Preface

Almost two decades have passed since the appearance of those graph theory texts that still set the agenda for most introductory courses taught today. The canon created by those books has helped to identify some main fields of study and research, and will doubtless continue to influence the development of the discipline for some time to come.

Yet much has happened in those 20 years, in graph theory no less than elsewhere: deep new theorems have been found, seemingly disparate methods and results have become interrelated, entire new branches have arisen. To name just a few such developments, one may think of how the new notion of list colouring has bridged the gulf between invariants such as average degree and chromatic number, how probabilistic methods and the regularity lemma have pervaded extremal graph theory and Ramsey theory, or how the entirely new field of graph minors and tree-decompositions has brought standard methods of surface topology to bear on long-standing algorithmic graph problems.

Clearly, then, the time has come for a reappraisal: *what are, today, the essential areas, methods and results that should form the centre of an introductory graph theory course aiming to equip its audience for the most likely developments ahead?*

I have tried in this book to offer material for such a course. In view of the increasing complexity and maturity of the subject, I have broken with the tradition of attempting to cover both theory and applications: this book offers an introduction to the theory of graphs as part of (pure) mathematics; it contains neither explicit algorithms nor 'real world' applications. My hope is that the potential for depth gained by this restriction in scope will serve students of computer science as much as their peers in mathematics: assuming that they prefer algorithms but will benefit from an encounter with pure mathematics of *some* kind, it seems an ideal opportunity to look for this close to where their heart lies!

In the selection and presentation of material, I have tried to accommodate two conflicting goals. On the one hand, I believe that an text.) There are two areas of graph theory which I find both fascinating and important, especially from the perspective of pure mathematics adopted here, but which are not covered in this book: these are algebraic graph theory and infinite graphs.

At the end of each chapter, there is a section with exercises and another with bibliographical and historical notes. Many of the exercises were chosen to complement the main narrative of the text: they illustrate new concepts, show how a new invariant relates to earlier ones, or indicate ways in which a result stated in the text is best possible. Particularly easy exercises are identified by the superscript $^-$, the more challenging ones carry a $^+$. The notes are intended to guide the reader on to further reading, in particular to any monographs or survey articles on the theme of that chapter. They also offer some historical and other remarks on the material presented in the text.

Ends of proofs are marked by the symbol □. Where this symbol is found directly below a formal assertion, it means that the proof should be clear after what has been said—a claim waiting to be verified! There are also some deeper theorems which are stated, without proof, as background information: these can be identified by the absence of both proof and □.

Almost every book contains errors, and this one will hardly be an exception. I shall try to post on the Web any corrections that become necessary. The relevant site may change in time, but will always be accessible via the following two addresses:

http://www.springer-ny.com/supplements/diestel/
http://www.springer.de/catalog/html-files/deutsch/math/3540609180.html

Please let me know about any errors you find.

Little in a textbook is truly original: even the style of writing and of presentation will invariably be influenced by examples. The book that no doubt influenced me most is the classic GTM graph theory text by Bollobás: it was in the course recorded by this text that I learnt my first graph theory as a student. Anyone who knows this book well will feel its influence here, despite all differences in contents and presentation.

I should like to thank all who gave so generously of their time, knowledge and advice in connection with this book. I have benefited particularly from the help of N. Alon, G. Brightwell, R. Gillett, R. Halin, M. Hintz, A. Huck, I. Leader, T. Łuczak, W. Mader, V. Rödl, A.D. Scott, P.D. Seymour, G. Simonyi, M. Škoviera, R. Thomas, C. Thomassen and P. Valtr. I am particularly grateful also to Tommy R. Jensen, who taught me much about colouring and all I know about k-flows, and who invested immense amounts of diligence and energy in his proofreading of the preliminary German version of this book.

March 1997 *RD*

Contents

1 The Basics

This chapter gives a gentle yet concise introduction to most of the terminology used later in the book. Fortunately, much of standard graph theoretic terminology is so intuitive that it is easy to remember; the few terms better understood in their proper setting will be introduced later, when their time has come.

Section 1.1 offers a brief but self-contained summary of the most basic definitions in graph theory, those centred round the notion of a graph. Most readers will have met these definitions before, or will have them explained to them as they begin to read this book. For this reason, Section 1.1 does not dwell on these definitions more than clarity requires: its main purpose is to collect the most basic terms in one place, for easy reference later.

From Section 1.2 onwards, all new definitions will be brought to life almost immediately by a number of simple yet fundamental propositions. Often, these will relate the newly defined terms to one another: the question of how the value of one invariant influences that of another underlies much of graph theory, and it will be good to become familiar with this line of thinking early.

By $\mathbb{N}$ we denote the set of natural numbers, including zero. The set $\mathbb{Z}/n\mathbb{Z}$ of integers modulo n is denoted by $\mathbb{Z}_n$; its elements are written as $\overline{i} := i + n\mathbb{Z}$. For a real number x we denote by $\lfloor x \rfloor$ the greatest integer $\leqslant x$, and by $\lceil x \rceil$ the least integer $\geqslant x$. Logarithms written as 'log' are taken at base 2; the natural logarithm will be denoted by 'ln'. A set $\mathcal{A} = \{ A_1, \dots, A_k \}$ of disjoint subsets of a set A is a *partition* of A if $A = \bigcup_{i=1}^{k} A_i$ and $A_i \neq \emptyset$ for every i. Another partition $\{ A'_1, \dots, A'_\ell \}$ of A *refines* the partition $\mathcal{A}$ if each A'_i is contained in some A_j. By $[A]^k$ we denote the set of all k-element subsets of A. Sets with k elements will be called *k-sets*; subsets with k elements are *k-subsets*.

$\mathbb{Z}_n$

$\lfloor x \rfloor, \lceil x \rceil$

log, ln

partition

$[A]^k$

k-set

adjacent *neighbour* *complete* K^n

Two vertices x, y of G are *adjacent*, or *neighbours*, if xy is an edge of G. Two edges $e \neq f$ are *adjacent* if they have an end in common. If all the vertices of G are pairwise adjacent, then G is *complete*. A complete graph on n vertices is a K^n; a K^3 is called a *triangle*.

independent

Pairwise non-adjacent vertices or edges are called *independent*. More formally, a set of vertices or of edges is *independent* (or *stable*) if no two of its elements are adjacent.

$\simeq$ *isomorphism* *invariant*

Let $G = (V, E)$ and $G' = (V', E')$ be two graphs. We call G and G' *isomorphic*, and write $G \simeq G'$, if there exists a bijection $\varphi: V \to V'$ with $xy \in E \Leftrightarrow \varphi(x)\varphi(y) \in E'$ for all $x, y \in V$. Such a map φ is called an *isomorphism*; if $G = G'$, it is called an *automorphism*. We do not normally distinguish between isomorphic graphs. Thus, we usually write $G = G'$ rather than $G \simeq G'$, speak of *the* complete graph on 17 vertices, and so on. A map taking graphs as arguments is called a *graph invariant* if it assigns equal values to isomorphic graphs. The number of vertices and the number of edges of a graph are two simple graph invariants; the greatest number of pairwise adjacent vertices is another.

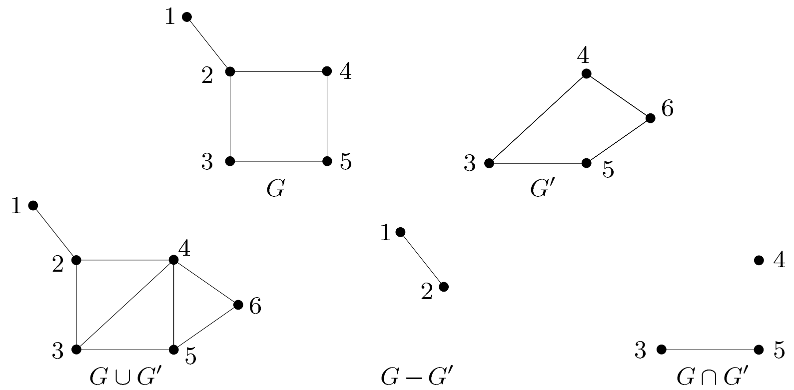

Fig. 1.1.2. Union, difference and intersection; the vertices 2,3,4 induce (or span) a triangle in $G \cup G'$ but not in G

$G \cap G'$ *subgraph* $G' \subseteq G$

We set $G \cup G' := (V \cup V', E \cup E')$ and $G \cap G' := (V \cap V', E \cap E')$. If $G \cap G' = \emptyset$, then G and G' are *disjoint*. If $V' \subseteq V$ and $E' \subseteq E$, then G' is a *subgraph* of G (and G a *supergraph* of G'), written as $G' \subseteq G$. Less formally, we say that G *contains* G'.

induced subgraph $G[U]$ *spanning*

If $G' \subseteq G$ and G' contains *all* the edges $xy \in E$ with $x, y \in V'$, then G' is an *induced subgraph* of G; we say that V' *induces* or *spans* G' in G, and write $G' =: G[V']$. Thus if $U \subseteq V$ is any set of vertices, then $G[U]$ denotes the graph on U whose edges are precisely the edges of G with both ends in U. If H is a subgraph of G, not necessarily induced, we abbreviate $G[V(H)]$ to $G[H]$. Finally, $G' \subseteq G$ is a *spanning* subgraph of G if V' spans all of G, i.e. if $V' = V$.

for $U \subseteq V$, the neighbours in $V \setminus U$ of vertices in U are called *neighbours of* U; their set is denoted by $N(U)$.

The *degree* (or *valency*) $d_G(v) = d(v)$ of a vertex v is the number $|E(v)|$ of edges at v; by our definition of a graph,[2] this is equal to the number of neighbours of v. A vertex of degree 0 is *isolated*. The number $\delta(G) := \min\{\, d(v) \mid v \in V \,\}$ is the *minimum degree* of G, the number $\Delta(G) := \max\{\, d(v) \mid v \in V \,\}$ its *maximum degree*. If all the vertices of G have the same degree k, then G is *k-regular*, or simply *regular*. A 3-regular graph is called *cubic*.

degree $d(v)$ *isolated* *$\delta(G)$* *$\Delta(G)$* *regular* *cubic*

The number

$$d(G) := \frac{1}{|V|} \sum_{v \in V} d(v)$$

$d(G)$

is the *average degree* of G. Clearly,

average degree

$$\delta(G) \leqslant d(G) \leqslant \Delta(G)\,.$$

The average degree quantifies globally what is measured locally by the vertex degrees: the number of edges of G per vertex. Sometimes it will be convenient to express this ratio directly, as $\varepsilon(G) := |E|/|V|$.

$\varepsilon(G)$

The quantities d and ε are, of course, intimately related. Indeed, if we sum up all the vertex degrees in G, we count every edge exactly twice: once from each of its ends. Thus

$$|E| = \tfrac{1}{2} \sum_{v \in V} d(v) = \tfrac{1}{2} d(G) \cdot |V|\,,$$

and therefore

$$\varepsilon(G) = \tfrac{1}{2} d(G)\,.$$

Proposition 1.2.1. *The number of vertices of odd degree in a graph is always even.* [10.3.3]

Proof. A graph on V has $\frac{1}{2}\sum_{v \in V} d(v)$ edges, so $\sum d(v)$ is an even number. □

If a graph has large minimum degree, i.e. everywhere, locally, many edges per vertex, it also has many edges per vertex globally: $\varepsilon(G) = \frac{1}{2}d(G) \geqslant \frac{1}{2}\delta(G)$. Conversely, of course, its average degree may be large even when its minimum degree is small. However, the vertices of large degree cannot be scattered completely among vertices of small degree: as the next proposition shows, every graph G has a subgraph whose average degree is no less than the average degree of G, and whose minimum degree is more than half its average degree:

[2] but not for multigraphs; see Section 1.10

For $0 \leqslant i \leqslant j \leqslant k$ we write *xPy, $\mathring{P}$*

$$Px_i := x_0 \dots x_i$$
$$x_iP := x_i \dots x_k$$
$$x_iPx_j := x_i \dots x_j$$

and

$$\mathring{P} := x_1 \dots x_{k-1}$$
$$P\mathring{x}_i := x_0 \dots x_{i-1}$$
$$\mathring{x}_iP := x_{i+1} \dots x_k$$
$$\mathring{x}_iP\mathring{x}_j := x_{i+1} \dots x_{j-1}$$

for the appropriate subpaths of P. We use similar intuitive notation for the concatenation of paths; for example, if the union $Px \cup xQy \cup yR$ of three paths is again a path, we may simply denote it by $PxQyR$. *$PxQyR$*

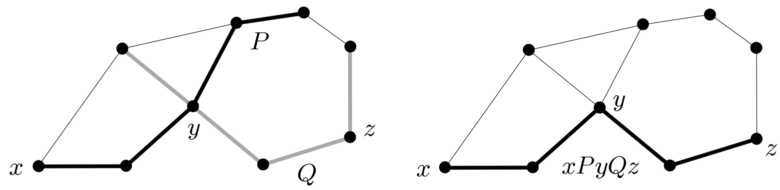

Fig. 1.3.2. Paths P, Q and $xPyQz$

Given sets A, B of vertices, we call $P = x_0 \dots x_k$ an *A–B path* if $V(P) \cap A = \{x_0\}$ and $V(P) \cap B = \{x_k\}$. As before, we write a–B path rather than $\{a\}$–B path, etc. Two or more paths are *independent* if none of them contains an inner vertex of another. Two a–b paths, for instance, are independent if and only if a and b are their only common vertices. *A–B path* *independent*

Given a graph H, we call P an *H-path* if P is non-trivial and meets H exactly in its ends. In particular, the edge of any H-path of length 1 is never an edge of H. *H-path*

If $P = x_0 \dots x_{k-1}$ is a path and $k \geqslant 3$, then the graph $C := P + x_{k-1}x_0$ is called a *cycle*. As with paths, we often denote a cycle by its (cyclic) sequence of vertices; the above cycle C might be written as $x_0 \dots x_{k-1}x_0$. The *length* of a cycle is its number of edges (or vertices); the cycle of length k is called a *k-cycle* and denoted by C^k. *cycle* *length* *C^k*

The minimum length of a cycle (contained) in a graph G is the *girth* $g(G)$ of G; the maximum length of a cycle in G is its *circumference*. (If G does not contain a cycle, we set the former to ∞, the latter to zero.) An edge which joins two vertices of a cycle but is not itself an edge of the cycle is a *chord* of that cycle. Thus, an *induced cycle* in G, a cycle in G forming an induced subgraph, is one that has no chords (Fig. 1.3.3). *girth $g(G)$* *circumference* *chord* *induced cycle*

A vertex is *central* in G if its greatest distance from any other vertex is as small as possible. This distance is the *radius* of G, denoted by $\mathrm{rad}(G)$. Thus, formally, $\mathrm{rad}(G) = \min_{x \in V(G)} \max_{y \in V(G)} d_G(x, y)$. As one easily checks (exercise), we have *central* *radius* $\mathrm{rad}(G)$

$$\mathrm{rad}(G) \leqslant \mathrm{diam}(G) \leqslant 2\,\mathrm{rad}(G)\,.$$

Diameter and radius are not directly related to the minimum or average degree: a graph can combine large minimum degree with large diameter, or small average degree with small diameter (examples?).

The maximum degree behaves differently here: a graph of large order can only have small radius and diameter if its maximum degree is large. This connection is quantified very roughly in the following proposition:

Proposition 1.3.3. *A graph G of radius at most k and maximum degree at most d has no more than $1 + kd^k$ vertices.* [9.4.1] [9.4.2]

Proof. Let z be a central vertex in G, and let D_i denote the set of vertices of G at distance i from z. Then $V(G) = \bigcup_{i=0}^{k} D_i$, and $|D_0| = 1$. Since $\Delta(G) \leqslant d$, we have $|D_i| \leqslant d\,|D_{i-1}|$ for $i = 1, \dots, k$, and thus $|D_i| \leqslant d^i$ by induction. Adding up these inequalities we obtain

$$|G| \leqslant 1 + \sum_{i=1}^{k} d^i \leqslant 1 + kd^k.$$

□

A *walk* (of *length* k) in a graph G is a non-empty alternating sequence $v_0 e_0 v_1 e_1 \dots e_{k-1} v_k$ of vertices and edges in G such that $e_i = \{\, v_i, v_{i+1} \,\}$ for all $i < k$. If $v_0 = v_k$, the walk is *closed*. If the vertices in a walk are all distinct, it defines an obvious path in G. In general, every walk between two vertices contains[4] a path between these vertices (proof?). *walk*

1.4 Connectivity

A non-empty graph G is called *connected* if any two of its vertices are linked by a path in G. If $U \subseteq V(G)$ and $G[\,U\,]$ is connected, we also call U itself connected (in G). *connected*

Proposition 1.4.1. *The vertices of a connected graph G can always be enumerated, say as $v_1, \dots, v_n$, so that $G_i := G[\,v_1, \dots, v_i\,]$ is connected for every i.* [1.5.2]

[4] We shall often use terms defined for graphs also for walks, as long as their meaning is obvious.

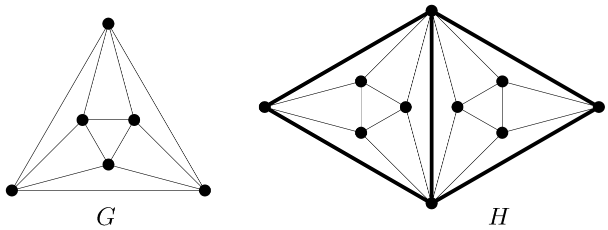

Fig. 1.4.3. The *octahedron* G (left) with $\kappa(G) = \lambda(G) = 4$, and a graph H with $\kappa(H) = 2$ but $\lambda(H) = 4$

such that G is ℓ-edge-connected is the *edge-connectivity* $\lambda(G)$ of G. In particular, we have $\lambda(G) = 0$ if G is disconnected. *edge-connectivity* $\lambda(G)$

For every non-trivial graph G we have

$$\kappa(G) \leqslant \lambda(G) \leqslant \delta(G)$$

(exercise), so in particular high connectivity requires a large minimum degree. Conversely, large minimum degree does not ensure high connectivity, not even high edge-connectivity (examples?). It does, however, imply the existence of a highly connected subgraph:

Theorem 1.4.2. (Mader 1972)
Every graph of average degree at least $4k$ *has a* k*-connected subgraph.* [8.1.1] [11.2.3]

Proof. For $k \in \{0, 1\}$ the assertion is trivial; we consider $k \geqslant 2$ and a graph $G = (V, E)$ with $|V| =: n$ and $|E| =: m$. For inductive reasons it will be easier to prove the stronger assertion that G has a k-connected subgraph whenever

(i) $n \geqslant 2k - 1$ and

(ii) $m \geqslant (2k-3)(n-k+1)+1$.

(This assertion is indeed stronger, i.e. (i) and (ii) follow from our assumption of $d(G) \geqslant 4k$: (i) holds since $n > \Delta(G) \geqslant d(G) \geqslant 4k$, while (ii) follows from $m = \frac{1}{2}d(G)n \geqslant 2kn$.)

