{v_0, V}}{n} \geq \ln(1+\varepsilon) > 0.$$

This contradicts the assumption that \mathcal{G} has sub-exponential growth. \square

LEMMA 18.7 (K. Whyte, [Why99]). *Suppose that \mathcal{G} is a non-amenable graph of finite valence.*

Let $r > 0$ and let V' be a subset in $V = V(\mathcal{G})$ that is r -dense in V , with the terminology in Definition 2.20.

Then there exists a bijection $f : V' \rightarrow V$ which is a bounded perturbation of the inclusion $V' \rightarrow V$: There exists $D < \infty$ such that

$$\text{dist}(x, f(x)) \leq D$$

for all $x \in V'$.

PROOF. Without loss of generality we may assume that r is an integer.

Let $m \in \mathbb{N}$ be the valence of the graph \mathcal{G} . Since \mathcal{G} is non-amenable, there exists a constant $C > 0$ such that for every finite non-empty subset Φ of V , $\bar{N}_C(\Phi) \cap V$ has cardinality at least m^{2r} times the cardinality of Φ . We take $D := C + 2r$ and the bipartite graph $\text{Bip}_D(V', V)$.

Clearly, for every finite subset $A \subset V'$,

$$|\partial^V A| \geq |A|,$$

Let B be an arbitrary finite subset in V and let $B' = V' \cap \bar{N}_r(B)$. Since $B \subset \bar{N}_r(B')$, $|B| \leq m^r |B'|$.

Theorem 18.4, (b'), implies that $\bar{N}_C(B') \cap V$ has cardinality at least $m^{2r} |B'| \geq m^r |B|$. The argument above, with B replaced by $\bar{N}_C(B') \cap V$ implies that $\bar{N}_{C+r}(B') \cap V'$ has cardinality at least $\frac{1}{m^r}$ times the cardinality of $\bar{N}_C(B') \cap V$, hence at least $|B|$. It follows that $\bar{N}_{C+2r}(B) \cap V'$ has cardinality at least $|B|$. In other words, in the bipartite graph $\text{Bip}_D(V', V)$,

$$|\partial^V B| \geq |B|.$$

By the Hall–Rado Marriage Theorem, there exists a bijection $f : V' \rightarrow V$ sending each $v' \in V'$ to a vertex $v = f(v')$ within distance $\leq D$ from v' . This map f is the required bijection. \square

REMARK 18.8. The map f in Lemma 18.7 is $(2D + 1)$ -bi-Lipschitz. Indeed

$$\begin{aligned} \text{dist}(f(a), f(b)) &\leq \text{dist}(a, b) + 2D \leq (2D + 1)\text{dist}(a, b); \\ \text{dist}(a, b) &\leq \text{dist}(f(a), f(b)) + 2D \leq (2D + 1)\text{dist}(f(a), f(b)). \end{aligned}$$

The next theorem is also due to K. Whyte, although it was implicit in [DSS95]:

THEOREM 18.9 (K. Whyte [Why99]). *Let $\mathcal{G}_i, i = 1, 2$, be two non-amenable graphs of bounded geometry. Then every quasiisometry $h : \mathcal{G}_1 \rightarrow \mathcal{G}_2$ is at bounded distance from a bi-Lipschitz map $g : V(\mathcal{G}_1) \rightarrow V(\mathcal{G}_2)$.*

PROOF. According to the proof of Proposition 8.17, there exist V'_i separated nets in $\mathcal{G}_i, i = 1, 2$, (with the terminology of Definition 2.21) such that, given the inclusion map $i : V'_1 \rightarrow \mathcal{G}_1$, the composition map $h \circ i$ is a bounded perturbation of a bi-Lipschitz map $h' : V'_1 \rightarrow V'_2$.

Lemma 18.7 implies the existence of bijections

$$f_i : V'_i \rightarrow V(\mathcal{G}_i), \quad i = 1, 2,$$

which are bounded perturbations of the corresponding inclusion maps. The composition

$$g := f_2 \circ h' \circ f_1^{-1} : V(\mathcal{G}_1) \rightarrow V(\mathcal{G}_2)$$

is the required bi-Lipschitz map. \square

For the sake of completeness we mention without proof two more properties equivalent to those in Theorem 18.4.

The first will turn out to be relevant to a discussion later on between non-amenable and existence of free sub-groups (the von Neumann-Day Question 18.71).

THEOREM 18.10 (Theorem 1.3 in [Why99]). *Let \mathcal{G} be an infinite connected graph of bounded geometry. The graph \mathcal{G} is non-amenable if and only if there exists a free action of a free group of rank two on \mathcal{G} by bi-Lipschitz maps which are at finite distance from the identity.*

The second property is related to probability on graphs. Let \mathcal{G} be an infinite locally finite connected graph with set of vertices V and set of edges E . For every vertex x of \mathcal{G} we denote by $\text{val}(x)$ the valency at the vertex x . We refer the reader to [Bre92, DS84, Woe00] for the definition of Markov chains and detailed treatment of random walks on graphs and groups.

A simple random walk on \mathcal{G} is a Markov chain with random variables

$$X_1, X_2, \dots, X_n, \dots$$

on V , with the transition probability $p(x, y) = \frac{1}{\text{val}(x)}$ if x and y are two vertices joined by an edge, and $p(x, y) = 0$ if x and y are not joined by an edge.

We denote by $p_n(x, y)$ the probability that a random walk starting in x will be at y after n steps. The spectral radius of the graph \mathcal{G} is defined by

$$\rho(\mathcal{G}) = \limsup_{n \rightarrow \infty} [p_n(x, y)]^{\frac{1}{n}}.$$

It can be easily checked that the spectral radius does not depend on x and y .

THEOREM 18.11 (J. Dodziuk, [Dod84]). *A graph of bounded geometry is non-amenable if and only if $\rho(\mathcal{G}) < 1$.*

Note that in the case of countable groups the corresponding theorem was proved by H. Kesten [Kes59].

COROLLARY 18.12. *In a non-amenable graph of bounded geometry, the simple random walk is transient, that is, for every $x, y \in V$,*

$$\sum_{n=1}^{\infty} p_n(x, y) < \infty.$$

18.2. Amenability and quasiisometry

THEOREM 18.13 (Graph amenability is QI invariant). *Suppose that \mathcal{G} and \mathcal{G}' are quasiisometric graphs of bounded geometry. Then \mathcal{G} is amenable if and only if \mathcal{G}' is.*

PROOF. We will show that non-amenableity is a quasiisometry invariant. We will assume that both \mathcal{G} and \mathcal{G}' are infinite, otherwise the assertion is clear. Note that according to Theorem 18.4, Part (b), nonamenability is equivalent to existence of a constant $C > 0$ such that for every finite non-empty set F of vertices, its closed neighborhood $\bar{\mathcal{N}}_C(\Phi)$ contains at least $2|\Phi|$ vertices.

Let V and V' be the vertex sets of graphs \mathcal{G} and \mathcal{G}' respectively. We assume that V, V' are endowed with the metrics obtained by restriction of the standard metrics on the respective graphs. Let $m < \infty$ be an upper bound on the valence of graphs $\mathcal{G}, \mathcal{G}'$. Let $f : V \rightarrow V'$ and $g : V' \rightarrow V$ be L -Lipschitz maps that are coarse inverses to each other:

$$\text{dist}(f \circ g, \text{Id}) \leq A, \quad \text{dist}(g \circ f, \text{Id}) \leq A.$$

Assume that \mathcal{G}' is amenable. Given a finite set F in V , consider

$$F \xrightarrow{f} F' = f(\Phi) \xrightarrow{g} F'' = g(F').$$

Since F'' is at Hausdorff distance $\leq A$ from F , it follows that $|\Phi| \leq b|F''|$, where $b = m^L$. In particular,

$$|f(\Phi)| \geq b^{-1}|\Phi|.$$

Likewise, for every finite set Φ' in V' we obtain

$$|g(\Phi')| \geq b^{-1}|\Phi'|.$$

By Theorem 18.4 (Part (b')), for every number $\alpha > b^2$, there exists $C \geq 1$ such that for an arbitrary finite set $F' \subset V'$, its neighborhood $\bar{\mathcal{N}}_C(F')$ contains at least $\alpha|F'|$ vertices. Therefore, for such C , the set $g(\bar{\mathcal{N}}_C(F'))$ contains at least

$$\frac{1}{b}|\mathcal{N}_C(F')| \geq \frac{\alpha}{b}|F'|$$

elements.

Pick a finite non-empty subset $\Phi \subset V$ and set $\Phi' := f(\Phi)$, $F'' = g f(\Phi)$. Then $|F''| \geq b^{-1}|\Phi|$ and, therefore,

$$|g(\bar{\mathcal{N}}_C(F'))| \geq \frac{\alpha}{b^2}|\Phi|.$$

Since g is L -Lipschitz,

$$g(\bar{\mathcal{N}}_C(F')) \subset \bar{\mathcal{N}}_{LC}(F'') \subset \bar{\mathcal{N}}_{LC+A}(\Phi).$$

We conclude that

$$|\bar{\mathcal{N}}_{LC+A}(\Phi)| \geq \frac{\alpha}{b^2} |\Phi|.$$

Setting $C' := LC + A$, and $\beta := \frac{\alpha}{b^2} > 1$, we conclude that \mathcal{G} satisfies the expansion property (b') in Theorem 18.4. Hence, \mathcal{G} is also non-amenable. \square

We will see below that this theorem generalizes in the context connected Riemannian manifolds M of bounded geometry and graphs \mathcal{G} obtained by discretization of M , and, thus, quasiisometric to M . More precisely, we will see that non-amenability of the graph is equivalent to positivity of the Cheeger constant of the manifold (see Definition 3.21). This may be seen as a version within the setting of amenability/isoperimetric problem of the Milnor–Efremovich–Schwartz Theorem 8.80 stating that the growth functions of M and \mathcal{G} are equivalent.

In what follows we use the terminology in Definitions 3.26 and 3.33 for the bounded geometry of a Riemannian manifold, respectively of a simplicial graph.

THEOREM 18.14. *Let M be a complete connected n -dimensional Riemannian manifold and \mathcal{G} a simplicial graph, both of bounded geometry. Assume that M is quasiisometric to \mathcal{G} . Then the Cheeger constant of M is positive if and only if the graph \mathcal{G} is non-amenable.*

REMARKS 18.15. (1) Theorem 18.14 was proved by R. Brooks [**Bro82a**], [**Bro81a**] in the special case when M is the universal cover of a compact Riemannian manifold and \mathcal{G} is the a Cayley graph of the fundamental group of this compact manifold.

(2) A more general version of Theorem 18.14 requires a weaker condition of bounded geometry for the manifold than the one used in this book. See for instance [**Gro93**], Proposition 0.5.A₅. A proof of that result can be obtained by combining the main theorem in [**Pan95**] and Proposition 11 in [**Pan07**].

PROOF. Since M has bounded geometry it follows that its sectional curvature is at least a and at most b , for some $b \geq a$. It also follows that the injectivity radius at every point of M is at least ρ , for some $\rho > 0$.

As in Theorem 3.23, we let $V_\kappa(r)$ denote the volume of ball of radius r in the n -dimensional space of constant curvature κ .

Choose ε so that $0 < \varepsilon < 2\rho$. Let N be a maximal ε -separated set in M .

It follows that $\mathcal{U} = \{B(x, \varepsilon) \mid x \in N\}$ is a covering of M , and by Lemma 3.31, (2), its multiplicity is at most

$$m = \frac{V_a\left(\frac{3\varepsilon}{2}\right)}{V_b\left(\frac{\varepsilon}{2}\right)}.$$

We now consider the restriction of the Riemannian distance function on M to the subset N . Define the Rips complex $\text{Rips}_{8\varepsilon}(N)$ (with respect to this metric on N), and the 1-dimensional skeleton of the Rips complex, the graph \mathcal{G}_ε . According to Theorem 8.52, the manifold M is quasiisometric to \mathcal{G}_ε . Furthermore, \mathcal{G}_ε has bounded geometry as well. This and Theorem 18.13 imply that \mathcal{G}_ε has positive

Cheeger constant if and only if \mathcal{G} has. Thus, it suffices to prove the equivalence in Theorem 18.14 for the graph $\mathcal{G} = \mathcal{G}_\epsilon$.

Assume that M has positive Cheeger constant. This means that there exists $h > 0$ such that for every open submanifold $\Omega \subset M$ with compact closure and smooth boundary,

$$\text{Area}(\partial\Omega) \geq h \text{Vol}(\Omega).$$

Our goal is to show that there exist uniform positive constants B and C such that for every finite subset $F \subset N$ there exists an open submanifold with compact closure and smooth boundary Ω , such that (with the notation in Definition 1.43),

$$(18.2) \quad |E(F, F^c)| \geq B \text{Area}(\partial\Omega) \quad \text{and} \quad C \text{Vol}(\Omega) \geq |F|.$$

Then it would follow that

$$|E(F, F^c)| \geq \frac{Bh}{C} |F|,$$

i.e. \mathcal{G} would be non-amenable. Here, as usual, $F^c = N \setminus F$.

Since M has bounded geometry, the open cover \mathcal{U} admits a smooth partition of unity $\{\varphi_x ; x \in N\}$ in the sense of Definition 3.7, such that all the functions φ_x are L -Lipschitz for some constant $L > 0$ independent of x , see Lemma 3.30. Let $F \subset N$ be a finite subset. Consider the smooth function $\varphi = \sum_{x \in F} \varphi_x$. By hypothesis and since \mathcal{U} has multiplicity at most m , the function Φ is Lm -Lipschitz. Furthermore, in view of Sard's Theorem, since the map φ has compact support, the set Θ of singular values of φ is compact and has Lebesgue measure zero.

For every $t \in (0, 1)$, the preimage

$$\Omega_t = \varphi^{-1}((t, \infty)) \subset M$$

is an open submanifold in M with compact closure. If we choose t to be a regular value of Φ , that is $t \notin \Theta$, then the hypersurface $\varphi^{-1}(t)$, which is the boundary of Ω_t , is smooth (Theorem 3.4).

Since N is ϵ -separated, the balls $B(x, \frac{\epsilon}{2})$, $x \in N$, are pairwise disjoint. Therefore, for every $x \in N$ the function φ_x restricted to $B(x, \frac{\epsilon}{2})$ is identically equal to 1. Hence, the union

$$\bigsqcup_{x \in F} B\left(x, \frac{\epsilon}{2}\right)$$

is contained in Ω_t for every $t \in (0, 1)$, and in view of Part 2 of Theorem 3.23 we get

$$\text{Vol}(\Omega_t) \geq \sum_{x \in F} \text{Vol}\left(x, \frac{\epsilon}{2}\right) \geq |F| \cdot V_b(\epsilon/2).$$

Therefore, for every $t \notin \Theta$, the domain Ω_t satisfies the second inequality in (18.2) with $C^{-1} = V_b(\epsilon/2)$. Our next goal is to find values of $t \notin \Theta$ so that the first inequality in (18.2) holds.

Fix a constant η in the open interval $(0, 1)$, and consider the open set $U = \Phi^{-1}((0, \eta))$.

Let F' be the set of points x in F such that $U \cap \overline{B(x, \epsilon)} \neq \emptyset$. Since for every $y \in U$ there exists $x \in F$ such that $\varphi_x(y) > 0$, it follows that the set of closed balls centered in points of F' and of radius ϵ cover U .

Since $\{\varphi_x : x \in N\}$ is a partition of unity for the cover \mathcal{U} of M , it follows that for every $y \in U$ there exists $z \in N \setminus F$ such that $\varphi_z(y) > 0$, whence $y \in \overline{B(z, \epsilon)}$.

Thus,

$$(18.3) \quad U \subset \left(\bigcup_{x \in F'} \overline{B(x, \varepsilon)} \right) \cap \left(\bigcup_{z \in N \setminus F} \overline{B(z, \varepsilon)} \right).$$

In particular, for every $x \in F'$ there exists $z \in N \setminus F$ such that $\overline{B(x, \varepsilon)} \cap \overline{B(z, \varepsilon)} \neq \emptyset$, whence x and z are connected by an edge in the graph \mathcal{G} .

Thus, every point $x \in F'$ belongs to the vertex-boundary $\partial_V F$ of the subset F of the vertex set of the graph \mathcal{G} . We conclude that $\text{card } F' \leq \text{card } E(F, F^c)$.

Since $|\nabla \Phi| \leq mL$, by the Coarea Theorem 3.14, with $g \equiv 1$, $f = \Phi$ and $U = \Phi^{-1}(0, \eta)$, we obtain:

$$\int_0^\eta \text{Area}(\partial \Omega_t) dt = \int_U |\nabla \varphi| dV \leq mL \text{Vol}(U) \leq mL \sum_{x \in F'} \text{Vol}(B(x, \varepsilon)).$$

The last inequality follows from the inclusion (18.3). At the same time, by applying Theorem 3.23, we obtain that for every $x \in M$

$$V_a(\varepsilon) \geq \text{Vol}(B(x, \varepsilon)).$$

By combining these inequalities, we obtain

$$\int_0^\eta \text{Area}(\partial \Omega_t) dt \leq mL V_a(\varepsilon) |F'| \leq mL V_a(\varepsilon) |E(F, F^c)|.$$

Since Θ has measure zero, it follows that for some $t \in (0, \eta) \setminus \Theta$,

$$\text{Area}(\partial \Omega_t) \leq 2 \frac{m}{\eta} L V_a(\varepsilon) |E(F, F^c)| = B |E(F, F^c)|.$$

This establishes the first inequality in (18.2) and, hence, shows that nonamenability of M implies nonamenability of the graph \mathcal{G} .

We now prove the converse implication. To that end, we assume that for some δ satisfying $2\rho > \delta > 0$, some maximal δ -separated set N and the corresponding graph (of bounded geometry) $\mathcal{G} = \mathcal{G}_\delta$ are constructed as above, so that \mathcal{G} has a positive Cheeger constant. Thus, there exists $h > 0$ such that for every finite subset F in N

$$\text{card } E(F, F^c) \geq h \text{card } F.$$

Let Ω be an arbitrary open bounded subset of M with smooth boundary. Our goal is to find a finite subset Φ_k in N such that for two constants P and Q independent of Ω , we have

$$(18.4) \quad \text{Area}(\partial \Omega) \geq P |E(\Phi_k, \Phi_k^c)| \quad \text{and} \quad |\Phi_k| \geq Q \text{Vol}(\Omega).$$

This would imply positivity of Cheeger constant of M . Note that, since the graph \mathcal{G} has finite valence, in the first inequality of (18.4) we may replace the edge boundary $E(\Phi_k, \Phi_k^c)$ by the vertex boundary $\partial_V \Phi_k$ (see Definition 1.43).

Consider the finite subset F of points $x \in N$ such that $\Omega \cap B(x, \delta) \neq \emptyset$. It follows that $\Omega \subseteq \bigcup_{x \in F} B(x, \delta)$. We split the set F in two parts:

$$(18.5) \quad F_1 = \left\{ x \in F : \text{Vol}[\Omega \cap B(x, \delta)] > \frac{1}{2} \text{Vol}[B(x, \delta)] \right\}$$

and

$$F_2 = \left\{ x \in F : \text{Vol}[\Omega \cap B(x, \delta)] \leq \frac{1}{2} \text{Vol}[B(x, \delta)] \right\}.$$

Set

$$v_k := \text{Vol} \left(\Omega \cap \bigcup_{x \in F_k} B(x, \delta) \right), k = 1, 2.$$

Thus,

$$\max(v_1, v_2) \geq \frac{1}{2} \text{Vol}(\Omega).$$

Case 1: $v_1 \geq \frac{1}{2} \text{Vol}(\Omega)$. In view of Theorem 3.23, this inequality implies that

$$(18.6) \quad \frac{1}{2} \text{Vol}(\Omega) \leq \sum_{x \in F_1} \text{Vol}(B(x, \delta)) \leq |F_1| V_a(\delta).$$

This gives the second inequality in (18.4). A point x in $\partial_V F_1$ is then a point in N satisfying (18.5), such that within distance 8δ of x there exists a point $y \in N$ satisfying the inequality opposite to (18.5). The (unique) shortest geodesic $xy \subset M$ will, therefore, intersect the set of points

$$\text{Half} = \left\{ x \in M ; \text{Vol}[B(x, \delta) \cap \Omega] = \frac{1}{2} \text{Vol}[B(x, \delta)] \right\}.$$

This implies that $\partial_V F_1$ is contained in the 8δ -neighborhood of the set $\text{Half} \subset M$. Given a maximal δ -separated subset H_δ of Half (with respect to the restriction of the Riemannian distance on M), $\partial_V F_1$ will then be contained in the 9δ -neighborhood of H_δ . In particular,

$$\bigsqcup_{x \in \partial_V F_1} B\left(x, \frac{\delta}{2}\right) \subseteq \bigcup_{y \in H_\delta} B(y, 10\delta),$$

whence

$$(18.7) \quad V_b(\delta/2) |\partial_V F_1| \leq \text{Vol} \left[\bigsqcup_{x \in \partial_V F_1} B\left(x, \frac{\delta}{2}\right) \right] \leq \sum_{y \in H_\delta} \text{Vol}[B(y, 10\delta)] \leq V_b(10\delta) |H_\delta|.$$

Since H_δ extends to a maximal δ -separated subset H' of M , Lemma 3.31, (2), implies that the multiplicity of the covering $\{B(x, \delta) \mid x \in H'\}$ is at most $\frac{V_a(\frac{3\delta}{2})}{V_b(\frac{\delta}{2})}$.

It follows that

$$m \cdot \text{Area}(\partial\Omega) \geq \sum_{y \in H_\delta} \text{Area}(\partial\Omega \cap B(y, \delta)).$$

We now apply Buser's Theorem 3.25 and deduce that there exists a constant $\lambda = \lambda(n, a, \delta)$ such that for all $y \in H_\delta$, we have,

$$\lambda \text{Area}(\partial\Omega \cap B(y, \delta)) \geq \text{Vol}[\Omega \cap B(y, \delta)] = \frac{1}{2} \text{Vol}[B(y, \delta)].$$

It follows that

$$\text{Area}(\partial\Omega) \geq \frac{1}{2\lambda m} \sum_{y \in H_\delta} \text{Vol}[B(y, \delta)] \geq \frac{1}{2\lambda m} V_b(\rho) |H_\delta|.$$

Combining this estimate with the inequality (18.7), we conclude that

$$\text{Area}(\partial\Omega) \geq P |\partial_V F_1|,$$

for some constant P independent of Ω .

This establishes the first inequality in (18.4) and, hence, proves positivity of the Cheeger constant of M in the Case 1.

Case 2. Assume now that v_2 is at least $\frac{1}{2}Vol(\Omega)$.

We obtain, using Buser's Theorem 3.25 for the second inequality below, that

$$mArea(\partial\Omega) \geq \sum_{y \in F_2} Area(\partial\Omega \cap B(y, \delta)) \geq \frac{1}{\lambda} \sum_{y \in F_2} Vol[\Omega \cap B(y, \delta)] \geq \frac{1}{2\lambda} Vol(\Omega).$$

Thus, in the Case 2 we obtain the required lower bound on $Area(\partial\Omega)$ directly. \square

COROLLARY 18.16. *Let M and M' be two complete connected Riemann manifolds of bounded geometry which are quasiisometric to each other. Then M has positive Cheeger constant if and only if M' has positive Cheeger constant.*

PROOF. Consider graphs of bounded geometry \mathcal{G} and \mathcal{G}' that are quasiisometric to M and M' respectively. Then $\mathcal{G}, \mathcal{G}'$ are also quasiisometric to each other. The result now follows by combining Theorem 18.14 with Theorem 18.13. \square

An interesting consequence of Corollary 18.16 is the quasiisometric invariance of the positivity of the *spectral gap* of Riemannian manifolds. Recall that $h(M) = 0 \iff \lambda_1(M) = 0$ (Theorem 3.53).

COROLLARY 18.17. *If M and M' are complete connected Riemann manifolds of bounded geometry which are quasiisometric to each other, then $\lambda_1(M) = 0 \iff \lambda_1(M') = 0$.*

18.3. Amenability of groups

Motivated by the Banach-Tarski Paradox, John von Neumann [vN29] studied properties of group actions that make paradoxical decompositions possible and on the contrary, impossible. He defined the notion of *amenable group* G , based on the existence of a mean/finitely additive measure invariant under the action of the group on itself¹, and equivalent to the *nonexistence* of a G -paradoxical decomposition for any space on which the group acts. One can ask furthermore that no subset has a paradoxical decomposition, for any space endowed with an action of the group. This defines a strictly smaller class, that of *supramenable groups*; such groups will be discussed in section 18.6.

In this section we define amenable actions and amenable groups, and prove that paradoxical behavior is equivalent to non-amenability. For simplicity of the discussion (and since it is the most relevant for geometric group theory), we only consider amenability in the context of discrete groups and group actions on sets (rather than continuous group actions on topological spaces). We refer the reader to Section 1.2.1 for the discussion on finitely additive probability measures (f.a.p. measures) on sets, and finitely additive integrals. Later in the chapter we relate amenability and paradoxical decompositions and prove (among other things) that, for finitely generated groups, amenability is equivalent to amenability of its Cayley graph (Theorem 18.50).

¹Von Neumann called these groups *measurable*.

Let $\mu : G \times X \rightarrow X$ be a left group action, $\mu(g, x) = g(x)$ on the a set X (for right group actions the discussion is very similar).

- DEFINITION 18.18. (1) A group action $G \curvearrowright X$ on a set X is *amenable* if there exists a G -invariant f.a.p. measure μ on $\mathcal{P}(X) = 2^X$, the set of all subsets of X .
- (2) A group G is *amenable* if the action of G on itself by left multiplication is amenable.

Yet another (more common) equivalent definition for amenability is formulated using the concept of *invariant mean*, which is responsible for the terminology ‘amenable’:

DEFINITION 18.19. A *mean* on a set X is a linear functional

$$m : B(X) \rightarrow \mathbb{R}$$

defined on the vector space $B(X)$ of bounded real-valued functions on X , satisfying the following properties:

- (M1) If $f \geq 0$ on X , then $m(f) \geq 0$.
(M2) $m(\mathbf{1}_X) = 1$.

Assume, moreover, that X is endowed with an action of a group G , $G \times X \rightarrow X$, $(g, x) \mapsto g \cdot x$. This induces an action of G on the vector space $B(X)$ defined by $g \cdot f(x) = f(g^{-1} \cdot x)$. A mean m on X is called *invariant* if $m(g \cdot f) = m(f)$ for every $f \in B(X)$ and $g \in G$.

PROPOSITION 18.20. A group action $G \curvearrowright X$ is amenable (in the sense of Definition 18.18) if and only if it admits an invariant mean.

PROOF. According to Theorem 1.12 each G -invariant f.a.p. measure μ on X defines a G -invariant integral

$$m : B(X) \rightarrow \mathbb{R}, \quad m(f) = \int_X f d\mu.$$

Since the integral \int_X is a linear functional nonnegative on nonnegative functions and satisfying

$$\int_X \mathbf{1}_X d\mu = 1,$$

the functional m is a G -invariant mean on X . Conversely, each G -invariant mean m on X , defines a G -invariant f.a.p. measure μ on X by $\mu(A) = m(\mathbf{1}_A)$. \square

EXAMPLE 18.21. If X is a finite nonempty set, then every group action $G \curvearrowright X$ is amenable. In particular, every finite group is amenable. Indeed, for a finite set X define $\mu : \mathcal{P}(X) \rightarrow [0, 1]$ by $\mu(A) = \frac{|A|}{|X|}$, where $|\cdot|$ denotes the cardinality of a subset.

We now specialize to the case $X = G$. A group G has two actions on itself, the action L by left multiplication and the action R by right multiplication

$$R : G \times G \rightarrow G, \quad R(g, x) = R_g(x) = xg.$$

DEFINITION 18.22. A *left-invariant mean* on G is a mean invariant under the action L ; a *right-invariant mean* on G is a mean invariant under the action R .

The following lemma shows that different notions of invariance for measures and means leads to the same class of groups:

PROPOSITION 18.23. *The following are equivalent:*

- (a) *G is amenable.*
- (b) *G has a right-invariant f.a.p. measure.*
- (c) *G has a right-invariant mean.*

PROOF. (a) \iff (b). Given a f.a.p. measure μ_L on G , we define a measure μ_R on G by

$$\mu_R(A) := \mu_L(A^{-1}), \quad A^{-1} = \{a^{-1} : a \in A\}.$$

Then μ_L is left-invariant iff μ_R is right-invariant. \square

In view of this proposition, by default, an invariant mean on a group G will mean a left-invariant mean.

LEMMA 18.24. *Every action $G \curvearrowright X$ of an amenable group is also amenable.*

PROOF. Let μ be an invariant measure on G . Given an action $G \curvearrowright X$, choose a point $x \in X$ and define a function $\nu : \mathcal{P}(X) \rightarrow [0, 1]$ by

$$\nu(A) = \mu(\{g \in G ; gx \in A\}).$$

We leave it to the reader to verify that ν is a G -invariant f.a.p. measure. \square

QUESTION 18.25. Suppose that G is a group which admits a mean $m : B(G) \rightarrow \mathbb{R}$ that is quasiinvariant, i.e. there exists a constant κ such that

$$|m(f \circ L_{g^{-1}}) - m(f)| \leq \kappa$$

for all functions $f \in B(G)$ and all group elements g . Is it true that G is amenable?

We refer the reader to Section 17.1 for the definitions of paradoxical sets, group actions and groups.

LEMMA 18.26. *A paradoxical action $G \curvearrowright X$ cannot be amenable.*

PROOF. Suppose to the contrary that X admits a G -invariant f.a.p. measure μ and

$$X = X_1 \sqcup \dots \sqcup X_k \sqcup Y_1 \sqcup \dots \sqcup Y_m$$

is a G -paradoxical decomposition, i.e. for some $g_1, \dots, g_k, h_1, \dots, h_m \in G$,

$$g_1(X_1) \sqcup \dots \sqcup g_k(X_k) = X \quad \text{and} \quad h_1(Y_1) \sqcup \dots \sqcup h_m(Y_m) = X.$$

Then

$$\mu(X_1 \sqcup \dots \sqcup X_k) = \mu(Y_1 \sqcup \dots \sqcup Y_m) = \mu(X),$$

which implies that $2\mu(X) = \mu(X)$, contradicting the fact that $\mu(X) = 1$. \square

COROLLARY 18.27. *A paradoxical group cannot be amenable.*

EXAMPLE 18.28. The free group of rank two F_2 is non-amenable since F_2 is paradoxical (F_2 acts paradoxically on itself), as explained in Section 17.2.

We will prove in Theorem 18.50 that a finitely generated group is amenable if and only if it is non-paradoxical.

The next theorem summarizes basic properties of amenable groups:

THEOREM 18.29. (1) *Each subgroup of an amenable group is amenable.*

- (2) Let N be a normal subgroup of a group G . The group G is amenable if and only if both N and G/N are amenable.
- (3) The direct limit G (see section 1.1) of a directed system $(H_i)_{i \in I}$ of amenable groups H_i , is amenable.

PROOF. (1) Let μ be a f.a.p. measure on an amenable group G , and let H be a subgroup. By the Axiom of Choice, there exists a subset D of G intersecting each right coset Hg in exactly one point. Then $\nu(A) := \mu(AD)$ defines a left-invariant f.a.p. measure on H .

(2) “ \Rightarrow ” Assume that G is amenable and let μ be a f.a.p. measure on G . The subgroup N is amenable according to (1). The amenability of G/N follows from Lemma 18.24, since G acts on G/N by left multiplication.

(2) “ \Leftarrow ” Let ν be a left-invariant f.a.p. measure on G/N , and λ a left-invariant f.a.p. measure on N . On every left coset gN one defines a f.a.p. measure by $\lambda_g(A) = \lambda(g^{-1}A)$. The H -left-invariance of λ implies that λ_g is independent of the representative g , i.e. $gN = g'N \Rightarrow \lambda_g = \lambda_{g'}$.

For every subset B in G define

$$\mu(B) = \int_{G/N} \lambda_g(B \cap gN) d\nu(gN).$$

Then μ is an invariant f.a.p. measure on G .

(3) Let $h_{ij} : H_i \rightarrow H_j, i \leq j$, be the homomorphisms defining the direct system of groups (H_i) and let G be the direct limit. Let $h_i : H_i \rightarrow G$ be the homomorphisms to the direct limit, as defined in Section 1.5. The set of functions

$$\{f : \mathcal{P}(G) \rightarrow [0, 1]\} = \prod_{\mathcal{P}(G)} [0, 1]$$

is compact according to Tychonoff’s theorem (see Remark 10.2, Part 5).

Note that each group H_i acts naturally on G by left multiplication *via* the homomorphism $h_i : H_i \rightarrow G$. For each $i \in I$ let \mathcal{M}_i be the set of H_i -left-invariant f.a.p. measures μ on $\mathcal{P}(G)$. Since H_i is amenable, Lemma 18.24 implies that the set \mathcal{M}_i is non-empty.

We claim that the subset \mathcal{M}_i is closed in $\prod_{\mathcal{P}(G)} [0, 1]$. Let $f : \mathcal{P}(G) \rightarrow [0, 1]$ be an element of $\prod_{\mathcal{P}(G)} [0, 1]$ in the closure of \mathcal{M}_i . Then, for every finite collection A_1, \dots, A_n of subsets of G and every $\epsilon > 0$ there exists μ in \mathcal{M}_i such that

$$|f(A_j) - \mu(A_j)| \leq \epsilon$$

for every $j \in \{1, 2, \dots, n\}$. This implies that for every $\epsilon > 0, |f(G) - 1| \leq \epsilon,$

$$|f(A \sqcup B) - f(A) - f(B)| \leq 3\epsilon$$

and

$$|f(gA) - f(A)| \leq 2\epsilon,$$

$\forall A, B \in \mathcal{P}(G)$ and $g \in H_i$. By letting $\epsilon \rightarrow 0$ we obtain that $f \in \mathcal{M}_i$. Thus, the subset \mathcal{M}_i is indeed closed.

By the definition of compactness, if $\{V_i : i \in I\}$ is a family of closed subsets of a compact space X such that $\bigcap_{j \in J} V_j \neq \emptyset$ for every finite subset $J \subseteq I$, then $\bigcap_{i \in I} V_i \neq \emptyset$.

Consider a finite subset J of I . Since I is a directed set, there exists $k \in I$ such that $j \leq k, \forall j \in J$. Hence, we have homomorphisms $h_{jk} : H_j \rightarrow H_k, \forall j \in J$, and

all homomorphisms $h_j : H_j \rightarrow G$ factor through $h_k : H_k \rightarrow G$. Thus, $\bigcap_{j \in J} \mathcal{M}_j$ contains \mathcal{M}_k , in particular, this intersection is non-empty. It follows from the above that $\bigcap_{i \in I} \mathcal{M}_i$ is non-empty. Every element μ of this intersection is clearly a f.a.p. measure, and μ is also G -left-invariant because

$$G = \bigcup_{i \in I} h_i(H_i). \quad \square$$

Below are several corollaries of this theorem.

COROLLARY 18.30. *Let G_1 and G_2 be two groups that are co-embeddable in the sense of Definition 5.13. Then G_1 is amenable if and only if G_2 is amenable.*

COROLLARY 18.31. *Any group containing a nonabelian free subgroup is non-amenable.*

The next corollary follows immediately from Part (2) of Theorem 18.29:

COROLLARY 18.32. *A semidirect product $N \rtimes H$ is amenable if and only if both N and H are amenable.*

COROLLARY 18.33. *1. If $G_i, i = 1, \dots, n$, are amenable groups, then the finite Cartesian product $G = G_1 \times \dots \times G_n$ is also amenable.*

2. Any direct sum $G = \bigoplus_{i \in I} G_i$ of amenable groups is again amenable.

PROOF. 1. The statement follows from an inductive application of Corollary 18.32.

2. This is a combination of Part 1 and the fact that G is isomorphic to a direct limit of finite direct products of the groups G_i . \square

EXAMPLE 18.34 (Infinite direct products of amenable groups need not be amenable). Let $F = F_2$ be the free group of rank 2. Recall (Corollary 7.112) that F is residually finite, hence, for every $g \in F \setminus \{1\}$ there exists a homomorphism $\varphi_g : F \rightarrow \Phi_g$ such that $\varphi_g(g) \neq 1$ and Φ_g is a finite group. Each Φ_g is, of course, amenable. Consider the direct product of these finite groups:

$$G = \prod_{g \in F - \{1\}} \Phi_g.$$

Then the product of homomorphisms $\varphi_g : F \rightarrow \Phi_g$, defines a homomorphism $\varphi : F \rightarrow G$. This homomorphism is injective since for every $g \neq 1$, $\varphi_g(g) \neq 1$. Thus, G cannot be amenable since it contains a nonamenable subgroup, namely, $\varphi(F)$.

COROLLARY 18.35. *Amenability is preserved by virtual isomorphisms of groups.*

PROOF. Suppose that $G/N \cong Q$ with finite normal subgroup $N \triangleleft G$. Since finite groups are amenable, Part (2) of Theorem 18.29 implies that G is amenable if and only if Q is.

Suppose that H is a finite-index subgroup of a group G . Then H contains a subgroup $N \triangleleft G$ which has finite index in G . Therefore, G is amenable if and only if N is. If G is amenable, so is H . If H is amenable, then N is amenable, which implies that G is amenable. \square

COROLLARY 18.36. *A group G is amenable if and only if all finitely generated subgroups of G are amenable.*

PROOF. The direct implication follows from Theorem 18.29, Part (1). The converse implication follows from (3), where, given the group G , we let I be the directed set of all the finite subsets in G (ordered by the inclusion), and for each $i \in I$, H_i is the subgroup of G generated by the elements in i . We define the directed system of groups (H_i) by letting $h_{ij} : H_i \rightarrow H_j$ be the natural inclusion map whenever $i \subset j$. Then G is the direct limit of the system (H_i) and the assertion follows from Theorem 18.29, Part (3). \square

COROLLARY 18.37. *Direct limits of direct systems of finite groups are amenable.*

PROOF. Since each finite group is amenable, the corollary follows from part (3) of Theorem 18.29.

In order to get more examples of amenable groups, we have to bring geometry into the discussion; this is done by introducing the *Følner sequence criterion of amenability* of groups, discussed in the next section and thereby connecting amenability of groups with amenability of graphs.

18.4. Følner property

Suppose that $R : X \times G \rightarrow X$ is a *right action* of a group G on a set X . Given subsets $E \subset X$, $K \subset G$ we let EK denote the subset

$$EK = \{R(x, g) = xg : x \in E, k \in K\} \subset X.$$

DEFINITION 18.38. A sequence of non-empty subsets $\Omega_n \subset X$ is called a *Følner sequence* for the right action $X \times G \rightarrow X$ if for every $g \in G$

$$(18.8) \quad \lim_{n \rightarrow \infty} \frac{|\Omega_n g \Delta \Omega_n|}{|\Omega_n|} = 0.$$

A sequence of subsets $\Omega_n \subset G$ is called a *Følner sequence* in G if it is a *Følner sequence* with respect to the right action of G on itself by the right multiplication:

$$R_g(x) = xg, g, x \in G.$$

For instance, suppose that $G \simeq \mathbb{Z} = X$ and G acts on itself via addition. Then the sequence of intervals $\Omega_n \subset \mathbb{Z}$ of length diverging to infinity is a *Følner sequence* for this group action.

EXERCISE 18.39. Prove that the subsets $\Omega_n = \mathbb{Z}^k \cap [-n, n]^k$ form a Følner sequence for \mathbb{Z}^k .

EXERCISE 18.40. The following are equivalent for a sequence of non-empty subsets $\Omega_n \subset X$:

- (1) Ω_n is a Følner sequence.
- (2) For every finite subset $K \subset G$

$$(18.9) \quad \lim_{n \rightarrow \infty} \frac{|\Omega_n K \Delta \Omega_n|}{|\Omega_n|} = 0.$$

- (3) For every symmetric finite subset $K \subset G$, (18.9) holds. (Recall that K is symmetric if $K = K^{-1}$.)

EXERCISE 18.41. A countable group G admits a Følner sequence if and only if G admits a Følner sequence Φ_n such that

$$\bigcup_{n \in \mathbb{N}} \Phi_n = G.$$

LEMMA 18.42. *Let G be finitely generated, let S and S' be two symmetric finite generating sets, and let \mathcal{G} and \mathcal{G}' be the Cayley graphs of G with respect to S and S' , respectively.*

A sequence of non-empty subsets $\Omega_n \subset G$ is a Følner sequence in the graph \mathcal{G} , in the sense of Definition 18.1, if and only if it is a Følner sequence in the graph \mathcal{G}' .

PROOF. We will prove only one implication, the other one is symmetric.

Let C be a positive number strictly larger than $|s|_{S'}$, $s \in S$, and $|s'|_S$, $s' \in S'$. Let Ω be a subset of G , i.e. a subset of the vertex sets of both \mathcal{G} and \mathcal{G}' .

In \mathcal{G} ,

$$(18.10) \quad \partial^V \Omega = \bigcup_{s \in S} (\Omega s \setminus \Omega).$$

It follows that in \mathcal{G}' , $\partial^V \Omega$ is contained in $\mathcal{N}_C(\Omega) \setminus \Omega$. In view of Exercise 18.5, this implies that if Ω_n is a Følner sequence in \mathcal{G}' , then

$$\lim_{n \rightarrow \infty} \frac{|\partial^V \Omega_n|}{|\Omega_n|} = 0,$$

and the latter is the definition of a Følner sequence in the graph \mathcal{G} . \square

LEMMA 18.43. *Suppose that G is finitely generated with symmetric finite generating set S and \mathcal{G} is the Cayley graph of G with respect to this generating set. Then the following properties of a sequence of non-empty subsets $\Omega_n \subset G$ are equivalent:*

(1)

$$(18.11) \quad \lim_{n \rightarrow \infty} \frac{|\Omega_n S \Delta \Omega_n|}{|\Omega_n|} = 0.$$

(2) Ω_n is a Følner sequence in the graph \mathcal{G} in the sense of Definition 18.1.

(3) Ω_n is a Følner sequence in G .

PROOF. Given Ω a subset of G , i.e. a set of vertices in \mathcal{G} , the equality (18.10) combined with

$$\partial_V \Omega = \bigcup_{s \in S} (\Omega \setminus \Omega s)$$

implies that

$$\Omega S \Delta \Omega = \partial^V \Omega \cup \partial_V \Omega.$$

Therefore (since \mathcal{G} has finite valence)

$$\lim_{n \rightarrow \infty} \frac{|\Omega_n S \Delta \Omega_n|}{|\Omega_n|} = 0$$

is equivalent to

$$\lim_{n \rightarrow \infty} \frac{|\partial^V \Omega_n|}{|\Omega_n|} = 0.$$

It remains to show that (1) implies that (18.8) holds for each $g \in G$. By Lemma 18.42, the sequence Ω_n is also Følner in the Cayley graph \mathcal{G}' with respect to the generating set $S' = S \cup \{g\}$, whence we obtain the desired conclusion. \square

DEFINITION 18.44. 1. A group action $X \times G \rightarrow X$ is said to satisfy the *Følner Property* if it admits a Følner sequence $\Omega_n \subset X$.

2. A group G is said to have the *Følner Property* if G contains a Følner sequence.

LEMMA 18.45. *A group G has the Følner Property if and only if for each $\epsilon > 0$ and each finite subset $K \subset G$ there exists a finite subset $F \subset G$ such that*

$$(18.12) \quad \frac{|KF \Delta F|}{|F|} \leq \epsilon.$$

PROOF. Applying the anti-automorphism $G \rightarrow G$ given by the inversion $g \mapsto g^{-1}$, we obtain:

$$\frac{|KF \Delta F|}{|F|} = \frac{|F^{-1}K^{-1} \Delta F^{-1}|}{|F^{-1}|}.$$

Lemma follows. \square

In view of Lemma 18.42, instead of defining the Følner Property for G by the right action of G on itself (by right multiplication), we can equivalently define it by the left action (by left multiplication).

EXERCISE 18.46. Show that the following are equivalent for a right action $X \times G \rightarrow X$:

- (1) $X \times G \rightarrow X$ satisfies the *Følner Property*.
- (2) For every $K \subset G$ there exists a sequence $\Omega_n \subset X$ such that

$$\lim_{n \rightarrow \infty} \frac{|\Omega_n K \Delta \Omega_n|}{|\Omega_n|} = 0.$$

Even though, as we will prove in the next section, the Følner Property is equivalent to the amenability, and the latter is inherited by subgroups, it is instructive to describe a construction of Følner sequences for a subgroup directly, in terms of Følner sequence on the ambient group.

PROPOSITION 18.47. *Let H be a subgroup of a group G satisfying the Følner Property, and let $(\Omega_n)_{n \in \mathbb{N}}$ be a Følner sequence for G . For every $n \in \mathbb{N}$ there exists $g_n \in G$ such that the intersections $g_n^{-1}\Omega_n \cap H = \Phi_n$ form a Følner sequence for H .*

PROOF. Consider a finite subset $K \subset H$, let s denote the cardinality of K . Since $(\Omega_n)_{n \in \mathbb{N}}$ is a Følner sequence for G , the ratios

$$(18.13) \quad \alpha_n = \frac{|\Omega_n K \Delta \Omega_n|}{|\Omega_n|}$$

converge to 0. We partition each subset Ω_n into intersections with left cosets of H :

$$\Omega_n = \Omega_n^{(1)} \sqcup \dots \sqcup \Omega_n^{(k_n)},$$

where

$$\Omega_n^{(i)} = \Omega_n \cap g_i H, \quad i = 1, \dots, k_n, \quad g_i H \neq g_j H, \quad \forall i \neq j.$$

Then $\Omega_n K \cap g_i H = \Omega_n^{(i)} K$. We have that

$$\Omega_n K \Delta \Omega_n = \left(\Omega_n^{(1)} K \Delta \Omega_n^{(1)} \right) \sqcup \dots \sqcup \left(\Omega_n^{(k_n)} K \Delta \Omega_n^{(k_n)} \right).$$

The inequality

$$\frac{|\Omega_n K \Delta \Omega_n|}{|\Omega_n|} \leq \alpha_n$$

implies that there exists $i \in \{1, 2, \dots, k_n\}$ such that

$$\frac{|\Omega_n^{(i)} K \Delta \Omega_n^{(i)}|}{|\Omega_n^{(i)}|} \leq \alpha_n.$$

In particular, $g_i^{-1} \Omega_n^{(i)} = \Phi_n$, with $\Phi_n \subseteq H$, and we obtain that

$$\frac{|\Phi_n K \Delta \Phi_n|}{|\Phi_n|} \leq \alpha_n. \quad \square$$

The following proposition complements Lemma 18.24.

PROPOSITION 18.48. *Let G be a group acting on a non-empty set X . The group G is amenable if and only if the action $G \curvearrowright X$ is amenable and for every $p \in X$ the stabilizer G_p of the point p is amenable.*

PROOF. The direct implication follows from Lemma 18.24 and from Part 1 of Theorem 18.29. Assume now that for every $p \in X$ its G -stabilizer G_p is amenable and let $m_X : B(X) \rightarrow \mathbb{R}$ and $m_p : B(G_p) \rightarrow \mathbb{R}$ be G -invariant and G_p -invariant means respectively. We define a left-invariant mean on $B(G)$ using a construction in the spirit of the construction of the product of two measures.

For each $p \in X$ and $F \in B(G)$ define a function F_p on the orbit Gp by

$$F_p(gp) = m_p(F|_{gG_p}).$$

Since m_p is G_p -invariant, $F_p(gp)$ depends only on $x = gp$ and not on g . Moreover, for each $q \in Gp$, the functions $F_p, F_q : Gp \rightarrow \mathbb{R}$ are equal. Therefore, we obtain a G -invariant function F_X on X whose restriction to each orbit Gp equals F_p . Since each m_p is a mean and F is bounded, the function F_X is bounded as well. We define

$$m(F) := m_X(F_X).$$

The linearity of m follows from the linearity of every m_p and of m_X , the properties (M1) and (M2) in Definition 18.19 follow from the fact that m_X and $m_p, p \in X$, are means. We will verify that m is G -invariant. Take an arbitrary element $h \in G$, and consider the pull-back function

$$h \cdot F : x \mapsto F(h^{-1} \cdot x) \quad x \in X.$$

Then

$$(h \cdot F)_p(gp) = m_p((h \cdot F)|_{gG_p}) = m_p(F|_{h^{-1}gG_p}) = F_p(h^{-1}gp).$$

We deduce from this that $(h \cdot F)_X = F_X \circ h^{-1} = h \cdot F_X$, whence

$$m(h \cdot F) = m_X((h \cdot F)_X) = m_X(h \cdot F_X) = m_X(F_X) = m(F). \quad \square$$

18.5. Amenability, paradoxality and the Følner property

In this section we will show that amenability of actions is equivalent to non-paradoxality and to the Følner property. According to Theorem 18.13, if one Cayley graph of a finitely generated group G is amenable then all the other Cayley graphs are. Thus, in what follows we fix a finite generating set S of G , the corresponding Cayley graph $\mathcal{G} = \text{Cayley}(G, S)$, and the corresponding word metric on G .

We will use a construction of *Cayley graphs of group actions*, generalizing the usual notion of Cayley graphs for groups. Let G be a group with a generating set S , and let $X \times G \rightarrow X$ be a right action. We define the *Cayley graph* of this action

(with respect to the generating set S) as the graph $\text{Cayley}(X, G, S)$ whose vertex set is X and whose edge set consists of unordered pairs $\{x, xs\}$, $x \in X, s \in S$.

REMARK 18.49. This construction explains why did we choose to define Følner sequences using *right actions* instead of *left actions*: One defines Cayley graphs using the *right action* of the generating sets S on the group G .

In the next theorem, given a (left) group action $L : G \times X \rightarrow X$, one we use the *right* group action $X \times G \rightarrow X$ defined by $R : (x, g) \mapsto L(g^{-1}, x)$. We note that the equivalence (1) \iff (2) in the next theorem is a special case of the Tarski's Alternative Theorem 18.51. Note also that the proof of Theorem 18.50 (namely, the implication (4) \implies (1)) uses the existence of ultrafilters on \mathbb{N} . One can show that ZF axioms of the set theory are insufficient to conclude that \mathbb{Z} has an invariant mean.

THEOREM 18.50. *The following three conditions are equivalent for a group action $G \curvearrowright X$:*

- (1) $G \curvearrowright X$ has an invariant mean.
- (2) $G \curvearrowright X$ is not paradoxical.
- (3) For every finitely generated subgroup $H \leq G$ and a generating set S of H , the Cayley graph (of the associated right H -action) $\mathcal{G} = \text{Cayley}(X, H, S)$ is amenable.
- (4) The associated right action $X \times G \rightarrow X$ satisfies the Følner property.

PROOF. (1) \implies (2). The implication (1) \implies (2) is established in Corollary 18.27.

(2) \implies (3). We will prove the contrapositive of the implication (2) \implies (3). Assume that the Cayley graph $\mathcal{G} = \mathcal{G}(X, H, S)$ is non-amenable for some finitely generated subgroup $H \leq G$. Equivalently, every component of \mathcal{G} is a non-amenable graph. According to Theorem 18.4, this implies that there exists a map $f : X \rightarrow X$ which is at finite distance from the identity map, such that $|f^{-1}(y)| = 2$ for every $y \in X$. Repeating the proof of Lemma 8.35 verbatim, we conclude that there exists a finite subset $S = \{h_1, \dots, h_n\} \subset G$ and a decomposition $X = T_1 \sqcup \dots \sqcup T_n$ such that f restricted to T_i coincides with $h_i|_{T_i}$.

For every $y \in X$ we have that $f^{-1}(y)$ consists of two elements, which we label as $\{y_1, y_2\}$. This gives a decomposition of X into $Y_1 \sqcup Y_2$. Now we decompose $Y_1 = A_1 \sqcup \dots \sqcup A_n$, where $A_i = Y_1 \cap T_i$, and likewise $Y_2 = B_1 \sqcup \dots \sqcup B_n$, where $B_i = Y_2 \cap T_i$. Clearly

$$A_1 h_1 \sqcup \dots \sqcup A_n h_n = X$$

and

$$B_1 h_1 \sqcup \dots \sqcup B_n h_n = X.$$

We have thus proved that $G \curvearrowright X$ is paradoxical.

(3) \iff (4) The proof of this equivalence is exactly the same as the one in Lemma 18.43. Let $K \subset G$ be a finite non-empty subset and let $H \leq G$ be the subgroup generated by K . It suffices to consider the case $K = K^{-1}$, see Exercise 18.40. The amenability of the Cayley graph $\mathcal{G} = \text{Cayley}(X, H, K)$ implies that there exists a sequence of subsets $\Omega_n \subset X$ such that

$$(18.14) \quad \lim_{n \rightarrow \infty} \frac{|\partial^V \Omega_n|}{|\Omega_n|} = 0.$$

As in the proof of Lemma 18.43,

$$\Omega_n K \triangleleft \Omega_n = \partial^V(\Omega_n K) \cup \partial_V(\Omega_n K).$$

Therefore, (18.14) implies that

$$\lim_{n \rightarrow \infty} \frac{|\Omega_n K \triangleleft \Omega_n|}{|\Omega_n|} = 0.$$

Lastly, Exercise 18.46 shows that the action $X \times G \rightarrow X$ satisfies the Følner property, proving that (3) \Rightarrow (4).

The reverse implication (4) \Rightarrow (3) is proven similarly and we leave details to the reader.

(4) \Rightarrow (1). We first illustrate the proof in the case $G = \mathbb{Z} = X$ and the Følner sequence

$$\Omega_n = [-n, n] \subset \mathbb{Z},$$

since the proof is more transparent in this case and illustrates the general argument. Let ω be a non-principal ultrafilter ω on \mathbb{N} (here we need a form of the Axiom of Choice). We define a function $\mu : 2^{\mathbb{Z}} \rightarrow [0, 1]$ by

$$\mu(A) := \omega\text{-lim} \frac{|A \cap \Omega_n|}{2n+1}, \quad A \subset \mathbb{Z}.$$

We leave it to the reader to check that μ is a f.a.p. measure. Let us show that μ is invariant under the translation $g : z \mapsto z + 1$. Note that

$$||A \cap \Omega_n| - |gA \cap \Omega_n|| \leq 2.$$

Thus,

$$|\mu(A) - \mu(gA)| \leq \omega\text{-lim} \frac{2}{2n+1} = 0.$$

This implies that μ is \mathbb{Z} -invariant.

Consider now the general case. We use a Følner sequence (Ω_n) of subsets of X to construct a G -invariant f.a.p. measure on $2^X = \mathcal{P}(X)$.

Let ω be a non-principal ultrafilter on \mathbb{N} . For every $A \subset X$ define

$$\mu(A) = \omega\text{-lim} \frac{|A \cap \Omega_n|}{|\Omega_n|}.$$

We claim that μ is a f.a.p. measure on X . Indeed, for any pair of disjoint subsets $A, B \subset X$, we have

$$\begin{aligned} \mu(A \sqcup B) &= \omega\text{-lim} \frac{|(A \sqcup B) \cap \Omega_n|}{|\Omega_n|} = \omega\text{-lim} \frac{|A \cap \Omega_n| + |B \cap \Omega_n|}{|\Omega_n|} = \\ &= \omega\text{-lim} \frac{|A \cap \Omega_n|}{|\Omega_n|} + \omega\text{-lim} \frac{|B \cap \Omega_n|}{|\Omega_n|} = \mu(A) + \mu(B). \end{aligned}$$

The condition that $\mu(X) = 1$ is equally clear. We will now verify that μ is G -invariant. Take an element $g \in G$. We have

$$|\mu(Ag) - \mu(A)| = \omega\text{-lim} \frac{||Ag \cap \Omega_n| - |A \cap \Omega_n||}{|\Omega_n|} = \omega\text{-lim} \frac{||A \cap \Omega_n g^{-1}| - |A \cap \Omega_n||}{|\Omega_n|}.$$

Furthermore,

$$||A \cap \Omega_n g^{-1}| - |A \cap \Omega_n|| \leq |A \cap (\Omega_n g^{-1} \triangleleft \Omega_n)|.$$

Since

$$\omega\text{-}\lim \frac{|A \cap (\Omega_n g^{-1} \Delta \Omega_n)|}{|\Omega_n|} \leq \omega\text{-}\lim \frac{|\Omega_n g^{-1} \Delta \Omega_n|}{|\Omega_n|} = 0$$

(as (Ω_n) be a Følner sequence), it follows that

$$\mu(Ag) = \mu(A),$$

i.e. μ is G -right-invariant. □

In particular, Theorem 18.50 shows that the nonexistence of a G -paradoxical decomposition of X is equivalent to the existence of a G -invariant f.a.p. measure on X .

A. Tarski proved ([Tar38], [Tar86, pp. 599–643], see also [Wag85, Corollary 9.2]) the following stronger form of this equivalence:

THEOREM 18.51 (Tarski’s Alternative). *Let G be a group acting on a space X and let E be a subset in X . Then E is not G -paradoxical if and only if there exists a G -invariant finitely additive measure $\mu : \mathcal{P}(X) \rightarrow [0, \infty]$ such that $\mu(E) = 1$.*

The equivalence in Theorem 18.50 gives another proof that the free group on two generators F_2 is paradoxical: Consider the map $f : F_2 \rightarrow F_2$ which is given by deleting the last letter in every non-empty reduced word and $f(1) = 1$. This map satisfies Gromov’s condition in Theorem 18.4. Hence, the Cayley graph of F_2 is non-amenable; thus, F_2 is non-amenable as well.

COROLLARY 18.52. *Each group is either paradoxical or amenable.*

COROLLARY 18.53. *Amenability is QI invariant for finitely generated groups.*

PROOF. This follows from the fact that amenability of graphs of finite valence is QI invariant, see Theorem 18.13. □

Now that we know that the group \mathbb{Z} is amenable, we can get a much larger class of amenable groups than direct limits of finite groups:

COROLLARY 18.54. *Every abelian group G is amenable.*

PROOF. Since every abelian group is a direct limit of finitely generated abelian subgroups, by Part (3) of Theorem 18.29, it suffices to prove amenability of finitely generated abelian groups. Since each \mathbb{Z}^k satisfies the Følner Property (see Exercise 18.39), it is amenable. Each finitely generated abelian group A is a product of a finite group and a free abelian group of finite rank; therefore, A is amenable, e.g. by Part (1) of Corollary 18.32. □

COROLLARY 18.55. *Every solvable group is amenable.*

PROOF. We argue by induction on the derived length. If $k = 1$ then G is abelian and, hence, amenable by Corollary 18.54. Assume that the assertion holds for k and take a group G such that $G^{(k+1)} = \{1\}$ and $G^{(i)} \neq \{1\}$ for every $i \leq k$. Then $G^{(k)}$ is abelian and $\bar{G} = G/G^{(k)}$ is solvable with derived length equal to k . Whence, by the induction hypothesis, \bar{G} is amenable. This and Theorem 18.29, Part (2), imply that G is amenable. □

Similarly, Theorem 18.29 proves amenability for a much larger class of groups, first introduced by M. Day [Day57]:

DEFINITION 18.56. The class of *elementary amenable groups* \mathcal{EA} is the smallest class of groups containing all finite groups, all abelian groups and closed under direct limits, taking subgroups, quotient groups and extensions

$$1 \rightarrow G_1 \rightarrow G_2 \rightarrow G_3 \rightarrow 1,$$

where both G_1, G_3 are elementary amenable.

This class contains all solvable groups and many groups which are not even virtually solvable. Some of these groups (elementary amenable but not virtually solvable) are finitely presented.

EXAMPLE 18.57. Let H_n be the n -th *Houghton group* H_n , $n \geq 2$. Every finite group embeds in every H_n , hence each H_n is not virtually solvable. The group H_n fits in a short exact sequence

$$1 \rightarrow F \rightarrow H_n \rightarrow \mathbb{Z}^{n-1} \rightarrow 1,$$

where F is isomorphic to the group of permutations of \mathbb{N} with finite support, i.e. each $g \in F$ acts as the identity on a complement of some finite subset of \mathbb{N} . In particular, F contains (up to an isomorphism) every finite group and is the direct limit of a system of finite groups. Finite presentability of Houghton groups (for $n \geq 3$) was proven by K. Brown [Bro87]. According to [Lee12], $H_n, n \geq 3$, has the following presentation:

$$\langle a, x_1, \dots, x_{n-1} \mid a^2 = 1, (aa^{x_1})^3 = 1, [a, a^{x_1}] = 1, a = [x_i, x_j], a^{x_i^{-1}} = a^{x_j^{-1}}, 1 \leq i < j \leq n-1 \rangle.$$

Thus, each $H_n, n \geq 3$, is finitely presentable, elementary amenable but not virtually solvable.

Theorem 18.29 and Corollary 18.54 imply that all elementary amenable groups are amenable. There are finitely presented amenable groups which are not elementary amenable; the first such examples were constructed by R. Grigorchuk [Gri98]. Namely, Grigorchuk proves that the group with the following presentation

$$\langle a, c, d, t \mid a^2 = c^2 = d^2 = (ad)^4 = (adacac)^4 = 1, t^{-1}at = aca, t^{-1}ct = dc, t^{-1}dt = c \rangle$$

is amenable but not elementary amenable. None of the elementary amenable groups have intermediate growth according to the following theorem of C. Chou [Cho80]:

THEOREM 18.58. *A finitely generated elementary amenable group either is virtually nilpotent or contains a nonabelian free subsemigroup.*

We note that J. Rosenblatt [Ros74] earlier proved this alternative for solvable groups.

18.6. Supramenability and weakly paradoxical actions

The following definition is formally similar to the one of amenable actions and groups. In order to motivate this definition we note that in the Banach-Tarski paradox, we had a paradoxical decomposition of the subset $\mathbb{B}^n \subset \mathbb{E}^n$ rather than of the entire Euclidean space, i.e. the action of the Euclidean isometry group was weakly paradoxical. While amenability obstructs G -paradoxical decompositions of the entire set X on which the group is acting, supramenability obstructs G -paradoxical decompositions of subsets of X .

DEFINITION 18.59. 1. A group action $G \curvearrowright X$ is said to be *supramenable* if for every non-empty subset $E \subset X$ there exists a f.a. G -invariant measure μ

$$\mu : \mathcal{P}(X) \rightarrow [0, \infty]$$

such that $\mu(E) = 1$.

2. A group G is said to be *supramenable* if the action $G \times G \rightarrow G$ of G on itself by left multiplication is supramenable.

EXERCISE 18.60. Show that in this definition it does not matter if G acts on itself by left or right multiplication.

It is immediate from the definition that each supramenable action is amenable and every supramenable group is amenable.

The following proposition is an analogue of Lemma 18.24 and Theorem 18.50 for supramenable groups.

PROPOSITION 18.61. *The following are equivalent for a group G :*

1. G is not weakly paradoxical.
2. There are no weakly paradoxical actions $G \curvearrowright X$.
3. Every action $G \curvearrowright X$ is supramenable.
4. G is supramenable.

PROOF. The implication (2) \Rightarrow (1) is immediate. The proof of the implication (3) \Rightarrow (2) is analogous to that of Lemma 18.26. Let $E \subset X$ be a non-empty subset and μ a f.a. G -invariant measure on X such that $\mu(E) = 1$. The existence of μ prevents G -paradoxical decompositions of E just as in the proof of Lemma 18.26.

The proof of the implication (4) \Rightarrow (3) is similar to the proof of Lemma 18.24. Consider an action $G \curvearrowright X$ and a non-empty subset E of X . Pick a point $x \in E$ and let G_E be the set of $g \in G$ such that $gx \in E$. This set is non-empty since $1 \in G_E$. Let μ be a G -left-invariant finitely additive measure $\mu_G : \mathcal{P}(G) \rightarrow [0, \infty]$ such that $\mu(G_E) = 1$. For $A \subset X$ let

$$\mu(A) := \mu_G(\{g \in G : g(x) \in A\}).$$

Then $\mu(E) = 1$ and μ is G -invariant f.a. measure on X . Lastly, the implication (1) \Rightarrow (4) is a special case of Tarski's Theorem 18.51. \square

PROPOSITION 18.62. *Each finitely generated weakly paradoxical group has exponential growth.*

PROOF. Let G be weakly paradoxical and let E be a G -paradoxical subset of G with the paradoxical decomposition

$$E = E' \sqcup E''$$

and bijections $\psi' : E \rightarrow E' \subset E, \psi'' : E \rightarrow E'' \subset E$ which are piecewise G -congruences:

$$E' = E'_1 \sqcup \dots \sqcup E'_k, \quad E'' = E''_1 \sqcup \dots \sqcup E''_l$$

and there exist $g'_1, \dots, g'_k \in G, g''_1, \dots, g''_l \in G$ such that

$$\psi' \big|_{g'_i E'_i} = (g'_i)^{-1} \big|_{g'_i E'_i}, \quad \psi'' \big|_{g''_j E''_j} = (g''_j)^{-1} \big|_{g''_j E''_j}, \quad 1 \leq i \leq k, 1 \leq j \leq l.$$

We now generate a semigroup H of (injective but not surjective) maps $E \rightarrow E$ by the maps ψ', ψ'' . Since the images of these maps are disjoint, the semigroup H is free on the generating set $\Psi := \{\psi', \psi''\}$, see Lemma 7.58: We obtain an injective

homomorphism $\rho : SF_2 \rightarrow \text{Map}(E, E)$ from the free semigroup of rank 2 on the generators s', s'' , sending s' to ψ' , s'' to ψ'' . Moreover, according to Lemma 7.58, given any two distinct elements $u, v \in SF_2$,

$$\rho(u)(E) \cap \rho(v)(E) = \emptyset.$$

In particular, given any $x \in E$, the subset

$$X_n := \{\rho(w)(x) : \ell_{SF_2}(w) \leq n\} \subset Gx$$

has cardinality

$$1 + 2 + 4 + \dots + 2^n = 2^{n+1} - 1.$$

Here ℓ_{SF_2} is the word metric on the free semigroup SF_2 . By the construction, for each $x \in G$ and every word $w \in SF_2$ of length m there exists a (positive) word \bar{w} of the same length m in the alphabet $A = \{g'_1, \dots, g'_k, g''_1, \dots, g''_l\}$, such that

$$\rho(w)(x) = \bar{w}x.$$

Taking $x = 1 \in G$, we conclude that the subsemigroup of G generated by the set A has exponential growth. It follows that the group G has exponential growth as well. \square

The following corollary of Proposition 18.62 is a strengthening of Proposition 18.6 in the group-theoretic setting.

COROLLARY 18.63. *Every group of subexponential growth is supramenable.*

It is not known if the converse of Proposition 18.62 is true or if, to the contrary, there exist supramenable groups of exponential growth. A weaker form of the converse of Proposition 18.62 is known though, and it runs as follows.

PROPOSITION 18.64. *A free two-generated subsemigroup S of a group G is always G -paradoxical, where the action $G \curvearrowright G$ is either by left or by right multiplication. In particular, a supramenable group G cannot contain a free two-generated subsemigroup.*

PROOF. Let a, b be the two elements in G generating the free sub-semigroup S , let S_a and S_b be the subsets of elements in S represented by words beginning in a , respectively by words beginning in b . Then $S = S_a \sqcup S_b$, with $a^{-1}S_a = S$ and $b^{-1}S_b = S$. \square

REMARK 18.65. The converse of Proposition 18.64, on the other hand, is not true: a weakly paradoxical group does not necessarily contain a nonabelian free subsemigroup. Namely, there exist paradoxical torsion groups (see Remark 18.78 in the next section).

Below we discuss basic properties of supramenable groups which parallel those of amenable groups, given in Theorem 18.29.

PROPOSITION 18.66. (1) *Any subgroup of a supramenable group is supramenable.*

(2) *Any finite extension of a supramenable group is supramenable.*

(3) *Any quotient of a supramenable group is supramenable.*

(4) *Any direct limit of a directed system of supramenable groups is supramenable.*

PROOF. (1) Let $H \leq G$ with G supramenable and let E be a non-empty subset of H . Let $\mu_G : \mathcal{P}(G) \rightarrow [0, \infty]$ be a G -left-invariant finitely additive measure such that $\mu(E) = 1$. Restricting μ_G to $\mathcal{P}(H)$ we obtain the required measure on G .

(2) Let $H \leq G$ with H supramenable and $G = \bigsqcup_{i=1}^m Hx_i$. Let E be a non-empty subset of G . The group H acts on G by left multiplication, according to Proposition 18.61, there exists $\mu : \mathcal{P}(G) \rightarrow [0, \infty]$, an H -left-invariant f.a. measure such that

$$\mu(\cup_{i=1}^m x_i E) = 1.$$

Define a measure $\nu : \mathcal{P}(G) \rightarrow [0, \infty]$ by

$$\nu(A) = \frac{\sum_{i=1}^m \mu(x_i A)}{\sum_{i=1}^m \mu(x_i E)}.$$

Note that the denominator in this fraction is positive, and it is also clear that ν is finitely additive and satisfies $\nu(E) = 1$. We need to verify the G -left-invariance of ν .

Let A be an arbitrary non-empty subset of G and g an arbitrary element in G . We have that

$$G = \bigsqcup_{i=1}^m Hx_i = \bigsqcup_{i=1}^m Hx_i g,$$

whence there exists a bijection $\varphi : \{1, \dots, m\} \rightarrow \{1, \dots, m\}$ depending on g such that $Hx_i g = Hx_{\varphi(i)}$.

We may then rewrite the numerator in the expression of $\nu(gA)$ as

$$\sum_{i=1}^m \mu(x_i g A) = \sum_{i=1}^m \mu(h_i x_{\varphi(i)} A) = \sum_{i=1}^m \mu(x_{\varphi(i)} A) = \sum_{j=1}^m \mu(x_j A).$$

For the second equality above we have used the H -invariance of μ . We conclude that ν is G -left-invariant.

(3) Let E be a non-empty subset of G/N . Proposition 18.61 applied to the action of G on G/N gives a G -left-invariant finitely additive measure $\mu : \mathcal{P}(G/N) \rightarrow [0, \infty]$ such that $\mu(E) = 1$. The same measure is also G/N -left-invariant.

(4) The proof is very similar to the one in Theorem 18.29, Part 4, and we present only a sketch. Let $h_{ij} : H_i \rightarrow H_j, i \leq j$, be the homomorphisms defining the direct system of groups (H_i) and let G be the direct limit. Let $h_i : H_i \rightarrow G$ be the homomorphisms to the direct limit, see Section 1.5. Consider a non-empty subset E of G . Without loss of generality we may assume that all intersections $h_i(H_i) \cap E$ are non-empty: There exists i_0 such that for every $i \geq i_0$, $h_i(H_i) \cap E \neq \emptyset$, and we then restrict to the set of indices $i \geq i_0$. The set of functions

$$\{f : \mathcal{P}(G) \rightarrow [0, \infty]\} = \prod_{\mathcal{P}(G)} [0, \infty]$$

is compact according to Tychonoff's theorem. For each $i \in I$ let \mathcal{M}_i be the set of H_i -left-invariant f.a. measures μ on $\mathcal{P}(G)$ such that $\mu(E) = 1$. Since H_i is supramenable, Proposition 18.61 implies that the set \mathcal{M}_i is non-empty. Then, as in the proof of Theorem 18.29 one verifies that each subset \mathcal{M}_i is closed in $\prod_{\mathcal{P}(G)} [0, \infty]$ and that the intersection $\bigcap_{i \in I} \mathcal{M}_i$ is non-empty. Every element μ

of this intersection is clearly a f.a. measure such that $\mu(E) = 1$; each μ in the intersection is also G -left-invariant because

$$G = \bigcup_{i \in I} h_i(H_i). \quad \square$$

REMARK 18.67. 1. Note that, unlike amenability, supramenability is not preserved by extensions: If a normal subgroup N in a group G is supramenable and $Q = G/N$ is supramenable then G might not be supramenable. Indeed, if G is metabelian but not virtually nilpotent, then both $G' = N$ and $G/G' = Q$ are abelian. However, each solvable group that is not virtually nilpotent contain a non-abelian free subsemigroup (Theorem 18.58) and, hence, cannot be supramenable according to Proposition 18.62.

2. Surprisingly, it is unknown if finite direct products of supramenable groups are supramenable.

As an example, the group of Euclidean isometries $G = \text{Isom}_+(\mathbb{E}^2)$ is solvable but not virtually nilpotent, therefore, it is not supramenable. Specifically:

PROPOSITION 18.68. $\text{Isom}_+(\mathbb{E}^2)$ contains a free subsemigroup S on two generators. In particular, according to Proposition 18.64, the subset $S \subset G$ is G -paradoxical.

PROOF. Let $\lambda \in \mathbb{C}$ be a transcendental number with $|\lambda| = 1$. Consider the rotation $g(z) = \lambda z$ in \mathbb{E}^2 (identified with the complex plane) and the translation $h : z \mapsto z + 1$. We claim that the semigroup $S \subset G$ generated by g and h is free two-generated. Indeed, consider the set X of all nonconstant integer polynomials with nonnegative coefficients in the variable λ . The semigroup S acts on X by the postcomposition

$$s \cdot p(\lambda) = s \circ p(\lambda), \quad s \in S.$$

Then for each $p \in X$, $h \cdot p$ has non-zero constant term, while $g \cdot p$ has zero constant term. Therefore, $g(X), h(X)$ are disjoint subsets of X . Hence, according to Lemma 7.58, the semigroup S is free of rank 2. \square

COROLLARY 18.69 (Sierpinski–Mazurkiewicz paradox). *There exists a countable G -paradoxical subset $E \subset \mathbb{E}^2$. In particular, the action of $G = \text{Isom}_+(\mathbb{E}^2)$ on \mathbb{E}^2 is weakly paradoxical.*

PROOF. Since the semigroup S is countable, for a generic choice of $z \in \mathbb{C}$ the orbit map $S \rightarrow \mathbb{C}, s \mapsto s(z)$ is injective. Now the claim follows from Lemma 17.7. \square

Neither one of the classes of supramenable and elementary amenable groups contains the other:

- solvable groups are all elementary amenable, while they are supramenable only if they are virtually nilpotent;
- there exist groups of intermediate growth (which are necessarily supramenable) that are not elementary amenable, see [Gri84a].

We are now able to relate amenability to the Banach–Tarski paradox.

PROPOSITION 18.70. (1) *The group of isometries $\text{Isom}(\mathbb{R})$ is supramenable and, hence, \mathbb{R} does not contain paradoxical subsets.*

- (2) *The group of isometries $G = \text{Isom}(\mathbb{E}^2)$ is amenable but not supramenable.*
- (3) *\mathbb{E}^2 contains paradoxical subsets.*
- (4) *\mathbb{E}^2 does not admit a paradoxical decomposition.*
- (5) *The group of isometries $\text{Isom}(\mathbb{E}^n)$ with $n \geq 3$ is non-amenable.*

PROOF. Part (1) follows from the fact that $\text{Isom}(\mathbb{R})$ contains the abelian subgroup $\text{Isom}_+(\mathbb{R})$ of index 2.

Part (2) follows from the fact that G is solvable (and, hence, amenable), but not supramenable since it contains a free subsemigroup of rank 2.

Part (3) is the Sierpinski–Mazurkiewicz paradox above.

Part (4) follows from the amenability of $G = \text{Isom}(\mathbb{E}^2)$, which implies the amenability and, hence, non-paradoxality, of any action of G .

Part (5) follows from the fact that $SO(3) < \text{Isom}(\mathbb{E}^n)$ contains rank 2 free subgroups. \square

Since many examples and counterexamples display a connection between amenability and the algebraic structure of a group, it is natural to ask whether there exists a purely algebraic criterion of amenability. J. von Neumann made the observation that the existence of a free subgroup excludes amenability in the very paper where he introduced the notion of amenable groups, [vN28]. It is this observation that has led to the following question:

QUESTION 18.71 (the von Neumann-Day problem). Does every non-amenable group contain a nonabelian free subgroup?

The question is implicit in [vN29], and it was formulated explicitly by M. Day [Day57, §4].

When restricted to the class of subgroups of Lie groups with finitely many components (in particular, subgroups of $GL(n, \mathbb{R})$), Question 18.71 has an affirmative answer, since, in view of the Tits’ alternative (see Chapter 15, Theorem 15.1), every such group either contains a nonabelian free subgroup or is virtually solvable. Since all virtually solvable groups are amenable, the claim follows. For the same reason, for all classes of groups satisfying the Tits Alternative (see section 15.7) Question 18.71 has an affirmative answer.

The first examples of finitely generated non-amenable groups with no (non-abelian) free subgroups were given by A. Ol’shanskiĭ in [Ol’80]. In [Ady82] it was shown that the free Burnside groups $B(n, m)$ with $n \geq 2$ and $m \geq 665$, m odd, are also non-amenable. The first examples of finitely presented non-amenable groups with no (non-abelian) free subgroups, were given by A. Ol’shanskiĭ and M. Sapir in [OS02]. A more recent example, due to Y. Lodha and J.T. Moore [LM16] is a subgroup of the group of piecewise projective homeomorphisms of the real projective line, **torsion-free** and with an explicit presentation with three generators and nine relators.

Still, metric versions of the von Neumann-Day Question 18.71 have positive answers. One of these versions is Whyte’s Theorem 18.10. Another metric version with positive answer was established by I. Benjamini and O. Schramm in [BS97b]. They proved

THEOREM 18.72. • Every infinite locally finite simplicial graph \mathcal{G} with positive Cheeger constant contains a tree with positive Cheeger constant.²

- If, moreover, the Cheeger constant of \mathcal{G} is at least an integer $n \geq 0$, then \mathcal{G} contains a spanning subgraph, where each connected component is a rooted tree with all vertices of valency n , except the root, which is of valency $n+1$.
- If X is either a graph or a Riemannian manifold with infinite diameter, bounded geometry and positive Cheeger constant (in particular, if X is the Cayley graph of a paradoxical group) then X contains a bi-Lipschitz embedding of the binary rooted tree.

Related to the above, the following is asked in [BS97b]:

QUESTION 18.73. Is it true that every Cayley graph of every finitely generated group with exponential growth contains a subtree with positive Cheeger constant?

We note that the existence of such subtrees does not contradict amenability, for instance, each finitely generated elementary amenable group G which is not virtually nilpotent contains a rank 2 free subsemigroup and, hence one of the Cayley graphs of G contains a 2-adic subtree.

18.7. Quantitative approaches to non-amenability and weak paradoxality

One can measure “how paradoxical” a group or a group action is *via* the *Tarski numbers*.

DEFINITION 18.74. (1) Given an action of a group G on a set X , and a subset $E \subset X$, which admits a G -paradoxical decomposition, the *Tarski number of the paradoxical decomposition* is the number $k + l$ of pieces of that decomposition (see Definition 17.4).

(2) The *Tarski number* $\text{Tar}_G(X, E)$ is the minimum of the Tarski numbers taken over all G -paradoxical decompositions of E . We set $\text{Tar}_G(X, E) = \infty$ in the case when there exists no G -paradoxical decomposition of the subset $E \subset X$. We use the notation $\text{Tar}_G(X)$ for $\text{Tar}_G(X, X)$.

(3) We define the *lower Tarski number* $\text{tar}(G)$ of a group G to be the minimum of the numbers $\text{Tar}_G(X, E)$ for all the actions $G \curvearrowright X$ and all the non-empty subsets E of X .

(4) When $X = G$ and the action is by left multiplication, we denote $\text{Tar}_G(G, G)$ simply by $\text{Tar}(G)$ and we call it the *Tarski number of G* .

Thus, for any G , X , $E \subset X$ and an action $G \curvearrowright X$, we have:

$$\text{tar}(G) \leq \text{Tar}_G(X, E) \leq \text{Tar}_G(X, X),$$

$$\text{tar}(G) \leq \text{Tar}(G)$$

Moreover, G is amenable if and only if $\text{Tar}(G) = \infty$, while G is supramenable if and only if $\text{tar}(G) = \infty$. Thus, the number $\text{tar}(G)$ measures the *weak paradoxality*

²Note that in this result no uniform bound on valence is assumed. The definition of the Cheeger constant is considered with the edge boundary.

of G , i.e. the degree of its failure to be supramenable. Similarly, the number $\text{Tar}(G)$ measures the *paradoxality* of G , i.e. the degree of its failure to be amenable.

REMARK 18.75. Of course, one could refine the discussion further and use other cardinal numbers besides the finite ones to quantify nonamenability. We do not pursue this direction here.

The following theorem was first proved by R. M. Robinson in [Rob47]:

THEOREM 18.76. *For $n \geq 3$, the Tarski number $\text{Tar}_G(\mathbb{E}^n, \mathbb{B}^n)$ of the unit ball $\mathbb{B}^n \subset \mathbb{E}^n$ (with respect to the action of the group of isometries $G = \text{Isom}(\mathbb{E}^n)$) is 5.*

PROPOSITION 18.77. *Let G be a group, $G \curvearrowright X$ be an action and $E \subset X$ be a non-empty subset.*

- (1) *For each subgroup $H \leq G$, $\text{Tar}_G(X, E) \leq \text{Tar}_H(X, E)$.*
- (2) *The lower Tarski number $\text{tar}(G)$ of every group is at least two. Moreover, $\text{tar}(G) = 2$ if and only if G contains a free two-generated subsemigroup S .*

PROOF. (1) This inequality follows immediately from the fact that for every $E \subset X$, each H -paradoxical decomposition of E is also G -paradoxical.

(2) The fact that every $\text{tar}_G(X, E)$ is at least two is clear from the definition of a paradoxical decomposition. We prove the direct part of the equivalence. Assume that $\text{tar}(G) = 2$. Then there exists an action $G \curvearrowright X$, a subset $E \subset X$ with a decomposition $E = E' \sqcup E''$ and two elements $g', g'' \in G$ such that $g'(E') = E$ and $g''(E'') = E$. Setting $g := (g')^{-1}$, $h := (g'')^{-1}$, we obtain injective maps

$$g : E \rightarrow E' \subset E, \quad h : E \rightarrow E'' \subset E$$

with disjoint images. Lemma 7.58 implies that g, h generate a rank 2 free subsemigroup in G .

Conversely, if x, y be two elements in G generating a free subsemigroup S , let S_x be the set of words beginning with x and S_y be the set of words beginning with y . Then $S = S_x \sqcup S_y$, with $x^{-1}S_x = S$ and $y^{-1}S_y = S$. Therefore, $\text{Tar}_G(G, S) = 2$. \square

REMARK 18.78. R. Grigorchuk constructed in [Gri87] examples of finitely generated amenable torsion groups G which are weakly paradoxical, thus answering Rosenblatt's conjecture [Wag85, Question 12.9.b]. Every such amenable group G satisfies

$$3 \leq \text{tar}(G) < \infty.$$

QUESTION 18.79. What are the possible values of $\text{tar}(G)$ for a weakly paradoxical group G ? How different can it be from $\text{Tar}(G)$?

We now move on to study the values of Tarski numbers $\text{Tar}_G(X)$ and $\text{Tar}(G)$

PROPOSITION 18.80. *For every group action $G \curvearrowright X$ on a non-empty set X we have:*

- (1) $\text{Tar}_G(X) \geq 4$.
- (2) *If G acts freely on X and G contains a free subgroup of rank two, then $\text{Tar}_G(X) = 4$.*

PROOF. (1) Since G acts on X via bijections, in every G -paradoxical decomposition of X one must have $k \geq 2$ and $l \geq 2$. Thus, the Tarski number $\text{Tar}_G(X)$ is always at least 4.

(2) This statement is the content of Lemma 17.20. □

The next proposition complements Part (2) of Proposition 18.80:

PROPOSITION 18.81. *1. If there exists a G -action $G \curvearrowright X$ for which X admits a G -paradoxical decomposition*

$$(18.15) \quad X = X' \sqcup X'', \quad X' = X'_1 \sqcup X'_2, \quad X'' = X''_1 \sqcup \dots \sqcup X''_l, \quad l \geq 2,$$

then G contains an element of infinite order.

2. If there exists an action $G \curvearrowright X$ with $\text{Tar}_G(X) = 4$, then G contains a non-abelian free subgroup.

PROOF. We let $\phi' : X' \rightarrow X$, $\phi'' : X'' \rightarrow X$ be piecewise G -congruences from (18.15); they restrict to X'_i as $g_i \in G$, $i = 1, 2$, and to X''_j as h_j , $j = 1, \dots, l$. We define products

$$g := g_1^{-1}g_2, \quad h := h_1^{-1}h_2.$$

We leave it to the reader to verify that

$$(18.16) \quad g(X'_1 \sqcup X'') \subset X'_1$$

and, therefore, by relabelling $1 \leftrightarrow 2$,

$$g^{-1}(X'_2 \sqcup X'') \subset X'_2,$$

1. Since X'' is non-empty, (18.16) implies that

$$g(X'_1) \subsetneq X'_1.$$

It now follows from Exercise 7.59 that $g \in G$ has infinite order.

2. Since $\text{Tar}_G(X) = 4$, there exists a G -paradoxical decomposition as in (18.15) with $l = 2$. Since the number of pieces in X' and X'' is now the same, we obtain (by relabelling):

$$h(X''_1 \sqcup X') \subset X''_1$$

and

$$h^{-1}(X''_2 \sqcup X') \subset X''_2.$$

It follows that the pair of elements $g, h \in G$ and the subsets X'_1, X'_2, X''_1, X''_2 satisfies the assumption of the Ping-pong Lemma (Lemma 7.60). Hence, g and h generate a free subgroup of rank 2 in G . □

COROLLARY 18.82. *If G is a torsion group then for every G -action on a set X , $\text{Tar}_G(X) \geq 6$.*

PROOF. Suppose that X admits a G -paradoxical decomposition with $k + l$ parts. Part 1 of Lemma 18.81 implies that if either k or l is ≤ 2 , then G contains an infinite order element, which contradicts our assumptions. Therefore $k \geq 3, l \geq 3$ and $\text{Tar}_G(X) \geq 6$. □

S. Wagon (Theorems 4.5 and 4.8 in [Wag85]) proved a stronger form of Proposition 18.81 and Proposition 18.80, part (2):

THEOREM 18.83 (S. Wagon). *Let G be a group acting on a non-empty set X . The Tarski number $\text{Tar}_G(X)$ is four if and only if G contains a nonabelian free subgroup F such that the stabilizer in F of each point in X is abelian.*

Below we describe the behavior of the Tarski number of groups with respect to certain group operations.

PROPOSITION 18.84. (1) *If H is a subgroup of G then $\text{Tar}(G) \leq \text{Tar}(H)$.*

(2) *Every paradoxical group G contains a finitely generated subgroup H such that $\text{Tar}(G) = \text{Tar}(H)$.*

(3) *$\text{Tar}(G) \leq \text{Tar}(Q)$ for every quotient group Q of G .*

PROOF. (1) If H is amenable then there is nothing to prove. Consider an H -paradoxical decomposition of H :

$$H = X_1 \sqcup \dots \sqcup X_k \sqcup Y_1 \sqcup \dots \sqcup Y_l$$

such that

$$H = g_1 X_1 \sqcup \dots \sqcup g_k X_k = h_1 Y_1 \sqcup \dots \sqcup h_l Y_l,$$

where g_i, h_j are elements of H and $k + l = \text{Tar}(H)$. Let $S \subset G$ denote the set of representatives of right H -cosets inside G : the product map

$$H \times S \rightarrow G, \quad (h, s) \mapsto hs,$$

is a bijection. Then the subsets

$$X'_i = X_i S, \quad 1 \leq i \leq k,$$

together with

$$Y'_j = Y_j S, \quad 1 \leq j \leq l,$$

form a paradoxical decomposition of G .

(2) Given a decomposition

$$G = X_1 \sqcup \dots \sqcup X_k \sqcup Y_1 \sqcup \dots \sqcup Y_l$$

such that

$$(18.17) \quad H = g_1 X_1 \sqcup \dots \sqcup g_k X_k = h_1 Y_1 \sqcup \dots \sqcup h_l Y_l$$

and $k + l = \text{Tar}(G)$, consider the subgroup $H \leq G$ generated by $g_1, \dots, g_k, h_1, \dots, h_l$. Then (18.17) is an H -paradoxical decomposition of G with respect to the action of H on G by left multiplication. Intersecting pieces of this decomposition with H , we obtain

$$G = (H \cap X_1) \sqcup \dots \sqcup (H \cap X_k) \sqcup (H \cap Y_1) \sqcup \dots \sqcup (H \cap Y_l),$$

an H -paradoxical decomposition of H . Thus $\text{Tar}(H) \leq \text{Tar}_H(G, G) \leq \text{Tar}(G)$; the opposite inequality is proven in Part 1.

(3) Suppose that $\pi : G \rightarrow Q$ is a quotient map with a set-theoretic cross-section $\sigma : Q \rightarrow G$. As before, we may assume, without loss of generality, that Q is paradoxical. Let

$$Q = X'_1 \sqcup \dots \sqcup X'_k \sqcup X''_1 \sqcup \dots \sqcup X''_l$$

be a paradoxical decomposition of Q with piecewise congruences given by restrictions of elements $q'_1, \dots, q'_k, q''_1, \dots, q''_l$ of Q . We assume that $\text{Tar}(Q) = k + l$. Then

$$G = \pi^{-1}(X'_1) \sqcup \dots \sqcup \pi^{-1}(X'_k) \sqcup \pi^{-1}(X''_1) \sqcup \dots \sqcup \pi^{-1}(X''_l)$$

is a paradoxical decomposition of G with $k+l$ pieces and with piecewise congruences defined by restrictions of the elements $\sigma(q'_i), 1 \leq i \leq k, \sigma(q''_j), 1 \leq j \leq l$, of the group G . \square

COROLLARY 18.85. *A group G contains a non-abelian free subgroup if and only if $\text{Tar}(G) = 4$.*

PROOF. If a group G contains a non-abelian free subgroup then the result follows by Proposition 18.80, (1), (2), and Proposition 18.84, (1). If a group G has $\text{Tar}(G) = 4$ then the claim follows from Proposition 18.81. \square

Thus, the Tarski number helps to classify the groups that are non-amenable and do not contain nonabelian free subgroups. This class of groups is not very well understood, its only known members are (small) “infinite monsters”. For torsion groups G , as we noted above, $\text{Tar}(G) \geq 6$. On the other hand, T. Ceccherini, R. Grigorchuk and P. de la Harpe proved:

THEOREM 18.86 (Theorem 2, [CSGdlH98]). *The Tarski number of every free Burnside group $B(n; m)$ with $n \geq 2$ and $m \geq 665$, m odd, is at most 14.*

Part (1) of Proposition 18.84 implies the following quantitative version of Corollary 18.30:

COROLLARY 18.87. *If two groups G_1, G_2 are co-embeddable then they have the same Tarski number: $\text{Tar}(G_1) = \text{Tar}(G_2)$.*

It is proven in [Šir76] and [Ady79, Theorem VI.3.7] that, for every odd $m \geq 665$, $n \geq 2, k \geq 2$, the two free Burnside groups $B(n; m)$ and $B(k; m)$ of exponent m are co-embeddable. Thus:

COROLLARY 18.88. *For every odd $m \geq 665$, and $n \geq 2$, the Tarski number of the free Burnside groups $B(n; m)$ is finite and independent of the number of generators n .*

Natural questions, in view of Corollary 18.88, are the following:

QUESTION 18.89. How does the Tarski number of a free Burnside group $B(n; m)$ depend on the exponent m ? What are its possible values?

The following appears as Question 22 [dlHGCS99] (asked also in [CSGdlH98]):

QUESTION 18.90 (Question 22 in [dlHGCS99], [CSGdlH98]). What are the possible values for the Tarski numbers of groups? Do they include 5 or 6? Are there groups which have arbitrarily large finite Tarski numbers?

An answer to the last question is given by examples of M. Ershov [Ers11], who proves that certain Golod-Shafarevich groups G have infinite quotients with Property (T) and, for every m large enough, G contains finite index subgroups H_m with the property that all their m -generated subgroups are finite. Indeed, according to Proposition 18.84, (2), every subgroup H_m has Tarski number at least $m + 1$.

The same examples show that Tarski numbers are not quasiisometry invariants, in fact not even commensurability invariants.

M. Ershov, G. Golan and M. Sapir [EGS15] proved moreover that there exist groups with Tarski numbers 5 and 6.

It would be still interesting to understand how much of the Tarski number is encoded in the large scale geometry of a group. In particular:

QUESTION 18.91. 1. Is it at least true that the existence of an (L, C) -quasiisometry between groups (with fixed finite generating sets) implies that their Tarski number differ at most by a constant $K = K(L, A)$?

2. Is the Tarski number a quasiisometry invariant, when it takes small values?

For instance, the quasiisometry invariance of the property of having Tarski number 4 is equivalent to a well-known open problem, which we describe below.

DEFINITION 18.92. A group G is called *small* if it contains no nonabelian free subgroups. Thus, G is small iff $\text{Tar}(G) > 4$. Accordingly, a group is called *large* if it contains a nonabelian free subgroup. Dually, a group G is *co-large* if it contains a finite-index subgroup $\Gamma' \leq \Gamma$, which admits an epimorphism to a nonabelian free group.

Thus, the last of the Questions 18.91, can be reformulated as the first of the Questions 18.93:

QUESTION 18.93. Is smallness invariant under quasiisometries of finitely generated groups? Is co-largeness a QI invariant for hyperbolic groups?

