Reinhard Diestel

Graph Theory

Electronic Edition 2000

This is an electronic version of the second (2000) edition of the above Springer book, from their series *Graduate Texts in Mathematics*, vol. 173. The cross-references in the text and in the margins are active links: click on them to be taken to the appropriate page.

The printed edition of this book can be ordered from your bookseller, or electronically from Springer through the Web sites referred to below.

Softcover $34.95, ISBN 0-387-98976-5
Hardcover $69.95, ISBN 0-387-95014-1

Further information (reviews, errata, free copies for lecturers etc.) and electronic order forms can be found on

http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/
http://www.springer-ny.com/supplements/diestel/

introductory text should be lean and concentrate on the essential, so as to offer guidance to those new to the field. As a graduate text, moreover, it should get to the heart of the matter quickly: after all, the idea is to convey at least an impression of the depth and methods of the subject. On the other hand, it has been my particular concern to write with sufficient detail to make the text enjoyable and easy to read: guiding questions and ideas will be discussed explicitly, and all proofs presented will be rigorous and complete.

A typical chapter, therefore, begins with a brief discussion of what are the guiding questions in the area it covers, continues with a succinct account of its classic results (often with simplified proofs), and then presents one or two deeper theorems that bring out the full flavour of that area. The proofs of these latter results are typically preceded by (or interspersed with) an informal account of their main ideas, but are then presented formally at the same level of detail as their simpler counterparts. I soon noticed that, as a consequence, some of those proofs came out rather longer in print than seemed fair to their often beautifully simple conception. I would hope, however, that even for the professional reader the relatively detailed account of those proofs will at least help to minimize reading time...

If desired, this text can be used for a lecture course with little or no further preparation. The simplest way to do this would be to follow the order of presentation, chapter by chapter: apart from two clearly marked exceptions, any results used in the proof of others precede them in the text.

Alternatively, a lecturer may wish to divide the material into an easy basic course for one semester, and a more challenging follow-up course for another. To help with the preparation of courses deviating from the order of presentation, I have listed in the margin next to each proof the reference numbers of those results that are used in that proof. These references are given in round brackets: for example, a reference (4.1.2) in the margin next to the proof of Theorem 4.3.2 indicates that Lemma 4.1.2 will be used in this proof. Correspondingly, in the margin next to Lemma 4.1.2 there is a reference [4.3.2] (in square brackets) informing the reader that this lemma will be used in the proof of Theorem 4.3.2. Note that this system applies between different sections only (of the same or of different chapters): the sections themselves are written as units and best read in their order of presentation.

The mathematical prerequisites for this book, as for most graph theory texts, are minimal: a first grounding in linear algebra is assumed for Chapter 1.9 and once in Chapter 5.5, some basic topological concepts about the Euclidean plane and 3-space are used in Chapter 4, and a previous first encounter with elementary probability will help with Chapter 11. (Even here, all that is assumed formally is the knowledge of basic definitions: the few probabilistic tools used are developed in the

About the second edition

Naturally, I am delighted at having to write this addendum so soon after this book came out in the summer of 1997. It is particularly gratifying to hear that people are gradually adopting it not only for their personal use but more and more also as a course text; this, after all, was my aim when I wrote it, and my excuse for agonizing more over presentation than I might otherwise have done.

There are two major changes. The last chapter on graph minors now gives a complete proof of one of the major results of the Robertson-Seymour theory, their theorem that excluding a graph as a minor bounds the tree-width if and only if that graph is planar. This short proof did not exist when I wrote the first edition, which is why I then included a short proof of the next best thing, the analogous result for path-width. That theorem has now been dropped from Chapter 12. Another addition in this chapter is that the tree-width duality theorem, Theorem 12.3.9, now comes with a (short) proof too.

The second major change is the addition of a complete set of hints for the exercises. These are largely Tommy Jensen's work, and I am grateful for the time he donated to this project. The aim of these hints is to help those who use the book to study graph theory on their own, but *not* to spoil the fun. The exercises, including hints, continue to be intended for classroom use.

Apart from these two changes, there are a few additions. The most noticable of these are the formal introduction of depth-first search trees in Section 1.5 (which has led to some simplifications in later proofs) and an ingenious new proof of Menger's theorem due to Böhme, Göring and Harant (which has not otherwise been published).

Finally, there is a host of small simplifications and clarifications of arguments that I noticed as I taught from the book, or which were pointed out to me by others. To all these I offer my special thanks.

The Web site for the book has followed me to

http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/

I expect this address to be stable for some time.

Once more, my thanks go to all who contributed to this second edition by commenting on the first—and I look forward to further comments!

December 1999 *RD*

1.1 Graphs

graph *vertex* *edge*

A *graph* is a pair $G = (V, E)$ of sets satisfying $E \subseteq [V]^2$; thus, the elements of E are 2-element subsets of V. To avoid notational ambiguities, we shall always assume tacitly that $V \cap E = \emptyset$. The elements of V are the *vertices* (or *nodes*, or *points*) of the graph G, the elements of E are its *edges* (or *lines*). The usual way to picture a graph is by drawing a dot for each vertex and joining two of these dots by a line if the corresponding two vertices form an edge. Just how these dots and lines are drawn is considered irrelevant: all that matters is the information which pairs of vertices form an edge and which do not.

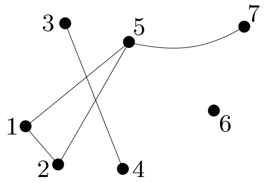

Fig. 1.1.1. The graph on $V = \{1, \dots, 7\}$ with edge set $E = \{\{1,2\},\{1,5\},\{2,5\},\{3,4\},\{5,7\}\}$

on $V(G), E(G)$

A graph with vertex set V is said to be a graph *on* V. The vertex set of a graph G is referred to as $V(G)$, its edge set as $E(G)$. These conventions are independent of any actual names of these two sets: the vertex set W of a graph $H = (W, F)$ is still referred to as $V(H)$, not as $W(H)$. We shall not always distinguish strictly between a graph and its vertex or edge set. For example, we may speak of a vertex $v \in G$ (rather than $v \in V(G)$), an edge $e \in G$, and so on.

order $|G|$, $\|G\|$

The number of vertices of a graph G is its *order*, written as $|G|$; its number of edges is denoted by $\|G\|$. Graphs are *finite* or *infinite* according to their order; unless otherwise stated, the graphs we consider are all finite.

$\emptyset$ *trivial graph*

For the *empty graph* $(\emptyset, \emptyset)$ we simply write $\emptyset$. A graph of order 0 or 1 is called *trivial*. Sometimes, e.g. to start an induction, trivial graphs can be useful; at other times they form silly counterexamples and become a nuisance. To avoid cluttering the text with non-triviality conditions, we shall mostly treat the trivial graphs, and particularly the empty graph $\emptyset$, with generous disregard.

incident *ends* $E(X, Y)$ $E(v)$

A vertex v is *incident* with an edge e if $v \in e$; then e is an edge *at* v. The two vertices incident with an edge are its *endvertices* or *ends*, and an edge *joins* its ends. An edge $\{x, y\}$ is usually written as xy (or yx). If $x \in X$ and $y \in Y$, then xy is an X–Y *edge*. The set of all X–Y edges in a set E is denoted by $E(X, Y)$; instead of $E(\{x\}, Y)$ and $E(X, \{y\})$ we simply write $E(x, Y)$ and $E(X, y)$. The set of all the edges in E at a vertex v is denoted by $E(v)$.

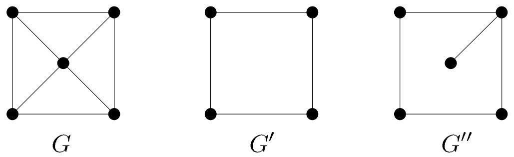

Fig. 1.1.3. A graph G with subgraphs G' and G'': G' is an induced subgraph of G, but G'' is not

$-$ If U is any set of vertices (usually of G), we write $G - U$ for $G[V \setminus U]$. In other words, $G - U$ is obtained from G by *deleting* all the vertices in $U \cap V$ and their incident edges. If $U = \{v\}$ is a singleton, we write $G - v$ rather than $G - \{v\}$. Instead of $G - V(G')$ we simply write $G - G'$. For a subset F of $[V]^2$ we write $G - F := (V, E \setminus F)$ and $+$ $G + F := (V, E \cup F)$; as above, $G - \{e\}$ and $G + \{e\}$ are abbreviated to $G - e$ and $G + e$. We call G *edge-maximal* with a given graph property if G itself has the property but no graph $G + xy$ does, for non-adjacent vertices $x, y \in G$.

edge-maximal

minimal
maximal

More generally, when we call a graph *minimal* or *maximal* with some property but have not specified any particular ordering, we are referring to the subgraph relation. When we speak of minimal or maximal sets of vertices or edges, the reference is simply to set inclusion.

$G * G'$

complement $\overline{G}$

line graph $L(G)$

If G and G' are disjoint, we denote by $G * G'$ the graph obtained from $G \cup G'$ by joining all the vertices of G to all the vertices of G'. For example, $K^2 * K^3 = K^5$. The *complement* $\overline{G}$ of G is the graph on V with edge set $[V]^2 \setminus E$. The *line graph* $L(G)$ of G is the graph on E in which $x, y \in E$ are adjacent as vertices if and only if they are adjacent as edges in G.

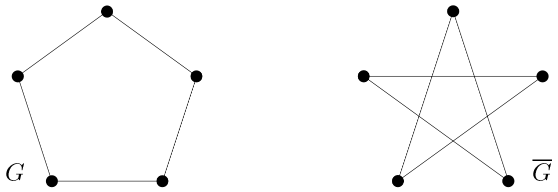

Fig. 1.1.4. A graph isomorphic to its complement

1.2 The degree of a vertex

$N(v)$

Let $G = (V, E)$ be a (non-empty) graph. The set of neighbours of a vertex v in G is denoted by $N_G(v)$, or briefly by $N(v)$.[1] More generally

---

[1] Here, as elsewhere, we drop the index referring to the underlying graph if the reference is clear.

[3.6.1] **Proposition 1.2.2.** *Every graph G with at least one edge has a subgraph H with $\delta(H) > \varepsilon(H) \geqslant \varepsilon(G)$.*

Proof. To construct H from G, let us try to delete vertices of small degree one by one, until only vertices of large degree remain. Up to which degree $d(v)$ can we afford to delete a vertex v, without lowering ε? Clearly, up to $d(v) = \varepsilon$: then the number of vertices decreases by 1 and the number of edges by at most ε, so the overall ratio ε of edges to vertices will not decrease.

Formally, we construct a sequence $G = G_0 \supseteq G_1 \supseteq \ldots$ of induced subgraphs of G as follows. If G_i has a vertex v_i of degree $d(v_i) \leqslant \varepsilon(G_i)$, we let $G_{i+1} := G_i - v_i$; if not, we terminate our sequence and set $H := G_i$. By the choices of v_i we have $\varepsilon(G_{i+1}) \geqslant \varepsilon(G_i)$ for all i, and hence $\varepsilon(H) \geqslant \varepsilon(G)$.

What else can we say about the graph H? Since $\varepsilon(K^1) = 0 < \varepsilon(G)$, none of the graphs in our sequence is trivial, so in particular $H \neq \emptyset$. The fact that H has no vertex suitable for deletion thus implies $\delta(H) > \varepsilon(H)$, as claimed. □

1.3 Paths and cycles

path A *path* is a non-empty graph $P = (V, E)$ of the form

$$V = \{ x_0, x_1, \ldots, x_k \} \qquad E = \{ x_0x_1, x_1x_2, \ldots, x_{k-1}x_k \},$$

where the x_i are all distinct. The vertices x_0 and x_k are *linked* by P and are called its *ends*; the vertices $x_1, \ldots, x_{k-1}$ are the *inner* vertices of P. The number of edges of a path is its *length*, and the path of length k is denoted by P^k. Note that k is allowed to be zero; thus, $P^0 = K^1$.

length
P^k

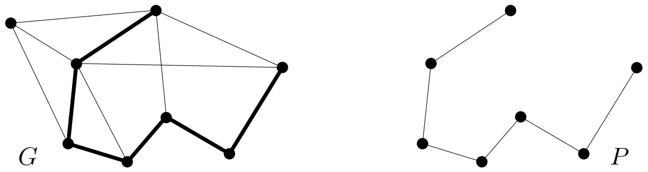

Fig. 1.3.1. A path $P = P^6$ in G

We often refer to a path by the natural sequence of its vertices,[3] writing, say, $P = x_0x_1 \ldots x_k$ and calling P a path *from* x_0 *to* x_k (as well as *between* x_0 and x_k).

---

[3] More precisely, by one of the two natural sequences: $x_0 \ldots x_k$ and $x_k \ldots x_0$ denote the same path. Still, it often helps to fix one of these two orderings of $V(P)$ notationally: we may then speak of things like the 'first' vertex on P with a certain property, etc.

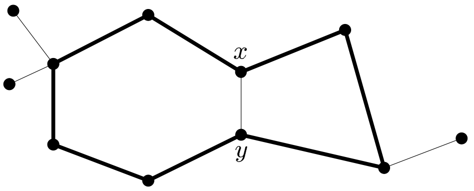

Fig. 1.3.3. A cycle C^8 with chord xy, and induced cycles C^6, C^4

If a graph has large minimum degree, it contains long paths and cycles:

[3.6.1] **Proposition 1.3.1.** *Every graph G contains a path of length $\delta(G)$ and a cycle of length at least $\delta(G)+1$ (provided that $\delta(G) \geqslant 2$).*

Proof. Let $x_0 \dots x_k$ be a longest path in G. Then all the neighbours of x_k lie on this path (Fig. 1.3.4). Hence $k \geqslant d(x_k) \geqslant \delta(G)$. If $i < k$ is minimal with $x_i x_k \in E(G)$, then $x_i \dots x_k x_i$ is a cycle of length at least $\delta(G)+1$. □

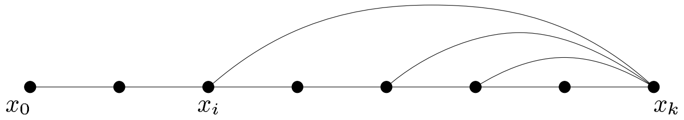

Fig. 1.3.4. A longest path $x_0 \dots x_k$, and the neighbours of x_k

Minimum degree and girth, on the other hand, are not related (unless we fix the number of vertices): as we shall see in Chapter 11, there are graphs combining arbitrarily large minimum degree with arbitrarily large girth.

distance $d_G(x,y)$

diameter diam(G)

The *distance* $d_G(x,y)$ in G of two vertices x, y is the length of a shortest x–y path in G; if no such path exists, we set $d(x,y) := \infty$. The greatest distance between any two vertices in G is the *diameter* of G, denoted by $\mathrm{diam}(G)$. Diameter and girth are, of course, related:

Proposition 1.3.2. *Every graph G containing a cycle satisfies $g(G) \leqslant 2\,\mathrm{diam}(G)+1$.*

Proof. Let C be a shortest cycle in G. If $g(G) \geqslant 2\,\mathrm{diam}(G)+2$, then C has two vertices whose distance in C is at least $\mathrm{diam}(G)+1$. In G, these vertices have a lesser distance; any shortest path P between them is therefore not a subgraph of C. Thus, P contains a C-path xPy. Together with the shorter of the two x–y paths in C, this path xPy forms a shorter cycle than C, a contradiction. □

Proof. Pick any vertex as v_1, and assume inductively that $v_1, \dots, v_i$ have been chosen for some $i < |G|$. Now pick a vertex $v \in G - G_i$. As G is connected, it contains a v–v_1 path P. Choose as v_{i+1} the last vertex of P in $G - G_i$; then v_{i+1} has a neighbour in G_i. The connectedness of every G_i follows by induction on i. □

component

Let $G = (V, E)$ be a graph. A maximal connected subgraph of G is called a *component* of G. Note that a component, being connected, is always non-empty; the empty graph, therefore, has no components.

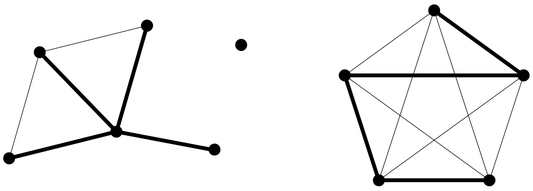

Fig. 1.4.1. A graph with three components, and a minimal spanning connected subgraph in each component

separate

cutvertex

bridge

If $A, B \subseteq V$ and $X \subseteq V \cup E$ are such that every A–B path in G contains a vertex or an edge from X, we say that X *separates* the sets A and B in G. This implies in particular that $A \cap B \subseteq X$. More generally we say that X *separates* G, and call X a *separating set* in G, if X separates two vertices of $G - X$ in G. A vertex which separates two other vertices of the same component is a *cutvertex*, and an edge separating its ends is a *bridge*. Thus, the bridges in a graph are precisely those edges that do not lie on any cycle.

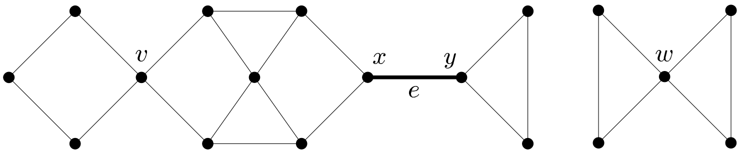

Fig. 1.4.2. A graph with cutvertices v, x, y, w and bridge $e = xy$

k-connected

G is called *k-connected* (for $k \in \mathbb{N}$) if $|G| > k$ and $G - X$ is connected for every set $X \subseteq V$ with $|X| < k$. In other words, no two vertices of G are separated by fewer than k other vertices. Every (non-empty) graph is 0-connected, and the 1-connected graphs are precisely the non-trivial connected graphs. The greatest integer k such that G is k-connected is the *connectivity* $\kappa(G)$ of G. Thus, $\kappa(G) = 0$ if and only if G is disconnected or a K^1, and $\kappa(K^n) = n - 1$ for all $n \geqslant 1$.

connectivity $\kappa(G)$

ℓ-edge-connected

If $|G| > 1$ and $G - F$ is connected for every set $F \subseteq E$ of fewer than ℓ edges, then G is called *ℓ-edge-connected.* The greatest integer ℓ

(completing the proof): if neither does, we have

$$\|G_i\| \leqslant (2k-3)(|G_i|-k+1)$$

for $i = 1, 2$, and hence

$$\begin{aligned} m &\leqslant \|G_1\| + \|G_2\| \\ &\leqslant (2k-3)\big(|G_1|+|G_2|-2k+2\big) \\ &\leqslant (2k-3)(n-k+1) \qquad (\text{by } |G_1 \cap G_2| \leqslant k-1) \end{aligned}$$

contradicting (ii). □

1.5 Trees and forests

forest *tree* *leaf*

An *acyclic* graph, one not containing any cycles, is called a *forest*. A connected forest is called a *tree*. (Thus, a forest is a graph whose components are trees.) The vertices of degree 1 in a tree are its *leaves*. Every nontrivial tree has at least two leaves—take, for example, the ends of a longest path. This little fact often comes in handy, especially in induction proofs about trees: if we remove a leaf from a tree, what remains is still a tree.

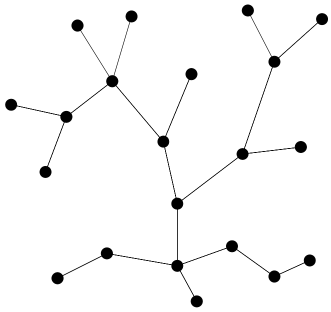

Fig. 1.5.1. A tree

[1.6.1] [1.9.6] [4.2.7]

Theorem 1.5.1. *The following assertions are equivalent for a graph T:*

(i) *T is a tree;*

(ii) *any two vertices of T are linked by a unique path in T;*

(iii) *T is minimally connected, i.e. T is connected but $T-e$ is disconnected for every edge $e \in T$;*

(iv) *T is maximally acyclic, i.e. T contains no cycle but $T+xy$ does, for any two non-adjacent vertices $x, y \in T$.* □

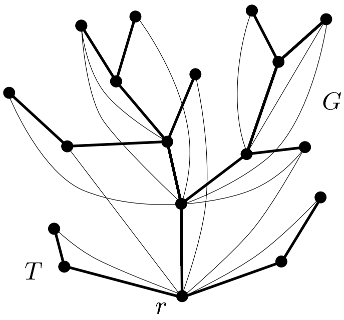

Fig. 1.5.2. A depth-first search tree with root r

Normal spanning trees provide a simple but powerful structural tool in graph theory. And they always exist:

[6.5.3] **Proposition 1.5.5.** *Every connected graph contains a normal spanning tree, with any specified vertex as its root.*

Proof. Let G be a connected graph and $r \in G$ any specified vertex. Let T be a maximal normal tree with root r in G; we show that $V(T) = V(G)$.

Suppose not, and let C be a component of $G - T$. As T is normal, $N(C)$ is a chain in T. Let x be its greatest element, and let $y \in C$ be adjacent to x. Let T' be the tree obtained from T by joining y to x; the tree-order of T' then extends that of T. We shall derive a contradiction by showing that T' is also normal in G.

Let P be a T'-path in G. If the ends of P both lie in T, then they are comparable in the tree-order of T (and hence in that of T'), because then P is also a T-path and T is normal in G by assumption. If not, then y is one end of P, so P lies in C except for its other end z, which lies in $N(C)$. Then $z \leqslant x$, by the choice of x. For our proof that y and z are comparable it thus suffices to show that $x < y$, i.e. that $x \in rT'y$. This, however, is clear since y is a leaf of T' with neighbour x. □

1.6 Bipartite graphs

r-partite

Let $r \geqslant 2$ be an integer. A graph $G = (V, E)$ is called *r-partite* if V admits a partition into r classes such that every edge has its ends in different classes: vertices in the same partition class must not be adjacent. Instead of '2-partite' one usually says *bipartite*.

bipartite

complete r-partite

An r-partite graph in which *every* two vertices from different partition classes are adjacent is called *complete*; the complete r-partite graphs for all r together are the *complete multipartite* graphs. The

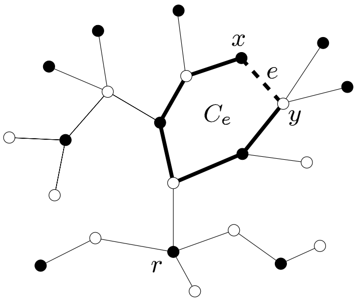

Fig. 1.6.3. The cycle C_e in $T+e$

1.7 Contraction and minors

In Section 1.1 we saw two fundamental containment relations between graphs: the subgraph relation, and the 'induced subgraph' relation. In this section we meet another: the minor relation.

G/e *contraction* Let $e = xy$ be an edge of a graph $G = (V, E)$. By G/e we denote the graph obtained from G by *contracting* the edge e into a new vertex v_e, which becomes adjacent to all the former neighbours of x and of y. Formally, G/e is a graph (V', E') with vertex set $V' := (V \setminus \{x, y\}) \cup \{v_e\}$ (where v_e is the 'new' vertex, i.e. $v_e \notin V \cup E$) and edge set

v_e

$$\begin{aligned} E' := &\Big\{ vw \in E \mid \{v, w\} \cap \{x, y\} = \emptyset \Big\} \\ &\cup \Big\{ v_e w \mid xw \in E \setminus \{e\} \text{ or } yw \in E \setminus \{e\} \Big\}. \end{aligned}$$

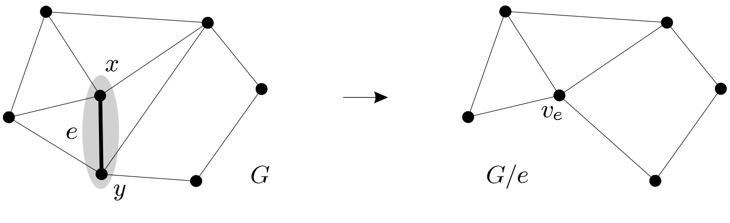

Fig. 1.7.1. Contracting the edge $e = xy$

More generally, if X is another graph and $\{V_x \mid x \in V(X)\}$ is a partition of V into connected subsets such that, for any two vertices $x, y \in X$, there is a V_x–V_y edge in G if and only if $xy \in E(X)$, we call G an MX and write[6] $G = MX$ (Fig. 1.7.2). The sets V_x are the *branch sets* of this MX. Intuitively, we obtain X from G by contracting every

MX *branch sets*

---

[6] Thus formally, the expression MX—where M stands for 'minor'; see below—refers to a whole class of graphs, and $G = MX$ means (with slight abuse of notation) that G belongs to this class.

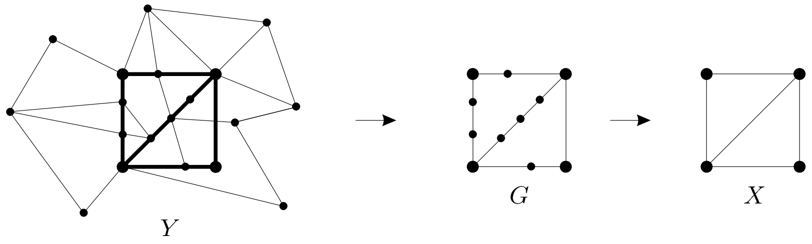

Fig. 1.7.3. $Y \supseteq G = TX$, so X is a topological minor of Y

branch vertices

If $G = TX$, we view $V(X)$ as a subset of $V(G)$ and call these vertices the *branch vertices* of G; the other vertices of G are its *subdividing vertices*. Thus, all subdividing vertices have degree 2, while the branch vertices retain their degree from X.

[4.4.2]
[8.3.1]

Proposition 1.7.2.

(i) *Every TX is also an MX (Fig. 1.7.4); thus, every topological minor of a graph is also its (ordinary) minor.*

(ii) *If $\Delta(X) \leqslant 3$, then every MX contains a TX; thus, every minor with maximum degree at most 3 of a graph is also its topological minor.* □

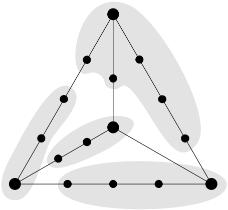

Fig. 1.7.4. A subdivision of K^4 viewed as an MK^4

[12.4.1]

Proposition 1.7.3. *The minor relation $\preccurlyeq$ and the topological-minor relation are partial orderings on the class of finite graphs, i.e. they are reflexive, antisymmetric and transitive.* □

1.8 Euler tours

Any mathematician who happens to find himself in the East Prussian city of Königsberg (and in the 18th century) will lose no time to follow the great Leonhard Euler's example and inquire about a round trip through

Conversely, let G be a connected graph with all degrees even, and let

$$W = v_0 e_0 \dots e_{\ell-1} v_\ell$$

be a longest walk in G using no edge more than once. Since W cannot be extended, it already contains all the edges at v_ℓ. By assumption, the number of such edges is even. Hence $v_\ell = v_0$, so W is a closed walk.

Suppose W is not an Euler tour. Then G has an edge e outside W but incident with a vertex of W, say $e = uv_i$. (Here we use the connectedness of G, as in the proof of Proposition 1.4.1.) Then the walk

$$u e v_i e_i \dots e_{\ell-1} v_\ell e_0 \dots e_{i-1} v_i$$

is longer than W, a contradiction. □

1.9 Some linear algebra

vertex space $\mathcal{V}(G)$

Let $G = (V, E)$ be a graph with n vertices and m edges, say $V = \{ v_1, \dots, v_n \}$ and $E = \{ e_1, \dots, e_m \}$. The *vertex space* $\mathcal{V}(G)$ of G is the vector space over the 2-element field $\mathbb{F}_2 = \{ 0, 1 \}$ of all functions $V \to \mathbb{F}_2$. Every element of $\mathcal{V}(G)$ corresponds naturally to a subset of V, the set of those vertices to which it assigns a 1, and every subset of V is uniquely represented in $\mathcal{V}(G)$ by its indicator function. We may thus think of $\mathcal{V}(G)$ as the power set of V made into a vector space: the sum $U + U'$ of two vertex sets $U, U' \subseteq V$ is their symmetric difference (why?), and $U = -U$ for all $U \subseteq V$. The zero in $\mathcal{V}(G)$, viewed in this way, is the empty (vertex) set $\emptyset$. Since $\{ \{ v_1 \}, \dots, \{ v_n \} \}$ is a basis of $\mathcal{V}(G)$, its *standard basis*, we have $\dim \mathcal{V}(G) = n$.

$+$

edge space $\mathcal{E}(G)$

standard basis

In the same way as above, the functions $E \to \mathbb{F}_2$ form the *edge space* $\mathcal{E}(G)$ of G: its elements are the subsets of E, vector addition amounts to symmetric difference, $\emptyset \subseteq E$ is the zero, and $F = -F$ for all $F \subseteq E$. As before, $\{ \{ e_1 \}, \dots, \{ e_m \} \}$ is the *standard basis* of $\mathcal{E}(G)$, and $\dim \mathcal{E}(G) = m$.

Since the edges of a graph carry its essential structure, we shall mostly be concerned with the edge space. Given two edge sets $F, F' \in \mathcal{E}(G)$ and their coefficients $\lambda_1, \dots, \lambda_m$ and $\lambda'_1, \dots, \lambda'_m$ with respect to the standard basis, we write

$\langle F, F' \rangle$

$$\langle F, F' \rangle := \lambda_1 \lambda'_1 + \dots + \lambda_m \lambda'_m \in \mathbb{F}_2 .$$

Note that $\langle F, F' \rangle = 0$ may hold even when $F = F' \neq \emptyset$: indeed, $\langle F, F' \rangle = 0$ if and only if F and F' have an even number of edges

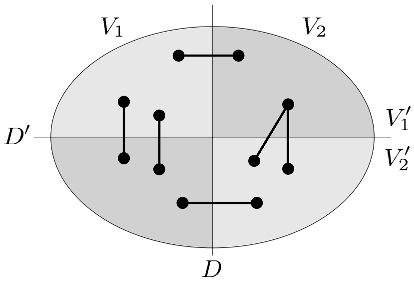

Fig. 1.9.1. Cut edges in $D + D'$

sets form another partition of V. Hence $D + D' \in \mathcal{C}^*$, and $\mathcal{C}^*$ is indeed a subspace of $\mathcal{E}(G)$.

Our second assertion, that the cuts $E(v)$ generate all of $\mathcal{C}^*$, follows from the fact that every edge $xy \in G$ lies in exactly two such cuts (in $E(x)$ and in $E(y)$); thus every partition $\{ V_1, V_2 \}$ of V satisfies $E(V_1, V_2) = \sum_{v \in V_1} E(v)$. □

cut space $\mathcal{C}^*(G)$

The subspace $\mathcal{C}^* =: \mathcal{C}^*(G)$ of $\mathcal{E}(G)$ from Proposition 1.9.3 will be called the *cut space* of G. It is not difficult to find among the cuts $E(v)$ an explicit basis for $\mathcal{C}^*(G)$, and thus to determine its dimension (exercise); together with Theorem 1.9.5 this yields an independent proof of Theorem 1.9.6.

The following lemma will be useful when we study the duality of plane graphs in Chapter 4.6:

[4.6.2] **Lemma 1.9.4.** *The minimal cuts in a connected graph generate its entire cut space.*

Proof. Note first that a cut in a connected graph $G = (V, E)$ is minimal if and only if both sets in the corresponding partition of V are connected in G. Now consider any connected subgraph $C \subseteq G$. If D is a component of $G - C$, then also $G - D$ is connected (Fig. 1.9.2); the edges between D

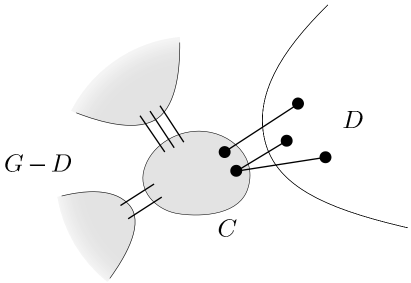

Fig. 1.9.2. $G - D$ is connected, and $E(C, D)$ a minimal cut

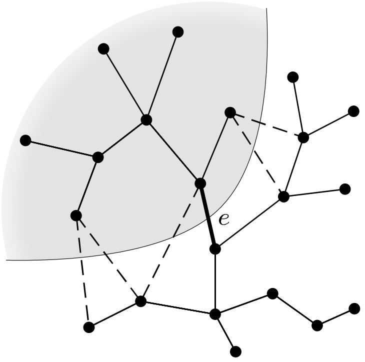

Fig. 1.9.3. The cut D_e

incidence matrix

The *incidence matrix* $B = (b_{ij})_{n \times m}$ of a graph $G = (V, E)$ with $V = \{ v_1, \dots, v_n \}$ and $E = \{ e_1, \dots, e_m \}$ is defined over $\mathbb{F}_2$ by

$$b_{ij} := \begin{cases} 1 & \text{if } v_i \in e_j \\ 0 & \text{otherwise.} \end{cases}$$

As usual, let B^t denote the transpose of B. Then B and B^t define linear maps $B\colon \mathcal{E}(G) \to \mathcal{V}(G)$ and $B^t\colon \mathcal{V}(G) \to \mathcal{E}(G)$ with respect to the standard bases.

Proposition 1.9.7.

(i) *The kernel of B is $\mathcal{C}(G)$.*

(ii) *The image of B^t is $\mathcal{C}^*(G)$.* □

adjacency matrix

The *adjacency matrix* $A = (a_{ij})_{n \times n}$ of G is defined by

$$a_{ij} := \begin{cases} 1 & \text{if } v_i v_j \in E \\ 0 & \text{otherwise.} \end{cases}$$

Our last proposition establishes a simple connection between A and B (now viewed as real matrices). Let D denote the real diagonal matrix $(d_{ij})_{n \times n}$ with $d_{ii} = d(v_i)$ and $d_{ij} = 0$ otherwise.

Proposition 1.9.8. *$BB^t = A + D$.* □

map $e' \mapsto \{\,\text{init}(e'), \text{ter}(e')\,\}$ of G has to be adjusted to the new vertex set in G/e. The notion of a minor adapts to multigraphs accordingly.

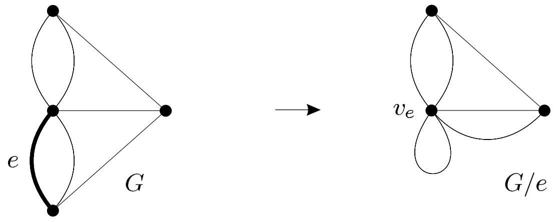

Fig. 1.10.1. Contracting the edge e in the multigraph corresponding to Fig. 1.8.1

Finally, it should be pointed out that authors who usually work with multigraphs tend to call them graphs; in their terminology, our graphs would be called simple graphs.

Exercises

1.$^-$ What is the number of edges in a K^n?

2. Let $d \in \mathbb{N}$ and $V := \{\,0, 1\,\}^d$; thus, V is the set of all 0–1 sequences of length d. The graph on V in which two such sequences form an edge if and only if they differ in exactly one position is called the *d-dimensional cube*. Determine the average degree, number of edges, diameter, girth and circumference of this graph.

   (Hint for circumference. Induction on d.)

3. Let G be a graph containing a cycle C, and assume that G contains a path of length at least k between two vertices of C. Show that G contains a cycle of length at least $\sqrt{k}$. Is this best possible?

4.$^-$ Is the bound in Proposition 1.3.2 best possible?

5. Show that $\text{rad}(G) \leqslant \text{diam}(G) \leqslant 2\,\text{rad}(G)$ for every graph G.

6.$^+$ Assuming that $d \geqslant 2$ and $k \geqslant 3$, improve the bound in Proposition 1.3.3 to d^k.

7.$^-$ Show that the components of a graph partition its vertex set. (In other words, show that every vertex belongs to exactly one component.)

8.$^-$ Show that every 2-connected graph contains a cycle.

9. (i)$^-$ Determine $\kappa(G)$ and $\lambda(G)$ for $G = P^k, C^k, K^k, K_{m,n}$ $(k, m, n \geqslant 3)$.

   (ii)$^+$ Determine the connectivity of the n-dimensional cube (defined in Exercise 2).

   (Hint for (ii). Induction on n.)

10. Show that $\kappa(G) \leqslant \lambda(G) \leqslant \delta(G)$ for every non-trivial graph G.

25. Given a graph G, find among all cuts of the form $E(v)$ a basis for the cut space of G.

26. Prove that the cycles and the cuts in a graph together generate its entire edge space, or find a counterexample.

27. Give a direct proof of the fact that the cycles C_e defined in the proof of Theorem 1.9.6 generate the cycle space.

28. Give a direct proof of the fact that the cuts D_e defined in the proof of Theorem 1.9.6 generate the cut space.

29. What are the dimensions of the cycle and the cut space of a graph with k components?

Notes

The terminology used in this book is mostly standard. Alternatives do exist, of course, and some of these are stated when a concept is first defined. There is one small point where our notation deviates slightly from standard usage. Whereas complete graphs, paths, cycles etc. of given order are mostly denoted by K_n, P_k, C_ℓ and so on, we use superscripts instead of subscripts. This has the advantage of leaving the variables K, P, C etc. free for ad-hoc use: we may now enumerate components as $C_1, C_2, \ldots$, speak of paths $P_1, \ldots, P_k$, and so on—without any danger of confusion.

Theorem[10] 1.4.2 is due to W. Mader, Existenz n-fach zusammenhängender Teilgraphen in Graphen genügend großer Kantendichte, *Abh. Math. Sem. Univ. Hamburg* **37** (1972) 86–97. Theorem 1.8.1 is from L. Euler, Solutio problematis ad geometriam situs pertinentis, *Comment. Acad. Sci. I. Petropolitanae* **8** (1736), 128–140.

Of the large subject of algebraic methods in graph theory, Section 1.9 does not claim to convey an adequate impression. The standard monograph here is N.L. Biggs, *Algebraic Graph Theory* (2nd edn.), Cambridge University Press 1993. Another comprehensive account is given by C.D. Godsil & G.F. Royle, *Algebraic Graph Theory*, in preparation. Surveys on the use of algebraic methods can also be found in the *Handbook of Combinatorics* (R.L. Graham, M. Grötschel & L. Lovász, eds.), North-Holland 1995.

---

[10] In the interest of readability, the end-of-chapter notes in this book give references only for Theorems, and only in cases where these references cannot be found in a monograph or survey cited for that chapter.

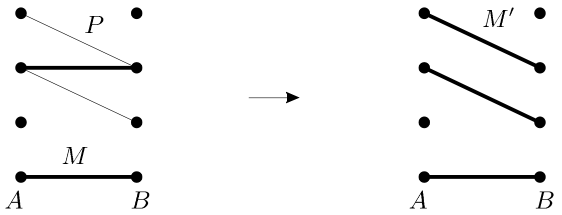

Fig. 2.1.1. Augmenting the matching M by the alternating path P

matching (consider the edges at a given vertex), and the set of matched vertices is increased by two, the ends of P.

Alternating paths play an important role in the practical search for large matchings. In fact, if we start with any matching and keep applying augmenting paths until no further such improvement is possible, the matching obtained will always be an optimal one, a matching with the largest possible number of edges (Exercise 1). The algorithmic problem of finding such matchings thus reduces to that of finding augmenting paths—which is an interesting and accessible algorithmic problem.

vertex cover

Our first theorem characterizes the maximal cardinality of a matching in G by a kind of duality condition. Let us call a set $U \subseteq V$ a *cover* of E (or a *vertex cover* of G) if every edge of G is incident with a vertex in U.

Theorem 2.1.1. (König 1931)
The maximum cardinality of a matching in G is equal to the minimum cardinality of a vertex cover.

M

U

Proof. Let M be a matching in G of maximum cardinality. From every edge in M let us choose one of its ends: its end in B if some alternating path ends in that vertex, and its end in A otherwise (Fig. 2.1.2). We shall prove that the set U of these $|M|$ vertices covers G; since any vertex cover of G must cover M, there can be none with fewer than $|M|$ vertices, and so the theorem will follow.

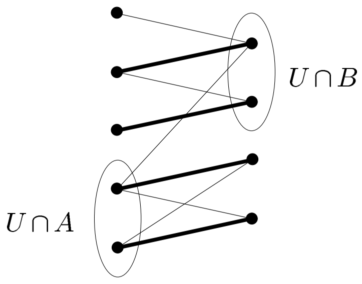

Fig. 2.1.2. The vertex cover U

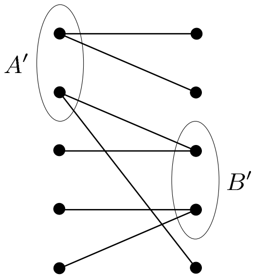

Fig. 2.1.3. A cover by fewer than $|A|$ vertices

M **Second proof.** Consider a matching M of G that leaves a vertex of A unmatched; we shall construct an augmenting path with respect to M. Let $a_0, b_1, a_1, b_2, a_2, \dots$ be a maximal sequence of distinct vertices $a_i \in A$ and $b_i \in B$ satisfying the following conditions for all $i \geqslant 1$ (Fig. 2.1.4):

(i) a_0 is unmatched;

$f(i)$ (ii) b_i is adjacent to some vertex $a_{f(i)} \in \{ a_0, \dots, a_{i-1} \}$;

(iii) $a_i b_i \in M$.

By the marriage condition, our sequence cannot end in a vertex of A: the i vertices $a_0, \dots, a_{i-1}$ together have at least i neighbours in B, so we can always find a new vertex $b_i \neq b_1, \dots, b_{i-1}$ that satisfies (ii). Let
k $b_k \in B$ be the last vertex of the sequence. By (i)–(iii),

P
$$P := b_k a_{f(k)} b_{f(k)} a_{f^2(k)} b_{f^2(k)} a_{f^3(k)} \dots a_{f^r(k)}$$

with $f^r(k) = 0$ is an alternating path.

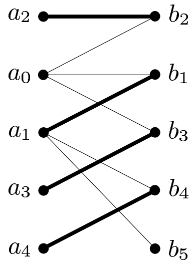

Fig. 2.1.4. Proving the marriage theorem by alternating paths

What is it that prevents us from extending our sequence further? If b_k is matched, say to a, we can indeed extend it by setting $a_k := a$, unless $a = a_i$ with $0 < i < k$, in which case (iii) would imply $b_k = b_i$ with a contradiction. So b_k is unmatched, and hence P is an augmenting path between a_0 and b_k. □

(1.8.1) *Proof*. Let G be any $2k$-regular graph ($k \geqslant 1$), without loss of generality connected. By Theorem 1.8.1, G contains an Euler tour $v_0 e_0 \dots e_{\ell-1} v_\ell$, with $v_\ell = v_0$. We replace every vertex v by a pair (v^-, v^+), and every edge $e_i = v_i v_{i+1}$ by the edge $v_i^+ v_{i+1}^-$ (Fig. 2.1.5). The resulting bipartite graph G' is k-regular, so by Corollary 2.1.4 it has a 1-factor. Collapsing every vertex pair (v^-, v^+) back into a single vertex v, we turn this 1-factor of G' into a 2-factor of G. $\square$

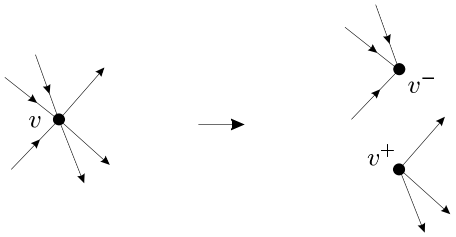

Fig. 2.1.5. Splitting vertices in the proof of Corollary 2.1.5

2.2 Matching in general graphs

$\mathcal{C}_G$
$q(G)$

Tutte's condition

Given a graph G, let us denote by $\mathcal{C}_G$ the set of its components, and by $q(G)$ the number of its *odd components*, those of odd order. If G has a 1-factor, then clearly

$$q(G-S) \leqslant |S| \qquad \text{for all } S \subseteq V(G),$$

since every odd component of $G - S$ will send a factor edge to S.

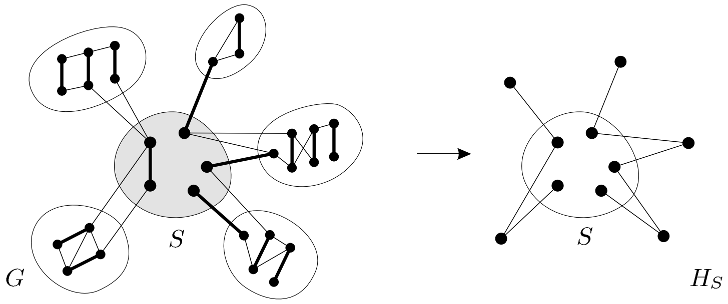

Fig. 2.2.1. Tutte's condition $q(G-S) \leqslant |S|$ for $q = 3$, and the contracted graph H_S from Theorem 2.2.3.

Again, this obvious necessary condition for the existence of a 1-factor is also sufficient:

Corollary 2.2.2. (Petersen 1891)
Every bridgeless cubic graph has a 1-factor.

Proof. We show that any bridgeless cubic graph G satisfies Tutte's condition. Let $S \subseteq V(G)$ be given, and consider an odd component C of $G - S$. Since G is cubic, the degrees (in G) of the vertices in C sum to an odd number, but only an even part of this sum arises from edges of C. So G has an odd number of S–C edges, and therefore has at least 3 such edges (since G has no bridge). The total number of edges between S and $G - S$ thus is at least $3q(G-S)$. But it is also at most $3|S|$, because G is cubic. Hence $q(G-S) \leqslant |S|$, as required. □

In order to shed a little more light on the techniques used in matching theory, we now give a second proof of Tutte's theorem. In fact, we shall prove a slightly stronger result, a result that places a structure interesting from the matching point of view on an *arbitrary* graph. If the graph happens to satisfy the condition of Tutte's theorem, this structure will at once yield a 1-factor.

factor-critical

matchable

H_S

A graph $G = (V, E)$ is called *factor-critical* if $G \neq \emptyset$ and $G - v$ has a 1-factor for every vertex $v \in G$. Then G itself has no 1-factor, because it has odd order. We call a vertex set $S \subseteq V$ *matchable to* $G - S$ if the (bipartite[1]) graph H_S, which arises from G by contracting the components $C \in \mathcal{C}_{G-S}$ to single vertices and deleting all the edges inside S, contains a matching of S. (Formally, H_S is the graph with vertex set $S \cup \mathcal{C}_{G-S}$ and edge set $\{\, sC \mid \exists\, c \in C : sc \in E \,\}$; see Fig. 2.2.1.)

Theorem 2.2.3. *Every graph $G = (V, E)$ contains a vertex set S with the following two properties:*

(i) *S is matchable to $G - S$;*

(ii) *every component of $G - S$ is factor-critical.*

Given any such set S, the graph G contains a 1-factor if and only if $|S| = |\mathcal{C}_{G-S}|$.

For any given G, the assertion of Tutte's theorem follows easily from this result. Indeed, by (i) and (ii) we have $|S| \leqslant |\mathcal{C}_{G-S}| = q(G-S)$ (since factor-critical graphs have odd order); thus Tutte's condition of $q(G-S) \leqslant |S|$ implies $|S| = |\mathcal{C}_{G-S}|$, and the existence of a 1-factor follows from the last statement of Theorem 2.2.3.

(2.1.3)

Proof of Theorem 2.2.3. Note first that the last assertion of the theorem follows at once from the assertions (i) and (ii): if G has a 1-factor, we have $q(G-S) \leqslant |S|$ and hence $|S| = |\mathcal{C}_{G-S}|$ as above;

---

[1] except for the—permitted—case that S or $\mathcal{C}_{G-S}$ is empty

H

equality, this implies that $\mathcal{C}$ too is non-empty. We now apply Corollary 2.1.3 to $H := H_S$, but 'backwards', i.e. with $A := \mathcal{C}$. Given $\mathcal{C}' \subseteq \mathcal{C}$, set $S' := N_H(\mathcal{C}') \subseteq S$. Since every $C \in \mathcal{C}'$ is an odd component also of $G - S'$, we have

$$|N_H(\mathcal{C}')| = |S'| \underset{(*)}{\geqslant} q(G - S') - d \geqslant |\mathcal{C}'| - d\,.$$

By Corollary 2.1.3, then, H contains a matching of cardinality

$$|\mathcal{C}| - d = q(G - S) - d = |S|\,,$$

which is therefore a matching of S. □

S
$\mathcal{C}$
$k_S, k_\mathcal{C}$

Let us consider once more the set S from Theorem 2.2.3, together with any matching M in G. As before, we write $\mathcal{C} := \mathcal{C}_{G-S}$. Let us denote by k_S the number of edges in M with at least one end in S, and by $k_\mathcal{C}$ the number of edges in M with both ends in $G - S$. Since each $C \in \mathcal{C}$ is odd, at least one of its vertices is not incident with an edge of the second type. Therefore every matching M satisfies

$$k_S \leqslant |S| \quad \text{and} \quad k_\mathcal{C} \leqslant \tfrac{1}{2}\Big(|V| - |S| - |\mathcal{C}|\Big)\,. \tag{1}$$

M_0

Moreover, G contains a matching M_0 with equality in both cases: first choose $|S|$ edges between S and $\bigcup \mathcal{C}$ according to (i), and then use (ii) to find a suitable set of $\frac{1}{2}\big(|C| - 1\big)$ edges in every component $C \in \mathcal{C}$. This matching M_0 thus has exactly

$$|M_0| = |S| + \tfrac{1}{2}\Big(|V| - |S| - |\mathcal{C}|\Big) \tag{2}$$

edges.