We apply induction on n. If $n = 2k-1$, then $k = \frac{1}{2}(n+1)$, and hence $m \geqslant \frac{1}{2}n(n-1)$ by (ii). Thus $G = K^n \supseteq K^{k+1}$, proving our claim. We now assume that $n \geqslant 2k$. If v is a vertex with $d(v) \leqslant 2k-3$, we can apply the induction hypothesis to $G - v$ and are done. So we assume that $\delta(G) \geqslant 2k-2$. If G is k-connected, there is nothing to show. We may therefore assume that G has the form $G = G_1 \cup G_2$ with $|G_1 \cap G_2| < k$ and $|G_1|, |G_2| < n$. As every edge of G lies in G_1 or in G_2, G has no edge between $G_1 - G_2$ and $G_2 - G_1$. Since each vertex in these subgraphs has at least $\delta(G) \geqslant 2k-2$ neighbours, we have $|G_1|, |G_2| \geqslant 2k-1$. But then at least one of the graphs G_1, G_2 must satisfy the induction hypothesis

The proof of Theorem 1.5.1 is straightforward, and a good exercise for anyone not yet familiar with all the notions it relates. Extending our notation for paths from Section 1.3, we write xTy for the unique path in a tree T between two vertices x, y (see (ii) above). xTy

A frequently used application of Theorem 1.5.1 is that every connected graph contains a spanning tree: by the equivalence of (i) and (iii), any minimal connected spanning subgraph will be a tree. Figure 1.4.1 shows a spanning tree in each of the three components of the graph depicted.

Corollary 1.5.2. *The vertices of a tree can always be enumerated, say as $v_1, \dots, v_n$, so that every v_i with $i \geqslant 2$ has a unique neighbour in $\{ v_1, \dots, v_{i-1} \}$.*

Proof. Use the enumeration from Proposition 1.4.1. □ (1.4.1)

[1.9.6] [3.5.1] [3.5.4] [4.2.7] [8.2.2]

Corollary 1.5.3. *A connected graph with n vertices is a tree if and only if it has $n-1$ edges.*

Proof. Induction on i shows that the subgraph spanned by the first i vertices in Corollary 1.5.2 has $i-1$ edges; for $i = n$ this proves the forward implication. Conversely, let G be any connected graph with n vertices and $n-1$ edges. Let G' be a spanning tree in G. Since G' has $n-1$ edges by the first implication, it follows that $G = G'$. □

[9.2.1] [9.2.3]

Corollary 1.5.4. *If T is a tree and G is any graph with $\delta(G) \geqslant |T|-1$, then $T \subseteq G$, i.e. G has a subgraph isomorphic to T.*

Proof. Find a copy of T in G inductively along its vertex enumeration from Corollary 1.5.2. □

Sometimes it is convenient to consider one vertex of a tree as special; such a vertex is then called the *root* of this tree. A tree with a fixed root is a *rooted tree*. Choosing a root r in a tree T imposes a partial ordering on $V(T)$ by letting $x \leqslant y$ if $x \in rTy$. This is the *tree-order* on $V(T)$ associated with T and r. Note that r is the least element in this partial order, every leaf $x \neq r$ of T is a maximal element, the ends of any edge of T are comparable, and every set of the form $\{ x \mid x \leqslant y \}$ (where y is any fixed vertex) is a *chain*, a set of pairwise comparable elements. (Proofs?) *root* *tree-order* *chain*

A rooted tree T contained in a graph G is called *normal* in G if the ends of every T-path in G are comparable in the tree-order of T. If T spans G, this amounts to requiring that two vertices of T must be comparable whenever they are adjacent in G; see Figure 1.5.2. Normal spanning trees are also called *depth-first search trees*, because of the way they arise in computer searches on graphs (Exercise 17). *normal tree*

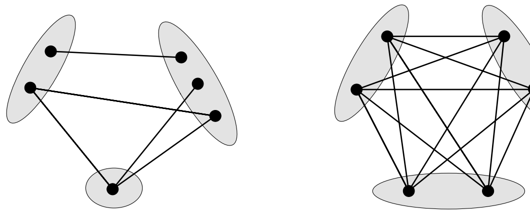

$K_{2,2,2} = K^3_2$

Fig. 1.6.1. Two 3-partite graphs

complete r-partite graph $\overline{K^{n_1}} * \ldots * \overline{K^{n_r}}$ is denoted by $K_{n_1,\ldots,n_r}$; if $n_1 = \ldots = n_r =: s$, we abbreviate this to K^r_s. Thus, K^r_s is the complete r-partite graph in which every partition class contains exactly s vertices.[5] (Figure 1.6.1 shows the example of the octahedron K^3_2; compare its drawing with that in Figure 1.4.3.) Graphs of the form $K_{1,n}$ are called *stars*. $K_{n_1,\ldots,n_r}$ K^r_s *star*

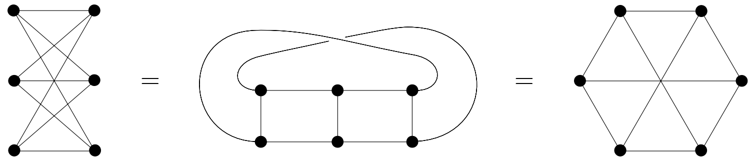

Fig. 1.6.2. Three drawings of the bipartite graph $K_{3,3} = K^2_3$

Clearly, a bipartite graph cannot contain an *odd cycle*, a cycle of odd length. In fact, the bipartite graphs are characterized by this property: *odd cycle*

Proposition 1.6.1. *A graph is bipartite if and only if it contains no odd cycle.* [5.3.1] [6.4.2]

Proof. Let $G = (V, E)$ be a graph without odd cycles; we show that G is bipartite. Clearly a graph is bipartite if all its components are bipartite or trivial, so we may assume that G is connected. Let T be a spanning tree in G, pick a root $r \in T$, and denote the associated tree-order on V by $\leqslant_T$. For each $v \in V$, the unique path rTv has odd or even length. This defines a bipartition of V; we show that G is bipartite with this partition. (1.5.1)

Let $e = xy$ be an edge of G. If $e \in T$, with $x <_T y$ say, then $rTy = rTxy$ and so x and y lie in different partition classes. If $e \notin T$ then $C_e := xTy + e$ is a cycle (Fig. 1.6.3), and by the case treated already the vertices along xTy alternate between the two classes. Since C_e is even by assumption, x and y again lie in different classes. □

5 Note that we obtain a K^r_s if we replace each vertex of a K^r by an independent s-set; our notation of K^r_s is intended to hint at this connection.

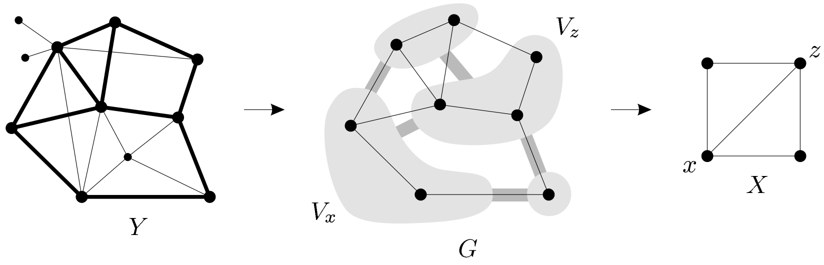

Fig. 1.7.2. $Y \supseteq G = MX$, so X is a minor of Y

branch set to a single vertex and deleting any 'parallel edges' or 'loops' that may arise.

G/U *v_U*

If $V_x = U \subseteq V$ is one of the branch sets above and every other branch set consists just of a single vertex, we also write G/U for the graph X and v_U for the vertex $x \in X$ to which U contracts, and think of the rest of X as an induced subgraph of G. The contraction of a single edge uu' defined earlier can then be viewed as the special case of $U = \{ u, u' \}$.

Proposition 1.7.1. *G is an MX if and only if X can be obtained from G by a series of edge contractions, i.e. if and only if there are graphs $G_0, \ldots, G_n$ and edges $e_i \in G_i$ such that $G_0 = G$, $G_n \simeq X$, and $G_{i+1} = G_i/e_i$ for all $i < n$.*

Proof. Induction on $|G| - |X|$. □

minor; $\preccurlyeq$

If $G = MX$ is a subgraph of another graph Y, we call X a *minor* of Y and write $X \preccurlyeq Y$. Note that every subgraph of a graph is also its minor; in particular, every graph is its own minor. By Proposition 1.7.1, any minor of a graph can be obtained from it by first deleting some vertices and edges, and then contracting some further edges. Conversely, any graph obtained from another by repeated deletions and contractions (in any order) is its minor: this is clear for one deletion or contraction, and follows for several from the transitivity of the minor relation (Proposition 1.7.3).

subdivision *TX* *topological minor*

If we replace the edges of X with independent paths between their ends (so that none of these paths has an inner vertex on another path or in X), we call the graph G obtained a *subdivision* of X and write $G = TX$.[7] If $G = TX$ is the subgraph of another graph Y, then X is a *topological minor* of Y (Fig. 1.7.3).

---

[7] So again TX denotes an entire class of graphs: all those which, viewed as a topological space in the obvious way, are homeomorphic to X. The T in TX stands for 'topological'.

Fig. 1.8.1. The bridges of Königsberg (anno 1736)

the old city that traverses each of the bridges shown in Figure 1.8.1 exactly once.

Thus inspired,[8] let us call a closed walk in a graph an *Euler tour* if it traverses every edge of the graph exactly once. A graph is *Eulerian* if it admits an Euler tour.

Eulerian

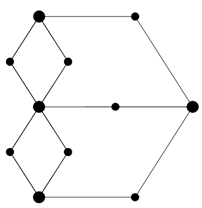

Fig. 1.8.2. A graph formalizing the bridge problem

Theorem 1.8.1. (Euler 1736)
A connected graph is Eulerian if and only if every vertex has even degree.

[2.1.5]
[10.3.3]

Proof. The degree condition is clearly necessary: a vertex appearing k times in an Euler tour (or $k+1$ times, if it is the starting and finishing vertex and as such counted twice) must have degree $2k$.

[8] Anyone to whom such inspiration seems far-fetched, even after contemplating Figure 1.8.2, may seek consolation in the *multigraph* of Figure 1.10.1.

in common. Given a subspace $\mathcal{F}$ of $\mathcal{E}(G)$, we write

$$\mathcal{F}^\perp := \{ D \in \mathcal{E}(G) \mid \langle F, D \rangle = 0 \text{ for all } F \in \mathcal{F} \} .$$ $\mathcal{F}^\perp$

This is again a subspace of $\mathcal{E}(G)$ (the space of all vectors solving a certain set of linear equations—which?), and we have

$$\dim \mathcal{F} + \dim \mathcal{F}^\perp = m .$$

The *cycle space* $\mathcal{C} = \mathcal{C}(G)$ is the subspace of $\mathcal{E}(G)$ spanned by all the cycles in G—more precisely, by their edge sets.[9] The dimension of $\mathcal{C}(G)$ is the *cyclomatic number* of G. *cycle space* $\mathcal{C}(G)$

Proposition 1.9.1. *The induced cycles in G generate its entire cycle space.* [3.2.3]

Proof. By definition of $\mathcal{C}(G)$ it suffices to show that the induced cycles in G generate every cycle $C \subseteq G$ with a chord e. This follows at once by induction on $|C|$: the two cycles in $C + e$ with e but no other edge in common are shorter than C, and their symmetric difference is precisely C. □

Proposition 1.9.2. *An edge set $F \subseteq E$ lies in $\mathcal{C}(G)$ if and only if every vertex of (V, F) has even degree.* [4.5.1]

Proof. The forward implication holds by induction on the number of cycles needed to generate F, the backward implication by induction on the number of cycles in (V, F). □

If $\{ V_1, V_2 \}$ is a partition of V, the set $E(V_1, V_2)$ of all the edges of G *crossing* this partition is called a *cut*. Recall that for $V_1 = \{ v \}$ this cut is denoted by $E(v)$. *cut*

Proposition 1.9.3. *Together with $\emptyset$, the cuts in G form a subspace $\mathcal{C}^*$ of $\mathcal{E}(G)$. This space is generated by cuts of the form $E(v)$.* [4.6.3]

Proof. Let $\mathcal{C}^*$ denote the set of all cuts in G, together with $\emptyset$. To prove that $\mathcal{C}^*$ is a subspace, we show that for all $D, D' \in \mathcal{C}^*$ also $D + D'$ $(= D - D')$ lies in $\mathcal{C}^*$. Since $D + D = \emptyset \in \mathcal{C}^*$ and $D + \emptyset = D \in \mathcal{C}^*$, we may assume that D and D' are distinct and non-empty. Let $\{ V_1, V_2 \}$ and $\{ V_1', V_2' \}$ be the corresponding partitions of V. Then $D + D'$ consists of all the edges that cross one of these partitions but not the other (Fig. 1.9.1). But these are precisely the edges between $(V_1 \cap V_1') \cup (V_2 \cap V_2')$ and $(V_1 \cap V_2') \cup (V_2 \cap V_1')$, and by $D \neq D'$ these two

[9] For simplicity, we shall not normally distinguish between cycles and their edge sets in connection with the cycle space.

and $G-D$ thus form a minimal cut. By choice of D, this cut is precisely the set $E(C,D)$ of all C–D edges in G.

To prove the lemma, let a partition $\{V_1,V_2\}$ of V be given, and consider a component C of $G[V_1]$. Then $E(C,V_2) = E(C,G-C)$ is the disjoint union of the edge sets $E(C,D)$ over all components D of $G-C$, and is thus the disjoint union of minimal cuts (see above). Now the disjoint union of all these edge sets $E(C,V_2)$, taken over all the components C of $G[V_1]$, is precisely our cut $E(V_1,V_2)$. So this cut is generated by minimal cuts, as claimed. □

Theorem 1.9.5. *The cycle space $\mathcal{C}$ and the cut space $\mathcal{C}^*$ of any graph satisfy*

$$\mathcal{C} = \mathcal{C}^{*\perp} \quad \textit{and} \quad \mathcal{C}^* = \mathcal{C}^{\perp}.$$

Proof. Let us consider a graph $G=(V,E)$. Clearly, any cycle in G has an even number of edges in each cut. This implies $\mathcal{C} \subseteq \mathcal{C}^{*\perp}$.

Conversely, recall from Proposition 1.9.2 that for every edge set $F \notin \mathcal{C}$ there exists a vertex v incident with an odd number of edges in F. Then $\langle E(v),F\rangle = 1$, so $E(v) \in \mathcal{C}^*$ implies $F \notin \mathcal{C}^{*\perp}$. This completes the proof of $\mathcal{C} = \mathcal{C}^{*\perp}$.

To prove $\mathcal{C}^* = \mathcal{C}^{\perp}$, it now suffices to show $\mathcal{C}^* = (\mathcal{C}^{*\perp})^{\perp}$. Here $\mathcal{C}^* \subseteq (\mathcal{C}^{*\perp})^{\perp}$ follows directly from the definition of $\perp$. But since

$$\dim \mathcal{C}^* + \dim \mathcal{C}^{*\perp} = m = \dim \mathcal{C}^{*\perp} + \dim (\mathcal{C}^{*\perp})^{\perp},$$

$\mathcal{C}^*$ has the same dimension as $(\mathcal{C}^{*\perp})^{\perp}$, so $\mathcal{C}^* = (\mathcal{C}^{*\perp})^{\perp}$ as claimed. □

Theorem 1.9.6. *Every connected graph G with n vertices and m edges satisfies* [4.5.1]

$$\dim \mathcal{C}(G) = m-n+1 \quad \textit{and} \quad \dim \mathcal{C}^*(G) = n-1.$$

Proof. Let $G=(V,E)$. As $\dim \mathcal{C} + \dim \mathcal{C}^* = m$ by Theorem 1.9.5, it suffices to find $m-n+1$ linearly independent vectors in $\mathcal{C}$ and $n-1$ linearly independent vectors in $\mathcal{C}^*$: since these numbers add up to m, neither the dimension of $\mathcal{C}$ nor that of $\mathcal{C}^*$ can then be strictly greater. (1.5.1) (1.5.3)

Let T be a spanning tree in G. By Corollary 1.5.3, T has $n-1$ edges, so $m-n+1$ edges of G lie outside T. For each of these $m-n+1$ edges $e \in E \setminus E(T)$, the graph $T+e$ contains a cycle C_e (see Fig. 1.6.3 and Theorem 1.5.1 (iv)). Since none of the edges e lies on $C_{e'}$ for $e' \neq e$, these $m-n+1$ cycles are linearly independent.

For each of the $n-1$ edges $e \in T$, the graph $T-e$ has exactly two components (Theorem 1.5.1 (iii)), and the set D_e of edges in G between these components form a cut (Fig.1.9.3). Since none of the edges $e \in T$ lies in $D_{e'}$ for $e' \neq e$, these $n-1$ cuts are linearly independent. □

1.10 Other notions of graphs

For completeness, we now mention a few other notions of graphs which feature less frequently or not at all in this book.

A *hypergraph* is a pair (V, E) of disjoint sets, where the elements of E are non-empty subsets (of any cardinality) of V. Thus, graphs are special hypergraphs. *hypergraph*

A *directed graph* (or *digraph*) is a pair (V, E) of disjoint sets (of *vertices* and *edges*) together with two maps $\text{init}: E \to V$ and $\text{ter}: E \to V$ assigning to every edge e an *initial vertex* $\text{init}(e)$ and a *terminal vertex* $\text{ter}(e)$. The edge e is said to be *directed from* $\text{init}(e)$ *to* $\text{ter}(e)$. Note that a directed graph may have several edges between the same two vertices x, y. Such edges are called *multiple edges*; if they have the same direction (say from x to y), they are *parallel*. If $\text{init}(e) = \text{ter}(e)$, the edge e is called a *loop*. *directed graph* $\text{init}(e)$ $\text{ter}(e)$ *loop*

A directed graph D is an *orientation* of an (undirected) graph G if $V(D) = V(G)$ and $E(D) = E(G)$, and if $\{\,\text{init}(e), \text{ter}(e)\,\} = \{\,x, y\,\}$ for every edge $e = xy$. Intuitively, such an *oriented graph* arises from an undirected graph simply by directing every edge from one of its ends to the other. Put differently, oriented graphs are directed graphs without loops or multiple edges. *orientation* *oriented graph*

A *multigraph* is a pair (V, E) of disjoint sets (of *vertices* and *edges*) together with a map $E \to V \cup [V]^2$ assigning to every edge either one or two vertices, its *ends*. Thus, multigraphs too can have loops and multiple edges: we may think of a multigraph as a directed graph whose edge directions have been 'forgotten'. To express that x and y are the ends of an edge e we still write $e = xy$, though this no longer determines e uniquely. *multigraph*

A graph is thus essentially the same as a multigraph without loops or multiple edges. Somewhat surprisingly, proving a graph theorem more generally for multigraphs may, on occasion, simplify the proof. Moreover, there are areas in graph theory (such as plane duality; see Chapters 4.6 and 6.5) where multigraphs arise more naturally than graphs, and where any restriction to the latter would seem artificial and be technically complicated. We shall therefore consider multigraphs in these cases, but without much technical ado: terminology introduced earlier for graphs will be used correspondingly.

