B8.5 Graph Theory Lecture Notes

Michaelmas Term 2024, 16 lectures
Lecturer: Paul Balister
Last updated: October 2, 2024

These notes are to accompany the lectures in Michaelmas Term 2024 on graph theory for Oxford Part B Mathematics, and are adapted from notes by Oliver Riordan, which in turn were adapted from notes by Alex Scott and Colin McDiarmid. They also owe much to the book *Modern Graph Theory*, Springer-Verlag, 1998 by Béla Bollobás. If you spot any errors or have any comments, email: `Paul.Balister@maths.ox.ac.uk`.

Relationship to Part A Graph Theory

Part A Graph Theory is recommended but not required as a prerequisite. The course as lectured should be self-contained, though a few key results covered in Part A will be stated as exercises to complete yourself if you did not do Part A Graph Theory.

1 Introduction

We need some preliminary definitions and notation (see the end of these notes for a summary). We write $[n]$ for the set $\{1, 2, \dots, n\}$. For any set S, we write $\binom{S}{k}$ for the set of subsets of S of size k, that is, $\binom{S}{k} = \{A \subseteq S : |A| = k\}$ (so has cardinality $\binom{|S|}{k}$). Some authors write $S^{(k)}$ instead of $\binom{S}{k}$.

A (***simple***) ***graph*** G is an ordered pair (V, E), where V is a set and $E \subseteq \binom{V}{2}$. In this course V will almost always be finite and non-empty – this is assumed unless stated otherwise. The elements of V are called the ***vertices*** of G and the elements of E the ***edges*** of G.

For brevity, we often write uv for the edge $\{u, v\}$ (so uv means the same as vu). We say that u and v are ***adjacent*** in G if uv is an edge of G. A vertex v and an edge e are ***incident*** if v is one of the ***end vertices*** of e, i.e., one of the two vertices in e. Two edges ***meet*** if they intersect, i.e., share a vertex. Graphically, we represent vertices as points (or more often blobs) and edges as lines or curves joining pairs of points (blobs); how a graph is drawn is irrelevant as far as the structure of the graph itself is concerned. The reason for using blobs is that it makes clear in the drawing where the vertices are: we may have to draw the lines/curves for two edges so that they cross even though the edges do not share a vertex.

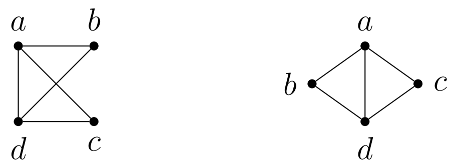

Figure 1: Two ways of drawing the same graph $G = (\{a, b, c, d\}, \{ab, ac, ad, bd, cd\})$.

If $G = (V, E)$, we write $V(G)$ for V and $E(G)$ for E. The ***order*** of a graph G, denoted by $|G|$ or $v(G)$, is the number of vertices, so $|G| = v(G) = |V(G)|$. The ***size*** of G is the number of edges, $e(G) = |E(G)|$; however, sometimes 'size' is used to mean 'order', so it is safest to avoid this term.

Graphs G and H are ***isomorphic*** if there exists a bijection $\varphi\colon V(G) \to V(H)$ such that, for each $x, y \in V(G)$, $xy \in E(G)$ iff $\varphi(x)\varphi(y) \in E(H)$. In this case we say that φ is an ***isomorphism***,[1] and write $G \cong H$. It is easy to check that isomorphism of graphs is an equivalence relation, and simply amounts to a 'relabeling' of the vertices. Often we do not make a distinction between isomorphic graphs, treating them as the same.[2]

A graph G is ***complete*** if $E(G) = \binom{V(G)}{2}$. We write $K_n = ([n], \binom{[n]}{2})$ for the complete graph on the vertex set $[n]$. Clearly any complete graph of order n is isomorphic to K_n. A graph G is ***empty*** if $E(G) = \emptyset$. We write $E_n = ([n], \emptyset)$ for the empty graph on the vertex set $[n]$. Any empty graph of order n is isomorphic to E_n. A graph G is a ***cycle*** on n vertices, or a cycle of ***length*** n, $(n \geqslant 3)$, if it is isomorphic to $C_n := ([n], \{12, 23, \dots, (n-1)n, n1\})$. A graph G is a ***path*** on n vertices, or a path of ***length*** $n-1$, $(n \geqslant 1)$, if it is isomorphic to $P_n := ([n], \{12, 23, \dots, (n-1)n\})$.

Warning. Some authors define P_n as a path with n edges, while in this course P_n has n vertices. Similarly, the length of a path is usually the number of edges, but some use it to mean the number of vertices. Always check which definitions are being used! Sometimes the term ***edge length*** is used to emphasise that it is the number of edges being counted.

Another commonly encountered graph is the ***complete bipartite*** graph $K_{a,b}$, which has vertex set $A \cup B$ with $|A| = a$, $|B| = b$, $A \cap B = \emptyset$, and all edges between A and B: $E = \{xy : x \in A,\ y \in B\}$. The special case when $a = 1$ is also called a ***star***.

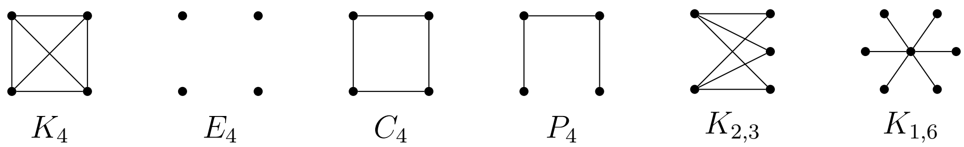

Figure 2: Some special graphs.

A graph H is a ***subgraph*** of G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. A ***spanning*** subgraph is one that includes all the vertices: $V(H) = V(G)$. An ***induced*** subgraph is one that includes all possible edges, i.e., all edges with both endvertices in $V(H)$: $E(H) = E(G) \cap \binom{V(H)}{2}$. Indeed, given any $W \subseteq V(G)$ there is a unique induced subgraph with vertex set W which we write as $G[W]$, and which we call the subgraph of G ***induced by*** W.

[1]More generally we can define a graph ***homomorphism*** as a map $\varphi\colon V(G) \to V(H)$ such that $xy \in E(G)$ implies $\varphi(x)\varphi(y) \in E(H)$. An isomorphism is then just an invertible map φ such that both φ and φ^{-1} are homomorphisms.

[2]Sometimes we talk of 'labelled graphs', meaning the actual values of the vertices are important, and 'unlabelled graphs' for graphs considered only up to isomorphism.

We often say that H is a subgraph of G (or more precisely, that G contains a ***copy*** of H) to mean that G has a subgraph isomorphic to H.[3]

The ***complement*** of a graph $G = (V, E)$ is $\overline{G} = (V, \binom{V}{2} \setminus E)$. Thus $\overline{K}_n = E_n$ and $\overline{E}_n = K_n$. For an edge e, we write[4] $G - e$ for the subgraph $(V, E \setminus \{e\})$, obtained by deleting the edge e from G. For $e \in E(\overline{G})$, $G + e = (V, E \cup \{e\})$ is the graph obtained by adding the edge e to G. For a vertex v, we write $G - v$ for the subgraph induced by $V \setminus \{v\}$, i.e., the subgraph obtained from G by deleting v and (as we must) all edges incident with v. We similarly define $G - S$ when S is a set of edges, or a set of vertices, in the obvious way.

In much of the following, unless otherwise indicated, the implicitly assumed setting is that of an arbitrary graph $G = (V, E)$.

Degrees

The ***degree*** of a vertex v is the number of incident edges,

$$d(v) = \big|\{w \in V : vw \in E\}\big|.$$

We write $d_G(v)$ if we want to specify the graph. A vertex w is a ***neighbour*** of v if v and w are adjacent, i.e., $vw \in E$. The set $N(v) = N_G(v) = \{w \in V : vw \in E\}$ is the ***neighbourhood*** of v, so $d(v) = |N(v)|$. Some authors write $\Gamma(v)$ instead of $N(v)$.

A graph G is r-***regular*** if every vertex has degree r. If $d(v) = 0$, v is an ***isolated vertex***. If $V = \{v_1, \ldots, v_n\}$, the ***degree sequence*** of G is the sequence $d(v_1), d(v_2), \ldots, d(v_n)$, often arranged in nondecreasing order. For example, K_n is $(n-1)$-regular and the degree sequence of P_5 is $1, 1, 2, 2, 2$.

Lemma 1.1 (Handshaking Lemma)**.** *For any graph G,*

$$\sum_{v \in V(G)} d(v) = 2e(G).$$

Proof. Consider the number of pairs (v, e) where v is a vertex of G and e is an edge of G incident with v. We count them in two different ways. Firstly, each vertex v is in exactly $d(v)$ such pairs, so there are $\sum_{v \in V(G)} d(v)$ pairs in total. Secondly, each edge e of G is in exactly two such pairs, so there are $2|E(G)| = 2e(G)$ pairs. □

Corollary 1.2. *For any graph G, the number of vertices with odd degree is even.* □

Paths, cycles and walks in graphs

A ***path on*** t ***vertices in*** G is a subgraph of G isomorphic to P_t; a ***cycle of length*** t ***in*** G is a subgraph isomorphic to C_t. We usually just list the vertices to describe a path or cycle. Thus $v_1 v_2 \cdots v_t$ is a path on t vertices (of length $t-1$) in G if and only

[3] Equivalently, there is an injective homomorphism from H to G.

[4] Some people write $G \setminus e$ for $G - e$. Not to be confused with G/e, defined later.

if $v_1, \ldots, v_t$ are *distinct* vertices of G and $v_1v_2, \ldots, v_{t-1}v_t$ are edges of G.[5] Similarly, $v_1v_2 \cdots v_tv_1$ is a ***cycle*** in G if and only if $t \geqslant 3$, $v_1, \ldots, v_t$ are *distinct* vertices of G, and $v_1v_2, \ldots, v_{t-1}v_t, v_tv_1$ are edges of G. A graph is ***acyclic*** if it contains no cycles.

More generally, we say $v_0v_1 \cdots v_t$ is a ***walk*** in G if $v_0, v_1, \ldots, v_t$ are (not necessarily distinct) vertices of G such that $v_iv_{i+1} \in E(G)$ for each $i = 0, 1, \ldots, t-1$. The ***length*** of a walk is the number of steps, here t. If $x = v_0$ and $y = v_t$ then we speak of a walk from x to y, or an x-y walk; an x-y path is defined similarly. A walk $v_0 \cdots v_t$ is ***closed*** if $v_t = v_0$.

Exercise. *Let G be a graph and $x, y \in V(G)$. Then G contains an x-y walk if and only if G contains an x-y path.*

In other words, if we want to get from x to y, then allowing ourselves to revisit vertices does not help. This simple observation is useful, allowing us to switch back and forth between using paths and walks to define connectedness, at any point using whichever definition is easiest to work with. The cleanest proof of the exercise is to consider a shortest x-y walk and show that it is a path (see Problem Sheet 0).

A graph G is ***connected*** if for all $x, y \in V(G)$ there is at least one x-y path (or walk) in G. The ***components*** of a general graph G are the maximal connected subgraphs. It is easy to check that G is the disjoint union of its components. Indeed, consider the relation $\sim$ on $V(G)$ defined by "$x \sim y$ iff there exists an x-y walk". It is easy to check that this is an equivalence relation, and that the components are the subgraphs induced by the equivalence classes.

Figure 3: (a) A connected graph. (b) A graph with 3 components.

We finish this section with a simple lemma giving a condition under which we are guaranteed that G contains a cycle.

Lemma 1.3. *Let G be a finite graph in which every vertex has degree at least 2. Then G contains a cycle.*

Proof. Pick $v_{\in}V(G)$ and $v_1 \in N(v_0)$. Now for each $i \geqslant 1$ we can successively pick $v_{i+1} \in N(v_i)$ such that $v_{i+1} \neq v_{i-1}$ (since $|N(v_i)| \geqslant 2$). Thus we have a sequence $v_0, v_1, v_2, \ldots$ such that $v_{i-1}v_i \in E(G)$ for every i, and v_{i-1}, v_i and v_{i+1} are always distinct. Since G has only finitely many vertices, some vertex must appear more than once. Pick $i < j$ such that $v_i = v_j$ and j is minimal. Then $j - i \geqslant 3$ and, by minimality of j, $v_i, \ldots, v_{j-1}$ are distinct. Thus $v_i \cdots v_{j-1}v_i$ is a cycle in G. □

[5]The two definitions of path in G are not quite the same: for existence, they are equivalent, but for counting paths, they differ by a factor of 2 for $t \geqslant 2$ as the subgraph corresponding to $v_1v_2 \cdots v_t$ is the same as for $v_tv_{t-1} \cdots v_1$. A similar comment applies to cycles with a different factor.

2 Trees

A ***tree*** is simply an acyclic connected graph. A general acyclic graph, or equivalently, a graph in which each component is a tree, is sometimes called a ***forest***.

Figure 4: (a) A tree. (b) A forest with 3 components.

Lemma 2.1. *The following are equivalent, where minimality/maximality is with respect to deleting/adding edges.*

(a) *T is a tree,*

(b) *T is a minimal connected graph,*

(c) *T is a maximal acyclic graph.*

The precise meaning of (b) is that $T = (V, E)$ is connected, but that for any strict subset $E' \subsetneq E$, (V, E') is not connected. Equivalently, T is connected, but for any edge e of T, $T - e$ is not connected. Similarly in (c), $T = (V, E)$ is acyclic, but (V, E') contains a cycle for any strict superset $E' \supsetneq E$. Equivalently, T is acyclic, but for any edge e of $\overline{T}$, $T + e$ is not acyclic.

Proof. Revision from Part A or exercise, as applicable. See Problem Sheet 0. □

An edge e in a connected graph is called a ***bridge*** if $G - e$ is disconnected. (In a general graph, e is a bridge if $G - e$ has more components than G.) Thus Lemma 2.1 implies that a connected graph G is a tree iff each edge is a bridge. Indeed, it is easy to see that for any edge e in a graph, e is a bridge iff e is not in any cycle.

A ***spanning tree*** of a graph G is a spanning subgraph of G that is a tree, i.e., a subgraph of G that is a tree containing all the vertices of G.

Corollary 2.2. *Every connected graph G has at least one spanning tree.*

Proof. Remove edges one by one, keeping the graph connected, until we can remove no more. The graph T that remains is a minimal connected graph with vertex set $V(G)$; by Lemma 2.1, T is a tree. □

A vertex v of a graph G is called a ***leaf*** if $d(v) = 1$. This term is most often used in the context of trees/forests.

Lemma 2.3. *Every tree on $n \geqslant 2$ vertices has at least one leaf.*

Proof. T is connected, so it has no isolated vertices (vertices of degree 0) as there can be no path from such a vertex to any other. But T has no cycle, so by Lemma 1.3 it must have a vertex of degree less than 2. Therefore it has a vertex v with degree 1. □

In fact, every tree with at least 2 vertices has at least two leaves; there are many proofs of this fact. One involves modifying the argument above slightly. Another way is to consider a longest path $v_0 \cdots v_t$ in the tree and show that $t > 0$ and that v_0 and v_t are both leaves (see Problem Sheet 0).

The significance of leaves is shown by the following simple result.

Lemma 2.4. *Let v be a leaf of a graph G. Then G is a tree iff $G - v$ is a tree.*

Proof. If G if connected and $x, y \in V(G) \setminus \{v\}$, then there is an x-y path P in G. But every vertex of P other than x and y has degree 2 in P, and hence degree at least 2 in G, so cannot be v. Thus P is an x-y path in $G-v$, so $G-v$ is connected. Conversely, if $G-v$ is connected then G is as every vertex in G (including v) has a path to the neighbour of v in G and so lies in the same component.

If $G - v$ has a cycle then obviously so does G. If G has a cycle C then all the vertices of C have degree at least 2 in G and so C cannot include v. Thus C is a cycle in $G - v$. Hence G is connected and acyclic iff $G - v$ is connected and acyclic. □

Lemma 2.5. *If T is a tree on n vertices, then $e(T) = n - 1$.*

Proof. We use induction on n; the case $n = 1$ is trivial. Let T be any tree with $n \geqslant 2$ vertices. By Lemma 2.3, T has a leaf v. By Lemma 2.4, $T' = T - v$ is a tree. Since T' has $n - 1$ vertices, by induction it has $n - 2$ edges. Thus T has $n - 1$ edges as removing v from T removes precisely $d(v) = 1$ edge. □

Combining Lemmas 2.1 and 2.5 gives some further characterisations of trees.

Corollary 2.6. *Let G be a graph with n vertices. The following are equivalent.*

(a) *G is a tree,*

(b) *G is connected and $e(G) \leqslant n - 1$,*

(c) *G is acyclic and $e(G) \geqslant n - 1$.*

Proof. (a) implies (b) and (c) by the definition of a tree and Lemma 2.5. Suppose that (b) holds. Then G has a spanning tree T which, by Lemma 2.5, has $n-1$ edges. A spanning subgraph includes all the vertices by definition, and since $e(T) = n-1 \geqslant e(G)$, in this case it includes all the edges too. Thus $T = G$ and so G is a tree, completing the proof that (b) implies (a). Now suppose that (c) holds. Each component C_i of G is connected and acyclic, so is a tree. Thus $e(C_i) = |C_i| - 1$ and so $e(G) = \sum e(C_i) = \sum(|C_i| - 1) = n - c$, where c is the number of components. Thus $c \leqslant 1$ and G is connected. □

Counting trees

Let's start with a simpler question: how many graphs $G = (V, E)$ are there with vertex set $[n]$? Each of the $\binom{n}{2}$ possible edges may or may not be included in E, with all possibilities allowed, so the answer is $2^{\binom{n}{2}}$. Note that we are not asking how many *isomorphism* classes there are: this is a much harder question. (Sometimes, counting graphs on a given vertex set is referred to as 'counting labelled graphs'; counting isomorphism classes is referred to as 'counting unlabelled graphs'.)