Now (1) and (2) together imply that *every* matching M of maximum cardinality satisfies both parts of (1) with equality: by $|M| \geqslant |M_0|$ and (2), M has at least $|S| + \frac{1}{2}\big(|V| - |S| - |\mathcal{C}|\big)$ edges, which implies by (1) that neither of the inequalities in (1) can be strict. But equality in (1), in turn, implies that M has the structure described above: by $k_S = |S|$, every vertex $s \in S$ is the end of an edge $st \in M$ with $t \in G - S$, and by $k_\mathcal{C} = \frac{1}{2}\big(|V| - |S| - |\mathcal{C}|\big)$ exactly $\frac{1}{2}\big(|C| - 1\big)$ edges of M lie in C, for every $C \in \mathcal{C}$. Finally, since these latter edges miss only one vertex in each C, the ends t of the edges st above lie in different components C for different s.

The seemingly technical Theorem 2.2.3 thus hides a wealth of structural information: it contains the essence of a detailed description of all maximum-cardinality matchings in all graphs.[2]

---

[2] A reference to the full statement of this structural result, known as the *Gallai-Edmonds matching theorem*, is given in the notes at the end of this chapter.

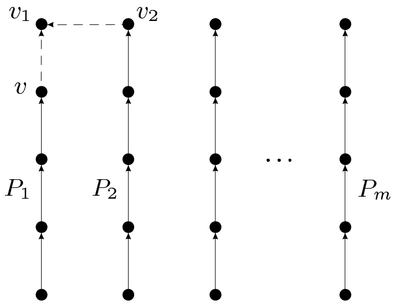

Fig. 2.3.1. The path cover $\mathcal{P}'$ of G'

$P \in \mathcal{P}''$ ends in v_2 (but none in v), we replace P in $\mathcal{P}''$ by Pv_2v_1, again contradicting the minimality of $\mathcal{P}$. Hence $\{\, \text{ter}(P) \mid P \in \mathcal{P}'' \,\} \subseteq \{\, v_3, \dots, v_m \,\}$, and in particular $|\mathcal{P}''| \leqslant |\mathcal{P}| - 2$. But now $\mathcal{P}''$ and the trivial path $\{\, v_1 \,\}$ together form a path cover of G that contradicts the minimality of $\mathcal{P}$.

Hence $\mathcal{P}'$ is minimal, as claimed. By the induction hypothesis, $\{\, V(P) \mid P \in \mathcal{P}' \,\}$ has an independent set of representatives. But this is also a set of representatives for $\mathcal{P}$, and $(*)$ is proved. □

As a corollary to Theorem 2.3.1 we now deduce a classic result from the theory of partial orders. Recall that a subset of a partially ordered set $(P, \leqslant)$ is a *chain* in P if its elements are pairwise comparable; it is an *antichain* if they are pairwise incomparable.

chain
antichain

Corollary 2.3.2. (Dilworth 1950)
In every finite partially ordered set $(P, \leqslant)$*, the minimum number of chains covering* P *is equal to the maximum cardinality of an antichain in* P*.*

Proof. If A is an antichain in P of maximum cardinality, then clearly P cannot be covered by fewer than $|A|$ chains. The fact that $|A|$ chains will suffice follows from Theorem 2.3.1 applied to the directed graph on P with the edge set $\{\, (x, y) \mid x < y \,\}$. □

Exercises

1. Let M be a matching in a bipartite graph G. Show that if M is suboptimal, i.e. contains fewer edges than some other matching in G, then G contains an augmenting path with respect to M. Does this fact generalize to matchings in non-bipartite graphs?

   (Hint. Symmetric difference.)

18.$^+$ Find a partially ordered set that has no infinite antichain but cannot be covered by finitely many chains.

(Hint. $\mathbb{N} \times \mathbb{N}$.)

Notes

There is a very readable and comprehensive monograph about matching in finite graphs: L. Lovász & M.D. Plummer, *Matching Theory*, Annals of Discrete Math. **29**, North Holland 1986. All the references for the results in this chapter can be found there.

As we shall see in Chapter 3, König's Theorem of 1931 is no more than the bipartite case of a more general theorem due to Menger, of 1929. At the time, neither of these results was nearly as well known as Hall's marriage theorem, which was proved even later, in 1935. To this day, Hall's theorem remains one of the most applied graph-theoretic results. Its special case that both partition sets have the same size was proved implicitly already by Frobenius (1917) in a paper on determinants.

Our proof of Tutte's 1-factor theorem is based on a proof by Lovász (1975). Our extension of Tutte's theorem, Theorem 2.2.3 (including the informal discussion following it) is a lean version of a comprehensive structure theorem for matchings, due to Gallai (1964) and Edmonds (1965). See Lovász & Plummer for a detailed statement and discussion of this theorem.

Theorem 2.3.1 is due to T. Gallai & A.N. Milgram, Verallgemeinerung eines graphentheoretischen Satzes von Rédei, *Acta Sci. Math. (Szeged)* **21** (1960), 181–186.

determined fully by that of its components, however, it is not captured completely by the structure of its blocks: since the blocks need not be disjoint, the way they intersect defines another structure, giving a coarse picture of G as if viewed from a distance.

The following proposition describes this coarse structure of G as formed by its blocks. Let A denote the set of cutvertices of G, and $\mathcal{B}$ the set of its blocks. We then have a natural bipartite graph on $A \cup \mathcal{B}$ formed by the edges aB with $a \in B$. This *block graph* of G is shown in Figure 3.1.1.

block graph

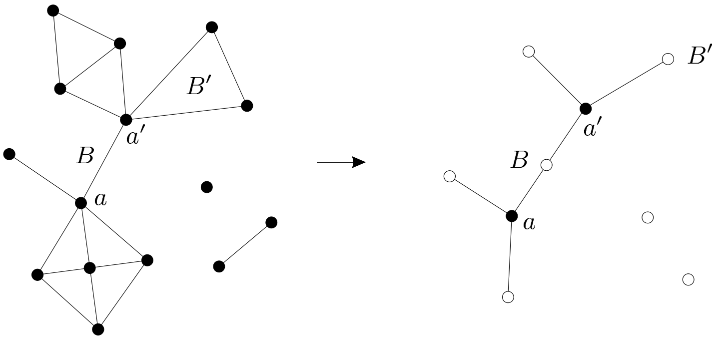

Fig. 3.1.1. A graph and its block graph

Proposition 3.1.1. *The block graph of a connected graph is a tree.* □

Proposition 3.1.1 reduces the structure of a given graph to that of its blocks. So what can we say about the blocks themselves? The following proposition gives a simple method by which, in principle, a list of all 2-connected graphs could be compiled:

[4.2.5] **Proposition 3.1.2.** *A graph is 2-connected if and only if it can be constructed from a cycle by successively adding H-paths to graphs H already constructed (Fig. 3.1.2).*

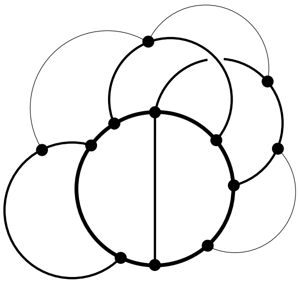

Fig. 3.1.2. The construction of 2-connected graphs

w

By assumption, G/zv is again not 3-connected, so again there is a vertex w such that $\{ z, v, w \}$ separates G, and as before every vertex in $\{ z, v, w \}$ has a neighbour in every component of $G - \{ z, v, w \}$.

As x and y are adjacent, $G - \{ z, v, w \}$ has a component D such that $D \cap \{ x, y \} = \emptyset$. Then every neighbour of v in D lies in C (since $v \in C$), so $D \cap C \neq \emptyset$ and hence $D \subsetneqq C$ by the choice of D. This contradicts the choice of xy, z and C. □

Theorem 3.2.2. (Tutte 1961)
A graph G is 3-connected if and only if there exists a sequence $G_0, \ldots, G_n$ of graphs with the following properties:

(i) *$G_0 = K^4$ and $G_n = G$;*

(ii) *G_{i+1} has an edge xy with $d(x), d(y) \geqslant 3$ and $G_i = G_{i+1}/xy$, for every $i < n$.*

Proof. If G is 3-connected, a sequence as in the theorem exists by Lemma 3.2.1. Note that all the graphs in this sequence are 3-connected.

xy
S
C_1, C_2

Conversely, let $G_0, \ldots, G_n$ be a sequence of graphs as stated; we show that if $G_i = G_{i+1}/xy$ is 3-connected then so is G_{i+1}, for every $i < n$. Suppose not, let S be a separating set of at most 2 vertices in G_{i+1}, and let C_1, C_2 be two components of $G_{i+1} - S$. As x and y are adjacent, we may assume that $\{ x, y \} \cap V(C_1) = \emptyset$ (Fig. 3.2.2). Then C_2 contains neither both vertices x, y nor a vertex $v \notin \{ x, y \}$: otherwise v_{xy} or v would be separated from C_1 in G_i by at most two vertices, a contradiction. But now C_2 contains only one vertex: either x or y. This contradicts our assumption of $d(x), d(y) \geqslant 3$. □

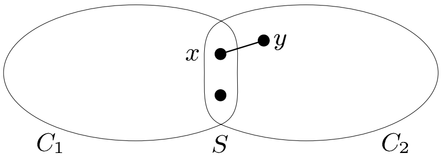

Fig. 3.2.2. The position of $xy \in G_{i+1}$ in the proof of Theorem 3.2.2

Theorem 3.2.2 is the essential core of a result of Tutte known as his *wheel theorem*.[1] Like Proposition 3.1.2 for 2-connected graphs, it enables us to construct all 3-connected graphs by a simple inductive process depending only on local information: starting with K^4, we pick a vertex v in a graph constructed already, split it into two adjacent vertices v', v'', and join these to the former neighbours of v as we please—provided only that v' and v'' each acquire at least 3 incident edges, and that every former neighbour of v becomes adjacent to at least one of v', v''.

---

wheel

[1] Graphs of the form $C^n * K^1$ are called *wheels*; thus, K^4 is the smallest wheel.

A fundamental triangle, $wxyw$ say, is clearly induced in G. If it separated G, then $\{ v_e, w \}$ would separate G', which contradicts the choice of e. This proves (1).

If $C \subseteq G$ is an induced cycle but not a fundamental triangle, then $C + C/e + D \in \{ \emptyset, \{e\} \}$ for some good $D \in \mathcal{C}(G)$. (2)

The gist of (2) is that, in terms of 'generatability', C and C/e differ only a little: after the addition of a permissible error term D, at most in the edge e. In which other edges, then, can C and C/e differ? Clearly at most in the two edges $e_u = uv_e$ and $e_w = v_e w$ incident with v_e in C/e; cf. Fig. 3.2.3. But these differences between the edge sets of C/e and C are levelled out precisely by adding the corresponding fundamental triangles uxy and xyw (which are basic by (1)). Indeed, let D_u denote the triangle uxy if $e_u \notin C$ and $\emptyset$ otherwise, and let D_w denote xyw if $e_w \notin C$ and $\emptyset$ otherwise. Then $D := D_u + D_w$ satisfies (2) as desired.

Next, we show how to lift basic cycles of G' back to G:

For every basic cycle $C' \subseteq G'$ there exists a basic cycle $C = C(C') \subseteq G$ with $C/e = C'$. (3)

If $v_e \notin C'$, then (3) is satisfied with $C := C'$. So we assume that $v_e \in C'$. Let u and w be the two neighbours of v_e on C', and let P be the u–w path in C' avoiding v_e (Fig. 3.2.4). Then $P \subseteq G$.

u, w
P

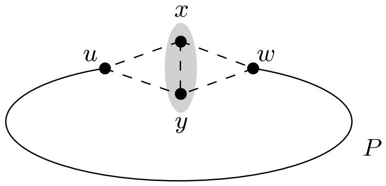

Fig. 3.2.4. The search for a basic cycle C with $C/e = C'$

C_x, C_y

We first assume that $\{ ux, uy, wx, wy \} \subseteq E(G)$, and consider (as candidates for C) the cycles $C_x := uPwxu$ and $C_y := uPwyu$. Both are induced cycles in G (because C' is induced in G'), and clearly $C_x/e = C_y/e = C'$. Moreover, neither of these cycles separates two vertices of $G - (V(P) \cup \{ x, y \})$ in G, since C' does not separate such vertices in G'. Thus, if C_x (say) is a separating cycle in G, then one of the components of $G - C_x$ consists just of y. Likewise, if C_y separates G then one of the arising components contains only x. However, this cannot happen for both C_x and C_y at once: otherwise $N_G(\{ x, y \}) \subseteq V(P)$ and hence $N_G(\{ x, y \}) = \{ u, w \}$ (since C' has no chord), which contradicts $\kappa(G) \geqslant 3$. Hence, at least one of C_x, C_y is a basic cycle in G.

3.3 Menger's theorem

The following theorem is one of the cornerstones of graph theory.

[3.6.2]
[8.1.2]
[12.3.9]
[12.4.4]
[12.4.5]

Theorem 3.3.1. (Menger 1927)
Let $G = (V, E)$ be a graph and $A, B \subseteq V$. Then the minimum number of vertices separating A from B in G is equal to the maximum number of disjoint A–B paths in G.

k

We offer three proofs. Whenever G, A, B are given as in the theorem, we denote by $k = k\,(G, A, B)$ the minimum number of vertices separating A from B in G. Clearly, G cannot contain more than k disjoint A–B paths; our task will be to show that k such paths exist.

First proof. We prove the following stronger statement:

> *If $\mathcal{P}$ is any set of fewer than k disjoint A–B paths in G then there is a set $\mathcal{Q}$ of $|\mathcal{P}| + 1$ disjoint A–B paths whose set of ends includes the set of ends of the paths in $\mathcal{P}$.*

Keeping G and A fixed, we let B vary and apply induction on $|G - B|$. Let R be an A–B path that avoids the (fewer than k) vertices of B that lie on a path in $\mathcal{P}$. If R avoids all the paths in $\mathcal{P}$, then $\mathcal{Q} := \mathcal{P} \cup \{R\}$ is as desired. (This will happen for $|G - B| = 0$ when all A–B paths are trivial.) If not, let x be the last vertex of R that lies on some $P \in \mathcal{P}$ (Fig. 3.3.1). Put $B' := B \cup V(xP \cup xR)$ and $\mathcal{P}' := \big(\mathcal{P} \setminus \{P\}\big) \cup \{Px\}$. Then $|\mathcal{P}'| = |\mathcal{P}|$ and $k(G, A, B') \geqslant k(G, A, B)$, so by the induction hypothesis there is a set $\mathcal{Q}'$ of $|\mathcal{P}| + 1$ disjoint A–B' paths whose ends include those of the paths in $\mathcal{P}'$. Then $\mathcal{Q}'$ contains a path Q ending in x, and a unique path Q' whose last vertex y is not among the last vertices of the paths in $\mathcal{P}'$. If $y \notin xP$, we let $\mathcal{Q}$ be obtained from $\mathcal{Q}'$ by adding xP to Q, and adding yR to Q' if $y \notin B$. Otherwise $y \in \mathring{x}P$, and we let $\mathcal{Q}$ be obtained from $\mathcal{Q}'$ by adding xR to Q and adding yP to Q'. □

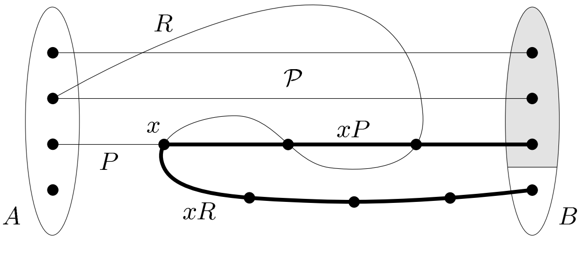

Fig. 3.3.1. Paths in the first proof of Menger's theorem

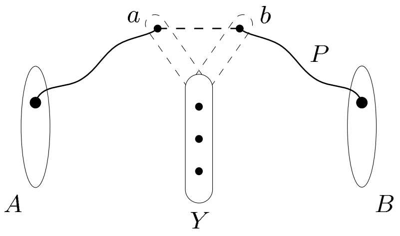

Fig. 3.3.3. Separating A from B in $G - ab$

Applied to a bipartite graph, Menger's theorem specializes to the assertion of König's theorem (2.1.1). For our third proof, we now adapt the alternating path proof of König's theorem to the more general set-up of Theorem 3.3.1. Let again G, A, B be given, and let $\mathcal{P}$ be a set of disjoint A–B paths in G. We write

$\mathcal{P}$

$$V[\mathcal{P}] := \bigcup \{ V(P) \mid P \in \mathcal{P} \}$$

$$E[\mathcal{P}] := \bigcup \{ E(P) \mid P \in \mathcal{P} \}.$$

A walk $W = x_0 e_0 x_1 e_1 \ldots e_{n-1} x_n$ in G with $e_i \neq e_j$ for $i \neq j$ is said to be *alternating* with respect to $\mathcal{P}$ if the following three conditions are satisfied for all $i < n$ (Fig. 3.3.4):

alternating walk

(i) if $e_i = e \in E[\mathcal{P}]$, then W traverses the edge e backwards, i.e. $x_{i+1} \in P \mathring{x}_i$ for some $P \in \mathcal{P}$;

(ii) if $x_i = x_j$ with $i \neq j$, then $x_i \in V[\mathcal{P}]$;

(iii) if $x_i \in V[\mathcal{P}]$, then $\{ e_{i-1}, e_i \} \cap E[\mathcal{P}] \neq \emptyset$.[2]

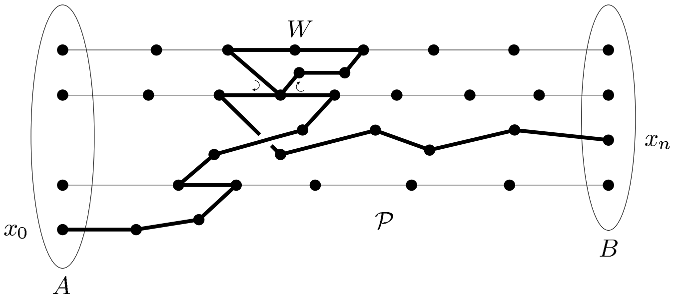

Fig. 3.3.4. An alternating walk from A to B

---

[2] For $i = 0$ we let $\{ e_{i-1}, e_i \} := \{ e_0 \}$.

y, P
x, W

let y be the last vertex of Q in $V[\mathcal{P}']$, let P be the path in $\mathcal{P}$ containing y, and let $x := x_P$. Finally, let W be an alternating walk from A_2 to x, as in the definition of x_P. By assumption, Q avoids X and hence x, so $y \in P\mathring{x}$, and $W \cup xPyQ$ is a walk from A_2 to B (Fig. 3.3.5). If this walk is alternating and ends in B_2, we are home: then G contains $|\mathcal{P}|+1$ disjoint A–B paths by Lemma 3.3.2, contrary to the maximality of $\mathcal{P}$.

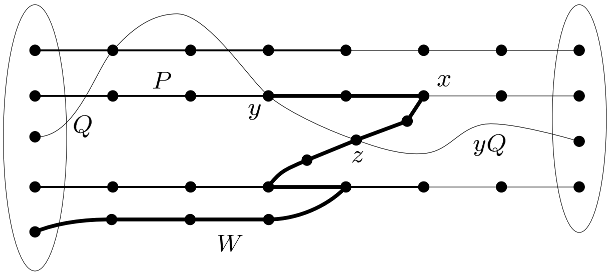

Fig. 3.3.5. Alternating walks in the third proof of Menger's theorem

x', W'

How could $W \cup xPyQ$ fail to be an alternating walk? For a start, W might already use an edge of xPy. But if x' is the first vertex of W on $xP\mathring{y}$, then $W' := Wx'Py$ is an alternating walk from A_2 to y. (By Wx' we mean the initial segment of W ending at the first occurrence of x' on W; from there onwards, W' follows P back to y.) Even our new walk $W'yQ$ need not yet be alternating: W' might still meet $\mathring{y}Q$. By definition of $\mathcal{P}'$ and W, however, and the choice of y on Q, we have

$$V(W') \cap V[\mathcal{P}] \subseteq V[\mathcal{P}'] \quad \text{and} \quad V(\mathring{y}Q) \cap V[\mathcal{P}'] = \emptyset\,.$$

Thus, W' and $\mathring{y}Q$ can meet only outside $\mathcal{P}$.

z
W''

If W' does indeed meet $\mathring{y}Q$, let z be the first vertex of W' on $\mathring{y}Q$. As z lies outside $V[\mathcal{P}]$, it occurs only once on W' (condition (ii)), and we let $W'' := W'zQ$. On the other hand if $W' \cap \mathring{y}Q = \emptyset$, we set $W'' := W' \cup yQ$. In both cases, W'' is alternating with respect to $\mathcal{P}'$, because W' is and $\mathring{y}Q$ avoids $V[\mathcal{P}']$. (Note that W'' satisfies condition (iii) at y in the second case, while in the first case (iii) is not applicable to z.) By definition of $\mathcal{P}'$, therefore, W'' avoids $V[\mathcal{P}] \smallsetminus V[\mathcal{P}']$; in particular, $V(\mathring{y}Q) \cap V[\mathcal{P}] = \emptyset$. Thus W'' is also alternating with respect to $\mathcal{P}$, and it ends in B_2. (Note that y cannot be the last vertex of W'', since $y \in P\mathring{x}$ and hence $y \notin B$.) Furthermore, W'' starts in A_2, because W does. We may therefore use W'' with Lemma 3.3.2 to obtain the desired contradiction to the maximality of $\mathcal{P}$. $\square$

3.4 Mader's theorem

In analogy to Menger's theorem we may consider the following question: given a graph G with an induced subgraph H, up to how many independent H-paths can we find in G?

In this section, we present without proof a deep theorem of Mader, which solves the above problem in a fashion similar to Menger's theorem. Again, the theorem says that an upper bound on the number of such paths that arises naturally from the size of certain separators is indeed attained by some suitable set of paths.

X F What could such an upper bound look like? Clearly, if $X \subseteq V(G-H)$ and $F \subseteq E(G-H)$ are such that every H-path in G has a vertex or an edge in $X \cup F$, then G cannot contain more than $|X \cup F|$ independent H-paths. Hence, the least cardinality of such a set $X \cup F$ is a natural upper bound for the maximum number of independent H-paths. (Note that every H-path meets $G-H$, because H is induced in G and edges of H do not count as H-paths.)

$\mathcal{C}_F$ ∂C In contrast to Menger's theorem, this bound can still be improved. Clearly, we may assume that no edge in F has an end in X: otherwise this edge would not be needed in the separator. Let $Y := V(G-H) \setminus X$, and denote by $\mathcal{C}_F$ the set of components of the graph (Y, F). Since every H-path avoiding X contains an edge from F, it has at least two vertices in ∂C for some $C \in \mathcal{C}_F$, where ∂C denotes the set of vertices in C with a neighbour in $G - X - C$ (Fig. 3.4.1). The number of independent H-paths in G is therefore bounded above by

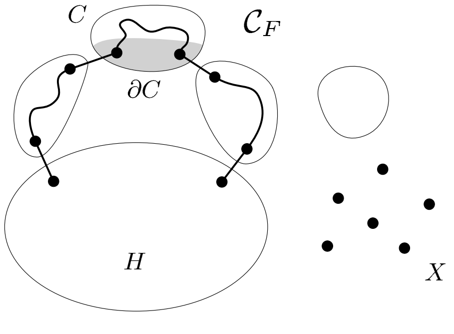

Fig. 3.4.1. An H-path in $G-X$

$M_G(H)$

$$M_G(H) := \min \Big(|X| + \sum_{C \in \mathcal{C}_F} \big\lfloor \tfrac{1}{2} |\partial C| \big\rfloor \Big),$$

X where the minimum is taken over all X and F as described above: $X \subseteq V(G-H)$ and $F \subseteq E(G-H-X)$ such that every H-path in G has a vertex or an edge in $X \cup F$.

3.5 Edge-disjoint spanning trees

The edge version of Menger's theorem tells us when a graph G contains k edge-disjoint paths between any two vertices. The actual routes of these paths within G may depend a lot on the choice of those two vertices: having found the paths for one pair of endvertices, we are not necessarily better placed to find them for another pair.

In a situation where quick access to a set of k edge-disjoint paths between any two vertices is desirable, it may be a good idea to ask for more than just k-edge-connectedness. For example, if G has *k edge-disjoint spanning trees*, there will be k canonical such paths between any two vertices, one in each tree.

When do such trees exist? If they do, the graph is clearly k-edge-connected. The converse is easily seen to be false; indeed, it is not even clear whether any edge-connectivity, however large, will imply the existence of k edge-disjoint spanning trees. Our first aim in this section will be to study conditions under which k edge-disjoint spanning trees exist.

As before, it is easy to write down some obvious necessary conditions for the existence of k edge-disjoint spanning trees. With respect to any partition of $V(G)$ into r sets, every spanning tree of G has at least $r-1$ *cross-edges*, edges whose ends lie in different partition sets (why?). Thus if G has k edge-disjoint spanning trees, it has at least $k\,(r-1)$ cross-edges.

cross-edges

Once more, this obvious necessary condition is also sufficient:

Theorem 3.5.1. (Tutte 1961; Nash-Williams 1961)
A multigraph contains k edge-disjoint spanning trees if and only if for every partition P of its vertex set it has at least $k\,(|P|-1)$ cross-edges.

Before we prove Theorem 3.5.1, let us note a surprising corollary: to ensure the existence of k edge-disjoint spanning trees, it suffices to raise the edge-connectivity to just $2k$:

[6.4.4]

Corollary 3.5.2. *Every $2k$-edge-connected multigraph G has k edge-disjoint spanning trees.*

Proof. Every set in a vertex partition of G is joined to other partition sets by at least $2k$ edges. Hence, for any partition into r sets, G has at least $\frac{1}{2}\sum_{i=1}^{r} 2k = kr$ cross-edges. The assertion thus follows from Theorem 3.5.1. □

$G=(V,E)$
$k,\mathcal{F}$

For the proof of Theorem 3.5.1, let a multigraph $G=(V,E)$ and $k \in \mathbb{N}$ be given. Let $\mathcal{F}$ be the set of all k-tuples $F=(F_1,\dots,F_k)$ of edge-disjoint spanning forests in G with the maximum total number of

$F := F^s$ in (1), we may think of e as a path of length 1 in $F_i' \cap C^0$. Successive applications of (1) to $F = F^s, \dots, F^0$ then give $xF_i^0 y \subseteq C^0$ as desired. □

(1.5.3) **Proof of Theorem 3.5.1.** We prove the backward implication by induction on $|G|$. For $|G| = 2$ the assertion holds. For the induction step, we now suppose that for every partition P of V there are at least $k\,(|P|-1)$ cross-edges, and construct k edge-disjoint spanning trees in G.

F^0 Pick a k-tuple $F^0 = (F_1^0, \dots, F_k^0) \in \mathcal{F}$. If every F_i^0 is a tree, we are done. If not, we have

$$\|F^0\| = \sum_{i=1}^{k} \|F_i^0\| < k\,(|G|-1)$$

by Corollary 1.5.3. On the other hand, we have $\|G\| \geqslant k\,(|G|-1)$ by assumption: consider the partition of V into single vertices. So there exists an edge $e^0 \in E \setminus E\,[\,F^0\,]$. By Lemma 3.5.3, there exists a set $U \subseteq V$ that is connected in every F_i^0 and contains the ends of e_0; in particular, $|U| \geqslant 2$. Since every partition of the contracted multigraph G/U induces a partition of G with the same cross-edges,[3] G/U has at least $k\,(|P|-1)$ cross-edges with respect to any partition P. By the induction hypothesis, therefore, G/U has k edge-disjoint spanning trees $T_1, \dots, T_k$. Replacing in each T_i the vertex v_U contracted from U by the spanning tree $F_i^0 \cap G\,[\,U\,]$ of $G\,[\,U\,]$, we obtain k edge-disjoint spanning trees in G. □

graph partition

Let us say that subgraphs $G_1, \dots, G_k$ of a graph G *partition* G if their edge sets form a partition of $E(G)$. Our spanning tree problem may then be recast as follows: into how *many connected* spanning subgraphs can we partition a given graph? The excuse for rephrasing our simple tree problem in this more complicated way is that it now has an obvious dual (cf. Theorem 1.5.1): into how *few acyclic* (spanning) subgraphs can we partition a given graph? Or for given k: which graphs can be partitioned into at most k forests?

An obvious necessary condition now is that every set $U \subseteq V(G)$ induces at most $k\,(|U|-1)$ edges, no more than $|U|-1$ for each forest. Once more, this condition turns out to be sufficient too. And surprisingly, this can be shown with the help of Lemma 3.5.3, which was designed for the proof of our theorem on edge-disjoint spanning trees:

Theorem 3.5.4. (Nash-Williams 1964)
A multigraph $G = (V, E)$ can be partitioned into at most k forests if and only if $\|G\,[\,U\,]\| \leqslant k\,(|U|-1)$ for every non-empty set $U \subseteq V$.

---

[3] see Chapter 1.10 on the contraction of multigraphs

$\varepsilon(G/U) \geqslant 2^{m-1}$; such a set U exists, because G itself has the form G/U with $|U| = 1$. Since G is connected, we have $N(U) \neq \emptyset$.

H

Let $H := G\,[\,N(U)\,]$. If H has a vertex v of degree $d_H(v) < 2^{m-1}$, we may add it to U and obtain a contradiction to the maximality of U: when we contract the edge vv_U in G/U, we lose one vertex and $d_H(v)+1 \leqslant 2^{m-1}$ edges, so ε will still be at least 2^{m-1}. Therefore $d(H) \geqslant \delta(H) \geqslant 2^{m-1}$. By the induction hypothesis, H contains a TY with $|Y| = r$ and $\|Y\| = m-1$. Let x, y be two branch vertices of this TY that are non-adjacent in Y. Since x and y lie in $N(U)$ and U is connected in G, G contains an x–y path whose inner vertices lie in U. Adding this path to the TY, we obtain the desired TX. □

How can Theorem 3.6.1 help with our aim to show that high connectivity will make a graph k-linked? Since high connectivity forces the average degree up (even the minimum degree), we may assume by the theorem that our graph contains a subdivision K of a large complete graph. Our plan now is to use Menger's theorem to link the given vertices s_i and t_i disjointly to branch vertices of K, and then to join up the correct pairs of those branch vertices inside K.

Theorem 3.6.2. (Jung 1970; Larman & Mani 1970)
There is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that every $f(k)$-connected graph is k-linked, for all $k \in \mathbb{N}$.

(3.3.1)
G
K
U

Proof. We prove the assertion for $f(k) = h(3k) + 2k$, where h is a function as in Theorem 3.6.1. Let G be an $f(k)$-connected graph. Then $d(G) \geqslant \delta(G) \geqslant \kappa(G) \geqslant h(3k)$; choose $K = TK^{3k} \subseteq G$ as in Theorem 3.6.1, and let U denote its set of branch vertices.

s_i, t_i
P_i, Q_i
$\mathcal{P}$

For the proof that G is k-linked, let distinct vertices $s_1, \dots, s_k$ and $t_1, \dots, t_k$ of G be given. By definition of $f(k)$, we have $\kappa(G) \geqslant 2k$. Hence by Menger's theorem (3.3.1), G contains disjoint paths $P_1, \dots, P_k$, $Q_1, \dots, Q_k$, such that each P_i starts in s_i, each Q_i starts in t_i, and all these paths end in U but have no inner vertices in U. Let the set $\mathcal{P}$ of these paths be chosen so that their total number of edges outside $E(K)$ is as small as possible.

u_i
L_i
v_i
M_i
w_i

Let $u_1, \dots, u_k$ be those k vertices in U that are not an end of a path in $\mathcal{P}$. For each $i = 1, \dots, k$, let L_i be the U-path in K (i.e., the subdivided edge of the K^{3k}) from u_i to the end of P_i in U, and let v_i be the first vertex of L_i on any path $P \in \mathcal{P}$. By definition of $\mathcal{P}$, P has no more edges outside $E(K)$ than $Pv_iL_iu_i$ does, so $v_iP = v_iL_i$ and hence $P = P_i$ (Fig. 3.6.1). Similarly, if M_i denotes the U-path in K from u_i to the end of Q_i in U, and w_i denotes the first vertex of M_i on any path in $\mathcal{P}$, then this path is Q_i. Then the paths $s_iP_iv_iL_iu_iM_iw_iQ_it_i$ are disjoint for different i and show that G is k-linked. □

5.$^-$ Prove the elementary properties of blocks mentioned at the beginning of Section 3.1.

6. Show that the block graph of any connected graph is a tree.

7. Show, without using Menger's theorem, that any two vertices of a 2-connected graph lie on a common cycle.

8. For edges $e, e' \in G$ write $e \sim e'$ if either $e = e'$ or e and e' lie on some common cycle in G. Show that $\sim$ is an equivalence relation on $E(G)$ whose equivalence classes are the edge sets of the non-trivial blocks of G.

9. Let G be a 2-connected graph but not a triangle, and let e be an edge of G. Show that either $G - e$ or G/e is again 2-connected.

10. Let G be a 3-connected graph, and let xy be an edge of G. Show that G/xy is 3-connected if and only if $G - \{x, y\}$ is 2-connected.

11. (i) Show that every cubic 3-edge-connected graph is 3-connected.

    (ii) Show that a graph is cubic and 3-connected if and only if it can be constructed from a K^4 by successive applications of the following operation: subdivide two edges by inserting a new vertex on each of them, and join the two new subdividing vertices by an edge.

12.$^-$ Show that Menger's theorem is equivalent to the following statement. For every graph G and vertex sets $A, B \subseteq V(G)$, there exist a set $\mathcal{P}$ of disjoint A–B paths in G and a set $X \subseteq V(G)$ separating A from B in G such that X has the form $X = \{x_P \mid P \in \mathcal{P}\}$ with $x_P \in P$ for all $P \in \mathcal{P}$.

13. Work out the details of the proof of Corollary 3.3.4 (ii).

14. Let $k \geqslant 2$. Show that every k-connected graph of order at least $2k$ contains a cycle of length at least $2k$.

15. Let $k \geqslant 2$. Show that in a k-connected graph any k vertices lie on a common cycle.

16. Derive the edge part of Corollary 3.4.2 from the vertex part.

    (Hint. Consider the H-paths in the graph obtained from the disjoint union of H and the line graph $L(G)$ by adding all the edges he such that h is a vertex of H and $e \in E(G) \setminus E(H)$ is an edge at h.)

17.$^-$ To the disjoint union of the graph $H = \overline{K^{2m+1}}$ with k copies of K^{2m+1} add edges joining H bijectively to each of the K^{2m+1}. Show that the resulting graph G contains at most $km = \frac{1}{2}\kappa_G(H)$ independent H-paths.

18. Find a bipartite graph G, with partition classes A and B say, such that for $H := G[A]$ there are at most $\frac{1}{2}\lambda_G(H)$ edge-disjoint H-paths in G.

(1969), their minimum degree is exactly k. The existence of a vertex of small degree can be particularly useful in induction proofs about k-connected graphs. Halin's theorem was the starting point for a series of more and more sophisticated studies of minimal k-connected graphs; see the books of Bollobás and Halin cited above, and in particular Mader's survey.

Our first proof of Menger's theorem is due to T. Böhme, F. Göring and J. Harant (manuscript 1999); the second to J.S. Pym, A proof of Menger's theorem, *Monatshefte Math.* **73** (1969), 81–88; the third to T. Grünwald (later Gallai), Ein neuer Beweis eines Mengerschen Satzes, *J. London Math. Soc.* **13** (1938), 188–192. The global version of Menger's theorem (Theorem 3.3.5) was first stated and proved by Whitney (1932).

Mader's Theorem 3.4.1 is taken from W. Mader, Über die Maximalzahl kreuzungsfreier H-Wege, *Arch. Math.* **31** (1978), 387–402. The theorem may be viewed as a common generalization of Menger's theorem and Tutte's 1-factor theorem (Exercise 19). Theorem 3.5.1 was proved independently by Nash-Williams and by Tutte; both papers are contained in *J. London Math. Soc.* **36** (1961). Theorem 3.5.4 is due to C.St.J.A. Nash-Williams, Decompositions of finite graphs into forests, *J. London Math. Soc.* **39** (1964), 12. Our proofs follow an account by Mader (personal communication). Both results can be elegantly expressed and proved in the setting of matroids; see § 18 in B. Bollobás, *Combinatorics*, Cambridge University Press 1986.

In Chapter 8.1 we shall prove that, in order to force a topological K^r minor in a graph G, we do not need an average degree of G as high as $h(r) = 2^{\binom{r}{2}}$ (as used in our proof of Theorem 3.6.1): the average degree required can be bounded above by a function quadratic in r (Theorem 8.1.1). The improvement of Theorem 3.6.2 mentioned in the text is due to B. Bollobás & A.G. Thomason, Highly linked graphs, *Combinatorica* **16** (1996), 313–320. N. Robertson & P.D. Seymour, Graph Minors XIII: The disjoint paths problem, *J. Combin. Theory B* **63** (1995), 65-110, showed that, for every fixed k, there is an $O(n^3)$ algorithm that decides whether a given graph of order n is k-linked. If k is taken as part of the input, the problem becomes NP-hard.

4.1 Topological prerequisites

In this section we briefly review some basic topological definitions and facts needed later. All these facts have (by now) easy and well-known proofs; see the notes for sources. Since those proofs contain no graph theory, we do not repeat them here: indeed our aim is to collect precisely those topological facts that we need but do *not* want to prove. Later, all proofs will follow strictly from the definitions and facts stated here (and be guided by but not rely on geometric intuition), so the material presented now will help to keep elementary topological arguments in those proofs to a minimum.

A *straight line segment* in the Euclidean plane is a subset of $\mathbb{R}^2$ that has the form $\{ p + \lambda(q-p) \mid 0 \leqslant \lambda \leqslant 1 \}$ for distinct points $p, q \in \mathbb{R}^2$.

polygon

A *polygon* is a subset of $\mathbb{R}^2$ which is the union of finitely many straight line segments and is homeomorphic to the unit circle. Here, as later, any subset of a topological space is assumed to carry the subspace topology. A *polygonal arc* is a subset of $\mathbb{R}^2$ which is the union of finitely many straight line segments and is homeomorphic to the closed unit interval $[0,1]$. The images of 0 and of 1 under such a homeomorphism are the *endpoints* of this polygonal arc, which *links* them and runs *between* them.

arc

Instead of 'polygonal arc' we shall simply say *arc*. If P is an arc between x and y, we denote the point set $P \setminus \{x, y\}$, the *interior* of P, by $\mathring{P}$.

region
separate
frontier

Let $O \subseteq \mathbb{R}^2$ be an open set. Being linked by an arc in O defines an equivalence relation on O. The corresponding equivalence classes are again open; they are the *regions* of O. A closed set $X \subseteq \mathbb{R}^2$ is said to *separate* O if $O \setminus X$ has more than one region. The *frontier* of a set $X \subseteq \mathbb{R}^2$ is the set Y of all points $y \in \mathbb{R}^2$ such that every neighbourhood of y meets both X and $\mathbb{R}^2 \setminus X$. Note that if X is open then its frontier lies in $\mathbb{R}^2 \setminus X$.

The frontier of a region O of $\mathbb{R}^2 \setminus X$, where X is a finite union of points and arcs, has two important properties. The first is accessibility: if $x \in X$ lies on the frontier of O, then x can be linked to some point in O by a straight line segment whose interior lies wholly inside O. As a consequence, any two points on the frontier of O can be linked by an arc whose interior lies in O (why?). The second notable property of the frontier of O is that it separates O from the rest of $\mathbb{R}^2$. Indeed, if $\varphi \colon [0,1] \to P \subseteq \mathbb{R}^2$ is continuous, with $\varphi(0) \in O$ and $\varphi(1) \notin O$, then P meets the frontier of O at least in the point $\varphi(y)$ for $y := \inf \{ x \mid \varphi(x) \notin O \}$, the *first point* of P in $\mathbb{R}^2 \setminus O$.

[4.2.1]
[4.2.4]
[4.2.5]
[4.2.10]
[4.3.1]
[4.5.1]
[4.6.1]
[5.1.2]

Theorem 4.1.1. (Jordan Curve Theorem for Polygons)
For every polygon $P \subseteq \mathbb{R}^2$, the set $\mathbb{R}^2 \setminus P$ has exactly two regions, of which exactly one is bounded. Each of the two regions has the entire polygon P as its frontier.

$\mathbb{R}^2 \setminus P$, let us call $C := \pi^{-1}(P)$ a *circle on* S^2, and the sets $\pi^{-1}(O)$ and $S^2 \setminus \pi^{-1}(P \cup O)$ the *regions* of C.

Our last tool is the theorem of Jordan and Schoenflies, again adapted slightly for our purposes:

[4.3.1] **Theorem 4.1.4.** *Let* $\varphi\colon C_1 \to C_2$ *be a homeomorphism between two circles on* S^2*, let* O_1 *be a region of* C_1*, and let* O_2 *be a region of* C_2*. Then* φ *can be extended to a homeomorphism* $C_1 \cup O_1 \to C_2 \cup O_2$.

4.2 Plane graphs

plane graph

A *plane graph* is a pair (V, E) of finite sets with the following properties (the elements of V are again called *vertices*, those of E *edges*):

(i) $V \subseteq \mathbb{R}^2$;
(ii) every edge is an arc between two vertices;
(iii) different edges have different sets of endpoints;
(iv) the interior of an edge contains no vertex and no point of any other edge.

A plane graph (V, E) defines a graph G on V in a natural way. As long as no confusion can arise, we shall use the name G of this abstract graph also for the plane graph (V, E), or for the point set $V \cup \bigcup E$; similar notational conventions will be used for abstract versus plane edges, for subgraphs, and so on.[1]

faces

F(G)

For every plane graph G, the set $\mathbb{R}^2 \setminus G$ is open; its regions are the *faces* of G. Since G is bounded—i.e., lies inside some sufficiently large disc D—exactly one of its faces is unbounded: the face that contains $\mathbb{R}^2 \setminus D$. This face is the *outer face* of G; the other faces are its *inner faces*. We denote the set of faces of G by $F(G)$. Note that if $H \subseteq G$ then every face of G is contained in a face of H.

In order to lay the foundations for the (easy but) rigorous introduction to plane graphs that this section aims to provide, let us descend once now into the realm of truly elementary topology of the plane, and prove what seems entirely obvious:[2] that the frontier of a face of a plane graph G is always a subgraph of G—not, say, half an edge. The following lemma states this formally, together with two similarly 'obvious' properties of plane graphs:

---

[1] However, we shall continue to use $\setminus$ for differences of point sets and $-$ for graph differences—which may help a little to keep the two apart.

[2] Note that even the best intuition can only ever be 'accurate', i.e., coincide with what the technical definitions imply, inasmuch as those definitions do indeed formalize what is intuitively intended. Given the complexity of definitions in elementary topology, this can hardly be taken for granted.

z $(D_0 \cup \ldots \cup D_n) \smallsetminus e$ to a point $z \in D_0 \smallsetminus e$ (Fig. 4.2.2); then y and z are equivalent in $\mathbb{R}^2 \smallsetminus G$. Hence, every point of $D_n \smallsetminus e$ lies in f_1 or in f_2, so x_1 cannot lie on the frontier of any other face of G. Since both half-discs of $D_0 \smallsetminus e$ can be linked to $D_n \smallsetminus e$ in this way (swap the roles of D_0 and D_n), we find that x_1 lies on the frontier of both f_1 and f_2. □

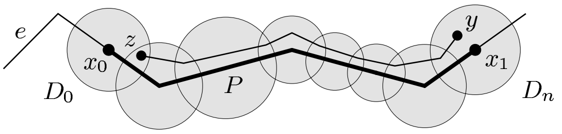

Fig. 4.2.2. An arc from y to D_0, close to P

Corollary 4.2.2. *The frontier of a face is always the point set of a subgraph.* □

boundary $G[f]$ The subgraph of G whose point set is the frontier of a face f is said to *bound* f and is called its *boundary*; we denote it by $G[f]$. A face is said to be *incident* with the vertices and edges of its boundary. Note that if $H \subseteq G$ then every face f of G is contained in a face f' of H. If $G[f] \subseteq H$ then $f = f'$ (why?); in particular, f is always also a face of its own boundary $G[f]$. These basic facts will be used frequently in the proofs to come.

[4.6.1] **Proposition 4.2.3.** *A plane forest has exactly one face.*

(4.1.3) *Proof.* Use induction on the number of edges and Lemma 4.1.3. □

With just one exception, different faces of a plane graph have different boundaries:

[4.3.1] **Lemma 4.2.4.** *If a plane graph has different faces with the same boundary, then the graph is a cycle.*

(4.1.1) *Proof.* Let G be a plane graph, and let $H \subseteq G$ be the boundary of distinct faces f_1, f_2 of G. Since f_1 and f_2 are also faces of H, Proposition 4.2.3 implies that H contains a cycle C. By Lemma 4.2.1 (ii), f_1 and f_2 are contained in different faces of C. Since f_1 and f_2 both have all of H as boundary, this implies that $H = C$: any further vertex or edge of H would lie in one of the faces of C and hence not on the boundary of the other. Thus, f_1 and f_2 are distinct faces of C. As C has only two faces, it follows that $f_1 \cup C \cup f_2 = \mathbb{R}^2$ and hence $G = C$. □

[4.3.1]
[4.4.3]
[4.5.1]
[4.5.2] **Proposition 4.2.5.** *In a 2-connected plane graph, every face is bounded by a cycle.*

Hence by Lemma 4.1.2 (ii), the plane edge v_2v_4 of $G[H]$ runs through f'_C rather than f_C (Fig. 4.2.3). Analogously, since $v_2, v_4 \in G[f]$, the edge v_1v_3 runs through f'_C. But the edges v_1v_3 and v_2v_4 are disjoint, so this contradicts Lemma 4.1.2 (ii). $\square$

The following classic result of Euler (1752)—here stated in its simplest form, for the plane—marks one of the common origins of graph theory and topology. The theorem relates the number of vertices, edges and faces in a plane graph: taken with the correct signs, these numbers always add up to 2. The general form of Euler's theorem asserts the same for graphs suitably embedded in other surfaces, too: the sum obtained is always a fixed number depending only on the surface, not on the graph, and this number differs for distinct (orientable closed) surfaces. Hence, any two such surfaces can be distinguished by a simple arithmetic invariant of the graphs embedded in them![3]

Let us then prove Euler's theorem in its simplest form:

Theorem 4.2.7. (Euler's Formula)
Let G be a connected plane graph with n vertices, m edges, and ℓ faces. Then

$$n - m + \ell = 2\,.$$

(1.5.1) (1.5.3) *Proof*. We fix n and apply induction on m. For $m \leqslant n-1$, G is a tree and $m = n-1$ (why?), so the assertion follows from Proposition 4.2.3.

e, G' Now let $m \geqslant n$. Then G has an edge e lying on a cycle; let $G' := G - e$. By Lemma 4.2.1 (ii), e lies on the boundary of exactly two faces f_1, f_2 of G; we put $f_{1,2} := f_1 \cup \mathring{e} \cup f_2$. We shall prove that
f_1, f_2
$f_{1,2}$

$$F(G) \setminus \{f_1, f_2\} = F(G') \setminus \{f_{1,2}\}\,, \qquad (*)$$

without assuming that $f_{1,2} \in F(G')$. However, since $\mathring{e}$ must lie in some face of G' and this will not be a face of G, by $(*)$ it can only be $f_{1,2}$. Thus again by $(*)$, G' has one face less than G. As G' also has one edge less than G, the assertion then follows from the induction hypothesis for G'.

For our proof of $(*)$ we first consider any $f \in F(G) \setminus \{f_1, f_2\}$. By Lemma 4.2.1 (i), we have $G[f] \subseteq G \setminus \mathring{e} = G'$. So f is also a face of G' (but obviously not equal to $f_{1,2}$) and hence lies in $F(G') \setminus \{f_{1,2}\}$.

f' Conversely, let a face $f' \neq f_{1,2}$ of G' be given. Since e lies on the boundary of both f_1 and f_2, we can link any two points of $f_{1,2}$ by an arc in $\mathbb{R}^2 \setminus G'$, so $f_{1,2}$ lies inside a face $f'_{1,2}$ of G'. Our assumption of $f' \neq f_{1,2}$ therefore implies $f' \not\subseteq f_{1,2}$ (as otherwise $f' \subseteq f_{1,2} \subseteq f'_{1,2}$
$f'_{1,2}$

---

[3] This fundamental connection between graphs and surfaces lies at the heart of the proof of the famous Robertson-Seymour *graph minor theorem*; see Chapter 12.5.

are linked by a C-path in G, because G is 3-connected. This path and e both run through the other face of C (not f) but do not intersect, a contradiction to Lemma 4.1.2 (ii).

It remains to show that C does not separate any two vertices $x, y \in G - C$. By Menger's theorem (3.3.5), x and y are linked in G by three independent paths. Clearly, f lies inside a face of their union, and by Lemma 4.1.2 (i) this face is bounded by only two of the paths. The third therefore avoids f and its boundary C. $\square$

4.3 Drawings

planar embedding

drawing

An embedding in the plane, or *planar embedding*, of an (abstract) graph G is an isomorphism between G and a plane graph $\tilde{G}$. The latter will be called a *drawing* of G. We shall not always distinguish notationally between the vertices and edges of G and of $\tilde{G}$.