Two differences, however, should be pointed out. First, a multigraph may have cycles of length 1 or 2: loops, and pairs of multiple edges (or *double edges*). Second, the notion of edge contraction is simpler in multigraphs than in graphs. If we contract an edge $e = xy$ in a multigraph $G = (V, E)$ to a new vertex v_e, there is no longer a need to delete any edges other than e itself: edges parallel to e become loops at v_e, while edges xv and yv become parallel edges between v_e and v (Fig. 1.10.1). Thus, formally, $E(G/e) = E \setminus \{\,e\,\}$, and only the incidence

11.$^-$ Is there a function $f:\mathbb{N}\to\mathbb{N}$ such that, for all $k\in\mathbb{N}$, every graph of minimum degree at least $f(k)$ is k-connected?

12. Let α,β be two graph invariants with positive integer values. Formalize the two statements below, and show that each implies the other:

 (i) α is bounded above by a function of β;

 (ii) β can be forced up by making α large enough.

 Show that the statement

 (iii) β is bounded below by a function of α

 is not equivalent to (i) and (ii). Which small change would make it so?

13.$^+$ What is the deeper reason behind the fact that the proof of Theorem 1.4.2 is based on an assumption of the form $m \geqslant cn - b$ rather than just on a lower bound for the average degree?

14. Prove Theorem 1.5.1.

15. Show that any tree T has at least $\Delta(T)$ leaves.

16. Show that the 'tree-order' associated with a rooted tree T is indeed a partial order on $V(T)$, and verify the claims made about this partial order in the text.

17. Let G be a connected graph, and let $r\in G$ be a vertex. Starting from r, move along the edges of G, going whenever possible to a vertex not visited so far. If there is no such vertex, go back along the edge by which the current vertex was first reached (unless the current vertex is r; then stop). Show that the edges traversed form a normal spanning tree in G with root r.

 (This procedure has earned those trees the name of *depth-first search* trees.)

18. Let $\mathcal{T}$ be a set of subtrees of a tree T. Assume that the trees in $\mathcal{T}$ have pairwise non-empty intersection. Show that their overall intersection $\bigcap\mathcal{T}$ is non-empty.

19. Show that every automorphism of a tree fixes a vertex or an edge.

20. Are the partition classes of a regular bipartite graph always of the same size?

21. Show that a graph is bipartite if and only if every *induced* cycle has even length.

22. Find a function $f:\mathbb{N}\to\mathbb{N}$ such that, for all $k\in\mathbb{N}$, every graph of average degree at least $f(k)$ has a bipartite subgraph of minimum degree at least k.

23. Show that the minor relation $\preccurlyeq$ defines a partial ordering on any set of (finite) graphs. Is the same true for infinite graphs?

24.$^-$ Show that the elements of the cycle space of a graph G are precisely the unions of the edges sets of edge-disjoint cycles in G.

2 Matching

Suppose we are given a graph and are asked to find in it as many independent edges as possible. How should we go about this? Will we be able to pair up all its vertices in this way? If not, how can we be sure that this is indeed impossible? Somewhat surprisingly, this basic problem does not only lie at the heart of numerous applications, it also gives rise to some rather interesting graph theory.

A set M of independent edges in a graph $G = (V, E)$ is called a *matching*. M is a matching *of* $U \subseteq V$ if every vertex in U is incident with an edge in M. The vertices in U are then called *matched* (by M); vertices not incident with any edge of M are *unmatched*. *matching* *matched*

A k-regular spanning subgraph is called a k-*factor*. Thus, a subgraph $H \subseteq G$ is a 1-factor of G if and only if $E(H)$ is a matching of V. The problem of how to characterize the graphs that have a 1-factor, i.e. a matching of their entire vertex set, will be our main theme in this chapter. *factor*

2.1 Matching in bipartite graphs

For this whole section, we let $G = (V, E)$ be a fixed bipartite graph with bipartition $\{A, B\}$. Vertices denoted as a, a' etc. will be assumed to lie in A, vertices denoted as b etc. will lie in B. $G = (V, E)$ A, B a, b *etc.*

How can we find a matching in G with as many edges as possible? Let us start by considering an arbitrary matching M in G. A path in G which starts in A at an unmatched vertex and then contains, alternately, edges from $E \setminus M$ and from M, is an *alternating path* with respect to M. An alternating path P that ends in an unmatched vertex of B is called an *augmenting path* (Fig. 2.1.1), because we can use it to turn M into a larger matching: the symmetric difference of M with $E(P)$ is again a *alternating path* *augmenting path*

Let $ab \in E$ be an edge; we show that either a or b lies in U. If $ab \in M$, this holds by definition of U, so we assume that $ab \notin M$. Since M is a maximal matching, it contains an edge $a'b'$ with $a = a'$ or $b = b'$. In fact, we may assume that $a = a'$: for if a is unmatched (and $b = b'$), then ab is an alternating path, and so the end of $a'b' \in M$ chosen for U was the vertex $b' = b$. Now if $a' = a$ is not in U, then $b' \in U$, and some alternating path P ends in b'. But then there is also an alternating path P' ending in b: either $P' := Pb$ (if $b \in P$) or $P' := Pb'a'b$. By the maximality of M, however, P' is not an augmenting path. So b must be matched, and was chosen for U from the edge of M containing it. $\square$

Let us return to our main problem, the search for some necessary and sufficient conditions for the existence of a 1-factor. In our present case of a bipartite graph, we may as well ask more generally when G contains a matching of A; this will define a 1-factor of G if $|A| = |B|$, a condition that has to hold anyhow if G is to have a 1-factor.

A condition clearly necessary for the existence of a matching of A is that every subset of A has enough neighbours in B, i.e. that

$$|N(S)| \geqslant |S| \qquad \text{for all } S \subseteq A.$$

marriage condition

The following *marriage theorem* says that this obvious necessary condition is in fact sufficient:

marriage theorem

Theorem 2.1.2. (Hall 1935)
G contains a matching of A if and only if $|N(S)| \geqslant |S|$ for all $S \subseteq A$.

We give three proofs for the non-trivial implication of this theorem, i.e. that the 'marriage condition' implies the existence of a matching of A. The first of these is based on König's theorem; the second is a direct constructive proof by augmenting paths; the third will be an independent proof from first principles.

First proof. If G contains no matching of A, then by Theorem 2.1.1 it has a cover U consisting of fewer than $|A|$ vertices, say $U = A' \cup B'$ with $A' \subseteq A$ and $B' \subseteq B$. Then

$$|A'| + |B'| = |U| < |A|,$$

and hence

$$|B'| < |A| - |A'| = |A \setminus A'|$$

(Fig. 2.1.3). By definition of U, however, G has no edges between $A \setminus A'$ and $B \setminus B'$, so

$$|N(A \setminus A')| \leqslant |B'| < |A \setminus A'|$$

and the marriage condition fails for $S := A \setminus A'$. $\square$

Third proof. We apply induction on $|A|$. For $|A| = 1$ the assertion is true. Now let $|A| \geqslant 2$, and assume that the marriage condition is sufficient for the existence of a matching of A when $|A|$ is smaller.

If $|N(S)| \geqslant |S|+1$ for every non-empty set $S \subsetneqq A$, we pick an edge $ab \in G$ and consider the graph $G' := G - \{a,b\}$. Then every non-empty set $S \subseteq A \smallsetminus \{a\}$ satisfies

$$|N_{G'}(S)| \geqslant |N_G(S)| - 1 \geqslant |S|,$$

so by the induction hypothesis G' contains a matching of $A \smallsetminus \{a\}$. Together with the edge ab, this yields a matching of A in G.

Suppose now that A has a non-empty proper subset A' with $|B'| = |A'|$ for $B' := N(A')$. By the induction hypothesis, $G' := G[A' \cup B']$ contains a matching of A'. But $G - G'$ satisfies the marriage condition too: for any set $S \subseteq A \smallsetminus A'$ with $|N_{G-G'}(S)| < |S|$ we would have $|N_G(S \cup A')| < |S \cup A'|$, contrary to our assumption. Again by induction, $G - G'$ contains a matching of $A \smallsetminus A'$. Putting the two matchings together, we obtain a matching of A in G. □ A', B' G'

Corollary 2.1.3. *If $|N(S)| \geqslant |S| - d$ for every set $S \subseteq A$ and some fixed $d \in \mathbb{N}$, then G contains a matching of cardinality $|A| - d$.* [2.2.3]

Proof. We add d new vertices to B, joining each of them to all the vertices in A. By the marriage theorem the new graph contains a matching of A, and at least $|A| - d$ edges in this matching must be edges of G. □

Corollary 2.1.4. *If G is k-regular with $k \geqslant 1$, then G has a 1-factor.*

Proof. If G is k-regular, then clearly $|A| = |B|$; it thus suffices to show by Theorem 2.1.2 that G contains a matching of A. Now every set $S \subseteq A$ is joined to $N(S)$ by a total of $k\,|S|$ edges, and these are among the $k\,|N(S)|$ edges of G incident with $N(S)$. Therefore $k\,|S| \leqslant k\,|N(S)|$, so G does indeed satisfy the marriage condition. □

Despite its seemingly narrow formulation, the marriage theorem counts among the most frequently applied graph theorems, both outside graph theory and within. Often, however, recasting a problem in the setting of bipartite matching requires some clever adaptation. As a simple example, we now use the marriage theorem to derive one of the earliest results of graph theory, a result whose original proof is not all that simple, and certainly not short:

Corollary 2.1.5. (Petersen 1891)
Every regular graph of positive even degree has a 2-factor.

Theorem 2.2.1. (Tutte 1947)
A graph G has a 1-factor if and only if $q(G-S) \leqslant |S|$ for all $S \subseteq V(G)$.

Proof. Let $G = (V, E)$ be a graph without a 1-factor. Our task is to find a *bad set* $S \subseteq V$, one that violates Tutte's condition. V, E *bad set*

We may assume that G is edge-maximal without a 1-factor. Indeed, if G' is obtained from G by adding edges and $S \subseteq V$ is bad for G', then S is also bad for G: any odd component of $G' - S$ is the union of components of $G - S$, and one of these must again be odd.

What does G look like? Clearly, *if* G contains a bad set S then, by its edge-maximality and the trivial forward implication of the theorem,

all the components of $G - S$ are complete and every vertex $s \in S$ is adjacent to all the vertices of $G - s$. $(*)$

But also conversely, if a set $S \subseteq V$ satisfies $(*)$ then either S or the empty set must be bad: if S is not bad we can join the odd components of $G - S$ disjointly to S and pair up all the remaining vertices—unless $|G|$ is odd, in which case $\emptyset$ is bad.

So it suffices to prove that G has a set S of vertices satisfying $(*)$. Let S be the set of vertices that are adjacent to every other vertex. If this set S does not satisfy $(*)$, then some component of $G - S$ has non-adjacent vertices a, a'. Let a, b, c be the first three vertices on a shortest a–a' path in this component; then $ab, bc \in E$ but $ac \notin E$. Since $b \notin S$, there is a vertex $d \in V$ such that $bd \notin E$. By the maximality of G, there is a matching M_1 of V in $G + ac$, and a matching M_2 of V in $G + bd$. S a, b, c d M_1, M_2

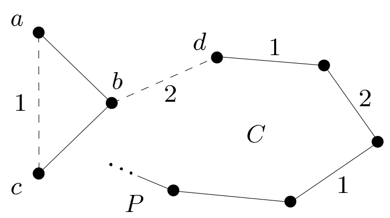

Fig. 2.2.2. Deriving a contradiction if S does not satisfy $(*)$

Let $P = d \dots v$ be a maximal path in G starting at d with an edge from M_1 and containing alternately edges from M_1 and M_2 (Fig. 2.2.2). If the last edge of P lies in M_1, then $v = b$, since otherwise we could continue P. Let us then set $C := P + bd$. If the last edge of P lies in M_2, then by the maximality of P the M_1-edge at v must be ac, so $v \in \{a, c\}$; then let C be the cycle $dPvbd$. In each case, C is an even cycle with every other edge in M_2, and whose only edge not in E is bd. Replacing in M_2 its edges on C with the edges of $C - M_2$, we obtain a matching of V contained in E, a contradiction. $\square$ v

conversely if $|S| = |\mathcal{C}_{G-S}|$, then the existence of a 1-factor follows straight from (i) and (ii).

We now prove the existence of a set S satisfying (i) and (ii). We apply induction on $|G|$. For $|G| = 0$ we may take $S = \emptyset$. Now let G be given with $|G| > 0$, and assume the assertion holds for graphs with fewer vertices.

Let d be the least non-negative integer such that d

$$q(G-T) \leqslant |T| + d \qquad \text{for every } T \subseteq V. \qquad (*)$$

Then there exists a set T for which equality holds in $(*)$: this follows from the minimality of d if $d > 0$, and from $q(G-\emptyset) \geqslant |\emptyset| + 0$ if $d = 0$. Let S be such a set T of maximum cardinality, and let $\mathcal{C} := \mathcal{C}_{G-S}$. $S, \mathcal{C}$

We first show that every component $C \in \mathcal{C}$ is odd. If $|C|$ is even, pick a vertex $c \in C$, and let $S' := S \cup \{c\}$ and $C' := C - c$. Then C' has odd order, and thus has at least one odd component. Hence, $q(G-S') \geqslant q(G-S)+1$. Since $T := S$ satisfies $(*)$ with equality, we obtain

$$q(G-S') \geqslant q(G-S)+1 = |S|+d+1 = |S'|+d \underset{(*)}{\geqslant} q(G-S')$$

with equality, which contradicts the maximality of S.

Next we prove the assertion (ii), that every $C \in \mathcal{C}$ is factor-critical. Suppose there exist $C \in \mathcal{C}$ and $c \in C$ such that $C' := C - c$ has no 1-factor. By the induction hypothesis (and the fact that, as shown earlier, for fixed G our theorem implies Tutte's theorem) there exists a set $T' \subseteq V(C')$ with

$$q(C'-T') > |T'|.$$

Since $|C|$ is odd and hence $|C'|$ is even, the numbers $q(C'-T')$ and $|T'|$ are either both even or both odd, so they cannot differ by exactly 1. We may therefore sharpen the above inequality to

$$q(C'-T') \geqslant |T'|+2.$$

For $T := S \cup \{c\} \cup T'$ we thus obtain

$$\begin{aligned} q(G-T) &= q(G-S) - 1 + q(C'-T') \\ &\geqslant |S| + d - 1 + |T'| + 2 \\ &= |T| + d \\ &\underset{(*)}{\geqslant} q(G-T) \end{aligned}$$

with equality, again contradicting the maximality of S.

It remains to show that S is matchable to $G-S$. If $S = \emptyset$, this is trivial, so we assume that $S \neq \emptyset$. Since $T := S$ satisfies $(*)$ with

2.3 Path covers

Let us return for a moment to König's duality theorem for bipartite graphs, Theorem 2.1.1. If we orient every edge of G from A to B, the theorem tells us how many disjoint directed paths we need in order to cover all the vertices of G: every directed path has length 0 or 1, and clearly the number of paths in such a 'path cover' is smallest when it contains as many paths of length 1 as possible—in other words, when it contains a maximum-cardinality matching.

In this section we put the above question more generally: how many paths in a given directed graph will suffice to cover its entire vertex set? Of course, this could be asked just as well for undirected graphs. As it turns out, however, the result we shall prove is rather more trivial in the undirected case (exercise), and the directed case will also have an interesting corollary.

A *directed path* is a directed graph $P \neq \emptyset$ with distinct vertices $x_0, \dots, x_k$ and edges $e_0, \dots, e_{k-1}$ such that e_i is an edge directed from x_i to x_{i+1}, for all $i < k$. We denote the last vertex x_k of P by $\mathrm{ter}(P)$. $\mathrm{ter}(P)$
In this section, *path* will always mean 'directed path'. A *path cover* of a directed graph G is a set of disjoint paths in G which together contain all the vertices of G. Let us denote the maximum cardinality of an independent set of vertices in G by $\alpha(G)$. *path* *path cover* $\alpha(G)$

Theorem 2.3.1. (Gallai & Milgram 1960)
Every directed graph G has a path cover by at most $\alpha(G)$ paths.

Proof. Given two path covers $\mathcal{P}_1, \mathcal{P}_2$ of a graph, we write $\mathcal{P}_1 < \mathcal{P}_2$ if $\{\, \mathrm{ter}(P) \mid P \in \mathcal{P}_1 \,\} \subseteq \{\, \mathrm{ter}(P) \mid P \in \mathcal{P}_2 \,\}$ and $|\mathcal{P}_1| < |\mathcal{P}_2|$. We shall prove the following: $\mathcal{P}_1 < \mathcal{P}_2$

> *If $\mathcal{P}$ is a $<$-minimal path cover of G, then G contains an independent set $\{\, v_P \mid P \in \mathcal{P} \,\}$ of vertices with $v_P \in P$ for every $P \in \mathcal{P}$.* $(*)$

Clearly, $(*)$ implies the assertion of the theorem.

We prove $(*)$ by induction on $|G|$. Let $\mathcal{P} = \{\, P_1, \dots, P_m \,\}$ be given as in $(*)$, and let $v_i := \mathrm{ter}(P_i)$ for every i. If $\{\, v_i \mid 1 \leqslant i \leqslant m \,\}$ is independent, there is nothing more to show; we may therefore assume that G has an edge from v_2 to v_1. Since $P_2 v_2 v_1$ is again a path, the minimality of $\mathcal{P}$ implies that v_1 is not the only vertex of P_1; let v be the vertex preceding v_1 on P_1. Then $\mathcal{P}' := \{\, P_1 v, P_2, \dots, P_m \,\}$ is a path cover of $G' := G - v_1$ (Fig. 2.3.1). We first show that $\mathcal{P}'$ is $<$-minimal with this property. $\mathcal{P}, P_i, m$ v_i v $\mathcal{P}'$ G'

Suppose that $\mathcal{P}'' < \mathcal{P}'$ is another path cover of G'. If a path $P \in \mathcal{P}''$ ends in v, we may replace P in $\mathcal{P}''$ by $P v v_1$ to obtain a smaller path cover of G than $\mathcal{P}$, a contradiction to the minimality of $\mathcal{P}$. If a path

2. Describe an algorithm that finds, as efficiently as possible, a matching of maximum cardinality in any bipartite graph.

3. Find an infinite counterexample to the statement of the marriage theorem.

4. Let k be an integer. Show that any two partitions of a finite set into k-sets admit a common choice of representatives.

5. Let A be a finite set with subsets $A_1, \dots, A_n$, and let $d_1, \dots, d_n \in \mathbb{N}$. Show that there are disjoint subsets $D_k \subseteq A_k$, with $|D_k| = d_k$ for all $k \leqslant n$, if and only if
$$\Big|\bigcup_{i \in I} A_i\Big| \geqslant \sum_{i \in I} d_i$$
for all $I \subseteq \{1, \dots, n\}$.

6.$^+$ Prove *Sperner's lemma*: in an n-set X there are never more than $\binom{n}{\lfloor n/2 \rfloor}$ subsets such that none of these contains another.
(Hint. Construct $\binom{n}{\lfloor n/2 \rfloor}$ chains covering the power set lattice of X.)

7. Find a set S for Theorem 2.2.3 when G is a forest.

8. Using (only) Theorem 2.2.3, show that a k-connected graph with at least $2k$ vertices contains a matching of size k. Is this best possible?