Counting trees is much harder than counting all graphs. The answer was found (but not really proved) by Cayley in 1889, though implicitly earlier by Borchardt in 1860; it is now known as ***Cayley's formula***.

Theorem 2.7. *For any $n \geqslant 1$ there are exactly n^{n-2} trees T with vertex set $[n]$.*

Proof. The result is trivial for $n = 1$ and 2, so fix $n \geqslant 3$. We shall map each tree on $[n]$ to its ***Prüfer code*** $\mathbf{c} = (c_1, c_2, \dots, c_{n-2})$, where $1 \leqslant c_i \leqslant n$. (The c_i need not be distinct.) Since there are n^{n-2} possible codes, it suffices to show that the map gives a bijection between trees on $[n]$ and codes.

Given a tree T on $[n]$ we construct its code as follows:

$T_1 := T$ has at least one leaf. Find the leaf v_1 with the smallest number, remove it, and write down the number c_1 of the (unique) vertex v_1 was adjacent to. Repeat until exactly two vertices remain. Thus, for example, v_2 is the smallest leaf of $T_2 := T - v_1$, and c_2 is the vertex of T_2 that v_2 is adjacent to. In general v_i is the smallest leaf of $T_i := T - v_1 - \cdots - v_{i-1}$ and c_i is the vertex of T_i it is adjacent to. Note that $c_1, \dots, c_{n-2}$ form the code, *not* $v_1, \dots, v_{n-2}$.

i	leaves of T_i	v_i	c_i
1	$\{2,3,5,6\}$	2	4
2	$\{3,5,6\}$	3	4
3	$\{4,5,6\}$	4	1
4	$\{5,6\}$	5	1

Prüfer Code = 4411

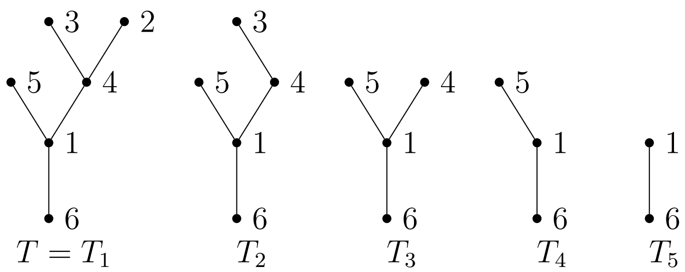

Figure 5: Example of a Prüfer code.

The following observation is key to the proof: a vertex w with degree d in T appears exactly $d-1$ times in the code $\mathbf{c}$. Indeed, we write w down in the code each time we delete a neighbour of w, i.e., each time its degree decreases. The final degree of w is always 1: either w is deleted when it is a leaf, or w is left at the end as one of the two final vertices, which are then leaves. More generally, if the degree of w in $T - v_1 - v_2 - \cdots - v_{i-1}$ is d, then w occurs $d-1$ times in $c_i, \dots, c_{n-2}$. It follows from this that

$$\begin{aligned} v_1 &= \min\big\{[n] \setminus \{c_1, \dots, c_{n-2}\}\big\} \\ v_2 &= \min\big\{[n] \setminus \{v_1, c_2 \dots, c_{n-2}\}\big\} \\ &\cdots \\ v_i &= \min\big\{[n] \setminus \{v_1, \dots, v_{i-1}, c_i, \dots, c_{n-2}\}\big\} \qquad (1 \leqslant i \leqslant n-2) \qquad (1) \end{aligned}$$

Let us write v_{n-1} and v_n (with wlog $v_{n-1} < v_n$) for the two vertices left at the end when we constructed the code, so

$$\{v_{n-1}, v_n\} = [n] \setminus \{v_1, \dots, v_{n-2}\}. \tag{2}$$

Then, since we deleted the edge v_ic_i at step i, and were left with the edge $v_{n-1}v_n$ between the final two vertices,

$$E(T) = \{v_1c_1, \dots, v_{n-2}c_{n-2}, v_{n-1}v_n\}. \tag{3}$$

The formulae above describe T, the tree that we started with, in terms of its code $\mathbf{c} = (c_1, \dots, c_{n-2})$. Does this mean that the proof is complete? No! We started by assuming that T was a tree, with code $\mathbf{c}$, and then showed that given $\mathbf{c}$, we could identify T (i.e., the map from trees to codes is injective). So for any code *coming from a tree*, there is a unique tree with that code. We still need to show that for every code $\mathbf{c}$, there is a tree with code $\mathbf{c}$ (i.e., the map from trees to codes is surjective).

The formulae above tell us where to look: if there is a tree with code $\mathbf{c}$, it *must* be as described above. So let us check.

Formally, let $\mathbf{c}$ be any possible code $(c_1, \dots, c_{n-2})$. Then we may use (1) to define $v_1, \dots, v_{n-2}$. (Each time we take the minimum of a non-empty set, which makes sense.) Also, from (1) we see that v_i is not equal to any of $v_1, \dots, v_{i-1}$. Thus $v_1, \dots, v_{n-2}$ are distinct.

Next, we define $v_{n-1} < v_n$ to be the two remaining elements of $[n]$, as in (2), so $v_1, \dots, v_n$ are distinct; they are $1, 2, \dots, n$ in some order.

Finally, we let T be the graph with vertex set $[n]$ and edge set given by (3). We need to check that T is indeed a tree, and that it has code $\mathbf{c}$. We do this step by step: first note that from our definition (1) of v_i, it is distinct from c_j, $j \geqslant i$. Thus c_j is distinct from v_i, $i \leqslant j$, so for each j, $c_j \in \{v_{j+1}, \dots, v_n\}$. Let T_i be the graph with

$$V(T_i) = \{v_i, \dots, v_n\} \quad \text{and} \quad E(T_i) = \{v_ic_i, \dots, v_{n-2}c_{n-2}, v_{n-1}v_n\}.$$

(This makes sense since the ends of the edges are distinct and lie in $V(T_i)$.) Then T_{n-1} is a tree with two vertices. Also, T_i is constructed from T_{i+1} by adding a new vertex v_i and one edge v_ic_i. So, by Lemma 2.3, T_i is a tree for $i = n-2, n-3, \dots, 2, 1$. In particular, $T = T_1$ is a tree. That the code of T is $\mathbf{c}$ is an exercise; see Problem Sheet 1. □

3 Long circuits, paths and cycles

Recall that a ***walk*** of length t in a graph G is a sequence of vertices $v_0v_1 \cdots v_t$, $t \geqslant 0$, where $v_{i-1}v_i \in E(G)$ for each $i \in [t]$. A walk is a ***path*** if all the vertices v_i are distinct. An intermediate concept is that of a ***trail***, which is a walk where the edges $v_{i-1}v_i$ are distinct, but the vertices need not be.

	Repeats ok	Edges distinct	Vertices distinct (⇒ edges distinct)
x-y	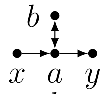 $xabay$ walk	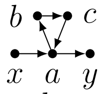 $xabcay$ trail	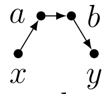 $xaby$ path
x-x	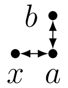 $xabax$ closed walk	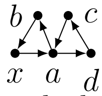 $xadcabx$ circuit	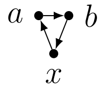 $xabx$ cycle

Figure 6: Various types of walk.

A walk $v_0v_1 \cdots v_t$ is ***closed*** if $v_0 = v_t$. A ***circuit*** is a closed trail, i.e., a walk that ends at the same vertex it starts at and does not repeat any edge. Note that walks, trails and circuits can have zero length. A ***cycle*** is a closed walk that does not repeat vertices except for $v_0 = v_t$. In the case of a cycle we insist that the length t is at least 3.

Warning. Some authors use 'circuit' to mean 'cycle' and vice versa.

An ***Euler circuit*** or ***Euler tour*** in a graph G is a circuit that uses all the edges of G, i.e., a closed walk that uses each edge of G precisely once. We call a graph ***Eulerian*** if it has an Euler circuit. (If $|G| = 1$ then G is considered to be Eulerian with the trivial Euler circuit of length 0.)

Theorem 3.1. *Let G be a connected graph. Then G is Eulerian if and only if the degree of every vertex is even.*

Proof. For the (easier) 'only if' direction, pick $v \in V(G)$. If an Euler circuit enters v k times then it leaves v k times, and so it uses $2k$ edges incident with v. Thus $d(v)$ must be even.

For the converse, we proceed by induction on $e(G)$, with the result being trivial for $e(G) = 0$. For the induction step, take any G with $e(G) > 0$ and assume the result holds for all graphs with fewer edges than G. Since G is connected, each vertex has degree at least 1. As all degrees are even, all vertices have degree at least 2. By Lemma 1.3, G contains a cycle, which is of course also a circuit. Let $C = v_1 \dots v_tv_1$ be a longest circuit in G and let $H = G - E(C)$ be the graph obtained by removing the edges of C from G. If H is an empty graph, then we are done as C is then an Euler circuit. Otherwise pick a non-trivial component H_0 of H (i.e., one with $e(H_0) > 0$). The degree of any vertex v in H_0 is even as $d_{H_0}(v) = d_H(v) = d_G(v) - d_C(v)$. Clearly $e(H_0) < e(G)$ so, by induction, H_0 has an Euler circuit $C' = u_1 \dots u_ku_1$. Now some vertex u_i must equal some vertex v_j. Indeed, consider a path from u_1 to v_1 in G. The first point at which it leaves H_0 must be via an edge xy that is not in H, so is in C, and so just before that point we have a vertex $x = u_i = v_j$ that is in both C and C'. But then $v_1 \dots v_ju_{i+1} \dots u_ku_1 \dots u_iv_{j+1} \dots v_tv_1$ is a longer circuit in G, contradicting our choice of C. Thus G has an Euler circuit. $\square$

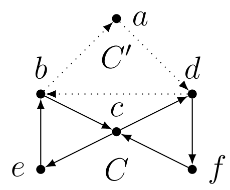

$$C = ebcdfce, \ C' = badb \implies ebadbcdfce$$

Figure 7: Combining circuits C and C' into a larger circuit.

A ***Hamilton cycle*** in a graph G is a cycle in G that contains every vertex; a graph is called ***Hamiltonian*** if it has a Hamilton cycle.

Superficially, the following two problems may seem similar: in a given graph G, is there a closed walk using every edge exactly once (Euler circuit), and is there a closed walk using every vertex exactly once (Hamilton cycle)? But it's easy to tell (using Theorem 3.1) what the answer to the first question is. The second is much harder; for those interested in complexity theory, it is an NP-complete problem. To deduce some sufficient conditions for the existence of Hamilton cycles, we first investigate how long a path in G we can find.

The ***minimum degree*** of a graph G is

$$\delta(G) = \min_{v \in V(G)} d(v),$$

the ***maximum degree*** is

$$\Delta(G) = \max_{v \in V(G)} d(v),$$

and the ***average degree*** is

$$\bar{d}(G) = \frac{1}{|G|} \sum_{v \in V(G)} d(v) = \frac{2\,e(G)}{|G|}.$$

It is not hard to see that any graph with $\delta(G) \geqslant d$ contains a path of edge length at least d: start at any v_0 and, given $v_0 \cdots v_i$ with $i < d$, choose v_{i+1} to be a neighbour of v_i not among $v_0, \ldots, v_{i-1}$. In fact, for connected graphs with many more than d vertices, we can find a path of roughly twice this length. To see this, we start with a couple of lemmas.

Lemma 3.2. *If G is a connected graph which is not Hamiltonian, then the edge length of a longest path in G is at least the length of a longest cycle.*

Proof. Let C be a longest cycle in G, with length ℓ. We have $\ell < n$ since G is not Hamiltonian, so there are vertices not on C. Since G is connected, there is at least one edge vw with $v \in V(C)$ and $w \notin V(C)$ (pick $x \in V(C)$, $y \notin V(C)$, and take the first edge of an x-y path that leaves $V(C)$). But then the edge vw and C between them contain a path of length ℓ. $\square$

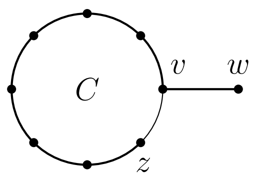

Path of length ℓ: $z \cdots vw$

Lemma 3.3. *Let G be a connected graph with $n \geqslant 3$ vertices and let $P = v_0 \cdots v_\ell$ be a longest path in G. Then either G is Hamiltonian, or $v_0 v_\ell \notin E(G)$ and $d(v_0) + d(v_\ell) \leqslant \ell$.*

Proof. Note that $n \geqslant 3$ and G connected imply there is a path of length 2, so $\ell \geqslant 2$. If G is not Hamiltonian, then by Lemma 3.2, G contains no cycle of length $> \ell$. In particular $v_0 v_\ell \notin E(G)$ as otherwise $v_0 \cdots v_\ell v_0$ would be a cycle of length $\ell + 1$. More generally, if for some $1 \leqslant i \leqslant \ell$ both $v_0 v_i$ and $v_{i-1} v_\ell$ were edges, then we would have a cycle $v_0 v_1 \cdots v_{i-1} v_\ell v_{\ell-1} \cdots v_i v_0$ of length $\ell + 1$.

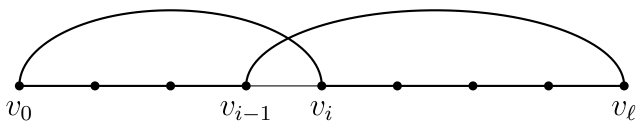

Hence $A = \{i \in [\ell] : v_0 v_i \in E(G)\}$ and $B = \{i \in [\ell] : v_{i-1} v_\ell \in E(G)\}$ are disjoint subsets of $[\ell]$. Thus, noting that all neighbours of v_0 and v_ℓ are on P (otherwise we would have a longer path),

$$d(v_0) + d(v_\ell) = |A| + |B| \leqslant |[\ell]| = \ell.$$ □

Corollary 3.4. *Let G be a connected graph with $n \geqslant 3$ vertices in which every pair v, w of non-adjacent vertices satisfies $d(v) + d(w) \geqslant k$. If $k < n$ then G contains a path of edge length k; if $k \geqslant n$ then G is Hamiltonian.*

Proof. If G has a Hamilton cycle, then it also has a path of edge length $n - 1$ (the maximum possible $k < n$) and we are done. If not, then by Lemma 3.3 there are two non-adjacent vertices u and v with $d(u) + d(v) \leqslant \ell$ where ℓ is the maximum length of a path in G. But by assumption $d(u) + d(v) \geqslant k$, so $\ell \geqslant k$ and G contains a path of length k. □

Corollary 3.5. *If G is connected, $|G| = n$, and $\delta(G) \geqslant d$, then G contains a path of edge length (at least) $\min\{2d, n-1\}$.*

Proof. Trivial for $n = 1$, 2. For $n \geqslant 3$, apply Corollary 3.4 with $k = 2d$. □

As another corollary we obtain the following result.

Theorem 3.6 (Dirac's Theorem)**.** *Let G be a graph with $n \geqslant 3$ vertices. If $\delta(G) \geqslant \frac{n}{2}$, then G contains a Hamilton cycle.*

Proof. If $\delta(G) \geqslant \frac{n}{2}$ then any two non-adjacent vertices have at least one common neighbour, so G is connected. (Or: if G is not connected then there is a component C with at most $n/2$ vertices. Then any $v \in V(C)$ has degree at most $|C| - 1 < n/2$.) Now Corollary 3.4 applies with $k = n$. □

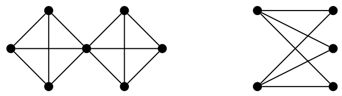

Figure 8: Two examples of non-Hamiltonian graphs with $\delta(G) = \frac{n-1}{2}$.

This result is best possible, in that we cannot replace the lower bound by $\lceil \frac{n}{2} \rceil - 1$. (For n even, consider the disjoint union of two complete graphs $K_{n/2}$. For n odd, consider two $K_{(n+1)/2}$'s sharing a single vertex, or the complete bipartite graph $K_{(n-1)/2,(n+1)/2}$.)

Corollary 3.4 of course implies a slightly stronger result than Dirac's Theorem, known as Ore's Theorem.

Theorem 3.7 (Ore's Theorem)**.** *If G has order $n \geqslant 3$, and if $d(x) + d(y) \geqslant n$ whenever $xy \notin E(G)$, then G has a Hamilton cycle.*

Lemma 3.3 also lets us relate the length of the longest path in G to the *average* degree $\bar{d}(G)$ of G or, equivalently, the number of edges in G.

Theorem 3.8. *In any graph G there exists a path with edge length at least $\bar{d}(G)$.*

Proof. We shall actually prove a stronger statement. Let $\ell(v) = \ell_G(v)$ be the edge length of the longest path starting at v. We shall show that

$$\sum_{v \in V(G)} \ell(v) \geqslant \sum_{v \in V(G)} d(v). \tag{4}$$

This clearly implies the result as then there is a path of length $\max_v \ell(v) \geqslant \frac{1}{n} \sum_v \ell(v) \geqslant \frac{1}{n} \sum_v d(v) = \bar{d}(G)$.