Note that co-largeness is not a QI invariant for finitely generated (and even $CAT(0)$) groups. The simplest examples of this phenomenon are given by uniform lattices acting on $\mathbb{H}^2 \times \mathbb{H}^2$: Among such lattices there are product groups $\Gamma = \Gamma_1 \times \Gamma_2$, where both Γ_1, Γ_2 are surface groups, as well as *irreducible lattices* Γ' . The groups Γ, Γ' are quasiisometric to each other, but Γ is co-large, while Γ' is not (as follows from Margulis' Superrigidity Theorem).

18.8. Uniform amenability and ultrapowers

In this section we discuss a *uniform version of amenability* and its relation to ultrapowers of groups.

Recall (Definition 18.44) that a group G has *the Følner Property* if for every finite subset K of G and every $\epsilon \in (0, 1)$ there exists a finite non-empty subset $F \subset G$ satisfying:

$$(18.18) \quad |KF \triangle F| < \epsilon|F|.$$

DEFINITION 18.94. A group G has the *uniform Følner Property* if, in addition, one can bound the size of F in terms of ϵ and $|K|$, i.e. there exists a function $\phi : (0, 1) \times \mathbb{N} \rightarrow \mathbb{N}$ such that for each $\epsilon \in (0, 1)$ and each finite subset $K \subset G$, there exists a finite subset $F \subset G$ satisfying the inequality (18.18) and

$$|F| \leq \phi(\epsilon, |K|).$$

EXAMPLES 18.95. (1) Nilpotent groups have the uniform Følner property [Bož80].

(2) A subgroup of a group with the uniform Følner Property also has this property [Bož80].

- (3) Let N be a normal subgroup of G . The group G has the uniform Følner Property if and only if both N and G/N have this property [Bož80].

THEOREM 18.96 (G. Keller [Kel72], [Wys88]). *The following are equivalent:*

- (1) G has the uniform Følner Property.
- (2) For some non-principal ultrafilter ω the ultrapower G^ω has the Følner Property.
- (3) For every non-principal ultrafilter ω , the ultrapower G^ω also has the uniform Følner property.

PROOF. The implication (3) \Rightarrow (2) is clear.

(2) \Rightarrow (1). We identify the group G with the “diagonal” subgroup \widehat{G} of G^ω , represented by constant sequences in G . It follows from Proposition 18.47 that G has the Følner property. Assume that G does not have the uniform Følner property. Then there exists $\varepsilon > 0$ and a sequence of subsets K_n in G of fixed cardinality k such that for every sequence of subsets $\Omega_n \subset G$

$$|K_n \Omega_n \triangle \Omega_n| < \varepsilon |\Omega_n| \Rightarrow \lim_{n \rightarrow \infty} |\Omega_n| = \infty.$$

Let $K_\omega = (K_n)^\omega$. According to Lemma 10.35, K has cardinality k . Since G^ω satisfies the Følner property, there exists a finite subset $U \subset G^\omega$ such that

$$|KU \triangle U| < \varepsilon |U|.$$

Let c denote the cardinality of U . According to Lemma 10.35, Part (3), $U = (U_n)^\omega$, where each $U_n \subset G$ has cardinality c . Moreover, ω -almost surely

$$|KU_n \triangle U_n| < \varepsilon |U_n|.$$

Contradiction. We, therefore, conclude that G has the uniform Følner Property.

(1) \Rightarrow (3). Let $k \in \mathbb{N}$ and $\varepsilon > 0$; define $m := \phi(\varepsilon, k)$ where ϕ is the function coming from the uniform Følner property of G . Let K be a subset of cardinality k in G^ω . Lemma 10.35 implies that $K = (K_n)^\omega$, for some sequence of subsets $K_n \subset G$ of cardinality k . The uniform Følner Property of G implies that for each n there exists a subset $\Omega_n \subset G$ of cardinality at most m such that

$$|K_n \Omega_n \triangle \Omega_n| < \varepsilon |\Omega_n|.$$

Let $F := (\Omega_n)^\omega$. The description of K and F given by Lemma 10.35, (1), implies that

$$KF \triangle F = (K_n \Omega_n \triangle \Omega_n)^\omega,$$

whence $|KF \triangle F| < \varepsilon |\Phi|$. Since $|F| \leq m$ according to Lemma 10.35, Part (1), the claim follows. \square

COROLLARY 18.97 (G. Keller, [Kel72], Corollary 5.9). *Every group with the uniform Følner property satisfies a law.*

PROOF. Indeed, by Theorem 18.96, if G has the uniform Følner Property then every ultrapower G^ω has the uniform Følner Property. Assume that G does not satisfy any law, i.e. in view of Lemma 10.42, the group G^ω contains a subgroup isomorphic to the free group F_2 . By Proposition 18.47 it would then follow that F_2 has the Følner Property, a contradiction. \square

EXAMPLE 18.98. Let $H = H_n, n \geq 3$, be the n th Houghton group, see Example 18.57. The group H is finitely presented, amenable and each finite group embeds into H . We claim that H cannot satisfy any law. Indeed, if H did satisfy a law $w(x_1, \dots, x_n) = 1$ then all finite groups would satisfy this law. Then the direct product

$$G = \prod_{\Phi \in \mathcal{F}} \Phi$$

would satisfy the same law. (Here \mathcal{F} denotes the set of isomorphism classes of all finite groups.) All subgroups of G would satisfy this law as well. However, since the free group F_2 is residually finite, it embeds in G . A contradiction. Therefore, H is amenable but not uniformly amenable.

18.9. Quantitative approaches to amenability

One quantitative approach to amenability (of finitely generated groups and of graphs of finite valence) is due to A.M. Vershik, who introduced in [Ver82] the notion of *Følner function*. Given an amenable graph \mathcal{G} of bounded geometry, its *Følner function* $F_o^{\mathcal{G}} : (0, \infty) \rightarrow \mathbb{N}$ is defined by the condition that $F_o^{\mathcal{G}}(t)$ is the minimal cardinality of a finite non-empty subset $F \subset V(\mathcal{G})$ satisfying the inequality

$$|\partial_V F| \leq \frac{1}{t}|F|.$$

According to the inequality (1.4) relating the cardinalities of the vertex and edge boundaries, if one replaces in this definition $\partial_V F$ with $E(F, F^c)$ or $\partial^V F$, one obtains a function asymptotically equal to the first, in the sense of Definition 1.4.

The following is a quantitative version of Theorem 18.13 (which establishes the quasiisometry invariance of the amenability for graphs).

PROPOSITION 18.99. *If two graphs of bounded geometry are quasiisometric then their Følner functions are asymptotically equal.*

PROOF. Let \mathcal{G} and \mathcal{G}' be two graphs of bounded geometry, and let $f : \mathcal{G} \rightarrow \mathcal{G}'$ and $g : \mathcal{G}' \rightarrow \mathcal{G}$ be two (L, C) -quasiisometries such that $f \circ g$ and $g \circ f$ are at distance at most C from the respective identity maps (in the sense of the inequalities (8.3)). Without loss of generality we may assume that both f and g send vertices to vertices. Let m be the maximal valency of a vertex in both \mathcal{G} and \mathcal{G}' .

We begin by some general considerations. We denote by α the maximal cardinality of $B(x, C) \cap V$, where $B(x, C)$ is an arbitrary ball of radius C in either \mathcal{G} or \mathcal{G}' . Since both graphs have bounded geometry, it follows that α is finite.

Let A be a finite subset in $V(\mathcal{G})$, let $A' = f(A)$ and $A'' = g(A')$. It is obvious that $|A''| \leq |A'| \leq |A|$. The Hausdorff distance between A'' and A is at most C , and, therefore, $|A| \leq \alpha|A''|$. Thus, we have the inequalities

$$(18.19) \quad \frac{1}{\alpha}|A| \leq |f(A)| \leq |A|,$$

and similar inequalities for finite subsets in $V(\mathcal{G}')$ and their images by g .

Assume now that both \mathcal{G} and \mathcal{G}' are amenable, and let $F_o^{\mathcal{G}}$ and $F_o^{\mathcal{G}'}$ be their respective Følner functions.

Fix $t \in (0, \infty)$, and let A be a finite subset in $V(\mathcal{G})$ such that $|A| = F_o^{\mathcal{G}}(t)$ and

$$|\partial_V A| \leq \frac{1}{t}|A|.$$

Let $A' = f(A)$ and $A'' = g(A')$. We fix the constant $R = L(2C + 1)$, and consider the set $B = \mathcal{N}_R(A')$. The vertex-boundary $\partial_V B$ consists of vertices u such that $R \leq \text{dist}(u, A') < R + 1$.

It follows that

$$\text{dist}(g(u), A) \geq \text{dist}(g(u), A'') - C \geq \frac{1}{L}R - 2C = 1$$

and that

$$\text{dist}(g(u), A) \leq L(R + 1) + C.$$

In particular, every vertex $g(u)$ is at distance at most $L(R + 1) + C - 1$ from $\partial_V A$ and it is not contained in A . We have, thus, proved that

$$g(\partial_V B) \subseteq \mathcal{N}_{L(R+1)+C-1}(\partial_V A) \setminus A.$$

It follows that if we denote $m^{L(R+1)+C-1}$ by λ , then we can write, using (18.19),

$$\begin{aligned} |\partial_V B| &\leq \alpha |g(\partial_V B)| \leq \alpha \lambda |\partial_V A| \leq \alpha \lambda \frac{1}{t} |A| \leq \\ &\alpha^2 \lambda \frac{1}{t} |A'| \leq \alpha^2 \lambda \frac{1}{t} |B|. \end{aligned}$$

We have thus obtained that, for $\kappa = \alpha^2 \lambda$ and an arbitrary $t > 0$, the value $F_o^{\mathcal{G}'}\left(\frac{t}{\kappa}\right)$ is at most $|B| \leq m^R |A'| \leq m^R |A| = m^R F_o^{\mathcal{G}}(t)$. We conclude that $F_o^{\mathcal{G}'} \preceq F_o^{\mathcal{G}}$.

The opposite inequality $F_o^{\mathcal{G}} \preceq F_o^{\mathcal{G}'}$ is obtained by relabelling. \square

Proposition 18.99 implies that, given a finitely generated amenable group G , any two of its Cayley graphs have asymptotically equal Følner functions. We will, therefore, write F_o^G , for the equivalence class of all these functions.

DEFINITIONS 18.100. (1) We say that a sequence (Φ_n) of finite subsets in a group *realizes the Følner function* of that group if for some generating set S , $|\Phi_n| = F_o^{\mathcal{G}}(n)$, where \mathcal{G} is the Cayley graph of G with respect to S , and

$$|E(\Phi_n, \Phi_n^c)| \leq \frac{1}{n} |\Phi_n|.$$

(2) We say that a sequence (A_n) of finite subsets in a group *quasirealizes the Følner function* of that group if $|A_n| \asymp F_o^{\mathcal{G}}(n)$ and there exists a constant $a > 0$ and a finite generating set S such that for every n ,

$$|E(A_n, A_n^c)| \leq \frac{a}{n} |A_n|,$$

where $|E(A_n, A_n^c)|$ is the edge boundary of A_n in the Cayley graph of G with respect to S .

LEMMA 18.101. *Let H be a finitely generated subgroup of a finitely generated amenable group G . Then $F_o^H \preceq F_o^G$.*

PROOF. Consider a generating set S of G containing a generating set T of H . We shall prove that the Følner functions defined with respect to these generating sets, satisfy the inequality

$$F_o^H(t) \leq F_o^G(t)$$

for every $t > 0$. Let Φ be a finite subset in G such that $|\Phi| = F_o^G(t)$ and $|\partial_V \Phi| \leq \frac{1}{t}|\Phi|$.

The set F intersects finitely many cosets of H, g_1H, \dots, g_kH . In particular $\Phi = \bigsqcup_{i=1}^k \Phi_i$, where $\Phi_i = F \cap g_iH$. We denote by $\partial_V^i \Phi_i$ the set of vertices in $\partial_V \Phi_i$ joined to vertices in the complementary set Φ_i^c by edges with labels in T . The sets $\partial_V^i \Phi_i$ are contained in g_iH for every $i \in \{1, 2, \dots, k\}$, hence, they are pairwise disjoint subsets of $\partial_V \Phi$. We, thus, obtain:

$$\sum_{i=1}^k |\partial_V^i \Phi_i| \leq |\partial_V \Phi| \leq \frac{1}{t}|\Phi| = \frac{1}{t} \sum_{i=1}^k |\Phi_i|.$$

It follows that there exists $i \in \{1, 2, \dots, k\}$ such that

$$|\partial_V^i \Phi_i| \leq \frac{1}{t}|\Phi_i|.$$

By the construction, $\Phi_i = g_iK_i$ with $K_i \subset H$, and the previous inequality is equivalent to

$$|\partial_V^H K_i| \leq \frac{1}{t}|K_i|,$$

where the vertex-boundary $\partial_V^H K_i$ is considered in the Cayley graph of H with respect to the generating set T . We then conclude that

$$F_o^H(t) \leq |K_i| \leq |\Phi| = F_o^G(t).$$

□

One may ask how do the Følner functions relate to the growth functions, and when do the sequences of balls of fixed centre quasirealize the Følner function, especially under the extra hypothesis of subexponential growth, see Proposition 18.6.

THEOREM 18.102. *Let G be an infinite finitely generated group. Then:*

- (1) $F_o^G(t) \succeq \mathfrak{G}_G(t)$.
- (2) *If the sequence of balls $B(1, n)$ quasirealizes the Følner function of G then G is virtually nilpotent.*

PROOF. (1) Consider a sequence (Φ_n) of finite subsets in G that realizes the Følner function of that group (for some finite generating set X). In particular

$$|E(\Phi_n, \Phi_n^c)| \leq \frac{1}{n}|\Phi_n|.$$

We let \mathfrak{G} denote the growth function of G with respect to the generating set X .

The inequality (8.11) in Proposition 8.91 implies that

$$\frac{|\Phi_n|}{2dk_n} \leq \frac{1}{n}|\Phi_n|,$$

where $d = |X|$ and k_n is such that $\mathfrak{G}(k_n - 1) \leq 2|\Phi_n| < \mathfrak{G}(k_n)$. Therefore,

$$k_n \geq \frac{n}{2d},$$

whence,

$$2F_o^G(n) \geq \mathfrak{G}(k_n - 1) \geq \mathfrak{G}\left(\frac{n}{2d}\right).$$

(2) The inequality in Part (2) of Definition 18.100 implies that for some $a > 0$,

$$|S(1, n + 1)| \leq \frac{a}{n} |B(1, n)| ,$$

where $S(1, k)$ is the sphere of radius k centered at $1 \in G$. In terms of the growth function, this inequality can be re-written as

$$(18.20) \quad \frac{\mathfrak{G}(n + 1) - \mathfrak{G}(n)}{\mathfrak{G}(n)} \leq \frac{a}{n} .$$

Let $f(t)$ be the piecewise-linear function on \mathbb{R}_+ whose restriction to \mathbb{N} equals \mathfrak{G} and which is linear on every interval $[n, n + 1], n \in \mathbb{N}$. Then the inequality (18.20) means that for all $t \notin \mathbb{N}, t > 0$,

$$\frac{f'(t)}{f(t)} \leq \frac{a}{t} .$$

which, by integration, implies that

$$\log |f(t)| \leq a \log(t) + b$$

and, hence,

$$f(t) \leq bt^a .$$

In particular, it follows that $\mathfrak{G}(t)$ is bounded by a polynomial in t , whence G is virtually nilpotent by Theorem 16.1. \square

In view of Theorem 18.102, (1), one may ask if there is a general upper bound for the Følner functions of a group, same as the exponential function is a general upper bound for the growth functions; related to this, one may ask how much can the Følner function and the growth function of a group differ. These questions are addressed in two papers of A. Erschler:

The first theorem shows that one cannot have an exponential upper bound for Følner functions:

THEOREM 18.103 (A. Erschler, [Ers03]). *Let G and H be two amenable groups and assume that some representative F of F_o^H has the property that for every $a > 0$ there exists $b > 0$ so that $aF(t) < F(bt)$ for every $t > 0$. Then the Følner function of the wreath product $G \wr H$ is asymptotically equal to $[F_o^B(x)]^{F_o^A(x)}$.*

The second theorem shows that there are no upper bounds for Følner functions whatsoever:

THEOREM 18.104 (A. Erschler, [Ers06]). *For every function $f : \mathbb{N} \rightarrow \mathbb{N}$, there exists a finitely generated group G , which is a subgroup of a group of intermediate growth (in particular, G is amenable), whose Følner function satisfies $F_o^G(n) \geq f(n)$ for all sufficiently large n .*

18.10. Summary of equivalent definitions of amenability

Below we present a (very much incomplete) list of equivalent definitions of amenability; most of these are theorems stated or proven earlier in this chapter, the exceptions are the characterizations in terms of bounded cohomology groups and of measure-equivalence. In order to streamline the discussion, G is assumed to be an infinite finitely generated group equipped with a word metric.

- (1) G is amenable iff it admits an invariant mean.
- (2) G is amenable iff it admits a Følner sequence.

- (3) A finitely presented group G is amenable iff G is the fundamental group of a closed Riemannian manifold whose universal cover \tilde{M} satisfies $\lambda_1(\tilde{M}) = 0$ (Theorem 18.14).
- (4) G is amenable iff for all Banach $\mathbb{Z}G$ -modules V and all $n \geq 1$, $H_b^n(G, V) = 0 \iff \forall V, H_b^1(G, V) = 0$ (G. Noskov, [Nos91]).
- (5) G is amenable iff it is non-paradoxical, i.e. if $\text{Tar}(G) = \infty$ (Theorem 18.50).
- (6) G is amenable iff it is measure-equivalent to \mathbb{Z} (D. Ornstein and B. Weiss, [OW80]).
- (7) G is nonamenable iff there exists a constant $C > 0$ such that for every finite non-empty subset $\Phi \subset G$, the set $\overline{\mathcal{N}}_C(\Phi) \subset G$ contains at least twice as many vertices as Φ (Theorem 18.4).
- (8) G is nonamenable iff there exists a map $f \in \mathcal{B}(G)$ such that for every $v \in V$ the pre-image $f^{-1}(v)$ contains at least two elements (Theorem 18.4).
- (9) G is nonamenable iff there exists a map $f \in \mathcal{B}(G)$ such that for every $v \in V$ the pre-image $f^{-1}(v)$ contains exactly two elements (Theorem 18.4).
- (10) G is nonamenable iff its Cayley graph \mathcal{G} has spectral radius $\rho(\mathcal{G}) < 1$ (Theorem 18.11).

18.11. Amenable hierarchy

We conclude this chapter with the following diagram illustrating the hierarchy of amenable groups:

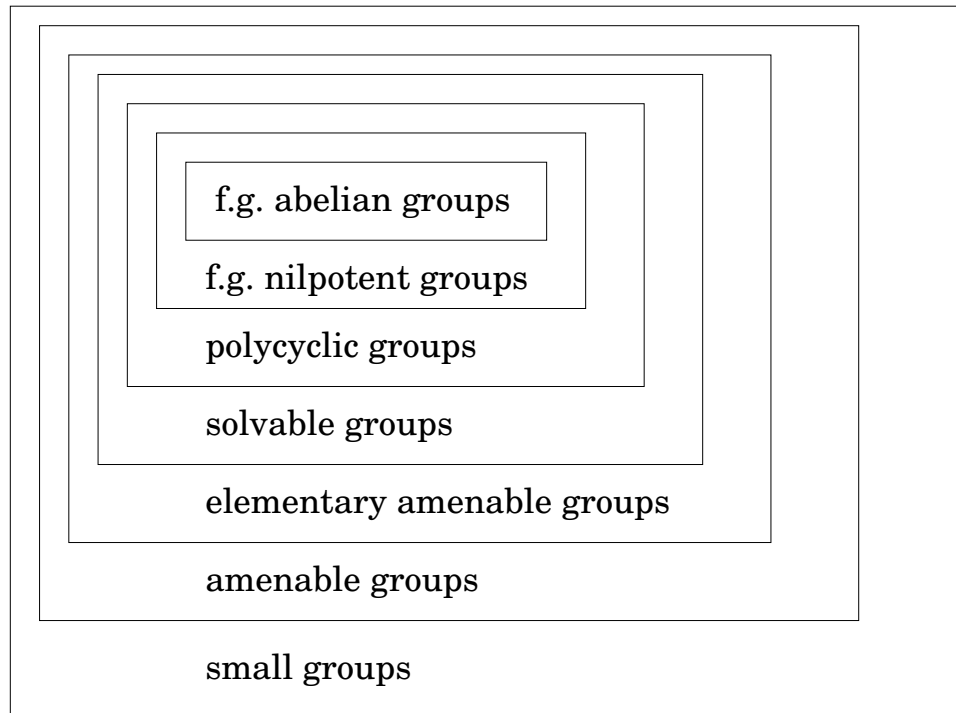


FIGURE 18.1. The hierarchy of amenable groups

Ultralimits, fixed point properties, proper actions

In this chapter we discuss various fixed-point properties, most notably, Kazhdan's Property (T), by comparison with properties of an opposite nature, such as amenability and a-T-menability.

19.1. Classes of Banach spaces stable with respect to ultralimits

We begin by discussing stability with respect to ultralimits of certain classes of Banach spaces. It is easy to see that ultralimits of Banach spaces are Banach spaces. Below, we will see that within the class of Banach spaces, Hilbert spaces and L^p -spaces can be distinguished by properties that are preserved under ultralimits. The main references for this section are [LT79], [Kak41] and [BDCK66].

CONVENTION 19.1. (a) Unless otherwise stated, for every ultralimit of Banach spaces, the base-points are the zero vectors. This assumption is harmless since translations are isometries of Banach spaces.

(b) We do not assume Hilbert spaces to be separable.

First we prove that ultralimits of Hilbert spaces are Hilbert spaces. While this can be done more quickly by using inner products and their limits, we prefer to provide another proof, which is a simplified version of the proof for the more general result that, for every $p \in [1, \infty)$, an ultralimit of L^p -spaces is an L^p -space. The main idea of these proofs is that, within the class of Banach spaces, Hilbert spaces, and more generally L^p -spaces, are characterized by certain (in)equalities. For Hilbert spaces, this characterization is as follows.

THEOREM 19.2 (Jordan–von Neumann [JvN35]). *A (real or complex) Banach space $(\mathcal{B}, \|\cdot\|)$ is Hilbert (i.e. the norm $\|\cdot\|$ comes from an inner product) if and only if every pair of vectors $x, y \in \mathcal{B}$ satisfies the parallelogram identity:*

$$\|x + y\|^2 + \|x - y\|^2 = 2\|x\|^2 + 2\|y\|^2.$$

PROOF. We claim that the inner/hermitian product on \mathcal{B} is given by the formula:

$$(x, y) := \frac{1}{4} (\|x + y\|^2 - \|x - y\|^2) = \frac{1}{4} \sum_{k=0}^1 (-1)^k \|x + (-1)^k y\|^2, \quad \text{if } \mathcal{B} \text{ is real}$$

and

$$(x, y) := \frac{1}{4} \sum_{k=0}^3 i^k \|x + i^k y\|^2, \quad \text{if } \mathcal{B} \text{ is complex,}$$

where $i = \sqrt{-1}$.

Note that it is clear that $(x, x) = \|x\|^2$ (real case), $(x, \bar{x}) = \|x\|^2$ (complex case). We will verify that (\cdot, \cdot) is a hermitian inner product in the complex case;

the real case is similar and it is left to the reader. We likewise leave to the reader to show that

$$(19.1) \quad (ix, y) = (x, -iy) = i(x, y), \quad (x, y) = \overline{(y, x)}$$

and that the parallelogram identity implies the equality

$$(19.2) \quad \|u + v\|^2 = -\|u\|^2 + \frac{1}{2}\|v\|^2 + 2\left\|u + \frac{1}{2}v\right\|^2.$$

By the definition of (\cdot, \cdot) , we have:

$$4(x/2, z) = \sum_{k=0}^3 i^k \left\| \frac{x}{2} + i^k z \right\|^2 =$$

(by applying the equation (19.2) to each term of this sum)

$$\sum_{k=0}^3 i^k 2 \left(\left\| \frac{x}{2} + i^k \frac{z}{2} \right\|^2 + 2 \left\| i^k \frac{z}{2} \right\|^2 - \|x/2\|^2 \right) =$$

(again, by the definition of (\cdot, \cdot))

$$\sum_{k=0}^3 2i^k \left(\left\| \frac{x}{2} + i^k \frac{z}{2} \right\|^2 + \left\| \frac{z}{2} \right\|^2 \right) = 2(x, z).$$

Thus, $(x/2, z) = \frac{1}{2}(x, z)$ and, clearly,

$$(19.3) \quad (2x, z) = 2(x, z)$$

By the symmetry of (\cdot, \cdot) it follows that

$$(19.4) \quad (x, 2z) = 2(x, z).$$

Instead of proving the multiplicative property for (\cdot, \cdot) for all scalars, we now prove the additivity property of (\cdot, \cdot) .

By the definition of (\cdot, \cdot) , we have

$$4(x + y, z) = \sum_{k=0}^3 \|(x + y) + i^k z\|^2 =$$

(by applying the parallelogram identity to each term of this sum)

$$\sum_{k=0}^3 i^k (2 \|x + i^k(z/2)\|^2 + \|y + i^k(z/2)\|^2) - \|x - y\|^2 =$$

$$\sum_{k=0}^3 i^k (2 \|x + i^k(z/2)\|^2 + \|y + i^k(z/2)\|^2) = 8(x, z/2) + 8(y, z/2) =$$

(by applying (19.4))

$$4(x, z) + 4(y, z).$$

Thus, $(x + y, z) = (x, z) + (y, z)$.

The additivity property of (\cdot, \cdot) applied inductively, yields

$$(nx, y) = n(x, y), \forall n \in \mathbb{N}.$$

For every $n \in \mathbb{N}$ we also have

$$(x, y) = \left(n \frac{1}{n} x, y \right) = n \left(\frac{1}{n} x, y \right) \Rightarrow \left(\frac{1}{n} x, y \right) = \frac{1}{n} (x, y).$$

Combined with the additivity property, this implies that $(rx, y) = r(x, y)$ holds for all $r \in \mathbb{Q}, r \geq 0$. Since $(-x, y) = -(x, y)$, we have the same multiplicative identity for all $r \in \mathbb{Q}$. Note that so far we did not use the triangle inequality in \mathcal{B} . Observe that the triangle inequality in \mathcal{B} implies that for all $x, y \in \mathcal{B}$ the function

$$t \mapsto (tx, y) = \frac{1}{4} (\|tx + y\|^2 - \|tx - y\|^2)$$

is continuous. Continuity implies that the identity $(tx, y) = t(x, y)$ holds for all $t \in \mathbb{R}$. This and (19.1) imply that the equality can be extended to $t \in \mathbb{C}$. Hence, by the symmetry of (\cdot, \cdot) as described in (19.1), it follows that (x, y) is indeed an inner product on \mathcal{B} . \square

COROLLARY 19.3. *Every ultralimit of a sequence of Hilbert spaces is a Hilbert space.*

EXERCISE 19.4. Every closed linear subspace of a Hilbert space is a Hilbert space.

EXERCISE 19.5. A Banach space is Hilbert if and only if it is a $CAT(0)$ -space.

Either Exercise 19.5 or, as mentioned before, a short argument involving inner products, provide alternative proofs of Corollary 19.3. Still, we preferred the use of the parallelogram identity because it is a piece of a proof for L^p -spaces and it illustrates in this simple case how the more general proof works.

A key feature of Banach function spaces is the existence of an order relation satisfying certain properties with respect to the algebraic operations and the norm.

DEFINITION 19.6. A *Banach lattice* $(\mathcal{B}, \|\cdot\|, \leq)$ is a real Banach space $(\mathcal{B}, \|\cdot\|)$ endowed with a partial order \leq such that:

- (1) if $x \leq y$ then $x + z \leq y + z$ for every $x, y, z \in \mathcal{B}$;
- (2) if $x \geq 0$ and λ is a non-negative real number then $\lambda x \geq 0$;
- (3) for every x, y in \mathcal{B} there exists a *least upper bound* (*l.u.b*), denoted by $x \vee y$, and a *greatest lower bound* (*g.l.b*), denoted by $x \wedge y$; this allows to define the *absolute value* of a vector $|x| = x \vee (-x)$;
- (4) if $|x| \leq |y|$ then $\|x\| \leq \|y\|$.

REMARKS 19.7. (a) It suffices to require the existence of one of the two bounds in Definition 19.6, (3). Either the relation $x \vee y + x \wedge y = x + y$ or the relation $x \wedge y = -[(-x) \vee (-y)]$ allows to deduce the existence of one bound from the existence of the other.

(b) The conditions (1), (2) and (3) in Definition 19.6 imply that

$$(19.5) \quad |x - y| = |x \vee z - y \vee z| + |x \wedge z - y \wedge z|.$$

This and the condition (4) imply that both operations \vee and \wedge on \mathcal{B} are continuous.

(c) The condition (4) applied to $x = u$ and $y = |u|$, and to $x = |u|$ and $y = u$ imply that $\|u\| = \||u|\|$.

DEFINITION 19.8. A *sublattice* in a Banach lattice $(\mathcal{B}, \|\cdot\|, \leq)$ is a linear subspace \mathcal{V} of \mathcal{B} such that if y, y' are elements of \mathcal{V} then $y \vee y'$ is in \mathcal{V} (hence $y \wedge y' = y + y' - y \vee y'$ is also in \mathcal{V}).

DEFINITION 19.9. Two elements $x, y \in \mathcal{B}$ of a Banach lattice are called *disjoint* if $x \wedge y = 0$.

EXERCISE 19.10. Prove that:

- (1) For every $p \in [1, \infty)$ and every measure space (X, μ) , the space $L^p(X, \mu)$ with the order defined by

$$f \leq g \Leftrightarrow f(x) \leq g(x), \mu\text{-almost surely,}$$

is a Banach lattice.

- (2) For every compact Hausdorff topological space K , the space $C(K)$ of continuous functions on K with the pointwise partial order and the sup-norm is a Banach lattice.
- (3) For both (1) and (2) prove that two functions are disjoint in the sense of Definition 19.9 if and only if both are non-negative functions with disjoint supports (up to a set of measure zero in the first case).

DEFINITION 19.11. Two Banach lattices $\mathcal{B}, \mathcal{B}'$ are *order isometric* if there exists a linear isometry $T : \mathcal{B} \rightarrow \mathcal{B}'$ which is also an order isomorphism. Such T is called an *order isometry*.

PROPOSITION 19.12 (Ultralimits of Banach lattices). *Any ultralimit of a family of Banach lattices has a canonical structure of a Banach lattice.*

PROOF. Let $(\mathcal{B}_i, \|\cdot\|_i, \leq_i), i \in I$, be a family of Banach lattices and let ω be a nonprincipal ultrafilter on I . Consider the ultralimit \mathcal{B}_ω endowed with the limit norm $\|\cdot\|_\omega$. We will need:

LEMMA 19.13. *Suppose that $a_i, b_i \in X_i$ are such that $u = \omega\text{-lim } a_i = \omega\text{-lim } b_i$. Then $u = \omega\text{-lim}(a_i \vee b_i) = \omega\text{-lim}(a_i \wedge b_i)$.*

PROOF. Equation (19.5) and Definition 19.6, (4), imply that

$$|x - y| \geq |x \vee z - y \vee z| \text{ and } |x - y| \geq |x \wedge z - y \wedge z|.$$

These inequalities applied to $x = a_i$ and $y = z = b_i$ imply that $|a_i \vee b_i - b_i| \leq |a_i - b_i|$ and $|a_i \wedge b_i - b_i| \leq |a_i - b_i|$. Part (4) of Definition 19.6 concludes the proof. \square

We define on \mathcal{B}_ω a relation \leq as follows:

Points $u, v \in \mathcal{B}_\omega$ satisfy $u \leq v$ if and only if there exist representatives $(x_i)_{i \in I}$ and $(y_i)_{i \in I}$ of u and v (i.e. $u = \omega\text{-lim } x_i$ and $v = \omega\text{-lim } y_i$) such that $x_i \leq y_i$ ω -almost surely.

We now verify that \leq is an order. Reflexivity of \leq is obvious. Let us check anti-symmetry. If $u \leq v$ and $v \leq u$ then we can write these vectors as

$$u = \omega\text{-lim } x_i = \omega\text{-lim } x'_i$$

and

$$v = \omega\text{-lim } y_i = \omega\text{-lim } y'_i,$$

so that ω -almost surely $x_i \leq y_i$ and $y'_i \leq x'_i$. The vectors $z_i = x_i - y'_i$ satisfy the inequalities $z_i \leq y_i - y'_i$ and $-z_i \leq x'_i - x_i$. This implies that

$$|z_i| \leq (y_i - y'_i) \vee (x'_i - x_i) \leq |y_i - y'_i| \vee |x'_i - x_i| \leq |y_i - y'_i| + |x'_i - x_i|.$$

Property (4) in Definition 19.6, the triangle inequality and Remark 19.7, (c), imply that

$$\|z_i\| \leq \|y_i - y'_i\| + \|x'_i - x_i\|.$$

It follows that $\omega\text{-lim } z_i = 0$, hence $u = v$.

We now check transitivity. Consider vectors

$$u = \omega\text{-lim } x_i, v = \omega\text{-lim } y_i = \omega\text{-lim } y'_i, w = \omega\text{-lim } z_i$$

such that ω -almost surely $x_i \leq y_i$ and $y'_i \leq z_i$. Then

$$x_i \leq z_i + y_i - y'_i.$$

Since $w = \omega\text{-lim}(z_i + y_i - y'_i)$, it follows that $u \leq w$.

We will now verify that \mathcal{B}_ω is a Banach lattice with respect to the order \leq . Properties (1) and (2) in Definition 19.6 are immediate.

Given $u = \omega\text{-lim } x_i$ and $v = \omega\text{-lim } y_i$ define $u \vee v$ as $\omega\text{-lim}(x_i \vee y_i)$. We claim that $u \vee v$ is well-defined, i.e. does not depend on the choice of representatives for u and v . Indeed, assume that $u = \omega\text{-lim } x'_i$ and $v = \omega\text{-lim } y'_i$ and take $w = \omega\text{-lim}(x_i \vee y_i)$ and $w' = \omega\text{-lim}(x'_i \vee y'_i)$. Let $a_i = x_i \wedge x'_i$ and $A_i = x_i \vee x'_i$; likewise, $b_i = y_i \wedge y'_i$ and $B_i = y_i \vee y'_i$. Clearly,

$$\omega\text{-lim}(a_i \vee b_i) \leq w \leq \omega\text{-lim}(A_i \vee B_i),$$

and the same for w' . The inequalities

$$a_i \vee b_i \leq A_i \vee B_i \leq a_i \vee b_i + A_i - a_i + B_i - b_i$$

imply that

$$\omega\text{-lim}(a_i \vee b_i) = \omega\text{-lim}(A_i \vee B_i)$$

and, therefore, $w = w'$. We conclude that the vector $u \vee v = \omega\text{-lim}(x_i \vee y_i)$ is well-defined. Clearly, $u \vee v \geq u$ and $u \vee v \geq v$. We need to verify that $u \vee v$ is the least upper bound for the vectors u, v .

Suppose that $z \geq u$, $z \geq v$, where $u = \omega\text{-lim } x_i$, $v = \omega\text{-lim } y_i$ and $z = \omega\text{-lim } z_i = \omega\text{-lim } z'_i$ such that ω -almost surely $z_i \geq x_i$ and $z'_i \geq y_i$. Lemma 19.13 implies that $z = \omega\text{-lim}(z_i \vee z'_i)$ and $z_i \vee z'_i \geq x_i \vee y_i$, whence, $z \geq (u \vee v)$.

Consider now $u, v \in X_\omega$ such that $|u| \leq |v|$. It follows that $u = \omega\text{-lim } x_i = \omega\text{-lim } x'_i$ and $v = \omega\text{-lim } y_i = \omega\text{-lim } y'_i$, where

$$x_i \vee (-x'_i) \leq y_i \vee (-y'_i).$$

Therefore:

$$|x_i| = x_i \vee (-x_i) \leq x_i \vee (-x'_i) + |x_i - x'_i| \leq |y_i| + |y_i - y'_i| + |x_i - x'_i|.$$

This inequality, part (4) of Definition 19.6 and the triangle inequality imply that

$$\|x_i\| \leq \|y_i\| + \|y_i - y'_i\| + \|x_i - x'_i\|.$$

In particular, $\|u\| \leq \|v\|$. □

It is a remarkable fact that L^p -spaces can be identified, up to order isometry, within the class of Banach lattices by a simple criterion that we will state below.

DEFINITION 19.14. Let $p \in [1, \infty)$. An *abstract L^p -space* is a Banach lattice \mathcal{B} such that for every pair of disjoint vectors $x, y \in \mathcal{B}$,

$$\|x + y\|^p = \|x\|^p + \|y\|^p.$$

Clearly, every space $L^p(X, \mu)$, with (X, μ) a measure space, is an abstract L^p -space. S. Kakutani proved that the converse is also true:

THEOREM 19.15 (Kakutani representation theorem [Kak41], see also Theorem 3 in [BDCK66] and Theorem 1.b.2 in [LT79]). *For every $p \in [1, \infty)$ every abstract L^p -space is order isometric to some space $L^p(X, \mu)$ for some measure space (X, μ) .*

COROLLARY 19.16. *For every $p \in [1, \infty)$ any closed sublattice of a space $L^p(X, \mu)$ is order isometric to some space $L^p(Y, \nu)$.*

COROLLARY 19.17. *Consider an indexed family of spaces $L^{p_i}(X_i, \mu_i)$, $i \in I$, such that $p_i \in [1, \infty)$. If ω is an ultrafilter on I such that $\omega\text{-lim } p_i = p$ then the ultralimit $\omega\text{-lim } L^{p_i}(X_i, \mu_i)$ is order isometric to some space $L^p(Y, \nu)$.*

COROLLARY 19.18. *For a fixed p , the family of spaces $L^p(X, \mu)$, where (X, μ) are measure spaces, is stable with respect to (rescaled) ultralimits.*

REMARK 19.19. The measure space (Y, ν) in Corollary 19.17 can be identified with the ultralimit of measure spaces (X_i, μ_i) . We refer to [Cut01] and [War12] for details of the construction of the *Loeb measure space*, which is the ultralimit of the measure spaces (X_i, μ_i) .

19.2. Limit actions and point-selection theorem

For finitely generated groups a *limit* of a family of actions may naturally occur in various settings, as noted in [Gro03] (see also [BFGM07], §3, c).

DEFINITION 19.20. We say that a topological action $\rho : G \curvearrowright X$ of a group G on a metric space X is a *Lipschitz action* if each $\rho(g), g \in G$, is a Lipschitz transformation.

In order to simplify the proofs, we will work, for the most part, with finitely generated groups equipped with the discrete topology. In this setting, it is useful to define a quantitative version of the notion of Lipschitz action, as follows.

DEFINITION 19.21. Let G be a finitely generated group, and let S be a finite generating subset of G ; it will be convenient to assume that $1 \in S$.

We say that a topological action $\rho : G \curvearrowright X$ of G on a metric space X is an *(L, S) -Lipschitz action*, for some $L \geq 1$, if each $\rho(s), s \in S$, is an L -bi-Lipschitz transformation.

We note that changing the finite generating set amounts to changing the Lipschitz constant L and, therefore, the particular choice of S is irrelevant.

Let $(X_i, \text{dist}_i), i \in I$, be a family of complete metric spaces and let

$$\rho_i : G \curvearrowright X_i$$

be non-trivial (L, S) -Lipschitz actions of the group G , with L independent of i .

For each i we define $F_i \subset X_i$, the set of points fixed by $\rho_i(G)$. Nontriviality of the action means that $F_i \neq X_i$ for each i .

THEOREM 19.22 (Point-selection theorem). *Let $x_i \in X_i \setminus F_i$ be a family of base-points. Assume that for some ultrafilter ω on I , either*

$$(19.6) \quad \omega\text{-lim } \frac{\text{dist}(x_i, F_i)}{\text{diam}(S_i x_i)} = \infty$$

or $F_i = \emptyset$ ω -almost surely. We take (δ_i) to be a sequence of positive numbers either equal to $\frac{\text{dist}(x_i, F_i)}{2 \text{diam}(S_i x_i)}$ in the former case, or arbitrary such that $\omega\text{-lim } \delta_i = +\infty$ in the latter case.

Then there exists an (L, S) -Lipschitz action $\rho_\omega : G \curvearrowright X_\omega$ on some rescaled ultralimit of the form

$$(19.7) \quad X_\omega = \omega\text{-lim}(X_i, y_i, \lambda_i \text{dist}_i), \text{ with } \lambda_i \geq \frac{2}{\text{diam}(S_i x_i)[1 + 2\delta_i(L + 1)]},$$

such that for every point z_ω in X_ω the diameter of $\rho_\omega(S)z_\omega$ is at least 1.

PROOF. We first note that, according to the choice of δ_i , (19.6) implies that

$$\omega\text{-lim } \delta_i = +\infty$$

when $F_i \neq \emptyset$ ω -almost surely. In what follows, for simplicity of notation, instead of writing $\rho_i(S)x$, we will write Sx for elements $x \in X_i, i \in I$.

PROPOSITION 19.23. For ω -almost every $i \in I$, there exists a point

$$y_i \in B(x_i, 2\delta_i \text{diam}(Sx_i))$$

such that for every point z in the ball $B\left(y_i, \frac{\delta_i \text{diam}(Sy_i)}{2}\right)$ the diameter of Sz is at least $\frac{\text{diam}(Sy_i)}{2}$.

PROOF. Assume to the contrary that ω -almost surely for every point y_i in $B(x_i, 2\delta_i \text{diam}(Sx_i))$ there exists

$$z_i \in B\left(y_i, \frac{\delta_i \text{diam}(Sy_i)}{2}\right)$$

such that the diameter of Sz is strictly less than $\frac{\text{diam}(Sy_i)}{2}$. Let $J \subset I$ be the set of indices such that the above holds, $\omega(J) = 1$, and let i be a fixed index in J . In what follows the argument is only for the index i and for simplicity we suppress the index i in our notation.

Set

$$D := 2\delta \text{diam}(Sx) = \text{dist}(x, F), \quad R := \frac{D}{2}.$$

Then for every point y in the ball $B(x, D) \subset X_i$, there exists

$$z \in B\left(y, \frac{\delta \text{diam}(Sy)}{2}\right)$$

such that $\text{diam}(Sz) < \frac{\text{diam}(Sy)}{2}$. Applied to $y = x$, it follows that there exists

$$u_1 \in B\left(x, \frac{R}{2}\right),$$

such that $\text{diam}(Su_1) < \frac{\text{diam}(Sx)}{2}$. Applied to u_1 , the same statement implies that there exists

$$u_2 \in B\left(x_1, \frac{\delta \text{diam}(Su_1)}{2}\right) \subset B\left(x, \frac{R}{2} + \frac{R}{4}\right)$$

such that

$$\text{diam}(Su_2) < \frac{\text{diam}(Su_1)}{2} < \frac{\text{diam}(Sx)}{2^2}.$$

Assume that we thus found points $u_1, u_2, \dots, u_n \in X_i$ such that for every $j \in \{1, 2, \dots, n\}$,

$$(19.8) \quad u_j \in B\left(x_{j-1}, \frac{\delta \text{diam}(Su_{j-1})}{2}\right) \subset B\left(x, \frac{R}{2} + \frac{R}{4} + \dots + \frac{R}{2^j}\right)$$

and $\text{diam}(Su_j) < \frac{\text{diam}(Sx)}{2^j}$. Then, by taking $y = u_n$, we conclude that there exists

$$u_{n+1} \in B\left(u_n, \frac{\delta \text{diam}(Su_n)}{2}\right) \subset B\left(x, \frac{R}{2} + \frac{R}{4} + \dots + \frac{R}{2^n} + \frac{R}{2^{n+1}}\right)$$

such that

$$(19.9) \quad \text{diam}(Su_{n+1}) < \frac{\text{diam}(Su_n)}{2} < \frac{\text{diam}(Su)}{2^{n+1}}.$$

We thus obtain a Cauchy sequence (u_n) in a complete metric space X_i ; therefore, (u_n) converges to a point u in X_i . By the inequalities (19.9), taking into account that the action of G on X_i is continuous, we conclude that

$$\text{diam}(Su) = 0,$$

and, hence, u is fixed by S , thus by the entire group G (since S generates G). Furthermore, (19.8) implies that $\text{dist}(u, x) \leq R$. On the other hand, $R = \frac{\text{dist}(x, F)}{2}$, where F is the set of points fixed by G , a contradiction. \square

Thus, ω -almost surely there exists y_i in $B(x_i, 2\delta_i \text{diam}(Sx_i))$ such that for every point

$$z \in B\left(y_i, \frac{\delta_i \text{diam}(Sy_i)}{2}\right),$$

the diameter of Sz is at least $\frac{\text{diam}(Sy_i)}{2}$. Define

$$(19.10) \quad \lambda_i := \frac{2}{\text{diam}(Sy_i)}.$$

Then triangle inequalities and the fact that each generator $s \in S$ acts on X_i as an L -bi-Lipschitz transformation, imply that

$$\lambda_i \geq \frac{2}{\text{diam}(Sx_i)(1 + 2\delta_i(L + 1))}.$$

Furthermore, for each $s \in S$ we have

$$\lambda_i \text{dist}(sy_i, y_i) \leq 2.$$

We therefore obtain an (L, S) -Lipschitz action ρ_ω of the group G on the ultralimit

$$X_\omega = \omega\text{-lim}(X_i, y_i, \lambda_i \text{dist}_i),$$

cf. Lemma 10.83. Since $\omega\text{-lim}_i \delta_i = \infty$, it follows that the natural inclusion map

$$\omega\text{-lim}\left(B\left(y_i, \frac{\delta_i \text{diam}(Sy_i)}{2}\right), y_i, \lambda_i \text{dist}_i\right) \rightarrow \omega\text{-lim}(X_i, y_i, \lambda_i \text{dist}_i)$$

is surjective. Hence, for every point z_ω in X_ω the diameter of Sz_ω is at least 1. \square

REMARK 19.24. Note that it could happen that the limit $\omega\text{-lim} \lambda_i$ is positive or even infinite. However, if

$$\omega\text{-lim} \inf_{x \in X_i} \text{diam}(\rho_i(S)(x)) = \infty,$$

then (19.10) implies that $\omega\text{-lim} \lambda_i = 0$. This situation appears frequently when one constructs group actions on real trees associated with divergent sequences of isometric actions $G \curvearrowright X$, where X is a δ -hyperbolic metric space, see Section 11.23, as well as [Kap01] for applications to the theory of Kleinian groups.

Our next goal is to sharpen a bit the conclusion of the Point Selection Theorem. Suppose, in addition, that the family I is the poset (\mathbb{N}, \leq) with the standard order. Assume also that $(N_i)_{i \in I}$ is a directed collection of normal subgroups in G , i.e. if $i \leq j$ then $N_i \leq N_j$. Accordingly, we obtain a collection of quotient groups $G_i = G/N_i$ and projection homomorphisms

$$p_i : G \rightarrow G_i, \quad f_{ij} : G_i \rightarrow G_j, \quad i \leq j.$$

The direct limit of the corresponding direct system of groups $(G_i, f_{ij})_{i, j \in I}$ is naturally isomorphic to the quotient group $\overline{G} = G/N$, where

$$N = \bigcup_{i \in I} N_i.$$

We let $p : G \rightarrow \overline{G}$ denote the quotient map.

For every infinite subset $J \subset I$, the natural homomorphism of direct limits

$$\varinjlim_{j \in J} G_j \longrightarrow \varinjlim_{i \in I} G_i \simeq \overline{G}$$

is an isomorphism, since the subset J is cofinal in (I, \leq) , when $I = \mathbb{N}$. We thus obtain the following addendum to Theorem 19.22:

THEOREM 19.25. *Assume that each representation ρ_i in Theorem 19.22 factors through the projection $p_i : G \rightarrow G_i$. Then for each nonprincipal ultrafilter ω , each limit action ρ_ω in Theorem 19.22 factors through a Lipschitz action $\overline{G} \curvearrowright X_\omega$.*

Theorem 19.22 allows one to prove certain fixed point properties for actions of groups using ultralimits. We will discuss such applications in what follows. Let \mathcal{C} be a collection of metric spaces, let $L \geq 1$ and let G be a (discrete) group.

DEFINITION 19.26. We say that a group G has the *fixed point property FC* if for every isometric action of G on every space $X \in \mathcal{C}$, the group G fixes a point in X .

REMARK 19.27. Note that one can further strengthen the property *FC* by considering Lipschitz actions, in the sense of Definition 19.20, instead of isometric actions. Some proofs in this and the following section go through in this setting without much change. However, we will not pursue this direction here.

The next corollary is an application of Theorem 19.25. Suppose that $G, (G_i)_{i \in \mathbb{N}}$ are as in Theorem 19.25. In particular, G is a finitely generated group and S is its finite generating set. We equip each quotient group G_i and the direct limit group \overline{G} with the finite generating sets S_i, \overline{S} , which are the images of S under the projections $p_i : G \rightarrow G_i, p : G \rightarrow \overline{G}$.

COROLLARY 19.28. *Let \mathcal{C} be a class of complete metric spaces stable with respect to rescaled ultralimits taken using countably infinite index sets I . If the group \overline{G} has Property *FC*, then there exists i_0 and $\varepsilon > 0$, such that for every $i > i_0$ the group G_i has Property *FC*; furthermore, for every isometric action φ_i of G_i on some space $(X_i, \text{dist}_i) \in \mathcal{C}$, if F_i is the set of points in X_i fixed by G_i , then for every point $x \in X_i$ the diameter of $\varphi_i(S_i)x$ is at least $\varepsilon \text{dist}_i(x, F_i)$ (and, obviously, at most $2 \text{dist}_i(x, F_i)$).*

PROOF. Assume to the contrary that for every $\varepsilon > 0$ and i_0 there exists an $i > i_0$ such that G_i has an isometric action ρ_i on some space $X_i \in \mathcal{C}$ either without fixed points or such that for some point $x_i \in X_i$ the diameter of $\rho_i(S_i)x_i$ is at most $\varepsilon \text{dist}_i(x, F_i)$. Then there exists a strictly increasing sequence of indices i_n and a sequence of actions φ_{i_n} of G_{i_n} on some $X_{i_n} \in \mathcal{C}$ either without fixed points or with points $x_n \in X_{i_n}$ satisfying

$$\text{diam}(\varphi_{i_n}(S_{i_n})x_n) < \frac{1}{n} \text{dist}_{i_n}(x_n, F_{i_n}).$$

We consider the case of non-empty fixed-point sets since the other case is analogous. Let ω be a nonprincipal ultrafilter on \mathbb{N} containing the sequence (i_n) . Then

$$\omega\text{-lim} \frac{\text{dist}_{i_n}(x_n, F_{i_n})}{\text{diam}(\varphi_{i_n}(S_{i_n})x_n)} = +\infty.$$

Theorem 19.25 then yields a contradiction, as it results in a fixed-point free action of \overline{G} on some space $X_\omega \in \mathcal{C}$. \square

In what follows we list several special cases of fixed point properties and proper action properties that are important in group theory. We begin by recalling some basic terminology and facts on isometric actions on Banach spaces. Note that the terminology established in the setting of actions on Banach spaces slightly deviates from the one used for actions on locally compact spaces topological spaces, as described in Section 5.6.1, even though in particular cases they coincide.

Given a Banach space $(\mathcal{B}, \|\cdot\|)$, its *unitary group* $U(\mathcal{B})$ is the group of linear invertible operators $U : \mathcal{H} \rightarrow \mathcal{H}$ that are isometries, i.e. $\|Ux\| = \|x\|$.

A *continuous unitary representation* of a topological group G on a Banach space \mathcal{B} is a homomorphism $\pi : G \rightarrow U(\mathcal{B})$ such that for every $x \in \mathcal{B}$, the map from G to \mathcal{B} defined by $g \mapsto gx$, is continuous.

By the Mazur-Ulam theorem (see e.g. [FJ03, p. 6]), every (not necessarily linear) isometry of a real Banach space \mathcal{B} is an affine transformation $v \mapsto Tv + b$, where $T \in U(\mathcal{B})$ and $b \in \mathcal{B}$. Thus, $\text{Isom}(\mathcal{B})$, the isometry group of \mathcal{B} , has the form $\mathcal{B} \rtimes U(\mathcal{B})$, where the first factor \mathcal{B} is identified to the group of translations on \mathcal{B} .

A *continuous isometric affine action* of a topological group G on a Banach space \mathcal{B} is a homomorphism

$$\alpha : G \rightarrow \mathcal{B} \rtimes U(\mathcal{B}) < \text{Isom}(\mathcal{B})$$

such that for every $x \in \mathcal{B}$, the orbit map $G \rightarrow \mathcal{B}$ defined by $g \mapsto gx$, is continuous. Such an action is *metrically proper* if for every $x \in \mathcal{B}$ and $D > 0$ there exists a compact subset $K \in G$ such that $g \notin K$ implies $\|gx\| > D$.

CONVENTION 19.29. Unless otherwise stated all the representations of topological groups that we consider are continuous and we omit to specify it each time.

Finitely generated groups are endowed, by default, with the discrete topology, therefore in that case the continuity condition is always satisfied.

19.3. Properties for actions on Hilbert spaces

For the rest of this chapter we focus on two of the most important properties for infinite groups in regard to their isometric affine actions on Hilbert spaces. In this section we introduce the two definitions, give several examples of groups satisfying these definitions and list a few corollaries of the two properties.

DEFINITION 19.30. A topological group G has *Property FH* if every continuous affine isometric action of G on a Hilbert space has a fixed point.

In view of the fixed-point theorem for isometric group actions on $CAT(0)$ spaces (Theorem 3.74), we obtain:

COROLLARY 19.31 (A. Guichardet). *A group G has Property FH if and only if every continuous affine isometric action of G on a Hilbert space has a bounded orbit.*

A strong negation of Property FH is the *a-T-menability property*:

DEFINITION 19.32. A topological group G is said to be *a-T-menable* if there exists a (metrically) proper affine isometric action $G \curvearrowright \mathcal{H}$ of G on a Hilbert space \mathcal{H} .

Examples of a-T-menable groups are:

- (1) Amenable groups, see Corollary 19.43.
- (2) Closed subgroups of $SO(n, 1)$ and $SU(n, 1)$. This is a theorem of J. Faraut and K. Harzallah [FH74]. See also [CCJ⁺01], Theorem 19.61 and the example following Definition 6.48. In particular, free groups and surface groups are a-T-menable.
- (3) Discrete groups G which admit isometric metrically proper actions on $CAT(0)$ cube complexes,¹ see [NR97b]. In particular, all Coxeter groups and all RAAGs are a-T-menable.
- (4) Various small cancelation groups, see [Wis04, OW11].

Examples of groups satisfying Property FH are:

- (1) All compact groups.
- (2) semisimple Lie groups whose Lie algebras do not contain factors isomorphic to the Lie algebra of $SO(n, 1)$ or $SU(n, 1)$, see [BdlHV08].
- (3) Lattices in Lie groups as in (2), e.g. $SL(n, \mathbb{Z})$, $n \geq 3$ see [BdlHV08].

Property FH is incompatible with non-trivial actions on certain hyperbolic spaces. In this vein, recall (see Definition 3.77) that a group G has *Property FA* if every isometric action of G on a complete real tree has a fixed point.² A link between the two fixed-point properties is the following theorem which was first proven independently by R. Alperin [Alp82] and Y. Watatani [Wat82]:

THEOREM 19.33 (R. Alperin; Y. Watatani). *FH \Rightarrow FA:
Every (discrete) group with Property FH also has Property FA.*

We will prove this theorem in section 19.5. The converse implication in Theorem 19.33 is not true in general: Coxeter groups with connected graphs have FA, as an easy consequence of Helly's Theorem in trees, while they are a-T-menable [BJS88]. Still, a variant of the converse of Theorem 19.33 holds, if instead of trees we take the larger class of median spaces, as described in Chapter 6. In Section 19.6 we thus provide characterizations of properties FH and of a-T-menability in terms of actions on median spaces.

¹The key here is that the vertex set of each cube complex admits a G -invariant structure of a space with measured walls, hence, of a median space on which G acts properly. Given this, Theorem 19.61 implies that G is a-T-menable.

²Some authors restrict to group actions on simplicial trees.

A similar result of incompatibility between property FH and non-trivial actions holds for actions on the real and complex hyperbolic spaces. The following is a consequence of the previously cited theorem of J. Faraut and K. Harzallah [FH74], that the group of isometries of a real-hyperbolic space or a complex hyperbolic space is a-T-menable, and of Corollary 19.52.

THEOREM 19.34. *If a group G has Property FH, then every isometric action of G on a real-hyperbolic space or on a complex hyperbolic space has a fixed point.*

At the same time, there are infinite hyperbolic groups satisfying Property FH. For instance, let $X = \mathbf{H}\mathbb{H}^n$ be the *quaternionic hyperbolic symmetric space* of quaternionic dimension $n \geq 2$, or $X = \mathbf{O}\mathbb{H}^2$ the octonionic hyperbolic plane, see Section 4.9. Let G denote the isometry group of X . We note that such G contains both uniform and non-uniform lattices; every uniform lattice $\Gamma < G$ acts isometrically properly discontinuously cocompactly on X , which implies that Γ is an infinite hyperbolic group. The group G and all its lattices have property FH [BdlHV08].

There are other examples of infinite hyperbolic groups satisfying Property FH, we will discuss them in more detail in section 19.8. It turns out that both in the triangular model and in Gromov’s model of randomness for groups “a majority of groups” are infinite hyperbolic with Property FH, when a certain parameter d , called *density*, of the randomness model, satisfies $\frac{1}{3} < d < \frac{1}{2}$, see [Žuk03]. On the other hand, for d varying in the interval $[0, 5/24]$, a majority of groups are infinite, hyperbolic and without Property FH, see [MP15].

We further note that a countable group satisfies both Property FH and a-T-menability if and only if it is finite (see Corollary 19.43).

On the other hand, there are finitely generated groups which fail both properties FH and a-T-menability. For instance the free product of two infinite groups $G_1 \star G_2$ where G_1, G_2 satisfy Property FH, e.g. $G_1 = G_2 = SL(3, \mathbb{Z})$. Indeed, since G is a non-trivial free product then G acts on a simplicial tree without fixed points, hence it cannot satisfy Property FH. On the other hand, G cannot act metrically properly on a Hilbert space \mathcal{H} since the groups G_1, G_2 do satisfy Property FH and, hence, have global fixed points in \mathcal{H} .

In section 19.4 we will discuss the relation between Property FH and Kazhdan’s Property (T), as well as the relation between a-T-menability and the Haagerup Property. We will see that for (discrete) groups, these two pairs of respective properties are equivalent. The equivalence is actually true for topological groups that are locally compact second countable, for a proof of this latter more general equivalences we refer to [BdlHV08] and [CCJ⁺01].

EXERCISE 19.35. Show that a discrete group G has Property FH if and only if $H^1(G, \mathcal{H}_\pi) = 0$ for every unitary representation $\pi : G \rightarrow U(\mathcal{H})$. *Hint:* Use Lemma 5.135.

19.4. Kazhdan’s Property (T) and the Haagerup property

In this section we recall the definitions of Kazhdan’s Property (T) and the Haagerup Property. We relate these properties to amenability. Furthermore, we also prove the equivalence (in the case of countable discrete groups) of Property (T) to Property FH, and of the Haagerup Property to a-T-menability. Our discussion is far from exhaustive, we refer the reader to [BdlHV08, CCJ⁺01] for an in-depth treatment.

DEFINITION 19.36. Let (π, \mathcal{H}) be a (continuous) unitary representation of a topological group G , where \mathcal{H} is a Hilbert space and $\pi : G \rightarrow U(\mathcal{H})$ is a continuous representation.

- (1) Given a subset $S \subseteq G$ and a number $\varepsilon > 0$, a unit vector x in \mathcal{H} is (S, ε) -invariant if

$$\sup_{g \in S} \|\pi(g)x - x\| \leq \varepsilon \|x\|.$$

- (2) The representation (π, \mathcal{H}) has almost invariant vectors if it has (K, ε) -invariant vectors for every compact subset K of G and every $\varepsilon > 0$.
- (3) The representation (π, \mathcal{H}) has invariant vectors if there exists a unit vector x in \mathcal{H} such that $\pi(g)x = x$ for all $g \in G$.

Clearly, the existence of invariant vectors implies the existence of almost invariant vectors. It is a remarkable fact that there are many groups for which the converse holds as well. This, in the language of the *Fell topology on the space of irreducible unitary representations*, means that the trivial representation is isolated, therefore the property is denoted by (T), to emphasize this isolation.

DEFINITION 19.37. A locally compact Hausdorff topological group G has *Kazhdan's Property (T)* if for every unitary representation π of G , if π has almost invariant vectors, then it also has invariant vectors.

Clearly, Property (T) is inherited by quotient groups.

A strong negation of Property (T), which under certain conditions will turn out to be equivalent to a-T-menability while (T) is equivalent to FH, is the *Haagerup Property*:

DEFINITION 19.38. A topological group G is said to have the *Haagerup Property* if either G is compact or there exists a unitary representation $\pi : G \rightarrow U(\mathcal{H})$ such that:

1. π is *mixing*, i.e. $\lim_{g \rightarrow \infty} \langle \pi(g)\xi, \eta \rangle = 0$ for all $\xi, \eta \in \mathcal{H}$, where $G \cup \{\infty\}$ is the one-point compactification of G .
2. G has almost invariant vectors.

In particular, the Haagerup Property is inherited by closed subgroups.

In order to simplify the discussion, in the sequel we will primarily limit ourselves to groups equipped with the discrete topology. Below, for a discrete group G , $\ell^2(G) = \ell^2(G, \mu)$ where μ is the counting measure (measure of each finite subset equals its cardinality).

THEOREM 19.39 (D. Kazhdan). *Each discrete group G satisfying Property (T) is finitely generated.*

PROOF. For each finitely generated subgroup $H \leq G$ we define the quotient G/H ; the group G acts on this space by the left multiplication. Accordingly, G has a unitary representation $\pi_{G/H}$ on the Hilbert space $\ell^2(G/H)$. Let $\mathbf{1}_H \in \ell^2(G/H)$ denote the indicator function of the coset H in G/H ; this is a unit vector fixed by the representation $\pi_{G/H}$. Now, consider the infinite direct sum

$$V := \bigoplus_{H \leq G} \ell^2(G/H)$$

taken over all finitely generated subgroups $H \leq G$. This is a pre-Hilbert space, we let \mathcal{H} denote its completion, a Hilbert space. The unitary representations $\pi_{G/H}$ yield a unitary representation π of G on \mathcal{H} . We regard each $\ell^2(G/H)$ as a subspace in \mathcal{H} . Then each vector $\mathbf{1}_H \in \ell^2(G/H) \subset \mathcal{H}$ is fixed by the subgroup H . Therefore, for every finite subset $K \subset G$ generating the subgroup $H = \langle K \rangle$, the vector $\mathbf{1}_H$ is fixed by the action of K . It follows that the representation (π, \mathcal{H}) has almost invariant vectors. Since G satisfies Property (T), π has a fixed unit vector $x \in \mathcal{H}$. Let x^* denote the (non-zero) linear functional on \mathcal{H} dual to x :

$$x^*(v) = \langle x, v \rangle.$$

This functional is G -invariant and has non-zero restriction to the dense subspace V , hence, non-zero G -invariant restriction to one of the subspaces $\ell^2(G/H)$. Therefore, G has a fixed non-zero vector u in $\ell^2(G/H)$. Since G acts transitively on G/H , the function $u \in \ell^2(G/H)$ has to be constant. It follows that the set G/H is finite. Thus, the finitely generated group H has finite index in G . Therefore, G itself is finitely generated. \square

PROPOSITION 19.40. *The action $G \curvearrowright \ell^2(G)$ is mixing for every infinite group G .*

PROOF. We consider the action on the left $G \curvearrowright \ell^2(G)$ defined by $g \cdot f(x) = f(g^{-1}x)$.

Fix $\epsilon > 0$. It suffices to consider functions $\xi, \eta \in \ell^2(G)$ with unit norm. Take a finite subset $F \subset G$ such that

$$\max \left[\left(\int_{G \setminus F} \xi^2 \right)^{1/2}, \left(\int_{G \setminus F} \eta^2 \right)^{1/2} \right] < \frac{\epsilon}{2}.$$

Let $g \in G$ be such that $g^{-1}(F) \cap F = \emptyset$. In what follows we denote by χ_A the characteristic function of a subset A of G .

We have

$$\begin{aligned} |\langle g \cdot \xi, \eta \rangle| &\leq |\langle g \cdot (\xi \chi_{G \setminus F}), \eta \rangle| + |\langle g \cdot (\xi \chi_F), \eta \rangle| \leq \\ &\frac{\epsilon}{2} + |\langle g \cdot (\xi \chi_F), \eta \chi_{G \setminus F} \rangle| + |\langle g \cdot (\xi \chi_F), \eta \chi_F \rangle| < \epsilon \end{aligned}$$

The conclusion of the Lemma follows. \square

COROLLARY 19.41. *If G is a discrete group that satisfies both Property (T) and Haagerup, then G is finite.*

PROOF. Assume that G is a group that satisfies the Haagerup Property. Consider the action $G \curvearrowright \ell^2(G)$. This action is mixing, and it has almost invariant vectors, since G satisfies the Haagerup Property. However, the only G -invariant functions on G are constant, hence there exist non-zero constant functions in $\ell^2(G)$. It follows that G is finite. \square

PROPOSITION 19.42 (A. Hulanicky [Hul66]; H. Reiter [Rei65]). *A discrete group G satisfies the Følner property if and only if the action of G on $\mathcal{H} = \ell^2(G)$ via left multiplication has almost invariant vectors.*

PROOF. We will prove only the direct implication, needed for the proof of Corollary 19.43. We refer the reader to [BdlHV08, Theorem G.3.2] for a proof of

the converse. Let $F_i \subset G$ be a Følner sequence. Let $f_i = \frac{1}{N_i} \mathbf{1}_{\Omega_i}$, where $N_i := |\Omega_i|^{1/2}$ and $\mathbf{1}_{\Omega_i}$ denotes the characteristic function of Ω_i . Then for every $g \in G$

$$\|g \cdot f_i - f_i\|^2 \leq \frac{|g\Omega_i \Delta \Omega_i|}{|\Omega_i|},$$

which converges to zero by the definition of a Følner sequence. \square

Since the action of an infinite group G on $\mathcal{H} = \ell^2(G)$ is mixing, we obtain the following:

COROLLARY 19.43. *For discrete groups the Følner property implies the Haagerup Property. In particular, a discrete amenable group which has Property (T) must be finite.*

For discrete groups (and more generally, for locally compact groups) one can define Property (T) in a more quantitative manner.

COROLLARY 19.44. *For each discrete group G satisfying Property (T) and equipped with a finite generating set S , there exists a number $\epsilon > 0$ such that whenever a unitary representation π has an (S, ϵ) -invariant vector, π has a non-zero invariant vector.*

PROOF. Suppose that this assertion fails and consider a sequence of positive numbers $\epsilon_i \rightarrow 0$ and unitary representations $\pi_i : G \rightarrow O(\mathcal{H}_i)$, such that π_i has an (S, ϵ_i) -invariant unit vector in \mathcal{H}_i but no invariant vectors. As in the proof of Theorem 19.39, consider the natural action π of G on the completion \mathcal{H} of the direct sum of Hilbert spaces

$$\bigoplus_i \mathcal{H}_i.$$

Then, as in the proof of Theorem 19.39, π has almost invariant vectors but no invariant vectors. This is a contradiction. \square

DEFINITION 19.45. Each pair $(S, \epsilon) \subset G \times \mathbb{R}$ satisfying this corollary is called a *Kazhdan pair* for G . A number $\epsilon > 0$ for which there exists $S \subset G$, such that (S, ϵ) is a Kazhdan pair, is called a *Kazhdan constant* of G .

The next theorem is due to Delorme [Del77] (who proved the implication $T \Rightarrow FH$) and Guichardet [Gui77] (who proved the opposite implication), see also [dlHV89] and [CCJ⁺01]. In these references the theorem is proven in greater generality, namely, in the setting of second countable, locally compact, Hausdorff topological groups; we limit ourselves to finitely generated groups with discrete topology.

However, first we need to define several notions pertaining to *kernels on groups*: *G*-invariant, bounded and proper kernels on groups.

DEFINITION 19.46. *G*-invariant kernel

- (1) A *left-invariant* (or, simply, *invariant*) kernel on a topological group G is a continuous kernel $\psi : G \times G \rightarrow \mathbb{R}$ such that $\psi(gh, gf) = \psi(h, f)$ for all $g, h, f \in G$.
- (2) A *G*-invariant kernel on a group is said to be *bounded* if the function $G \rightarrow \mathbb{R}$,

$$g \mapsto \psi(1, g)$$

is bounded.

(3) A G -invariant kernel on a group is said to be *proper* if the function

$$g \mapsto \psi(1, g)$$

is proper, i.e. when $g \rightarrow \infty$ (in the one-point compactification of G), $\psi(1, g) \rightarrow \infty$.

THEOREM 19.47. *Let G be a finitely generated group. The following are equivalent:*

- (1) G has Property FH;
- (2) G has Property (T);
- (3) Every conditionally negative semidefinite G -invariant kernel $\psi : G \times G \rightarrow \mathbb{R}$ is bounded.

PROOF. We will deduce the implication (1) \Rightarrow (2) from Theorem 19.22. Our proof follows [Sil] and [Gro03].

Let G be a finitely generated group with Property FH and assume that it does not satisfy (2). Fix a finite generating set S of G . Then, for every $n \in \mathbb{N}$ there exists a unitary representation $\pi_n : G \rightarrow U(\mathcal{H}_n)$ with an $(S, \frac{1}{n})$ -invariant (unit) vector x_n and no invariant vectors. Let X_n be the unit sphere $\{u \in \mathcal{H}_n : \|u\| = 1\}$ with the induced path metric dist_n . Theorem 19.22 applied to the sequence of isometric actions of G on X_n and a choice of δ_n such that $\omega\text{-lim } \delta_n = +\infty$ and $\omega\text{-lim}[\delta_n \text{diam}(Sx_n)] = 0$, implies that G acts by isometries on a rescaled ultralimit

$$X_\omega = \omega\text{-lim}(X_n, x_n, \lambda_n \text{dist}_n), \text{ with } \lambda_n \geq \frac{2}{(1 + 2\delta_n) \text{diam}(Sx_n)}.$$

Note that $\omega\text{-lim } \lambda_n = +\infty$. Moreover, for every point $z_\omega \in X_\omega$ the diameter of Sz_ω is at least 1. Since $\omega\text{-lim } \lambda_n = +\infty$, Example 10.65 shows that the ultralimit X_ω is isometric to a Hilbert space \mathcal{H} . We thus obtain an isometric action of G on a Hilbert space \mathcal{H} without a global fixed point, contradicting Property FH.

In order to prove the implication (2) \Rightarrow (1), we will need a construction described in the following lemma.

LEMMA 19.48. *Let \mathcal{H} be a Hilbert space and $s > 0$ an arbitrary positive number. Let \mathcal{H}'_s be the space of finitely supported real-valued functions on \mathcal{H} endowed with the inner product:*

$$\langle f, g \rangle = \sum_{x, y \in \mathcal{H}} e^{-s\|x-y\|} f(x)g(y),$$

and let \mathcal{H}_s denote the completion of \mathcal{H}'_s with respect to the norm defined by this inner product.

- (1) The map F_s from \mathcal{H} to the unit sphere in \mathcal{H}_s defined by $F_s(x) := \mathbf{1}_x$ is equivariant with respect to the representation $\rho_s : \text{Isom}(\mathcal{H}) \rightarrow U(\mathcal{H}_s)$ extending the action of $\text{Isom}(\mathcal{H})$ on \mathcal{H}_s by pre-composition. Moreover, this map satisfies the property:

$$\lim_{\|x-y\| \rightarrow \infty} \langle F_s(x), F_s(y) \rangle = 0.$$

- (2) Suppose that G is a (discrete) group acting isometrically on \mathcal{H} with unbounded orbits. Then $\rho_s(G)$ has no non-zero fixed vectors in \mathcal{H}_s .

PROOF. The proof of (1) is a straightforward calculation.

(2) Suppose that $v \in \mathcal{H}_s$ is a vector fixed by $\rho_s(G)$. Consider a sequence $g_n \in G$ such that $\|g_n(x)\| \rightarrow \infty$ for (one/all) $x \in \mathcal{H}$. By (1) and the definition of \mathcal{H}_s

$$\lim_{n \rightarrow \infty} \langle v, F_s(g_n x) \rangle = 0.$$

We then have

$$\langle v, F(g_n x) \rangle = \langle g_n v, F(g_n x) \rangle = \langle g_n v, g_n F(x) \rangle = \langle v, F(x) \rangle.$$

Hence $\langle v, F(x) \rangle = 0$ for all $x \in \mathcal{H}$. Since the vectors $F(x), x \in \mathcal{H}$, span a dense subset in \mathcal{H}_s , it follows that $v = 0$. \square

We now return to the proof of the theorem. Suppose that G is finitely generated and satisfies Property (T). Let ϵ be a Kazhdan constant of G with respect to a finite generating set S . For an arbitrary affine isometric action on a Hilbert space $G \curvearrowright \mathcal{H}$, consider a parameter s such that for a unit vector $u \in \mathcal{H}_s$, the diameter of the set $\rho_s(S)(u)$ is $< \epsilon$. It follows that $\rho_s(G)$ has to fix a non-zero vector in \mathcal{H}_s . By Lemma 19.48, the action $G \curvearrowright \mathcal{H}$ has to have bounded orbits, hence, G must fix a point in \mathcal{H} .

Lastly, we will prove the equivalence of (3) and (1). If ψ is an unbounded conditionally negative semidefinite G -invariant kernel on G , then, according to Theorem 2.90, there exists a Hilbert space \mathcal{H} and a representation

$$\rho : G \rightarrow \text{Isom}(\mathcal{H})$$

such that

$$(19.11) \quad \psi(g, h) = \|\rho(g)(0) - \rho(h)(0)\|^2.$$

(Here in order to apply Theorem 2.90 we let $X = G$ and let G act on itself by the left multiplication.) Since ψ is unbounded, taking $h = 1 \in G$, we conclude that the action of G on \mathcal{H} has unbounded orbits. This contradicts Property FH. Conversely, given a representation $\rho : G \rightarrow \text{Isom}(\mathcal{H})$ of G to the isometry group of a Hilbert space, we define the kernel ψ on G by the formula (19.11). This kernel is clearly G -invariant; it is also conditionally negative semidefinite according to Theorem 2.90. Therefore, this kernel has to be bounded. Taking $h = 1$, we conclude that the action $G \curvearrowright \mathcal{H}$ has bounded orbits and, therefore, has a fixed point (see Corollary 19.31). \square

REMARK 19.49. Yves de Cornulier observed in [dC06] that there are uncountable discrete groups with Property FH that do not satisfy Property (T).

A theorem similar to Theorem 19.47 establishes the equivalence of the a-T-menability and the Haagerup Property.

THEOREM 19.50. *Let G be a discrete group. The following are equivalent:*

- (1) G is a-T-menable.
- (2) G has the Haagerup Property.
- (3) There exists a conditionally negative semidefinite G -invariant proper kernel $\psi : G \times G \rightarrow \mathbb{R}$.

PROOF. The proof of the equivalence (1) \Leftrightarrow (2) is similar to the proof of the equivalence between (T) and FH in Theorem 19.47. In what follows we sketch the proof of the equivalence (1) \Leftrightarrow (3). We refer to [CCJ⁺01] and [AW81] for the details.

If ψ is a conditionally negative semidefinite G -invariant kernel on G , then, according to Theorem 2.90, there exists a Hilbert space \mathcal{H} and a representation

$$\rho : G \rightarrow \text{Isom}(\mathcal{H})$$

such that

$$(19.12) \quad \psi(g, h) = \|\rho(g)(0) - \rho(h)(0)\|^2.$$

Assuming that ψ is proper and taking $h = 1 \in G$, we see that the sublevel sets in G of the function

$$g \mapsto \|\rho(g)(0) - 0\|$$

are relatively compact in G . Hence, the action ρ is metrically proper, cf. Exercise 5.41. Conversely, given a representation $\rho : G \rightarrow \text{Isom}(\mathcal{H})$ of G to the (affine) isometry group of a Hilbert space, we define the kernel ψ on G by the formula (19.12). This kernel is clearly G -invariant; it is also conditionally negative semidefinite according to Theorem 2.90. Properness of the kernel ψ is clear in view of Exercise 5.41. \square

A consequence of Theorem 19.50 and Corollary 19.43 is the following.

COROLLARY 19.51. *Every discrete amenable group is a - T -menable.*

Another immediate consequence either of the characterizations with kernels, or of Corollary 19.41 and of the equivalences in Theorems 19.47 and 19.50, is

COROLLARY 19.52. *Let G be a countable group. The following properties are equivalent:*

- (1) G has both Property FH and a - T -menability;
- (2) G is finite.

Theorems 19.47 and 19.50 also allow to prove a characterization of the Haagerup property in terms of isometric actions on median spaces (see Section 19.6).

Further properties of groups with Property (T).

Recall that Property (T) is inherited by quotient groups. Since a (discrete) amenable group has Property (T) if and only if it is finite, it follows that every amenable quotient of a group with Property (T) has to be finite. In particular, every discrete group with Property (T) has finite abelianization. For instance, free groups and surface groups never have Property (T). On the other hand, unlike Haagerup, Property (T) is not inherited by subgroups. For instance, $SL(3, \mathbb{Z})$ has Property (T) and contains non-trivial free subgroups, hence, contains subgroups which do not satisfy Property (T).

LEMMA 19.53. *Property (T) is a VI-invariant.*

PROOF. 1. Suppose that a group H has Property (T) and G is a group containing H as a finite-index subgroup. Suppose that $G \curvearrowright \mathcal{H}$ is an isometric affine action of G on a Hilbert space. Since H has Property (T), there exists $x \in \mathcal{H}$ fixed

by H . Therefore, the G -orbit of x is finite. Therefore, by Theorem 3.74, G fixes a point in \mathcal{H} as well.

2. Suppose that $H \leq G$ is a finite-index subgroup and G has Property (T). Let $H \curvearrowright \mathcal{H}$ be an isometric affine action. Define the *induced* action Ind_H^G of G on the space V :

$$V = \{\phi : G \rightarrow \mathcal{H} : \phi(gh^{-1}) = h\phi(g), \forall h \in H, g \in G\}.$$

Every such function is, of course, determined by its values on $\{g_1, \dots, g_n\}$, coset representatives for G/H . The group G acts on V by the left multiplication $g : \phi(x) \mapsto \phi(gx)$. Therefore, as a vector space, V is naturally isomorphic to the n -fold sum of \mathcal{H} . We equip V with the inner product

$$\langle \phi, \psi \rangle := \sum_{i=1}^n \langle \phi(g_i), \psi(g_i) \rangle,$$

making it a Hilbert space. We leave it to the reader to verify that the action of G on V is affine and isometric. The initial Hilbert space \mathcal{H} embeds diagonally in V ; this embedding is H -equivariant, linear and isometric. Since G has Property (T), it has a fixed vector $\psi \in V$. Therefore, the orthogonal projection of ψ to the diagonal in V is fixed by H . Hence, H also has Property (T).

3. Consider a short exact sequence

$$1 \rightarrow F \rightarrow G \rightarrow H \rightarrow 1.$$

If G has Property (T), then so does H (as a quotient of G).

Conversely, suppose that H and F both have Property (T) (we will use it in the case where F is a finite group). Consider an affine isometric action $G \curvearrowright \mathcal{H}$ on a Hilbert space. Since F has Property (T), it has non-empty fixed-point set $V \subset \mathcal{H}$. Then V is a closed affine subspace in \mathcal{H} , which implies that V (with the restriction of the metric from \mathcal{H}) is isometric to a Hilbert space. The group G preserves V and the affine isometric action $G \curvearrowright V$ factors through the group H . Since H has Property (T), it has a fixed point $v \in V$. Thus, v is fixed by the entire group G . In particular, every coextension of a group with Property (T) with finite kernel, also has Property (T).

Putting all these facts together, we conclude that Property (T) is invariant under virtual isomorphisms. \square

Moreover (see e.g. [BdlHV08]):

THEOREM 19.54. *Let G be a locally compact Hausdorff group and $\Gamma < G$ a lattice in it, equipped with the discrete topology. Then G has the (topological) Property (T) if and only if Γ does.*

19.5. Groups acting on trees do not have Property (T)

The main result of this section is the following theorem proven independently by R. Alperin and Y. Watatani. Although this result is a special case of a more general theorem about group actions on median spaces proven in the next section, we present a direct proof here, since it is simpler than that of the more general result and illustrates nicely the ideas behind the general proof.

THEOREM 19.55 (R. Alperin and Y. Watatani). *Each group with Property FH satisfies Property FA.*

PROOF. Let G be a group with Property FH, and $G \curvearrowright X$ an isometric action on a real tree (X, dist) . It suffices to prove that the function $\text{dist}(x, y)$ is a conditionally negative semidefinite G -invariant kernel on X , since according to Theorem 19.47 this will imply that G has bounded orbits on X .

The only statement that is not immediate is that dist is conditionally negative semidefinite. Since this statement needs to be verified for each finite subset $Y = \{y_1, \dots, y_n\}$ of X , it suffices to prove it for the finite metric sub-tree $T \subset X$ that is the convex hull of Y in X . We will assume that Y consists of at least two points since the statement is clear otherwise. Being a finite metric tree, T is a finite simply-connected graph equipped with a path-metric. We orient each edge e of this graph in an arbitrary fashion. Let $V \subset T$ denote the vertex set of the tree. For each point $p \in T \setminus V$ we define a function $f_p : T \rightarrow \{0, 1\}$ as follows. The point p separates T in two connected components. Let $e = uv$ be the oriented edge of T containing p . If $x \in T$ is contained in the same connected component of $T \setminus \{p\}$ as the vertex v , we set $f_p(x) = 1$. For all other points $x \in T$ we set $f_p(x) = 0$. We equip the tree T with the measure μ without atoms, whose restriction to each edge e of T is the Lebesgue measure, so that $\mu(e)$ equals the length of e . Define the function

$$\psi(x, y) = \int_T (1 - f_p(x))f_p(y)d\mu(p).$$

This function can be regarded as a *nonsymmetric pseudo-metric* on T : It is non-negative and satisfies the triangle inequality, but in general, $\psi(x, y) \neq \psi(y, x)$. We leave it to the reader to verify that for all points $x, y \in T$,

$$\text{dist}(x, y) = \psi(x, y) + \psi(y, x).$$

We are now ready to verify that dist is conditionally negative semidefinite. Let $Y = \{y_1, \dots, y_n\} \subset T \subset X$ be as above (T is the convex hull of Y in X) and take a vector $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathbb{R}^n$ such that

$$\lambda_1 + \dots + \lambda_n = 0.$$

We have:

$$\begin{aligned} \sum_{i,j=1}^n \lambda_i \lambda_j \psi(y_i, y_j) &= \sum_{i,j=1}^n \lambda_i \lambda_j \int_T (1 - f_p(y_i))f_p(y_j)d\mu(p) = \\ &= \sum_{i,j=1}^n \lambda_i \lambda_j \int_T f_p(y_j)d\mu(p) - \sum_{i,j=1}^n \lambda_i \lambda_j \int_T f_p(y_i)f_p(y_j)d\mu(p) = \\ &= (\lambda_1 + \dots + \lambda_n) \sum_{j=1}^n \lambda_j \int_T f_p(y_j)d\mu(p) - \int_T \left(\sum_{i=1}^n \lambda_i f_p(y_i) \right)^2 d\mu(p) = \\ &= - \int_T \left(\sum_{i=1}^n \lambda_i f_p(y_i) \right)^2 d\mu(p) \leq 0. \end{aligned}$$

□

COROLLARY 19.56. *Each group G which admits a non-trivial amalgamated free product decomposition $G \simeq G_1 \star_{G_3} G_2$ or a non-trivial HNN-decomposition $G \simeq G_1 \star_{G_3}$, does not have Property FH.*

As an application of this corollary, we obtain:

COROLLARY 19.57. *If M is a connected 3-dimensional manifold with infinite fundamental group G , then G does not satisfy Property (T).*

PROOF. First of all, if G is not finitely generated, it cannot satisfy Property (T), see Theorem 19.39. Thus, we will assume that G is finitely generated. According to the Scott Compact Core Theorem (see [Sco73]), there exists a compact submanifold (possibly with boundary) $M_1 \subset M$ such that the inclusion map $M_1 \rightarrow M$ induces an isomorphism of fundamental groups. Thus, the problem reduces to the case of compact 3-dimensional manifolds. Since Property (T) is a virtual isomorphism invariant, we can assume that the manifold M_1 is oriented. Attach 3-dimensional balls to each spherical boundary component of M_1 ; this results in a compact 3-dimensional manifold M_2 such that each boundary component of M_2 has genus ≥ 1 . The Euler characteristic $\chi(M_2)$ of the manifold M_2 equals

$$\frac{1}{2}\chi(\partial M_2) \leq 0.$$

If one of the boundary components of M_2 has genus ≥ 2 , then

$$\chi(M_2) = b_0(M_2) - b_1(M_2) + b_2(M_2) < 0.$$

It follows that $b_1(M_2) \geq 1$, i.e. there exists an epimorphism

$$\pi_1(M_2) \rightarrow \mathbb{Z}.$$

Since \mathbb{Z} does not satisfy Property (T), the group $\pi_1(M_2)$ does not satisfy it either.

Suppose, therefore, that each boundary component of $N = M_2$ is a torus (this includes the case $\partial M_2 = \emptyset$). Since $\pi_1(N)$ is assumed to be infinite, N is not homeomorphic to S^3 (this is the 3-dimensional Poincaré Conjecture, proven by Perelman). We now apply *Thurston's Geometrization Conjecture/Perelman's Theorem* to the manifold N . According to this theorem, N admits a 2-step decomposition as follows. First of all, N splits as a connected sum

$$N = N_1 \# \dots \# N_n, \quad n \geq 1,$$

where each manifold N_i is *prime*, i.e. does not have a non-trivial connected sum decomposition, and is not simply-connected. If $n \geq 2$, then the group $\pi_1(N)$ admits a non-trivial free product decomposition and, hence, cannot have Property (T). Assume, therefore, that the manifold N itself is prime, $n = 1$. Then N admits a splitting along a system of pairwise disjoint π_1 -injective tori T^2 into submanifolds K_1, \dots, K_m with toral boundary (which could be empty if $N = K_1$ and $\partial N = \emptyset$). Each piece K_i of this decomposition is *geometric*. If the secondary decomposition of N is non-trivial, then $\pi_1(N)$ is isomorphic to the fundamental group of a non-trivial graph of groups and, therefore, again, $\pi_1(N)$ admits a non-trivial action on a simplicial tree. We are thus left with the case when N itself is geometric, i.e. admits a geometric structure modeled on one of Thurston's eight 3-dimensional geometries. By looking at these geometries one-by-one, it is clear that either:

1. $\pi_1(N)$ is virtually solvable and, thus amenable (this happens in the case of the geometries \mathbb{E}^3 , *Nil*, *Sol* and $\mathbb{S}^2 \times \mathbb{R}$). The group $\pi_1(N)$ cannot satisfy Property (T) in these cases.

2. $\pi_1(N)$ admits an isometric action on the hyperbolic space \mathbb{H}^3 with unbounded orbits (this happens in the case of the geometries \mathbb{H}^3 , $\mathbb{H}^2 \times \mathbb{R}$ and $\widetilde{SL}(2, \mathbb{R})$). In these cases $\pi_1(N)$ cannot satisfy Property (T) according to Theorem 19.34. \square

REMARK 19.58. Actually, according to the recent solution of the *Virtual First Betti Number Problem* by Ian Agol [Ago13], if M is a compact 3-dimensional manifold with infinite fundamental group, then M has a finite cover $N \rightarrow M$ such that $b_1(N) > 0$. Since Property (T) is a virtual isomorphism invariant, one again concludes that $\pi_1(M)$ cannot satisfy Property (T). However, the proof given above is much more elementary than the solution of the Virtual First Betti Number Problem.

19.6. Property FH, a-T-menability, and group actions on median spaces

The goal of this section is to characterize properties FH and a-T-menability using actions on median spaces and on spaces with measured walls (see sections 6.1 and 6.2 for the definitions). In this setting, it is important to be able to associate to every action by isometries on a Hilbert space an invariant structure of measured walls. It is straightforward to endow each Hilbert space \mathcal{H} with a structure of measured walls inducing the metric defined via the standard norm on \mathcal{H} . Indeed, according to Theorem 2.100, (1), a Hilbert space embeds isometrically into an L^1 -space. By Lemma 6.20, each L^1 -space is median, hence its set of convex walls is endowed with a measure, according to Theorem 6.57. Therefore by Lemma 6.49 the set of convex walls of a Hilbert space is endowed with a measure. On the other hand, it is *a priori* not guaranteed that this structure of a space with measured walls is invariant under affine isometries. Indeed, it is impossible to exclude beforehand the case where an affine isometry sends one quadruple of equidistant points in the Hilbert space to another, so that the two quadruples have different median completions in the L^1 -space containing the Hilbert space, as described in Remark 6.22.

Fortunately, the problem can be solved by constructing the structure of space with measured walls directly on separable Hilbert spaces:

PROPOSITION 19.59. *Every separable real Hilbert space \mathcal{H} has a structure of a space with measured walls invariant with respect to its affine isometries.*

PROOF. Define the set of walls \mathcal{W} in \mathcal{H} to be the collection of closed *cooriented* affine hyperplanes in \mathcal{H} , more precisely, the set of partitions of \mathcal{H} into open/closed half-spaces defined by affine hyperplanes.

Following Proposition 6.45, in order to define a structure of a space with measured walls on $(\mathcal{H}, \mathcal{W})$, it suffices to define a premeasure μ on the ring \mathcal{R} of disjoint unions

$$\bigsqcup_{i=1}^n \mathcal{W}(F_i | G_i), \quad n \in \mathbb{N} \cup \{\infty\},$$

where F_i, G_i are finite non-empty subsets of \mathcal{H} ($i \in \mathbb{N}$), such that for every $x, y \in X$, $\mu(\mathcal{W}(x|y))$ is finite. Moreover, as noted in Remark 6.58, the measure induced by the premeasure μ is unique if for some countable subset

$$\{x_n : n \in \mathbb{N}\} \subset \mathcal{H},$$

we have

$$\mathcal{W} = \bigcup_{m,n \in \mathbb{N}} \mathcal{W}(x_m | x_n).$$

This condition is satisfied in our case since \mathcal{H} is separable.

Given finite subsets $F, G \subset \mathcal{H}$, we consider an n -dimensional ($n < \infty$) affine subspace V of \mathcal{H} containing $F \cup G$. The pull-back by the inclusion map $V \hookrightarrow \mathcal{H}$ of the set of walls of \mathcal{H} consists of cooriented hyperplanes in V . The space of cooriented hyperplanes can be identified with the quotient space

$$\text{Isom}(\mathbb{E}^n)/\text{Isom}(\mathbb{E}^{n-1}).$$

As the Lie group $\text{Isom}(\mathbb{E}^{n-1})$ is unimodular, every Haar measure on $\text{Isom}(\mathbb{E}^n)$ descends to a unique measure on $\text{Isom}(\mathbb{E}^n)/\text{Isom}(\mathbb{E}^{n-1})$; [Bou63, Chapter 7, §2, Proposition 4]. This measure has a unique normalization such that $\mathcal{W}(x \mid y) = \|x - y\|$. The existence of such a normalization follows from a formula of the Haar measure on the group of affine transformations that can be found, for instance, in [Bou63, §3, Example 2]. Thus, on the space V with the set of walls \mathcal{W}_V composed of cooriented hyperplanes, there is a natural measure μ_V with dist_{μ_v} equal to the Euclidean metric.

Given two finite dimensional affine subspaces $V \subset U$ of \mathcal{H} , the uniqueness of the measure μ_V implies that the measured walls structure induced on V by the inclusion in U and by the structure of U coincides with the structure of V . This implies that $\mu_V(\mathcal{W}_V(F \mid G))$ is independent of the choice of the finite dimensional affine subspace V .

Likewise, the uniqueness of the measure μ_V implies that given any affine isometry φ of the Hilbert space \mathcal{H} , the measure $\mu_{\varphi(V)}(\mathcal{W}_{\varphi(V)}(\varphi F \mid \varphi G))$ is the same as the measure $\mu_V(\mathcal{W}_V(F \mid G))$. We define $\mu(\mathcal{W}(F \mid G))$ to be $\mu_V(\mathcal{W}_V(F \mid G))$, and observe that this premeasure is invariant under affine isometries of \mathcal{H} .

As we noted earlier, separability of \mathcal{H} implies that μ has a unique extension to the σ -algebra generated by \mathcal{R} , and therefore that this extension is invariant with respect to the group of affine isometries of \mathcal{H} . □

THEOREM 19.60. *Let G be a finitely generated group. The following are equivalent:*

- (1) G satisfies Property (T);
- (2) every isometric G -action on a submedian space has bounded orbits;
- (3) every G -action by automorphisms on a space with measured walls has bounded orbits (with respect to the measured walls pseudo-metric);
- (4) every affine isometric G -action on a median metric space has bounded orbits;
- (5) every affine isometric G -action on a space $L^1(X, \mu)$ has bounded orbits.