9. A graph G is called (vertex-) *transitive* if, for any two vertices $v, w \in G$, there is an automorphism of G mapping v to w. Using the observations following the proof of Theorem 2.2.3, show that every transitive connected graph is either factor-critical or contains a 1-factor.
(Hint. Consider the cases of $S = \emptyset$ and $S \neq \emptyset$ separately.)

10. Show that a graph G contains k independent edges if and only if $q(G - S) \leqslant |S| + |G| - 2k$ for all sets $S \subseteq V(G)$.
(Hint. For the 'if' direction, suppose that G has no k independent edges, and apply Tutte's 1-factor theorem to the graph $G * K^{|G|-2k}$. Alternatively, use Theorem 2.2.3.)

11.$^-$ Find a cubic graph without a 1-factor.

12. Derive the marriage theorem from Tutte's theorem.

13.$^-$ Prove the undirected version of the theorem of Gallai & Milgram (without using the directed version).

14. Derive the marriage theorem from the theorem of Gallai & Milgram.

15.$^-$ Show that a partially ordered set of at least $rs + 1$ elements contains either a chain of size $r + 1$ or an antichain of size $s + 1$.

16. Prove the following dual version of Dilworth's theorem: in every finite partially ordered set $(P, \leqslant)$, the minimum number of antichains covering P is equal to the maximum cardinality of a chain in P.

17. Derive König's theorem from Dilworth's theorem.

3 Connectivity

Our definition of k-connectedness, given in Chapter 1.4, is somewhat unintuitive. It does not tell us much about 'connections' in a k-connected graph: all it says is that we need at least k vertices to *dis*connect it. The following definition—which, incidentally, implies the one above—might have been more descriptive: 'a graph is *k-connected* if any two of its vertices can be joined by k independent paths'.

It is one of the classic results of graph theory that these two definitions are in fact equivalent, are dual aspects of the same property. We shall study this theorem of Menger (1927) in some depth in Section 3.3.

In Sections 3.1 and 3.2, we investigate the structure of the 2-connected and the 3-connected graphs. For these small values of k it is still possible to give a simple general description of how these graphs can be constructed.

In the remaining sections of this chapter we look at other concepts of connectedness, more recent than the standard one but no less important: the number of H-paths in a graph for a given subgraph H; the number of edge-disjoint spanning trees; and the existence of disjoint paths linking up several given pairs of vertices.

3.1 2-Connected graphs and subgraphs

A maximal connected subgraph without a cutvertex is called a *block*. *block*
Thus, every block of a graph G is either a maximal 2-connected subgraph, or a bridge (with its ends), or an isolated vertex. Conversely, every such subgraph is a block. By their maximality, different blocks of G overlap in at most one vertex, which is then a cutvertex of G. Hence, every edge of G lies in a unique block, and G is the union of its blocks.

In a sense, blocks are the 2-connected analogues of components, the maximal connected subgraphs of a graph. While the structure of G is

Proof. Clearly, every graph constructed as described is 2-connected. Conversely, let a 2-connected graph G be given. Then G contains a cycle, and hence has a maximal subgraph H constructible as above. H
Since any edge $xy \in E(G) \setminus E(H)$ with $x, y \in H$ would define an H-path, H is an induced subgraph of G. Thus if $H \neq G$, then by the connectedness of G there is an edge vw with $v \in G - H$ and $w \in H$. As G is 2-connected, $G - w$ contains a v–H path P. Then wvP is an H-path in G, and $H \cup wvP$ is a constructible subgraph of G larger than H. This contradicts the maximality of H. □

3.2 The structure of 3-connected graphs

We start this section with the analogue of Proposition 3.1.2 for 3-connectedness: our first theorem describes how every 3-connected graph can be obtained from a K^4 by a succession of elementary operations preserving 3-connectedness. We then prove a deep result of Tutte about the algebraic structure of the cycle space of 3-connected graphs; this will play an important role again in Chapter 4.5.

Lemma 3.2.1. *If G is 3-connected and $|G| > 4$, then G has an edge e such that G/e is again 3-connected.* [4.4.3]

Proof. Suppose there is no such edge e. Then, for every edge $xy \in G$, xy
the graph G/xy contains a separating set S of at most 2 vertices. Since $\kappa(G) \geqslant 3$, the contracted vertex v_{xy} of G/xy (see Chapter 1.7) lies in S and $|S| = 2$, i.e. G has a vertex $z \notin \{x, y\}$ such that $\{v_{xy}, z\}$ z
separates G/xy. Then any two vertices separated by $\{v_{xy}, z\}$ in G/xy are separated in G by $T := \{x, y, z\}$. Since no proper subset of T separates G, every vertex in T has a neighbour in every component C C
of $G - T$.

We choose the edge xy, the vertex z, and the component C so that $|C|$ is as small as possible, and pick a neighbour v of z in C (Fig. 3.2.1). v

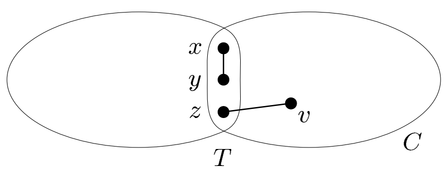

Fig. 3.2.1. Separating vertices in the proof of Lemma 3.2.1

Theorem 3.2.3. (Tutte 1963)
The cycle space of a 3-connected graph is generated by its non-separating induced cycles. [4.5.2]

Proof. We apply induction on the order of the graph G considered. (1.9.1) In K^4, every cycle is a triangle or (in terms of edges) the symmetric difference of triangles. As these are both induced and non-separating, the assertion holds for $|G| = 4$.

For the induction step, let $e = xy$ be an edge of G for which $G' := G/e$ is again 3-connected; cf. Lemma 3.2.1. Then every edge $e' \in E(G') \setminus E(G)$ is of the form $e' = uv_e$, where at least one of the two edges ux and uy lies in G. We pick one that does (either ux or uy), and identify it notationally with the edge e'; thus e' now denotes both the edge uv_e of G' and one of the two edges ux, uy. In this way we may regard $E(G')$ as a subset of $E(G)$, and $\mathcal{E}(G')$ as a subspace of $\mathcal{E}(G)$; thus all vector operations will take place unambiguously in $\mathcal{E}(G)$. $e = xy$ G'

Let us consider an induced cycle $C \subseteq G$. If $e \in C$ and $C = C^3$, we call C a *fundamental triangle*; then $C/e = K^2$. If $e \in C$ but $C \neq C^3$, then C/e is a cycle in G'. Finally if $e \notin C$, then at most one of x, y lies on C (otherwise e would be a chord), so the vertices of C in order also form a cycle in G' if we replace x or y by v_e; this cycle, too, will be denoted by C/e. Thus, as long as C is not a fundamental triangle, C/e will always denote a unique cycle in G'. Note, however, that in the case of $e \notin C$ the edge set of C/e when viewed as a subset of $E(G)$ need not coincide with $E(C)$, or even be a cycle at all; an example is shown in Figure 3.2.3. *fundamental triangle* C/e

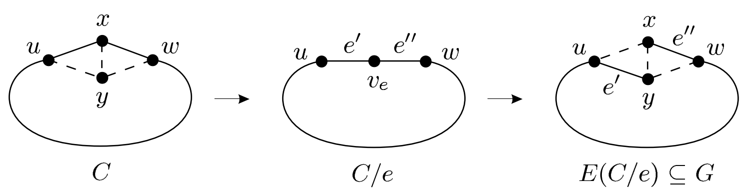

Fig. 3.2.3. One of the four possibilities for $E(C/e)$ when $e \notin C$

Let us refer to the non-separating induced cycles in G or G' as *basic* cycles. An element of $\mathcal{C}(G)$ will be called *good* if it is a linear combination of basic cycles in G; we thus want to show that every element of $\mathcal{C}(G)$ is good. The basic idea of our proof is to contract a given cycle $C \in \mathcal{C}(G)$ to C/e, generate C/e in $\mathcal{C}(G')$ by induction, and try to lift the generators back to basic cycles in G that generate C. *basic cycles* *good*

We start by proving three auxiliary facts.

$$\textit{Every fundamental triangle is a basic cycle in } G. \tag{1}$$

It remains to consider the case that $\{ux, uy, wx, wy\} \not\subseteq E(G)$, say $ux \notin E(G)$. Then, as above, either $uPwyu$ or $uPwxyu$ is a basic cycle in G, according as wy is an edge of G or not. This completes the proof of (3).

We now come to the main part of our proof, the proof that every $C \in \mathcal{C}(G)$ is good. By Proposition 1.9.1 we may assume that C is an induced cycle in G. By (1) we may further assume that C is not a fundamental triangle; so C/e is a cycle. Our aim is to argue as follows. By (2), C differs from C/e at most by some good error term D (and possibly in e); by (3), the basic cycles C'_i of G' that sum to C/e by induction can be contracted from basic cycles of G, which likewise differ from the C'_i only by a good error term D_i (and possibly in e); hence these basic cycles of G and all the error terms together sum to C—except that the edge e will need some special attention. C

By the induction hypothesis, C/e has a representation

$$C/e = C'_1 + \ldots + C'_k$$

$C'_1, \ldots, C'_k$

in $\mathcal{C}(G')$, where every C'_i is a basic cycle in G'. For each i, we obtain from (3) a basic cycle $C(C'_i) \subseteq G$ with $C(C'_i)/e = C'_i$ (in particular, $C(C'_i)$ is not a fundamental triangle), and from (2) some good $D_i \in \mathcal{C}(G)$ such that

$$C(C'_i) + C'_i + D_i \in \{\emptyset, \{e\}\}. \tag{4}$$

We let

$$C_i := C(C'_i) + D_i;$$

$C_1, \ldots, C_k$

then C_i is good, and by (4) it differs from C'_i at most in e. Again by (2), we have

$$C + C/e + D \in \{\emptyset, \{e\}\}$$

for some good $D \in \mathcal{C}(G)$, i.e. $C + D$ differs from C/e at most in e. But then $C + D + C_1 + \ldots + C_k$ differs from $C/e + C'_1 + \ldots + C'_k = \emptyset$ at most in e, that is, D

$$C + D + C_1 + \ldots + C_k \in \{\emptyset, \{e\}\}.$$

Since $C + D + C_1 + \ldots + C_k \in \mathcal{C}(G)$ but $\{e\} \notin \mathcal{C}(G)$, this means that in fact

$$C + D + C_1 + \ldots + C_k = \emptyset,$$

so $C = D + C_1 + \ldots + C_k$ is good. □

Second proof. We show by induction on $|G| + \|G\|$ that G contains k disjoint A–B paths. For all G, A, B with $k \in \{0, 1\}$ this is true. For the induction step let G, A, B with $k \geqslant 2$ be given, and assume that the assertion holds for graphs with fewer vertices or edges.

If there is a vertex $x \in A \cap B$, then $G - x$ contains $k - 1$ disjoint A–B paths by the induction hypothesis. (Why?) Together with the trivial path $\{x\}$, these form the desired paths in G. We shall therefore assume that

$$A \cap B = \emptyset\,. \tag{1}$$

We first construct the desired paths for the case that A and B are separated by a set $X \subseteq V$ with $|X| = k$ and $X \neq A, B$. Let C_A be the union of all the components of $G - X$ meeting A; note that $C_A \neq \emptyset$, since $|A| \geqslant k = |X|$ but $A \neq X$. The subgraph C_B defined likewise is not empty either, and $C_A \cap C_B = \emptyset$. Let us write $G_A := G[V(C_A) \cup X]$ and $G_B := G[V(C_B) \cup X]$. Since every A–B path in G contains an A–X path in G_A, we cannot separate A from X in G_A by fewer than k vertices. Thus, by the induction hypothesis, G_A contains k disjoint A–X paths (Fig. 3.3.2). In the same way, there are k disjoint X–B paths in G_B. As $|X| = k$, we can put these paths together to form k disjoint A–B paths.

X

C_A, C_B

G_A, G_B

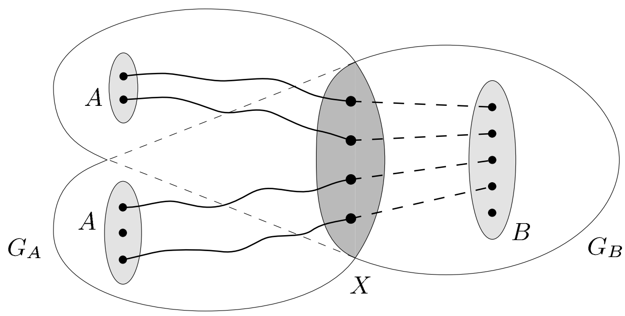

Fig. 3.3.2. Disjoint A–X paths in G_A

For the general case, let P be any A–B path in G. By (1), P has an edge ab with $a \notin B$ and $b \notin A$. Let Y be a set of as few vertices as possible separating A from B in $G - ab$ (Fig. 3.3.3). Then $Y_a := Y \cup \{a\}$ and $Y_b := Y \cup \{b\}$ both separate A from B in G, and by definition of k we have

ab

Y

Y_a, Y_b

$$|Y_a|, |Y_b| \geqslant k\,.$$

If equality holds here, we may assume by the case already treated that $\{Y_a, Y_b\} \subseteq \{A, B\}$, so $\{Y_a, Y_b\} = \{A, B\}$ since $a \notin B$ and $b \notin A$. Thus, $Y = A \cap B$. Since $|Y| \geqslant k - 1 \geqslant 1$, this contradicts (1).

We therefore have either $|Y_a| > k$ or $|Y_b| > k$, and hence $|Y| \geqslant k$. By the induction hypothesis, then, there are k disjoint A–B paths even in $G - ab \subseteq G$. □

Let us consider a walk $W = x_0 e_0 x_1 e_1 \dots e_{n-1} x_n$ from $A \setminus V[\mathcal{P}]$ to $B \setminus V[\mathcal{P}]$, alternating with respect to $\mathcal{P}$. By (ii), any vertex outside $V[\mathcal{P}]$ occurs at most once on W. Since the edges e_i of W are all distinct, (iii) implies that any vertex in $V[\mathcal{P}]$ occurs at most twice on W. This can happen in two ways: if $x_i = x_j$ with $0 < i < j < n$, say, then W, x_i, e_i

$$\begin{aligned} &\textit{either } e_{i-1}, e_j \in E[\mathcal{P}] \text{ and } e_i, e_{j-1} \notin E[\mathcal{P}] \\ &\quad\textit{or } e_i, e_{j-1} \in E[\mathcal{P}] \text{ and } e_{i-1}, e_j \notin E[\mathcal{P}]. \end{aligned}$$

Lemma 3.3.2. *If such a walk W exists, then G contains $|\mathcal{P}|+1$ disjoint A–B paths.*

Proof. Let H be the graph on $V[\mathcal{P}] \cup \{x_0, \dots, x_n\}$ whose edge set is the symmetric difference of $E[\mathcal{P}]$ with $\{e_0, \dots, e_{n-1}\}$. In H, the ends of the paths in $\mathcal{P}$ and of W have degree 1 (or 0, if the path or W is trivial), and all other vertices have degree 0 or 2. For each of the $|\mathcal{P}|+1$ vertices $a \in (A \cap V[\mathcal{P}]) \cup \{x_0\}$, therefore, the component of H containing a is a path, $P = v_0 \dots v_k$ say, which starts in a and ends in A or B. Using conditions (i) and (iii), one easily shows by induction on $i = 0, \dots, k-1$ that P traverses each of its edges $e = v_i v_{i+1}$ in the forward direction with respect to $\mathcal{P}$ or W. (Formally: if $e \in P'$ with $P' \in \mathcal{P}$, then $v_i \in P' \mathring{v}_{i+1}$; if $e = e_j \in W$, then $v_i = x_j$ and $v_{i+1} = x_{j+1}$.) Hence, P ends in B. As we have $|\mathcal{P}|+1$ disjoint such paths P, this completes the proof. □ P

Third proof of Menger's theorem. Let $\mathcal{P}$ be a set of as many disjoint A–B paths in G as possible. Unless otherwise stated, all alternating walks considered are alternating with respect to $\mathcal{P}$. We set $\mathcal{P}$

$$A_1 := A \cap V[\mathcal{P}] \quad \text{and} \quad A_2 := A \setminus A_1 ,$$ A_1, A_2

and

$$B_1 := B \cap V[\mathcal{P}] \quad \text{and} \quad B_2 := B \setminus B_1 .$$ B_1, B_2

For every path $P \in \mathcal{P}$, let x_P be the last vertex of P that lies on some alternating walk starting in A_2; if no such vertex exists, let x_P be the first vertex of P. Clearly, the set x_P

$$X := \{ x_P \mid P \in \mathcal{P} \}$$ X

has cardinality $|\mathcal{P}|$; it thus suffices to show that X separates A from B.

Let Q be any A–B path in G; we show that Q meets X. Suppose not. By the maximality of $\mathcal{P}$, the path Q meets $V[\mathcal{P}]$. Since the A–$V[\mathcal{P}]$ path in Q is trivially an alternating walk, Q also meets the vertex set $V[\mathcal{P}']$ of Q

$$\mathcal{P}' := \{ P x_P \mid P \in \mathcal{P} \};$$ $\mathcal{P}'$

A set of a–B paths is called an *a–B fan* if any two of the paths have only a in common. *fan*

Corollary 3.3.3. *For $B \subseteq V$ and $a \in V \setminus B$, the minimum number of vertices $\neq a$ separating a from B in G is equal to the maximum number of paths forming an a–B fan in G.* [10.1.2]

Proof. Apply Theorem 3.3.1 with $A := N(a)$. □

Corollary 3.3.4. *Let a and b be two distinct vertices of G.*

(i) *If $ab \notin E$, then the minimum number of vertices $\neq a, b$ separating a from b in G is equal to the maximum number of independent a–b paths in G.*

(ii) *The minimum number of edges separating a from b in G is equal to the maximum number of edge-disjoint a–b paths in G.*

Proof. (i) Apply Theorem 3.3.1 with $A := N(a)$ and $B := N(b)$.

(ii) Apply Theorem 3.3.1 to the line graph of G, with $A := E(a)$ and $B := E(b)$. □

Theorem 3.3.5. (Global Version of Menger's Theorem) [4.2.10] [6.6.1] [9.4.2]

(i) *A graph is k-connected if and only if it contains k independent paths between any two vertices.*

(ii) *A graph is k-edge-connected if and only if it contains k edge-disjoint paths between any two vertices.*

Proof. (i) If a graph G contains k independent paths between any two vertices, then $|G| > k$ and G cannot be separated by fewer than k vertices; thus, G is k-connected.

Conversely, suppose that G is k-connected (and, in particular, has more than k vertices) but contains vertices a, b not linked by k independent paths. By Corollary 3.3.4 (i), a and b are adjacent; let $G' := G - ab$. Then G' contains at most $k-2$ independent a–b paths. By Corollary 3.3.4 (i), we can separate a and b in G' by a set X of at most $k-2$ vertices. As $|G| > k$, there is at least one further vertex $v \notin X \cup \{a, b\}$ in G. Now X separates v in G' from either a or b—say, from a. But then $X \cup \{b\}$ is a set of at most $k-1$ vertices separating v from a in G, contradicting the k-connectedness of G. a, b G' X v

(ii) follows straight from Corollary 3.3.4 (ii). □

Now Mader's theorem says that this upper bound is always attained by some set of independent H-paths:

Theorem 3.4.1. (Mader 1978)
Given a graph G with an induced subgraph H, there are always $M_G(H)$ independent H-paths in G.