In this section we investigate how two planar embeddings of a graph can differ. For this to make sense, we first have to agree when two embeddings are to be considered the same: for example, if we compose one embedding with a simple rotation of the plane, the resulting embedding will hardly count as a genuinely different way of drawing that graph.

G; V, E, F

G'; V', E', F'

σ

To prepare the ground, let us first consider three possible notions of equivalence for plane graphs (refining abstract isomorphism), and see how they are related. Let $G = (V, E)$ and $G' = (V', E')$ be two plane graphs, with face sets $F(G) =: F$ and $F(G') =: F'$. Assume that G and G' are isomorphic as abstract graphs, and let $\sigma: V \to V'$ be an isomorphism. Setting $xy \mapsto \sigma(x)\sigma(y)$, we may extend σ in a natural way to a bijection $V \cup E \to V' \cup E'$ which maps V to V' and E to E', and which preserves incidence (and non-incidence) between vertices and edges.

π

topological isomorphism

Our first notion of equivalence between plane graphs is perhaps the most natural one. Intuitively, we would like to call our isomorphism σ 'topological' if it is induced by a homeomorphism from the plane $\mathbb{R}^2$ to itself. To avoid having to grant the outer faces of G and G' a special status, however, we take a detour via the homeomorphism $\pi: S^2 \setminus \{ (0,0,1) \} \to \mathbb{R}^2$ chosen in Section 4.1: we call σ a *topological isomorphism* between the plane graphs G and G' if there exists a homeomorphism $\varphi: S^2 \to S^2$ such that $\psi := \pi \circ \varphi \circ \pi^{-1}$ induces σ on $V \cup E$. (More formally: we ask that ψ agree with σ on V, and that it map every plane edge $e \in G$ onto the plane edge $\sigma(e) \in G'$. Unless φ fixes the point $(0,0,1)$, the map ψ will be undefined at $\pi(\varphi^{-1}(0,0,1))$.)

It can be shown that, up to topological isomorphism, inner and outer faces are indeed no longer different: if we choose as φ a rotation of S^2 mapping the π^{-1}-image of a point of some inner face of G to the north pole $(0,0,1)$ of S^2, then ψ maps the rest of this face to the outer some face or not, and we require that σ map the subgraphs that do onto each other. At first glance, this third notion of equivalence may appear a little less natural than the previous two. However, it has the practical advantage of being formally weaker and hence easier to verify, and moreover, it will turn out to be equivalent to the other two notions in most cases.

As we have seen, every topological isomorphism between two plane graphs is also combinatorial, and every combinatorial isomorphism is also graph-theoretical. The following theorem shows that, for most graphs, the converse is true as well:

Theorem 4.3.1.

(i) *Every graph-theoretical isomorphism between two plane graphs is combinatorial. Its extension to a face bijection is unique if and only if the graph is not a cycle.*

(ii) *Every combinatorial isomorphism between two 2-connected plane graphs is topological.*

(4.1.1)
(4.1.4)
(4.2.4)
(4.2.5)

Proof. Let $G = (V, E)$ and $G' = (V', E')$ be two plane graphs, put $F(G) =: F$ and $F(G') =: F'$, and let $\sigma: V \cup E \to V' \cup E'$ be an isomorphism between the underlying abstract graphs.

(i) If G is a cycle, the assertion follows from the Jordan curve theorem. We now assume that G is not a cycle. Let $\mathcal{H}$ and $\mathcal{H}'$ be the sets of all face boundaries in G and G', respectively. If σ is a graph-theoretical isomorphism, then the map $H \mapsto \sigma(H)$ is a bijection between $\mathcal{H}$ and $\mathcal{H}'$. By Lemma 4.2.4, the map $f \mapsto G[f]$ is a bijection between F and $\mathcal{H}$, and likewise for F' and $\mathcal{H}'$. The composition of these three bijections is a bijection between F and F', which we choose as $\sigma: F \to F'$. By construction, this extension of σ to $V \cup E \cup F$ preserves incidences (and is unique with this property), so σ is indeed a combinatorial isomorphism.

σ (ii) Let us assume that G is 2-connected, and that σ is a combinatorial isomorphism. We have to construct a homeomorphism $\varphi: S^2 \to S^2$ which, for every vertex or plane edge $x \in G$, maps $\pi^{-1}(x)$ to $\pi^{-1}(\sigma(x))$.

$\tilde{\sigma}$ Since σ is a combinatorial isomorphism, $\tilde{\sigma}: \pi^{-1} \circ \sigma \circ \pi$ is an incidence preserving bijection from the vertices, edges and faces[4] of $\tilde{G} := \pi^{-1}(G)$

$\tilde{G}, \tilde{G}'$ to the vertices, edges and faces of $\tilde{G}' := \pi^{-1}(G')$.

We construct φ in three steps. Let us first define φ on the vertex set of $\tilde{G}$, setting $\varphi(x) := \tilde{\sigma}(x)$ for all $x \in V(\tilde{G})$. This is trivially a homeomorphism between $V(\tilde{G})$ and $V(\tilde{G}')$.

As the second step, we now extend φ to a homeomorphism between $\tilde{G}$ and $\tilde{G}'$ that induces $\tilde{\sigma}$ on $V(\tilde{G}) \cup E(\tilde{G})$. We may do this edge by

---

[4] By the ‘vertices, edges and faces’ of $\tilde{G}$ and $\tilde{G}'$ we mean the images under π^{-1} of the vertices, edges and faces of G and G' (plus $(0, 0, 1)$ in the case of the outer face). Their sets will be denoted by $V(\tilde{G})$, $E(\tilde{G})$, $F(\tilde{G})$ and $V(\tilde{G}')$, $E(\tilde{G}')$, $F(\tilde{G}')$, and incidence is defined as inherited from G and G'.

(4.2.10) *Proof.* Let G be a 3-connected graph with planar embeddings $\sigma_1 : G \to G_1$ and $\sigma_2 : G \to G_2$. By Theorem 4.3.1 it suffices to show that $\sigma_2 \circ \sigma_1^{-1}$ is a graph-theoretical isomorphism, i.e. that $\sigma_1(C)$ bounds a face of G_1 if and only if $\sigma_2(C)$ bounds a face of G_2, for every subgraph $C \subseteq G$. This follows at once from Proposition 4.2.10. □

4.4 Planar graphs: Kuratowski's theorem

planar A graph is called *planar* if it can be embedded in the plane: if it is isomorphic to a plane graph. A planar graph is *maximal*, or *maximally planar*, if it is planar but cannot be extended to a larger planar graph by adding an edge (but no vertex).

Drawings of maximal planar graphs are clearly maximally plane. The converse, however, is not obvious: when we start to draw a planar graph, could it happen that we get stuck half-way with a proper subgraph that is already maximally plane? Our first proposition says that this can never happen, that is, a plane graph is never maximally plane just because it is badly drawn:

Proposition 4.4.1.

(i) *Every maximal plane graph is maximally planar.*

(ii) *A planar graph with $n \geqslant 3$ vertices is maximally planar if and only if it has $3n-6$ edges.*

(4.2.6) (4.2.8) *Proof.* Apply Proposition 4.2.6 and Corollary 4.2.8. □

Which graphs are planar? As we saw in Corollary 4.2.9, no planar graph contains K^5 or $K_{3,3}$ as a topological minor. Our aim in this section is to prove the surprising converse, a classic theorem of Kuratowski: any graph without a topological K^5 or $K_{3,3}$ minor is planar.

Before we prove Kuratowski's theorem, let us note that it suffices to consider ordinary minors rather than topological ones:

Proposition 4.4.2. *A graph contains K^5 or $K_{3,3}$ as a minor if and only if it contains K^5 or $K_{3,3}$ as a topological minor.*

(1.7.2) *Proof.* By Proposition 1.7.2 it suffices to show that every graph G with a K^5 minor contains either K^5 as a topological minor or $K_{3,3}$ as a minor. So suppose that $G \succcurlyeq K^5$, and let $K \subseteq G$ be minimal such that $K = MK^5$. Then every branch set of K induces a tree in K, and between any two branch sets K has exactly one edge. If we take the tree induced by a branch set V_x and add to it the four edges joining it to other branch sets, we obtain another tree, T_x say. By the minimality

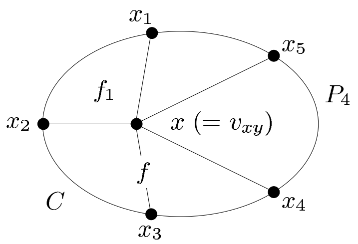

Fig. 4.4.2. $\tilde{G}'$ as a drawing of $G - y$: the vertex x is represented by the point v_{xy}

f_i denote the *other* face of C_i by f_i. Since f_i contains points of f (close to x) but no points of C, we have $f_i \subseteq f$. Moreover, the plane edges xx_j with $j \notin \{ i, i+1 \}$ meet C_i only in x and end outside f_i in $C \smallsetminus P_i$, so f_i meets none of those edges. Hence $f_i \subseteq \mathbb{R}^2 \smallsetminus \tilde{G}'$, that is, f_i is contained in a face of $\tilde{G}'$. (In fact, f_i is a face of $\tilde{G}'$, but we do not need this.)

In order to turn $\tilde{G}'$ into a drawing of G, let us try to find an i such that $Y \subseteq V(P_i)$; we may then embed y into f_i and link it up to its neighbours by arcs inside f_i. Suppose there is no such i: how then can the vertices of Y be distributed around C? If y had a neighbour in some $\mathring{P}_i$, it would also have one in $C - P_i$, so G would contain a $TK_{3,3}$ (with branch vertices x, y, x_i, x_{i+1} and those two neighbours of y). Hence $Y \subseteq X$. Now if $|Y| = |Y \cap X| \geqslant 3$, we have a TK^5 in G. So $|Y| \leqslant 2$; in fact, $|Y| = 2$, because $d(y) \geqslant \kappa(G) \geqslant 3$. Since the two vertices of Y lie on no common P_i, we can once more find a $TK_{3,3}$ in G, a contradiction. □

Compared with other proofs of Kuratowski's theorem, the above proof has the attractive feature that it can easily be adapted to produce a drawing in which every inner face is convex (exercise); in particular, every edge can be drawn straight. Note that 3-connectedness is essential here: a 2-connected planar graph need not have a drawing with all inner faces convex (example?), although it always has a straight-line drawing (Exercise 12).

It is not difficult, in principle, to reduce the general Kuratowski theorem to the 3-connected case by manipulating and combining partial drawings assumed to exist by induction. For example, if $\kappa(G) = 2$ and $G = G_1 \cup G_2$ with $V(G_1 \cap G_2) = \{ x, y \}$, and if G has no TK^5 or $TK_{3,3}$ subgraph, then neither $G_1 + xy$ nor $G_2 + xy$ has such a subgraph, and we may try to combine drawings of these graphs to one of $G + xy$. (If xy is already an edge of G, the same can be done with G_1 and G_2.) For $\kappa(G) \leqslant 1$, things become even simpler. However, the geometric operations involved require some cumbersome shifting and scaling, even if all the plane edges occurring are assumed to be straight.

(4.2.9) *Proof.* We apply induction on $|G|$. For $|G| = 4$, we have $G = K^4$ and the assertion holds. Now let $|G| > 4$, and let G be edge-maximal without a TK^5 or $TK_{3,3}$. Suppose $\kappa(G) \leqslant 2$, and choose G_1 and G_2 as in Lemma 4.4.4. For $\mathcal{X} := \{ K^5, K_{3,3} \}$, the lemma says that $G_1 \cap G_2$ is a K^2, with vertices x, y say. By Lemmas 4.4.4, 4.4.3 and the induction hypothesis, G_1 and G_2 are planar. For each $i = 1, 2$ separately, choose a drawing of G_i, a face f_i with the edge xy on its boundary, and a vertex $z_i \neq x, y$ on the boundary of f_i. Let K be a TK^5 or $TK_{3,3}$ in the abstract graph $G + z_1 z_2$ (Fig. 4.4.4).

G_1, G_2
x, y
f_i
z_i
K

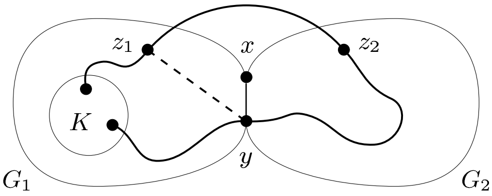

Fig. 4.4.4. A TK^5 or $TK_{3,3}$ in $G + z_1 z_2$

If all the branch vertices of K lie in the same G_i, then either $G_i + xz_i$ or $G_i + yz_i$ (or G_i itself, if z_i is already adjacent to x or y, respectively) contains a TK^5 or $TK_{3,3}$; this contradicts Corollary 4.2.9, since these graphs are planar by the choice of z_i. Since $G + z_1 z_2$ does not contain four independent paths between $(G_1 - G_2)$ and $(G_2 - G_1)$, these subgraphs cannot both contain a branch vertex of a TK^5, and cannot both contain two branch vertices of a $TK_{3,3}$. Hence K is a $TK_{3,3}$ with only one branch vertex v in, say, $G_2 - G_1$. But then also the graph $G_1 + v + \{ vx, vy, vz_1 \}$, which is planar by the choice of z_1, contains a $TK_{3,3}$. This contradicts Corollary 4.2.9. □

[4.5.1]
[12.4.3]

Theorem 4.4.6. (Kuratowski 1930; Wagner 1937)
The following assertions are equivalent for graphs G:

(i) *G is planar;*

(ii) *G contains neither K^5 nor $K_{3,3}$ as a minor;*

(iii) *G contains neither K^5 nor $K_{3,3}$ as a topological minor.*

(4.2.9) *Proof.* Combine Corollary 4.2.9 and Proposition 4.4.2 with Lemmas 4.4.3 and 4.4.5. □

Corollary 4.4.7. *Every maximal planar graph with at least four vertices is 3-connected.*

Proof. Apply Lemma 4.4.5 and Theorem 4.4.6. □

those of their subdivisions) do *not* have a simple basis: then G cannot contain a TK^5 or $TK_{3,3}$, and so is planar by Kuratowski's theorem.

Let us consider K^5 first. By Theorem 1.9.6, $\dim \mathcal{C}(K^5) = 6$; let $\mathcal{B} = \{ C_1, \dots, C_6 \}$ be a simple basis, and put $C_0 := C_1 + \ldots + C_6$. As $\mathcal{B}$ is linearly independent, none of the sets $C_0, \dots, C_6$ is empty, and so each of them contains at least three edges (cf. Proposition 1.9.2). The simplicity of $\mathcal{B}$ therefore implies

$$\begin{aligned} 18 = 6 \cdot 3 &\leqslant |C_1| + \ldots + |C_6| \\ &\leqslant 2\,\|K^5\| - |C_0| \\ &\leqslant 2 \cdot 10 - 3 = 17\,, \end{aligned}$$

a contradiction; for the middle inequality note that every edge in C_0 lies in just one of the sets $C_1, \dots, C_6$.

For $K_{3,3}$, Theorem 1.9.6 gives $\dim \mathcal{C}(K_{3,3}) = 4$; let $\mathcal{B} = \{ C_1, \dots, C_4 \}$ be a simple basis, and put $C_0 := C_1 + \ldots + C_4$. Since $K_{3,3}$ has girth 4, each C_i contains at least four edges, so

$$\begin{aligned} 16 = 4 \cdot 4 &\leqslant |C_1| + \ldots + |C_4| \\ &\leqslant 2\,\|K_{3,3}\| - |C_0| \\ &\leqslant 2 \cdot 9 - 4 = 14\,, \end{aligned}$$

a contradiction. □

It is one of the hidden beauties of planarity theory that two such abstract and seemingly unintuitive results about generating sets in cycle spaces as MacLane's theorem and Tutte's theorem 3.2.3 conspire to produce a very tangible planarity criterion for 3-connected graphs:

Theorem 4.5.2. (Tutte 1963)
A 3-connected graph is planar if and only if every edge lies on at most (equivalently: exactly) two non-separating induced cycles.

(3.2.3)
(4.2.1)
(4.2.5)
(4.2.10)

Proof. The forward implication follows from Propositions 4.2.10 and 4.2.1 (and Proposition 4.2.5 for the 'exactly two' version); the backward implication follows from Theorems 3.2.3 and 4.5.1. □

(ii) $|e^* \cap G| = |\mathring{e}^* \cap \mathring{e}| = |e \cap G^*| = 1$ for all $e \in E$;

(iii) $v \in f^*(v)$ for all $v \in V$.

The existence of such bijections implies that both G and G^* are connected (exercise). Conversely, every connected plane multigraph G has a plane dual G^*: if we pick from each face f of G a point $v^*(f)$ as a vertex for G^*, we can always link these vertices up by independent arcs as required by condition (ii), and there is always a bijection $V \to F^*$ satisfying (iii) (exercise).

If G_1^* and G_2^* are two plane duals of G, then clearly $G_1^* \simeq G_2^*$; in fact, one can show that the natural bijection $v_1^*(f) \mapsto v_2^*(f)$ is a topological isomorphism between G_1^* and G_2^*. In this sense, we may speak of *the* plane dual G^* of G.

Finally, G is in turn a plane dual of G^*. Indeed, this is witnessed by the inverse maps of the bijections from the definition of G^*: setting $v^*(f^*(v)) := v$ and $f^*(v^*(f)) := f$ for $f^*(v) \in F^*$ and $v^*(f) \in V^*$, we see that conditions (i) and (iii) for G^* transform into (iii) and (i) for G, while condition (ii) is symmetrical in G and G^*. Thus, the term 'dual' is also formally justified.

Plane duality is fascinating not least because it establishes a connection between two natural but very different kinds of edge sets in a multigraph, between cycles and cuts:

[6.5.2] **Proposition 4.6.1.** *For any connected plane multigraph G, an edge set $E \subseteq E(G)$ is the edge set of a cycle in G if and only if $E^* := \{ e^* \mid e \in E \}$ is a minimal cut in G^*.*

(4.1.1)
(4.2.3)

Proof. By conditions (i) and (ii) in the definition of G^*, two vertices $v^*(f_1)$ and $v^*(f_2)$ of G^* lie in the same component of $G^* - E^*$ if and only if f_1 and f_2 lie in the same region of $\mathbb{R}^2 \setminus \bigcup E$: every $v^*(f_1)$–$v^*(f_2)$ path in $G^* - E^*$ is an arc between f_1 and f_2 in $\mathbb{R}^2 \setminus \bigcup E$, and conversely every such arc P (with $P \cap V(G) = \emptyset$) defines a walk in $G^* - E^*$ between $v^*(f_1)$ and $v^*(f_2)$.

Now if $C \subseteq G$ is a cycle and $E = E(C)$ then, by the Jordan curve theorem and the above correspondence, $G^* - E^*$ has exactly two components, so E^* is a minimal cut in G^*.

Conversely, if $E \subseteq E(G)$ is such that E^* is a cut in G^*, then, by Proposition 4.2.3 and the above correspondence, E contains the edges of a cycle $C \subseteq G$. If E^* is minimal as a cut, then E cannot contain any further edges (by the implication shown before), so $E = E(C)$. □

Proposition 4.6.1 suggests the following generalization of plane duality to a notion of duality for abstract multigraphs. Let us call a multigraph G^* an *abstract dual* of a multigraph G if $E(G^*) = E(G)$ and the minimal cuts in G^* are precisely the edge sets of cycles in G. Note that any abstract dual of a multigraph is connected.

abstract dual

6. Let $G_1, G_2, \ldots$ be an infinite sequence of pairwise non-isomorphic graphs. Show that if $\limsup \varepsilon(G_i) > 3$ then the graphs G_i have unbounded genus—that is to say, there is no (closed) surface S in which all the G_i can be embedded.

   (Hint. You may use the fact that for every surface S there is a constant $\chi(S) \leqslant 2$ such that every graph embedded in S satisfies the generalized Euler formula of $n - m + \ell \geqslant \chi(S)$.)

7. Find a direct proof for planar graphs of Tutte's theorem on the cycle space of 3-connected graphs (Theorem 3.2.3).

8.$^-$ Show that the two plane graphs in Fig. 4.3.1 are not combinatorially (and hence not topologically) isomorphic.

9. Show that the two graphs in Fig. 4.3.2 are combinatorially but not topologically isomorphic.

10.$^-$ Show that our definition of equivalence for planar embeddings does indeed define an equivalence relation.

11. Find a 2-connected planar graph whose drawings are all topologically isomorphic but whose planar embeddings are not all equivalent.

12.$^+$ Show that every plane graph is combinatorially isomorphic to a plane graph whose edges are all straight.

   (Hint. Given a plane triangulation, construct inductively a graph-theoretically isomorphic plane graph whose edges are straight. Which additional property of the inner faces could help with the induction?)

Do not use Kuratowski's theorem in the following two exercises.

13. Show that any minor of a planar graph is planar. Deduce that a graph is planar if and only if it is the minor of a grid. (*Grids* are defined in Chapter 12.3.)

14. (i) Show that the planar graphs can in principle be characterized as in Kuratowski's theorem, i.e., that there exists a set $\mathcal{X}$ of graphs such that a graph G is planar if and only if G has no topological minor in $\mathcal{X}$.

   (ii) More generally, which graph properties can be characterized in this way?

15.$^-$ Does every planar graph have a drawing with all inner faces convex?

16. Modify the proof of Lemma 4.4.3 so that all inner faces become convex.

17. Does every minimal non-planar graph G (i.e., every non-planar graph G whose proper subgraphs are all planar) contain an edge e such that $G - e$ is maximally planar? Does the answer change if we define 'minimal' with respect to minors rather than subgraphs?

18. Show that adding a new edge to a maximal planar graph of order at least 6 always produces both a TK^5 and a $TK_{3,3}$ subgraph.

33. Show that a connected graph $G = (V, E)$ is planar if and only if there exists a connected multigraph $G' = (V', E)$ (i.e. with the same edge set) such that the following holds for every set $F \subseteq E$: the graph (V, F) is a tree if and only if $(V', E \setminus F)$ is a tree.

Notes

There is an excellent monograph on the embedding of graphs in surfaces, including the plane: B. Mohar & C. Thomassen, *Graphs on Surfaces*, Johns Hopkins University Press, to appear. Proofs of the results cited in Section 4.1, as well as all references for this chapter, can be found there. A good account of the Jordan curve theorem, both polygonal and general, is given also in J. Stillwell, *Classical topology and combinatorial group theory*, Springer 1980.

The short proof of Corollary 4.2.8 uses a trick that deserves special mention: the so-called *double counting* of pairs, illustrated in the text by a bipartite graph whose edges can be counted alternatively by summing its degrees on the left or on the right. Double counting is a technique widely used in combinatorics, and there will be more examples later in the book.

The material of Section 4.3 is not normally standard for an introductory graph theory course, and the rest of the chapter can be read independently of this section. However, the results of Section 4.3 are by no means unimportant. In a way, they have fallen victim to their own success: the shift from a topological to a combinatorial setting for planarity problems which they achieve has made the topological techniques developed there dispensable for most of planarity theory.

In its original version, Kuratowski's theorem was stated only for topological minors; the version for general minors was added by Wagner in 1937. Our proof of the 3-connected case (Lemma 4.4.3) can easily be strengthened to make all the inner faces convex (exercise); see C. Thomassen, Planarity and duality of finite and infinite graphs, *J. Combin. Theory B* **29** (1980), 244–271. The existence of such 'convex' drawings for 3-connected planar graphs follows already from the theorem of Steinitz (1922) that these graphs are precisely the 1-skeletons of 3-dimensional convex polyhedra. Compare also W.T. Tutte, How to draw a graph, *Proc. London Math. Soc.* **13** (1963), 743–767.

As one readily observes, adding an edge to a maximal planar graph (of order at least 6) produces not only a topological K^5 *or* $K_{3,3}$, but both. In Chapter 8.3 we shall see that, more generally, every graph with n vertices and more than $3n - 6$ edges contains a TK^5 and, with one easily described class of exceptions, also a $TK_{3,3}$. Seymour conjectures that every 5-connected non-planar graph contains a TK^5 (unpublished).

The simple cycle space basis constructed in the proof of MacLane's theorem, which consists of the inner face boundaries, is canonical in the following sense: for every simple basis $\mathcal{B}$ of the cycle space of a 2-connected planar graph there exists a drawing of that graph in which $\mathcal{B}$ is precisely the set of inner face boundaries. (This is proved in Mohar & Thomassen, who also mention some further planarity criteria.) Our proof of the backward direction of MacLane's theorem is based on Kuratowski's theorem. A more direct approach, in which

above, which leads to the problem of determining the maximum chromatic number of planar graphs. The committee scheduling problem, too, can be phrased as a vertex colouring problem—how?

edge colouring

chromatic index
$\chi'(G)$

An *edge colouring* of $G = (V, E)$ is a map $c\colon E \to S$ with $c(e) \neq c(f)$ for any adjacent edges e, f. The smallest integer k for which a k*-edge-colouring* exists, i.e. an edge colouring $c\colon E \to \{1, \dots, k\}$, is the *edge-chromatic number*, or *chromatic index*, of G; it is denoted by $\chi'(G)$. The third of our introductory questions can be modelled as an edge colouring problem in a bipartite multigraph (how?).

Clearly, every edge colouring of G is a vertex colouring of its line graph $L(G)$, and vice versa; in particular, $\chi'(G) = \chi(L(G))$. The problem of finding good edge colourings may thus be viewed as a restriction of the more general vertex colouring problem to this special class of graphs. As we shall see, this relationship between the two types of colouring problem is reflected by a marked difference in our knowledge about their solutions: while there are only very rough estimates for χ, its sister χ' always takes one of two values, either Δ or $\Delta + 1$.

5.1 Colouring maps and planar graphs

If any result in graph theory has a claim to be known to the world outside, it is the following *four colour theorem* (which implies that every map can be coloured with at most four colours):

Theorem 5.1.1. (Four Colour Theorem)
Every planar graph is 4-colourable.

Some remarks about the proof of the four colour theorem and its history can be found in the notes at the end of this chapter. Here, we prove the following weakening:

Proposition 5.1.2. (Five Colour Theorem)
Every planar graph is 5-colourable.

(4.1.1)
(4.2.8)

n, m

Proof. Let G be a plane graph with $n \geqslant 6$ vertices and m edges. We assume inductively that every plane graph with fewer than n vertices can be 5-coloured. By Corollary 4.2.8,

$$d(G) = 2m/n \leqslant 2\,(3n-6)/n < 6\,;$$

v
H
c

let $v \in G$ be a vertex of degree at most 5. By the induction hypothesis, the graph $H := G - v$ has a vertex colouring $c\colon V(H) \to \{1, \dots, 5\}$. If c uses at most 4 colours for the neighbours of v, we can extend it to a 5-colouring of G. Let us assume, therefore, that v has exactly 5 neighbours, and that these have distinct colours.

5.2 Colouring vertices

How do we determine the chromatic number of a given graph? How can we *find* a vertex-colouring with as few colours as possible? How does the chromatic number relate to other graph invariants, such as average degree, connectivity or girth?

Straight from the definition of the chromatic number we may derive the following upper bound:

Proposition 5.2.1. *Every graph G with m edges satisfies*

$$\chi(G) \leqslant \tfrac{1}{2} + \sqrt{2m + \tfrac{1}{4}}\,.$$

Proof. Let c be a vertex colouring of G with $k = \chi(G)$ colours. Then G has at least one edge between any two colour classes: if not, we could have used the same colour for both classes. Thus, $m \geqslant \frac{1}{2}k(k-1)$. Solving this inequality for k, we obtain the assertion claimed. □

greedy algorithm

One obvious way to colour a graph G with not too many colours is the following *greedy algorithm*: starting from a fixed vertex enumeration $v_1, \dots, v_n$ of G, we consider the vertices in turn and colour each v_i with the first available colour—e.g., with the smallest positive integer not already used to colour any neighbour of v_i among $v_1, \dots, v_{i-1}$. In this way, we never use more than $\Delta(G) + 1$ colours, even for unfavourable choices of the enumeration $v_1, \dots, v_n$. If G is complete or an odd cycle, then this is even best possible.

In general, though, this upper bound of $\Delta + 1$ is rather generous, even for greedy colourings. Indeed, when we come to colour the vertex v_i in the above algorithm, we only need a supply of $d_{G[\,v_1,\dots,v_i\,]}(v_i) + 1$ rather than $d_G(v_i) + 1$ colours to proceed; recall that, at this stage, the algorithm ignores any neighbours v_j of v_i with $j > i$. Hence in most graphs, there will be scope for an improvement of the $\Delta + 1$ bound by choosing a particularly suitable vertex ordering to start with: one that picks vertices of large degree early (when most neighbours are ignored) and vertices of small degree last. Locally, the number $d_{G[\,v_1,\dots,v_i\,]}(v_i) + 1$ of colours required will be smallest if v_i has minimum degree in $G\,[\,v_1, \dots, v_i\,]$. But this is easily achieved: we just choose v_n first, with $d(v_n) = \delta(G)$, then choose as v_{n-1} a vertex of minimum degree in $G - v_n$, and so on.

colouring number $\mathrm{col}(G)$

The least number k such that G has a vertex enumeration in which each vertex is preceded by fewer than k of its neighbours is called the *colouring number* $\mathrm{col}(G)$ of G. The enumeration we just discussed shows that $\mathrm{col}(G) \leqslant \max_{H \subseteq G} \delta(H) + 1$. But for $H \subseteq G$ clearly also $\mathrm{col}(G) \geqslant \mathrm{col}(H)$ and $\mathrm{col}(H) \geqslant \delta(H) + 1$, since the 'back-degree' of the last vertex in any enumeration of H is just its ordinary degree in H, which is at least $\delta(H)$. So we have proved the following:

contrary to (1). Hence the neighbour of v_i on P is its only neighbour in $C_{i,j}$, and similarly for v_j. Thus if $C_{i,j} \neq P$, then P has an inner vertex with three identically coloured neighbours in H; let u be the first such vertex on P (Fig. 5.2.1). Since at most $\Delta - 2$ colours are used on the neighbours of u, we may recolour u. But this makes $P\mathring{u}$ into a component of $H_{i,j}$, contradicting (2).

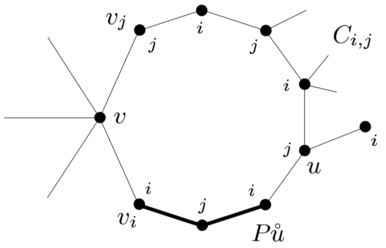

Fig. 5.2.1. The proof of (3) in Brooks's theorem

For distinct i, j, k, the paths $C_{i,j}$ and $C_{i,k}$ meet only in v_i. (4)

For if $v_i \neq u \in C_{i,j} \cap C_{i,k}$, then u has two neighbours coloured j and two coloured k, so we may recolour u. In the new colouring, v_i and v_j lie in different components of $H_{i,j}$, contrary to (2).

$v_1, \dots, v_\Delta$
c
u
c'

The proof of the theorem now follows easily. If the neighbours of v are pairwise adjacent, then each has Δ neighbours in $N(v) \cup \{ v \}$ already, so $G = G[N(v) \cup \{ v \}] = K^{\Delta+1}$. As G is complete, there is nothing to show. We may thus assume that $v_1 v_2 \notin G$, where $v_1, \dots, v_\Delta$ derive their names from some fixed Δ-colouring c of H. Let $u \neq v_2$ be the neighbour of v_1 on the path $C_{1,2}$; then $c(u) = 2$. Interchanging the colours 1 and 3 in $C_{1,3}$, we obtain a new colouring c' of H; let v'_i, $H'_{i,j}$, $C'_{i,j}$ etc. be defined with respect to c' in the obvious way. As a neighbour of $v_1 = v'_3$, our vertex u now lies in $C'_{2,3}$, since $c'(u) = c(u) = 2$. By (4) for c, however, the path $\mathring{v}_1 C_{1,2}$ retained its original colouring, so $u \in \mathring{v}_1 C_{1,2} \subseteq C'_{1,2}$. Hence $u \in C'_{2,3} \cap C'_{1,2}$, contradicting (4) for c'. □

As we have seen, a graph G of large chromatic number must have large maximum degree: at least $\chi(G) - 1$. What else can we say about the structure of graphs with large chromatic number?

One obvious possible cause for $\chi(G) \geqslant k$ is the presence of a K^k subgraph. This is a local property of G, compatible with arbitrary values of global invariants such as ε and κ. Hence, the assumption of $\chi(G) \geqslant k$ does not tell us anything about those invariants for G itself. It does, however, imply the existence of a subgraph where those invariants are large: by Corollary 5.2.3, G has a subgraph H with $\delta(H) \geqslant k - 1$, and hence by Theorem 1.4.2 a subgraph H' with $\kappa(H') \geqslant \lfloor \frac{1}{4}(k-1) \rfloor$.

(iii) If G_1, G_2 are k-constructible and there are vertices x, y_1, y_2 such that $G_1 \cap G_2 = \{x\}$, $xy_1 \in E(G_1)$ and $xy_2 \in E(G_2)$, then also $(G_1 \cup G_2) - xy_1 - xy_2 + y_1y_2$ is k-constructible (Fig. 5.2.2).

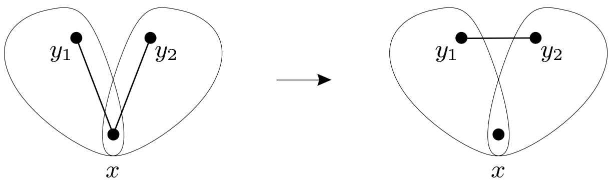

Fig. 5.2.2. The *Hajós construction* (iii)

One easily checks inductively that all k-constructible graphs—and hence their supergraphs—are at least k-chromatic. Indeed, if $(G + xy)/xy$ as in (ii) has a colouring with fewer than k colours, then this defines such a colouring also for G, a contradiction. Similarly, in any colouring of the graph constructed in (iii), the vertices y_1 and y_2 do not both have the same colour as x, so this colouring induces a colouring of either G_1 or G_2 and hence uses at least k colours.

It is remarkable, though, that the converse holds too:

Theorem 5.2.5. (Hajós 1961)
Let G be a graph and $k \in \mathbb{N}$. Then $\chi(G) \geqslant k$ if and only if G has a k-constructible subgraph.

Proof. Let G be a graph with $\chi(G) \geqslant k$; we show that G has a k-constructible subgraph. Suppose not; then $k \geqslant 3$. Adding some edges if necessary, let us make G edge-maximal with the property that none of its subgraphs is k-constructible. Now G is not a complete r-partite graph for any r: for then $\chi(G) \geqslant k$ would imply $r \geqslant k$, and G would contain the k-constructible graph K^k.

Since G is not a complete multipartite graph, non-adjacency is not an equivalence relation on $V(G)$. So there are vertices y_1, x, y_2 such that $y_1x, xy_2 \notin E(G)$ but $y_1y_2 \in E(G)$. Since G is edge-maximal without a k-constructible subgraph, each edge xy_i lies in some k-constructible subgraph H_i of $G + xy_i$ $(i = 1, 2)$.

y_1x, xy_2

H_1, H_2

Let H_2' be an isomorphic copy of H_2 that contains x and $H_2 - H_1$ but is otherwise disjoint from G, together with an isomorphism $v \mapsto v'$ from H_2 to H_2' that fixes $H_2 \cap H_2'$ pointwise. Then $H_1 \cap H_2' = \{x\}$, so

H_2'

v' *etc.*

$$H := (H_1 \cup H_2') - xy_1 - xy_2' + y_1y_2'$$

is k-constructible by (iii). One vertex at a time, let us identify in H each vertex $v' \in H_2' - G$ with its partner v; since vv' is never an edge of H,

colouring
α-edge

given, and assume that the assertion holds for graphs with fewer edges. Instead of '$(\Delta+1)$-edge-colouring' let us just say 'colouring'. An edge coloured α will again be called an *α-edge*.

missing
α/β-path

For every edge $e \in G$ there exists a colouring of $G - e$, by the induction hypothesis. In such a colouring, the edges at a given vertex v use at most $d(v) \leqslant \Delta$ colours, so some colour $\beta \in \{1, \dots, \Delta+1\}$ is *missing at v*. For any other colour α, there is a unique maximal walk (possibly trivial) starting at v, whose edges are coloured alternately α and β. This walk is a path; we call it the *α/β-path from v*.

Suppose that G has no colouring. Then the following holds:

Given $xy \in E$, and any colouring of $G - xy$ in which the colour α is missing at x and the colour β is missing at y, the α/β-path from y ends in x. (1)

Otherwise we could interchange the colours α and β along this path and colour xy with α, obtaining a colouring of G (contradiction).

xy_0
G_0, c_0, α
$y_1, \dots, y_k$
G_i
c_i

Let $xy_0 \in G$ be an edge. By induction, $G_0 := G - xy_0$ has a colouring c_0. Let α be a colour missing at x in this colouring. Further, let $y_0, y_1, \dots, y_k$ be a maximal sequence of distinct neighbours of x in G, such that $c_0(xy_i)$ is missing in c_0 at y_{i-1} for each $i = 1, \dots, k$. For each of the graphs $G_i := G - xy_i$ we define a colouring c_i, setting

$$c_i(e) := \begin{cases} c_0(xy_{j+1}) & \text{for } e = xy_j \text{ with } j \in \{0, \dots, i-1\} \\ c_0(e) & \text{otherwise;} \end{cases}$$

note that in each of these colourings the same colours are missing at x as in c_0.

β
i

Now let β be a colour missing at y_k in c_0. Clearly, β is still missing at y_k in c_k. If β were also missing at x, we could colour xy_k with β and thus extend c_k to a colouring of G. Hence, x is incident with a β-edge (in every colouring). By the maximality of k, therefore, there is an $i \in \{1, \dots, k-1\}$ such that

$$c_0(xy_i) = \beta\,.$$

P
P'

Let P be the α/β-path from y_k in G_k (with respect to c_k; Fig. 5.3.1). By (1), P ends in x, and it does so on a β-edge, since α is missing at x. As $\beta = c_0(xy_i) = c_k(xy_{i-1})$, this is the edge xy_{i-1}. In c_0, however, and hence also in c_{i-1}, β is missing at y_{i-1} (by (2) and the choice of y_i); let P' be the α/β-path from y_{i-1} in G_{i-1} (with respect to c_{i-1}). Since P' is uniquely determined, it starts with $y_{i-1}Py_k$; note that the edges of $P\mathring{x}$ are coloured the same in c_{i-1} as in c_k. But in c_0, and hence in c_{i-1}, there is no β-edge at y_k (by the choice of β). Therefore P' ends in y_k, contradicting (1). □

In spite of these inequalities, many of the known upper bounds for the chromatic number have turned out to be valid for the choice number, too. Examples for this phenomenon include Brooks's theorem and Proposition 5.2.2; in particular, graphs of large choice number still have subgraphs of large minimum degree. On the other hand, it is easy to construct graphs for which the two invariants are wide apart (Exercise 24). Taken together, these two facts indicate a little how far those general upper bounds on the chromatic number may be from the truth.

The following theorem shows that, in terms of its relationship to other graph invariants, the choice number differs fundamentally from the chromatic number. As mentioned before, there are 2-chromatic graphs of arbitrarily large minimum degree, e.g. the graphs $K_{n,n}$. The choice number, however, will be forced up by large values of invariants like δ, ε or κ:

Theorem 5.4.1. (Alon 1993)
There exists a function $f\colon \mathbb{N}\to\mathbb{N}$ *such that, given any integer* k*, all graphs* G *with average degree* $d(G) \geqslant f(k)$ *satisfy* $\mathrm{ch}(G) \geqslant k$*.*

The proof of Theorem 5.4.1 uses probabilistic methods as introduced in Chapter 11.

Empirically, the choice number's different character is highlighted by another phenomenon: even in cases where known bounds for the chromatic number could be transferred to the choice number, their proofs have tended to be rather different.

One of the simplest and most impressive examples for this is the list version of the five colour theorem: every planar graph is 5-choosable. This had been conjectured for almost 20 years, before Thomassen found a very simple induction proof. This proof does not use the five colour theorem—which thus gets reproved in a very different way.

Theorem 5.4.2. (Thomassen 1994)
Every planar graph is 5-choosable.

(4.2.6) *Proof.* We shall prove the following assertion for all plane graphs G with at least 3 vertices:

> *Suppose that every inner face of* G *is bounded by a triangle and its outer face by a cycle* $C = v_1 \dots v_k v_1$*. Suppose further that* v_1 *has already been coloured with the colour 1, and* v_2 *has been coloured 2. Suppose finally that with every other vertex of* C *a list of at least* 3 *colours is associated, and with every vertex of* $G - C$ *a list of at least* 5 *colours. Then the colouring of* v_1 *and* v_2 *can be extended to a colouring of* G *from the given lists.* $(*)$

of C are bounded by triangles, $P := v_1 u_1 \dots u_m v_{k-1}$ is a path in G, and $C' := P \cup (C - v_k)$ a cycle.

C'

We now choose two different colours $j, \ell \neq 1$ from the list of v_k and delete these colours from the lists of all the vertices u_i. Then every list of a vertex on C' still has at least 3 colours, so by induction we may colour C' and its interior, i.e. the graph $G - v_k$. At least one of the two colours j, ℓ is not used for v_{k-1}, and we may assign that colour to v_k. □

As is often the case with induction proofs, the trick of the proof above lies in the delicately balanced strengthening of the assertion proved. Note that the proof uses neither traditional colouring arguments (such as swapping colours along a path) nor the Euler formula implicit in the standard proof of the five colour theorem. This suggests that maybe in other unsolved colouring problems too it might be of advantage to aim straight for their list version, i.e. to prove an assertion of the form $\mathrm{ch}(G) \leqslant k$ instead of the formally weaker $\chi(G) \leqslant k$. Unfortunately, this approach fails for the four colour theorem: planar graphs are *not* in general 4-choosable.

As mentioned before, the chromatic number of a graph and its choice number may differ a lot. Surprisingly, however, no such examples are known for edge colourings. Indeed it has been conjectured that none exist:

List colouring conjecture. *Every graph G satisfies* $\mathrm{ch}'(G) = \chi'(G)$.

We shall prove the list colouring conjecture for bipartite graphs. As a tool we shall use orientations of graphs, defined in Chapter 1.10. If D is a directed graph and $v \in V(D)$, we denote by $N^+(v)$ the set, and by $d^+(v)$ the number, of vertices w such that D contains an edge directed from v to w.

$N^+(v)$
$d^+(v)$

To see how orientations come into play in the context of colouring, let us recall the greedy algorithm from Section 5.2. In order to apply the algorithm to a graph G, we first have to choose a vertex enumeration $v_1, \dots, v_n$ of G. The enumeration chosen defines an orientation of G: just orient every edge $v_i v_j$ 'backwards', from v_i to v_j if $i > j$. Then, for each vertex v_i to be coloured, the algorithm considers only those edges at v_i that are directed away from v_i: if $d^+(v) < k$ for all vertices v, it will use at most k colours. Moreover, the first colour class U found by the algorithm has the following property: it is an independent set of vertices to which every other vertex sends an edge. The second colour class has the same property in $G - U$, and so on.

The following lemma generalizes this to orientations D of G that do not necessarily come from a vertex enumeration, but may contain some directed cycles. Let us call an independent set $U \subseteq V(D)$ a *kernel* of D

kernel

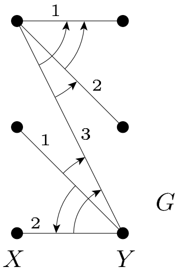

Fig. 5.4.3. Orienting the line graph of G

U c-value, and let U denote the set of all those edges e_x. Then every edge $e' \in E' \setminus U$ meets some $e \in U$ in X, and the edge $ee' \in D'$ is directed from e' to e. If U is independent, it is thus a kernel of D' and we are home; let us assume, therefore, that U is not independent.

e, e' Let $e, e' \in U$ be adjacent, and assume that $c(e) < c(e')$. By definition of U, e and e' meet in Y, so the edge $ee' \in D'$ is directed from e to e'. U' By the induction hypothesis, $D' - e$ has a kernel U'. If $e' \in U'$, then U' is also a kernel of D', and we are done. If not, there exists an $e'' \in U'$ such that D' has an edge directed from e' to e''. If e' and e'' met in X, then $c(e'') < c(e')$ by definition of D, contradicting $e' \in U$. Hence e' and e'' meet in Y, and $c(e') < c(e'')$. Since e and e' meet in Y, too, also e and e'' meet in Y, and $c(e) < c(e') < c(e'')$. So the edge ee'' is directed from e towards e'', so again U' is also a kernel of D'. □

By Proposition 5.3.1, we now know the exact list-chromatic index of bipartite graphs:

Corollary 5.4.5. *Every bipartite graph G satisfies* $\mathrm{ch}'(G) = \Delta(G)$. □

5.5 Perfect graphs

As discussed in Section 5.2, a high chromatic number may occur as a purely global phenomenon: even when a graph has large girth, and thus locally looks like a tree, its chromatic number may be arbitrarily high. Since such 'global dependence' is obviously difficult to deal with, one may become interested in graphs where this phenomenon does not occur, i.e. whose chromatic number is high only when there is a local reason for it.

Before we make this precise, let us note two definitions for a graph G. $\omega(G)$ The greatest integer r such that $K^r \subseteq G$ is the *clique number* $\omega(G)$ of G, and the greatest integer r such that $\overline{K^r} \subseteq G$ (induced) is the *independence number* $\alpha(G)$ of G. $\alpha(G)$ Clearly, $\alpha(G) = \omega(\overline{G})$ and $\omega(G) = \alpha(\overline{G})$.

X C G_1, G_2 S

be two non-adjacent vertices, and let $X \subseteq V(G) \setminus \{a, b\}$ a minimal set of vertices separating a from b. Let C denote the component of $G - X$ containing a, and put $G_1 := G[V(C) \cup X]$ and $G_2 := G - C$. Then G arises from G_1 and G_2 by pasting these graphs together along $S := G[X]$.

s, t

Since G_1 and G_2 are both chordal (being induced subgraphs of G) and hence constructible by induction, it suffices to show that S is complete. Suppose, then, that $s, t \in S$ are non-adjacent. By the minimality of $X = V(S)$ as an a–b separator, both s and t have a neighbour in C. Hence, there is an X-path from s to t in G_1; we let P_1 be a shortest such path. Analogously, G_2 contains a shortest X-path P_2 from s to t. But then $P_1 \cup P_2$ is a chordless cycle of length $\geqslant 4$ (Fig. 5.5.1), contradicting our assumption that G is chordal. □

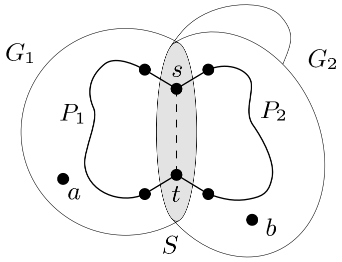

Fig. 5.5.1. If G_1 and G_2 are chordal, then so is G

Proposition 5.5.2. *Every chordal graph is perfect.*

Proof. Since complete graphs are perfect, it suffices by Proposition 5.5.1 to show that any graph G obtained from perfect graphs G_1, G_2 by pasting them together along a complete subgraph S is again perfect. So let $H \subseteq G$ be an induced subgraph; we show that $\chi(H) \leqslant \omega(H)$.

Let $H_i := H \cap G_i$ for $i = 1, 2$, and let $T := H \cap S$. Then T is again complete, and H arises from H_1 and H_2 by pasting along T. As an induced subgraph of G_i, each H_i can be coloured with $\omega(H_i)$ colours. Since T is complete and hence coloured injectively, two such colourings, one of H_1 and one of H_2, may be combined into a colouring of H with $\max\{\omega(H_1), \omega(H_2)\} \leqslant \omega(H)$ colours—if necessary by permuting the colours in one of the H_i. □

We now come to the main result in the theory of perfect graphs, the *perfect graph theorem*:

perfect graph theorem

Theorem 5.5.3. (Lovász 1972)
A graph is perfect if and only if its complement is perfect.

Proof of Theorem 5.5.3. Applying induction on $|G|$, we show that the complement $\overline{G}$ of any perfect graph $G = (V,E)$ is again perfect. For $|G| = 1$ this is trivial, so let $|G| \geqslant 2$ for the induction step. Let $\mathcal{K}$ denote the set of all vertex sets of complete subgraphs of G. Put $\alpha(G) =: \alpha$, and let $\mathcal{A}$ be the set of all independent vertex sets A in G with $|A| = \alpha$.

$G = (V,E)$
$\mathcal{K}$
α
$\mathcal{A}$

Every proper induced subgraph of $\overline{G}$ is the complement of a proper induced subgraph of G, and is hence perfect by induction. For the perfection of $\overline{G}$ it thus suffices to prove $\chi(\overline{G}) \leqslant \omega(\overline{G})$ $(= \alpha)$. To this end, we shall find a set $K \in \mathcal{K}$ such that $K \cap A \neq \emptyset$ for all $A \in \mathcal{A}$; then

$$\omega(\overline{G} - K) = \alpha(G - K) < \alpha = \omega(\overline{G}),$$

so by the induction hypothesis

$$\chi(\overline{G}) \leqslant \chi(\overline{G} - K) + 1 = \omega(\overline{G} - K) + 1 \leqslant \omega(\overline{G})$$

as desired.