PROOF. (1) \Rightarrow (2): Let G be a group with Property (T) acting by isometries on a submedian space (X, dist) . Theorem 6.51, Part (3), and Proposition 2.98 imply that the left-invariant kernel $\psi : G \times G \rightarrow \mathbb{R}_+$ defined by $\psi(g, h) = \text{dist}(gx, hx)$, for some $x \in X$, is conditionally negative semidefinite. It follows from Theorem 19.47 that ψ is bounded, equivalently that the orbit Gx is bounded.

The implication (2) \Rightarrow (3) is an immediate consequence of Theorem 6.51, Part (1).

The implication (3) \Rightarrow (4) follows from Theorem 6.51, Part (2), while (4) \Rightarrow (5) is a consequence of Lemma 6.20.

(5) \Rightarrow (1). Consider an isometric affine action of the group G on a real Hilbert space \mathcal{H} . Since G is countable, without loss of generality we may assume that \mathcal{H} is separable. By Proposition 19.59, \mathcal{H} has structure of a space with measured walls $(\mathcal{H}, \mathcal{W}, \mu)$ such that pdist_μ is the Hilbert space metric and such that G acts by automorphisms on $(\mathcal{H}, \mathcal{W}, \mu)$.

Let \mathcal{D} denote the set of (open) half-spaces defined by the walls in \mathcal{H} , let $\mu_{\mathcal{D}}$ the measure induced by μ on \mathcal{D} . Fix $x \in \mathcal{H}$ and let σ_x denote the set of half spaces containing x . According to Proposition 6.50, the map $b : G \rightarrow L^1(\mathcal{D}, \mu_{\mathcal{D}})$ defined by $b(g) = \chi_{\sigma_{gx}} - \chi_{\sigma_x}$ is a 1-cocycle with respect to the unitary representation π of G on $L^1(\mathcal{D}, \mu_{\mathcal{D}})$. We obtain an affine isometric action ρ of G on $L^1(\mathcal{D}, \mu_{\mathcal{D}})$ defined by:

$$g \cdot f = \pi(g)f + b(g).$$

Property (5) implies that $b(g)$ is bounded, whence the orbit of $\rho(G)x$ is bounded in \mathcal{H} and, hence, $\rho(G)$ fixes a point in \mathcal{H} . \square

Note that by [BGM12] in Part (5) the statement can be replaced by the fixed point property for actions on L^1 -spaces. This cannot be deduced from Part (5) only, since L^1 -spaces are not convex, and a further argument is needed. The next theorem is an analogue of Theorem 19.60 in the context of a-T-menability:

THEOREM 19.61. *Let G be a finitely generated group. The following are equivalent:*

- (1) G is a-T-menable (has the Haagerup property);
- (2) there exists a metrically proper isometric action of G on a submedian space;
- (3) G acts by automorphisms on a space with measured walls such that the action is metrically proper (with respect to the measured walls pseudo-metric);
- (4) G has a metrically proper isometric action on a median metric space;
- (5) G has a metrically proper affine isometric action on a space $L^1(X, \mu)$.

PROOF. The implications (5) \Rightarrow (4) \Rightarrow (3) \Rightarrow (2) \Rightarrow (1) follow from the same arguments as in the proof of Theorem 19.60.

(1) \Rightarrow (5). Consider a metrically proper affine isometric action of the group G on a Hilbert space \mathcal{H} . As in the proof of Theorem 19.60, we may assume that \mathcal{H} is separable, and apply Proposition 19.59 to deduce that \mathcal{H} has a G -invariant structure of space with measured walls $(\mathcal{H}, \mathcal{W}, \mu)$ such that pdist_μ is the Hilbert metric.

With the same notation as in the proof of the implication (5) \Rightarrow (1) of Theorem 19.60, we define an affine isometric action ρ of G on $L^1(\mathcal{D}, \mu_{\mathcal{D}})$, by:

$$g \cdot f = \pi(g)f + b(g).$$

Since the action of G on \mathcal{H} was metrically proper, the cocycle b is proper and, hence, the isometric action ρ of G on $L^1(\mathcal{D}, \mu_{\mathcal{D}})$ is also metrically proper. \square

19.7. Fixed point property and proper actions for L^p -spaces

In this section we discuss a generalization of Properties FH and a-T-menability in the context of other classical Banach spaces.

CONVENTION 19.62. For the remainder of this section, p is a real number in $(0, +\infty)$

DEFINITION 19.63. Let G be a (discrete) group. We say G has *Property FL^p* if for every measure space (X, μ) , every affine isometric action of G on $L^p(X, \mu)$ has bounded orbits. We say G is *a- FL^p -menable* if G has a metrically proper affine isometric action on some space $L^p(X, \mu)$.

For $p > 1$ each space $\mathcal{B} = L^p(X, \mu)$ is *uniformly convex*, see [LT79, p. 128] and, hence, every nonempty bounded subset $C \subset \mathcal{B}$ is contained in a unique ball $B(p, R)$ of the least diameter, see [BL00, 1.4]. The center p of this ball is, therefore, invariant under the stabilizer of C in $\text{Isom}(\mathcal{B})$, which, therefore, has a fixed point in \mathcal{B} . It follows that if a group G acts by affine isometries on a space $L^p(X, \mu)$ with bounded orbits, then G has a fixed point in $L^p(X, \mu)$, cf. the proof of Theorem 3.74.

For $p = 1$ one cannot use convexity. Still, in that case too, the bounded orbit property for isometric actions on L^1 -spaces implies the fixed point property, *via* a more intricate argument, see [BGM12].

For $p \in (0, 1)$ boundedness of G -orbits no longer implies the existence of a fixed point; in view of this exception, Property FL^p should probably be called BL^p , but we follow the established terminology.

To summarize, for $p \in [1, \infty)$, Property FL^p is equivalent to the property FC in Definition 19.26, where \mathcal{C} is the class of L^p -spaces.

Below is yet another application of limits of actions. This theorem was proved by Y. Shalom [Sha00, Theorem p. 5] in the case $p = 2$ (i.e. Property FH), answering a question of R. Grigorchuk and A. Zuk.

THEOREM 19.64. *Let G be a finitely generated group satisfying Property FL^p , for some $p \geq 1$. Then G can be written as H/N , where H is a finitely presented group with Property FL^p and N is a normal subgroup in H .*

PROOF. Consider an infinite presentation of G , $G = \langle S \mid r_1, \dots, r_n, \dots \rangle$, where S is a finite set generating G and (r_i) is a sequence of relators in S . Let $F(S)$ be the free group in the alphabet S and N_i the normal closure in $F(S)$ of the finite set $\{r_1, \dots, r_i\}$. The groups $G_i = F(S)/N_i$ are all finitely presented, and form a direct system whose direct limit is G . Assume that none of these groups has Property FL^p . It follows that for each i there exists some space $L^p(Y_i, \mu_i)$ and an affine isometric action of G_i on $L^p(Y_i, \mu_i)$ without a fixed point. Theorem 19.22 and Corollary 19.17 imply that G acts by affine isometries and without a global fixed point on some space $L^p(Z, \nu)$, contradicting the hypothesis. \square

We now compare Properties FL^p for various values of p . To begin with, whatever the choice of p in $(0, \infty)$, Property FL^p always implies Property FH, while a-T-menability always implies a- FL^p -menability. This is a consequence of the characterization of both FH and respectively a-T-menability using actions on spaces with measured walls.

THEOREM 19.65. *Let G be a (discrete) group and let $p > 0$ be an arbitrary fixed positive number.*

- (1) *If G has Property FL^p then G has Property FH .*
- (2) *If G is a - T -menable then G is a - FL^p -menable.*

PROOF. (1) According to Theorem 19.60 it suffices to consider an arbitrary action by automorphisms of G on a space with measured walls $(X, \mathcal{W}, \mathcal{B}, \mu)$, and to prove that it has bounded orbits. Proposition 6.50 implies that this defines an action by affine isometries on $L^p(\mathcal{D}, \mu_{\mathcal{D}})$, where \mathcal{D} is the space of half-spaces and $\mu_{\mathcal{D}}$ the measure induced by it. The group G has Property FL^p , therefore the action has bounded orbits, equivalently the cocycle $b : G \rightarrow L^p(\mathcal{D}, \mu_{\mathcal{D}})$, $b(g) = \chi_{\sigma_{gx}} - \chi_{\sigma_x}$, is bounded, where x is an arbitrary point in X , and σ_x is the set of half-spaces containing x . The latter implies that the orbit of x is bounded with respect to the wall metric.

(2) A - T -menability implies, by Theorem 19.61, that G has a metrically proper action by automorphisms on a space with measured walls $(X, \mathcal{W}, \mathcal{B}, \mu)$. Hence, the action by affine isometries on $L^p(\mathcal{D}, \mu_{\mathcal{D}})$ described in Proposition 6.50 is also metrically proper. \square

The converse implications in Theorem 19.65 hold only for certain values of $p > 0$.

THEOREM 19.66 ([Del77], [AW81], [WW75]). *Let G be a discrete group.*

- (1) *If G satisfies Property FH , then it also satisfies Property FL^p for every $p \in (0, 2]$.*
- (2) *If G is a - FL^p -menable for some (for every) $p \in (0, 2]$ then G is a - T -menable.*

PROOF. (1) follows from Proposition 2.98, and from the implication (1) \Rightarrow (3) in Theorem 19.47.

(2) follows from the same Proposition and the implication (3) \Rightarrow (1) in Theorem 19.50. \square

We thus obtain:

THEOREM 19.67. *For every $p \in (0, 2]$,*

$$FL^p \iff FH \text{ and } a\text{-}FL^p\text{-menability} \iff a\text{-}T\text{-menability}$$

The implication $FH \Rightarrow FL^p$ extends a little bit beyond the interval $[1, 2]$ according to the following theorem:

THEOREM 19.68 (D. Fisher and G. Margulis, see [BFGM07], §3.c). *For every discrete group G with Property FH there exists $\varepsilon = \varepsilon(G)$ such that G has Property FL^p for every $p \in (0, 2 + \varepsilon)$.*

We will prove a slightly stronger form of this result below. For each discrete group G define the subset $\mathcal{FP}_G \subset (0, \infty)$ consisting of those p such that G satisfies Property FL^p .

THEOREM 19.69. *For each finitely generated group G the set \mathcal{FP}_G is open.*

PROOF. In view of Theorem 19.67, it suffices to show that the set of $p \in [2, \infty)$ such that G does not satisfy Property FL^p is closed. Let $p_n \in [2, \infty)$ be a sequence converging to $p < \infty$, such that for every n , G has an isometric action on some space $L^{p_n}(X_n, \mu_n)$ without a fixed point. Theorem 19.22 and Corollary 19.17 imply that, for some set Y and measure ν on Y , the group G also acts isometrically on the space $L^p(Y, \nu)$ without a fixed point. \square

For p much larger than 2, Properties FH and FL^p are no longer equivalent, nor are a-T-menability and a- FL^p -menability equivalent. Below we mention some results illustrating both.

THEOREM 19.70 (P. Pansu, [Pan95] and Y. de Cornulier, R. Tessera and A. Valette [dCTV08]). *Let G denote the isometry group of the quaternionic-hyperbolic space $\mathbf{H}\mathbf{H}^n$. Then every lattice $\Gamma < G$ admits a proper isometric action on some $L^p(X, \mu)$, for every $p > 4n + 2$.*

Thus, each lattice in G has Property FH and is a- FL^p -menable for all $p > 4n + 2$. Furthermore:

THEOREM 19.71 (M. Bourdon, H. Pajot [BP03, Bou16], G. Yu [Yu05]). *Every infinite hyperbolic group is a- FL^p -menable for every p larger than the conformal dimension of the boundary at infinity of G .*

We refer the reader to [Nic13] for an alternative proof of this theorem. Recall that there are many examples of hyperbolic groups with Property (T), besides lattices in the isometry groups of the quaternionic hyperbolic spaces, coming from the theory of random groups (see Theorem 19.75).

On the other hand, lattices in Lie groups of higher rank exhibit a different behavior with respect to Property FL^p :

THEOREM 19.72 (U. Bader, A. Furman, T. Gelander and N. Monod, [BFGM07]). *Let G be a semisimple Lie group with all non-compact factors of rank ≥ 2 . Then every lattice $\Gamma < G$ satisfies Property FL^p for all $p \in (0, \infty)$.*

19.8. Groups satisfying Property (T) and the spectral gap

Recall that Property FH (equivalently, Property (T)) for a discrete countable group G can be reformulated as the vanishing of $H^1(G, \mathcal{H}_\pi) = 0$ for all unitary representations $\pi : G \rightarrow U(\mathcal{H})$. There is an old technique (going back to Bochner) for proving vanishing theorems of this type; namely, if $G \simeq \pi_1(M)$, where M is a closed Riemannian manifold,

$$H^1(G, \mathcal{H}_\pi) \cong H_{DR}^1(M, \mathcal{V}),$$

where the right hand side is the *de Rham cohomology* of M with coefficients in a flat vector bundle over M with fibers isometric to \mathcal{H} . Then one uses the Hodge Theorem to represent the de Rham cohomology classes $\omega \in H_{DR}^1(M, \mathcal{V})$ by *harmonic 1-forms*. Lastly, one uses some geometric properties of M to show that each harmonic form as above has to vanish. Harmonic 1-forms on M (with coefficients in \mathcal{V}) lift to G -invariant forms on the universal cover \tilde{M} of M and can be interpreted as G -equivariant harmonic maps $\tilde{M} \rightarrow \mathcal{H}$, where G acts on \tilde{M} by covering transformations and on \mathcal{H} by an affine isometric action ρ whose linear part is the representation π and the translational part is given by the 1-cocycle $c : G \rightarrow \mathcal{H}$

representing the cohomology class ω . We note that the existence of equivariant harmonic maps with respect to group actions on $CAT(0)$ spaces was established by Korevaar and Schoen in much greater generality than the one of Hilbert space targets; see [KS97].

This line of reasoning was extended by H. Garland [Gar73] and, later, in greater generality, by P. Pansu [Pan96], by A. Żuk [Żuk96, Żuk03] and by W. Ballmann and J. Swiatkowski, [BŚ99], to the setting when one replaces the free properly discontinuous isometric action $G \curvearrowright \tilde{M}$ with a properly discontinuous, (not necessarily free) isometric and cocompact action on a piecewise-Euclidean (or hyperbolic) simplicial complex, $G \curvearrowright X$ (see also [Bou00] and [Bou16]). We now describe the combinatorial conditions on X (replacing the geometric conditions on M in the classical setting of Bochner technique), which lead to the vanishing of H^1 and, hence, to combinatorial examples of groups satisfying Property (T).

We begin with a combinatorial replacement of the Bochner technique; our discussion follows the papers by A. Żuk, [Żuk03], and by M. Bourdon, [Bou16].

Let \mathcal{G} be a graph, possibly with loops and multiple edges (a multigraph). Let V be its set of vertices and E its set of edges. For every edge $e \in E$, we denote by $\mathcal{V}(e)$ its set of endpoints, enumerated with multiplicity.

The *graph Laplacian* on \mathcal{G} is an operator from $\mathbb{R}^{|V|}$ to $\mathbb{R}^{|V|}$ defined by

$$(\Delta f)(u) = \frac{1}{\text{val}(u)} \sum_{e \in E, \mathcal{V}(e)=\{u,v\}} [f(u) - f(v)] \text{ for every } u \in V,$$

where $\text{val}(u)$ is the valency of the vertex u . The operator Δ is linear, therefore one can define its eigenvalues and eigenfunctions in the usual manner.

We denote by $\lambda_1(\mathcal{G})$ the smallest eigenvalue of Δ which corresponds to a non-constant eigenfunction, sometimes also called the second eigenvalue of the Laplacian. Bounding $\lambda_1(\mathcal{G})$ away from zero corresponds to measuring the expansion in the graph \mathcal{G} . For instance, suppose that \mathcal{G} is the incidence graph of the finite projective plane $\mathbb{F}_q\mathbb{P}^2$, where \mathbb{F}_q is the field of order q . Then

$$\lambda_1(\mathcal{G}) = 1 - \frac{\sqrt{q}}{q+1}.$$

A sufficient condition for Property FH, described in the theorem below, can then be obtained, using Garland's method of harmonic maps, [Pan96, BŚ97a, Bou16], or by a direct geometric argument, [Żuk03].

THEOREM 19.73 ([Żuk03], [Bou16]). *Let X be a simplicial 2-complex where the link $L(x)$ of every vertex x has $\lambda_1(L(x)) > \frac{1}{2}$. If some group G acts on X simplicially, properly, and cocompactly, then G has Property FH.*

Theorem 19.73 is a criterion for Property (T) that can be applied to every finitely presented group. Indeed, every finitely presented group has a finite triangular presentation, i.e. a presentation with all relators of length three. If $G = \langle S|R \rangle$ is a triangular finite presentation, then G acts on a simplicial 2-complex X which is the Cayley complex. The link of a vertex in X is the graph $L(S)$ with the vertex set $S \cup S^{-1}$ and, for each relator of the form $s_x s_y s_z$ in R , edges (s_x^{-1}, s_y) , (s_y^{-1}, s_z) , and (s_z^{-1}, s_x) .

In this particular case Theorem 19.73 has the following effective version.

THEOREM 19.74 (A. Żuk, [Żuk03]). *If $L(S)$ is connected and $\lambda = \lambda_1(L) > \frac{1}{2}$ then $G = \langle S|R \rangle$ has Property (T). Moreover,*

$$\frac{2}{\sqrt{3}} \left(2 - \frac{1}{\lambda} \right)$$

is a Kazhdan constant of G with respect to S .

The spectral gap condition in Theorems 19.73 and 19.74 allows one to prove that, in some models, “generic” finitely presented groups satisfy Property (T). This can be done, roughly, by combining Theorem 19.74 with a theorem of Friedman stating that random graphs typically have second eigenvalues of the Laplacian larger than $\frac{1}{2}$ [Fri91, Theorem B].

The following description of random groups can thus be formulated, due to the work of Gromov [Gro93, Section 9.B]; Żuk [Żuk03]; Olliver ([Oll05], [Oll04, I.3.b]); Dahmani, Guirardel, and Przytycki [DGP11], Kotowski and Kotowski [KK13], and Antoniuk, Łuczak and Świątkowski [ALS15].

THEOREM 19.75. *Both in the triangular model and in the Gromov density model, a random finitely presented group G at density $\frac{1}{3} < d < \frac{1}{2}$ is, with overwhelming probability, infinite, word-hyperbolic, with aspherical presentation complex, has Euler characteristic $\gg 1$, has boundary homeomorphic to the Menger curve and satisfies Property (T).*

In particular, such a group has $H^2(G, \mathbb{Q}) \neq 0$, as

$$1 + b_2(G) = b_0(G) - b_1(G) + b_2(G) = \chi(G) \gg 1.$$

While “generic” finitely presented groups are infinite and satisfy Property (T), finding explicit and reasonably short presentations presents a bit of a challenge. Explicit presentations of some infinite $CAT(0)$ groups satisfying Property (T) were computed by J. Świątkowski [Świ01b, Świ01a], M. Ershov [Ers08] and J. Essert [Ess13]; the following two examples are taken from [Ess13]:

$$G_1 = \langle s, t, x \mid s^7 = t^7 = x^7 = 1, st = x, s^3t^3 = x^3 \rangle,$$

$$G_2 = \langle s, t, x \mid s^7 = t^7 = x^7 = 1, st = x^3, s^3t^3 = x \rangle.$$

These two groups act geometrically on some \tilde{A}_2 -Euclidean buildings.

Explicit presentations of infinite hyperbolic groups satisfying Property (T) can be extracted from M. Bourdon’s paper [Bou00]; Bourdon’s groups act geometrically on 2-dimensional hyperbolic buildings. See the recent preprint by P.-E. Caprace, [Cap17], for an explicit presentation of an infinite hyperbolic group satisfying Property (T), with four generators and sixteen relators of lengths ranging between 2 and 138.

19.9. Failure of quasiisometric invariance of Property (T)

THEOREM 19.76. *Property (T) is not QI invariant for finitely generated groups.*

PROOF. This theorem should be probably attributed to S. Gersten and M. Raghunathan; the example below is a variation on the Raghunathan’s example discussed in [Ger92].

Let Γ be a hyperbolic group which satisfies Property (T) and such that $H^2(\Gamma, \mathbb{Q}) \neq 0$, see the previous section. Next, pick an infinite order element $\omega \in H^2(\Gamma, \mathbb{Z})$ and consider the central coextension

$$1 \rightarrow \mathbb{Z} \rightarrow \tilde{\Gamma} \rightarrow \Gamma \rightarrow 1$$

with the extension class ω , see Section 5.9.6. Since the group Γ is hyperbolic, Theorem 11.159 implies that the groups $\tilde{\Gamma}$ and $\mathbb{Z} \times \Gamma$ are quasiisometric (see also [Ger92] for a more general version of this argument in the case of central coextensions defined by *bounded cohomology classes*). The group $\mathbb{Z} \times \Gamma$ does not satisfy Property (T), since it surjects to \mathbb{Z} . On the other hand, the group $\tilde{\Gamma}$ satisfies Property (T), see [dlHV89, 2.c, Theorem 12]. \square

According to a very recent result by M. Carette [Car14]:

THEOREM 19.77. *The Haagerup property is not QI invariant for finitely generated groups.*

QUESTION 19.78. Are Properties (T) and Haagerup QI invariant within the class of hyperbolic groups?

19.10. Summary of examples

Below we list some examples and non-examples of groups with Property (T):

Groups with Property (T)	Groups without Property (T)
All Lie groups with simple factors of rank ≥ 2	$O(n, 1)$ and $U(n, 1)$
Lattices in simple Lie groups of rank ≥ 2	Unbounded subgroups of $O(n, 1)$ and $U(n, 1)$
$SL(n, \mathbb{Z}), n \geq 3$	$SL(2, \mathbb{Z})$
Lattices in the isometry group of $\mathbf{H}\mathbb{H}^n, n \geq 2$	Lattices in $O(n, 1)$ and $U(n, 1)$
$SL(n, \mathbb{Z}[t]), n \geq 3$, [Sha06]	Thompson group
	All finitely generated infinite Coxeter groups
	Infinite 3-manifold groups
Some hyperbolic groups	Some hyperbolic groups
	Groups which admit non-trivial splittings as amalgams
	Infinite amenable groups
	Infinite fundamental groups of closed conformally-flat manifolds

REMARK 19.79. (1) A theorem that infinite mapping class groups do not satisfy Property (T) appears in the preprint [And07] by J. E. Andersen. We note that Andersen's preprint is still unpublished.

- (2) Each infinite finitely generated Coxeter group is a-T-menable [BJS88]. Even though the theorem in [BJS88] states only that such groups do not have property (T), what is actually proven there is a-T-menability.
- (3) If M is a compact connected 3-manifold then $\pi_1(M)$ satisfies Property (T) if and only if $\pi_1(M)$ is finite, see Corollary 19.57.
- (4) If M is an n -dimensional connected conformally-flat manifold then there exists a homomorphism $\rho : \pi_1(M) \rightarrow PO(n, 1)$. If the image of ρ is relatively compact and M is closed then $\pi_1(M)$ is finite, see [Kui50]. If

the image of ρ is not relatively compact then $\pi_1(M)$ has the Haagerup property, since $PO(n, 1)$ does.

Property (T) is unclear for the following groups:

- $Out(F_n)$, $n \geq 4$.
- Infinite Burnside groups.
- Shephard groups. (Property (T) fails at least for some of these groups.)
- Generalized von Dyck groups. (Property (T) fails at least for some of these groups, e.g. for infinite von Dyck groups.)
- Hyperbolic Kähler groups. (Property (T) fails at least for some of these groups, e.g. for surface groups and for the fundamental groups of compact complex-hyperbolic manifolds.)

We conclude with a diagram illustrating relationship between different classes of groups discussed in the book:




FIGURE 19.1. The world of infinite finitely generated groups.

Stallings Theorem and accessibility

The goal of this chapter is to prove Stallings Theorem (Theorem 9.24) on ends of groups in the class of (almost) finitely presented groups and Dunwoody's Accessibility Theorem for finitely presented groups. As a corollary we obtain QI rigidity of the class of virtually free groups. Our proofs are a geometric combination of arguments due to Dunwoody [Dun85], Swarup [Swa93] and Jaco and Rubinstein [JR88], which are inspired by the theory of minimal surfaces. One advantage of this approach is that in the process we fill in some of the details the theory of *PL minimal surfaces* developed by Jaco and Rubinstein. The definition of *almost finitely presented groups* (abbreviated as *afp groups*) will be given in Definition 20.27, for now it suffices to note that the class of *afp groups* contains all finitely presented groups.

THEOREM 20.1 (The Stallings ends of groups theorem for *afp groups*). *Let G be an *afp group* with at least 2 ends. Then G splits as the fundamental group of a finite graph of finitely generated groups with finite edge-groups.*

The Stallings theorem allows one to start the decomposition process (using graphs of groups with finite edge groups) of groups with at least two ends. A group is called *accessible* if any such decomposition process terminates after finitely many steps:

THEOREM 20.2 (Dunwoody accessibility theorem). *Every *afp group* is accessible.*

As a combination of these two fundamental theorems one obtains:

COROLLARY 20.3. *Suppose that G is an *afp group* with at least 2 ends. Then G splits as the fundamental group of a finite graph of finitely generated groups with finite edge-groups, such that each vertex group is either finite or 1-ended.*

The Stallings theorem, unlike the one by Dunwoody, holds for all finitely generated groups. In the next chapter we prove the Stallings theorem for finitely generated groups using harmonic functions following an idea proposed by Gromov.

20.1. Maps to trees and hyperbolic metrics on 2-dimensional simplicial complexes

Collapsing maps. Let Δ be a 2-dimensional simplex with the vertices x_i , $i = 1, 2, 3$. Our goal is to define a class of maps $\Delta \rightarrow Y$, where Y is a simplicial tree with the standard metric (the same could be done when targets are arbitrary real trees but we will not need it). The construction of f is, as usual, by induction on skeleta. This construction is analogous to the construction of collapsing maps κ

in Section 11.8. (The difference with the maps κ is that the maps f will not be isometric on edges, only linear.) Let $f : \Delta^{(0)} = \{x_1, x_2, x_3\} \rightarrow Y$ be given. If the image of this map is contained in a geodesic segment α in Y , then we extend f to be a linear map $f : \Delta \rightarrow \alpha$. Otherwise, the points of $f(\Delta^{(0)})$ span a tripod T in Y with the center o and extreme vertices $y_i := f(x_i)$. We extend f to a map $f : \Delta^{(1)} \rightarrow Y$ by sending edges $[x_i, x_{i+1}]$ of Δ to the geodesics $y_i y_{i+1} \subset T$ by linear maps. The preimage $f^{-1}(o)$ consists of three interior points x_{ij} of the edges $[x_i, x_j]$ of Δ , called *center points* of Δ (with respect to f). The 1-dimensional triangle $T(x_{12}, x_{23}, x_{31})$ (called *middle triangle*) splits Δ in four solid sub-triangles $\blacktriangle_i, i = 0, 1, 2, 3$ (\blacktriangle_0 is spanned by the center points while each \blacktriangle_i contains x_i as a vertex). Then f sends the vertices of each \blacktriangle_i to points in one of the legs of T . We then extend f to a linear map on each of these four sub-triangles; clearly, $f(\blacktriangle_0) = \{o\}$.

DEFINITION 20.4. The resulting map $f : \Delta \rightarrow T \subset Y$ is called a *canonical collapsing map*.

It is clear that if X is a simplicial complex and $f : X^{(0)} \rightarrow Y$ is a map, then f admits a unique extension to $f : X \rightarrow Y$ which is linear on every edge of X and is a canonical collapsing map on each 2-simplex. We refer to the map $f : X \rightarrow Y$ as a *canonical map* $X \rightarrow Y$ (it depends, of course, on the initial map $f : X^{(0)} \rightarrow Y$). Suppose that G is a group acting simplicially on X and isometrically on Y . By uniqueness of the extension of f from $X^{(0)}$ to X , if $f : X^{(0)} \rightarrow Y$ is a G -equivariant map, then its extension $f : X \rightarrow Y$ is also G -equivariant. Such an equivariant map $f : X \rightarrow Y$ is called a *canonical resolution* of the G -tree Y .

Existence of resolutions of simplicial G -trees. Recall that every finite group acting isometrically on a real tree T has a fixed point (Corollary 3.75 and Exercise 3.76). If T is a simplicial tree with the standard metric and the action is without inversions, then G has to fix a vertex of T (since a fixed point in the interior of an edge implies that the edge is fixed pointwise).

Let T be a simplicial tree and $G \curvearrowright T$ be a cocompact simplicial action (without inversions). Let X be a connected simplicial 2-dimensional complex on which G acts properly discontinuously and cocompactly (possibly non-freely). We construct a resolution $f : X \rightarrow T$ as follows. Let $v \in X^{(0)}$ be a vertex. This vertex has finite stabilizer G_v in G , therefore, this stabilizer fixes a vertex w in T . We then set $f(v) := w$. (If the fixed vertex is not unique then we choose it arbitrarily.) We then extend this map to the orbit $G \cdot v$ by equivariance. Repeating this for each vertex-orbit we obtain an equivariant map $f : X^{(0)} \rightarrow T^{(0)}$. Note that without loss of generality, by subdividing X barycentrically if necessary, we may assume that $f : X^{(0)} \rightarrow T^{(0)}$ is onto (all that we need for this is that X/G has more vertices than T/G). We then extend f to the rest of X by the canonical collapsing map, therefore obtaining the resolution.

Piecewise-canonical maps. In the proof of Theorem 20.40 we will need a mild generalization of the canonical maps and resolutions. Suppose that in the 2-simplex Δ we are given a subdivision into the solid triangles $\blacktriangle_i, i = 0, \dots, 3$ with vertices x_i, x_{jk} . Suppose we are also given a structure of a polygonal cell complex P on Δ such that:

- (1) Every vertex belongs to the boundary of Δ .
- (2) Every edge is geodesic.
- (3) Every geodesic segment $x_{ij}x_{jk}$ is an edge.

- (4) Every vertex has valence 3 except for $x_{jk}, x_i, i, j, k = 1, 2, 3$.




FIGURE 20.1. Polygonal subdivision of a simplex.

Edges of P not contained in the boundary of Δ are called *interior edges*.

DEFINITION 20.5. A map $f : \Delta \rightarrow Y$ is called *piecewise-canonical* (PC) if it is constant on every interior edge and linear on each 2-cell. Note that the map f could be constant on some 2-faces of P (for instance, it is always constant on the solid middle triangle).

Clearly, a map f of the 1-skeleton of P which is constant on every interior edge, admits a unique PC extension to Δ . A map $f : X \rightarrow Y$ from a simplicial complex to a tree is *piecewise-canonical* (PC) if it is PC on every 2-simplex of X and piecewise-linear on every edge not contained in a 2-face. Every canonical map $f : \Delta \rightarrow Y$ is also PC: The vertices of P are the points x_i, x_{jk} .

Let X be a simplicial complex, Y a simplicial tree and $f : X \rightarrow Y$ a PC map. We say that a point $y \in Y$ is a *regular value* of f if for every 2-simplex Δ in X we have:

- a. $f^{-1}(y)$ is disjoint from the vertex set of Δ .
- b. $f^{-1}(y) \cap \Delta$ is either empty or is a single topological arc (which necessarily connects distinct edges of Δ).

A point $y \in Y$ which is not a regular value of f is called a *critical value* of f . The following is an analogue of Sard's Theorem in the context of PC maps.

LEMMA 20.6. *Let X be a countable simplicial complex, Y a simplicial tree and $f : X \rightarrow Y$ a PC map. Then almost every point $y \in Y$ is a regular point of f .*

PROOF. Let $\Delta \subset X$ be a 2-simplex and P its polygonal cell complex structure. Then there are only finitely many critical values of f , namely the images of the vertices of Δ and of all the 2-faces of P where f is constant. Since X is countable, this means that the set of critical values of f is at most countable. \square

Complete hyperbolic metrics on punctured 2-dimensional simplicial complexes. Our next goal is to introduce a path metric on $X' := X \setminus X^{(0)}$, such that each 2-simplex (minus vertices) is isometric to a solid ideal hyperbolic triangle.

PROPOSITION 20.7. *Let X be a locally finite 2-dimensional simplicial complex. Then there exists a proper path-metric on $X' := X \setminus X^{(0)}$ such that each 2-simplex in X is isometric to the ideal hyperbolic triangle. Moreover, this metric is invariant under all automorphisms of X .*

PROOF. We identify each 2-simplex s in X with the solid ideal hyperbolic triangle \blacktriangle (so that vertices of s correspond to the ideal vertices of the hyperbolic triangle). We now would like to glue edges of the solid triangles isometrically according to the combinatorics of the complex X . However, this identification is not unique since each complete geodesic in \mathbb{H}^2 is invariant under a group of translations. Moreover, some of the identifications will yield incomplete hyperbolic metrics. (Even if we glue two ideal triangles along their boundaries!) Therefore, we have to choose gluing isometries appropriately.

The ideal triangle \blacktriangle admits a unique inscribed circle; the points of tangency of this circle and the sides τ_k of \blacktriangle are the *central points* $\xi_{ij} \in \tau_k, k \notin \{i, j\}$, see Section 11.8.

Now, given two solid ideal triangles $\blacktriangle_i, i = 1, 2$ and oriented sides $\tau_i, i = 1, 2$ of these triangles, there is a unique isometry $\tau_1 \rightarrow \tau_2$ which sends center-point to center-point and preserves orientation. We use these gluings to obtain a path-metric on X' . Clearly, this metric is invariant under all automorphisms of X in the following sense:

If $g \in \text{Aut}(X)$ then the restriction of g to $X^{(0)}$ admits a unique extension $\widehat{g} : X \rightarrow X$ which is an isometry of X' .

We claim that X' is proper. The proof relies upon a certain collection of functions b_ξ on X' defined below, $\xi \in X^{(0)}$.

We first define three functions b_1, b_2, b_3 on the ideal triangle \blacktriangle . Let $\xi_i, i = 1, 2, 3$, be the ideal vertices of \blacktriangle , τ_k the ideal edge connecting ξ_i to ξ_j ; $\xi_{ij} \in \tau_k$ be the central point ($k = i + j \pmod{3}$).

Each pair of central points ξ_{ij}, ξ_{jk} belongs to a unique horocycle H_j in \mathbb{H}^2 with the ideal center ξ_j . One can see this using the upper half-plane model of \mathbb{H}^2 so that $\xi_j = \infty$. Then H_j is the horizontal line passing through the points ξ_{ij} and ξ_{jk} .

Consider the circular arcs $\alpha_i := H_i \cap \blacktriangle$. The arcs $\alpha_1, \alpha_2, \alpha_3$ cut out a solid triangle \blacktriangledown (with horocyclic arcs α_i 's as its edges) from \blacktriangle . We refer to the complementary components C_i of $\blacktriangle \setminus \blacktriangledown$ as *corners* of \blacktriangle with the ideal vertices $\xi_i, i = 1, 2, 3$. Their closures in Δ are the *closed corners* \overline{C}_i . We then define a 1-Lipschitz function $b_i : \overline{C}_i \rightarrow \mathbb{R}_+$ by

$$b_i(x) = \text{dist}(x, \alpha_i), \quad i = 1, 2, 3.$$

The level sets of $b_i : C_i \rightarrow \mathbb{R}_+$ are arcs of horocycles in C_i . (The functions b_i are the negatives of *Busemann functions*, see [Bal95].) We extend each b_i by zero to $\blacktriangle \setminus C_i$.

For each vertex ξ of X we define \overline{C}_ξ to be the union of closed corners \overline{C}_i (with the vertex $\xi = \xi_i$) of 2-simplices $s \subset X$ which have ξ as a vertex. Then the functions $b_i : \overline{C}_i \rightarrow \mathbb{R}_+$ match on intersections of their domains (since the central points do), thus, we obtain a collection of 1-Lipschitz functions $b_\xi : X' \rightarrow \mathbb{R}_+$. It is clear from the construction that each b_ξ is proper on \overline{C}_ξ .




FIGURE 20.2. Geometry of an ideal hyperbolic triangle.

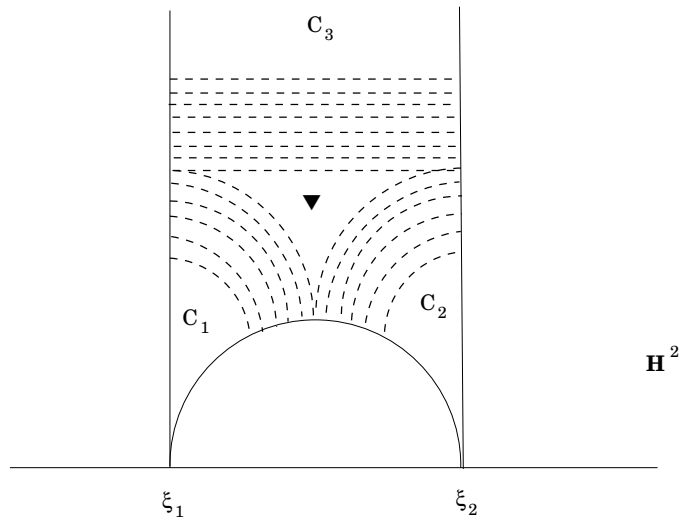


FIGURE 20.3. Partial foliation of corners of an ideal triangle by the level sets of the functions b_i .

Set

$$\blacktriangle_r := \{x \in \blacktriangle : \forall i \quad b_i(x) \leq r\}, \quad X_r := \{x \in X' : \forall \xi \in X^{(0)}, b_\xi(x) \leq r\}.$$

Since each b_ξ is 1-Lipschitz, for every path \mathbf{p} in X' of length $\leq r$, if $p = \text{Image}(\mathbf{p}) \cap X_0 \neq \emptyset$ then $p \subset X_r$.

For $x, y \in X'$ we let $\rho(x, y)$ be the minimal number of edges that a path in X' from x to y has to intersect. Since X is locally finite, for every $x \in X$, $k \in \mathbb{N}$, the set $\{y \in X' : \rho(x, y) \leq k\}$ is a union of finitely many cells.

Every ideal side τ_i of \blacktriangle intersects \blacktriangle_r in a compact subset. Thus, there exists $D(r) > 0$ such that the minimal distance between the geodesics

$$\tau_i \cap \blacktriangle_r, \tau_j \cap \blacktriangle_r \quad (i \neq j)$$

is at least $D(r)$. Therefore, if \mathfrak{p} is a path in X_r connecting x to y , then its length is at least

$$D(r)\rho(x, y).$$

Thus, for every $x \in X_0$ the metric ball $B(x, r) \subset X'$ intersects only finitely many cells in X and is contained in X_r . Since intersection of X_r with any finite subcomplex in X is compact, it is now immediate that X' is a proper metric space. \square

20.2. Transversal graphs and Dunwoody tracks

We continue with the notation of the previous section.

Our goal is to introduce for X' notions analogous to transversality in the theory of smooth manifolds. We define the *vertex-complexity* of a finite graph Γ , denoted $\nu(\Gamma)$, to be the cardinality of the vertex set $V(\Gamma)$. We say that a properly embedded graph $\Gamma \subset X'$ is *transversal* if the following hold:

1. $\Gamma \cap X^{(1)} = V(\Gamma) = \Gamma^{(0)}$.
2. For every edge $e \subset X^{(1)}$, for every 2-face s of X containing e , for every vertex $v \in \Gamma \cap e$, there is exactly one edge γ of Γ in s which has v as its vertex.

Transversal graphs generalize the concept of properly embedded *smooth codimension 1 submanifolds* in a smooth manifold.




FIGURE 20.4. Dunwoody track.

If a transversal graph Γ satisfies the property:

3. For every edge γ of Γ , the end-points of γ belong to distinct edges of $X^{(1)}$, then Γ is called a *Dunwoody track*, or simply a *track*.

EXERCISE 20.8. Let $f : X \rightarrow Y$ be a PC map from a simplicial 2-complex X to a simplicial tree Y and $y \in Y$ a regular value of f . Then $f^{-1}(y)$ is a Dunwoody track in X .

The following lemma is left as an exercise to the reader, it shows that every Dunwoody track in X behaves like a codimension one smooth submanifold in a differentiable manifold.

LEMMA 20.9. *Let Γ be a Dunwoody track. Then for every $x \in \Gamma$ there exists a neighborhood U of x which is naturally homeomorphic to the product*

$$\Gamma_U \times [-1, 1],$$

where $\Gamma_U = \Gamma \cap U$ and the above homeomorphism sends Γ_U to $\Gamma_U \times \{0\}$. We will refer to the neighborhoods U as product neighborhoods of points of Γ .

Note that the entire track need not have a product neighborhood. For instance, let Γ be a non-separating loop in the Moebius band X . Triangulate X so that Γ is a track. Then every regular neighborhood of Γ in X is again a Moebius band. However, the neighborhoods Γ_U combine in a neighborhood N_Γ of Γ in X which is an interval bundle over Γ , where the product neighborhoods $U \cong \Gamma_U \times [-1, 1]$ above serve as coordinate neighborhoods in the fibration.

We say that the track Γ is *1-sided* if the interval bundle $N_\Gamma \rightarrow \Gamma$ is non-trivial and *2-sided* otherwise.

EXERCISE 20.10. Suppose that Γ is connected and 1-sided. Then $N_\Gamma \setminus \Gamma$ is connected.

For each transversal graph $\Gamma \subset X$ we define the *counting function* $m_\Gamma : \text{Edges}(X) \rightarrow \mathbb{Z}$:

$$m_\Gamma(e) := |\Gamma \cap e|.$$

The \mathbb{Z}_2 -cocycle of a transversal graph. Recall that, by the Poincaré duality, every proper codimension k embedding of smooth manifolds

$$N \hookrightarrow M$$

defines an element $[N]$ of $H^k(M, \mathbb{Z}_2)$. Our goal is to introduce a similar concept for transversal graphs Γ in X . Observe that for every 2-face s in X

$$\sum_{i=1}^3 m_\Gamma(e_i) = 0, \quad \text{mod } 2,$$

where e_1, e_2, e_3 are the edges of s (since every edge γ of $\Gamma \cap s$ contributes zero to this sum. Therefore, m_Γ determines an element of $Z^1(X, \mathbb{Z}_2)$. If Γ is finite, then $m_\Gamma \in Z_c^1(X, \mathbb{Z}_2)$ since the cocycle m_Γ is supported only on the finitely many edges which cross Γ . We let $[\Gamma]$ denote the cohomology class in $H_c^1(X, \mathbb{Z}_2)$ determined by m_Γ . It is clear that $[\Gamma]$ depends only on the isotopy class of Γ .

LEMMA 20.11. *Suppose that Γ is 1-sided. Then m_Γ represents a non-trivial class in $H^1(X, \mathbb{Z}_2)$.*

PROOF. We first subdivide X so that N_Γ is a subcomplex in X and will compute $H^1(X, \mathbb{Z}_2)$ using the new cell complex denoted X' . We will use the notation m for the cocycle on Y defined by Γ . Since Γ is 1-sided, there are vertices u, v of X' which belong to $\partial N_\Gamma \cap X^{(1)}$, such that:

1. The edge $\tau = [u, v]$ of Y connecting u and v is contained in an edge e of X .
2. The edge τ intersects Γ in exactly one point.




FIGURE 20.5. Nontriviality of a cocycle.

Hence,

$$(20.1) \quad m(u) + m(v) = 1 \pmod{2}$$

Since $N_\Gamma \setminus \Gamma$ is connected, there exists a path $\alpha \subset \partial N_\Gamma$ connecting u to v . Suppose that $m = \delta\eta$, where $\eta \in C^0(Y, \mathbb{Z}_2)$. In other words, $\eta : (Y)^{(0)} \rightarrow \mathbb{Z}_2$ and for every pair of vertices p, q of Y connected by the edge $[p, q]$, we have:

$$\eta(p) - \eta(q) = m([p, q]).$$

In particular, if p, q are connected by an edge-path in Y which is disjoint from Γ , then $\eta(p) = \eta(q)$. Since the path α connecting u to v is disjoint from Γ , we obtain

$$\eta(u) = \eta(v).$$

On the other hand, in view of (20.1), we also have

$$\eta(u) + \eta(v) = 1$$

Contradiction. □

LEMMA 20.12. *Suppose that $H^1(X, \mathbb{Z}_2) = 0$. Then a connected finite transversal graph Γ separates X into at least two unbounded components if and only if $[\Gamma]$ is a non-trivial class in $H_c^1(X, \mathbb{Z}_2)$. Such a graph Γ is said to be essential.*

PROOF. The proof is similar to the argument of Lemma 20.11.

1. Suppose that $X \setminus \Gamma$ contains at least two unbounded complementary components U and V , but $[\Gamma] = 0$ in $H_c^1(X, \mathbb{Z}_2)$. Then there exists a compactly-supported function $\sigma : X^{(0)} \rightarrow \mathbb{Z}_2$ such that $\delta(\sigma) = m_\Gamma$, mod 2. Since σ is compactly supported, there exists a compact subset $C \subset X$ such that $\sigma = 0$ on $U \setminus C, V \setminus C$. Let $\alpha \subset X^{(1)}$ be a path connecting $u \in U \cap X^{(0)}$ to $v \in V \cap X^{(0)}$. We leave it to the reader to verify that if an edge $e = [x, y]$ of X crosses Γ in an even number of points then x, y belong to the same connected component of $X \setminus \Gamma$ (this is the only place where we use the assumption that Γ is connected). Therefore, the path α crosses Γ in an odd number of points, which implies that

$$\langle m_\Gamma, \alpha \rangle = 1 \in \mathbb{Z}_2.$$

However,

$$\langle m_\Gamma, \alpha \rangle = \langle \sigma, \partial\alpha \rangle = \sigma(u) + \sigma(v) = 0.$$

Contradiction.

2. Suppose that $[\Gamma] \neq 0$ in $H_c^1(X, \mathbb{Z}_2)$. Since $H^1(X, \mathbb{Z}_2) = 0$, there exists a 0-cochain $\sigma : X^{(0)} \rightarrow \mathbb{Z}_2$ such that

$$\delta\sigma = m_\Gamma.$$

Since m_Γ takes non-zero value on some edge $e = [u, v]$ of X , we obtain $\sigma(u) = 0, \sigma(v) = 1$. If the set $\sigma^{-1}(1) \subset X^{(0)}$ is finite, then $\sigma \in C_c^0(X, \mathbb{Z}_2)$ and, hence $[\Gamma] = 0$, which is a contradiction. Therefore, the set of such vertices is unbounded. Consider another 0-cochain $\sigma+1$ (which equals to $\sigma(x)+1$ on every vertex $x \in X^{(0)}$). Then $\delta(\sigma+1) = \delta\sigma = m$ and

$$\{w \in X^{(0)} : \sigma(w) = 0\} = \{w \in X^{(0)} : \sigma(w) + 1 = 1\}.$$

Therefore, by the above argument, the set $\sigma^{-1}(0) \subset X^{(0)}$ is also unbounded. Thus, since Γ is a finite graph, there are unbounded connected subsets $U, V \subset X \setminus \Gamma$ such that

$$\forall u \in U \cap X^{(0)}, \sigma(u) = 0, \quad \forall v \in V \cap X^{(0)}, \sigma(v) = 1.$$

These are the required unbounded complementary components of Γ . \square

EXERCISE 20.13. If $H^1(X, \mathbb{Z}_2) = 0$, then every connected essential Dunwoody track $\Gamma \subset X$ has exactly two complementary components, both of which are unbounded. We will use the notation Γ^\pm for these components.

The following key lemma due to Dunwoody is a direct generalization of the Kneser–Haken finiteness theorem for triangulated 3-dimensional manifolds, see e.g. [Hem78].

LEMMA 20.14 (M. Dunwoody). *Suppose that X has F faces and $H^1(X, \mathbb{Z}_2) \cong \mathbb{Z}_2^r$. Suppose that $\Gamma_1, \dots, \Gamma_k$ are pairwise disjoint pairwise non-isotopic connected tracks in X . Then*

$$k \leq 6F + r.$$

PROOF. The union Γ of tracks Γ_i cuts each 2-simplex s in X in triangles, rectangles and hexagons. Note that some of the complementary rectangles might contain vertices of X . In what follows, we regard such rectangles as *degenerate hexagons* (and not as rectangles). The boundary of each complementary rectangle has two disjoint edges contained in $X^{(1)}$, we call these edges *vertical*. Consider an

edge of Γ which is contained in the boundary of a complementary triangle or a (possibly degenerate) hexagon. The number of such edges is at most $6F$. Thus, the number of tracks Γ_i containing such edges is at most $6F$ as well. We now remove from X the union of closures of all components of $X \setminus \Gamma$ which contain complementary triangles and (possibly degenerate) hexagons.

The remainder R is a union of rectangles Q_j glued together along their vertical edges. Therefore, R is homeomorphic to an open interval bundle over a track $\Lambda \subset X$: The edges of Λ are geodesics connecting midpoints of vertical edges of Q_j 's. If a component R_i of R is a trivial interval bundle then the boundary of R_i is the union of tracks Γ_j, Γ_k which are therefore isotopic. This contradicts our assumption on the tracks Γ_i . Therefore, each component of R is a non-trivial interval bundle. For each R_i we define the cohomology class $[\Lambda_i] \in H^1(X, \mathbb{Z}_2) = H_c^1(X, \mathbb{Z}_2)$ (using the counting function m_Λ). We claim that these classes are linearly independent. Suppose to the contrary that

$$\sum_{i=1}^{\ell} [\Lambda_i] = 0.$$

This means that the track $\Lambda' := \Lambda_1 \cup \dots \cup \Lambda_\ell$ determines a trivial cohomology class $[\Lambda'] = 0$. Since Λ' is 1-sided, we obtain a contradiction with Lemma 20.11. Therefore, the number of components of R is at most r , the dimension of $H^1(X, \mathbb{Z}_2)$. Each component of R is bounded by a track Γ_i . Therefore, the total number of tracks Γ_i is at most $6F + r$. \square

20.3. Existence of minimal Dunwoody tracks

Our next goal is to deform finite transversal graphs to Dunwoody tracks, so that the cohomology class is preserved and so that the counting function $m_\Gamma : \text{Edges}(X) \rightarrow \mathbb{Z}$ decreases as the result of the deformation. To this end, we define the operation *pull* on transversal graphs $\Gamma \subset X$.

Pull. Suppose that v_1, v_2 are distinct vertices of Γ which belong to a common edge e of X and which are not separated by any vertex of $\Gamma \cap e$ on e . We call such vertex pair $\{v_1, v_2\}$ *innermost*. Then for every 2-face s of X containing e and every pair of distinct edges $\gamma_i = [u_i, v_i], i = 1, 2$ of Γ we perform the following operation. We replace $\gamma_1 \cup \gamma_2 \subset \Gamma$ by a single edge $\gamma = [u_1, u_2] \subset s$, keeping the rest of $\Gamma' = \Gamma \setminus \gamma_1 \cup \gamma_2$ unchanged, so that γ intersects Γ' only at the end-points u_1, u_2 . In case $\gamma_1 = \gamma_2$ we simply eliminate this edge from Γ . Let $\text{pull}(\Gamma)$ denote the resulting graph. It is clear that $\nu(\text{pull}(\Gamma)) < \nu(\Gamma)$ and $\text{pull}(\Gamma)$ is again a transversal graph. Note that *a priori*, $\text{pull}(\Gamma)$ need not be connected even if Γ is. See Figures 20.6 and 20.7.

EXERCISE 20.15. Verify that $[\text{pull}(\Gamma)] = [\Gamma]$; actually, the functions m_Γ and $m_{\text{pull}(\Gamma)}$ are equal as \mathbb{Z}_2 -cochains.

LEMMA 20.16. *Given a finite transversal graph $\Gamma \subset X$, there exists a finite sequence of pull-operations which transforms Γ to a new graph Γ' ; the graph Γ' is a track such that for every edge e , $m_{\Gamma'}(e) \in \{0, 1\}$. Moreover, $[\Gamma] = [\Gamma']$.*

PROOF. We apply the pull-operation to Γ as long as possible; since the vertex-complexity under *pull* is decreasing, this process terminates at some transversal graph Γ' . If $m_{\Gamma'}(e) \geq 2$ for some edge e of X , we can again perform *pull* using an




FIGURE 20.6. Pull.




FIGURE 20.7. Eliminating the edge $\gamma_1 = \gamma_2$. In this example, $u_1 = v_2, u_2 = v_1$

innermost pair of vertices of Γ' on e , which is a contradiction. Since $pull$ preserves the cohomology class of a transversal graph, $[\Gamma] = [\Gamma']$. \square

LEMMA 20.17. Assume that $H^1(X, \mathbb{Z}_2) = 0$. If $|Ends(X)| > 1$ then there exists a connected essential transversal graph $\Gamma \subset X$.

PROOF. Let ϵ_+, ϵ_- be distinct ends of X . We claim that there exists a proper 1-Lipschitz function $h : X \rightarrow \mathbb{R}$ such that

$$\lim_{x \rightarrow \epsilon_{\pm}} h = \pm\infty.$$

Indeed, let K be a compact which separates the ends ϵ_+, ϵ_- . We define h to be constant on K . We temporarily re-metrize X by equipping it with the standard metric (every simplex is isometric to the standard Euclidean simplex with unit edges). Let U_{\pm} be the unbounded components of $X \setminus K$ which are neighborhoods of the ends ϵ_{\pm} . We then set

$$h|_{U_{\pm}} := \pm \text{dist}(\cdot, K).$$

For every other component V of $X \setminus K$ we set

$$h|_V := \text{dist}(\cdot, K).$$

It is immediate that this function satisfies the required properties. We give \mathbb{R} the structure of a simplicial tree T , where integers serve as vertices. We next approximate h by a proper canonical map $f : X \rightarrow T$. Namely, for every vertex v of X we let $f(v)$ be a vertex of T within distance ≤ 1 from $h(v)$. We extend $f : X^{(0)} \rightarrow T$ to a canonical map $f : X \rightarrow T$. Then $\text{dist}(f, h) \leq 3$ and, hence, f is

again proper. By Lemma 20.6, $\Gamma := f^{-1}(y)$ is a finite transversal graph separating ϵ_+ from ϵ_- for almost every y . Hence, $[\Gamma] \neq 0$ in $H_c^1(X, \mathbb{Z}_2)$. The graph Γ need not be connected, let $\Gamma_1, \dots, \Gamma_n$ be its connected components: They are still transversal graphs. Thus,

$$[\Gamma] = \sum_{i=1}^n [\Gamma_i],$$

which implies that at least one of the graphs Γ_i is essential. □

Note, that the graph Γ constructed in the above proof need not be a Dunwoody track. However, Lemma 20.16 implies that we can replace Γ with a essential Dunwoody track Γ' which intersects every edge in at most one point. The graph Γ' need not be connected, but it has an essential connected component. Therefore:

COROLLARY 20.18. *Assume that $H^1(X, \mathbb{Z}_2) = 0$ and $|\text{Ends}(X)| > 1$. Then there exists a connected essential Dunwoody track $\Gamma \subset X$. Moreover,*

$$m_\Gamma : \text{Edges}(X) \rightarrow \{0, 1\}.$$

We define the *complexity* of a transversal graph $\Gamma \subset X$, denoted $c(\Gamma)$, to be the pair $(\nu(\Gamma), \ell(\Gamma))$, where $\nu(\Gamma)$ is the number of vertices in Γ and $\ell(\Gamma)$ is the total length of Γ , with respect to the metric on X' defined in Proposition 20.7. We give the set of complexities the lexicographic order. It is clear that $c(\Gamma)$ is preserved by isometric actions $G \curvearrowright X'$.

An essential connected Dunwoody track $\Gamma \subset X$ is said to be *minimal* if it has minimal complexity among all connected essential Dunwoody tracks in X .

DEFINITION 20.19. A vertex v of X is said to be a *cut-vertex* if $X \setminus \{v\}$ contains at least two unbounded components. (Note that our definition is slightly stronger than the usual definition of a cut-vertex, where it is only assumed that $X \setminus \{v\}$ is not connected.)

LEMMA 20.20. *Suppose that X admits a cocompact simplicial action $G \curvearrowright X$, $H^1(X, \mathbb{Z}_2) = 0$, $|\text{Ends}(X)| > 1$ and X has no cut-vertices. Then there exists a (connected and essential) minimal track Γ_{\min} .*

PROOF. By Corollary 20.18, the set of connected essential tracks in X is non-empty. Let Γ_i be a sequence of such graphs whose complexity converges to the infimum. Without loss of generality, we can assume that each Γ_i has minimal vertex-complexity $\nu = \nu(\Gamma)$ among all connected essential tracks in X . Since X is a simplicial complex, it is easy to see that each Γ_i is also a simplicial complex. Therefore, the number of edges of the graphs Γ_i is also uniformly bounded (by $\frac{\nu(\nu-1)}{2}$). In particular, there are only finitely many combinatorial types of these graphs; therefore, after passing to a subsequence, we can assume that the graphs Γ_i are combinatorially isomorphic to a fixed graph Γ .

Replace each edge of Γ_i with the hyperbolic geodesic (in the appropriate 2-simplex of X). This does not increase the complexity of Γ_i , keeps the graph embedded and preserves all the properties of Dunwoody tracks. Therefore, we will assume that each edge of Γ_i is geodesic. We let $h_i : \Gamma \rightarrow \Gamma_i$ be graph isomorphisms. Since $\ell(\Gamma_i)$ are uniformly bounded from above, there exists a path-metric on Γ such that all the maps h_i are 1-Lipschitz. We let \bar{h}_i denote the composition of h_i with the quotient map $X \rightarrow Y = X/G$. If there exists a compact set $C \subset Y' := X'/G$ such

that $\bar{h}_i(\Gamma) \cap C \neq \emptyset$, then the Arzela–Ascoli Theorem implies that the sequence (\bar{h}_i) subconverges to a 1-Lipschitz map $\Gamma \rightarrow Y'$. On the other hand, if such compact does not exist, then, since the edges of Y' have infinite length and Γ is connected, the sequence of maps \bar{h}_i subconverges to a constant map sending Γ to one of the vertices of Y . Hence, in this case, by post-composing the maps h_i with $g_i \in G$, we conclude that the sequence $g_i \circ h_i$ subconverges to a constant map whose image is one of the vertices of X . Recall that, by our assumption, X has no cut-vertices. Therefore, every sufficiently small neighborhood of a vertex v of X does not separate X into several unbounded components. This contradicts the assumption that each Γ_i is essential.

Therefore, by replacing h_i with $g_i \circ h_i$ (and preserving the notation h_i for the resulting maps), we conclude that the maps h_i subconverge to a 1-Lipschitz map $h : \Gamma \rightarrow X'$. In view of Lemma 20.16 (and the fact that Γ_i 's have minimal vertex-complexity), for every face s and edge $e \subset s$ of X there exists at most one edge of Γ_i contained in s and intersecting e . Therefore, the map h is injective and $\Gamma_{min} = h(\Gamma)$ is a track in X . Moreover, for each sufficiently large i , the graph Γ_{min} is isotopic to Γ_i as they have the same counting function $m_\Gamma = m_{\Gamma_i}$. Thus, Γ_{min} is essential. Therefore, it is the required minimal track. \square

20.4. Properties of minimal tracks

20.4.1. Stationarity. The following discussion is local and does not require any assumptions on $H^1(X, \mathbb{Z}_2)$.

We say that a transversal graph Γ is *stationary* if for every small smooth isotopy Γ_t of Γ (through transversal graphs), with $\Gamma_0 = \Gamma$, we have

$$\frac{d}{dt} \ell(\Gamma_t)|_{t=0} = 0.$$

In particular, every edge of Γ is geodesic.

EXAMPLE 20.21. Every minimal essential Dunwoody track is stationary.

Let Γ be a Dunwoody track with geodesic edges. Let v be a vertex of Γ which belongs to an edge e of X and γ an edge of Γ incident to v . We assume that $\gamma = \gamma(t)$ is parameterized by its arc-length so that $\gamma(0) = v$. We define $\pi_e(\gamma')$ to be the orthogonal projection of the vector $\gamma'(0) \in T_e \mathbb{H}^2$ to the tangent line of e at v .

LEMMA 20.22. *If Γ is stationary then for every vertex v as above we have*

$$(20.2) \quad \sum_{\gamma} \pi_e(\gamma') = 0$$

where the sum is taken over all edges $\gamma_1, \dots, \gamma_k$ of Γ incident to v .

PROOF. We construct a small isotopy Γ_t of Γ by fixing all the vertices and edges of Γ except for the vertex v which is moved along e , so that $v(t), t \in [0, 1]$, is a smooth function. We assume that all edges of Γ_t are geodesic. This variation of v uniquely determines Γ_t . It is clear that

$$0 = \frac{d}{dt} \ell(\Gamma_t)|_{t=0} = \sum_{i=1}^k \frac{d}{dt} \ell(\gamma_i(t))|_{t=0}.$$

By the first variation formula (4.11), we conclude that

$$0 = \sum_{i=1}^k \frac{d}{dt} \ell(\gamma_i(t))|_{t=0} = \sum_{i=1}^k \pi_e(\gamma'). \quad \square$$

REMARK 20.23. The proof of the above lemma also shows the following. Suppose that Γ fails the stationarity condition (20.2) at a vertex v . Orient the edge e and assume that the vector

$$\sum_{\gamma} \pi_e(\gamma')$$

points to the “right” of zero. Construct a small isotopy Γ_t , $\Gamma_0 = \Gamma$, $t \in [0, 1]$, so that all edges of Γ_t are geodesic, vertices of Γ_t except for v stay fixed, while the vertex $v(t)$ moves to the “right” of $v = v(0)$. Then

$$\ell(\Gamma_t) < \ell(\Gamma)$$

for all small $t > 0$.

LEMMA 20.24 (The Maximum Principle). *Let Λ_1, Λ_2 be stationary Dunwoody tracks. Then in a small product neighborhood U of every common vertex u of these graphs, either the graphs Λ_1, Λ_2 coincide, or one “crosses” the other. The latter means that*

$$\Lambda_1 \cap U_+ \neq \emptyset, \Lambda_2 \cap U_- \neq \emptyset.$$

Here $U_{\pm} = \Lambda_{1,U} \times (0, \pm 1]$ where we identify the product neighborhood U with $\Lambda_{1,U} \times [-1, 1]$, $\Lambda_{1,U} = \Lambda_1 \cap U$. In other words, if $h : U = \Lambda_{1,U} \times [-1, 1] \rightarrow [-1, 1]$ is the projection to the second factor, then $h|_{\Lambda_2}$ cannot have maximum or minimum at u , unless $h|_{\Lambda_2}$ is identically zero.

PROOF. Let e be the edge of X containing u . Since Λ_1, Λ_2 are tracks, every 2-simplex s of X adjacent to e contains (unique) edges $\gamma_{i,s} \subset \Lambda_i$, $i = 1, 2$, which are incident to u . Suppose that Λ_2 does not cross Λ_1 . Then either for every s , $\gamma_i = \gamma_{i,s}$, $i = 1, 2$ as above,

$$\pi_e(\gamma'_1) \geq \pi_e(\gamma'_2)$$

or for every s , γ_1, γ_2

$$\pi_e(\gamma'_1) \geq \pi_e(\gamma'_2).$$

Since, by the previous lemma,

$$\sum_s \pi_e(\gamma'_{i,s}) = 0, \quad i = 1, 2,$$

we conclude that $\pi_e(\gamma'_{1,s}) = \pi_e(\gamma'_{2,s})$. Therefore, since any geodesic is uniquely determined by its derivative at a point, it follows that $\gamma_{1,s} = \gamma_{2,s}$ for every 2-simplex s containing e . Thus, $\Lambda_1 \cap U = \Lambda_2 \cap U$. \square

20.4.2. Disjointness of essential minimal tracks. The following proposition is the key for the proof of Stallings Theorem presented in the next section:

PROPOSITION 20.25. *If $H^1(X, \mathbb{Z}_2) = 0$ then any two (connected, essential) minimal tracks in X are either equal or disjoint.*

PROOF. Our proof follows [JR89]. The central ingredients in the proof are the exchange and round-off arguments as well as the Meeks–Yau trick. All three come from the theory of least area surfaces in 3-dimensional Riemannian manifolds.

Suppose that Λ, M are distinct connected essential minimal tracks which have non-empty intersection.

Step 1: Transversal case. We first present an argument that this is impossible under the assumption that the graphs Λ and M are *transverse to each other*, i.e. $\Lambda \cap M$ is disjoint from the 1-skeleton of X . Since both Λ, M are essential and connected, each of them separates X into exactly two components, denoted Λ^\pm, M^\pm ; all of these components are unbounded, see Exercise 20.13. We consider the four sets

$$\Lambda^+ \cap M^+, \Lambda^+ \cap M^-, \Lambda^- \cap M^+, \Lambda^- \cap M^-.$$

Since both Λ, M separate ends of X , at least two of the above sets are unbounded. After relabeling, we obtain that the intersections

$$\Lambda^+ \cap M^+, \quad \Lambda^- \cap M^-$$

are unbounded. Observe that

$$\Lambda \cup M = \partial(\Lambda^+ \cap M^+) \cup \partial(\Lambda^- \cap M^-).$$

Set

$$\Gamma_+ := \partial(\Lambda^+ \cap M^+), \Gamma_- := \partial(\Lambda^- \cap M^-),$$

see Figure 20.8. It is immediate that both graphs are transversal (here and below we disregard valency 2 vertices of Γ_\pm contained in the interiors of 2-simplices of X). Note that, at this point, we do not yet know if the graphs Γ_\pm are connected.

We now compare the complexity of Λ (which is the same as the complexity of M) and complexities of the graphs Γ_+, Γ_- . After relabeling, $\ell(\Gamma_+) \leq \ell(\Gamma_-)$. We leave it to the reader to verify that for both Γ_+, Γ_- , the number of edges is the same as the number of edges of Λ . Clearly, the total length of $\Gamma_+ \cup \Gamma_-$ is the same as $2\ell(\Lambda) = 2\ell(M)$. Therefore,

$$\ell(\Gamma_+) \leq \ell(\Lambda).$$

Hence, $c(\Gamma_+) \leq c(\Lambda)$. The transition from Λ to the graph Γ_+ is called the *exchange argument*: We replaced parts of Λ with parts of M in order to get Γ_+ .




FIGURE 20.8. Exchange argument.

By the assumption, the intersection $\Lambda^+ \cap M^+$ is unbounded. The complement to this intersection contains $\Lambda^- \cap M^-$, which is also assumed to be unbounded. Therefore, the graph Γ_+ separates ends of X and, hence, $[\Gamma_+] \neq 0$ in $H_c^1(X, \mathbb{Z}_2)$, see Lemma 20.12. It is then immediate that at least one connected component of Γ_+ represents a non-trivial element of $H_c^1(X, \mathbb{Z}_2)$. Since $H^1(X, \mathbb{Z}_2) = 0$, this component

is essential. By minimality of the graph Λ , it follows that Γ_+ is connected (otherwise we replace it with the above essential component thereby decreasing the vertex-complexity). Since Λ, M cross at a point $x \notin X^{(1)}$, there exists an edge γ of Γ_+ which is a broken geodesic containing x in its interior. (Recall that $X \setminus X^{(0)}$ is equipped with a certain path-metric which is hyperbolic on each 2-dimensional simplex.) Replacing the broken edge γ with a shorter path (and keeping the end-points) we get a new graph $\widehat{\Lambda}$ whose total length is strictly smaller than the one of Γ_+ . (This part of the proof is called the “round-off” argument.) We obtain a contradiction with minimality of Λ . This finishes the proof in the case of transversal intersections of Λ and M .

Step 2: Weakly transversal case: Meeks–Yau trick. We assume now that $\Lambda \cap M$ contains at least one point p of transversal intersection which is not in the 1-skeleton of X . We say that in this situation Λ, M are *weakly transversal* to each other. Note that by doing the “exchange and round-off” at p we have some definite reduction in the complexity of the tracks, which depended only on the intersection angle α between Λ, M at p . Then the weakly transversal case is handled *via* the “original Meeks-Yau trick” [MY81], which reduces the proof to the transversal case. This trick was introduced in the work by Meeks and Yau in the context of minimal surfaces in 3-dimensional manifolds and generalized by Jaco and Rubinstein in the context of PL minimal surfaces in 3-manifolds, see [JR88]. The idea is to isotope Λ to a (non-minimal) geodesic graph Λ_t , whose total length is slightly larger than Λ but which is transversal to M : $\ell(\Lambda_t) = \ell(\Lambda) + o(t)$.

The intersection angle α_t between Λ_t and M near p can be made arbitrarily close to the original angle α . Therefore, by taking t small, one can make the complexity loss ϵ to be higher than the length gain $\ell(\Lambda_t) - \ell(\Lambda)$. Then, as in Case 1, we obtain a contradiction with minimality of Λ and M .

Step 3: Non-weakly transversal case. We, thus assume that $\Lambda \cap M$ contains no points of transversal intersection. (This case does not occur in the context of minimal surfaces in 3-dimensional Riemannian manifolds.) The idea is again to isotope Λ to Λ_t , so that $\ell(\Lambda_t) = \ell(\Lambda) + o(t)$. One then repeats the arguments from Step 1 (exchange and round off) and verifies that the new graph $\widehat{\Lambda}_t$ satisfies

$$\ell(\Lambda_t) - \ell(\widehat{\Lambda}_t) \geq O(t).$$

It will then follow that $\ell(\widehat{\Lambda}_t) < \ell(\Lambda)$ when t is sufficiently small, contradicting minimality of Λ .




FIGURE 20.9. Meeks-Yau trick: Initially, the graphs Λ, M had a common edge. After isotopy of Λ , this edge is no longer common. The isotopy Λ_t is through geodesic graphs, which no longer satisfy the balancing condition (20.2).

We now provide the details of the Meeks–Yau trick in this setting. We first push the graph Λ in the direction of Λ^+ , so that the result is a smooth family of isotopic Dunwoody tracks $\Lambda_t, t \in [0, t_0], \Lambda_0 = \Lambda$, where each Λ_t has geodesic edges and so that each vertex of Λ_t is within distance t from the corresponding vertex of Λ . Since Λ was stationary, we have

$$\ell(L_t) = \ell(\Lambda) + ct^2 + o(t^2).$$

It follows from the Maximum Principle (Lemma 20.24) that the graphs Λ_t and M have to intersect. For sufficiently small values of $t \neq 0$, the intersection is necessarily disjoint from $X^{(1)}$. We now apply the exchange argument and obtain a graph Γ_{t+} , such that

$$\ell(\Gamma_{t+}) \leq \ell(\Lambda) + ct^2 + o(t^2).$$

Let $\widehat{\Lambda}_t$ be obtained from Γ_{t+} by the round-off argument (straightening the broken edges). Lastly, we need to estimate from below the difference

$$\ell(\widehat{\Lambda}_t) - \ell(\Gamma_{t+}).$$

It suffices to analyze what happens within a single 2-simplex s of X where the graphs Λ_t and M intersect. We will consider only the most difficult case:

The geodesic segments λ and μ in s , which are components of $\Lambda \cap s$ and $M \cap s$ respectively and such that λ and μ share only their end-point A .

In particular, the point A belongs to an edge e of s . Furthermore, a component $\lambda_t \subset \Lambda_t \cap s$, such that

$$\lim_{t \rightarrow 0} \lambda_t = \lambda_0 := \lambda,$$

has non-empty intersection with M for small $t > 0$.

REMARK 20.26. If such face s and segments λ, μ do not exist, then $\Lambda = M$. In this situation, the geodesic segments λ_t and μ will be disjoint for small $t > 0$ and nothing interesting happens during the exchange and round-off argument.

We introduce the notation:

$$\lambda = [A, B], \quad \mu = [A, C], \quad \lambda_t = [A_t, B_t], \quad t \in [0, t_0].$$

By the construction, $\text{dist}(A, A_t) = t, \text{dist}(B, B_t) = t$. Set $D_t := \mu \cap \lambda_t$. There are several possibilities for the intersection $\Gamma_{t+} \cap s$. If this intersection contains the broken geodesics AD_tA_t or B_tD_tC , then the round-off of Γ_{t+} will result in reduction of the number of edges, contradicting minimality of Λ . We, therefore, consider the case when $\Gamma_{t+} \cap s$ contains the broken geodesic A_tD_tC , as the case of the path AD_tB_t is similar.

Consider the triangle $\Delta(A, D_t, A_t)$. We note that the angles of this triangle are bounded away from zero and π if t_0 is sufficiently small. Therefore, the Sine Law for hyperbolic triangles (4.4) implies that $\text{dist}(A_t, D_t) \sim t$ as $t \rightarrow 0$. Consider then the triangle $\Delta(A_t, D_t, C)$ (see Figure 20.10). Again, the angles of this triangle are bounded away from zero and π if t_0 is sufficiently small. Therefore, Lemma 4.41 implies that

$$\text{dist}(A_t, D_t) + \text{dist}(D_t, C) - \text{dist}(A_t, C) \geq c_1 \text{dist}(A_t, D_t) \geq c_2 t = O(t)$$

if t_0 is sufficiently small. Here c_1, c_2 are positive constants. Observe that when we replace Γ_{t+} with $\widehat{\Lambda}_t$ (the round-off), the path $[A_t, D_t] \cup [D_t, C]$ is replaced with the geodesic $[A_t, C]$. Therefore,

$$\ell(\Gamma_{t+}) - \ell(\widehat{\Lambda}_t) = O(t).$$




FIGURE 20.10. Meeks-Yau trick: Isotope the edge λ so that $D_t = \lambda_t \cap \mu$ is no longer on the edge e .

Since

$$\ell(\Gamma_{t+}) - \ell(\Lambda) = o(t),$$

we conclude that

$$\ell(\widehat{\Lambda}_t) - \ell(\Lambda) < 0$$

if t is sufficiently small. This contradicts minimality of Λ . \square

20.5. Stallings Theorem for almost finitely presented groups

DEFINITION 20.27. A group G is said to be *almost finitely presented* (afp) if it admits a properly discontinuous cocompact simplicial action on a 2-dimensional simplicial complex X such that $H^1(X, \mathbb{Z}_2) = 0$.

Note that every free simplicial action is properly discontinuous. Furthermore, in view of Lemma 5.103, in the definition of an afp group one can replace a complex X with a new simplicial complex \widehat{X} which is 2-dimensional, has $H^1(\widehat{X}, \mathbb{Z}_2) = 0$, and the action $G \curvearrowright \widehat{X}$ is free and cocompact.

LEMMA 20.28. *Every finitely presented group G is also afp.*

PROOF. Let Y be a finite presentation complex of G , subdivide it to obtain a triangulated complex W , then let X be the universal cover of W . \square

We are now ready to prove

THEOREM 20.29. *Let G be an almost finitely presented group with at least 2 ends. Then G splits as the fundamental group of a finite graph of finitely generated groups with finite edge-groups.*

PROOF. Since G is afp, it admits a properly discontinuous cocompact simplicial action on a (locally finite) 2-dimensional simplicial complex X with $H^1(X, \mathbb{Z}_2) = 0$. We give $X' := X \setminus X^{(0)}$ the piecewise-hyperbolic path metric as in Section 20.1.

DEFINITION 20.30. A subset $Z \subset X$ is called *precisely-invariant* (under its G -stabilizer) if for every $g \in G$ either $gZ = Z$ or $gZ \cap Z = \emptyset$.

PROPOSITION 20.31. *There exists an finite connected subgraph $\Lambda \subset X$ which separates ends of X and is precisely-invariant.*

PROOF. If X has a cut-vertex, then we take Λ to be this vertex. Suppose, therefore, that X contains no such vertices. Then, by Lemma 20.20, X contains a minimal (essential) Dunwoody track $\Lambda \subset X$. By Proposition 20.25, for every

$g \in G$, the track $M := g\Lambda$ (which is also minimal) is either disjoint from Λ or equal to Λ . \square

The proof of Theorem 20.29 then reduces to:

PROPOSITION 20.32. *Every finite subgraph $\Lambda \subset X$ as in Proposition 20.31 gives rise to a non-trivial action of G on a simplicial tree T with finite edge-stabilizers and finitely generated vertex groups.*

PROOF. Let Λ be either a cut-vertex of X or a finite connected essential precisely-invariant track $\Gamma \subset X$ (see Definition 20.30). We first consider the more interesting case of when Γ is a Dunwoody track.

We partition of X into components of $G \cdot \Gamma$ and components of $X \setminus G \cdot \Gamma$, which we will refer to as *complementary regions*. Each complementary region C_v is declared to be a *vertex* v of the partition and each $\Gamma_e := g \cdot \Gamma$ is declared to be an *edge* e . Since Γ is a Dunwoody track and $H_1(X, \mathbb{Z}_2) = 0$, the complement $X \setminus \Gamma_e$ consists of exactly two components Γ_e^\pm ; therefore, each edge of the partition is incident to exactly two (distinct) complementary regions. These regions represent vertices incident to e . Thus, we obtain a graph T . Since the action of G preserves the above partition of X , the group G acts on the graph T .

LEMMA 20.33. *The group G does not fix any vertices of T and does not stabilize any edges.*

PROOF. Suppose that G fixes a vertex v of T . Let E_v denote the set of edges of T incident to v . By relabeling, we can assume that the corresponding component C_v of $X \setminus G \cdot \Gamma$ equals

$$\bigcap_{e \in E_v} \Gamma_e^+.$$

Therefore, for every $x \in C_v$, $g \in G$, and $e \in E_v$ we have

$$g(x) \notin \Gamma_e^-.$$

Recall that the action $\Gamma \curvearrowright X$ is cocompact. Therefore, there exists a finite subcomplex $K \subset X$ whose G -orbit is the entire X . Clearly, $x \in K$ for some $x \in C_v$. On the other hand, by the above observation, the intersection

$$G \cdot K \cap \Gamma_e^-$$

is a finite subcomplex. This contradicts the fact that Γ_e^- is unbounded. Thus, G does not fix any vertex in T . Similarly, we see that G does not preserve any edge of T . \square

LEMMA 20.34. *The graph T is a tree.*

PROOF. Connectedness of T immediately follows from connectedness of X . If T were to contain a circuit, it would follow that some Γ_e did not separate X , which is a contradiction. \square

Lastly, we observe that compactness of Γ_e 's and proper discontinuity of the action $G \curvearrowright X$ imply that the stabilizer G_e of every edge e in G is finite. Note that, a priori, G acts on T with inversions since $g \in G$ can preserve Γ_e and interchange Γ_e^+, Γ_e^- .

Since the closure \overline{C}_v of each vertex-space C_v is connected and \overline{C}_v/G_v is compact, it follows that the stabilizer G_v of each vertex $v \in T$ is finitely generated (this is a special case of the Milnor-Schwartz Lemma).

Suppose now that Λ is a single vertex v . If Λ were to separate X into exactly two components, we would be done by repeating the arguments above. Otherwise, we modify X by replacing the vertex v with an edge e whose mid-point m separates X into exactly two components both of which are unbounded. We repeat this for every point in $G \cdot v$ in G -equivariant fashion. The result is a new complex \widehat{X} with a cocompact action $G \curvearrowright \widehat{X}$. Clearly, $\Lambda := \{m\}$ is precisely-invariant, and, hence, we are done as above. Proposition 20.32 follows. \square

In both cases, the quotient graph $\Gamma = T/G$ is finite since the action $G \curvearrowright X$ is cocompact.

We can now finish the proof of Theorem 20.29. In view of Proposition 20.32, Bass-Serre correspondence (Section 7.5.6), implies that the group G is the fundamental group of a non-trivial finite graph of groups \mathcal{G} with finite edge groups and finitely generated vertex groups. \square

EXERCISE 20.35. 1. Show that every vertex group of the graph of groups in constructed in Theorem 20.29 is afp. Hint: First prove that for every vertex-space C_v , $H^1(\overline{C}_v, \partial C_v; \mathbb{Z}_2) = H^1(X_v; \mathbb{Z}_2) = 0$, where the complex X_v is obtained by collapsing every boundary track of C_v to a point.

2. Show that the tree T defined in the proof of Theorem 20.29 satisfies

$$|\partial_\infty T| = |\epsilon(G)|.$$

20.6. Accessibility

Let G be a finitely generated group which splits non-trivially as an amalgam $G_1 \star_H G_2$ or $G_1 \star_H$ with finite edge-group H . Sometimes, this decomposition process can be iterated, by decomposing the groups G_i as amalgams with finite edge groups, etc. The key issue that we will be addressing in this section is:

Does the decomposition process terminate after finitely many steps?

If this decomposition process of G terminates then the group G is isomorphic to the fundamental group of a graph of groups, where all edge groups are finite and all vertex groups have at most one end. This leads to

DEFINITION 20.36. A group G is said to be *accessible* if it admits a (finite) graph of groups decomposition with finite edge groups and 1-ended vertex groups.

C. Thomassen and W. Woess in [TW93] prove:

THEOREM 20.37. *A finitely generated group G is accessible if and only if one (equivalently, every) Cayley graph Γ of G satisfies the following property:*

There exists a number D so that every two ends of Γ_G can be separated by a bounded subset of Γ of diameter $\leq D$.

In particular, accessibility is QI invariant.

Our first goal is to show that all finitely generated torsion-free groups are accessible. Recall that the rank of a finitely generated group is the least number of its generators.

THEOREM 20.38 (Grushko’s theorem). *Suppose that G is finitely generated and $G = G_1 \star G_2$ is a non-trivial free product decomposition. Then*

$$\text{rank}(G_i) < \text{rank}(G), \quad i = 1, 2.$$

We refer the reader to the paper by Scott and Wall [SW79] for a topological proof of this classical theorem in group theory. We can now prove:

THEOREM 20.39. *All torsion-free finitely generated groups are accessible.*

PROOF. If G is torsion-free, then all (inductively constructed) decompositions $G_1 \star_H G_2$ or $G_1 \star_H$ are just free products $G_1 \star G_2$ and $G_1 \star \mathbb{Z}$ respectively. Then, by Grushko’s theorem, $\text{rank}(G_i) < \text{rank}(G), i = 1, 2$ and, hence, the decomposition process terminates after at most $\text{rank}(G)$ steps. \square

M. Dunwoody ([Dun93]) constructed an example of a finitely generated group which is not accessible. The main result of this section is

THEOREM 20.40 (M. Dunwoody, [Dun85]). *Every almost finitely presented group is accessible.*

Our goal below is to give a proof of Dunwoody’s theorem, mostly following papers by M. Dunwoody [Dun85] and G. Swarup [Swa93]. Before proving Dunwoody’s theorem, we will establish several technical facts.

Refinements of graphs of groups. Let \mathcal{G} be a graph of groups with the underlying graph Γ , let $H = G_v$ be one of its vertex groups. Let \mathcal{H} be a graph of groups decomposition of H with the underlying graph Λ . Suppose that:

ASSUMPTION 20.41. For every edge $e \subset \Gamma$ incident to v , the subgroup $G_e < H$ is conjugate in H to a subgroup of one of the vertex groups H_w of \mathcal{H} , $w = w(e)$ (this vertex need not be unique). For instance, if every G_e is finite, then, in view of Property FA for finite groups, G_e will fix a vertex in the tree corresponding to \mathcal{H} . Thus, our assumption will hold in this case.

Under this assumption, we can construct a new graph of groups decomposition \mathcal{F} of G as follows. Cut Γ open at v , i.e. remove v from Γ and then replace each open or half-open edge of the resulting space with a closed edge. The resulting graph Γ' could be disconnected. We have the natural map $r : \Gamma' \rightarrow \Gamma$. Let Φ denote the graph obtained from the union $\Gamma' \sqcup \Lambda$ by identifying each vertex $v'_i \in r^{-1}(v) \in e'_i \subset \Gamma'$ with the vertex $w(e) \in \Lambda$ as in the above assumption. Then Φ is connected. We retain for Φ the vertex and edge groups and the group homomorphisms coming from Γ and Φ . The only group homomorphisms which need to be defined are for the edges $e = [e^-, e^+]$, where $e^- = w(e) = w$. In this case, the embedding $G_e \rightarrow G_w$ is the one given by the conjugation of G_e to the corresponding subgroup of G_w .

We leave it to the reader to verify (using Seifert – Van Kampen theorem) that $\pi_1(\Phi) \simeq G$.

DEFINITION 20.42. The new graph of groups \mathcal{F} is called the *refinement* of \mathcal{G} via \mathcal{H} . A refinement is said to be *trivial* if \mathcal{H} is a trivial graph of groups. We use the notation $\mathcal{G} \prec \mathcal{F}$ for a refinement.

PROPOSITION 20.43. *Let \mathcal{G} be a finite graph of finitely generated groups with finite edge-subgroups (with the underlying graph Λ). Then:*

1. *Every vertex subgroup G_v is QI embedded in $G = \pi_1(\mathcal{G})$. (Note that finite generation of the vertex groups implies that G is itself finitely generated.)*




FIGURE 20.11. Cut Γ open and glue Φ from Γ' and Λ .

2. If, in addition, the group G is finitely presented, then every vertex group of \mathcal{G} is also finitely presented.

PROOF. The proofs of both statements are very similar. We first construct (as in Section 7.5.4) a tree of graphs \mathcal{Z} , corresponding to \mathcal{G} .

Namely, let $G \curvearrowright T$ be the action of G on a tree corresponding to the graph of groups \mathcal{G} (see Section 7.5.4). For every edge-group G_e in \mathcal{G} we take $S_e := G_e \setminus \{1\}$ as the generating set of G_e . For every vertex-group G_v in \mathcal{G} we pick a finite generating set S_v of G_v , such that for every edge $e = [v, w]$, the sets S_v, S_w contain the images of S_e under the embeddings $G_e \rightarrow G_v, G_e \rightarrow G_w$.

Then, using the projection $p : T \rightarrow \Lambda = T/G$, we define generating sets $S_{\tilde{v}}, S_{\tilde{e}}$ of $G_{\tilde{v}}, G_{\tilde{e}}$ using the isomorphisms $G_{\tilde{v}} \rightarrow G_v, G_{\tilde{e}} \rightarrow G_e$, where \tilde{v}, \tilde{e} project to v, e under the map $T \rightarrow \Lambda$. Let $Z_{\tilde{v}}, Z_{\tilde{e}}$ denote the Cayley graphs of the groups $G_{\tilde{v}}, G_{\tilde{e}}$ ($\tilde{v} \in V(T), \tilde{e} \in E(T)$) with respect to the generating sets $S_{\tilde{v}}, S_{\tilde{e}}$.

Now, to simplify the notation, we will use the notation v and e for vertices and edges of T . Let Z_v, Z_e denote the Cayley graphs of the groups G_v, G_e ($v \in V(T), e \in E(T)$) with respect to the generating sets S_v, S_e defined above. We have natural injective maps of graphs $f_{ev} : Z_e \hookrightarrow Z_v$, whenever v is incident to e . The resulting collection of graphs Z_v, Z_e and embeddings $Z_e \hookrightarrow Z_v$, defines a *tree of graphs* \mathcal{Z} with the underlying tree T . For each Z_e we consider the product $Z_e \times [-1, 1]$ with the standard triangulation of the product of simplicial complexes. Let \tilde{Z}_e be the 1-skeleton of this product. Lastly, let Z denote the graph obtained by identifying vertices and edges of each \tilde{Z}_e with their images in Z_v under the maps $f_{ev} \times \{\pm 1\}$. We endow Z with the standard metric.

Clearly, the group G acts on Z freely. The quotient graph Z/G has only finitely many vertices; this quotient graph is finite if each G_v is finitely generated.

For every $v \in V(T)$ define the map $\rho_v : Z^0 \rightarrow Z_v^0 = G_v$ to be the G_v -equivariant nearest-point projection. For every edge $e = [v, w]$, the subgraph $f_v(Z_e) \subset Z_v$ separates Z , and every two distinct vertices in $f_v(Z_e)$ are connected by an edge in this graph. It follows that for $x, y \in Z^0$ within unit distance from each other,

$$\text{dist}(\rho_v(x), \rho_v(y)) \leq 1.$$

Hence, the map ρ_v is 1-Lipschitz. We then extend, G_v -equivariantly, each ρ_v to the entire graph Z .

We now can prove the assertions of the proposition.

1. Since each G_v is finitely generated, the action $G \curvearrowright Z_v$ is geometric and, hence, the action $G \curvearrowright Z$ is geometric as well. Thus, the space Z is QI to the group G and Z_v is QI to G_v for every vertex v . Let $x, y \in Z_v$ be two vertices and $\alpha \subset Z$ be the shortest edge-path connecting them. Then $\rho(\alpha) \subset Z$ still connects x to y and its length is at most the length of α . It follows that Z_v is isometrically embedded in Z . Hence, each G_v is QI embedded in G . This proves (1).

2. Since G is finitely presented and $G \curvearrowright Z$ is geometric, the space Z is coarsely simply-connected by Corollaries 9.36 and 9.55. Our goal is to show that each vertex space Z_v of Z is also coarsely simply-connected. Let α be a loop in a vertex space Z_v . Since Z is coarsely simply-connected, there exists a constant C (independent of α) and a collection of (oriented) loops α_i in the 1-skeleton of Z such that

$$\alpha = \prod_i \alpha_i$$

and each α_i has length $\leq C$. We then apply the retraction ρ to the loops α_i . Then

$$\alpha = \prod_i \rho(\alpha_i)$$

and $\text{length}(\rho(\alpha_i)) \leq \text{length}(\alpha_i) \leq C$ for each i . Thus, Z_v is coarsely simply-connected and, therefore, G_v is finitely presented. \square

We are now ready to prove Dunwoody's accessibility theorem.

Proof of Theorem 20.40. We will construct inductively a finite chain of refinements

$$\mathcal{G}_1 \prec \mathcal{G}_2 \prec \mathcal{G}_3 \dots$$

which are graphs of groups with finite edge groups, afp vertex groups, and so that the terminal graph of groups has only finite and 1-ended vertex groups.

We need a notion of *complexity* (which will be denoted $\sigma(G)$) for afp groups G which generalizes the notion of *rank* used for the proof of accessibility in the torsion-free case. Note that if we drop the assumption $H^1(X, \mathbb{Z}_2) = 0$ below and assume instead that X is a graph, then $\sigma(G)$ equals $\text{rank}(G) + 1$.

DEFINITION 20.44 (σ -complexity). Suppose that X be a 2-dimensional simplicial complex with $H^1(X, \mathbb{Z}_2) = 0$, such that X admits a simplicial properly discontinuous, cocompact action $G \curvearrowright X$. We let $\sigma(G, X)$ denote the total number of cells in the cell-complex X/G (the quotient need not be a simplicial complex). Accordingly, we will use the notation $\sigma(G)$ for the minimum of the numbers $\sigma(G, X)$ where the minimum is taken over all complexes X and group actions $G \curvearrowright X$ as above. If \mathcal{G} is a finite graph with afp vertex groups, then $\sigma(\mathcal{G})$ is defined to be the maximum of complexities $\sigma(G_v)$, taken over all vertex groups G_v of \mathcal{G} .

We will show that *some* process of refinement results in the strict reduct of the σ -complexity. Such refinement process necessarily terminates.

Let G be an afp group and X be a G -complex which realizes the complexity $\sigma(G)$.

First, suppose that such X has a cut-vertex v (see Definition 20.19). Then, as in the proof of Theorem 20.29, we split G as a graph of groups (with the edge-groups stabilizing v) so that each vertex-group G_i acts on a subcomplex $X_i \subset X$, where

the frontier of X_i in X is contained in $G \cdot v$. It follows from the Mayer-Vietoris sequence that $H^1(X_i, \mathbb{Z}_2) = 0$ for each i . Thus, $\sigma(G_i) < \sigma(G)$ for every i .

Hence, without loss of generality, we may assume that X has no cut-vertices. If the group G has at most one end, we are done. Suppose that G has at least 2 ends. Then, by Propositions 20.31 and 20.32, there exists a (connected) finite precisely-invariant track $\tilde{\tau}_1 \subset X_1 := X$ which determines a non-trivial graph of groups decomposition \mathcal{G}_1 of $G_1 := G$ with finite edge groups. Our assumption that $H^1(X, \mathbb{Z}_2) = 0$ implies that τ_1 is 2-sided in $X_1 := X$. Let X_2 be the closure of a connected component of $X \setminus G \cdot \tilde{\tau}_1$. By compactness of X/G and Milnor-Schwartz Lemma, the stabilizer G_2 of X_2 in G is finitely generated. Since $H^1(X, \mathbb{Z}_2) = 0$, it follows by excision and Mayer-Vietoris sequence that

$$H^1(X_2, \partial X_2; \mathbb{Z}_2) = 0,$$

where ∂X_2 is the frontier of X_2 in X_1 .

Therefore, if define W_2 by pinching each boundary component of X_2 to a point, then $H^1(W_2, \mathbb{Z}_2) = 0$. The stabilizer G_2 of X_2 in G acts on W_2 properly discontinuously and cocompactly. Therefore, each vertex group of \mathcal{G}_1 is again afp.

