Hand-out notes for Geometric Group Theory

Cornelia Druţu

January 16, 2025

Abstract

These notes collect material seen in courses from previous years, and extra information on concepts used in the course, and are to be used as a reference only. The material in these notes is not examinable.

1 Groups and their actions

1.1 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 $K \lhd G$ 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 $\phi : G \to G$, $\phi(K) = K$. Note that every characteristic subgroup is normal (since conjugation is an automorphism). But not every normal subgroup is characteristic:

Example 1.1. 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 $\phi : \mathbb{Z}^2 \to \mathbb{Z}^2$, $\phi(m, n) = (n, m)$.

Definition 1.2. 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.

The following is a basic result in group theory:

Lemma 1.3. *If G is a group, $N \lhd G$, and $A \lhd B < G$, then BN/AN is isomorphic to $B/A(B \cap N)$.*

Definition 1.4. 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 $\operatorname{Tor} G = \{g \in G \mid g \text{ of finite order}\}$ of the group G, sometimes called the *torsion* of G, is in general not a subgroup.

Definition 1.5. 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 1.6. 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 1.7. *If $N \lhd H$ and $H \lhd G$, 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 $H \to F$ 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 $N, N_1, .., N_k$. The subgroup $K = \bigcap_{g \in G} gNg^{-1} = N \cap N_1 \cap \cdots \cap N_k$ is normal in G and has finite index in N, since each of the subgroups $N_1, \ldots, N_k$ has finite index in H. $\square$

Proposition 1.8. *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 $H \leqslant G$ be a subgroup of index n. We list the left cosets of H:

$$H = g_1 \cdot H, g_2 \cdot H, \ldots, g_n \cdot H,$$

and label these cosets by the numbers $\{1, \ldots, n\}$. The action by left multiplication of G on the set of left cosets of H defines a homomorphism $\phi : G \to S_n$ such that $\phi(G)$ acts transitively on $\{1, 2, \ldots, n\}$ 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 $\phi : G \to S_n$ is such that $\phi(G)$ acts transitively on $\{1, 2, \dots, n\}$, then $G/\phi^{-1}(\mathrm{Stab}\,(1))$ has cardinality n.

Since the group G is finitely generated, a homomorphism $\phi : G \to S_n$ 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 $\phi : G \to S_n$ 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 $\varphi : G \to G$, $\varphi(H)$ is a subgroup of index n. According to (1) the set $\{\varphi(H) \mid \varphi \in \mathrm{Aut}\,(G)\}$ is finite, equal $\{H, H_1, \dots, H_k\}$. It follows that

$$K = \bigcap_{\varphi \in \mathrm{Aut}\,(G)} \varphi(H) = H \cap H_1 \cap \ldots \cap H_k .$$

Then K is a characteristic subgroup of finite index in H hence in G. $\square$

Exercise 1.9. Does the conclusion of Proposition 1.8 still hold for groups which are not finitely generated?

Let S be a subset in a group G, and let $H \leqslant G$ be a subgroup. The following are equivalent:

1. H is the smallest subgroup of G containing S ;

2. $H = \bigcap_{S \subset G_1 \leqslant G} G_1$;

3. $H = \left\{ s_1 s_2 \cdots s_n \ : \ n \in \mathbb{N}, s_i \in S \text{ or } s_i^{-1} \in S \text{ for every } i \in \{1, 2, \dots, n\} \right\}$.

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 $K \lhd G$ is *normally generated* by a set $R \subset K$ if K is the smallest normal subgroup of G which contains R, i.e.

$$K = \bigcap_{R \subset N \lhd G} N .$$

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 $\langle R \rangle^G$.

1.2 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 $N_i \lhd G, i = 1, \ldots, k$, if one of the following equivalent statements holds:

- $G = N_1 \cdots N_k$ and

$$N_i \cap N_1 \cdot \ldots \cdot N_{i-1} \cdot N_{i+1} \cdot \ldots \cdot N_k = \{1\} \text{ for all } i;$$

- for every element g of G there exists a unique k-tuple

$$(n_1, \ldots, n_k), n_i \in N_i, i = 1, \ldots, k$$

  such that $g = n_1 \cdots n_k$.