Suppose there is no such K; thus, for every $K \in \mathcal{K}$ there exists a set $A_K \in \mathcal{A}$ with $K \cap A_K = \emptyset$. Let us replace in G every vertex x by a complete graph G_x of order

A_K
G_x

$k(x)$

$$k(x) := \bigl|\{\, K \in \mathcal{K} \mid x \in A_K \,\}\bigr|,$$

joining all the vertices of G_x to all the vertices of G_y whenever x and y are adjacent in G. The graph G' thus obtained has vertex set $\bigcup_{x \in V} V(G_x)$, and two vertices $v \in G_x$ and $w \in G_y$ are adjacent in G' if and only if $x = y$ or $xy \in E$. Moreover, G' can be obtained by repeated vertex expansion from the graph $G[\{\, x \in V \mid k(x) > 0 \,\}]$. Being an induced subgraph of G, this latter graph is perfect by assumption, so G' is perfect by Lemma 5.5.4. In particular,

G'

$$\chi(G') \leqslant \omega(G'). \tag{1}$$

In order to obtain a contradiction to (1), we now compute in turn the actual values of $\omega(G')$ and $\chi(G')$. By construction of G', every maximal complete subgraph of G' has the form $G'[\bigcup_{x \in X} G_x]$ for some $X \in \mathcal{K}$. So there exists a set $X \in \mathcal{K}$ such that

X

$$\begin{aligned}
\omega(G') &= \sum_{x \in X} k(x) \\
&= \bigl|\{\, (x,K) : x \in X,\ K \in \mathcal{K},\ x \in A_K \,\}\bigr| \\
&= \sum_{K \in \mathcal{K}} |X \cap A_K| \\
&\leqslant |\mathcal{K}| - 1;
\end{aligned} \tag{2}$$

Indeed, while the first equality is immediate from the perfection of $G-U$, the second is easy: '$\leqslant$' is obvious, while $\chi(G-U) < \omega$ would imply $\chi(G) \leqslant \omega$, so G would be perfect contrary to our assumption.

Let us apply (1) to a singleton $U = \{ u \}$ and consider an ω-colouring of $G-u$. Let K be the vertex set of any K^ω in G. Clearly,

$$\textit{if } u \notin K \textit{ then } K \textit{ meets every colour class of } G-u; \tag{2}$$

$$\textit{if } u \in K \textit{ then } K \textit{ meets all but exactly one colour class of } G-u. \tag{3}$$

A_0
A_i
K_i

Let $A_0 = \{ u_1, \dots, u_\alpha \}$ be an independent set in G of size α. Let $A_1, \dots, A_\omega$ be the colour classes of an ω-colouring of $G-u_1$, let $A_{\omega+1}, \dots, A_{2\omega}$ be the colour classes of an ω-colouring of $G-u_2$, and so on; altogether, this gives us $\alpha\omega+1$ independent sets $A_0, A_1, \dots, A_{\alpha\omega}$ in G. For each $i = 0, \dots, \alpha\omega$, there exists by (1) a $K^\omega \subseteq G - A_i$; we denote its vertex set by K_i.

Note that if K is the vertex set of any K^ω in G, then

$$K \cap A_i = \emptyset \textit{ for exactly one } i \in \{ 0, \dots, \alpha\omega+1 \}. \tag{4}$$

Indeed, if $K \cap A_0 = \emptyset$ then $K \cap A_i \neq \emptyset$ for all $i \neq 0$, by definition of A_i and (2). Similarly if $K \cap A_0 \neq \emptyset$, then $|K \cap A_0| = 1$, so $K \cap A_i = \emptyset$ for exactly one $i \neq 0$: apply (3) to the unique vertex $u \in K \cap A_0$, and (2) to all the other vertices $u \in A_0$.

J
A
B

Let J be the real $(\alpha\omega+1) \times (\alpha\omega+1)$ matrix with zero entries in the main diagonal and all other entries 1. Let A be the real $(\alpha\omega+1) \times n$ matrix whose rows are the incidence vectors of the subsets $A_i \subseteq V$: if $a_{i1}, \dots, a_{in}$ denote the entries of the ith row of A, then $a_{ij} = 1$ if $v_j \in A_i$, and $a_{ij} = 0$ otherwise. Similarly, let B denote the real $n \times (\alpha\omega+1)$ matrix whose columns are the incidence vectors of the subsets $K_i \subseteq V$. Now while $|K_i \cap A_i| = 0$ for all i by the choice of K_i, we have $K_i \cap A_j \neq \emptyset$ and hence $|K_i \cap A_j| = 1$ whenever $i \neq j$, by (4). Thus,

$$AB = J.$$

Since J is non-singular, this implies that A has rank $\alpha\omega+1$. In particular, $n \geqslant \alpha\omega+1$, which contradicts $(*)$ for $H := G$. $\square$

By definition, every induced subgraph of a perfect graph is again perfect. The property of perfection can therefore be characterized by forbidden induced subgraphs: there exists a set $\mathcal{H}$ of imperfect graphs such that any graph is perfect if and only if it has no induced subgraph isomorphic to an element of $\mathcal{H}$. (For example, we may choose as $\mathcal{H}$ the set of all imperfect graphs with vertices in $\mathbb{N}$.)

Naturally, it would be desirable to keep $\mathcal{H}$ as small as possible. In fact, one of the best known conjectures in graph theory says that $\mathcal{H}$

8. Show that the bound of Proposition 5.2.2 is always at least as sharp as that of Proposition 5.2.1.

9. Find a function f such that every graph of arboricity at least $f(k)$ has colouring number at least k, and a function g such that every graph of colouring number at least $g(k)$ has arboricity at least k, for all $k \in \mathbb{N}$. (The *arboricity* of a graph is defined in Chapter 3.5.)

10.$^-$ A k-chromatic graph is called *critically k-chromatic*, or just *critical*, if $\chi(G - v) < k$ for every $v \in V(G)$. Show that every k-chromatic graph has a critical k-chromatic induced subgraph, and that any such subgraph has minimum degree at least $k-1$.

11. Determine the critical 3-chromatic graphs.

12.$^+$ Show that every critical k-chromatic graph is $(k-1)$ - edge-connected.

13. Given $k \in \mathbb{N}$, find a constant $c_k > 0$ such that every graph G with $|G| \geqslant 3k$ and $\alpha(G) \leqslant k$ contains a cycle of length at least $c_k\,|G|$.

14.$^-$ Find a graph G for which Brooks's theorem yields a significantly weaker bound on $\chi(G)$ than Proposition 5.2.2.

15.$^+$ Show that, in order to prove Brooks's theorem for a graph $G = (V, E)$, we may assume that $\kappa(G) \geqslant 2$ and $\Delta(G) \geqslant 3$. Prove the theorem under these assumptions, showing first the following two lemmas.

   (i) Let $v_1, \dots, v_n$ be an enumeration of V. If every v_i $(i < n)$ has a neighbour v_j with $j > i$, and if $v_1v_n, v_2v_n \in E$ but $v_1v_2 \notin E$, then the greedy algorithm uses at most $\Delta(G)$ colours.

   (ii) If G is not complete and v_n has maximum degree in G, then v_n has neighbours v_1, v_2 as in (i).

16. Given a graph G and $k \in \mathbb{N}$, let $P_G(k)$ denote the number of vertex colourings $V(G) \to \{1, \dots, k\}$. Show that P_G is a polynomial in k of degree $n := |G|$, in which the coefficient of k^n is 1 and the coefficient of k^{n-1} is $-\|G\|$. (P_G is called the *chromatic polynomial* of G.)

    (Hint. Apply induction on $\|G\|$. In the induction step, compare the values of $P_G(k)$, $P_{G-e}(k)$ and $P_{G/e}(k)$.)

17.$^+$ Determine the class of all graphs G for which $P_G(k) = k\,(k-1)^{n-1}$. (As in the previous exercise, let $n := |G|$, and let P_G denote the chromatic polynomial of G.)

18. In the definition of k-constructible graphs, replace the axiom (ii) by

    (ii)$'$ *Every supergraph of a k-constructible graph is k-constructible*;

    and the axiom (iii) by

    (iii)$'$ *If G is a graph with vertices x, y_1, y_2 such that $y_1y_2 \in E(G)$ but $xy_1, xy_2 \notin E(G)$, and if both $G + xy_1$ and $G + xy_2$ are k-constructible, then G is k-constructible.*

    Show that a graph is k-constructible with respect to this new definition if and only if its chromatic number is at least k.

37. A graph G is called an *interval graph* if there exists a set $\{ I_v \mid v \in V(G) \}$ of real intervals such that $I_u \cap I_v \neq \emptyset$ if and only if $uv \in E(G)$.

    (i) Show that every interval graph is chordal.

    (ii) Show that the complement of any interval graph is a comparability graph.

    (Conversely, a chordal graph is an interval graph if its complement is a comparability graph; this is a theorem of Gilmore and Hoffman (1964).)

38. Show that $\chi(H) \in \{ \omega(H) , \omega(H) + 1 \}$ for every line graph H.

39.$^+$ Characterize the graphs whose line graphs are perfect.

40. Show that a graph G is perfect if and only if every non-empty induced subgraph H of G contains an independent set $A \subseteq V(H)$ such that $\omega(H - A) < \omega(H)$.

41.$^+$ Consider the graphs G for which every induced subgraph H has the property that every maximal complete subgraph of H meets every maximal independent vertex set in H.

    (i) Show that these graphs G are perfect.

    (ii) Show that these graphs G are precisely the graphs not containing an induced copy of P^3.

42.$^+$ Show that in every perfect graph G one can find a set $\mathcal{A}$ of independent vertex sets and a set $\mathcal{O}$ of vertex sets of complete subgraphs such that $\bigcup \mathcal{A} = V(G) = \bigcup \mathcal{O}$ and every set in $\mathcal{A}$ meets every set in $\mathcal{O}$.

    (Hint. Lemma 5.5.4.)

43.$^+$ Let G be a perfect graph. As in the proof of Theorem 5.5.3, replace every vertex x of G with a perfect graph G_x (not necessarily complete). Show that the resulting graph G' is again perfect.

44. Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be two sets of imperfect graphs, each minimal with the property that a graph is perfect if and only if it has no induced subgraph in $\mathcal{H}_i$ $(i = 1, 2)$. Do $\mathcal{H}_1$ and $\mathcal{H}_2$ contain the same graphs, up to isomorphism?

Notes

The authoritative reference work on all questions of graph colouring is T.R. Jensen & B. Toft, *Graph Coloring Problems*, Wiley 1995. Starting with a brief survey of the most important results and areas of research in the field, this monograph gives a detailed account of over 200 open colouring problems, complete with extensive background surveys and references. Most of the remarks below are discussed comprehensively in this book, and all the references for this chapter can be found there.

The *four colour problem*, whether every map can be coloured with four colours so that adjacent countries are shown in different colours, was raised by a certain Francis Guthrie in 1852. He put the question to his brother Frederick, who was then a mathematics undergraduate in Cambridge. The problem was classical graph invariants (including a proof of Theorem 5.4.1) is given by N. Alon, Restricted colorings of graphs, in (K. Walker, ed.) *Surveys in Combinatorics*, LMS Lecture Notes **187**, Cambridge University Press 1993. Both the list colouring conjecture and Galvin's proof of the bipartite case are originally stated for multigraphs. Kahn (1994) proved that the conjecture is asymptotically correct, as follows: given any $\epsilon > 0$, every graph G with large enough maximum degree satisfies $\mathrm{ch}'(G) \leqslant (1+\epsilon)\Delta(G)$.

The total colouring conjecture was proposed around 1965 by Vizing and by Behzad; see Jensen & Toft for details.

A gentle introduction to the basic facts about perfect graphs and their applications is given by M.C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press 1980. Our first proof of the perfect graph theorem follows L. Lovász's survey on perfect graphs in (L.W. Beineke and R.J. Wilson, eds.) *Selected Topics in Graph Theory 2*, Academic Press 1983. The theorem was also proved independently, and only a little later, by Fulkerson. Our second proof, the proof of Theorem 5.5.5, is due to G.S. Gasparian, Minimal imperfect graphs: a simple approach, *Combinatorica* **16** (1996), 209–212. The approximate proof of the perfect graph conjecture is due to H.J. Prömel & A. Steger, Almost all Berge graphs are perfect, *Combinatorics, Probability and Computing* **1** (1992), 53–79.

6.1 Circulations

$G = (V, E)$

In the context of flows, we have to be able to speak about the 'directions' of an edge. Since, in a multigraph $G = (V, E)$, an edge $e = xy$ is not identified uniquely by the pair (x, y) or (y, x), we define directed edges as triples:

$\vec{E}$

$$\vec{E} := \{\, (e, x, y) \mid e \in E;\ x, y \in V;\ e = xy \,\}\,.$$

direction (e, x, y)

$\overleftarrow{e}$

Thus, an edge $e = xy$ with $x \neq y$ has the two *directions* (e, x, y) and (e, y, x); a loop $e = xx$ has only one direction, the triple (e, x, x). For given $\vec{e} = (e, x, y) \in \vec{E}$, we set $\overleftarrow{e} := (e, y, x)$, and for an arbitrary set $\vec{F} \subseteq \vec{E}$ of edge directions we put

$\overleftarrow{F}$

$$\overleftarrow{F} := \{\, \overleftarrow{e} \mid \vec{e} \in \vec{F} \,\}\,.$$

Note that $\vec{E}$ itself is symmetrical: $\overleftarrow{E} = \vec{E}$. For $X, Y \subseteq V$ and $\vec{F} \subseteq \vec{E}$, define

$\vec{F}(X, Y)$

$$\vec{F}(X, Y) := \{\, (e, x, y) \in \vec{F} \mid x \in X;\ y \in Y;\ x \neq y \,\}\,,$$

$\vec{F}(x, Y)$

abbreviate $\vec{F}(\{x\}, Y)$ to $\vec{F}(x, Y)$ etc., and write

$\vec{F}(x)$

$$\vec{F}(x) := \vec{F}(x, V) = \vec{F}(\{x\}, \overline{\{x\}})\,.$$

$\overline{X}$

Here, as below, $\overline{X}$ denotes the complement $V \smallsetminus X$ of a vertex set $X \subseteq V$. Note that any loops at vertices $x \in X \cap Y$ are disregarded in the definitions of $\vec{F}(X, Y)$ and $\vec{F}(x)$.

0

f

Let H be an abelian semigroup,[2] written additively with zero 0. Given vertex sets $X, Y \subseteq V$ and a function $f\colon \vec{E} \to H$, let

$f(X, Y)$

$$f(X, Y) := \sum_{\vec{e} \in \vec{E}(X, Y)} f(\vec{e})\,.$$

$f(x, Y)$

circulation

Instead of $f(\{x\}, Y)$ we again write $f(x, Y)$, etc.

From now on, we assume that H is a group. We call f a *circulation* on G (with values in H), or an H-*circulation*, if f satisfies the following two conditions:

(F1) $f(e, x, y) = -f(e, y, x)$ for all $(e, x, y) \in \vec{E}$ with $x \neq y$;

(F2) $f(v, V) = 0$ for all $v \in V$.

---

[2] This chapter contains no group theory. The only semigroups we ever consider for H are the natural numbers, the integers, the reals, the cyclic groups $\mathbb{Z}_k$, and (once) the Klein four-group.

(F1) $f(e,x,y) = -f(e,y,x)$ for all $(e,x,y) \in \vec{E}$ with $x \neq y$;
(F2′) $f(v,V) = 0$ for all $v \in V \smallsetminus \{\, s,t \,\}$;
(F3) $f(\vec{e}) \leqslant c(\vec{e})$ for all $\vec{e} \in \vec{E}$.

integral

We call f *integral* if all its values are integers.

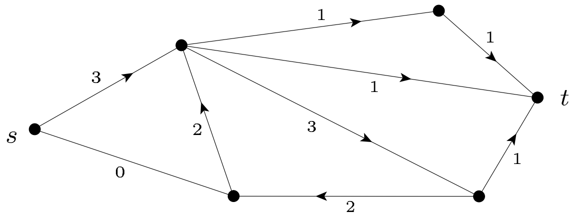

Fig. 6.2.1. A network flow in short notation: all values refer to the direction indicated (capacities are not shown)

f
cut in N
capacity

Let f be a flow in N. If $S \subseteq V$ is such that $s \in S$ and $t \in \overline{S}$, we call the pair $(S, \overline{S})$ a *cut in* N, and $c(S, \overline{S})$ the *capacity* of this cut.

Since f now has to satisfy only (F2′) rather than (F2), we no longer have $f(X, \overline{X}) = 0$ for all $X \subseteq V$ (as in Proposition 6.1.1). However, the value is the same for all cuts:

Proposition 6.2.1. *Every cut* $(S, \overline{S})$ *in* N *satisfies* $f(S, \overline{S}) = f(s, V)$.

Proof. As in the proof of Proposition 6.1.1, we have

$$\begin{aligned} f(S,\overline{S}) &= f(S,V) - f(S,S) \\ &\underset{\text{(F1)}}{=} f(s,V) + \sum_{v \in S \smallsetminus \{\, s \,\}} f(v,V) \; - \; 0 \\ &\underset{\text{(F2′)}}{=} f(s,V)\,. \end{aligned}$$

□

total value
$|f|$

The common value of $f(S, \overline{S})$ in Proposition 6.2.1 will be called the *total value* of f and denoted by $|f|$;[3] the flow shown in Figure 6.2.1 has total value 3.

By (F3), we have

$$|f| = f(S, \overline{S}) \leqslant c(S, \overline{S})$$

for every cut $(S, \overline{S})$ in N. Hence the total value of a flow in N is never larger than the smallest capacity of a cut. The following *max-flow min-cut* theorem states that this upper bound is always attained by some flow:

---

[3] Thus, formally, $|f|$ may be negative. In practice, however, we can change the sign of $|f|$ simply by swapping the roles of s and t.

Clearly, f_{n+1} is again an integral flow in N. Let us compute its total value $|f_{n+1}| = f_{n+1}(s,V)$. Since W contains the vertex s only once, $\vec{e}_0$ is the only triple (e,x,y) with $x = s$ and $y \in V$ whose f-value was changed. This value, and hence that of $f_{n+1}(s,V)$ was raised. Therefore $|f_{n+1}| > |f_n|$ as desired.

If $t \notin S_n$, then $(S_n, \overline{S_n})$ is a cut in N. By (F3) for f_n, and the definition of S_n, we have

$$f_n(\vec{e}) = c(\vec{e})$$

for all $\vec{e} \in \vec{E}(S_n, \overline{S_n})$, so

$$|f_n| = f_n(S_n, \overline{S_n}) = c(S_n, \overline{S_n})$$

as desired. □

Since the flow constructed in the proof of Theorem 6.2.2 is integral, we have also proved the following:

Corollary 6.2.3. *In every network (with integral capacity function) there exists an integral flow of maximum total value.* □

6.3 Group-valued flows

Let $G = (V,E)$ be a multigraph and H an abelian group. If f and g are two H-circulations then, clearly, $(f+g)\colon \vec{e} \mapsto f(\vec{e}) + g(\vec{e})$ and $-f\colon \vec{e} \mapsto -f(\vec{e})$ are again H-circulations. The H-circulations on G thus form a group in a natural way.

$f+g$
$-f$

nowhere zero

H-flow

A function $f\colon \vec{E} \to H$ is *nowhere zero* if $f(\vec{e}) \neq 0$ for all $\vec{e} \in \vec{E}$. An H-circulation that is nowhere zero is called an *H-flow*.[4] Note that the set of H-flows on G is not closed under addition: if two H-flows add up to zero on some edge $\vec{e}$, then their sum is no longer an H-flow. By Corollary 6.1.2, a graph with an H-flow cannot have a bridge.

For finite groups H, the number of H-flows on G—and, in particular, their existence—surprisingly depends only on the order of H, not on H itself:

Theorem 6.3.1. (Tutte 1954)
For every multigraph G there exists a polynomial P such that, for any finite abelian group H, the number of H-flows on G is $P\big(|H|-1\big)$.

---

[4] This terminology seems simplest for our purposes but is not standard; see the notes.

and the set of H-flows on G_2. Given f, let g be the restriction of f to $\vec{E}' \smallsetminus \vec{E}'(y,x)$. (As the x–y edges $e \in E'$ become loops in G_2, they have only the one direction (e, v_0, v_0) there; as its g-value, we choose $f(e,x,y)$.) Then g is indeed an H-flow on G_2; note that (F2) holds at v_0 by Proposition 6.1.1 for G, with $X := \{x, y\}$.

It remains to show that the map $f \mapsto g$ is a bijection. If we are given an H-flow g on G_2 and try to find an $f \in F_2$ with $f \mapsto g$, then $f(\vec{e})$ is already determined as $f(\vec{e}) = g(\vec{e})$ for all $\vec{e} \in \vec{E}' \smallsetminus \vec{E}'(y,x)$; by (F1), we further have $f(\vec{e}) = -f(\overleftarrow{e})$ for all $\vec{e} \in \vec{E}'(y,x)$. Thus our map $f \mapsto g$ is bijective if and only if for given g there is always a unique way to define the remaining values of $f(\vec{e}_0)$ and $f(\overleftarrow{e}_0)$ so that f satisfies (F1) in e_0 and (F2) in x and y.

V'

This is indeed the case. Let $V' := V \smallsetminus \{x, y\}$. As g satisfies (F2), the f-values fixed already are such that

$$f(x, V') + f(y, V') = g(v_0, V') = 0\,. \tag{2}$$

With

$$h := \sum_{\vec{e} \in \vec{E}'(x,y)} f(\vec{e}) \quad \Big(= \sum_{e \in E'(x,y)} g(e, v_0, v_0) \Big),$$

(F2) for f requires

$$0 = f(x, V) = f(\vec{e}_0) + h + f(x, V')$$

and

$$0 = f(y, V) = f(\overleftarrow{e}_0) - h + f(y, V')\,,$$

so we have to set

$$f(\vec{e}_0) := -f(x, V') - h \quad \text{and} \quad f(\overleftarrow{e}_0) := -f(y, V') + h\,.$$

Then $f(\vec{e}_0) + f(\overleftarrow{e}_0) = 0$ by (2), so f also satisfies (F1) in e_0. □

flow polynomial

The polynomial P of Theorem 6.3.1 is known as the *flow polynomial* of G.

[6.4.5]

Corollary 6.3.2. *If H and H' are two finite abelian groups of equal order, then G has an H-flow if and only if G has an H'-flow.* □

Corollary 6.3.2 has fundamental implications for the theory of algebraic flows: it indicates that crucial difficulties in existence proofs of H-flows are unlikely to be of a group-theoretic nature. On the other hand, being able to choose a convenient group can be quite helpful; we shall see a pretty example for this in Proposition 6.4.5.

K

$$K(f) := \sum_{x\in V} |f(x,V)|$$

of all deviations from Kirchhoff's law is least possible. We shall prove that $K(f) = 0$; then, clearly, $f(x,V) = 0$ for every x, as desired.

x

Suppose $K(f) \neq 0$. Since f satisfies (F1), and hence $\sum_{x\in V} f(x,V) = f(V,V) = 0$, there exists a vertex x with

$$f(x,V) > 0\,. \tag{1}$$

X X'

Let $X \subseteq V$ be the set of all vertices x' for which G contains a walk $x_0e_0\dots e_{\ell-1}x_\ell$ from x to x' such that $f(e_i,x_i,x_{i+1}) > 0$ for all $i < \ell$; furthermore, let $X' := X \smallsetminus \{x\}$.

We first show that X' contains a vertex x' with $f(x',V) < 0$. By definition of X, we have $f(e,x',y) \leqslant 0$ for all edges $e = x'y$ such that $x' \in X$ and $y \in \overline{X}$. In particular, this holds for $x' = x$. Thus, (1) implies $f(x,X') > 0$. Then $f(X',x) < 0$ by (F1), as well as $f(X',X') = 0$. Therefore

$$\sum_{x'\in X'} f(x',V) = f(X',V) = f(X',\overline{X}) + f(X',x) + f(X',X') < 0\,,$$

x'

so some $x' \in X'$ must indeed satisfy

$$f(x',V) < 0\,. \tag{2}$$

W f'

As $x' \in X$, there is an x–x' walk $W = x_0e_0\dots e_{\ell-1}x_\ell$ such that $f(e_i,x_i,x_{i+1}) > 0$ for all $i < \ell$. We now modify f by sending some flow back along W, letting $f'\colon \vec{E} \to \mathbb{Z}$ be given by

$$f'\colon \vec{e} \mapsto \begin{cases} f(\vec{e}) - k & \text{for } \vec{e} = (e_i,x_i,x_{i+1}),\ \ i = 0,\dots,\ell-1;\\ f(\vec{e}) + k & \text{for } \vec{e} = (e_i,x_{i+1},x_i),\ \ i = 0,\dots,\ell-1;\\ f(\vec{e}) & \text{for } e \notin W. \end{cases}$$

By definition of W, we have $|f'(\vec{e})| < k$ for all $\vec{e} \in \vec{E}$. Hence f', like f, lies in F.

How does the modification of f affect K? At all inner vertices v of W, as well as outside W, the deviation from Kirchhoff's law remains unchanged:

$$f'(v,V) = f(v,V) \qquad \text{for all } v \in V \smallsetminus \{x,x'\}. \tag{3}$$

For x and x', on the other hand, we have

$$f'(x,V) = f(x,V) - k \quad \text{and} \quad f'(x',V) = f(x',V) + k\,. \tag{4}$$

(1.6.1)
(6.3.3)

Proof. Let $G = (V, E)$ be a cubic graph. Let us assume first that G has a 3-flow, and hence also a $\mathbb{Z}_3$-flow f. We show that any cycle $C = x_0 \ldots x_\ell x_0$ in G has even length (cf. Proposition 1.6.1). Consider two consecutive edges on C, say $e_{i-1} := x_{i-1}x_i$ and $e_i := x_i x_{i+1}$. If f assigned the same value to these edges in the direction of the forward orientation of C, i.e. if $f(e_{i-1}, x_{i-1}, x_i) = f(e_i, x_i, x_{i+1})$, then f could not satisfy (F2) at x_i for any non-zero value of the third edge at x_i. Therefore f assigns the values $\overline{1}$ and $\overline{2}$ to the edges of C alternately, and in particular C has even length.

Conversely, let G be bipartite, with vertex bipartition $\{X, Y\}$. Since G is cubic, the map $\vec{E} \to \mathbb{Z}_3$ defined by $f(e, x, y) := \overline{1}$ and $f(e, y, x) := \overline{2}$ for all edges $e = xy$ with $x \in X$ and $y \in Y$ is a $\mathbb{Z}_3$-flow on G. By Theorem 6.3.3, then, G has a 3-flow. □

What are the flow numbers of the complete graphs K^n? For odd $n > 1$, we have $\varphi(K^n) = 2$ by Proposition 6.4.1. Moreover, $\varphi(K^2) = \infty$, and $\varphi(K^4) = 4$; this is easy to see directly (and it follows from Propositions 6.4.2 and 6.4.5). Interestingly, K^4 is the only complete graph with flow number 4:

Proposition 6.4.3. *For all even $n > 4$, $\varphi(K^n) = 3$.*

(6.3.3)

Proof. Proposition 6.4.1 implies that $\varphi(K^n) \geqslant 3$ for even n. We show, by induction on n, that every $G = K^n$ with even $n > 4$ has a 3-flow.

For the induction start, let $n = 6$. Then G is the edge-disjoint union of three graphs G_1, G_2, G_3, with $G_1, G_2 = K^3$ and $G_3 = K_{3,3}$. Clearly G_1 and G_2 each have a 2-flow, while G_3 has a 3-flow by Proposition 6.4.2. The union of all these flows is a 3-flow on G.

Now let $n > 6$, and assume the assertion holds for $n-2$. Clearly, G is the edge-disjoint union of a K^{n-2} and a graph $G' = (V', E')$ with $G' = \overline{K^{n-2}} * K^2$. The K^{n-2} has a 3-flow by induction. By Theorem 6.3.3, it thus suffices to find a $\mathbb{Z}_3$-flow on G'. For every vertex z of the $\overline{K^{n-2}} \subseteq G'$, let f_z be a $\mathbb{Z}_3$-flow on the triangle $zxyz \subseteq G'$, where $e = xy$ is the edge of the K^2 in G'. Let $f\colon \vec{E'} \to \mathbb{Z}_3$ be the sum of these flows. Clearly, f is nowhere zero, except possibly in (e, x, y) and (e, y, x). If $f(e, x, y) \neq \overline{0}$, then f is the desired $\mathbb{Z}_3$-flow on G'. If $f(e, x, y) = \overline{0}$, then $f + f_z$ (for any z) is a $\mathbb{Z}_3$-flow on G'. □

Proposition 6.4.4. *Every 4-edge-connected graph has a 4-flow.*

(3.5.2)

Proof. Let G be a 4-edge-connected graph. By Corollary 3.5.2, G has two edge-disjoint spanning trees T_i, $i = 1, 2$. For each edge $e \notin T_i$, let $C_{i,e}$ be the unique cycle in $T_i + e$, and let $f_{i,e}$ be a $\mathbb{Z}_4$-flow of value $\overline{i}$ around $C_{i,e}$—more precisely: a $\mathbb{Z}_4$-circulation on G with values $\overline{i}$ and $-\overline{i}$ on the edges of $C_{i,e}$ and zero otherwise. $f_{1,e}, f_{2,e}$

6.5 Flow-colouring duality

In this section we shall see a surprising connection between flows and colouring: every k-flow on a plane multigraph gives rise to a k-vertex-colouring of its dual, and vice versa. In this way, the investigation of k-flows appears as a natural generalization of the familiar map colouring problems in the plane.

$G = (V, E)$
G^*

Let $G = (V, E)$ and $G^* = (V^*, E^*)$ be dual plane multigraphs. For simplicity, let us assume that G and G^* have neither bridges nor loops and are non-trivial. For edge sets $F \subseteq E$, let us write

F^*

$$F^* := \{\, e^* \in E^* \mid e \in F \,\}.$$

Conversely, if a subset of E^* is given, we shall usually write it immediately in the form F^*, and thus let $F \subseteq E$ be defined implicitly via the bijection $e \mapsto e^*$.

Suppose we are given a circulation g on G^*: how can we employ the duality between G and G^* to derive from g some information about G? The most general property of all circulations is Proposition 6.1.1, which says that $g(X, \overline{X}) = 0$ for all $X \subseteq V^*$. By Proposition 4.6.1, the minimal cuts $E^*(X, \overline{X})$ in G^* correspond precisely to the cycles in G. Thus if we take the composition f of the maps $e \mapsto e^*$ and g, and sum its values over the edges of a cycle in G, then this sum should again be zero.

Of course, there is still a technical hitch: since g takes its arguments not in E^* but in $\vec{E^*}$, we cannot simply define f as above: we first have to refine the bijection $e \mapsto e^*$ into one from $\vec{E}$ to $\vec{E^*}$, i.e. assign to every $\vec{e} \in \vec{E}$ canonically one of the two directions of e^*. This will be the purpose of our first lemma. After that, we shall show that f does indeed sum to zero along any cycle in G.

If $C = v_0 \dots v_{\ell-1} v_0$ is a cycle with edges $e_i = v_i v_{i+1}$ (and $v_\ell := v_0$), we shall call

$\vec{C}$

$$\vec{C} := \{\, (e_i, v_i, v_{i+1}) \mid i < \ell \,\}$$

cycle with orientation

a *cycle with orientation*. Note that this definition of $\vec{C}$ depends on the vertex enumeration chosen to denote C: every cycle has two orientations. Conversely, of course, C can be reconstructed from the set $\vec{C}$. In practice, we shall therefore speak about C freely even when, formally, only $\vec{C}$ has been defined.

Lemma 6.5.1. *There exists a bijection* $^*\colon \vec{e} \mapsto \vec{e}^{\,*}$ *from* $\vec{E}$ *to* $\vec{E^*}$ *with the following properties.*

(i) *The underlying edge of* $\vec{e}^{\,*}$ *is always* e^*, *i.e.* $\vec{e}^{\,*}$ *is one of the two directions* $\vec{e^*}, \overleftarrow{e^*}$ *of* e^*.

(ii) *If* $C \subseteq G$ *is a cycle,* $F := E(C)$, *and if* $X \subseteq V^*$ *is such that* $F^* = E^*(X, \overline{X})$, *then there exists an orientation* $\vec{C}$ *of* C *with* $\{\, \vec{e}^{\,*} \mid \vec{e} \in \vec{C} \,\} = \vec{E^*}(X, \overline{X})$.

the corresponding value for our given orientation of C must be zero.

For the backward implication it suffices by (i) to show that g satisfies (F2), i.e. that $g(x, V^*) = 0$ for every $x \in V^*$. We shall prove that $g(x, V(B)) = 0$ for every block B of G^* containing x; since every edge of G^* at x lies in exactly one such block, this will imply $g(x, V^*) = 0$.

B

F^*, F

So let $x \in V^*$ be given, and let B be any block of G^* containing x. Since G^* is a non-trivial plane dual, and hence connected, we have $B - x \neq \emptyset$. Let F^* be the set of all edges of B at x (Fig. 6.5.2),

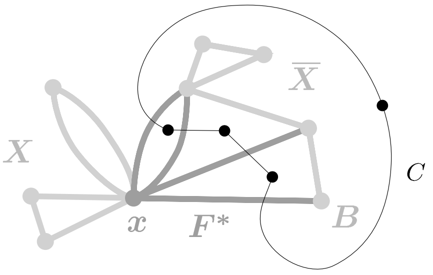

Fig. 6.5.2. The cut F^* in G^*

X

and let X be the vertex set of the component of $G^* - F^*$ containing x. Then $\emptyset \neq V(B - x) \subseteq \overline{X}$, by the maximality of B as a cutvertex-free subgraph. Hence

$$F^* = E^*(X, \overline{X}) \tag{1}$$

by definition of X, i.e. F^* is a cut in G^*. As a dual, G^* is connected, so $G^*[\,\overline{X}\,]$ too is connected. Indeed, every vertex of $\overline{X}$ is linked to x by a path $P \subseteq G^*$ whose last edge lies in F^*. Then $P - x$ is a path in $G^*[\,\overline{X}\,]$ meeting B. Since x does not separate B, this shows that $G^*[\,\overline{X}\,]$ is connected.

C

Thus, X and $\overline{X}$ are both connected in G^*, so F^* is even a minimal cut in G^*. Let $C \subseteq G$ be the cycle with $E(C) = F$ that exists by Proposition 4.6.1. By Lemma 6.5.1 (ii), C has an orientation $\vec{C}$ such that $\{\, \vec{e}^{\,*} \mid \vec{e} \in \vec{C}\,\} = \vec{E^*}(X, \overline{X})$. By (1), however, $\vec{E^*}(X, \overline{X}) = \vec{E^*}(x, V(B))$, so

$$g(x, V(B)) = g(X, \overline{X}) = f(\vec{C}) = 0$$

by definition of f and g. □

With the help of Lemma 6.5.2, we can now prove our colouring-flow duality theorem for plane multigraphs. If $P = v_0 \dots v_\ell$ is a path with edges $e_i = v_i v_{i+1}$ ($i < \ell$), we set (depending on our vertex enumeration of P)

$\vec{P}$

$$\vec{P} := \{\, (e_i, v_i, v_{i+1}) \mid i < \ell \,\}$$

$v_0 \to v_\ell$ *path*

and call $\vec{P}$ a $v_0 \to v_\ell$ *path*. Again, P may be given implicitly by $\vec{P}$.

6.6 Tutte's flow conjectures

How can we determine the flow number of a graph? Indeed, does every (bridgeless) graph have a flow number, a k-flow for some k? Can flow numbers, like chromatic numbers, become arbitrarily large? Can we characterize the graphs admitting a k-flow, for given k?

Of these four questions, we shall answer the second and third in this section: we prove that every bridgeless graph has a 6-flow. In particular, a graph has a flow number if and only if it has no bridge. The question asking for a characterization of the graphs with a k-flow remains interesting for $k = 3, 4, 5$. Partial answers are suggested by the following three conjectures of Tutte, who initiated algebraic flow theory.

The oldest and best known of the Tutte conjectures is his *5-flow conjecture*:

Five-Flow Conjecture. (Tutte 1954)
Every bridgeless multigraph has a 5-flow.

Which graphs have a 4-flow? By Proposition 6.4.4, the 4-edge-connected graphs are among them. The Petersen graph (Fig. 6.6.1), on the other hand, is an example of a bridgeless graph without a 4-flow: since it is cubic but not 3-edge-colourable (Ex. 19, Ch. 5), it cannot have a 4-flow by Proposition 6.4.5 (ii).

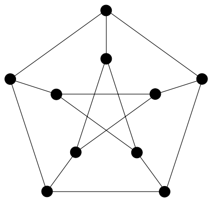

Fig. 6.6.1. The Petersen graph

Tutte's *4-flow conjecture* states that the Petersen graph must be present in every graph without a 4-flow:

Four-Flow Conjecture. (Tutte 1966)
Every bridgeless multigraph not containing the Petersen graph as a minor has a 4-flow.

By Proposition 1.7.2, we may replace the word 'minor' in the 4-flow conjecture by 'topological minor'.

$H_0, \dots, H_n$ $F_1, \dots, F_n$ V_i, E_i

We shall construct a sequence $H_0, \dots, H_n$ of disjoint connected and even subgraphs of G, together with a sequence $F_1, \dots, F_n$ of non-empty sets of edges between them. The sets F_i will each contain only one or two edges, between H_i and $H_0 \cup \dots \cup H_{i-1}$. We write $H_i =: (V_i, E_i)$,

H^i

$$H^i := (H_0 \cup \dots \cup H_i) + (F_1 \cup \dots \cup F_i)$$

V^i, E^i

and $H^i =: (V^i, E^i)$. Note that each $H^i = (H^{i-1} \cup H_i) + F_i$ is connected (induction on i). Our assumption that H_i is even implies by Proposition 6.4.1 (or directly by Proposition 1.2.1) that H_i has no bridge.

n X_i

As H_0 we choose any K^1 in G. Now assume that $H_0, \dots, H_{i-1}$ and $F_1, \dots, F_{i-1}$ have been defined for some $i > 0$. If $V^{i-1} = V$, we terminate the construction and set $i - 1 =: n$. Otherwise, we let $X_i \subseteq \overline{V^{i-1}}$ be minimal such that $X_i \neq \emptyset$ and

$$\left| E(X_i, \overline{V^{i-1}} \setminus X_i) \right| \leqslant 1 \tag{1}$$

F_i

(Fig. 6.6.2); such an X_i exists, because $\overline{V^{i-1}}$ is a candidate. Since G is 2-edge-connected, (1) implies that $E(X_i, V^{i-1}) \neq \emptyset$. By the minimality of X_i, the graph $G[X_i]$ is connected and bridgeless, i.e. 2-edge-connected or a K^1. As the elements of F_i we pick one or two edges from $E(X_i, V^{i-1})$, if possible two. As H_i we choose any connected even subgraph of $G[X_i]$ containing the ends in X_i of the edges in F_i.

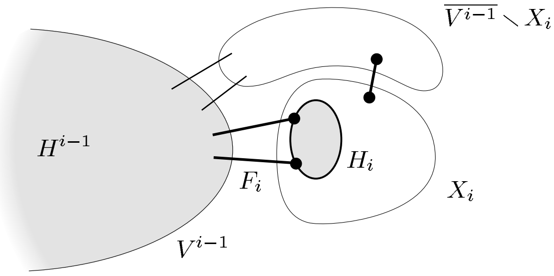

Fig. 6.6.2. Constructing the H_i and F_i

H E'

When our construction is complete, we set $H^n =: H$ and $E' := E \setminus E(H)$. By definition of n, H is a spanning connected subgraph of G.

$f_n, \dots, f_0$ $\vec{C_e}$ f_e f_n

We now define, by 'reverse' induction, a sequence $f_n, \dots, f_0$ of $\mathbb{Z}_3$-circulations on G. For every edge $e \in E'$, let $\vec{C_e}$ be a cycle (with orientation) in $H + e$ containing e, and f_e a positive flow around $\vec{C_e}$; formally, we let f_e be a $\mathbb{Z}_3$-circulation on G such that $f_e^{-1}(\overline{0}) = \vec{E} \setminus (\vec{C_e} \cup \overleftarrow{C_e})$. Let f_n be the sum of all these f_e. Since each $e' \in E'$ lies on just one of the cycles C_e (namely, on $C_{e'}$), we have $f_n(\vec{e}) \neq \overline{0}$ for all $\vec{e} \in \vec{E'}$.

Exercises

1.$^-$ Prove Proposition 6.2.1 by induction on $|S|$.

2. (i)$^-$ Given $n \in \mathbb{N}$, find a capacity function for the network below such that the algorithm from the proof of the max-flow min-cut theorem will need more than n augmenting paths W if these are badly chosen.

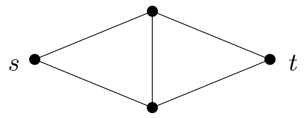

(ii)$^+$ Show that, if all augmenting paths are chosen as short as possible, their number is bounded by a function of the size of the network.

3.$^+$ Derive Menger's Theorem 3.3.4 from the max-flow min-cut theorem.

(Hint. The edge version is easy. For the vertex version, apply the edge version to a suitable auxiliary graph.)

4.$^-$ Let f be an H-circulation on G and $g\colon H \to H'$ a group homomorphism. Show that $g \circ f$ is an H'-circulation on G. Is $g \circ f$ an H'-flow if f is an H-flow?

5.$^-$ Given $k \geqslant 1$, show that a graph has a k-flow if and only if each of its blocks has a k-flow.

6.$^-$ Show that $\varphi(G/e) \leqslant \varphi(G)$ whenever G is a multigraph and e an edge of G. Does this imply that, for every k, the class of all multigraphs admitting a k-flow is closed under taking minors?

7.$^-$ Work out the flow number of K^4 directly, without using any results from the text.

8. Let H be a finite abelian group, G a graph, and T a spanning tree of G. Show that every mapping from the directions of $E(G) \setminus E(T)$ to H that satisfies (F1) extends uniquely to an H-circulation on G.

Do not use the 6-flow Theorem 6.6.1 for the following three exercises.

9. Show that $\varphi(G) < \infty$ for every bridgeless multigraph G.

10. Assume that a graph G has m spanning trees such that no edge of G lies in all of these trees. Show that $\varphi(G) \leqslant 2^m$.

11.$^+$ Let G be a bridgeless connected graph with n vertices and m edges. By considering a normal spanning tree of G, show that $\varphi(G) \leqslant m-n+2$.

12. Show that every graph with a Hamilton cycle has a 4-flow. (A *Hamilton cycle* of G is a cycle in G that contains all the vertices of G.)

13. A family of (not necessarily distinct) cycles in a graph G is called a *cycle double cover* of G if every edge of G lies on exactly two of these cycles. The *cycle double cover conjecture* asserts that every bridgeless multigraph has a cycle double cover. Prove the conjecture for graphs with a 4-flow.

covered also in F. Jaeger's survey, Nowhere-zero[7] flow problems, in (L.W. Beineke & R.J. Wilson, eds.) *Selected Topics in Graph Theory 3*, Academic Press 1988. For the flow conjectures, see also T.R. Jensen & B. Toft, *Graph Coloring Problems*, Wiley 1995. Seymour's 6-flow theorem is proved in P.D. Seymour, Nowhere-zero 6-flows, *J. Combin. Theory B* **30** (1981), 130–135. This paper also indicates how Tutte's 5-flow conjecture reduces to snarks. In 1998, Robertson, Sanders, Seymour and Thomas announced a proof of the 4-flow conjecture for cubic graphs.

Finally, Tutte discovered a 2-variable polynomial associated with a graph, which generalizes both its chromatic polynomial and its flow polynomial. What little is known about this *Tutte polynomial* can hardly be more than the tip of the iceberg: it has far-reaching, and largely unexplored, connections to areas as diverse as knot theory and statistical physics. See D.J.A. Welsh, *Complexity: knots, colourings and counting* (LMS Lecture Notes **186**), Cambridge University Press 1993.

---

[7] In the literature, the term 'flow' is often used to mean what we have called 'circulation', i.e. flows are not required to be nowhere zero unless this is stated explicitly.

given any graph H that contains at least one cycle, there are graphs of arbitrarily large chromatic number not containing H as a subgraph (Theorem 11.2.2). By Corollary 5.2.3 and Theorem 1.4.2, such graphs have subgraphs of arbitrarily large average degree and connectivity, so these invariants too can be large without the presence of an H subgraph.

Thus, unless H is a forest, the only way to force the presence of an H subgraph in an arbitrary graph G by global assumptions on G is to raise $\|G\|$ substantially above any value implied by large values of the above invariants. If H is not bipartite, then any function f such that $f(n)$ edges on n vertices force an H subgraph must even grow quadratically with n: since complete bipartite graphs can have $\frac{1}{4}n^2$ edges, $f(n)$ must exceed $\frac{1}{4}n^2$.

dense

edge density

Graphs with a number of edges roughly[1] quadratic in their number of vertices are usually called *dense*; the number $\|G\|/\binom{|G|}{2}$—the proportion of its potential edges that G actually has—is the *edge density* of G. The question of exactly which edge density is needed to force a given subgraph is the archetypal extremal graph problem in its original (narrower) sense; it is the topic of this chapter. Rather than attempting to survey the wide field of (dense) extremal graph theory, however, we shall concentrate on its two most important results and portray one powerful general proof technique.

The two results are Turán's classic extremal graph theorem for $H = K^r$, a result that has served as a model for countless similar theorems for other graphs H, and the fundamental Erdős-Stone theorem, which gives precise asymptotic information for all H at once (Section 7.1). The proof technique, one of increasing importance in the extremal theory of dense graphs, is the use of the Szemerédi *regularity lemma*. This lemma is presented and proved in Section 7.2. In Section 7.3, we outline a general method for applying the regularity lemma, and illustrate this in the proof of the Erdős-Stone theorem postponed from Section 7.1. Another application of the regularity lemma will be given in Chapter 9.2.

7.1 Subgraphs

Let H be a graph and $n \geqslant |H|$. How many edges will suffice to force an H subgraph in any graph on n vertices, no matter how these edges are arranged? Or, to rephrase the problem: which is the greatest possible number of edges that a graph on n vertices can have *without* containing a copy of H as a subgraph? What will such a graph look like? Will it be unique?

---

[1] Note that, formally, the notions of sparse and dense make sense only for families of graphs whose order tends to infinity, not for individual graphs.

Theorem 7.1.1. (Turán 1941)
[7.1.2] [9.2.2] *For all integers* r, n *with* $r > 1$, *every graph* $G \not\supseteq K^r$ *with* n *vertices and* $\mathrm{ex}(n, K^r)$ *edges is a* $T^{r-1}(n)$.

Proof. We apply induction on n. For $n \leqslant r-1$ we have $G = K^n = T^{r-1}(n)$ as claimed. For the induction step, let now $n \geqslant r$.

Since G is edge-maximal without a K^r subgraph, G has a subgraph
K $K = K^{r-1}$. By the induction hypothesis, $G - K$ has at most $t_{r-1}(n-r+1)$ edges, and each vertex of $G - K$ has at most $r-2$ neighbours in K. Hence,

$$\|G\| \leqslant t_{r-1}(n-r+1) + (n-r+1)(r-2) + \binom{r-1}{2} = t_{r-1}(n)\,; \quad (1)$$

the equality on the right follows by inspection of the Turán graph $T^{r-1}(n)$ (Fig. 7.1.3).

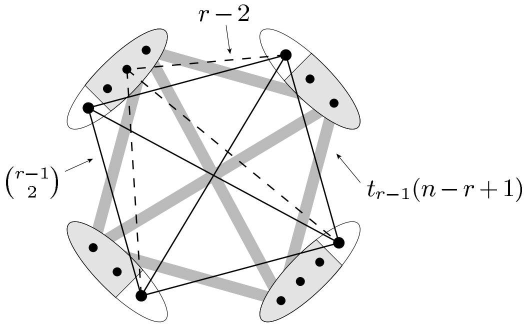

Fig. 7.1.3. The equation from (1) for $r = 5$ and $n = 14$

Since G is extremal for K^r (and $T^{r-1}(n) \not\supseteq K^r$), we have equality in (1). Thus, every vertex of $G - K$ has *exactly* $r-2$ neighbours in K—
$x_1, \dots, x_{r-1}$ just like the vertices $x_1, \dots, x_{r-1}$ of K itself. For $i = 1, \dots, r-1$ let

$V_1, \dots, V_{r-1}$
$$V_i := \{\, v \in V(G) \mid vx_i \notin E(G) \,\}$$

be the set of all vertices of G whose $r-2$ neighbours in K are precisely the vertices other than x_i. Since $K^r \not\subseteq G$, each of the sets V_i is independent, and they partition $V(G)$. Hence, G is $(r-1)$-partite. As $T^{r-1}(n)$ is the unique $(r-1)$-partite graph with n vertices and the maximum number of edges, our claim that $G = T^{r-1}(n)$ follows from the assumed extremality of G. □

The Turán graphs $T^{r-1}(n)$ are dense: in order of magnitude, they have about n^2 edges. More exactly, for every n and r we have

$$t_{r-1}(n) \leqslant \tfrac{1}{2} n^2 \,\frac{r-2}{r-1}\,,$$

r **Proof of Corollary 7.1.3.** Let $r := \chi(H)$. Since H cannot be coloured with $r-1$ colours, we have $H \not\subseteq T^{r-1}(n)$ for all $n \in \mathbb{N}$, and hence

$$t_{r-1}(n) \leqslant \mathrm{ex}(n,H)\,.$$

On the other hand, $H \subseteq K^r_s$ for all sufficiently large s, so

$$\mathrm{ex}(n,H) \leqslant \mathrm{ex}(n,K^r_s)$$

s for all those s. Let us fix such an s. For every $\epsilon > 0$, Theorem 7.1.2 implies that eventually (i.e. for large enough n)

$$\mathrm{ex}(n,K^r_s) < t_{r-1}(n) + \epsilon n^2.$$

Hence for n large,

$$\begin{aligned} t_{r-1}(n)/\tbinom{n}{2} &\leqslant \mathrm{ex}(n,H)/\tbinom{n}{2} \\ &\leqslant \mathrm{ex}(n,K^r_s)/\tbinom{n}{2} \\ &< t_{r-1}(n)/\tbinom{n}{2} + \epsilon n^2/\tbinom{n}{2} \\ &= t_{r-1}(n)/\tbinom{n}{2} + 2\epsilon/(1-\tfrac{1}{n}) \\ &\leqslant t_{r-1}(n)/\tbinom{n}{2} + 4\epsilon \qquad \text{(assume } n \geqslant 2\text{).} \end{aligned}$$

Therefore, since $t_{r-1}(n)/\binom{n}{2}$ converges to $\frac{r-2}{r-1}$ (Lemma 7.1.4), so does $\mathrm{ex}(n,H)/\binom{n}{2}$. Thus

$$\lim_{n\to\infty} \mathrm{ex}(n,H)\binom{n}{2}^{-1} = \frac{r-2}{r-1}$$

as claimed. □

For bipartite graphs H, Corollary 7.1.3 says that substantially fewer than $\binom{n}{2}$ edges suffice to force an H subgraph. It turns out that

$$c_1 n^{2-\frac{2}{r+1}} \leqslant \mathrm{ex}(n,K_{r,r}) \leqslant c_2 n^{2-\frac{1}{r}}$$

for suitable constants c_1, c_2 depending on r; the lower bound is obtained by random graphs,[2] the upper bound is calculated in Exercise 13. If H is a forest, then $H \subseteq G$ as soon as $\varepsilon(G)$ is large enough, so $\mathrm{ex}(n,H)$ is at most linear in n (Exercise 5). Erdős and Sós conjectured in 1963 that $\mathrm{ex}(n,T) \leqslant \frac{1}{2}(k-1)n$ for all trees with $k \geqslant 2$ edges; as a general bound for all n, this is best possible for every T. See Exercises 15–18 for details.