If each vertex group of \mathcal{G}_1 is 1-ended, we are again done. Suppose therefore that the closure X_2 of some component $X_1 \setminus G \cdot \tilde{\tau}_1$ as above has stabilizer $G_2 < G_1$ which has at least two ends. According to Theorem 9.24, G_2 splits (non-trivially) as a graph of groups with finite edge groups. Let $G_2 \curvearrowright T_2$ be a non-trivial action of G_2 on a simplicial tree (without inversions) which corresponds to this decomposition. Since each edge-group of \mathcal{G}_1 is finite, if such a group is contained in G_2 , it has to fix a vertex in T_2 , see Corollary 3.75. Recall that the edge-groups of \mathcal{G}_1 are conjugate to the stabilizers of components of $G \cdot \tilde{\tau}_1$ in G . Therefore, every such stabilizer fixes a vertex in T_2 . We let X_2^+ denote the union of X_2 with all simplices in X which have non-empty intersection ∂X_2 . Clearly, G_2 still acts properly discontinuously cocompactly on X_2^+ . The stabilizer of each component of $X_2^+ \setminus X_2$ is an edge group of \mathcal{G}_1 .

We then construct a (G_1 -equivariant) PC map (see Definition 20.5) $f_2 : X_2^+ \rightarrow T_2$ such that:

1. f_2 sends components of $X_2^+ \setminus \text{int}(X_2)$ to the corresponding fixed vertices in T_2 .
2. f_2 sends vertices to vertices of T .
3. f_2 is linear on each edge of the cell-complex X_2 .

If the image of the map f_2 is bounded then the action $G_2 \curvearrowright T_2$ has a bounded orbit. By Cartan's Theorem (Theorem 3.74), G_2 fixes a point in T_2 , which contradicts the assumption that the action $G_2 \curvearrowright T_2$ is non-trivial.

Therefore, the image of f_2 contains an edge $e \subset T_2$ which separates T_2 into at least two unbounded subsets.

Then, by Lemma 20.6, there exists a point $t \in e$ which is a regular value of f_2 . Thus, by Exercise 20.8, $f^{-1}(t)$ is a track. It is immediate that $f^{-1}(t)$ is precisely-invariant in X_2 with finite G_2 -stabilizer. By the choice of e , the graph $f^{-1}(t)$ separates X into at least two unbounded components. Let $\tilde{\tau}_2$ be an essential component of $f^{-1}(t)$.

Thus, by Proposition 20.32, the track $\tilde{\tau}_2$ determines a decomposition of G_2 as a graph of groups \mathcal{G}_3 with finite edge groups. We continue this decomposition

inductively. We obtain a collection of pairwise disjoint connected tracks $\tau_1, \tau_2, \dots \subset Y = X/G$ which are projections of the tracks $\tilde{\tau}_i \subset X$.

Suppose that the number of tracks τ_i is $> 6F + r$, where F is the number of faces in X and r is the dimension of $H^1(X, \mathbb{Z}_2)$. Then, by Lemma 20.14, every τ_k , $k > 6F + r$ is isotopic to some graph $\tau_{i(k)}$, $i = i(k) \leq 6F + r$. Let R be the product region in Y bounded by such tracks. Lifting this region in X we again obtain a product region \tilde{R} bounded by tracks $g_i \tilde{\tau}_i, g_k \tilde{\tau}_k$, $g_i, g_k \in G$. Therefore, the stabilizers of $g_i \tilde{\tau}_i, g_k \tilde{\tau}_k$ and R in G have to be the same. It follows that every X_k , $k > 6F + r$ is a product region whose stabilizer fixes its boundary components. The corresponding tree T_k is just the union of two edges which are fixed by the entire group G_k . This contradicts the assumption that each refinement $\mathcal{G}_k \prec \mathcal{G}_k$ is non-trivial. Therefore, the decomposition process of G terminates after $6F + r$ steps and G is accessible. \square

20.7. QI rigidity of virtually free groups and free products

THEOREM 20.45. *If G is virtually free, then every group G' which is QI to G is also virtually free.*

PROOF. If the group G is finite, the assertion is clear. If G is virtually cyclic, then G' and G are 2-ended, which, by Part 3 of Theorem 9.22, implies that G' is also virtually cyclic.

Suppose now that G has infinitely many ends. Since G is finitely presented, by Corollary 9.55, the group G' is finitely presented as well. The group G acts geometrically on a locally finite simplicial tree T with infinitely many ends, therefore, the groups G and G' are QI to T . Since G' is finitely presented, by Theorem 20.40, the group G' splits as a graph of groups where every edge group is finite, every vertex group is finitely generated and each vertex group has 0 or 1 ends.

By Proposition 20.43, every vertex group G'_v of this decomposition is QI embedded in G' . In particular, every G'_v is quasi-isometrically embedded in a simplicial tree of finite valence. The image of such an embedding is coarsely-connected (with respect to the restriction of the metric on T) and, therefore, is within finite distance from a subtree $T'_v \subset T$. Thus, each G'_v is QI to a locally-finite simplicial tree (embedded in T).

LEMMA 20.46. *T'_v cannot have one end.*

PROOF. Suppose that T'_v has one end. The group G'_v is infinite and finitely generated. Therefore, its Cayley graph contains a bi-infinite geodesic (see Exercise 7.84). Such geodesic γ cannot embed quasi-isometrically in a 1-ended tree (since both ends of γ would have to map to the same end of T'_v). \square

Thus, every vertex group G'_v has zero ends and, hence, is finite. By Theorem 7.51, the group G is virtually free. \square

We conclude this chapter by generalizing Theorem 20.45 to quasi-isometries of arbitrary finitely presented groups. The following theorem is due to P. Papasoglu and K. Whyte, [PW02], who proved it for finitely generated groups. We refer the reader to C. Cashen's paper [Cas10] for further analysis.

THEOREM 20.47 (QI rigidity for graphs of groups with finite edge groups). *Suppose that G, G' are quasi-isometric finitely presented groups and $f : G' \rightarrow G$ is a quasi-isometry. Assume also that the group G has a decomposition as a finite*

graph of groups \mathcal{G} with finite edge groups and finitely generated vertex groups which have at most one end. Then G' also admits a decomposition \mathcal{G}' , such that for every 1-ended vertex group G'_v in \mathcal{G}' , $f(G'_v)$ is Hausdorff-close to a conjugate of a 1-ended vertex group of \mathcal{G} . In particular, every 1-ended vertex group of \mathcal{G} is QI to a vertex group of \mathcal{G}' .

PROOF. Finite presentability implies that both groups G and G' are accessible. As in the proof of Theorem 20.45, we conclude that:

The group G' also splits as a finite graph of groups \mathcal{G}' with finite edge groups, with finitely generated vertex groups each of which has at most two ends. We associate with the graph \mathcal{G} a tree of spaces \mathfrak{X} as in Section 7.5.6, with the underlying cell complex X . Recall that there exists a G -equivariant projection $p : X \rightarrow T$ to the associated Bass-Serre tree of \mathcal{G} .

Each edge e of T splits the tree T in two subtrees T_e^\pm , which are the two components of the subgraph of T obtained by erasing the edge e . (We will see below how to assign the \pm labels to these subtrees, depending on a vertex v of \mathcal{G}' .) We define two subspaces $X_e^\pm \subset X$, the unions of all vertex spaces X_v with $v \in V(T_e^\pm)$. Consider the Cayley graph Γ'_v of a 1-ended vertex group G'_v of \mathcal{G}' . Since the edge-spaces of \mathfrak{X} have bounded diameter, they cannot coarsely separate the image $Y := f(\Gamma'_v)$ in X . Thus, there exists a number D such that for every edge e in $E := E(T)$, exactly one of the two subspaces X_e^+, X_e^- coarsely contains Y ; we denote this subspace X_e^+ :

$$Y \subset \mathcal{N}_D(X_e^+).$$

REMARK 20.48. Note that since G'_v has infinitely many ends, Y cannot be coarsely contained in X_e^- . Thus, the choice of the vertex v defines a label \pm for each edge e of \mathcal{G} . Accordingly, we obtain a collection of subtrees T_e^+ , $e \in E$.

We claim that the intersection of all these subtrees,

$$T^+ := \bigcap_{e \in E} T_e^+,$$

consists of a single vertex w . It then would immediately follow that

$$Y \subset \mathcal{N}_D\left(\bigcap_{e \in E} X_e^+\right).$$

Observe that it is not even clear that the intersection T^+ is non-empty. We note that, since no edge-space X_e coarsely separates Y , the projection of Y to T is contained in a subtree $A \subset T$ of finite diameter (otherwise, there exists an edge $e \in E$ such that Y contains points in X_e^+, X_e^- arbitrarily far from X_e). Since the graph Γ'_v is connected, and has infinite diameter, it follows that Γ'_v contains an embedded half-line H , see Exercise 2.18. Lemma 2.19 then shows that there exists a vertex $w \in V(T)$ such that $p^{-1}(w) \cap Y$ is unbounded. If w were a vertex of some T_e^- , then the intersection

$$X_w \cap Y \subset p^{-1}(T_e^-) \cap Y$$

would be bounded. Therefore, w belongs to the intersection of all subtrees T_e^+ :

$$w \in T^+.$$

Lastly, if T^+ contains another vertex $u \neq w$, then it also contains an edge e separating these two vertices. Since $w \in T_e^+$, $u \in T_e^-$, which means that u is not in T^+ . This is a contradiction. Therefore, $T^+ = \{w\}$.

To summarize: There is $D \in \mathbb{R}$ such that for each 1-ended vertex subgroup G'_v of \mathcal{G}' , there exists a vertex $w \in T$ such that $f(\Gamma'_v) \subset \mathcal{N}_D(X_v)$. Applying the coarse inverse map $\bar{f} : G \rightarrow G'$, we conclude that $f(G'_v)$ and G_w are Hausdorff-close to each other. \square

Proof of Stallings' Theorem using harmonic functions

In this chapter we will prove Stallings' theorem in full generality:

THEOREM 21.1. *If G is a finitely generated group with infinitely many ends then G admits a non-trivial decomposition as a graph of groups with finite edge groups.*

In his essay [Gro87, Pages 228–230], Gromov outlined a proof of Stallings' theorem using harmonic functions. The goal of this chapter is to provide details for Gromov's arguments. One advantage is that this proof works for finitely generated groups without the finite presentability assumption. However, the geometry behind the proof is not as transparent as in Chapter 20. The proof that we present in this chapter is a simplified form of the arguments which appear in [Kap14]. The simplification presented here (in the proof of Theorem 21.5) is due to Bruce Kleiner.

We refer the reader to the material of Section 3.9 for the definition of harmonic functions on Riemannian manifolds and to Section 3.4 for the discussion of the coarea formula. Both will be key ingredients in the proof.

Every finitely generated group G admits an isometric free properly discontinuous cocompact action $G \curvearrowright M$ on a Riemannian manifold M , which, then, necessarily has bounded geometry (since it covers a compact Riemannian manifold).

EXAMPLE 21.2. If G is k -generated, and N is a Riemann surface of genus k , we have an epimorphism

$$\phi : \pi_1(N) \rightarrow G.$$

Then G acts isometrically and cocompactly on the covering space M of N with $\pi_1(M) = \text{Ker}(\phi)$.

The space $\epsilon(M)$ of ends of M is naturally homeomorphic to the space of ends of G , see Proposition 9.18. Let $\bar{M} := M \cup \epsilon(M)$ denote the compactification of M by its space of ends; the action of the group G extends to a topological action of G on \bar{M} .

We will see in Section 21.3 that every continuous function $\chi : \epsilon(M) \rightarrow \{0, 1\}$, admits a unique continuous extension

$$h = h_\chi : \bar{M} \rightarrow [0, 1],$$

such that the function $h|_M$ is harmonic (the function h is the energy minimizer among all extensions of χ lying in a suitable Sobolev space). Furthermore, unless χ is constant, the extension h restricts to a proper function $h : M \rightarrow (0, 1)$.

REMARK 21.3. This theorem fails if M has one or two ends. For instance, every harmonic function $h : \mathbb{R} \rightarrow (0, 1)$ has to be linear and, hence, constant.

We will refer to the function h as the *harmonic extension* of χ (even though, only its restriction to M is harmonic). Let $H(M)$ denote the space of functions $M \rightarrow \mathbb{R}$ which are harmonic extensions nonconstant functions $\chi : \epsilon(M) \rightarrow \{0, 1\}$. We equip $H(M)$ with the topology of uniform convergence on compacts in M . For each $h \in H(M)$ we define its *energy*

$$E(h) := E(h|_M),$$

see Section 3.9 for the definition of energy of functions $M \rightarrow \mathbb{R}$. We will see in Section 21.3 that $E(h)$ is always finite.

DEFINITION 21.4. We define the *energy gap* $e(M)$ of M as

$$e(M) := \inf\{E(h) : h \in H(M)\}.$$

The isometric group action $G \curvearrowright M$ yields a linear action $G \curvearrowright H(M)$

$$g \cdot h = h \circ g^{-1}$$

which preserves the functional E . Therefore E projects to a lower semi-continuous (see Theorem 3.41) functional $E : H(M)/G \rightarrow \mathbb{R}_+$, where we give $H(M)/G$ the quotient topology. The main technical result needed for the proof of Stallings' theorem is:

THEOREM 21.5. 1. $e(M) > 0$.

2. The functional $E : H(M)/G \rightarrow \mathbb{R}_+$ is proper in the sense that the preimage

$$E^{-1}([0, T])$$

is compact for every $T \in \mathbb{R}_+$. In particular, $e(M)$ is attained.

We now sketch our proof of Stallings' theorem. Let $h \in H(M)$ be an energy-minimizing harmonic function guaranteed by Theorem 21.5, $E(h) = e(M)$. We then verify that the set $S = \{h(x) = \frac{1}{2}\}$ is *precisely-invariant* with respect to the action of G (see Definition 20.30), it also separates the ends of M .

REMARK 21.6. The set S , as each level set of a nonconstant harmonic function, has zero measure (as it is a union of a smooth hypersurface and a subset of Hausdorff dimension $< \dim(M) - 1$, see [Bär97]), but we will not need this property.

If the level set S is connected, we can use the standard construction of a dual simplicial tree T using the tiling of M by the components of $M \setminus G \cdot S$, cf. Section 7.5.6 and the proof of Theorem 20.29. The edges of T in this case are the "walls", i.e. hypersurfaces $g(S)$, $g \in G$. Every wall then lies in the boundary of exactly two components of $M \setminus G \cdot S$, which are the adjacent vertices.

This construction does not work in the case when S is not connected. We still use hypersurfaces $g(S)$ as the edges, but the definition of vertices has to be modified. Instead of the topological separation (which is meaningless if we separate disconnected subsets of M), we define separation via functions from the set $\mathcal{M} = \{g^*h, g^*(1-h) : g \in G\}$. A union V of components of $M \setminus G \cdot S$ is called *indecomposable* if V cannot be separated by any function $f \in \mathcal{M}$. These indecomposable sets are the vertices of T . We then show that each $g(S)$ lies in the boundary of exactly two indecomposable subsets of $M \setminus G \cdot S$, thereby defining a graph on which G is acting. Since the action of G on M is proper and S is compact, the edge stabilizers for the action $G \curvearrowright T$ are finite. Lastly, we verify that the graph T is a tree and that the action of G on T is non-trivial, i.e. G does not fix a point in T .

21.1. Proof of Stallings' theorem

The goal of this section is to prove Stallings' theorem assuming the results of Section 21.3 and Theorem 21.5.

Let $H(M)$ denote the space of harmonic functions $h : M \rightarrow [0, 1]$ as above. According to Theorem 21.5, there exists a function $h = h_\chi \in H(M)$ with minimal energy $E(h) = e(M) > 0$. We will refer to the harmonic function h as *minimal*. Then, for every $g \in G$, the function

$$g^*h := h \circ g$$

has the same energy as h and (in view of uniqueness of the harmonic extension) equals

$$h_{g^*(\chi)}.$$

Following Gromov, for $g \in G$, define two new functions

$$g_+(h) := \max(h, g^*(h)), \quad g_-(h) := \min(h, g^*(h)).$$

Clearly,

$$g_-(h)(x) = g_+(h)(x) \iff h(x) = g^*h(x).$$

We will see (Lemma 21.20) that

$$E(g_+(h)) + E(g_-(h)) = 2E(h).$$

Both functions $g_+(h), g_-(h)$ admit continuous extension to \bar{M} : The function $g_+(h)$ extends to $\chi_+ := \max(\chi, g^*(\chi))$ and $g_-(h)$ extends to $\chi_- := \min(\chi, g^*(\chi))$. The functions χ_\pm take only the values 0 and 1 on $\epsilon(M)$. Define

$$h_\pm := h_{\chi_\pm},$$

the harmonic extensions of χ_\pm . Since harmonic functions are energy-minimizers,

$$E(h_\pm) \leq E(g_\pm(h)),$$

and, hence,

$$(21.1) \quad E(h_+) + E(h_-) \leq E(g_+(h)) + E(g_-(h)) = 2E(h) = 2e(M).$$

REMARK 21.7. The functions h are *functional analogues* of the minimal tracks in Chapter 20. The definition of the functions $g_\pm(h)$ is an analogue of the “exchange” argument and the definition of the functions h_\pm is an analogue of the “round off” argument.

Note that it is, *a priori*, possible that χ_- or χ_+ is constant. Set

$$G^c := \{g \in G : \chi_- \text{ or } \chi_+ \text{ is constant}\}.$$

We first analyze the set $G \setminus G^c$. For $g \notin G^c$, both h_- and h_+ belong to $H(M)$ and, hence, by (21.1),

$$E(h_+) = E(h_-) = E(h) = e(M),$$

and

$$E(g_+(h)) = E(h_+), \quad E(g_-(h)) = E(h_-).$$

It follows that both functions $g_\pm(h)$ are harmonic. Since

$$g_-(h) \leq g_+(h),$$

the maximum principle (see Corollary 3.47) implies that either $g_-(h) = g_+(h)$ or $g_-(h) < g_+(h)$. Indeed, if $g_-(h)(x) = g_+(h)(x)$ at some $x \in M$, then the difference

$$g_+(h) - g_-(h)$$

is harmonic and attains its absolute maximum at $x \in M$. The maximum principle then implies that the difference $g_+(h) - g_-(h)$ is constant, hence, equals to zero.

Hence, the set $\{h = g^*h\}$ is either empty or equals the entire M , in which case $g^*(h) = h$. Therefore, for every $g \in G \setminus G^c$ one of the following holds:

1. $g^*h = h$.
2. $g^*h(x) < h(x)$, $\forall x \in M$.
3. $g^*h(x) > h(x)$, $\forall x \in M$.

In particular, the level set

$$S := h^{-1}\left(\frac{1}{2}\right)$$

is *precisely-invariant* under the elements of $G \setminus G^c$: For every $g \in G \setminus G^c$, either

$$g(S) = S$$

or

$$g(S) \cap S = \emptyset.$$

(The equality case occurs iff $g^*h = h$ or $g^*(h) = 1 - h$.)

We now consider elements $g \in G^c$: For such g 's either χ_- is constant or χ_+ is constant. Since $\chi_- \leq \chi_+$ and both functions only take the values 0 and 1, either $\chi_- \equiv 0$ or $\chi_+ \equiv 1$.

Case 1: $g \in G^c$ is such that $\chi_- \equiv 0$. In other words, whenever $\chi(\xi) = 1$, we also have $g^*(\xi) = 0$, $\xi \in \epsilon(M)$. It follows that

$$g^*\chi \leq 1 - \chi.$$

We claim that

$$g^*h \leq 1 - h.$$

Indeed, suppose that $h_1 = h_{\chi_1}, h_2 = h_{\chi_2} \in H(M)$ are such that $\chi_1 \leq \chi_2$, but

$$h_{\chi_1}(x) > h_{\chi_2}(x)$$

for some $x \in M$. Then the difference

$$f = h_{\chi_1} - h_{\chi_2}$$

is harmonic, positive at x and $\lim_{n \rightarrow \infty} f(x_n) \leq 0$ for each sequence (x_n) in M diverging to infinity. This means that there exists an open bounded subset $\Omega \subset M$ with smooth boundary, which contains x and such that $f|_{\Omega}$ attains its maximum at the point x (and not on the boundary). This contradicts the Maximum Principle for harmonic functions. Applying this to the harmonic functions $h_1 = g^*h$ and $h_2 = 1 - h$, we conclude that

$$g^*h \leq 1 - h.$$

The same argument shows that if $g^*h(x) = 1 - h(x)$ for some $x \in M$, then $g^*h = 1 - h$. The latter implies that

$$g(S) = S,$$

where $S = \{h = \frac{1}{2}\}$ as before.

This leaves us with the case

$$g^*h < 1 - h.$$

Then, clearly, $g(S) \cap S = \emptyset$.

Case 2: $\chi_+ \equiv 1$. We then replace χ with $\chi' = 1 - \chi$, h with $h_{\chi'}$, and conclude that $\chi'_- \equiv 0$. Then, by appealing to the Case 1, we see that either $g^*h = h$ or $g^*h > h$.

We thus proved an analogue of the Proposition 20.31 in the proof of Stallings' theorem for almost finitely presented groups in Chapter 20:

LEMMA 21.8. *Each minimal harmonic function h satisfies:*

1. *For every $g \in G$ one of the following occurs:*

$$(21.2) \quad g^*h = h, \quad g^*h < h, \quad g^*h > h, \quad g^*h = 1 - h, \quad g^*h < 1 - h, \quad g^*h > 1 - h.$$

2. *The subset $h^{-1}(1/2) = S \subset M$ is compact and precisely-invariant under the action of the group G . Moreover, for each $g \in G$, if $g(S) = S$ then either $g^*h = h$ or $g^*h = 1 - h$.*

We let G_S denote the stabilizer of S in G . Since S is compact and the action $G \curvearrowright M$ is properly discontinuous, the group G_S is finite.

By the construction, the subset S separates M into at least two unbounded components: One where $h > 1/2$ and the other where $h < 1/2$.

We now show that G splits non-trivially over a subgroup of G_S . (As we noted in the beginning of the chapter, the proof is straightforward under the assumption that S is connected and has two complementary components, but requires extra work in general.) We proceed by constructing a simplicial G -tree T on which G acts non-trivially, *possibly with inversions* and with finite edge-stabilizers.

Construction of T . Given a minimal harmonic function h , define the set of minimal functions

$$\mathcal{M} = \{g^*h, g^*(1 - h) : g \in G\}.$$

Each function $f \in \mathcal{M}$ defines the *wall* $W_f = \{x : f(x) = 1/2\}$ and the *half-spaces* $W_f^+ := \{x : f(x) > 1/2\}$, $W_f^- := \{x : f(x) < 1/2\}$ (these spaces are not necessarily connected). Note that

$$W_f^+ = W_{1-f}^-.$$

Let \mathcal{E} denote the set of walls. We say that a wall W_f *separates* $x, y \in M$ if

$$x \in W_f^+, \quad y \in W_f^-.$$

EXERCISE 21.9. For $f_1, f_2 \in \mathcal{M}$, $W_{f_1} \cap W_{f_2} = \emptyset$ unless $f_1 = f_2$ or $f_1 + f_2 = 1$.

EXERCISE 21.10. No two points on the same wall W_{f_1} can be separated by a wall W_{f_2} . Hint: See Lemma 21.8.

Maximal subsets V of

$$M^\circ := M \setminus \bigcup_{f \in \mathcal{M}} W_f$$

consisting of points which cannot be separated from each other by a wall, are called *indecomposable* subsets of M° . Similarly to the walls, these sets need not be connected. Set

$$\mathcal{V} := \{\text{indecomposable subsets of } M^\circ\}.$$

We will refer to the elements of \mathcal{V} as *vertex spaces* and to the walls W_f as *edge-spaces*. We say that a wall W is *adjacent* to $V \in \mathcal{V}$ if $W \cap cl(V) \neq \emptyset$.

LEMMA 21.11. *If a vertex space V is contained in $W_{f_1}^- \cap W_{f_2}^+$ and W_{f_1}, W_{f_2} are both adjacent to V then $f_1 < f_2$ on M .*

PROOF. We have

$$(21.3) \quad f_1|_V < \frac{1}{2} < f_2|_V.$$

Since f_1, f_2 are both in \mathcal{M} and $f_1 \neq f_2$, $f_1 + f_2 \neq 1$, Lemma 21.8 shows that either $f_i > f_{3-i}$ for some $i \in \{1, 2\}$ or $f_1 + f_2 < 1$ or $f_1 + f_2 > 1$ on the entire manifold M . If the former occurs then the inequality (21.3) implies that $f_1 < f_2$.

Consider the case $f_1 + f_2 < 1$. Then, taking a point $x \in cl(V) \cap W_{f_1}$ and taking into account that $f_1(x) = 1/2, f_2(x) > 1/2$, we obtain

$$f_1(x) + f_2(x) > 1,$$

a contradiction. In the remaining case $f_1 + f_2 > 1$ we take $x \in cl(V) \cap W_{f_2}$ and reach a similar contradiction. \square

LEMMA 21.12. *Each wall $W = W_f$ is adjacent to exactly two indecomposable sets $V^+, V^- \in \mathcal{V}$ (contained in W_f^+, W_f^- respectively).*

PROOF. We first construct two vertex spaces adjacent to W .

There exist sequences $x_k^\pm \in W_f^\pm$ which converge to $x^\pm \in W$. The inequalities (21.2) imply that for large k the points x_k^+ belong to the same vertex space, which we denote by V^+ and the points x_k^- belong to the same vertex space, which we denote by V^- .

Now, let us prove that V^\pm are the only vertex spaces adjacent to W . Suppose that $V \subset W_f^+$ is adjacent to W . Considering a sequence $y_k^+ \in V$ converging to some $y \in W$, we conclude that no wall separates y_k^+ from x_k^+ for large k (since this holds for the limit points y, x). \square

We define a graph T with the vertex set \mathcal{V} and the edge set \mathcal{E} , where $W \in \mathcal{E}$ connects V^+ and V^- if and only if W is adjacent to both.

LEMMA 21.13. *The graph T is a tree.*

PROOF. By the construction, every point of M belongs to a wall or to an indecomposable set. Hence, connectedness of T follows from connectedness of M .

Suppose that T contains a circuit with the consecutive vertices

$$V_1, V_2, \dots, V_k, V_{k+1} = V_1,$$

and the edges $W_{f_1}, \dots, W_{f_k}, W_{f_{k+1}} = W_{f_1}$, where for each $i = 1, \dots, k+1$, $W_{f_i} = [V_i, V_{i+1}]$ and $V_i \neq V_j$ unless $i = j$ or $|i - j| = k$. By replacing the functions f_i with $1 - f_i$ if necessary, we may assume that

$$V_i \subset W_{f_i}^- \cap W_{f_{i+1}}^+, i = 1, \dots, k.$$

In view of Lemma 21.11, we have

$$f_1 < f_2 < f_3 < \dots < f_k < f_{k+1}.$$

Note that $W_{f_1} = W_{f_{k+1}}$ and we either have $f_1 = f_{k+1}$ or $f_1 + f_{k+1} = 1$. The former clearly contradicts the inequality $f_1 < f_{k+1}$. In the later case,

$$V_{k+1} \subset W_{f_k}^- \cap W_{f_{k+1}}^+ = W_{f_k}^- \cap W_{f_1}^- \subset W_{f_1}^-.$$

However, $V_{k+1} = V_1 \subset W_{f_1}^+$, a contradiction. \square

We next note that G acts naturally on T since the sets \mathcal{M} , \mathcal{E} and \mathcal{V} are G -invariant and G preserves adjacency. If $g(W_f) = W_f$, then $g^*f = f$ or $g^*(f) = 1-f$, which implies that g preserves both W_f^+ and W_f^- or swaps them. Thus, it is possible that the G -action on T is not without inversions. The stabilizer of an edge in T corresponding to a wall W is finite, since W is compact and G acts on M properly discontinuously.

It remains to verify non-triviality of the action of G on T . Suppose that $G \curvearrowright T$ has a fixed vertex. This means that the indecomposable subset $V \subset M$ defining this vertex is G -invariant. Since G acts cocompactly on M , it follows that $M = \mathcal{N}_r(V)$ for some $r \in \mathbb{R}_+$. The indecomposable subset V is contained in the half-space W_f^+ for some wall W_f . Since W_f is compact and W_f^- is not, the subset W_f^- is not contained in $\mathcal{N}_r(W_f)$. Thus $W_f^- \setminus \mathcal{N}_r(V) \neq \emptyset$. This is a contradiction.

Lastly, in order to obtain a G -action on a tree without inversions, we barycentrically subdivide the tree T . Therefore we obtain a non-trivial graph of groups decomposition of G where the edge groups are conjugate to subgroups of the finite group G_S . \square

21.2. Nonamenability

The goal of the section is to show that each group G with infinitely many ends is nonamenable; accordingly, the manifold M (on which G acts geometrically) has $\lambda_1(M) > 0$, equivalently, M has positive Cheeger constant. This property will be used in constructing harmonic extensions h_χ of functions $\chi : \epsilon(M) \rightarrow \{0, 1\}$ and proving Part 1 of Theorem 21.5.

Let X be a metric space. A metric ball $B(x, r) \subset X$ is a *neck* (more precisely, an r -neck) if

$$X \setminus B(x, r)$$

has at least three unbounded components. The point x is the *center* of the r -neck. The following theorem was proven by C. Pittet in [Pit98]:

THEOREM 21.14. *Let X be a connected graph (equipped with the standard metric) such that there exists $r > 0$ for which every vertex $x \in V(X)$ is the center of an r -neck in X . Then X is nonamenable.*

PROOF. Let m be an integer such that $m > 4r + 2$. Define $V \subset V(X)$ as a maximal m -separated subset of $V(X)$. We will prove the theorem by constructing a map

$$f : V \rightarrow V$$

such that $d(v, f(v)) \leq 2m + 1$ and $f^{-1}(u)$ has cardinality ≥ 2 for every $u \in V$. Then the theorem will follow from Theorem 18.4.

The construction of f is somewhat reminiscent of the proof of the Milnor-Schwarz Theorem. Fix a vertex $v_0 \in V$. For $v \in V$ such that $d(v_0, v) \leq m$ set $f(v) = v_0$. Otherwise, take a geodesic $g \subset X$ connecting v_0 to v and let $x \in g$ be the vertex of X with $d(x, v) = m + 1$. Then let $f(v)$ be a point $w \in V$ closest to x . (If there are several such points, pick one at random.) By maximality of V , $d(x, w) \leq m$. Therefore,

$$d(v, f(v)) \leq 2m + 1.$$

Before proving the statement about cardinality of $f^{-1}(u)$, we will need a technical lemma.

LEMMA 21.15. *Pick $u \in V$ and consider the ball $B = B(u, r) \subset X$. Let C, C' be distinct components of $X \setminus B$. Then for $v \in V \cap C, v' \in V \cap C'$ we have*

$$d(v, v') > m + 1.$$

PROOF. Every geodesic g connecting v to v' is the union

$$vx \cup xx' \cup x'v'$$

where $x \in g \cap B, x' \in g \cap B$ and $xx' \subset B$. Then

$$d(v, v') \geq d(v, x) + d(v', x') \geq d(v, u) - r + d(v', u) - r \geq 2m - 2r > m + 1$$

by our choice of m . □

We will now proceed to proving the inequality on the cardinality of $f^{-1}(u)$. For $u \in V$ let C be one of the (at least two) unbounded components of $X \setminus B(u, r)$ which does not contain v_0 . Let $v \in C \cap V$ be a point closest to v_0 . We claim that $f(v) = u$. Since there are at least two such components C , we will then conclude that $f^{-1}(u)$ contains at least two points.

We let g denote a geodesic in X connecting v_0 to v . This geodesic necessarily passes through the ball $B = B(u, r)$.

Case 1: $d(v_0, v) \leq m$ (which implies that $f(v) = v_0$ by the definition of f). The vertices v_0, v cannot belong to distinct components of $X \setminus B$, as it would contradict Lemma 21.15. This means that v_0 has to belong to the ball B , i.e. $v_0 = u$. Thus, in this case, $f(v) = u$, as required.

Case 2: $d(v_0, v) \geq m + 1$. Let $x \in g$ be the vertex with $d(x, v) = m + 1$.

Subcase 2a: $x \notin C \cup B$. Pick $y \in g \cap B$. Then

$$m + 1 = d(x, v) \geq d(x, x') + (d(u, v) - d(y, u)) \geq d(x, y) + m - r,$$

which implies that

$$d(x, y) \leq r + 1.$$

Therefore,

$$d(x, u) \leq d(x, y) + r \leq 2r + 1.$$

If $w \in V$ is a vertex with $d(x, w) \leq 2r + 1$ then

$$d(u, w) \leq 4r + 2 < m$$

implying that $u = w$, as the set V is m -separated. Therefore, in this case, $f(v) = u$.

Subcase 2b: $x \in B$. Then $d(x, u) \leq r$ and, hence, for every $u' \in V \setminus \{u\}$, $d(x, u') \geq 2m - r > 2$. Therefore, in this case, again, $f(v) = u$.

Subcase 2c: $x \in C$. We leave it to the reader to verify that for every component C' of $X \setminus (B \cup C)$, if $v' \in C' \cap V$ then $d(u, x) < d(v', x)$, implying that

$$f(v) \in \{u\} \cup C.$$

Suppose, that $f(v) = w \in V \cap C$. Then $d(x, w) \leq m$ and, hence,

$$d(v_0, w) \leq d(v_0, x) + m < d(v_0, v)$$

as $d(v_0, x) = m$. This, however, contradicts the choice of v as the point in $V \cap C$ closest to v_0 . This leaves only one possibility: $f(v) = u$. □

COROLLARY 21.16. *Suppose that M is a Riemannian manifold which admits an isometric properly discontinuous cocompact action of a group G with infinitely many ends. Then G is nonamenable and $\lambda_1(M) > 0$.*

PROOF. Since amenability is QI invariant, G is amenable if and only if its Cayley graph X is. Since the graph X is nonamenable by Theorem 21.14, its Cheeger constant is positive, $h(M) > 0$ (see Theorem 18.14), equivalently, $\lambda_1(M) > 0$, see Theorem 3.53. \square

21.3. An existence theorem for harmonic functions

Theorem 21.17 below was originally proven by Kaimanovich and Woess in Theorem 5 of [KW92] using probabilistic methods (they also proved it for arbitrary continuous functions with values in $[0, 1]$). At the same time, an analytical proof of this result was given by Li and Tam [LT92], see also [Li12, Chapter 21] for a detailed and more general treatment.

Let M be a Riemannian manifold as in the beginning of this chapter (M admits a geometric action of a group G with infinitely many ends). We owe the following proof to Mohan Ramachandran:

THEOREM 21.17. *Let $\chi : \epsilon(M) \rightarrow \{0, 1\}$ be a continuous function. Then:*

1. χ admits a continuous harmonic extension to M .
2. This harmonic extension h has finite energy.

PROOF. We let $C_c^\infty(M)$ denote the space of C^∞ functions $M \rightarrow \mathbb{R}$ with compact support. For $u, v \in C_c^\infty(M)$ define the inner product

$$\langle u, v \rangle = \int_M uv \, dV$$

where dV is the Riemannian volume density on M . In what follows we will use this volume density to define a measure (of the Lebesgue class) on the sigma-algebra of Borel subsets of M . We let

$$\|u\|_{L_2} = \langle u, u \rangle^{1/2}$$

denote the norm of u with respect to this inner product. Since the differential of each function $u \in C_c^\infty(M)$ is also compactly supported, the energy $E(u)$ of the function u is finite.

We leave it to the reader to verify that the quantity

$$\|u\| := \|u\|_{L_2} + \sqrt{E(u)}$$

is also a norm on $C_c^\infty(M)$. We define a *Sobolev space* $W_o^{1,2}(M)$ as the completion of $C_c^\infty(M)$ with respect to the norm $\|u\|$. The space $W_o^{1,2}(M)$ sits naturally in the Hilbert space $L^2(M)$. Furthermore, by the construction, the energy functional extends continuously to $W_o^{1,2}(M)$.

Recall that every function χ extends to a continuous function $\varphi : \bar{M} \rightarrow \mathbb{R}$ which is smooth on M , see Lemma 9.26.

We let $L_{loc}^2(M)$ denote the space of functions of M which are *locally* in L^2 , i.e. functions whose restrictions to compact subsets $K \subset M$ are in $L^2(K)$. By continuity, every extension φ above, belongs to $L_{loc}^2(M)$. Thus, for a fixed function φ we define the affine subspace of functions

$$\mathcal{G} := \varphi + W_o^{1,2}(M) \subset L_{loc}^2(M).$$

Then the energy is a continuous (nonlinear) functional on \mathcal{G} and we set $E := \inf_{f \in \mathcal{G}} E(f)$.

Note that, since \mathcal{G} is affine, for $u, v \in \mathcal{G}$ we also have

$$\frac{u+v}{2} \in \mathcal{G},$$

in particular,

$$E\left(\frac{u+v}{2}\right) \geq E.$$

We set

$$E(u, v) := 2E\left(\frac{u+v}{2}\right) - \frac{E(u) + E(v)}{2}.$$

The latter equals

$$(21.4) \quad E(u, v) := \int_M \langle \nabla u, \nabla v \rangle dV$$

in the case when u, v are smooth. We thus obtain,

$$E(u, v) \geq 2E - \frac{E(u) + E(v)}{2}$$

for all $u, v \in \mathcal{G}$. Hence, in view of (21.4), by continuity of E ,

$$(21.5) \quad E(u - v) = E(u) + E(v) - 2E(u, v) \leq 2E(u) + 2E(v) - 4E.$$

Pick a sequence $u_n \in \mathcal{G}$ such that

$$\lim_{n \rightarrow \infty} E(u_n) = E.$$

Then, according to (21.5),

$$E(u_m - u_n) \leq 2E(u_n) + 2E(u_m) - 4E = 2(E(u_n) - E) + 2(E(u_m) - E).$$

We now come to the first, and only, point where the assumption that the number of ends M is infinite (and not 2) is used:

Since $\lambda = \lambda_1(M) > 0$ (Theorem 21.14), we obtain, by the definition of the bottom of the spectrum (3.6),

$$(21.6) \quad \lambda \int_M f^2 dV \leq E(f)$$

for all $f \in C_c^\infty(M)$ and, hence, by continuity, for all $f \in W_o^{1,2}(M)$. Therefore, the functions $v_n := u_n - \varphi \in W_o^{1,2}(M)$ satisfy

$$\|v_n - v_m\| \leq (2 + \lambda^{-1})(E(u_n) - E + E(u_m) - E).$$

Hence, the sequence (v_n) is Cauchy in $W_o^{1,2}(M)$. Set

$$v := \lim_n v_n, u := \varphi + v \in \mathcal{G}.$$

By semicontinuity of energy (Theorem 3.41), $E(u) = E$. Since u minimizes energy among all functions in \mathcal{G} , it is necessarily harmonic and, hence, u is smooth (see Section 3.9). Since $d\varphi$ is compactly supported (its support K is contained in the support of φ), the function v is also harmonic away from the compact subset $K \subset M$. By the inequality (21.6), we have

$$(21.7) \quad \int_M v^2 dV \leq \lambda^{-1} E(v) < \infty.$$

Let $r > 0$ denote the injectivity radius of M . Pick a base-point $o \in M$. Then (21.7) implies that there exists a function $\rho : M \rightarrow \mathbb{R}_+$ which converges to 0 as $d(x, o) \rightarrow \infty$, such that

$$\int_{B(x,r)} v^2(x) dV \leq \rho(x)$$

for all $x \in M$. By the mean value inequality (Corollary 3.49 in Section 3.9), there exists $C_1 < \infty$, such that

$$\sup_{x \in B(x,r)} v^2(x) \leq C_1 \inf_{B(x,r)} v^2$$

provided that $d(x, K) \geq r$. Therefore,

$$v^2(x) \leq \frac{C_1}{\text{Vol}(B(x,r))} \int_{B(x,r)} v^2 \leq C_2 \rho(x),$$

and, thus,

$$\lim_{d(x,o) \rightarrow \infty} v(x) = 0.$$

We conclude that the harmonic function u extends to the function χ on $\epsilon(M)$. \square

Properties of harmonic extensions.

PROPOSITION 21.18. 1. For each continuous function $\chi : \epsilon(M) \rightarrow \{0, 1\}$ its harmonic extension $h : M \rightarrow \mathbb{R}$ is unique.

2. The unique harmonic extension takes values in the interval $[0, 1]$.

3. If $h(x) = 0$ or $h(x) = 1$, for some $x \in M$, then h is constant.

PROOF. We prove all three properties by appealing to the Maximum Principle for harmonic functions. Since the proofs are analogous to the ones which appear in Lemma 21.8, our arguments will be somewhat brief.

1. Suppose that $h_1, h_2 : M \rightarrow \mathbb{R}$ are harmonic extensions of $\chi : \epsilon(M) \rightarrow \{0, 1\}$. Then the difference $h = h_1 - h_2$ is harmonic and for every sequence $x_n \in M$ diverging to infinity,

$$\lim_{n \rightarrow \infty} h(x_n) = 0.$$

Hence, h attains its maximum or minimum at a point $x \in M$. By the Maximum Principle, h is constant, implying that $h_1 = h_2$.

2. Let h be the unique harmonic extension of χ . Suppose that there exists $x \in M$ such that $h(x) \geq 1$. Then h again attains its maximum at a point of M , implying that h is constant. This is impossible if $h(x) > 1$. The same argument, with $1 - h$ replacing h , handles the case $h(x) \leq 0$. This proves (2) as well as (3). \square

21.4. Energy of minimum and maximum of two smooth functions

The arguments here are again due to Mohan Ramachandran.

Let M be a smooth manifold and f be a C^1 -smooth function on M . Define the function $f^+ := \max(f, 0)$ and the closed set

$$\Gamma := \{x \in M : f(x) = 0, df(x) = 0\}.$$

Set

$$\Omega := \{x \in M : f(x) = 0, df(x) \neq 0\} = f^{-1}(0) \setminus \Gamma.$$

By the implicit function theorem, Ω is a smooth submanifold in M and, hence, has measure zero. Clearly, f^+ is smooth on $M \setminus \{f = 0\}$.

LEMMA 21.19. *Under the above conditions, a.e. on M we have: $df^+(x) = df(x)$ if $f(x) > 0$ and $df^+(x) = 0$ if $f(x) \leq 0$.*

PROOF. The claim is clear at the points x where $f(x) \neq 0$. Since Ω has measure zero, it suffices to prove the assertion for points $x_0 \in \Gamma$. Choose local coordinates on M at a point $x_0 \in \Gamma$, so that $x_0 = 0$. Since f has zero derivative at 0, we have:

$$\lim_{v \rightarrow 0} \frac{|f(v)|}{\|v\|} = 0.$$

Since $0 \leq |f^+| \leq |f|$, it follows that

$$\lim_{v \rightarrow 0} \frac{|f^+(v)|}{\|v\|} = 0.$$

Therefore, f^+ is differentiable at x_0 and $df^+(x_0) = 0$. □

Consider now two C^1 -smooth functions f_1, f_2 on M . Define

$$f_{max} := \max(f_1, f_2), \quad f_{min} := \min(f_1, f_2), \quad f := f_1 - f_2.$$

LEMMA 21.20. $E(f_1) + E(f_2) = E(f_{max}) + E(f_{min})$.

PROOF. Set

$$M_1 := \{f_1 > f_2\}, M_2 := \{f_2 > f_1\}, M_0 := \{f_1 = f_2\}.$$

Since

$$f_{max} = f_2 - f^+, \quad f_{min} = f_1 - f^+,$$

by the above lemma we have:

$$\nabla f_{max} = \nabla f_2, \quad \nabla f_{min} = \nabla f_1 \quad \text{a.e. on } M_0.$$

Clearly,

$$\nabla f_{max} = \nabla f_i|_{M_i}, \quad \nabla f_{min} = \nabla f_{i+1}|_{M_{i+1}}, \quad i = 1, 2.$$

Hence,

$$\int_{M_i} (|\nabla f_{max}|^2 + |\nabla f_{min}|^2) dV = \int_{M_i} (|\nabla f_1|^2 + |\nabla f_2|^2) dV, \quad i = 0, 1, 2.$$

Therefore,

$$E(f_1) + E(f_2) = E(f_{max}) + E(f_{min}). \quad \square$$

21.5. A compactness theorem for harmonic functions

21.5.1. Positive energy gap implies existence of an energy minimizer.

Let M be a bounded geometry Riemannian manifold with infinitely many ends and positive Cheeger constant $\geq c > 0$, and let $\bar{M} = M \cup \epsilon(M)$ be the end compactification of M .

We state several definitions and notations used in what follows. For an m -dimensional Riemannian manifold N (possibly with boundary), we let $|N|$ denote the m -dimensional volume of N . Given a function $f : N \rightarrow \mathbb{R}$, we set $Var(f) := \sup(f) - \inf(f)$, the *variation* of f on N . For a function f on N , we define the *average* of f ,

$$\int_N f = \frac{\int_N f dV}{Vol(N)}.$$

In order to simplify the notation, we will frequently omit dV in the notation for integrals. Let $U \subset M$ be a smooth codimension 0 submanifold with compact boundary K . The *capacitance* $\text{cap}(U, K)$ of the pair (U, K) is the infimum of energies of compactly supported functions $u : U \rightarrow [0, 1]$, which are equal to 1 on K . We refer to Section 21.2 for the definition of R -necks in M . Note that each proper function $f : M \rightarrow (t_1, t_2)$ admits a continuous extension to \bar{M} : We will always retain the name f for this extension. The same convention will be used for functions defined on subsets with compact boundary in M .

Let \mathcal{F} denote the collection of continuous functions u on \bar{M} , whose restriction to $\epsilon(M)$ is nonconstant, and takes values in $\{0, 1\}$, while u is differentiable almost everywhere on M .

In Section 21.5.4 we will prove Part 1 of Theorem 21.5:

THEOREM 21.21. *There exists $\mu > 0$ such that every $u \in \mathcal{F}$ has energy at least μ , i.e. $e(M) > 0$, M has positive energy gap.*

Our goal below is to derive Part 2 of Theorem 21.5 from Part 1. We first state several corollaries of Theorem 21.21.

COROLLARY 21.22. *For each $U \subset M$ and K as above, $\text{cap}(U, K) \geq \mu$.*

PROOF. Given a function $u : U \rightarrow [0, 1]$ which equals 1 on K , we extend u by 1 to the rest of M . Then, clearly, the extension \tilde{u} has the same energy as u (since $\nabla \tilde{u}$ vanishes on $M \setminus U$) and $u \in \mathcal{F}$. Therefore, $E(u) = E(\tilde{u}) \geq \mu$. \square

As an application, we prove:

PROPOSITION 21.23. *Assume that every point in M belongs to an R -neck. Then for all $0 < a < b < 1$, $E \in [0, \infty)$, there is an $r = r(a, b, E) \in (0, \infty)$ with the following property. If $u : M \rightarrow (0, 1)$ is a proper a.e. smooth map, and $p \in M$, then either:*

- (1) $u(B(p, r))$ is not contained in the interval $[a, b]$, or
- (2) the energy of u is at least E .

PROOF. Define $s = \min(a^2, 1 - b^2)$. Since every point of M belongs to an R -neck, there exists r_0 , such that the complement of every ball $B(p, r_0)$ has more than

$$k = \frac{E}{s\mu}$$

unbounded components.

We claim that the desired property holds for $r = r_0$ (and, hence, for all greater values of r as well). Suppose that this fails. For a point $p \in M$ the distance function $d(p, \cdot)$ is smooth away from p and a.e. $r \in \mathbb{R}_+$ is its regular value. Thus, for generic $r \geq r_0$ (which we fix from now on), the metric ball $B(p, r) \subset M$ has smooth boundary. We let \mathcal{C} denote the collection of unbounded components of $M \setminus B(p, r)$. Let $u : M \rightarrow (0, 1)$ be a proper a.e. smooth map such that $u(B(p, r)) \subset [a, b]$, while u has energy $\leq E$. For each $U \in \mathcal{C}$, the function u takes the values in $[a, b]$ on $K = \partial U$. Consider the two functions $u^+ = \max\{b, u\}$ and $u^- = \min\{a, u\}$ on U . Then

$$E(u^\pm) \leq E(u|_U)$$

and $u^+|_K \equiv b, u^-|_K \equiv a$. Let \tilde{u}^\pm denote the extension of u^\pm to the rest of M such that

$$\tilde{u}^\pm|_{M \setminus U} \equiv u^\pm|_K.$$

Then

$$E(\tilde{u}^\pm) = E(u^\pm) \leq E(u|_U).$$

Note that u^-, \tilde{u}^- and u^+, \tilde{u}^+ are proper functions to intervals $(0, a)$ and $(b, 1)$ respectively.

Consider the function \tilde{u}^- : Its values on $\epsilon(M)$ belong to $\{0, a\}$ and a is one of its values. If \tilde{u}^- does not take zero value on $\epsilon(M)$, then $u|_{\epsilon(U)}$ takes only the value 1.

Case 1: The function \tilde{u}^- takes both values 0 and a on $\epsilon(M)$. Then the rescaled function $\frac{1}{a}\tilde{u}^-$ belongs to \mathcal{F} and, hence,

$$E(u|_U) \geq E(\tilde{u}^-) \geq a^2\mu$$

by Theorem 21.21.

Case 2: The function $u|_{\epsilon(U)}$ takes only the value 1. (Then \tilde{u}^- is constant, equal to a , on $\epsilon(M)$ and we obtain no lower energy bound from the above arguments.) Since u is nonconstant on $\epsilon(M)$, it has to take the zero value somewhere on $\epsilon(M \setminus U)$, which means that the function \tilde{u}^+ takes both values b and 1 on $\epsilon(M)$.

Consider the function

$$\tilde{v} := 1 - \tilde{u}^+$$

and, similarly to the Case 1 argument, obtain

$$E(u|_U) \geq E(\tilde{v}) \geq (1 - b^2)\mu.$$

In either case, we conclude that

$$E(u|_U) \geq s\mu > 0,$$

where, as we recall, $s = \min(a^2, 1 - b^2)$.

By the definition of r_0 , the number of unbounded components of $M \setminus B(p, r)$ is greater than

$$k = \frac{E}{s^2\mu}.$$

The restriction of u to each of these ends is at least $s^2\mu$, which implies that the energy of u is greater than E . This is a contradiction. \square

COROLLARY 21.24. *If $u : M \rightarrow (0, 1)$ is a proper a.e. smooth function of energy $\leq E$ and u is nearly constant on a large ball $B = B(p, r)$, then it nearly equals to 0 or 1 on B . More precisely, the supremum-norm of $u|_B$ or of $(u - 1)|_B$ converges to zero as $\text{Var}(u|_B) \rightarrow 0$.*

We next prove that harmonic functions of bounded energy have small variation on “most” balls in M .

LEMMA 21.25. *Suppose that $h : M \rightarrow [0, 1]$ is a harmonic function of finite energy. Fix $r > 0$ and let $x_i \in M$ be a sequence diverging to infinity, i.e.*

$$\lim_{i \rightarrow \infty} d(x_1, x_i) = \infty.$$

Then

$$\lim_{i \rightarrow \infty} \text{Var}(h|_{B(x_i, r)}) = 0.$$

PROOF. Suppose to the contrary that there exists a sequence (x_i) such that the variation of h on $B(x_i, r)$ does not converge to zero. After passing to a subsequence, we can assume that the balls $B(x_i, r)$ are pairwise disjoint and there exist $\delta > 0$ and points $y_i \in B(x_i, r)$ such that $|h(x_i) - h(y_i)| \geq \delta$ for all i . For each i we pick a geodesic segment $\gamma_i \subset B(x_i, r)$ of length $\leq r$ connecting x_i to y_i .

By the Mean Value Theorem, for each i there exists $z_i \in \gamma_i$ such that

$$|\nabla u(z_i)| \geq \frac{|h(x_i) - h(y_i)|}{r} \geq \frac{\delta}{r}.$$

Hence, we obtain a lower energy-density bound at one point:

$$|\nabla h(z_i)|^2 \geq \frac{\delta^2}{r^2}.$$

We next promote this to a lower energy bound for h . According to Theorem 3.51, there exists a constant L depending only on the geometry bounds of M , such that

$$|\nabla|\nabla h(x)||^2 \leq L$$

at each $x \in M$ where $\nabla h(x) \neq 0$. By appealing to the Mean Value Theorem again, for all $x \in B(z_i, \rho)$, we obtain:

$$|\nabla h(x)|^2 \geq \eta := \frac{\delta^2}{r^2} - L\rho.$$

We fix $\rho > 0$ such that $\eta > 0$ and observe that there exists $V > 0$ for which

$$\text{Vol}(B(z_i, \rho)) \geq V,$$

for all i . Therefore,

$$E(h|_{B(z_i, \rho)}) \geq \eta V.$$

Since the balls $B(z_i, \rho)$ are pairwise disjoint, we conclude that h has infinite energy. A contradiction. \square

We can now prove Part 2 of Theorem 21.5. Recall that $H = H(M)$ is the space of functions $f \in \mathcal{F}$ which are harmonic on M . Consider a sequence $f_n \in H$ whose energy is bounded above by a number E . Since each f_n takes both values 0 and 1 on $\epsilon(M)$, there exist a sequence $x_n \in f_n^{-1}(1/2)$. By applying elements of the group G , we can assume that the points x_n belong to a fixed compact $K \subset M$. After passing to a subsequence, we may assume that $\lim_{n \rightarrow \infty} x_n = x \in K$. In view of Compactness Theorem 3.52, the sequence of functions f_n subconverges uniformly on compacts to a harmonic function f which attains the value $1/2$ at $x \in K$. We have to show that $f \in \mathcal{F}$.

1. By lower semicontinuity of energy (Theorem 3.41), f has energy $\leq E$.

2. Suppose that f is constant on M . Then for each $\delta > 0$ and $r > 0$ there exists n such that

$$\text{Var}(f_n|_{B(x, r)}) < \delta.$$

By taking r sufficiently large and taking into account Corollary 21.24, we conclude that f_n approximately equals to 0 or 1 on $B(x, r)$. This contradicts the assumption that $f_n(x_n) = 1/2$. Therefore, f cannot be constant.

3. Suppose now that f either does not extend continuously to a point $\xi \in \text{Ends}(M)$ or that the extension $f(\xi)$ exists but $f(\xi)$ is different from 0 and 1.

Then there exist $a, b \in (0, 1)$ and a sequence $p_i \in M$ converging to ξ in the topology of \bar{M} such that for all i ,

$$0 < a \leq f(p_i) \leq b < 1.$$

REMARK 21.26. Note that this also includes the case when there are sequences $x_i \rightarrow \xi, y_i \rightarrow \xi$ with

$$\lim_{i \rightarrow \infty} f(x_i) = 0, \quad \lim_{i \rightarrow \infty} f(y_i) = 1.$$

In this case, for large i , f takes the value $\frac{1}{2}$ at a point p_i in a path connecting x_i to y_i and close to ξ in the topology of \bar{M} .

Take $r = r(a, b, E)$ as in Proposition 21.23. Since $E(f) < \infty$, $\text{Var}(f|_{B_r(p_i)})$ converges to 0 as $i \rightarrow \infty$, see Lemma 21.25. Since for each fixed i

$$\lim_{n \rightarrow \infty} f_n|_{B(p_i, r)} = f|_{B(p_i, r)},$$

we conclude that (for large n and i) the function f_n contradicts Proposition 21.23. \square

REMARK 21.27. One could remove the cocompactness assumption by saying that any sequence $u_i \in H$ has a pointed limit living in a pointed Gromov-Hausdorff limit of a sequence (M, x_n) (which will be another bounded geometry manifold with a linear isoperimetric inequality and ubiquitous R -necks).

Thus, it remains to prove Theorem 21.21; the proof occupies the rest of the chapter.

21.5.2. Some coarea estimates. Recall that if $u : M \rightarrow \mathbb{R}$ is a smooth function on a Riemannian manifold M , then for a.e. $t \in \mathbb{R}$, the level set $u^{-1}(\{t\}) = \{u = t\}$ is a smooth hypersurface, and for any measurable function $\phi : M \rightarrow \mathbb{R}$ such that $\phi|\nabla u|$ is integrable, we have the *coarea formula*

$$(21.8) \quad \int_M \phi|\nabla u| = \int_{\mathbb{R}} \left(\int_{\{u=t\}} \phi \right) dt,$$

where the integration $\int_{\{u=t\}} \phi$ is with respect to the Riemannian measure on the hypersurface, see Theorem 3.14.

The two applications of this appearing below are:

$$(21.9) \quad \int_{\{t_1 \leq u \leq t_2\}} |\nabla u|^2 = \int_{t_1}^{t_2} \left(\int_{\{u=t\}} |\nabla u| \right) dt,$$

where we take $\phi = |\nabla u|$ on $\{t_1 \leq u \leq t_2\}$ and zero otherwise, and

$$(21.10) \quad |\{t_1 \leq u \leq t_2\}| = \int_{\{t_1 \leq u \leq t_2\}} 1 = \int_{t_1}^{t_2} \left(\int_{\{u=t\}} \frac{1}{|\nabla u|} \right) dt,$$

where we take $\phi = \frac{1}{|\nabla u|}$ under the assumption that $\nabla u \neq 0$ a.e. on M .

We first combine these in the following general inequality:

LEMMA 21.28. *Suppose that $u : M \rightarrow [t_1, t_2]$ is a smooth function on a compact Riemannian manifold with boundary, such that $A(t) = |\{u = t\}| \geq A_0 > 0$ for a.e. t . Then*

$$(21.11) \quad \int_M |\nabla u|^2 \geq \frac{A_0^2(t_2 - t_1)^2}{|M|}.$$

PROOF. The argument combines (21.9), (21.10), and Jensen's inequality. We decompose M as $M = M_0 \sqcup M_+$, where $M_0 = \{x \in M : \nabla u(x) = 0\}$. Of course,

$$\int_M |\nabla u|^2 = \int_{M_+} |\nabla u|^2.$$

By Sard's Theorem, a.e. $t \in [t_1, t_2]$, is a regular value of u . Furthermore, since u is proper, the set of critical values of u is a closed nowhere dense subset of $[t_1, t_2]$. For a.e. $t \in [t_1, t_2]$ we have

$$\begin{aligned} \int_{\{u=t\}} |\nabla u| &= A(t) \int_{\{u=t\}} |\nabla u| \\ &\geq A(t) \frac{1}{\int_{\{u=t\}} \frac{1}{|\nabla u|}} \quad \text{by Jensen's inequality} \\ &= \frac{A^2(t)}{\int_{\{u=t\}} \frac{1}{|\nabla u|}}, \end{aligned} \tag{21.12}$$

with the equality in the case when $|\nabla u|$ is constant on M .

Since ∇u is non-zero on almost every hypersurface $\{u = t\}$, in the following computation we can consider only non-zero values of ∇u :

$$\begin{aligned} \int_M |\nabla u|^2 &= \int_{t_1}^{t_2} \left(\int_{\{u=t\}} |\nabla u| \right) dt \quad \text{by (21.9)} \\ &\geq \int_{t_1}^{t_2} \frac{A(t)}{\left(\int_{\{u=t\}} \frac{1}{|\nabla u|} \right)} dt \quad \text{by (21.12)} \\ &\geq A_0(t_2 - t_1) \int_{t_1}^{t_2} \frac{dt}{\left(\int_{\{u=t\}} \frac{1}{|\nabla u|} \right)} \\ &\geq A_0(t_2 - t_1) \frac{1}{\int_{t_1}^{t_2} \left(\int_{\{u=t\}} \frac{1}{|\nabla u|} \right) dt} \quad \text{by Jensen's inequality} \\ &\geq \frac{A_0^2(t_2 - t_1)^2}{\int_{t_1}^{t_2} \left(\int_{\{u=t\}} \frac{1}{|\nabla u|} \right) dt} = \frac{A_0^2(t_2 - t_1)^2}{|M_+|} \geq \frac{A_0^2(t_2 - t_1)^2}{|M|} \end{aligned} \tag{21.13}$$

by (21.10). □ □

We note, furthermore, that continuity of u implies that the volume function

$$V(t) = |\{u \geq t\}|$$

is continuous for each t . The fact that the set of critical of u is closed and has zero measure, in conjunction with (21.10), implies that the function $V(t)$ is differentiable a.e. in $[t_1, t_2]$: For every regular value $t \in [t_1, t_2]$ of u we have:

$$\frac{d}{dt} V(t) = \int_{\{u=t\}} \frac{1}{|\nabla u|}.$$

21.5.3. Energy comparison in the case of a linear isoperimetric inequality. Recall that in Section 3.11.1 for each κ we defined X_κ , the unique complete simply-connected surface of the curvature κ . In the case when

$$\kappa = -c < 0,$$

this surface is the upper half-plane \mathbf{U}^2 equipped with the Riemannian metric

$$\frac{dx^2 + dy^2}{c^2 y^2}, \quad c > 0.$$

Consider now a cyclic parabolic subgroup $\Gamma < \text{Isom}(X_\kappa)$ and its quotient surface

$$\hat{N} := X_\kappa/\Gamma.$$

The group Γ preserves horoballs in X_κ with a common center in $\partial_\infty X_\kappa$ fixed by Γ . We will denote projections of these horoballs to \hat{N} by $D_s, s \in \mathbb{R}_+$, with the convention that the length of the boundary of D_s equals s .

EXERCISE 21.29. Show that each $D = D_s$ satisfies the isoperimetric inequality

$$s = \text{length}(\partial D) = c \text{Area}(D).$$

A function $\hat{u} : \hat{N} \rightarrow \mathbb{R}$ is said to be *radial* if it is constant on the circles $\partial D_s, s \in \mathbb{R}_+$.

Consider a Riemannian manifold N with compact boundary ∂N . We will assume that N has Cheeger constant $\geq c > 0$, i.e. N satisfies the linear isoperimetric inequality

$$(21.14) \quad |\partial \Omega| \geq c |\Omega|,$$

where $\Omega \subset N$ is any compact domain with smooth boundary. Our goal is to estimate from below the energy of smooth proper functions $u : N \rightarrow (0, 1]$, satisfying $u^{-1}(\{1\}) = \partial N$. We will do so by comparing the energy of u with that of a suitable proper radial function $\hat{u} : \hat{N} \rightarrow (0, 1)$.

Given u , we define a proper radial function

$$\hat{u} : \hat{N} \rightarrow (0, 1)$$

such that the superlevel sets of \hat{u} have the same volume as the corresponding superlevel sets of u :

$$|u^{-1}([t, 1])| = |\{\hat{u} \geq t\}| = |\{u \geq t\}| \quad \text{for all } t \in (0, 1).$$

Since the function $V(t) = |\{\hat{u} \geq t\}|$ is continuous and differentiable a.e., so is the function \hat{u} . For $t \in (0, 1)$, define $\hat{V}(t) = |\{t \leq \hat{u} < 1\}|$, $\hat{A}(t) = |\{\hat{u} = t\}|$ and $A(t) = |\{u = t\}|$. As we noted above,

$$\hat{A}(t) = c \hat{V}(t)$$

for each t .

LEMMA 21.30 (Energy comparison lemma). *Suppose that for some $T \in (0, 1]$, we have*

$$V(T) \geq \frac{2}{c} A(1) = \frac{2}{c} |\partial N|.$$

Then

$$\int_{\{0 < u \leq T\}} |\nabla u|^2 \geq \frac{1}{4} \int_{\{0 < \hat{u} \leq T\}} |\nabla \hat{u}|^2.$$




FIGURE 21.1. Functions u and \hat{u} .

PROOF. Since $V(t) = \hat{V}(t)$, differentiating

$$V(t) = \int_t^1 \int_{u=\tau} \frac{1}{|\nabla u|} d\tau$$

with respect to t , we get that (for a.e. $t \in [0, 1]$)

$$(21.15) \quad \int_{\{u=t\}} \frac{1}{|\nabla u|} = \int_{\{\hat{u}=t\}} \frac{1}{|\nabla \hat{u}|}.$$

For all $t \leq T$, in view of the isoperimetric inequality (21.14), we have

$$|\partial\{u \geq t\}| = |\partial N| + A(t) \geq cV(t),$$

while

$$|\partial N| \leq \frac{c}{2}V(T) \leq \frac{c}{2}V(t).$$

By combining these inequalities, we obtain:

$$(21.16) \quad A(t) \geq cV(t) - |\partial N| \geq cV(t) - \frac{c}{2}V(t) \geq \frac{c}{2}V(t) = \frac{c}{2}\hat{V}(t) = \frac{\hat{A}(t)}{2}.$$

Now, for each regular value t of u ,

$$\int_{\{u=t\}} |\nabla u| \geq \frac{A^2(t)}{\int_{\{u=t\}} \frac{1}{|\nabla u|}} \quad \text{see (21.12)}$$

$$\begin{aligned}
&\geq \frac{\hat{A}^2(t)}{4 \int_{\{\hat{u}=t\}} \frac{1}{|\nabla \hat{u}|}} \quad \text{by (21.15) and (21.16)} \\
(21.17) \quad &= \frac{c^2}{2} \int_{\{\hat{u}=t\}} |\nabla \hat{u}|
\end{aligned}$$

because $|\nabla \hat{u}|$ is constant on $\{u = t\}$ and so the equality case of (21.12) applies.

The lemma now follows from (21.9) and (21.17). \square

21.5.4. Proof of positivity of the energy gap. Consider $v \in \mathcal{F}$ with the set of regular values $R \subset (0, 1)$. Every regular level set of v defines a non-trivial homology class in M , since it separates positive and negative ends of M . As we saw in Theorem 3.39,

$$\inf_{\tau \in R} |\{v = \tau\}| \geq A_0 > 0,$$

for a certain constant A_0 .

REMARK 21.31. This is another place where the proof simplifies considerably in the case when M is a surface: Then every homologically non-trivial cycle in M has length $\geq A_0$, the injectivity radius of M .

Choose a regular value $t_1 \in (0, 1)$ of v where $A(t), t \in R$, almost attains its infimum, i.e.

$$A(t_1) \geq \inf_{\tau \in R} |\{v = \tau\}| \geq A(t_1)/2.$$

We may assume that $t_1 \geq \frac{1}{2}$ (otherwise, we use the function $1 - v$ instead) and we focus attention on the codimension 0 submanifold with boundary $N \subset M$ given by the sublevel set $\{v \leq t_1\}$. Replacing v with $u = \frac{1}{t_1}v$, we get a proper function $u : N \rightarrow (0, 1]$ which is 1 on ∂N , such that all the level sets $\{u = t\}$ have area at least $\frac{1}{2} \text{Area}(\partial N)$. Clearly,

$$E(v) \geq t_1^2 E(u) \geq \frac{E(u)}{4}.$$

Thus, it suffices to get a lower bound on $E(u)$. We will see below that

$$E(u) \geq \frac{cA_0}{32}.$$

Since the volume $V(t) = |\{u \geq t\}|$ is a continuous function of t which vanishes at $t = 1$ and satisfies

$$\lim_{t \rightarrow 0} |\{u \geq t\}| = \infty,$$

there exists a superlevel set $\{u \geq T\} \subset N$ whose volume equals $\frac{2}{c}|\partial N|$, where c is the Cheeger constant of M .

Case 1: $T \leq \frac{1}{2}$. Applying Lemma 21.28, we get

$$\begin{aligned}
E(u) &\geq \int_{\{T \leq u \leq 1\}} |\nabla u|^2 \geq \frac{(1-T)^2 A_0^2}{\frac{2}{c}|\partial N|}, \\
&\geq \frac{\left(\frac{(1-T)|\partial N|}{2}\right)^2}{\frac{2}{c}|\partial N|} = \frac{(1-T)^2}{8} c |\partial N| \geq \frac{(1-T)^2}{8} c A_0 \geq \frac{cA_0}{32}.
\end{aligned}$$

Therefore, we obtain a lower energy bound for u (and, hence, v) in this case.

Case 2: $T \leq \frac{1}{2}$. Lemma 21.30 shows that the energy of u is at least

$$\frac{1}{4} \int_{\{\hat{u} \geq T\}} |\nabla \hat{u}|^2,$$

where \hat{u} is the radial comparison function defined on the surface $\hat{N} = X_\kappa/\Gamma$.

By our choice of T ,

$$\hat{V}(T) = |\{T \leq \hat{u} \leq 1\}| = |\{T \leq u \leq 1\}| = \frac{2}{c} |\partial N| \geq \frac{2}{c} A_0.$$

LEMMA 21.32.

$$E(\hat{u}|_{\{\hat{u} \leq T\}}) = \int_{0 < \hat{u} \leq T} |\nabla \hat{u}|^2 \geq \frac{cA_0}{2}.$$

PROOF. We will identify the subsurface $\{0 < \hat{u} \leq T\} \subset \hat{N}$ with the rectangle

$$Q = \{(x, y) : 0 < y \leq 1, 0 \leq x \leq a\},$$

whose vertical sides are identified via the translation $(x, y) \mapsto (x + a, y)$. Since energy is a conformal invariant, it suffices to do the computation of energy with respect to the Euclidean metric. Accordingly, below, $|\nabla \hat{u}|$ is the Euclidean norm of the Euclidean gradient. Since the function \hat{u} is radial, $\hat{u}(x, y) = f(y)$ and, hence,

$$|\nabla \hat{u}|^2 = f'^2.$$

We obtain:

$$(21.18) \quad E(\hat{u}|_{\{0 < \hat{u} \leq T\}}) = a \int_0^1 (f')^2 dy \quad (\text{by Cauchy's inequality}) \\ \geq a \left(\int_0^1 f' dy \right)^2 = aT^2 \geq \frac{a}{4} \quad (\text{since } T \geq \frac{1}{2}).$$

In order to estimate the number a , note that the area $\hat{V}(T)$ equals the area of the strip

$$P = \{(x, y) : 1 < y < \infty, 0 \leq x \leq a\} \subset X_\kappa.$$

The latter equals $\frac{a}{c^2}$ and, hence,

$$\frac{a}{c^2} = \hat{V}(T) \geq \frac{2}{c} A_0,$$

$$(21.19) \quad a \geq 2cA_0.$$

Combining the inequalities (21.18) and (21.19), we obtain:

$$E(\hat{u}|_{\{\hat{u} \leq T\}}) = \int_{0 < \hat{u} \leq T} |\nabla \hat{u}|^2 \geq \frac{a}{4} \geq \frac{cA_0}{2}. \quad \square$$

Lemmata 21.30 and 21.32 imply that, in the Case 2:

$$E(u) \geq \frac{1}{4} E(\hat{u}|_{\{\hat{u} \leq T\}}) \geq \frac{cA_0}{8},$$

which is a higher bound than in the Case 1.

Therefore, the energy of the function $v : M \rightarrow (0, 1)$ is at least $\frac{1}{4} E(u) \geq 2^{-7} cA_0$. This completes the proof of Theorem 21.21. \square

Quasiconformal mappings

Quasiconformal and quasisymmetric maps play prominent role in geometric analysis and Geometric Group Theory. (The classes of quasiconformal and quasisymmetric maps coincide in the case of maps $\mathbb{S}^n \rightarrow \mathbb{S}^n$, $n \geq 2$.) Their importance in Geometric Group Theory comes from the fundamental fact that quasisymmetric maps appear as *boundary extensions* of quasiisometries between Gromov–hyperbolic spaces: Each quasiisometry

$$f : X \rightarrow Y$$

extends to a unique quasisymmetric homeomorphism

$$f_\infty : \partial_\infty X \rightarrow \partial_\infty Y$$

of the Gromov boundaries of the spaces X and Y . Conversely, each quasisymmetric homeomorphism $\partial_\infty X \rightarrow \partial_\infty Y$ extends to a quasiisometry $X \rightarrow Y$ and any two such quasiisometric extensions are within bounded distance from each other. In the case when $X = Y = \mathbb{H}^{n+1}$, the extensions f_∞ are the *classical quasiconformal maps*. This remarkable interaction between quasiconformal analysis and hyperbolic geometry is somewhat akin to fruitful relation between complex analysis and hyperbolic geometry.

The intuition of classical quasiconformal maps comes from the theory of holomorphic functions of one complex variable: Conformal maps are characterized (locally) by the property that they send infinitesimal circles to infinitesimal circles. Accordingly, *classical quasiconformal maps* are defined by the condition that they send infinitesimal spheres to infinitesimal ellipsoids of uniformly bounded eccentricity.

We refer the reader to the books [Res89], [Vuo88], [Väi71] and [IM01] for the detailed discussion of classical quasiconformal maps and to [HK95], [HK98] and [Hei01] for the treatment of quasiconformal and quasisymmetric maps between more general metric spaces.

In this book we will be using only classical quasiconformal maps, whose basic analytical and geometric properties will be established in this chapter. The main applications of classical quasiconformal maps in the book are Mostow Rigidity Theorem (for lattices in the isometry groups $SO(n+1, 1)$ of \mathbb{H}^{n+1}), Tukia’s QI Rigidity Theorem for the class uniform lattices in $SO(n+1, 1)$ and Schwartz’ QI Rigidity Theorem for non-uniform lattices in $SO(n+1, 1)$. These theorems will be proven in Chapters 23 and 24.

HISTORICAL REMARK 22.1. Quasiconformal mappings between open subsets of the complex plane were introduced in the 1920s by Herbert Grötzsch [Gro28] as a generalization of conformal mappings. Quasiconformal mappings in higher dimensions were defined by Mikhail Lavrentiev in the 1930s as a tool in applied mathematics (hydrodynamics), [Lav38]. The discovery of the relation between

quasiisometries of hyperbolic spaces and quasiconformal mappings was made independently by Vadim Efremovich and Ekaterina Tihomirova [ET64] and George Mostow [Mos65] in the 1960-s.

22.1. Linear algebra and eccentricity of ellipsoids

Suppose that $M \in GL(n, \mathbb{R})$ is an invertible linear transformation. We would like to measure the deviation of M from being a conformal linear transformation, i.e. from being an element of $\mathbb{R}_+ \cdot O(n)$. Geometrically speaking, we are interested in measuring the deviation of the ellipsoid $E = M(\mathbb{D}) \subset \mathbb{R}^n$ from a round ball, where $\mathbb{D} = \mathbb{D}^n$ is the unit ball in \mathbb{R}^n .

In the case $n = 2$, there is only one way to define such a measurement, namely, *eccentricity* of the ellipsoid E , which is the ratio of major to minor axes of E . In higher dimensions, there are several invariants which are useful in different situations. This reflects the simple fact that the matrix M has n singular values, while the invariants we are looking for are single real numbers.

Recall that every invertible $n \times n$ matrix M has the *singular value decomposition* (see Theorem 2.68)

$$M = UDV = U \text{Diag}(\lambda_1, \dots, \lambda_n)V,$$

where the (positive) diagonal entries $\lambda_1 \leq \dots \leq \lambda_n$ of the diagonal matrix $D = \text{Diag}(\lambda_1, \dots, \lambda_n)$ are the *singular values* of M . Here U, V are orthogonal matrices. Equivalently, if we symmetrize M (i.e. set $A = M^T A$), then the numbers λ_i are square roots of the eigenvalues of A . Geometrically speaking, the singular values λ_i are the half-lengths of the axes of the ellipsoid $E = M(\mathbb{D})$.

We define the following *distortion quantities* for the matrix M :

- Linear dilatation:

$$H(M) := \frac{\lambda_n}{\lambda_1} = \|M\| \cdot \|M^{-1}\|,$$

where $\|M\|$ is the operator norm of a matrix M :

$$\max_{\mathbf{v} \in \mathbb{R}^n \setminus \{0\}} \frac{|M\mathbf{v}|}{|\mathbf{v}|}.$$

Thus, $H(M) = \epsilon(E)$ is the *eccentricity* of the ellipsoid E , the ratio of lengths of major and minor axes of E . This is the invariant that we will be using most of the time.

- Inner dilatation:

$$H_I(M) := \frac{\lambda_1 \dots \lambda_n}{\lambda_1^n} = |\det(M)| \cdot \|M^{-1}\|^n.$$

- Outer dilatation:

$$H_O(M) := \frac{\lambda_n^n}{\lambda_1 \dots \lambda_n} = \|M\|^n |\det(M)|^{-1}.$$

- Maximal dilatation:

$$K(M) := \max(H_I(M), H_O(M)).$$

Thus, geometrically speaking, the inner and outer dilatations compute volume ratios of E and inscribed/circumscribed balls, while the linear dilatation compares the radii of inscribed/circumscribed balls. Note that all four dilatations agree for $n = 2$.

EXERCISE 22.2. M is conformal $\iff H(M) = 1 \iff H_I(M) = 1 \iff H_O(M) = 1 \iff K(M) = 1$.

EXERCISE 22.3. Logarithms of linear and maximal dilatations are comparable:

$$(H(M))^{n/2} \leq K(M) \leq (H(M))^{n-1}.$$

Hint: It suffices to consider the case when M is a diagonal matrix $Diag(\lambda_1, \dots, \lambda_n)$.

EXERCISE 22.4. 1. $H(M) = H(M^{-1})$ and $H(M_1 \cdot M_2) \leq H(M_1) \cdot H(M_2)$.

2. $K(M) = K(M^{-1})$ and $K(M_1 \cdot M_2) \leq K(M_1) \cdot K(M_2)$.

Hint: Use geometric interpretation of the four dilatations.

22.2. Quasisymmetric maps

Our next goal is to generalize the dilatation constants of linear maps to non-linear maps. The linear dilatation is the easiest to generalize, since it deals only with distances. Recall the geometric meaning of the linear dilatation $H(M)$: If E is the image of the round ball \mathbb{D} , then $H(M)$ is the ratio of the “outer radius” of E by its “inner radius.” Such ratio makes sense not only for ellipsoids but also for arbitrary (closed) topological balls $D \subset \mathbb{R}^n$, where we have chosen a “center” , a point x' in the interior of D : Then we have two real numbers r and R , such that $\bar{B}(x', r)$ is the largest metric ball (centered at x') contained in D and $\bar{B}(x', R)$ is the smallest metric ball containing D . Then the numbers r and R can be regarded as the inner and outer radii of D . In other words,

$$\frac{R}{r} = \max \frac{|y' - x'|}{|z' - x'|},$$

where the maximum is taken over all points $y', z' \in \partial D$. This ratio is the “eccentricity” of the topological ball $D \subset \mathbb{R}^n$. The idea then is to consider homeomorphisms f which send round balls $B(x, \rho)$ to topological balls of uniformly bounded eccentricity with respect to the “centers” $x' = f(x)$.

This leads to

DEFINITION 22.5. A homeomorphism $f : \Omega \rightarrow \Omega'$ between two domains in \mathbb{R}^n is *c-weakly quasisymmetric* if

$$(22.1) \quad \frac{|f(x) - f(y)|}{|f(x) - f(z)|} \leq c$$

for all $x, y, z \in \Omega$, such that $|x - y| = |y - z| > 0$.

Note that we do not assume that f preserves orientation. We will be mostly interested in the case $\Omega = \Omega' = \mathbb{R}^n$.

The name *quasisymmetric* comes from the case $n = 1$ (and quasisymmetric maps were originally introduced only for $n = 1$ by Lars Ahlfors and Arne Beurling [AB56]). Namely, a homeomorphism $f : \mathbb{R} \rightarrow \mathbb{R}$, satisfying $f(0) = 0$, is *symmetric* at the origin if it sends any pair of points symmetric about 0 to points symmetric about 0, i.e. these homeomorphisms are *odd* functions: $f(-y) = -f(y)$. In the case of *c-weakly quasisymmetric* maps, the exact symmetry is lost, but is replaced by a uniform bound on the ratio of absolute values.

EXERCISE 22.6. Show that 1-weakly quasisymmetric homeomorphisms $f : \mathbb{R} \rightarrow \mathbb{R}$ are compositions of dilations and isometries of \mathbb{R} .

It turns out that there is a slightly stronger condition, which is a bit easier to work with and which generalizes naturally to metric spaces other than \mathbb{R}^n :

DEFINITION 22.7. Fix a homeomorphism $\eta : [0, \infty) \rightarrow [0, \infty)$. A homeomorphism $f : \Omega \subset \mathbb{R}^n \rightarrow \Omega' \subset \mathbb{R}^n$ is called η -quasisymmetric if for all $x, y, z \in \mathbb{R}^n$ we have

$$(22.2) \quad \frac{|f(x) - f(y)|}{|f(x) - f(z)|} \leq \eta \left(\frac{|x - y|}{|x - z|} \right).$$

Thus, if we take $c = \eta(1)$, then every η -quasisymmetric map is also c -weakly quasisymmetric. It is a non-trivial theorem (see e.g. [Hei01]) that for $\Omega = \Omega' = \mathbb{R}^n$, the two concepts are equivalent.

EXERCISE 22.8. Show that:

1. Every invertible affine transformation $L : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is η -quasisymmetric with $\eta(t) = H(L)t$.
2. L -Lipschitz homeomorphisms are η -quasisymmetric with $\eta(t) = L^2t$.

As in the case of quasiisometries, we will say that a homeomorphism is (weakly) quasisymmetric if it is η -quasisymmetric (respectively c -weakly quasisymmetric) for some η or $c < \infty$.

The following exercise requires nothing but the definition of quasisymmetry:

EXERCISE 22.9. Show that the composition of quasisymmetric maps is again quasisymmetric. Show that the inverse of a quasisymmetric map is also quasisymmetric.

Recall that we think of \mathbb{S}^n as the 1-point compactification of \mathbb{R}^n . Accordingly, we can define quasisymmetric homeomorphisms of $\mathbb{R}^n \cup \{\infty\}$ as extensions of quasisymmetric homeomorphisms $\mathbb{R}^n \rightarrow \mathbb{R}^n$. The drawback of this definition of quasisymmetric maps is that we are restricted to maps sending the point ∞ to itself. In particular, we cannot apply this definition to Moebius transformations.

DEFINITION 22.10. A homeomorphism of \mathbb{S}^n is called *quasimoebius* if it is a composition of a Moebius transformation with a quasisymmetric map.

Recall (Theorem 4.4) that Moebius transformations of \mathbb{S}^n can be characterized by the property that they preserve the cross-ratios

$$[x, y, z, w] := \frac{|x - y| \cdot |z - w|}{|y - z| \cdot |w - x|}, \quad x, y, z, w \in \mathbb{S}^n.$$

Similarly, one can prove (see [Väi85]) that a homeomorphism f of \mathbb{S}^n is quasimoebius if and only if there exists a homeomorphism $\kappa : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that

$$[f(x), f(y), f(z), f(w)] \leq \kappa([x, y, z, w])$$

for all $x, y, z, w \in \mathbb{S}^n$. While the notion of quasimoebius maps is esthetically appealing, we will be working mostly with quasisymmetric and *quasiconformal maps*, which will be introduced in the next section.

22.3. Quasiconformal maps

The idea of quasiconformality is very natural: We take the definition of weakly quasiconformal maps *via* the ratio (22.1) and then take the limit in this ratio as $\rho = |x - y| = |y - z| \rightarrow 0$.

For a homeomorphism $f : \Omega \subset \mathbb{R}^n \rightarrow \Omega' \subset \mathbb{R}^n$ between two domains in \mathbb{R}^n and $x \in \Omega$ we define the quantity

$$(22.3) \quad H_x(f) := \limsup_{\rho \rightarrow 0} \left(\max_{y,z} \frac{|f(x) - f(y)|}{|f(x) - f(z)|} \right),$$

where, for each $\rho > 0$, the maximum is taken over all $y, z \in \Omega$ with $\rho = |x - y| = |x - z|$. For instance, if f is c -weakly quasisymmetric, then $H_x(f) \leq c$ for every $x \in \Omega$.

DEFINITION 22.11. A homeomorphism $f : \Omega \rightarrow \Omega'$ is called *quasiconformal* if

$$\sup_{x \in \Omega} H_x(f)$$

finite.

The function $H_x(f)$ is called *the (linear) dilatation function* of f ; a quasiconformal map f is said to have *dilatation* $\leq H$ if

$$H(f) := \text{ess sup}_{x \in \Omega} H_x(f) \leq H.$$

Note that the essential supremum is the L^∞ -norm, thus, it ignores subsets of measure zero. We will see the reason for this discrepancy between the definition of quasiconformality (where $H_x(f)$ is required to be uniformly bounded) and the definition of the dilatation $H(f)$, in the next section.

Thus, the intuitive meaning of quasiconformality is that *quasiconformal maps send infinitesimal spheres to infinitesimal ellipsoids of uniformly bounded eccentricity*.

EXERCISE 22.12. Let $f : \mathbb{S}^n \rightarrow \mathbb{S}^n$ be a Moebius transformation, $p = f^{-1}(\infty)$. Then the restriction

$$f|_{\mathbb{R}^n \setminus \{p\}}$$

is 1-quasiconformal, i.e. conformal. Hint: It suffices to verify conformality only for the inversion in the unit sphere.

Note that here and in what follows *we do not assume that conformal maps preserve orientation*. For instance, in this terminology, complex conjugation is a conformal map $\mathbb{C} \rightarrow \mathbb{C}$.

EXERCISE 22.13. 1. Suppose that $f : \Omega \rightarrow \Omega'$ is a C^1 -diffeomorphism such that $\|D_x(f)\|$ is uniformly bounded above and $|J_x(f)|$ is uniformly bounded below. Show, using the definition of differentiability, that f is quasiconformal. Namely, verify that $H_x(f) = H(D_x(f))$ for every $x \in \Omega$.

2. Show that every C^1 -diffeomorphism $\mathbb{S}^n \rightarrow \mathbb{S}^n$ is quasiconformal.

22.4. Analytical properties of quasiconformal mappings

In this section we list certain analytical properties of quasiconformal (quasisymmetric) mappings used in the book. We will prove most of them with two notable exceptions: Gehring's version of the Liouville's theorem and Tukia's Strong Convergence Property. (The Measurable Riemann Mapping Theorem in the next chapter is another exception.) Proving these theorems would go well beyond the scope of this book.

22.4.1. Some notion and results from real analysis. For a subset $E \subset \mathbb{R}^n$ we let $\text{mes}(E)$ denote the n -dimensional Lebesgue measure of E . In what follows, Ω is an open subset in \mathbb{R}^n .

22.4.1.A. *Derivatives of measures.* Let μ be a measure on Ω of the Lebesgue class, i.e. μ -measurable sets are in the Borel σ -algebra. The *derivative* of μ at $x \in \Omega$, denoted $\mu'(x)$, is defined as

$$\mu'(x) := \limsup \frac{\mu(B)}{\text{mes}(B)},$$

where the limit is taken over all balls B containing x whose radii tend to zero. The key fact that we will need is the following theorem (see e.g. [Fol199, Theorem 3.22]):

THEOREM 22.14 (Lebesgue–Radon–Nikodym differentiation theorem). *The function $\mu'(x)$ is Lebesgue-measurable and is finite a.e. in Ω . Furthermore, $\mu'(x)$ is the Radon–Nikodym derivative of the component of μ which is absolutely continuous with respect to the Lebesgue measure.*

For a continuous map $f : \Omega \rightarrow \mathbb{R}^m$ we define the pull-back measure $\mu = \mu_f$ by

$$\mu(E) := \text{mes}(f(E)).$$

22.4.1.B. *Approximate continuity.* A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is called *approximately continuous* at a point $x \in \mathbb{R}^n$ if for every $\epsilon > 0$

$$(22.4) \quad \lim_{r \rightarrow 0} \frac{\text{mes}(\{y \in B(x, r) : |f(x) - f(y)| > \epsilon\})}{\text{mes}(B(x, r))} = 0.$$

(Here, as before, mes denotes the Lebesgue measure.) In other words, as we “zoom into” the point x , “most” points $y \in B(x, r)$, have value $f(y)$ close to $f(x)$, i.e. the rescaled functions $f_r(x) := f(rx)$ converge (as $r \rightarrow 0$) in measure to a constant function.

LEMMA 22.15. *Every L_∞ -function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is approximately continuous at almost every point.*

PROOF. The proof is an application of *The Lebesgue Density theorem* (see e.g. [SS05, p. 106]): For every measurable function h on \mathbb{R}^n and almost every x ,

$$\lim_{r \rightarrow 0} \frac{1}{\text{mes}(B_r)} \int_{B_r} |h(y) - h(x)| dy = 0.$$

Here and below, we set $B_r = B(x, r)$.

Fix $\epsilon > 0$ and let $E_r \subset B_r$ denote the subset consisting of $y \in B_r$ with

$$|f(y) - f(x)| > \epsilon.$$

If the equality (22.4) fails, then

$$\lim_{r \rightarrow 0} \frac{\text{mes}(E_r)}{\text{mes}(B_r)} > 0.$$

By the definition of the subset E_r we have the inequality:

$$\frac{1}{\text{mes}(B_r)} \int_{B_r} |f(y) - f(x)| dy \geq \epsilon \frac{\text{mes}(E_r)}{\text{mes}(B_r)}.$$

Since

$$\lim_{r \rightarrow 0} \frac{\text{mes}(E_r)}{\text{mes}(B_r)} > 0,$$

we conclude that

$$\liminf_{r \rightarrow 0} \frac{1}{\text{mes}(B_r)} \int_{B_r} |f(y) - f(x)| dy \neq 0,$$

contradicting the Lebesgue Density Theorem. \square

22.4.1.C. *Rademacher–Stepanov Theorem.* Rademacher–Stepanov theorem is a strengthening of Rademacher’s theorem (Theorem 2.28); we will need it in order to prove differentiability a.e. of quasiconformal mappings, among other things.

Recall that a map $f : \Omega \rightarrow \mathbb{R}^m$ is called *differentiable* at $x \in \Omega$ with the derivative $D_x f$ at x equal to the matrix A , if

$$\lim_{h \rightarrow 0} \frac{|f(x+h) - f(x) - Ah|}{|h|} = 0.$$

It follows directly from the definition that, for $n = m$, at every point x of differentiability of f , the measure derivative of μ_f equals the absolute value of the Jacobian of f :

$$\mu'_f(x) = |\det(A)| = |J_f(x)|.$$

The other key result that we will use is:

THEOREM 22.16 (Rademacher and Stepanov, see e.g. Theorem 3.4 in [Hei05]). *Let $f : \Omega \rightarrow \mathbb{R}^m$. For every $x \in \Omega$ define*

$$\|D_x^+(f)\| := \limsup_{h \rightarrow 0} \frac{|f(x+h) - f(x)|}{|h|}.$$

Then f is differentiable a.e. in the set $E = \{x \in \Omega : \|D_x^+(f)\| < \infty\}$.

A special case of this theorem is Rademacher’s theorem (Theorem 2.28), since for L -Lipschitz maps

$$\|D_x^+(f)\| \leq L.$$

22.4.1.D. *Absolutely continuous functions.* Informally, absolutely continuous functions are those which map sets of small measure to sets of small measure. More precisely, suppose that f is a real-valued function defined on an interval I in \mathbb{R} . The function f is called *absolutely continuous* (AC) if for every $\epsilon > 0$ there exists $\delta > 0$ such that: For every collection of pairwise disjoint subintervals

$$(a_1, b_1), (a_2, b_2), \dots, (a_k, b_k) \subset I$$

with

$$\sum_{i=1}^k |b_i - a_i| < \delta,$$

we have

$$(22.5) \quad \sum_{i=1}^k |f(b_i) - f(a_i)| < \epsilon.$$

In particular, AC functions send subsets of zero measure to subsets of zero measure (this is nearly clear from the definition, the reader can find the details in [Fol99, Proposition 3.32]).

Clearly, every absolutely continuous function is uniformly continuous, but the converse is false. For instance, the Cantor function, defined using a Cantor set C of zero measure, sends C to the unit interval.

AC functions are characterized by the fact that the Fundamental Theorem of Calculus holds for them:

THEOREM 22.17. *The following properties are equivalent for a function $f : [a, b] \rightarrow \mathbb{R}$:*

1. f is absolutely continuous.
2. There exists a function $h \in L^1([a, b])$ such that

$$f(x) = \int_a^x h(t)dt$$

for all $x \in [a, b]$.

3. The function f is differentiable almost everywhere in $[a, b]$ with measurable derivative $f'(x)$, such that

$$f(x) = \int_a^x f'(t)dt.$$

for all $x \in [a, b]$.

We refer the reader to [Fol199, Theorem 3.35] or [SS05, Theorem 3.11] for a proof.

The notion of absolutely continuous function generalizes readily to functions of one variable with values in \mathbb{R}^n where in the formula (22.5) instead of the absolute value we use the norm in \mathbb{R}^n .

EXERCISE 22.18. Show that a function $f = (f_1, \dots, f_m) : I \rightarrow \mathbb{R}^m$ is absolutely continuous if and only if each component f_i of f is absolutely continuous.

We will need a sufficient condition for absolute continuity of vector-valued functions:

LEMMA 22.19. *Suppose that $f : I \rightarrow \mathbb{R}^m$ is a continuous function for which there exists a constant C such that for every measurable subset $E \subset I$ we have*

$$\mu_1(f(E)) \leq C \text{mes}(E).$$

Then f is absolutely continuous.

PROOF. It follows immediately from the definition of the 1-dimensional (normalized) Hausdorff measure μ_1 that each component f_i of the function f also satisfies

$$\text{mes}(f_i(E)) \leq C \text{mes}(E).$$

Absolute continuity of f_i follows by taking in the definition of absolute continuity

$$\delta = \epsilon/C. \quad \square$$

We will use the following generalization of the notion of absolute continuity to the case of functions of several variables:

DEFINITION 22.20. A map $f : \Omega \rightarrow \mathbb{R}^m$ defined on an open subset $\Omega \subset \mathbb{R}^n$, is called ACL, *absolutely continuous on lines*, if the restriction of f to almost every coordinate line segment in Ω is an absolutely continuous function of one variable. Here and in what follows, a coordinate line segment is a compact straight line segment parallel to one of the coordinate axes in \mathbb{R}^n .