We use induction on n. We may assume G is connected, as otherwise we can just sum (4) over each component separately by induction. If $n \leqslant 2$ or G is Hamiltonian then the result holds as all $\ell(v) = n - 1 \geqslant d(v)$. Otherwise pick a longest path $P = v_0 v_1 \dots v_\ell$ in G and note that Lemma 3.3 implies that $d(v_0) + d(v_\ell) \leqslant \ell$. Wlog $d(v_0) \leqslant d(v_\ell)$, so $d(v_0) \leqslant \frac{\ell}{2}$. Now let $G' = G - v_0$. Then by induction $\sum_{v \neq v_0} d_{G'}(v) \leqslant \sum_{v \neq v_0} \ell_{G'}(v)$. Hence

$$\sum_v d_G(v) = 2e(G) = 2(d(v_0) + e(G')) \geqslant 2d(v_0) + \sum_{v \neq v_0} d_{G'}(v) \leqslant \ell + \sum_{v \neq v_0} \ell_{G'}(v) \leqslant \sum_v \ell_G(v),$$

as $\ell = \ell_G(v_0)$ and clearly $\ell_{G'}(v) \leqslant \ell_G(v)$ for all $v \neq v_0$. □

Note that we do not get the extra factor of 2 here as we had in Corollary 3.5, but we are assuming something only about the *average* degree, not about the degree of *every* vertex.

4 Vertex colourings

A ***proper vertex colouring*** (or simply a ***colouring***) of a graph G is an assignment of a colour to each vertex such that adjacent vertices receive different colours. The least number of colours in such a colouring is the ***chromatic number*** $\chi(G)$. For example $\chi(K_n) = n$, $\chi(E_n) = 1$, $\chi(C_4) = 2$ and $\chi(C_5) = 3$. In fact, any *even* cycle (cycle of even length) has chromatic number 2, and any *odd* cycle has chromatic number 3.

Often we use positive integers as the colours: a (***proper***) ***k-colouring*** of G is a function[6] $f\colon V(G) \to [k]$ so that $f(u) \neq f(v)$ whenever $uv \in E(G)$. Note that not all colours are necessarily used, i.e., we do *not* assume f is surjective. We say that G is ***k-colourable*** if it has a k-colouring, so $\chi(G)$ is the least k for which G is k-colourable.

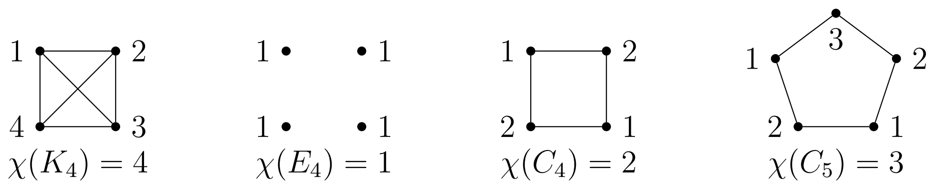

Figure 9: Chromatic numbers of some small graphs and example colourings.

As an example of an application, suppose we have to schedule exams, where each exam takes one time period. Construct a graph G with a vertex for each exam and an edge uv whenever one or more students need to take both exams u and v. Then a feasible exam schedule corresponds to a colouring of G, and the least number of time periods possible to schedule all the exams without conflicts is $\chi(G)$.

Certainly, if H is a subgraph of G, then $\chi(H) \leqslant \chi(G)$ (just restrict a $\chi(G)$-colouring of G to $V(H)$). Clearly, a disconnected graph is k-colourable if and only if all of its components are, so the chromatic number of G is the maximum of the chromatic numbers of its components. In fact, we can extend this to graphs overlapping in certain ways.

The ***union*** of two graphs G_1 and G_2 is the graph with vertex set $V(G_1) \cup V(G_2)$ and edge set $E(G_1) \cup E(G_2)$.

Lemma 4.1. *Let G_1 and G_2 be graphs with $V(G_1) \cap V(G_2) = W$ such that $G_1[W]$ and $G_2[W]$ are complete. Then $\chi(G_1 \cup G_2) = \max\{\chi(G_1), \chi(G_2)\}$.*

Proof. As each G_i is a subgraph of $G_1 \cup G_2$, $\chi(G_1 \cup G_2) \geqslant \max\{\chi(G_1), \chi(G_2)\}$. Now let $k = \max\{\chi(G_1), \chi(G_2)\}$ so that both G_1 and G_2 are k-colourable; we must show that $G_1 \cup G_2$ is also k-colourable. Let c_i be a k-colouring of G_i, and let $W = \{w_1, \dots, w_r\}$. Since c_1 assigns distinct colours to $w_1, \dots, w_r$, we may permute the colours (i.e., keep fixed which sets of vertices get the same colour, but assign different colours to these sets) to obtain a new k-colouring $\tilde{c}_1$ of G_1 in which $w_1, \dots, w_r$ get colours $1, 2, \dots, r$ in this order. Do the same for G_2, and then combine the colourings $\tilde{c}_1$ and $\tilde{c}_2$, which agree on W,

[6]We can also consider a k-colouring as a homomorphism of G to K_k, see footnote [1] on page 2.

to obtain a colouring of $G_1 \cup G_2$. As the colouring is proper on both G_1 and G_2 and there are no edges from $G_1 - W$ to $G_2 - W$, this is a proper k-colouring of $G_1 \cup G_2$. □

A ***cut vertex*** v in a connected graph G is a vertex such that $G - v$ is disconnected. (In a general graph, it's a vertex whose deletion increases the number of components of the graph.) Lemma 4.1 may be applied in particular to any graph G with a cut vertex.

Now we consider what happens when G is k-colourable for small k.

A graph G has $\chi(G) = 1$ if and only if G has no edges, so is an empty graph.

A graph $G = (V, E)$ is ***bipartite*** if V can be split into disjoint sets X and Y such that $E \subseteq \{xy : x \in X, y \in Y\}$. (We allow one of X or Y to be empty, so K_1 is bipartite.) The ***complete bipartite graph*** $K_{m,n}$ has $V = \{a_1, \ldots, a_m, b_1, \ldots, b_n\}$ and $E = \{a_i b_j : i = 1, \ldots, m,\ j = 1, \ldots, n\}$. The connection to colouring is that $\chi(G) \leqslant 2$ if and only if G is bipartite: consider $X = \{v : c(v) = 1\}$ and $Y = \{v : c(v) = 2\}$.

Deciding whether a (connected) graph is 2-colourable (i.e., bipartite) is very easy: start somewhere with one colour (it doesn't matter which) and work outwards from there – having coloured a vertex, the colours of its neighbours are forced, and we either get stuck or we don't. The next simple lemma gives a criterion.

Lemma 4.2. *A graph G is 2-colourable (bipartite) if and only if it contains no odd cycles.*

Proof. If G is 2-colourable then, in any 2-colouring, the colours around any cycle C in G alternate, implying that C has even length.

For the reverse implication we use induction on $|G|$; the base case $|G| = 1$ is trivial. For the induction step let G be a graph with $n \geqslant 2$ vertices with no odd cycle. We may assume that G is connected (else colour its components). It follows that there is (at least) one vertex v such that $G - v$ is connected (take v to be a leaf of a spanning tree of G). By induction we may 2-colour $G - v$. If all neighbours of v have the same colour in this colouring, then we may extend the colouring to G by using the opposite colour for v. So we may suppose that v has neighbours x and y with different colours. As $G - v$ is connected, there is a path P in $G - v$ from x to y. Along this path the colours alternate, so P has odd length, and together with vx and vy forms an odd cycle in G, contradicting our assumption. □

In general, finding the chromatic number of a graph is very hard; even the question 'is $\chi(G) \leqslant 3$' is hard (NP-complete). However, we can give some general bounds on $\chi(G)$.

A ***copy of*** K_k (i.e., a subgraph isomorphic to K_k) in a graph G is called a ***complete subgraph*** or a ***clique***. The ***clique number*** $\omega(G)$ of G is the largest k such that G contains a copy of K_k. A set S of vertices is an ***independent set*** (or ***stable set***) in G if $G[S]$ has no edges, i.e., no two vertices of S are adjacent in G. Thus, a (proper) colouring of G corresponds to a partition of $V(G)$ into independent sets. The ***independence number*** $\alpha(G)$ is the maximum size of an independent set in G. For example, $\omega(C_5) = \alpha(C_5) = 2$. Note that $\alpha(G) = \omega(\overline{G})$.

Lemma 4.3. $\chi(G) \geqslant \max\left\{\omega(G), \frac{|G|}{\alpha(G)}\right\}$.

Proof. All vertices in a clique must get different colours in any colouring, so $\chi(G) \geqslant \omega(G)$. Also, since the vertices of each colour form an independent set, each colour is used on at most $\alpha(G)$ vertices, so we need at least $\frac{|G|}{\alpha(G)}$ colours. □

It should be emphasized that the lower bound in Lemma 4.3 is far from tight in general. See, for example, Problem Sheet 2, question 13.

Given an ordering $v_1, \dots, v_n$ of the vertices of a graph G, the ***greedy algorithm*** constructs a (proper) colouring of G with positive integers by colouring the vertices in order: each vertex receives the least colour not already assigned to one of its neighbours.

Lemma 4.4. $\chi(G) \leqslant \Delta(G) + 1$.

Proof. Take any ordering of the vertices and apply the greedy algorithm: each vertex has at most $\Delta(G)$ forbidden colours, and so will get a colour from $\{1, 2, \dots, \Delta(G) + 1\}$. □

This bound is tight in some cases: in particular if G is complete or an odd cycle. But usually we can do better; we start with two simple lemmas.

Lemma 4.5. *Let G be a connected graph with n vertices and let $v \in V(G)$. Then we may order the vertices as $v_1, \dots, v_{n-1}, v_n = v$ so that each vertex other than v has at least one neighbour coming after it.*

Proof. See Problem Sheet 2. □

Lemma 4.6. *Let G be a connected graph with $\Delta(G) \leqslant d$ and $\delta(G) < d$. Then $\chi(G) \leqslant d$.*

Proof. Pick a vertex v with $d(v) < d$, take an ordering as in Lemma 4.5, and apply the greedy algorithm: each vertex has at most $d - 1$ forbidden colours. □

(We won't need this lemma in the proof that follows, but it encapsulates one key idea of that proof.)

Dealing with the d-regular case will be significantly harder, though we now have the tools we need.

Theorem 4.7 (Brooks' Theorem)**.** *Let G be a connected graph. If G is neither an odd cycle nor a complete graph then $\chi(G) \leqslant \Delta(G)$.*

Proof. For $\Delta(G) \leqslant 2$ the result is easy as the graph is either a cycle or a path, so suppose $\Delta(G) \geqslant 3$. It is convenient to restate the result slightly as follows:

Theorem 4.7′. *Let $d \geqslant 3$ and let G be any graph with $\Delta(G) \leqslant d$ which does not contain a copy of K_{d+1}. Then $\chi(G) \leqslant d$.*

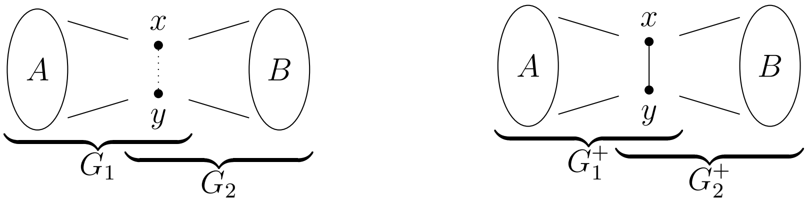

Figure 10: Decomposition of G into G_1, G_2, overlapping in $\{x, y\}$, and definition of G_1^+, G_2^+.

Since a connected graph with maximum degree d that contains a copy of K_{d+1} must be K_{d+1}, this restatement (applied with $d = \Delta(G)$) implies Brooks' theorem. We prove the restated result by induction on $n = |G|$.

If G is disconnected we are done by induction applied to each component, so suppose G is connected. If G has a cut vertex v, then we may write $G = G_1 \cup G_2$ where G_1 and G_2 overlap precisely in v and $|G_1|$, $|G_2| < n$. By induction $\chi(G_1) \leqslant d$ and $\chi(G_2) \leqslant d$ so, by Lemma 4.1, $\chi(G) \leqslant d$. Hence we may assume G has no cut vertex.

Let v be a vertex of G with degree d. (If there is none, then $\chi(G) \leqslant \Delta(G) + 1 \leqslant d$ by Lemma 4.4.) Since G contains no K_{d+1}, we can find neighbours x and y of v such that $xy \notin E(G)$. Suppose that $G - x - y$ is connected. Then we may order the vertices of $G - x - y$ as in Lemma 4.5, ending at v. Putting x and y at the beginning of this ordering, we obtain an ordering of the vertices of G in which each vertex apart from v precedes at least one of its neighbours. Moreover, the greedy algorithm gives x and y the same colour, so when it comes to assign a colour to v, at most $d - 1$ colours are forbidden. Therefore the greedy algorithm uses at most d colours with this ordering.

Suppose instead that $G - x - y$ is not connected. Then $V(G) \setminus \{x, y\}$ can be partitioned into non-empty sets A and B with[7] $e(A, B) = 0$. Let $G_1 = G[A \cup \{x, y\}]$ and $G_2 = G[B \cup \{x, y\}]$, so G consists of its subgraphs G_1 and G_2 overlapping in the non-adjacent vertices x and y (see Figure 10). Both x and y must have neighbours in each of A and B (if say x had no neighbours in A then $G - y$ would be disconnected, so G would have a cut vertex). Hence x and y have degree at most $d-1$ in G_1 and in G_2. Let $G_j^+ = G_j + xy$. Then $\Delta(G_j^+) \leqslant d$. If neither G_1^+ nor G_2^+ contains K_{d+1} then by Lemma 4.1 and induction

$$\chi(G) \leqslant \chi(G + xy) = \chi(G_1^+ \cup G_2^+) = \max\left\{\chi(G_1^+), \chi(G_2^+)\right\} \leqslant d.$$

So suppose that one, say G_1^+, contains a copy of K_{d+1}. Note that this copy must include x and y, since $G_1 \subseteq G$ contains no K_{d+1}. Since G is connected, in fact G_1^+ is isomorphic to K_{d+1}. Since x and y have degree $d - (d-1) = 1$ in G_2, we can d-colour G_2 with x and y having the same colour. Indeed, by induction we can d-colour $G[B]$, and each of x and y has only one colour ruled out, so since $d \geqslant 3$ we can choose the same colour for both. But now we can extend this colouring to G_1, and hence to all of G. □

[7] $e(A, B)$ is the number of edges ab of G with $a \in A$ and $b \in B$.

The chromatic polynomial

Given a graph G, for $k = 1, 2, \dots$, let $P_G(k)$ be the number of (proper) k-colourings of G, i.e., colourings with $[k]$ as the set of available colours (not all colours have to be used). For example, $P_{K_n}(k) = k(k-1)(k-2)\cdots(k-n+1)$ and, trivially, $P_{E_n}(k) = k^n$. It turns out that with k fixed we can calculate $P_G(k)$ inductively, using two operations on graphs.

If $e = uv$ is an edge in a graph G, we let G/e denote the graph obtained by ***contracting*** e; that is, G/e is obtained from G by deleting the vertices u and v and adding a new vertex adjacent to each vertex in $(N(u) \cup N(v)) \setminus \{u, v\}$. In other words, we 'merge' u and v into a single vertex and each vertex $z \neq u, v$ is joined to this merged vertex if and only if either zu or zv was originally an edge of G.[8] (There is a slightly different notion of contraction for multigraphs where we keep all edges from z to u or v and thus may end up with multiple edges to the merged vertex.)

Lemma 4.8. *For each edge e of G and $k \geqslant 0$, $P_{G-e}(k) = P_G(k) + P_{G/e}(k)$.*

Proof. Suppose that $e = uv$. Let S be the set of k-colourings of $G - e$, let $S_1 = \{c \in S : c(u) \neq c(v)\}$ and let $S_2 = \{c \in S : c(u) = c(v)\}$. Clearly $|S| = |S_1| + |S_2|$. Also, $P_{G-e}(k) = |S|$, $P_G(k) = |S_1|$ (since these are exactly the colourings of G), and $P_{G/e}(k) = |S_2|$ (since these correspond to the colourings of G/e, taking the common colour of u and v for the new vertex and vice versa). □

Theorem 4.9. *For every graph G there is a unique polynomial $p_G(x) \in \mathbb{Z}[x]$, the* ***chromatic polynomial*** *of G, such that*

$$p_G(k) = P_G(k) \quad \text{for each } k = 0, 1, 2, \dots \tag{5}$$

Moreover, for every edge e of G we have $p_G(x) = p_{G-e}(x) - p_{G/e}(x)$.

Proof. Uniqueness is immediate since two polynomials that agree on all non-negative integers must be the same. For existence we use induction on $e(G)$. For the base case $e(G) = 0$, $G \cong E_n$ for some n, so $P_G(k) = k^n$ for every k and the polynomial x^n has the required properties.

For the inductive step, pick any edge e of G and note that $G - e$ and G/e have fewer edges than G. So by induction there are polynomials p_{G-e} and $p_{G/e}$ satisfying (5) for the corresponding graphs. Consider $p_G = p_{G-e} - p_{G/e}$; this is a polynomial. By Lemma 4.8, for every positive integer k we have $p_G(k) = p_{G-e}(k) - p_{G/e}(k) = P_{G-e}(k) - P_{G/e}(k) = P_G(k)$, as required. The final statement follows immediately: we have shown that there is a polynomial p_G satisfying (5), and know that it is unique. We have also shown that for any edge e, $p_{G-e} - p_{G/e}$ is such a polynomial, so $p_G = p_{G-e} - p_{G/e}$. □

[8] One can think of this as taking a quotient $G/\sim$ where $\sim$ is an equivalence relation on $V(G)$. We collapse each equivalence class down to a single vertex and join two equivalence classes if there was any edge between them in the original graph. G/e is then just $G/\sim$ where $\sim$ just equates the end vertices of e.