---

[2] see Chapter 11

may think of V_0 as a kind of bin: its vertices are disregarded when the uniformity of the partition is assessed, but there are only few such vertices.

Lemma 7.2.1. (Regularity Lemma)

[9.2.2] *For every $\epsilon > 0$ and every integer $m \geqslant 1$ there exists an integer M such that every graph of order at least m admits an ϵ-regular partition $\{ V_0, V_1, \dots, V_k \}$ with $m \leqslant k \leqslant M$.*

The regularity lemma thus says that, given any $\epsilon > 0$, every graph has an ϵ-regular partition into a bounded number of sets. The upper bound M on the number of partition sets ensures that for large graphs the partition sets are large too; note that ϵ-regularity is trivial when the partition sets are singletons, and a powerful property when they are large. In addition, the lemma allows us to specify a lower bound m on the number of partition sets; by choosing m large, we may increase the proportion of edges running between different partition sets (rather than inside one), i.e. the proportion of edges that are subject to the regularity assertion.

Note that the regularity lemma is designed for use with dense graphs:[4] for sparse graphs it becomes trivial, because all densities of pairs—and hence their differences—tend to zero (Exercise 22).

The remainder of this section is devoted to the proof of the regularity lemma. Although the proof is not difficult, a reader meeting the regularity lemma here for the first time is likely to draw more insight from seeing how the lemma is typically applied than from studying the technicalities of its proof. Any such reader is encouraged to skip to the start of Section 7.3 now and come back to the proof at his or her leisure.

We shall need the following inequality for reals $\mu_1, \dots, \mu_k > 0$ and $e_1, \dots, e_k \geqslant 0$:

$$\sum \frac{e_i^2}{\mu_i} \geqslant \frac{(\sum e_i)^2}{\sum \mu_i} . \tag{1}$$

This follows from the Cauchy-Schwarz inequality $\sum a_i^2 \sum b_i^2 \geqslant (\sum a_i b_i)^2$ by taking $a_i := \sqrt{\mu_i}$ and $b_i := e_i / \sqrt{\mu_i}$.

$G = (V, E)$
n

Let $G = (V, E)$ be a graph and $n := |V|$. For disjoint sets $A, B \subseteq V$ we define

$q(A, B)$

$$q(A, B) := \frac{|A|\,|B|}{n^2}\, d^2(A, B) = \frac{\|A, B\|^2}{|A|\,|B|\, n^2} .$$

For partitions $\mathcal{A}$ of A and $\mathcal{B}$ of B we set

$q(\mathcal{A}, \mathcal{B})$

$$q(\mathcal{A}, \mathcal{B}) := \sum_{A' \in \mathcal{A};\ B' \in \mathcal{B}} q(A', B') ,$$

---

[4] Sparse versions do exist, though; see the notes.

Proof. (i) Let $\mathcal{C} =: \{ C_1, \dots, C_k \}$ and $\mathcal{D} =: \{ D_1, \dots, D_\ell \}$. Then

$$\begin{aligned} q(\mathcal{C}, \mathcal{D}) &= \sum_{i,j} q(C_i, D_j) \\ &= \frac{1}{n^2} \sum_{i,j} \frac{\|C_i, D_j\|^2}{|C_i|\,|D_j|} \\ &\underset{(1)}{\geqslant} \frac{1}{n^2} \frac{\left(\sum_{i,j} \|C_i, D_j\|\right)^2}{\sum_{i,j} |C_i|\,|D_j|} \\ &= \frac{1}{n^2} \frac{\|C, D\|^2}{\left(\sum_i |C_i|\right)\left(\sum_j |D_j|\right)} \\ &= q(C, D)\,. \end{aligned}$$

(ii) Let $\mathcal{P} =: \{ C_1, \dots, C_k \}$, and for $i = 1, \dots, k$ let $\mathcal{C}_i$ be the partition of C_i induced by $\mathcal{P}'$. Then

$$\begin{aligned} q(\mathcal{P}) &= \sum_{i<j} q(C_i, C_j) \\ &\underset{\text{(i)}}{\leqslant} \sum_{i<j} q(\mathcal{C}_i, \mathcal{C}_j) \\ &\leqslant q(\mathcal{P}')\,, \end{aligned}$$

since $q(\mathcal{P}') = \sum_i q(\mathcal{C}_i) + \sum_{i<j} q(\mathcal{C}_i, \mathcal{C}_j)$. □

Next, we show that refining a partition by subpartitioning an irregular pair of partition sets increases the value of q a little; since we are dealing here with a single pair only, the amount of this increase will still be less than any constant.

Lemma 7.2.3. *Let $\epsilon > 0$, and let $C, D \subseteq V$ be disjoint. If (C, D) is not ϵ-regular, then there are partitions $\mathcal{C} = (C_1, C_2)$ of C and $\mathcal{D} = (D_1, D_2)$ of D such that*

$$q(\mathcal{C}, \mathcal{D}) \geqslant q(C, D) + \epsilon^4 \frac{|C|\,|D|}{n^2}\,.$$

Proof. Suppose (C, D) is not ϵ-regular. Then there are sets $C_1 \subseteq C$ and $D_1 \subseteq D$ with $|C_1| > \epsilon\,|C|$ and $|D_1| > \epsilon\,|D|$ such that

$$|\eta| > \epsilon \tag{2}$$

η for $\eta := d(C_1, D_1) - d(C, D)$. Let $\mathcal{C} := \{ C_1, C_2 \}$ and $\mathcal{D} := \{ D_1, D_2 \}$, where $C_2 := C \setminus C_1$ and $D_2 := D \setminus D_1$.

$\mathcal{C}_{ij}$

Proof. For all $1 \leqslant i < j \leqslant k$, let us define a partition $\mathcal{C}_{ij}$ of C_i and a partition $\mathcal{C}_{ji}$ of C_j, as follows. If the pair (C_i, C_j) is ϵ-regular, we let $\mathcal{C}_{ij} := \{ C_i \}$ and $\mathcal{C}_{ji} := \{ C_j \}$. If not, then by Lemma 7.2.3 there are partitions $\mathcal{C}_{ij}$ of C_i and $\mathcal{C}_{ji}$ of C_j with $|\mathcal{C}_{ij}| = |\mathcal{C}_{ji}| = 2$ and

$$q(\mathcal{C}_{ij}, \mathcal{C}_{ji}) \geqslant q(C_i, C_j) + \epsilon^4 \frac{|C_i|\,|C_j|}{n^2} = q(C_i, C_j) + \frac{\epsilon^4 c^2}{n^2}\,. \tag{3}$$

$\mathcal{C}_i$

For each $i = 1, \ldots, k$, let $\mathcal{C}_i$ be the unique minimal partition of C_i that refines every partition $\mathcal{C}_{ij}$ with $j \neq i$. (In other words, if we consider two elements of C_i as equivalent whenever they lie in the same partition set of $\mathcal{C}_{ij}$ for every $j \neq i$, then $\mathcal{C}_i$ is the set of equivalence classes.) Thus, $|\mathcal{C}_i| \leqslant 2^{k-1}$. Now consider the partition

$\mathcal{C}$

$$\mathcal{C} := \{ C_0 \} \cup \bigcup_{i=1}^{k} \mathcal{C}_i$$

of V, with C_0 as exceptional set. Then $\mathcal{C}$ refines $\mathcal{P}$, and

$$k \leqslant |\mathcal{C}| \leqslant k2^k. \tag{4}$$

$\mathcal{C}_0$

Let $\mathcal{C}_0 := \big\{ \{ v \} : v \in C_0 \big\}$. Now if $\mathcal{P}$ is not ϵ-regular, then for more than ϵk^2 of the pairs (C_i, C_j) with $1 \leqslant i < j \leqslant k$ the partition $\mathcal{C}_{ij}$ is non-trivial. Hence, by our definition of q for partitions with exceptional set, and Lemma 7.2.2 (i),

$$\begin{aligned}
q(\mathcal{C}) &= \sum_{1 \leqslant i < j} q(\mathcal{C}_i, \mathcal{C}_j) + \sum_{1 \leqslant i} q(\mathcal{C}_0, \mathcal{C}_i) + \sum_{0 \leqslant i} q(\mathcal{C}_i) \\
&\geqslant \sum_{1 \leqslant i < j} q(\mathcal{C}_{ij}, \mathcal{C}_{ji}) + \sum_{1 \leqslant i} q\big(\mathcal{C}_0, \{ C_i \}\big) + q(\mathcal{C}_0) \\
&\underset{(3)}{\geqslant} \sum_{1 \leqslant i < j} q(C_i, C_j) + \epsilon k^2 \frac{\epsilon^4 c^2}{n^2} + \sum_{1 \leqslant i} q\big(\mathcal{C}_0, \{ C_i \}\big) + q(\mathcal{C}_0) \\
&= q(\mathcal{P}) + \epsilon^5 \left(\frac{kc}{n} \right)^2 \\
&\geqslant q(\mathcal{P}) + \epsilon^5/2\,.
\end{aligned}$$

(For the last inequality, recall that $|C_0| \leqslant \epsilon n \leqslant \frac{1}{4}n$, so $kc \geqslant \frac{3}{4}n$.)

In order to turn $\mathcal{C}$ into our desired partition $\mathcal{P}'$, all that remains to do is to cut its sets up into pieces of some common size, small enough that all remaining vertices can be collected into the exceptional set without making this too large. Let $C'_1, \ldots, C'_\ell$ be a maximal collection of disjoint sets of size $d := \lfloor c/4^k \rfloor$ such that each C'_i is contained in some

d

of Lemma 7.2.4, where in each iteration this number may grow from its current value r to at most $r4^r$. So let f be the function $x \mapsto x4^x$, and take $M := \max\{\, f^s(k), 2k/\epsilon \,\}$; the second term in the maximum ensures that any $n \geqslant M$ is large enough to satisfy (5). M

We finally have to show that every graph $G = (V, E)$ of order at least m has an ϵ-regular partition $\{\, V_0, V_1, \dots, V_k \,\}$ with $m \leqslant k \leqslant M$. So let G be given, and let $n := |G|$. If $n \leqslant M$, we partition G into $k := n$ singletons, choosing $V_0 := \emptyset$ and $|V_1| = \ldots = |V_k| = 1$. This partition of G is clearly ϵ-regular. Suppose now that $n > M$. Let $C_0 \subseteq V$ be minimal such that k divides $|V \setminus C_0|$, and let $\{\, C_1, \dots, C_k \,\}$ be any partition of $V \setminus C_0$ into sets of equal size. Then $|C_0| < k$, and hence $|C_0| \leqslant \epsilon n$ by (5). Starting with $\{\, C_0, C_1, \dots, C_k \,\}$ we apply Lemma 7.2.4 again and again, until the partition of G obtained is ϵ-regular; this will happen after at most s iterations, since by (5) the size of the exceptional set in the partitions stays below ϵn, so the lemma could indeed be reapplied up to the theoretical maximum of s times. □ n

7.3 Applying the regularity lemma

The purpose of this section is to illustrate how the regularity lemma is typically applied in the context of (dense) extremal graph theory. Suppose we are trying to prove that a certain edge density of a graph G suffices to force the occurrence of some given subgraph H, and that we have an ϵ-regular partition of G. The edges inside almost all the pairs (V_i, V_j) of partition sets are distributed uniformly, although their density may depend on the pair. But since G has many edges, this density cannot be zero for all the pairs: some sizeable proportion of the pairs will have positive density. Now if G is large, then so are the pairs: recall that the number of partition sets is bounded, and they have equal size. But any large enough bipartite graph with equal partition sets, fixed positive edge density (however small!) and a uniform distribution of edges will contain any given bipartite subgraph[5]—this will be made precise below. Thus if enough pairs in our partition of G have positive density that H can be written as the union of bipartite graphs each arising in one of those pairs, we may hope that $H \subseteq G$ as desired.

These ideas will be formalized by Lemma 7.3.2 below. We shall then use this and the regularity lemma to prove the Erdős-Stone theorem from Section 7.1; another application will be given later, in the proof of Theorem 9.2.2.

Before we state Lemma 7.3.2, let us note a simple consequence of the ϵ-regularity of a pair (A, B): for any subset $Y \subseteq B$ that is not too

---

[5] Readers already acquainted with random graphs may find it instructive to compare this statement with Proposition 11.3.1.

u_i, h σ v_i

of R_s (not just isomorphic to one), with vertices $u_1, \dots, u_h$ say. Each vertex u_i lies in one of the s-sets V_j^s of R_s; this defines a map $\sigma\colon i \mapsto j$. Our aim is to define an embedding $u_i \mapsto v_i \in V_{\sigma(i)}$ of H in G; thus, $v_1, \dots, v_h$ will be distinct, and $v_i v_j$ will be an edge of G whenever $u_i u_j$ is an edge of H.

Our plan is to choose the vertices $v_1, \dots, v_h$ inductively. Throughout the induction, we shall have a 'target set' $Y_i \subseteq V_{\sigma(i)}$ assigned to each i; this contains the vertices that are still candidates for the choice of v_i. Initially, Y_i is the entire set $V_{\sigma(i)}$. As the embedding proceeds, Y_i will get smaller and smaller (until it collapses to $\{ v_i \}$ when v_i is chosen): whenever we choose a vertex v_j with $j < i$ and $u_j u_i \in E(H)$, we delete all those vertices from Y_i that are not adjacent to v_j. The set Y_i thus evolves as

$$V_{\sigma(i)} = Y_i^0 \supseteq \dots \supseteq Y_i^i = \{ v_i \},$$

where Y_i^j denotes the version of Y_i current after the definition of v_j (and any corresponding deletion of vertices from Y_i^{j-1}).

In order to make this approach work, we have to ensure that the target sets Y_i do not get too small. When we come to embed a vertex u_j, we consider all the indices $i > j$ with $u_j u_i \in E(H)$; there are at most Δ such i. For each of these i, we wish to select v_j so that

$$Y_i^j = N(v_j) \cap Y_i^{j-1} \tag{2}$$

is large, i.e. not much smaller than Y_i^{j-1}. Now this can be done by Lemma 7.3.1 (with $A = V_{\sigma(j)},\ B = V_{\sigma(i)}$ and $Y = Y_i^{j-1}$): unless Y_i^{j-1} is tiny (of size less than $\epsilon\ell$), all but at most $\epsilon\ell$ choices of v_j will be such that (2) implies

$$|Y_i^j| \geqslant (d-\epsilon)|Y_i^{j-1}|\,. \tag{3}$$

Doing this simultaneously for all of the at most Δ values of i considered, we find that all but at most $\Delta\epsilon\ell$ choices of v_j from $V_{\sigma(j)}$, and in particular from $Y_j^{j-1} \subseteq V_{\sigma(j)}$, satisfy (3) for all i.

It remains to show that the sets Y considered for Lemma 7.3.1 above are indeed never tiny, and that $|Y_j^{j-1}| - \Delta\epsilon\ell \geqslant s$ to ensure that a suitable choice for v_j exists: since $\sigma(j') = \sigma(j)$ for at most $s-1$ of the vertices $u_{j'}$ with $j' < j$, a choice between s suitable candidates for v_j will suffice to keep v_j distinct from $v_1, \dots, v_{j-1}$. But all this follows from our choice of ϵ_0. Indeed, the initial target sets Y_i^0 have size ℓ, and each Y_i has vertices deleted from it only when some v_j with $j < i$ and $u_j u_i \in E(H)$ is defined, which happens at most Δ times. Thus,

$$|Y_i^j| - \Delta\epsilon\ell \underset{(3)}{\geqslant} (d-\epsilon)^\Delta \ell - \Delta\epsilon\ell \geqslant (d-\epsilon_0)^\Delta \ell - \Delta\epsilon_0\ell \underset{(1)}{\geqslant} \epsilon_0\ell \geqslant s$$

whenever $j < i$, so in particular $|Y_i^j| \geqslant \epsilon_0\ell \geqslant \epsilon\ell$ and $|Y_j^{j-1}| - \Delta\epsilon\ell \geqslant s$. □

that $K^r_s \subseteq G$, all that remains to be checked is that $K^r \subseteq R$ (and hence $K^r_s \subseteq R_s$).

Our plan was to show $K^r \subseteq R$ by Turán's theorem. We thus have to check that R has enough edges, i.e. that enough ϵ-regular pairs (V_i, V_j) have density at least d. This should follow from our assumption that G has at least $t_{r-1}(n) + \gamma n^2$ edges, i.e. an edge density of about $\frac{r-2}{r-1} + 2\gamma$: this lies substantially above the approximate edge density $\frac{r-2}{r-1}$ of the Turán graph $T^{r-1}(k)$, and hence substantially above any density that G could have if no more than $t_{r-1}(k)$ of the pairs (V_i, V_j) had density $\geqslant d$—even if all those pairs had density 1!

Let us then estimate $\|R\|$ more precisely. How many edges of G lie outside ϵ-regular pairs? At most $\binom{|V_0|}{2}$ edges lie inside V_0, and by condition (i) in the definition of ϵ-regularity these are at most $\frac{1}{2}(\epsilon n)^2$ edges. At most $|V_0|k\ell \leqslant \epsilon nk\ell$ edges join V_0 to other partition sets. The at most ϵk^2 other pairs (V_i, V_j) that are not ϵ-regular contain at most ℓ^2 edges each, together at most $\epsilon k^2 \ell^2$. The ϵ-regular pairs of insufficient density ($< d$) each contain no more than $d\ell^2$ edges, altogether at most $\frac{1}{2}k^2 d\ell^2$ edges. Finally, there are at most $\binom{\ell}{2}$ edges inside each of the partition sets $V_1, \dots, V_k$, together at most $\frac{1}{2}\ell^2 k$ edges. All *other* edges of G lie in ϵ-regular pairs of density at least d, and thus contribute to edges of R. Since each edge of R corresponds to at most ℓ^2 edges of G, we thus have in total

$$\|G\| \leq \tfrac{1}{2}\epsilon^2 n^2 + \epsilon nk\ell + \epsilon k^2 \ell^2 + \tfrac{1}{2}k^2 d\ell^2 + \tfrac{1}{2}\ell^2 k + \|R\|\,\ell^2 .$$

Hence, for all sufficiently large n,

$$\begin{aligned}
\|R\| &\geq \tfrac{1}{2}k^2 \, \frac{\|G\| - \frac{1}{2}\epsilon^2 n^2 - \epsilon nk\ell - \epsilon k^2\ell^2 - \frac{1}{2}dk^2\ell^2 - \frac{1}{2}k\ell^2}{\frac{1}{2}k^2\ell^2} \\
&\underset{(1,2)}{\geq} \tfrac{1}{2}k^2 \left(\frac{t_{r-1}(n) + \gamma n^2 - \frac{1}{2}\epsilon^2 n^2 - \epsilon nk\ell}{n^2/2} - 2\epsilon - d - \frac{1}{k} \right) \\
&\underset{(2)}{\geq} \tfrac{1}{2}k^2 \left(\frac{t_{r-1}(n)}{n^2/2} + 2\gamma - \epsilon^2 - 4\epsilon - d - \frac{1}{m} \right) \\
&= \tfrac{1}{2}k^2 \left(t_{r-1}(n) \binom{n}{2}^{-1} \left(1 - \frac{1}{n}\right) + \delta \right) \\
&> \tfrac{1}{2}k^2 \, \frac{r-2}{r-1} \\
&\geqslant t_{r-1}(k)\,.
\end{aligned}$$

(The strict inequality follows from Lemma 7.1.4.) Therefore $K^r \subseteq R$ by Theorem 7.1.1, as desired. □

14. The *upper density* of an infinite graph G is the infimum of all reals α such that the finite graphs $H \subseteq G$ with $\|H\| \binom{|H|}{2}^{-1} > \alpha$ have bounded order. Show that this number always takes one of the countably many values $0, 1, \frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \ldots$.

    (Hint. Erdős-Stone.)

15. Prove the following weakening of the Erdős-Sós conjecture (stated at the end of Section 7.1): given integers $2 \leqslant k < n$, every graph with n vertices and at least $(k-1)n$ edges contains every tree with k edges as a subgraph.

16. Show that, as a general bound for arbitrary n, the bound on $\text{ex}(n,T)$ claimed by the Erdős-Sós conjecture is best possible for every tree T. Is it best possible even for every n and every T?

17.$^-$ Prove the Erdős-Sós conjecture for the case when the tree considered is a star.

18. Prove the Erdős-Sós conjecture for the case when the tree considered is a path.

    (Hint. Use the result of the next exercise.)

19. Show that every connected graph G contains a path of length at least $\min\{\, 2\delta(G), |G|-1 \,\}$.

20.$^-$ In the definition of an ϵ-regular pair, what is the purpose of the requirement that $|X| > \epsilon\,|A|$ and $|Y| > \epsilon\,|B|$?

21.$^-$ Show that any ϵ-regular pair in G is also ϵ-regular in $\overline{G}$.

22. Prove the regularity lemma for sparse graphs, that is, for every sequence $(G_n)_{n \in \mathbb{N}}$ of graphs G_n of order n such that $\|G_n\|/n^2 \to 0$ as $n \to \infty$.

Notes

The standard reference work for results and open problems in extremal graph theory (in a very broad sense) is still B. Bollobás, *Extremal Graph Theory*, Academic Press 1978. A kind of update on the book is given by its author in his chapter of the *Handbook of Combinatorics* (R.L. Graham, M. Grötschel & L. Lovász, eds.), North-Holland 1995. An instructive survey of extremal graph theory in the narrower sense of our chapter is given by M. Simonovits in (L.W. Beineke & R.J. Wilson, eds.) *Selected Topics in Graph Theory 2*, Academic Press 1983. This paper focuses among other things on the particular role played by the Turán graphs. A more recent survey by the same author can be found in (R.L. Graham & J. Nešetřil, eds.) *The Mathematics of Paul Erdős*, Vol. 2, Springer 1996.

Turán's theorem is not merely one extremal result among others: it is the result that sparked off the entire line of research. Our proof of Turán's theorem is essentially the original one; the proof indicated in Exercise 10 is due to Zykov.

suggests that there is indeed a gap here: while a value of $c = c'r\sqrt{\log r}$ (where c' is independent of both n and r) is best possible for $d(G) \geqslant c$ to imply $H \preccurlyeq G$ (Section 8.2), the conjecture says that $\chi(G) \geqslant r$ will do the same! Thus, if true, then Hadwiger's conjecture shows that the effect of a large chromatic number on the occurrence of minors somehow goes beyond that part which is well-understood: its effect via mere edge density. We shall consider Hadwiger's conjecture in Section 8.3.

8.1 Topological minors

In this section we prove that an average degree of cr^2 suffices to force the occurrence of a topological K^r minor in a graph; complete bipartite graphs show that, up to the constant c, this is best possible (Exercise 5).

The following theorem was proved independently around 1996 by Bollobás & Thomason and by Komlós & Szemerédi.

Theorem 8.1.1. *There exists a $c \in \mathbb{R}$ such that, for every $r \in \mathbb{N}$, every graph G of average degree $d(G) \geqslant cr^2$ contains K^r as a topological minor.*

The proof of this theorem, in which we follow Bollobás & Thomason, will occupy us for the remainder of this section. A set $U \subseteq V(G)$ will be called *linked* (in G) if for any distinct vertices $u_1, \dots, u_{2h} \in U$ there are h disjoint paths $P_i = u_{2i-1} \dots u_{2i}$ in G, $i = 1, \dots, h$.[2] The graph G itself is (k, ℓ)-*linked* if every k-set of its vertices contains a linked ℓ-set.

linked

(k,ℓ)-*linked*

How can we hope to find the TK^r in G claimed to exist by Theorem 8.1.1? Our basic approach will be to identify first some r-set X as a set of branch vertices, and to choose for each $x \in X$ a set Y_x of $r-1$ neighbours, one for every edge incident with x in the K^r. If the constant c from the theorem is large enough, the $r + r(r-1) = r^2$ vertices of $X \cup \bigcup Y_x$ can be chosen distinct: by Proposition 1.2.2, G has a subgraph of minimum degree at least $\varepsilon(G) = \frac{1}{2}d(G) \geqslant \frac{1}{2}cr^2$, so we can choose X and its neighbours inside this subgraph. Having fixed X and the sets Y_x, we then have to link up the correct pairs of vertices in $Y := \bigcup Y_x$ by disjoint paths in $G - X$, to obtain the desired TK^r.

This would be possible at once if Y were linked in $G - X$. Unfortunately, this is unrealistic to hope for: no average degree, however large, will force *every* $r(r-1)$-set to be linked. (Why not?) However, if we pick for X significantly more than the r vertices needed eventually, and for each $x \in X$ significantly more than $r-1$ neighbours as Y_x, then Y might become so large that the high average degree of G guarantees the

---

[2] Thus, in a k-linked graph—see Chapter 3.6—every set of up to $2k+1$ vertices is linked.

(3.3.1)
$\mathcal{V}, V_x$
$v_1, \dots, v_k$
linkage
link

Proof. Let $\mathcal{V} := \{\, V_x \mid x \in V(H) \,\}$ be the set of branch sets in G corresponding to the vertices of H. For our proof that G is $(k, \lceil k/2 \rceil)$-linked, let k distinct vertices $v_1, \dots, v_k \in G$ be given. Let us call a sequence $P_1, \dots, P_k$ of disjoint paths in G a *linkage* if the P_i each start in v_i and end in pairwise distinct sets $V \in \mathcal{V}$; the paths P_i themselves will be called *links*. Since our assumptions about H imply that $|H| \geqslant k$, and G is k-connected, such linkages exist: just pick k vertices from pairwise distinct sets $V \in \mathcal{V}$, and link them disjointly to $\{\, v_1, \dots, v_k \,\}$ by Menger's theorem.

$\mathcal{P}$
$P_1, \dots, P_k$
$f(P)$

Now let $\mathcal{P} = (P_1, \dots, P_k)$ be a linkage whose total number of edges outside $\bigcup_{V \in \mathcal{V}} G[V]$ is as small as possible. Thus, if $f(P)$ denotes the number of edges of P not lying in any $G[V_x]$, we choose $\mathcal{P}$ so as to minimize $\sum_{i=1}^{k} f(P_i)$. Then for every $V \in \mathcal{V}$ that meets a path $P_i \in \mathcal{P}$ there exists one such path that ends in V: if not, we could terminate P_i in V and reduce $f(P_i)$. Thus, exactly k of the branch sets of H meet a link. Let us divide these sets into two classes:

$\mathcal{U}$
$\mathcal{W}$

$$\mathcal{U} := \{\, V \in \mathcal{V} \mid V \text{ meets exactly one link} \,\}$$
$$\mathcal{W} := \{\, V \in \mathcal{V} \mid V \text{ meets more than one link} \,\}.$$

Since H is dense and each $U \in \mathcal{U}$ meets only one link, it will be easy to show that the starting vertices v_i of those links form a linked set in G. Hence, our aim is to show that $|\mathcal{U}| \geqslant \lceil k/2 \rceil$, i.e. that $\mathcal{U}$ is no smaller than $\mathcal{W}$. (Recall that $|\mathcal{U}| + |\mathcal{W}| = k$.) To this end, we first prove the following:

Every $V \in \mathcal{W}$ is met by some link which leaves V again and next meets a set from $\mathcal{U}$ (where it ends). (1)

x

Suppose $V_x \in \mathcal{W}$ is a counterexample to (1). Since

$$2\delta(H) \geqslant |H| + \tfrac{3}{2}k \geqslant \delta(H) + \tfrac{3}{2}k\,,$$

P_i
P_i'

we have $\delta(H) \geqslant \frac{3}{2}k$. As $|\mathcal{U} \cup \mathcal{W}| = k$, this implies that x has a neighbour y in H with $V_y \in \mathcal{V} \smallsetminus (\mathcal{U} \cup \mathcal{W})$; let $w_x w_y$ be an edge of G with $w_x \in V_x$ and $w_y \in V_y$. Let $Q = w \dots w_x w_y$ be a path in $G[V_x \cup \{\, w_y \,\}]$ of whose vertices only w lies on any link, say on P_i (Fig. 8.1.2). Replacing P_i in $\mathcal{P}$ by $P_i' := P_i w Q$ then yields another linkage.

If P_i is not the link ending in V_x, then $f(P_i') \leqslant f(P_i)$. The choice of $\mathcal{P}$ then implies that $f(P_i') = f(P_i)$, i.e. that P_i ends in the branch set W it enters immediately after V_x. Since V_x is a counterexample to (1) we have $W \notin \mathcal{U}$, i.e. $W \in \mathcal{W}$. Let $P \neq P_i$ be another link meeting W. Then P does not end in W (because P_i ends there); let $P' \subseteq P$ be the (minimal) initial segment of P that ends in W. If we now replace P_i and P by P_i' and P' in $\mathcal{P}$, we obtain a linkage contradicting the choice of $\mathcal{P}$.

$|H| + \frac{3}{2}k$, we can find for any two sets $V_x, V_y \in \mathcal{U}$ at least $\frac{3}{2}k$ sets V_z such that $xz, yz \in E(H)$. At least $k/2$ of these sets V_z do not lie in $\mathcal{U} \cup \mathcal{W}$. Thus whenever $U_1, \dots, U_{2h}$ are distinct sets in $\mathcal{U}$ (so $h \leqslant u/2 \leqslant k/2$), we may find inductively h distinct sets $V^i \in \mathcal{V} \setminus (\mathcal{U} \cup \mathcal{W})$ $(i = 1, \dots, h)$ such that V^i is joined in G to both U_{2i-1} and U_{2i}. For each i, any vertex of U_{2i-1} can be linked by a path through V^i to any desired vertex of U_{2i}, and these paths will be disjoint for different i. Joining up the appropriate pairs of paths from $\mathcal{P}$ in this way, we see that the set $\{ v_1, \dots, v_u \}$ is linked in G, and the lemma is proved. □

Lemma 8.1.3. *Let $k \geqslant 6$ be an integer. Then every graph G with $\varepsilon(G) \geqslant k$ has a minor H such that $2\delta(H) \geqslant |H| + \frac{1}{6}k$.*

G_0

Proof. We begin by choosing a ($\preccurlyeq$-)minimal minor G_0 of G with $\varepsilon(G_0) \geqslant k$. The minimality of G_0 implies that $\delta(G_0) > k$ and $\varepsilon(G_0) = k$ (otherwise we could delete a vertex or an edge of G_0), and hence

$$k + 1 \leqslant \delta(G_0) \leqslant d(G_0) = 2k\,.$$

x_0

Let $x_0 \in G_0$ be a vertex of minimum degree.

If k is odd, let $m := (k+1)/2$ and

$$G_1 := G_0\,[\,\{ x_0 \} \cup N_{G_0}(x_0)\,]\,.$$

Then $|G_1| = \delta(G_0) + 1 \leqslant 2k + 1 \leqslant 2(k+1) = 4m$. By the minimality of G_0, contracting any edge x_0y of G_0 will result in the loss of at least $k+1$ edges. The vertices x_0 and y thus have at least k common neighbours, so $\delta(G_1) \geqslant k + 1 = 2m$ (Fig. 8.1.4).

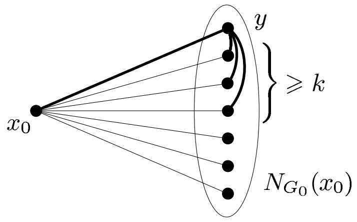

Fig. 8.1.4. The graph $G_1 \preccurlyeq G$: a first approximation to the desired minor H

If k is even, we let $m := k/2$ and

$$G_1 := G_0\,[\,N_{G_0}(x_0)\,]\,.$$

Then $|G_1| = \delta(G_0) \leqslant 2k = 4m$, and $\delta(G_1) \geqslant k = 2m$ as before.

For every $H \in \mathcal{H}$, any graph obtained from H by one of the following three operations will again be in $\mathcal{H}$:

(i) deletion of an edge, if $\|H\| \geqslant m\,|H| + f(|H|) - \binom{m}{2} + 1$;

(ii) deletion of a vertex of degree at most $\frac{7}{6}m$;

(iii) contraction of an edge $xy \in H$ such that x and y have at most $\frac{7}{6}m - 1$ common neighbours in H.

H_0 Starting with G_1, let us apply these operations as often as possible, and let $H_0 \in \mathcal{H}$ be the graph obtained eventually. Since

$$\|K^m\| = m\,|K^m| - m - \binom{m}{2}$$

and

$$f(m) = -\tfrac{5}{6}m > -m\,,$$

x_1 K^m does not have enough edges to be in $\mathcal{H}$; thus, $\mathcal{H}$ contains no graph on m vertices. Hence $|H_0| > m$, and in particular $H_0 \neq \emptyset$. Let $x_1 \in H_0$ be a vertex of minimum degree, and put

H_1
$$H_1 := H_0\,[\,\{\,x_1\,\} \cup N_{H_0}(x_1)\,]\,.$$

We shall prove that the minimum degree of $H := H_1$ is as large as claimed in the lemma.

Note first that

$$\delta(H_1) > \tfrac{7}{6}m\,. \tag{3}$$

Indeed, since H_0 is minimal with respect to (ii) and (iii), we have $d(x_1) > \frac{7}{6}m$ in H_0 (and hence in H_1), and every vertex $y \neq x_1$ of H_1 has more than $\frac{7}{6}m - 1$ common neighbours with x_1 (and hence more than $\frac{7}{6}m$ neighbours in H_1 altogether). In order to convert (3) into the desired inequality of the form

$$2\delta(H_1) \geqslant |H_1| + \alpha m\,,$$

we need an upper bound for $|H_1|$ in terms of m. Since H_0 lies in $\mathcal{H}$ but is minimal with respect to (i), we have

$$\begin{aligned}\|H_0\| &< m\,|H_0| + \left(\tfrac{1}{6}m\,|H_0| - \tfrac{1}{6}m^2 - \tfrac{5}{6}m\right) - \binom{m}{2} + 1\\ &= \tfrac{7}{6}m\,|H_0| - \tfrac{4}{6}m^2 - \tfrac{1}{3}m + 1\\ &\underset{(2)}{\leqslant} \tfrac{7}{6}m\,|H_0| - \tfrac{4}{6}m^2.\end{aligned} \tag{4}$$

Although Theorem 8.1.1 already gives a good estimate, it seems very difficult to determine the exact average degree needed to force a TK^r subgraph, even for small r. We shall come back to the case of $r = 5$ in Section 8.3; more results and conjectures are given in the notes.

The following almost counter-intuitive result of Mader implies that the existence of a topological K^r minor can be forced essentially by large girth. In the next section, we shall prove the analogue of this for ordinary minors.

Theorem 8.1.4. (Mader 1997)
For every graph H of maximum degree $d \geqslant 3$ there exists an integer k such that every graph G of minimum degree at least d and girth at least k contains H as a topological minor.

As discussed already in Chapter 5.2 and the introduction to Chapter 7, no constant average degree, however large, will force an arbitrary graph to contain a given graph H as a subgraph—as long as H contains at least one cycle. By Proposition 1.2.2 and Corollary 1.5.4, on the other hand, any graph G contains all trees on up to $\varepsilon(G) + 2$ vertices. Large average degree therefore does ensure the occurrence of any fixed tree T as a subgraph. What can we say, however, if we would like T to occur as an *induced* subgraph?

Here, a large average degree appears to do as much harm as good, even for graphs of bounded clique number. (Consider, for example, complete bipartite graphs.) It is all the more remarkable, then, that the assumption of a large chromatic number rather than a large average degree seems to make a difference here: according to a conjecture of Gyárfás, any graph of large enough chromatic number contains either a large complete graph or any given tree as an induced subgraph. (Formally: for every integer r and every tree T, there exists an integer k such that every graph G with $\chi(G) \geqslant k$ and $\omega(G) < r$ contains an induced copy of T.)

The weaker topological version of this is indeed true:

Theorem 8.1.5. (Scott 1997)
For every integer r and every tree T there exists an integer k such that every graph with $\chi(G) \geqslant k$ and $\omega(G) < r$ contains an induced copy of some subdivision of T.

at least $|V_i| + 2 \geqslant 2k$ edges joining V_i to $V \setminus V_i$. We shall prove that every V_i sends at most two edges to each of the other V_j; then V_i must send edges to at least k of those V_j, so the V_i are the branch sets of an $MH \subseteq G$ with $\delta(H) \geqslant k$.

Suppose, without loss of generality, that G has three V_1–V_2 edges. Then there are vertices $v_1 \in V_1$ and $v_2 \in V_2$ such that $G[V_1 \cup V_2]$ contains three independent v_1–v_2 paths P_1, P_2, P_3 (Fig. 8.2.1). At most one of

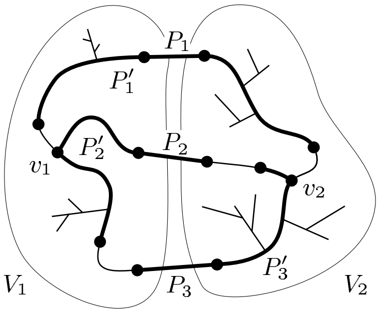

Fig. 8.2.1. Three edges between V_1 and V_2

these paths can be shorter than $\frac{1}{2}g(G)$. We assume that P_1 has length at least $\lceil \frac{1}{2}g(G) \rceil \geqslant 2k-1$ and let $P_1' := \mathring{P}_1$; then $|P_1'| \geqslant 2k-2$. Since $P_2 \cup P_3$ is a cycle of length at least $4k-3$, we can further find disjoint paths $P_2', P_3' \subseteq P_2 \cup P_3$ with $2k-2$ vertices each. Since $G[V_1 \cup V_2]$ is connected, there exists a partition of $V_1 \cup V_2$ into three connected sets V_1', V_2', V_3' such that $V(P_i') \subseteq V_i'$ for $i = 1, 2, 3$. Replacing the two sets V_1, V_2 in our partition of V with the three sets V_1', V_2', V_3', we obtain a partition of V that contradicts the maximality of m. □

The following combination of Theorems 8.2.1 and 8.2.2 brings out the paradoxical character of the latter particularly well:

Corollary 8.2.3. *There exists a $c \in \mathbb{R}$ such that, for every $r \in \mathbb{N}$, every graph G with girth $g(G) \geqslant c\,r\sqrt{\log r}$ and $\delta(G) \geqslant 3$ has a K^r minor.*

Proof. We prove the corollary for $c := 4c'$, where c' is the constant from Theorem 8.2.1. Let G be given as stated. By Theorem 8.2.2, G has a minor H with $\delta(H) \geqslant c' r\sqrt{\log r}$. By Theorem 8.2.1, H (and hence G) has a K^r minor. □

is also a simple structural characterization of the graphs without a K^r minor:

[12.4.2] **Proposition 8.3.1.** *A graph with at least three vertices is edge-maximal without a K^4 minor if and only if it can be constructed recursively from triangles by pasting*[4] *along K^2s.*

(1.7.2)
(4.4.4)

Proof. Recall first that every MK^4 contains a TK^4, because $\Delta(K^4) = 3$ (Proposition 1.7.2); the graphs without a K^4 minor thus coincide with those without a topological K^4 minor. The proof that any graph constructible as described is edge-maximal without a K^4 minor is left as an easy exercise; in order to deduce Hadwiger's conjecture for $r = 4$, we only need the converse implication anyhow. We prove this by induction on $|G|$.

Let G be given, edge-maximal without a K^4 minor. If $|G| = 3$ then G is itself a triangle, so let $|G| \geqslant 4$ for the induction step. Then G is not complete; let $S \subseteq V(G)$ be a separating set with $|S| = \kappa(G)$, and let C_1, C_2 be distinct components of $G - S$. Since S is a minimal separator, every vertex in S has a neighbour in C_1 and another in C_2. If $|S| \geqslant 3$, this implies that G contains three independent paths P_1, P_2, P_3 between a vertex $v_1 \in C_1$ and a vertex $v_2 \in C_2$. Since $\kappa(G) = |S| \geqslant 3$, the graph $G - \{ v_1, v_2 \}$ is connected and contains a (shortest) path P between two different P_i. Then $P \cup P_1 \cup P_2 \cup P_3 = TK^4$, a contradiction.

Hence $\kappa(G) \leqslant 2$, and the assertion follows from Lemma 4.4.4[5] and the induction hypothesis. □

One of the interesting consequences of Proposition 8.3.1 is that all the edge-maximal graphs without a K^4 minor have the same number of edges, and are thus all 'extremal':

Corollary 8.3.2. *Every edge-maximal graph G without a K^4 minor has $2\,|G| - 3$ edges.*

Proof. Induction on $|G|$. □

Corollary 8.3.3. *Hadwiger's conjecture holds for $r = 4$.*

Proof. If G arises from G_1 and G_2 by pasting along a complete graph, then $\chi(G) = \max\{ \chi(G_1), \chi(G_2) \}$ (see the proof of Proposition 5.5.2). Hence, Proposition 8.3.1 implies by induction on $|G|$ that all edge-maximal (and hence all) graphs without a K^4 minor can be 3-coloured. □

---

[4] This was defined formally in Chapter 5.5.

[5] The proof of this lemma is elementary and can be read independently of the rest of Chapter 4.

By Corollary 8.3.5, any graph with n vertices and more than $3n-6$ edges contains an MK^5. In fact, it even contains a TK^5. This inconspicuous improvement is another deep result that had been conjectured for over 30 years:

Theorem 8.3.8. (Mader 1998)
Every graph with n vertices and more than $3n-6$ edges contains K^5 as a topological minor.

No structure theorem for the graphs without a TK^5, analogous to Proposition 8.3.1 and Theorem 8.3.4, is known. However, Mader has characterized those with the greatest possible number of edges:

Theorem 8.3.9. (Mader 1997)
A graph is extremal without a TK^5 if and only if it can be constructed recursively from maximal planar graphs by pasting along triangles.

Exercises

1. Prove, from first principles, the theorem of Wagner (1964) that every graph of chromatic number at least 2^r contains K^r as a minor.

   (Hint. Apply induction on r.)

2. Prove, from first principles, the result of Mader (1967) that every graph of average degree at least 2^{r-2} contains K^r as a minor.

   (Hint. Induction on r.)

3.$^-$ Derive Wagner's theorem (Ex. 1) from Mader's theorem (Ex. 2).

4.$^+$ Given an integer $r > 0$, find an integer k such that every grid with k additional edges has a K^r minor, provided that all the ends of the new edges have distance at least k in the grid both from each other and from the grid boundary. (*Grids* are defined in Chapter 12.3.)

5.$^+$ Show that any function h as in Theorem 3.6.1 satisfies the inequality $h(r) > \frac{1}{8}r^2$ for all even r, and hence that Theorem 8.1.1 is best possible up to the value of the constant c.

6. Prove the statement of Lemma 8.1.3 for $k < 6$.

7. Explain how exactly the term of $\frac{1}{6}k$ in the statement of Lemma 8.1.3 is used in the proof of Theorem 8.1.1. Could it be replaced by $k/1000$, or by zero?

8. Explain how exactly the number $\frac{7}{6}$ in the proof of Lemma 8.1.3 was arrived at. Could it be replaced by $\frac{3}{2}$?

Notes

The investigation of graphs not containing a given graph as a minor, or topological minor, has a long history. It probably started with Wagner's 1935 PhD thesis, in which he sought to 'detopologize' the four colour problem by classifying the graphs without a K^5 minor. His hope was to be able to show abstractly that all those graphs were 4-colourable; since the graphs without a K^5 minor include the planar graphs, this would amount to a proof of the four colour conjecture involving no topology whatsoever. The result of Wagner's efforts, Theorem 8.3.4, falls tantalizingly short of this goal: although it succeeds in classifying the graphs without a K^5 minor in structural terms, planarity re-emerges as one of the criteria used in the classification. From this point of view, it is instructive to compare Wagner's K^5 theorem with similar classification theorems, such as his analogue for K^4 (Proposition 8.3.1), where the graphs are decomposed into parts from a *finite* set of irreducible graphs. See R. Diestel, *Graph Decompositions*, Oxford University Press 1990, for more such classification theorems.

Despite its failure to resolve the four colour problem, Wagner's K^5 structure theorem had consequences for the development of graph theory like few others. To mention just two: it prompted Hadwiger to make his famous conjecture; and it inspired the notion of a tree-decomposition, which is fundamental to the work of Robertson and Seymour on minors (see Chapter 12). Wagner himself responded to Hadwiger's conjecture with a proof that, in order to force a K^r minor, it does suffice to raise the chromatic number of a graph to *some* value depending only on r (Exercise 1). This theorem then, along with its analogue for topological minors proved independently by Dirac and by Jung, prompted the question of which average degree suffices to force the desired minor.

The deepest contribution in this field of research was no doubt made by Mader, in a series of papers from the late sixties. Our proof of Lemma 8.1.3 is presented intentionally in a step-by-step fashion, to bring out some of Mader's ideas. Mader's own proof—not to mention that of Thomason's best possible version of the lemma, as used in the original proof of Theorem 8.1.1—is wrapped up so elegantly that it becomes hard to see the ideas behind it. Except for this lemma, our proof of Theorem 8.1.1 follows B. Bollobás & A.G. Thomason, Proof of a conjecture of Mader, Erdős and Hajnal on topological complete subgraphs, *Europ. J. Combinatorics* **19** (1998), 883–887. The constant c from the theorem was shown by J. Komlós & E. Szemerédi, Topological cliques in graphs II, *Combinatorics, Probability and Computing* **5** (1996), 79–90, to be no greater than about $\frac{1}{2}$, which is not far from the lower bound of $\frac{1}{8}$ given in Exercise 5.

Theorem 8.1.4 is from W. Mader, Topological subgraphs in graphs of large girth, *Combinatorica* **18** (1998), 405–412. For $H = K^r$, the theorem says that every graph G with $\delta(G) \geqslant r-1$ and $g(G)$ large contains a TK^r. For $r = 5$, Mader conjectured that $g(G) \geqslant 5$ should be enough, and that the requirement of $\delta(G) \geqslant 4$ could be weakened further: he conjectured that any graph of girth at least 5, large enough order n, and $2n-4$ or more edges has a topological K^5 minor. (To see that this implies the minimum degree version of the conjecture even for small order, consider enough disjoint copies of the given graph.) For

9.1 Ramsey's original theorems

In its simplest version, Ramsey's theorem says that, given an integer $r \geqslant 0$, every large enough graph G contains either K^r or $\overline{K^r}$ as an induced subgraph. At first glance, this may seem surprising: after all, we need about $(r-2)/(r-1)$ of all possible edges to force a K^r subgraph in G (Cor. 7.1.3), but neither G nor $\overline{G}$ can be expected to have more than half the total number of edges. However, as the Turán graphs illustrate well, squeezing many edges into G without creating a K^r imposes additional structure on G, which may help us find an induced $\overline{K^r}$.

So how could we go about proving Ramsey's theorem? Let us try to build a K^r or $\overline{K^r}$ in G inductively, starting with an arbitrary vertex $v_1 \in V_1 := V(G)$. If $|G|$ is large, there will be a large set $V_2 \subseteq V_1 \smallsetminus \{ v_1 \}$ of vertices that are either all adjacent to v_1 or all non-adjacent to v_1. Accordingly, we may think of v_1 as the first vertex of a K^r or $\overline{K^r}$ whose other vertices all lie in V_2. Let us then choose another vertex $v_2 \in V_2$ for our K^r or $\overline{K^r}$. Since V_2 is large, it will have a subset V_3, still fairly large, of vertices that are all 'of the same type' with respect to v_2 as well: either all adjacent or all non-adjacent to it. We then continue our search for vertices inside V_3, and so on (Fig. 9.1.1).