22.4.2. Differentiability properties of quasiconformal mappings. We now return to quasiconformal maps. Recall that the dilatation $H_x(f)$ of a homeomorphism f at a point x is defined as

$$H_x(f) := \limsup_{\rho \rightarrow 0} \frac{R(x, \rho)}{r(x, \rho)},$$

where

$$(22.6) \quad R(x, \rho) = \max_{|h|=\rho} |f(x+h) - f(x)|, \quad r(x, \rho) = \min_{|h|=\rho} |f(x+h) - f(x)|.$$

22.4.2.A. *Differentiability a.e. of quasiconformal homeomorphisms.*

THEOREM 22.21 (F. Gehring, see [Väi71]). *Every quasiconformal map $f : \Omega \rightarrow \mathbb{R}^n$ is differentiable a.e. in Ω and*

$$\|D_x f\| \leq H_x(f) |J_x(f)|^{1/n}$$

for a.e. x in Ω .

PROOF. By the definition of $|D_x^+(f)|$ and $H_x(f)$:

$$\|D_x^+(f)\| = \limsup_{\rho \rightarrow 0} \frac{R(x, \rho)}{\rho} = H_x(f) \limsup_{\rho \rightarrow 0} \frac{r(x, \rho)}{\rho}.$$

Notice that for $r = r(x, \rho)$, $B(f(x), r) \subset f(B(x, \rho))$, which implies that

$$\omega_n r^n = \text{mes}(B(f(x), r)) \leq \text{mes}(f(B(x, \rho))),$$

where ω_n is the volume of the unit n -ball. Therefore,

$$\frac{\text{mes}(f(B(x, \rho)))}{\text{mes}(B(x, \rho))} \geq \frac{r^n}{\rho^n}$$

and, thus,

$$\mu'_f(x) = \limsup_{\rho \rightarrow 0} \frac{\text{mes}(f(B(x, \rho)))}{\text{mes}(B(x, \rho))} \geq \limsup_{\rho \rightarrow 0} \frac{r^n}{\rho^n} = \left(\frac{1}{H_x(f)} \|D_x^+(f)\| \right)^n.$$

It follows that

$$\|D_x^+(f)\| \leq H_x(f) (\mu'_f(x))^{1/n}.$$

The right-hand side of this inequality is finite for a.e. x (by Theorem 22.14). Thus, f is differentiable at a.e. x by Rademacher-Stepanov theorem. We also obtain (for a.e. $x \in \Omega$)

$$\|D_x(f)\| = \|D_x^+(f)\| \leq H_x(f) (\mu'_f(x))^{1/n} = H_x(f) |J_x(f)|^{1/n} \quad \square$$

22.4.2.B. *ACL property of quasiconformal mappings.* Gehring's differentiability theorem is strengthened as follows:

THEOREM 22.22 (F. Gehring, J. Väisälä, see [Väi71] and [Mos73]). *For $n \geq 2$, quasiconformal maps $f : \Omega \rightarrow \mathbb{R}^n$ belong to the Sobolev class $W_{loc}^{1,n}$, i.e. their first partial distributional derivatives are locally in $L^n(\Omega)$. This, in particular, implies that quasiconformal maps are ACL.*

We will prove a slightly weaker, but sufficient for our purposes, version of this important theorem:

THEOREM 22.23. *Every weakly quasisymmetric homeomorphism $f : \Omega \rightarrow \Omega'$, between domains in \mathbb{R}^n , $n \geq 2$, is ACL. In particular, the matrix of partial derivatives $D_x(f)$ is a measurable matrix-valued function on Ω .*

PROOF. Our proof closely follows the one given by Mostow in [Mos73]. Until the very end of the proof, we will not be using the assumption $n \geq 2$. Let κ be the quasisymmetry constant of f . Then:

$$r(p, t) \leq R(p, t) \leq \kappa r(p, t)$$

for all $t > 0$ and all $p \in \Omega$, where $R(p, t)$ and $r(p, t)$ are defined by the equations (22.6). Let $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^{n-1}$ be the orthogonal projection, where \mathbb{R}^{n-1} is one of the coordinate hyperplanes in \mathbb{R}^n , defined by the equation $x_i = 0$. Fix a bounded open subset $Q \subset \Omega$ and for $y \in \mathbb{R}^{n-1}$ define

$$Q(y) := Q \cap \pi^{-1}(y).$$

For $t > 0$ set

$$Q(y, t) := \mathcal{N}_t(Q(y)) \cap Q.$$

Clearly, the subset $\pi^{-1}(B(y, t)) \cap Q$ contains $Q(y, t)$.

LEMMA 22.24. *For almost all $y \in \mathbb{R}^{n-1}$, the limit*

$$\tau(y) := \limsup_{t \rightarrow 0} \frac{\text{mes}(f(Q(y, t)))}{t^{n-1}}$$

is finite.

PROOF. For each measurable subset $E \subset \mathbb{R}^{n-1}$, define

$$\phi(E) = \text{mes}(f(\pi^{-1}(E) \cap Q)).$$

Then ϕ is a measure of the Lebesgue class on \mathbb{R}^{n-1} and by Theorem 22.14 we obtain that for almost all $y \in \mathbb{R}^{n-1}$

$$\limsup_{t \rightarrow 0} \frac{\phi(B(y, t))}{\text{mes}(B(y, t))} < \infty.$$

Clearly,

$$\text{mes}(f(Q(y, t))) \leq \phi(B(y, t)).$$

On the other hand, t^{n-1} is (up to a constant factor) equal to $\text{mes}(B(y, t))$. The lemma follows. \square

We now claim that for every y such that $\tau(y) < \infty$, the function f is absolutely continuous on the intersection of the line $\pi^{-1}(y)$ with the set Q . Recall that μ_1 denotes the 1-dimensional Hausdorff measure of subsets in \mathbb{R}^n . The following lemma is the key to absolute continuity:

LEMMA 22.25. *There exists a constant C , depending only on n and κ , such that for every y satisfying $\tau(y) < \infty$ and for every compact subset $E \subset Q(y)$, we have*

$$\mu_1(f(E))^n \leq C\tau(y)(\mu_1(E))^{n-1}.$$

PROOF. It suffices to prove the lemma in the case when E is a closed interval. Let $L = \mu_1(E)$ denote the length of the interval E . For each compact $K \subset \mathbb{R}^n$ define

$$\Lambda(K, a) = \inf_{\mathcal{U}} \sum_{U_i \in \mathcal{U}} \text{diam}(U_i),$$

where \mathcal{U} 's are finite coverings of K by balls U_i satisfying $\text{diam}(U_i) \leq 2a$. By the definition of the Hausdorff measure,

$$\liminf_{a \rightarrow 0} \Lambda(K, a) = \mu_1(K),$$

the 1-dimensional Hausdorff measure of K . Therefore, in order to prove the lemma, we need to establish a suitable upper bound on $\Lambda(K, a)$ for small values of a .

Now, pick $a > 0$ such that $\mathcal{N}_a(E) \subset Q$. By continuity of f , there exists a constant $c > 0$ such that for all p in E , we have $R(p, c) < a$. Next, pick $t > 0$ such that $t < c$ and $L/t = N \in \mathbb{N}$. Choose points $p_1, \dots, p_N \in E$ such that

$$|p_i - p_j| = t|i - j|.$$

Then

$$E \subset \bigcup_{i=1}^N B(p_i, t).$$

Set

$$s_i = R(p_i, t), i = 1, \dots, N.$$

Clearly, for all i we have $s_i < a$ and

$$\kappa^{-1}s_i \leq r(p_i, t).$$

Furthermore, by the definition of the quasisymmetry constant κ ,

$$B(f(p_i), \kappa^{-1}s_i) \subset B(f(p_i), r(p_i, t)) \subset f(B(p_i, t)) \subset B(f(p_i), s_i).$$

Since

$$E \subset \bigcup_{i=1}^N B(p_i, t),$$

we obtain

$$K := f(E) \subset \bigcup_{i=1}^N f(B(p_i, t)) \subset \bigcup_{i=1}^N B(f(p_i), s_i).$$

Therefore, for $B_i := B(f(p_i), s_i)$, the set

$$\mathcal{B} = \{B_i : i = 1, \dots, N\}$$

is a covering of K and the radius of each ball B_i is less than a . In particular, by the definition of Λ , we have:

$$\Lambda(K, a) \leq \sum_i 2s_i.$$

In addition to the balls B_i , we define smaller concentric balls D_i :

$$D_i := \kappa^{-1}B_i = B(f(p_i), \kappa^{-1}s_i) \subset f(B(p_i, t)) \subset B_i, i = 1, \dots, N.$$

Since the covering $\{B(p_i, t) : i = 1, \dots, N\}$ has multiplicity 3 and f is 1-1, it follows that the collection of balls $\mathcal{D} = \{D_i : i = 1, \dots, N\}$ has multiplicity ≤ 3 as well.

Let q be such that

$$\frac{1}{n} + \frac{1}{q} = 1.$$

Then, by the Hölder inequality,

$$\sum_{i=1}^N 2s_i \leq \left(\sum_{i=1}^N 2^q \right)^{1/q} \left(\sum_{i=1}^N s_i^n \right)^{1/n} = 2N^{1/q} \cdot \left(\sum_{i=1}^N s_i^n \right)^{1/n}.$$

Hence, for ω_n , the volume of the unit ball in \mathbb{R}^n ,

$$\Lambda(K, a) \leq 2\omega_n^{1/n} N^{1/q} \left(\sum_{i=1}^n \text{mes}(B_i) \right)^{1/n}.$$

Since \mathcal{D} has multiplicity ≤ 3 , we obtain

$$\sum_m \mu_n(D_i) \leq 3\mu_n\left(\bigcup_i D_i\right).$$

Recall that we need to estimate the sum of the volumes of the balls B_i from above (this would lead to an upper bound on $\Lambda(K, a)$).

We have:

$$B_i = \kappa D_i$$

and, hence,

$$\begin{aligned} \text{mes}(B_i) &= \kappa^n \text{mes}(D_i), \\ \sum_i \text{mes}(B_i) &= \kappa^n \sum_i \text{mes}(D_i) \leq 3\kappa^n \text{mes}\left(\bigcup_i D_i\right) \leq \\ &3\kappa^n \text{mes}\left(f\left(\bigcup_i B(p_i, t)\right)\right) \leq 3\kappa^n \text{mes}(f(Q(y, t))), \end{aligned}$$

since

$$B(p_i, t) \subset Q(y, t), i = 1, \dots, N.$$

By combining the inequalities we obtain

$$\Lambda(f(E), a)^n \leq 2\omega_n^{1/n} N^{n/q} \cdot 3\kappa^n \text{mes}(f(Q(y, t))) = C' N^{n-1} \text{mes}(f(Q(y, t))),$$

where $C' = 6\omega_n^{1/n} \kappa^n$. Therefore,

$$\Lambda(f(E), a)^n \leq C'(Nt)^{n-1} \frac{\text{mes}(f(Q(y, t)))}{t^{n-1}} = C' L^{n-1} \frac{\text{mes}(f(Q(y, t)))}{t^{n-1}},$$

and, by taking the limit as $t \rightarrow 0$, we get:

$$\Lambda(f(E), a)^n \leq C' L^{n-1} \tau(y)$$

for all sufficiently small t . Since this inequality holds for all sufficiently small $a > 0$, we obtain:

$$\mu_1(f(E))^n \leq C' 2^{n-1} (\mu_1(E))^{n-1} \tau(y) = C\tau(y) (\mu_1(E))^{n-1}.$$

The lemma follows. \square

We now can finish the proof of Theorem 22.23. Lemma 22.25 implies that for a.e. $y \in \mathbb{R}^{n-1}$,

$$(22.7) \quad \mu_1(f(E)) \leq (C\tau(y))^{1/n} \mu_1(E)^{1-\frac{1}{n}}.$$

Since $n \geq 2$, this inequality implies that the function

$$f|_{\pi^{-1}(y)}$$

satisfies the hypothesis of Lemma 22.19, which, in turn, implies that this restriction is absolutely continuous. We conclude that the map f is ACL. \square

REMARK 22.26. Note that for $n = 1$, the inequality (22.7) says nothing interesting about the measure of $f(E)$. Furthermore, some quasisymmetric maps fail to be absolutely continuous.

Theorem 22.23 has an important corollary:

COROLLARY 22.27 (F. Gehring, J. Väisälä, see [Väi71]). *For $n \geq 2$, every quasiconformal mapping $f : \Omega \rightarrow \mathbb{R}^n$ has a.e. non-vanishing Jacobian: $J_x(f) \neq 0$ a.e. in Ω .*

PROOF. We will prove a weaker property that will suffice for our purposes, which is that $J_x(f) \neq 0$ on a subset of positive measure, under the assumption that f is weakly quasisymmetric. Suppose to the contrary that $J_x(f) = 0$ a.e. in Ω . The inequality

$$\|D_x f\| \leq H_x(f) |J_x(f)|^{1/n},$$

established in Theorem 22.21, then implies that $D_x f = 0$ a.e. in Ω . Thus, all partial derivatives of f vanish a.e. in Ω . Let $J = [p, q = p + Te_1]$ be a nondegenerate coordinate line segment (parallel to the x_1 -axis), connecting p to q , on which f is absolutely continuous. This means that the Fundamental Theorem of Calculus applies to $f|_J$:

$$f(q) - f(p) = \int_J \frac{\partial}{\partial x_1} f(x) dx_1 = \int_0^T \frac{d}{dt} f(p_1 + te_1, p_2, \dots, p_n) dt = 0.$$

Hence, $f(p) = f(q)$ contradicting injectivity of f . \square

22.4.2.C. *Analytical definition of quasiconformality.* Since quasiconformal maps are differentiable a.e., it is natural to ask if quasiconformality of a map could be defined analytically, in terms of its derivatives. Theorem below gives two alternative analytical definitions of quasiconformality. Even though we will not need this result, it provides a useful prospective on the nature of quasiconformal mappings. We remind the reader that $K(A)$ is the maximal dilatation of the linear transformation A .

THEOREM 22.28. *Suppose that $f : \Omega \rightarrow \Omega' \subset \mathbb{R}^n$ is a homeomorphism. Then the following are equivalent:*

1. f is a quasiconformal mapping.
2. $D_x(f)$ is in $W_{loc}^{1,n}(\Omega)$ and

$$(22.8) \quad K(f) := \operatorname{ess\,sup}_{x \in \Omega} K(D_x(f)) < \infty.$$

3. The mapping f is ALC and satisfies (22.8).

Lastly, analytically and geometrically defined dilatations of f are related by:

$$H_x(f) = H(D_x f)$$

for a.e. $x \in \Omega$.

We refer the reader to [Väi71] for the proof of this theorem. In view of Theorem 22.28, we arrive to

DEFINITION 22.29. A quasiconformal homeomorphism f is called K -quasiconformal if $K(f) \leq K$. The number $K(f)$ is called the *quasiconformality constant* of f .

The reason for defining K -quasiconformality in terms of maximal dilatation $K(D_x(f))$ instead of $H_x(f)$ is that K -quasiconformality is equivalent to yet another, more geometric, definition, in terms of the extremal length (the modulus) of families of curves. The latter definition, for historic reasons, is the main definition of quasiconformality, see [Väi71].

According to Exercise 22.3, the two key measures of quasiconformality, $H(f)$ and $K(f)$ are log-comparable, therefore, using one or the other is only a matter of convenience. What is most important is that $K(f) = 1$ if and only if $H(f) = 1$. If $n = 2$, then, of course, $K_x(f) = H_x(f)$ and $K(f) = H(f)$.

REMARK 22.30. We can now explain the discrepancy in the definition of dilatations of quasiconformal maps: The condition that $H_x(f)$ is bounded is needed in order to ensure that $f : \Omega \rightarrow \mathbb{R}^n$ belongs to $W_{loc}^{1,n}(\Omega)$; on the other hand, the actual bound on dilatation is computed only almost everywhere in Ω . This makes sense since derivatives of f exist only almost everywhere.

22.4.2.D. *Liouville theorem.* Recall that the classical Liouville’s theorem which states that smooth conformal maps between domains in $\mathbb{S}^n, n \geq 3$, are restrictions of Moebius transformations. Gehring’s theorem below shows how one can relax the smoothness assumption in Liouville’s theorem:

THEOREM 22.31 (F. Gehring). *Every 1-quasiconformal homeomorphism of an open connected domain in \mathbb{S}^n ($n \geq 3$) is the restriction of a Moebius transformation.*

Proofs of this theorem can be found in [IM01], [Res89] and [Väi71].

Liouville’s theorem fails, of course, in dimension 2. We will see, however, in Section 23.5.1 that an orientation-preserving quasiconformal homeomorphism $f : \Omega \rightarrow \Omega'$ of two domains in $\mathbb{S}^2 = \mathbb{C} \cup \{\infty\}$, is 1-quasiconformal if and only if it is conformal. Composing with complex conjugation, we conclude that every 1-quasiconformal map is either holomorphic or antiholomorphic. In particular:

THEOREM 22.32. *$f : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ is 1-quasiconformal if and only if f is a Moebius transformation.*

22.4.2.E. *Quasiconformal and quasisymmetric maps.* So far, we have the implications

$$\text{quasisymmetric} \Rightarrow \text{weakly quasisymmetric} \Rightarrow \text{quasiconformal}$$

for maps between domains in \mathbb{R}^n . It turns out that these arrows can be reversed:

THEOREM 22.33 (See e.g. [Hei01]). *Every quasiconformal homeomorphism defined on the entire \mathbb{R}^n is quasiconformal if and only if it is quasisymmetric.*

22.4.2.F. *Convergence property.* The convergence property of quasiconformal mappings is an analogue of the Arzela–Ascoli theorem for uniformly continuous families of maps between metric spaces. In fact, once Theorems 22.36 and 22.38 are established, one can derive the convergence property from the Coarse Arzela–Ascoli theorem (Proposition 8.34) applied to quasiisometries of the hyperbolic space \mathbb{H}^{n+1} . (The three-point normalization for quasiconformal mappings corresponds to a 1-point normalization of quasiisometries.)

Let $z_1, z_2, z_3 \in \mathbb{S}^n$ be three distinct points. A sequence of quasiconformal maps $f_i : \mathbb{S}^n \rightarrow \mathbb{S}^n$ is said to be *normalized at* $\{z_1, z_2, z_3\}$ if the limits

$$\lim_{i \rightarrow \infty} f_i(z_k), \quad k = 1, 2, 3,$$

exist and are pairwise distinct.

THEOREM 22.34 (See [Väi71]). *Let $\Omega \subset \mathbb{S}^n, n \geq 2$, be a connected open subset and $f_i : \Omega \rightarrow f_i(\Omega) \subset \mathbb{S}^n$ be a sequence of K -quasiconformal homeomorphisms normalized at three points in Ω . Then (f_i) contains a subsequence which converges to a K -quasiconformal map.*

The same theorem holds for $n = 1$, except one replaces *quasiconformal* with *quasimöbius*.

We note that the convergence property is usually stated with our normalization condition replaced by the assumption that the three values $w_{ik} = f_i(z_k)$, $k = 1, 2, 3$, of f_i 's are fixed. Recall, however, that Möbius transformations act transitively on three-point subsets of \mathbb{S}^n (see Exercise 4.49). Therefore, there exists a sequence $\gamma_i \in \text{Mob}(\mathbb{S}^n)$ satisfying

$$\gamma_i(w_k) = z_k, \quad k = 1, 2, 3.$$

Since the three limits

$$\lim_{i \rightarrow \infty} w_{ik} = w'_k, \quad k = 1, 2, 3,$$

are all distinct, the sequence (γ_i) subconverges to a Möbius transformation (Lemma 4.50). Composing the quasiconformal mappings f_i with the conformal mappings γ_i , of course, does not change the quasiconformality constants. Therefore, the normalization used in Theorem 22.34 is equivalent to the standard normalization.

22.4.2.G. Strong convergence property. The following strengthening of the convergence property is the key analytical ingredient needed for the proof of Tukia's theorem 23.17 in the next chapter. Fix an open subset $\Omega \subset \mathbb{S}^n$; we will use the notation *mes* for the Lebesgue measure on \mathbb{S}^n restricted to Ω . For a subset $E \subset \Omega$ we let E^c denote the complement $\Omega \setminus E$.

THEOREM 22.35 (Tukia's Strong Convergence Property, [Tuk86]; see also [IM01] for a stronger version). *Consider a sequence $f_i : \Omega \rightarrow \mathbb{S}^n$ of K -quasiconformal maps defined on an open subset $\Omega \subset \mathbb{S}^n$. Suppose that:*

1. *The sequence (f_i) converges to a quasiconformal map f uniformly on compacts in Ω .*
2. *There exists a sequence of subsets $E_i \subset \Omega$, satisfying*

$$\limsup_{i \rightarrow \infty} H(f_i|_{E_i^c}) = H,$$

while

$$\lim_{i \rightarrow \infty} \text{mes}(E_i) = 0.$$

Then $H(f) \leq H$. In particular, if $H = 1$, then f is conformal.

This is a non-trivial theorem since its hypothesis is merely C^0 (uniform convergence of mappings), while the conclusion is about infinitesimal quantities (dilations of quasiconformal mappings).

22.5. Quasisymmetric maps and hyperbolic geometry

The last goal of this chapter is to relate quasisymmetric mappings and quasiisometries of hyperbolic spaces. We will identify the hyperbolic space \mathbb{H}^{n+1} , $n \geq 1$, with the upper half-space

$$\mathbf{U}^{n+1} = \{(x_1, \dots, x_{n+1}) : x_{n+1} > 0\}.$$

We will be also using the notation

$$e_{n+1} = (0, \dots, 0, 1) \in \mathbb{H}^{n+1},$$

and δ for the hyperbolicity constant of \mathbb{H}^{n+1} .

Let $f : \mathbb{H}^{n+1} \rightarrow \mathbb{H}^{n+1}$ be an (L, A) -quasiisometry and let

$$f_\infty : \mathbb{S}^n = \mathbb{R}^n \cup \{\infty\} \rightarrow \mathbb{S}^n$$

be its homeomorphic extension to the boundary sphere of the hyperbolic space given by Theorem 11.108. To simplify the notation, we retain the name f for f_∞ . After composing f with an isometry of \mathbb{H}^{n+1} , we can assume that $f(\infty) = \infty$.

THEOREM 22.36. *There exists $C = C(L, A)$, such that the restriction $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is η -quasisymmetric, with*

$$\eta(t) = e^{2C+At^L}.$$

PROOF. Pick a point $x \in \mathbb{R}^n$ and consider an annulus $\mathbb{A} \subset \mathbb{R}^n$,

$$\mathbb{A} = \{v \in \mathbb{R}^n : R_1 \leq |v - x| \leq R_2\},$$

where $0 < R_1 \leq R_2 < \infty$. We will refer to the ratio $t = \frac{R_2}{R_1}$ as the *eccentricity* of \mathbb{A} . In other words, the eccentricity of \mathbb{A} equals the ratio

$$\frac{|y - x|}{|z - x|}$$

for points y, z which belong to the outer and inner boundaries of \mathbb{A} respectively. Consider the smallest annulus \mathbb{A}' centered at $x' = f(x)$, which contains the topological annulus $f(\mathbb{A})$. Let t' denote the eccentricity of \mathbb{A}' . Then, by the definition of \mathbb{A}' and t' ,

$$\frac{|f(y) - f(x)|}{|f(z) - f(x)|} \leq t',$$

for all points y and z as above. In order to verify that f is η -quasisymmetric, we need to show that $t' \leq \eta(t)$.

After precomposing and postcomposing f with translations of \mathbb{R}^n , we can assume that $x = x' = f(x) = 0$. Let $\alpha \subset \mathbb{H}^{n+1}$ denote the vertical geodesic, connecting 0 to ∞ , i.e. α is the x_{n+1} -axis in \mathbb{H}^{n+1} . Let $\pi_\alpha : \mathbb{H}^{n+1} \rightarrow \alpha$ denote the orthogonal projection to α : For every $p \in \mathbb{H}^{n+1}$, $\pi_\alpha(p) = q \in \alpha$, where $q \in \alpha$ is the unique point such that the geodesic pq is orthogonal to α . The map π_α is the nearest-point projection to α . This projection extends continuously to a map

$$\pi_\alpha : \mathbb{H}^{n+1} \cup (\mathbb{R}^n \setminus \{0\}) \rightarrow \alpha.$$

Then $\pi_\alpha(\mathbb{A})$ is the interval

$$\sigma = [R_1 e_{n+1}, R_2 e_{n+1}] \subset \alpha,$$

whose hyperbolic length equals $\ell = \log(R_2/R_1)$, see Exercise 4.14.

By Lemma 11.105, the (L, A) -quasigeodesic $f(\alpha)$ lies within distance $D(L, A, \delta)$ from the $\alpha \subset \mathbb{H}^{n+1}$, since we are assuming that $f(0) = 0, f(\infty) = \infty$. According to Proposition 11.107, quasiisometries “almost commute” with nearest-point projections and, thus, we obtain

$$d(f\pi_\alpha(x), \pi_\alpha f(x)) \leq C = C(L, A, \delta), \quad \forall x \in \mathbb{H}^{n+1} \cup (\mathbb{R}^n \setminus \{0\}).$$

LEMMA 22.37.

$$\text{diam}(\pi_\alpha(f(\mathbb{A}))) \leq 2C + L\ell + A.$$

PROOF. The ideal boundary of the spherical half-shell

$$\tilde{\mathbb{A}} := \pi_\alpha^{-1}(\sigma) \cap \mathbb{H}^{n+1}$$

is the annulus \mathbb{A} . Therefore, in view of continuity of

$$f : \mathbb{H}^{n+1} \cup \partial_\infty \mathbb{H}^{n+1} \rightarrow \mathbb{H}^{n+1} \cup \partial_\infty \mathbb{H}^{n+1}$$




FIGURE 22.1. Proving quasismetry of f_∞ .

at ideal boundary points, it suffices to verify the inequality in the lemma with $f(\mathbb{A})$ replaced by $f(\tilde{\mathbb{A}})$.

For any two points $p, q \in \tilde{\mathbb{A}}$ we have

$$d(f\pi_\alpha(p), \pi_\alpha f(p)) \leq C, \quad d(f\pi_\alpha(q), \pi_\alpha f(q)) \leq C.$$

Since $d(\pi_\alpha(p), \pi_\alpha(q)) \leq \ell$,

$$d(f\pi_\alpha(p), f\pi_\alpha(q)) \leq L\ell + A,$$

and, by the triangle inequality, we obtain

$$d(\pi_\alpha f(p), \pi_\alpha f(q)) \leq 2C + L\ell + A.$$

The lemma follows. \square

Now we can finish the proof of the theorem. Lemma 22.37 implies that

$$f(\mathbb{A}) \subset \pi_\alpha^{-1}(\sigma'),$$

where $\sigma' \subset \alpha$ has length $\leq \ell' = 2C + L\ell + A$. The eccentricity of the annulus

$$\pi_\alpha^{-1}(\sigma') \cap \mathbb{S}^n$$

is at most $\leq e^{\ell'}$. Thus, the eccentricity of the annulus \mathbb{A}' is also at most

$$e^{\ell'} = e^{2C+A} \cdot e^{L\ell} = e^{2C+A} \cdot e^{L \log(t)} = e^{2C+A} t^L,$$

where $t = R_2/R_1$. \square

The following converse theorem was first proven by Pekka Tukia in the case of hyperbolic spaces and then extended by Frederic Paulin to the case of more general Gromov–hyperbolic spaces.

THEOREM 22.38 (P. Tukia [Tuk94], F. Paulin [Pau96]). *Every quasismetric homeomorphism $\mathbb{R}^n \rightarrow \mathbb{R}^n$ extends to a quasiisometric map $\mathbb{H}^{n+1} \rightarrow \mathbb{H}^{n+1}$. More precisely: For every η -quasismetric homeomorphism $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ there exists an (L, A) -quasiisometric map F of the hyperbolic space \mathbb{H}^{n+1} , such that*

$$F_\infty = f,$$

where F_∞ is the boundary extension of the quasiisometry F given by Theorem 11.108. Moreover, the constants L, A depend only on $\eta(1)$ and $\eta(2)$.

PROOF. We let $\Pi : \mathbb{H}^{n+1} = \mathbb{R}_+^{n+1} \rightarrow \mathbb{R}^n$ denote the coordinate projection. We define the extension F as follows. For every $p \in \mathbb{H}^{n+1}$, let $\alpha = \alpha_p$ be the complete vertical geodesic through p . This geodesic is asymptotic to the points ∞ and $x = x_p = \Pi(p) \in \mathbb{R}^n$. Let $y \in \mathbb{R}^n$ be a point such that $\pi_\alpha(y) = p$. (The point y is non-unique, of course: Every point $y \in S(x, R)$ in the sphere centered at x and of the radius $R = |x - p|$ would work.) Let $x' := f(x), y' := f(y)$ and let $\alpha' \subset \mathbb{H}^{n+1}$ denote the vertical geodesic through x' and let $p' := \pi_{\alpha'}(y')$. Lastly, set $F(p) := p'$.

We will prove that F is an (L, A) -coarse Lipschitz map. The coarse inverse to F will be the map \bar{F} defined *via* extension of the map f^{-1} following the same procedure. We will leave it as an exercise to the reader to verify that \bar{F} is indeed a coarse inverse to F and to estimate $d(\bar{F} \circ F, \text{Id})$.




FIGURE 22.2. Triangle $\Delta(p_1, p_2, p_3)$.

Suppose that $d(p_1, p_2) \leq 1/4$. We are looking for a uniform upper bound on the distance

$$d(p'_1, p'_2) \leq A,$$

with A depending only on $\eta(1)$ and $\eta(2)$: Existence of such upper bound will imply that F is $(4A, A)$ -coarse Lipschitz (see Lemma 8.8).

Consider the triangle $\Delta(p_1, p_2, p_3)$ as in Figure 22.2, where the points p_1, p_3 belong to the same vertical geodesic while p_2, p_3 belong to the same horosphere Σ centered at the point $\infty \in \mathbb{S}^n$ (after swapping p_1, p_2 if necessary we may assume that the point p_1 does not belong to the horoball bounded by Σ). Then

$$d(p_1, p_2) \leq \frac{1}{4} \Rightarrow d(p_1, p_3) \leq \frac{1}{4} \Rightarrow d(p_2, p_3) \leq \frac{1}{2}.$$

The uniform upper bounds

$$d(F(p_1), F(p_2)) \leq C_1, \quad d(F(p_3), F(p_2)) \leq C_3$$

would imply

$$d(F(p_1), F(p_2)) \leq A := C_1 + C_2,$$

as required. Therefore, our problem reduces to the two special cases:

Case 1. p_1, p_2 belong to the common vertical geodesic α ; in particular, $x_1 = x_2 = x$. We will assume, for concreteness, that $|x - y_1| \leq |x - y_2|$. Hence,

$$d(p_1, p_2) = \log \left(\frac{|y_2 - x|}{|y_1 - x|} \right)$$

and the assumption $d(p_1, p_2) \leq 1/2$ implies that

$$\left(\frac{|y_2 - x|}{|y_1 - x|} \right) \leq \sqrt{e} < 2.$$

Since the map f is η -quasisymmetric,

$$\frac{|y'_2 - x'|}{|y'_1 - x'|} \leq \eta \left(\frac{|y_2 - x|}{|y_1 - x|} \right) \leq \eta(2).$$

In particular,

$$d(p'_1, p'_2) \leq C_1 := \log(\eta(2)).$$

This estimate also shows that different choices of the point $y \in \mathbb{R}^n$ in the definition of F lead to maps which are within distance $\leq C_1$ from each other.

Case 2. Suppose that the points p_1, p_2 have the same last coordinate, i.e. they belong to the same horosphere centered at the point $\infty \in \mathbb{S}^n$. As before, for $i = 1, 2$, we set $x_i := \Pi(p_i)$, let α_i denote the vertical hyperbolic geodesic through p_i and pick arbitrarily $y_i \in \mathbb{R}^n$ with $\pi_{\alpha_i}(y_i) = p_i$. Then

$$R_i := |y_i - x_i| = |p_i - x_i|$$

and $R = R_1 = R_2$ (as p_1, p_2 have the same last coordinate). Set $R_3 := |x_1 - x_2|$.

The reader will verify, using the formula (4.15), that

$$d(p_1, p_2) \leq 1/2 \Rightarrow t := |x_1 - x_2| < R \Rightarrow R_3 \leq 2R.$$

In particular, if $n \geq 2$, we could have chosen, if we wish, $y_1 = y_2$.

The points $p'_i = F(x_i)$, $i = 1, 2$, belong to the vertical geodesics α'_i , such that

$$\pi_{\alpha'_i}(y'_i) = p'_i,$$

$y'_i = f(y_i)$, $i = 1, 2$. We define the points $x'_i := \Pi(p'_i) \in \mathbb{R}^n$, and set

$$R'_i := |y'_i - x'_i| = |p'_i - x'_i|, i = 1, 2.$$

Lastly, set

$$t' := |x'_1 - x'_2|, \quad R'_3 := |x'_1 - y'_2|.$$

After switching the roles of p_1 and p_2 , we can assume that $R'_1 \leq R'_2$. Then

$$(22.9) \quad d(p'_1, p'_2) \leq \frac{t'}{R'_1} + \log \left(\frac{R'_2}{R'_1} \right),$$

as we can first travel from p'_1 to the line α'_2 horizontally, along a path of the hyperbolic length

$$\frac{t'}{R'_1} = \frac{|x'_2 - x'_1|}{R'_1},$$

and then vertically, along α'_2 , along a path of the length $\log(R'_2/R'_1)$. We now apply the η -quasisymmetry condition (equation (22.2)) to the triple of points x_1, y_1, x_2 , (with x_1 playing the role of the center) and get:

$$(22.10) \quad \frac{t'}{R'_1} \leq \eta \left(\frac{t}{R} \right) \leq \eta(1).$$

Setting $R_3 := |x_1 - y_2|$, $R'_3 := |x'_1 - y'_2|$ and applying the η -quasisymmetry condition to the triple of points x_1, y_1, y_2 (with x_1 again playing the role of the center), we obtain

$$(22.11) \quad \frac{R'_3}{R'_1} \leq \eta \left(\frac{R_3}{R_1} \right) \leq \eta \left(\frac{2R}{R} \right) = \eta(2).$$

The inequalities $R'_2 \leq t' + R'_3$, (22.10) and (22.11) imply

$$\frac{R'_2}{R'_1} \leq \frac{t' + R'_3}{R'_1} \leq \eta(1) + \eta(2).$$

Combining this inequality with (22.9), we conclude that

$$d(p'_1, p'_2) \leq C_2 := \eta(1) + \log(\eta(1) + \eta(2)).$$

Thus, in general, for $p_1, p_2 \in \mathbb{H}^{n+1}$, $d(p_1, p_2) \leq 1/4$, we have:

$$d(F(p_1), F(p_2)) \leq C_1 + C_2 = A.$$

It follows that F is an (A, A) -coarse Lipschitz map, where $A = C_1 + C_2$.

The last thing to observe is that since F , by the construction, sends vertical geodesics to vertical geodesics, its extension $F_\infty : \mathbb{S}^n \rightarrow \mathbb{S}^n$, defined by Theorem 11.108, equals f . \square

EXERCISE 22.39. Consider a linear quasiconformal mapping $f : \mathbf{x} \mapsto A\mathbf{x}$, $\mathbf{x} \in \mathbb{R}^n$, $A \in GL(n, \mathbb{R})$. Define the linear mapping

$$\tilde{f} : (\mathbf{x}, t) \mapsto (A\mathbf{x}, t), \quad \mathbf{x} \in \mathbb{R}^n, t > 0.$$

Show that \tilde{f} is a quasiisometry of \mathbb{H}^{n+1} .

Groups quasi-isometric to \mathbb{H}^n

The main result of this chapter is the following theorem, due to Pekka Tukia, see [Tuk86] and [Tuk94], as well as the paper by Jim Cannon and Daryl Cooper [CC92]:

THEOREM 23.1 (P. Tukia). *If G is a finitely generated group QI to \mathbb{H}^{n+1} (with $n \geq 2$), then G acts geometrically on \mathbb{H}^{n+1} . In particular, G is virtually isomorphic to a uniform lattice in the Lie group $\text{Isom}(\mathbb{H}^{n+1})$.*

REMARK 23.2. The same result also holds for $n = 1$, but the proof in this case is completely different, see Section 23.7.

Recall that if a group G is QI to \mathbb{H}^{n+1} , then it quasiacts on \mathbb{H}^{n+1} , see Lemma 8.63. Furthermore (by Theorem 11.135), every such quasiaction φ determines a topological action

$$\varphi_\infty : G \curvearrowright \mathbb{S}^n$$

on the boundary sphere of \mathbb{H}^{n+1} . Since the quasiaction of G is by uniform quasi-isometries, the action of $G \curvearrowright \mathbb{S}^n$ is by uniformly quasiconformal homeomorphisms, see Theorem 22.36. Such group actions are called *uniformly quasiconformal*. According to Lemma 8.63, the quasiaction $G \curvearrowright \mathbb{H}^{n+1}$ is geometric and, by Lemma 11.118, every point $\xi \in \mathbb{S}^n$ is a conical limit point of $G \curvearrowright \mathbb{S}^n$. Lastly, by Theorem 11.135, the fact that the quasiaction $G \curvearrowright \mathbb{H}^{n+1}$ is geometric translates to:

The action $G \curvearrowright \text{Trip}(\mathbb{S}^n)$ is properly discontinuous and cocompact. In particular:

1. The kernel of the homomorphism $\varphi_\infty : G \rightarrow \text{Homeo}(\mathbb{S}^n)$ is a finite normal subgroup of G .
2. The image $\bar{G} = \varphi_\infty(G)$ is a discrete subgroup of the group of homeomorphisms $\text{Homeo}(\mathbb{S}^n)$, where $\text{Homeo}(\mathbb{S}^n)$ is equipped with the topology of uniform convergence.

We refer the reader to the Notation 11.88 for the definition of the space $\text{Trip}(\mathbb{S}^n)$.

Our goal, and this is the main result of Dennis Sullivan (for $n = 2$) and Pekka Tukia (for all $n \geq 2$), is to show that, under the above hypothesis, there exists a quasiconformal homeomorphism $f : \mathbb{S}^n \rightarrow \mathbb{S}^n$ which conjugates \bar{G} to a group of Moebius transformations, whose action on \mathbb{H}^{n+1} is geometric. Once the existence of such f is established, Theorem 23.1 would follow. We see that in order to prove Theorem 23.1, one is naturally lead to study *uniformly quasiconformal group actions on \mathbb{S}^n* . Our treatment of quasiconformal groups mostly follows the arguments in [Tuk86] and in [IM01]. A different, but related, proof is given by Peter Haïssinsky [Haï09].

23.1. Uniformly quasiconformal groups

Let $G < \text{Homeo}(\mathbb{S}^n)$ be a group consisting of quasiconformal homeomorphisms. The group G is called *uniformly quasiconformal*, if there exists $K < \infty$ such that $K(g) \leq K$ for all $g \in G$. Recall that $K(g)$ is the quasiconformality constant of the homeomorphism $g : \mathbb{S}^n \rightarrow \mathbb{S}^n$, see Equation (22.8). Trivial examples of uniformly quasiconformal groups are given by subgroups $\Gamma < \text{Mob}(\mathbb{S}^n)$ of Moebius transformations and their quasiconformal conjugates

$$\Gamma^f = f\Gamma f^{-1},$$

where f is k -quasiconformal. Then for every $g \in \Gamma^f$,

$$K(g) = K(f\gamma f^{-1}) \leq k^2 = K.$$

We say that a uniformly quasiconformal subgroup $G < \text{Homeo}(\mathbb{S}^n)$ is *exotic* if it is not quasiconformally conjugate to a group of Moebius transformations. The following theorem is a fundamental fact of quasiconformal analysis in dimension $n = 2$, observed first by D. Sullivan in [Sul81]:

THEOREM 23.3. *There are no exotic uniformly quasiconformal subgroups in $\text{Homeo}(\mathbb{S}^2)$.*

This theorem fails rather badly for $n \geq 3$. The first examples of exotic uniformly quasiconformal subgroups $G < \text{Homeo}(\mathbb{S}^n)$, $n \geq 3$, were constructed by P. Tukia [Tuk81]. Tukia's subgroups G are non-discrete, isomorphic to certain connected solvable Lie groups, which do not admit embeddings into $\text{Isom}(\mathbb{H}^m)$ for any m . Algebraically, Tukia's examples are semidirect products $\mathbb{R}^k \rtimes \mathbb{R}^2$, where $(a, b) \in \mathbb{R}^2$ acts on \mathbb{R}^k via a diagonal matrix $D(a, b)$ that has (generically) two distinct eigenvalues $\neq \pm 1$. Further examples of *discrete* exotic uniformly quasiconformal subgroups of $\text{Homeo}(\mathbb{S}^3)$ were constructed in [FS88], [Mar86] (these groups have torsion) and in [Kap92] (these are certain surface groups acting on \mathbb{S}^3). An example of a discrete uniformly quasiconformal subgroup of $\text{Homeo}(\mathbb{S}^3)$ which is not *isomorphic* to subgroup of $\text{Isom}(\mathbb{H}^4)$ was constructed in [Isa90].

PROBLEM 23.4. Suppose that $G < \text{Homeo}(\mathbb{S}^n)$ is a discrete uniformly quasiconformal subgroup. Is it true that G is isomorphic to a subgroup of $\text{Isom}(\mathbb{H}^m)$ for some m ?

The answer to this questions is probably negative. One can, nevertheless, ask which algebraic properties of discrete groups of Moebius transformations are shared by discrete uniformly quasiconformal subgroups, e.g.:

PROBLEM 23.5. Suppose that $G < \text{Homeo}(\mathbb{S}^n)$ is an infinite discrete uniformly quasiconformal subgroup, $n \geq 3$.

1. Is it true that G does not have the Property (T)?
2. Is it true that G has the Haagerup property?
3. Is it true that the action $G \curvearrowright \mathbb{S}^n$ extends to a uniformly quasiconformal action $G \curvearrowright \mathbb{H}^{n+1}$?

Note that, in view of Theorems 23.7 and 8.66, there exists a Gromov-hyperbolic space X quasiisometric to \mathbb{H}^{n+1} , such that G acts *isometrically* on X and the actions of G on $\partial_\infty X$ and \mathbb{S}^n are topologically conjugate.

Another problem, open since Tukia's examples of exotic connected solvable uniformly quasiconformal subgroups of $\text{Homeo}(\mathbb{S}^n)$ is:

PROBLEM 23.6. Suppose that $N < \text{Homeo}(\mathbb{S}^n)$ is a uniformly quasiconformal connected nilpotent subgroup. Is it true that N is abelian?

23.2. Hyperbolic extension of uniformly quasiconformal groups

As we saw, every quasiaction $G \curvearrowright \mathbb{H}^{n+1}$ extends to a uniformly quasiconformal action $G \curvearrowright \mathbb{S}^n$. Our first goal is to prove the converse:

THEOREM 23.7 (P. Tukia, [Tuk94]). *Every uniformly quasiconformal action $\rho : G \curvearrowright \mathbb{S}^n$ extends to a quasiaction $\varphi : G \curvearrowright \mathbb{H}^{n+1}$ in the sense that*

$$\varphi(g)_\infty = \rho(g), \quad \forall g \in G,$$

where $h_\infty : \mathbb{S}^n \rightarrow \mathbb{S}^n$ is the extension of a quasiisometry $h : \mathbb{H}^{n+1} \rightarrow \mathbb{H}^{n+1}$ given by Theorem 22.36.

PROOF. For every $g \in G$ we let $\varphi(g) : \mathbb{H}^{n+1} \rightarrow \mathbb{H}^{n+1}$ denote the quasiisometric extension of $\rho(g)$ constructed in Theorem 22.38. In view of the same extension theorem, since every $\rho(g)$ is K -quasiconformal, every $\varphi(g)$ is an (L, A) -quasiisometry, where L and A depend only on K . We need to show that the extension φ defines a quasiaction, i.e. there exists $C = C(L, A)$ such that:

(1) For all $g_1, g_2 \in G$

$$\text{dist}(\varphi(g_1) \circ \varphi(g_2), \varphi(g_1 g_2)) \leq C,$$

(2)

$$\text{dist}(\varphi(1_G), \text{Id}) \leq C.$$

It follows immediately from the construction of the quasiisometric extension in the proof of Theorem 22.38 that

$$\varphi(1_G) = \text{Id}.$$

In order to verify (1), we note that for all $g_1, g_2 \in G$, the composition

$$f' = \varphi(g_1) \circ \varphi(g_2)$$

is an $(L^2, LA + A)$ -quasiisometry, while

$$f'' = \varphi(g_1 g_2)$$

is an (L, A) -quasiisometry. Furthermore,

$$f'_\infty = \rho(g_1) \circ \rho(g_2) = \rho(g_1 g_2) = f''_\infty.$$

By homogeneity of \mathbb{H}^{n+1} , every point of the hyperbolic space is a centroid of an ideal triangle. Therefore, Lemma 11.112 implies that

$$\text{dist}(f', f'') \leq C(L, A) = D(L, A, 0, \delta),$$

where δ is the hyperbolicity constant of \mathbb{H}^{n+1} . □

This theorem shows that the study of uniformly quasiconformal groups is equivalent to the study of quasiactions on \mathbb{H}^{n+1} . In particular, we can define *conical limit points* for uniformly quasiconformal subgroups $G < \text{Homeo}(\mathbb{S}^n)$ as conical limit points of the extended quasiactions.

Our goal, thus, is to prove the following theorem which was first established by D. Sullivan [Sul81] for $n = 2$ (without restrictions on conical limit points) and then by P. Tukia in full generality:

THEOREM 23.8 (P. Tukia, [Tuk86]). *Suppose that $G < \text{Homeo}(\mathbb{S}^n)$ is a countable uniformly quasiconformal subgroup. Assume also that $n \geq 2$ and that almost every point of \mathbb{S}^n is a conical limit point of G . Then G is quasiconformally conjugate to a subgroup of the Moebius group $\text{Mob}(\mathbb{S}^n)$.*

Before proving this theorem, we will need a few technical tools.

23.3. Least volume ellipsoids

Observe that a closed ellipsoid centered at 0 in \mathbb{R}^n can be described as

$$E = E_A = \{x \in \mathbb{R}^n : \varphi_A(x) = x^T A x \leq 1\},$$

where A is some positive-definite symmetric $n \times n$ matrix. The volume of such an ellipsoid is given by the formula

$$\text{Vol}(E_A) = \omega_n (\det(A))^{-1/2},$$

where ω_n is the volume of the unit ball in \mathbb{R}^n . A subset $X \subset \mathbb{R}^n$ is called *centrally-symmetric* if $X = -X$.

THEOREM 23.9 (F. John, [Joh48]). *For every compact centrally-symmetric subset $X \subset \mathbb{R}^n$ with non-empty interior, there exists a unique ellipsoid $E(X)$ of least volume containing X . The ellipsoid $E(X)$ is called the John–Loewner ellipsoid of X .*

PROOF. The existence of $E(X)$ is clear by compactness. We need to prove uniqueness. Consider the function f on the space P_n of positive definite symmetric $n \times n$ matrices, given by

$$f(A) = -\frac{1}{2} \log(\det(A)).$$

LEMMA 23.10. *The function $f : P_n \rightarrow \mathbb{R}$ is strictly convex, in the sense that for every family of matrices $C_t \in P_n$,*

$$C_t = tA + (1-t)B \in P_n, \quad 0 \leq t \leq 1,$$

the function $g(t) = f(C_t)$ is strictly convex.

PROOF. The matrices A and B in P_n can be simultaneously diagonalized by a matrix $M \in GL(n, \mathbb{R})$:

$$MAM^T = D_A, \quad MBM^T = D_B,$$

where D_A, D_B are diagonal matrices. The matrices

$$D_t = tD_A + (1-t)D_B = M(tA + (1-t)B)M^T$$

are, of course, also diagonal and

$$f(D_t) = f(MC_tM^T) = -\log \det(M) - \frac{1}{2} \log \det(C_t) = -\log \det(M) + f(C_t).$$

Therefore, it suffices to prove strict convexity of f on the space Diag_n^+ of positive-definite diagonal $n \times n$ matrices. For each diagonal matrix $D = \text{Diag}(x_1, \dots, x_n)$ with the diagonal entries x_1, \dots, x_n ,

$$f(D) = -\frac{1}{2} \sum_{i=1}^n \log(x_i).$$

Lastly, the function $f : \text{Diag}_n^+ \rightarrow \mathbb{R}$ is strictly convex since \log is strictly concave. \square

In particular, whenever $V \subset P_n$ is a convex subset and $f|_V$ is proper, f attains a unique minimum on V . Since \log is a strictly increasing function, the same uniqueness assertion holds for the function $\det^{-1/2}$ on P_n . Let $V = V_X$ denote the set of matrices $C \in P_n$ such that $X \subset E_C$. Since $\varphi_A(x)$ is linear as a function of A for any fixed $x \in X$, it follows that V convex. Thus, the least volume ellipsoid containing X is unique. \square

23.4. Invariant measurable conformal structure

Throughout this section, we assume that n is at least 2 (the discussion becomes meaningless otherwise). Recall (see Section 3.3) that a measurable Riemannian metric on $\mathbb{S}^n = \mathbb{R}^n \cup \{\infty\}$ is a measurable map g from \mathbb{S}^n to the space P_n of positive definite symmetric $n \times n$ matrices. (Since we are working in the measurable category, we can and will ignore the point ∞ .)

A *measurable conformal structure* on \mathbb{S}^n is a measurable Riemannian metric defined up to multiplication by a positive measurable function. In order to avoid the ambiguity with the choice of the conformal factor, one can normalize the measurable metric $g : \mathbb{R}^n \rightarrow P_n$ so that $\det(g(x)) = 1$ for all $x \in \mathbb{S}^n$. We will refer to such g as a *normalized measurable Riemannian metric*.

Every quasiconformal mapping $f : \mathbb{S}^n \rightarrow \mathbb{S}^n$ acts on measurable Riemannian metrics *via* the pull-back by the usual formula:

$$f^*(g) = h, \quad h(x) = (D_x f)^T g(f(x)) (D_x f).$$

Here we are using the fact that f is differentiable almost everywhere in \mathbb{S}^n and its derivative is a measurable matrix-valued function on \mathbb{R}^n , see Theorem 22.23. Measurability of the function $x \mapsto D_x f$ explains why considering measurable Riemannian metrics is the right thing to do in the context of quasiconformal mappings.

REMARK 23.11. The reader might have noticed that in the book we proved the ACL property only for quasisymmetric rather than quasiconformal mappings. For the purposes of quasiisometric rigidity this does not matter, since extensions of quasiisometries are quasimöbius mappings and, hence, we can use the analytical properties of quasisymmetric mappings proven in Chapter 22. Furthermore, every quasiconformal mapping of \mathbb{R}^n is also quasisymmetric, Section 22.4.2.E.

If we consider normalized Riemannian metrics, then the appropriate action is given by the formula:

$$f^\bullet(g) = h, \quad h(x) = (J_x)^{-2n} (D_x f)^T g(f(x)) (D_x f)$$

in order for h to be normalized as well. Here J_x is the Jacobian determinant of f at x . We will think of normalized measurable Riemannian metrics as measurable conformal structures on \mathbb{S}^n .

A measurable conformal structure μ on \mathbb{S}^n is called *bounded* if it is represented by a bounded normalized measurable Riemannian metric, i.e. a bounded map

$$\mathbb{S}^n \rightarrow P_n \cap \{\det = 1\}.$$

Below, we interpret boundedness of μ in terms of eigenvalues.

Given a measurable Riemannian metric $\mu(x) = A_x$, we define its *linear dilatation* $H(\mu)$ as the essential supremum of the ratios

$$H(x) := \frac{\sqrt{\lambda_n(x)}}{\sqrt{\lambda_1(x)}},$$

where $\lambda_1(x) \leq \dots \leq \lambda_n(x)$ are the eigenvalues of A_x . Geometrically speaking, if $E_x \subset T_x \mathbb{R}^n$ is the unit ball with respect to A_x , then $H(x)$ is the eccentricity of the ellipsoid E_x , i.e. the ratio of the largest to the smallest axis of E_x . In particular, $H(x)$ and $H(\mu)$ are conformal invariants of μ .

EXERCISE 23.12. 1. A measurable conformal structure μ is bounded if and only if $H(\mu) < \infty$.

2. A subset \mathcal{M} in the space normalized measurable conformal structures is bounded if and only if

$$\sup_{\mu \in \mathcal{M}} H(\mu) < \infty.$$

We say that a measurable conformal structure $\mu(x) = A_x$ on \mathbb{R}^n is *invariant under a quasiconformal subgroup* $G < \text{Homeo}(\mathbb{S}^n)$ if

$$g^\bullet \mu = \mu, \forall g \in G.$$

In detail:

$$\forall g \in G, \quad (J_{g,x})^{-\frac{1}{2n}} (D_x g)^T \cdot A_{g,x} \cdot D_x g = A_x$$

a.e. in \mathbb{R}^n .

The following was first proven by Sullivan in [Sul81] for $n = 2$ and, then, by Tukia [Tuk86] for an arbitrary n :

PROPOSITION 23.13. *Every countable uniformly quasiconformal subgroup $G < \text{Homeo}(\mathbb{S}^n)$ admits an invariant measurable conformal structure λ on \mathbb{S}^n .*

PROOF. Let μ_0 be the Euclidean metric on \mathbb{R}^n , it is given by the constant matrix function $x \mapsto I$. Consider the orbit $G \cdot \mu_0$ in the space of normalized measurable Riemannian metrics. The idea is to take the ‘‘average’’ of all the measurable conformal structures in this orbit.

Since G is countable, there exists a subset of full measure in \mathbb{S}^n on which we have matrix-valued functions

$$A_{g,x} = g^\bullet \mu_0 = (J_{g,x})^{-\frac{1}{2n}} (D_x g)^T \cdot D_x g, \quad g \in G.$$

With this definition, $H(A_{g,x}) = H_g(x)$, is the linear dilatation of g at x , see Definition 22.11. Therefore, the assumption that G is uniformly quasiconformal is equivalent to the assumption that the family of measurable conformal structures $G \cdot \mu_0$ is uniformly bounded:

$$H := \sup_{g \in G} H(g^\bullet \mu_0) < \infty.$$

Geometrically, one can think of this as follows. For a.e. x we let $E_{g,x}$ denote the unit ball in $T_x \mathbb{R}^n$ with respect to $g^\bullet(\mu_0)$. From the Euclidean viewpoint, $E_{g,x}$ is just an ellipsoid of the volume ω_n (since $g^\bullet(\mu_0)$ is normalized). This ellipsoid (up to scaling) is the image of the unit ball under the inverse of the derivative $D_x g$. Then uniform boundedness of the conformal structures $g^\bullet(\mu_0)$ simply means that the eccentricities of the ellipsoids $E_{g,x}$ are bounded by the number H , which is independent of g and x . Since the volume of each $E_{g,x}$ is fixed, it follows that the

diameters of these ellipsoids are uniformly bounded above and below: There exists $0 < R < \infty$ such that

$$B(0, R^{-1}) \subset E_{g,x} \subset B(0, R), \forall g \in G$$

for a.e. $x \in \mathbb{R}^n$.

Let U_x denote the union of the ellipsoids

$$\bigcup_{g \in G} E_{g,x}.$$

Since each ellipsoid $E_{g,x}$ is centrally-symmetric, so is U_x . By the construction, the family of sets $\{U_x, x \in \mathbb{R}^n\}$ is invariant under the group G :

$$(J_{g,x})^{-1/n} D_x g(U_x) = U_{g(x)}, \quad \forall g \in G.$$

For each x (where U_x is defined, which is a subset of full measure), we define an ellipsoid E_x , the John-Loewner ellipsoid of the set U_x . Since the group G preserves the family of sets U_x and since, after normalization, the action of $D_x g$ on the tangent space is volume-preserving, it follows (by uniqueness of the John-Loewner ellipsoid, Theorem 23.9) that G also preserves the family of ellipsoids E_x .

Clearly,

$$B(0, R^{-1}) \subset E_x \subset B(0, R)$$

a.e. in \mathbb{R}^n , and, hence, the eccentricities of the ellipsoids E_x are uniformly bounded above and below. Let $\mu(x)$ denote the (a.e. defined) function $\mathbb{R}^n \rightarrow P_n$ which sends x to the matrix A_x such that E_x is the unit ball with respect to the quadratic form defined by A_x . **Then** $H(A_x) \leq R^2$ a.e..

LEMMA 23.14. *The function $\mu : x \rightarrow A_x$ is measurable.*

PROOF. Since G is countable, we can represent G as an increasing union of finite subsets $G_i \subset G$. For each i we define the sets

$$U_{x,i} = \bigcup_{g \in G_i} E_{g,x}$$

and the corresponding John-Loewner ellipsoids $E_{x,i}$. We leave it to the reader to check that since each ellipsoid $E_{g,x}$ is measurable as a function of y , then $E_{x,i}$ is also measurable. Note also that

$$E_x = \bigcup_{i \in \mathbb{N}} E_{x,i}.$$

Let $\mu_i : \mathbb{R}^n \rightarrow P_n$ denote the measurable functions defined by the ellipsoids $E_{x,i}$. We will think of these functions as functions $\mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}_+$,

$$(x, v) \mapsto v^T \mu_i(x) v \in \mathbb{R}_+.$$

Then the fact that $E_i \subset E_{i+1}$ means that

$$\mu_i(x, v) \geq \mu_{i+1}(x, v).$$

Furthermore,

$$\mu = \lim_i \mu_i.$$

Now, the lemma follows from the Lebesgue monotone convergence theorem (Beppo Levi's theorem), see e.g. [SS05]. \square

This also concludes the proof of the proposition. \square

The above proposition also holds for uncountable uniformly quasiconformal groups, see [Tuk86], but we will not need this fact.

23.5. Quasiconformality in dimension 2

In this section we reformulate quasiconformality of a map in the 2-dimensional case in terms of the *Beltrami equation* and explain the relation between measurable conformal structures on domains in $\mathbb{C} = \mathbb{R}^2$ and *Beltrami differentials*. We refer to [Ahl06] and [Leh87] for further details.

23.5.1. Beltrami equation. For computational purposes, we will use the complex differentials $dz = dx + idy$ and $d\bar{z} = dx - idy$. These differentials define coordinates on the complexification of the real tangent space $T_z\Omega$ of open subsets $\Omega \subset \mathbb{R}^2$. Accordingly,

$$\begin{aligned}\frac{\partial}{\partial z} &= \frac{1}{2} \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right), \\ \frac{\partial}{\partial \bar{z}} &= \frac{1}{2} \left(\frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right).\end{aligned}$$

To simplify the notation, we let ∂f denote $\frac{\partial f}{\partial z} = f_z$ and let $\bar{\partial} f$ denote $\frac{\partial f}{\partial \bar{z}} = f_{\bar{z}}$, the *holomorphic* and *antiholomorphic* derivatives respectively.

Consider a function $f(z)$ which is differentiable at a point $z \in \mathbb{C}$. Writing $f = u + iv$, we obtain a formula for the (real) Jacobian determinant of f :

$$J_f = u_x v_y - u_y v_x = |\partial f|^2 - |\bar{\partial} f|^2.$$

We will assume from now on that f is orientation-preserving at z , i.e. $|\partial f(z)| > |\bar{\partial} f(z)|$.

For $\alpha \in [0, 2\pi]$, the directional derivative of f at z in the direction $e^{i\alpha}$ equals

$$\partial_\alpha f = \partial f + e^{-2i\alpha} \bar{\partial} f.$$

We now can compute lengths of the major and minor semi-axes of the ellipse, which is the image of the unit tangent circle under $D_z f$:

$$\begin{aligned}\max_\alpha |\partial_\alpha f| &= |\partial f| + |\bar{\partial} f|, \\ \min_\alpha |\partial_\alpha f| &= |\partial f| - |\bar{\partial} f|.\end{aligned}$$

Thus,

$$H_z(f) = \max_{\alpha, \beta} \frac{|\partial_\alpha f|}{|\partial_\beta f|} = \frac{|\partial f| + |\bar{\partial} f|}{|\partial f| - |\bar{\partial} f|}$$

is the linear dilatation of f at z . Setting $\mu(z) = \frac{\bar{\partial} f}{\partial f}$, we obtain

$$H_z(f) = \frac{1 + |\mu(z)|}{1 - |\mu(z)|}.$$

Suppose now that $f : \Omega \rightarrow \mathbb{C}$ and $f \in W_{loc}^{1,2}(\Omega)$; in particular, f is differentiable a.e. in Ω , its derivatives are locally square-integrable in Ω and $J_z(f) > 0$ in Ω , i.e. f is orientation-preserving. **Then** we have a measurable function

$$(23.1) \quad \mu = \mu(z) = \frac{f_{\bar{z}}}{f_z},$$

called the *Beltrami differential* of f ; the equation (23.1) is called the *Beltrami equation*. Let $k = k_f = \|\mu\|$ be the L^∞ -norm of μ in Ω . **Then**

$$K(f) = \sup_{z \in \Omega} H_z(f) = \frac{1+k}{1-k}$$

is the coefficient of quasiconformality of f .

We conclude that the following are equivalent for a function f :

1. f is K -quasiconformal, where $K = \frac{1+k}{1-k}$.
2. f satisfies the Beltrami equation and $k = \|\mu\| < 1$.

In particular, an (orientation-preserving) quasiconformal map is 1-quasiconformal if and only if $k_f = 0$, i.e. $\mu = 0$, equivalently, $\bar{\partial}f = 0$ (almost everywhere). A theorem of Weyl (see e.g. [Ahl06]) then states that such maps are holomorphic.

23.5.2. Measurable Riemannian metrics. Let $f : \Omega \rightarrow \Omega' \subset \mathbb{C}$ be an orientation-preserving quasiconformal homeomorphism, $w = f(z)$, with the Beltrami differential μ . For $w = u + iv$ it is useful to compute the pull-back of the Euclidean metric $du^2 + dv^2 = |dw|^2$ by the map f :

$$\begin{aligned} |dw|^2 &= |\partial f dz + \bar{\partial} f d\bar{z}|^2 = \\ &= |\partial f|^2 \cdot |dz + \frac{\bar{\partial} f}{\partial f} d\bar{z}|^2 = |\partial f|^2 \cdot |dz + \mu(z) d\bar{z}|^2. \end{aligned}$$

Therefore, up to the conformal multiple $|\partial f|^2$, the pull-back metric $f^*(|dw|^2)$ equals the measurable Riemannian metric

$$ds_\mu^2 := |dz + \mu(z) d\bar{z}|^2.$$

Our next goal is to show that an arbitrary measurable Riemannian metric ds^2 on a domain (an open connected subset) $\Omega \subset \mathbb{C}$ is conformal to a metric of the form ds_μ^2 for some μ . Consider a measurable Riemannian metric

$$ds^2 = E dx^2 + 2F dx dy + G dy^2.$$

We will do the computation in the tangent space at each point $z \in \Omega$. Then, by a change of variables $z = e^{i\theta} w$, we convert a general form ds^2 to the one with $F(z) = 0$; the same change of variables converts $|dz + \mu(z) d\bar{z}|^2$ to $|dw + \mu(z) e^{-2i\theta} d\bar{w}|^2$. Therefore, below we assume that $F = 0$. The condition that ds_μ^2 is a multiple of ds^2 translates to

$$1 + \mu = t\sqrt{E}, \quad 1 - \mu = t\sqrt{G},$$

for some $t = t(z) \in (0, \infty)$. Solving this system of equations, we obtain that $\mu(z)$ is real,

$$\mu = \frac{\sqrt{E} - \sqrt{G}}{\sqrt{E} + \sqrt{G}}.$$

Clearly, $|\mu| < 1$. Furthermore, $\lim_{z \rightarrow z_0} |\mu(z)| = 1$ if and only if

$$\lim_{z \rightarrow z_0} \frac{E(z)}{G(z)} \in \{0, \infty\}.$$

Thus, the condition that the measurable conformal structure $[ds^2]$ defined by ds_μ^2 is bounded is equivalent to the inequality

$$\|\mu\|_\infty < 1.$$

To summarize these computations: The correspondence $\mu \mapsto ds_\mu^2$, establishes an equivalence of Beltrami differentials μ with norm < 1 and bounded measurable conformal structures. Furthermore, if f is a quasiconformal map solving the *Beltrami equation* (23.1), then the measurable Riemannian metric $f^*(|dz|^2)$ is conformal to the metric ds_μ^2 .

The following fundamental theorem goes one step further; it will be used for the proof of nonexistence of exotic uniformly quasiconformal groups acting on \mathbb{S}^2 .

THEOREM 23.15 (Measurable Riemann Mapping Theorem). *For every measurable function $\mu(z)$ on a domain $\Omega \subset \mathbb{S}^2$ satisfying $\|\mu\|_\infty < 1$, there exists a quasiconformal homeomorphism $f : \Omega \rightarrow \Omega' \subset \mathbb{S}^2$ with the Beltrami differential μ . Equivalently, every bounded measurable conformal structure $[ds^2]$ on Ω is equivalent to the standard conformal structure on a domain $\Omega' \subset \mathbb{S}^2$ via a quasiconformal map $f : \Omega' \rightarrow \Omega$.*

HISTORICAL REMARK 23.16. In the case of smooth Riemannian metric ds^2 , a local version of this theorem was proven by Gauss, it is called *Gauss' theorem on isothermal coordinates*. In full generality it was established by Morrey [Mor38]. Modern proofs can be found, for instance, in [Ahl06] and [Leh87].

23.6. Proof of Tukia's theorem on uniformly quasiconformal groups

We are now ready to prove Tukia's theorem. Recall that the notion of *approximate continuity* was defined in Section 22.4.1.B.

THEOREM 23.17 (P. Tukia, [Tuk86]). *Let $G < \text{Homeo}(\mathbb{S}^n)$ be a uniformly quasiconformal group and $n \geq 2$. Assume also that μ is a G -invariant bounded measurable conformal structure on \mathbb{S}^n , which is approximately continuous at a conical limit point ξ of G . Then there exists a quasiconformal homeomorphism $f : \mathbb{S}^n \rightarrow \mathbb{S}^n$ which sends μ to the standard conformal structure on \mathbb{S}^n and conjugates G to a group of Moebius transformations.*

PROOF. As before, we will identify \mathbb{S}^n with $\widehat{\mathbb{R}}^n = \mathbb{R}^n \cup \infty$. We first explain Sullivan's proof of this theorem in the case $n = 2$ since it is easier and does not use the conical limit points assumption.

In view of the Measurable Riemann Mapping Theorem for \mathbb{S}^2 , the bounded measurable conformal structure μ on \mathbb{S}^2 is equivalent to the standard conformal structure μ_0 on \mathbb{S}^2 , i.e. there exists a quasiconformal map $f : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ which sends μ to μ_0 :

$$f^\bullet \mu_0 = \mu.$$

Since the quasiconformal group G preserves the conformal structure μ on \mathbb{S}^2 , it follows that the conjugate group $G^f = fGf^{-1}$ preserves the conformal structure μ_0 . Therefore, each $h \in G^f$ is a 1-quasiconformal homeomorphism of \mathbb{S}^2 , hence, a Moebius transformation, see Section 22.4. Thus, G^f acts as a group of Moebius automorphisms of the round sphere. This proves theorem for $n = 2$.

We now consider the general case. Without loss of generality, we may assume that the conical limit point ξ in the statement of the theorem is the origin in \mathbb{R}^n and (by conjugating G via an affine transformation if necessary) that $\mu(0) = \mu_0(0)$ is the standard conformal structure on \mathbb{R}^n . We will identify \mathbb{H}^{n+1} with the upper half-space \mathbb{R}_+^{n+1} . Define the point

$$\mathbf{e} = \mathbf{e}_{n+1} = (0, \dots, 0, 1) \in \mathbb{H}^{n+1}.$$

Let $\varphi : G \curvearrowright \mathbb{H}^{n+1}$ be the quasiaction, extending the action $G \curvearrowright \mathbb{S}^n$, see Theorem 22.38. (In the context of the proof of Theorem 23.1, which is our main goal, we can use the original quasiaction, of course.) Let (L, A) be the quasiisometry constants for this quasiaction. Every element $g \in G$ is a K -quasiconformal homeomorphism of \mathbb{S}^n for some $K < \infty$.

By the definition of a conical limit point, there exists a sequence $g_i \in G$ and a number $c \in \mathbb{R}$, such that

$$\lim_{i \rightarrow \infty} \varphi(g_i)(\mathbf{e}) = 0$$

and

$$(23.2) \quad d(\varphi(g_i)(\mathbf{e}), t_i \mathbf{e}) \leq c,$$

where d is the hyperbolic metric on \mathbb{H}^{n+1} and (t_i) is a sequence of positive numbers converging to zero.

Let γ_i denote the hyperbolic isometry given by

$$\mathbf{x} \mapsto t_i \mathbf{x}, \quad \mathbf{x} \in \mathbb{R}^{n+1}.$$

The composition

$$f_i := g_i^{-1} \circ \gamma_i$$

is a K -quasiconformal homeomorphism of \mathbb{S}^n . This homeomorphism has the (L, A) -quasiisometric extension

$$\varphi(g_i^{-1}) \circ \gamma_i$$

to the hyperbolic space \mathbb{H}^{n+1} . Using the inequality (23.2), we obtain the estimate

$$d(\varphi(g_i^{-1})\gamma_i(\mathbf{e}), \mathbf{e}) \leq Lc + A.$$

We claim that the sequence (f_i) contains a subsequence which converges to a K -quasiconformal mapping. One way to prove it is to appeal to the Coarse Arzela–Ascoli theorem (Proposition 8.34). We will use the Convergence Property of quasiconformal mappings instead.

Let T be an ideal hyperbolic triangle with the centroid \mathbf{e} and the set of ideal vertices $Z = \{\zeta_1, \zeta_2, \zeta_3\}$. By the Extended Morse Lemma (Lemma 11.105), the quasi-geodesic triangles $\phi(g_i^{-1})\gamma_i(T)$ are uniformly close to ideal geodesic triangles T_i in \mathbb{H}^{n+1} , such that the distances from centroids of T_i 's to the point \mathbf{e} are uniformly bounded (cf. the proof of Proposition 11.107). After passing to a subsequence (which we suppress) vertex sets Z_i of ideal triangles T_i converge to a three-point set $Z' \subset \mathbb{S}^n$. In particular, the K -quasiconformal maps f_i restricted to the set Z , subconverge to a bijection

$$Z \rightarrow Z' \subset \mathbb{S}^n.$$

Therefore, by Theorem 22.34, the sequence (f_i) subconverges to a quasiconformal mapping $f : \mathbb{S}^n \rightarrow \mathbb{S}^n$.

We now record the transformations of measurable conformal structures:

$$\mu_i := f_i^\bullet(\mu) = (\gamma_i)^\bullet(g_i)^{-1\bullet}(\mu) = (\gamma_i)^\bullet\mu,$$

since $g^\bullet(\mu) = \mu$. Putting it all together:

$$\mu_i(x) = \mu(\gamma_i x) = \mu(t_i x).$$

In other words, the measurable conformal structure μ_i is obtained by “zooming into” the point 0. Since, by the hypothesis of Theorem 23.17, $\xi = 0$ is an approximate

continuity point for μ , the sequence of functions (μ_i) converges (in measure) to the constant function $\mu_0 = \mu(0)$. This leads to the diagram:

$$\begin{array}{ccc} \mu & \xrightarrow{f_i} & \mu_i \\ & & \downarrow \\ \mu & \xrightarrow{f} & \mu_0 \end{array}$$

If we knew that the derivatives Df_i subconverge (in measure) to the derivative Df , then we would conclude that

$$f^\bullet \mu = \mu_0.$$

Then f would conjugate the group G (preserving μ) to a group G^f preserving μ_0 and, hence, acting conformally on \mathbb{S}^n .

However, derivatives of quasiconformal maps (in general), converge only in the "biting" sense (see [IM01]), which does not suffice for our purposes. Thus, we have to use a less direct argument below.

We claim that every element of G^f is 1-quasiconformal. Since it suffices to verify 1-quasiconformality locally, we restrict to a certain round ball $B = B(0, R)$ in \mathbb{R}^n . Since μ is approximately continuous at 0, for every $\epsilon \in (0, \frac{1}{2})$,

$$\|\mu_i(x) - \mu(0)\| < \epsilon,$$

away from a subset $W_i \subset B$ of measure $< \epsilon_i$, where $\lim_{i \rightarrow \infty} \epsilon_i = 0$. Thus, for $x \in B \setminus W_i$,

$$1 - \epsilon < \lambda_1(x) \leq \dots \leq \lambda_n(x) < 1 + \epsilon,$$

where $\lambda_k(x)$'s are the eigenvalues of the matrix $A_{i,x}$ of the normalized metric $\mu_i(x)$. It follows that

$$H_x(\mu_i) < \frac{\sqrt{1+\epsilon}}{\sqrt{1-\epsilon}} \leq \sqrt{1+4\epsilon} \leq 1+2\epsilon,$$

away from the subset W_i . For every $g \in G$, each map $h_i := f_i g f_i^{-1}$ is conformal with respect to the structure μ_i and, hence, $(1+2\epsilon)$ -quasiconformal away from the set W_i . Since measures of the subsets W_i converge to zero, we conclude, by the Strong Convergence Property (Theorem 22.35), that each $h := \lim h_i$ is $(1+2\epsilon)$ -quasiconformal. As this holds for arbitrary $\epsilon > 0$ and arbitrary $R > 0$, we conclude that each h is 1-quasiconformal (with respect to the standard conformal structure on \mathbb{S}^n). By Liouville's Theorem for quasiconformal mappings (Theorem 22.31), it follows that h is Moebius.

This proves that the group $G^f = f G f^{-1}$ consists of Moebius transformations and concludes the proof of Theorem 23.17. \square

REMARK 23.18. The key idea of the above proof is the *zooming* argument : The fraction appearing in the definition of the derivative of a function f of several real variables is nothing but a pre- and post-composition of f with some Moebius transformations. This argument will be used again in the proofs of Mostow and Schwartz Rigidity Theorems (sections 24.3) and 24.4).

Proof of Theorem 23.8. According to Proposition 23.13, there exists a G -invariant measurable conformal structure μ on the sphere \mathbb{S}^n . By Lemma 22.15, almost every point of \mathbb{S}^n is a point of approximate continuity of μ . Therefore, Theorem 23.17 applies and the action of G is conjugate to a Moebius action. \square

HISTORICAL REMARK 23.19. Theorem 23.17 was first stated by Gromov in [Gro81b] in the same volume where Sullivan proved it for $n = 2$, [Sul81]. Gromov’s sketch of the proof includes the zooming argument; this seems to be the first time when this argument appeared in the literature. However, Gromov did not have Theorem 22.35, which is the key analytical ingredient in the proof.

Proof of QI rigidity of groups acting geometrically on \mathbb{H}^{n+1} . We now can conclude the proof of Theorem 23.1. Let G be a finitely generated group quasi-isometric to \mathbb{H}^{n+1} , $n \geq 2$. Then there exists a quasi-action $\varphi : G \curvearrowright \mathbb{H}^{n+1}$ and this quasi-action extends to a uniformly quasiconformal action $\varphi_\infty : G \curvearrowright \mathbb{S}^n$. By Lemma 11.118, every point of \mathbb{S}^n is a conical limit point for this action. Since the quasi-action $G \curvearrowright \mathbb{H}^{n+1}$ is geometric, the action $G \curvearrowright \mathbb{S}^n$ is a *uniform convergence action*, see Theorem 11.135. Note that the action $G \curvearrowright \mathbb{S}^n$ is not necessarily faithful, but, by the same theorem, it has to have finite kernel. We will ignore the kernel and identify G with its image in the group of homeomorphisms of \mathbb{S}^n . By Proposition 23.13, there exists a G -invariant bounded measurable conformal structure μ on \mathbb{S}^n . By Lemma 22.15, almost every point of \mathbb{S}^n is a point of approximate continuity of μ . Lastly, by Theorem 23.8, the action $G \curvearrowright \mathbb{S}^n$ is quasiconformally conjugate to a Moebius action $G^f \curvearrowright \mathbb{S}^n$.

Being a uniform convergence group is a purely topological concept invariant under homeomorphic conjugation. Thus, the group G^f also acts on \mathbb{S}^n as a uniform convergence group. Recall that the Moebius group $Mob(\mathbb{S}^n)$ is isomorphic to the isometry group $Isom(\mathbb{H}^{n+1})$ via the extension map from hyperbolic space to the boundary sphere, see Corollary 4.21. Therefore, by applying Theorem 11.132, we conclude that the isometric action $G^f \curvearrowright \mathbb{H}^{n+1}$ is again geometric. It follows that the group G admits a geometric action on \mathbb{H}^{n+1} , which finishes the proof of Theorem 23.1. \square

23.7. QI rigidity for surface groups

Note that the proof of Tukia’s theorem given above fails in the case $n = 1$, i.e. for groups G quasi-isometric to the hyperbolic plane. However, Theorem 11.135 still implies that G acts on \mathbb{S}^1 as a uniform convergence group. It was proven as a result of the combined efforts of Tukia, Gabai, Casson and Jungreis in 1988–1994 (see [Tuk88, Gab92, CJ94]) that every uniform convergence group acting on \mathbb{S}^1 is isomorphic to a Fuchsian group, i.e. a discrete cocompact subgroup of $Isom(\mathbb{H}^2)$. Below we outline an alternative argument, which relies, however, on Thurston’s Geometrization Conjecture for 3-dimensional manifolds/Perelman’s Theorem (see [KL08] and [BBB⁺10] for the detailed proofs). For the required background (related to the statement of Thurston’s Geometrization Conjecture/Perelman’s Theorem) we refer the reader to [Sco83], [Kap01] and [Thu97].

THEOREM 23.20. *If a group G is QI to the hyperbolic plane then G admits a geometric action on \mathbb{H}^2 .*

PROOF. Let $\tilde{M} \subset Trip(\mathbb{S}^1)$ denote the set of positively oriented ordered triples of distinct points on \mathbb{S}^1 , i.e. points ξ_1, ξ_2, ξ_3 in \mathbb{S}^1 which appear in the counter-clockwise order on the circle. Thus, \tilde{M} is a connected 3-dimensional manifold, an open subset of $\mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$. The group G acts as a uniform convergence group on \mathbb{S}^1 ; i.e. it acts properly discontinuously and cocompactly on $Trip(\mathbb{S}^1)$; in particular,

the restricted action $G \curvearrowright \tilde{M}$ is also properly discontinuous and cocompact. (See Theorem 11.135.)

LEMMA 23.21. *If $g \in G$ fixes a point in \tilde{M} then it fixes the entire \tilde{M} .*

PROOF. Assume that $g \in G$ fixes three distinct points ξ_1, ξ_2, ξ_3 in \mathbb{S}^1 . In particular, g preserves each component of $\mathbb{S}^1 \setminus \{\xi_1, \xi_2, \xi_3\}$. These components are arcs $\alpha_i, i = 1, 2, 3$. Since g fixes points ξ_i , it also preserves orientation on each α_i . Proper discontinuity of the action $G \curvearrowright \tilde{M}$ implies that the element g has finite order. We claim that g fixes each arc α_i pointwise. We identify each α_i with \mathbb{R} ; then $g : \mathbb{R} \rightarrow \mathbb{R}$ is an orientation-preserving homeomorphism of finite order. Pick a point $x \in \mathbb{R}$ not fixed by g and suppose that $y = g(x) > x$. Then, since g preserves orientation, $g(y) > y$; similarly, $g^i(x) > g^{i-1}(x)$ for every $i \in \mathbb{Z}$. Thus, g cannot have finite order. Contradiction. The same argument applies if $y < x$. \square

Let, therefore, \bar{G} denote the quotient of G by the (finite) kernel of the action $G \curvearrowright \mathbb{S}^1$. According to Lemma 23.21, the group \bar{G} acts freely on \tilde{M} .

LEMMA 23.22. *\tilde{M} is homeomorphic to $\mathbb{H}^2 \times \mathbb{S}^1$.*

PROOF. Given a triple $\xi = (\xi_1, \xi_2, \xi_3) \in \tilde{M}$ of distinct points in \mathbb{S}^1 ; we let T_ξ denote the a unique ideal hyperbolic triangle with ideal vertices $\xi_i, i = 1, 2, 3$. Let p_ξ denote the center of this triangle, i.e. the center of the inscribed circle.

Clearly, the map $\xi \rightarrow p_\xi$ is continuous as a map $\tilde{M} \rightarrow \mathbb{H}^2$. Furthermore, let ρ_i denote the geodesic rays emanating from p_ξ and asymptotic to $\xi_i, i = 1, 2, 3$. These rays meet at the angles equal to $2\pi/3$ at the points p_ξ . Thus, the ray ρ_1 uniquely determines the rays ρ_2, ρ_3 (since the triple ξ is positively oriented). Let v_ξ be the derivative of ρ_1 at p_ξ . Thus, we obtain a continuous map

$$c : \tilde{M} \rightarrow U\mathbb{H}^2, \quad c : \xi = (\xi_1, \xi_2, \xi_3) \mapsto v_\xi \in T_p\mathbb{H}^2,$$

where $U\mathbb{H}^2$ is the unit tangent bundle of \mathbb{H}^2 . The map c also has a continuous inverse: Given $(p, v) \in U\mathbb{H}^2, v \in T_p\mathbb{H}^2$, we let ρ_1 be the geodesic ray emanating from p with the derivative v . From this ray ρ_1 we construct rays ρ_2, ρ_3 (meeting ρ_1 at angles $2\pi/3$) and, therefore, the points $\xi_i, i = 1, 2, 3 \in \mathbb{S}^1$. Since \mathbb{H}^2 is contractible, the unit tangent bundle $U\mathbb{H}^2$ is trivial and, hence, \tilde{M} is homeomorphic to $U\mathbb{H}^2 \cong \mathbb{H}^2 \times \mathbb{S}^1$. \square

In particular, $\pi_i(\tilde{M}) = 0, i \geq 2$, and $\pi_1(\tilde{M}) \cong \mathbb{Z}$. We now consider the quotient $M = \tilde{M}/\bar{G}$. Since the action $\bar{G} \curvearrowright \tilde{M}$ is free, properly discontinuous and cocompact, M is a compact 3-dimensional manifold. Furthermore, $C = \pi_1(\tilde{M}) < \pi_1(M)$ is a normal infinite cyclic subgroup and

$$\bar{G} \simeq \pi_1(M)/C.$$

Hence, the normal subgroup C has infinite index in $\pi_1(M)$. Since $\pi_i(\tilde{M}) = 0, i \geq 2$, the manifold M also has trivial homotopy groups $\pi_i(M), i \geq 2$, i.e. the manifold M is aspherical.

We next review, briefly, the classification of closed 3-dimensional manifolds given by Perelman's Geometrization Theorem (Thurston's Conjecture). The description of closed connected oriented 3-dimensional manifolds starts with the connected sum decomposition of a closed 3-manifold into *prime 3-manifolds*:

$$M = M_1 \# \dots \# M_k,$$

where each manifold M_i does not admit a non-trivial connected sum decomposition. If $M \neq M_1$, then each 2-sphere, along which M splits as a connected sum, defines a non-trivial element of $\pi_2(M)$. Thus, in our case, $M = M_1$.

We now consider the case of (closed, connected and oriented) prime 3-manifolds. Every prime manifold M is either *geometric* or is obtained by gluing geometric manifolds along boundary tori and Klein bottles, which are *incompressible surfaces*.

1. If this decomposition is not void, then the manifold M is *Haken* and classification of Haken manifolds was known before Perelman, primarily, due to work of Waldhausen, Jaco, Shalen, Johannson and Thurston.

2. Otherwise, if this secondary decomposition of M is void, then M is geometric and we explain below what this means.

There are eight types of closed 3-dimensional geometric manifolds, they are homeomorphic to quotients X/Γ of certain simply-connected homogeneous 3-dimensional Riemannian manifolds X . The groups Γ are discrete subgroups of $\text{Isom}(X)$, acting on X freely and cocompactly.

The list of homogeneous manifolds X is:

$$\mathbb{H}^3, \mathbb{E}^3, \mathbb{S}^3, \mathbb{H}^2 \times \mathbb{R}, \mathbb{S}^2 \times \mathbb{R}, \widetilde{SL}(2, \mathbb{R}), Nil_3, Sol_3.$$

Note that the first five of these homogeneous manifolds are symmetric spaces (three of which have nonpositive curvature); the remaining three are Lie groups equipped with left-invariant Riemannian metrics.

In case when $\pi_2(M) = \pi_3(M) = 0$, as in our situation, the manifold X cannot be \mathbb{S}^3 and $\mathbb{S}^2 \times \mathbb{R}$. This leaves only the spaces X isometric to:

In the case $X \cong \mathbb{H}^3$, the quotient manifold $M = X/\Gamma$ is hyperbolic, and, hence, its fundamental group $\pi_1(M)$ is Gromov-hyperbolic. In particular, $\pi_1(M)$ cannot contain a normal infinite cyclic subgroup, see Section 11.14. This excludes hyperbolic manifolds ($X \cong \mathbb{H}^3$). Similarly, every cocompact lattice Γ acting on Sol_3 is isomorphic to the semidirect product

$$\Gamma \cong \mathbb{Z}^2 \rtimes_A \mathbb{Z},$$

where the matrix $A \in GL(2, \mathbb{Z})$ has both eigenvalues different from ± 1 . This shows that Γ cannot contain a normal infinite cyclic subgroup. Thus, *Sol*-manifolds are also excluded.

Closed manifolds M homeomorphic to quotients of the remaining homogeneous spaces ($\mathbb{H}^2 \times \mathbb{R}, \widetilde{SL}(2, \mathbb{R}), \mathbb{E}^3, Nil_3$) have an important common feature, they are *Seifert manifolds*. Fundamental groups of all aspherical closed Seifert manifolds M admit short exact sequences:

$$(23.3) \quad 1 \rightarrow \mathbb{Z} \rightarrow \pi_1(M) \xrightarrow{\psi} F \rightarrow 1,$$

where F 's are groups acting faithfully and geometrically either on the Euclidean plane or on the hyperbolic plane. The former occurs in the case of the geometries $X = \mathbb{E}^3$ and $X = Nil$. In both cases, the group $\pi_1(M)$ is amenable; thus, all its quotients are amenable as well. However, a group G quasi-isometric to the hyperbolic plane cannot be amenable. This leaves only the cases $X \cong \mathbb{H}^2 \times \mathbb{R}$ and $X \cong \widetilde{SL}(2, \mathbb{R})$. The infinite cyclic normal subgroup $C \triangleleft \pi_1(M)$ described above, projects to a normal subgroup of the group F . Since the latter cannot have any non-trivial normal cyclic subgroups (Corollary 12.21), the group C has to be contained in the kernel of the homomorphism ψ in the sequence (23.3). We

conclude, therefore, that

$$\bar{G}/\Phi \simeq F,$$

where $\Phi \simeq \text{Ker}(\psi)/C$ is a finite cyclic group. Thus, the group \bar{G} (and, hence G) admits a geometric action on the hyperbolic plane, as required.

REMARK 23.23. With a bit more work, one shows that $C = \text{Ker}(\psi)$, and, hence, $F \simeq \bar{G}$. Furthermore, one verifies that $X \cong \widetilde{SL}(2, \mathbb{R})$.

We are now done with the case when the manifold M itself is geometric. It remains to rule out the case when M is obtained by gluing geometric 3-manifolds along their boundary tori and Klein bottles. Such a manifold M is necessarily Haken and, hence, Seifert (since $\pi_1(M)$ contains an infinite cyclic normal subgroup): A proof of this theorem can be found for instance in Hempel's book [Hem78]. An alternative to this reference is to argue that existence of a non-trivial infinite cyclic normal subgroup of $\pi_1(M)$ implies that the manifold M is obtained by gluing only Seifert manifolds (as hyperbolic ones are excluded by the same argument we used to rule out the entire M from being hyperbolic). Then, similarly to the proof in Hempel's book, one argues that the gluing has to preserve (up to isotopy) *Seifert fibrations* of the geometric pieces and, hence, the manifold M itself is Seifert. \square

COROLLARY 23.24. *The class of fundamental groups of closed surfaces is QI rigid.*

PROOF. Suppose that S is a closed connected surface. Since we are interested in VI invariance, we can assume that S is oriented. If $S = \mathbb{S}^2$, then its fundamental group is obviously QI rigid. For surfaces of genus ≥ 2 , QI rigidity follows from Theorem 23.20. Lastly, suppose that $S = T^2$, is the torus. Then any group G which is QI to $\pi_1(S)$ is virtually abelian of rank 2, see Theorem 16.26. \square

Quasiisometries of non-uniform lattices in \mathbb{H}^n

Suppose that G is either a Lie group or a finitely generated group and $\Gamma \leq G$ is a finitely generated subgroup. For each element of the commensurator $g \in \text{Comm}_G(\Gamma)$, the Hausdorff distance between Γ and $g\Gamma g^{-1}$ is finite. Therefore, g defines a quasiisometry $q = q_g : \Gamma \rightarrow \Gamma$, which sends $\gamma \in \Gamma$ to an element $\gamma' \in \Gamma$ nearest to $g\gamma g^{-1}$.

The main goal of this chapter is to prove a converse to this elementary observation, as well as QI rigidity, for non-uniform lattices in $PO(n, 1)$, $n \geq 3$. Along the way, we give a proof of the Mostow Rigidity Theorem.

THEOREM 24.1 (R. Schwartz [Sch96b]). *Let $\Gamma < G = PO(n, 1)$ be a non-uniform lattice, $n \geq 3$. Then:*

(a) *For each quasiisometry $f : \Gamma \rightarrow \Gamma$ there exists $g \in \text{Comm}_G(\Gamma)$, defining a quasiisometry q_g , which is within finite distance from f . The distance between these two quasiisometries depends only on Γ (and its word-metric) and on the quasiisometry constants of f .*

(b) *Suppose that $\Gamma, \Gamma' < G$ are non-uniform lattices quasiisometric to each other. Then there exists an isometry $g \in \text{Isom}(\mathbb{H}^n)$ such that the groups Γ' and $g\Gamma g^{-1}$ are commensurable¹.*

(c) *Suppose that Γ' is a finitely generated group which is quasiisometric to a non-uniform lattice Γ as above. Then the groups Γ, Γ' are virtually isomorphic.*

Note that the above theorem fails in the case of the hyperbolic plane (except for the last part). Indeed, every free group F_r of rank ≥ 2 can be realized as a non-uniform lattice Γ acting on \mathbb{H}^2 . In view of the thick-thin decomposition (see Section 12.6.3) of the hyperbolic surface $M = \mathbb{H}^2/\Gamma$, Γ contains only finitely many Γ -conjugacy classes of maximal parabolic subgroups: Every such class corresponds to a component of $M_{thin} = M \setminus M_c$. Suppose now that $r \geq 3$. Then there are atoroidal automorphisms ϕ of F_r , such that for every non-trivial cyclic subgroup $C < F_r$ and every m , $\phi^m(C)$ is not conjugate to C , see e.g. [BFH97]. Therefore, such ϕ cannot send parabolic subgroups of Γ to parabolic subgroups of Γ . Hence, the quasiisometry of F_r given by ϕ cannot extend to a quasiisometry $\mathbb{H}^2 \rightarrow \mathbb{H}^2$. It follows that Part (a) fails for $n = 2$. Similarly, one can show that Part (b) also fails, since commensurability preserves arithmeticity and there are both arithmetic and non-arithmetic lattices in $\text{Isom}(\mathbb{H}^2)$. All these lattices are virtually free, hence, virtually isomorphic.

Our proof of Theorem 24.1 mostly follows [Sch96b].

¹See Definition 5.17.

24.1. Coarse topology of truncated hyperbolic spaces

Suppose that $\Gamma < \text{Isom}(\mathbb{H}^n)$ is a non-uniform lattice. In Section 12.6.3 we defined the *truncated hyperbolic space* $\Omega \subset \mathbb{H}^n$ associated with Γ . The space Ω is a certain manifold with boundary; its boundary components are *peripheral horospheres* Σ_j . We equip the truncated hyperbolic space Ω with the path-metric $d = d_\Omega$, induced by the restriction of the Riemannian metric of \mathbb{H}^n to Ω :

$$d(x, y) = \inf_p \text{length}(p)$$

where the infimum is taken over all the paths p in Ω connecting x to y . The metric d is invariant under Γ and, since the quotient Ω/Γ is compact, (Ω, d) is quasiisometric to the group Γ . We will use the notation dist for the hyperbolic distance function in \mathbb{H}^n .