In order to obtain direct analogues to the vertex and edge version of Menger's theorem, let us consider the two special cases of the above problem where either F or X is required to be empty. Given an induced subgraph $H \subseteq G$, we denote by $\kappa_G(H)$ the least cardinality of a vertex set $X \subseteq V(G-H)$ that meets every H-path in G. Similarly, we let $\lambda_G(H)$ denote the least cardinality of an edge set $F \subseteq E(G)$ that meets every H-path in G. $\kappa_G(H)$ $\lambda_G(H)$

Corollary 3.4.2. *Given a graph G with an induced subgraph H, there are at least $\frac{1}{2}\kappa_G(H)$ independent H-paths and at least $\frac{1}{2}\lambda_G(H)$ edge-disjoint H-paths in G.*

Proof. To prove the first assertion, let k be the maximum number of independent H-paths in G. By Theorem 3.4.1, there are sets $X \subseteq V(G-H)$ and $F \subseteq E(G-H-X)$ with k

$$k = |X| + \sum_{C \in \mathcal{C}_F} \left\lfloor \tfrac{1}{2} |\partial C| \right\rfloor$$

such that every H-path in G has a vertex in X or an edge in F. For every $C \in \mathcal{C}_F$ with $\partial C \neq \emptyset$, pick a vertex $v \in \partial C$ and let $Y_C := \partial C \setminus \{ v \}$; if $\partial C = \emptyset$, let $Y_C := \emptyset$. Then $\left\lfloor \frac{1}{2} |\partial C| \right\rfloor \geqslant \frac{1}{2} |Y_C|$ for all $C \in \mathcal{C}_F$. Moreover, for $Y := \bigcup_{C \in \mathcal{C}_F} Y_C$ every H-path has a vertex in $X \cup Y$. Hence Y

$$k \geqslant |X| + \sum_{C \in \mathcal{C}_F} \tfrac{1}{2} |Y_C| \geqslant \tfrac{1}{2} |X \cup Y| \geqslant \tfrac{1}{2} \kappa_G(H)$$

as claimed.

The second assertion follows from the first by considering the line graph of G (Exercise 16). □

It may come as a surprise to see that the bounds in Corollary 3.4.2 are best possible (as general bounds): one can find examples for G and H where G contains no more than $\frac{1}{2}\kappa_G(H)$ independent H-paths or no more than $\frac{1}{2}\lambda_G(H)$ edge-disjoint H-paths (Exercises 17 and 18).

edges, i.e. such that $\|F\| := \big|E[F]\big|$ with $E[F] := E(F_1) \cup \ldots \cup E(F_k)$ is as large as possible. $E[F]$, $\|F\|$

If $F = (F_1, \ldots, F_k) \in \mathcal{F}$ and $e \in E \setminus E[F]$, then every $F_i + e$ contains a cycle ($i = 1, \ldots, k$): otherwise we could replace F_i by $F_i + e$ in F and obtain a contradiction to the maximality of $\|F\|$. Let us consider an edge $e' \neq e$ of this cycle, for some fixed i. Putting $F'_i := F_i + e - e'$, and $F'_j := F_j$ for all $j \neq i$, we see that $F' := (F'_1, \ldots, F'_k)$ is again in $\mathcal{F}$; we say that F' has been obtained from F by the *replacement* of the edge e' with e. Note that the component of F_i containing e' keeps its vertex set when it changes into a component of F'_i. Hence for every path $x \ldots y \subseteq F'_i$ there is a unique path xF_iy in F_i; this will be used later. *edge replacement* xF_iy

We now consider a fixed k-tuple $F^0 = (F^0_1, \ldots, F^0_k) \in \mathcal{F}$. The set of all k-tuples in $\mathcal{F}$ that can be obtained from F^0 by a series of edge replacements will be denoted by $\mathcal{F}^0$. Finally, we let F^0 $\mathcal{F}^0$

$$E^0 := \bigcup_{F \in \mathcal{F}^0} (E \setminus E[F])$$ E^0

and $G^0 := (V, E^0)$. G^0

Lemma 3.5.3. *For every $e^0 \in E \setminus E[F^0]$ there exists a set $U \subseteq V$ that is connected in every F^0_i ($i = 1, \ldots, k$) and contains the ends of e^0.*

Proof. As $F^0 \in \mathcal{F}^0$, we have $e^0 \in E^0$; let C^0 be the component of G^0 containing e^0. We shall prove the assertion for $U := V(C^0)$. C^0 U

Let $i \in \{1, \ldots, k\}$ be given; we have to show that U is connected in F^0_i. To this end, we first prove the following: i

> *Let $F = (F_1, \ldots, F_k) \in \mathcal{F}^0$, and let $(F'_1, \ldots, F'_k)$ have been obtained from F by the replacement of an edge of F_i. If x, y are the ends of a path in $F'_i \cap C^0$, then also $xF_iy \subseteq C^0$.* (1)

Let $e = vw$ be the new edge in $E(F'_i) \setminus E[F]$; this is the only edge of F'_i not lying in F_i. We assume that $e \in xF'_iy$: otherwise we would have $xF_iy = xF'_iy$ and nothing to show. It suffices to show that $vF_iw \subseteq C^0$: then $(xF'_iy - e) \cup vF_iw$ is a connected subgraph of $F_i \cap C^0$ that contains x, y, and hence also xF_iy. Let e' be any edge of vF_iw. Since we could replace e' in $F \in \mathcal{F}^0$ by e and obtain an element of $\mathcal{F}^0$ not containing e', we have $e' \in E^0$. Thus $vF_iw \subseteq G^0$, and hence $vF_iw \subseteq C^0$ since $v, w \in xF'_iy \subseteq C^0$. This proves (1).

In order to prove that $U = V(C^0)$ is connected in F^0_i we show that, for every edge $xy \in C^0$, the path xF^0_iy exists and lies in C^0. As C^0 is connected, the union of all these paths will then be a connected spanning subgraph of $F^0_i[U]$.

So let $e = xy \in C^0$ be given. As $e \in E^0$, there exist an $s \in \mathbb{N}$ and k-tuples $F^r = (F^r_1, \ldots, F^r_k)$ for $r = 1, \ldots, s$ such that each F^r is obtained from F^{r-1} by edge replacement and $e \in E \setminus E[F^s]$. Setting

Proof. The forward implication was shown above. Conversely, we show that every k-tuple $F = (F_1, \dots, F_k) \in \mathcal{F}$ partitions G, i.e. that $E[F] = E$. If not, let $e \in E \setminus E[F]$. By Lemma 3.5.3, there exists a set $U \subseteq V$ that is connected in every F_i and contains the ends of e. Then $G[U]$ contains $|U| - 1$ edges from each F_i, and in addition the edge e. Thus $\|G[U]\| > k(|U| - 1)$, contrary to our assumption. □ (1.5.3)

The least number of forests forming a partition of a graph G is called the *arboricity* of G. By Theorem 3.5.4, the arboricity is a measure for the maximum local density: a graph has small arboricity if and only if it is 'nowhere dense', i.e. if and only if it has no subgraph H with $\varepsilon(H)$ large. *arboricity*

3.6 Paths between given pairs of vertices

A graph with at least $2k$ vertices is said to be *k-linked* if for every $2k$ distinct vertices $s_1, \dots, s_k, t_1, \dots, t_k$ it contains k disjoint paths $P_1, \dots, P_k$ with $P_i = s_i \dots t_i$ for all i. Thus unlike in Menger's theorem, we are not merely asking for k disjoint paths between two *sets* of vertices: we insist that each of these paths shall link a specified pair of endvertices. *k-linked*

Clearly, every k-linked graph is k-connected. The converse, however, is far from true: being k-linked is generally a much stronger property than k-connectedness. But still, the two properties are related: our aim in this section is to prove the existence of a function $f: \mathbb{N} \to \mathbb{N}$ such that every $f(k)$-connected graph is k-linked.

As a lemma, we need a result that would otherwise belong in Chapter 8:

Theorem 3.6.1. (Mader 1967)
There is a function $h: \mathbb{N} \to \mathbb{N}$ such that every graph with average degree at least $h(r)$ contains K^r as a topological minor, for every $r \in \mathbb{N}$.

Proof. For $r \leqslant 2$, the assertion holds with $h(r) = 1$; we now assume that $r \geqslant 3$. We show by induction on $m = r, \dots, \binom{r}{2}$ that every graph G with average degree $d(G) \geqslant 2^m$ has a topological minor X with r vertices and m edges; for $m = \binom{r}{2}$ this implies the assertion with $h(r) = 2^{\binom{r}{2}}$. (1.2.2) (1.3.1)

If $m = r$ then, by Propositions 1.2.2 and 1.3.1, G contains a cycle of length at least $\varepsilon(G) + 1 \geqslant 2^{r-1} + 1 \geqslant r + 1$, and the assertion follows with $X = C^r$.

Now let $r < m \leqslant \binom{r}{2}$, and assume the assertion holds for smaller m. Let G with $d(G) \geqslant 2^m$ be given; thus, $\varepsilon(G) \geqslant 2^{m-1}$. Since G has a component C with $\varepsilon(C) \geqslant \varepsilon(G)$, we may assume that G is connected. Consider a maximal set $U \subseteq V(G)$ such that U is connected in G and *U*

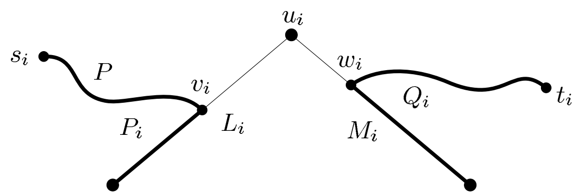

Fig. 3.6.1. Constructing an s_i–t_i path via u_i

In our proof of Theorem 3.6.2 we did not try to find any particularly good bound on the connectivity needed to force a graph to be k-linked; the function f we used grows exponentially in k. Not surprisingly, this is far from being best possible. It is still remarkable, though, that f can in fact be chosen linear: as Bollobás & Thomason (1996) have shown, every $22k$-connected graph is k-linked.

Exercises

For the first three exercises, let G be a graph and $a, b \in V(G)$. Suppose that $X \subseteq V(G) \setminus \{a, b\}$ separates a from b in G. We say that X separates a from b *minimally* if no proper subset of X separates a from b in G.

1.$^-$ Show that X separates a from b minimally if and only if every vertex in X has a neighbour in the component C_a of $G - X$ containing a, and another in the component C_b of $G - X$ containing b.

2. Let $X' \subseteq V(G) \setminus \{a, b\}$ be another set separating a from b, and define C'_a and C'_b correspondingly. Show that both

$$Y_a := (X \cap C'_a) \cup (X \cap X') \cup (X' \cap C_a)$$

and

$$Y_b := (X \cap C'_b) \cup (X \cap X') \cup (X' \cap C_b)$$

separate a from b (see figure).

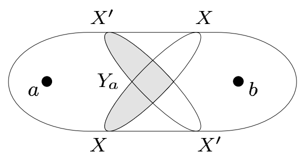

3. Do Y_a and Y_b separate a from b minimally if X and X' do? Are $|Y_a|$ and $|Y_b|$ minimal for vertex sets separating a from b if $|X|$ and $|X'|$ are?

4. Let X and X' be minimal separating vertex sets in G such that X meets at least two components of $G - X'$. Show that X' meets all the components of $G - X$, and that X meets all the components of $G - X'$.

19.$^+$ Derive Tutte's 1-factor theorem (2.2.1) from Mader's theorem.

(Hint. Extend the given graph G to a graph G' by adding, for each vertex $v \in G$, a new vertex v' and joining v' to v. Choose $H \subseteq G'$ so that the 1-factors in G correspond to the large enough sets of independent H-paths in G'.)

20. Find the error in the following short 'proof' of Theorem 3.5.1. Call a partition *non-trivial* if it has at least two classes and at least one of the classes has more than one element. We show by induction on $|V| + |E|$ that $G = (V, E)$ has k edge-disjoint spanning trees if every non-trivial partition of V into r sets (say) has at least $k(r-1)$ cross-edges. The induction starts trivially with $G = K^1$ if we allow k copies of K^1 as a family of k edge-disjoint spanning trees of K^1. We now consider the induction step. If every non-trivial partition of V into r sets (say) has more than $k(r-1)$ cross-edges, we delete any edge of G and are done by induction. So V has a non-trivial partition $\{ V_1, \dots, V_r \}$ with exactly $k(r-1)$ cross-edges. Assume that $|V_1| \geqslant 2$. If $G' := G[V_1]$ has k disjoint spanning trees, we may combine these with k disjoint spanning trees that exist in G/V_1 by induction. We may thus assume that G' has no k disjoint spanning trees. Then by induction it has a non-trivial vertex partition $\{ V'_1, \dots, V'_s \}$ with fewer than $k(s-1)$ cross-edges. Then $\{ V'_1, \dots, V'_s, V_2, \dots, V_r \}$ is a non-trivial vertex partition of G into $r + s - 1$ sets with fewer than $k(r-1) + k(s-1) = k((r+s-1)-1)$ cross-edges, a contradiction.

21.$^-$ Show that every k-linked graph is $(2k-1)$-connected.

Notes

Although connectivity theorems are doubtless among the most natural, and also the most applicable, results in graph theory, there is still no comprehensive monograph on this subject. Some areas are covered in B. Bollobás, *Extremal Graph Theory*, Academic Press 1978, in R. Halin, *Graphentheorie*, Wissenschaftliche Buchgesellschaft 1980, and in A. Frank's chapter of the *Handbook of Combinatorics* (R.L. Graham, M. Grötschel & L. Lovász, eds.), North-Holland 1995. A survey specifically of techniques and results on minimally k-connected graphs (see below) is given by W. Mader, On vertices of degree n in minimally n-connected graphs and digraphs, in (D. Miklós, V.T. Sós & T. Szőnyi, eds.) *Paul Erdős is 80*, Vol. 2, Proc. Colloq. Math. Soc. János Bolyai, Budapest 1996.

Our proof of Tutte's Theorem 3.2.3 is due to C. Thomassen, Planarity and duality of finite and infinite graphs, *J. Combin. Theory B* **29** (1980), 244–271. This paper also contains Lemma 3.2.1 and its short proof from first principles. (The lemma's assertion, of course, follows from Tutte's wheel theorem—its significance lies in its independent proof, which has shortened the proofs of both of Tutte's theorems considerably.)

An approach to the study of connectivity not touched upon in this chapter is the investigation of *minimal* k-connected graphs, those that lose their k-connectedness as soon as we delete an edge. Like all k-connected graphs, these have minimum degree at least k, and by a fundamental result of Halin

4 Planar Graphs

When we draw a graph on a piece of paper, we naturally try to do this as transparently as possible. One obvious way to limit the mess created by all the lines is to avoid intersections. For example, we may ask if we can draw the graph in such a way that no two edges meet in a point other than a common end.

Graphs drawn in this way are called *plane graphs*; abstract graphs that *can* be drawn in this way are called *planar*. In this chapter we study both plane and planar graphs—as well as the relationship between the two: the question of how an abstract graph might be drawn in fundamentally different ways. After collecting together in Section 4.1 the few basic topological facts that will enable us later to prove all results rigorously without too much technical ado, we begin in Section 4.2 by studying the structural properties of plane graphs. In Section 4.3, we investigate how two drawings of the same graph can differ. The main result of that section is that 3-connected planar graphs have essentially only one drawing, in some very strong and natural topological sense. The next two sections are devoted to the proofs of all the classical planarity criteria, conditions telling us when an abstract graph is planar. We complete the chapter with a section on *plane duality*, a notion with fascinating links to algebraic, colouring, and flow properties of graphs (Chapters 1.9 and 6.5).

The traditional notion of a graph drawing is that its vertices are represented by points in the Euclidean plane, its edges are represented by curves between these points, and different curves meet only in common endpoints. To avoid unnecessary topological complication, however, we shall only consider curves that are piecewise linear; it is not difficult to show that any drawing can be straightened out in this way, so the two notions come to the same thing.

With the help of Theorem 4.1.1, it is not difficult to prove the following lemma.

[4.2.5] [4.2.6] [4.2.10]

Lemma 4.1.2. *Let P_1, P_2, P_3 be three arcs, between the same two endpoint but otherwise disjoint.*

(i) *$\mathbb{R}^2 \setminus (P_1 \cup P_2 \cup P_3)$ has exactly three regions, with frontiers $P_1 \cup P_2$, $P_2 \cup P_3$ and $P_1 \cup P_3$.*

(ii) *If P is an arc between a point in $\mathring{P}_1$ and a point in $\mathring{P}_3$ whose interior lies in the region of $\mathbb{R}^2 \setminus (P_1 \cup P_3)$ that contains $\mathring{P}_2$, then $\mathring{P} \cap \mathring{P}_2 \neq \emptyset$.*

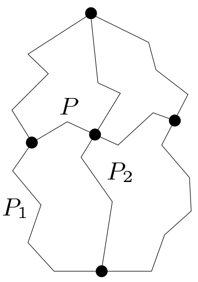

Fig. 4.1.1. The arcs in Lemma 4.1.2 (ii)

Our next lemma complements the Jordan curve theorem by saying that an arc does *not* separate the plane. For easier application later, we phrase this a little more generally:

[4.2.1] [4.2.3]

Lemma 4.1.3. *Let $X_1, X_2 \subseteq \mathbb{R}^2$ be disjoint sets, each the union of finitely many points and arcs, and let P be an arc between a point in X_1 and one in X_2 whose interior $\mathring{P}$ lies in a region O of $\mathbb{R}^2 \setminus (X_1 \cup X_2)$. Then $O \setminus \mathring{P}$ is a region of $\mathbb{R}^2 \setminus (X_1 \cup P \cup X_2)$.*

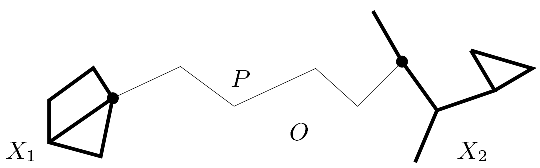

Fig. 4.1.2. P does not separate the region O of $\mathbb{R}^2 \setminus (X_1 \cup X_2)$

It remains to introduce a few terms and facts that will be used only once, when we consider notions of equivalence for graph drawings in Chapter 4.3.