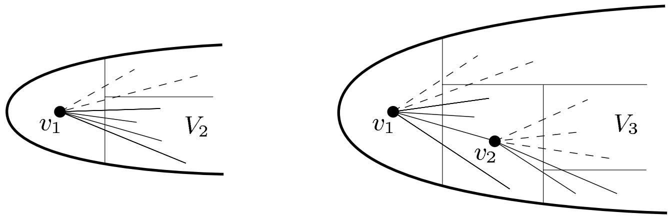

Fig. 9.1.1. Choosing the sequence $v_1, v_2, \ldots$

How long can we go on in this way? This depends on the size of our initial set V_1: each set V_i has at least half the size of its predecessor V_{i-1}, so we shall be able to complete s construction steps if G has order about 2^s. As the following proof shows, the choice of $s = 2r-3$ vertices v_i suffices in order to find among them the vertices of a K^r or $\overline{K^r}$.

Theorem 9.1.1. (Ramsey 1930)
[9.2.2] *For every $r \in \mathbb{N}$ there exists an $n \in \mathbb{N}$ such that every graph of order at least n contains either K^r or $\overline{K^r}$ as an induced subgraph.*

Proof. The assertion is trivial for $r \leqslant 1$; we assume that $r \geqslant 2$. Let $n := 2^{2r-3}$, and let G be a graph of order at least n. We shall define a sequence $V_1, \ldots, V_{2r-2}$ of sets and choose vertices $v_i \in V_i$ with the following properties:

(i) $|V_i| = 2^{2r-2-i}$ $(i = 1, \ldots, 2r-2)$;

[12.1.1] **Theorem 9.1.2.** *Let k, c be positive integers, and X an infinite set. If $[X]^k$ is coloured with c colours, then X has an infinite monochromatic subset.*

Proof. We prove the theorem by induction on k, with c fixed. For $k = 1$ the assertion holds, so let $k > 1$ and assume the assertion for smaller values of k.

Let $[X]^k$ be coloured with c colours. We shall construct an infinite sequence $X_0, X_1, \dots$ of infinite subsets of X and choose elements $x_i \in X_i$ with the following properties (for all i):

(i) $X_{i+1} \subseteq X_i \smallsetminus \{ x_i \}$;

(ii) all k-sets $\{ x_i \} \cup Z$ with $Z \in [X_{i+1}]^{k-1}$ have the same colour, which we *associate* with x_i.

We start with $X_0 := X$ and pick $x_0 \in X_0$ arbitrarily. By assumption, X_0 is infinite. Having chosen an infinite set X_i and $x_i \in X_i$ for some i, we c-colour $[X_i \smallsetminus \{ x_i \}]^{k-1}$ by giving each set Z the colour of $\{ x_i \} \cup Z$ from our c-colouring of $[X]^k$. By the induction hypothesis, $X_i \smallsetminus \{ x_i \}$ has an infinite monochromatic subset, which we choose as X_{i+1}. Clearly, this choice satisfies (i) and (ii). Finally, we pick $x_{i+1} \in X_{i+1}$ arbitrarily.

Since c is finite, one of the c colours is associated with infinitely many x_i. These x_i form an infinite monochromatic subset of X. □

To deduce the finite version of Theorem 9.1.2, we make use of a standard graph-theoretical tool in combinatorics:

Lemma 9.1.3. (König's Infinity Lemma)
Let $V_0, V_1, \dots$ be an infinite sequence of disjoint non-empty finite sets, and let G be a graph on their union. Assume that every vertex v in a set V_n with $n \geqslant 1$ has a neighbour $f(v)$ in V_{n-1}. Then G contains an infinite path $v_0 v_1 \dots$ with $v_n \in V_n$ for all n.

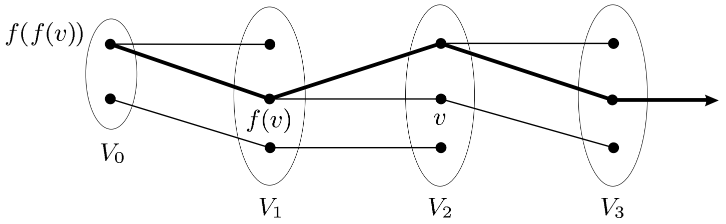

Fig. 9.1.2. König's infinity lemma

Proof. Let $\mathcal{P}$ be the set of all paths of the form $v\, f(v)\, f(f(v)) \dots$ ending in V_0. Since V_0 is finite but $\mathcal{P}$ is infinite, infinitely many of the paths in $\mathcal{P}$ end at the same vertex $v_0 \in V_0$. Of these paths, infinitely many also agree on their penultimate vertex $v_1 \in V_1$, because V_1 is finite. Of those

H_1, H_2, only very rough estimates are known for $R(H_1, H_2)$. Interestingly, lower bounds given by random graphs (as in Theorem 11.1.3) are often sharper than even the best bounds provided by explicit constructions.

The following proposition describes one of the few cases where exact Ramsey numbers are known for a relatively large class of graphs:

Proposition 9.2.1. *Let s, t be positive integers, and let T be a tree of order t. Then $R(T, K^s) = (s-1)(t-1)+1$.*

(5.2.3)
(1.5.4)

Proof. The disjoint union of $s-1$ graphs K^{t-1} contains no copy of T, while the complement of this graph, the complete $(s-1)$-partite graph K^{s-1}_{t-1}, does not contain K^s. This proves $R(T, K^s) \geqslant (s-1)(t-1)+1$.

Conversely, let G be any graph of order $n = (s-1)(t-1)+1$ whose complement contains no K^s. Then $s > 1$, and in any vertex colouring of G (in the sense of Chapter 5) at most $s-1$ vertices can have the same colour. Hence, $\chi(G) \geqslant \lceil n/(s-1) \rceil = t$. By Corollary 5.2.3, G has a subgraph H with $\delta(H) \geqslant t-1$, which by Corollary 1.5.4 contains a copy of T. □

As the main result of this section, we shall now prove one of those rare general theorems providing a relatively good upper bound for the Ramsey numbers of a large class of graphs, a class defined in terms of a standard graph invariant. The theorem deals with the Ramsey numbers of sparse graphs: it says that the Ramsey number of graphs H with bounded maximum degree grows only linearly in $|H|$—an enormous improvement on the exponential bound from the proof of Theorem 9.1.1.

Theorem 9.2.2. (Chvátal, Rödl, Szemerédi & Trotter 1983)
For every positive integer Δ there is a constant c such that

$$R(H) \leqslant c\,|H|$$

for all graphs H with $\Delta(H) \leqslant \Delta$.

(7.1.1)
(7.2.1)
(7.3.2)
(9.1.1)

Proof. The basic idea of the proof is as follows. We wish to show that $H \subseteq G$ or $H \subseteq \overline{G}$ if $|G|$ is large enough (though not too large). Consider an ϵ-regular partition of G, as provided by the regularity lemma. If enough of the ϵ-regular pairs in this partition have positive density, we may hope to find a copy of H in G. If most pairs have zero or low density, we try to find H in $\overline{G}$. Let R, R' and R'' be the 'regularity graphs'[2] of G whose edges correspond to the pairs of density $\geqslant 0$; $\geqslant 1/2$; $< 1/2$; respectively. Then R is the edge-disjoint union of R' and R''.

Now to obtain $H \subseteq G$ or $H \subseteq \overline{G}$, it suffices by Lemma 7.3.2 to ensure that H is contained in a suitable 'inflated regularity graph' R'_s

---

[2] Later, we shall define R'' a little differently, so that it complies with our formal definition of a regularity graph.

the spanning subgraph of R formed by the green edges and those whose corresponding pair has density exactly 1/2. Then R' is a regularity graph of G with parameters ϵ, ℓ and 1/2. And R'' is a regularity graph of $\overline{G}$, with the same parameters: as one easily checks, every pair (V_i, V_j) that is ϵ-regular for G is also ϵ-regular for $\overline{G}$.

By definition of m, our graph K contains a red or a green K^r, for $r := \chi(H) \leqslant \Delta + 1$. Correspondingly, $H \subseteq R'_s$ or $H \subseteq R''_s$. Since $\epsilon \leqslant \epsilon_0$ and $\ell \geqslant s/\epsilon_0$ by (2), both R' and R'' satisfy the requirements of Lemma 7.3.2, so $H \subseteq G$ or $H \subseteq \overline{G}$ as desired. □

r

So far in this section, we have been asking what is the least order of a graph G such that every 2-colouring of its edges yields a monochromatic copy of some given graph H. Rather than focusing on the order of G, we might alternatively try to minimize G itself, with respect to the subgraph relation. Given a graph H, let us call a graph G *Ramsey-minimal* for H if G is minimal with the property that every 2-colouring of its edges yields a monochromatic copy of H.

Ramsey-minimal

What do such Ramsey-minimal graphs look like? Are they unique? The following result, which we include for its pretty proof, answers the second question for some H:

Proposition 9.2.3. *If T is a tree but not a star, then infinitely many graphs are Ramsey-minimal for T.*

(1.5.4)
(5.2.3)
(11.2.2)

Proof. Let $|T| =: r$. We show that for every $n \in \mathbb{N}$ there is a graph of order at least n that is Ramsey-minimal for T.

Let us borrow the assertion of Theorem 11.2.2 from Chapter 11: by that theorem, there exists a graph G with chromatic number $\chi(G) > r^2$ and girth $g(G) > n$. If we colour the edges of G red and green, then the red and the green subgraph cannot both have an r-(vertex-)colouring in the sense of Chapter 5: otherwise we could colour the vertices of G with the pairs of colours from those colourings and obtain a contradiction to $\chi(G) > r^2$. So let $G' \subseteq G$ be monochromatic with $\chi(G') > r$. By Corollary 5.2.3, G' has a subgraph of minimum degree at least r, which contains a copy of T by Corollary 1.5.4.

Let $G^* \subseteq G$ be Ramsey-minimal for T. Clearly, G^* is not a forest: the edges of any forest can be 2-coloured (partitioned) so that no monochromatic subforest contains a path of length 3, let alone a copy of T. (Here we use that T is not a star, and hence contains a P^3.) So G^* contains a cycle, which has length $g(G) > n$ since $G^* \subseteq G$. In particular, $|G^*| > n$ as desired. □

$G[U \to H]$ is the graph on

$$(U \times V(H)) \cup ((V \setminus U) \times \{\emptyset\})$$

in which two vertices (v, w) and (v', w') are adjacent if and only if either $vv' \in E$, or else $v = v' \in U$ and $ww' \in E(H)$.[3]

We prove the following formal strengthening of Theorem 9.3.1:

$G(H_1, H_2)$

For any two graphs H_1, H_2 there exists a graph $G = G(H_1, H_2)$ such that every edge colouring of G with the colours 1 and 2 yields either an induced $H_1 \subseteq G$ with all its edges coloured 1 or an induced $H_2 \subseteq G$ with all its edges coloured 2. (∗)

This formal strengthening makes it possible to apply induction on $|H_1| + |H_2|$, as follows.

If either H_1 or H_2 has no edges (in particular, if $|H_1| + |H_2| \leqslant 1$), then (∗) holds with $G = \overline{K^n}$ for large enough n. For the induction step, we now assume that both H_1 and H_2 have at least one edge, and that (∗) holds for all pairs (H_1', H_2') with smaller $|H_1'| + |H_2'|$.

x_i

H_i', H_i''

For each $i = 1, 2$, pick a vertex $x_i \in H_i$ that is incident with an edge. Let $H_i' := H_i - x_i$, and let H_i'' be the subgraph of H_i' induced by the neighbours of x_i.

We shall construct a sequence $G^0, \dots, G^n$ of disjoint graphs; G^n will be the desired Ramsey graph $G(H_1, H_2)$. Along with the graphs G_i, we shall define subsets $V^i \subseteq V(G^i)$ and a map

$$f : V^1 \cup \dots \cup V^n \to V^0 \cup \dots \cup V^{n-1}$$

such that

$$f(V^i) = V^{i-1} \tag{1}$$

f^i

origin

for all $i \geqslant 1$. Writing $f^i := f \circ \dots \circ f$ for the i-fold composition of f whenever it is defined, and f^0 for the identity map on $V^0 = V(G^0)$, we thus have $f^i(v) \in V^0$ for all $v \in V^i$. We call $f^i(v)$ the *origin* of v.

The subgraphs $G^i[V^i]$ will reflect the structure of G^0 as follows:

Vertices in V^i with different origins are adjacent in G^i if and only if their origins are adjacent in G^0. (2)

Assertion (2) will not be used formally in the proof below. However, it can help us to visualize the graphs G^i: every G^i (more precisely, every $G^i[V^i]$—there will also be some vertices $x \in G^i - V^i$) is essentially an inflated copy of G^0 in which every vertex$w \in G^0$ has been replaced by

---

[3] The replacement of $V \setminus U$ by $(V \setminus U) \times \{\emptyset\}$ is just a formal device to ensure that all vertices of $G[U \to H]$ have the same form (v, w), and that $G[U \to H]$ is formally disjoint from G.

$i-1$, then (2) holds again for i (in $\tilde{G}^{i-1}$). The graph $\tilde{G}^{i-1}$ is already the 'essential part' of G^i: the part that looks like an inflated copy of G^0.

$\mathcal{F}$

In the second step we now extend $\tilde{G}^{i-1}$ to the desired graph G^i by adding some further vertices $x \notin V^i$. Let $\mathcal{F}$ denote the set of all families F of the form

$$F = \big(H_1'(u) \mid u \in U^{i-1}\big)\,,$$

$H_1'(u)$

$x(F)$

$H_1''(u)$

G^i

where each $H_1'(u)$ is an induced subgraph of $G_2(u)$ isomorphic to H_1'. (Less formally: $\mathcal{F}$ is the collection of ways to select from each $G_2(u)$ exactly one induced copy of H_1'.) For each $F \in \mathcal{F}$, add a vertex $x(F)$ to $\tilde{G}^{i-1}$ and join it to all the vertices of $H_1''(u)$ for every $u \in U^{i-1}$, where $H_1''(u)$ is the image of H_1'' under some isomorphism $H_1' \to H_1'(u)$ (Fig. 9.3.2). Denote the resulting graph by G^i. This completes the inductive definition of the graphs $G^0, \dots, G^n$.

Let us now show that $G := G^n$ satisfies $(*)$. To this end, we prove the following assertion $(**)$ about G^i for $i = 0, \dots, n$:

> *For every edge colouring with the colours 1 and 2, G^i contains either an induced H_1 coloured 1, or an induced H_2 coloured 2, or an induced subgraph H coloured 2 such that $V(H) \subseteq V^i$ and the restriction of f^i to $V(H)$ is an isomorphism between H and $G^0\,[\,W_k'\,]$ for some $k \in \{\,i, \dots, n-1\,\}$.* $(**)$

Note that the third of the above cases cannot arise for $i = n$, so $(**)$ for n is equivalent to $(*)$ with $G := G^n$.

For $i = 0$, $(**)$ follows from the choice of G^0 as a copy of $G_1 = G(H_1, H_2')$ and the definition of the sets W_k'. Now let $1 \leqslant i \leqslant n$, and assume $(**)$ for smaller values of i.

Let an edge colouring of G^i be given. For each $u \in U^{i-1}$ there is a copy of G_2 in G^i:

$$G^i \supseteq G_2(u) \simeq G(H_1', H_2)\,.$$

x

y_u

If $G_2(u)$ contains an induced H_2 coloured 2 for some $u \in U^{i-1}$, we are done. If not, then every $G_2(u)$ has an induced subgraph $H_1'(u) \simeq H_1'$ coloured 1. Let F be the family of these graphs $H_1'(u)$, one for each $u \in U^{i-1}$, and let $x := x(F)$. If, for some $u \in U^{i-1}$, all the x–$H_1''(u)$ edges in G^i are also coloured 1, we have an induced copy of H_1 in G^i and are again done. We may therefore assume that each $H_1''(u)$ has a vertex y_u for which the edge xy_u is coloured 2. Let

$\hat{U}^{i-1}$

$$\hat{U}^{i-1} := \{\, y_u \mid u \in U^{i-1} \,\} \subseteq V^i.$$

Then f defines an isomorphism from

$\hat{G}^{i-1}$

$$\hat{G}^{i-1} := G^i \left[\,\hat{U}^{i-1} \cup \big\{(v, \emptyset) \mid v \in V(G^{i-1}) \smallsetminus U^{i-1}\big\}\,\right]$$

Let us return once more to the reformulation of Ramsey's theorem considered at the beginning of this section: for every graph H there exists a graph G such that every 2-colouring of the edges of G yields a monochromatic $H \subseteq G$. The graph G for which this follows at once from Ramsey's theorem is a sufficiently large complete graph. If we ask, however, that G shall not contain any complete subgraphs larger than those in H, i.e. that $\omega(G) = \omega(H)$, the problem again becomes difficult—even if we do not require H to be induced in G.

Our second proof of Theorem 9.3.1 solves both problems at once: given H, we shall construct a Ramsey graph for H with the same clique number as H.

bipartite

For this proof, i.e. for the remainder of this section, let us view bipartite graphs P as triples (V_1, V_2, E), where V_1 and V_2 are the two vertex classes and $E \subseteq V_1 \times V_2$ is the set of edges. The reason for this more explicit notation is that we want embeddings between bipartite graphs to respect their bipartitions: given another bipartite graph $P' = (V_1', V_2', E')$, an injective map $\varphi\colon V_1 \cup V_2 \to V_1' \cup V_2'$ will be called an *embedding* of P in P' if $\varphi(V_i) \subseteq V_i'$ for $i = 1, 2$ and $\varphi(v_1)\varphi(v_2)$ is an edge of P' if and only if $v_1 v_2$ is an edge of P. (Note that such embeddings are 'induced'.) Instead of $\varphi\colon V_1 \cup V_2 \to V_1' \cup V_2'$ we may simply write $\varphi\colon P \to P'$.

embedding $P \to P'$

We need two lemmas.

E

Lemma 9.3.2. *Every bipartite graph can be embedded in a bipartite graph of the form* $(X, [X]^k, E)$ *with* $E = \{ xY \mid x \in Y \}$.

Proof. Let P be any bipartite graph, with vertex classes $\{ a_1, \dots, a_n \}$ and $\{ b_1, \dots, b_m \}$, say. Let X be a set with $2n + m$ elements, say

$$X = \{ x_1, \dots, x_n,\ y_1, \dots, y_n,\ z_1, \dots, z_m \};$$

we shall define an embedding $\varphi\colon P \to (X, [X]^{n+1}, E)$.

Let us start by setting $\varphi(a_i) := x_i$ for all $i = 1, \dots, n$. Which $(n+1)$-sets $Y \subseteq X$ are suitable candidates for the choice of $\varphi(b_i)$ for a given vertex b_i? Clearly those adjacent exactly to the images of the neighbours of b_i, i.e. those satisfying

$$Y \cap \{ x_1, \dots, x_n \} = \varphi(N_P(b_i)). \tag{1}$$

Since $d(b_i) \leqslant n$, the requirement of (1) leaves at least one of the $n+1$ elements of Y unspecified. In addition to $\varphi(N_P(b_i))$, we may therefore include in each $Y = \varphi(b_i)$ the vertex z_i as an 'index'; this ensures that $\varphi(b_i) \neq \varphi(b_j)$ for $i \neq j$, even when b_i and b_j have the same neighbours in P. To specify the sets $Y = \varphi(b_i)$ completely, we finally fill them up with 'dummy' elements y_j until $|Y| = n + 1$. □

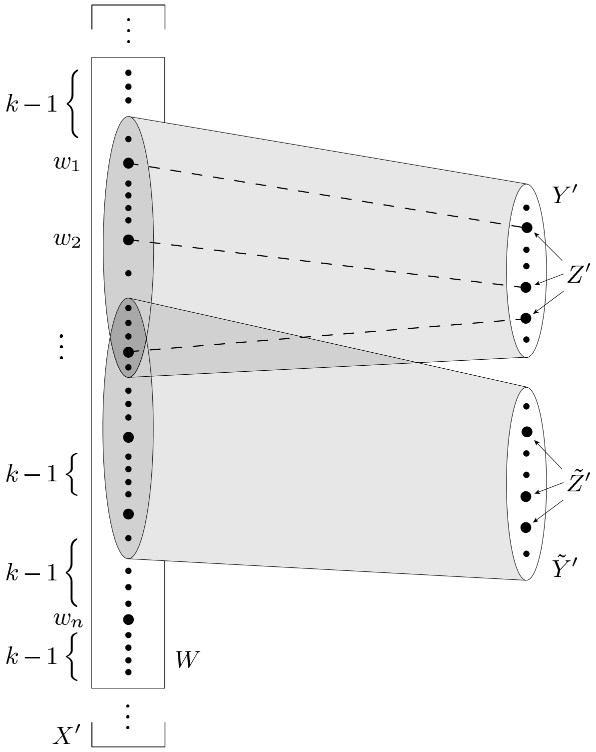

Fig. 9.3.4. The graph of Lemma 9.3.3

$\varphi|_{[X]^k}$

We now define φ on $[X]^k$. Given $Y \in [X]^k$, we wish to choose $\varphi(Y) =: Y' \in [X']^{k'}$ so that the neighbours of Y' among the vertices in $\varphi(X)$ are precisely the images of the neighbours of Y in P, i.e. the vertices $\varphi(x)$ with $x \in Y$, and so that all these edges at Y' are coloured α. To find such a set Y', we first fix its subset Z' as $\{\, \varphi(x) \mid x \in Y \,\}$ (these are k vertices of type w_i) and then extend Z' by $k' - k$ further vertices $u \in W \setminus \varphi(X)$ to a set $Y' \in [W]^{k'}$, in such a way that Z' lies correctly within Y', i.e. so that $\sigma_{Y'}(Z') = S$. This can be done, because $k - 1 = k' - k$ other vertices of W lie between any two w_i. Then

$$Y' \cap \varphi(X) = Z' = \{\, \varphi(x) \mid x \in Y \,\},$$

so Y' has the correct neighbours in $\varphi(X)$, and all the edges between Y' and these neighbours are coloured α (because those neighbours lie in Z' and Y' is associated with α). Finally, φ is injective on $[X]^k$: the images Y' of different vertices Y are distinct, because their intersections with $\varphi(X)$ differ. Hence, our map φ is indeed an embedding of P in P'. □

Formally, we shall define a sequence $G^0, \dots, G^m$ of n-partite graphs G^k, with n-partition $\{ V_1^k, \dots, V_n^k \}$ say, and then let $G := G^m$. The graph G^0 has been defined above; let $V_1^0, \dots, V_n^0$ be its rows:

V_i^0

$$V_i^0 := \bigl\{ (i,j) \mid j = 1, \dots, \tbinom{n}{r} \bigr\} .$$

e_k, m — i_1, i_2 — P — P' — W_1, W_2 — φ_p, q

Now let $e_1, \dots, e_m$ be an enumeration of the edges of K. For $k = 0, \dots, m-1$, construct G^{k+1} from G^k as follows. If $e_{k+1} = i_1 i_2$, say, let $P = (V_{i_1}^k, V_{i_2}^k, E)$ be the bipartite subgraph of G^k induced by its i_1th and i_2th row. By Lemma 9.3.3, P has a bipartite Ramsey graph $P' = (W_1, W_2, E')$. We wish to define $G^{k+1} \supseteq P'$ in such a way that every (monochromatic) embedding $P \to P'$ can be extended to an embedding $G^k \to G^{k+1}$. Let $\{ \varphi_1, \dots, \varphi_q \}$ be the set of all embeddings of P in P', and let

$$V(G^{k+1}) := V_1^{k+1} \cup \ldots \cup V_n^{k+1},$$

where

$$V_i^{k+1} := \begin{cases} W_1 & \text{for } i = i_1 \\ W_2 & \text{for } i = i_2 \\ \bigcup_{p=1}^{q} (V_i^k \times \{ p \}) & \text{for } i \notin \{ i_1, i_2 \}. \end{cases}$$

(Thus for $i \neq i_1, i_2$, we take as V_i^{k+1} just q disjoint copies of V_i^k.) We now define the edge set of G^{k+1} so that the obvious extensions of φ_p to all of $V(G^k)$ become embeddings of G^k in G^{k+1}: for $p = 1, \dots, q$, let $\psi_p \colon V(G^k) \to V(G^{k+1})$ be defined by

$$\psi_p(v) := \begin{cases} \varphi_p(v) & \text{for } v \in P \\ (v, p) & \text{for } v \notin P \end{cases}$$

and let

$$E(G^{k+1}) := \bigcup_{p=1}^{q} \{ \psi_p(v)\psi_p(v') \mid vv' \in E(G^k) \} .$$

Now for every 2-colouring of its edges, G^{k+1} contains an induced copy $\psi_p(G^k)$ of G^k whose edges in P, i.e. those between its i_1th and i_2th row, have the same colour: just choose p so that $\varphi_p(P)$ is the monochromatic induced copy of P in P' that exists by Lemma 9.3.3.

We claim that $G := G^m$ satisfies the assertion of the theorem. So let a 2-colouring of the edges of G be given. By the construction of G^m from G^{m-1}, we can find in G^m an induced copy of G^{m-1} such that for $e_m = ii'$ all edges between the ith and the i'th row have the same colour. In the same way, we find inside this copy of G^{m-1} an induced copy of G^{m-2} whose edges between the ith and the i'th row have the same colour also for $ii' = e_{m-1}$. Continuing in this way, we finally arrive at an induced copy of G^0 in G such that, for each pair (i, i'), all the edges between V_i^0 and $V_{i'}^0$ have the same colour. As shown earlier, this G^0 contains a monochromatic induced copy H_j of H. □

Assume now that G has a vertex v of degree $> d$. Since G is 2-connected, $G - v$ is connected and thus has a spanning tree; let T be a minimal tree in $G - v$ that contains all the neighbours of v. Then every leaf of T is a neighbour of v. By the choice of d, either T has a vertex of degree $\geqslant r$ or T contains a path of length $\geqslant r$, without loss of generality linking two leaves. Together with v, such a path forms a cycle of length $\geqslant r$. A vertex u of degree $\geqslant r$ in T can be joined to v by r independent paths through T, to form a $TK_{2,r}$. □

Theorem 9.4.3. (Oporowski, Oxley & Thomas 1993)
For every $r \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that every 3-connected graph of order at least n contains a wheel of order r or a $K_{3,r}$ as a minor.

Let us call a graph of the form $C^n * \overline{K^2}$ $(n \geqslant 4)$ a *double wheel*, the 1-skeleton of a triangulation of the cylinder as in Fig. 9.4.1 a *crown*, and the 1-skeleton of a triangulation of the Möbius strip a *Möbius crown*.

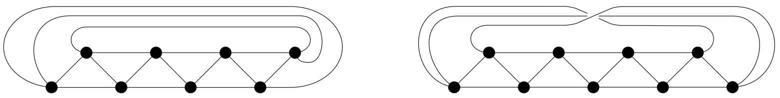

Fig. 9.4.1. A crown and a Möbius crown

Theorem 9.4.4. (Oporowski, Oxley & Thomas 1993)
For every $r \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that every 4-connected graph with at least n vertices has a minor of order $\geqslant r$ that is a double wheel, a crown, a Möbius crown, or a $K_{4,s}$.

Note that the minors occurring in Theorems 9.4.3 and 9.4.4 are themselves 3- and 4-connected, respectively, and are not minors of one another. Thus in each case, the collection of minors is minimal in the sense discussed earlier.

Exercises

1.$^-$ Determine the Ramsey number $R(3)$.

2. Deduce the case $k = 2$ (but c arbitrary) of Theorem 9.1.4 directly from Theorem 9.1.1.

3.$^+$ Construct a graph on $\mathbb{R}$ that has neither a complete nor an edgeless induced subgraph on $|\mathbb{R}| = 2^{\aleph_0}$ vertices. (So Ramsey's theorem does not extend to uncountable sets.)

4.$^+$ Use Ramsey's theorem to show that for any $k, \ell \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that every sequence of n distinct integers contains an increasing subsequence of length $k+1$ or a decreasing subsequence of length $\ell+1$. Find an example showing that $n > k\ell$. Then prove the theorem of Erdős and Szekeres that $n = k\ell + 1$ will do.

Notes

Due to increased interaction with research on random and pseudo-random[4] structures (the latter being provided, for example, by the regularity lemma), the Ramsey theory of graphs has recently seen a period of major activity and advance. Theorem 9.2.2 is an early example of this development.

For the more classical approach, the introductory text by R.L. Graham, B.L. Rothschild & J.H. Spencer, *Ramsey Theory* (2nd edn.), Wiley 1990, makes stimulating reading. This book includes a chapter on graph Ramsey theory, but is not confined to it. A more recent general survey is given by J. Nešetřil in the *Handbook of Combinatorics* (R.L. Graham, M. Grötschel & L. Lovász, eds.), North-Holland 1995. The Ramsey theory of infinite sets forms a substantial part of combinatorial set theory, and is treated in depth in P. Erdős, A. Hajnal, A. Máté & R. Rado, *Combinatorial Set Theory*, North-Holland 1984. An attractive collection of highlights from various branches of Ramsey theory, including applications in algebra, geometry and point-set topology, is offered in B. Bollobás, *Graph Theory*, Springer GTM 63, 1979.

König's infinity lemma, or *König's lemma* for short, is contained in the first-ever book on the subject of graph theory: D. König, *Theorie der endlichen und unendlichen Graphen*, Akademische Verlagsgesellschaft, Leipzig 1936. The *compactness* technique for deducing finite results from infinite (or vice versa), hinted at in Section 9.1, is less mysterious than it sounds. As long as 'infinite' means 'countably infinite', it is precisely the art of applying the infinity lemma (as in the proof of Theorem 9.1.4), no more no less. For larger infinite sets, the same argument becomes equivalent to the well-known theorem of Tychonov that arbitrary products of compact spaces are compact—which has earned the compactness argument its name. Details can be found in Ch. 6, Thm. 10 of Bollobás, and in Graham, Rothschild & Spencer, Ch. 1, Thm. 4. Another frequently used version of the general compactness argument is *Rado's selection lemma*; see A. Hajnal's chapter on infinite combinatorics in the Handbook cited above.

Theorem 9.2.2 is due to V. Chvátal, V. Rödl, E. Szemerédi & W.T. Trotter, The Ramsey number of a graph with bounded maximum degree, *J. Combin. Theory B* **34** (1983), 239–243. Our proof follows the sketch in J. Komlós & M. Simonovits, Szemerédi's Regularity Lemma and its applications in graph theory, in (D. Miklós, V.T. Sós & T. Szőnyi, eds.) *Paul Erdős is 80*, Vol. 2, Proc. Colloq. Math. Soc. János Bolyai (1996). The theorem marks a breakthrough towards a conjecture of Burr and Erdős (1975), which asserts that the Ramsey numbers of graphs with bounded *average* degree in every subgraph are linear: for every $d \in \mathbb{N}$, the conjecture says, there exists a constant c such that $R(H) \leqslant c\,|H|$ for all graphs H with $d(H') \leqslant d$ for all $H' \subseteq H$. This conjecture has been verified also for the class of planar graphs (Chen & Schelp 1993) and, more generally, for the class of graphs not containing K^r (for any fixed r) as a topological minor (Rödl & Thomas 1996). See Nešetřil's Handbook chapter for references.

---

[4] Concrete graphs whose structure resembles the structure expected of a random graph are called *pseudo-random*. For example, the bipartite graphs spanned by an ϵ-regular pair of vertex sets in a graph are pseudo-random.

The following classic result derives its significance from this background:

Theorem 10.1.1. (Dirac 1952)
Every graph with $n \geqslant 3$ vertices and minimum degree at least $n/2$ has a Hamilton cycle.

Proof. Let $G = (V, E)$ be a graph with $|G| = n \geqslant 3$ and $\delta(G) \geqslant n/2$. Then G is connected: otherwise, the degree of any vertex in the smallest component C of G would be less than $|C| \leqslant n/2$.

Let $P = x_0 \dots x_k$ be a longest path in G. By the maximality of P, all the neighbours of x_0 and all the neighbours of x_k lie on P. Hence at least $n/2$ of the vertices $x_0, \dots, x_{k-1}$ are adjacent to x_k, and at least $n/2$ of these same $k < n$ vertices x_i are such that $x_0 x_{i+1} \in E$. By the pigeon hole principle, there is a vertex x_i that has both properties, so we have $x_0 x_{i+1} \in E$ and $x_i x_k \in E$ for some $i < k$ (Fig. 10.1.1).

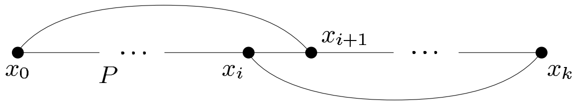

Fig. 10.1.1. Finding a Hamilton cycle in the proof Theorem 10.1.1

We claim that the cycle $C := x_0 x_{i+1} P x_k x_i P x_0$ is a Hamilton cycle of G. Indeed, since G is connected, C would otherwise have a neighbour in $G - C$, which could be combined with a spanning path of C into a path longer than P. □

Theorem 10.1.1 is best possible in that we cannot replace the bound of $n/2$ with $\lfloor n/2 \rfloor$: if n is odd and G is the union of two copies of $K^{\lceil n/2 \rceil}$ meeting in one vertex, then $\delta(G) = \lfloor n/2 \rfloor$ but $\kappa(G) = 1$, so G cannot have a Hamilton cycle. In other words, the high level of the bound of $\delta \geqslant n/2$ is needed to ensure, if nothing else, that G is 2-connected: a condition just as trivially necessary for hamiltonicity as a minimum degree of at least 2. It would seem, therefore, that prescribing some high (constant) value for κ rather than for δ stands a better chance of implying hamiltonicity. However, this is not so: although k-connected graphs contain long cycles in terms of k (Ex. 14, Ch. 3), the graphs $K_{n,k}$ show that their circumference need not grow with n.

There is another invariant with a similar property: a low independence number $\alpha(G)$ ensures that G has long cycles (Ex. 13, Ch. 5), though not necessarily a Hamilton cycle. Put together, however, the two assumptions of high connectivity and low independence number surprisingly complement each other to produce a sufficient condition for hamiltonicity:

10.2 Hamilton cycles and degree sequences

Historically, Dirac's theorem formed the point of departure for the discovery of a series of weaker and weaker degree conditions, all sufficient for hamiltonicity. The development culminated in a single theorem that encompasses all the earlier results: the theorem we shall prove in this section.

degree sequence

hamiltonian sequence

pointwise greater

If G is a graph with n vertices and degrees $d_1 \leqslant \ldots \leqslant d_n$, then the n-tuple $(d_1, \ldots, d_n)$ is called the *degree sequence* of G. Note that this sequence is unique, even though G has several vertex enumerations giving rise to its degree sequence. Let us call an arbitrary integer sequence $(a_1, \ldots, a_n)$ *hamiltonian* if every graph with n vertices and a degree sequence pointwise greater than $(a_1, \ldots, a_n)$ is hamiltonian. (A sequence $(d_1, \ldots, d_n)$ is *pointwise greater* than $(a_1, \ldots, a_n)$ if $d_i \geqslant a_i$ for all i.)

The following theorem characterizes all hamiltonian sequences:

Theorem 10.2.1. (Chvátal 1972)
An integer sequence $(a_1, \ldots, a_n)$ *such that* $0 \leqslant a_1 \leqslant \ldots \leqslant a_n < n$ *and* $n \geqslant 3$ *is hamiltonian if and only if the following holds for every* $i < n/2$*:*

$$a_i \leqslant i \Rightarrow a_{n-i} \geqslant n-i\,.$$

$(a_1, \ldots, a_n)$

$(d_1, \ldots, d_n)$

Proof. Let $(a_1, \ldots, a_n)$ be an arbitrary integer sequence such that $0 \leqslant a_1 \leqslant \ldots \leqslant a_n < n$ and $n \geqslant 3$. We first assume that this sequence satisfies the condition of the theorem and prove that it is hamiltonian. Suppose not; then there exists a graph $G = (V, E)$ with a degree sequence $(d_1, \ldots, d_n)$ such that

$$d_i \geqslant a_i \qquad \text{for all } i \tag{1}$$

$G = (V, E)$

$v_1, \ldots, v_n$

but G has no Hamilton cycle. Let G be chosen with the maximum number of edges, and let $(v_1, \ldots, v_n)$ be an enumeration of V with $d(v_i) = d_i$ for all i. By (1), our assumptions for $(a_1, \ldots, a_n)$ transfer to $(d_1, \ldots, d_n)$, i.e.,

$$d_i \leqslant i \Rightarrow d_{n-i} \geqslant n-i \qquad \text{for all } i < n/2. \tag{2}$$

x, y

$x_1, \ldots, x_n$

Let x, y be distinct and non-adjacent vertices in G, with $d(x) \leqslant d(y)$ and $d(x) + d(y)$ as large as possible. One easily checks that the degree sequence of $G + xy$ is pointwise greater than $(d_1, \ldots, d_n)$, and hence than $(a_1, \ldots, a_n)$. Hence, by the maximality of G, the new edge xy lies on a Hamilton cycle H of $G + xy$. Then $H - xy$ is a Hamilton path $x_1, \ldots, x_n$ in G, with $x_1 = x$ and $x_n = y$ say.

As in the proof of Dirac's theorem, we now consider the index sets

$$I := \{\, i \mid xx_{i+1} \in E \,\} \quad \text{and} \quad J := \{\, j \mid x_j y \in E \,\}\,.$$

i.e. the union of a K^{n-h} on the vertices $v_{h+1}, \ldots, v_n$ and a $K_{h,h}$ with partition sets $\{ v_1, \ldots, v_h \}$ and $\{ v_{n-h+1}, \ldots, v_n \}$. $\square$

By applying Theorem 10.2.1 to $G * K^1$, one can easily prove the following adaptation of the theorem to Hamilton paths. Let an integer sequence be called *path-hamiltonian* if every graph with a pointwise greater degree sequence has a Hamilton path.

Corollary 10.2.2. *An integer sequence* $(a_1, \ldots, a_n)$ *such that* $n \geqslant 2$ *and* $0 \leqslant a_1 \leqslant \ldots \leqslant a_n < n$ *is path-hamiltonian if and only if every* $i \leqslant n/2$ *is such that* $a_i < i \Rightarrow a_{n+1-i} \geqslant n-i$. $\square$

10.3 Hamilton cycles in the square of a graph

G^d

Given a graph G and a positive integer d, we denote by G^d the graph on $V(G)$ in which two vertices are adjacent if and only if they have distance at most d in G. Clearly, $G = G^1 \subseteq G^2 \subseteq \ldots$ Our goal in this section is to prove the following fundamental result:

Theorem 10.3.1. (Fleischner 1974)
If G is a 2-connected graph, then G^2 has a Hamilton cycle.

bridges

We begin with three simple lemmas. Let us say that an edge $e \in G^2$ *bridges* a vertex $v \in G$ if its ends are neighbours of v in G.

Lemma 10.3.2. *Let* $P = v_0 \ldots v_k$ *be a path* ($k \geqslant 1$)*, and let G be the graph obtained from P by adding two vertices u, w, together with the edges uv_1 and wv_k (Fig. 10.3.1).*

(i) P^2 *contains a path Q from v_0 to v_1 with $V(Q) = V(P)$ and $v_{k-1}v_k \in E(Q)$, such that each of the vertices $v_1, \ldots, v_{k-1}$ is bridged by an edge of Q.*

(ii) G^2 *contains disjoint paths Q from v_0 to v_k and Q' from u to w, such that $V(Q) \cup V(Q') = V(G)$ and each of the vertices $v_1, \ldots, v_k$ is bridged by an edge of Q or Q'.*

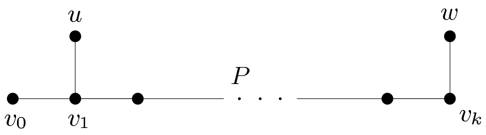

Fig. 10.3.1. The graph G in Lemma 10.3.2

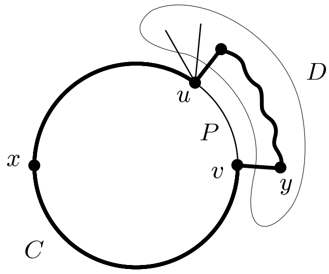

Fig. 10.3.3. The proof of Lemma 10.3.4

Proof of Theorem 10.3.1. We show by induction on $|G|$ that, given any vertex $x^* \in G$, there is a Hamilton cycle H in G^2 with the following property:

$$\text{Both edges of } H \text{ at } x^* \text{ lie in } G. \qquad (*)$$

For $|G| = 3$, we have $G = K^3$ and the assertion is trivial. So let $|G| \geqslant 4$, assume the assertion for graphs of smaller order, and let $x^* \in V(G)$ be given. By Lemma 10.3.4, there is a cycle $C \subseteq G$ that contains both x^* and a vertex $y^* \neq x^*$ whose neighbours in G all lie on C. x^* C y^*

If C is a Hamilton cycle of G, there is nothing to show; so assume that $G - C \neq \emptyset$. Consider a component D of $G - C$. Let $\tilde{D}$ denote the graph $G/(G-D)$ obtained from G by contracting $G-D$ into a new vertex $\tilde{x}$. If $|D| = 1$, set $\mathcal{P}(D) := \{D\}$. If $|D| > 1$, then $\tilde{D}$ is again 2-connected. Hence, by the induction hypothesis, $\tilde{D}^2$ has a Hamilton cycle $\tilde{C}$ whose edges at $\tilde{x}$ both lie in $\tilde{D}$. Note that the path $\tilde{C} - \tilde{x}$ may have some edges that do not lie in G^2: edges joining two neighbours of $\tilde{x}$ that have no common neighbour in G (and are themselves non-adjacent in G). Let $\tilde{E}$ denote the set of these edges, and let $\mathcal{P}(D)$ denote the set of components of $(\tilde{C} - \tilde{x}) - \tilde{E}$; this is a set of paths in G^2 whose ends are adjacent to $\tilde{x}$ in $\tilde{D}$ (Fig. 10.3.4). $\mathcal{P}(D)$ $\tilde{C}$ $\mathcal{P}(D)$

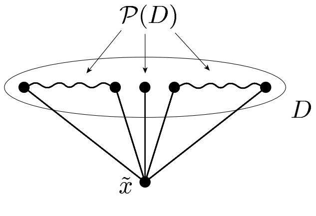

Fig. 10.3.4. $\mathcal{P}(D)$ consists of three paths, one of which is trivial

$\mathcal{P}$ Let $\mathcal{P}$ denote the union of the sets $\mathcal{P}(D)$ over all components D of $G - C$. Clearly, $\mathcal{P}$ has the following properties:

paths from $\mathcal{P}_1$. For the construction of W we assume that $\mathcal{P}_2 \neq \emptyset$; the case of $\mathcal{P}_2 = \emptyset$ is much simpler and will be treated later.

v^+, *right*
v^-, *left*

We start by choosing a fixed cyclic orientation of C, a bijection $i \mapsto v_i$ from $\mathbb{Z}_{|C|}$ to $V(C)$ with $v_i v_{i+1} \in E(C)$ for all $i \in \mathbb{Z}_{|C|}$. Let us think of this orientation as clockwise; then every vertex $v_i \in C$ has a *right* neighbour $v_i^+ := v_{i+1}$ and a *left* neighbour $v_i^- := v_{i-1}$. Accordingly, the edge v^-v lies to the *left* of v, the edge vv^+ lies on its *right*, and so on.

interval
$[\,v,w\,]$ *etc.*
I^*, P^*
Q^*

A non-trivial path $P = v_i v_{i+1} \dots v_{j-1} v_j$ in C such that $V(P) \cap X_2 = \{\,v_i, v_j\,\}$ will be called an *interval*, with *left end* v_i and *right end* v_j. Thus, C is the union of $|X_2| = 2\,|\mathcal{P}_2|$ intervals. As usual, we write $P =: [\,v_i, v_j\,]$ and set $(v_i, v_j) := \mathring{P}$ as well as $[\,v_i, v_j) := P\mathring{v}_j$ and $(v_i, v_j\,] := \mathring{v}_i P$. For intervals $[\,u,v\,]$ and $[\,v,w\,]$ with a common end v we say that $[\,u,v\,]$ lies to the *left* of $[\,v,w\,]$, and $[\,v,w\,]$ lies to the *right* of $[\,u,v\,]$. We denote the unique interval $[\,v,w\,]$ with $x^* \in (v,w]$ as I^*, the path in $\mathcal{P}_2$ with foot w as P^*, and the path I^*wP^* as Q^*.

W

For the construction of W, we may think of $\tilde{G}$ as a multigraph M on X_2 whose edges are the intervals on C and the paths in $\mathcal{P}_2$ (with their feet as ends). By (2), M is cubic, so we may apply Lemma 10.3.3 with $e := I^*$ and $f := P^*$. The lemma provides us with a closed walk W in $\tilde{G}$ which traverses I^* once, every other interval of C once or twice, and every path in $\mathcal{P}_2$ once. Moreover, W contains Q^* as a subpath. The two edges at x^* of this path lie in G; in this sense, W already satisfies $(*)$.

$e(v)$

Let us now modify W so that W passes through every vertex of C exactly once. Simultaneously, we shall prepare for the later inclusion of the paths from $\mathcal{P}_1$ by defining a map $v \mapsto e(v)$ that is injective on X_1 and assigns to every $v \in X_1$ an edge $e(v)$ of the modified W with the following property:

> *The edge $e(v)$ either bridges v or is incident with it. In the latter case, $e(v) \in C$ and $e(v) \neq vx^*$.* $(**)$

For simplicity, we shall define the map $v \mapsto e(v)$ on all of $V(C) \setminus X_2$, a set that includes X_1 by (2). To ensure injectivity on X_1, we only have to make sure that no edge $vw \in C$ is chosen both as $e(v)$ and as $e(w)$. Indeed, since $|X_1| \geqslant 2$ if injectivity is a problem, and $\mathcal{P}_2 \neq \emptyset$ by assumption, we have $|C - y^*| \geqslant |X_1| + 2\,|\mathcal{P}_2| \geqslant 4$ and hence $|C| \geqslant 5$; thus, no edge of G^2 can bridge more than one vertex of C, or bridge a vertex of C and lie on C at the same time.

For our intended adjustments of W at the vertices of C, we consider the intervals of C one at a time. By definition of W, every interval is of one of the following three types:

Type 1: W traverses I once;

Type 2: W traverses I twice, in one direction and back immediately afterwards (formally: W contains a triple (e, x, e) with $x \in X_2$ and $e \in E(I)$);

edge. Since W traverses I^+ no more than twice, it must traverse the edge x_2y_2 immediately after one of its two subpaths y_1x_1I and $x_1^-x_1I$. Take the starting vertex of this subpath (y_1 or x_1^-) as the vertex u in Lemma 10.3.2 (ii), and the other vertex in $\{y_1, x_1^-\}$ as v_0; moreover, set $v_k := x_2$ and $w := y_2$. Then the lemma enables us to replace these two subpaths of W between $\{y_1, x_1^-\}$ and $\{x_2, y_2\}$ by disjoint paths in G^2 (Fig. 10.3.6), and furthermore assigns to every vertex $v \in \mathring{I}$ an edge $e(v)$ of one of those paths, bridging v.

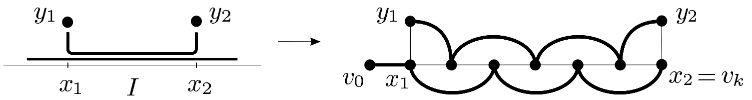

Fig. 10.3.6. A type 3 modification for the case $u = y_1$ and k odd

Following the above modifications, W is now a closed walk in $\tilde{G}^2$. Let us check that, moreover, W contains every vertex of $\tilde{G}$ exactly once. For vertices of the paths in $\mathcal{P}_2$ this is clear, because W still traverses every such path once and avoids it otherwise. For the vertices of $C - X_2$, it follows from the above modifications by Lemma 10.3.2. So how about the vertices in X_2?

Let $x \in X_2$ be given, and let y be its neighbour on a path in $\mathcal{P}_2$. Let I_1 denote the interval I that satisfied $yxI \subseteq W$ before the modification of W, and let I_2 denote the other interval ending in x. If I_1 is of type 1, then I_2 is of type 2 with dead end x. In this case, x was retained in W when W was modified on I_1 but skipped when W was modified on I_2, and is thus contained exactly once in W now. If I_1 is of type 2, then x is not its dead end, and I_2 is of type 1. The subwalk of W that started with yx and then went along I_1 and back, was replaced with a y–x path. This path is now followed on W by the unchanged interval I_2, so in this case too the vertex x is now contained in W exactly once. Finally, if I_1 is of type 3, then x was contained in one of the replacement paths Q, Q' from Lemma 10.3.2 (ii); as these paths were disjoint by the assertion of the lemma, x is once more left on W exactly once.

We have thus shown that W, after the modifications, is a closed walk in $\tilde{G}^2$ containing every vertex of $\tilde{G}$ exactly once, so W defines a Hamilton cycle $\tilde{H}$ of $\tilde{G}^2$. Since W still contains the path Q^*, $\tilde{H}$ satisfies $(*)$.