LEMMA 24.2. *The identity map $\iota : (\Omega, d) \rightarrow (\Omega, \text{dist})$ is 1-Lipschitz and uniformly proper.*

PROOF. If p is a path in Ω , then p has the same length with respect to the metrics d and dist . This immediately implies that ι is 1-Lipschitz. Uniform properness follows from the fact that the group Γ acts geometrically on both (Ω, d) , (Ω, dist) and that the map ι is Γ -equivariant, see Lemma 8.43. \square

LEMMA 24.3. *The restriction of d to each peripheral horosphere $\Sigma \subset \partial\Omega$ equals the Riemannian distance function defined by the restriction of the hyperbolic Riemannian metric to Σ . In particular, (Σ, d_Σ) is isometric to the Euclidean space \mathbb{E}^{n-1} .*

PROOF. Without loss of generality, we may assume that (in the upper half-space model of \mathbb{H}^n), $\Sigma = \{(x_1, \dots, x_n) : x_n = 1\}$. Hence,

$$\Omega \subset \{(x_1, \dots, x_n) : 0 < x_n \leq 1\}.$$

The hyperbolic Riemannian metric restricted to Σ equals the flat metric on Σ . Therefore, it is enough to show that for every path p in Ω , connecting points $x, y \in \Sigma$, there exists a path q in Σ (still, connecting x to y), such that $\text{length}(q) \leq \text{length}(p)$. Consider the vertical projection

$$\pi : \Omega \rightarrow \Sigma, \quad \pi(x_1, \dots, x_{n-1}, x_n) = (x_1, \dots, x_{n-1}, 1).$$

According to Exercise 4.62, $\|d\pi\| \leq 1$ (with respect to the hyperbolic metric). Therefore, setting $q := \pi \circ p$, we obtain

$$\text{length}(q) \leq \text{length}(p). \quad \square$$

LEMMA 24.4. *For every horoball $B \subset \mathbb{H}^n$, the R -neighborhood $\mathcal{N}_R(B)$ of B in \mathbb{H}^n is also an open horoball $B' \subset \mathbb{H}^n$.*

PROOF. We again work in the upper half-space model where

$$B = \{(x_1, \dots, x_n) : x_n > 1\}.$$

We let Σ denote the boundary of B and let $\pi : \mathbb{H}^n \setminus B \rightarrow \Sigma$ denote the vertical projection as in the proof of the previous lemma. We leave it to the reader to check that

$$\text{dist}(x, \Sigma) = \text{dist}(x, \pi(x)).$$

It follows, in view of Exercise 4.14, that $\mathcal{N}_R(B)$ is the horoball

$$\{(x_1, \dots, x_n) : x_n > e^{-R}\}. \quad \square$$

We refer the reader to Section 9.6 for the notion of *coarse separation, deep/shallow components* and *inradii* of shallow components, used below. The following lemma is the key for distinguishing the case of the hyperbolic plane from the higher-dimensional hyperbolic spaces (of dimension ≥ 3):

LEMMA 24.5. *Let Ω is a truncated hyperbolic space of dimension ≥ 3 . Then each peripheral horosphere $\Sigma \subset \Omega$ does not coarsely separate Ω .*

PROOF. Let $B \subset \mathbb{H}^n$ denote the horoball bounded by Σ . For each R , the R -neighborhood $\mathcal{N}_R(B)$ of B in \mathbb{H}^n is again a horoball B'_R . We claim that B'_R does not separate Ω . Indeed, the horoball B' does not separate \mathbb{H}^n . Therefore, for each pair of points $x, y \in \Omega \setminus B'_R$, there exists a piecewise-geodesic path p connecting them within $\mathbb{H}^n \setminus B'_R$. If the path p is entirely contained in Ω , we are done. Otherwise, we subdivide p into finitely many subpaths, each of which is either contained in Ω or connects a pair of points in the boundary of one of the complementary horoballs $B_j \subset \mathbb{H}^n \setminus \Omega$.

According to Lemma 4.60, the intersection of B'_R with $\Sigma_j = \partial B_j$ is isometric to a metric ball in the Euclidean space (Σ_j, d) .

Note that a round ball cannot separate \mathbb{R}^{n-1} , provided that $n - 1 \geq 2$. Thus, we can replace $p_j = p \cap B_j$ with a new path p'_j which connects the end-points of p_j within the complement $\Sigma_j \setminus B'_R$. By making these replacements for each j , we get a new path q connecting x to y within $\Omega \setminus B'_R$. Therefore, B'_R does not separate Ω .

We are now ready to show that Σ cannot coarsely separate (Ω, d) . We will use the notation $\mathcal{N}_{R,d}$ for the R -neighborhood with respect to the metric d . Suppose that for some R , $Y := \Omega \setminus \mathcal{N}_{R,d}(B)$ contains at least two deep components C_1, C_2 . Let $x_i \in C_i, i = 1, 2$. By the definition of a deep component of Y , there are continuous proper paths $\alpha_i : \mathbb{R}_+ \rightarrow C_i, \alpha_i(0) = x_i, i = 1, 2$. Thus,

$$\lim_{t \rightarrow \infty} d(\alpha_i(t), \Sigma) = \infty.$$

By Lemma 24.2, there exists T such that $y_i := \alpha_i(T) \notin B' = \mathcal{N}_{R,d}(B), i = 1, 2$. Therefore, as we proved in Lemma 24.5, we can connect y_1 to y_2 by a path in $\Omega \setminus B' \subset Y$. We conclude that $C_1 = C_2$, which is a contradiction. \square

EXERCISE 24.6. Show that Lemma 24.5 fails for $n = 2$. Hint: Use the fact that each Cayley graph of a free nonabelian group of finite rank has infinitely many ends.

In order to appreciate the difficulty of the proof of Proposition 24.8 below, we encourage the reader to do first the following exercise:

EXERCISE 24.7. Suppose that α is an isometry of $\mathbb{H}^n, n \geq 2$, such that

$$\text{dist}_{\text{Haus}}(\Omega, \alpha(\Omega)) \leq C.$$

Show that for each peripheral horosphere $\Sigma \subset \partial\Omega$, there exists a peripheral horosphere $\Sigma' \subset \partial\Omega$ satisfying

$$\text{dist}_{\text{Haus}}(\Sigma', \alpha(\Sigma)) \leq R,$$

where R depends only on C and not on Σ .

Now, suppose that Ω, Ω' are truncated hyperbolic spaces for lattices $\Gamma, \Gamma' < \text{Isom}(\mathbb{H}^n)$, and $f : \Omega \rightarrow \Omega'$ is an (L, A) -quasiisometry. Let Σ be a peripheral horosphere of Ω . Recall that we are assuming that $n \geq 3$.

PROPOSITION 24.8. *There exists a peripheral horosphere $\Sigma' \subset \partial\Omega'$ which is within finite Hausdorff distance from $f(\Sigma)$.*

PROOF. We start with the idea of the proof. Suppose that $h : M \rightarrow M'$ is a homeomorphism of compact connected n -dimensional manifolds with boundary, satisfying $H_{n-1}(M) = 0, H_{n-1}(M') = 0$. Then $h(\partial M) = \partial M'$. Of course, one can prove it in many ways (and without using our homological assumption), but the following, admittedly, somewhat silly, proof is a model of the proof of the proposition. We first note that no boundary component of M separates M , while a connected hypersurface which is not contained in the interior of M' has to separate M' (due to our homological assumptions). The proof below is a coarse version of this argument, where we use coarse separation arguments. The case 1 in this proof corresponds to the possibility that the entire boundary of M' is contained in one component of $M' \setminus f(\partial M)$, while the case 2 corresponds to the possibility that $f(\partial M)$ separates two boundary components of M' .

We now proceed with the actual proof. Since Ω'/Γ' is compact, there exists $D < \infty$, such that for every $x \in \Omega'$,

$$(24.1) \quad \text{dist}(x, \partial\Omega') \leq D.$$

The horosphere Σ , being isometric to \mathbb{R}^{n-1} (with respect to the metric d), has bounded geometry and is uniformly contractible. Therefore, according to Theorem 9.73, $f(\Sigma)$ coarsely separates \mathbb{H}^n . However $f(\Sigma)$ cannot coarsely separate Ω' , since f is a quasiisometry and Σ does not coarsely separate Ω (Lemma 24.5).

EXERCISE 24.9. Suppose that $Y \subset X$ coarsely separates subsets $X_1, X_2 \subset X$. Then, under any quasiisometry $f : X \rightarrow X'$, the set $f(Y)$ coarsely separates $f(X_1)$ from $f(X_2)$.

Let $r < \infty$ be such that $\mathcal{N}_r(f(\Sigma))$ separates \mathbb{H}^n into (two) deep components X_1, X_2 . We define a new truncated hyperbolic space

$$\Omega'' := \mathcal{N}_r(\Omega').$$

We will use the notation $B_j'' := B_j' \setminus \Omega''$ for the complementary horoballs of Ω'' .

Case 1. Suppose first that for each B_j' , the smaller horoball B_j'' is contained in the complementary region X_1 . In view of (24.1), for every $x \in \Omega'$, we obtain:

$$\text{dist}(x, X_1) \leq r + D,$$

since the hyperbolic distance from x to some point of Ω'' is at most $D + r$. Furthermore, every $x \in \mathbb{H}^n \setminus \Omega'$ is within a distance $\leq r$ from Ω'' . Therefore, for every $x \in \mathbb{H}^n$,

$$\text{dist}(x, X_1) \leq 2r + D.$$

In particular, if we take any point $x \in X_2$, there exists a path p of length $\leq 2r + D$ connecting it to X_1 . This path has to cross the neighborhood $\mathcal{N}_r(f(\Sigma))$ separating X_1 from X_2 . Therefore, the entire set X_2 is *shallow*: It is contained within distance $\leq 2r + D$ from $\mathcal{N}_r(f(\Sigma))$. This contradicts the property that the set X_2 is deep.

Similarly (renaming X_1 and X_2), we rule out the possibility that all horoballs B_j'' are contained in X_2 .

Case 2. There are two complementary horoballs B'_1, B'_2 of Ω' such that

$$B''_1 \subset X_1, \quad B''_2 \subset X_2.$$

Set $\Sigma'_i := \partial B'_i, i = 1, 2$. If both intersections

$$T'_i := \Sigma'_i \cap X_i, i = 1, 2,$$

contain points which are arbitrarily far from $f(\Sigma)$, then $f(\Sigma)$ coarsely separates Ω' , which is again a contradiction. Therefore, say, for $i = 1$, there exists $r' < \infty$ such that $\Sigma' := \Sigma'_1$ satisfies

$$(24.2) \quad \Sigma' \subset \mathcal{N}_{r'}(f(\Sigma)).$$

Our goal is to show that $f(\Sigma) \subset \mathcal{N}_R(\Sigma')$ for some $R < \infty$.

The inclusion (24.2) implies that the nearest-point projection $\Sigma' \rightarrow f(\Sigma)$ defines a quasiisometric embedding $h : \Sigma' \rightarrow \Sigma$, see Exercise 8.12. However, Lemma 10.84 proves that a quasiisometric embedding between two Euclidean spaces of the same dimension is a quasiisometry. Thus, there exists $R' < \infty$ such that $f(\Sigma) \subset \mathcal{N}_{R'}(\Sigma')$. Hence, $f(\Sigma)$ is Hausdorff-close to Σ' . □

EXERCISE 24.10. Show that the horosphere Σ' in Proposition 24.8 is unique.

We note that there are alternative proofs of Proposition 24.8, which use asymptotic cones instead of coarse topology; see for instance, [KL97] or [BDM09] (Theorem 25.40 in the next chapter).

We now improve Proposition 24.8 and establish uniform control on the distance from $f(\Sigma)$ to the boundary horosphere $\Sigma' \subset \Omega'$ in this proposition.

LEMMA 24.11. *In Proposition 24.8, for all peripheral horospheres $\Sigma \subset \partial\Omega$,*

$$\text{dist}_{\text{Haus}}(f(\Sigma), \Sigma') \leq c(L, A),$$

where $c(L, A)$ is independent of Σ .

PROOF. The proof is by inspection of the arguments in the proof of Proposition 24.8. First of all, the constant r depends only on the quasiisometry constants of the mapping f and the uniform geometry/uniform contractibility bounds for \mathbb{R}^{n-1} and \mathbb{H}^n . The *inradii* of the *shallow* components of

$$\Omega' \setminus \mathcal{N}_r(f(\Sigma))$$

again depend only on the above data. Therefore, there exists a uniform constant r' such that one of the horospheres Σ'_1 or Σ'_2 in the proof of Proposition 24.8 is contained in $\mathcal{N}_{r'}(f(\Sigma))$. Finally, an upper bound on R' such that

$$\mathcal{N}_{R'}(\text{Image}(h)) = \Sigma'$$

(coming from Lemma 10.84) again depends only on the quasiisometry constants of the projection $h : \Sigma' \rightarrow \Sigma$. □

REMARK 24.12. Proposition 24.8 and Lemma 24.11 combine into the *Quasi-flat Lemma* from [Sch96b], §3.2. This lemma can be seen as a version of the Morse Lemma 11.40 for truncated hyperbolic spaces. The spaces Ω, Ω' are, in fact, *relatively hyperbolic* in the strong sense. See Section 11.26 for further details.

24.2. Hyperbolic extension

Let $\Omega, \Omega' \subset \mathbb{H}^n$ be truncated hyperbolic spaces ($n \geq 3$) of lattices $\Gamma, \Gamma' < \text{Isom}(\mathbb{H}^n)$ and let C, C' denote the sets whose elements are peripheral horospheres of Ω, Ω' respectively. The main result of this section is

THEOREM 24.13 (Horoball QI extension theorem). *Each quasiisometry $f : \Omega \rightarrow \Omega'$ admits a quasiisometric extension $\tilde{f} : \mathbb{H}^n \rightarrow \mathbb{H}^n$. Moreover, the extension \tilde{f} satisfies the following equivariance property:*

Suppose that $f : X \rightarrow X'$ is quasiequivariant with respect to an isomorphism

$$\rho : \Gamma \rightarrow \Gamma'.$$

Then the extension \tilde{f} is also ρ -quasiequivariant.

PROOF. By Lemma 24.11, for every peripheral horosphere $\Sigma \subset \Omega$ there exists a peripheral horosphere Σ' of Ω' such that $\text{dist}_{\text{Haus}}(f(\Sigma), \Sigma') \leq c < \infty$, where c depends only on the quasiisometry constants of f . By the uniqueness of the horosphere Σ' , we obtain a map

$$(24.3) \quad \theta : C \rightarrow C', \theta : \Sigma \mapsto \Sigma',$$

which is ρ -equivariant, provided that f was ρ -quasiequivariant. We will use the notation B' for the horoball bounded by Σ' .

We first alter f on $\partial\Omega$ by postcomposing $f|_{\Sigma}$ with the nearest-point projection to Σ' for every $\Sigma \in C$. The new map is again a quasiisometry, since its distance from f is finite. The modification clearly preserves the ρ -quasiequivariance. We retain the notation f for the new quasiisometry, which now satisfies

$$f(\Sigma) \subset \Sigma', \forall \Sigma \in C.$$

We construct an extension $\tilde{f} : B \rightarrow B'$ of $f|_{\Sigma}$ for each complementary horoball $B \subset \mathbb{H}^n \setminus \Omega$ as follows.

For points $x \in \Sigma, x' \in \Sigma'$ and $t \in \mathbb{R}_+$ we define $x_t \in B, x'_t \in B'$, so that the maps

$$t \mapsto x_t, t \mapsto x'_t, t \in \mathbb{R}_+$$

are geodesic rays of the origin x , respectively x' , asymptotic to the centers of the horoballs B , respectively, B' . Of course, every point $y \in B$ has the form $y = x_t$ for unique $x \in \Sigma, t > 0$. Then we define the extension $\tilde{f} : B \rightarrow B'$ by the formula:

$$x_t \mapsto x'_t, \quad x' = f(x), x \in \Sigma.$$

By construction, this extension is quasiequivariant if f is.

We will now verify that this extension is coarsely Lipschitz. Since being coarse Lipschitz is a local property, it suffices to show that for each horoball B , the map $\tilde{f} : B \rightarrow B'$ is coarse Lipschitz. By composing \tilde{f} with isometries of \mathbb{H}^n , the problem reduces to the case when $B = B'$ is given by the inequality $x_n \geq 1$ (in the upper half-space model of \mathbb{H}^n). It suffices to consider points $z, w \in B$ within unit distance from each other.

We need to show that

$$\text{dist}(\tilde{f}(z), \tilde{f}(w)) \leq \text{Const}.$$

If $z, w \in B$ have the form $z = x_t, w = x_s$ for some $x \in \Sigma, s$ and t , then, by the construction,

$$\text{dist}(\tilde{f}(z), \tilde{f}(w)) = \text{dist}(z, w) = |t - s|.$$

Therefore, by the triangle inequality, the problem reduces to getting a uniform upper bound on the distances $\text{dist}(\tilde{f}(z), \tilde{f}(w))$ for points z and w belonging to the same horosphere $\Sigma_t \subset B$:

$$z = x_t, w = y_t, \quad x \in \Sigma, y \in \Sigma.$$

We will use the notation

$$\text{dist}_{\Sigma_t}(z, w)$$

for the distance between z and w computed with respect to the Riemannian distance function on Σ_t , where the Riemannian metric on Σ_t is the restriction of the hyperbolic Riemannian metric. In other words,

$$\text{dist}_{\Sigma_t}(z, w) = e^{-t}|z - w|.$$

It follows from Exercise 4.55 that

$$\text{dist}_{\Sigma_t}(z, w) \leq \epsilon := \sqrt{2(e-1)},$$

since we are assuming that $\text{dist}(z, w) \leq 1$. Accordingly,

$$\text{dist}_{\Sigma}(x, y) \leq \epsilon e^t.$$

Since $f : (\Omega, d_{\Omega}) \rightarrow (\Omega', d_{\Omega'})$ is (L, A) -coarse Lipschitz,

$$\text{dist}_{\Sigma}(f(x), f(y)) \leq e^t L \epsilon + A.$$

It follows that

$$d(\tilde{f}(z), \tilde{f}(w)) \leq \text{dist}_{\Sigma_t}(\tilde{f}(z), \tilde{f}(w)) \leq L \epsilon + A e^{-t} \leq L \epsilon + A.$$

This proves that the extension \tilde{f} is coarse Lipschitz in the horoball B and, hence, in the entire \mathbb{H}^n . The same argument applies to the extension \tilde{f}' of the coarse inverse f' to the mapping f . We leave it to the reader to verify that the inequality

$$d_{\Omega}(f' \circ f, \text{Id}) \leq A$$

implies

$$\text{dist}(\tilde{f}' \circ \tilde{f}, \text{Id}) \leq A.$$

Thus, \tilde{f} is a quasiisometry. □

Since \tilde{f} is a quasiisometry of \mathbb{H}^n , it admits a quasiconformal extension $h : \partial_{\infty} \mathbb{H}^n \rightarrow \partial_{\infty} \mathbb{H}^n$ (see Theorems 11.108 and 22.36). By Corollary 11.111, the homeomorphism h is ρ -equivariant, provided that f is quasiequivariant.

Let Λ, Λ' denote the sets of the centers of the peripheral horospheres of Ω, Ω' respectively. Since $f(\Sigma) = \Sigma'$ for every peripheral horosphere of Ω , continuity of the extension also implies that $h(\Lambda) = \Lambda'$.

24.3. Mostow Rigidity Theorem

The Mostow Rigidity Theorem that we will prove in this section was a precursor and inspiration for the Schwartz Rigidity Theorem. We prove this theorem first, since its proof is technically simpler and also illustrates some of the ideas behind Schwartz' proof. Our arguments are inspired by the ones of P. Tukia [Tuk85] and N. Ivanov [Iva96].

THEOREM 24.14 (Mostow Rigidity Theorem). *Suppose that $n \geq 3$, that $\Gamma, \Gamma' < \text{Isom}(\mathbb{H}^n)$ are lattices and $\rho : \Gamma \rightarrow \Gamma'$ is an isomorphism. Then ρ is induced by an isometry, i.e. there exists an isometry $\alpha \in \text{Isom}(\mathbb{H}^n)$ such that*

$$\alpha \circ \gamma = \rho(\gamma) \circ \alpha$$

for all $\gamma \in \Gamma$.

PROOF. Step 1. We first observe that Γ is uniform if and only if Γ' is uniform. Indeed, if Γ is uniform, it is Gromov–hyperbolic and, hence, cannot contain a noncyclic free abelian group. On the other hand, if Γ' is non-uniform then Corollary 12.28 implies that Γ' contains free abelian subgroups of rank $n - 1 > 1$.

LEMMA 24.15. *There exists a ρ -equivariant quasiisometry $\tilde{f} : \mathbb{H}^n \rightarrow \mathbb{H}^n$.*

PROOF. As in the proof of Theorem 24.1, we choose truncated hyperbolic spaces $\Omega \subset \mathbb{H}^n, \Omega' \subset \mathbb{H}^n$ for the lattices Γ and Γ' respectively. (If Γ is a uniform lattice, we take, of course, $\Omega = \Omega' = \mathbb{H}^n$.) Lemma 8.45 implies that there exists a ρ -quasiequivariant quasiisometry

$$f : \Omega \rightarrow \Omega'.$$

Therefore, according to Theorem 24.13, f extends to a ρ -equivariant quasiisometry $\tilde{f} : \mathbb{H}^n \rightarrow \mathbb{H}^n$. \square

REMARK 24.16. The most difficult part of the proof of Theorem 24.13 was to show that f sends peripheral horospheres uniformly close to peripheral horospheres. In the equivariant setting the proof is much easier: The homomorphism ρ sends maximal abelian subgroups of Γ to maximal abelian subgroups of Γ' . The stabilizers of peripheral horospheres are virtually \mathbb{Z}^{n-1} . Therefore, ρ sends stabilizers of peripheral horospheres to stabilizers of peripheral horospheres. From this, it is immediate that peripheral horospheres map uniformly close to peripheral horospheres.

Step 2. Let h denote the ρ -equivariant quasiconformal homeomorphism

$$h : \mathbb{S}^{n-1} \rightarrow \mathbb{S}^{n-1}$$

extending the quasiequivariant quasiisometry \tilde{f} . Our goal is to show that h is Moebius. We argue as in the proof of Theorem 24.31. We will identify \mathbb{S}^{n-1} with the extended Euclidean space $\mathbb{R}^{n-1} \cup \{\infty\}$. Accordingly, we will identify \mathbb{H}^n with the upper half-space. The key to the proof is the fact that h is differentiable almost everywhere on \mathbb{R}^{n-1} and that its Jacobian determinant is non-zero for almost every $z \in \mathbb{R}^{n-1}$. (In fact, we need only uncountably many points of differentiability, where the Jacobian determinant is non-zero.)

In Theorem 12.29 we proved that every point of \mathbb{S}^{n-1} is either a conical limit point of Γ or is a parabolic fixed point. Since Γ has only countably many parabolic elements and each has only one fixed point, almost every point of \mathbb{S}^{n-1} is a conical limit point of Γ . Hence, we find a conical limit point $z \in \mathbb{S}^{n-1} \setminus \{\infty\}$, which is a point of differentiability of h , where $J_z(h) \neq 0$. After applying a Moebius change of coordinates, we can assume that $z = h(z) = 0 \in \mathbb{R}^{n-1}$ and that $h(\infty) = \infty$.

The following proof is yet another version of the *zooming argument*. Let $L \subset \mathbb{H}^n$ be the vertical geodesic emanating from 0; pick a base-point $y_0 \in L$. Since z is a conical limit point, there is a sequence of elements $\gamma_i \in \Gamma$ such that

$$\lim_{i \rightarrow \infty} \gamma_i(y_0) = z$$

and

$$\text{dist}(\gamma_i(y_0), L) \leq \text{Const},$$

for each i . Let y_i denote the nearest-point projection of $\gamma_i(y_0)$ to L . Take the sequence of hyperbolic translations $T_i : y \mapsto t_i y$ with the axis L , such that $T_i(y_0) = y_i$. Then the sequence $k_i := \gamma_i^{-1} T_i$ lies in a compact $C \subset G = \text{Isom}(\mathbb{H}^n)$.

Now, the proof turns from mappings of the hyperbolic n -space to mappings of \mathbb{R}^{n-1} . We form a sequence of quasiconformal homeomorphisms

$$h_i(\mathbf{x}) := t_i^{-1} h(t_i \mathbf{x}) = T_i^{-1} \circ h \circ T_i(\mathbf{x}), \mathbf{x} \in \mathbb{R}^{n-1},$$

$$\lim_{i \rightarrow \infty} t_i = 0.$$

Since the mapping h is assumed to have invertible derivative at the origin, there is a linear transformation $A \in GL(n-1, \mathbb{R})$ such that

$$\lim_{i \rightarrow \infty} h_i(\mathbf{x}) = A\mathbf{x}$$

for all $\mathbf{x} \in \mathbb{R}^{n-1}$. Since $h(\infty) = \infty$, it follows that

$$\lim_{i \rightarrow \infty} h_i = A$$

pointwise on \mathbb{S}^{n-1} .

By construction, h_i conjugates the group $\Gamma_i := T_i^{-1} \Gamma T_i \subset \text{Mob}(\mathbb{S}^{n-1})$ into the group of Moebius transformations $\text{Mob}(\mathbb{S}^{n-1})$. We have

$$\Gamma_i = T_i^{-1} \Gamma T_i = (k_i^{-1} \gamma_i) \Gamma (k_i^{-1} \gamma_i)^{-1} = k_i^{-1} \Gamma k_i.$$

After passing to a subsequence, we can assume that

$$\lim_{i \rightarrow \infty} k_i = k \in \text{Isom}(\mathbb{H}^n).$$

Therefore the sequence of subsets $\Gamma_i \subset G = \text{Mob}(\mathbb{S}^{n-1})$ converges to $\Gamma_\infty := k^{-1} \Gamma k$; here the convergence is understood in the Chabauty topology on the set of closed subsets of G , see Section 2.4. For each sequence $\beta_i \in \Gamma_i$ which converges to some $\beta \in G$ we have

$$\lim_{i \rightarrow \infty} h_i \beta_i h_i^{-1} = A \beta A^{-1}.$$

Since the subgroup G is closed in $\text{Homeo}(\mathbb{S}^{n-1})$ (with respect to the topology of pointwise convergence, see Corollary 4.5), it follows that the limit $A \beta A^{-1}$ of the sequence of Moebius transformations $(h_i \beta_i h_i^{-1})$, is again a Moebius transformation. This shows that $A \beta A^{-1} \in G$, for each $\beta \in \Gamma_\infty$. Thus,

$$A \Gamma_\infty A^{-1} \subset G.$$

The subgroup $\Gamma_\infty < G$ is conjugate to the lattice Γ and, hence, it cannot have a finite orbit in \mathbb{S}^{n-1} , see Corollary 12.20. In particular, the Γ_∞ -orbit of ∞ is infinite, which implies that Γ_∞ contains an element γ such that $\gamma(\infty) \notin \{\infty, 0\}$.

LEMMA 24.17. *Suppose that $\gamma \in G = \text{Mob}(\mathbb{S}^{n-1})$ is such that $\gamma(\infty) \neq \infty, 0$, and that $A \in GL(n-1, \mathbb{R})$ is an element which conjugates γ to $A \gamma A^{-1} \in G$. Then A is a Euclidean similarity, i.e. it belongs to $\mathbb{R}_+ \times O(n-1)$.*

PROOF. Suppose that A is not a similarity. Let P be a hyperplane in \mathbb{R}^{n-1} which contains the origin 0 but does not contain $A \gamma^{-1}(\infty)$. Then $\gamma \circ A^{-1}(P)$ does not contain ∞ and, hence, is a round sphere S in \mathbb{R}^{n-1} . Since A is not a similarity, the image $A(S)$ is an ellipsoid, which is not a round sphere. Hence, the composition

$A\gamma A^{-1}$ does not send planes to round spheres and, therefore, it is not Moebius. This is a contradiction. \square

We conclude that the derivative of h at 0 is a similarity $A \in \mathbb{R}_+ \times O(n-1)$. Thus, h is conformal at a.e. point of \mathbb{R}^n . One option now is to use Liouville's theorem for quasiconformal maps (Theorem 22.31). Instead, we give a direct argument.

Step 3. We will be using the notation of the Step 2. Consider the quotient

$$Q = G \backslash \text{Homeo}(\mathbb{S}^{n-1})$$

consisting of the cosets $[\varphi] = \{g \circ \varphi : g \in G\}$. We equip this quotient with the quotient topology, where we endow $\text{Homeo}(\mathbb{S}^{n-1})$ with the topology of pointwise convergence. Since G is a closed subgroup in $\text{Homeo}(\mathbb{S}^{n-1})$, it follows that every point in Q is closed. (Actually, Q is Hausdorff, but we will not need this.) The group $\text{Homeo}(\mathbb{S}^{n-1})$ acts on Q by the formula

$$[\varphi] \mapsto [\varphi \circ g], g \in \text{Homeo}(\mathbb{S}^{n-1}).$$

It is clear from the definition of the quotient topology on Q , that this action is continuous, i.e. the map

$$Q \times \text{Homeo}(\mathbb{S}^{n-1}) \rightarrow Q$$

is continuous.

Since h is a ρ -equivariant homeomorphism, we have

$$[h] \circ \gamma = [h], \quad \forall \gamma \in \Gamma.$$

Recall that we have two sequences: $\gamma_i \in \Gamma$, $k_i \in G$, such that

$$\lim_{i \rightarrow \infty} k_i = k \in G.$$

We also have a sequence of dilations $T_i = \gamma_i \circ k_i$ (fixing the origin in \mathbb{R}^{n-1}). Furthermore,

$$\lim_{i \rightarrow \infty} h_i = A \in \mathbb{R}_+ \times O(n-1) \subset G,$$

where

$$h_i = T_i^{-1} \circ h \circ T_i.$$

Therefore

$$[h_i] = [h\gamma_i k_i] = [h] \circ k_i,$$

$$[1] = [A] = \lim_{i \rightarrow \infty} [h_i] = \lim_{i \rightarrow \infty} ([h_i] \circ k_i) = [h] \circ \lim_{i \rightarrow \infty} k_i = [h] \circ k.$$

(Recall that every point in Q is closed.) Thus, $[h] = [1] \circ k^{-1} = [1]$, which implies that h is in G , i.e. h is a Moebius transformation, which we now denote by α . Regarding α as an isometry of \mathbb{H}^n and taking into account that the map $\alpha : \mathbb{S}^{n-1} \rightarrow \mathbb{S}^{n-1}$ is ρ -equivariant, we conclude that the isometry

$$\alpha : \mathbb{H}^n \rightarrow \mathbb{H}^n$$

is also ρ -equivariant. The Mostow Rigidity Theorem follows. \square

24.4. Zooming in

We now return to the proof of the Schwartz Rigidity Theorem. In Section 24.2, given a quasiisometry $f : \Omega \rightarrow \Omega'$, we constructed its quasiisometric extension $\tilde{f} : \mathbb{H}^n \rightarrow \mathbb{H}^n$; the latter, in turn, has a quasiconformal extension $h : \partial_\infty \mathbb{H}^n \rightarrow \partial_\infty \mathbb{H}^n$. Our main goal is to show that h is Moebius. By the Liouville's theorem for quasiconformal mappings (Theorem 22.31), h is Moebius provided that it is 1-quasiconformal, i.e. for a.e. point $\xi \in \mathbb{S}^{n-1}$, the derivative $D_\xi h$ of h at ξ is a Euclidean similarity.

We will continue to work with the upper half-space model of the hyperbolic space \mathbb{H}^n .

PROPOSITION 24.18. *Suppose that h is not Moebius. Then there exist lattices $\Gamma_\infty, \Gamma'_\infty$ conjugate to Γ, Γ' respectively, with truncated hyperbolic spaces $\Omega_\infty, \Omega'_\infty$ and a quasiisometry $\tilde{F} : \mathbb{H}^n \rightarrow \mathbb{H}^n$, such that:*

1. $\tilde{F}(\Omega_\infty) = \Omega'_\infty$.
2. The extension A of \tilde{F} to $\partial_\infty \mathbb{H}^n = \mathbb{R}^{n-1} \cup \{\infty\}$ fixes the points $0, \infty$.
3. The mapping A is a linear map, but not a similarity.

PROOF. Our arguments follow the ones in the proof of the Mostow Rigidity Theorem. Since h is differentiable a.e. and is not Moebius, there is a point $\xi \in \mathbb{S}^{n-1}$ such that $D_\xi h$ exists, is invertible, but is not a similarity. Since the subset $\Lambda \subset \mathbb{S}^{n-1}$ consisting of fixed points of parabolic element of Γ is countable, we can assume that $\xi \notin \Lambda$, i.e. is not the center of a complementary horoball of Ω . By pre- and post-composing with isometries of \mathbb{H}^n we can assume that $\xi = 0 = h(\xi)$. We will use the notation

$$A = D_\xi h \in GL(n-1, \mathbb{R})$$

for the derivative of h at 0 .

Let $L \subset \mathbb{H}^n$ denote the vertical geodesic asymptotic to 0 . Since 0 is not the center of a complementary horoball of Ω , there exists a sequence

$$y_i \in L \cap \Omega, \lim_{i \rightarrow \infty} y_i = 0.$$

We now break the symmetry between the lattices Γ, Γ' and, instead of taking points in $\Omega' \cap L$, we take the images $y'_i = f(y_i) \in \Omega'$.

For each i there exists a hyperbolic isometry

$$T_i(y) = t_i y, \quad y \in \mathbb{H}^n, t_i > 0,$$

which maps y_1 to y_i . The compositions

$$\tilde{f}_i := T_i^{-1} \circ \tilde{f} \circ T_i$$

are uniform quasiisometries of \mathbb{H}^n . The quasiconformal extensions of these quasiisometries to \mathbb{R}^{n-1} are given by

$$h_i(x) = \frac{h(t_i x)}{t_i}, \quad x \in \mathbb{R}^{n-1}.$$

By the definition of the derivative of h at the origin,

$$(24.4) \quad \lim_{i \rightarrow \infty} h_i = A,$$

where the convergence above is uniform on compact subsets of \mathbb{R}^{n-1} . We now claim that the sequence of quasiisometries \tilde{f}_i coarsely subconverges to a quasiisometry of

\mathbb{H}^n . As in the proof of Theorem 23.17, there are two ways to argue. One argument is that the claim follows from the convergence property of the quasiconformal mappings h_i (in conjunction with the extension Theorem 22.38). Alternatively, the claim follows from the coarse Arzela-Ascoli theorem (Theorem 8.34), since the limit (24.4) forces the quasiisometries \tilde{f}_i to send a base-point $p \in \mathbb{H}^n$ to points $q_i \in \mathbb{H}^n$ satisfying

$$\sup_i \text{dist}(p, q_i) < \infty.$$

In either case, the sequence of quasiisometries (\tilde{f}_i) has a subsequence which coarsely converges to a quasiisometry of \mathbb{H}^n . Thus, without loss of generality we suppose that the sequence (\tilde{f}_i) itself converges. In view of the limit (24.4), we can take the linear extension \tilde{A} of the linear mapping A (defined in Exercise 22.39) as the (coarse) limit quasiisometry.

At this point, there is no reason for \tilde{A} to send Ω to a subset of \mathbb{H}^n within a finite Hausdorff distance from Ω' (after all, we were composing with the mappings $T_i^{\pm 1}$ which do not preserve Ω and Ω'). The reader who went through the proofs of Tukia's and Mostow's theorems probably already knows what to do: We need to compose the hyperbolic isometries $T_i^{\pm 1}$ of \mathbb{H}^n with suitable elements of the groups Γ and Γ' . Recall that the quotients Ω/Γ and Ω'/Γ' are compact. Therefore, there exist $R < \infty$ and sequences $\gamma_i \in \Gamma, \gamma'_i \in \Gamma'$ such that for all i ,

$$\text{dist}(\gamma_i(y_1), y_i) \leq R, \quad \text{dist}(\gamma'_i(y_1), y'_i) \leq R.$$

(This is where we are using the fact that $y_i \in \Omega$ and $y'_i \in \Omega'$.) Hence, both sequences

$$k_i := T_i^{-1} \circ \gamma_i, \quad k'_i := T_i^{-1} \circ \gamma'_i$$

belong to a compact subset in $\text{Isom}(\mathbb{H}^n)$. After passing to a subsequence (which we again suppress from our notation), we obtain

$$\lim_{i \rightarrow \infty} k_i = k \in \text{Isom}(\mathbb{H}^n), \quad \lim_{i \rightarrow \infty} k'_i = k' \in \text{Isom}(\mathbb{H}^n).$$

For $i \in \mathbb{N}$ we define truncated hyperbolic spaces

$$\Omega_i := T_i^{-1}\Omega = k_i \circ \gamma_i^{-1}\Omega = k_i\Omega,$$

and

$$\Omega'_i := T_i^{-1}\Omega' = k'_i\Omega',$$

for the lattices $\Gamma_i = k_i^{-1}\Gamma k_i$ and $\Gamma'_i = k'_i{}^{-1}\Gamma' k'_i$ respectively. By the definition of these truncated hyperbolic spaces, the quasiisometry \tilde{f}_i sends Ω_i to Ω'_i . Since (k_i) converges to k and (k'_i) converges to k' , we have limits (in the Chabauty topology on the set of closed subsets of $\text{Isom}(\mathbb{H}^n)$, Section 2.4)

$$\lim_{i \rightarrow \infty} \Gamma_i = \Gamma_\infty, \quad \lim_{i \rightarrow \infty} \Gamma'_i = \Gamma'_\infty.$$

Since the groups $\Gamma_\infty, \Gamma'_\infty$ are conjugate to the lattices Γ, Γ' respectively, they are lattices themselves. We leave it to the reader to verify that the sets

$$\Omega_\infty := k(\Omega), \quad \Omega'_\infty := k'(\Omega')$$

are truncated hyperbolic spaces for the lattices Γ_∞ and Γ'_∞ respectively and that

$$\lim_{i \rightarrow \infty} \Omega_i = \Omega_\infty, \quad \lim_{i \rightarrow \infty} \Omega'_i = \Omega'_\infty,$$

again, in the Chabauty topology. Since the sequence (\tilde{f}_i) coarsely converges to \tilde{A} , it follows that the affine map \tilde{A} defines a quasiisometry $\Omega_\infty \rightarrow \Omega'_\infty$ in the sense that the sets $\tilde{A}(\Omega_\infty)$ and Ω'_∞ are Hausdorff-close to each other. Since the lattice Γ is conjugate to Γ_∞ and Γ' is conjugate to Γ'_∞ , the proposition follows. \square

We are aiming for a contradiction, therefore, from now on, we rename Γ_∞ to Γ and Γ'_∞ to Γ' , etc. The situation when we have a linear mapping (that is not a similarity!) sending Λ to Λ' seems, at the first glance, impossible. Here, however, is an example:

EXAMPLE 24.19. Let $\Gamma := PSL(2, \mathbb{Z}[i]), \Gamma' := PSL(2, \mathbb{Z}[\sqrt{-2}])$ be Bianchi subgroups of $PSL(2, \mathbb{C})$. Then

$$\Lambda = \mathbb{Q}(i) \cup \{\infty\}, \quad \Lambda' = \mathbb{Q}(\sqrt{-2}) \cup \{\infty\}.$$

Take the real linear mapping $A : \mathbb{C} \rightarrow \mathbb{C}$ sending 1 to 1 and i to $\sqrt{-2}$. Then A is invertible, it is not a similarity, but $A(\Lambda) = \Lambda'$.

Thus, in order to get a contradiction, we have to exploit the fact that the linear map A we constructed is the quasiconformal extension of a quasiisometry \tilde{F} , $\tilde{F}(\Omega) \subset \Omega'$. We will show (Theorem 24.30) the following:

For every peripheral horosphere $\Sigma \subset \partial\Omega$ whose center is not ∞ , there exists a sequence of peripheral horospheres $\Sigma_k \subset \partial\Omega$ such that:

$$\text{dist}(\Sigma, \Sigma_k) \leq \text{Const}, \quad \lim_{k \rightarrow \infty} \text{dist}(\Sigma', \Sigma'_k) = \infty.$$

(We remind the reader that $\text{dist}(\cdot, \cdot)$ denotes the minimal distance between the horospheres and $\theta(\Sigma) = \Sigma', \theta(\Sigma_k) = \Sigma'_k, k \in \mathbb{N}$, see (24.3).)

Of course, this means that \tilde{F} cannot be coarse Lipschitz. We will prove the above statement by conjugating \tilde{F} by an inversion which interchanges a horosphere with the center at ∞ and the horospheres Σ, Σ' above. This will amount to replacing the linear map A (from Proposition 24.18) with an *inverted linear mapping*. Inverted linear mappings are defined and analyzed in the next section.

24.5. Inverted linear mappings

Let $A : \mathbb{R}^{n-1} \rightarrow \mathbb{R}^{n-1}$ be an (invertible) linear mapping. Recall that the inversion J in the unit sphere \mathbb{S}^{n-2} is given by the formula

$$J(\mathbf{x}) = \frac{\mathbf{x}}{|\mathbf{x}|^2}, \quad \mathbf{x} \in \mathbb{R}^{n-1}.$$

DEFINITION 24.20. An *inverted linear map* is the conjugate of an invertible linear map A by the inversion J , i.e. the composition

$$h := J \circ A \circ J, \quad h(\mathbf{x}) = \frac{|\mathbf{x}|^2}{|A\mathbf{x}|^2} A(\mathbf{x}).$$

We will introduce the notation

$$\phi = \phi_h = \frac{|\mathbf{x}|^2}{|A\mathbf{x}|^2}$$

and will refer to this function as the *nonlinear factor* of the inverted linear map h .

LEMMA 24.21. The function $\phi(\mathbf{x}) = \frac{|\mathbf{x}|^2}{|A\mathbf{x}|^2}$ is asymptotically constant, in the sense that

$$|\nabla\phi(\mathbf{x})| = O(|\mathbf{x}|^{-1}), \quad \|Hess(\phi(\mathbf{x}))\| = O(|\mathbf{x}|^{-2})$$

as $|\mathbf{x}| \rightarrow \infty$. Here $Hess(\phi(\mathbf{x}))$ stands for the Hessian of the function $\phi(\mathbf{x})$.

PROOF. The function ϕ is a rational function of zero degree, hence, its gradient is a rational function of the degree -1 , while every component of its Hessian is a rational function of degree -2 . \square

EXERCISE 24.22. The following are equivalent:

1. The function ϕ is constant.
2. The mapping h is linear.
3. The linear transformation A is a similarity.

EXERCISE 24.23. The mapping h is differentiable at 0 if and only if A is a similarity.

Since each inverted linear mapping h is, clearly, differentiable everywhere on $\mathbb{R}^{n-1} \setminus \{0\}$, it follows that h determines the origin 0 in the Euclidean space. Hence, h also determines its nonlinear factor ϕ_h (up to a scalar multiple). The next exercise also shows that h determines the origin in \mathbb{R}^n :

EXERCISE 24.24. Suppose that h is an inverted non-linear map with non-constant factor ϕ_h . Show that for each Euclidean hyperplane $P \subset \mathbb{E}^{n-1}$ not passing through the origin, $h(P)$ is not a Euclidean hyperplane. In contrast, show that h sends each linear subspace in \mathbb{R}^{n-1} to a linear subspace in \mathbb{R}^{n-1} .

COROLLARY 24.25. Fix a positive real number R , and let (\mathbf{v}_k) be a sequence diverging to infinity in \mathbb{R}^{n-1} . Then the sequence of maps

$$h_k(\mathbf{x}) := h(\mathbf{x} + \mathbf{v}_k) - h(\mathbf{v}_k)$$

subconverges (uniformly on the R -ball $B = B(0, R) \subset \mathbb{R}^{n-1}$) to an affine map, as $k \rightarrow \infty$.

PROOF. We have:

$$\begin{aligned} h(\mathbf{x} + \mathbf{v}_k) - h(\mathbf{v}_k) &= \phi(\mathbf{x} + \mathbf{v}_k)A(\mathbf{x} + \mathbf{v}_k) - \phi(\mathbf{v}_k)A(\mathbf{v}_k) = \\ &= \phi(\mathbf{x} + \mathbf{v}_k)A(\mathbf{x}) + (\phi(\mathbf{x} + \mathbf{v}_k) - \phi(\mathbf{v}_k))A(\mathbf{v}_k). \end{aligned}$$

Since $\phi(\mathbf{y})$ is asymptotically constant, $\lim_{k \rightarrow \infty} \phi(\mathbf{x} + \mathbf{v}_k)A(\mathbf{x}) = c \cdot A(\mathbf{x})$ for some constant c (uniformly on $B(0, R)$). Since

$$|\phi(\mathbf{x} + \mathbf{v}_k) - \phi(\mathbf{v}_k)| = O(|\mathbf{v}_k|^{-1})$$

(as $k \rightarrow \infty$), the sequence of vectors

$$(\phi(\mathbf{x} + \mathbf{v}_k) - \phi(\mathbf{v}_k))A(\mathbf{v}_k)$$

is uniformly bounded for $\mathbf{x} \in B(0, R)$. Furthermore, for every pair of indices $1 \leq i, j \leq n-1$,

$$\frac{\partial^2}{\partial x_i \partial x_j} (\phi(\mathbf{x} + \mathbf{v}_k) - \phi(\mathbf{v}_k))A(\mathbf{v}_k) = \frac{\partial^2}{\partial x_i \partial x_j} \phi(\mathbf{x} + \mathbf{v}_k) \cdot A(\mathbf{v}_k) =$$

$$O(|\mathbf{v}_k|^{-2})A(\mathbf{v}_k) = O(|\mathbf{v}_k|^{-1}).$$

Therefore, the Hessians of $h_k|_B$ uniformly converge to zero as $k \rightarrow \infty$. \square

We would like to strengthen the assertion that ϕ is not constant, unless A is a similarity, which we assume not to be the case. Let Φ be a group of Euclidean isometries acting cocompactly on \mathbb{E}^{n-1} . Fix a point $\mathbf{v} \in \mathbb{E}^{n-1}$. We say that a function ψ defined on a subset E of \mathbb{R}^{n-1} is linear if it extends from E to a linear function on \mathbb{R}^{n-1} .

LEMMA 24.26. *Suppose that the linear transformation A is not a similarity. Then there exists a number R and a sequence of points $\mathbf{v}_k \in \Phi\mathbf{v}$ diverging to infinity, such that the restrictions of h to $B(\mathbf{v}_k, R) \cap \Phi\mathbf{v}$ are nonlinear for all k .*

PROOF. Let R be such that

$$\bigcup_{g \in \Phi} B(g\mathbf{v}, R) = \mathbb{E}^{n-1}$$

and that the intersection $B(\mathbf{v}, R) \cap \Phi \cdot \mathbf{v}$ contains a subset

$$V_R := \{\mathbf{v}_1, \dots, \mathbf{v}_{n-1}\}$$

such that the vectors

$$\mathbf{u}_i = \mathbf{v}_i - \mathbf{v}, i = 1, \dots, n-1,$$

span \mathbb{R}^{n-1} .

Suppose that a sequence (\mathbf{v}_k) as required by the lemma does not exist. This means that there exists $r < \infty$ such that for every $\mathbf{v}_k \in \Phi\mathbf{v} \setminus B(\mathbf{v}, r)$, the restriction of ϕ to every subset $B(\mathbf{v}_k, 4R) \cap \Phi\mathbf{v}$, has a linear extension ϕ_k to \mathbb{R}^{n-1} .

Whenever $\mathbf{v}_k, \mathbf{v}_l \in \Phi\mathbf{v}$ satisfy $|\mathbf{v}_k - \mathbf{v}_l| \leq R$, the intersection

$$B(\mathbf{v}_k, 4R) \cap B(\mathbf{v}_l, 4R)$$

contains the subset

$$\mathbf{v}_k + V_R.$$

Hence, in view of the ‘spanning assumption’ on the subset V_R , the linear extensions ϕ_k, ϕ_l have to be the same.

Define the subset $Y = \{\mathbf{v}_k : |\mathbf{v}_k - \mathbf{v}| > 4R + r\}$. The collection \mathcal{U} of the balls

$$B(\mathbf{v}_k, R), \quad \mathbf{v}_k \in Y,$$

defines an open cover of the complement $\mathbb{R}^{n-1} \setminus B(\mathbf{v}, 4R+r)$. The latter is connected since \mathbb{R}^{n-1} has dimension > 1 . It follows that the nerve of \mathcal{U} is also connected. It follows that the linear functions ϕ_k have to be the same for all $\mathbf{v}_k \in Y$. Since the function ϕ is bounded on $\mathbb{R}^{n-1} \setminus B(\mathbf{v}, r)$, we conclude that the function ϕ_k is constant and, hence, the function ϕ is constant on the subset $\Phi\mathbf{v} \setminus B(\mathbf{v}, 4R+r)$.

According to the Bieberbach Theorem (see e.g. [Rat06, Theorem 7.5.2]), the group Φ contains a free abelian subgroup of rank $n-1$ acting on \mathbb{E}^{n-1} by translations. Up to an affine conjugation, this subgroup is $\mathbb{Z}^{n-1} \subset \mathbb{R}^{n-1}$. The projection of \mathbb{Z}^{n-1} to $\mathbb{R}P^{n-2}$ is dense in the latter. Therefore, the set

$$\left\{ \frac{\mathbf{y}}{|\mathbf{y}|} : \mathbf{y} \in Y \right\}$$

is also dense in the unit sphere. Since $\phi(\mathbf{u}/|\mathbf{u}|) = \phi(\mathbf{u})$ for all non-zero vectors $\mathbf{u} \in \mathbb{R}^{n-1}$, it follows that the function $\phi : \mathbb{R}^{n-1} \rightarrow \mathbb{R}$ is constant. This is a contradiction. \square

EXERCISE 24.27. Reprove the results of this section for inverted affine maps $J \circ F \circ J$, where

$$F : \mathbf{x} \mapsto A\mathbf{x} + \mathbf{b}.$$

24.6. Scattering

We now return to the discussion of quasiisometries. We continue with the notation of Section 24.4. In particular, we have a linear transformation (that is not a similarity) $A \in GL(n-1, \mathbb{R})$, $A(\Lambda) = \Lambda'$, where $\Lambda, \Lambda' \subset \mathbb{E}^{n-1}$ are the sets of centers of peripheral horospheres of the truncated hyperbolic spaces Ω, Ω' . Moreover, after composing A with hyperbolic isometries, we can assume that the origin in \mathbb{R}^n belongs to $\Lambda \cap \Lambda'$. (Recall that earlier, in Section 24.4 we were carefully choosing the point of differentiability, the origin, not to be in Λ .) Lastly, A extends to a quasiisometry \tilde{f} of \mathbb{H}^n , which restricts to a quasiisometry $\Omega \rightarrow \Omega'$.

Let $J : \mathbb{E}^{n-1} \cup \{\infty\} \rightarrow \mathbb{E}^{n-1} \cup \{\infty\}$ be the inversion in the unit sphere \mathbb{S}^{n-2} . Then $\infty = J(0)$ belongs to both $J(\Lambda)$ and $J(\Lambda')$. The conjugate quasiisometry

$$J \circ \tilde{f} \circ J : \mathbb{H}^n \rightarrow \mathbb{H}^n$$

sends $J(\Omega)$ to $J(\Omega')$ and $J \circ A \circ J$ is the boundary extension of this quasiisometry. Since ∞ now belongs to $J(\Lambda) \cap J(\Lambda')$, we have two horoballs B_∞, B'_∞ (with centers at ∞) in the complements of $J(\Omega), J(\Omega')$. The latter are the truncated hyperbolic spaces of the lattices $J\Gamma J, J\Gamma' J$ respectively.

In order to simplify the notation, we now set

$$\Gamma := J\Gamma J, \quad \Gamma' := J\Gamma' J, \quad \Omega := J(\Omega), \quad \Omega' := J(\Omega'), \quad \Lambda := J(\Lambda), \quad \Lambda' := J(\Lambda'),$$

use h for the inverted linear map $J \circ A \circ J$ and \tilde{h} for its quasiisometric extension to \mathbb{H}^n that sends Ω to Ω' . Further, let $\Gamma_\infty, \Gamma'_\infty$ denote the stabilizers of ∞ in Γ, Γ' respectively. Then the groups $\Gamma_\infty, \Gamma'_\infty$ act cocompactly on the boundaries of the horoballs B_∞, B'_∞ , since the quotients $\Omega/\Gamma, \Omega'/\Gamma'$ are both compact. Therefore, the groups $\Gamma_\infty, \Gamma'_\infty$ also act cocompactly on the Euclidean space \mathbb{E}^{n-1} .

REMARK 24.28. We realize that all this is very inconsistent with the notation from Section 24.4, but we no longer need the notation used there.

Lastly, given a point $\mathbf{x} \in \mathbb{E}^{n-1}$, we define the subset $h_*(\mathbf{x}) := h(\Gamma_\infty \mathbf{x}) \subset \mathbb{E}^{n-1}$.

LEMMA 24.29 (Scattering lemma). *Suppose that the nonlinear factor $\phi = \phi_h$ of h is nonconstant. Then for each $\mathbf{x} \in \mathbb{E}^{n-1}$, the set $h_*(\mathbf{x})$ is not contained in the union of finitely many Γ'_∞ -orbits.*

PROOF. Suppose that $h_*(\mathbf{x})$ is contained in the union of finitely many Γ'_∞ -orbits. Each Γ'_∞ -orbit is a discrete subset of \mathbb{E}^{n-1} ; the same is true for a finite union of such orbits. Therefore, for every Euclidean ball $B(\mathbf{x}, R) \subset \mathbb{E}^{n-1}$, the intersection

$$(\Gamma'_\infty \cdot h_*(\mathbf{x})) \cap B(\mathbf{x}, R)$$

is finite. We will show that this cannot be the case.

We apply Lemma 24.26 to the discrete group $\Phi := \Gamma_\infty$ and the point $\mathbf{x} = \mathbf{v}$. The lemma gives us a positive number R and an infinite sequence $\mathbf{x}_k, \mathbf{x}_k = \gamma_k(\mathbf{x}), \gamma_k \in \Phi$, satisfying

$$\lim_{k \rightarrow \infty} |\mathbf{x}_k| = \infty.$$

Since the group Γ'_∞ also acts cocompactly on \mathbb{E}^{n-1} , there exists a sequence $\gamma'_k \in \Gamma'_\infty$, such that the set

$$\{\gamma'_k h(\mathbf{x}_k) : k \in \mathbb{N}\}$$

is relatively compact in \mathbb{E}^{n-1} .

According to Lemma 24.26, the restriction

$$h|_{B(\mathbf{x}_k, R) \cap \Phi \mathbf{x}}$$

is nonlinear for each k . Therefore, the maps

$$h_k := \gamma'_k \circ h \circ \gamma_k$$

cannot be affine on $B(\mathbf{x}, R) \cap \Phi \mathbf{x}$. On the other hand, Corollary 24.25 implies that the sequence of maps

$$h_k|_{B(\mathbf{x}, R)}$$

subconverges to an affine mapping h_∞ . Since each h_k is not affine on $B(\mathbf{x}, R) \cap \Phi \mathbf{x}$, this subconvergence cannot be eventually constant. In other words, there exists $\mathbf{y} \in B(\mathbf{x}, R) \cap \Phi \mathbf{x}$ such that the set $\{h_k(\mathbf{y}) : k \in \mathbb{N}\}$ is infinite. We conclude that the union

$$\bigcup_{k \in \mathbb{N}} h_k(\Phi \mathbf{x} \cap B(\mathbf{x}, R)) \subset (\Gamma'_\infty \cdot h_*(\mathbf{x})) \cap B(\mathbf{x}, R)$$

is an infinite set. This is a contradiction. The lemma follows. \square

THEOREM 24.30. *Suppose that h is an inverted linear map that admits a quasi-isometric extension $\tilde{h} : \mathbb{H}^n \rightarrow \mathbb{H}^n$ sending Ω to Ω' . Then ϕ_h is constant, i.e. h is a similarity map.*

PROOF. Let \mathbf{x} be the center of a complementary horoball B of Ω , $B \neq B_\infty$. Suppose that ϕ_h is nonconstant.

According to the Scattering Lemma, $h_*(\mathbf{x})$ is *not contained* in a finite union of Γ'_∞ -orbits. Let $\gamma_k \in \Gamma_\infty$ be a sequence such that the Γ'_∞ -orbits of the points $h\gamma_k(\mathbf{x})$ are all distinct. Since Γ'_∞ acts on \mathbb{E}^{n-1} cocompactly, there exists an infinite sequence (k_i) and elements

$$\gamma'_{k_i} \in \Gamma'_\infty,$$

such that the sequence

$$\mathbf{x}'_{k_i} := \gamma'_{k_i} h\gamma_{k_i}(\mathbf{x})$$

converges to a point $\mathbf{x}' \in \mathbb{E}^{n-1}$. According to our assumption, all the points \mathbf{x}'_{k_i} are distinct. Let B'_{k_i} denote the complementary horoball to Ω' whose center is \mathbf{x}'_{k_i} . All these horoballs are distinct since their centers are. As the horoballs B'_{k_i} are also pairwise disjoint, we obtain

$$\lim_{i \rightarrow \infty} \text{diam}_{\mathbb{E}^n}(B'_{k_i}) = 0.$$

Let B_k be the complementary horoball to Ω whose center is $\gamma_k \mathbf{x}$. Then

$$D := \text{dist}(B_k, B_\infty) = \text{dist}(B_1, B_\infty) = -\log(\text{diam}_{\mathbb{E}^n}(B_1)).$$

At the same time,

$$\lim_{i \rightarrow \infty} \text{dist}(B'_{k_i}, B'_\infty) = -\lim_{i \rightarrow \infty} \log(\text{diam}(B'_{k_i})) = \infty.$$

Recall that we are assuming that there exists an (L, A) quasiisometric extension \tilde{h} of h such that $\tilde{h} : \Omega \rightarrow \Omega'$. According to Lemma 24.11,

$$\text{dist}(B'_j, B'_\infty) \leq R(L, A) + LD + A.$$

This is a contradiction. □

By combining all these results, we conclude:

THEOREM 24.31. *Suppose that $f : \Omega \rightarrow \Omega'$ is a quasiisometry of the truncated hyperbolic spaces, $n \geq 3$. Then the extension of f to $\partial_\infty \mathbb{H}^n$ is a Moebius transformation.*

24.7. Schwartz Rigidity Theorem

Before proving Theorem 24.1 we will need two technical assertions concerning isometries of \mathbb{H}^n which “almost preserve” the truncated hyperbolic spaces.

Let Ω be the truncated hyperbolic space of a non-uniform lattice $\Gamma < G = \text{Isom}(\mathbb{H}^n)$. We will say that a subset $A \subset G$ *almost preserves* Ω if there exists $C < \infty$ such that

$$\text{dist}_{\text{Haus}}(\Omega, \alpha\Omega) \leq C, \forall \alpha \in A.$$

Note that each $\alpha \in A$ determines an (L, A) -quasiisometry $\Omega \rightarrow \Omega$, defined by composing α with the nearest-point projection $\pi_\Omega : \alpha(\Omega) \rightarrow \Omega$, see Exercise 8.12, Part 1.

LEMMA 24.32. *Suppose that $\beta_k \in G$ is a sequence almost preserving Ω and $\lim_k \beta_k = \beta \in G$. Then the sequence (β_k) consists of finitely many elements of G .*

PROOF. Assume to the contrary that the sequence (β_k) consists of distinct elements. The lattice Γ cannot preserve a proper round sphere in $\partial_\infty \mathbb{H}^n$. Therefore, there exists a finite subset $\Lambda_o \subset \Lambda$ not contained in a proper round sphere in \mathbb{S}^{n-1} . It follows that (after eventually passing to a subsequence), there exists a peripheral horosphere Σ centered at a point $\xi \in \Lambda_o$, such that all the elements of the sequence $\xi_k := \beta_k(\xi)$ are distinct. Since $\lim_{k \rightarrow \infty} \beta_k = \beta$, the horospheres $\beta_k(\Sigma)$ converge to the horosphere $\beta(\Sigma)$. For each k we have a unique horosphere $\widehat{\Sigma}_k \subset \partial\Omega$ whose Hausdorff distance from $\beta_k(\Sigma)$ is uniformly bounded. It follows that the horospheres $\widehat{\Sigma}_k$ have to have non-empty intersections for all large k . This forces the equality

$$\widehat{\Sigma}_k = \widehat{\Sigma}_{k+1}, \forall k \geq k_o.$$

Hence, the centers ξ_k of these horospheres are also equal, which is a contradiction. □

PROPOSITION 24.33. *Let Γ, Γ' be non-uniform lattices in $G = \text{Isom}(\mathbb{H}^n)$ such that Γ' almost preserves Ω , the truncated hyperbolic space of Γ . Then the groups Γ, Γ' are commensurable.*

PROOF. Suppose that the assertion fails. Then the projection of Γ' to $\Gamma \backslash G$ is infinite. Therefore, there exists an infinite sequence (ψ_k) of elements of Γ' whose projections to $\Gamma \backslash G$ are all distinct. Since $\partial\Omega/\Gamma$ is compact, there are only finitely many Γ -orbits of peripheral horospheres of Ω . The set Λ/Γ , consisting of the Γ -orbits of their centers, is also finite. Therefore, after passing to a subsequence in (ψ_k) , we can assume that for some horosphere $\Sigma \subset \partial\Omega$, for every k , the centers of all the horospheres $\psi_k(\Sigma)$ lie in the same Γ -orbit. In other words, there are elements $\gamma_k \in \Gamma$ such that every $\alpha_k := \gamma_k \psi_k$ fixes the center ξ of Σ .

Since all γ_k 's preserve Ω and all ψ_k 's almost preserve Ω , the infinite set

$$A = \{\alpha_k : k \in \mathbb{N}\} \subset G$$

also almost preserves Ω . The projections of all the ψ_k 's to $\Gamma \backslash G$ are pairwise distinct, thus, A projects injectively into $\Gamma \backslash G$.

Without loss of generality, we may assume that $\xi = \infty$ in the upper half-space model of \mathbb{H}^n and Σ is given by the equation $\{x_n = 1\}$. Then the elements of A are Euclidean similarities (they all fix the point ξ). Since the stabilizer Γ_∞ of ∞ in Γ acts cocompactly on the Euclidean space \mathbb{E}^{n-1} , there exists a constant C' and a sequence $\tau_k \in \Gamma_\infty$ such that the compositions $\alpha'_k := \tau_k \alpha_k$ satisfy,

$$|\beta_k(0)| \leq C'.$$

Set $A' := \{\alpha'_k : k \in \mathbb{N}\}$. As before, the subset $A' \subset G$ is infinite, almost preserves Ω and projects injectively to $\Gamma \backslash G$. Since every $\alpha' \in A'$ determines a uniform quasiisometry $\Omega \rightarrow \Omega$, there exists $C < \infty$ such that for every $\alpha' \in A'$,

$$\text{dist}_{\text{Haus}}(\Sigma, \alpha'\Sigma) \leq C.$$

(This is a special case of Proposition 24.8. Cf. Exercise 24.7.) In other words, the value of the coordinate x_n on $\alpha'\Sigma$ satisfies the inequality

$$e^{-C} \leq x_n \leq e^C.$$

Thus, the subset A' is contained in the compact set of similarities

$$\{\beta : \mathbf{x} \mapsto tU\mathbf{x} + \mathbf{v} : e^{-C} \leq t \leq e^C, U \in O(n-1), |\mathbf{v}| \leq C'\}.$$

Therefore, the set A' is infinite and has compact closure in G . This contradicts Lemma 24.32. \square

Proof of Theorem 24.1. Suppose that $\Gamma < G = \text{Isom}(\mathbb{H}^n), n \geq 3$, is a non-uniform lattice.

(a) For each (L, A) -quasiisometry $f : \Gamma \rightarrow \Gamma$, there exists $\alpha \in \text{Comm}_G(\Gamma)$, satisfying

$$\text{dist}(f, \alpha) < \infty.$$

PROOF. The quasiisometry f extends to a quasiisometry of the hyperbolic space $\tilde{f} : \mathbb{H}^n \rightarrow \mathbb{H}^n$ (Theorem 24.13). This quasiisometry extends to a quasiconformal mapping $h : \partial_\infty \mathbb{H}^n \rightarrow \partial_\infty \mathbb{H}^n$. The quasiconformal mapping h has to be Moebius according to Theorem 24.31. Therefore, \tilde{f} is within finite distance from an isometry α of \mathbb{H}^n (which is the unique isometric extension of h to \mathbb{H}^n), see Lemma 11.112.

EXERCISE 24.34. Verify that $\text{dist}(\tilde{f}, \alpha)$ depends only on Γ and on the quasiisometry constants L, A .

It remains to show that α belongs to $\text{Comm}_G(\Gamma)$. We note that f (and, hence, α) almost preserves Ω , the truncated hyperbolic space of Γ . Since Γ preserves Ω , the entire group

$$\Gamma' := \alpha\Gamma\alpha^{-1}$$

almost preserves Ω . By Proposition 24.33, the groups Γ, Γ' are commensurable. Thus, α belongs to the commensurator $\text{Comm}_G(\Gamma)$. \square

(b) Suppose that $\Gamma, \Gamma' < G = \text{Isom}(\mathbb{H}^n)$ are non-uniform lattices quasiisometric to each other. Then there exists $\alpha \in \text{Isom}(\mathbb{H}^n)$ such that the groups Γ and $\alpha\Gamma'\alpha^{-1}$ are commensurable.

PROOF. The proof is analogous to (a): The quasiisometry $f : \Omega' \rightarrow \Omega$ of the truncated hyperbolic spaces of the lattices Γ', Γ is within finite distance from an isometry α . The group $\Gamma'' := \alpha\Gamma'\alpha^{-1}$ again almost preserves Ω . By Proposition 24.33, the groups Γ, Γ'' are commensurable. \square

(c) Suppose that Γ' is a finitely generated group quasiisometric to a non-uniform lattice $\Gamma < \text{Isom}(\mathbb{H}^n)$. Then the groups Γ, Γ' are virtually isomorphic; more precisely, there exists a finite normal subgroup $K \triangleleft \Gamma'$ such that the groups $\Gamma, \Gamma'/K$ contain isomorphic subgroups of finite index.

PROOF. Let $f : \Gamma \rightarrow \Gamma'$ be a quasiisometry and let $f' : \Gamma' \rightarrow \Gamma$ be its coarse inverse. Then, by Lemma 8.63, we have a quasiaction $\Gamma' \curvearrowright \Omega$ via

$$\gamma' \mapsto \rho(\gamma') := f' \circ \gamma' \circ f \in QI(\Omega).$$

According to Part (a), each quasiisometry $g = \rho(\gamma')$ is within a (uniformly) bounded distance from a quasiisometry of Ω induced by an element g^* of $Comm_G(\Gamma)$. We obtain a map

$$\psi : \gamma' \mapsto \rho(\gamma') = g \mapsto g^* \in Comm_G(\Gamma).$$

We claim that this map is a homomorphism with finite kernel. For each quasiisometry $h : \mathbb{H}^n \rightarrow \mathbb{H}^n$ we let h_∞ denote its extension to $\partial_\infty \mathbb{H}^n$. Then, since ρ is a quasiaction, ψ induces a homomorphism

$$\psi_\infty : \gamma' \mapsto g_\infty = g_\infty^*, \quad \psi_\infty : \Gamma' \rightarrow Comm_G(\Gamma),$$

see Theorem 11.135. Since the quasiaction $\rho : \Gamma' \curvearrowright \Omega'$ is geometric (see Lemma 8.63), by Lemma 11.117 the kernel K of the quasiaction $\Gamma' \curvearrowright \Omega'$ is quasifinite. The subgroup $K \triangleleft \Gamma'$ is also the kernel of the homomorphism ψ_∞ ; by Lemma 11.112, this subgroup K is finite.

The rest of the proof is the same as for (a) and (b): The group $\Gamma'' := \psi(\Gamma')$ almost preserves Ω , hence, it is commensurable to Γ . \square

A survey of quasiisometric rigidity

In this chapter we review results and open problems on quasiisometric rigidity of groups and metric spaces. We refer the reader to Section 8.6 for the basic terminology that we will be using. Our survey covers three types of problems within the theme of quasiisometric rigidity:

- (1) **The description of the group of quasiisometries $QI(X)$ of specific metric spaces X and $QI(G)$ for specific finitely generated groups G .** For instance, in some cases, $QI(X)$ coincides with the subgroup of isometries of X , or with the subgroup of virtual automorphisms of G , or with the commensurator of G , either abstract or considered in a larger group.
- (2) **The identification of the classes of groups \mathcal{G} that are QI rigid.** This problem was formulated for the first time (with slightly different terminology) by M. Gromov in [Gro83]. It is sometimes related to the first problem. Indeed, if a group G' is quasiisometric to G , then there exists a homomorphism $G' \rightarrow QI(G)$, which, in many cases, has finite kernel (see Lemma 8.64). If $QI(G)$ is either very close to G , or very close to an ambient group in which all groups in the class \mathcal{G} lie, then one is halfway through a proof of QI rigidity of the class \mathcal{G} .
- (3) **The quasiisometric classification within a given class of groups.** This can be achieved either by a complete description of the equivalence classes or by using QI invariants. An extreme case is when the QI class of a group G contains only finite-index subgroups of G , their quotients by finite normal subgroups and finite extensions of these quotients, i.e. the group G is QI rigid.

We refer the reader to [Bes04], [Sap07], [GBS12] and [MK14] for other open problems in group theory.

25.1. Rigidity of symmetric spaces, lattices and hyperbolic groups

25.1.1. Uniform lattices. The oldest QI rigidity theorem in the context of symmetric spaces was proven by P. Pansu:

THEOREM 25.1 (P. Pansu, [Pan89]). *Let X be a quaternionic hyperbolic space $\mathbb{H}\mathbb{H}^n$ ($n \geq 2$) or the octonionic hyperbolic plane $\mathbb{O}\mathbb{H}^2$. Then X is strongly QI rigid.*

Even though real and complex hyperbolic spaces are not strongly QI rigid, the classes of uniform lattices in their isometry groups are QI rigid. We saw that the class of uniform lattices in the group $PO(n, 1)$ is QI rigid (see Chapter 23), with rigidity results primarily due to P. Tukia, D. Gabai, A. Casson and D. Jungreis.

An analogous QI rigidity theorem was proven by R. Chow [Cho96] for complex-hyperbolic spaces $\mathbb{C}\mathbb{H}^n$, $n \geq 2$, by methods similar to the proof of Tukia's Theorem. We summarize these results as follows:

THEOREM 25.2. *Let X be a symmetric space of negative curvature. Then the class of uniform lattices in X is QI rigid.*

For higher rank symmetric spaces, strong QI rigidity follows from a series of results of B. Kleiner and B. Leeb, which were later also obtained by A. Eskin and B. Farb in [EF97b] by a different method.

THEOREM 25.3 (B. Kleiner, B. Leeb, [KL98b]). *Let X be a symmetric space of non-compact type, without rank one de Rham factors. Assume that if two irreducible factors of X are homothetic to each other then they are isometric. Then X is strongly QI rigid.*

As an application of this rigidity theorem, Kleiner and Leeb, as well as Eskin and Farb, obtained:

THEOREM 25.4 (B. Kleiner, B. Leeb, [KL98b]). *Let X be a symmetric space of non-compact type, without rank one de Rham factors. Then the class of uniform lattices in $\text{Isom}(X)$ is QI rigid.*

Furthermore, Kleiner and Leeb proved that even if the de Rham decomposition of X

$$X = \prod_{i=1}^n X_i$$

does have rank one de Rham factors, the associated quasiaction of Γ on X is within finite distance from another quasiaction, which preserves the de Rham factors X_i (except, it might permute them). Here Γ is a group QI to X .

Kleiner and Leeb also established that strong QI rigidity holds for Euclidean buildings:

THEOREM 25.5 (B. Kleiner, B. Leeb, [KL98b]). *Let X be a Euclidean building such that each de Rham factor of X is a Euclidean building of rank ≥ 2 . Assume that if two irreducible factors of X are homothetic to each other then they are isometric. Then X is strongly QI rigid.*

The overall QI rigidity result for uniform lattices reads as follows:

THEOREM 25.6. *Suppose that X is a symmetric space of non-compact type. Then the class of uniform lattices in $\text{Isom}(X)$ is rigid.*

25.1.2. Non-uniform lattices. Turning to non-uniform lattices, one should first note that Theorem 24.1 of R. Schwartz (see Chapter 24) in its most general form holds even when the space \mathbb{H}^n , $n \geq 3$, is replaced by an arbitrary negatively curved symmetric space of dimension > 2 . This theorem answers the three types of problems described in the beginning of this chapter, and can be stated as follows.

THEOREM 25.7 (R. Schwartz [Sch96a]). *Suppose that X is a negatively curved symmetric space of dimension > 2 ; we let G denote the isometry group of X . Then:*

- (1) *Each non-uniform lattice $\Gamma < G$ is strongly QI rigid: The natural homomorphism $\text{Comm}_G(\Gamma) \rightarrow \text{QI}(\Gamma)$ is an isomorphism.*

- (2) *The class of non-uniform lattices in G is QI rigid.*
- (3) *If Γ and Γ' are quasiisometric non-uniform lattices of isometry groups of negatively curved symmetric spaces X and respectively X' , then $X \cong X'$ (up to rescaling the metric) and the lattices Γ and Γ' are commensurable in G .*

A non-uniform lattice of \mathbb{H}^2 contains a finite-index subgroup which is free non-abelian. In this case we may therefore apply QI rigidity of virtually free groups (Theorem 20.45) and conclude:

THEOREM 25.8. *Each non-uniform lattice in $\text{Isom}(\mathbb{H}^2)$ is QI rigid.*

In the special case when X is the 3-dimensional hyperbolic space, the group of orientation-preserving isometries of X , $\text{Isom}_+(X)$, is isomorphic to $PSL(2, \mathbb{C})$. Schwartz's result has the following arithmetic version. Let \mathbb{F}_1 and \mathbb{F}_2 be imaginary quadratic extensions of \mathbb{Q} and let \mathcal{O}_1 and \mathcal{O}_2 be their respective rings of integers. Then the arithmetic lattices $PSL(2, \mathcal{O}_1)$ and $PSL(2, \mathcal{O}_2)$ (Bianchi groups) are commensurable (quasiisometric) if and only if the fields \mathbb{F}_1 and \mathbb{F}_2 are isomorphic.

When instead of imaginary quadratic extensions, one takes totally real quadratic extensions, the corresponding groups $PSL(2, \mathcal{O}_i)$ become non-uniform \mathbb{Q} -rank one lattices in $\text{Isom}(\mathbb{H}^2 \times \mathbb{H}^2)$, a rank two semisimple Lie group. In general, when \mathbb{F}_i 's are algebraic extensions of \mathbb{Q} , the groups $PSL(2, \mathcal{O}_i)$ are isomorphic to non-uniform \mathbb{Q} -rank one lattices in $\text{Isom}(X)$, where the symmetric space X is isometric to a product of several copies of \mathbb{H}^2 and \mathbb{H}^3 . The QI rigidity theorem in this context was first proven by B. Farb and R. Schwartz [FS96] in the case of fields \mathbb{F}_i of degree 2 and by R. Schwartz [Sch96a] in full generality:

- THEOREM 25.9.**
- (1) *Let \mathbb{F} be an algebraic extension of \mathbb{Q} , let \mathcal{O} be the ring of integers of \mathbb{F} , let $\Gamma = PSL(2, \mathcal{O})$ and let X be the product of hyperbolic spaces on which Γ acts as a lattice. Then the group Γ is strongly QI rigid.*
 - (2) *Let $\mathbb{F}_1, \mathbb{F}_2$ be two algebraic extensions of \mathbb{Q} , and let \mathcal{O}_i be their corresponding rings of integers. Then the lattices $PSL(2, \mathcal{O}_1)$ and $PSL(2, \mathcal{O}_2)$ are quasiisometric if and only if the fields \mathbb{F}_1 and \mathbb{F}_2 are isomorphic.*

It is known that every irreducible lattice Γ in a semisimple group of \mathbb{R} -rank at least 2 and at least one factor of \mathbb{R} -rank one, is an arithmetic \mathbb{Q} -rank one lattice [Pra73, Lemma 1.1]. However, such lattices Γ are, in general, quite different from the arithmetic group $PSL(2, \mathcal{O})$. The case of *higher \mathbb{Q} -rank groups* was settled by A. Eskin:

THEOREM 25.10 (A. Eskin [Esk98]). *Let X, X_1, X_2 be symmetric spaces of non-compact type with all the de Rham factors of rank at least 2. Then:*

- (1) *Every non-uniform lattice $\Gamma < \text{Isom}(X)$ is strongly QI rigid.*
- (2) *The class of non-uniform irreducible lattices in $\text{Isom}(X)$ is QI rigid.*
- (3) *If Γ_1 and Γ_2 are quasiisometric non-uniform irreducible lattices in isometry groups of symmetric spaces X_1 and X_2 , then X_1 is isometric to X_2 (up to rescaling the metrics on the de Rham factors). Moreover, the lattices Γ_1 and Γ_2 are commensurable in $\text{Isom}(X_1) \simeq \text{Isom}(X_2)$.*

An alternative proof of Theorem 25.10, using asymptotic cones, was later provided by C. Druţu in [Dru00]. Some aspects of this alternate approach played an important part in the work of Fisher–Whyte and Fisher–Nguyen mentioned in Theorem 25.77 and in the paragraph following it.

Theorems of Schwartz and Eskin, therefore, prove QI rigidity for irreducible non-uniform lattices in semisimple Lie groups with simple factors of rank $\geq r$ and in semisimple Lie groups with simple factors locally isomorphic to $SL(2, \mathbb{R})$ and $SL(2, \mathbb{C})$. As it turns out, the arguments of Schwartz and Eskin extend to cover irreducible lattices in semisimple Lie groups with other factors of rank one (e.g. products of several copies of $SU(2, 1)$) and a mix of rank one and higher rank factors (such as the product $SU(2, 1) \times SL(3, \mathbb{R})$), as was observed by B. Farb in [Far97], resulting in:

THEOREM 25.11. *The class of irreducible non-uniform lattices in each connected semisimple Lie group G is QI rigid. Two such lattices are quasiisometric if and only if they are virtually isomorphic.*

The class of groups which are left out from this classification include, for instance, the group

$$PSL(2, \mathbb{Z}[\sqrt{-1}]) \times SL(3, \mathbb{Z}).$$

PROBLEM 25.12 (I. Belegardek). Prove QI rigidity of the class of non-uniform, possibly reducible, lattices, for a general symmetric space X of non-compact type.

The QI rigidity results for non-uniform lattices in semisimple Lie groups were extended by K. Wortman [Wor07] to *S-arithmetic lattices*, which are lattices in products of some semisimple Lie groups and some p -adic Lie groups.

Comparison of QI rigidity properties of uniform and non-uniform lattices. Assume, for simplicity, that X is an irreducible symmetric space of non-positive curvature, not isometric to the hyperbolic plane. The results we described above, show that two non-uniform lattices in $G = \text{Isom}(X)$ are commensurable if and only if they are quasiisometric. This fails in the case of uniform lattices, as all such lattices are quasiisometric to each other and it is known that G contains infinitely many virtual isomorphism classes of uniform arithmetic lattices. Restricting to non-arithmetic lattices in $PO(n, 1)$ still leads to infinitely many VI equivalence classes. (We refer the reader to Section 12.3 for the references.)

25.1.3. Symmetric spaces with Euclidean de Rham factors and Lie groups with nilpotent normal subgroups. So far, we considered only non-positively symmetric spaces X of *non-compact type*. The naive QI rigidity fails for uniform lattices in isometry groups of non-positively curved symmetric spaces which are not of non-compact type:

THEOREM 25.13. *Suppose that $G = PO(n, 1) \times \mathbb{R}$, $G = PU(n, 1) \times \mathbb{R}$ or $G = SO(n, 2) \times \mathbb{R}$. Then there are uniform lattices $\Lambda < G = \text{Isom}(X)$, quasiisometric to groups which are not VI to lattices in G . In other words, the class of uniform lattices in G is not QI rigid.*

PROOF. Each of the Lie groups in the theorem has the form $G = G_1 \times \mathbb{R}$. According to Corollary 12.33, there exists a lattice $\Gamma < G_1$ which admits a central coextension

$$1 \rightarrow \mathbb{Z} \rightarrow \tilde{\Gamma} \rightarrow \Gamma \rightarrow 1$$

such that $\tilde{\Gamma}$ is not VI to any product group $\mathbb{Z} \times \Gamma'$ and, at the same time, $\tilde{\Gamma}$ is QI to the product lattice

$$\Lambda = \Gamma \times \mathbb{Z} < G_1 \times \mathbb{R}.$$

Furthermore, every lattice $\Lambda' < G_1 \times \mathbb{R}$ is virtually isomorphic to the direct product $\Gamma' \times \mathbb{Z}$, where Γ' is the projection of Λ' to G_1 . \square

On the other hand, the following theorem shows that the central coextension construction is the only source of failure of QI rigidity for uniform lattices in symmetric spaces of nonpositive curvature.

THEOREM 25.14 (B. Kleiner, B. Leeb, [KL01]). *Suppose that G is a connected Lie group, which fits into a short exact sequence*

$$1 \rightarrow N \rightarrow G \rightarrow \bar{G} \rightarrow 1,$$

where the group N is connected and nilpotent, the group \bar{G} is semisimple and acts via the trivial representation on N . Equip G with a left-invariant Riemannian metric and the associated Riemannian distance function. Then every finitely generated group Γ quasiisometric to G fits into a short exact sequence

$$1 \rightarrow K \rightarrow \Gamma \rightarrow \bar{\Gamma} \rightarrow 1,$$

where K is quasiisometric to N and the group $\bar{\Gamma}$ is virtually isomorphic to a uniform lattice in \bar{G} .

An example of the situation covered by this theorem is a symmetric space $X = Y \times \mathbb{E}^k$, where Y is a symmetric space of non-compact type. The group $G = \text{Isom}_o(Y) \times \mathbb{R}^k = \bar{G} \times N$ acts transitively and isometrically on X .

COROLLARY 25.15. *Each finitely generated group Γ quasiisometric to X fits into a short exact sequence*

$$1 \rightarrow K \rightarrow \Gamma \rightarrow \bar{\Gamma} \rightarrow 1,$$

where the group K is virtually abelian and $\bar{\Gamma}$ acts as a cocompact lattice on Y .

These results leave out the case of general connected Lie groups G ; these groups admit the Levi-Mal'cev decomposition

$$1 \rightarrow S \rightarrow G \rightarrow \bar{G} \rightarrow 1,$$

where S is a solvable Lie group and \bar{G} is semisimple.

The following problem is known to be quite difficult. Note, however, that it bypasses the notoriously difficult problem of QI rigidity for polycyclic groups.

PROBLEM 25.16. Prove an analogue of Theorem 25.14 for all connected Lie groups G , fitting into that exact sequences, where N is solvable and \bar{G} is semisimple.

25.1.4. QI rigidity for hyperbolic spaces and groups. We turn now to QI rigidity in the context of Gromov-hyperbolic spaces and groups. As we saw before, proofs of QI rigidity theorems for lattices in rank one Lie groups are, to large extent, based on a well-developed theory of quasiconformal mappings of ideal boundaries of rank one symmetric spaces. Such a theory was developed by M. Bourdon and H. Pajot in [BP00], and extended further by X. Xie [Xie06], for *2-dimensional hyperbolic buildings* and resulted in strong QI rigidity for these buildings. Instead of giving precise definitions of hyperbolic buildings, we note here only that n -dimensional hyperbolic buildings are certain CAT(-1) spaces X , covered by isometric copies of \mathbb{H}^n (called ‘‘apartments’’), which, in turn, are tiled by some compact convex hyperbolic polyhedra, called *fundamental domains*.

THEOREM 25.17 (M. Bourdon, H. Pajot [BP00]). *Suppose that X is a thick 2-dimensional hyperbolic building, whose links are complete bipartite graphs and whose fundamental domains are right-angled polygons. Then X is strongly QI rigid.*

Some of the restrictions in this theorem were removed later on by X. Xie, using techniques similar to the ones of Bourdon and Pajot:

THEOREM 25.18 (X. Xie [Xie06]). *Each thick 2-dimensional hyperbolic building is strongly QI rigid.*

COROLLARY 25.19. *The class \mathcal{C}_X of groups acting geometrically on a thick 2-dimensional hyperbolic building X , is QI rigid.*

Note that there are many examples of hyperbolic groups acting geometrically on thick 2-dimensional hyperbolic buildings. However, such groups need not be commensurable to each other, which leads to:

PROBLEM 25.20. Construct examples of QI rigid hyperbolic groups whose boundaries are homeomorphic to the Menger curve. (The ideal boundary of each thick 2-dimensional hyperbolic building is a Menger curve.)

The restriction to 2-dimensional hyperbolic buildings also appears to be unnatural, since higher-dimensional hyperbolic spaces tend to be more rigid than low-dimensional ones.

CONJECTURE 25.21. *Each thick hyperbolic building is strongly QI rigid.*

The QI rigidity problem is wide-open for Kleinian groups, i.e. discrete groups of isometries of real-hyperbolic spaces of dimensions $n \geq 3$.

CONJECTURE 25.22. *The class of finitely generated discrete subgroups of $PO(3, 1)$ is QI rigid.*

Recently, this conjecture was proven by P. Haïssinsky in the case of the class of Gromov–hyperbolic discrete subgroups of $PO(3, 1)$:

THEOREM 25.23 (P. Haïssinsky [Haï15]). *1. The class of Gromov–hyperbolic discrete subgroups of $PO(3, 1)$ is QI rigid.*

2. Moreover, if Γ is a finitely generated group which admits a QI embedding into \mathbb{H}^3 , then Γ is virtually isomorphic to a Gromov–hyperbolic discrete subgroup of $PO(3, 1)$.

One can ask for a stronger *topological* rigidity property in this regard, namely:

CONJECTURE 25.24. *Let Γ be a hyperbolic group whose ideal boundary is planar, i.e. topologically embeds in the 2-sphere. Then Γ is virtually isomorphic to a discrete subgroup of $PO(3, 1)$.*

Note that this conjecture includes two well-known problems as special cases:

1. *Cannon’s conjecture:* A hyperbolic group whose ideal boundary is homeomorphic to \mathbb{S}^2 , acts geometrically on \mathbb{H}^3 .

2. *Sierpinsky carpet conjecture [KK00]:* A hyperbolic group whose ideal boundary is homeomorphic to the Sierpinsky carpet, is virtually isomorphic to a discrete subgroup in $PO(3, 1)$.

We refer the reader to Haïssinsky’s paper [Haï15] for the most recent results in this direction.

Very little is known about strong QI rigidity of hyperbolic groups with higher-dimensional boundary. Below we list what is known and some open problems.

THEOREM 25.25 (M. Kapovich, B. Kleiner, [KK00]). *There are hyperbolic groups G with 2-dimensional boundaries, which are strongly QI rigid. Furthermore, in these examples, all homeomorphisms of $\partial_\infty G$ are restrictions of elements of G .*

Suppose that M is a compact n -dimensional Riemannian manifold of constant curvature -1 and non-empty totally-geodesic boundary. We will refer to the fundamental groups of boundary components of M as *peripheral subgroups* of $G = \pi_1(M)$. The universal cover \tilde{M} is isometric to a certain convex subset $C \subset \mathbb{H}^n$ bounded by pairwise disjoint hyperbolic subspaces of codimension 1. The group G acts geometrically on C and, hence, embeds as a discrete subgroup of $PO(n, 1)$. The ideal boundary of C in \mathbb{S}^{n-1} is a *round Sierpinsky carpet*: It is nowhere dense in \mathbb{S}^{n-1} and its complement is a union of open round balls with pairwise disjoint closures. Such groups G are called *round Sierpinsky carpet groups*. The following theorem was known to various people, its proof is sketched by J.-F. Lafont in [Laf04]:

THEOREM 25.26. *1. The convex sets C above are strongly QI rigid. 2. In particular, each round Sierpinsky carpet group is QI rigid.*

It is a corollary of Thurston’s Geometrization Theorem (see [Kap01]) that each Gromov–hyperbolic discrete subgroup of $PO(3, 1)$ whose ideal boundary is homeomorphic to a Sierpinsky carpet is isomorphic to a round Sierpinsky carpet group.

PROBLEM 25.27. Suppose that $G < PO(n, 1)$ is a discrete subgroup, which is Gromov–hyperbolic and $\partial_\infty G$ is homeomorphic to a Sierpinsky carpet. Is it true that G is QI rigid?

Round Sierpinsky carpet groups can be amalgamated along their peripheral subgroups. More specifically, consider an n -dimensional complex C obtained by gluing finitely many n -dimensional compact hyperbolic manifolds with geodesic boundary in such a way that along each boundary component we glue at least three n -dimensional manifolds. The fundamental groups of the complexes C are hyperbolic, e.g. by Bestvina-Feighn Combination Theorem [BF92]. J.-F. Lafont in [Laf04] proved that the groups $\pi_1(C)$ are QI rigid.

Fundamental groups of complexes of simplicial negative curvature (see [JŚ03, JŚ06]) provide a major new class of higher dimensional hyperbolic groups.

PROBLEM 25.28. Are any of these higher-dimensional hyperbolic groups QI rigid?

Recall that, in view of Thurston’s Hyperbolization Theorem, for all closed hyperbolic surfaces S and pseudo-Anosov homeomorphisms $f : S \rightarrow S$, the mapping tori M_f of $f : S \rightarrow S$ are hyperbolic. The fundamental group of every such M , is a semidirect product

$$\pi_1(S) \rtimes_\phi \mathbb{Z},$$

where $\phi \in \text{Aut}(\pi_1(S))$ is the automorphism induced by f (it is defined up to an inner automorphism of the fundamental group). Then all of the groups $\pi_1(M)$ are quasiisometric to each other, since they embed as uniform lattices in $\text{Isom}(\mathbb{H}^3)$. The question is what happens if we replace closed surface groups in such extensions by finitely generated free groups F of finite rank $r \geq 3$. The natural generalization of the notion of a pseudo-Anosov homeomorphism in this setting is the one of an

atoroidal fully irreducible (iwip) automorphism ϕ of F : These are automorphisms so that no power of ϕ preserves (up to conjugation) a free factor of F and so that ϕ has no periodic conjugacy classes. The corresponding semidirect products

$$F_\phi := F \rtimes_\phi \mathbb{Z}$$

are hyperbolic groups and their boundaries are Menger curves, see [Bri00]. One can imagine several possibilities of quasiisometric behavior of such semidirect products, the question is which (if any) of them actually occurs:

QUESTION 25.29 (P. Sardar). Which (if any) of the following hold:

1. All groups F_ϕ are quasiisometric to each other.
2. The groups F_ϕ are quasiisometric if and only if they are virtually isomorphic.
3. If a group G is quasiisometric to some F_ϕ , then G is VI to F_ϕ .
4. Every hyperbolic group G with Menger curve boundary is quasiisometric to one of the groups F_ϕ .

Note that (3) and (4) are mutually exclusive, since there are hyperbolic groups with Menger curve boundary which have the Property (T), while none of the free-by-cyclic groups has this property.

Lastly, we note that numerous results about the structure of quasiisometries of solvable Lie groups with negatively curved left-invariant Riemannian metrics were obtained by X. Xie, [SX12, Xie12, Xie13, Xie14, Xie16].

25.1.5. Failure of QI rigidity. So far we discussed QI rigidity in various forms. Below are examples of failure of QI rigidity.

Central coextensions. As we saw in Theorem 25.13, there are numerous example of groups Γ which admit central extensions

$$1 \rightarrow \mathbb{Z} \rightarrow \tilde{\Gamma} \rightarrow \Gamma \rightarrow 1,$$

such that $\tilde{\Gamma}$ is QI to the product $\Gamma \times \mathbb{Z}$, but it is not VI to it. For instance:

EXAMPLE 25.30. Let S be a closed hyperbolic surface and let M be the unit tangent bundle of S . Then we have an exact sequence

$$1 \rightarrow \mathbb{Z} \rightarrow G = \pi_1(M) \rightarrow Q := \pi_1(S) \rightarrow 1.$$

This sequence does not split even after passage to a finite index subgroup in G , hence, G is not virtually isomorphic to $Q \times \mathbb{Z}$. However, since Q is hyperbolic, the group G is quasiisometric to $Q \times \mathbb{Z}$, by Theorem 11.159. Note that the group $Q \times \mathbb{Z}$ is CAT(0), while the group G is not (see e.g. [BH99] or [KL98a]). In particular, the class of CAT(0) groups is not QI rigid.

Groups quasiisometric to products of trees. In [BM00], M. Burger and S. Mozes constructed examples of *simple* groups G acting geometrically on products of locally finite simplicial trees $T_1 \times T_2$. In their examples, each tree T_i was has infinitely many ends and large group of automorphisms (it is transitive on the set of all embedded edge-paths of length n , for each n). In particular, the trees T_i have constant valence ≥ 3 and, hence, are quasiisometric to the free group F_2 . Therefore, in these examples, the group G is quasiisometric to the product group $F_2 \times F_2$.

COROLLARY 25.31. *The product of free groups $F_n \times F_m$, ($n, m \geq 2$) is not QI rigid: It is quasiisometric to a group G as in the Burger-Mozes construction mentioned above, but it is not virtually isomorphic to it.*

The group $F_2 \times F_2$ is co-large and not simple; therefore:

COROLLARY 25.32 (M. Burger, S. Mozes). *Virtual simplicity and co-largeness are not QI invariant.*

Here a group G is *virtually simple* if it is VI to a simple group.

These examples, of course, have the same *geometric model*, Definition 8.55. According to Theorem 8.56, there are commensurable groups without a common geometric model.

PROBLEM 25.33. Find an example of a pair of groups G_1, G_2 which are QI to each other, but they are not VI to groups with a common geometric model.

Virtual torsion-freeness. There are example of virtually isomorphic groups G_1, G_2 such that G_1 is virtually **torsion-free** (even linear), while G_2 is not: Millson [Mil79] constructed examples of lattices G_1 in a linear Lie group and central coextensions

$$0 \rightarrow \mathbb{Z}_2 \rightarrow G_2 \rightarrow G_1 \rightarrow 1$$

such that the groups G_2 are not virtually **torsion-free**. Being linear, the lattice G_1 is, of course, virtually torsion-free by Selberg's Lemma. Further examples like this were constructed by Raghunathan [Rag84].