From now on we write $p_G(k)$ for the number of k-colourings of G, since this number is an evaluation of the chromatic polynomial. In general, identities for numbers of k-colourings valid for all k give polynomial identities. As a simple example, if G has components $G_1, \ldots, G_j$ then $p_G(x) = p_{G_1}(x) \cdots p_{G_j}(x)$; this is valid for each $x \in \mathbb{N}$, and both sides are polynomials. We also note that since any graph on $n \geqslant 1$ vertices has no 0-colourings, $p_G(0) = 0$ and so $x \mid p_G(x)$.

Theorem 4.10. *Let G be a graph with n vertices and m edges. Then*

$$p_G(x) = \sum_{i=0}^{n-1} (-1)^i a_i x^{n-i} = a_0 x^n - a_1 x^{n-1} + \cdots + (-1)^{n-1} a_{n-1} x,$$

where $a_0 = 1$, $a_1 = m$ and $a_i \geqslant 0$ for all i.

Proof. We argue by induction on m. For $m = 0$ we have $G \cong E_n$, so $p_G(x) = x^n$, and we are done. For $m > 0$, pick an edge e of G. Then $|G - e| = n$ and $e(G - e) = m - 1$, so by the induction hypothesis,

$$p_{G-e}(x) = x^n - (m-1)x^{n-1} + \sum_{i=2}^{n-1} (-1)^i a_i x^{n-i},$$

where each $a_i \geqslant 0$. Also $|G/e| = n - 1$ and $e(G/e) \leqslant m - 1$ so, by induction,

$$p_{G/e}(x) = x^{n-1} + \sum_{j=1}^{n-2} (-1)^j b_j x^{n-1-j} = x^{n-1} + \sum_{i=2}^{n-1} (-1)^{i-1} b_{i-1} x^{n-i},$$

where each $b_j \geqslant 0$. By the last part of Theorem 4.9,

$$p_G(x) = p_{G-e}(x) - p_{G/e}(x) = x^n - m x^{n-1} + \sum_{i=2}^{n-1} (-1)^i (a_i + b_{i-1}) x^{n-i},$$

and $a_i + b_{i-1} \geqslant 0$ for each i. □

5 Edge colourings

A function $f\colon E(G) \to [k]$ is a (***proper***) k-***edge-colouring*** of G if edges that meet (i.e., share an end vertex) always receive distinct colours. The ***edge-chromatic number*** $\chi'(G)$ (also called the ***chromatic index***) is the smallest k such that G has a k-edge-colouring.

Proposition 5.1. *If $e(G) > 0$, then $\Delta(G) \leqslant \chi'(G) \leqslant 2\Delta(G) - 1$.*

Proof. Since the edges incident with a given vertex must get different colours we have $\chi'(G) \geqslant \Delta(G)$. For the upper bound, list the edges in any order and apply the greedy algorithm to colour the edges. When we come to colour an edge uv, the number of colours unavailable is at most $d(u) - 1 + d(v) - 1 \leqslant 2\Delta(G) - 2$. □

Amazingly, given the maximum degree Δ of a graph, there are only two possible values for the edge-chromatic number, Δ and $\Delta + 1$. The proof involves a 'colour chasing' argument. (More precisely, the proof combines two simple colour chasing arguments in a clever way.)

Theorem 5.2 (Vizing's Theorem). *$\chi'(G) = \Delta(G)$ or $\Delta(G) + 1$.*

Proof. We need to prove that $\chi'(G) \leqslant \Delta(G) + 1$. We argue by induction on $e(G)$. The result is trivial if $e(G) = 0$, so let G be a graph with $e(G) > 0$. Let xy_1 be any edge of G, and assume (applying the induction hypothesis to $G - xy_1$) that we have coloured every edge of G except xy_1 with colours $1, \ldots, \Delta(G - xy_1) + 1 \leqslant \Delta(G) + 1$. Our aim is to show that we can recolour so that we can colour the edge xy_1 as well.

For any vertex v, since $d(v) < \Delta(G) + 1$, there is at least one colour *missing* at v, i.e., not appearing on any edges incident with v. Let t_1 be a colour missing at y_1.

If colour t_1 is missing at x, colour xy_1 with t_1 and we are done. If not, there is an edge xy_2 with colour t_1, and some colour t_2 ($\neq t_1$) is missing at y_2. If t_2 is missing at x, recolour xy_2 with t_2 and colour xy_1 with t_1, and we are done. Otherwise, there is an edge xy_3 with colour t_2, and there is a colour t_3 missing at y_3. If t_3 is missing at x we can recolour as above; otherwise there is an edge xy with colour t_3. This could be a 'new' edge, but it could instead be xy_2.

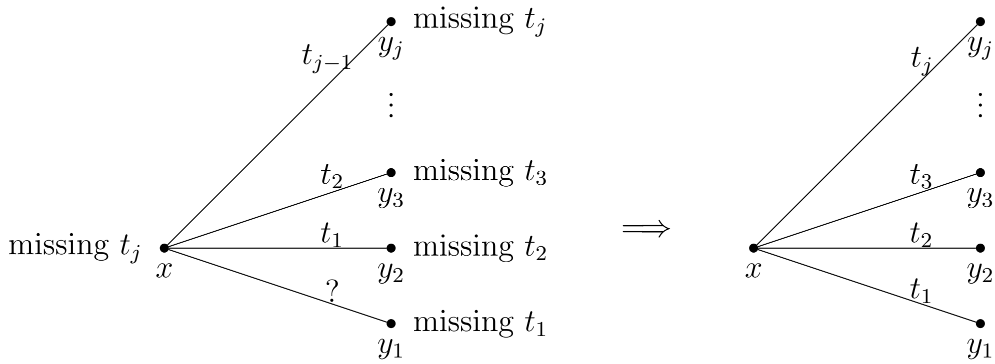

Figure 11: Recolouring a fan when some t_j is missing at x.

In general, suppose that we have distinct neighbours $y_1, \ldots, y_j$ of x and distinct colours $t_1, \ldots, t_{j-1}$ such that t_i is missing at y_i for each $i = 1, \ldots, j-1$, the edge xy_1 is uncoloured, and xy_i has colour t_{i-1} for each $i = 2, \ldots, j$. We call this a ***fan*** of size j. Note that there is a fan of size 1, consisting of the uncoloured edge xy_1. Let t_j be a colour missing at y_j.

If t_j is missing at x, then recolour xy_i with t_i for each $i = 1, \ldots, j$, and we are done. (There are no conflicts at the y_i since t_i was missing at y_i, and no conflicts at x since $t_1, \ldots, t_{j-1}$ were already present on $xy_2, \ldots, xy_j$ and t_j was missing.) Otherwise there is an edge xy with colour t_j. Note that if $t_j \notin \{t_1, \ldots, t_{j-1}\}$ then $y \notin \{y_1, \ldots, y_j\}$ (since xy_1 is uncoloured and the xy_i is coloured t_{i-1} for $i = 2, \ldots, j$. Then y is a 'new' vertex, so let $y_{j+1} = y$ – we now have a fan of size $j + 1$.

The process must terminate (consider a fan of maximal size), so eventually we reach a case where the missing colour t_j at y_j is the same as some t_i, $i < j$. Let s be a colour missing at x so, by construction, we may assume $s, t_1, \dots, t_{j-1}$ are distinct and $t_j = t_i = t$, say.

Let H be the spanning subgraph of G consisting of all edges coloured s or t. Then we can swap s and t within any component of H without causing conflicts. Since $\Delta(H) \leqslant 2$, H consists of paths and cycles. At each of x, y_i and y_j, at least one of s and t is missing, so each of these vertices has degree $\leqslant 1$ in H. Hence the components of H containing x, y_i and y_j are paths (possibly of length 0), with each of x, y_i and y_j being an end of one of these paths. Since a path has at most two ends, at least one of y_i or y_j is not connected to x in H. Thus if we swap the colors s and t in the component of this y_i or y_j in H, we obtain a colouring in which x is still missing s, and all xy_k keep the same color (as no edge incident to x is recoloured), but now either y_i or y_j is also missing colour s.

If y_i and x are not in the same component of H and we perform this recolouring. Note that each y_k, $k < i$, is still missing t_k as $t_k \notin \{s, t\}$ and only edges coloured s or t have been recoloured. Thus we obtain a fan of size i with s missing at both x and y_i. We can then recolour as above to obtain a proper edge colouring of G.

On the other hand, if y_i and x are in the same component of H, y_j is in a different component and we can recolor the component of H containing y_j by swapping colours s and t. Again, y_k, $k < j$, is still missing t_k as either $t_k \notin \{s, t\}$ or $k = i$ and y_i is in a component of H that was not recoloured. Thus we obtain a fan of size j with s missing at both x and y_j. We can then recolour as above to obtain a proper edge colouring of G. □

Proper edge colourings of any graph G correspond exactly to proper vertex colourings of the ***line graph*** $L(G)$ of G. This is (as it must be for the previous sentence to be true) the graph with a vertex for each edge of G in which two vertices are adjacent if and only if the corresponding edges of G meet. So in a sense, edge colouring is a special case of vertex colouring, though this viewpoint is not likely to be helpful in proving results such as Vizing's Theorem.

6 Planar Graphs

The graph K_4 may be drawn in the plane with no edges crossing. What about $K_{3,3}$ (Dudeney's problem), or K_5? We first need to define clearly what we mean by such a drawing of a graph.

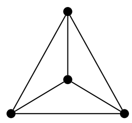

Figure 12: A plane drawing of K_4.

A ***simple curve*** in the plane is an image of a continuous injective function $\gamma\colon [0,1] \to \mathbb{R}^2$.

Its ***endpoints*** are $\gamma(0)$ and $\gamma(1)$. A ***simple closed curve*** is the image of a continuous map $\gamma\colon [0,1] \to \mathbb{R}^2$ that is injective except that $\gamma(0) = \gamma(1)$. A curve is ***polygonal*** if it is formed from a *finite* number of straight-line segments, i.e., γ is piecewise linear.

A ***drawing*** of a graph $G = (V, E)$ in the plane is a representation consisting of distinct points x_v for the vertices $v \in V$, and simple polygonal curves γ_{uv} for the edges $uv \in E$, such that γ_{uv} has x_u and x_v as its endpoints, and the interiors of the curves (i.e., the curves without their endpoints) are disjoint from each other and from the x_v. In other words, the points and curves meet only 'as they should' according to the incidence relation of the graph. We often identify a vertex or edge with its image in $\mathbb{R}^2$.

In fact, the usual definition allows the edges to be drawn as simple curves that need not be polygonal; it is an exercise in analysis (that we will not do) to show that the two definitions coincide: a general drawing can be 'converted' into a polygonal drawing.[9]

A graph together with a drawing in the plane is often called a ***plane*** graph. We tend to use the notation G for a plane graph without explicitly indicating the drawing. A graph is ***planar*** if it has a drawing in the plane.

Given a plane graph, if we omit from the plane the points corresponding to the vertices and edges, what remains falls into open connected components, the ***faces***, exactly one of which is unbounded. To study plane (and planar) graphs we need surprisingly little topology. We start with a simple and seemingly trivial observation.

Lemma 6.1. *An edge lies in the boundary of a face if and only if any one of its interior points does. In particular, an edge can lie in the boundary of at most two faces, and the boundary of a face is just the union of some number of edges and vertices.*

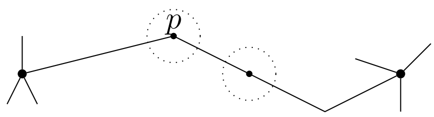

Proof. Suppose the edge e is drawn using $\gamma_e\colon [0,1] \to \mathbb{R}^2$. For any point $p = \gamma_e(t)$ in the interior of e (i.e., $t \in (0,1)$), there is some $\delta > 0$ such that the only parts of the drawing of G intersecting a ball $B_\delta(p)$ of radius δ about p are the (one or two) line segments of γ_e meeting p. Indeed, the union of all other line segments and points of the drawing of G forms a closed set not containing p, so is at a positive distance from p. But now it is clear that the points of γ_e near p are in the boundaries of exactly the same faces as p. Thus every set $U_{\mathcal{F}} = \{t \in (0,1) : \mathcal{F}$ is the set of faces whose boundary contains $\phi_e(t)\}$ is open. As these partition $(0,1)$ and $(0,1)$ is connected, all interior points of e meet exactly the same set of faces. Moreover it is clear that they can meet at most two faces as $B_\delta(p) \setminus \gamma_e([0,1])$ has only two components. Also, if the boundary of a face intersects any edge, it must contain the entire edge (including endpoints as these are arbitrarily close to interior points of e that are arbitrarily close to the face), so the boundary must be formed as a union of edges and possibly some extra (isolated) vertices. □

[9] In fact drawings of simple graphs can be converted to ones where all the edges are single straight line segments – this is Fáry's Theorem.

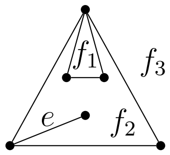

Figure 13: A plane graph with three faces. Note that e is only in the boundary of face f_2.

This allows us to identify the ***boundary*** ∂f of a face f with the subgraph of G consisting of all the edges and vertices of G that lie in the (topological) boundary of f.

Note that it is possible for an edge to be in the boundary of a single face. As a tool to show certain faces are distinct, we shall use the Jordan Curve Theorem (or at least the special case when γ is assumed to be polygonal).

Theorem 6.2 (Jordan Curve Theorem)**.** *Let γ be a simple closed curve in $\mathbb{R}^2$. Then $\mathbb{R}^2 \setminus \gamma$ consists of exactly two connected components. One of these components is bounded (the interior) and the other is unbounded (the exterior), and the curve γ is the boundary of each component.*

It is not hard to prove this in the case when γ is polygonal. The proof of Lemma 6.1 implies there are *at most* two components (as each component must have some point of γ on its boundary), and their boundary is exactly γ. To show that there are *at least* two components one can, for example, use the idea of the winding number $I(\gamma, p)$ defined in the Part A *Metric Spaces and Complex Analysis* course. This is a continuous integer-valued function on $\mathbb{R}^2 \setminus \gamma$, and so is constant on the components of $\mathbb{R}^2 \setminus \gamma$. But it is easy to see that it changes by ± 1 when p crosses one of the line segments of γ. Hence there must be at least two distinct components of $\mathbb{R}^2 \setminus \gamma$.

There is one other result proved in *Metric Spaces and Complex Analysis* that is often useful here, namely that any two points in a connected open subset $U \subseteq \mathbb{R}^2$ can be joined by a simple polygonal curve in U. This easily implies the following.

Lemma 6.3. *Any two vertices x and y in the boundary of the same face f of a plane graph G can be joined by a simple polygonal curve path with interior in f. In particular, if x and y are also non-adjacent then $G + xy$ remains planar.*

Proof. For any vertex x, there is some $\delta > 0$ such that the only parts of the drawing of G meeting $B_\delta(x)$ are x and the initial line segments from x of the edges incident with x. Each face f with x in the boundary includes one or more of the resulting sectors of the disk $B_\delta(x)$ between these line segments. One can thus easily construct short line segments ℓ_x and ℓ_y from x and y into f that avoids the rest of the drawing of G. Join their ends in f by a polygonal path γ. Now γ could potentially cross ℓ_x and ℓ_y, but then just truncate ℓ_x and ℓ_y at the first points $\gamma(t_0)$ and $\gamma(t_1)$ where they hit γ (note $\{t : \gamma(t) \in \ell_x\}$ is a closed set) and combine with the sub-curve $\gamma_{|[t_0,t_1]}$ to form a polygonal x-y curve. □

We now determine when an edge e in a plane drawing is in the boundary of one face, and when it is in he boundary of two faces. A ***bridge*** in a (simple, abstract) graph G is an edge whose deletion would disconnect the component of G that it lies in. Note that e is a bridge if and only if e is not in any cycle.

Lemma 6.4. *Let e be an edge of a plane graph G. Then e is in the boundary of a single face if and only if e is a bridge, otherwise e lies in the boundary of precisely two distinct faces. Moreover, if G is not a forest, then the boundary of every face contains a cycle.*

Proof. Suppose e is not a bridge. Then e lies in some cycle C, and the drawing of C is a simple closed (polygonal) curve in the plane which separates the plane into its inside and outside by the Jordan curve theorem. The inside must contain a face of G whose boundary meets e, and the outside contains a different face that meets e (as C is the boundary of both the inside and the outside of C). Thus e lies in the boundary of two faces.

In the other direction, let f and f' be distinct faces with e in the boundary. Let H be the subgraph of ∂f consisting of all edges h such that h is in the boundary of f and some other face, i.e., h separates f from non-f. Note that $e \in E(H)$. Going around a small circle centred at a vertex v, so that we cross each of the $d_H(v)$ edges incident with v exactly once and do not cross any other edges, we enter and leave f the same number of times. Thus $d_H(v)$ is even. It is easy to check (see Problem Sheet 3) that in a graph with all degrees even, every edge is in a cycle. So e is in a cycle in H, and hence in G.

For the last part, if G is not a forest, then it contains a cycle and so has more than one face. For any face f define H as above; then H contains a cycle which consists of edges in ∂f. □

Remark. The above proof also shows that an edge of ∂f is a bridge in G if and only if it is a bridge in ∂f. Indeed, the subgraph H contains just the non-bridges of G in ∂f, and each lies in a cycle in $H \subseteq \partial f$, and hence is not a bridge in ∂f.