As usual, we denote by S^n the n-dimensional sphere, the set of points in $\mathbb{R}^{n+1}$ at distance 1 from the origin. The 2-sphere minus its 'north pole' $(0,0,1)$ is homeomorphic to the plane; let us choose a fixed such homeomorphism $\pi \colon S^2 \setminus \{(0,0,1)\} \to \mathbb{R}^2$ (for example, stereographic projection). If $P \subseteq \mathbb{R}^2$ is a polygon and O is the bounded region of

S^n π

Lemma 4.2.1. *Let G be a plane graph and e an edge of G.* [4.5.1] [4.5.2]

(i) *If X is the frontier of a face of G, then either $e \subseteq X$ or $X \cap \mathring{e} = \emptyset$.*

(ii) *If e lies on a cycle $C \subseteq G$, then e lies on the frontier of exactly two faces of G, and these are contained in distinct faces of C.*

(iii) *If e lies on no cycle, then e lies on the frontier of exactly one face of G.*

Proof. We prove all three assertions together. Let us start by considering one point $x_0 \in \mathring{e}$. We show that x_0 lies on the frontier of either exactly two faces or exactly one, according as e lies on a cycle in G or not. We then show that every other point in $\mathring{e}$ lies on the frontier of exactly the same faces as x_0. Then the endpoints of e will also lie on the frontier of these faces—simply because every neighbourhood of an endpoint of e is also the neighbourhood of an inner point of e. (4.1.1) (4.1.3)

G is the union of finitely many straight line segments; we may assume that any two of these intersect in at most one point. Around every point $x \in \mathring{e}$ we can find an open disc D_x, with centre x, which meets only those (one or two) straight line segments that contain x. D_x

Let us pick an inner point x_0 from a straight line segment $S \subseteq e$. Then $D_{x_0} \cap G = D_{x_0} \cap S$, so $D_{x_0} \setminus G$ is the union of two open half-discs. Since these half-discs do not meet G, they each lie in a face of G. Let us denote these faces by f_1 and f_2; they are the only faces of G with x_0 on their frontier, and they may coincide (Fig. 4.2.1). x_0, S f_1, f_2

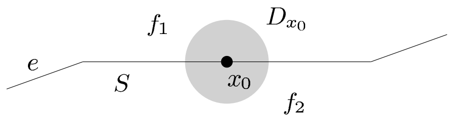

Fig. 4.2.1. Faces f_1, f_2 of G in the proof of Lemma 4.2.1

If e lies on a cycle $C \subseteq G$, then D_{x_0} meets both faces of C (Theorem 4.1.1). The faces f_1, f_2 of G are therefore contained in distinct faces of C—since $C \subseteq G$, every face of G is a subset of a face of C—and in particular $f_1 \neq f_2$. If e does not lie on any cycle, then e is a bridge and thus links two disjoint point sets X_1, X_2 with $X_1 \cup X_2 = G \setminus \mathring{e}$. Clearly, $f_1 \cup \mathring{e} \cup f_2$ is the subset of a face f of $G - e$. (Why?) By Lemma 4.1.3, $f \setminus \mathring{e}$ is a face of G. But $f \setminus \mathring{e}$ contains f_1 and f_2 by definition of f, so $f_1 = f \setminus \mathring{e} = f_2$ since f_1, f_2 and f are all faces of G.

Now consider any other point $x_1 \in \mathring{e}$. Let P be the arc from x_0 to x_1 contained in e. Since P is compact, finitely many of the discs D_x with $x \in P$ cover P. Let us enumerate these discs as $D_0, \dots, D_n$ in the natural order of their centres along P; adding D_{x_0} or D_{x_1} as necessary, we may assume that $D_0 = D_{x_0}$ and $D_n = D_{x_1}$. By induction on n, one easily proves that every point $y \in D_n \setminus e$ can be linked by an arc inside x_1 P $D_0, \dots, D_n$ y

Proof. Let f be a face in a 2-connected plane graph G. We show by induction on $|G|$ that $G[f]$ is a cycle. If G is itself a cycle, this holds by Theorem 4.1.1; we therefore assume that G is not a cycle. (3.1.2) (4.1.1) (4.1.2)

By Proposition 3.1.2, there exist a 2-connected plane graph $H \subseteq G$ and a plane H-path P such that $G = H \cup P$. The interior of P lies in a face f' of H, which by the induction hypothesis is bounded by a cycle C. *H* *P* *f′, C*

If f is also a face of H, we are home by the induction hypothesis. If not, then the frontier of f meets $P \setminus H$, so $f \subseteq f'$. Therefore f is a face of $C \cup P$, and is hence bounded by a cycle (Lemma 4.1.2 (i)). □

A plane graph G is called *maximally plane*, or just *maximal*, if we cannot add a new edge to form a plane graph $G' \supsetneq G$ with $V(G') = V(G)$. We call G a *plane triangulation* if every face of G (including the outer face) is bounded by a triangle. *maximal plane graph* *plane triangulation*

Proposition 4.2.6. *A plane graph of order at least* 3 *is maximally plane if and only if it is a plane triangulation.* [4.4.1] [5.4.2]

Proof. Let G be a plane graph of order at least 3. It is easy to see that if every face of G is bounded by a triangle, then G is maximally plane. Indeed, any additional edge e would have its interior inside a face of G and its ends on the boundary of that face. Hence these ends are already adjacent in G, so $G \cup e$ cannot satisfy condition (iii) in the definition of a plane graph. (4.1.2)

Conversely, assume that G is maximally plane and let $f \in F(G)$ be a face; let us write $H := G[f]$. Since G is maximal as a plane graph, $G[H]$ is complete: any two vertices of H that are not already adjacent in G could be linked by an arc through f, extending G to a larger plane graph. Thus $G[H] = K^n$ for some n—but we do not know yet which edges of $G[H]$ lie in H. *f* *H* *n*

Let us show first that H contains a cycle. If not, then $G \setminus H \neq \emptyset$: by $G \supseteq K^n$ if $n \geqslant 3$, or else by $|G| \geqslant 3$. On the other hand we have $f \cup H = \mathbb{R}^2$ by Proposition 4.2.3 and hence $G = H$, a contradiction.

Since H contains a cycle, it suffices to show that $n \leqslant 3$: then $H = K^3$ as claimed. Suppose $n \geqslant 4$, and let $C = v_1v_2v_3v_4v_1$ be a cycle in $G[H]$ ($= K^n$). By $C \subseteq G$, our face f is contained in a face f_C of C; let f'_C be the other face of C. Since the vertices v_1 and v_3 lie on the boundary of f, they can be linked by an arc whose interior lies in f_C and avoids G. *C, v_i* *f_C, f'_C*

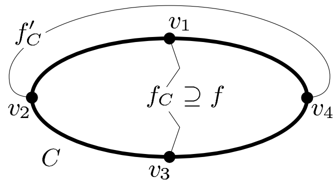

Fig. 4.2.3. The edge v_2v_4 of G runs through the face f'_C

and hence $f' = f_{1,2} = f'_{1,2}$); let x be a point in $f' \setminus f_{1,2}$. Then x lies in some face $f \neq f_1, f_2$ of G. As shown above, f is also a face of G'. Hence $x \in f \cap f'$ implies $f = f'$, and we have $f' \in F(G) \setminus \{ f_1, f_2 \}$ as desired. □ x f

[4.4.1] [5.1.2] [8.3.5]

Corollary 4.2.8. *A plane graph with $n \geqslant 3$ vertices has at most $3n-6$ edges. Every plane triangulation with n vertices has $3n-6$ edges.*

Proof. By Proposition 4.2.6 it suffices to prove the second assertion. In a plane triangulation G, every face boundary contains exactly three edges, and every edge lies on the boundary of exactly two faces (Lemma 4.2.1). The bipartite graph on $E(G) \cup F(G)$ with edge set $\{ ef \mid e \subseteq G[f] \}$ thus has exactly $2|E(G)| = 3|F(G)|$ edges. According to this identity we may replace ℓ with $2m/3$ in Euler's formula, and obtain $m = 3n-6$. □

Euler's formula can be useful for showing that certain graphs cannot occur as plane graphs. The graph K^5, for example, has $10 > 3 \cdot 5 - 6$ edges, more than allowed by Corollary 4.2.8. Similarly, $K_{3,3}$ cannot be a plane graph. For since $K_{3,3}$ is 2-connected but contains no triangle, every face of a plane $K_{3,3}$ would be bounded by a cycle of length $\geqslant 4$ (Proposition 4.2.5). As in the proof of Corollary 4.2.8 this implies $2m \geqslant 4\ell$, which yields $m \leqslant 2n-4$ when substituted in Euler's formula. But $K_{3,3}$ has $9 > 2 \cdot 6 - 4$ edges.

Clearly, along with K^5 and $K_{3,3}$ themselves, their subdivisions cannot occur as plane graphs either:

Corollary 4.2.9. *A plane graph contains neither K^5 nor $K_{3,3}$ as a topological minor.* □ [4.4.5] [4.4.6]

Surprisingly, it turns out that this simple property of plane graphs identifies them among all other graphs: as Section 4.4 will show, an arbitrary graph can be drawn in the plane if and only if it has no (topological) K^5 or $K_{3,3}$ minor.

As we have seen, every face boundary in a 2-connected plane graph is a cycle. In a 3-connected graph, these cycles can be identified combinatorially:

Proposition 4.2.10. *The face boundaries in a 3-connected plane graph are precisely its non-separating induced cycles.* [4.3.2] [4.5.2]

(3.3.5) (4.1.1) (4.1.2)

Proof. Let G be a 3-connected plane graph, and let $C \subseteq G$. If C is a non-separating induced cycle, then by the Jordan curve theorem its two faces cannot both contain points of $G \setminus C$. Therefore it bounds a face of G.

Conversely, suppose that C bounds a face f. By Proposition 4.2.5, C is a cycle. If C has a chord $e = xy$, then the components of $C - \{ x, y \}$ C, f

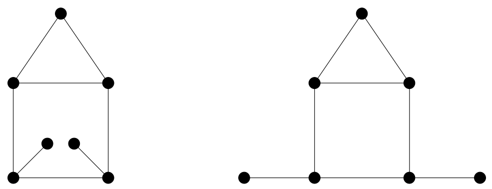

Fig. 4.3.1. Two drawings of a graph that are not topologically isomorphic—why not?

face of $\psi(G)$. (To ensure that the edges of $\psi(G)$ are again piecewise linear, however, one may have to adjust φ a little.)

If σ is a topological isomorphism as above, then—except possibly for a pair of missing points where ψ or ψ^{-1} is undefined—ψ maps the faces of G onto those of G' (proof?). In this way, σ extends naturally to a bijection $\sigma: V \cup E \cup F \to V' \cup E' \cup F'$ which preserves incidence of vertices, edges and faces.

Let us single out this last property of a topological isomorphism as the defining property for our second notion of equivalence for plane graphs: let us call our given isomorphism σ between the abstract graphs G and G' a *combinatorial isomorphism* of the plane graphs G and G' if it can be extended to a bijection $\sigma: V \cup E \cup F \to V' \cup E' \cup F'$ that preserves incidence not only of vertices with edges but also of vertices and edges with faces. (Formally: we require that a vertex or edge $x \in G$ shall lie on the boundary of a face $f \in F$ if and only if $\sigma(x)$ lies on the boundary of the face $\sigma(f)$.)

combinatorial isomorphism

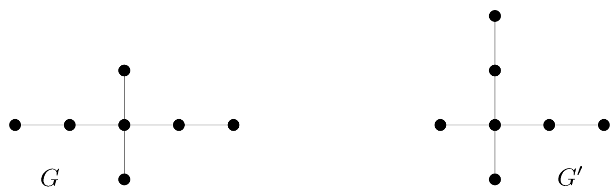

Fig. 4.3.2. Two drawings of a graph that are combinatorially isomorphic but not topologically—why not?

If σ is a combinatorial isomorphism of the plane graphs G and G', it maps the face boundaries of G to those of G'. Let us raise this property to our third definition of equivalence for plane graphs: we call our isomorphism σ of the abstract graphs G and G' a *graph-theoretical isomorphism* of the plane graphs G and G' if

graph-theoretical isomorphism

$$\{\, \sigma(G\,[\,f\,]) : f \in F \,\} = \{\, G'\,[\,f'\,] : f' \in F' \,\}\,.$$

Thus, we no longer keep track of *which* face is bounded by a given subgraph: the only information we keep is whether a subgraph bounds

$$\begin{array}{ccc} S^2 \supseteq \tilde{G} & \xrightarrow{\tilde{\sigma}} & \tilde{G}' \subseteq S^2 \\ \pi \big\downarrow & & \big\downarrow \pi \\ \mathbb{R}^2 \supseteq G & \xrightarrow[\sigma]{} & G' \subseteq \mathbb{R}^2 \end{array}$$

Fig. 4.3.3. Defining $\tilde{\sigma}$ via σ

edge, as follows. Every edge xy of $\tilde{G}$ is homeomorphic to the edge $\tilde{\sigma}(xy) = \varphi(x)\varphi(y)$ of $\tilde{G}'$, by a homeomorphism mapping x to $\varphi(x)$ and y to $\varphi(y)$. Then the union of all these homeomorphisms, one for every edge of $\tilde{G}$, is indeed a homeomorphism between $\tilde{G}$ and $\tilde{G}'$—our desired extension of φ to $\tilde{G}$: all we have to check is continuity at the vertices (where the edge homeomorphisms overlap), and this follows at once from our assumption that the two graphs and their individual edges all carry the subspace topology in $\mathbb{R}^3$.

In the third step we now extend our homeomorphism $\varphi: \tilde{G} \to \tilde{G}'$ to all of S^2. This can be done analogously to the second step, face by face. By Proposition 4.2.5, all face boundaries in $\tilde{G}$ and $\tilde{G}'$ are cycles. Now if f is a face of $\tilde{G}$ and C its boundary, then $\tilde{\sigma}(C) := \bigcup\{\, \tilde{\sigma}(e) \mid e \in E(C)\,\}$ bounds the face $\tilde{\sigma}(f)$ of $\tilde{G}'$. By Theorem 4.1.4, we may therefore extend the homeomorphism $\varphi: C \to \tilde{\sigma}(C)$ defined so far to a homeomorphism from $C \cup f$ to $\tilde{\sigma}(C) \cup \tilde{\sigma}(f)$. We finally take the union of all these homeomorphisms, one for every face f of $\tilde{G}$, as our desired homeomorphism $\varphi: S^2 \to S^2$; as before, continuity is easily checked. □

So far, we have considered ways of comparing plane graphs. We now come to our actual goal, the definition of equivalence for planar embeddings. Let us call two planar embeddings σ_1, σ_2 of a graph G *topologically* (respectively, *combinatorially*) *equivalent* if $\sigma_2 \circ \sigma_1^{-1}$ is a topological (respectively, combinatorial) isomorphism between $\sigma_1(G)$ and $\sigma_2(G)$. If G is 2-connected, the two definitions coincide by Theorem 4.3.1, and we simply speak of *equivalent* embeddings. Clearly, this is indeed an equivalence relation on the set of planar embeddings of any given graph. *equivalent embeddings*

Note that two drawings of G resulting from inequivalent embeddings may well be topologically isomorphic (exercise): for the equivalence of two embeddings we ask not only that some (topological or combinatorial) isomorphism exist between the their images, but that the canonical isomorphism $\sigma_2 \circ \sigma_1^{-1}$ be a topological or combinatorial one.

Even in this strong sense, 3-connected graphs have only one embedding up to equivalence:

Theorem 4.3.2. (Whitney 1932)
Any two planar embeddings of a 3-connected graph are equivalent.

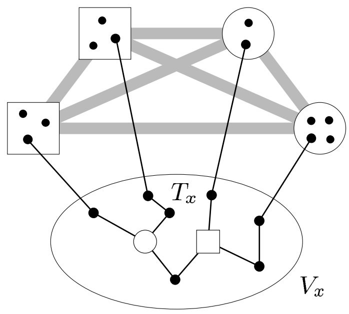

Fig. 4.4.1. Every MK^5 contains a TK^5 or $MK_{3,3}$

of K, T_x has exactly 4 leaves, the 4 neighbours of V_x in other branch sets (Fig. 4.4.1).

If each of the five trees T_x is a $TK_{1,4}$ then K is a TK^5, and we are done. If one of the T_x is not a $TK_{1,4}$ then it has exactly two vertices of degree 3. Contracting V_x onto these two vertices, and every other branch set to a single vertex, we obtain a graph on 6 vertices containing a $K_{3,3}$. Thus, $G \succcurlyeq K_{3,3}$ as desired. □

We first prove Kuratowski's theorem for 3-connected graphs. This is the heart of the proof: the general case will then follow easily.

Lemma 4.4.3. *Every 3-connected graph G without a K^5 or $K_{3,3}$ minor is planar.*

Proof. We apply induction on $|G|$. For $|G| = 4$ we have $G = K^4$, and the assertion holds. Now let $|G| > 4$, and assume the assertion is true for smaller graphs. By Lemma 3.2.1, G has an edge xy such that G/xy is again 3-connected. Since the minor relation is transitive, G/xy has no K^5 or $K_{3,3}$ minor either. Thus, by the induction hypothesis, G/xy has a drawing $\tilde{G}$ in the plane. Let f be the face of $\tilde{G} - v_{xy}$ containing the point v_{xy}, and let C be the boundary of f. Let $X := N_G(x) \setminus \{ y \}$ and $Y := N_G(y) \setminus \{ x \}$; then $X \cup Y \subseteq V(C)$, because $v_{xy} \in f$. Clearly, (3.2.1) (4.2.5) xy $\tilde{G}$ f, C X, Y

$$\tilde{G}' := \tilde{G} - \{ v_{xy}v \mid v \in Y \setminus X \}$$ $\tilde{G}'$

may be viewed as a drawing of $G - y$, in which the vertex x is represented by the point v_{xy} (Fig. 4.4.2). Our aim is to add y to this drawing to obtain a drawing of G.

Since $\tilde{G}$ is 3-connected, $\tilde{G} - v_{xy}$ is 2-connected, so C is a cycle (Proposition 4.2.5). Let $x_1, \dots, x_k$ be an enumeration along this cycle of the vertices in X, and let $P_i = x_i \dots x_{i+1}$ be the X-paths on C between them ($i = 1, \dots, k$; with $x_{k+1} := x_1$). For each i, the set $C \setminus P_i$ is contained in one of the two faces of the cycle $C_i := x x_i P_i x_{i+1} x$; we $x_1, \dots, x_k$ P_i C_i

The following more combinatorial route is just as easy, and may be a welcome alternative.

Lemma 4.4.4. *Let $\mathcal{X}$ be a set of 3-connected graphs. Let G be a graph with $\kappa(G) \leqslant 2$, and let G_1, G_2 be proper induced subgraphs of G such that $G = G_1 \cup G_2$ and $|G_1 \cap G_2| = \kappa(G)$. If G is edge-maximal without a topological minor in $\mathcal{X}$, then so are G_1 and G_2, and $G_1 \cap G_2 = K^2$.* [8.3.1]

Proof. Note first that every vertex $v \in S := V(G_1 \cap G_2)$ has a neighbour in every component of $G_i - S$, $i = 1, 2$: otherwise $S \setminus \{v\}$ would separate G, contradicting $|S| = \kappa(G)$. By the maximality of G, every edge e added to G lies in a $TX \subseteq G + e$ with $X \in \mathcal{X}$. For all the choices of e considered below, the 3-connectedness of X will imply that the branch vertices of this TX all lie in the same G_i, say in G_1. (The position of e will always be symmetrical with respect to G_1 and G_2, so this assumption entails no loss of generality.) Then the TX meets G_2 at most in a path P corresponding to an edge of X. S X P

If $S = \emptyset$, we obtain an immediate contradiction by choosing e with one end in G_1 and the other in G_2. If $S = \{v\}$ is a singleton, let e join a neighbour v_1 of v in $G_1 - S$ to a neighbour v_2 of v in $G_2 - S$ (Fig. 4.4.3). Then P contains both v and the edge v_1v_2; replacing vPv_1 with the edge vv_1, we obtain a TX in $G_1 \subseteq G$, a contradiction.