PROBLEM 25.34. 1. Is it true that every finitely generated group is QI to a torsion-free group?

2. Construct examples of groups G which are QI to torsion-free groups but not VI to torsion-free groups.

Note that a positive answer to the first question (which seems unlikely) would be an ultimate form of Selberg's Lemma. As for the question 2, natural candidates would be simple groups acting geometrically on products of trees. However, it appears that all the currently known examples of simple groups acting geometrically on products of trees are torsion-free.

Hopfian and cohopfian properties. Both properties are not preserved by virtual isomorphisms, see Section 7.12.

Unbounded group actions on trees. We will say that a group Γ *virtually splits* if it is virtually isomorphic to a group which admits a non-trivial graph of groups decomposition. Recall that, in view of the Bass-Serre theory, a group Γ admits a non-trivial decomposition as a graph of groups if and only if Γ acts on a simplicial tree (without inversions) without a fixed vertex. The Lie group $G = SL(2, \mathbb{R}) \times SL(2, \mathbb{R})$ contains uniform irreducible lattices Γ , for instance $\Gamma \cong SL(2, \mathbb{Z}[\sqrt{2}])$. According to Theorem 12.13 of Margulis, every action of such irreducible lattice on a simplicial tree has a fixed point. Therefore, irreducible lattices in G do not virtually split. On the other hand, G also contains reducible uniform lattices Λ , e.g., discrete subgroups isomorphic to $\pi_1(S) \times \pi_1(S)$, where S is a compact hyperbolic surface. Such lattices do split non-trivially as graphs of groups, since $\pi_1(S)$ does. Therefore, we obtain examples of quasiisometric groups Γ and Λ such that Γ does not virtually split, while Λ does.

The examples of failure of quasiisometric invariance of Property (T), see Theorem 19.76, also show that the property FA is not a quasiisometric invariant either.

QUESTION 25.35. 1. Are Properties T and Haagerup quasiisometrically invariant in the class of hyperbolic groups?

2. Suppose that G_1, G_2 are quasiisometric hyperbolic groups and G_1 virtually splits. Is it true that G_2 virtually splits as well?

The second part of the problem is open even for arithmetic lattices in $PU(2, 1)$. On the other hand, both parts have positive answers in some notable examples, like lattices in $\text{Isom}(\mathbf{H}\mathbb{H}^n)$ ($n \geq 2$) and in isometry groups of 2-dimensional thick hyperbolic buildings.

25.1.6. Rigidity of random groups. At this time, it is far from clear how common is the phenomenon of QI rigidity for finitely presented groups. The reader might have noticed that all the examples of QI rigid groups and classes of groups are quite special: They come from groups acting discretely and isometrically on some highly homogeneous metric spaces (symmetric spaces and buildings). On the other hand, all the examples of groups failing QI rigidity are rather special as well. This leads to:

PROBLEM 25.36. Let G be an infinite *random*¹ finitely presented group. Is G QI rigid?

We refer the reader to Section 11.25 for the description of some models of *randomness* among finitely presented groups. Here we recall only that, according to all these models, *random groups* are Gromov–hyperbolic. Generic QI rigidity is an open problem for all these models (except for the case of density larger than $\frac{1}{2}$, in which random groups are finite.) One reason to expect random groups to be QI rigid is that I. Kapovich and P. Schupp proved in [KS08] that, in one of the models, random groups are “algebraically rigid”, i.e. their isomorphisms are induced by Nielsen transformations.

25.2. Rigidity of relatively hyperbolic groups

As we know, see Corollary 11.43, the class of hyperbolic groups is QI rigid. In this section we discuss QI rigidity properties of relatively hyperbolic groups. In the following discussion, by a relatively hyperbolic group we always mean a group G which admits a relatively hyperbolic structure where each peripheral subgroup has infinite index in G . Thus, a group G is *not relatively hyperbolic (NRH)* if G contains no finite collection of infinite index subgroups with respect to which it is relatively hyperbolic.

THEOREM 25.37 (C. Druţu, [Dru09]). *The class of relatively hyperbolic groups is QI rigid. More precisely, if a group G_1 is relatively hyperbolic and a group G_2 is quasiisometric to G_1 , then G_2 is also relatively hyperbolic.*

Other theorems appearing in this chapter emphasize that various subclasses of relatively hyperbolic groups are, likewise, QI rigid: Non-uniform lattices in rank one symmetric spaces (Theorem 25.7), fundamental groups of non-geometric Haken manifolds (Theorem 25.69), fundamental groups of graphs of groups with finite edge groups [PW02].

Theorem 25.37 suggests the following natural question.

¹In any currently existing model of randomness.

PROBLEM 25.38 (P. Papasoglou). Is there a geometric criterion allowing to recognize whether a finitely generated group is relatively hyperbolic (without any reference to peripheral subgroups or subsets)?

Concerning the proof of Theorem 25.37, it is not difficult to see that if $f : X \rightarrow Y$ is a quasiisometry between two metric spaces and X is hyperbolic relative to \mathcal{A} then Y is hyperbolic relative to $\{f(A) : A \in \mathcal{A}\}$. Thus, the main step towards proving Theorem 25.37 is to show that if a group G is hyperbolic relative to some collection of subsets \mathcal{A} , then it is also hyperbolic relative to some collection of subgroups H_1, \dots, H_n , such that each H_i is contained in a metric neighborhood of some A_i , [Dru09]. A variation of the same argument, appears in [BDM09]:

THEOREM 25.39. *Let X be a metric space which is hyperbolic relative to a collection of subsets \mathcal{A} . Suppose that $f : G \rightarrow X$ is a quasiisometric embedding of a finitely generated group G . Then G is hyperbolic relative to the collection of pre-images $f^{-1}(\mathcal{N}_C(A))$, $A \in \mathcal{A}$, for some $C < \infty$.*

This result and the above argument imply that either $f(G)$ lies in an M -tubular neighborhood of some set $A \in \mathcal{A}$, or G is hyperbolic relative to finitely many (proper) subgroups H_i with each $f(H_i)$ contained in an M -tubular neighborhood of some set $A_i \in \mathcal{A}$. Thus, we have the following generalization of the Quasi-Flat Lemma of R. Schwartz [Sch96b] (see Proposition 24.8, Lemma 24.11 as well as Remark 24.12 in Chapter 24). This generalization was proven by J. Behrstock, C. Druțu, L. Mosher in [BDM09], Theorem 4.1:

THEOREM 25.40 (NRH subgroups are always peripheral). *Let X be a metric space hyperbolic relative to a collection \mathcal{A} of subsets. For every $L \geq 1$ and $C \geq 0$ there exists $R = R(L, C, X, \mathcal{A})$ such that the following holds:*

If G is a finitely generated group and G is NRH, then the image of any (L, C) -quasiisometric embedding $f : G \rightarrow X$ is contained in the R -neighborhood of some set $A \in \mathcal{A}$.

In this theorem, the constant R does not depend on the group G . In [DS05b] the same theorem was proved under the stronger hypothesis that the group G has one asymptotic cone without global cut-points.

As in the case of the proof of Theorem 24.1, Theorem 25.40 is a step towards a QI classification of relatively hyperbolic groups:

THEOREM 25.41 (J. Behrstock, C. Druțu, L. Mosher, [BDM09], Theorem 4.8). *Let G be a finitely generated group hyperbolic relative to a finite collection of finitely generated subgroups \mathcal{H} , such that each $H \in \mathcal{H}$ is NRH. If G' is a finitely generated group quasiisometric to G , then G' is hyperbolic relative to a finite collection of finitely generated subgroups \mathcal{H}' , where each subgroup in \mathcal{H}' is quasiisometric to one of the subgroups in \mathcal{H} .*

When working in full generality, it is impossible to establish a relation between peripheral subgroups of QI relatively hyperbolic groups; hence, this is not mentioned in Theorem 25.37. For instance, when $G = G' = A \star B \star C$, the group G is hyperbolic relative to $\{A, B, C\}$, and it is also hyperbolic relative to $\{A \star B, C\}$.

By the results in [PW02], the QI classification of relatively hyperbolic groups reduces to the classification of one-ended relatively hyperbolic groups. Theorem 25.41 points out a fundamental necessary condition for two one-ended relatively

hyperbolic groups (with NRH peripheral subgroups) to be quasiisometric: Their peripheral subgroups have to define the same collection of quasiisometry classes.

Related to this, one may ask whether every relatively hyperbolic group G admits a relatively hyperbolic structure $(G; \mathcal{P})$, such that all peripheral subgroups $P_i \in \mathcal{P}$ are NRH. The answer is negative in general, a counter-example is Dunwoody's inaccessible group [Dun93]. Since finitely presented groups are accessible, this raises the following natural question:

PROBLEM 25.42 (J. Behrstock, C. Druţu, L. Mosher, [BDM09]). Is there an example of a finitely presented relatively hyperbolic group G such that for every relatively hyperbolic structure $(G; \mathcal{P})$ at least one group $P_i \in \mathcal{P}$ is a relatively hyperbolic group?

25.3. Rigidity of classes of amenable groups

The class of amenable groups is QI rigid, see Theorem 18.13. Recall that, by Corollary 18.52, the set of finitely generated groups splits into amenable groups and paradoxical groups. This implies that the class of paradoxical groups is also QI rigid. Since the latter class is characterized by the fact that the Cheeger constant is positive (Theorem 18.4), it follows that having a positive Cheeger constant is a QI invariant property. As noted, the property of having positive Cheeger constant is QI invariant not only among groups, but also among graphs and manifolds of locally bounded geometry.

Various subclasses of amenable groups behave quite differently with respect to QI rigidity, and relatively little is known about their QI classification and the description of groups of quasiisometries.

The class of virtually nilpotent groups is QI rigid by Theorem 16.25. Concerning the QI classification of nilpotent groups, the following is known:

THEOREM 25.43 (P. Pansu [Pan89]). *If G and H are finitely generated quasiisometric nilpotent groups, then the graded Lie groups associated with $G/\text{Tor}(G)$ and $H/\text{Tor}(H)$ are isomorphic.*

One of the steps in the proof of Theorem 25.43 is that all the asymptotic cones of a finitely generated nilpotent group G with a canonically chosen word metric are isometric to the graded Lie group associated to $G/\text{Tor}(G)$, endowed with a Carnot-Caratheodory metric, see Theorem 16.28.

Theorem 25.43 establishes other quasiisometry invariants in the class of nilpotent groups: The nilpotency class of $\bar{G} = G/\text{Tor}(G)$ and ranks of the abelian groups $C^i \bar{G}/C^{i+1} \bar{G}$, where $C^i \bar{G}$ is the i -th group in the lower central series of \bar{G} .

Other QI invariants in the class of nilpotent group that help to distinguish nilpotent groups with the same associated nilpotent graded Lie groups are the virtual Betti numbers [Sha04, Theorem 1.2]. Recall that Pansu also proves the QI rigidity of abelian groups, see Theorem 16.26.

Unlike abelian groups, nilpotent groups are not completely classified up to QI. In particular, the following remains an open problem:

PROBLEM 25.44. Is it true that two nilpotent simply-connected Lie groups (endowed with left-invariant Riemannian metrics) are quasiisometric if and only if they are isomorphic?

The answer to this problem is positive for graded nilpotent Lie groups, according to Theorem 25.43.

The group of quasiisometries is very large already for abelian groups, see Examples 8.15 and 8.16. However, it is a reasonable and not well-understood problem to classify uniformly quasiisometric discrete subgroups of quasiisometries of Euclidean spaces and of nilpotent groups. For instance:

PROBLEM 25.45. Is there a discrete quasiaction $G \curvearrowright \mathbb{E}^n$ of a finitely generated nilpotent group G , which is not virtually abelian?

In view of Erschler's Theorem 14.40, the class of (virtually) solvable groups is not QI rigid. Note, however, that:

1. the groups constructed in the proof are not finitely presented;
2. both groups G_A and G_B in the proof are elementary amenable.

This leads to:

- PROBLEM 25.46.**
1. Is the class of finitely presented solvable groups QI rigid?
 2. Is the class of finitely presented metabelian groups (i. e., solvable groups of derived length 2) QI rigid?
 3. Is the class of elementary amenable groups QI rigid?

Below we review partial results and open problems in this direction.

THEOREM 25.47 (B. Farb and L. Mosher [FM00]). *The class of finitely presented non-polycyclic abelian-by-cyclic groups is QI rigid.*

The starting point in the proof of this theorem is to consider torsion-free finite-index subgroups and to apply a result of R. Bieri and R. Strebel [BS78]. The latter states that for every torsion-free finitely presented abelian-by-cyclic group G , there exists $n \in \mathbb{N}$ and a matrix $M = (m_{ij}) \in M(n, \mathbb{Z})$ with non-zero determinant, such that the group G has the presentation

$$(25.1) \quad \langle a_1, a_2, \dots, a_n, t \mid [a_i, a_j], ta_i t^{-1} a_1^{m_{1i}} a_2^{m_{2i}} \dots a_n^{m_{ni}} \rangle .$$

Let Γ_M be the group with the presentation in (25.1) for the integer matrix M . The group Γ_M is polycyclic if and only if $|\det(M)| = 1$, see [BS80].

In [FM00], Farb and Mosher prove that if a finitely generated group G is quasiisometric to the group Γ_M , for an integer matrix M with $|\det(M)| > 1$, then G , is virtually isomorphic to a group Γ_N defined by an integer matrix N with $|\det(N)| > 1$.

THEOREM 25.48 (B. Farb, L. Mosher, [FM00]). *Let M_1 and M_2 be integer matrices with $|\det(M_i)| > 1$, $i = 1, 2$. The groups Γ_{M_1} and Γ_{M_2} are quasiisometric if and only if there exist two positive integers k_1 and k_2 such that $M_1^{k_1}$ and $M_2^{k_2}$ have the same absolute Jordan form.*

The absolute Jordan form of a matrix is obtained from the Jordan form over \mathbb{C} by replacing the diagonal entries with their absolute values and arranging the Jordan blocks in a canonical way.

In the case of solvable Baumslag–Solitar groups, which form a subclass in the class of groups in Theorem 25.48, more can be said:

THEOREM 25.49 (B. Farb, L. Mosher, [FM98], [FM99]). *Each solvable Baumslag–Solitar group*

$$BS(1, m) = \langle x, y \mid xyx^{-1} = y^m \rangle$$

is *QI rigid*.

This theorem is complemented by:

THEOREM 25.50 (K. Whyte, [Why01]). *1. All non-solvable Baumslag–Solitar groups*

$$BS(n, m) = \langle x, y \mid xy^n x^{-1} = y^m \rangle,$$

$|n| \neq 1, |m| \neq 1$ are QI to each other.

2. If $(a, b) = (c, d) = 1$ and $a/b = c/d$, then the groups $BS(a, b)$ and $BS(c, d)$ are not commensurable.

Since non-solvable Baumslag–Solitar groups are nonamenable, these results complete the QI classification of Baumslag–Solitar groups.

The fact that polycyclic groups were excluded from theorems of Farb and Mosher is not an accident: These groups are much harder to handle since the tools of coarse topology do not apply to them.

PROBLEM 25.51. (1) Is the class of finitely generated polycyclic groups QI rigid?

(2) What is the QI classification of finitely generated polycyclic groups?

Every virtually polycyclic group of course has a finite-index subgroup with infinite abelianization. Shalom in [Sha04] proved (among other QI rigidity properties for various classes of amenable groups) the following:

THEOREM 25.52. *Suppose that G is a group QI to a polycyclic group. Then G contains a finite-index subgroup with infinite abelianization.*

Even the problem of QI rigidity for finitely generated polycyclic abelian-by-cyclic groups has remained open for some time. The papers of Eskin, Fisher and Whyte [EFW13, EFW12] made a major progress in this direction. In particular, they prove:

THEOREM 25.53. *Consider the class $Poly_3$ of groups G which are not virtually nilpotent and admit short exact sequences*

$$1 \rightarrow \mathbb{Z}^2 \rightarrow G \rightarrow \mathbb{Z} \rightarrow 1.$$

Then the class $Poly_3$ is QI rigid.

Note that the groups in $Poly_3$ play an important role in 3-dimensional topology. Namely, there exists a 3-dimensional simply-connected solvable Lie group Sol_3 , such that each $\Gamma \in Poly_3$ is isomorphic to a uniform lattice in Sol_3 . Accordingly, Γ is isomorphic to the fundamental group of a closed 3-dimensional manifold $M = \Gamma \backslash Sol_3$. The manifolds M of this form (and manifolds which are covered by such M 's) are called *Sol_3 -manifolds*, they appear in the classification theory of 3-dimensional manifolds.

In view of the Problem 25.51, Part (1), one may ask about the QI classification of solvable Lie groups.

Y. Cornuier proved that each connected Lie group is quasiisometric to a closed connected subgroup of the group of real upper triangular matrices [dC08, Lemma 6.7]. This has lead him to ask:

PROBLEM 25.54 (Y. Cornuier [dC11]). Suppose that G_1, G_2 are closed connected subgroups of the group of real upper triangular matrices, endowed with left-invariant Riemannian metrics. Is it true that G_1, G_2 are quasiisometric if and only if they are isomorphic?

Groups QI to abelian-by-abelian solvable groups. Generalizing the results of [EFW13, EFW12], I. Peng in [Pen11a, Pen11b] considered quasiisometries of lattices in solvable Lie groups G of the type

$$G = G_\varphi = \mathbb{R}^n \rtimes_\varphi \mathbb{R}^m,$$

where $\varphi : \mathbb{R}^m \rightarrow GL(n, \mathbb{R})$ is an action of \mathbb{R}^m on \mathbb{R}^n . The number m is called the *rank* of G . The group G is clearly a Lie group and we equip G with a left-invariant Riemannian metric. Then G admits *horizontal* and *vertical* foliations by the left translates of \mathbb{R}^n and \mathbb{R}^m .

DEFINITION 25.55. 1. G_φ is *unimodular* if $\varphi(\mathbb{R}^m) \subset SL(n, \mathbb{R})$.

2. G_φ is *nondegenerate* if for each $\mathbf{x} \in \mathbb{R}^m \setminus \{0\}$, $\varphi(\mathbf{x})$ has at least one eigenvalue with the absolute value $\neq 1$.

3. G_φ is *split* if there are no non-zero $\varphi(\mathbb{R}^m)$ -invariant subspaces $V \subset \mathbb{R}^n$ such that the image group $\varphi(\mathbb{R}^m) \ltimes GL(V)$ is a bounded subgroup.

Note that this definition is given in [Pen11b] in terms of root systems, but it is easy to see that our definitions are equivalent to hers.

The main result of [Pen11a, Pen11b] is to establish control of quasiisometries $G_\varphi \rightarrow G_\psi$ between the solvable groups, the key being that quasiisometries send leaves of horizontal and vertical foliations uniformly close to leaves of horizontal and vertical foliations. The precise result is too technical to be stated here (see [Pen11b, Theorem 5.2]); below is its main corollary, which is a combination of the work of I. Peng and T. Dymarz (Corollary 1.0.2 in [Pen11b]):

THEOREM 25.56 (I. Peng, T. Dymarz). *Suppose that G_φ is non-degenerate and unimodular, and G' is a finitely generated group quasiisometric to G_φ . Then the group G' is virtually polycyclic.*

As a special case, consider a group $G = G_\varphi$ of rank 1, i.e. $m = 1$. T. Dymarz in [Dym10] proved the following:

THEOREM 25.57 (T. Dymarz). *Every finitely generated group QI to G is virtually isomorphic to a lattice in G .*

25.4. Bilipschitz vs. quasiisometric

The question about the difference between quasiisometries and bilipschitz maps between finitely generated groups is both very basic and interesting. At the first glance, there should not be any need for passing to a subnet in order to go from quasiisometries to bilipschitz maps of finitely generated groups. Gromov asked in [Gro93, § 1.A0] if this is really the case, as the situation was unclear even for separated nets in Euclidean spaces and for free groups of different (finite) ranks:

QUESTION 25.58 (M. Gromov). 1. Suppose that $X_1, X_2 \subset \mathbb{E}^n$ are separated nets. Is there a bilipschitz homeomorphism $X_1 \rightarrow X_2$?

2. Suppose that F_m, F_n , $2 \leq m, n < \infty$, are free groups of ranks m and n respectively, equipped with the word metrics associated with their free generating sets. Is F_m bilipschitz to F_n ?

The case of free groups was settled quickly by P. Papasoglu [Pap95a]:

THEOREM 25.59. *Any two nonabelian free groups of finite ranks are bilipschitz to each other.*

In view of his theorem it was reasonable to expect that any two separated nets in \mathbb{E}^n are bilipschitz to each other and that any two finitely generated quasiisometric groups are also bilipschitz equivalent. In a surprising development, D. Burago and B. Kleiner [BK02b] and C. McMullen [McM98], independently constructed examples of separated nets in \mathbb{R}^2 which are not bi-Lipschitz equivalent.

QUESTION 25.60 (D. Burago, B. Kleiner, [BK02b]). 1. When placing a point in the barycenter of each tile of a *Penrose tiling* in \mathbb{E}^2 , is the resulting separated net bi-Lipschitz equivalent to \mathbb{Z}^2 ?

2. More generally: Embed \mathbb{R}^2 into \mathbb{R}^n as a plane P with irrational slope and take B , a bounded subset of \mathbb{R}^n with non-empty interior. Consider the subset $Z \subset \mathbb{Z}^n$ of all z 's such that $z + B$ intersects P . The orthogonal projection Z to P composes a separated net in \mathbb{R}^2 . Is such a net bilipschitz equivalent to \mathbb{Z}^2 ?

Part 1 of this question was answered by Y. Solomon in [Sol11], see also [APCG13] for an improvement of his results.

This has left open the case of finitely generated groups. The case of non-amenable groups was settled by K. Whyte:

THEOREM 25.61 (K. Whyte, [Why99]; see Theorem 18.9 of this book). *Suppose that G_1, G_2 are finitely generated quasiisometric non-amenable groups. Then G_1, G_2 are bilipschitz equivalent. More generally, for every quasiisometry $f : \Gamma_1 \rightarrow \Gamma_2$ between nonamenable graphs of bounded geometry, the restriction*

$$f|_{V(\Gamma_1)} : V(\Gamma_1) \rightarrow V(\Gamma_2)$$

is at a bounded distance from a bilipschitz bijection.

We also note that the theorem about graphs is implicitly contained in the earlier paper [DSS95] of W. A. Deuber, M. Simonovits and V. T. Sós.

The case of amenable groups was settled (in the negative) by T. Dymarz in [Dym10]. She constructed certain *lamplighter groups* which are quasi-isometric but not bilipschitz equivalent. Her examples, however, are commensurable. This leads to:

PROBLEM 25.62. Generate an equivalence relation CLIP on finitely generated groups by virtual isomorphism and bilipschitz equivalence. Is CLIP equal to the quasiisometry equivalence relation?

We note that Dymarz' examples are merely finitely generated; finitely presented examples were constructed by Dymarz, Peng and Taback [DPT15].

25.5. Various other QI rigidity results and problems

The following theorem was first proven by R. Grigorchuk in [Gri84a], who proved in [Gri84a] that there are uncountably many equivalence classes of growth

functions of groups of intermediate growth. B. Bowditch in [Bow98a] gave a different argument, not based on the growth of groups.

THEOREM 25.63 (R. Grigorchuk [Gri84a]). *There are uncountably many QI classes of finitely generated groups.*

We note that most of the progress in establishing QI rigidity was achieved in the context of lattices in Lie groups or certain solvable groups. Below we review some QI rigidity results for groups which do not belong to these classes.

The following rigidity theorem was proven by J. Behrstock, B. Kleiner, Y. Minsky and L. Mosher in [BKMM12]:

THEOREM 25.64. *Let S be a closed surface of genus g with n punctures, so that $3g - 3 + n \geq 2$ and $(g, n) \neq (1, 2)$. Then the Mapping Class group $\Gamma = \text{Map}(S)$ of S is strongly QI rigid. Moreover, quasisisometries of Γ are uniformly close to automorphisms of Γ .*

Note that for a closed surface S , the group $\text{Map}(S)$ is isomorphic to the group of outer automorphisms $\text{Out}(\pi)$, where $\pi = \pi_1(S)$, see [FM11]. Furthermore, N. Ivanov [Iva88] proved that $\text{Out}(\text{Map}(S))$ is trivial if $3g - 3 + n \geq 2$, $(g, n) \neq (2, 0)$ and $\text{Out}(\text{Map}(S)) \cong \mathbb{Z}_2 \times \mathbb{Z}_2$ if $(g, n) = (2, 0)$. Recall that for a group π , the group of outer automorphisms $\text{Out}(\pi)$ is the quotient

$$\text{Out}(\pi) = \text{Aut}(\pi)/\text{Inn}(\pi)$$

where $\text{Inn}(\pi)$ consists of automorphisms of π given by conjugations *via* elements of π .

PROBLEM 25.65. Is the group $\text{Out}(F_n)$ QI rigid?

Artin groups and Coxeter groups are prominent classes of groups which appear frequently in Geometric Group Theory. Note that some of these groups are not QI rigid, e.g., the group $F_2 \times F_2$, see the above-mentioned examples of Burger and Mozes. In particular, if G is a Coxeter or Artin group which splits as the fundamental group of graph of groups with finite edge groups, where one of the vertex groups G_v is virtually $F_2 \times F_2$, then G cannot be QI rigid. The same applies if one takes a direct product of such G with a Coxeter/Artin group. Also, there are many Coxeter groups which appear as uniform lattices in $O(n, 1)$ (for relatively small n). Such Coxeter groups are QI to non-Coxeter lattices in $O(n, 1)$. This leads to

PROBLEM 25.66. (a) Suppose that G is an Artin group, which does not contain $F_2 \times F_2$. Is such a group G QI rigid?

(b) Suppose that G is a non-hyperbolic 1-ended Coxeter group, which does not contain $F_2 \times F_2$. Is G QI rigid?

Note that Theorem 25.64 implies QI rigidity of *Artin Braid groups* B_n : The quotient of B_n by its center is isomorphic to the mapping class group $\text{Map}(S)$, where S is the 2-sphere with $n + 1$ punctures. QI rigidity results for other classes of Artin groups were obtained by J. Behrstock, T. Januszkiewicz and W. Neumann [BJN09, BJN10], M. Bestvina, B. Kleiner, M. Sageev [BKS08], and J. Huang [Hua14], [Hua16].

M. Gromov and W. Thurston [GT87] constructed interesting examples of closed negatively curved manifolds. The fundamental group of such a manifold

is not isomorphic to a lattice in a Lie group (with finitely many components). We will refer to the manifolds constructed in [GT87] as *Gromov-Thurston manifolds*. Some of these manifolds are obtained as ramified covers over closed hyperbolic n -manifolds ($n \geq 4$), ramified over totally-geodesic submanifolds.

PROBLEM 25.67. Are the fundamental groups of Gromov-Thurston n -manifolds QI rigid?

The reason to be hopeful that these groups Γ are QI rigid is the following. Each Γ is associated with a uniform lattice $\Gamma' < O(n, 1)$ and a sublattice

$$\Gamma'' = \Gamma' \cap O(n - 2, 1).$$

The sublattice Γ'' yields a Γ' -invariant collection of $(n - 2)$ -dimensional hyperbolic subspaces $X_i \subset \mathbb{H}^n$, where X_1 is Γ'' -invariant. (For instance, a Gromov-Thurston manifold can appear as a ramified cover over \mathbb{H}^n/Γ' which is ramified over the submanifold X_1/Γ'' .) While the entire hyperbolic n -space is highly non-rigid, R. Schwartz proved in [Sch97] that the pair $(\mathbb{H}^n, \cup_i X_i)$ is QI rigid.

PROBLEM 25.68. Let S be a closed hyperbolic surface. Let M be the 4-dimensional manifold obtained by taking the 2-fold ramified cover over $S \times S$, which is ramified over the diagonal, see [BGS85, Exercise 1]. Is $\pi_1(M)$ QI rigid?

3-manifold groups. Another class of groups whose QI rigidity properties are relative well (but not completely) understood are fundamental groups of compact 3-manifolds.

THEOREM 25.69 (M. Kapovich, B. Leeb, [KL97]). *The class of fundamental groups G of 3-dimensional closed Haken 3-manifolds, which are not Sol_3 -manifolds, is QI rigid.*

This rigidity theorem was generalized to higher dimensional graph-manifolds by R. Frigerio, J.-F. Lafont and A. Sisto [RF15].

The combination of several rigidity results for 3-manifold groups, leads to:

THEOREM 25.70 (Gromov, Cannon–Cooper, Eskin–Fisher–Whyte, Kapovich–Leeb, Pansu, Papasoglu–Whyte, Schwartz). *The class of fundamental groups of closed connected 3-manifolds is QI rigid.*

PROOF. Suppose that G' is a group QI to $G = \pi_1(M)$, where M is a closed connected 3-dimensional manifold. Without loss of generality, we may assume that M is oriented. Recall that, according to Thurston's Geometrization Conjecture/Perelman's Theorem, the manifold M has the following structure: M splits as a connected sum $M = M_1 \# \dots \# M_k$, where each manifold M_i is either geometric, or is obtained by gluing compact 3-dimensional geometric manifolds with boundary along boundary tori. (See Section 23.7 for more details.)

1. Suppose that $M = M_1$ and M is non-geometric. Then G is VI to the fundamental group of a closed 3-manifold (Theorem 25.69).

2. Suppose that $M = M_1$ and M is geometric. The case when M is hyperbolic is covered by the theorem of Cannon and Cooper (Theorem 23.1). If M has spherical geometric structure, then $\pi_1(M)$ and, hence, G , is finite. If M has the $\mathbb{S}^2 \times \mathbb{R}$ -structure, then G is infinite cyclic or infinite dihedral, i.e. it is 2-ended. Hence,

the group G' is also 2-ended and, therefore, it is VI to \mathbb{Z} , see Proposition 9.23. Of course,

$$\mathbb{Z} \simeq \pi_1(\mathbb{S}^2 \times \mathbb{S}^1)$$

is a 3-manifold group.

The cases of manifolds with Euclidean geometry and Nil_3 -geometry are covered by Gromov–Pansu’s rigidity theorems as follows. Since $\pi_1(M)$ is virtually nilpotent in both cases, by applying Gromov’s theorem, we can assume that G' is also torsion-free nilpotent.

a. If M is a Euclidean manifold, its fundamental group is VI to \mathbb{Z}^3 . Therefore, by Pansu’s theorem 13.11, G' is VI to $\mathbb{Z}^3 \cong \pi_1(T^3)$.

b. If M has Nil_3 -geometry, then the group G (after passing to a finite-index subgroup) is 2-step nilpotent and $C^1(G)/C^2(G) \cong \mathbb{Z}^2$, $C^2(G) \cong \mathbb{Z}$. By Theorem 25.43, the group G' has the same properties. Since \mathbb{Z} -central coextensions of \mathbb{Z}^2 are classified by $H^2(\mathbb{Z}^2, \mathbb{Z}) \cong \mathbb{Z}$, every coextension is isomorphic to a group

$$G_k = \langle a, b, c \mid [a, b] = c^k, [a, c] = 1, [b, c] = 1 \rangle.$$

All these groups (with $k \neq 1$) are commensurable to each other and, hence, to the integer Heisenberg group $H_3(\mathbb{Z})$, which is a cocompact lattice in $Nil_3 = H_3(\mathbb{R})$. If $k = 0$, then $G_k \cong \mathbb{Z}^3$ and, thus, has the \mathbb{E}^3 -geometry.

Finally, QI rigidity of the class of Sol_3 -manifold groups is the content of Theorem 25.53.

3. Suppose that $k \geq 2$, i.e. $\pi_1(M)$ has infinitely many ends. Then the group G splits as an graph of groups with finite edge-groups and vertex groups G_v which have at most one end. This decomposition corresponds to the connected sum decomposition of the manifold M :

$$M = M_1 \# \dots \# M_k,$$

where each M_i either has finite fundamental group, or 2-ended fundamental group (which is either infinite cyclic or infinite dihedral) or 1-ended fundamental group.

According to Theorem 20.47, the group G' also splits as a graph of groups, with finite edge groups, where each vertex group G'_w is either finite or QI to one of the 1-ended vertex groups G_v .

By combining 1, 2 and 3, we conclude that G' splits as a finite graph of groups with finite edge groups and vertex groups G'_w which are VI to the fundamental groups of closed 3-manifolds. Up to passing to a finite-index subgroup, each group G_w has the form

$$1 \rightarrow K_w \rightarrow G'_w \rightarrow \bar{G}_w \rightarrow 1,$$

where \bar{G}_w is a closed 3-manifold group. It is observed in [Kap07] that such groups G'_w are virtually torsion-free and, hence, each contains a finite-index subgroup which is the fundamental group of a closed 3-manifold. Lastly, as in the proof of Theorem 7.51, one assembles all these finite-index subgroups in the vertex groups G'_w , into a finite-index subgroup $H < G'$, which is a free product of fundamental groups of closed 3-manifolds M'_u . The connected sum of the 3-manifolds M'_u is a closed 3-manifold M' with $\pi_1(M') \cong H$. \square

This leaves open the internal QI classification of fundamental groups of closed 3-manifolds.

PROBLEM 25.71. Classify fundamental groups of closed non-geometric irreducible 3-dimensional manifolds up to quasiisometry.

Partial progress towards this problem is achieved in several papers of J. Behrstock and W. Neumann. In the paper [BN08] they proved that the fundamental groups of all *nongeometric* 3-dimensional *graph-manifolds* are QI to each other. They obtained further QI rigidity results for manifolds with hyperbolic components in [BN12].

Most QI rigidity results for fundamental groups of 3-manifolds are obtained under the assumption that the 3-manifolds in question have either empty boundary or a boundary consisting of tori and Klein bottles. However, the problem is also interesting in the case of compact 3-dimensional manifolds with more complex boundary surfaces. In view of the recent results of P. Haissinsky [Hai15], it is reasonable to ask:

QUESTION 25.72. Is the class of fundamental groups of compact 3-dimensional manifolds with non-empty boundary QI rigid?

Haissinsky's results settle the problem in the case of 3-manifolds with Gromov-hyperbolic fundamental groups.

Quasiisometric invariance of group decompositions.

Two sets of theorems below show that, under certain conditions, quasiisometries respect certain graph of groups decompositions and direct product decompositions. In order to state the first of these results, we note that each 1-ended finitely presented group G admits a special decomposition as a graph of groups with 2-ended edge groups. This special decomposition is called the *JSJ decomposition* of G . Decompositions of this type first appeared in the context of 3-dimensional manifolds (and their fundamental groups), in the work of Jaco, Shalen and Johannson, hence the name JSJ. Each finitely presented group G has a unique JSJ decomposition; this decomposition, furthermore, is a refinement of any splitting of G with 2-ended edge groups. We refer the reader to the works of Rips, Sela, Dunwoody, Sageev, Fujiwara and Papasoglu, [RS97, DS99, FP06] for the detailed treatment of group-theoretic JSJ decompositions.

THEOREM 25.73 (P. Papasoglu, [Pap05]). *The class of one-ended finitely presented groups which split over 2-ended groups G is QI rigid. Moreover, the quasiisometries of such a group G preserve the JSJ decomposition of G .*

The second part of this theorem means that if X is a simply-connected tree of spaces associated with the JSJ decomposition of G , then each quasiisometry $X \rightarrow X$ sends each vertex space uniformly close to a vertex space, sends each edge space uniformly close to an edge space, and induces an automorphism of the corresponding tree.

In the case when vertex and edge groups are 2-ended, Mosher, Sageev and Whyte obtained a bit more precise result:

THEOREM 25.74 (L. Mosher, M. Sageev and K. Whyte [MSW03]). *The class of groups which are fundamental groups of finite graphs of 2-ended groups is QI rigid. In particular, each group which is QI to a Baumslag-Solitar group is isomorphic to the fundamental group of such a graph of groups.*

In the follow-up paper [MSW11] L. Mosher, M. Sageev and K. Whyte prove that the class of groups which are fundamental groups of coarse Poincaré Duality groups satisfying certain conditions is QI rigid.

The next theorem deals with the QI invariance of direct products of groups in the context of fundamental groups of closed manifolds M of nonpositive curvature. The de Rham decomposition of the universal cover X of such an M is a canonical decomposition of X into a Riemannian direct product of manifolds of nonpositive curvature

$$X = \mathbb{E}^m \times X_1 \times \dots \times X_k.$$

None of the factors X_i is further decomposable as a Riemannian direct product.

THEOREM 25.75 (M. Kapovich, B. Kleiner, B. Leeb, [KKL98]). *Quasiisometries $X \rightarrow X$ preserve the de Rham decomposition. More precisely:*

1. *Each (L, A) -quasiisometry $f : X \rightarrow X$ sends each Euclidean leaf $\mathbb{E}^m \times \{x\}$ uniformly Hausdorff-close to another leaf $\mathbb{E}^m \times \{x'\}$, where x, x' belong to*

$$\bar{X} = \prod_{i=1}^k X_i.$$

In particular, f induces an (\bar{L}, \bar{A}) -quasiisometry $\bar{f} : \bar{X} \rightarrow \bar{X}$.

2. *Suppose that $f : X \rightarrow X$ is an (L, A) quasiisometry, where X does not have a Euclidean de Rham factor. Then, after composing f with a permutation of the factors X_i if necessary, the map f is uniformly close to a product map*

$$f_1 \times \dots \times f_k,$$

where each $f_i : X_i \rightarrow X_i$ is a quasiisometry.

On the group-theoretic side:

COROLLARY 25.76. *Suppose that X does not have a Euclidean de Rham factor ($m = 0$) and that the manifold M splits as a direct product $M = M_1 \times \dots \times M_k$, where each M_i has the universal cover X_i . Accordingly, the group G splits as a direct product $G = G_1 \times \dots \times G_k$. Then each quasiisometry $G \rightarrow G$ preserves the direct product decomposition of G in the same sense as a quasiisometry $X \rightarrow X$ preserves the de Rham decomposition.*

Note that a group acting geometrically on X (say, $G = \pi_1(M)$) need not contain a finite-index subgroup which splits as a direct product. This happens, for instance, in the case of irreducible lattices in semisimple Lie groups.

Rigidity for quasiisometric embeddings. Sometimes not only quasiisometries between metric spaces, but even quasiisometric embeddings, exhibit surprising rigidity properties. For instance, it was proven by D. Fisher and K. Whyte [FW14], that for some classes of (nonisometric) symmetric spaces of equal rank X, Y , every quasiisometric embedding $X \rightarrow Y$ is within finite distance from a totally geodesic embedding. For instance, they prove:

THEOREM 25.77. *Suppose that X, Y are non-positively curved irreducible symmetric spaces of equal rank with non-exceptional root systems. Then, unless the root system of X is of type A_n and Y is of type B_n, C_n or BC_n , every quasiisometric embedding $X \rightarrow Y$ is within finite distance from a totally geodesic embedding.*

D. Fisher and T. Nguyen proved in [FN15] that for certain classes of higher rank non-uniform lattices, every quasiisometric embedding is within finite distance from a group homomorphism.

Recall that a group G is *co-hopfian* if every injective homomorphism $G \rightarrow G$ is surjective.

DEFINITION 25.78 (I. Kapovich). A metric space X is *coarsely co-hopfian* if every quasiisometric embedding $X \rightarrow X$ is a quasiisometry.

For instance, one can show that Poincaré Duality groups are coarsely co-hopfian. The same holds for some classes of relative Poincaré Duality groups, see [KL12]. Recently, the coarse co-hopfian property was verified by S. Merenkov [Mer10] for some classes of Gromov–hyperbolic spaces (whose ideal boundaries are homeomorphic to Sierpinsky carpets). One can ask what are other “interesting” examples of coarsely co-hopfian spaces and groups. Here is a concrete open problem:

PROBLEM 25.79. Let X be a thick hyperbolic building. Is it true that X is coarsely co-hopfian?

We conclude with a table of algebraic and geometric properties/invariants/classes of finitely generated groups in relation to their QI invariance. Most of the definitions in this table were introduced earlier in the book. The missing ones are:

1. A group G is *noetherian* if every subgroup of G is finitely generated (sometimes, such groups are called *slender*). Notable examples of noetherian groups are polycyclic groups and Tarski monsters (finitely generated groups where every proper subgroup is cyclic, see [Ol’91a]).
2. A group G is *coherent* if every finitely generated subgroup of G is finitely presented (sometimes, such groups are called *artinian*). Notable examples of such groups are the fundamental groups of compact 3-dimensional manifolds and polycyclic groups.
3. A group G is *hyperbolike* if it is a direct limit of a sequence of epimorphisms of hyperbolic groups:

$$G_1 \twoheadrightarrow G_2 \twoheadrightarrow G_3 \twoheadrightarrow \dots$$

Many examples of such groups are *lacunary hyperbolic*, i.e. finitely generated groups for which one asymptotic cone is a tree, see [OOS09]. Hyperbolike groups appear frequently in constructions of group-theoretic *monsters*, which are finitely generated groups satisfying some exotic properties (see e.g. [Ol’91a]).

4. A finitely generated group G is said to satisfy *Yu’s Property A* if G admits a uniformly proper embedding in a Hilbert space, see [Yu00] (there are many other definitions).

5. *Shalom’s Property H_{FD}* : A group G is said to satisfy the Property H_{FD} if for every unitary representation $\pi : G \rightarrow U(\mathcal{H})$ of G with $H^1(G, \mathcal{H}_\pi) \neq 0$, there exists a G -invariant finite-dimensional non-zero subspace $\mathcal{H}' \subset \mathcal{H}$. Shalom proved in [Sha04] that Property H_{FD} is a QI invariant among amenable groups.

6. It would take too much space here to define *Poincaré duality groups* and *duality groups*; we refer the reader instead to Brown’s book [Bro82b]. The same applies to *semihyperbolic* and *automatic groups*; the reader is referred to [JA95], [BH99] and [ECH⁺92]. Both properties capture some features of nonpositive curvature.

QI invariant	Not QI invariant	Unknown
hyperbolic, Corollary 11.43		CAT(-1)
semihyperbolic	CAT(0), Example 25.30	automatic
relatively hyperbolic, Theorem 25.37		hyperbolike
virtually nilpotent, Theorem 16.25	virtually solvable, Theorem 14.40	virtually polycyclic
solvable word problem	solvable conjugacy problem see [CM77]	solvable isomorphism problem
	simple	
virtually free, Theorem 20.45	colarge	small
finite	residually finite	torsion
	virtually torsion-free	bounded generation property
amenable, Theorem 18.13	virtually metabelian, Theorem 14.40	elementary amenable
	hopfian	noetherian
	co-hopfian	coherent
Yu's Property A	Properties T & Haagerup, Theorem 19.76 and [Car14]	
Property H_{FD} for amenable groups		
amalgam/HNN with finite edge groups	amalgam/HNN	contains proper infinite subgroups of infinite index
virtually splits with virtually cyclic edge groups	virtually splits	
Type \mathbf{F}_n , Theorem 9.56		
virtually a closed surface group, Corollary 23.24		
VI to a closed 3-manifold group, Theorem 25.70		VI to a compact 3-manifold group
cohomological dimension n over \mathbb{Q} , Theorem 9.64		Poincaré duality group over \mathbb{Q}
		duality group over \mathbb{Q}

Appendix by Bogdan Nica: Three theorems on linear groups

Bogdan Nica

Mathematisches Institut, Georg-August Universität Göttingen, Germany.

Email: bogdan.nica@gmail.com

26.1. Introduction

Recall that a group is *linear* if it is (isomorphic to) a subgroup of $GL_n(\mathbb{K})$, where \mathbb{K} is a field. If we want to specify a field, we say that a group is linear over \mathbb{K} . The following theorems are fundamental, at least from the perspective of the combinatorial group theory.

THEOREM 26.1 (A. I. Mal'cev, 1940). *Every finitely generated linear group is residually finite.*

THEOREM 26.2 (A. Selberg, 1960). *Every finitely generated linear group over a field of zero characteristic is virtually torsion-free.*

A group is *residually finite* if its elements are distinguished by the finite quotients of the group, i.e. if each non-trivial element of the group remains non-trivial in a finite quotient. A group is *virtually torsion-free* if some finite-index subgroup is torsion-free. As a matter of further terminology, Selberg's theorem is usually referred to as Selberg's lemma, and Mal'cev is alternatively transliterated as Malcev or Maltsev.

Residual finiteness and virtual torsion-freeness are related to a third property — roughly speaking, a “*p*-adic” refinement of residual finiteness. A theorem due to V. Platonov (1968) gives such refined residual properties for finitely generated linear groups. Both Mal'cev's theorem and Selberg's lemma are consequences of this more powerful, but lesser known, theorem of Platonov.

Once we have Platonov's theorem and its proof, we are not too far from our third theorem of interest. In order to formulate it, let us first observe that every non-trivial torsion element in a group G gives rise to a non-trivial idempotent in the complex group algebra $\mathbb{C}G$. Namely, if $g \in G$ has order $n > 1$, then

$$e = \frac{1}{n}(1 + g + \dots + g^{n-1}) \in \mathbb{C}G$$

satisfies $e^2 = e$, and $e \neq 0, 1$. The *Idempotent Conjecture* is the bold statement that the converse also holds:

CONJECTURE 26.3 (Idempotent Conjecture). *If G is a torsion-free group, then the group algebra $\mathbb{C}G$ has no non-trivial idempotents.*

While not yet settled in general, this conjecture is known for many classes of groups. A particularly important partial result is proven by H. Bass in [Bas76]:

THEOREM 26.4. *Torsion-free linear groups satisfy the Idempotent Conjecture.*

26.2. Virtual and residual properties of groups

Virtual torsion-freeness and residual finiteness are instances of the following terminology. Let \mathcal{P} be a group-theoretic property. A group is *virtually* \mathcal{P} if it has a finite-index subgroup enjoying \mathcal{P} . A group is *residually* \mathcal{P} if each non-trivial element of the group remains non-trivial in some quotient group enjoying \mathcal{P} . The virtually \mathcal{P} groups and the residually \mathcal{P} groups contain the \mathcal{P} groups. It may certainly happen that a property is virtually stable (e.g., finiteness) or residually stable (e.g., torsion-freeness).

Besides virtual torsion-freeness and residual finiteness, we are interested in the hybrid notion of *virtual residual p -finiteness* where p is a prime. This is obtained by residualizing the property of being a finite p -group, followed by the virtual extension. The notion of virtual residual p -finiteness has, in fact, a leading role in this account for it relates both to residual finiteness and to virtual torsion-freeness.

Observe the following:

(Going down) If \mathcal{P} is inherited by subgroups, then both virtually \mathcal{P} and residually \mathcal{P} are inherited by subgroups. In particular, virtual torsion-freeness, residual finiteness, and virtual residual p -finiteness are inherited by subgroups.

(Going up) Virtually \mathcal{P} passes also to finite-index supergroups. In particular, both virtual torsion-freeness and virtual residual p -finiteness pass to finite-index supergroups. Residual finiteness passes to finite-index supergroups.

Residual p -finiteness trivially implies residual finiteness. Going up, we obtain:

LEMMA 26.5. *Virtual residual p -finiteness for some prime p implies residual finiteness.*

On the other hand, residual p -finiteness imposes torsion restrictions. Namely, in a residually p -finite group, the order of a torsion element must be a p -th power. Hence, if a group is residually p -finite and residually q -finite for two different primes p and q , then it is torsion-free. Virtualizing this statement, we obtain:

LEMMA 26.6. *Virtual residual p -finiteness and virtual residual q -finiteness for two primes $p \neq q$ imply virtual torsion-freeness.*

26.3. Platonov's theorem

In light of Lemmas 26.5 and 26.6, we see that Mal'cev's theorem and Selberg's lemma are consequences of the following:

THEOREM 26.7 (Platonov, 1968). *Let G be a finitely generated linear group over a field \mathbb{K} . If $\text{char } \mathbb{K} = 0$, then G is virtually residually p -finite for all but finitely many primes p . If $\text{char } \mathbb{K} = p$, then G is virtually residually p -finite.*

Actually, the zero characteristic part of Platonov's theorem had been proved slightly earlier by Kargapolov (1967) and, independently, Merzlyakov (1967).

EXAMPLE 26.8. (Cf. Example 7.111) $SL_n(\mathbb{Z})$, where $n \geq 2$, is a finitely generated linear group over \mathbb{Q} . Reduction modulo a positive integer N defines a group homomorphism $SL_n(\mathbb{Z}) \rightarrow SL_n(\mathbb{Z}/N)$, whose kernel

$$\Gamma(N) := \text{Ker}(SL_n(\mathbb{Z}) \rightarrow SL_n(\mathbb{Z}/N)) = \{X \in SL_n(\mathbb{Z}) : X \equiv 1_n \pmod{N}\}$$

is the *principal congruence subgroup of level N* . The principal congruence subgroups are finite-index, normal subgroups of $SL_n(\mathbb{Z})$. They are organized according to the divisibility of their levels: $\Gamma(M) \supseteq \Gamma(N)$ if and only if $M|N$, that is, “to contain is to divide”. Hence the prime stratum $\{\Gamma(p) : p \text{ prime}\}$, and each descending chain $\{\Gamma(p^k) : k \geq 1\}$ corresponding to fixed prime p , stand out.

Elements of $SL_n(\mathbb{Z})$ can be distinguished both along the prime stratum,

$$\bigcap_p \Gamma(p) = \{1_n\},$$

as well as along each descending p -chain,

$$\bigcap_k \Gamma(p^k) = \{1_n\}.$$

We thus have two ways of seeing that $SL_n(\mathbb{Z})$ is residually finite.

There is no prime p for which $SL_n(\mathbb{Z})$ is residually p -finite, simply because the matrix

$$\begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix}$$

has order 6. However, $SL_n(\mathbb{Z})$ is virtually residually p -finite for each prime p . The reason is that $\Gamma(p)$ is residually p -finite, and this is easily seen by noting that each successive quotient $\Gamma(p^k)/\Gamma(p^{k+1})$ in the descending chain $\{\Gamma(p^k) : k \geq 1\}$ is a p -group: for $X \in \Gamma(p^k)$ we have

$$X^p = 1_n + \sum_{i=1}^p \binom{p}{i} (X - 1_n)^i \in \Gamma(p^{k+1}).$$

EXAMPLE 26.9. $SL_n(\mathbb{F}_p[t])$, where $n \geq 2$, is linear over $\mathbb{F}_p(t)$ and finitely generated for $n \geq 3$ (though not for $n = 2$). A similar argument to the one of the previous example, this time involving the principal congruence subgroups corresponding to the descending chain of ideals (t^k) for $k \geq 1$, shows that $SL_n(\mathbb{F}_p[t])$ is virtually residually p -finite. On the other hand, $SL_n(\mathbb{F}_p[t])$ contains a copy of the infinite torsion group $(\mathbb{F}_p[t], +)$, and this prevents $SL_n(\mathbb{F}_p[t])$ from being virtually torsion-free. Consequently, $SL_n(\mathbb{F}_p[t])$ cannot be virtually residually q -finite for any prime $q \neq p$.

Platonov’s theorem implies the following “ p -adic” refinement of Mal’cev’s theorem.

COROLLARY 26.10. *A finitely generated linear group is virtually residually p -finite for some prime p .*

This corollary, combined with Example 26.9, leads us to a simple example of a finitely generated group which is non-linear but residually finite:

$$SL_n(\mathbb{F}_p[t]) \times SL_n(\mathbb{F}_q[t]),$$

where p and q are different primes, and $n \geq 3$.

26.4. Proof of Platonov’s theorem

Let G be a finitely generated linear group over a field \mathbb{K} , say $G \leq GL_n(\mathbb{K})$. In \mathbb{K} , consider the subring A generated by the multiplicative identity 1 and the matrix entries of a finite, symmetric set of generators for G . Thus A is a finitely generated

domain, and G is a subgroup of $GL_n(A)$. Platonov's theorem is then a consequence of the following:

THEOREM 26.11. *Let A be a finitely generated domain. If $\text{char } A = 0$, then $GL_n(A)$ is virtually residually p -finite for all but finitely many primes p . If $\text{char } A = p$, then $GL_n(A)$ is virtually residually p -finite.*

Here, and for the remainder of the section, rings are commutative and unital. The proof of Theorem 26.11 is a straightforward variation on the example of $SL_n(\mathbb{Z})$, as soon as we know the following facts:

LEMMA 26.12. *Let A be a finitely generated domain. Then the following hold:*

- i. A is noetherian.
- ii. $\bigcap_k I^k = 0$ for any ideal $I \neq A$.
- iii. If A is a field, then A is finite.
- iv. The intersection of all maximal ideals of A is 0.
- v. If $\text{char } A = 0$, then only finitely many primes $p = p \cdot 1$ are invertible in A .

Let us postpone the proof of Lemma 26.12 for the moment, and focus instead on deriving Theorem 26.11. The principal congruence subgroup of $GL_n(A)$ corresponding to an ideal I of A is defined by

$$\Gamma(I) = \text{Ker}(GL_n(A) \rightarrow GL_n(A/I)).$$

If π is a maximal ideal then A/π is a finite field, by Lemma 26.12 iii, so $\Gamma(\pi)$ has finite index in $GL_n(A)$. Also,

$$\bigcap_{\pi} \Gamma(\pi) = \{1_n\}$$

as π runs over the maximal ideals of A , by Lemma 26.12 iv. This shows that $GL_n(A)$ is residually finite, thereby proving Mal'cev's theorem.

For each $k \geq 1$, the quotient π^k/π^{k+1} is naturally an A/π -module. It inherits finite generation from the finite generation of the A -module π^k , the latter due to A being noetherian. As A/π is finite, π^k/π^{k+1} is finite as well. It follows that the ring A/π^k is finite, and so $\Gamma(\pi^k)$ has finite index in $GL_n(A)$. Furthermore,

$$\bigcap_k \Gamma(\pi^k) = \{1_n\}$$

by Lemma 26.12 ii, which shows once again that $GL_n(A)$ is residually finite. Now let p denote the characteristic of A/π , so $p = p \cdot 1 \in \pi$. Then $\Gamma(\pi^k)/\Gamma(\pi^{k+1})$ is a p -group: for $X \in \Gamma(\pi^k)$ we have

$$X^p = 1_n + \sum_{i=1}^p \binom{p}{i} (X - 1_n)^i \in \Gamma(\pi^{k+1}).$$

To conclude, $GL_n(A)$ is virtually residually p -finite for each prime p not invertible in A . By Lemma 26.12 v, this happens for all but finitely many primes p in the zero characteristic case. In characteristic p , there is only such prime, namely p itself. Theorem 26.11 is proved. \square

We now return to the proof of the lemma.

PROOF OF LEMMA 26.12. The first two points are standard: i) follows from the Hilbert Basis Theorem, and ii) is the Krull Intersection Theorem for domains; see e.g. [Abh06], p. 223.

iii) We claim the following: If $F \subseteq F(u)$ is a field extension with $F(u)$ finitely generated as a ring, then $F \subseteq F(u)$ is a finite extension and F is finitely generated as a ring.

We use the claim as follows. Let F be the prime field of A and let a_1, \dots, a_k be generators of A as a ring. Thus $A = F(a_1, \dots, a_k)$. Going down the chain

$$A = F(a_1, \dots, a_k) \supseteq F(a_1, \dots, a_{k-1}) \supseteq \dots \supseteq F,$$

we obtain that $F \subseteq A$ is a finite extension, and that F is finitely generated as a ring. Then F is a finite field, as \mathbb{Q} is not finitely generated as a ring, and so A is finite.

Now let us prove the claim. Assume that u is transcendental over F , i.e. $F(u)$ is the field of rational functions in u . Let $P_1/Q_1, \dots, P_k/Q_k$ generate $F(u)$ as a ring, where $P_i, Q_i \in F[u]$. The multiplicative inverse of $1 + u \cdot \prod Q_i$ is a polynomial expression in the P_i/Q_i 's, which can be written as $R/\prod Q_i^{s_i}$. Therefore,

$$\prod Q_i^{s_i} = (1 + u \cdot \prod Q_i)R$$

in $F[u]$. But this is impossible, since $\prod Q_i^{s_i}$ is relatively prime to $1 + u \cdot \prod Q_i$.

Thus u is algebraic over F . Let

$$X^d + \alpha_1 X^{d-1} + \dots + \alpha_d$$

be the minimal polynomial of u over F . Let also a_1, \dots, a_k be ring generators of $F(u) = F[u]$. We may write each a_i as

$$\sum_{0 \leq m \leq d-1} \beta_{i,m} u^m$$

with $\beta_{i,m} \in F$. We claim that the α_j 's and the $\beta_{i,m}$'s are ring generators of F . Let $c \in F$. Then c is a polynomial in a_1, \dots, a_k over F , hence a polynomial in u over the subring of F generated by the $\beta_{i,m}$'s, hence a polynomial in u of degree less than d over the subring of F generated by the α_j 's and the $\beta_{i,m}$'s. By the linear independence of $\{1, u, \dots, u^{d-1}\}$, the latter polynomial is actually of degree 0. Hence c ends up in the subring of F generated by the α_j 's and the $\beta_{i,m}$'s.

iv) Let $a \neq 0$ in A . To find a maximal ideal of A not containing a , we rely on the basic avoidance: maximal ideals do not contain invertible elements. Consider the localization $A' = A[1/a]$. Let π' be a maximal ideal in A' , so $a \notin \pi'$. The restriction $\pi = \pi' \cap A$ is an ideal in A , and $a \notin \pi$. We show that π is maximal. The embedding $A \hookrightarrow A'$ induces an embedding $A/\pi \hookrightarrow A'/\pi'$. As A'/π' is a field which is finitely generated as a ring, it follows from iii) that A'/π' is finite field. Therefore the subring A/π is a finite domain, hence a field as well.

v) We shall use Noether's Normalization Theorem: If R is a finitely generated algebra over a field $F \subseteq R$, then there are elements $x_1, \dots, x_k \in R$ algebraically independent over F such that R is integral over $F[x_1, \dots, x_k]$; see e.g. [Abh06], p. 248.

In our case, \mathbb{Z} is a subring of A , and A is an integral domain which is finitely generated as a \mathbb{Z} -algebra. Extending to rational scalars, we have that $A_{\mathbb{Q}} = \mathbb{Q} \otimes_{\mathbb{Z}} A$ is a finitely generated \mathbb{Q} -algebra. By the Normalization Theorem, there exist elements x_1, \dots, x_k in $A_{\mathbb{Q}}$ which are algebraically independent over \mathbb{Q} , and such that $A_{\mathbb{Q}}$ is integral over $\mathbb{Q}[x_1, \dots, x_k]$. Up to replacing each x_i by an integral multiple of itself, we may assume that x_1, \dots, x_k are in A . There is some positive $m \in \mathbb{Z}$ such that each ring generator of A is integral over $\mathbb{Z}[1/m][x_1, \dots, x_k]$. Thus $A[1/m]$ is integral

over the subring $\mathbb{Z}[1/m][x_1, \dots, x_k]$. If a prime p is invertible in A , then it is also invertible in $A[1/m]$ while at the same time $p \in \mathbb{Z}[1/m][x_1, \dots, x_k]$.

Now we use the following general fact. Let R be a ring which is integral over a subring S . If $s \in S$ is invertible in R , then s is already invertible in S . The proof is easy. Let $r \in R$ with $rs = 1$. We have

$$r^d + s_1 r^{d-1} + \dots + s_{d-1} r + s_d = 0$$

for some $s_i \in S$, since r is integral over S . Multiplying through by s^{d-1} yields $r \in S$.

Returning to our proof, we infer that p is invertible in $\mathbb{Z}[1/m][x_1, \dots, x_k]$. By the algebraic independence of x_1, \dots, x_k , it follows that p is actually invertible in $\mathbb{Z}[1/m]$. But only finitely many primes have this property, namely the prime factors of m . \square

26.5. The Idempotent Conjecture for linear groups

Our approach to Bass's theorem relies on the following criterion of E. Formanek [For73], whose proof is postponed till the next section.

THEOREM 26.13 (E. Formanek, 1973). *Let G be a torsion-free group with the property that, for infinitely many primes p , G has no p -self-similar elements. Then the Idempotent Conjecture holds for G .*

Given a group G , we say that a non-trivial element $g \in G$ is *self-similar* if g is conjugate in G to a proper power g^N , where $N \geq 2$. Clearly, torsion elements are self-similar. It turns out that the converse holds for linear groups in positive characteristic.

LEMMA 26.14. *In a linear group over a field of positive characteristic, every self-similar element is torsion.*

PROOF. Let $\text{char } \mathbb{K} = p$, and consider the relation $g^N = x^{-1}gx$ in $GL_n(\mathbb{K})$, where $N \geq 2$. Without loss of generality, \mathbb{K} is algebraically closed and g is in Jordan normal form. Each Jordan block is of the form $\lambda \cdot 1_k + \Delta_k$, where Δ_k is the $k \times k$ -matrix with 1's on the super-diagonal and 0's everywhere else. Since

$$(\lambda \cdot 1_k + \Delta_k)^{p^s} = \lambda^{p^s} \cdot 1_k + \Delta_k^{p^s}$$

and $\Delta_k^{p^s} = 0$ for large enough s , it follows that g^{p^s} is diagonal for large enough s . Thus, up to replacing g by g^{p^s} , we may assume that g is diagonal. So let g have $\lambda_1, \dots, \lambda_n \in \mathbb{K}$ along the diagonal, and write out the relation $gx = xg^N$ in matrix form: $(x_{ij} \lambda_i) = (x_{ij} \lambda_j^N)$. Compare the i -th row on the two sides. At least one of $x_{i1}, x_{i2}, \dots, x_{in}$ is non-zero, hence $\lambda_i = \lambda_{\sigma(i)}^N$ for some $\sigma(i) \in \{1, \dots, n\}$. Since $\sigma^s = \sigma^{s+t}$ for some positive integers s and t , it follows that

$$\lambda_i = \lambda_{\sigma^{s+t}(i)}^{N^{s+t}} = (\lambda_{\sigma^s(i)}^{N^s})^{N^t} = \lambda_i^{N^t}$$

for each i . We conclude that $g^{N^t-1} = 1$ in $GL_n(\mathbb{K})$. \square

In characteristic zero, a linear group may contain self-similar elements of infinite order. A simple example in, say, $GL_2(\mathbb{R})$ is provided by the matrix

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix},$$

which is conjugate into its N -th power by

$$\begin{pmatrix} 1 & 0 \\ 0 & N \end{pmatrix}.$$

EXERCISE 26.15. Show that the entire subgroup generated by these two matrices is torsion-free.

The analogue of Lemma 26.14 in characteristic zero involves the following refined notion of self-similarity. Given a group G and a prime p , let us say that a non-trivial element $g \in G$ is *p -self-similar* if g is conjugate in G to a proper p -th power g^{p^k} , where $k \geq 1$.

LEMMA 26.16. *In a finitely generated linear group over a field of characteristic zero, the following holds for all but finitely many primes p : Every p -self-similar element is torsion.*

PROOF. The characteristic zero case of Platonov's theorem reduces the claim to showing that, in a virtually residually p -finite group, every p -self-similar element is torsion. This easily follows from the observation that a residually p -finite group has no p -self-similar elements. \square

The upshot of Lemmas 26.14 and 26.16 is that a finitely generated, torsion-free linear group comfortably meets the requirement of Formanek's criterion, and so it satisfies the Idempotent Conjecture. The theorem of Bass follows.

26.6. Proof of Formanek's criterion

The proof of Theorem 26.13 uses tracial methods. Let us first recall that a *trace* on a \mathbb{K} -algebra \mathcal{A} is a \mathbb{K} -linear map $\tau : \mathcal{A} \rightarrow \mathbb{K}$ with the property that $\tau(ab) = \tau(ba)$ for all $a, b \in \mathcal{A}$. In short, traces are linear functionals which vanish on commutators. The ersatz commutativity afforded by a trace is extremely valuable in a noncommutative world.

On a group algebra $\mathbb{K}G$, the *standard trace* $\text{tr} : \mathbb{K}G \rightarrow \mathbb{K}$ is the linear functional which records the coefficient of the identity element:

$$\text{tr}\left(\sum a_g g\right) = a_1.$$

In general, traces on $\mathbb{K}G$ are in bijective correspondence with maps $G \rightarrow \mathbb{K}$ which are constant on conjugacy classes. The characteristic map $1_C : G \rightarrow \mathbb{K}$ of a conjugacy class $C \subseteq G$ defines the trace

$$\tau_C\left(\sum a_g g\right) = \sum_{g \in C} a_g,$$

thus, $\text{tr} = \tau_{\{1\}}$ with this notation. The traces τ_C , where C runs over the conjugacy classes of G , provide a natural basis for the \mathbb{K} -linear space formed by the traces of $\mathbb{K}G$. Another distinguished trace is the *augmentation map* $\epsilon : \mathbb{K}G \rightarrow \mathbb{K}$, given by

$$\epsilon\left(\sum a_g g\right) = \sum a_g.$$

This is the trace on $\mathbb{K}G$ defined by the constant map $1 : G \rightarrow \mathbb{K}$. The augmentation map is in fact a unital \mathbb{K} -algebra homomorphism, hence ϵ is a trace which is $\{0, 1\}$ -valued on idempotents.

Understanding the range of the standard trace on idempotents is much more difficult. The following theorem addresses this problem in the case of complex group algebras.

THEOREM 26.17 (I. Kaplansky, 1969). *Let e be an idempotent in $\mathbb{C}G$. Then $\text{tr}(e) \in [0, 1]$. Furthermore, $\text{tr}(e) = 0$ if and only if $e = 0$, and $\text{tr}(e) = 1$ if and only if $e = 1$.*

Now let us return to the proof of Formanek's criterion. It consists of two steps.

(Positive characteristic claim) Fix a prime p . If G has no p -self-similar elements and \mathbb{K} is a field of characteristic p , then the standard trace is $\{0, 1\}$ -valued on the idempotents of $\mathbb{K}G$.

It is a familiar fact that the identity $(a+b)^p = a^p + b^p$ holds in any commutative \mathbb{K} -algebra. Its noncommutative generalization, somewhat lesser known, says that, in a \mathbb{K} -algebra, $(a+b)^p - a^p - b^p$ is a sum of commutators. Indeed, we may assume that we are in the free \mathbb{K} -algebra on a and b . We expand $(a+b)^p$ into monomials of degree p in a and b , and we let the cyclic group of order p act on these monomials by cyclic permutations. We see orbits of size p , except for a^p and b^p , which are fixed by the action. Now we observe that the sum of monomials corresponding to each orbit of size p is a sum of commutators. This follows from the identity

$$\begin{aligned} & x_1 x_2 \dots x_{p-1} x_p + x_2 x_3 \dots x_p x_1 + \dots + x_p x_1 \dots x_{p-2} x_{p-1} \\ &= p \cdot x_1 x_2 \dots x_{p-1} x_p - [x_1, x_2 \dots x_p] - [x_1 x_2, x_3 \dots x_p] - \dots - [x_1 \dots x_{p-1}, x_p]. \end{aligned}$$

Next, let us iterate: We show by induction that $(a+b)^{p^k} - a^{p^k} - b^{p^k}$ is a sum of commutators for every positive integer k . For the induction step we write

$$(a+b)^{p^{k+1}} = \left(a^{p^k} + b^{p^k} + \sum [u_i, v_i] \right)^p = a^{p^{k+1}} + b^{p^{k+1}} + \sum [u_i, v_i]^p + \sum [u'_j, v'_j]$$

and

$$[u, v]^p = (uv)^p - (vu)^p + \sum [y_l, z_l] = [(uv)^{p-1} u, v] + \sum [y_l, z_l].$$

In particular, a trace τ on a \mathbb{K} -algebra has the property that

$$\tau((a+b)^{p^k}) = \tau(a^{p^k}) + \tau(b^{p^k})$$

for every positive integer k . For a basic trace τ_C , where $C \neq \{1\}$, and an idempotent $e \in \mathbb{K}G$, we obtain

$$\tau_C(e) = \tau_C(e^{p^k}) = \tau_C\left(\left(\sum e_g g\right)^{p^k}\right) = \sum \tau_C((e_g g)^{p^k}) = \sum e_g^{p^k} 1_C(g^{p^k})$$

for each positive integer k . The hypothesis that G has no p -self-similar elements implies that, for each g in the support of e , there is at most one k so that $g^{p^k} \in C$. Thus, taking k large enough, we see that $\tau_C(e) = 0$. Using the relation

$$\epsilon = \text{tr} + \sum_{C \neq \{1\}} \tau_C,$$

we conclude that tr is $\{0, 1\}$ -valued on the idempotents of $\mathbb{K}G$.

(Zero characteristic claim) Assume that, for infinitely many primes p , the following holds: The standard trace is $\{0, 1\}$ -valued on the idempotents of $\mathbb{K}G$, whenever \mathbb{K} is a field of characteristic p . Then the standard trace is $\{0, 1\}$ -valued on the idempotents of $\mathbb{C}G$.

Arguing by contradiction, we assume that e is an idempotent in $\mathbb{C}G$ with $e_1 = \text{tr}(e) \notin \{0, 1\}$. Let $A \subseteq \mathbb{C}$ be the subring generated by the support of e together with $1/e_1$ and $1/(1 - e_1)$, and view e as an idempotent in the group ring AG . By Lemma 26.12 v, for all but finitely many primes p there is a quotient map $A \rightarrow \mathbb{K}$, $a \mapsto \bar{a}$, onto a field of characteristic p . Note that $\bar{e}_1 \neq 0, 1$ in \mathbb{K} , since e_1 and $1 - e_1$ are invertible in A . The induced ring homomorphism $AG \rightarrow \mathbb{K}G$ sends e to an idempotent \bar{e} in $\mathbb{K}G$ with $\text{tr}(\bar{e}) \neq 0, 1$, thereby contradicting our hypothesis.

The proof of Theorem 26.13 is concluded by invoking Kaplansky's theorem. \square

26.7. Notes

Platonov's theorem. Besides the Russian original [Pla68], the only other source in the literature for Platonov's theorem appears to be the presentation by B. A. F. Wehrfritz in [Weh73]. The proof presented herein seems considerably simpler. It is mainly influenced by the discussion of Mal'cev's theorem in lecture notes by Stallings [Sta00], and it has a certain degree of similarity with Platonov's own arguments in [Pla68].

Selberg's lemma. It is important to note that Selberg's lemma is just a minor step in Selberg's paper [Sel60], whose true importance is that it started the rich stream of rigidity results for lattices in higher rank semisimple Lie groups. An alternative road to Selberg's lemma is to use valuations. This is the approach taken by J. W. Cassels in [Cas86] and by J. Ratcliffe in [Rat06].

The Idempotent Conjecture. The Idempotent Conjecture is usually attributed to Kaplansky, but a reference seems elusive. What Kaplansky did state on more than one occasion (Problem 1, p. 122 in [Kap69], and Problem 6, p. 448 in [Kap70]) is a problem nowadays referred to as the

CONJECTURE 26.18 (Zero-Divisor Conjecture). *If G is a torsion-free group and \mathbb{K} is a field, then the group algebra $\mathbb{K}G$ has no zero-divisors, i.e. $ab \neq 0$ whenever $a, b \neq 0$ in $\mathbb{K}G$.*

The Zero-Divisor Conjecture over the complex field, which clearly implies the Idempotent Conjecture, is still not settled for the class of (torsion-free) linear groups.

Kaplansky's theorem. We refer to M. Burger and A. Valette [BV98] for a proof, as well as for a nice complementary reading. The main insight of Kaplansky's analytic proof is to pass from the group algebra $\mathbb{C}G$ to a completion afforded by the regular representation on $\ell^2 G$. One can use the weak completion, that is the von Neumann algebra LG , or the norm completion, the so-called reduced C^* -algebra C_r^*G . Kaplansky's proof, while remarkable in itself, is perhaps more important for suggesting what came to be known as the **Kadison Conjecture**:

CONJECTURE 26.19 (Kadison Conjecture). *For every torsion-free group G , the reduced C^* -algebra C_r^*G has no non-trivial idempotents.*

At the time of writing, the Kadison Conjecture for the class of (torsion-free) linear groups is still open.

Bass's theorem. As we have seen, the step from Formanek's criterion to the theorem of Bass is rather short, and it uses results on linear groups which were known — certainly on the eastern side of the Iron Curtain, but probably also on

its western side — at the time of [For73]. Ascribing the theorem to Bass *and* Formanek is therefore not entirely unwarranted. The hard facts, however, are that Bass [Bas76] actually proves much more whereas Formanek [For73] states less.

Bibliography

- [AB56] L. V. Ahlfors and A. Beurling, *The boundary correspondence under quasiconformal mappings*, Acta Math. **96** (1956), 125–142.
- [ABD⁺13] A. Abrams, N. Brady, P. Dani, M. Duchin, and R. Young, *Pushing fillings in right-angled Artin groups*, J. Lond. Math. Soc. (2) **87** (2013), no. 3, 663–688.
- [ABDY13] A. Abrams, N. Brady, P. Dani, and R. Young, *Homological and homotopical Dehn functions are different*, Proc. Natl. Acad. Sci. USA **110** (2013), no. 48, 19206–19212.
- [Abh06] S. S. Abhyankar, *Lectures on algebra. Vol. I*, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2006.
- [AD] G.N. Arzhantseva and T. Delzant, *Examples of random groups*, preprint ARXIV:0711.4238, to appear in Journal of Topology.
- [AD82] P. Assouad and M. Deza, *Metric subspaces of L^1* , Publications mathématiques d’Orsay **3** (1982), iv+47.
- [Ado36] I. D. Ado, *Note on the representation of finite continuous groups by means of linear substitutions*, Bull. Soc. Phys.-Math. Kazan **7** (1936), 3–43.
- [Ady79] S. I. Adyan, *The Burnside problem and identities in groups*, Springer Verlag, 1979.
- [Ady82] ———, *Random walks on free periodic groups*, Izv. Akad. Nauk SSSR Ser. Mat. **46** (1982), no. 6, 1139–1149.
- [Aea91] J. M. Alonso and et al., *Notes on word hyperbolic groups*, Group theory from a geometrical viewpoint (Trieste, 1990), World Sci. Publ., River Edge, NJ, 1991, Edited by H. Short, pp. 3–63.
- [AG99] D. Allcock and S. Gersten, *A homological characterization of hyperbolic groups*, Invent. Math. **135** (1999), 723–742.
- [Ago13] I. Agol, *The virtual Haken conjecture*, Doc. Math. **18** (2013), 1045–1087, With an appendix by Agol, Daniel Groves, and Jason Manning.
- [Ahl06] L. V. Ahlfors, *Lectures on quasiconformal mappings*, second ed., University Lecture Series, vol. 38, American Mathematical Society, Providence, RI, 2006, With supplemental chapters by C. J. Earle, I. Kra, M. Shishikura and J. H. Hubbard.
- [AK11] Y. Algom-Kfir, *Strongly contracting geodesics in outer space*, Geom. Topol. **15** (2011), no. 4, 2181–2233.
- [Alo94] J. M. Alonso, *Finiteness conditions on groups and quasi-isometries*, J. Pure Appl. Algebra **95** (1994), no. 2, 121–129.
- [Alp82] R. Alperin, *Locally compact groups acting on trees and Property T*, Monatsh. Math. **93** (1982), 261–265.
- [ALŚ15] S. Antoniuk, T. Łuczak, and J. Świątkowski, *Random triangular groups at density 1/3*, Compos. Math. **151** (2015), no. 1, 167–178.
- [AM15] Y. Antolín and A. Minasyan, *Tits alternatives for graph products*, J. Reine Angew. Math. **704** (2015), 55–83.
- [And05] J. W. Anderson, *Hyperbolic geometry*, Springer Verlag, 2005.
- [And07] J. E. Andersen, *Mapping Class Groups do not have Kazhdan’s property (T)*, preprint, ARXIV:0706.2184, 2007.
- [APCG13] J. Aliste-Prieto, D. Coronel, and J.-M. Gambaudo, *Linearly repetitive Delone sets are rectifiable*, Ann. Inst. H. Poincaré Anal. Non Linéaire **30** (2013), no. 2, 275–290.
- [Are46] R. Arens, *Topologies for homeomorphism groups*, Amer. J. Math. **68** (1946), 593–610.
- [Ass80] P. Assouad, *Plongements isométriques dans L^1 : aspect analytique*, Initiation Seminar on Analysis: G. Choquet-M. Rogalski-J. Saint-Raymond, 19th Year: 1979/1980, Publ. Math. Univ. Pierre et Marie Curie, vol. 41, Univ. Paris VI, Paris, 1980, pp. Exp. No. 14, 23.

- [Ass81] ———, *Comment on: “Isometric embeddings in L^1 : analytic aspect”*, Initiation Seminar on Analysis: G. Choquet-M. Rogalski-J. Saint-Raymond, 20th Year: 1980/1981, Publ. Math. Univ. Pierre et Marie Curie, vol. 46, Univ. Paris VI, Paris, 1981, pp. Comm. No. C4, 10.
- [Ass84] ———, *Sur les inégalités valides dans L^1* , Europ. J. Combinatorics **5** (1984), 99–112.
- [Att94] O. Attie, *Quasi-isometry classification of some manifolds of bounded geometry*, Math. Z. **216** (1994), 501–527.
- [AVS57] G. M. Adelson-Velskii and Yu. A. Sreider, *The Banach mean on groups*, Uspehi Mat. Nauk. (N.S.) **126** (1957), 131–136.
- [AW81] C. A. Akemann and M. E. Walter, *Unbounded negative definite functions*, Canad. J. Math. **33** (1981), 862–871.
- [AWP99] J. M. Alonso, X. Wang, and S. J. Pride, *Higher-dimensional isoperimetric (or Dehn) functions of groups*, J. Group Theory **2:1** (1999), 81–122.
- [Bal95] W. Ballmann, *Lectures on spaces of nonpositive curvature*, DMV Seminar, Band 25, Birkhäuser, 1995.
- [Bär97] C. Bär, *On nodal sets for Dirac and Laplace operators*, Comm. Math. Phys. **188** (1997), no. 3, 709–721.
- [Bas72] H. Bass, *The degree of polynomial growth of finitely generated nilpotent groups*, Proc. London Math. Soc. **25** (1972), 603–614.
- [Bas76] ———, *Euler characteristics and characters of discrete groups*, Invent. Math. **35** (1976), 155–196.
- [Bas01] S. A. Basarab, *The dual of the category of generalized trees*, Ann. Ştiinţ. Univ. Ovidius Constanţa Ser. Mat. **9** (2001), 1–20.
- [Bat99] M. Batty, *Groups with a sublinear isoperimetric inequality*, IMS Bulletin **42** (1999), 5–10.
- [Bau01] H. Bauer, *Measure and integration theory*, Studies in Math., vol. 26, de Gruyter, Berlin, 2001.
- [Bav91] C. Bavard, *Longueur stable des commutateurs*, Enseign. Math. (2) **37** (1991), no. 1-2, 109–150.
- [BB95] W. Ballmann and M. Brin, *Orbihedra of nonpositive curvature*, Inst. Hautes Études Sci. Publ. Math. (1995), no. 82, 169–209 (1996).
- [BB97] M. Bestvina and N. Brady, *Morse theory and finiteness properties of groups*, Invent. Math. **129** (1997), 445–470.
- [BBB⁺10] L. Bessières, G. Besson, M. Boileau, S. Maillot, and J. Porti, *Geometrisation of 3-manifolds*, EMS Tracts in Mathematics, vol. 13, European Mathematical Society, Zurich, 2010.
- [BBFS09] N. Brady, M. Bridson, M. Forester, and K. Shankar, *Snowflake groups, Perron-Frobenius eigenvalues and isoperimetric spectra*, Geom. Topol. **13** (2009), no. 1, 141–187.
- [BBI01] D. Burago, Y. Burago, and S. Ivanov, *A course in metric geometry*, Graduate Studies in Math., vol. 33, American Mathematical Society, Providence, RI, 2001.
- [BC01] R. Bishop and R. Crittenden, *Geometry of manifolds*, AMS Chelsea Publishing, Providence, RI, 2001, Reprint of the 1964 original.
- [BC12] J. Behrstock and R. Charney, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. **352** (2012), no. 2, 339–356.
- [BCG96] G. Besson, G. Courtois, and S. Gallot, *Minimal entropy and Mostow’s rigidity theorems*, Ergodic Theory Dynam. Systems **16** (1996), 623–649.
- [BCG98] ———, *A real Schwarz lemma and some applications*, Rend. Mat. Appl. (7) **18** (1998), no. 2, 381–410.
- [BCM12] J. Brock, R. Canary, and Y. Minsky, *The classification of Kleinian surface groups, II: The Ending Lamination Conjecture*, Ann. of Math. (2) **176** (2012), no. 1, 1–149.
- [BDCK66] J. Bretagnolle, D. Dacunha-Castelle, and J.-L. Krivine, *Lois stables et espaces L^p* , Ann. IHP, section B **2** (1966), no. 3, 231–259.
- [BdlHV08] B. Bekka, P. de la Harpe, and A. Valette, *Kazhdan’s property (T)*, New Mathematical Monographs, vol. 11, Cambridge University Press, Cambridge, 2008.
- [BDM09] J. Behrstock, C. Druţu, and L. Mosher, *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*, Math. Annalen **344** (2009), 543–595.

- [BE12] L. Bartholdi and A. Erschler, *Growth of permutational extensions*, Invent. Math. **189** (2012), no. 2, 431–455.
- [Bea83] A.F. Beardon, *The geometry of discrete groups*, Graduate Texts in Mathematics, vol. 91, Springer, 1983.
- [Bel03] I. Belegradek, *On cohopfian nilpotent groups*, Bull. London Math. Soc. **35** (2003), 805–811.
- [Ber68] G. Bergman, *On groups acting on locally finite graphs*, Annals of Math. **88** (1968), 335–340.
- [Bes88] M. Bestvina, *Degenerations of the hyperbolic space*, Duke Math. J. **56** (1988), 143–161.
- [Bes04] ———, *Problem list in geometric group theory*, 2004, <http://www.math.utah.edu/~bestvina/eprints/questions-updated.pdf>.
- [Bes14] M. Bestvina, *Geometric group theory and 3-manifolds hand in hand: the fulfillment of Thurston’s vision*, Bull. Amer. Math. Soc. (N.S.) **51** (2014), no. 1, 53–70.
- [BF92] M. Bestvina and M. Feighn, *A combination theorem for negatively curved groups*, J. Diff. Geom. **35** (1992), no. 1, 85–101.
- [BF95] ———, *Stable actions of groups on real trees*, Invent. Math. **121** (1995), no. 2, 287–321.
- [BF10] N. Brady and M. Forester, *Density of isoperimetric spectra*, Geom. Topol. **14** (2010), no. 1, 435–472.
- [BFGM07] U. Bader, A. Furman, T. Gelander, and N. Monod, *Property (T) and rigidity for actions on Banach spaces*, Acta Math. **198** (2007), 57–105.
- [BFH97] M. Bestvina, M. Feighn, and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups*, Geom. Funct. Anal. **7** (1997), 215–244.
- [BFH00] ———, *The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms.*, Ann. of Math. **151** (2000), no. 2, 517–623.
- [BFH04] ———, *Solvable subgroups of $\text{Out}(F_n)$ are virtually Abelian*, Geom. Dedicata **104** (2004), 71–96.
- [BFH05] ———, *The Tits alternative for $\text{Out}(F_n)$. II. A Kolchin type theorem*, Ann. of Math. (2) **161** (2005), no. 1, 1–59.
- [BG03] E. Breuillard and T. Gelander, *On dense free subgroups of Lie groups*, J. Algebra **261** (2003), no. 2, 448–467.
- [BG08a] J. Bourgain and A. Gamburd, *Uniform expansion bounds for Cayley graphs of $\text{SL}_2(\mathbb{F}_p)$* , Ann. of Math. (2) **167** (2008), no. 2, 625–642.
- [BG08b] E. Breuillard and T. Gelander, *Uniform independence in linear groups*, Invent. Math. **173** (2008), 225–263.
- [BG12] E. Breuillard and B. Green, *Approximate groups III: the unitary case*, Turkish J. Math. **36** (2012), no. 2, 199–215.
- [BGM12] U. Bader, T. Gelander, and N. Monod, *A fixed point theorem for L^1 spaces*, Invent. Math. **189** (2012), no. 1, 143–148.
- [BGS85] W. Ballmann, M. Gromov, and V. Schroeder, *Manifolds of non-positive curvature*, Progress in Math., vol. 61, Birkhauser, 1985.
- [BGT11] E. Breuillard, B. Green, and T. Tao, *Approximate subgroups of linear groups*, Geom. Funct. Anal. **21** (2011), no. 4, 774–819.
- [BGT12] ———, *The structure of approximate groups*, Publ. Math. Inst. Hautes Études Sci. **116** (2012), 115–221.
- [BH83] H.-J. Bandelt and J. Hedlikova, *Median algebras*, Discrete Math. **45** (1983), 1–30.
- [BH99] M. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Springer-Verlag, Berlin, 1999.
- [BH05] M. Bridson and J. Howie, *Conjugacy of finite subsets in hyperbolic groups*, Internat. J. Algebra Comput. **5** (2005), 725–756.
- [Bie76a] R. Bieri, *Homological dimension of discrete groups*, Mathematical Notes, Queen Mary College, 1976.
- [Bie76b] ———, *Normal subgroups in duality groups and in groups of cohomological dimension 2*, J. Pure Appl. Algebra **1** (1976), 35–51.
- [Bie79] ———, *Finitely presented soluble groups*, Séminaire d’Algèbre Paul Dubreil 31ème année (Paris, 1977–1978), Lecture Notes in Math., vol. 740, Springer, Berlin, 1979, pp. 1–8.

- [Big01] S. Bigelow, *Braid groups are linear*, J. Amer. Math. Soc. **14** (2001), no. 2, 471–486.
- [BJN09] J. Behrstock, T. Januszkiewicz, and W. Neumann, *Commensurability and QI classification of free products of finitely generated abelian groups*, Proc. Amer. Math. Soc. **137** (2009), no. 3, 811–813.
- [BJN10] ———, *Quasi-isometric classification of some high dimensional right-angled Artin groups*, Groups Geom. Dyn. **4** (2010), no. 4, 681–692.
- [BJS88] M. Bożejko, T. Januszkiewicz, and R. Spatzier, *Infinite Coxeter groups do not have Property (T)*, J. Op. Theory **19** (1988), 63–67.
- [BK47] G. Birkhoff and S. A. Kiss, *A ternary operation in distributive lattices*, Bull. Amer. Math. Soc. **53** (1947), 749–752.
- [BK81] P. Buser and H. Karcher, *Gromov’s almost flat manifolds*, Astérisque, vol. 81, Société Mathématique de France, Paris, 1981.
- [BK98] D. Burago and B. Kleiner, *Separated nets in Euclidean space and Jacobians of bi-Lipschitz maps*, Geom. Funct. Anal. **8** (1998), no. 2, 273–282.
- [BK02a] M. Bonk and B. Kleiner, *Quasisymmetric parametrizations of two-dimensional metric spheres*, Invent. Math. **150** (2002), no. 1, 127–183.
- [BK02b] D. Burago and B. Kleiner, *Rectifying separated nets*, Geom. Funct. Anal. **12** (2002), 80–92.
- [BK05] M. Bonk and B. Kleiner, *Conformal dimension and Gromov hyperbolic groups with 2-sphere boundary*, Geom. Topol. **9** (2005), 219–246 (electronic).
- [BK13] M. Bourdon and B. Kleiner, *Combinatorial modulus, the combinatorial Loewner property, and Coxeter groups*, Groups Geom. Dyn. **7** (2013), no. 1, 39–107.
- [BKMM12] J. Behrstock, B. Kleiner, Y. Minsky, and L. Mosher, *Geometry and rigidity of mapping class groups*, Geom. Topol. **16** (2012), no. 2, 781–888.
- [BKS08] M. Bestvina, B. Kleiner, and M. Sageev, *The asymptotic geometry of right-angled Artin groups. I*, Geom. Topol. **12** (2008), no. 3, 1653–1699.
- [BL00] Y. Benyamini and J. Lindenstrauss, *Geometric nonlinear functional analysis. Vol. 1*, American Mathematical Society Colloquium Publications, vol. 48, American Mathematical Society, Providence, RI, 2000.
- [BL12] A. Bartels and W. Lück, *The Borel conjecture for hyperbolic and CAT(0)-groups*, Ann. of Math. (2) **175** (2012), no. 2, 631–689.
- [BLW10] A. Bartels, W. Lück, and S. Weinberger, *On hyperbolic groups with spheres as boundary*, J. Differential Geom. **86** (2010), no. 1, 1–16.
- [BM91] M. Bestvina and G. Mess, *The boundary of negatively curved groups*, Jour. Amer. Math. Soc. **4** (1991), 469–481.
- [BM00] M. Burger and S. Mozes, *Lattices in product of trees*, Inst. Hautes Études Sci. Publ. Math. (2000), no. 92, 151–194.
- [BN08] J. Behrstock and W. Neumann, *Quasi-isometric classification of graph manifold groups*, Duke Math. J. **141** (2008), 217–240.
- [BN12] ———, *Quasi-isometric classification of non-geometric 3-manifold groups*, Crelle Journal **669** (2012), 101–120.
- [Bol79] B. Bollobás, *Graph theory, an introductory course*, Graduate Texts in Mathematics, vol. 63, Springer, 1979.
- [Bon11] M. Bonk, *Uniformization of Sierpinski carpets in the plane*, Invent. Math. **186** (2011), 559–665.
- [Boo57] W. Boone, *Certain simple, unsolvable problems of group theory. V, VI*, Nederl. Akad. Wetensch. Proc. Ser. A. 60, Indag. Math. **19** (1957), 22–27, 227–232.
- [Bor60] A. Borel, *Density properties for certain subgroups of semisimple groups without compact components*, Ann. of Math. **72** (1960), 179–188.
- [Bor63] ———, *Compact Clifford-Klein forms of symmetric spaces*, Topology **2** (1963), 111–122.
- [Bou63] N. Bourbaki, *Éléments de mathématique. Fascicule XXIX. Livre VI: Intégration. Chapitre 7: Mesure de Haar. Chapitre 8: Convolution et représentations*, Actualités Scientifiques et Industrielles, No. 1306, Hermann, Paris, 1963.
- [Bou65] ———, *Topologie générale*, Hermann, Paris, 1965.
- [Bou00] M. Bourdon, *Sur les immeubles fuchsien et leur type de quasi-isométrie*, Ergodic Th. & Dyn. Sys. **20** (2000), no. 2, 343–364.

- [Bou02] N. Bourbaki, *Lie groups and Lie algebras. Chapters 4–6*, Elements of Mathematics (Berlin), Springer-Verlag, Berlin, 2002, Translated from the 1968 French original by Andrew Pressley.
- [Bou16] M. Bourdon, *Cohomologie et actions isométriques propres sur les espaces L_p* , Geometry, topology, and dynamics in negative curvature, London Math. Soc. Lecture Note Ser., vol. 425, Cambridge Univ. Press, Cambridge, 2016, pp. 84–109.
- [Bow91] B. H. Bowditch, *Notes on Gromov’s hyperbolicity criterion for path-metric spaces*, “Group theory from a geometrical point of view”; ICTP, Trieste April 1990 (E. Ghys, A. Haefliger, and A. Verjovsky, eds.), 1991, pp. 64–167.
- [Bow94] ———, *Some results on the geometry of convex hulls in manifolds of pinched negative curvature*, Comment. Math. Helv. **69** (1994), no. 1, 49–81.
- [Bow95a] ———, *A short proof that a sub-quadratic isoperimetric inequality implies a linear one*, Mich. J. Math. **42** (1995), 103–107.
- [Bow95b] ———, *Geometrical finiteness with variable negative curvature*, Duke Math. J. **77** (1995), no. 1, 229–274.
- [Bow98a] ———, *Continuously many quasi-isometry classes of 2-generator groups*, Comment. Math. Helv. **73** (1998), no. 2, 232–236.
- [Bow98b] ———, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186.
- [Bow98c] ———, *A topological characterisation of hyperbolic groups*, Jour. Amer. Math. Soc. **11** (1998), no. 3, 643–667.
- [Bow06a] ———, *A course on geometric group theory*, MSJ Memoirs, vol. 16, Mathematical Society of Japan, Tokyo, 2006.
- [Bow06b] ———, *Intersection numbers and the hyperbolicity of the curve complex*, J. Reine Angew. Math. **598** (2006), 105–129.
- [Bow12] ———, *Relatively hyperbolic groups*, Internat. J. Algebra Comput. **22** (2012), no. 3, 1250016, 66.
- [Bow16] ———, *Some properties of median metric spaces*, Groups, Geom. and Dyn. **10** (2016), 279–317.
- [Bož80] M. Božejko, *Uniformly amenable discrete groups*, Math. Ann. **251** (1980), 1–6.
- [BP92] R. Benedetti and C. Petronio, *Lectures on hyperbolic geometry*, Springer, 1992.
- [BP00] M. Bourdon and H. Pajot, *Rigidity of quasi-isometries for some hyperbolic buildings*, Comment. Math. Helv. **75** (2000), no. 4, 701–736.
- [BP03] ———, *Cohomologie l_p et espaces de Besov*, J. Reine Angew. Math. **558** (2003), 85–108.
- [Bra99] N. Brady, *Branched coverings of cubical complexes and subgroups of hyperbolic groups*, J. London Math. Soc. (2) **60** (1999), no. 2, 461–480.
- [Bre92] L. Breiman, *Probability*, Classics in Applied Mathematics, vol. 7, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992, Corrected reprint of the 1968 original.
- [Bre09] W. Breslin, *Thick triangulations of hyperbolic n -manifolds*, Pacific J. Math. **241** (2009), no. 2, 215–225.
- [Bre14] E. Breuillard, *Expander graphs, property (τ) and approximate groups*, Geometric group theory, IAS/Park City Math. Ser., vol. 21, Amer. Math. Soc., Providence, RI, 2014, pp. 325–377.
- [Bri00] P. Brinkmann, *Hyperbolic automorphisms of free groups*, Geom. Funct. Anal. **10** (2000), no. 5, 1071–1089.
- [Bro81a] R. Brooks, *The fundamental group and the spectrum of the Laplacian*, Comment. Math. Helv. **56** (1981), no. 4, 581–598.
- [Bro81b] ———, *Some remarks on bounded cohomology*, Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978), Ann. of Math. Stud., vol. 97, Princeton Univ. Press, Princeton, N.J., 1981, pp. 53–63.
- [Bro82a] ———, *Amenability and the spectrum of the Laplacian*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), no. 1, 87–89.
- [Bro82b] K. Brown, *Cohomology of groups*, Graduate Texts in Math., vol. 87, Springer, 1982.
- [Bro87] ———, *Finiteness properties of groups*, J. Pure Appl. Algebra **44** (1987), 45–75.
- [BS69] J. Bell and A. Slomson, *Models and ultraproducts*, North-Holland, Amsterdam, 1969.