Theorem 6.5 (Euler's Formula)**.** *Let $G = (V, E)$ be a connected plane graph and let F be the set of its faces. Then*

$$|V| - |E| + |F| = 2.$$

More generally, if G is a planar graph with c components then

$$|V| - |E| + |F| = 1 + c.$$

Proof. By induction on $|F|$. If $|F| = 1$ then G does not contain a cycle, so it is a forest, and, by summing over component trees, $|E| = |V| - c$ when G has c components.

Suppose now that $|F| \geqslant 2$ and the result holds for smaller values of $|F|$. Pick an edge e in the boundary of two faces. By Lemma 6.4, $G - e$ has the same number of components as G. When we delete e from the drawing, two faces join up to form a new face, while all other faces remain unchanged. So by induction $|V| - (|E| - 1) + (|F| - 1) = 1 + c$, and hence $|V| - |E| + |F| = 1 + c$ as required. □

Corollary 6.6. *Let G be a planar graph with $n \geqslant 3$ vertices. Then $e(G) \leqslant 3n - 6$.*

Proof. We may assume G is contains a cycle as otherwise $e(G) \leqslant n-1 \leqslant 3n-6$ for $n \geqslant 3$. Let $m = e(G)$. Consider a drawing of G in the plane, with faces $f_1, \dots, f_f$. Let $d(f_i)$ be the number of edges in the boundary of f_i *counting any bridges twice.* Since each non-bridge is in the boundary of two faces and each bridge of only one, we have $\sum_i d(f_i) = 2m$. By the last part of Lemma 6.4, $d(f_i) \geqslant 3$ for every face f_i, so $3f \leqslant 2m$. Hence $2 \leqslant 1 + c = n - m + f \leqslant n - m + 2m/3 = n - m/3$ and the result follows. □

We now see that K_5 is not planar, since $e(K_5) = 10 > 9 = 3 \times 5 - 6$. It is an exercise to show that any triangle-free planar graph with $n \geqslant 3$ vertices has at most $2n - 4$ edges (see Problem Sheet 3); this shows that $K_{3,3}$ is not planar.

The following results relate some properties of the drawings of a plane graph to abstract properties of the graph.

Lemma 6.7. *A plane graph G is connected iff the boundary ∂f of every face f of G is connected.*

Proof. Suppose G is disconnected and consider a shortest straight line segment in $\mathbb{R}^2$ joining distinct components of G. (Such a line exists by compactness.) This must traverse a face of G whose boundary meets edges or vertices (at the endpoints of the line segement) that are in distinct components of G. Thus ∂f is disconnected.

Conversely, if some ∂f is disconnected, we can join two components of ∂f by an edge e keeping the graph planar (Lemma 6.3). As e is a bridge in the graph $\partial f + e$, it does not disconnect f in G (so for example one can find a closed curve in f joining opposite sides of the midpoint of e). But then e is a bridge in $G + e$, so G was disconnected. □

Lemma 6.8. *A vertex v in a plane graph is a cut vertex if and only if it is in the boundary of strictly fewer than $d(v)$ faces.*

Proof. Removing v and its incident edges merges all faces with v on the boundary and leaves all other faces unaffected. Thus if v was incident to t faces, then $|V| - |E| + |F|$ increases by $(-1) + d(v) - (t-1) = d(v) - t$ on removing v. Hence by Theorem 6.5 the number of components increases by $d(v) - t$. Thus v is a cut vertex iff $t < d(v)$. □

Lemma 6.9. *The following are equivalent.*

(a) *G is connected with no cut vertex and $|G| \geqslant 3$.*

(b) *The boundary of every face is a cycle.*

Remark. If condition (a) holds we call G ***2-connected***, see later.

Proof. If G is disconnected, then some ∂f is disconnected, so not a cycle. Suppose now that G is connected with cut vertex v. Then two of the segments of a small disk around v belong to the same face, whose boundary is now clearly not a cycle (either because $d_{\partial f}(v) \geqslant 4$ or because ∂f contains a bridge).

Conversely, suppose G has a face f for which ∂f is not a cycle. If ∂f is disconnected, then so is G. Thus ∂f must contain a vertex v of degree $\neq 2$. If $d_{\partial f}(v) = 0$ then v is an isolated vertex as any edge incident to it would be in ∂f. Thus G is not connected. If $d_{\partial f}(v) = 1$, then the edge incident to v is a bridge. Thus as $|G| \geqslant 3$ and G is connected, the other end of the bridge is a cut vertex. Otherwise $d_{\partial f}(v) \geqslant 3$ and f must occupy at least two segments of a small disk around v. Thus v is incident to at most $d(v) - 1$ faces and so is a cut vertex by Lemma 6.8. □

Dual graphs

Slightly informally, a ***multigraph*** consists of a set V of vertices and a set E of edges, where each $e \in E$ either joins some vertex v to itself (such an edge is called a ***loop***) or joins some (unordered) pair $\{u, v\}$ of vertices. There may be several edges joining the same pair of vertices, and there may be several loops at a given vertex v. (Formally, we may define a multigraph as a triple (V, E, ϕ), where V and E are finite sets, and $\phi\colon E \to \binom{V}{2} \cup \binom{V}{1}$ encodes the ends (or end for a loop) of an edge $e \in E$.) It is clear how to extend the definition of a drawing in the plane to multigraphs; for example, a loop at v is drawn as a (polygonal) simple closed curve from x_v to itself meeting the other edges only at x_v.

If G is a plane (multi-)graph then G has a ***dual*** G^* obtained as follows: take one vertex f^* for each face f of G, and one edge e^* for each edge e of G, joining the vertices f_1^* and f_2^* corresponding to the faces f_1 and f_2 of G on the two sides of e. (For a bridge e, $f_1 = f_2$, so e^* is a loop.)

Remark. The dual of a connected simple graph (i.e., a graph – no loops or multiple edges) may be a multigraph.

Remark. The dual of a graph depends on the drawing of that graph! Two distinct drawings of the same abstract planar graph can give rise to non-isomorphic dual graphs.

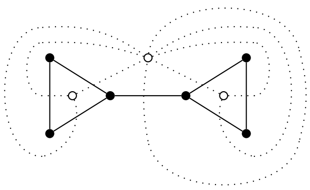

Figure 14: The dual of a plane graph.

Lemma 6.10. *We can draw G^* in the plane so that each vertex f^* is a point in the corresponding face f of G, and each edge e^* crosses the corresponding edge e of G at one point, and is otherwise disjoint from G. Moreover, every face of G^* contains a vertex of G.*

Proof. Pick a point p_e in the interior of each edge e and, for each face f, pick f^* to be any point in f. We can join f^* to each p_e, $e \in E(\partial f)$, with polygonal curves whose interiors are disjoint and all lie in f. Indeed, we join them one by one. At each stage f may be subdivided, but all subdivisions still have f^* in their boundary and so at each stage all remaining p_e are in the boundary of a face that also has f^* in its boundary, so can be joined to f^* by Lemma 6.3. Moreover, each time a p_e is joined to f^*, any subdivided face still contains at least one of the end vertices of e in its boundary. This procedure does not affect any other face of G, so we can do this independently with every face. The e^* are now formed by concatenating the two curves incident to p_e. It is easy to check that the result is a plane drawing of G^*, and each face of G^* must contain at least one vertex of G. □

We note that G^* is always connected, and if G is connected, then it is easy to check that G^* has one face for every vertex of G, and indeed that $(G^*)^*$ is isomorphic to G (thus motivating the name 'dual').

Lemma 6.11. *If G is a connected plane graph then $(G^*)^*$ is isomorphic to G.*

Proof. It is clear from Lemma 6.10 that the drawing of G gives a representation of the dual of G^*, provided every face of G^* contains just one vertex of G. (Indeed, this fails if G is not connected.) However, both G and G^* are connected, $|E| = |E^*|$ and $|V^*| = |F|$. But by Euler's Formula, $|V| - |E| + |F| = 2 = |V^*| - |E^*| + |F^*|$, so $|F^*| = |V|$. As each face of G^* contains at least one vertex of G, it must contain exactly one. □

A ***map*** is a connected bridgeless plane (multi-)graph. One of the most famous problems in graph theory, posed in 1852, is: can the faces of every map be coloured with 4 colours so that faces sharing an edge get different colours? Taking duals, it is not hard to check that this is equivalent to asking whether every planar (simple) graph G has $\chi(G) \leqslant 4$. The answer is yes; the result is known as the ***Four Colour Theorem***.

If G is planar and has $n \geqslant 3$ vertices, then $e(G) \leqslant 3n - 6$, so $\sum_v d(v) = 2e(G) < 6n$, and G must have $\delta(G) \leqslant 5$. It follows easily by induction on $|G|$ that every planar graph G has $\chi(G) \leqslant 6$. With not too much work, we can improve this by one, to obtain the 'Five Colour Theorem'.

Theorem 6.12 (Heawood, 1890). *If G is planar then $\chi(G) \leqslant 5$.*

Proof. We argue by induction on $n = |G|$. If $n \leqslant 5$ then the result is trivial, so suppose G is planar and has $n \geqslant 6$ vertices, and every planar graph with fewer vertices is 5-colourable.

As shown above, G has some vertex v with $d(v) \leqslant 5$. Draw G in the plane, and let c be a 5-colouring of the plane graph $G - v$. If any of the 5 colours does not appear on a neighbour of v we can extend the colouring to G, and we are done. So we may assume that $d(v) = 5$ and that the colours of the neighbours of v are distinct. Let the neighbours of v be $v_1, v_2, \ldots, v_5$ in cyclic order, and without loss of generality suppose that $c(v_i) = i$.

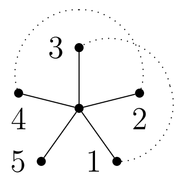

Let H be the subgraph of $G - v$ induced by the vertices with colour 1 or 3. If v_1 and v_3 are in different components of H then, swapping colours 1 and 3 in the component of H containing v_3, say, we find a new 5-colouring c' of $G - v$ in which $c'(v_1) = c'(v_3) = 1$; this colouring extends to all of G and we are done. Thus we may assume there exists a path P_1 in $G - v$ joining v_1 to v_3 in which all vertices have colour (in c) 1 or 3. Similarly, there exists a path P_2 in $G - v$ joining v_2 to v_4 in which all vertices have colour (in c) 2 or 4. The paths P_1 and P_2 are vertex disjoint, so in the drawing they do not cross. Since the cycle vP_1v separates the plane and P_2 starts and ends on different sides of this cycle, this gives a contradiction. □

The paths described above are often called ***Kempe chains***. Kempe thought he had proved the *four* colour theorem in 1879. The theorem was actually first proved by Appel and Haken in 1977 making extensive use of computers. A simpler, but still computer-based, proof was given by Robertson, Sanders, Seymour and Thomas in 1997. As of today there is no simple proof known.

7 Flows, connectivity and matchings

Imagine a road network in which each road has a certain 'capacity', or maximum flow in cars/hour. How can we work out the maximum traffic flow from one or more 'sources' to one or more 'sinks' or target destinations? Since the capacity of a road may not be the same in the two directions (for example if it is one-way) it makes sense to consider this question in the context of ***directed graphs***.

Formally, a ***directed graph*** or ***digraph*** $\overrightarrow{G} = (V, \overrightarrow{E})$ consists of a set V, the set of vertices, and a set $\overrightarrow{E}$ of *ordered* pairs of distinct elements of V, the (***directed***) ***edges***. We write $\overrightarrow{E}$ to remind ourselves the graph is directed; often the edge-set is just denoted E. We think of $(x, y) \in \overrightarrow{E}$ as an edge from x to y, and write $\overrightarrow{xy}$ or simply xy. Note that a directed graph cannot contain more than one edge from x to y, but can contain both edges xy and yx, which are now considered as distinct edges. We write

$$N^+(x) = \{y \in V : xy \in \overrightarrow{E}\}$$

for the ***out-neighbourhood*** of $x \in V$, and

$$N^-(x) = \{y \in V : yx \in \overrightarrow{E}\}$$

for its ***in-neighbourhood***.

A ***flow*** in G with ***source*** s and ***sink*** t is a function $f\colon \overrightarrow{E} \to [0,\infty)$ such that for every $x \in V \setminus \{s,t\}$ we have

$$\sum_{y\in N^+(x)} f(xy) = \sum_{y\in N^-(x)} f(yx),$$

i.e., the flow out of x is equal to the flow into x. Here s and t are distinct vertices.

Given any function $f\colon \overrightarrow{E} \to \mathbb{R}$, for $x \in V$ let

$$I_f(x) = \sum_{y\in N^+(x)} f(xy) - \sum_{y\in N^-(x)} f(yx).$$

We may think of $I_f(x)$ as the amount of flow that must be injected into the graph at x to maintain balance; in a flow, $I_f(x) = 0$ for $x \in V \setminus \{s,t\}$.

For any flow (or, indeed, any function on $\overrightarrow{E}$),

$$\sum_{x\in V} I_f(x) = \sum_{x\in V}\Bigg(\sum_{y\in N^+(x)} f(xy) - \sum_{y\in N^-(x)} f(yx)\Bigg) = 0,$$

since for every $uv \in \overrightarrow{E}$, $f(uv)$ appears exactly twice, once with $x = u$ and once with $x = v$. For a flow, the terms with $x \neq s,t$ are zero, so $I_f(s) = -I_f(t)$. In other words, the net flow leaving s equals the net flow arriving at t. This common value is called the ***value*** of f, and written $v(f)$. (Usually, $v(f) = I_f(s) = -I_f(t)$ is positive – otherwise we would regard the flow as having t as source and s as sink.) We can think of flow as being 'conserved' at every vertex, but with flow $v(f)$ injected into the graph at s and flow $v(f)$ extracted at t.

A ***capacity function*** on a directed graph $G = (V, \overrightarrow{E})$ is just a function $c\colon \overrightarrow{E} \to [0,\infty)$. A flow f is ***feasible*** (w.r.t. c) if $f(xy) \leqslant c(xy)$ for every $xy \in \overrightarrow{E}$. The key question in the theory of flows is: what is the maximum value of a feasible flow in a given graph with given source s, sink t and capacity function c? To avoid repeating the definitions, we shall call a directed graph with a given sink, source and capacity function a ***network***. (Of course, the word 'network' has many different meanings, depending on the context.) When we say f is a flow in a given network, it is always understood that f is feasible.

Given sets S and T of vertices of a directed graph $(V, \overrightarrow{E})$, let

$$\overrightarrow{E}(S,T) = \big\{xy \in \overrightarrow{E} : x \in S,\ y \in T\big\}$$

be the set of edges from S to T.

A ***cut*** in a network is a partition of the vertex set into disjoint sets S and T with $s \in S$ and $t \in T$. (Alternatively, we may say that a corresponding set $\overrightarrow{E}(S,T)$ of edges is a cut.) The ***capacity*** of a cut (S,T) is

$$c(S,T) = \sum_{xy\in \overrightarrow{E}(S,T)} c(xy),$$

i.e., the maximum conceivable flow from S to T (ignoring what happens within S and T). Clearly, in any feasible flow f, $v(f) \leqslant c(S,T)$. Indeed,

$$\begin{aligned} v(f) = I_f(s) = \sum_{x\in S} I_f(x) &= \sum_{x\in S}\bigg(\sum_{y\in N^+(x)} f(xy) - \sum_{y\in N^-(x)} f(yx)\bigg) \\ &= \sum_{xy\in \overrightarrow{E}(S,T)} f(xy) + \sum_{xy\in \overrightarrow{E}(S,S)} f(xy) - \sum_{yx\in \overrightarrow{E}(T,S)} f(yx) - \sum_{yx\in \overrightarrow{E}(S,S)} f(yx) \\ &= \sum_{xy\in \overrightarrow{E}(S,T)} f(xy) - \sum_{yx\in \overrightarrow{E}(T,S)} f(yx) \leqslant \sum_{xy\in \overrightarrow{E}(S,T)} c(xy) - 0 = c(S,T). \end{aligned} \tag{6}$$

Thus the maximum value of a feasible flow is at most the minimum capacity of a cut. The remarkable ***max-flow min-cut theorem*** tells us that we have equality.

Theorem 7.1 (Max-flow min-cut). *In any network $(\overrightarrow{G}, s, t, c)$ we have*

$$\sup\big\{v(f) : f \text{ is a feasible flow}\big\} = \min\big\{c(S,T) : (S,T) \text{ is a cut}\big\}.$$

Moreover, the supremum is attained.

The key ingredient of the proof is the notion of an ***augmenting path***, or ***slack path***. Let f be a flow in a network. We say that an ordered pair (x,y) is ε***-slack*** if either $xy \in \overrightarrow{E}$ and $f(xy) \leqslant c(xy) - \varepsilon$, or $yx \in \overrightarrow{E}$ and $f(yx) \geqslant \varepsilon$ (or both). A path $x_0x_1\cdots x_r$ in the undirected graph associated to $\overrightarrow{G}$ is ε***-slack*** if $x_{i-1}x_i$ is ε-slack for $1 \leqslant i \leqslant r$, and ***slack*** (or ***augmenting***) if it is slack for some $\varepsilon > 0$.