Then G is isomorphic to the direct product $N_1 \times \ldots \times N_k$. 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 1.10. 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* $G \to H$, i.e. a homomorphism which restricts to the identity on H, and whose kernel is N.

Observe that the map $\varphi : H \to \mathrm{Aut}\,(N)$ 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 $\varphi : H \to \mathrm{Aut}\,(N)$, 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 1.11. 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 $N \rtimes_\phi H$ 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

$$\ldots G_{n-1} \xrightarrow{\varphi_{n-1}} G_n \xrightarrow{\varphi_n} G_{n+1} \ldots$$

such that $\mathrm{Im}\,\varphi_{n-1} = \mathrm{Ker}\,\varphi_n$ for every n. A *short exact sequence* is an exact sequence of the form:

$$\{1\} \longrightarrow N \xrightarrow{\varphi} G \xrightarrow{\psi} H \longrightarrow \{1\}\,. \tag{1}$$

In other words, φ is an isomorphism from N to a normal subgroup $N' \lhd G$ and ψ descends to an isomorphism $G/N' \simeq H$.

Definition 1.12. A short exact sequence *splits* if there exists a homomorphism $\sigma : H \to G$ (called a *section*) such that

$$\psi \circ \sigma = \mathrm{Id}\,.$$

When the sequence splits we shall sometimes write it as

$$1 \to N \to G \overset{\curvearrowleft}{\to} H \to 1.$$

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 $\psi : G \to H$ is the retraction.

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 $\mathbb{S}^1 \subset \mathbb{C}$ 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 $s(x) = -x$.

Examples 1.13.

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 $\mathbb{Z}_2 = \{\mathrm{Id}, (12)\}$.
4. The group $(\mathrm{Aff}(\mathbb{R})\,, \circ)$ of affine maps $f : \mathbb{R} \to \mathbb{R}, f(x) = ax + b\,$, 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 1.14.

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 $\mathrm{Isom}(\mathbb{R}^n)$ 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 $0 < \lambda \leqslant 1$) $\mathbf{w} = \lambda\mathbf{v}$ 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.

Exercise 1.15. 1. Prove the statement (2) of Proposition 1.14. 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 $\mathrm{Isom}(\mathbb{R}^n)$. 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 $O(n)$?

1.3 Group actions

Let G be a group and X be a set. An *action of G on X on the left* is a map

$$\mu : G \times X \to X, \quad \mu(g, a) = g(a),$$

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 1.16. 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

$$\mu : X \times G \to X, \quad \mu(a, g) = (a)g,$$

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.

We often simply write gx instead of $\mu(g, x)$ or $g(x)$ (respectively xg instead of $\mu(x, g)$ or $(x)g$).

If not specified, an action of a group G on a set X is always on the left, and it is often denoted $G \curvearrowright X$.

An equivalent definition of a left action of a group is as a homomorphism from G to the group $Bij(X)$ of bijections of X.

Indeed, an action on the left

$$\mu : G \times X \to X, \quad \mu(g, a) = g(a),$$

defines $\varphi : G \to Bij(X)$ by $\varphi(g)(x) = \mu(g, x)$.

Property (2) of μ implies that $\varphi(g_1 g_2) = \varphi(g_1) \circ \varphi(g_2)$.

Conversely, given a group homomorphism $\varphi : G \to Bij(X)$, we define

$$\mu : G \times X \to X, \quad \mu(g, a) = \varphi(g)(a),$$

and check that it satisfies the required properties.

An action is called *effective* or *faithful* if this homomorphism is injective.