- [BS78] R. Bieri and R. Strebel, *Almost finitely presented soluble groups*, Comm. Math. Helv. **53** (1978), 258–278.
- [BS80] ———, *Valuations and finitely presented metabelian groups*, Proc. London Math. Soc. **41** (1980), no. 3, 439–464.
- [BS97a] W. Ballmann and J. Świątkowski, *On l^2 -cohomology and property (T) for automorphism groups of polyhedral cell complexes*, Geom. Funct. Anal. **7** (1997), no. 4, 615–645.
- [BS97b] I. Benjamini and O. Schramm, *Every graph with a positive Cheeger constant contains a tree with a positive Cheeger constant*, Geom. Funct. Anal. **7** (1997), 403–419.
- [BS99] W. Ballmann and J. Świątkowski, *On groups acting on nonpositively curved cubical complexes*, Enseignement Math. **45** (1999), 51–81.
- [BS00] M. Bonk and O. Schramm, *Embeddings of Gromov hyperbolic spaces*, Geom. Funct. Anal. **10** (2000), no. 2, 266–306.
- [BT24] S. Banach and A. Tarski, *Sur la décomposition des ensembles de points en parties respectivement congruentes*, Fundamenta Mathematicae **6** (1924), 244–277.
- [BT02] J. Burillo and J. Taback, *Equivalence of geometric and combinatorial Dehn functions*, New York J. Math. **8** (2002), 169–179 (electronic).
- [Bur99] J. Burillo, *Dimension and fundamental groups of asymptotic cones*, J. London Math. Soc. (2) **59** (1999), 557–572.
- [Bus65] H. Busemann, *Extremals on closed hyperbolic space forms*, Tensor (N.S.) **16** (1965), 313–318.
- [Bus82] P. Buser, *A note on the isoperimetric constant*, Ann. Sci. École Norm. Sup. (4) **15** (1982), no. 2, 213–230.
- [Bus10] ———, *Geometry and spectra of compact Riemann surfaces*, Modern Birkhäuser Classics, Birkhäuser Boston Inc., Boston, MA, 2010, Reprint of the 1992 edition.
- [BV98] M. Burger and A. Valette, *Idempotents in complex group rings: theorems of Zalesskii and Bass revisited*, Journal of Lie Theory (1998), 219–228.
- [BW11] M. Bridson and H. Wilton, *On the difficulty of presenting finitely presentable groups*, Groups Geom. Dyn. **5** (2011), no. 2, 301–325.
- [BZ88] Yu. Burago and V. Zalgaller, *Geometric inequalities*, Grundlehren der mathematischen Wissenschaften, vol. 285, Springer-Verlag, 1988.
- [Cai61] S. S. Cairns, *A simple triangulation method for smooth manifolds*, Bull. Amer. Math. Soc. **67** (1961), 389–390.
- [Cal08] D. Calegari, *Length and stable length*, Geom. Funct. Anal. **18** (2008), no. 1, 50–76.
- [Cal09] ———, *scl*, MSJ Memoirs, vol. 20, Mathematical Society of Japan, Tokyo, 2009.
- [Cap17] P.-E. Caprace, *A sixteen-relator presentation of an infinite hyperbolic kazhdan group*, Preprint, 2017.
- [Car14] M. Carlette, *The Haagerup property is not invariant under quasi-isometry*, Preprint, arXiv:1403.5446, 2014.
- [Cas86] J.W. Cassels, *Local fields*, Cambridge University Press, 1986.
- [Cas10] C. Cashen, *Quasi-isometries between tubular groups*, Groups, Geometry and Dynamics **4** (2010), 473–516.
- [CC92] J. Cannon and D. Cooper, *A characterization of cocompact hyperbolic and finite-volume hyperbolic groups in dimension three*, Trans. Amer. Math. Soc. **330** (1992), 419–431.
- [CCH81] I. M. Chiswell, D. J. Collins, and J. Huebschmann, *Aspherical group presentations*, Math. Z. **178** (1981), no. 1, 1–36.
- [CCJ⁺01] P.-A. Cherix, M. Cowling, P. Jolissaint, P. Julg, and A. Valette, *Groups with the Haagerup property. Gromov’s a - T -menability*, Progress in Mathematics, vol. 197, Birkhäuser, 2001.
- [CD17] I. Chatterji and C. Druţu, *Median geometry for spaces with measured walls and for groups*, preprint, arXiv:1708.00254, 2017.
- [CDP90] M. Coornaert, T. Delzant, and A. Papadopoulos, *Géométrie et théorie des groupes. Les groupes hyperboliques de Gromov*, Lec. Notes Math., vol. 1441, Springer-Verlag, 1990.
- [CGT82] J. Cheeger, M. Gromov, and M. Taylor, *Finite propagation speed, kernel estimates for functions of the Laplace operator, and the geometry of complete Riemannian manifolds*, J. Differential Geom. **17** (1982), no. 1, 15–53.

- [Cha01] I. Chavel, *Isoperimetric inequalities*, Cambridge Tracts in Mathematics, vol. 145, Cambridge University Press, Cambridge, 2001, Differential geometric and analytic perspectives.
- [Cha06] ———, *Riemannian geometry*, second ed., Cambridge Studies in Advanced Mathematics, vol. 108, Cambridge University Press, Cambridge, 2006, A modern introduction.
- [Che70] J. Cheeger, *A lower bound for the smallest eigenvalue of the Laplacian*, Problems in analysis (Papers dedicated to Salomon Bochner, 1969), Princeton Univ. Press, Princeton, N. J., 1970, pp. 195–199.
- [Che95] S. Chern, *Complex manifolds without potential theory (with an appendix on the geometry of characteristic classes)*, second ed., Universitext, Springer-Verlag, New York, 1995.
- [Che00] V. Chepoi, *Graphs of some CAT(0) complexes*, Adv. in Appl. Math. **24** (2000), 125–179.
- [Cho80] C. Chou, *Elementary amenable groups*, Illinois J. Math. **24** (1980), no. 3, 396–407.
- [Cho96] R. Chow, *Groups quasi-isometric to complex hyperbolic space*, Trans. Amer. Math. Soc. **348** (1996), no. 5, 1757–1769.
- [CJ94] A. Casson and D. Jungreis, *Convergence groups and Seifert-fibered 3-manifolds*, Invent. Math. **118** (1994), 441–456.
- [CL83] D. J. Collins and F. Levin, *Automorphisms and Hopficity of certain Baumslag-Solitar groups*, Arch. Math. (Basel) **40** (1983), no. 5, 385–400.
- [CM77] D. J. Collins and C. F. Miller, III, *The conjugacy problem and subgroups of finite index*, Proc. London Math. Soc. (3) **34** (1977), no. 3, 535–556.
- [CMV04] P. A. Cherix, F. Martin, and A. Valette, *Spaces with measured walls, the Haagerup property and property (T)*, Ergod. Th. & Dynam. Sys. **24** (2004), 1895–1908.
- [CN05] I. Chatterji and G. Niblo, *From wall spaces to CAT(0) cube complexes*, Internat. J. Algebra Comput. **15** (2005), no. 5-6, 875–885.
- [CN07] ———, *A characterization of hyperbolic spaces*, Groups, Geometry and Dynamics **1** (2007), no. 3, 281–299.
- [Com84] W. W. Comfort, *Topological groups*, Handbook of set-theoretic topology, North-Holland, Amsterdam, 1984, pp. 1143–1263.
- [Cor92] K. Corlette, *Archimedean superrigidity and hyperbolic geometry*, Ann. of Math. (2) **135** (1992), no. 1, 165–182.
- [CR13] L. Carbone and E. Rips, *Reconstructing group actions*, Internat. J. Algebra Comput. **23** (2013), no. 2, 255–323.
- [CSC93] Th. Coullhon and L. Saloff-Coste, *Isopérimétrie pour les groupes et variétés*, Rev. Math. Iberoamericana **9** (1993), 293–314.
- [CSGdlH98] T. Ceccherini-Silberstein, R. Grigorchuk, and P. de la Harpe, *Décompositions paradoxales des groupes de Burnside*, C. R. Acad. Sci. Paris Sér. I Math. **327** (1998), no. 2, 127–132.
- [Cut01] N. J. Cutland, *Loeb measures in practice: Recent advances*, Lecture Notes in Mathematics, vol. 1751, Springer-Verlag, Berlin, 2001.
- [CY75] S. Y. Cheng and S. T. Yau, *Differential equations on Riemannian manifolds and their geometric applications*, Comm. Pure Appl. Math. **28** (1975), no. 3, 333–354.
- [Dah03a] F. Dahmani, *Combination of convergence groups*, Geometry and Topology **7** (2003), 933–963.
- [Dah03b] ———, *Les groupes relativement hyperboliques et leurs bords*, Ph.D. thesis, University of Strasbourg, 2003.
- [Dav08] M. Davis, *The geometry and topology of Coxeter groups*, London Mathematical Society Monographs Series, vol. 32, Princeton University Press, Princeton, NJ, 2008.
- [Day50] M. M. Day, *Means for the bounded functions and ergodicity of the bounded representations of semi-groups*, Trans. Amer. Math. Soc. **69** (1950), 276–291.
- [Day57] ———, *Amenable semigroups*, Illinois J. Math. **1** (1957), 509–544.
- [dC92] M. P. do Carmo, *Riemannian geometry*, Birkhäuser Boston Inc., Boston, MA, 1992.
- [dC06] Y. de Cornulier, *Strongly bounded groups and infinite powers of finite groups*, Comm. Algebra **34** (2006), 2337–2345.
- [dC08] ———, *Dimension of asymptotic cones of Lie groups*, J. Topology **1** (2008), no. 2, 343–361.

- [dC11] ———, *Asymptotic cones of Lie groups and cone equivalences*, Illinois J. Math. **55** (2011), no. 1, 237–259.
- [dCTV08] Y. de Cornulier, R. Tessera, and A. Valette, *Isometric group actions on Banach spaces and representations vanishing at infinity*, Transform. Groups **13** (2008), no. 1, 125–147.
- [DD89] W. Dicks and M. J. Dunwoody, *Groups acting on graphs*, Cambridge University Press, Cambridge-New York, 1989.
- [DD16] K. Dekimpe and J. Deré, *Expanding maps and non-trivial self-covers on infranilmanifolds*, Topol. Methods Nonlinear Anal. **47** (2016), no. 1, 347–368.
- [dDW84] L. Van den Dries and A. Wilkie, *Gromov’s theorem on groups of polynomial growth and elementary logic*, J. Algebra **89** (1984), 349–374.
- [Del77] P. Delorme, *1-cohomologie des représentations unitaires des groupes de Lie semi-simples et résolubles. Produits tensoriels et représentations*, Bull. Soc. Math. France **105** (1977), 281–336.
- [Del96] T. Delzant, *Sous-groupes distingués et quotients des groupes hyperboliques*, Duke Math. J. **83** (1996), no. 3, 661–682.
- [DF04] D. S. Dummit and R. M. Foote, *Abstract algebra*, third ed., John Wiley & Sons, Inc., Hoboken, NJ, 2004.
- [DG] T. Delzant and M. Gromov, *Groupes de Burnside et géométrie hyperbolique*, preprint.
- [DG08] F. Dahmani and D. Groves, *The isomorphism problem for toral relatively hyperbolic groups*, Publ. Math. of IHES **107** (2008), 211–290.
- [DG11] F. Dahmani and V. Guirardel, *The isomorphism problem for all hyperbolic groups*, Geom. Funct. Anal. **21** (2011), 223–300.
- [DGLY02] A. N. Dranishnikov, G. Gong, V. Lafforgue, and G. Yu, *Uniform embeddings into Hilbert space and a question of Gromov*, Canad. Math. Bull. **45** (2002), no. 1, 60–70.
- [DGP11] F. Dahmani, V. Guirardel, and P. Przytycki, *Random groups do not split*, Mathematische Annalen **349** (2011), 657–673.
- [dlH00] P. de la Harpe, *Topics in geometric group theory*, Chicago Lectures in Mathematics, University of Chicago Press, 2000.
- [dlHGCS99] P. de la Harpe, R. I. Grigorchuk, and T. Ceccherini-Sil’berstein, *Amenability and paradoxical decompositions for pseudogroups and discrete metric spaces*, Proc. Steklov Inst. Math. **224** (1999), 57–97.
- [dlHV89] P. de la Harpe and A. Valette, *La propriété (T) de Kazhdan pour les groupes localement compacts*, Astérisque, vol. 175, Société Mathématique de France, 1989.
- [DM16] C. Druţu and J. Mackay, *Random groups, random graphs and eigenvalues of p -Laplacians*, preprint, arXiv:1607.04130, 2016.
- [Dod84] J. Dodziuk, *Difference equations, isoperimetric inequality and transience of certain random walks*, Trans. Amer. Math. Soc. **284** (1984), no. 2, 787–794.
- [Dol80] A. Dold, *Lectures on algebraic topology*, second ed., Grundlehren der Math. Wissenschaften, vol. 200, Springer, 1980.
- [DP98] A. Dyubina and I. Polterovich, *Structures at infinity of hyperbolic spaces*, Russian Math. Surveys **53** (1998), no. 5, 1093–1094.
- [DP01] ———, *Explicit constructions of universal \mathbb{R} -trees and asymptotic geometry of hyperbolic spaces*, Bull. London Math. Soc. **33** (2001), 727–734.
- [DPT15] T. Dymarz, I. Peng, and J. Taback, *Bilipschitz versus quasi-isometric equivalence for higher rank lamplighter groups*, New York J. Math. **21** (2015), 129–150.
- [DR09] M. Duchin and K. Rafi, *Divergence of geodesics in Teichmüller space and the Mapping Class group*, GAFA **19** (2009), 722–742.
- [DR13] W. Dison and T. Riley, *Hydra groups*, Comment. Math. Helv. **88** (2013), no. 3, 507–540.
- [Dre84] A. W. M. Dress, *Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups: a note on combinatorial properties of metric spaces*, Adv. in Math. **53** (1984), no. 3, 321–402.
- [Dru00] C. Druţu, *Quasi-isometric classification of non-uniform lattices in semisimple groups of higher rank*, Geom. Funct. Anal. **10** (2000), no. 2, 327–388.
- [Dru01] ———, *Cônes asymptotiques et invariants de quasi-isométrie pour des espaces métriques hyperboliques*, Ann. Inst. Fourier (Grenoble) **51** (2001), no. 1, 81–97.

- [Dru09] ———, *Relatively hyperbolic groups: geometry and quasi-isometric invariance*, Comment. Math. Helv. **84** (2009), no. 3, 503–546.
- [DS84] P. G. Doyle and J. L. Snell, *Random walks and electric networks*, Carus Mathematical Monographs, vol. 22, Mathematical Association of America, Washington, DC, 1984.
- [DS88] N. Dunford and J. T. Schwartz, *Linear operators. Part I. General theory.*, John Wiley & Sons, Inc., New York, 1988.
- [DS99] M. J. Dunwoody and M. Sageev, *JSJ splittings for finitely presented groups over slender groups*, Invent. Math. **135** (1999), no. 1, 25–44.
- [DS05a] C. Druţu and M. Sapir, *Relatively hyperbolic groups with rapid decay property*, Int. Math. Res. Notices **19** (2005), 1181–1194.
- [DS05b] ———, *Tree-graded spaces and asymptotic cones of groups*, Topology **44** (2005), 959–1058, with an appendix by D. Osin and M. Sapir.
- [DS07] ———, *Groups acting on tree-graded spaces and splittings of relatively hyperbolic groups*, Adv. Math. **217** (2007), 1313–1367.
- [DSS95] W. A. Deuber, M. Simonovits, and V. T. Sós, *A note on paradoxical metric spaces*, Studia Sci. Math. Hungar. **30** (1995), no. 1-2, 17–23.
- [Dun85] M. J. Dunwoody, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985), 449–457.
- [Dun93] ———, *An inaccessible group*, Geometric group theory, Vol. 1 (Sussex, 1991) (G. A. Niblo and M. A. Roller, eds.), London Math. Soc. Lecture Note Ser., vol. 181, Cambridge Univ. Press, Cambridge, 1993, pp. 75–78.
- [dV84] M. Van de Vel, *Dimension of binary convex structures*, Proc. London Math. Soc. **48** (1984), no. 3, 24–54.
- [Dym10] T. Dymarz, *Bilipschitz equivalence is not equivalent to quasi-isometric equivalence for finitely generated groups*, Duke Math. J. **154** (2010), no. 3, 509–526.
- [Dyu00] A. Dyubina, *Instability of the virtual solvability and the property of being virtually torsion-free for quasi-isometric groups*, Internat. Math. Res. Notices **21** (2000), 1097–1101.
- [Ebe96] P. Eberlein, *Geometry of nonpositively curved manifolds*, Chicago Lectures in Mathematics, University of Chicago Press, Chicago, IL, 1996.
- [ECH⁺92] D. B. A. Epstein, J. Cannon, D. F. Holt, S. Levy, M. S. Paterson, and W. P. Thurston, *Word processing and group theory*, Jones and Bartlett, 1992.
- [EF97a] D. B. A. Epstein and K. Fujiwara, *The second bounded cohomology of word-hyperbolic groups*, Topology **36** (1997), no. 6, 1275–1289.
- [EF97b] A. Eskin and B. Farb, *Quasi-flats and rigidity in higher rank symmetric spaces*, J. Amer. Math. Soc. **10** (1997), no. 3, 653–692.
- [Efr53] V. A. Efremovič, *The proximity geometry of riemannian manifolds*, Uspehi Matem. Nauk (N.S.) **8** (1953), 189.
- [EFW12] A. Eskin, D. Fisher, and K. Whyte, *Coarse differentiation of quasi-isometries I: Spaces not quasi-isometric to Cayley graphs*, Ann. of Math. (2) **176** (2012), no. 1, 221–260.
- [EFW13] ———, *Coarse differentiation of quasi-isometries II: Rigidity for Sol and lamplighter groups*, Ann. of Math. (2) **177** (2013), no. 3, 869–910.
- [EGS15] M. Ershov, G. Golan, and M. Sapir, *The Tarski numbers of groups*, Adv. Math. **284** (2015), 21–53.
- [Eng95] R. Engelking, *General Topology*, Sigma Series in Pure Mathematics, vol. 6, Heldermann Verlag, 1995.
- [EO73] P. Eberlein and B. O’Neill, *Visibility manifolds*, Pacific J. Math. **46** (1973), 45–109.
- [Ers03] A. Erschler, *On isoperimetric profiles of finitely generated groups*, Geometriae Dedicata **100** (2003), 157–171.
- [Ers04] ———, *Not residually finite groups of intermediate growth, commensurability and non-geometricity*, J. Algebra **272** (2004), no. 1, 154–172.
- [Ers06] ———, *Piecewise automatic groups*, Duke Math. J. **134** (2006), no. 3, 591–613.
- [Ers08] M. Ershov, *Golod-Shafarevich groups with property (T) and Kac-Moody groups*, Duke Math. J. **145** (2008), no. 2, 309–339.
- [Ers11] ———, *Kazhdan quotients of Golod-Shafarevich groups*, Proc. Lond. Math. Soc. (3) **102** (2011), no. 4, 599–636.

- [Esk98] A. Eskin, *Quasi-isometric rigidity of nonuniform lattices in higher rank symmetric spaces*, J. Amer. Math. Soc. **11** (1998), no. 2, 321–361.
- [Ess13] J. Essert, *A geometric construction of panel-regular lattices for buildings of types \tilde{A}_2 and \tilde{C}_2* , Algebr. Geom. Topol. **13** (2013), no. 3, 1531–1578.
- [ET64] V. Efremovich and E. Tihomirova, *Equimorphisms of hyperbolic spaces*, Izv. Akad. Nauk SSSR **28** (1964), 1139–1144.
- [Far97] B. Farb, *The quasi-isometry classification of lattices in semisimple Lie groups*, Math. Res. Letters **4** (1997), no. 5, 705–717.
- [Far98] ———, *Relatively hyperbolic groups*, Geom. Funct. Anal. **8** (1998), no. 5, 810–840.
- [Fed69] H. Federer, *Geometric measure theory*, Die Grundlehren der mathematischen Wissenschaften, Band 153, Springer-Verlag New York Inc., New York, 1969.
- [FF60] H. Federer and W. H. Fleming, *Normal and integral currents*, Ann. of Math. (2) **72** (1960), 458–520.
- [FH74] J. Faraut and K. Harzallah, *Distances hilbertiennes invariantes sur un espace homogène*, Ann. Institut Fourier **3** (1974), no. 24, 171–217.
- [FH94] W. Fulton and J. Harris, *Representation theory: A first course*, Springer, 1994.
- [Fio17] E. Fioravanti, *Roller boundaries for median spaces and algebras*, preprint, arXiv:1708.01005, 2017.
- [FJ03] R. Fleming and J. Jamison, *Isometries on Banach spaces: Function spaces*, Monographs and Surveys in Pure and Applied Mathematics, CRC Press, 2003.
- [FK16] K. Fujiwara and M. Kapovich, *On quasihomomorphisms with noncommutative targets*, Geom. Funct. Anal. **26** (2016), no. 2, 478–519.
- [FM98] B. Farb and L. Mosher, *A rigidity theorem for the solvable Baumslag–Solitar groups*, Invent. Math. **131** (1998), no. 2, 419–451, With an appendix by Daryl Cooper.
- [FM99] ———, *Quasi-isometric rigidity for the solvable Baumslag–Solitar groups, II*, Invent. Math. **137** (1999), no. 3, 613–649.
- [FM00] ———, *On the asymptotic geometry of abelian-by-cyclic groups*, Acta Math. **184** (2000), no. 2, 145–202.
- [FM11] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton University Press, 2011.
- [FN15] D. Fisher and T. Nguyen, *Quasi-isometric embeddings of non-uniform lattices*, preprint, arXiv:1512.07285, 2015.
- [Fol99] G. Folland, *Real analysis: Modern techniques and their applications*, Pure and Applied Mathematics, Wiley–Interscience, 1999.
- [For73] E. Formanek, *Idempotents in noetherian group rings*, Canad. J. Math. **25** (1973), 366–369.
- [Fox53] R. H. Fox, *Free differential calculus. I. Derivation in the free group ring*, Ann. of Math. (2) **57** (1953), 547–560.
- [FP06] K. Fujiwara and P. Papasoglu, *JSJ-decompositions of finitely presented groups and complexes of groups*, Geom. Funct. Anal. **16** (2006), no. 1, 70–125.
- [Fri60] A. A. Fridman, *On the relation between the word problem and the conjugacy problem in finitely defined groups*, Trudy Moskov. Mat. Obšč. **9** (1960), 329–356.
- [Fri91] J. Friedman, *On the second eigenvalue and random walks in random d -regular graphs*, Combinatorica **11** (1991), no. 4, 331–362.
- [FS88] M. Freedman and R. Skora, *Strange actions of groups on spheres. II*, “Holomorphic functions and moduli”, Vol. II (Berkeley, CA, 1986), Springer, New York, 1988, pp. 41–57.
- [FS96] B. Farb and R. Schwartz, *The large-scale geometry of Hilbert modular groups*, J. Differential Geom. **44** (1996), no. 3, 435–478.
- [FW91] M. Foreman and F. Wehrung, *The Hahn-Banach theorem implies the existence of a non-Lebesgue measurable set*, Fundam. Math. **138** (1991), 13–19.
- [FW14] D. Fisher and K. Whyte, *Quasi-isometric embeddings of symmetric spaces*, preprint, arXiv:1407.0445, 2014.
- [G60] P. Günther, *Einige Sätze über das Volumenelement eines Riemannschen Raumes*, Publ. Math. Debrecen **7** (1960), 78–93.
- [Gab92] D. Gabai, *Convergence groups are Fuchsian groups*, Annals of Math. **136** (1992), 447–510.

- [Gab05] D. Gaboriau, *Examples of groups that are measure equivalent to the free group*, Ergodic Theory Dynam. Systems **25** (2005), no. 6, 1809–1827.
- [Gab10] ———, *Orbit equivalence and measured group theory*, Proceedings of the International Congress of Mathematicians. Volume III (New Delhi), Hindustan Book Agency, 2010, pp. 1501–1527.
- [Gar73] H. Garland, *p-adic curvature and cohomology of discrete subgroups of p-adic groups*, Ann. of Math. **97** (1973), 375–423.
- [Gau73] C. F. Gauß, *Werke. Band VIII*, Georg Olms Verlag, Hildesheim, 1973, Reprint of the 1900 original.
- [GBS12] A. Myasnikov G. Baumslag and V. Shpilrain, *Open problems in combinatorial and geometric group theory*, <http://www.sci.cny.cuny.edu/~shpil/gworld/problems/oproblems.html>, 2012.
- [GdlH90] E. Ghys and P. de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov*, Progress in Mathematics, vol. 83, Birkhäuser, 1990.
- [GdlHJ89] F. M. Goodman, P. de la Harpe, and V. Jones, *Coxeter graphs and towers of algebras*, Mathematical Sciences Research Institute Publications, vol. 14, Springer-Verlag, New York, 1989.
- [Gel11] T. Gelander, *Volume versus rank of lattices*, J. Reine Angew. Math. **661** (2011), 237–248.
- [Gel14] ———, *Lectures on lattices and locally symmetric spaces*, Geometric group theory, IAS/Park City Math. Ser., vol. 21, Amer. Math. Soc., Providence, RI, 2014, pp. 249–282.
- [Geo08] R. Geoghegan, *Topological methods in group theory*, Graduate Texts in Mathematics, vol. 243, Springer, New York, 2008.
- [Ger87] S. Gersten, *Reducible diagrams and equations over groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 15–73.
- [Ger92] ———, *Bounded cocycles and combings of groups*, Internat. J. Algebra Comput. **2** (1992), no. 3, 307–326.
- [Ger93a] ———, *Isoperimetric and isodiametric functions of finite presentations*, Geometric group theory, Vol. 1 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 181, Cambridge Univ. Press, Cambridge, 1993, pp. 79–96.
- [Ger93b] ———, *Quasi-isometry invariance of cohomological dimension*, C. R. Acad. Sci. Paris Sér. I Math. **316** (1993), no. 5, 411–416.
- [Ger94] ———, *Quadratic divergence of geodesics in CAT(0)-spaces*, Geom. Funct. Anal. **4** (1994), no. 1, 37–51.
- [Ger09] V. Gerasimov, *Expansive convergence groups are relatively hyperbolic*, Geom. Funct. Anal. **19** (2009), no. 1, 137–169.
- [Ger12] ———, *Floyd maps for relatively hyperbolic groups*, Geom. Funct. Anal. **22** (2012), no. 5, 1361–1399.
- [GHL04] S. Gallot, D. Hulin, and J. Lafontaine, *Riemannian geometry*, third ed., Universitext, Springer-Verlag, Berlin, 2004.
- [Ghy04] E. Ghys, *Groupes aléatoires (d'après Misha Gromov, . . .)*, Astérisque (2004), no. 294, viii, 173–204.
- [Gol98] R. Goldblatt, *Lectures on the hyperreals. an introduction to nonstandard analysis*, Graduate Texts in Mathematics, vol. 188, Springer-Verlag, New York, 1998.
- [GP10] V. Guillemin and A. Pollack, *Differential topology*, AMS Chelsea Publishing, Providence, RI, 2010, Reprint of the 1974 original.
- [GP13] V. Gerasimov and L. Potyagailo, *Quasi-isometric maps and Floyd boundaries of relatively hyperbolic groups*, J. Eur. Math. Soc. (JEMS) **15** (2013), no. 6, 2115–2137.
- [GPS88] M. Gromov and I. Piatetski-Shapiro, *Non-arithmetic groups in Lobachevski spaces*, Publ. Math. IHES **66** (1988), 93–103.
- [Gre62] L. Greenberg, *Discrete subgroups of the Lorentz group*, Math. Scand. **10** (1962), 85–107.
- [Gri83] R. I. Grigorchuk, *On the Milnor problem of group growth*, Dokl. Akad. Nauk SSSR **271** (1983), no. 1, 30–33.
- [Gri84a] ———, *Degrees of growth of finitely generated groups and the theory of invariant means*, Izv. Akad. Nauk SSSR Ser. Mat. **48** (1984), no. 5, 939–985.

- [Gri84b] ———, *The growth rate of finitely generated groups and the theory of invariant means*, Inv. Akad. Nauk. **45** (1984), 939–986.
- [Gri87] ———, *Superamenability and the problem of occurrence of free semigroups*, Functional Analysis and its Applications **21** (1987), 64–66.
- [Gri98] ———, *An example of a finitely presented amenable group that does not belong to the class EG*, Mat. Sb. **189** (1998), no. 1-2, 75–95.
- [Gro28] H. Groetzsch, *Über einige Extremalprobleme der konformen Abbildung. I, II.*, Berichte Leipzig **80** (1928), 367–376, 497–502.
- [Gro81a] M. Gromov, *Groups of polynomial growth and expanding maps*, Publ. Math. IHES **53** (1981), 53–73.
- [Gro81b] ———, *Hyperbolic manifolds, groups and actions*, Annals of Mathematics Studies, vol. 97, Princeton University Press, Princeton, 1981.
- [Gro82] ———, *Volume and bounded cohomology*, Publ. Math. IHES (1982), no. 56, 5–99.
- [Gro83] ———, *Infinite groups as geometric objects*, Proceedings of the International Congress of Mathematicians, Warsaw, Amer. Math. Soc., 1983, pp. 385–392.
- [Gro86] ———, *Isoperimetric inequalities in Riemannian manifolds*, “Asymptotic Theory of Finite Dimensional Normed Spaces”, Lecture Notes Math., vol. 1200, Springer-Verlag, Berlin, 1986, pp. 114–129.
- [Gro87] ———, *Hyperbolic groups*, “Essays in group theory”, Math. Sci. Res. Ins. Publ., vol. 8, Springer, 1987, pp. 75–263.
- [Gro93] ———, *Asymptotic invariants of infinite groups*, Geometric Group Theory, Vol. 2 (Sussex, 1991) (G. Niblo and M. Roller, eds.), LMS Lecture Notes, vol. 182, Cambridge Univ. Press, 1993, pp. 1–295.
- [Gro00] ———, *Spaces and questions*, Geom. Funct. Anal. **Special Volume, Part I** (2000), 118–161.
- [Gro03] ———, *Random walk in random groups*, Geom. Funct. Anal. **13** (2003), no. 1, 73–146.
- [Gro07] ———, *Metric structures for Riemannian and non-Riemannian spaces*, Modern Birkhäuser Classics, Birkhäuser Boston Inc., Boston, MA, 2007, based on the 1981 French original, with appendices by M. Katz, P. Pansu and S. Semmes.
- [Gro09] C. Groft, *Generalized Dehn functions I*, arXiv:0901.2303, 2009.
- [Gru57] K. W. Gruenberg, *Residual properties of infinite soluble groups*, Proc. London Math. Soc. **7** (1957), 29–62.
- [GS90] S. Gersten and H. Short, *Small cancellation theory and automatic groups*, Invent. Math. **102** (1990), 305–334.
- [GS92] M. Gromov and R. Schoen, *Harmonic maps into singular spaces and p -adic super-rigidity for lattices in groups of rank one*, Publ. Math. of IHES **76** (1992), 165–246.
- [GT83] D. Gilbarg and N. S. Trudinger, *Elliptic partial differential equations of second order*, vol. 224, Springer Verlag, Grundlehren der Mathematischen Wissenschaften, 1983.
- [GT87] M. Gromov and W. Thurston, *Pinching constants for hyperbolic manifolds*, Invent. Math. **89** (1987), 1–12.
- [Gui70] Y. Guivarc’h, *Groupes de Lie à croissance polynomiale*, C. R. Acad. Sci. Paris, Ser. A—B **271** (1970), A237–A239.
- [Gui73] ———, *Croissance polynomiale et périodes des fonctions harmonique*, Bull. Soc. Math. France **101** (1973), 333–379.
- [Gui77] A. Guichardet, *Étude de la 1-cohomologie et de la topologie du dual pour les groupes de Lie à radical abélien*, Math. Ann. **228** (1977), no. 1, 215–232.
- [Hae92] A. Haefliger, *Extension of complexes of groups*, Ann. Inst. Fourier (Grenoble) **42** (1992), no. 1-2, 275–311.
- [Hai09] P. Haïssinsky, *Géométrie quasiconforme, analyse au bord des espaces métriques hyperboliques et rigidités [d’après Mostow, Pansu, Bourdon, Pajot, Bonk, Kleiner...]*, Astérisque (2009), no. 326, Exp. No. 993, ix, 321–362 (2010), Séminaire Bourbaki. Vol. 2007/2008.
- [Hai15] ———, *Hyperbolic groups with planar boundaries*, Invent. Math. **201** (2015), no. 1, 239–307.
- [Hal49] M. Hall, *Coset representations in free groups*, Trans. Amer. Math. Soc. **67** (1949), 421–432.

- [Hal64] J. Halpern, *The independence of the axiom of choice from the boolean prime ideal theorem*, Fund. Math. **55** (1964), 57–66.
- [Hal76] M. Hall, *The theory of groups*, Chelsea Publishing Co., New York, 1976, Reprinting of the 1968 edition.
- [Ham07] U. Hamenstädt, *Geometry of complex of curves and Teichmüller spaces*, Handbook of Teichmüller Theory, vol. 1, EMS, 2007, pp. 447–467.
- [Hat02] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.
- [Hau14] F. Hausdorff, *Bemerkung ber den inhalt von Punktmengen*, Mathematische Annalen **75** (1914), 428–434.
- [Hei01] J. Heinonen, *Lectures on analysis on metric spaces*, Universitext, Springer-Verlag, New York, 2001.
- [Hei05] ———, *Lectures on Lipschitz analysis*, Report. University of Jyväskylä Department of Mathematics and Statistics, vol. 100, University of Jyväskylä, Jyväskylä, 2005.
- [Hel01] S. Helgason, *Differential geometry, Lie groups, and symmetric spaces*, Graduate Studies in Mathematics, vol. 34, Amer. Math. Soc., 2001.
- [Hem78] J. Hempel, *3-manifolds*, Annals of Mathematics Studies, Princeton University Press, 1978.
- [Hem87] ———, *Residual finiteness for 3-manifolds*, “Combinatorial group theory and topology” (S.M. Gersten and J.R. Stallings, eds.), Annals of Mathematics Studies, vol. 111, Princeton University Press, Princeton, 1987, pp. 379–396.
- [Hig40] G. Higman, *The units of group-rings*, Proc. London Math. Soc. (2) **46** (1940), 231–248.
- [HIK99] J. Hagauer, W. Imrich, and S. Klavžar, *Recognizing median graphs in subquadratic time*, Theoretical Computer Science **215** (1999), 123–136.
- [Hir38] K. A. Hirsch, *On infinite soluble groups, II*, Proc. Lond. Math. Soc. **44** (1938), 336–344.
- [Hir76] M. Hirsch, *Differential topology*, Graduate Texts in Mathematics, vol. 33, Springer, 1976.
- [HJ99] K. Hrbacek and T. Jech, *Introduction to set theory*, third ed., Monographs and Textbooks in Pure and Applied Mathematics, vol. 220, Marcel Dekker Inc., New York, 1999.
- [HK95] J. Heinonen and P. Koskela, *Definitions of quasiconformality*, Invent. Math. **120** (1995), 61–79.
- [HK98] ———, *Quasiconformal maps in metric spaces with controlled geometry*, Acta Math. **181** (1998), 1–61.
- [HNN49] G. Higman, B. H. Neumann, and H. Neumann, *Embedding theorems for groups*, J. London Math. Soc. **24** (1949), 247–254.
- [HP74] E. Hille and R. Phillips, *Functional analysis and semi-groups*, American Mathematical Society, Providence, R. I., 1974, Third printing of the revised edition of 1957, American Mathematical Society Colloquium Publications, Vol. XXXI.
- [HP98] F. Haglund and F. Paulin, *Simplicité de groupes d’automorphismes d’espaces à courbure négative*, The Epstein Birthday Schrift (C. Rourke, I. Rivin, and C. Series, eds.), Geometry and Topology Monographs, vol. 1, International Press, 1998, pp. 181–248.
- [Hru12] E. Hrushovski, *Stable group theory and approximate subgroups*, Jour. Amer. Math. Soc. (2012), 189–243.
- [HS53] G. Hochschild and J.-P. Serre, *Cohomology of group extensions*, Trans. Amer. Math. Soc. **74** (1953), 110–134.
- [HŠ08] F. Haglund and J. Świątkowski, *Separating quasi-convex subgroups in 7-systolic groups*, Groups Geom. Dyn. **2** (2008), no. 2, 223–244.
- [Hua14] J. Huang, *Quasi-isometric classification of right-angled Artin groups I: the finite out case*, arXiv:1410.8512, 2014.
- [Hua16] ———, *Quasi-isometry classification of right-angled Artin groups II: several infinite out cases*, arXiv:1603.02372, 2016.
- [Hul66] A. Hulanicki, *Means and Følner condition on locally compact groups*, Studia Math. **27** (1966), 87–104.
- [Hum75] J. E. Humphreys, *Linear algebraic groups*, Springer-Verlag, New York, 1975, Graduate Texts in Mathematics, No. 21.

- [Hum97] ———, *Reflection groups and Coxeter groups*, Cambridge University Press, 1997, Cambridge Studies in Advanced Mathematics, Vol. 29.
- [Hun80] T. W. Hungerford, *Algebra*, Graduate Texts in Mathematics, vol. 73, Springer-Verlag, New York, 1980.
- [HW41] W. Hurewicz and H. Wallman, *Dimension theory*, Princeton University Press, 1941.
- [HW12] F. Haglund and D. Wise, *A combination theorem for special cube complexes*, Ann. of Math. (2) **176** (2012), no. 3, 1427–1482.
- [HY88] J. G. Hocking and G. S. Young, *Topology*, second ed., Dover Publications, Inc., New York, 1988.
- [IM01] T. Iwaniec and G. Martin, *Geometric function theory and non-linear analysis*, Oxford Mathematical Monographs, Oxford University Press, New York, 2001.
- [IS98] S. V. Ivanov and P. E. Schupp, *On the hyperbolicity of small cancellation groups and one-relator groups*, Trans. Amer. Math. Soc. **350** (1998), no. 5, 1851–1894.
- [Isa90] N. A. Isachenko, *Uniformly quasiconformal discontinuous groups that are not isomorphic to Möbius groups*, Dokl. Akad. Nauk SSSR **313** (1990), no. 5, 1040–1043.
- [Isb80] J. R. Isbell, *Median algebra*, Trans. Amer. Math. Soc. (1980), no. 260, 319–362.
- [Iva88] N. Ivanov, *Automorphisms of Teichmüller modular groups*, Topology and geometry—Rohlin Seminar, Lecture Notes in Math., vol. 1346, Springer, Berlin, 1988, pp. 199–270.
- [Iva92] ———, *Subgroups of Teichmüller modular groups*, Translations of Math. Monographs, vol. 115, AMS, 1992.
- [Iva94] S. V. Ivanov, *The free Burnside groups of sufficiently large exponents*, Internat. J. Algebra Comput. **4** (1994), no. 1-2, ii+308 pp.
- [Iva96] N. Ivanov, *Actions of Moebius transformations on homeomorphisms: stability and rigidity*, Geom. Funct. Anal. **6** (1996), 79–119.
- [JA95] M. Bridson J. Alonso, *Semihyperbolic groups*, Proc. London Math. Soc. (3) **70** (1995), no. 1, 56–114.
- [Jec03] Th. Jech, *Set theory: The third millennium edition, Revised and Expanded*, Springer, 2003.
- [Joh48] F. John, *Extremum problems with inequalities as subsidiary conditions*, Studies and Essays Presented to R. Courant on his 60th Birthday, January 8, 1948, Interscience Publishers, Inc., New York, N. Y., 1948, pp. 187–204.
- [JR88] W. Jaco and J. H. Rubinstein, *PL minimal surfaces in 3-manifolds*, J. Diff. Geom. **27** (1988), no. 3, 493–524.
- [JR89] ———, *PL equivariant surgery and invariant decomposition of 3-manifolds*, Advances in Math. **73** (1989), 149–191.
- [JŚ03] T. Januszkiewicz and J. Świątkowski, *Hyperbolic Coxeter groups of large dimension*, Comment. Math. Helv. **78** (2003), no. 3, 555–583.
- [JŚ06] ———, *Simplicial nonpositive curvature*, Publ. Math. Inst. Hautes Études Sci. (2006), no. 104, 1–85.
- [JvN35] P. Jordan and J. von Neumann, *On inner products in linear metric spaces*, Ann. of Math. **36** (1935), 719–723.
- [Kak41] S. Kakutani, *Concrete representation of abstract (L)-spaces and the mean ergodic theorem*, Ann. of Math. (2) **42** (1941), 523–537.
- [Kap69] I. Kaplansky, *Fields and rings*, The University of Chicago Press, 1969.
- [Kap70] ———, *“Problems in the theory of rings” revisited*, Amer. Math. Monthly **77** (1970), 445–454.
- [Kap92] M. Kapovich, *Intersection pairing on hyperbolic 4-manifolds*, preprint, 1992.
- [Kap96] I. Kapovich, *Detecting quasiconvexity: algorithmic aspects*, Geometric and computational perspectives on infinite groups (Minneapolis, MN and New Brunswick, NJ, 1994), DIMACS Ser. Discrete Math. Theoret. Comput. Sci., vol. 25, Amer. Math. Soc., Providence, RI, 1996, pp. 91–99.
- [Kap01] M. Kapovich, *Hyperbolic manifolds and discrete groups*, Birkhäuser Boston Inc., Boston, MA, 2001.
- [Kap05] ———, *Representations of polygons of finite groups*, Geom. Topol. **9** (2005), 1915–1951.

- [Kap07] ———, *Kleinian groups in higher dimensions*, “Geometry and Dynamics of Groups and Spaces. In memory of Alexander Reznikov”, vol. 265, Birkhauser, Progress in Mathematics, 2007, pp. 485–562.
- [Kap09] ———, *Homological dimension and critical exponent of Kleinian groups*, GAFA. **18** (2009), no. 6, 2017–2054.
- [Kap14] ———, *Energy of harmonic functions and Gromov’s proof of the Stallings’ theorem*, Georgian Math. Journal **21** (2014), 281–296.
- [Kaz75] D. Kazhdan, *On arithmetic varieties*, Lie groups and their representations, Halsted, 1975, pp. 151–216.
- [KB02] I. Kapovich and N. Benakli, *Boundaries of hyperbolic groups*, Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001), Contemp. Math., vol. 296, Amer. Math. Soc., Providence, RI, 2002, pp. 39–93.
- [Kei76] J. Keisler, *Foundations of infinitesimal calculus*, Prindel-Weber-Schmitt, Boston, 1976.
- [Kei10] ———, *The ultraproduct construction. Ultrafilters across mathematics*, Contemp. Math., vol. 530, Amer. Math. Soc., Providence, RI, 2010, pp. 163–179.
- [Kel72] G. Keller, *Amenable groups and varieties of groups*, Illinois J. Math. **16** (1972), 257–268.
- [Kes59] H. Kesten, *Full Banach mean values on countable groups*, Math. Scand. **7** (1959), 146–156.
- [KK00] M. Kapovich and B. Kleiner, *Hyperbolic groups with low-dimensional boundary*, Ann. Sci. École Norm. Sup. (4) **33** (2000), no. 5, 647–669.
- [KK05] ———, *Coarse Alexander duality and duality groups*, Journal of Diff. Geometry **69** (2005), 279–352.
- [KK09] ———, *Appendix to “Lacunary hyperbolic groups”, by A. Olshanskii, D. Osin, M. Sapir*, Geometry and Topology **13** (2009), 2132–2137.
- [KK13] M. Kotowski and M. Kotowski, *Random groups and property (T): Żuk’s theorem revisited*, J. Lond. Math. Soc. (2) **88** (2013), no. 2, 396–416.
- [KKL98] M. Kapovich, B. Kleiner, and B. Leeb, *Quasi-isometries and the de Rham decomposition*, Topology **37** (1998), 1193–1212.
- [KL95] M. Kapovich and B. Leeb, *On asymptotic cones and quasi-isometry classes of fundamental groups of nonpositively curved manifolds*, Geom. Funct. Anal. **5** (1995), no. 3, 582–603.
- [KL97] ———, *Quasi-isometries preserve the geometric decomposition of haken manifolds*, Invent. Math. **128** (1997), no. 2, 393–416.
- [KL98a] ———, *3-manifold groups and nonpositive curvature*, Geom. Funct. Anal. **8** (1998), no. 5, 841–852.
- [KL98b] B. Kleiner and B. Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, Math. Publ. of IHES **86** (1998), 115–197.
- [KL01] ———, *Groups quasi-isometric to symmetric spaces*, Comm. Anal. Geom. **9** (2001), no. 2, 239–260.
- [KL08] B. Kleiner and J. Lott, *Notes on Perelman’s papers*, Geometry & Topology **12** (2008), no. 5, 2587–2855.
- [KL09] B. Kleiner and B. Leeb, *Induced quasi-actions: A remark*, Proc. Amer. Math. Soc. **137** (2009), no. 5, 1561–1567.
- [KL12] I. Kapovich and A. Lukyanenko, *Quasi-isometric co-Hopficity of non-uniform lattices in rank-one semi-simple Lie groups*, Conform. Geom. Dyn. **16** (2012), 269–282.
- [Kle10] B. Kleiner, *A new proof of Gromov’s theorem on groups of polynomial growth*, J. Amer. Math. Soc. **23** (2010), no. 3, 815–829.
- [KLP14] M. Kapovich, B. Leeb, and J. Porti, *A Morse lemma for quasigeodesics in symmetric spaces and euclidean buildings*, Preprint, arXiv:1411.4176, 2014, 2014.
- [KM98a] O. Kharlampovich and A. Myasnikov, *Irreducible affine varieties over a free group. II. Systems in triangular quasi-quadratic form and description of residually free groups*, J. Algebra **200** (1998), no. 2, 517–570.
- [KM98b] ———, *Irreducible affine varieties over a free group. Irreducible affine varieties over a free group. I. Irreducibility of quadratic equations and Nullstellensatz*, J. Algebra (1998), no. 2, 472–516.

- [KM98c] ———, *Tarski's problem about the elementary theory of free groups has a positive solution*, Electron. Res. Announc. Amer. Math. Soc. **4** (1998), 101–108 (electronic).
- [KM05] ———, *Implicit function theorem over free groups*, J. Algebra **290** (2005), no. 1, 1–203.
- [KM12] J. Kahn and V. Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. of Math. (2) **175** (2012), no. 3, 1127–1190.
- [KPR84] N. J. Kalton, N. T. Peck, and J. W. Roberts, *An F -space sampler*, London Math. Soc. Lecture Note Series, vol. 89, Cambridge University Press, Cambridge, 1984, xii+240 pp.
- [Kro57] L. Kronecker, *Zwei Sätze Über Gleichungen mit ganzzahligen Coefficienten*, J. für Reine und Angewandte Mathematik **53** (1857), 173–175.
- [Krö10] B. Krön, *Cutting up graphs revisited—a short proof of Stallings' structure theorem*, Groups Complex. Cryptol. **2** (2010), no. 2, 213–221.
- [KS97] N. Korevaar and R. Schoen, *Global existence theorems for harmonic maps to non-locally compact spaces*, Comm. Anal. Geom. **5** (1997), no. 2, 333–387.
- [KS08] I. Kapovich and P. Schupp, *On group-theoretic models of randomness and genericity*, Groups Geom. Dyn. **2** (2008), no. 3, 383–404.
- [KSS06] I. Kapovich, P. Schupp, and V. Shpilrain, *Generic properties of Whitehead's algorithm and isomorphism rigidity of random one-relator groups*, Pacific J. Math. **223** (2006), no. 1, 113–140.
- [KSTT05] L. Kramer, S. Shelah, K. Tent, and S. Thomas, *Asymptotic cones of finitely presented groups*, Adv. Math. **193** (2005), no. 1, 142–173.
- [Kui50] N. H. Kuiper, *On compact conformally Euclidean spaces of dimension > 2* , Ann. of Math. (2) **52** (1950), 478–490.
- [Kun80] K. Kunen, *Set theory: An introduction to independence proofs*, Elsevier, 1980.
- [KW92] V. Kaimanovich and W. Woess, *The Dirichlet problem at infinity for random walks on graphs with a strong isoperimetric inequality*, Probab. Theory Related Fields **91** (1992), 445–466.
- [KY12] G. Kasparov and G. Yu, *The Novikov conjecture and geometry of Banach spaces*, Geom. Topol. **16** (2012), no. 3, 1859–1880.
- [Laf04] J.-F. Lafont, *Rigidity results for certain 3-dimensional singular spaces and their fundamental groups*, Geom. Dedicata **109** (2004), 197–219.
- [Lan64] S. Lang, *Algebraic numbers*, Addison-Wesley Publishing Co., Inc., Reading, Mass.-Palo Alto-London, 1964.
- [Lan00] U. Lang, *Higher-dimensional linear isoperimetric inequalities in hyperbolic groups*, Internat. Math. Res. Notices (2000), no. 13, 709–717.
- [Lan02] S. Lang, *Algebra*, Addison-Wesley Publishing Company, 2002.
- [Lav38] M. Lavrentieff, *Sur un critère différentiel des transformations homéomorphes des domaines à trois dimensions.*, C. R. (Dokl.) Acad. Sci. URSS, n. Ser. **20** (1938), 241–242.
- [Lee12] S. R. Lee, *Geometry of Houghton's groups*, Preprint, arXiv:1212.0257, 2012.
- [Leh87] O. Lehto, *Univalent functions and Teichmüller spaces*, Springer, 1987.
- [Li12] P. Li, *Geometric analysis*, Cambridge Studies in Advanced Mathematics, vol. 134, Cambridge University Press, Cambridge, 2012.
- [LM16] Y. Lodha and J. T. Moore, *A nonamenable finitely presented group of piecewise projective homeomorphisms*, Groups Geom. Dyn. **10** (2016), no. 1, 177–200.
- [ŁRN51] J. Łoś and C. Ryll-Nardzewski, *On the application of Tychonoff's theorem in mathematical proofs*, Fund. Math. **38** (1951), 233–237.
- [LS77] R. C. Lyndon and P. E. Schupp, *Combinatorial group theory*, Springer-Verlag, Berlin, 1977, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.
- [LT79] J. Lindenstrauss and L. Tzafriri, *Classical Banach spaces. II: Function spaces*, Ergebnisse der Mathematik und ihrer Grenzgebiete, vol. 97, Springer-Verlag, Berlin, 1979.
- [LT92] P. Li and L. Tam, *Harmonic functions and the structure of complete manifolds*, J. Differential Geom. **35** (1992), 359–383.
- [Lux62] W. A. J. Luxemburg, *Two applications of the method of construction by ultrapowers to analysis*, Bull. Amer. Math. Soc. **68** (1962), 416–419.
- [Lux67] ———, *Beweis des satzes von Hahn-Banach*, Arch. Math (Basel) **18** (1967), 271–272.

- [Lux69] ———, *Reduced powers of the real number system and equivalents of the Hahn-Banach extension theorem*, Applications of Model Theory to Algebra, Analysis, and Probability (Internat. Sympos., Pasadena, Calif., 1967), Holt, Rinehart and Winston, New York, 1969, pp. 123–137.
- [Lys96] I. G. Lysenok, *Infinite Burnside groups of even exponent*, Izvestiya Akad. Nauk. SSSR Ser., Mat **60** (1996), no. 3, 453–654.
- [Mac13] N. Macura, *CAT(0) spaces with polynomial divergence of geodesics*, Geom. Dedicata **163** (2013), 361–378.
- [Mag39] W. Magnus, *On a theorem of Marshall Hall*, Ann. of Math. (2) **40** (1939), 764–768.
- [Mal40] A. I. Mal'cev, *On isomorphic matrix representations of infinite groups*, Mat. Sbornik **8** (1940), 405–422.
- [Mal49a] ———, *Generalized nilpotent algebras and their associated groups*, Mat. Sbornik **25** (1949), 347–366.
- [Mal49b] ———, *On a class of homogeneous spaces*, Izvestiya Akad. Nauk. SSSR Ser., Mat **13** (1949), 9–32, Amer. Math. Soc. Translation No. 39 (1951).
- [Mal51] ———, *On certain classes of infinite soluble groups*, Mat. Sbornik **28** (1951), 567–588.
- [Mar86] G. Martin, *Discrete quasiconformal groups which are not quasiconformal conjugates of Mobius groups*, Ann. Ac. Sci. Fenn. **11** (1986), 179–202.
- [Mar91] G. A. Margulis, *Discrete subgroups of semisimple Lie groups*, Springer-Verlag, Berlin, 1991.
- [Mar13] V. Markovic, *Criterion for Cannon's conjecture*, Geom. Funct. Anal. **23** (2013), no. 3, 1035–1061. MR 3061779
- [Mas91] W. S. Massey, *A basic course in algebraic topology*, Graduate Texts in Mathematics, vol. 127, Springer, 1991.
- [McM98] C. T. McMullen, *Lipschitz maps and nets in Euclidean space*, Geom. Funct. Anal. **8** (1998), no. 2, 304–314.
- [Mer10] S. Merenkov, *A Sierpinsky carpet with the co-hopfian property*, Invent. Math. **180** (2010), 361–388.
- [Mes72] S. Meskin, *Nonresidually finite one-relator groups*, Trans. Amer. Math. Soc. **164** (1972), 105–114.
- [Mil68a] J. Milnor, *Growth of finitely generated solvable groups*, J. Diff. Geom. **2** (1968), 447–449.
- [Mil68b] ———, *A note on curvature and fundamental group*, J. Diff. Geom. **2** (1968), 1–7.
- [Mil68c] ———, *Problem 5603*, Amer. Math. Monthly **75** (1968), 685–686.
- [Mil76] J. J. Millson, *On the first betti number of a constant negatively curved manifold*, Ann. of Math. **104** (1976), no. 2, 235–247.
- [Mil79] ———, *Real vector bundles with discrete structure group*, Topology **18** (1979), no. 1, 83–89.
- [Mil12] J. S. Milne, *Group theory*, <http://www.jmilne.org/math/CourseNotes/GT.pdf>, 2012.
- [Min01] I. Mineyev, *Straightening and bounded cohomology of hyperbolic groups*, Geom. Funct. Anal. **11** (2001), no. 4, 807–839.
- [Min10] Y. Minsky, *The classification of Kleinian surface groups. I. Models and bounds*, Ann. of Math. (2) **171** (2010), no. 1, 1–107.
- [Mj14a] M. Mj, *Cannon-Thurston maps for surface groups*, Ann. of Math. (2) **179** (2014), no. 1, 1–80.
- [Mj14b] ———, *Ending laminations and Cannon-Thurston maps*, Geom. Funct. Anal. **24** (2014), no. 1, 297–321. With an appendix by Shubhabrata Das and Mj.
- [MK14] V. D. Mazurov and E. I. Khukhro, *Unsolved problems in group theory. The Kourovka notebook*, arXiv:1401.0300, 2014.
- [MN82] J. Meakin and K. S. S. Nambooripad, *Coextensions of regular semigroups by rectangular bands. I*, Trans. Amer. Math. Soc. **269** (1982), no. 1, 197–224.
- [MNO92] J. C. Mayer, J. Nikiel, and L. G. Oversteegen, *Universal spaces for \mathbb{R} -trees*, Trans. Amer. Math. Soc. **334** (1992), no. 1, 411–432.
- [Mon06] N. Monod, *An invitation to bounded cohomology*, International Congress of Mathematicians. Vol. II, Eur. Math. Soc., Zürich, 2006, pp. 1183–1211.
- [Mor24] M. Morse, *A fundamental class of geodesics on any closed surface of genus greater than one*, Transactions of AMS **26** (1924), 25–60.

- [Mor38] C. B. Morrey, *On the solution of quasilinear elliptic partial differential equations*, Transactions of AMS **43** (1938), 126–166.
- [Mos65] G. D. Mostow, *Quasiconformal mappings in n -space and the rigidity of hyperbolic space-forms*, Publ. Math. IHES **34** (1965), 53–104.
- [Mos73] ———, *Strong Rigidity of Locally Symmetric Spaces*, Annals of Math. Studies, no. 78, Princeton Univ. Press, 1973.
- [Mou88] G. Moussong, *Hyperbolic Coxeter groups*, Ph.D. thesis, The Ohio State University, 1988.
- [MP15] J. M. Mackay and P. Przytycki, *Balanced walls for random groups*, Michigan Math. J. **64** (2015), 397–419.
- [MR81] J. J. Millson and M. S. Raghunathan, *Geometric construction of cohomology for arithmetic groups. I*, Proc. Indian Acad. Sci. Math. Sci. **90** (1981), no. 2, 103–123.
- [MR03] C. Maclachlan and A. Reid, *The arithmetic of hyperbolic 3-manifolds*, Graduate Texts in Mathematics, vol. 219, Springer-Verlag, New York, 2003.
- [MR08] M. Mj and L. Reeves, *A combination theorem for strong relative hyperbolicity*, Geom. Topol. **12** (2008), no. 3, 1777–1798.
- [MS86] V. D. Milman and G. Schechtman, *Asymptotic theory of finite-dimensional normed spaces*, Lecture Notes in Mathematics, vol. 1200, Springer-Verlag, Berlin, 1986, With an appendix by M. Gromov.
- [MSW03] L. Mosher, M. Sageev, and K. Whyte, *Quasi-actions on trees. I. Bounded valence*, Ann. of Math. (2) **158** (2003), no. 1, 115–164.
- [MSW11] ———, *Quasi-actions on trees II: Finite depth Bass-Serre trees*, Mem. Amer. Math. Soc. **214** (2011), no. 1008, vi+105.
- [Mul80] H. M. Mulder, *The interval function of a graph*, Mathematical Centre Tracts, vol. 132, Mathematisch Centrum, Amsterdam, 1980.
- [Mun75] J. R. Munkres, *Topology: a first course*, Prentice-Hall Inc., 1975.
- [Mur05] A. Muranov, *Diagrams with selection and method for constructing boundedly generated and boundedly simple groups*, Comm. Algebra **33** (2005), no. 4, 1217–1258.
- [MY81] W. Meeks and S.T. Yau, *The equivariant Dehn’s lemma and equivariant loop theorem*, Comment. Math. Helv. **56** (1981), 225–239.
- [MY02] I. Mineyev and G. Yu, *The Baum-Connes conjecture for hyperbolic groups*, Invent. Math. **149** (2002), no. 1, 97–122.
- [MZ74] D. Montgomery and L. Zippin, *Topological transformation groups*, Robert E. Krieger Publishing Co., Huntington, N.Y., 1974, Reprint of the 1955 original.
- [Nac65] L. Nachbin, *The Haar Integral*, van Nostrand, Princeton, New Jersey, 1965.
- [Nag83] J. Nagata, *Modern dimension theory*, revised ed., Sigma Series in Pure Mathematics, vol. 2, Heldermann Verlag, Berlin, 1983.
- [Nag85] ———, *Modern general topology*, second ed., North Holland Mathematical Library, vol. 33, North Holland, 1985.
- [Nér64] A. Néron, *Modèles minimaux des variétés abéliennes sur les corps locaux et globaux*, Math. Publ. of IHES **21** (1964), 5–128.
- [New68] B. B. Newman, *Some results on one-relator groups*, Bull. Amer. Math. Soc. **74** (1968), 568–571.
- [Nib04] G. Niblo, *A geometric proof of Stallings’ theorem on groups with more than one end*, Geometriae Dedicata **105** (2004), no. 1, 61–76.
- [Nic04] B. Nica, *Cubulating spaces with walls*, Alg. & Geom. Topology **4** (2004), 297–309.
- [Nic08] ———, *Group actions on median spaces*, preprint, arXiv:0809.4099, 2008.
- [Nic13] ———, *Proper isometric actions of hyperbolic groups on L^p -spaces*, Compos. Math. **149** (2013), no. 5, 773–792.
- [Nie78] J. Nieminen, *The ideal structure of simple ternary algebras*, Colloq. Math. **40** (1978), 23–29.
- [Nos91] G. Noskov, *Bounded cohomology of discrete groups with coefficients*, Leningrad Math. Journal **2** (1991), no. 5, 1067–1084.
- [Nov58] P. S. Novikov, *On the algorithmic insolvability of the word problem in group theory*, American Mathematical Society Translations, Ser 2, Vol. 9, American Mathematical Society, Providence, R. I., 1958, pp. 1–122.
- [NR97a] W. Neumann and L. Reeves, *Central extensions of word hyperbolic groups*, Ann. of Math. **145** (1997), 183–192.

- [NR97b] G. A. Niblo and L. D. Reeves, *Groups acting on $CAT(0)$ cube complexes*, *Geometry and Topology* **1** (1997), 7 pp.
- [NŠ16] B. Nica and J. Špakula, *Strong hyperbolicity*, *Groups Geom. Dyn.* **10** (2016), no. 3, 951–964.
- [NY11] P. Nowak and G. Yu, *Large scale geometry*, European Mathematical Society, 2011.
- [OŠ80] A. Yu. Ol’shanskiĭ, *On the question of the existence of an invariant mean on a group*, *Uspekhi Mat. Nauk* **35** (1980), no. 4, 199–200.
- [OŠ91a] ———, *Geometry of defining relations in groups*, *Mathematics and its Applications (Soviet Series)*, vol. 70, Kluwer Academic Publishers Group, Dordrecht, 1991.
- [OŠ91b] ———, *Hyperbolicity of groups with subquadratic isoperimetric inequalities*, *Intl. J. Alg. Comp.* **1** (1991), 282–290.
- [OŠ91c] ———, *Periodic quotient groups of hyperbolic groups*, *Mat. Sb.* **182** (1991), no. 4, 543–567.
- [OŠ92] ———, *Almost every group is hyperbolic*, *Internat. J. Algebra Comput.* **2** (1992), no. 1, 1–17.
- [OŠ95] ———, *SQ-universality of hyperbolic groups*, *Mat. Sb.* **186** (1995), no. 8, 119–132.
- [Oll04] Y. Ollivier, *Sharp phase transition theorems for hyperbolicity of random groups*, *Geom. Funct. Anal.* **14** (2004), no. 3, 595–679.
- [Oll05] ———, *A January 2005 invitation to random groups*, *Ensaos Matemáticos [Mathematical Surveys]*, vol. 10, Sociedade Brasileira de Matemática, Rio de Janeiro, 2005.
- [Oll07] ———, *Some small cancellation properties of random groups*, *Internat. J. Algebra Comput.* **17** (2007), no. 1, 37–51.
- [OOS09] A. Yu. Ol’shanskiĭ, D. V. Osin, and M. V. Sapir, *Lacunary hyperbolic groups*, *Geom. Topol.* **13** (2009), no. 4, 2051–2140, With an appendix by M. Kapovich and B. Kleiner.
- [Ore51] O. Ore, *Some remarks on commutators*, *Proc. Amer. Math. Soc.* **2** (1951), 307–314.
- [OS01] A. Yu. Ol’shanskiĭ and M. V. Sapir, *Length and area functions on groups and quasi-isometric Higman embeddings*, *International Journal of Algebra and Computation* **11** (2001), 137–170.
- [OS02] ———, *Non-amenable finitely presented torsion-by-cyclic groups*, *Publ. Math. IHES* **96** (2002), 43–169.
- [Osa13] D. Osajda, *A construction of hyperbolic Coxeter groups*, *Comment. Math. Helv.* **88** (2013), no. 2, 353–367.
- [Osa14] ———, *Small cancellation labellings of some infinite graphs and applications*, preprint [ARXIV:1406.5015](https://arxiv.org/abs/1406.5015), 2014.
- [Osi01] D. Osin, *Subgroup distortions in nilpotent groups*, *Comm. Algebra* **29** (2001), no. 12, 5439–5463.
- [Osi06] ———, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, *Mem. Amer. Math. Soc.* **179** (2006), no. 843, vi+100pp.
- [Osi10] ———, *Small cancellations over relatively hyperbolic groups and embedding theorems*, *Ann. of Math.* **172** (2010), 1–39.
- [OV90] A. Onishchik and E. Vinberg, *Lie groups and algebraic groups*, Springer, 1990.
- [OW80] D. Ornstein and B. Weiss, *Ergodic theory of amenable group actions. I. The Rohlin lemma*, *Bull. Amer. Math. Soc. (N.S.)* **2** (1980), no. 1, 161–164.
- [OW11] Y. Ollivier and D. Wise, *Cubulating random groups at density less than 1/6*, *Trans. Amer. Math. Soc.* **363** (2011), no. 9, 4701–4733.
- [Oza15] N. Ozawa, *A functional analysis proof of Gromov’s polynomial growth theorem*, Preprint, 2015.
- [Pan83] P. Pansu, *Croissance des boules et des géodésiques fermées dans les nilvariétés*, *Ergodic Th. & Dyn. Sys.* **3** (1983), 415–455.
- [Pan89] ———, *Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un*, *Ann. of Math. (2)* **129** (1989), no. 1, 1–60.
- [Pan95] ———, *Cohomologie L^p : invariance sous quasi-isométrie*, preprint, <http://www.math.u-psud.fr/~pansu/liste-prepub.html>, 1995.
- [Pan96] ———, *Formule de Matsushima, de Garland, et propriété (T) pour des groupes agissant sur des espaces symétriques ou des immeubles*, Preprint, Orsay, 1996.
- [Pan07] ———, *Cohomologie L^p en degré 1 des espaces homogènes*, *Potential Anal.* **27** (2007), no. 2, 151–165.

- [Pap95a] P. Papasoglu, *Homogeneous trees are bi-Lipschitz equivalent*, *Geom. Dedicata* **54** (1995), no. 3, 301–306.
- [Pap95b] ———, *On the subquadratic isoperimetric inequality*, *Geometric group theory*, vol. 25, de Gruyter, Berlin-New-York, 1995, R. Charney, M. Davis, M. Shapiro (eds), pp. 193–200.
- [Pap95c] ———, *Strongly geodesically automatic groups are hyperbolic*, *Invent. Math.* **121** (1995), no. 2, 323–334.
- [Pap96] ———, *An algorithm detecting hyperbolicity*, *Geometric and computational perspectives on infinite groups* (Minneapolis, MN and New Brunswick, NJ, 1994), vol. 25, Amer. Math. Soc., Providence, RI, 1996, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., pp. 193–200.
- [Pap98] ———, *Quasi-flats in semihyperbolic groups*, *Proc. Amer. Math. Soc.* **126** (1998), 1267–1273.
- [Pap00] ———, *Isodiametric and isoperimetric inequalities for complexes and groups*, *J. London Math. Soc. (2)* **62** (2000), no. 1, 97–106.
- [Pap03] ———, *Notes on hyperbolic and automatic groups*, *Lecture Notes*, based on the notes of M. Batty, 2003.
- [Pap05] ———, *Quasi-isometry invariance of group splittings*, *Ann. of Math.* **161** (2005), no. 2, 759–830.
- [Par08] J. Parker, *Hyperbolic spaces*, vol. 2, Jyväskylä Lectures in Mathematics, 2008.
- [Pau88] F. Paulin, *Topologie de Gromov équivariante, structures hyperboliques et arbres réels*, *Invent. Math.* **94** (1988), 53–80.
- [Pau91a] ———, *Outer automorphisms of hyperbolic groups and small actions on \mathbf{R} -trees*, *Arboreal group theory* (Berkeley, CA, 1988), *Math. Sci. Res. Inst. Publ.*, vol. 19, Springer, New York, 1991, pp. 331–343.
- [Pau91b] ———, *Outer automorphisms of hyperbolic groups and small actions on \mathbb{R} -trees*, *Arboreal group theory* (R. Alperin, ed.), *Publ. M.S.R.I.*, vol. 19, Springer, 1991, pp. 331–343.
- [Pau96] ———, *Un groupe hyperbolique est déterminé par son bord*, *J. London Math. Soc. (2)* **54** (1996), no. 1, 50–74.
- [Paw91] J. Pawlikowski, *The Hahn-Banach theorem implies the Banach-Tarski paradox*, *Fundamenta Mathematicae* **138** (1991), no. 1, 21–22.
- [Ped95] E. K. Pedersen, *Bounded and continuous control*, *Proceedings of the conference “Novikov conjectures, index theorems and rigidity” volume II*, Oberwolfach 1993, *LMS Lecture Notes Series*, vol. 227, Cambridge University Press, 1995, pp. 277–284.
- [Pen11a] I. Peng, *Coarse differentiation and quasi-isometries of a class of solvable Lie groups I*, *Geom. Topol.* **15** (2011), no. 4, 1883–1925.
- [Pen11b] ———, *Coarse differentiation and quasi-isometries of a class of solvable Lie groups II*, *Geom. Topol.* **15** (2011), no. 4, 1927–1981.
- [Pet16] P. Petersen, *Riemannian geometry*, third ed., *Graduate Texts in Mathematics*, vol. 171, Springer, Cham, 2016.
- [Pin72] D. Pincus, *Independence of the prime ideal theorem from the Hahn Banach theorem*, *Bull. Amer. Math. Soc.* **78** (1972), 766–770.
- [Pin74] ———, *The strength of the Hahn-Banach theorem*, *Victoria Symposium on Non-standard Analysis* (Univ. Victoria, Victoria, B.C., 1972), Springer, Berlin, 1974, pp. 203–248. *Lecture Notes in Math.*, Vol. 369.
- [Pit98] C. Pittet, *On the isoperimetry of graphs with many ends*, *Colloq. Math.* **78** (1998), 307–318.
- [Pla68] V. P. Platonov, *A certain problem for finitely generated groups*, *Dokl. Akad. Nauk BSSR* **12** (1968), 492–494.
- [Poi95] F. Point, *Groups of polynomial growth and their associated metric spaces*, *J. Algebra* **175** (1995), no. 1, 105–121.
- [Pos37] B. Pospíšil, *Remark on bicomact spaces*, *Ann. of Math. (2)* **38** (1937), no. 4, 845–846.
- [Pra73] G. Prasad, *Strong rigidity of Q -rank one lattices*, *Invent. Math.* **21** (1973), 255–286.
- [PW02] P. Papasoglu and K. Whyte, *Quasi-isometries between groups with infinitely many ends*, *Comment. Math. Helv.* **77** (2002), no. 1, 133–144.

- [Rab58] M. O. Rabin, *Recursive unsolvability of group theoretic problems*, Ann. of Math. (2) **67** (1958), 172–194.
- [Rag72] M. Raghunathan, *Discrete subgroups of Lie groups*, Springer, 1972.
- [Rag84] ———, *Torsion in cocompact lattices in coverings of $Spin(2, n)$* , Math. Ann. **266** (1984), no. 4, 403–419.
- [Rat06] J. G. Ratcliffe, *Foundations of hyperbolic manifolds*, Springer, Second Edition, 2006.
- [Rei65] H. Reiter, *On some properties of locally compact groups*, Nederl. Akad. Wetensch. Proc. Ser. A 68=Indag. Math. **27** (1965), 697–701.
- [Res89] Yu. G. Reshetnyak, *Space mappings with bounded distortion*, Translations of Mathematical Monographs, vol. 73, American Mathematical Society, Providence, RI, 1989, Translated from the Russian by H. H. McFaden.
- [RF15] A. Sisto R. Frigerio, J.-F. Lafont, *Rigidity of high dimensional graph manifolds*, Astérisque (2015), no. 372, xxi+177.
- [Ril03] T. R. Riley, *Higher connectedness of asymptotic cones*, Topology **42** (2003), no. 6, 1289–1352.
- [Rin61] W. Rinow, *Die innere Geometrie der metrischen Räume*, Die Grundlehren der mathematischen Wissenschaften, Bd. 105, Springer-Verlag, Berlin, 1961.
- [Rip82] E. Rips, *Subgroups of small cancellation groups*, Bull. London Math. Soc. **14** (1982), no. 1, 45–47.
- [Rob47] R. M. Robinson, *On the decomposition of spheres*, Fund. Math. **34** (1947), 246–260.
- [Rob72] D. Robinson, *Finiteness conditions and generalized soluble groups, I*, vol. 62, Springer Verlag, 1972, Ergebnisse der Mathematik und ihrer Grenzgebiete.
- [Rob98] G. Robertson, *Crofton formulae and geodesic distance in hyperbolic spaces*, J. Lie Theory **8** (1998), 163–172.
- [Roe03] J. Roe, *Lectures on coarse geometry*, University Lecture Series, vol. 31, American Mathematical Society, Providence, RI, 2003.
- [Rol16] M. Roller, *Poc sets, median algebras and group actions*, preprint, arXiv:1607.0774, 2016.
- [Ros74] J. Rosenblatt, *Invariant measures and growth conditions*, Trans. Amer. Math. Soc. **193** (1974), 33–53.
- [Roy68] H. Royden, *Real analysis*, Macmillan, New York, 1968.
- [RR85] H. Rubin and J. E. Rubin, *Equivalents of the axiom of choice. II*, Studies in Logic and the Foundations of Mathematics, vol. 116, North-Holland Publishing Co., Amsterdam, 1985.
- [RS94] E. Rips and Z. Sela, *Structure and rigidity in hyperbolic groups. I*, Geom. Funct. Anal. **4** (1994), no. 3, 337–371.
- [RS97] ———, *Cyclic splittings of finitely presented groups and the canonical JSJ decomposition*, Ann. of Math. **146** (1997), 53–109.
- [Rud87] W. Rudin, *Real and complex analysis*, McGraw-Hill International editions, 1987.
- [Sal12] A.W. Sale, *Short conjugators and compression exponents in free solvable groups*, preprint, arXiv:1202.5343v1, 2012.
- [Sal15] ———, *Metric behaviour of the Magnus embedding*, Geometriae Dedicata **176** (2015), 303–315.
- [Sap07] M. V. Sapir, *Some group theory problems*, Internat. J. Algebra Comput. **17** (2007), no. 5-6, 1189–1214.
- [Sap14] M. Sapir, *Combinatorial algebra: Syntax and semantics*, Springer Verlag, 2014.
- [Sap15] ———, *On groups with locally compact asymptotic cones*, International Journal of Algebra and Computation **25** (2015), 37–40.
- [Sau06] R. Sauer, *Homological invariants and quasi-isometry*, Geom. Funct. Anal. **16** (2006), no. 2, 476–515.
- [SBR02] M. Sapir, J.-C. Birget, and E. Rips, *Isoperimetric and isodiametric functions of groups*, Ann. of Math. (2) **156** (2002), no. 2, 345–466.
- [Sch38] I. J. Schoenberg, *Metric spaces and positive definite functions*, Trans. Amer. Math. Soc. **44** (1938), 522–536.
- [Sch85] R. Scharlau, *A characterization of Tits buildings by metrical properties*, J. London Math. Soc. **32** (1985), no. 2, 317–327.
- [Sch96a] R. Schwartz, *Quasi-isometric rigidity and Diophantine approximation*, Acta Math. **177** (1996), no. 1, 75–112.

- [Sch96b] ———, *The quasi-isometry classification of rank one lattices*, Publ. Math. IHES **82** (1996), 133–168.
- [Sch97] ———, *Symmetric patterns of geodesics and automorphisms of surface groups*, Invent. Math. **128** (1997), no. 1, 177–199.
- [Sch13] V. Schur, *A quantitative version of the Morse lemma and quasi-isometries fixing the ideal boundary*, Journal of Functional Analysis **264** (2013), 815–836.
- [Sco73] P. Scott, *Compact submanifolds of 3-manifolds*, J. London Math. Soc. **7** (1973), 246–250.
- [Sco78] ———, *Subgroups of surface groups are almost geometric*, J. London Math. Soc. (2) **17** (1978), no. 3, 555–565.
- [Sco83] ———, *The geometry of 3-manifolds*, Bull. of the LMS **15** (1983), 401–487.
- [Sel60] A. Selberg, *On discontinuous groups in higher-dimensional symmetric spaces*, Contributions to Function Theory (K. Chandrasekharan, ed.), Tata Inst. of Fund. Research, Bombay, 1960, pp. 147–164.
- [Sel92] Z. Sela, *Uniform embeddings of hyperbolic groups in Hilbert spaces*, Israel J. Math. **80** (1992), no. 1-2, 171–181.
- [Sel95] ———, *The isomorphism problem for hyperbolic groups. I*, Ann. of Math. (2) **141** (1995), no. 2, 217–283.
- [Sel97a] ———, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups. II*, Geom. Funct. Anal. **7** (1997), no. 3, 561–593.
- [Sel97b] ———, *Structure and rigidity in hyperbolic groups and discrete groups in rank 1 Lie groups, II*, Geom. Funct. Anal. **7** (1997), 561–593.
- [Sel99] ———, *Endomorphisms of hyperbolic groups. I. The Hopf property*, Topology **38** (1999), no. 2, 301–321.
- [Sel01] ———, *Diophantine geometry over groups. I. Makanin-Razborov diagrams*, Publ. Math. IHES **93** (2001), 31–105.
- [Sel03] Z. Sela, *Diophantine geometry over groups. II. Completions, closures and formal solutions*, Israel J. Math. **134** (2003), 173–254.
- [Sel04] ———, *Diophantine geometry over groups. IV. An iterative procedure for validation of a sentence*, Israel J. Math. **143** (2004), 1–130.
- [Sel05a] ———, *Diophantine geometry over groups. III. Rigid and solid solutions*, Israel J. Math. **147** (2005), 1–73.
- [Sel05b] ———, *Diophantine geometry over groups. V₁. Quantifier elimination. I*, Israel J. Math. **150** (2005), 1–197.
- [Sel06a] ———, *Diophantine geometry over groups. V₂. Quantifier elimination. II*, Geom. Funct. Anal. **16** (2006), no. 3, 537–706.
- [Sel06b] ———, *Diophantine geometry over groups. VI. The elementary theory of a free group*, Geom. Funct. Anal. **16** (2006), no. 3, 707–730.
- [Sel09] ———, *Diophantine geometry over groups. VII. The elementary theory of a hyperbolic group*, Proc. Lond. Math. Soc. (3) **99** (2009), no. 1, 217–273.
- [Sel13] ———, *Diophantine geometry over groups VIII: Stability*, Ann. of Math. (2) **177** (2013), no. 3, 787–868.
- [Ser80] J. P. Serre, *Trees*, Springer, New York, 1980.
- [Sha98] Y. Shalom, *The growth of linear groups*, J. Algebra **199** (1998), no. 1, 169–174.
- [Sha00] Y. Shalom, *Rigidity of commensurators and irreducible lattices*, Invent. Math. **141** (2000), no. 1, 1–54.
- [Sha04] ———, *Harmonic analysis, cohomology, and the large-scale geometry of amenable groups*, Acta Math. **192** (2004), no. 2, 119–185.
- [Sha06] ———, *The algebraization of Kazhdan’s property (T)*, International Congress of Mathematicians. Vol. II, Eur. Math. Soc., Zürich, 2006, pp. 1283–1310.
- [She78] S. Shelah, *Classification theory and the number of non-isomorphic models*, Studies in Logic and the Foundations of Mathematics, vol. 92, North-Holland Publishing Co., 1978.
- [Sho54a] M. Sholander, *Medians and betweenness*, Proc. Amer. Math. Soc. **5** (1954), 801–807.
- [Sho54b] ———, *Medians, lattices and trees*, Proc. Amer. Math. Soc. **5** (1954), 808–812.
- [Sil] L. Silbermann, *Equivalence between properties (T) and (FH)*, unpublished notes.
- [Šir76] V. L. Širvanjan, *Imbedding of the group $B(\infty, n)$ in the group $B(2, n)$* , Izv. Akad. Nauk SSSR Ser. Mat. **40** (1976), no. 1, 190–208, 223.

- [Siu80] Y. T. Siu, *The complex-analyticity of harmonic maps and the strong rigidity of compact Kähler manifolds*, Ann. of Math. (2) **112** (1980), 73–111.
- [Sol11] Y. Solomon, *Substitution tilings and separated nets with similarities to the integer lattice*, Israel J. Math. **181** (2011), 445–460.
- [SS05] E. Stein and R. Shakarchi, *Real analysis. measure theory, integration, and hilbert spaces*, Princeton Lectures in Analysis, III, Princeton University Press, 2005.
- [ST10] Y. Shalom and T. Tao, *A finitary version of Gromov’s polynomial growth theorem*, Geom. Funct. Anal. **20** (2010), no. 6, 1502–1547.
- [Sta68] J. Stallings, *On torsion-free groups with infinitely many ends*, Ann. of Math. **88** (1968), 312–334.
- [Sta83] ———, *Topology of finite graphs*, Inv. Math. **71** (1983), 551–565.
- [Sta00] ———, *Commutative rings and groups*, UC Berkeley Lecture Notes, 2000.
- [Str74] R. A. Struble, *Metrics in locally compact groups*, Compositio Math. **28** (1974), 217–222.
- [Str06] G. Strang, *Linear algebra and its applications*, Thomson, Brooks/Cole, 2006.
- [Sul81] D. Sullivan, *On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions*, Riemann surfaces and related topics, Proceedings of the 1978 Stony Brook Conference, Ann. Math. Studies, vol. 97, Princeton University Press, 1981, pp. 465–496.
- [Sul14] H. Sultan, *Hyperbolic quasi-geodesics in CAT(0) spaces*, Geom. Dedicata **169** (2014), 209–224.
- [Šva55] A. S. Švarc, *A volume invariant of coverings*, Dokl. Akad. Nauk SSSR (N.S.) **105** (1955), 32–34.
- [SW79] P. Scott and T. Wall, *Topological methods in group theory*, “Homological Group Theory”, London Math. Soc. Lecture Notes, vol. 36, 1979, pp. 137–204.
- [SW05] M. Sageev and D. Wise, *The Tits alternative for CAT(0) cubical complexes*, Bull. London Math. Soc. **37** (2005), no. 5, 706–710.
- [Swa93] G. A. Swarup, *A note on accessibility*, Geometric group theory, Vol. 1 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 181, Cambridge Univ. Press, Cambridge, 1993, pp. 204–207.
- [Świ01a] J. Świątkowski, *A class of automorphism groups of polygonal complexes*, Q. J. Math. **52** (2001), no. 2, 231–247.
- [Świ01b] ———, *Some infinite groups generated by involutions have Kazhdan’s property (T)*, Forum Math. **13** (2001), no. 6, 741–755.
- [SX12] N. Shanmugalingam and X. Xie, *A rigidity property of some negatively curved solvable Lie groups*, Commentarii Mathematici Helvetici **87** (2012), 805–823.
- [SY94] R. Schoen and S.-T. Yau, *Lectures on differential geometry*, Conference Proceedings and Lecture Notes in Geometry and Topology, I, International Press, Cambridge, MA, 1994.
- [Tao08] T. Tao, *Product set estimates for non-commutative groups*, Combinatorica **28** (2008), no. 5, 547–594.
- [Tar38] A. Tarski, *Algebraische fassung des massproblems*, Fund. Math. **31** (1938), 47–66.
- [Tar86] ———, *Collected papers. Vol. 1*, Contemporary Mathematicians, Birkhäuser Verlag, Basel, 1986, 1921–1934, Edited by Steven R. Givant and Ralph N. McKenzie.
- [Tho02] B. Thornton, *Asymptotic cones of symmetric spaces*, Ph.D. Thesis, University of Utah, <http://www.math.utah.edu/theses/2002/thornton/blake-thornton-dissertation.pdf>, 2002.
- [Thu97] W. Thurston, *Three-dimensional geometry and topology, I*, Princeton Mathematical Series, vol. 35, Princeton University Press, 1997.
- [Tit72] J. Tits, *Free subgroups in linear groups*, Journal of Algebra **20** (1972), 250–270.
- [Tuk81] P. Tukia, *A quasiconformal group not isomorphic to a Möbius group*, Ann. Acad. Sci. Fenn. Ser. A I Math. **6** (1981), 149–160.
- [Tuk85] ———, *Differentiability and rigidity of Moebius groups*, Invent. Math. **82** (1985), 555–578.
- [Tuk86] ———, *On quasiconformal groups*, J. Analyse Math. **46** (1986), 318–346.
- [Tuk88] ———, *Homeomorphic conjugates of Fuchsian groups*, J. Reine Angew. Math. **391** (1988), 1–54.

- [Tuk94] ———, *Convergence groups and Gromov's metric hyperbolic spaces*, New Zealand J. Math. **23** (1994), no. 2, 157–187.
- [TW93] C. Thomassen and W. Woess, *Vertex-transitive graphs and accessibility*, J. Combin. Theory Ser. B **58** (1993), no. 2, 248–268.
- [Uml91] K. A. Umlauf, *über Die Zusammensetzung Der Endlichen Continuierlichen Transformationsgruppen, Insbesondere Der Gruppen Vom Range Null*, Nabu Press, 2010, 1891, Inaugural-Dissertation, Universität Leipzig.
- [Väi71] J. Väisälä, *Lectures on n -dimensional quasiconformal mappings*, Lecture Notes in Math., vol. 229, Springer, 1971.
- [Väi05] ———, *Gromov hyperbolic spaces*, Expo. Math. **23** (2005), no. 3, 187–231.
- [Väi85] ———, *Quasi-Möbius maps*, J. Analyse Math. **44** (1984/85), 218–234.
- [Var99] N. Varopoulos, *Distance distortion on Lie groups*, Random walks and discrete potential theory (Cortona, 1997), Sympos. Math., XXXIX, Cambridge Univ. Press, Cambridge, 1999, pp. 320–357.
- [Vas12] S. Vassileva, *The Magnus embedding is a quasi-isometry*, Internat. J. Algebra Comput. **22** (2012), no. 8, 1240005.
- [vdV93] M. L. J. van de Vel, *Theory of convex structures*, North-Holland Mathematical Library, vol. 50, North-Holland Publishing Co., Amsterdam, 1993.
- [Ver82] A. Vershik, *Amenability and approximation of infinite groups*, Selecta Math. Soviet. **2** (1982), no. 4, 311–330.
- [Ver93] E. R. Verheul, *Multimedians in metric and normed spaces*, CWI Tract, vol. 91, Stichting Mathematisch Centrum voor Wiskunde en Informatica, Amsterdam, 1993.
- [Vie27] L. Vietoris, *Über den höheren Zusammenhang kompakter Räume und eine Klasse von zusammenhangstreuen Abbildungen*, Mathematische Annalen **97** (1927), 454–472.
- [Vit05] G. Vitali, *Sul problema della misura dei gruppi di punti di una retta*, Gamberini e Parmeggiani, 1905, Bologna.
- [vN28] J. von Neumann, *Über die Definition durch transfiniten Induktion und verwandte Fragen der allgemeinen Mengenlehre*, Math. Ann. **99** (1928), 373–391.
- [vN29] ———, *Zur allgemeinen theorie des masses*, Fund. math. **13** (1929), 73–116.
- [VSCC92] N. Th. Varopoulos, L. Saloff-Coste, and T. Coulhon, *Analysis and geometry on groups*, Cambridge Tracts in Mathematics, vol. 100, Cambridge University Press, Cambridge, 1992.
- [Vuo88] M. Vuorinen, *Conformal geometry and quasiregular mappings*, Lecture Notes in Math, vol. 1319, Springer, 1988.
- [Wag85] S. Wagon, *The Banach-Tarski paradox*, Cambridge Univ. Press, 1985.
- [War83] F. W. Warner, *Foundations of differentiable manifolds and Lie groups*, Graduate Texts in Mathematics, 94, Springer-Verlag, 1983.
- [War12] E. Warner, *Ultraproducts and foundations of higher order Fourier analysis*, Thesis, Princeton University, 2012.
- [Wat82] Y. Watatani, *Property (T) of Kazhdan implies property (FA) of Serre*, Math. Japonica **27** (1982), 97–103.
- [Weh73] B. A. F. Wehrfritz, *Infinite linear groups*, Ergebnisse der Mathematik und ihrer Grenzgebiete, vol. 76, Springer, 1973.
- [Wei84] B. Weisfeiler, *On the size and structure of finite linear groups*, arXiv:1203.1960, 1984.
- [Wen05] S. Wenger, *Isoperimetric inequalities of Euclidean type in metric spaces*, Geom. Funct. Anal. **15** (2005), 534–554.
- [Wen08] ———, *Gromov hyperbolic spaces and the sharp isoperimetric constant*, Invent. Math. **171** (2008), no. 1, 227–255.
- [Wen11] ———, *Nilpotent groups without exactly polynomial Dehn function*, J. Topol. **4** (2011), no. 1, 141–160.
- [Why99] K. Whyte, *Amenability, bilipschitz equivalence, and the von Neumann conjecture*, Duke Math. J. **99** (1999), 93–112.
- [Why01] ———, *The large scale geometry of the higher Baumslag-Solitar groups*, Geom. Funct. Anal. **11** (2001), no. 6, 1327–1343.
- [Wis04] D. T. Wise, *Cubulating small cancellation groups*, Geom. Funct. Anal. **14** (2004), no. 1, 150–214.

- [Woe00] W. Woess, *Random walks on infinite graphs and groups*, Cambridge University Press, 2000.
- [Wol68] J. Wolf, *Growth of finitely generated solvable groups and curvature of riemannian manifolds*, *J. Diff. Geom.* **2** (1968), 421–446.
- [Wor07] K. Wortman, *Quasi-isometric rigidity of higher rank S -arithmetic lattices*, *Geometry & Topology* **11** (2007), 995–1048.
- [WW75] J. H. Wells and L. R. Williams, *Embeddings and extensions in analysis*, Springer-Verlag, New York-Heidelberg, 1975.
- [Wys88] J. Wysoczński, *On uniformly amenable groups*, *Proc. Amer. Math. Soc.* **102** (1988), no. 4, 933–938.
- [Xie06] X. Xie, *Quasi-isometric rigidity of Fuchsian buildings*, *Topology* **45** (2006), 101–169.
- [Xie12] ———, *Quasisymmetric maps on the ideal boundary of a negatively curved solvable Lie group*, *Mathematische Annalen* **353** (2012), 727–746.
- [Xie13] ———, *Quasisymmetric maps on reducible Carnot groups*, *Pacific Journal of Mathematics* **265** (2013), 113–122.
- [Xie14] ———, *Large scale geometry of negatively curved \mathbb{R}^n* , *Geometry & Topology* **18** (2014), 831–872.
- [Xie16] ———, *Some examples of quasiisometries of nilpotent Lie groups*, *J. Reine Angew. Math. (Crelle's Journal)* **718** (2016), 25–38.
- [Yam53] H. Yamabe, *A generalization of a theorem of Gleason*, *Ann. of Math. (2)* **58** (1953), 351–365.
- [Yam04] A. Yaman, *A topological characterisation of relatively hyperbolic groups*, *J. Reine Angew. Math. (Crelle's Journal)* **566** (2004), 41–89.
- [Yu00] G. Yu, *The coarse Baum-Connes conjecture for spaces which admit a uniform embedding into Hilbert space*, *Invent. Math.* **139** (2000), no. 1, 201–240.
- [Yu05] ———, *Hyperbolic groups admit proper isometric actions on ℓ^p -spaces*, *Geom. Funct. Anal.* **15** (2005), no. 5, 1144–1151.
- [Zas38] H. Zassenhaus, *Beweis eines satzes über diskrete gruppen*, *Abh. Math. Sem. Hansisch. Univ.* **12** (1938), 289–312.
- [Zer04] E. Zermelo, *Beweis, dass jede Menge wohlgeordnet werden kann*, *Math. Ann.* **59** (1904), 514–516.
- [Zim84] R. Zimmer, *Ergodic theory and semisimple groups*, *Monographs in Math*, vol. 81, Birkhauser, 1984.
- [Žuk96] A. Žuk, *La propriété (T) de Kazhdan pour les groupes agissant sur les polyèdres*, *C. R. Acad. Sci. Paris Sér. I Math.* **323** (1996), no. 5, 453–458.
- [Žuk03] ———, *Property (T) and Kazhdan constants for discrete groups*, *Geom. Funct. Analysis* **13** (2003), 643–670.