Lemma 7.2. *Let f be a flow in a network. If $x_0x_1\cdots x_r$ is an ε-slack path with $x_0 = s$ and $x_r = t$ then $v(f)$ is not maximal; in particular, there is a flow f' with $v(f') = v(f)+\varepsilon$.*

Proof. For each i we can either increase the flow along $x_{i-1}x_i$ by ε, or decrease the flow along x_ix_{i-1} by ε. Doing either increases $I_f(x_{i-1})$ by ε and decreases $I_f(x_i)$ by the same amount. Making such a change for each $i = 1, 2, \ldots, r$, we see that $I_f(x)$ is unchanged for every $x \neq s,t$ (so we still have a flow), and that $v(f) = I_f(s) = -I_f(t)$ is increased by ε. □

Proof of Theorem 7.1. First, we show that the supremum is attained. As noted earlier, for any flow f and cut (S,T) we have $v(f) \leqslant c(S,T)$. In particular,

$$v(f) \leqslant c\big(\{s\}, V\setminus\{s\}\big) = \sum_{y\in N^+(s)} c(sy) < \infty,$$

so the set $\{v(f) : f \text{ a flow}\}$ is bounded. So there are flows f_i with $v(f_i) \to M < \infty$, where $M = \sup\{v(f) : f \text{ a flow}\}$. Let $xy \in \overrightarrow{E}$. Then, passing to a subsequence, we may assume (by Bolzano–Weierstrass, given that $f_i(xy) \in [0, c(xy)]$ is bounded) that $f_i(xy)$ converges. Repeating this for each edge, we find a (sub)sequence of flows with $v(f_i) \to M$ such that $f_i(xy)$ converges for each $xy \in \overrightarrow{E}$. But then $f(xy) = \lim_{i\to\infty} f_i(xy)$ defines a flow with value $\lim_{i\to\infty} v(f_i) = M$.

Let f be a flow attaining the supremum. It suffices to find a cut with capacity $v(f)$. Let

$$S = \{x \in V : \text{ there is a slack path from } s \text{ to } x\},$$

and let $T = V \setminus S$. Clearly $s \in S$ (consider a path of length 0). By Lemma 7.2, $t \notin S$. Thus (S,T) is a cut. Suppose $x \in S$ and $y \in T$ with (x,y) slack. Then taking a slack path $s = x_0 \cdots x_r = x$ and appending $x_{r+1} = y$ gives a slack path ending at y, contradicting $y \in T$. Hence, for every $xy \in \overrightarrow{E}(S,T)$ we have $f(xy) = c(xy)$, and for every $yx \in \overrightarrow{E}(T,S)$ we have $f(yx) = 0$. I.e., equality holds in (6), so $c(S,T) = v(f)$. □

A ***maximal*** flow is one with maximum value. A function (here f or c) is ***integral*** if all its values are integers.

Theorem 7.3. *Let $(\overrightarrow{G}, s, t, c)$ be a network in which the capacity function c is integral. Then there is a maximal flow f which is integral.*

Proof. We have essentially described an algorithm to find such an f; the key point is that if the capacity function c and flow f are integral, then any slack path is 1-slack. Start with the flow with $f(xy) = 0$ for all edges, and repeat the following: if there is a slack (and hence 1-slack) path from s to t, augment the flow along this path by 1 as above, obtaining a new integral flow with larger value; repeat. Otherwise, by the last part of the proof above, there is a cut (S,T) with $v(f) = c(S,T)$, so f is maximal. □

The algorithm defined above is in fact reasonably efficient: it is easy to check for the existence of slack paths by (for example) breadth-first search.

A ***directed path***[10] in a directed graph $\overrightarrow{G} = (V, \overrightarrow{E})$ is a sequence $x_0x_1 \cdots x_r$ of distinct vertices such that $x_0x_1, \ldots, x_{r-1}x_r \in \overrightarrow{E}$, i.e., a path in which all the edges are directed the 'right way'. A set $\overrightarrow{X} \subseteq \overrightarrow{E}$ of edges ***separates*** s ***from*** t if $\overrightarrow{G} - \overrightarrow{X}$ contains no directed path (or, equivalently, no directed walk) from s to t. If (S,T) is a cut, then $\overrightarrow{E}(S,T)$ separates s from t. Conversely, if $\overrightarrow{X}$ separates s from t then it contains $\overrightarrow{E}(S,T)$ for some cut (S,T) – for example, take S to be the set of vertices x such that $\overrightarrow{G} - \overrightarrow{X}$ contains a directed s-x path. Let $c(\overrightarrow{X}) = \sum_{xy \in \overrightarrow{X}} c(xy)$. Then we see that

$$\min\{c(S,T) : (S,T) \text{ is a cut}\} = \min\{c(\overrightarrow{X}) : \overrightarrow{X} \text{ separates } s \text{ from } t\}. \tag{7}$$

This gives an alternative formulation of the max-flow min-cut theorem. Note, however, that cuts arise in the proof in an essential way, and it is necessary to consider reducing flow along backwards edges as well as increasing it along forwards ones.

The max-flow min-cut theorem has many variants, some of which we leave as exercises. For example, we may consider several sources $s_1, \ldots, s_k$ and several sinks $t_1, \ldots, t_\ell$. In this context, a cut (S,T) is a partition of the vertices with all sources in S and all sinks

[10] ***Directed walks***, ***directed cycles***, etc. are defined similarly. One can have directed 2-cycles if $uv, vu \in \overrightarrow{E}$.

in T. A flow must satisfy $I_f(x) = 0$ for every vertex that is neither a source nor a sink, and its value is $\sum_{i=1}^{k} I_f(s_i)$. Theorems 7.1 and 7.3 apply *mutatis mutandis* to this setting.

Another important variation allows some edges to have infinite capacity, meaning that the flow along xy can take any finite value. The results hold in this setting too, with the proviso that if there is no cut with finite capacity, then $\{v(f)\}$ is unbounded, so of course there is no flow with maximum value.

One more substantial variant is to impose capacity restrictions on the *vertices* rather than edges. Let $\overrightarrow{G}$ be a directed graph with source s and sink t, and let c be a (vertex) capacity function assigning every vertex $x \neq s, t$ a capacity $c(x) \in [0, \infty)$. A flow in $\overrightarrow{G}$ is feasible if for each vertex $x \neq s, t$ we have

$$\sum_{y \in N^-(x)} f(yx) = \sum_{y \in N^+(x)} f(xy) \leqslant c(x),$$

i.e., the flow through x is at most $c(x)$. (The equality is the definition of a flow; feasibility is the inequality.)

A ***vertex-cut*** is a set $S \subseteq V \setminus \{s, t\}$ of vertices such that in $\overrightarrow{G} - S$ there is no directed path from s to t. The ***capacity*** of S is $\sum_{x \in S} c(x)$.

Theorem 7.4. *Let $\overrightarrow{G}$ be a directed graph with source s, sink t and vertex capacity function c. Then the maximum value of a feasible flow from s to t is the minimum capacity of any vertex-cut. Furthermore, if c is integral, then there is a flow with maximum value that is integral.*

Proof. Rather than modify the proof of Theorem 7.1, we modify the network so that we can apply that result.

Form a directed graph $\overrightarrow{H}$ with source s and sink t by replacing each vertex $x \neq s, t$ by two vertices x^- and x^+ joined by an edge x^-x^+ with capacity $c(x)$. For each edge of $\overrightarrow{G}$ there is an edge of $\overrightarrow{H}$ with infinite (or very large) capacity; edges that start/end at s/t in $\overrightarrow{G}$ do so in $\overrightarrow{H}$; every edge of $\overrightarrow{G}$ ending at $x \neq s, t$ now ends at x^-, and every edge starting at x now starts at x^+. It is easy to check that feasible flows in $\overrightarrow{G}$ are in 1-to-1 correspondence with feasible flows in $\overrightarrow{H}$. In $\overrightarrow{H}$, a set $\overrightarrow{X}$ of edges with $c(\overrightarrow{X})$ finite must be of the form $\overrightarrow{X}_S = \{x^-x^+ : x \in S\}$ for some $S \subseteq V \setminus \{s, t\}$. Moreover, $\overrightarrow{X}_S$ is separating if and only if S is a vertex-cut. The result thus follows from Theorem 7.1 and (7) and, for integrality, Theorem 7.3. □

Connectivity and Menger's Theorem

Let G be an (undirected) graph and $S \subseteq V(G)$. We say that S ***separates*** G (or is a ***vertex-cut*** of G) if $G - S$ is disconnected. For vertices x, y of G, S ***separates*** x ***and*** y if they are in different components of $G - S$.

For a non-negative integer k, a graph G is k***-connected*** if $|G| \geqslant k + 1$ and no set with fewer than k vertices separates G. For example, every graph is 0-connected. A graph G

is 1-connected iff it is connected and $|G| \geqslant 2$. A graph G is 2-connected iff it is connected with no cut-vertex and $|G| \geqslant 3$. The only k-connected graph with $|G| = k + 1$ is K_{k+1}.

The (***vertex***) ***connectivity*** of a graph G is defined as

$$\kappa(G) = \max \{k : G \text{ is } k\text{-connected}\}.$$

Equivalently, $\kappa(G)$ is the minimum number of vertices that must be deleted to either disconnect G, or reduce it to a single vertex. It follows easily from the definition that $\kappa(G - x) \geqslant \kappa(G) - 1$, and that if H is a spanning subgraph of G then $\kappa(G) \geqslant \kappa(H)$. It is an exercise to check that if e is an edge of G then $\kappa(G - e) \geqslant \kappa(G) - 1$.

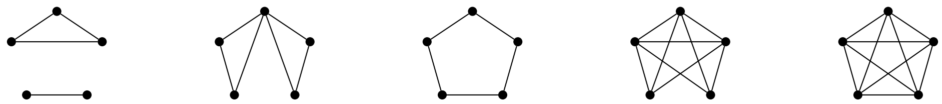

Figure 15: Examples of graphs with various connectivities.

We now define a 'local' version of connectivity. For distinct *non-adjacent* vertices x and y of G we write

$$\kappa(x, y) = \kappa_G(x, y) = \min \{|S| : S \text{ separates } x \text{ and } y\}.$$

Note that adjacent vertices can never be separated by deleting other vertices. Also, it is easy to check that for any non-complete graph G,

$$\kappa(G) = \min_{xy \in E(\overline{G})} \kappa_G(x, y).$$

Two distinct x-y paths are ***independent*** (or ***internally vertex-disjoint***) if the only vertices they share are x and y. A set of x-y paths is independent if the paths are pairwise independent.

Theorem 7.5 (Menger's Theorem)**.** *Let x and y be distinct non-adjacent vertices of G. Then the maximum size of an independent set of x-y paths is $\kappa_G(x, y)$.*

Proof. If there are k independent x-y paths then $\kappa_G(x, y) \geqslant k$ as we must remove at least one vertex from each path to separate x from y. So it is enough to show that there are $\kappa_G(x, y)$ independent paths.

Turn G into a network with source x and sink y by replacing each edge uv by two directed edges $\overrightarrow{uv}$ and $\overrightarrow{vu}$, and assigning each vertex other than x and y capacity 1. Then a vertex-cut S is simply a set of vertices separating x and y, and its capacity is just $|S|$. Hence, by Theorem 7.4, there is an *integral* flow f from x to y with value $\kappa_G(x, y)$. Given the vertex capacities, f can only take values 0 and 1, so f corresponds to a set of edges consisting of independent x-y paths and perhaps some (vertex disjoint) directed cycles. The value of f is the number of paths, so there are $\kappa_G(x, y)$ paths as required. □

Corollary 7.6. *A graph G is k-connected iff $|G| \geqslant k+1$ and every pair of non-adjacent vertices is joined by k independent paths.*

We can also define ***edge connectivity***, and prove a form of Menger's Theorem for edge-disjoint paths. More specifically, we say a graph with at least 2 vertices is k-***edge-connected*** if removing at most $k-1$ edges cannot disconnect the graph. We define the edge connectivity as

$$\lambda(G) = \max\big\{k : G \text{ is } k\text{-edge-connected}\big\}.$$

(By convention $\lambda(K_1) = 0$.) It should be noted that $\lambda(G) \neq \kappa(G)$ in general. For example, the 2nd graph in Figure 15 has $\kappa(G) = 1$ but $\lambda(G) = 2$. However, in general we have the following inequalities.

Lemma 7.7. *For any graph G, $\kappa(G) \leqslant \lambda(G) \leqslant \delta(G)$.*

Proof. All parameters are 0 when $|G| = 1$, so assume $|G| \geqslant 2$. Removing all edges incident to a vertex of minimum degree disconnects the graph, so $\lambda(G) \leqslant \delta(G)$. To show $\kappa(G) \leqslant \lambda(G)$, suppose S is a minimal set of edges such that $G - S$ is disconnected. We may assume $|S| \leqslant |G| - 2$ as otherwise $\kappa(G) \leqslant |G| - 1 \leqslant |S| = \lambda(G)$. Let A be the vertices of some component of $G - S$ and $B = V(G) \setminus A$. Then $E(A, B) \subseteq S$ disconnects G so $S = E(A, B)$ by minimality of $|S|$. Now for each edge $e \in S$ remove one of the end vertices of e in such a way that we do not remove the whole of A or the whole of B. We can do this as each edge in S has one endpoint in A and one endpoint in B, and $|S| \leqslant |G|-2$. Now we have removed $|S|$ vertices and disconnected the graph as the remaining vertices of A are not connected to the remaining vertices of B. Hence $\kappa(G) \leqslant |S| = \lambda(G)$. □

Given distinct vertices x and y, define $\lambda_G(x, y)$ to be the minimum number of edges one needs to remove from G so that x and y are no longer in the same component (i.e., the minimum size of a separating set of edges). Clearly, for any graph G with $|G| \geqslant 2$,

$$\lambda(G) = \min_{x \neq y} \lambda_G(x, y).$$

Theorem 7.8 (Menger's Theorem, edge version)**.** *Let x and y be distinct vertices of G. Then the maximum number of pairwise edge-disjoint x-y paths is $\lambda_G(x, y)$.*

Proof. If there are k edge-disjoint x-y paths then $\lambda_G(x, y) \geqslant k$ as we must remove at least one edge from each path to separate x from y. So it is enough to show that there are at least $\lambda_G(x, y)$ edge-disjoint paths.

Turn G into a network with source x and sink y by replacing each edge uv by two directed edges $\overrightarrow{uv}$ and $\overrightarrow{vu}$, and assigning each edge capacity 1. Then any cut (S, T) gives rise to a set $E(S, T)$ of edges that separates x from y, and $c(S, T) = |\overrightarrow{E}(S, T)| = |E(S, T)| \geqslant \lambda_G(x, y)$. Hence, by Theorem 7.3, there is an *integral* flow f from x to y with value at least $\lambda_G(x, y)$. Given the edge capacities, f can only take values 0 and 1, so f corresponds to a set of directed edges. If there is a directed cycle, we can remove it (set f to zero on its edges) and this leaves $v(f)$ and the fact that f is a flow unaffected. Now each undirected edge can appear at most once (as $\vec{uv}$ and $\vec{vu}$ would form a directed 2-cycle). Also if we just 'follow directed edges' starting at x using the directed edges where $f = 1$ then we must follow a path (as there are no directed cycles), and end at y (as at any other vertex we have as many edges out as in). Removing this path and repeating gives $v(f) \geqslant \lambda_G(x, y)$ edge-disjoint paths from x to y. □

Hall's Theorem

A ***matching*** M in a graph G is a set of pairwise disjoint edges of G; its ***size*** $|M|$ is the number of edges. Let G be a bipartite graph with vertex classes V_1 and V_2. A ***complete matching*** from V_1 to V_2 is a matching such that every vertex in V_1 is incident with some edge in the matching, i.e., a matching of size $|V_1|$.

Given a set S of vertices in a graph G, we write $N(S) = \bigcup_{s \in S} N(s)$ for the set of vertices with at least one neighbour in S.[11]

Theorem 7.9 (Hall's Marriage Theorem)**.** *Let G be a bipartite graph with bipartition (V_1, V_2). Then G contains a complete matching from V_1 to V_2 iff $|N(S)| \geqslant |S|$ for each $S \subseteq V_1$.*

The condition that $|N(S)| \geqslant |S|$ for each $S \subseteq V_1$ is called ***Hall's condition***. It is trivially necessary as each element of S must be matched to a different element of $N(S)$. We give two proofs of sufficiency.

Proof. We can deduce the result from Menger's Theorem. Instead, here is an outline of a proof directly from Theorem 7.3.

Form a directed graph by orienting every edge from V_1 to V_2, and adding a new vertex s with an edge sx for every $x \in V_1$ and a new vertex t with edges xt, $x \in V_2$. Assign all the new edges capacity 1, and the edges from V_1 to V_2 some very large (or infinite) capacity; $|V_1| + 1$ is large enough. Let (S, T) be a cut, and let $S_i = S \cap V_i$ and $T_i = T \cap V_i$. Either (i) $\overrightarrow{E}(S, T)$ contains some edge from V_1 to V_2. Then $c(S, T) > |V_1|$. Or (ii) not, i.e., $N(S_1) \subseteq S_2$. Then

$$\begin{aligned} c(S,T) &= |T_1| + 0 + |S_2| \\ &= |V_1| - |S_1| + |S_2| \\ &\geqslant |V_1| - |S_1| + |N(S_1)| \\ &\geqslant |V_1|, \end{aligned}$$

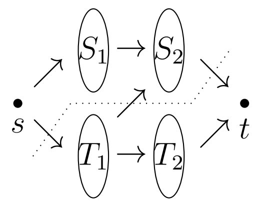

by Hall's condition. Hence the capacity of any cut is at least $|V_1|$, so by Theorem 7.3 there is an integral flow f with value $|V_1|$. But it is easy to check that f can only take

[11] Some authors exclude elements of S from $N(S)$. In the present context, where G is bipartite and $S \subseteq V_1$, it makes no difference.

the values 0 and 1, and that the edges e from V_1 to V_2 with $f(e) = 1$ correspond to a complete matching in G. $\square$

Here is a direct proof of Hall's Theorem.