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

$$G \to \mathrm{Isom}(X).$$

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 $f : X \to Y$ 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 $f : X \to Y$ is called *G–equivariant* if

$$f(\mu(g, x)) = \nu(g, f(x)), \quad \forall g \in G, x \in X.$$

2 Metric spaces and graphs

2.1 General metric spaces

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

(M1) $\mathrm{dist}(x, y) \geqslant 0$ for all $x, y \in X$; $\mathrm{dist}(x, y) = 0$ if and only if $x = y$;

(M2) (Symmetry) for all $x, y \in X$, $\mathrm{dist}(y, x) = \mathrm{dist}(x, y)$;

(M3) (Triangle inequality) for all $x, y, z \in X$, $\mathrm{dist}(x, z) \leqslant \mathrm{dist}(x, y) + \mathrm{dist}(y, z)$.

The function dist is called *metric* or *distance function*.

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 $\mathrm{dist}(x, A)$ for the *minimal distance* from x to A, i.e.

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

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

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

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. $\mathcal{N}_R(A) = \{x \in X : \mathrm{dist}(x, A) < R\}$. 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

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

Given two metric spaces (X, dist_X), (Y, dist_Y), a map $f : X \to Y$ is an *isometric embedding* if for every $x, x' \in X$

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

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 $f : X \to Y$ is called *a similarity with factor λ* if for all $x, x' \in X$,

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

The group of isometries of a metric space X is denoted $\mathrm{Isom}(X)$.

2.2 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*.

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

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

[2] Also known as *digons*.

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 bigons[3].

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 $E \to V$ is the choice of an *orientation* of Γ: An orientation of Γ is a choice of two maps

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

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 2.1. 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 $\overline{\Gamma}$, its edge-set by $\overline{E}$, and oriented edges by $\bar{e}$.

Convention 2.2. 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 well-known graphs.

Example 2.3 (n-rose). This graph, denoted R_n, has one vertex and n edges connecting this vertex to itself.

Example 2.4. [i-star or i-pod] This graph, denoted T_i, has $i+1$ vertices, $v_0, v_1, \ldots, v_i$. 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 2.5 (n-circle). This graph, denoted C_n, has n vertices which are identified with the n–th roots of unity:

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

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

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

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

Example 2.6 (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 2.7 (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 2.8 (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

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

A *morphism* of graphs $f : \Gamma \to \Gamma'$ is a pair of maps $f_V : V(\Gamma) \to V(\Gamma')$, $f_E : E(\Gamma) \to E(\Gamma')$ 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. A *monomorphism* of graphs is a morphism such that the corresponding maps f_V, f_E are injective. The image of a monomorphism $\Gamma \to \Gamma'$ 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 $f : \Gamma \to \Gamma'$ 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 $\Gamma' \to \Gamma$. In other words, an isomorphism of graphs is an isomorphism of the corresponding cell complexes.

Exercise 2.9. Isomorphisms of graphs are morphisms such that the corresponding vertex and edge maps are bijective.

We use the notation $\mathrm{Aut}(\Gamma)$ 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 $[v_1, v_2], [v_2, v_3], \ldots, [v_n, v_{n+1}]$ 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 $v_i, i = 1, \ldots, n$, 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 2.10. 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* $f : \Gamma \to \Gamma'$ consists of a pair of maps (g, h):

1. A map $g : V(\Gamma) \to V(\Gamma')$ sending adjacent vertices to adjacent or equal vertices;

2. A *partially defined* map of the edge-sets:

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

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 $\Gamma_o \to \Gamma'$, 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 $V(\Gamma) \to V(\Gamma')$ 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_0. If an edge e connects two vertices in W^c, it is also retained and connects the same vertices in Γ'.

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

Exercise 2.11. If Γ is a tree and Λ is a subtree, then Γ' is again a tree.

2.3 Connected graphs as metric spaces

Let Γ be a connected graph. We introduce a metric dist on 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 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 $x, y \in \Gamma$,

$$\operatorname{dist}(x, y) = \inf_{\mathfrak{p}} (\operatorname{length}(\mathfrak{p})),$$

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