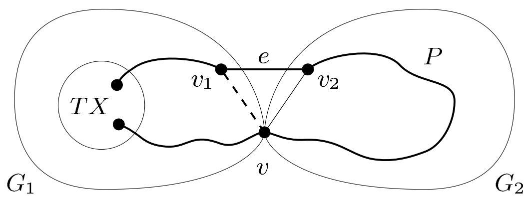

Fig. 4.4.3. If $G + e$ contains a TX, then so does G_1 or G_2

So $|S| = 2$, say $S = \{x, y\}$. If $xy \notin G$, we let $e := xy$, and in the arising TX replace e by an x–y path through G_2; this gives a TX in G, a contradiction. Hence $xy \in G$, and $G[S] = K^2$ as claimed. x, y

It remains to show that G_1 and G_2 are edge-maximal without a topological minor in $\mathcal{X}$. So let e' be an additional edge for G_1, say. Replacing xPy with the edge xy if necessary, we obtain a TX either in $G_1 + e'$ (which shows the edge-maximality of G_1, as desired) or in G_2 (which contradicts $G_2 \subseteq G$). □

Lemma 4.4.5. *If $|G| \geqslant 4$ and G is edge-maximal with $TK^5, TK_{3,3} \not\subseteq G$, then G is 3-connected.*

4.5 Algebraic planarity criteria

In this section we show that planarity can be characterized in purely algebraic terms, by a certain abstract property of its cycle space. Theorems relating such seemingly distant graph properties are rare, and their significance extends beyond their immediate applicability. In a sense, they indicate that both ways of viewing a graph—in our case, the topological and the algebraic way—are not just formal curiosities: if both are natural enough that, quite unexpectedly, each can be expressed in terms of the other, the indications are that they have the power to reveal some genuine insights into the structure of graphs and are worth pursuing.

Let $G = (V, E)$ be a graph. We call a subset $\mathcal{F}$ of its edge space $\mathcal{E}(G)$ *simple* if every edge of G lies in at most two sets of $\mathcal{F}$. For example, the cut space $\mathcal{C}^*(G)$ has a simple basis: according to Proposition 1.9.3 it is generated by the cuts $E(v)$ formed by all the edges at a given vertex v, and an edge $xy \in G$ lies in $E(v)$ only for $v = x$ and for $v = y$. *simple*

Theorem 4.5.1. (MacLane 1937)
A graph is planar if and only if its cycle space has a simple basis. [4.6.3]

Proof. The assertion being trivial for graphs of order at most 2, we consider a graph G of order at least 3. If $\kappa(G) \leqslant 1$, then G is the union of two proper induced subgraphs G_1, G_2 with $|G_1 \cap G_2| \leqslant 1$. Then $\mathcal{C}(G)$ is the direct sum of $\mathcal{C}(G_1)$ and $\mathcal{C}(G_2)$, and hence has a simple basis if and only if both $\mathcal{C}(G_1)$ and $\mathcal{C}(G_2)$ do (proof?). Moreover, G is planar if and only if both G_1 and G_2 are: this follows at once from Kuratowski's theorem, but also from easy geometrical considerations. The assertion for G thus follows inductively from those for G_1 and G_2. For the rest of the proof, we now assume that G is 2-connected. (1.9.2) (1.9.6) (4.1.1) (4.2.1) (4.2.5) (4.4.6)

We first assume that G is planar and choose a drawing. By Lemma 4.2.5, the face boundaries of G are cycles, so they are elements of $\mathcal{C}(G)$. We shall show that the face boundaries generate all the cycles in G; then $\mathcal{C}(G)$ has a simple basis by Lemma 4.2.1. Let $C \subseteq G$ be any cycle, and let f be its inner face. By Lemma 4.2.1, every edge e with $\mathring{e} \subseteq f$ lies on exactly two face boundaries $G[f']$ with $f' \subseteq f$, and every edge of C lies on exactly one such face boundary. Hence the sum in $\mathcal{C}(G)$ of all those face boundaries is exactly C.

Conversely, let $\{C_1, \ldots, C_k\}$ be a simple basis of $\mathcal{C}(G)$. Then, for every edge $e \in G$, also $\mathcal{C}(G - e)$ has a simple basis. Indeed, if e lies in just one of the sets C_i, say in C_1, then $\{C_2, \ldots, C_k\}$ is a simple basis of $\mathcal{C}(G - e)$; if e lies in two of the C_i, say in C_1 and C_2, then $\{C_1 + C_2, C_3, \ldots, C_k\}$ is such a basis. (Note that the two bases are indeed subsets of $\mathcal{C}(G - e)$ by Proposition 1.9.2.) Thus every subgraph of G has a cycle space with a simple basis. For our proof that G is planar, it thus suffices to show that the cycle spaces of K^5 and $K_{3,3}$ (and hence

4.6 Plane duality

In this section we shall use MacLane's theorem to uncover another connection between planarity and algebraic structure: a connection between the duality of plane graphs, defined below, and the duality of the cycle and cut space hinted at in Chapters 1.9 and 3.5.

plane multigraph

A *plane multigraph* is a pair $G = (V, E)$ of finite sets (of *vertices* and *edges*, respectively) satisfying the following conditions:

(i) $V \subseteq \mathbb{R}^2$;

(ii) every edge is either an arc between two vertices or a polygon containing exactly one vertex (its *endpoint*);

(iii) apart from its own endpoint(s), an edge contains no vertex and no point of any other edge.

We shall use terms defined for plane graphs freely for plane multigraphs. Note that, as in abstract multigraphs, both loops and double edges count as cycles.

Let us consider the plane multigraph G shown in Figure 4.6.1. Let us place a new vertex inside each face of G and link these new vertices up to form another plane multigraph G^*, as follows: for every edge e of G we link the two new vertices in the faces incident with e by an edge e^* crossing e; if e is incident with only one face, we attach a loop e^* to the new vertex in that face, again crossing the edge e. The plane multigraph G^* formed in this way is then dual to G in the following sense: if we apply the same procedure as above to G^*, we obtain a plane multigraph very similar to G; in fact, G itself may be reobtained from G^* in this way.

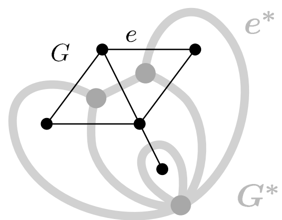

Fig. 4.6.1. A plane graph and its dual

plane dual G^*

To make this idea more precise, let $G = (V, E)$ and (V^*, E^*) be any two plane multigraphs, and put $F(G) =: F$ and $F((V^*, E^*)) =: F^*$. We call (V^*, E^*) a *plane dual* of G, and write $(V^*, E^*) =: G^*$, if there are bijections

$$\begin{array}{ccc} F \to V^* & E \to E^* & V \to F^* \\ f \mapsto v^*(f) & e \mapsto e^* & v \mapsto f^*(v) \end{array}$$

satisfying the following conditions:

(i) $v^*(f) \in f$ for all $f \in F$;

Proposition 4.6.2. *If G^* is an abstract dual of G, then the cut space of G^* is the cycle space of G, i.e.*

$$\mathcal{C}^*(G^*) = \mathcal{C}(G).$$

Proof. By Lemma 1.9.4,[5] $\mathcal{C}^*(G^*)$ is the subspace of $\mathcal{E}(G^*) = \mathcal{E}(G)$ generated by the minimal cuts in G^*. By assumption, these are precisely the edge sets of the cycles in G, and these generate $\mathcal{C}(G)$ in $\mathcal{E}(G)$. □ (1.9.4)

We finally come to one of the highlights of classical planarity theory: the planar graphs are characterized by the fact that they have an abstract dual. Although less obviously intuitive, this duality is just as fundamental a property as planarity itself; indeed the following theorem may well be seen as a topological characterization of the graphs that have a dual:

Theorem 4.6.3. (Whitney 1933)
A graph is planar if and only if it has an abstract dual.

Proof. Let G be a graph. If G is plane, then every component C of G has a plane dual C^*. Let us consider these C^* as abstract multigraphs, pick a vertex in each of them, and identify these vertices. In the connected multigraph G^* obtained, the set of minimal cuts is the union of the sets of minimal cuts in the multigraphs C^*. By Proposition 4.6.1, these cuts are precisely the edge sets of the cycles in G, so G^* is an abstract dual of G. (1.9.3) (4.5.1)

Conversely, suppose that G has an abstract dual G^*. By Theorem 4.5.1 and Proposition 4.6.2 it suffices to show that $\mathcal{C}^*(G^*)$ has a simple basis, which it has by Proposition 1.9.3. □

Exercises

1. Show that every graph can be embedded in $\mathbb{R}^3$ with all edges straight.
2.$^-$ Show directly by Lemma 4.1.2 that $K_{3,3}$ is not planar.
3.$^-$ Find an Euler formula for disconnected graphs.
4. Show that every connected planar graph with n vertices, m edges and finite girth g satisfies $m \leqslant \frac{g}{g-2}(n-2)$.
5. Show that every planar graph is a union of three forests.

[5] Although the lemma was stated for graphs only, its proof remains the same for multigraphs.

19. Prove the general Kuratowski theorem from its 3-connected case by manipulating plane graphs, i.e. avoiding Lemma 4.4.5.

    (This is not intended as an exercise in elementary topology; for the topological parts of the proof, a rough sketch will do.)

20. A graph is called *outerplanar* if it has a drawing in which every vertex lies on the boundary of the outer face. Show that a graph is outerplanar if and only if it contains neither K^4 nor $K_{2,3}$ as a minor.

21. Let $G = G_1 \cup G_2$, where $|G_1 \cap G_2| \leqslant 1$. Show that $\mathcal{C}(G)$ has a simple basis if both $\mathcal{C}(G_1)$ and $\mathcal{C}(G_2)$ have one.

22.$^+$ Find a cycle space basis among the face boundaries of a 2-connected plane graph.

23. Show that a 2-connected plane graph is bipartite if and only if every face is bounded by an even cycle.

24.$^-$ Let G be a connected plane multigraph, and let G^* be its plane dual. Prove the following two statements for every edge $e \in G$:

    (i) If e lies on the boundary of two distinct faces f_1, f_2 of G, then $e^* = v^*(f_1)\, v^*(f_2)$.

    (ii) If e lies on the boundary of exactly one face f of G, then e^* is a loop at $v^*(f)$.

25.$^-$ What does the plane dual of a plane tree look like?

26.$^-$ Show that the plane dual of a plane multigraph is connected.

27.$^+$ Show that a plane multigraph has a plane dual if and only if it is connected.

28. Let G, G^* be mutually dual plane multigraphs, and let $e \in E(G)$. Prove the following statements (with a suitable definition of G/e):

    (i) If e is not a bridge, then G^*/e^* is a plane dual of $G - e$.

    (ii) If e is not a loop, then $G^* - e^*$ is a plane dual of G/e.

29. Show that any two plane duals of a plane multigraph are combinatorially isomorphic.

30. Let G, G^* be mutually dual plane graphs. Prove the following statements:

    (i) If G is 2-connected, then G^* is 2-connected.

    (ii) If G is 3-connected, then G^* is 3-connected.

    (iii) If G is 4-connected, then G^* need not be 4-connected.

31. Let G, G^* be mutually dual plane graphs. Let $B_1, \dots, B_n$ be the blocks of G. Show that $B_1^*, \dots, B_n^*$ are the blocks of G^*.

32. Show that if G^* is an abstract dual of a multigraph G, then G is an abstract dual of G^*.

a planar embedding is actually constructed from a simple basis, is adopted in K. Wagner, *Graphentheorie*, BI Hochschultaschenbücher 1972.

The proper setting for duality phenomena between cuts and cycles in abstract graphs (and beyond) is the theory of *matroids*; see J.G. Oxley, *Matroid Theory*, Oxford University Press 1992.

5 Colouring

How many colours do we need to colour the countries of a map in such a way that adjacent countries are coloured differently? How many days have to be scheduled for committee meetings of a parliament if every committee intends to meet for one day and some members of parliament serve on several committees? How can we find a school timetable of minimum total length, based on the information of how often each teacher has to teach each class?

A *vertex colouring* of a graph $G = (V, E)$ is a map $c: V \to S$ such that $c(v) \neq c(w)$ whenever v and w are adjacent. The elements of the set S are called the available *colours*. All that interests us about S is its size: typically, we shall be asking for the smallest integer k such that G has a *k-colouring*, a vertex colouring $c: V \to \{1, \dots, k\}$. This k is the *(vertex-) chromatic number* of G; it is denoted by $\chi(G)$. A graph G with $\chi(G) = k$ is called *k-chromatic*; if $\chi(G) \leqslant k$, we call G *k-colourable*.

vertex colouring

chromatic number $\chi(G)$

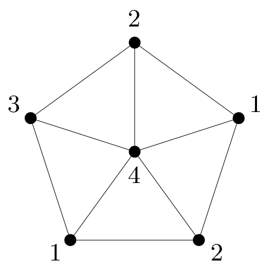

Fig. 5.0.1. A vertex colouring $V \to \{1, \dots, 4\}$

Note that a k-colouring is nothing but a vertex partition into k independent sets, now called *colour classes*; the non-trivial 2-colourable graphs, for example, are precisely the bipartite graphs. Historically, the colouring terminology comes from the map colouring problem stated

colour classes

Let D be an open disc around v, so small that it meets only those five straight edge segments of G that contain v. Let us enumerate these segments according to their cyclic position in D as $s_1, \dots, s_5$, and let vv_i be the edge containing s_i ($i = 1, \dots, 5$; Fig. 5.1.1). Without loss of generality we may assume that $c(v_i) = i$ for each i. D $s_1, \dots, s_5$ $v_1, \dots, v_5$

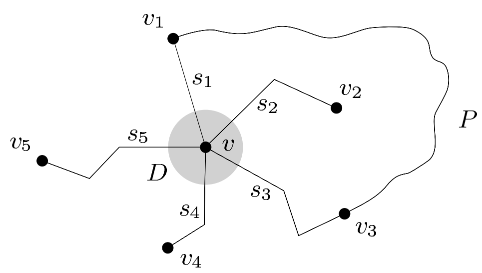

Fig. 5.1.1. The proof of the five colour theorem

Let us show first that every v_1–v_3 path $P \subseteq H$ separates v_2 from v_4 in H. Clearly, this is the case if and only if the cycle $C := vv_1Pv_3v$ separates v_2 from v_4 in G. We prove this by showing that v_2 and v_4 lie in different faces of C. P C

Consider the two regions of $D \smallsetminus (s_1 \cup s_3)$. One of these regions meets s_2, the other s_4. Since $C \cap D \subseteq s_1 \cup s_3$, the two regions are each contained within a face of C. Moreover, these faces are distinct: otherwise, D would meet only one face of C, contrary to the fact that v lies on the boundary of both faces (Theorem 4.1.1). Thus $D \cap s_2$ and $D \cap s_4$ lie in distinct faces of C. As C meets the edges $vv_2 \supseteq s_2$ and $vv_4 \supseteq s_4$ only in v, the same holds for v_2 and v_4.

Given $i, j \in \{1, \dots, 5\}$, let $H_{i,j}$ be the subgraph of H induced by the vertices coloured i or j. We may assume that the component C_1 of $H_{1,3}$ containing v_1 also contains v_3. Indeed, if we interchange the colours 1 and 3 at all the vertices of C_1, we obtain another 5-colouring of H; if $v_3 \notin C_1$, then v_1 and v_3 are both coloured 3 in this new colouring, and we may assign colour 1 to v. Thus, $H_{1,3}$ contains a v_1–v_3 path P. As shown above, P separates v_2 from v_4 in H. Since $P \cap H_{2,4} = \emptyset$, this means that v_2 and v_4 lie in different components of $H_{2,4}$. In the component containing v_2, we now interchange the colours 2 and 4, thus recolouring v_2 with colour 4. Now v no longer has a neighbour coloured 2, and we may give it this colour. □ $H_{i,j}$

As a backdrop to the two famous theorems above, let us cite another well-known result:

Theorem 5.1.3. (Grötzsch 1959)
Every planar graph not containing a triangle is 3-colourable.

Proposition 5.2.2. *Every graph G satisfies*

$$\chi(G) \leqslant \operatorname{col}(G) = \max\{\, \delta(H) \mid H \subseteq G \,\} + 1 .$$

□

[9.2.1] [9.2.3] [11.2.3]

Corollary 5.2.3. *Every graph G has a subgraph of minimum degree at least $\chi(G) - 1$.* □

The colouring number of a graph is closely related to its arboricity; see the remark following Theorem 3.5.4.

As we have seen, every graph G satisfies $\chi(G) \leqslant \Delta(G) + 1$, with equality for complete graphs and odd cycles. In all other cases, this general bound can be improved a little:

Theorem 5.2.4. (Brooks 1941)
Let G be a connected graph. If G is neither complete nor an odd cycle, then

$$\chi(G) \leqslant \Delta(G) .$$

Proof. We apply induction on $|G|$. If $\Delta(G) \leqslant 2$, then G is a path or a cycle, and the assertion is trivial. We therefore assume that $\Delta :=$ Δ
$\Delta(G) \geqslant 3$, and that the assertion holds for graphs of smaller order. Suppose that $\chi(G) > \Delta$.

Let $v \in G$ be a vertex and $H := G - v$. Then $\chi(H) \leqslant \Delta$: by v, H
induction, every component H' of H satisfies $\chi(H') \leqslant \Delta(H') \leqslant \Delta$ unless H' is complete or an odd cycle, in which case $\chi(H') = \Delta(H') + 1 \leqslant \Delta$ as every vertex of H' has maximum degree in H' and one such vertex is also adjacent to v in G.

Since H can be Δ-coloured but G cannot, we have the following:

> *Every Δ-colouring of H uses all the colours $1, \dots, \Delta$ on the neighbours of v; in particular, $d(v) = \Delta$.* (1)

Given any Δ-colouring of H, let us denote the neighbour of v coloured i by v_i, $i = 1, \dots, \Delta$. For all $i \neq j$, let $H_{i,j}$ denote the subgraph $v_1, \dots, v_\Delta$
of H spanned by all the vertices coloured i or j. $H_{i,j}$

> *For all $i \neq j$, the vertices v_i and v_j lie in a common component $C_{i,j}$ of $H_{i,j}$.* (2) $C_{i,j}$

Otherwise we could interchange the colours i and j in one of those components; then v_i and v_j would be coloured the same, contrary to (1).