Proof. We argue by induction on $n = |V_1|$. If $n = 1$, the result is trivial. For the induction step, let $n \geqslant 2$ and suppose that the result holds for all graphs with $|V_1| < n$. Consider a graph G with $|V_1| = n$ and assume that Hall's condition holds. There are two cases.
(a) Suppose first that $|N(S)| > |S|$ for each $\emptyset \neq S \subsetneq V_1$. Let xy be any edge of G with $x \in V_1$ and $y \in V_2$. Form G' by deleting the vertices x and y from G. Then G' satisfies Hall's condition (since if $\emptyset \neq S \subseteq V_1 \backslash \{x\}$ then $|N'(S)| = |N(S) \backslash \{y\}| \geqslant |N(S)| - 1 \geqslant |S|$), and so by induction G' has a complete matching from $V_1 \setminus \{x\}$ to $V_2 \setminus \{y\}$. Now adding the edge xy gives the required matching.
(b) If case (a) does not hold then $|N(S)| = |S|$ for some $\emptyset \neq S \subsetneq V_1$. The bipartite subgraph induced by $S \cup N(S)$ still satisfies Hall's condition, so by induction there is a complete matching M_1 from S to $N(S)$. Now consider $T = V_1 \backslash S$ and $U = V_2 \backslash N(S)$. We shall see that the bipartite subgraph H induced by $T \cup U$ also satisfies Hall's condition. For each $A \subseteq T$ we have

$$\begin{aligned} |N_H(A)| &= |N(A) \cap U| \\ &= |N(A \cup S) \setminus N(S)| \\ &= |N(A \cup S)| - |N(S)| \\ &\geqslant |A \cup S| - |S| = |A|, \end{aligned}$$

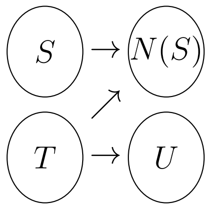

since $|N(A \cup S)| \geqslant |A \cup S|$ and $|N(S)| = |S|$. So Hall's condition holds in H, and by induction there is a complete matching M_2 from T to U. Then $M_1 \cup M_2$ is the required matching from V_1 to V_2. $\square$

Tutte's 1-factor Theorem

Although especially natural in bipartite graphs, it makes perfect sense to consider matchings in general graphs. A ***k*-factor** in a graph G is a spanning k-regular subgraph. Thus a 1-factor is exactly the same as a matching covering all vertices.

We call a component of a graph G ***odd*** if it has an odd number of vertices, and ***even*** otherwise. Let $q(G)$ denote the number of odd components of G, and note that $q(G) \equiv |G| \bmod 2$.

Theorem 7.10 (Tutte's 1-factor theorem)**.** *A graph G has a 1-factor if and only if, for every $S \subseteq V(G)$, we have*

$$q(G - S) \leqslant |S|. \tag{8}$$

Proof. In any 1-factor (complete matching) M, every odd component C of $G - S$ contains at least one vertex paired with some vertex outside C. Since the only edges leaving C in G go to S, a vertex of C must be paired with a vertex of S, and so $|S| \geqslant q(G - S)$. This shows that (8) is necessary. We prove sufficiency by induction on $|G|$. The case $|G| = 1$ (or $|G| = 2$) is trivial.

Suppose then that G satisfies (8), and that the result holds for all smaller graphs. Taking $S = \emptyset$ in (8) we see that $q(G) = 0$, and in particular $|G|$ is even. Also, for any vertex v of G, $q(G - v)$ is odd (since $|G - v|$ is). Hence $q(G - v) \geqslant 1$ and (since we are assuming (8)), for $S = \{v\}$ we have $q(G - S) = |S|$.

Let S be a subset of $V(G)$ for which $q(G - S) = |S|$ with $s = |S|$ maximal. From the above, $s \geqslant 1$, so S is not empty. Let $O_1, \ldots, O_s$ be the odd components of $G - S$ and $E_1, \ldots, E_k$, $k \geqslant 0$, the even components (if there are any). We shall prove the following three statements.

(a) each E_i has a 1-factor,

(b) if v is any vertex of any O_i, then $O_i - v$ has a 1-factor, and

(c) there is a matching $s_1v_1, \ldots, s_sv_s$ in G such that $\{s_1, \ldots, s_s\} = S$ and $v_i \in O_i$ for $1 \leqslant i \leqslant s$.

Clearly, if (a)–(c) hold then G has a 1-factor: apply (c) first, then (b) and (a). It remains to prove (a)–(c).

To see (a), let $A \subseteq V(E_i)$. The components of $G - (A \cup S)$ are $O_1, \ldots, O_s$, all E_j other than E_i, and the components of $E_i - A$, so $q(G - (A \cup S)) = s + q(E_i - A)$ and

$$q(E_i - A) = q(G - (A \cup S)) - s \leqslant |A \cup S| - s = |A| + s - s = |A|.$$

Hence E_i satisfies (8) and by induction E_i has a 1-factor.

For (b), let v be a vertex of O_i. Let $A \subseteq V(O_i - v)$. Then the components of $G - (A \cup \{v\} \cup S)$ are the E_j, all O_j other than O_i, and the components of $(O_i - v) - A$. Hence

$$q((O_i - v) - A) = q(G - (A \cup \{v\} \cup S)) - (s - 1) < |A \cup \{v\} \cup S| - s + 1 = |A| + 2,$$

where the inequality is from (8) and the maximality of $|S|$. Modulo 2,

$$q((O_i - v) - A) \equiv |(O_i - v) - A| = |O_i| - 1 - |A| \equiv |A|,$$

since O_i is odd. Since $x < y + 2$ and $x \equiv y \bmod 2$ imply $x \leqslant y$, it follows that $q((O_i - v) - A) \leqslant |A|$, so $O_i - v$ satisfies (8), and (b) follows by induction.

Finally, for (c) let H be the bipartite graph with $V_1 = S$ and $V_2 = \{o_1, \ldots, o_s\}$, with an edge xo_i whenever $x \in S$ and there is at least one edge in G from x to O_i. It suffices to find a complete matching in H, so we check Hall's condition. Let $A \subseteq V_1 = S$. If $o_i \in V_2 \setminus N_H(A)$ then in G there are no edges from A to O_i, so O_i is a component of $G - (S \setminus A)$. Hence, $q(G - (S \setminus A)) \geqslant |V_2 \setminus N_H(A)| = s - |N_H(A)|$. Thus, by (8),

$$s - |N_H(A)| \leqslant q(G - (S \setminus A)) \leqslant |S \setminus A| = s - |A|.$$

Hence $|N_H(A)| \geqslant |A|$, so Hall's condition holds in H, and by Hall's Theorem H has the required complete matching. □

8 Extremal Problems

If G has a subgraph isomorphic to H we say G ***contains*** (***a copy of***) H, and sometimes write[12] $G \supseteq H$. If G does not contain a copy of H then we say G is H-***free***. For a graph H with $e(H) > 0$ and $n \geqslant 1$ an integer, define

$$\mathrm{ex}(n, H) = \max\big\{e(G) : |G| = n,\ G \text{ is } H\text{-free}\big\},$$

i.e., $\mathrm{ex}(n, H)$ is the maximum number of edges an n-vertex H-free graph can have. Define

$$\mathrm{EX}(n, H) = \big\{G : |G| = n,\ e(G) = \mathrm{ex}(n, H),\ G \text{ is } H\text{-free}\big\}$$

to be the set of n-vertex, H-free graphs with this maximum number of edges. The graphs in $\mathrm{EX}(n, H)$ are called the ***extremal graphs***; we often describe $\mathrm{EX}(n, H)$ by listing one graph from each isomorphism class. We call $\mathrm{ex}(n, H)$ the ***extremal number*** for H (a function of n, of course).

For example, if G is P_3-free then all edges must be disjoint. Thus $\mathrm{ex}(n, P_3) = \lfloor \frac{n}{2} \rfloor$, and $\mathrm{EX}(n, P_3)$ is $\{\frac{n}{2}K_2\}$ if n is even and $\{\frac{n-1}{2}K_2 \cup K_1\}$ if n is odd, where mK denotes the vertex disjoint union of m copies of K. A generalisation of this is the following.

Theorem 8.1 (Erdős and Gallai)**.** *For all n and k, $\mathrm{ex}(n, P_{k+1}) \leqslant \frac{k-1}{2}n$.*

Proof. This is just a reformulation of Theorem 3.8, since if a graph has no P_{k+1}, then the longest path in G has edge length at most $k-1$, and so $\bar{d}(G) \leqslant k-1$. Thus

$$e(G) = \frac{n\bar{d}(G)}{2} \leqslant \frac{k-1}{2}n. \qquad \square$$

The bound in Theorem 8.1 is tight in the case when $k \mid n$ as we can take G to be $\frac{n}{k}$ disjoint copies of K_k. In general however, the bound is not tight.

What about $\mathrm{ex}(n, K_3)$? Good candidate extremal graphs are the complete bipartite graphs $K_{k,n-k}$; to maximize the number $k(n-k)$ of edges we take $k = \lfloor n/2 \rfloor$. More generally, what is $\mathrm{ex}(n, K_{r+1})$?

A graph G is r-***partite*** if $V(G)$ is the disjoint union of r sets $V_1, \ldots, V_r$ (the ***vertex classes***) with $e(G[V_i]) = 0$ for each i, i.e., no edges within each V_i. In other words, all edges go between parts. This is exactly the same as saying that G is r-colourable. A graph G is ***complete*** r-***partite*** if, in addition, every possible edge between parts is present.

Note that empty parts are allowed: the key point is that inside any part with at least two vertices, edges are forbidden. Clearly, any r-partite graph does not contain a copy of K_{r+1} as two of the vertices of the K_{r+1} would have to lie in the same part.

[12] Perhaps we should not write this, since in other contexts it means that H itself is a subgraph of G, i.e., that the particular vertices and edges of H are present in G. But usually it is clear from context whether or not we are considering isomorphic copies.

Before continuing we make a trivial observation: if $a_1, \dots, a_r$ are integers with average $\bar{a} = \frac{1}{r}\sum a_i$ then all a_i are within 1 of each other (i.e., $\max a_i \leqslant \min a_i + 1$) if and only if every a_i is equal to $\lfloor \bar{a} \rfloor$ or $\lceil \bar{a} \rceil$. (There are two cases: all $a_i = m = \bar{a}$ for some integer m, or some $a_i = m$, some $a_i = m+1$; then $m < \bar{a} < m+1$.) Moreover, given r and $\sum a_i$, there is only one way (up to reordering) to choose the a_i so that these conditions hold. I.e., there is only one way to divide a given number of (indivisible) objects among r people 'as fairly as possible'.

For $n, r \geqslant 1$, the ***Turán graph*** $T_r(n)$ is the complete r-partite graph on n vertices with the vertex class sizes as equal as possible, i.e., each is $\lfloor \frac{n}{r} \rfloor$ or $\lceil \frac{n}{r} \rceil$. The ***Turán number*** $t_r(n)$ is $e(T_r(n))$. Note that if $n \leqslant r$ then $T_r(n) = K_n$.

For example, $T_2(n)$ is $K_{\lfloor n/2 \rfloor, \lceil n/2 \rceil}$. Thus $t_2(n)$ is $\frac{n^2}{4}$ if n is even, and $\frac{n^2-1}{4}$ if n is odd: $t_2(n) = \lfloor \frac{n^2}{4} \rfloor$.

$T_3(10)$ has class sizes 3, 3 and 4, and $9 + 12 + 12 = 33$ edges, so $t_3(10) = 33$. $T_1(n)$ has no edges, so $t_1(n) = 0$.

Facts about Turán graphs

1. Among all r-partite graphs with n vertices, $T_r(n)$ is the unique (up to isomorphism) one with the largest number of edges. Indeed, only complete r-partite graphs are candidates, and if two classes differ in size by 2 or more, moving a vertex from the larger to the smaller gains at least one edge (since it reduces the number of vertices it is *not* adjacent to).

2. Since each vertex class has size $\lfloor \frac{n}{r} \rfloor$ or $\lceil \frac{n}{r} \rceil$, $\delta(T_r(n)) = n - \lceil \frac{n}{r} \rceil$ and $\Delta(T_r(n)) = n - \lfloor \frac{n}{r} \rfloor$, so $\Delta - \delta \leqslant 1$. Hence $\delta(T_r(n)) = \lfloor \bar{d}(T_r(n)) \rfloor$ and $\Delta(T_r(n)) = \lceil \bar{d}(T_r(n)) \rceil$, where $\bar{d}(G) = 2e(G)/|G|$ is the average degree of a graph G.

3. To get from $T_r(n)$ to $T_r(n-1)$ we delete any vertex from a largest vertex class, i.e., any vertex of minimum degree. So $t_r(n) - \delta(T_r(n)) = t_r(n-1)$.

Theorem 8.2 (Turán's Theorem)**.** *For all positive integers n and r we have*

$$\mathrm{ex}(n, K_{r+1}) = t_r(n) \qquad \textit{and} \qquad \mathrm{EX}(n, K_{r+1}) = \{T_r(n)\}.$$

Proof. We fix r and use induction on n. If $n \leqslant r$ then $\mathrm{ex}(n, K_{r+1}) = \binom{n}{2} = t_r(n)$, and $\mathrm{EX}(n, K_{r+1}) = \{K_n\} = \{T_r(n)\}$, as required.

Now let $n > r$ and suppose that the result holds for $n-1$. Let G be a graph with n vertices and $t_r(n)$ edges containing no copy of K_{r+1}. We will show that $G \cong T_r(n)$, from which the result follows. (If H had $|H| = n$, $e(H) > t_r(n)$ and contained no copy of K_{r+1}, then some spanning subgraph G would have $t_r(n)$ edges; so $G \cong T_r(n)$ and then H would contain a copy of K_{r+1} as adding any edge to $T_r(n)$ forms a K_{r+1}.)

First note that by Fact 2

$$\delta(G) \leqslant \lfloor \bar{d}(G) \rfloor = \lfloor \bar{d}(T_r(n)) \rfloor = \delta(T_r(n)).$$

Let $v \in V(G)$ have degree $d(v) = \delta(G)$. Then for $H = G - v$ we have

$$e(H) = e(G) - d(v) = t_r(n) - \delta(G) \geqslant t_r(n) - \delta(T_r(n)) = t_r(n-1),$$

using Fact 3. But H contains no K_{r+1} so by the induction hypothesis, $H \cong T_r(n-1)$.

Now v cannot have a neighbour in each vertex class of H (or we would get a copy of K_{r+1}), so its neighbours must miss some vertex class V_i completely. Adding v to this class, we see that G is r-partite. Now by Fact 1, $G \cong T_r(n)$. □

The ***density*** of a graph G is $e(G)/\binom{|G|}{2} = \bar{d}(G)/(|G|-1)$. For r fixed and $n \to \infty$, $T_r(n)$ has density $1 - \frac{1}{r} + o(1)$. Thus by Turán's Theorem, $\mathrm{ex}(n, K_{r+1})/\binom{n}{2} = 1 - \frac{1}{r} + o(1)$. We say that K_{r+1} 'appears' at density $1 - \frac{1}{r}$. Thus K_3 appears at density $\frac{1}{2}$, K_4 at density $\frac{2}{3}$, K_5 at density $\frac{3}{4}$, and so on.

What about other graphs? If $\chi(H) \geqslant r+1$, then, since H is not r-partite, $T_r(n)$ contains no copy of H. Thus for any H, letting $r = \chi(H) - 1$ we have

$$\mathrm{ex}(n, H) \geqslant t_r(n) = \left(1 - \tfrac{1}{r} + o(1)\right)\binom{n}{2}.$$

In particular, H cannot appear before the density $1 - \frac{1}{r}$ at which K_{r+1} appears. Are there 'big' graphs with a given chromatic number that appear significantly later? Amazingly, the answer turns out to be no.

For $s, t \geqslant 1$ let $K_s(t)$ be the complete s-partite graph with t vertices in each class. For example, $K_1(t)$ is the empty graph E_t, $K_2(t)$ is the complete bipartite graph $K_{t,t}$, and $K_s(1) = K_s$. In general $K_s(t) = T_s(st)$.

Theorem 8.3 (Erdős–Stone Theorem)**.** *Let $r, t \geqslant 1$ be integers and let $\varepsilon > 0$. There is a constant $n_0 = n_0(r, t, \varepsilon)$ such that every graph G with $n \geqslant n_0$ vertices and*

$$e(G) \geqslant \left(1 - \tfrac{1}{r} + \varepsilon\right)\binom{n}{2}$$

contains a copy of $K_{r+1}(t)$.

Note that for $t = 1$ this follows (and is somewhat weaker than) Turán's theorem. We prove this theorem in two steps. We first show that any graph with a given density contains a subgraph with a relatively large minimum degree.

Lemma 8.4. *Let $0 \leqslant \alpha < \beta \leqslant 1$. If G is a graph with $|G| = n$ and $e(G) \geqslant \beta\binom{n}{2}$ then G has an (induced) subgraph H with*

$$\delta(H) \geqslant \alpha(|H| - 1)$$

and $|H| \geqslant \sqrt{\varepsilon} n$, where $\varepsilon = \beta - \alpha$.