$\tilde{H}$

Up until now, we have assumed that $\mathcal{P}_2$ is non-empty. If $\mathcal{P}_2 = \emptyset$, let us set $\tilde{H} := \tilde{G} = C$; then, again, $\tilde{H}$ satisfies $(*)$. It remains to turn $\tilde{H}$ into a Hamilton cycle H of G^2 by incorporating the paths from $\mathcal{P}_1$. In order to be able to treat the case of $\mathcal{P}_2 = \emptyset$ along with the case of $\mathcal{P}_2 \neq \emptyset$, we define a map $v \mapsto e(v)$ also when $\mathcal{P}_2 = \emptyset$, as follows: for

$\tilde{H}$

So v is indeed incident with e. Hence $v \in \{ x^*, z \}$ by definition of e, while $e \neq vx^*$ by $(**)$. Thus $v = x^*$, and e was replaced by a path of the form $x^* y_1 P y_2 z$. Since $x^* y_1$ is an edge of G, this replacement again preserves $(*)$. Therefore H does indeed satisfy $(*)$, and our induction is complete. □

We close the chapter with a far-reaching conjecture generalizing Dirac's theorem:

Conjecture. (Seymour 1974)
Let G be a graph of order $n \geqslant 3$, and let k be a positive integer. If G has minimum degree

$$\delta(G) \geqslant \frac{k}{k+1} n \,,$$

then G has a Hamilton cycle H such that $H^k \subseteq G$.

For $k = 1$, this is precisely Dirac's theorem. The case $k = 2$ had already been conjectured by Pósa in 1963 and was proved for large n by Komlós, Sárközy & Szemerédi (1996).

Exercises

1. Show that every uniquely 3-edge-colourable cubic graph is hamiltonian. ('Unique' means that all 3-edge-colourings induce the same edge partition.)

2. Prove or disprove the following strengthening of Proposition 10.1.2: 'Every k-connected graph G with $|G| \geqslant 3$ and $\chi(G) \geqslant |G|/k$ has a Hamilton cycle.'

3. Given a graph G, consider a maximal sequence $G_0, \ldots, G_k$ such that $G_0 = G$ and $G_{i+1} = G_i + x_i y_i$ for $i = 0, \ldots, k-1$, where x_i, y_i are two non-adjacent vertices of G_i satisfying $d_{G_i}(x_i) + d_{G_i}(y_i) \geqslant |G|$. The last graph of the sequence, G_k, is called the *Hamilton closure* of G. Show that this graph depends only on G, not on the choice of the sequence $G_0, \ldots, G_k$.

4. Let x, y be two nonadjacent vertices of a connected graph G, with $d(x) + d(y) \geqslant |G|$. Show that G has a Hamilton cycle if and only if $G + xy$ has one. Using the previous exercise, deduce the following strengthening of Dirac's theorem: if $d(x) + d(y) \geqslant |G|$ for every two non-adjacent vertices $x, y \in G$, then G has a Hamilton cycle.

5. Given an even positive integer k, construct for every $n \geqslant k$ a k-regular graph of order $2n+1$.

6.$^-$ Find a hamiltonian graph whose degree sequence is not hamiltonian.

this subject in 1883, he seems to have been aware that he could not—really—prove the above statement about Hamilton cycles. His account in P.G. Tait, Listing's topologie, *Phil. Mag.* **17** (1884), 30–46, makes some entertaining reading.

A shorter proof of Tutte's theorem that 4-connected planar graphs are hamiltonian was given by C. Thomassen, A theorem on paths in planar graphs, *J. Graph Theory* **7** (1983), 169–176. Tutte's counterexample to Tait's assumption that even 3-connectedness suffices (at least for cubic graphs) is shown in Bollobás, and in J.A. Bondy & U.S.R. Murty, *Graph Theory with Applications*, Macmillan 1976 (where Tait's attempted proof is discussed in some detail).

Proposition 10.1.2 is due to Chvátal & Erdős (1972). Our proof of Fleischner's theorem is based on S. Říha, A new proof of the theorem by Fleischner, *J. Combin. Theory B* **52** (1991), 117–123. Seymour's conjecture is from P.D. Seymour, Problem 3, in (T.P. McDonough and V.C. Mavron, eds.) *Combinatorics*, Cambridge University Press 1974. Pósa's conjecture was proved for large n by J. Komlós, G.N. Sárközy & E. Szemerédi, On the square of a Hamiltonian cycle in dense graphs, *Random Structures and Algorithms* **9** (1996), 193–211.

last section, we give a detailed proof of a theorem of Erdős and Rényi that illustrates a proof technique frequently used in random graphs, the so-called *second moment method*.

11.1 The notion of a random graph

V $\mathcal{G}$

Let V be a fixed set of n elements, say $V = \{0, \dots, n-1\}$. Our aim is to turn the set $\mathcal{G}$ of all graphs on V into a probability space, and then to consider the kind of questions typically asked about random objects: What is the probability that a graph $G \in \mathcal{G}$ has this or that property? What is the expected value of a given invariant on G, say its expected girth or chromatic number?

p q

Intuitively, we should be able to generate G randomly as follows. For each $e \in [V]^2$ we decide by some random experiment whether or not e shall be an edge of G; these experiments are performed independently, and for each the probability of success—i.e. of accepting e as an edge for G—is equal to some fixed[1] number $p \in [0,1]$. Then if G_0 is some fixed graph on V, with m edges say, the elementary event $\{G_0\}$ has a probability of $p^m q^{\binom{n}{2}-m}$ (where $q := 1-p$): with this probability, our randomly generated graph G is this particular graph G_0. (The probability that G is *isomorphic* to G_0 will usually be greater.) But if the probabilities of all the elementary events are thus determined, then so is the entire probability measure of our desired space $\mathcal{G}$. Hence all that remains to be checked is that such a probability measure on $\mathcal{G}$, one for which all individual edges occur independently with probability p, does indeed exist.[2]

Ω_e P_e $\mathcal{G}(n,p)$

In order to construct such a measure on $\mathcal{G}$ formally, we start by defining for every potential edge $e \in [V]^2$ its own little probability space $\Omega_e := \{0_e, 1_e\}$, choosing $P_e(\{1_e\}) := p$ and $P_e(\{0_e\}) := q$ as the probabilities of its two elementary events. As our desired probability space $\mathcal{G} = \mathcal{G}(n,p)$ we then take the product space

Ω

$$\Omega := \prod_{e \in [V]^2} \Omega_e \,.$$

---

[1] Often, the value of p will depend on the cardinality n of the set V on which our random graphs are generated; thus, p will be the value $p = p(n)$ of some function $n \mapsto p(n)$. Note, however, that V (and hence n) is fixed for the definition of $\mathcal{G}$: for each n separately, we are constructing a probability space of the graphs G on $V = \{0, \dots, n-1\}$, and within each space the probability that $e \in [V]^2$ is an edge of G has the same value for all e.

[2] Any reader ready to believe this may skip ahead now to the end of Proposition 11.1.1, without missing anything.

contrast, the probability that H is an *induced* subgraph of G is $p^\ell q^{\binom{k}{2}-\ell}$: now the edges missing from H are required to be missing from G too, and they do so independently with probability q.

The probability P_H that G has an induced subgraph *isomorphic* to H is usually more difficult to compute: since the possible instances of H on subsets of V overlap, the events that they occur in G are not independent. However, the sum (over all k-sets $U \subseteq V$) of the probabilities $P[\,H \simeq G[\,U\,]\,]$ is always an upper bound for P_H, since P_H is the measure of the union of all those events. For example, if $H = \overline{K^k}$, we have the following trivial upper bound on the probability that G contains an induced copy of H:

[11.2.1] [11.3.4] **Lemma 11.1.2.** *For all integers n, k with $n \geqslant k \geqslant 2$, the probability that $G \in \mathcal{G}(n,p)$ has a set of k independent vertices is at most*

$$P[\,\alpha(G) \geqslant k\,] \leqslant \binom{n}{k} q^{\binom{k}{2}}.$$

Proof. The probability that a fixed k-set $U \subseteq V$ is independent in G is $q^{\binom{k}{2}}$. The assertion thus follows from the fact that there are only $\binom{n}{k}$ such sets U. □

Analogously, the probability that $G \in \mathcal{G}(n,p)$ contains a K^k is at most

$$P[\,\omega(G) \geqslant k\,] \leqslant \binom{n}{k} p^{\binom{k}{2}}.$$

Now if k is fixed, and n is small enough that these bounds for the probabilities $P[\,\alpha(G) \geqslant k\,]$ and $P[\,\omega(G) \geqslant k\,]$ sum to less than 1, then $\mathcal{G}$ contains graphs that have neither property: graphs which contain neither a K^k nor a $\overline{K^k}$ induced. But then any such n is a lower bound for the Ramsey number of k!

As the following theorem shows, this lower bound is quite close to the upper bound of 2^{2k-3} implied by the proof of Theorem 9.1.1:

Theorem 11.1.3. (Erdős 1947)
For every integer $k \geqslant 3$, the Ramsey number of k satisfies

$$R(k) > 2^{k/2}.$$

Proof. For $k = 3$ we trivially have $R(3) \geqslant 3 > 2^{3/2}$, so let $k \geqslant 4$. We show that, for all $n \leqslant 2^{k/2}$ and $G \in \mathcal{G}(n, \frac{1}{2})$, the probabilities $P[\,\alpha(G) \geqslant k\,]$ and $P[\,\omega(G) \geqslant k\,]$ are both less than $\frac{1}{2}$.

Since $p = q = \frac{1}{2}$, Lemma 11.1.2 and the analogous assertion for $\omega(G)$ imply the following for all $n \leqslant 2^{k/2}$ (use that $k! > 2^k$ for $k \geqslant 4$):

$$\begin{aligned} &\geqslant \sum_{\substack{G\in\mathcal{G}(n,p)\\ X(G)\geqslant a}} P(\{G\})\cdot X(G)\\ &\geqslant \sum_{\substack{G\in\mathcal{G}(n,p)\\ X(G)\geqslant a}} P(\{G\})\cdot a\\ &= P[X\geqslant a]\cdot a\,. \end{aligned}$$

□

Since our probability spaces are finite, the expectation can often be computed by a simple application of *double counting*, a standard combinatorial technique we met before in the proofs of Corollary 4.2.8 and Theorem 5.5.3. For example, if X is a random variable on $\mathcal{G}(n,p)$ that counts the number of subgraphs of G in some fixed set $\mathcal{H}$ of graphs on V, then $E(X)$, by definition, counts the number of pairs (G,H) such that $H\subseteq G$, each weighted with the probability of $\{G\}$. Algorithmically, we compute $E(X)$ by going through the graphs $G\in\mathcal{G}(n,p)$ in an 'outer loop' and performing, for each G, an 'inner loop' that runs through the graphs $H\in\mathcal{H}$ and counts '$P(\{G\})$' whenever $H\subseteq G$. Alternatively, we may count the same set of weighted pairs with H in the outer and G in the inner loop: this amounts to adding up, over all $H\subseteq\mathcal{H}$, the probabilities $P[H\subseteq G]$.

To illustrate this once in detail, let us compute the expected number of cycles of some given length $k\geqslant 3$ in a random graph $G\in\mathcal{G}(n,p)$. So let $X\colon\mathcal{G}(n,p)\to\mathbb{N}$ be the random variable that assigns to every random graph G its number of k-cycles, the number of subgraphs isomorphic to C^k. Let us write

X

$(n)_k$

$$(n)_k := n\,(n-1)(n-2)\cdots(n-k+1)$$

for the number of sequences of k distinct elements of a given n-set.

[11.2.2]
[11.4.3]

Lemma 11.1.5. *The expected number of k-cycles in $G\in\mathcal{G}(n,p)$ is*

$$E(X)=\frac{(n)_k}{2k}\,p^k.$$

Proof. For every k-cycle C with vertices in $V=\{0,\dots,n-1\}$, the vertex set of the graphs in $\mathcal{G}(n,p)$, let $X_C\colon\mathcal{G}(n,p)\to\{0,1\}$ denote the *indicator random variable* of C:

$$X_C\colon\ G\mapsto\begin{cases}1 & \text{if } C\subseteq G;\\ 0 & \text{otherwise.}\end{cases}$$

Since X_C takes only 1 as a positive value, its expectation $E(X_C)$ equals the measure $P[X_C=1]$ of the set of all graphs in $\mathcal{G}(n,p)$ that contain C. But this is just the probability that $C\subseteq G$:

$$E(X_C)=P[C\subseteq G]=p^k. \tag{1}$$

$\mathcal{G}(n,p)$ will contain at least one graph without either short cycles or big independent sets.

Unfortunately, such a choice of p is impossible: the two ranges of p do not overlap! As we shall see in Section 11.4, we must keep p below n^{-1} to make the occurrence of short cycles in G unlikely—but for any such p there will most likely be no cycles in G at all (Exercise 19), so G will be bipartite and hence have at least $n/2$ independent vertices.

But all is not lost. In order to make big independent sets unlikely, we shall fix p above n^{-1}, at $n^{\epsilon-1}$ for some $\epsilon > 0$. Fortunately, though, if ϵ is small enough then this will produce only *few* short cycles in G, even compared with n (rather than, more typically, with n^k). If we then delete a vertex in each of those cycles, the graph H obtained will have no short cycles, and its independence number $\alpha(H)$ will be at most that of G. Since H is not much smaller than G, its chromatic number will thus still be large, so we have found a graph with both large girth and large chromatic number.

To prepare for the formal proof of Erdős's theorem, we first show that an edge probability of $p = n^{\epsilon-1}$ is indeed always large enough to ensure that $G \in \mathcal{G}(n,p)$ 'almost surely' has no big independent set of vertices. More precisely, we prove the following slightly stronger assertion:

Lemma 11.2.1. *Let $k > 0$ be an integer, and let $p = p(n)$ be a function of n such that $p \geqslant (6k\ln n)n^{-1}$ for n large. Then*

$$\lim_{n\to\infty} P\,[\,\alpha \geqslant \tfrac{1}{2}n/k\,] = 0\,.$$

(11.1.2) *Proof.* For all integers n, r with $n \geqslant r \geqslant 2$, and all $G \in \mathcal{G}(n,p)$, Lemma 11.1.2 implies

$$\begin{aligned} P\,[\,\alpha \geqslant r\,] &\leqslant \binom{n}{r} q^{\binom{r}{2}} \\ &\leqslant n^r q^{\binom{r}{2}} \\ &= \left(nq^{(r-1)/2}\right)^r \\ &\leqslant \left(ne^{-p(r-1)/2}\right)^r ; \end{aligned}$$

here, the last inequality follows from the fact that $1-p \leqslant e^{-p}$ for all p. (Compare the functions $x \mapsto e^x$ and $x \mapsto x+1$ for $x = -p$.) Now if $p \geqslant (6k\ln n)n^{-1}$ and $r \geqslant \frac{1}{2}n/k$, then the term under the exponent satisfies

$$\begin{aligned} ne^{-p(r-1)/2} &= ne^{-pr/2\,+\,p/2} \\ &\leqslant ne^{-(3/2)\ln n\,+\,p/2} \\ &\leqslant n\,n^{-3/2}\,e^{1/2} \\ &= \sqrt{e}/\sqrt{n} \underset{n\to\infty}{\longrightarrow} 0\,. \end{aligned}$$

11.3 Properties of almost all graphs

property

A *graph property* is a class of graphs that is closed under isomorphism, one that contains with every graph G also the graphs isomorphic to G. If $p = p(n)$ is a fixed function (possibly constant), and $\mathcal{P}$ is a graph property, we may ask how the probability $P[\,G \in \mathcal{P}\,]$ behaves for $G \in \mathcal{G}(n,p)$ as $n \to \infty$. If this probability tends to 1, we say that $G \in \mathcal{P}$ for *almost all* (or *almost every*) $G \in \mathcal{G}(n,p)$, or that $G \in \mathcal{P}$ *almost surely*; if it tends to 0, we say that *almost no* $G \in \mathcal{G}(n,p)$ has the property $\mathcal{P}$. (For example, in Lemma 11.2.1 we proved that, for a certain p, almost no $G \in \mathcal{G}(n,p)$ has a set of more than $\frac{1}{2}n/k$ independent vertices.)

almost all etc.

To illustrate the new concept let us show that, for constant p, every fixed abstract[3] graph H is an induced subgraph of almost all graphs:

Proposition 11.3.1. *For every constant $p \in (0,1)$ and every graph H, almost every $G \in \mathcal{G}(n,p)$ contains an induced copy of H.*

Proof. Let H be given, and $k := |H|$. If $n \geqslant k$ and $U \subseteq \{\,0,\dots,n-1\,\}$ is a fixed set of k vertices of G, then $G[\,U\,]$ is isomorphic to H with a certain probability $r > 0$. This probability r depends on p, but not on n (why not?). Now G contains a collection of $\lfloor n/k \rfloor$ disjoint such sets U. The probability that none of the corresponding graphs $G[\,U\,]$ is isomorphic to H is $(1-r)^{\lfloor n/k \rfloor}$, since these events are independent by the disjointness of the edges sets $[U]^2$. Thus

$$P[\,H \not\subseteq G \text{ induced}\,] \leqslant (1-r)^{\lfloor n/k \rfloor} \xrightarrow[n\to\infty]{} 0\,,$$

which implies the assertion. □

The following lemma is a simple device enabling us to deduce that quite a number of natural graph properties (including that of Proposition 11.3.1) are shared by almost all graphs. Given $i,j \in \mathbb{N}$, let $\mathcal{P}_{i,j}$ denote the property that the graph considered contains, for any disjoint vertex sets U, W with $|U| \leqslant i$ and $|W| \leqslant j$, a vertex $v \notin U \cup W$ that is adjacent to all the vertices in U but to none in W.

$\mathcal{P}_{i,j}$

Lemma 11.3.2. *For every constant $p \in (0,1)$ and $i,j \in \mathbb{N}$, almost every graph $G \in \mathcal{G}(n,p)$ has the property $\mathcal{P}_{i,j}$.*

---

[3] The word ‘abstract’ is used to indicate that only the isomorphism type of H is known or relevant, not its actual vertex and edge sets. In our context, it indicates that the word ‘subgraph’ is used in the usual sense of ‘isomorphic to a subgraph’.

Proposition 11.3.4. *For every constant $p \in (0,1)$ and every $\epsilon > 0$, almost every graph $G \in \mathcal{G}(n,p)$ has chromatic number*

$$\chi(G) > \frac{\log(1/q)}{2+\epsilon} \cdot \frac{n}{\log n}\,.$$

(11.1.2) *Proof.* For any fixed $n \geqslant k \geqslant 2$, Lemma 11.1.2 implies

$$\begin{aligned} P\,[\,\alpha \geqslant k\,] &\leqslant \binom{n}{k} q^{\binom{k}{2}} \\ &\leqslant n^k q^{\binom{k}{2}} \\ &= q^{k\frac{\log n}{\log q} + \frac{1}{2}k(k-1)} \\ &= q^{\frac{k}{2}\left(-\frac{2\log n}{\log(1/q)} + k - 1\right)}. \end{aligned}$$

For

$$k := (2+\epsilon)\,\frac{\log n}{\log(1/q)}$$

the exponent of this expression tends to infinity with n, so the expression itself tends to zero. Hence, almost every $G \in \mathcal{G}(n,p)$ is such that in any vertex colouring of G no k vertices can have the same colour, so every colouring uses more than

$$\frac{n}{k} = \frac{\log(1/q)}{2+\epsilon} \cdot \frac{n}{\log n}$$

colours. □

By a result of Bollobás (1988), Proposition 11.3.4 is sharp in the following sense: if we replace ϵ by $-\epsilon$, then the lower bound given for χ turns into an upper bound.

Most of the results of this section have the interesting common feature that the values of p played no role whatsoever: if almost every graph in $\mathcal{G}(n,\frac{1}{2})$ had the property considered, then the same was true for almost every graph in $\mathcal{G}(n, 1/1000)$. How could this happen?

Such insensitivity of our random model to changes of p was certainly not intended: after all, among all the graphs with a certain property $\mathcal{P}$ it is often those having $\mathcal{P}$ 'only just' that are the most interesting—for those graphs are most likely to have different properties too, properties to which $\mathcal{P}$ might thus be set in relation. (The proof of Erdős's theorem is a good example.) For most properties, however—and this explains the above phenomenon—the critical order of magnitude of p around which the property will 'just' occur or not occur lies far below any constant value of p: it is typically a function of n tending to zero as $n \to \infty$.

11.4 Threshold functions and second moments

Consider a graph property of the form

$$\mathcal{P} = \{\, G \mid X(G) > 0 \,\},$$

$X \geqslant 0$

where $X \geqslant 0$ is a random variable on $\mathcal{G}(n,p)$. Countless properties can be expressed naturally in this way; if X denotes the number of spanning trees, for example, then $\mathcal{P}$ corresponds to connectedness.

How could we prove that $\mathcal{P}$ has a threshold function t? Any such proof will consist of two parts: a proof that almost no $G \in \mathcal{G}(n,p)$ has $\mathcal{P}$ when p is small compared with t, and one showing that almost every G has $\mathcal{P}$ when p is large.

If X is integral, we may use Markov's inequality for the first part of the proof and find an upper bound for $E(X)$ instead of $P\,[\,X > 0\,]$: if $E(X)$ is small then $X(G)$ can be positive—and hence at least 1—only for few $G \in \mathcal{G}(n,p)$. Besides, the expectation is much easier to calculate than probabilities: without worrying about such things as independence or incompatibility of events, we may compute the expectation of a sum of random variables—for example, of indicator random variables—simply by adding up their individual expected values.

For the second part of the proof, things are more complicated. In order to show that $P\,[\,X > 0\,]$ is large, it is not enough to bound $E(X)$ from below: since X is not bounded above, $E(X)$ may be large simply because X is very large on just a few graphs G—so X may still be zero for most $G \in \mathcal{G}(n,p)$.[6] In order to prove that $P\,[\,X > 0\,] \to 1$, we thus have to show that this cannot happen, i.e. that X does not deviate a lot from its mean too often.

The following elementary tool from probability theory achieves just that. As is customary, we write

μ

$$\mu := E(X)$$

and define $\sigma \geqslant 0$ by setting

σ^2

$$\sigma^2 := E\big((X-\mu)^2\big)\,.$$

This quantity σ^2 is called the *variance* or *second moment* of X; by definition, it is a (quadratic) measure of how much X deviates from its mean. Since E is linear, the defining term for σ^2 expands to

$$\sigma^2 = E(X^2 - 2\mu X + \mu^2) = E(X^2) - \mu^2.$$

---

[6] For some p between n^{-1} and $(\log n)n^{-1}$, for example, almost every $G \in \mathcal{G}(n,p)$ has an isolated vertex (and hence no spanning tree), but its expected number of spanning trees tends to infinity with n! See the Exercise 13 for details.

(11.1.4)
(11.1.5)
X

Proof. Let $X(G)$ denote the number of subgraphs of G isomorphic to H. Given $n \in \mathbb{N}$, let $\mathcal{H}$ denote the set of all graphs isomorphic to H whose vertices lie in $\{0,\dots,n-1\}$, the vertex set of the graphs $G \in \mathcal{G}(n,p)$:

$\mathcal{H}$

$$\mathcal{H} := \bigl\{ H' \mid H' \simeq H,\ V(H') \subseteq \{0,\dots,n-1\} \bigr\}.$$

Given $H' \in \mathcal{H}$ and $G \in \mathcal{G}(n,p)$, we shall write $H' \subseteq G$ to express that H' itself—not just an isomorphic copy of H'—is a subgraph of G.

h By h we denote the number of isomorphic copies of H on a fixed k-set; clearly, $h \leqslant k!$. As there are $\binom{n}{k}$ possible vertex sets for the graphs in $\mathcal{H}$, we thus have

$$|\mathcal{H}| = \binom{n}{k} h \leqslant \binom{n}{k} k! \leqslant n^k. \tag{1}$$

p, γ Given $p = p(n)$, we set $\gamma := p/t$; then

$$p = \gamma n^{-k/\ell}. \tag{2}$$

We have to show that almost no $G \in \mathcal{G}(n,p)$ lies in $\mathcal{P}_H$ if $\gamma \to 0$ as $n \to \infty$, and that almost all $G \in \mathcal{G}(n,p)$ lie in $\mathcal{P}_H$ if $\gamma \to \infty$ as $n \to \infty$.

For the first part of the proof, we find an upper bound for $E(X)$, the expected number of subgraphs of G isomorphic to H. As in the proof of Lemma 11.1.5, double counting gives

$$E(X) = \sum_{H' \in \mathcal{H}} P[\, H' \subseteq G \,]. \tag{3}$$

For every fixed $H' \in \mathcal{H}$, we have

$$P[\, H' \subseteq G \,] \; = \; p^{\ell}, \tag{4}$$

because $\|H\| = \ell$. Hence,

$$E(X) \underset{(3,4)}{=} |\mathcal{H}| p^{\ell} \underset{(1,2)}{\leqslant} n^k (\gamma n^{-k/\ell})^{\ell} = \gamma^{\ell}. \tag{5}$$

Thus if $\gamma \to 0$ as $n \to \infty$, then

$$P[\, G \in \mathcal{P}_H \,] \; = \; P[\, X \geqslant 1 \,] \; \leqslant \; E(X) \; \leqslant \; \gamma^{\ell} \xrightarrow[n \to \infty]{} 0$$

by Markov's inequality (11.1.4), so almost no $G \in \mathcal{G}(n,p)$ lies in $\mathcal{P}_H$.

$$\begin{aligned}
&= \textstyle\sum_0 P\,[\,H' \subseteq G\,] \cdot P\,[\,H'' \subseteq G\,] \\
&\leqslant \sum_{(H',H'')\in\mathcal{H}^2} P\,[\,H' \subseteq G\,] \cdot P\,[\,H'' \subseteq G\,] \\
&= \Big(\sum_{H'\in\mathcal{H}} P\,[\,H' \subseteq G\,]\Big) \cdot \Big(\sum_{H''\in\mathcal{H}} P\,[\,H'' \subseteq G\,]\Big) \\
&\underset{(3)}{=} \mu^2.
\end{aligned} \tag{9}$$

Let us now estimate A_i for $i \geqslant 1$. Writing $\sum'$ for $\sum_{H'\in\mathcal{H}}$ and $\sum''$ for $\sum_{H''\in\mathcal{H}}$, we note that $\sum_i$ can be written as $\sum'\sum''_{|H'\cap H''|=i}$. For fixed H' (corresponding to the first sum $\sum'$), the second sum ranges over

$\sum'$

$$\binom{k}{i}\binom{n-k}{k-i}h$$

summands: the number of graphs $H'' \in \mathcal{H}$ with $|H'' \cap H'| = i$. Hence, for all $i \geqslant 1$ and suitable constants c_1, c_2 independent of n,

$$\begin{aligned}
A_i &= \textstyle\sum_i P\,[\,H' \cup H'' \subseteq G\,] \\
&\underset{(8)}{\leqslant} \textstyle\sum' \binom{k}{i}\binom{n-k}{k-i} h\, p^{2\ell} p^{-i\ell/k} \\
&\underset{(2)}{=} |\mathcal{H}| \binom{k}{i}\binom{n-k}{k-i} h\, p^{2\ell} \big(\gamma\, n^{-k/\ell}\big)^{-i\ell/k} \\
&\leqslant |\mathcal{H}|\, p^{\ell} c_1\, n^{k-i}\, h\, p^{\ell} \gamma^{-i\ell/k}\, n^{i} \\
&\underset{(5)}{=} \mu\, c_1 n^k h\, p^{\ell} \gamma^{-i\ell/k} \\
&\underset{(6)}{\leqslant} \mu\, c_2 \binom{n}{k} h\, p^{\ell} \gamma^{-i\ell/k} \\
&\underset{(1,5)}{=} \mu^2 c_2 \gamma^{-i\ell/k} \\
&\leqslant \mu^2 c_2 \gamma^{-\ell/k}
\end{aligned}$$

if $\gamma \geqslant 1$. By definition of the A_i, this implies with $c_3 := k c_2$ that

$$E(X^2)/\mu^2 \underset{(7)}{=} \Big(A_0/\mu^2 + \sum_{i=1}^{k} A_i/\mu^2\Big) \underset{(9)}{\leqslant} 1 + c_3 \gamma^{-\ell/k}$$

and hence

$$\frac{\sigma^2}{\mu^2} = \frac{E(X^2) - \mu^2}{\mu^2} \leqslant c_3 \gamma^{-\ell/k} \underset{\gamma\to\infty}{\longrightarrow} 0\,.$$

6. Show that if almost all $G \in \mathcal{G}(n,p)$ have a graph property $\mathcal{P}_1$ and almost all $G \in \mathcal{G}(n,p)$ have a graph property $\mathcal{P}_2$, then almost all $G \in \mathcal{G}(n,p)$ have both properties, i.e. have the property $\mathcal{P}_1 \cap \mathcal{P}_2$.

7.$^-$ Show that, for constant $p \in (0,1)$, almost every graph in $\mathcal{G}(n,p)$ has diameter 2.

8. Show that, for constant $p \in (0,1)$, almost no graph in $\mathcal{G}(n,p)$ has a separating complete subgraph.

9. Derive Proposition 11.3.1 from Lemma 11.3.2.

10.$^+$ (i) Show that with probability 1 an infinite random graph $G \in \mathcal{G}(\aleph_0, p)$ has all the properties $\mathcal{P}_{i,j}$ $(i,j \in \mathbb{N})$.

(ii) Show that any two (infinite) graphs having all the properties $\mathcal{P}_{i,j}$ are isomorphic.

(Thus, up to isomorphism, there is only one countably infinite random graph.)

11. Let $\epsilon > 0$ and $p = p(n) > 0$, and let $r \geqslant (1+\epsilon)(2 \ln n)/p$ be an integer-valued function of n. Show that almost no graph in $\mathcal{G}(n,p)$ contains r independent vertices.

12. Show that for every graph H there exists a function $p = p(n)$ such that $\lim_{n\to\infty} p(n) = 0$ but almost every $G \in \mathcal{G}(n,p)$ contains an induced copy of H.

13.$^+$ (i) Show that, for every $0 < \epsilon \leqslant 1$ and $p = (1-\epsilon)(\ln n)n^{-1}$, almost every $G \in \mathcal{G}(n,p)$ has an isolated vertex.

(ii) Find a probability $p = p(n)$ such that almost every $G \in \mathcal{G}(n,p)$ is disconnected but the expected number of spanning trees of G tends to infinity as $n \to \infty$.

(Hint for (ii): A theorem of Cayley states that K^n has exactly n^{n-2} spanning trees.)

14.$^+$ Given $r \in \mathbb{N}$, find a $c > 0$ such that, for $p = cn^{-1}$, almost every $G \in \mathcal{G}(n,p)$ has a K^r minor. Can c be chosen independently of r?

15. Find an increasing graph property without a threshold function, and a property that is not increasing but has a threshold function.

16.$^-$ Let H be a graph of order k, and let h denote the number of graphs isomorphic to H on some fixed set of k elements. Show that $h \leqslant k!$. For which graphs H does equality hold?

17.$^-$ For every $k \geqslant 1$, find a threshold function for $\{\, G \mid \Delta(G) \geqslant k \,\}$.

18.$^-$ Given $d \in \mathbb{N}$, is there a threshold function for the property of containing a d-dimensional cube (see Ex. 2, Ch. 1)? If so, which; if not, why not?

19. Show that $t(n) = n^{-1}$ is also a threshold function for the property of containing *any* cycle.

20. Does the property of containing any tree of order k (for $k \geqslant 2$ fixed) have a threshold function? If so, which?

There is another way of defining a random graph G, which is just as natural and common as the model we considered. Rather than choosing the edges of G independently, we choose the entire graph G uniformly at random from among all the graphs on $\{0, \ldots, n-1\}$ that have exactly $M = M(n)$ edges: then each of these graphs occurs with the same probability of $\binom{N}{M}$, where $N := \binom{n}{2}$. Just as we studied the likely properties of the graphs in $\mathcal{G}(n,p)$ for different functions $p = p(n)$, we may investigate how the properties of G in the other model depend on the function $M(n)$. If M is close to pN, the expected number of edges of a graph in $\mathcal{G}(n,p)$, then the two models behave very similarly. It is then largely a matter of convenience which of them to consider; see Bollobás for details.

In order to study threshold phenomena in more detail, one often considers the following *random graph process*: starting with a $\overline{K^n}$ as stage zero, one chooses additional edges one by one (uniformly at random) until the graph is complete. This is a simple example of a Markov chain, whose Mth stage corresponds to the 'uniform' random graph model described above. A survey about threshold phenomena in this setting is given by T. Łuczak, The phase transition in a random graph, in (D. Miklós, V.T. Sós & T. Szőnyi, eds.) *Paul Erdős is 80*, Vol. 2, Proc. Colloq. Math. Soc. János Bolyai (1996).

good pair

good/bad sequence

there are indices $i < j$ such that $x_i \leqslant x_j$. Then (x_i, x_j) is a *good pair* of this sequence. A sequence containing a good pair is a *good sequence*; thus, a quasi-ordering on X is a well-quasi-ordering if and only if every infinite sequence in X is good. An infinite sequence is *bad* if it is not good.

Proposition 12.1.1. *A quasi-ordering $\leqslant$ on X is a well-quasi-ordering if and only if X contains neither an infinite antichain nor an infinite strictly decreasing sequence $x_0 > x_1 > \ldots$.*

(9.1.2) *Proof.* The forward implication is trivial. Conversely, let $x_0, x_1, \ldots$ be any infinite sequence in X. Let K be the complete graph on $\mathbb{N} = \{0, 1, \ldots\}$. Colour the edges ij $(i < j)$ of K with three colours: green if $x_i \leqslant x_j$, red if $x_i > x_j$, and amber if x_i, x_j are incomparable. By Ramsey's theorem (9.1.2), K has an infinite induced subgraph whose edges all have the same colour. If there is neither an infinite antichain nor an infinite strictly decreasing sequence in X, then this colour must be green, i.e. $x_0, x_1, \ldots$ has an infinite subsequence in which every pair is good. In particular, the sequence $x_0, x_1, \ldots$ is good. □

In the proof of Proposition 12.1.1, we showed more than was needed: rather than finding a single good pair in $x_0, x_1, \ldots$, we found an infinite increasing subsequence. We have thus shown the following:

Corollary 12.1.2. *If X is well-quasi-ordered, then every infinite sequence in X has an infinite increasing subsequence.* □

The following lemma, and the idea of its proof, are fundamental to the theory of well-quasi-ordering. Let $\leqslant$ be a quasi-ordering on a set X. For finite subsets $A, B \subseteq X$, write $A \leqslant B$ if there is an injective mapping $f\colon A \to B$ such that $a \leqslant f(a)$ for all $a \in A$. This naturally extends $\leqslant$ to a quasi-ordering on $[X]^{<\omega}$, the set of all finite subsets of X.

$\leqslant$

$[X]^{<\omega}$

[12.2.1] **Lemma 12.1.3.** *If X is well-quasi-ordered by $\leqslant$, then so is $[X]^{<\omega}$.*

Proof. Suppose that $\leqslant$ is a well-quasi-ordering on X but not on $[X]^{<\omega}$. We start by constructing a bad sequence $(A_n)_{n\in\mathbb{N}}$ in $[X]^{<\omega}$, as follows. Given $n \in \mathbb{N}$, assume inductively that A_i has been defined for every $i < n$, and that there exists a bad sequence in $[X]^{<\omega}$ starting with $A_0, \ldots, A_{n-1}$. (This is clearly true for $n = 0$: by assumption, $[X]^{<\omega}$ contains a bad sequence, and this has the empty sequence as an initial segment.) Choose $A_n \in [X]^{<\omega}$ so that some bad sequence in $[X]^{<\omega}$ starts with $A_0, \ldots, A_n$ and $|A_n|$ is as small as possible.

Clearly, $(A_n)_{n\in\mathbb{N}}$ is a bad sequence in $[X]^{<\omega}$; in particular, $A_n \neq \emptyset$ for all n. For each n pick an element $a_n \in A_n$ and set $B_n := A_n \setminus \{a_n\}$.

(12.1.3) **Proof of Theorem 12.2.1.** We show that the rooted trees are well-quasi-ordered by the relation $\leqslant$ defined above; this clearly implies the theorem.

Suppose not. To derive a contradiction, we proceed as in the proof of Lemma 12.1.3. Given $n \in \mathbb{N}$, assume inductively that we have chosen a sequence $T_0, \dots, T_{n-1}$ of rooted trees such that some bad sequence of T_n rooted trees starts with this sequence. Choose as T_n a minimum-order rooted tree such that some bad sequence starts with $T_0, \dots, T_n$. For each r_n $n \in \mathbb{N}$, denote the root of T_n by r_n.

A_n Clearly, $(T_n)_{n \in \mathbb{N}}$ is a bad sequence. For each n, let A_n denote the set of components of $T_n - r_n$, made into rooted trees by choosing the neighbours of r_n as their roots. Note that the tree-order of these trees A is that induced by T_n. Let us prove that the set $A := \bigcup_{n \in \mathbb{N}} A_n$ of all these trees is well-quasi-ordered.

T^k Let $(T^k)_{k \in \mathbb{N}}$ be any sequence of trees in A. For every $k \in \mathbb{N}$ choose $n(k)$ an $n = n(k)$ such that $T^k \in A_n$. Pick a k with smallest $n(k)$. Then

$$T_0, \dots, T_{n(k)-1}, T^k, T^{k+1}, \dots$$

is a good sequence, by the minimal choice of $T_{n(k)}$ and $T^k \subsetneqq T_{n(k)}$. Let (T, T') be a good pair of this sequence. Since $(T_n)_{n \in \mathbb{N}}$ is bad, T cannot be among the first $n(k)$ members $T_0, \dots, T_{n(k)-1}$ of our sequence: then T' would be some T^i with $i \geqslant k$, i.e.

$$T \leqslant T' = T^i \leqslant T_{n(i)}\,;$$

since $n(k) \leqslant n(i)$ by the choice of k, this would make $(T, T_{n(i)})$ a good pair in the bad sequence $(T_n)_{n \in \mathbb{N}}$. Hence (T, T') is a good pair also in $(T^k)_{k \in \mathbb{N}}$, completing the proof that A is well-quasi-ordered.

By Lemma 12.1.3,[1] the sequence $(A_n)_{n \in \mathbb{N}}$ in $[A]^{<\omega}$ has a good pair i, j (A_i, A_j); let $f\colon A_i \to A_j$ be an injection such that $T \leqslant f(T)$ for all $T \in A_i$. We now extend the union of the embeddings $T \to f(T)$ to a map φ from $V(T_i)$ to $V(T_j)$ by letting $\varphi(r_i) := r_j$. This map φ preserves the tree-order of T_i, and it defines an embedding to show that $T_i \leqslant T_j$, since the edges $r_i r \in T_i$ map naturally to the paths $r_j T_j \varphi(r)$. Hence (T_i, T_j) is a good pair in our original bad sequence of rooted trees, a contradiction.

□

---

[1] Any readers worried that we might need the lemma for sequences or multisets rather than just sets here, please note that isomorphic elements of A_n are not identified: we always have $|A_n| = d(r_n)$.

Lemma 12.3.1. *Let t_1t_2 be any edge of T and let T_1, T_2 be the components of $T - t_1t_2$, with $t_1 \in T_1$ and $t_2 \in T_2$. Then $V_{t_1} \cap V_{t_2}$ separates $U_1 := \bigcup_{t \in T_1} V_t$ from $U_2 := \bigcup_{t \in T_2} V_t$ in G (Fig. 12.3.2).*

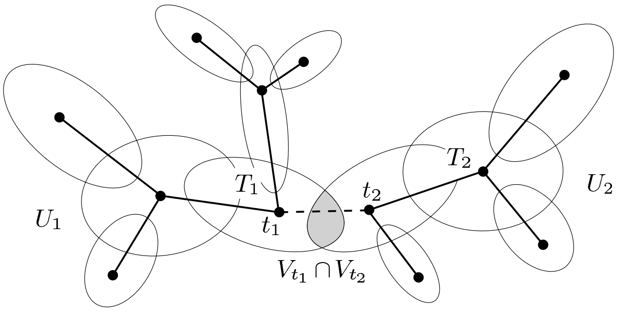

Fig. 12.3.2. $V_{t_1} \cap V_{t_2}$ separates U_1 from U_2 in G

Proof. Both t_1 and t_2 lie on every t–t' path in T with $t \in T_1$ and $t' \in T_2$. Therefore $U_1 \cap U_2 \subseteq V_{t_1} \cap V_{t_2}$ by (T3), so all we have to show is that G has no edge u_1u_2 with $u_1 \in U_1 \setminus U_2$ and $u_2 \in U_2 \setminus U_1$. If u_1u_2 is such an edge, then by (T2) there is a $t \in T$ with $u_1, u_2 \in V_t$. By the choice of u_1 and u_2 we have neither $t \in T_2$ nor $t \in T_1$, a contradiction. □

Note that tree-decompositions are passed on to subgraphs:

[12.4.2] **Lemma 12.3.2.** *For every $H \subseteq G$, the pair $\big(T, (V_t \cap V(H))_{t \in T}\big)$ is a tree-decomposition of H.* □

Similarly for contractions:

Lemma 12.3.3. *Suppose that G is an MH with branch sets U_h, $h \in V(H)$. Let $f\colon V(G) \to V(H)$ be the map assigning to each vertex of G the index of the branch set containing it. For all $t \in T$ let $W_t := \{ f(v) \mid v \in V_t \}$, and put $\mathcal{W} := (W_t)_{t \in T}$. Then $(T, \mathcal{W})$ is a tree-decomposition of H.*

Proof. The assertions (T1) and (T2) for $(T, \mathcal{W})$ follow immediately from the corresponding assertions for $(T, \mathcal{V})$. Now let $t_1, t_2, t_3 \in T$ be as in (T3), and consider a vertex $h \in W_{t_1} \cap W_{t_3}$ of H; we show that $h \in W_{t_2}$. By definition of W_{t_1} and W_{t_3}, there are vertices $v_1 \in V_{t_1} \cap U_h$ and $v_3 \in V_{t_3} \cap U_h$. Since U_h is connected in G and V_{t_2} separates v_1 from v_3 in G by Lemma 12.3.1, V_{t_2} has a vertex in U_h. By definition of W_{t_2}, this implies $h \in W_{t_2}$. □

Here is another useful consequence of Lemma 12.3.1:

In order to make use of Theorem 12.3.7 for a proof of the general minor theorem, we should be able to say something about the graphs it does not cover, i.e. to deduce some information about a graph from the assumption that its tree-width is large. Our next theorem achieves just that: it identifies a canonical obstruction to small tree-width, a structural phenomenon that occurs in a graph if and only if its tree-width is large.

touch Let us say that two subsets of $V(G)$ *touch* if they have a vertex in common or G contains an edge between them. A set of mutually touching *bramble* connected vertex sets in G is a *bramble*. Extending our terminology of *cover* Chapter 2.1, we say that a subset of $V(G)$ *covers* (or is a *cover* of) a bramble $\mathcal{B}$ if it meets every element of $\mathcal{B}$. The least number of vertices *order* covering a bramble is the *order* of that bramble.

The following simple observation will be useful:

Lemma 12.3.8. *Any set of vertices separating two covers of a bramble also covers that bramble.*

Proof. Since each set in the bramble is connected and meets both of the covers, it also meets any set separating these covers. □

A typical example of a bramble is the set of crosses in a grid. The *grid* $k \times k$ *grid* is the graph on $\{1,\dots,k\}^2$ with the edge set

$$\{\,(i,j)(i',j') : |i-i'|+|j-j'| = 1\,\}.$$

The *crosses* of this grid are the k^2 sets

$$C_{ij} := \{\,(i,\ell) \mid \ell = 1,\dots,k\,\} \cup \{\,(\ell,j) \mid \ell = 1,\dots,k\,\}.$$

Thus, the cross C_{ij} is the union of the grid's ith column and its jth row. Clearly, the crosses of the $k \times k$ grid form a bramble of order k: they are covered by any row or column, while any set of fewer than k vertices misses both a row and a column, and hence a cross.

The following result is sometimes called the *tree-width duality theorem*:

Theorem 12.3.9. (Seymour & Thomas 1993)
Let $k \geqslant 0$ be an integer. A graph has tree-width $\geqslant k$ if and only if it contains a bramble of order $> k$.

(3.3.1) *Proof.* For the backward implication, let $\mathcal{B}$ be any bramble in a graph G. We show that every tree-decomposition $(T,(V_t)_{t\in T})$ of G has a part that meets every set in $\mathcal{B}$.

As in the proof of Lemma 12.3.4 we start by orienting the edges t_1t_2 of T. If $X := V_{t_1} \cap V_{t_2}$ meets every $B \in \mathcal{B}$, we are done. If not, then

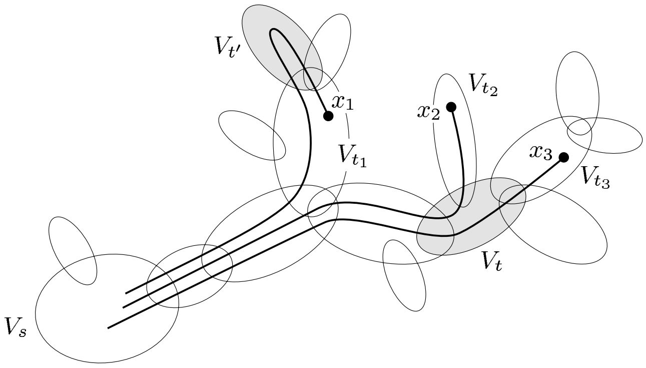

Fig. 12.3.3. W_t contains x_2 and x_3 but not x_1; $W_{t'}$ contains no x_i

x_i have been added to some of the parts. Despite these additions, we still have $|W_t| \leqslant |V_t|$ for all t: for each $x_i \in W_t \smallsetminus V_t$ we have $t \in sTt_i$, so V_t contains some other vertex of P_i (Lemma 12.3.1); that vertex does not lie in W_t, because P_i meets H only in x_i. Moreover, $(T, (W_t)_{t\in T})$ clearly satisfies (T3), because each x_i is added to every part along some path in T, so it is again a tree-decomposition.

As $W_s = X$, all that is left to show for $(*)$ is that this decomposition is $\mathcal{B}$-admissible. Consider any W_t of order $> k$. Then W_t meets C, because $|X| = \ell \leqslant k$. Since $(T, (V_t)_{t\in T})$ is $\mathcal{B}'$-admissible and $|V_t| \geqslant |W_t| > k$, we know that V_t fails to meet some $B \in \mathcal{B}$; let us show that W_t does not meet this B either. If it does, it must do so in some $x_i \in W_t \smallsetminus V_t$. Then B is a connected set meeting both V_s and V_{t_i} but not V_t. As $t \in sTt_i$ by definition of W_t, this contradicts Lemma 12.3.1. □

Often, Theorem 12.3.9 is stated in terms of the *bramble number* of a graph, the largest order of any bramble in it. The theorem then says that the tree-width of a graph is exactly one less than its bramble number (Exercise 15).

How useful even the easy backward direction of Theorem 12.3.9 can be is exemplified once more by our example of the crosses bramble in the $k \times k$ grid: this bramble has order k, so by the theorem the $k \times k$ grid has tree-width at least $k-1$. (Try to show this without the theorem!)

In fact, the $k \times k$ grid has tree-width k (Exercise 16). But more important than its precise value is the fact that the tree-width of grids tends to infinity with their size. For as we shall see, large grid minors pose another canonical obstruction to small tree-width: not only do large grids (and hence all graphs containing large grids as minors; cf. Proposition 12.3.6) have large tree-width, but conversely every graph of large tree-width has a large grid minor (Theorem 12.4.4).

Yet another canonical obstruction to small tree-width is described in Exercise 30.

Wagner graph W, and similarly for graphs without K^4 minors (see Proposition 12.4.2).

Tree-decompositions may thus lead to intuitive structural characterizations of graph properties. A particularly simple example is the following characterization of chordal graphs:

[12.4.2] **Proposition 12.3.11.** *G is chordal if and only if G has a tree-decomposition into complete parts.*

(5.5.1) *Proof.* We apply induction on $|G|$. We first assume that G has a tree-decomposition $(T, \mathcal{V})$ such that $G[V_t]$ is complete for every $t \in T$; let us choose $(T, \mathcal{V})$ with $|T|$ minimal. If $|T| \leqslant 1$, then G is complete and hence chordal. So let $t_1t_2 \in T$ be an edge, and for $i = 1, 2$ define T_i and $G_i := G[U_i]$ as in Lemma 12.3.1. Then $G = G_1 \cup G_2$ by (T1) and (T2), and $V(G_1 \cap G_2) = V_{t_1} \cap V_{t_2}$ by the lemma; thus, $G_1 \cap G_2$ is complete. Since $(T_i, (V_t)_{t \in T_i})$ is a tree-decomposition of G_i into complete parts, both G_i are chordal by the induction hypothesis. (By the choice of $(T, \mathcal{V})$, neither G_i is a subgraph of $G[V_{t_1} \cap V_{t_2}] = G_1 \cap G_2$, so both G_i are indeed smaller than G.) Since $G_1 \cap G_2$ is complete, any induced cycle in G lies in G_1 or in G_2 and hence has a chord, so G too is chordal.