Proof. Define a sequence $G_n, G_{n-1}, \ldots$ of graphs with $|G_t| = t$ as follows. Set $G_n = G$. If $\delta(G_t) \geqslant \alpha(t-1)$ then stop. Otherwise, remove a vertex of G_t with minimum degree to get G_{t-1}. The construction must stop at some point (G_1 at the latest); let G_k be the final graph, so $\delta(G_k) \geqslant \alpha(|G_k|-1)$ by construction. Now

$$\begin{aligned} e(G_k) &= e(G_n) - \delta(G_n) - \delta(G_{n-1}) - \cdots - \delta(G_{k+1}) \\ &\geqslant \beta\binom{n}{2} - \alpha\big((n-1)+(n-2)+\cdots+k\big) \\ &= \beta\binom{n}{2} - \alpha\binom{n}{2} + \alpha\binom{k}{2} \geqslant \varepsilon\binom{n}{2}. \end{aligned}$$

Since $e(G_k) \leqslant \binom{k}{2}$ it follows that $\binom{k}{2} \geqslant \varepsilon\binom{n}{2}$. Multiplying by $2k/(k-1) \geqslant 2n/(n-1)$ gives $k^2 \geqslant \varepsilon n^2$, so $k \geqslant \sqrt{\varepsilon}n$. □

Lemma 8.5. *Let $r, t \geqslant 1$ be integers and let $\varepsilon > 0$. There is a constant $n_1 = n_1(r,t,\varepsilon)$ such that every graph G with $n \geqslant n_1$ vertices and*

$$\delta(G) \geqslant \big(1 - \tfrac{1}{r} + \varepsilon\big)(n-1)$$

contains a copy of $K_{r+1}(t)$.

Note that Theorem 8.3 follows easily: apply Lemma 8.4 with $\beta = 1 - \frac{1}{r} + \varepsilon$ and $\alpha = 1 - \frac{1}{r} + \frac{\varepsilon}{2}$, and then Lemma 8.5 with $\varepsilon/2$ in place of ε.

Proof. We use induction on r, proving the base case and induction step together. More precisely, for $r \geqslant 1$ let $\mathbb{H}_r$ be the statement that for every $t \geqslant 1$ and $\varepsilon > 0$ there is an n_1 such that .... We shall prove $\mathbb{H}_r$ assuming, for $r \geqslant 2$, that $\mathbb{H}_{r-1}$ holds. Then we will have shown that $\mathbb{H}_1$ holds, and that $\mathbb{H}_{r-1}$ implies $\mathbb{H}_r$ for all $r \geqslant 2$, so by induction $\mathbb{H}_r$ holds for all $r \geqslant 1$. A key point is that in proving $\mathbb{H}_r$, $r \geqslant 2$, we must consider all t and all $\varepsilon > 0$; but for a given t we may use the fact that $\mathbb{H}_{r-1}$ holds for some other, perhaps *much* larger value of t.

Let $r \geqslant 1$. To prove $\mathbb{H}_r$, let $t \geqslant 1$ and $\varepsilon > 0$ be given, and let G be a graph with $|G| = n$ and $\delta(G) \geqslant (1 - \frac{1}{r} + \varepsilon)(n-1)$. Set

$$T = \lceil 2t/(\varepsilon r) \rceil.$$

If $r \geqslant 2$ then since

$$\delta(G) \geqslant (1 - \tfrac{1}{r} + \varepsilon)(n-1) \geqslant (1 - \tfrac{1}{r-1} + \varepsilon)(n-1),$$

we know by the induction hypothesis $\mathbb{H}_{r-1}$ that, if n is large enough (depending on r, t, ε), G must contain a copy of $K_r(T)$. If $r = 1$ then $K_r(T)$ is an empty graph with T vertices, so if n is large enough (i.e., $n \geqslant T$) then G certainly contains a copy of this graph.

In either case, let H be a subgraph of G isomorphic to $K_r(T)$. Denote its vertex classes by $S_1, \ldots, S_r$ and let $S = V(H)$ be their union. Let U be the set of vertices in $V(G) \setminus S$ which have at least t neighbours in each class S_i; these vertices are the *useful* ones.

If $|U| > (t-1)\binom{T}{t}^r$ then there are at least t vertices in U that have at least t *common* neighbours in each S_i, giving a $K_{r+1}(t)$; to see this assign to each $u \in U$ an r-tuple

$(A_1, \ldots, A_r)$ where $A_i \subseteq S_i \cap N(u)$ and $|A_i| = t$. As there are $\binom{T}{t}^r$ possible r-tuples, the average number of times an r-tuple is chosen is $|U|/\binom{T}{r}^t > t-1$, and so some r-tuple is chosen $\geqslant t$ times, i.e., we have t vertices in U all joined to the same copy $H' \subseteq H$ of $K_r(t)$.

So we may suppose that $|U| \leqslant (t-1)\binom{T}{t}^r$. Let $B = V(G) \setminus (S \cup U)$. We count the number N of edges of $\overline{G}$ between S and B in two different ways. Firstly, the degree in $\overline{G}$ of any vertex v is

$$n - 1 - d_G(v) \leqslant n - 1 - \delta(G) \leqslant (n-1)\left(\tfrac{1}{r} - \varepsilon\right) \leqslant \left(\tfrac{1}{r} - \varepsilon\right)n,$$

so counting from S we find that

$$N \leqslant |S|\left(\tfrac{1}{r} - \varepsilon\right)n = rT\left(\tfrac{1}{r} - \varepsilon\right)n = (T - \varepsilon rT)n.$$

On the other hand, for each $u \in B$ there is a vertex class S_i of H such that, in G, the vertex u has at most $t-1 \leqslant t$ neighbours in S_i. Then, in $\overline{G}$, u has at least $T-t$ neighbours in $S_i \subseteq S$. Hence, counting from B, we see that

$$N \geqslant |B|(T-t).$$

Since we chose T so that $\varepsilon rT \geqslant 2t$, it follows that

$$|B| \leqslant \frac{T - \varepsilon rT}{T-t}n \leqslant \frac{T-2t}{T-t}n = (1-c)n$$

for some constant $c > 0$ (depending on r, t, ε). But now in total there are

$$n = |S| + |B| + |U| \leqslant rT + (1-c)n + (t-1)\binom{T}{t}^r = (1-c)n + O(1)$$

vertices, a contradiction if n is large enough. □

Corollary 8.6. *Let H be any graph with $e(H) > 0$. Then*

$$\mathrm{ex}(n, H) = \left(1 - \tfrac{1}{\chi(H)-1} + o(1)\right)\binom{n}{2}$$

as $n \to \infty$, where $\chi(H)$ is the chromatic number of H.

Proof. Let $r = \chi(H) - 1$, so H has chromatic number $r+1$. Since H is not r-partite, $T_r(n)$ does not contain any copies of H, so

$$\mathrm{ex}(n, H) \geqslant t_r(n) = \left(1 - \tfrac{1}{r} + o(1)\right)\binom{n}{2}.$$

On the other hand, for large enough t (e.g., $t = |H|$), the graph $K_{r+1}(t)$ does contains a copy of H. Therefore

$$\mathrm{ex}(n, H) \leqslant \mathrm{ex}(n, K_{r+1}(t)) \leqslant \left(1 - \tfrac{1}{r} + o(1)\right)\binom{n}{2},$$

where the second inequality is from Theorem 8.3. □

Corollary 8.6 answers, at some level, the basic extremal question for any graph H. However, there is a weak point: while for $\chi(H) \geqslant 3$ it tells us asymptotically what value $\mathrm{ex}(n, H)$ has, for $\chi(H) = 2$ it only tells us that $\mathrm{ex}(n, H) = o(n^2)$, leaving a wide range of possible functions (e.g., roughly $n^2/\log n$, roughly n, roughly $n^{3/2}$ etc). Can we say something more precise in this case?

The Zarankiewicz Problem

Let G be a bipartite graph where the vertex classes have a given size n. How many edges can G have if it does not contain a copy of some given graph H? This makes sense only if H is bipartite, and in particular we consider $H = K_2(t) = K_{t,t}$, i.e., the bipartite analogue of the Turán problem. Formally, let

$$z(n,t) = \max\big\{e(G) : G \subseteq K_{n,n} \text{ and } G \text{ is } K_{t,t}\text{-free}\big\}.$$

Theorem 8.7. *If $n \geqslant t \geqslant 2$ then*

$$z(n,t) \leqslant (t-1)^{1/t} n^{2-1/t} + (t-1)n$$

In particular, as $n \to \infty$ with t fixed we have $z(n,t) = O(n^{2-1/t})$.

Proof. Let G be a maximal $K_{t,t}$-free bipartite graph with vertex classes X and Y, $|X| = |Y| = n$. Note that by maximality, $d(v) \geqslant t-1$ for every vertex v. (Otherwise, add a new edge incident with v. The degree of v would still be less than t, so the new edge cannot be in any $K_{t,t}$.)

We say that a vertex v ***covers*** a set S of vertices if $S \subseteq N(v)$. A vertex $v \in X$ covers exactly $\binom{d(v)}{t}$ t-element subsets of Y. On the other hand, a t-element subset of Y is covered by at most $t-1$ vertices in X; otherwise we have a $K_{t,t}$. Hence,

$$\sum_{v\in X} \binom{d(v)}{t} \leqslant (t-1)\binom{n}{t}.$$

From here on it is just calculation. Firstly, the polynomial $x(x-1)\cdots(x-t+1)$ is convex on $[t-1,\infty)$, so $\binom{x}{t} = x(x-1)\cdots(x-t+1)/t!$ is convex as a function of the real variable x when $x \geqslant t-1$. Let $d = \frac{1}{n}\sum_{v\in X} d(v)$ be the average degree in X. Then, by Jensen's inequality,

$$\binom{d}{t} \leqslant \frac{1}{n}\sum_{v\in X}\binom{d(v)}{t} \leqslant \frac{t-1}{n}\binom{n}{t}. \tag{9}$$

Hence

$$\frac{t-1}{n} \geqslant \binom{d}{t}\Big/\binom{n}{t} = \frac{d(d-1)\cdots(d-t+1)}{n(n-1)\cdots(n-t+1)} \geqslant \Big(\frac{d-t+1}{n}\Big)^t.$$

Rearranging gives

$$d - t + 1 \leqslant (t-1)^{1/t} n^{1-1/t}$$

and so

$$e(G) = dn \leqslant (t-1)^{1/t} n^{2-1/t} + (t-1)n. \qquad \square$$

Remark. The same method of proof works to show that $\mathrm{ex}(n, K_{t,t}) = O(n^{2-1/t})$; we count the number of copies of $K_{1,t}$ in two ways. (As in the proof above, but in the bipartite case we had the extra restriction that the special vertex of $K_{1,t}$ should be in X.) Alternatively, if we have a graph on n vertices with no copy of $K_{t,t}$, then by splitting the vertex set into two sets of size $n/2$ and considering the number of edges between and within each set, we have $\mathrm{ex}(n, K_{t,t}) \leqslant 2\,\mathrm{ex}(n/2, K_{t,t}) + z(n/2, t)$. This (for $t > 1$) and induction gives $\mathrm{ex}(n, K_{t,t}) = O(n^{2-1/t})$.

The special case $t = 2$ is the same but with simpler calculations.

Theorem 8.8. *For $n \geqslant 2$ we have*

$$z(n, 2) \leqslant \frac{n}{2}\bigl(1 + \sqrt{4n-3}\bigr) \sim n^{3/2}$$

as $n \to \infty$.

Proof. We have (9) as before. With $t = 2$ this becomes $d(d-1) \leqslant n-1$, rearranging and noting as before that $e(G) = nd$ gives the result. □

The bounds just given are only upper bounds. Are they close to the truth? In general, this is an open question. The case $t = 2$ is particularly nice. Here we have equality if and only if G is regular, any two vertices in Y have exactly one common neighbour in X, and vice versa. A structure with these properties is called a ***finite projective plane***: think of the vertices in X as points, those in Y as lines, and edges of G as representing incidence.

It turns out that, except for some degenerate cases, for equality we must have $n = q^2 + q + 1$, each point on $q + 1$ lines and each line having $q + 1$ points. Is this possible? For q any prime power the answer is yes: take the projective plane over a finite field with q elements. (This is enough to show that in fact $z(n, 2) \sim n^{3/2}$). However, for n not of this form, even the value of $z(n, 2)$ is unknown in general.

Summary of Notation

$[n] = \{1, 2, \dots, n\}$.

$\binom{X}{k} = \{A \subseteq X : |A| = k\}$ is the set of k-element subsets of X.

$V(G)$, the vertex set of the graph G.

$E(G)$, $\overrightarrow{E}(G)$, the edge set a graph or directed graph G.

$uv = vu = \{u, v\}$, an edge of a graph.

$|G| = |V(G)|$, the order, or number of vertices in G

$e(G) = |E(G)|$, the size, or number of edges in G

$e(A) = e(G[A])$, the number of edges of G with both endpoints inside A.

$E(A, B)$, $\overrightarrow{E}(A, B)$, the set of edges ab in a graph or digraph with $a \in A$, $b \in B$.

$e(A, B) = |E(A, B)|$, the number of edges ab of G with $a \in A$ and $b \in B$.

$N(v) = N_G(v) = \{u : uv \in E\}$, the neighbourhood of $v \in V(G)$.

$N^+(v) = \{u : vu \in \overrightarrow{E}\}$, the out-neighbourhood of a vertex in a directed graph.

$N^-(v) = \{u : uv \in \overrightarrow{E}\}$, the in-neighbourhood of a vertex in a directed graph.

$N_A(v) = A \cap N(v)$, the set of neighbours of v that lie in $A \subseteq V(G)$.

$N(S) = \bigcup_{x \in S} N(x)$, the neighbourhood of a set of vertices.

$d(v) = d_G(v) = |N(v)|$, the degree of v.

$d_A(v) = |N_A(v)|$, the degree of v into A, $A \subseteq V(G)$.

$\delta(G) = \min_{v \in V(G)} d(v)$, the minimum degree of the vertices of G.

$\Delta(G) = \max_{v \in V(G)} d(v)$, the maximum degree of the vertices of G.

$\bar{d}(G) = \frac{1}{|G|} \sum_{v \in V(G)} d(v)$, the average degree of the vertices of G.

$G \cong H$, graphs G and H are isomorphic.

$G[S] = (S, E(G) \cap \binom{S}{2})$, the subgraph of G induced by $S \subseteq V(G)$.

$\overline{G} = (V, \binom{V}{2} \setminus E)$, the complement of $G = (V, E)$.

$G - e = (V, E \setminus \{e\})$, the graph obtained by deleting $e \in E(G)$ from G.

$G - S = (V, E \setminus S)$, the graph obtained by removing all edges in S, $S \subseteq E(G)$.

$G + e = (V, E \cup \{e\})$, the graph obtained by adding $e \in E(\overline{G})$ to G.

$G - v = G[V \setminus \{v\}]$, the graph obtained by deleting the vertex v and any incident edges.

$G - S = G[V(G) \setminus S]$, the graph with all vertices in $S \subseteq V(G)$ removed.

G/e, the graph with e contracted (endvertices merged into a single vertex).

$G \cup H = (V \cup V', E \cup E')$, the union of two graphs.

$L(G)$, the line graph of G with vertex set $E(G)$ and edges ef iff e, f meet.

G^*, the dual of a plane graph G.

∂f, the boundary of a face of a plane graph G as a subgraph of G.

$d(f)$, the number of edges in ∂f, counting bridges twice.

K_n, the complete graph on $n \geqslant 1$ vertices, $K_n = ([n], \binom{[n]}{2})$.

E_n, the empty graph on $n \geqslant 1$ vertices, $E_n = ([n], \emptyset)$.

P_n, the path on $n \geqslant 1$ vertices ($n - 1$ edges), $P_n = ([n], \{12, 23, \ldots, (n-1)n\})$.

C_n, the cycle on $n \geqslant 3$ vertices (also n edges), $C_n = ([n], \{12, 23, \ldots, (n-1)n, n1)$.

$K_{a,b}$, the complete bipartite graph with a vertices in one part and b in the other.

$K_r(t)$, the complete r-partite graph with t vertices in each of the r partite classes.

$T_r(n)$, the Turán graph – complete r-partite graph with n vertices partitioned as equitably as possible.

$t_r(n) = e(T_r(n))$, the Turán number – the number of edges in the Turán graph.

$\chi(G) = \min\{k : \exists k\text{-colouring of } G\}$, the chromatic number of G.

$p_G(x)$, the chromatic polynomial of G: $p_G(k) = |\{k\text{-colourings of } G\}|$.

$\chi'(G) = \min\{k : \exists k\text{-edge-colouring of } G\}$, the edge-chromatic number of G.

$\omega(G)$, the clique number, or largest order of a complete subgraph of G.

$\alpha(G)$, the independence number, or size of the largest set of vertices, none of whom are adjacent in G.

$\kappa(G) = \max\{k : G \text{ is } k\text{-connected}\}$, the (vertex) connectivity of G.

$\kappa_G(x, y) = \min\{|S| : S \subseteq V(G) \setminus \{x, y\}, S \text{ separates } x \text{ and } y\}$.

$\lambda(G) = \max\{k : G \text{ is } k\text{-edge-connected}\}$, the edge connectivity of G.

$\lambda_G(x, y) = \min\{|S| : S \subseteq E(G), S \text{ separates } x \text{ and } y\}$.

$I_f(v) = \sum_{u \in N^+(v)} f(uv) - \sum_{u \in N^-(v)} f(vu)$, net flow out of a vertex.

$v(f) = I_f(s) = -I_f(t)$, the value of a flow.

$c(S, T) = \sum_{uv \in \vec{E}(S,T)} c(uv)$, the capacity of a cut.

$q(G)$, the number of odd components in G (in Tutte's 1-factor theorem).

$\mathrm{ex}(n, H)$, the maximum number of edges in an H-free graph on n vertices.

$\mathrm{EX}(n, H)$, the set of H-free graphs on n vertices with $\mathrm{ex}(n, H)$ edges.

$z(n, t)$, the maximum number of edges in a $K_{t,t}$-free subgraph of $K_{n,n}$.

If you find an error please check the website, and if it has not already been corrected, e-mail: `Paul.Balister@maths.ox.ac.uk`.