Conversely, assume that G is chordal. If G is complete, there is nothing to show. If not then, by Proposition 5.5.1, G is the union of smaller chordal graphs G_1, G_2 with $G_1 \cap G_2$ complete. By the induction hypothesis, G_1 and G_2 have tree-decompositions $(T_1, \mathcal{V}_1)$ and $(T_2, \mathcal{V}_2)$ into complete parts. By Lemma 12.3.5, $G_1 \cap G_2$ lies inside one of those parts in each case, say with indices $t_1 \in T_1$ and $t_2 \in T_2$. As one easily checks, $((T_1 \cup T_2) + t_1t_2, \mathcal{V}_1 \cup \mathcal{V}_2)$ is a tree-decomposition of G into complete parts. $\square$

Corollary 12.3.12. $\mathrm{tw}(G) = \min \left\{ \omega(H) - 1 \mid G \subseteq H;\ H \text{ chordal} \right\}$.

Proof. By Lemma 12.3.5 and Proposition 12.3.11, each of the graphs H considered for the minimum has a tree-decomposition of width $\omega(H) - 1$. Every such tree-decomposition induced one of G by Lemma 12.3.2, so $\mathrm{tw}(G) \leqslant \omega(H) - 1$ for every H.

Conversely, let us construct an H as above with $\omega(H) - 1 \leqslant \mathrm{tw}(G)$. Let $(T, \mathcal{V})$ be a tree-decomposition of G of width $\mathrm{tw}(G)$. For every $t \in T$ let K_t denote the complete graph on V_t, and put $H := \bigcup_{t \in T} K_t$. Clearly, $(T, \mathcal{V})$ is also a tree-decomposition of H. By Proposition 12.3.11, H is chordal, and by Lemma 12.3.5, $\omega(H) - 1$ is at most the width of $(T, \mathcal{V})$, i.e. at most $\mathrm{tw}(G)$. $\square$

A question converse to the above is to ask for which H (other than K^3 and K^4) the tree-width of the graphs in $\mathrm{Forb}_{\preccurlyeq}(H)$ is bounded. Interestingly, it is not difficult to show that any such H must be planar. (4.4.6) Indeed, as all grids and their minors are planar (why?), every class $\mathrm{Forb}_{\preccurlyeq}(H)$ with non-planar H contains all grids; yet as we saw after Theorem 12.3.9, the grids have unbounded tree-width.

The following deep and surprising theorem says that, conversely, the tree-width of the graphs in $\mathrm{Forb}_{\preccurlyeq}(H)$ is bounded for every planar H:

Theorem 12.4.3. (Robertson & Seymour 1986)
Given a graph H, the graphs without an H minor have bounded tree-width if and only if H is planar.

The rest of this section is devoted to the proof of Theorem 12.4.3.

To prove Theorem 12.4.3 we have to show that forbidding any planar graph H as a minor bounds the tree-width of a graph. In fact, we only have to show this for the special cases when H is a grid, because every planar graph is a minor of some grid. (To see this, take a drawing of the graph, fatten its vertices and edges, and superimpose a sufficiently fine plane grid.) It thus suffices to show the following:

Theorem 12.4.4. (Robertson & Seymour 1986)
For every integer r there is an integer k such that every graph of tree-width at least k has an $r \times r$ grid minor.

Our proof of Theorem 12.4.4, which is much shorter than the original proof, proceeds as follows. Let r be given, and let G be any graph of large enough tree-width (depending on r). We first show that G contains a large family $\mathcal{A} = \{ A_1, \dots, A_m \}$ of disjoint connected vertex sets such that each pair $A_i, A_j \in \mathcal{A}$ can be linked in G by a family $\mathcal{P}_{ij}$ of many disjoint A_i–A_j paths avoiding all the other sets in $\mathcal{A}$. We then consider all the pairs $(\mathcal{P}_{ij}, \mathcal{P}_{i'j'})$ of these path families. If we can find a pair among these such that many of the paths in $\mathcal{P}_{ij}$ meet many of the paths in $\mathcal{P}_{i'j'}$, we shall think of the paths in $\mathcal{P}_{ij}$ as horizontal and the paths in $\mathcal{P}_{i'j'}$ as vertical and extract a subdivision of an $r \times r$ grid from their union. (This will be the difficult part of the proof, because these paths will in general meet in a less orderly way than they do in a grid.) If not, then for every pair $(\mathcal{P}_{ij}, \mathcal{P}_{i'j'})$ many of the paths in $\mathcal{P}_{ij}$ avoid many of the paths in $\mathcal{P}_{i'j'}$. We can then select one path $P_{ij} \in \mathcal{P}_{ij}$ from each family so that these selected paths are pairwise disjoint. Contracting each of the connected sets $A \in \mathcal{A}$ will then give us a K^m minor in G, which contains the desired $r \times r$ grid if $m \geqslant r^2$.

externally k-connected To implement these ideas formally, we need a few definitions. Let us call a set $X \subseteq V(G)$ *externally k-connected* in G if $|X| \geq k$ and for all disjoint subsets $Y, Z \subseteq X$ with $|Y| = |Z| \leq k$ there are $|Y|$ disjoint

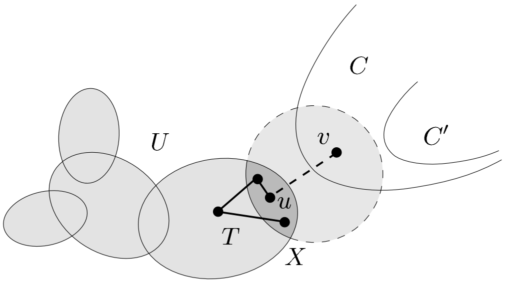

Fig. 12.4.1. Extending U and $\mathcal{D}$ when $|X| < h$

check that $(G - C', \tilde{C}')$ is again a premesh of order $\leq h$, contrary to the maximality of U.

Thus $|X| = h$, so by assumption our premesh $(G - C, \tilde{C})$ cannot be a k-mesh; let $Y, Z \subseteq X$ be sets to witness this. Let $\mathcal{P}$ be a set of as many disjoint Y–Z paths in $H := G[V(C) \cup Y \cup Z] - E(G[Y \cup Z])$ as possible. Since all these paths are 'external' to X in $\tilde{C}$, we have $k' := |\mathcal{P}| < |Y| = |Z| \leqslant k$ by the choice of Y and Z. By Menger's theorem (3.3.1), Y and Z are separated in H by a set S of k' vertices. Clearly, S has exactly one vertex on each path in $\mathcal{P}$; we denote the path containing the vertex $s \in S$ by P_s (Fig. 12.4.2).

Y, Z

k'

S

P_s

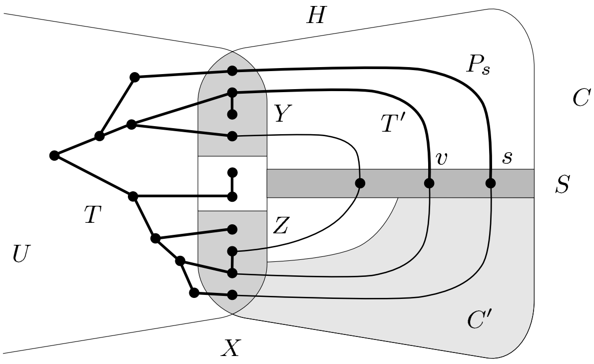

Fig. 12.4.2. S separates Y from Z in H

Let $X' := X \cup S$ and $U' := U \cup S$, and let $\mathcal{D}'$ be the tree-decomposition of $G[U']$ obtained from $\mathcal{D}$ by adding X' as a new part. Clearly, $|X'| \leq |X| + |S| \leq h + k - 1$. We show that U' contradicts the maximality of U.

Since $Y \cup Z \subseteq N(C)$ and $|S| < |Y| = |Z|$ we have $S \cap C \neq \emptyset$, so U' is larger than U. Let C' be a component of $G - U'$. If $C' \cap C = \emptyset$, we argue as earlier. So $C' \subseteq C$ and $N(C') \subseteq X'$. As before, C' has at least one neighbour v in $S \cap C$, since X cannot separate $C' \subseteq C$ from $S \cap C$. By definition of S, C' cannot have neighbours in both $Y \setminus S$

v

Lemma 12.4.9. *Let $d, r \geq 2$ be integers such that $d \geq r^{2r+2}$. Let G be a graph containing a set $\mathcal{H}$ of $r^2 - 1$ disjoint paths and a set $\mathcal{V} = \{ V_1, \ldots, V_d \}$ of d disjoint paths. Assume that every path in $\mathcal{V}$ meets every path in $\mathcal{H}$, and that each path $H \in \mathcal{H}$ consists of d consecutive (vertex-disjoint) segments such that V_i meets H only in its ith segment, for every $i = 1, \ldots, d$ (Fig. 12.4.3). Then G has an $r \times r$ grid minor.*

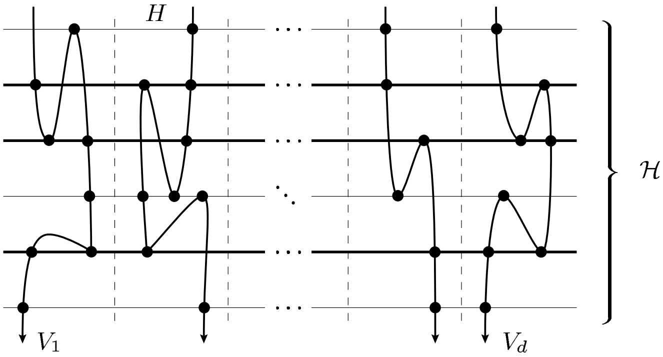

Fig. 12.4.3. Paths intersecting as in Lemma 12.4.9

T_i

$H^1, \ldots, H^r$

I, i_k

$\mathcal{H}'$

Proof. For each $i = 1, \ldots, d$, consider the graph with vertex set $\mathcal{H}$ in which two paths are adjacent whenever V_i contains a subpath between them that meets no other path in $\mathcal{H}$. Since V_i meets every path in $\mathcal{H}$, this is a connected graph; let T_i be a spanning tree in it. Since $|\mathcal{H}| \geq r(r-1)$, Lemma 12.4.8 implies that each of these $d \geq r^2(r^2)^r$ trees T_i has a good r-tuple of vertices. Since there are no more than $(r^2)^r$ distinct r-tuples on $\mathcal{H}$, some r^2 of the trees T_i have a common good r-tuple $(H^1, \ldots, H^r)$. Let $I = \{ i_1, \ldots, i_{r^2} \}$ be the index set of these trees (with $i_j < i_k$ for $j < k$) and put $\mathcal{H}' := \{ H^1, \ldots, H^r \}$.

Here is an informal description of how we construct our $r \times r$ grid. Its 'horizontal' paths will be the paths $H^1, \ldots, H^r$. Its 'vertical' paths will be pieced together edge by edge, as follows. The $r-1$ edges of the first vertical path will come from the first $r-1$ trees T_i, trees with their index i among the first r elements of I. More precisely, its 'edge' between H^j and H^{j+1} will be the sequence of subpaths of V_{i_j} (together with some connecting horizontal bits taken from paths in $\mathcal{H} \setminus \mathcal{H}'$) induced by the edges of an H^j–H^{j+1} path in T_{i_j} that has no inner vertices in $\mathcal{H}'$; see Fig. 12.4.4. (This is why we need $(H^1, \ldots, H^r)$ to be a good r-tuple in every tree T_i.) Similarly, the jth edge of the second vertical path will come from an H^j–H^{j+1} path in $T_{i_{r+j}}$, and so on.[6] To merge these individual edges into r vertical paths, we then contract in each horizontal

---

[6] Although we need only $r-1$ edges for each vertical path, we reserve r rather than just $r-1$ of the paths V_i for each vertical path to make the indexing more lucid. The paths $V_{i_r}, V_{i_{2r}}, \ldots$ are left unused.

path the initial segment that meets the first r paths V_i with $i \in I$, then contract the segment that meets the following r paths V_i with $i \in I$, and so on.

H^j_k
$\hat{H}^j$
v^j_k

Formally, we proceed as follows. Consider all $j, k \in \{1, \dots, r\}$. (We shall think of the index j as counting the horizontal paths, and of the index k as counting the vertical paths of the grid to be constructed.) Let H^j_k be the minimal subpath of H^j that contains the ith segment of H^j for all i with $i_{(k-1)r} < i \leq i_{kr}$ (put $i_0 := 0$). Let $\hat{H}^j$ be obtained from H^j by first deleting any vertices following its i_{r^2}th segment and then contracting every subpath H^j_k to one vertex v^j_k. Thus, $\hat{H}^j = v^j_1 \dots v^j_r$.

V^j_k

Given $j \in \{1, \dots, r-1\}$ and $k \in \{1, \dots, r\}$, we have to define a path V^j_k that will form the subdivided 'vertical edge' $v^j_k v^{j+1}_k$. This path will consist of segments of the path V_i together with some otherwise unused segments of paths from $\mathcal{H} \setminus \mathcal{H}'$, for $i := i_{(k-1)r+j}$; recall that, by definition of $\hat{H}^j$ and $\hat{H}^{j+1}$, this V_i does indeed meet H^j and H^{j+1} precisely in vertices that were contracted into v^j_k and v^{j+1}_k, respectively. To define V^j_k, consider an H^j–H^{j+1} path $P = H_1 \dots H_t$ in T_i that has no inner vertices in $\mathcal{H}'$. (Thus, $H_1 = H^j$ and $H_t = H^{j+1}$.) Every edge $H_s H_{s+1}$ of P corresponds to an H_s–H_{s+1} subpath of V_i that has no inner vertex on any path in $\mathcal{H}$. Together with (parts of) the ith segments of $H_2, \dots, H_{t-1}$, these subpaths of V_i form an H^j–H^{j+1} path P' in G that has no inner vertices on any of the paths $H^1, \dots, H^r$ and meets no path from $\mathcal{H}$ outside its ith segment. Replacing the ends of P' on H^j and H^{j+1} with v^j_k and v^{j+1}_k, respectively, we obtain our desired path V^j_k forming the jth (subdivided) edge of the kth 'vertical' path of our grid. Since the paths P' are disjoint for different i and different pairs (j, k) give rise to different i, the paths V^j_k are disjoint except for possible common ends v^j_k. Moreover, they have no inner vertices on any of the paths $H^1, \dots, H^r$, because none of these H^j is an inner vertex of any of the paths $P \subseteq T_i$ used in the construction of V^j_k. □

Proof of Theorem 12.4.4. We are now ready to prove the following quantitative version of our theorem (which clearly implies it):

> *Let $r, m > 0$ be integers, and let G be a graph of tree-width at least $r^{4m^2(r+2)}$. Then G contains either the $r \times r$ grid or K^m as a minor.*

c, k

Since K^{r^2} contains the $r \times r$ grid as a subgraph we may assume that $2 \leq m \leq r^2$. Put $c := r^{4(r+2)}$, and let $k := c^{2\binom{m}{2}}$. Then $c \geqslant 2^{16}$ and hence $2m + 3 \leq c^m$, so G has tree-width at least

$$c^{m^2} = c^m k \geqslant (2m+3)k \geqslant (m+1)(2k-1) + k - 1\,,$$

(A, B)
T

enough for Lemma 12.4.5 to ensure that G contains a k-mesh (A, B) of order $(m+1)(2k-1)$. Let $T \subseteq A$ be a tree associated with the

Let us show that $\ell^* > 0$. Let $pq := \sigma^{-1}(0)$ and put $\mathcal{P}^0_{pq} := \mathcal{P}_{pq}$. To define $\mathcal{P}^0_{ij}$ for $\sigma(ij) > 0$ put $H_{ij} := \bigcup \mathcal{P}_{ij}$, let $F \subseteq E(H_{ij}) \setminus E(H^0_{pq})$ be maximal such that $(H^0_{pq} \cup H_{ij}) - F$ still contains k/c disjoint A_i–A_j paths, and let $\mathcal{P}^0_{ij}$ be such a set of paths. Since the vertices from $A_p \cup A_q$ have degree 1 in $H^0_{pq} \cup H_{ij}$ unless they also lie in $A_i \cup A_j$, these paths have no inner vertices in A. Our choices of $\mathcal{P}^0_{ij}$ therefore satisfy (i)–(v) for $\ell = 0$.

ℓ

pq

Having shown that $\ell^* > 0$, let us now consider $\ell := \ell^* - 1$. Thus, conditions (i)–(v) are satisfied for ℓ but cannot be satisfied for $\ell + 1$. Let $pq := \sigma^{-1}(\ell)$. If $\mathcal{P}^\ell_{pq}$ contains a path P that avoids a set $\mathcal{Q}_{ij}$ of some $|\mathcal{P}^\ell_{ij}|/c$ of the paths in $\mathcal{P}^\ell_{ij}$ for all ij with $\sigma(ij) > \ell$, then we can define $\mathcal{P}^{\ell+1}_{ij}$ for all ij as before (with a contradiction). Indeed, let $st := \sigma^{-1}(\ell+1)$ and put $\mathcal{P}^{\ell+1}_{st} := \mathcal{Q}_{st}$. For $\sigma(ij) > \ell+1$ write $H_{ij} := \bigcup \mathcal{Q}_{ij}$, let $F \subseteq E(H_{ij}) \setminus E(H^{\ell+1}_{st})$ be maximal such that $(H^{\ell+1}_{st} \cup H_{ij}) - F$ still contains at least $|\mathcal{P}^\ell_{ij}|/c^2$ disjoint A_i–A_j paths, and let $\mathcal{P}^{\ell+1}_{ij}$ be such a set of paths. Setting $\mathcal{P}^{\ell+1}_{pq} := \{ P \}$ and $\mathcal{P}^{\ell+1}_{ij} := \mathcal{P}^\ell_{ij} = \{ P^\ell_{ij} \}$ for $\sigma(ij) < \ell$ then gives us a family of sets $\mathcal{P}^{\ell+1}_{ij}$ that contradicts the maximality of ℓ^*.

$\mathcal{P}$

ij

Thus for every path $P \in \mathcal{P}^\ell_{pq}$ there exists a pair ij with $\sigma(ij) > \ell$ such that P avoids fewer than $|\mathcal{P}^\ell_{ij}|/c$ of the paths in $\mathcal{P}^\ell_{ij}$. For some $\lceil |\mathcal{P}^\ell_{pq}|/\binom{m}{2} \rceil$ of these P that pair ij will be the same; let $\mathcal{P}$ denote the set of those P, and keep ij fixed from now on. Note that $|\mathcal{P}| \geq |\mathcal{P}^\ell_{pq}|/\binom{m}{2} = c\,|\mathcal{P}^\ell_{ij}|/\binom{m}{2}$ by (iii) and (iv).

Let us use Lemma 12.4.7 to find sets $\mathcal{V} \subseteq \mathcal{P} \subseteq \mathcal{P}^\ell_{pq}$ and $\mathcal{H} \subseteq \mathcal{P}^\ell_{ij}$ such that

$$\begin{aligned} |\mathcal{V}| &\geqslant \tfrac{1}{2}|\mathcal{P}| \quad \Big(\geq \frac{c}{m^2}|\mathcal{P}^\ell_{ij}| \Big) \\ |\mathcal{H}| &= r^2 \end{aligned}$$

and every path in $\mathcal{V}$ meets every path in $\mathcal{H}$. We have to check that the bipartite graph with vertex sets $\mathcal{P}$ and $\mathcal{P}^\ell_{ij}$ in which $P \in \mathcal{P}$ is adjacent to $Q \in \mathcal{P}^\ell_{ij}$ whenever $P \cap Q = \emptyset$ does not have too many edges. Since every $P \in \mathcal{P}$ has fewer than $|\mathcal{P}^\ell_{ij}|/c$ neighbours (by definition of $\mathcal{P}$), this graph indeed has at most

$$\begin{aligned} |\mathcal{P}||\mathcal{P}^\ell_{ij}|/c &\leqslant |\mathcal{P}||\mathcal{P}^\ell_{ij}|/6r^2 \\ &\leqslant \lfloor |\mathcal{P}|/2 \rfloor\, |\mathcal{P}^\ell_{ij}|/2r^2 \\ &\leqslant \lfloor |\mathcal{P}|/2 \rfloor \big(|\mathcal{P}^\ell_{ij}|/r^2 - 1 \big) \\ &= \big(|\mathcal{P}| - \lceil |\mathcal{P}|/2 \rceil \big) \big(|\mathcal{P}^\ell_{ij}| - r^2 \big)/r^2 \end{aligned}$$

$\mathcal{V}, \mathcal{H}$

edges, as required. Hence, $\mathcal{V}$ and $\mathcal{H}$ exist as claimed.

Although all the ('vertical') paths in $\mathcal{V}$ meet all the ('horizontal') paths in $\mathcal{H}$, these paths do not necessarily intersect in such an orderly way as required for Lemma 12.4.9. In order to divide the paths from $\mathcal{H}$ into segments, and to select paths from $\mathcal{V}$ meeting them only in the

$n = 1, \ldots, d$ (and hence in particular that $P'_n \neq \emptyset$, ie. that $P_{n-1} \subset P_n$). Indeed V_n cannot meet P_{n-1}, because $P_{n-1} \cup V_n \cup (Q - Q_{n-1})$ would then contain an A_i–A_j path in $(H^\ell_{pq} \cup H^\ell_{ij}) - e_{n-1} - S_{n-1}$, and likewise (consider S_n) V_n cannot meet $P - P_n$. Thus for all $n = 1, \ldots, d$, the path V_n meets every path $P \in \mathcal{H} \setminus \{ Q \}$ precisely in its nth segment P'_n. Applying Lemma 12.4.9 to the path systems $\mathcal{H} \setminus \{ Q \}$ and $\{ V_1, \ldots, V_d \}$ now yields the desired grid minor. □

12.5 The graph minor theorem

Hereditary graph properties, those that are closed under taking minors, occur frequently in graph theory. Among the most natural examples are the properties of being embeddable in some fixed surface, such as planarity.

By Kuratowski's theorem, planarity can be expressed by forbidding the minors K^5 and $K_{3,3}$. This is a *good characterization* of planarity in the following sense. Suppose we wish to persuade someone that a certain graph is planar: this is easy (at least intuitively) if we can produce a drawing of the graph. But how do we persuade someone that a graph is non-planar? By Kuratowski's theorem, there is also an easy way to do that: we just have to exhibit an MK^5 or $MK_{3,3}$ in our graph, as an easily checked 'certificate' for non-planarity. Our simple Proposition 12.4.2 is another example of a good characterization: if a graph has tree width < 3, we can prove this by exhibiting a suitable tree-decomposition; if not, we can produce an MK^4 as evidence.

Theorems that characterize a hereditary property $\mathcal{P}$ by a set $\mathcal{H}$ of forbidden minors are doubtless among the most attractive results in graph theory. As we saw in the proof of Proposition 12.4.1, there is always some such characterization: that where $\mathcal{H}$ is the complement $\overline{\mathcal{P}}$ of $\mathcal{P}$. However, one naturally seeks to make $\mathcal{H}$ as small as possible. And as it turns out, there is indeed a unique smallest such set $\mathcal{H}$: the set

$$\mathcal{H}_\mathcal{P} := \{ H \mid H \text{ is } \preccurlyeq\text{-minimal in } \overline{\mathcal{P}} \}$$

satisfies $\mathcal{P} = \text{Forb}_\preccurlyeq(\mathcal{H})$ and is contained in every other such set $\mathcal{H}$.

Proposition 12.5.1. *$\mathcal{P} = \text{Forb}_\preccurlyeq(\mathcal{H}_\mathcal{P})$, and $\mathcal{H}_\mathcal{P} \subseteq \mathcal{H}$ for every set $\mathcal{H}$ with $\mathcal{P} = \text{Forb}_\preccurlyeq(\mathcal{H})$.* □

Clearly, the elements of $\mathcal{H}_\mathcal{P}$ are incomparable under the minor relation $\preccurlyeq$. Now the *graph minor theorem* of Robertson & Seymour says that any set of $\preccurlyeq$-incomparable graphs must be finite:

one of the surfaces $S \in \mathcal{S}$. (The 'nearly' hides a measure of disorderliness that depends on n but not on the graph to be embedded.) By a generalization of Theorem 12.3.7—and hence of Kruskal's theorem—it now suffices, essentially, to prove that the set of all the parts in these tree-decompositions is well-quasi-ordered: then the graphs decomposing into these parts are well-quasi-ordered, too. Since $\mathcal{S}$ is finite, every infinite sequence of such parts has an infinite subsequence whose members are all (nearly) embeddable in the same surface $S \in \mathcal{S}$. Thus all we have to show is that, given any closed surface S, all the graphs embeddable in S are well-quasi-ordered by the minor relation.

This is shown by induction on the genus of S (more precisely, on $2 - \chi(S)$, where $\chi(S)$ denotes the Euler characteristic of S) using the same approach as before: if $H_0, H_1, H_2, \ldots$ is an infinite sequence of graphs embeddable in S, we may assume that none of the graphs $H_1, H_2, \ldots$ contains H_0 as a minor. If $S = S^2$ we are back in the case that H_0 is planar, so the induction starts. For the induction step we now assume that $S \neq S^2$. Again, the exclusion of H_0 as a minor constrains the structure of the graphs $H_1, H_2, \ldots$, this time topologically: each H_i with $i \geqslant 1$ has an embedding in S which meets some non-contractible closed curve $C_i \subseteq S$ in no more than a bounded number of vertices (and no edges), say in $X_i \subseteq V(H_i)$. (The bound on $|X_i|$ depends on H_0, but not on H_i.) Cutting along C_i, and sewing a disc on to each of the one or two closed boundary curves arising from the cut, we obtain one or two new closed surfaces of larger Euler characteristic. If the cut produces only one new surface S_i, then our embedding of $H_i - X_i$ still counts as a near-embedding of H_i in S_i (since X_i is small). If this happens for infinitely many i, then infinitely many of the surfaces S_i are also the same, and the induction hypothesis gives us a good pair among the corresponding graphs H_i. On the other hand, if we get two surfaces S_i' and S_i'' for infinitely many i (without loss of generality the same two surfaces), then H_i decomposes accordingly into subgraphs H_i' and H_i'' embedded in these surfaces, with $V(H_i' \cap H_i'') = X_i$. The set of all these subgraphs taken together is again well-quasi-ordered by the induction hypothesis, and hence so are the pairs (H_i', H_i'') by Lemma 12.1.3. Using a sharpening of the lemma that takes into account not only the graphs H_i' and H_i'' themselves but also how X_i lies inside them, we finally obtain indices i, j not only with $H_i' \preccurlyeq H_j'$ and $H_i'' \preccurlyeq H_j''$, but also such that these minor embeddings extend to the desired minor embedding of H_i in H_j—completing the proof of the minor theorem.

In addition to its impact on 'pure' graph theory, the graph minor theorem has had far-reaching algorithmic consequences. Using their tree structure theorem for the graphs in $\text{Forb}_{\preccurlyeq}(K^n)$, Robertson & Seymour have shown that testing for any fixed minor is 'fast': for every graph H

4. Given a quasi-ordering $(X, \leqslant)$ and subsets $A, B \subseteq X$, write $A \leqslant' B$ if there exists an *order preserving* injection $f\colon A \to B$ with $a \leqslant f(a)$ for all $a \in A$. Does Lemma 12.1.3 still hold if the quasi-ordering considered for $[X]^{<\omega}$ is $\leqslant'$?

5.$^-$ Show that the relation $\leqslant$ between rooted trees defined in the text is indeed a quasi-ordering.

6. Show that the finite trees are not well-quasi-ordered by the subgraph relation.

7. The last step of the proof of Kruskal's theorem considers a 'topological' embedding of T_m in T_n that maps the root of T_m to the root of T_n. Suppose we assume inductively that the trees of A_m are embedded in the trees of A_n in the same way, with roots mapped to roots. We thus seem to obtain a proof that the finite rooted trees are well-quasi-ordered by the subgraph relation, even with roots mapped to roots. Where is the error?

8.$^+$ Show that the finite graphs are not well-quasi-ordered by the topological minor relation.

9.$^+$ Given $k \in \mathbb{N}$, is the class $\{\, G \mid G \not\supseteq P^k \,\}$ well-quasi-ordered by the subgraph relation?

10. Show that a graph has tree-width at most 1 if and only if it is a forest.

11. Let G be a graph, T a set, and $(V_t)_{t \in T}$ a family of subsets of $V(G)$ satisfying (T1) and (T2) from the definition of a tree-decomposition. Show that there exists a tree on T that makes (T3) true if and only if there exists an enumeration $t_1, \dots, t_n$ of T such that for every $k = 2, \dots, n$ there is a $j < k$ satisfying $V_{t_k} \cap \bigcup_{i<k} V_{t_i} \subseteq V_{t_j}$.

    (The new condition tends to be more convenient to check than (T3). It can help, for example, with the construction of a tree-decomposition into a given set of parts.)

12. Prove the following converse of Lemma 12.3.1: if $(T, \mathcal{V})$ satisfies condition (T1) and the statement of the lemma, then $(T, \mathcal{V})$ is a tree-decomposition of G.

13. Can the tree-width of a subdivision of a graph G be smaller than $\mathrm{tw}(G)$? Can it be larger?

14. Let $(T, (V_t)_{t \in T})$ be a tree-decomposition of a graph G. For each vertex $v \in G$, set $T_v := \{\, t \in T \mid v \in V_t \,\}$. Show that T_v is always connected in T. More generally, for which subsets $U \subseteq V(G)$ is the set $\{\, t \in T \mid V_t \cap U \neq \emptyset \,\}$ always connected in T (i.e. for all tree-decompositions)?

15.$^-$ Show that the tree-width of a graph is one less than its bramble number.

16. Apply Theorem 12.3.9 to show that the $k \times k$ grid has tree-width at least k, and find a tree-decomposition of width exactly k.

27. Let $\mathcal{P}$ be a hereditary graph property. Show that strengthening the notion of a minor (for example, to that of topological minor) increases the set of forbidden minors required to characterize $\mathcal{P}$.

28. Deduce from the minor theorem that every hereditary property can be expressed by forbidding finitely many topological minors. Is the same true for every property that is closed under taking topological minors?

29. Show that every horizontal path in the $k \times k$ grid is externally k-connected in that grid.

30.$^+$ Show that the tree-width of a graph is large if and only if it contains a large externally k-connected set of vertices, with k large. For example, show that graphs of tree-width $< k$ contain no externally $(k+1)$-connected set of $3k$ vertices, and that graphs containing no externally $(k+1)$-connected set of $3k$ vertices have tree-width $< 4k$.

31.$^+$ (continued)

Find an $\mathbb{N} \to \mathbb{N}^2$ function $k \mapsto (h, \ell)$ such that every graph with an externally ℓ-connected set of h vertices contains a bramble of order at least k. Deduce the weakening of Theorem 12.3.9 that, given k, every graph of large enough tree-width contains a bramble of order at least k.

32. Without using the minor theorem, show that the chromatic number of the graphs in any $\preccurlyeq$-antichain is bounded.

33. Seymour's *self-minor conjecture* asserts that 'every countably infinite graph is a proper minor of itself'. Make this assertion precise, and deduce the minor theorem from it.

34. Given an orientable surface S of genus g, find a lower bound in terms of g for the number of forbidden minors needed to characterize embeddability in S.

(Hint. The smallest genus of an orientable surface in which a given graph can be embedded is called the (orientable) *genus* of that graph. Use the theorem that the genus of a graph is equal to the sum of the genera of its blocks.)

Notes

Kruskal's theorem on the well-quasi-ordering of finite trees was first published in J.A. Kruskal, Well-quasi ordering, the tree theorem, and Vászonyi's conjecture, *Trans. Amer. Math. Soc.* **95** (1960), 210–225. Our proof is due to Nash-Williams, who introduced the versatile proof technique of choosing a 'minimal bad sequence'. This technique was also used in our proof of Higman's Lemma 12.1.3.

Nash-Williams generalized Kruskal's theorem to infinite graphs. This extension is much more difficult than the finite case; it is one of the deepest theorems in infinite graph theory. The general graph minor theorem becomes false for arbitrary infinite graphs, as shown by R. Thomas, A counterexample to 'Wagner's conjecture' for infinite graphs, *Math. Proc. Camb. Phil. Soc.* **103** the number of forbidden minors needed for an arbitrary closed surface is given in P.D. Seymour, A bound on the excluded minors for a surface, *J. Combin. Theory B* (to appear). B. Mohar, Embedding graphs in an arbitrary surface in linear time, *Proc. 28th Ann. ACM STOC* (Philadelphia 1996), 392–397, has developed a set of algorithms, one for each surface, that decide embeddability in that surface in linear time. As a corollary, Mohar obtains an independent and constructive proof of the 'generalized Kuratowski theorem', Corollary 12.5.4. Another independent and short proof of this corollary, which builds on Theorem 12.4.3 and Graph Minors IV but on no other papers of the Graph Minors series, was found by C. Thomassen, A simpler proof of the excluded minor theorem for higher surfaces, *J. Combin. Theory B* **70** (1997), 306–311. A survey of the classical forbidden minor theorems is given in Chapter 6.1 of R. Diestel, *Graph Decompositions*, Oxford University Press 1990. More recent developments are surveyed in R. Thomas, Recent excluded minor theorems, in (J.D. Lamb & D.A. Preece, eds) *Surveys in Combinatorics 1999*, Cambridge University Press 1999, 201–222.

For every graph X, Graph Minors XIII gives an explicit algorithm that decides in cubic time for every input graph G whether $X \preccurlyeq G$. The constants in the cubic polynomials bounding the running time of these algorithms depend on X but are constructively bounded from above. For an overview of the algorithmic implications of the Graph Minors series, see Johnson's NP-completeness column in *J. Algorithms* **8** (1987), 285–303.

The concept of a 'good characterization' of a graph property was first suggested by J. Edmonds, Minimum partition of a matroid into independent subsets, *J. Research of the National Bureau of Standards (B)* **69** (1965) 67–72. In the language of complexity theory, a characterization is *good* if it specifies two assertions about a graph such that, given any graph G, the first assertion holds for G if and only if the second fails, and such that each assertion, if true for G, provides a certificate for its truth. Thus every good characterization has the corollary that the decision problem corresponding to the property it characterizes lies in NP ∩ co-NP.

9. (i) Straightforward from the definitions.

   (ii) Prove $\kappa \geqslant n$ by induction on n: partition the n-dimensional cube into cubes of lower dimension, and show inductively that the deletion of $< n$ vertices leaves a connected subgraph.

10. For the first inequality, consider the endvertices of a set of $\lambda(G)$ edges whose deletion disconnects G. Use the definition of $\lambda(G)$ to show the second inequality.

11.$^-$ Try to find counterexamples for $k = 1$.

12. Rephrase (i) and (ii) as statements about the existence of two $\mathbb{N} \to \mathbb{N}$ functions. To show the equivalence, express each of these functions in terms of the other. Show that (iii) may hold even if (i) and (ii) do not, and strengthen (iii) to remedy this.

13.$^+$ Try to imitate the proof assuming $\varepsilon(G) \geqslant 2k$ instead of condition (ii). Why does this fail, and why does condition (ii) remedy the problem?

14. Show (i) $\Rightarrow$ (ii) $\Rightarrow$ (iii) $\Rightarrow$ (iv) $\Rightarrow$ (i) from the definitions of the relevant concepts.

15. Consider paths emanating from a vertex of maximum degree.

16. Theorem 1.5.1.

17. Induction.

18. The easiest solution is to apply induction on $|T|$. What kind of vertex of T will be best to delete in the induction step?

19. Induction on $|T|$ is a possibility, but not the only one.

20. Count the edges.

21. Show that if a graph contains any odd cycle at all it also contains an induced one.

22. Apply Proposition 1.2.2. Split the subgraph thus found into two sides so that every vertex has many neighbours on the opposite side.

23. Try to carry the proof for finite graphs over to the infinite case. Where does it fail?

24.$^-$ Use Proposition 1.9.2.

25. Why do *all* the cuts $E(v)$ generate the cut space? Will they still do so if we omit one of them? Or even two?

26. Start with the case that the graph considered is a cycle.

27. Induction on $|F \setminus E(T)|$ for given $F \in \mathcal{C}(G)$.

28. Induction on $|D \cap E(T)|$ for a given cut D.

29. Apply Theorem 1.9.6.

Hints for Chapter 3

1.$^-$ Recall the definitions of ‘separate’ and ‘component’.

2. Describe in words what the picture suggests.

3. Use Exercise 1 to answer the first question. The second requires an elementary calculation, which the figure may already suggest.

4. Only the first part needs arguing; the second then follows by symmetry. So suppose a component of $G - X$ is not met by X', and refer to Exercise 1. Where does X' lie? Are all our assumptions about X' consistent?

5.$^-$ How can a block fail to be a maximal 2-connected subgraph? And what else follows then?

6. Deduce the connectedness of the block graph from that of the graph itself, and its acyclicity from the maximality of each block.

7. Prove the statement inductively using Proposition 3.1.2. Alternatively, choose a cycle through one of the two vertices and with minimum distance from the other vertex. Show that this distance cannot be positive.

8. Belonging to the same block is an equivalence relation on the edge set; see Exercise 5.

9. Induction along Proposition 3.1.2.

10. Assuming that G/xy is not 3-connected, distinguish the cases when v_{xy} lies inside or outside a separating set of at most 2 vertices.

11. (i) Consider the edges incident with a smaller separator.

    (ii) Induction shows that all the graphs obtained by the construction are cubic and 3-connected. For the converse, consider a maximal subgraph $TH \subseteq G$ such that H is constructible as stated; then show that $H = G$.

12.$^-$ Can *any* choice of X and $\mathcal{P}$ as suggested by Menger’s theorem fail?

13. Choose the disjoint A–B paths in $L(G)$ minimal.

14. Consider a longest cycle C. How are the other vertices joined to C?

15. Consider a cycle through as many of the k given vertices as possible. If one them is missed, can you re-route the cycle through it?

16. Consider the graph of the hint. Show that any subset of its vertices that meets all H-paths (but not H) corresponds to a similar subset of $E(G) \setminus E(H)$. What does a pair of independent H-paths in the auxiliary graph correspond to in G?

17.$^-$ How many paths can a single K^{2m+1} accomodate?

18. Choose suitable degrees for the vertices in B.

19.$^+$ Let H be the (edgeless) graph on the new vertices. Consider the sets X and F that Mader’s theorem provides if G' does not contain $|G|/2$ independent H-paths. If G has no 1-factor, use these to find a suitable set that can play the role of S in Tutte’s theorem.

20. Think small.

21.$^-$ If two vertices s, t are separated by fewer than $2k - 1$ vertices, extend $\{ s \}$ and $\{ t \}$ to k-sets S and T showing that G is not k-linked.

27.$^+$ To show existence, define the required bijections $F \to V^*$, $E \to E^*$, $V \to F^*$ successively in this order, while at the same time constructing G^*. Show that connectedness is necessary to ensure that these three functions can all be made bijective.

28. Solve the previous exercise first.

29. Use the bijections that come with the two duals to define the desired isomorphism and to prove that it is combinatorial.

30. Apply Menger's theorem and Proposition 4.6.1. For (iii), consider a 4-connected graph with six vertices.

31. Apply induction on n, starting with part (i) of the previous exercise.

32. Theorem 1.9.5.

33. For the forward implication, consider $G' := G^*$. For the converse, apply a suitable planarity criterion.

Hints for Chapter 5

1.$^-$ Duality.

2.$^-$ Whenever more than three countries have some point in common, apply a small local change to the map where this happens.

3. Where does the five colour proof use the fact that v has no more neighbours than there are colours?

4. How can the colourings of different blocks interfere with each other?

5.$^-$ Use a colouring of G to derive a suitable ordering.

6. Consider how the removal of certain edges may lead the greedy algorithm to use more colours.

7. Describe more precisely how to implement this alternative algorithm. Then, where is the difference to the traditional greedy algorithm?

8. Compare the number of edges in a subgraph H as in 5.2.2 with the number m of edges in G.

9. To find f, consider a given graph of small colouring number and partition it inductively into a small number of forest. For g, use Proposition 5.2.2 and the easy direction of Theorem 3.5.4.

10.$^-$ Remove vertices successively until the graph becomes critically k-chromatic. What can you say about the degree of any vertex that remains?

11. Proposition 1.6.1.

12.$^+$ Modify colourings of the two sides of a hypothetical cut of fewer than $k-1$ edges so that they combine to a $(k-1)$-colouring of the entire graph (with a contradiction).

13. Proposition 1.3.1.

14.$^-$ For which graphs with large maximum degree does Proposition 5.2.2 give a particularly small upper bound?

35. Look at the complement.

36. Define the colour classes of a given induced subgraph $H \subseteq G$ inductively, starting with the class of all minimal elements.

37. (i) Can the vertices on an induced cycle contain each other as intervals?

    (ii) Use the natural ordering of the reals.

38. Compare $\omega(H)$ with $\Delta(G)$ (where $H = L(G)$).

39.$^+$ Which graphs are such that their line graphs contain no induced cycles of odd length $\geqslant 5$? To prove that the edges of such a graph G can be coloured with $\omega(L(G))$ colours, imitate the proof of Vizing's theorem.

40. Use A as a colour class.

41.$^+$ (i) Induction.

    (ii) Assume that G contains no induced P^3. Suppose some H has a maximal complete subgraph K and a maximal set A of independent vertices disjoint from K. For each vertex $v \in K$, consider the set of neighbours of v in A. How do these sets intersect? Is there a smallest one?

42.$^+$ Start with a candidate for the set $\mathcal{O}$, i.e. a set of maximal complete subgraphs covering the vertex set of G. If all the elements of $\mathcal{O}$ happen to have order $\omega(G)$, how does the existence of $\mathcal{A}$ follow from the perfection of G? If not, can you expand G (maintaining perfection) so that they do and adapt the $\mathcal{A}$ for the expanded graph to G?

43.$^+$ Reduce the general case to the case when all but one of the G_x are trivial; then imitate the proof of Lemma 5.5.4.

44. Apply the property of $\mathcal{H}_1$ to the graphs in $\mathcal{H}_2$, and vice versa.

Hints for Chapter 6

1.$^-$ Move the vertices, one by one, from $\overline{S}$ to S. How does the value of $f(S, \overline{S})$ change each time?

2. (i) Trick the algorithm into repeatedly using the middle edge in alternating directions.

   (ii) At any given time during the algorithm, consider for each vertex v the shortest s–v walk that qualifies as an initial segment of an augmenting path. Show for each v that the length of this s–v walk never decreases during the algorithm. Now consider an edge which is used twice for an augmenting path, in the same direction. Show that the second of these paths must have been longer than the first. Now derive the desired bound.

3.$^+$ For the edge version, define the capacity function so that a flow of maximum value gives rise to sufficiently many edge-disjoint paths. For the vertex version, split every vertex x into two adjacent vertices x^-, x^+. Define the edges of the new graph and their capacities in such a way that positive flow through an edge x^-x^+ corresponds to the use of x by a path in G.

6.$^+$ What is the maximum number of edges in a graph of the structure given by Theorem 2.2.3 if it has no matching of size k? What is the optimal distribution of vertices between S and the components of $G - S$? Is there always a graph whose number of edges attains the corresponding upper bound?

7. Consider a vertex $x \in G$ of maximum degree, and count the edges in $G - x$.

8. Choose k and i so that $n = (r-1)k + i$ with $0 \leqslant i < r - 1$. Treat the case of $i = 0$ first, and then show for the general case that $t_{r-1}(n) = \frac{1}{2}\frac{r-2}{r-1}(n^2 - i^2) + \binom{i}{2}$.

9. The bounds given in the hint are the sizes of two particularly simple Turán graphs—which ones?

10.$^+$ How can you choose v so that the number of edges does not decrease? Where in the graph can the operation be repeated, and what does the situation look like when nothing new happens?

11. Choose among the m vertices a set of s vertices that are still incident with as many edges as possible.

12. For the first inequality, double the vertex set of an extremal graph for $K_{s,t}$ to obtain a bipartite graph with twice as many edges but still not containing a $K_{s,t}$.

13.$^+$ For the displayed inequality, count the pairs (x, Y) such that $x \in A$ and $Y \subseteq B$, with $|Y| = r$ and x adjacent to all of Y. For the bound on $\mathrm{ex}(n, K_{r,r})$, use the estimate $(s/t)^t \leq \binom{s}{t} \leq s^t$ and the fact that the function $z \mapsto z^r$ is convex.

14. Assume that the upper density is larger than $1 - \frac{1}{r-1}$. What does this mean precisely, and what does the Erdős-Stone theorem then imply?

15. Corollary 1.5.4 and Proposition 1.2.2.

16. Complete graphs.

17.$^-$ Average degree.

18. Do $\frac{1}{2}(k-1)n$ edges force a subgraph of suitable minimum degree?

19. Consider a longest path P in G. Where do its endvertices have their neighbours? Can $G[P]$ contain a cycle on $V(P)$?

20.$^-$ Why would it be impractical to include, say, 1-element sets X, Y in the comparison?

21.$^-$ Apply the definition of an ϵ-regular pair.

22. Sparse graphs have few edges. How does that affect the average degree condition in the definition of ϵ-regularity?

Hints for Chapter 9

1.$^-$ Can you colour the edges of K^5 red and green without creating a red or a green triangle? Can you do the same for a K^6?

2. Induction on c. In the induction step, unite two of the colour classes.

3.$^+$ Choose a well-ordering of $\mathbb{R}$, and compare it with the natural ordering. Use the fact that countable unions of countable sets are countable.

4.$^+$ The first and second question are easy. To prove the theorem of Erdős and Szekeres, use induction on k for fixed ℓ, and consider in the induction step the last elements of increasing subsequences of length k. Alternatively, apply Dilworth's Theorem.

5. Use the fact that $n \geqslant 4$ points span a convex polygon if and only if every four of them do.

6. Translate the given k-partition of $\{ 1, 2, \dots, n \}$ into a k-colouring of the edges of K^n.

7. (i) is easy. For (ii) use the existence of $R(2, k, 3)$.

8. Begin by finding infinitely many sets whose pairwise intersections all have the same size.

9. The exercise offers more information than you need. Consult Chapter 8.1 to see what is relevant.

10. Consider an auxiliary graph whose vertices are coloured finite subgraphs of the given graph.

11. Imitate the proof of Proposition 9.2.1.

12. The lower bound is easy. Given a colouring for the upper bound, consider a vertex and the neighbours joined to it by suitably coloured edges.

13.$^-$ Given H_1 and H_2, construct a graph H for which the G of Theorem 9.3.1 satisfies $(*)$.

14. Show inductively for $k = 0, \dots, m$ that $\omega(G^k) = \omega(H)$.

15. For the induction step, construct $G(H_1, H_2)$ from the disjoint union of $G(H_1, H_2')$ and $G(H_1', H_2)$ by joining some new vertices in a suitable way.

16. Infinity lemma.

17.$^-$ How exactly does Proposition 9.4.1 fail if we delete K^r from the statement?

Hints for Chapter 10

1. Consider the union of two colour classes.

2. Will the proof of Proposition 10.1.2 go through if we assume $\chi(G) \geqslant |G|/k$ instead of $\alpha(G) \leqslant k$? What do k-connected graphs look like that satisfy the first condition but not the second?

3. Examine an edge that gets added in one sequence but not in another.

15. For the first problem modify an increasing property slightly, so that it ceases to be increasing but keeps its threshold function. For the second, look for an increasing property whose probability does not really depend on p.
16.$^-$ Permutations of $V(H)$.
17.$^-$ This is a result from the text in disguise.
18.$^-$ Balance.
19. For $p/t \to 0$ apply Lemmas 11.1.4 and 11.1.5. For $p/t \to \infty$ apply Corollary 11.4.4.
20. There are only finitely many trees of order k.
21.$^+$ Show first that no such threshold function $t = t(n)$ can tend to zero as $n \to \infty$. Then use Exercise 12.
22.$^+$ Examine the various steps in the proof of Theorem 11.4.3, and note which changes will be needed. In the final steps of the proof, how are the sums A_F defined, and why is the sum of all the A_F with $||F|| = \emptyset$ equal to A_0? For $||F|| \neq \emptyset$, calculate a bound on A_F, and show that each A_F/μ^2 tends to zero as $n \to \infty$, as before.

Hints for Chapter 12

1.$^-$ Antisymmetry.
2. Proposition 12.1.1.
3. To prove Proposition 12.1.1, consider an infinite sequence in which every strictly decreasing subsequence is finite. How does the last element of a maximal decreasing subsequence compare with the elements that come after it? For Corollary 12.1.2, start by proving that at least one element forms a good pair with infinitely many later elements.
4. An obvious approach is to try to imitate the proof of Lemma 12.1.3 for $\leqslant'$; if it fails, what is the reason? Alternatively, you might try to modify the injective map produced by Lemma 12.1.3 into an order-preserving one, without losing the property of $a \leqslant f(a)$ for all a.
5.$^-$ This is an exercise in precision: 'easy to see' is not a proof...
6. Start by finding two trees T, T' with $|T| < |T'|$ but $T \not\leqslant T'$; then iterate.
7. Does the original proof ever map the root of a tree to an ordinary vertex of another tree?
8.$^+$ When we try to embed a graph TG in another graph H, the branch vertices of the TG can be mapped only to certain vertices of H. Enlarge G to a similar graph H that does not contain G as a topological minor because these vertices of H are inconveniently positioned in H. Then iterate this example to obtain an infinite antichain.
9.$^+$ It is. One possible proof uses normal spanning trees with labels, and imitates the proof of Kruskal's theorem.
10. Why are there no cycles of tree-width 1?

32. Consult Chapter 8.2 for substructures to be found in graphs of large chromatic number.
33. Derive the minor theorem first for connected graphs.
34. K^5.