> *$C_{i,j}$ is always a v_i–v_j path.* (3)

Indeed, let P be a v_i–v_j path in $C_{i,j}$. As $d_H(v_i) \leqslant \Delta - 1$, the neighbours of v_i have pairwise different colours: otherwise we could recolour v_i, So are those somewhat denser subgraphs the ‘cause’ for the large value of χ? Do they, in turn, necessarily contain a graph of high chromatic number—maybe even one from some small collection of *canonical* such graphs, such as K^k? Interestingly, this is not so: those subgraphs of large but ‘constant’ average degree—bounded below only by a function of k, not of $|G|$—are not nearly dense enough to contain (necessarily) any particular graph of high chromatic number, let alone K^k.[1]

Yet even if the above local structures do not appear to help, it might still be the case that, somehow, a high chromatic number forces the existence of certain canonical highly chromatic subgraphs. That this is in fact not the case will be our main result in Chapter 11: according to a classic result of Erdős, proved by probabilistic methods, *there are graphs of arbitrarily large chromatic number and yet arbitrarily large girth* (Theorem 11.2.2). Thus given any graph H that is not a forest, for every $k \in \mathbb{N}$ there are graphs G with $\chi(G) \geqslant k$ but $H \not\subseteq G$.[2]

Thus, contrary to our initial guess that a large chromatic number might always be caused by some dense local substructure, it can in fact occur as a purely global phenomenon: after all, locally (around each vertex) a graph of large girth looks just like a tree, and is in particular 2-colourable there!

So far, we asked what a high chromatic number implies: it forces the invariants δ, d, Δ and κ up in some subgraph, but it does not imply the existence of any concrete subgraph of large chromatic number. Let us now consider the converse question: from what assumptions could we deduce that the chromatic number of a given graph is large?

Short of a concrete subgraph known to be highly chromatic (such as K^k), there is little or nothing in sight: no values of the invariants studied so far imply that the graph considered has a large chromatic number. (Recall the example of $K_{n,n}$.) So what exactly can cause high chromaticity as a global phenomenon largely remains a mystery!

Nevertheless, there exists a simple—though not always short—procedure to construct all the graphs of chromatic number $\geqslant k$. For each $k \in \mathbb{N}$, let us define the class of *k-constructible* graphs recursively as follows: *k-constructible*

(i) K^k is k-constructible.

(ii) If G is k-constructible and $x, y \in V(G)$ are non-adjacent, then also $(G + xy)/xy$ is k-constructible.

[1] This is obvious from the examples of $K_{n,n}$, which are 2-chromatic but whose connectivity and average degree n exceeds any constant bound. Which (non-constant) average degree exactly will force the existence of a given subgraph will be the topic of Chapter 7.

[2] By Corollaries 5.2.3 and 1.5.4, of course, every graph of sufficiently high chromatic number will contain any given forest.

each of these identifications amounts to a construction step of type (ii). Eventually, we obtain the graph

$$(H_1 \cup H_2) - xy_1 - xy_2 + y_1y_2 \subseteq G\,;$$

this is the desired k-constructible subgraph of G. □

5.3 Colouring edges

Clearly, every graph G satisfies $\chi'(G) \geqslant \Delta(G)$. For bipartite graphs, we have equality here:

Proposition 5.3.1. (König 1916)
Every bipartite graph G satisfies $\chi'(G) = \Delta(G)$.

Proof. We apply induction on $\|G\|$. For $\|G\| = 0$ the assertion holds. (1.6.1)
Now assume that $\|G\| \geqslant 1$, and that the assertion holds for graphs with fewer edges. Let $\Delta := \Delta(G)$, pick an edge $xy \in G$, and choose a Δ-edge-colouring of $G - xy$ by the induction hypothesis. Let us refer to the edges coloured α as *α-edges*, etc. Δ, xy α-edge

In $G - xy$, each of x and y is incident with at most $\Delta - 1$ edges. Hence there are $\alpha, \beta \in \{1, \ldots, \Delta\}$ such that x is not incident with an α-edge and y is not incident with a β-edge. If $\alpha = \beta$, we can colour the edge xy with this colour and are done; so we may assume that $\alpha \neq \beta$, and that x is incident with a β-edge. α, β

Let us extend this edge to a maximal walk W whose edges are coloured β and α alternately. Since no such walk contains a vertex twice (why not?), W exists and is a path. Moreover, W does not contain y: if it did, it would end in y on an α-edge (by the choice of β) and thus have even length, so $W + xy$ would be an odd cycle in G (cf. Proposition 1.6.1). We now recolour all the edges on W, swapping α with β. By the choice of α and the maximality of W, adjacent edges of $G - xy$ are still coloured differently. We have thus found a Δ-edge-colouring of $G - xy$ in which neither x nor y is incident with a β-edge. Colouring xy with β, we extend this colouring to a Δ-edge-colouring of G. □

Theorem 5.3.2. (Vizing 1964)
Every graph G satisfies

$$\Delta(G) \leqslant \chi'(G) \leqslant \Delta(G) + 1\,.$$

Proof. We prove the second inequality by induction on $\|G\|$. For $\|G\| = 0$ it is trivial. For the induction step let $G = (V, E)$ with $\Delta := \Delta(G) > 0$ be V, E Δ

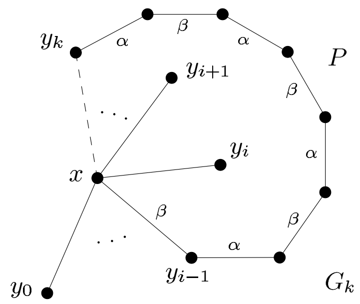

Fig. 5.3.1. The α/β-path P in G_k

Vizing's theorem divides the finite graphs into two classes according to their chromatic index; graphs satisfying $\chi' = \Delta$ are called (imaginatively) *class 1*, those with $\chi' = \Delta + 1$ are *class 2*.

5.4 List colouring

In this section, we take a look at a relatively recent generalization of the concepts of colouring studied so far. This generalization may seem a little far-fetched at first glance, but it turns out to supply a fundamental link between the classical (vertex and edge) chromatic numbers of a graph and its other invariants.

Suppose we are given a graph $G = (V, E)$, and for each vertex of G a list of colours permitted at that particular vertex: when can we colour G (in the usual sense) so that each vertex receives a colour from its list? More formally, let $(S_v)_{v \in V}$ be a family of sets. We call a vertex colouring c of G with $c(v) \in S_v$ for all $v \in V$ a colouring *from the lists* S_v. The graph G is called *k-list-colourable*, or *k-choosable*, if, for every family $(S_v)_{v \in V}$ with $|S_v| = k$ for all v, there is a vertex colouring of G from the lists S_v. The least integer k for which G is k-choosable is the *list-chromatic number*, or *choice number* $\mathrm{ch}(G)$ of G. *k-choosable* *choice number* $\mathrm{ch}(G)$

List-colourings of edges are defined analogously. The least integer k such that G has an edge colouring from any family of lists of size k is the *list-chromatic index* $\mathrm{ch}'(G)$ of G; formally, we just set $\mathrm{ch}'(G) := \mathrm{ch}(L(G))$, where $L(G)$ is the line graph of G. $\mathrm{ch}'(G)$

In principle, showing that a given graph is k-choosable is more difficult than proving it to be k-colourable: the latter is just the special case of the former where all lists are equal to $\{1, \ldots, k\}$. Thus,

$$\mathrm{ch}(G) \geqslant \chi(G) \quad \text{and} \quad \mathrm{ch}'(G) \geqslant \chi'(G)$$

for all graphs G.

Let us check first that (∗) implies the assertion of the theorem. Let any plane graph be given, together with a list of 5 colours for each vertex. Add edges to this graph until it is a maximal plane graph G. By Proposition 4.2.6, G is a plane triangulation; let $v_1v_2v_3v_1$ be the boundary of its outer face. We now colour v_1 and v_2 (differently) from their lists, and extend this colouring by (∗) to a colouring of G from the lists given.

Let us now prove (∗), by induction on $|G|$. If $|G| = 3$, then $G = C$ and the assertion is trivial. Now let $|G| \geqslant 4$, and assume (∗) for smaller graphs. If C has a chord vw, then vw lies on two unique cycles vw $C_1, C_2 \subseteq C + vw$ with $v_1v_2 \in C_1$ and $v_1v_2 \notin C_2$. For $i = 1, 2$, let G_i denote the subgraph of G induced by the vertices lying on C_i or in its inner face (Fig. 5.4.1). Applying the induction hypothesis first to G_1 and then—with the colours now assigned to v and w—to G_2 yields the desired colouring of G.

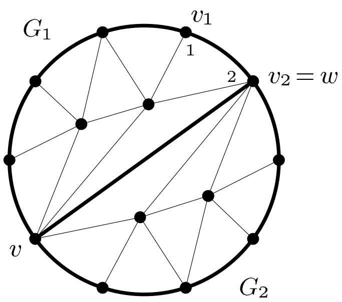

Fig. 5.4.1. The induction step with a chord vw; here the case of $w = v_2$

If C has no chord, let $v_1, u_1, \dots, u_m, v_{k-1}$ be the neighbours of v_k in $u_1, \dots, u_m$ their natural cyclic order order around v_k;[3] by definition of C, all those neighbours u_i lie in the inner face of C (Fig. 5.4.2). As the inner faces

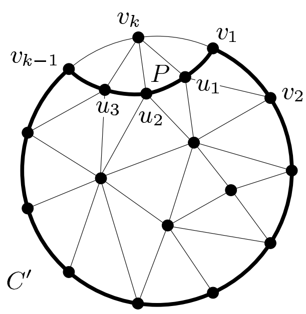

Fig. 5.4.2. The induction step without a chord

[3] as in the first proof of the five colour theorem

if, for every vertex $v \in D - U$, there is an edge in D directed from v to a vertex in U. Note that kernels of non-empty directed graphs are themselves non-empty.

Lemma 5.4.3. *Let H be a graph and $(S_v)_{v \in V(H)}$ a family of lists. If H has an orientation D with $d^+(v) < |S_v|$ for every v, and such that every induced subgraph of D has a kernel, then H can be coloured from the lists S_v.*

Proof. We apply induction on $|H|$. For $|H| = 0$ we take the empty colouring. For the induction step, let $|H| > 0$. Let α be a colour occurring in one of the lists S_v, and let D be an orientation of H as stated. The vertices v with $\alpha \in S_v$ span a non-empty subgraph D' in D; by assumption, D' has a kernel $U \neq \emptyset$. α D' U

Let us colour the vertices in U with α, and remove α from the lists of all the other vertices of D'. Since each of those vertices sends an edge to U, the modified lists S'_v for $v \in D - U$ again satisfy the condition $d^+(v) < |S'_v|$ in $D - U$. Since $D - U$ is an orientation of $H - U$, we can thus colour $H - U$ from those lists by the induction hypothesis. As none of these lists contains α, this extends our colouring $U \to \{\alpha\}$ to the desired list colouring of H. □

Theorem 5.4.4. (Galvin 1995)
Every bipartite graph G satisfies $\mathrm{ch}'(G) = \chi'(G)$.

Proof. Let $G =: (X \cup Y, E)$, where $\{X, Y\}$ is a vertex bipartition of G. Let us say that two edges of G *meet in X* if they share an end in X, and correspondingly for Y. Let $\chi'(G) =: k$, and let c be a k-edge-colouring of G. X, Y, E k c

Clearly, $\mathrm{ch}'(G) \geqslant k$; we prove that $\mathrm{ch}'(G) \leqslant k$. Our plan is to use Lemma 5.4.3 to show that the line graph H of G is k-choosable. To apply the lemma, it suffices to find an orientation D of H with $d^+(v) < k$ for every vertex v, and such that every induced subgraph of D has a kernel. To define D, consider adjacent $e, e' \in E$, say with $c(e) < c(e')$. If e and e' meet in X, we orient the edge $ee' \in H$ from e' towards e; if e and e' meet in Y, we orient it from e to e' (Fig 5.4.3). H D

Let us compute $d^+(e)$ for given $e \in E = V(D)$. If $c(e) = i$, say, then every $e' \in N^+(e)$ meeting e in X has its colour in $\{1, \dots, i-1\}$, and every $e' \in N^+(e)$ meeting e in Y has its colour in $\{i+1, \dots, k\}$. As any two neighbours e' of e meeting e either both in X or both in Y are themselves adjacent and hence coloured differently, this implies $d^+(e) < k$ as desired.

It remains to show that every induced subgraph D' of D has a kernel. We show this by induction on $|D'|$. For $D' = \emptyset$, the empty set is a kernel; so let $|D'| \geqslant 1$. Let $E' := V(D') \subseteq E$. For every $x \in X$ at which E' has an edge, let $e_x \in E'$ be the edge at x with minimum D' E'

A graph is called *perfect* if every induced subgraph $H \subseteq G$ has chromatic number $\chi(H) = \omega(H)$, i.e. if the trivial lower bound of $\omega(H)$ colours always suffices to colour the vertices of H. Thus, while proving an assertion of the form $\chi(G) > k$ may in general be difficult, even in principle, for a given graph G, it can always be done for a perfect graph simply by exhibiting some K^{k+1} subgraph as a 'certificate' for non-colourability with k colours. *perfect*

At first glance, the structure of the class of perfect graphs appears somewhat contrived: although it is closed under induced subgraphs (if only by explicit definition), it is not closed under taking general subgraphs or supergraphs, let alone minors (examples?). However, perfection is an important notion in graph theory: the fact that several fundamental classes of graphs are perfect (as if by fluke) may serve as a superficial indication of this.[4]

What graphs, then, are perfect? Bipartite graphs are, for instance. Less trivially, the complements of bipartite graphs are perfect, too—a fact equivalent to König's duality theorem 2.1.1 (Exercise 34). The so-called *comparability graphs* are perfect, and so are the *interval graphs* (see the exercises); both these turn up in numerous applications.

In order to study at least one such example in some detail, we prove here that the chordal graphs are perfect: a graph is *chordal* (or *triangulated*) if each of its cycles of length at least 4 has a chord, i.e. if it contains no induced cycles other than triangles. *chordal*

To show that chordal graphs are perfect, we shall first characterize their structure. If G is a graph with induced subgraphs G_1, G_2 and S, such that $G = G_1 \cup G_2$ and $S = G_1 \cap G_2$, we say that G arises from G_1 and G_2 by *pasting* these graphs together along S. *pasting*

Proposition 5.5.1. *A graph is chordal if and only if it can be constructed recursively by pasting along complete subgraphs, starting from complete graphs.* [12.3.11]

Proof. If G is obtained from two chordal graphs G_1, G_2 by pasting them together along a complete subgraph, then G is clearly again chordal: any induced cycle in G lies in either G_1 or G_2, and is hence a triangle by assumption. Since complete graphs are chordal, this proves that all graphs constructible as stated are chordal.

Conversely, let G be a chordal graph. We show by induction on $|G|$ that G can be constructed as described. This is trivial if G is complete. We therefore assume that G is not complete, in particular $|G| > 1$, and that all smaller chordal graphs are constructible as stated. Let $a, b \in G$ *a, b*

[4] The class of perfect graphs has duality properties with deep connections to optimization and complexity theory, which are far from understood. Theorem 5.5.5 shows the tip of an iceberg here; for more, the reader is referred to Lovász's survey cited in the notes.

We shall give two proofs of Theorem 5.5.3. The first of these is Lovász's original proof, which is still unsurpassed in its clarity and the amount of 'feel' for the problem it conveys. Our second proof, due to Gasparian (1996), is in fact a very short and elegant linear algebra proof of another theorem of Lovász's (Theorem 5.5.5), which easily implies Theorem 5.5.3.

Let us prepare our first proof of the perfect graph theorem by a lemma. Let G be a graph and $x \in G$ a vertex, and let G' be obtained from G by adding a vertex x' and joining it to x and all the neighbours of x. We say that G' is obtained from G by *expanding* the vertex x to an edge xx' (Fig. 5.5.2). *expanding a vertex*

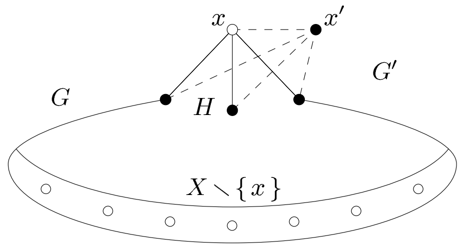

Fig. 5.5.2. Expanding the vertex x in the proof of Lemma 5.5.4

Lemma 5.5.4. *Any graph obtained from a perfect graph by expanding a vertex is again perfect.*

Proof. We use induction on the order of the perfect graph considered. Expanding the vertex of K^1 yields K^2, which is perfect. For the induction step, let G be a non-trivial perfect graph, and let G' be obtained from G by expanding a vertex $x \in G$ to an edge xx'. For our proof that G' is perfect it suffices to show $\chi(G') \leqslant \omega(G')$: every proper induced subgraph H of G' is either isomorphic to an induced subgraph of G or obtained from a proper induced subgraph of G by expanding x; in either case, H is perfect by assumption and the induction hypothesis, and can hence be coloured with $\omega(H)$ colours. x, x'

Let $\omega(G) =: \omega$; then $\omega(G') \in \{\omega, \omega+1\}$. If $\omega(G') = \omega+1$, then ω

$$\chi(G') \leqslant \chi(G) + 1 = \omega + 1 = \omega(G')$$

and we are done. So let us assume that $\omega(G') = \omega$. Then x lies in no $K^\omega \subseteq G$: together with x', this would yield a $K^{\omega+1}$ in G'. Let us colour G with ω colours. Since every $K^\omega \subseteq G$ meets the colour class X of x but not x itself, the graph $H := G - (X \setminus \{x\})$ has clique number $\omega(H) < \omega$ (Fig. 5.5.2). Since G is perfect, we may thus colour H with $\omega - 1$ colours. Now X is independent, so the set $(X \setminus \{x\}) \cup \{x'\} = V(G' - H)$ is also independent. We can therefore extend our $(\omega-1)$-colouring of H to an ω-colouring of G', showing that $\chi(G') \leqslant \omega = \omega(G')$ as desired. □ X H

the last inequality follows from the fact that $|X \cap A_K| \leqslant 1$ for all K (since A_K is independent but $G[X]$ is complete), and $|X \cap A_X| = 0$ (by the choice of A_X). On the other hand,

$$\begin{aligned} |G'| &= \sum_{x \in V} k(x) \\ &= \big|\{ (x,K) : x \in V,\ K \in \mathcal{K},\ x \in A_K \}\big| \\ &= \sum_{K \in \mathcal{K}} |A_K| \\ &= |\mathcal{K}| \cdot \alpha \,. \end{aligned}$$

As $\alpha(G') \leqslant \alpha$ by construction of G', this implies

$$\chi(G') \geqslant \frac{|G'|}{\alpha(G')} \geqslant \frac{|G'|}{\alpha} = |\mathcal{K}|\,. \tag{3}$$

Putting (2) and (3) together we obtain

$$\chi(G') \geqslant |\mathcal{K}| > |\mathcal{K}| - 1 \geqslant \omega(G')\,,$$

a contradiction to (1). □

Since the following characterization of perfection is symmetrical in G and $\overline{G}$, it clearly implies Theorem 5.5.3. As our proof of Theorem 5.5.5 will again be from first principles, we thus obtain a second and independent proof of the perfect graph theorem.

Theorem 5.5.5. (Lovász 1972)
A graph G is perfect if and only if

$$|H| \leqslant \alpha(H) \cdot \omega(H) \tag{$*$}$$

for all induced subgraphs $H \subseteq G$.

Proof. Let us write $V(G) =: V =: \{ v_1, \dots, v_n \}$, and put $\alpha := \alpha(G)$ and $\omega := \omega(G)$. The necessity of $(*)$ is immediate: if G is perfect, then every induced subgraph H of G can be partitioned into at most $\omega(H)$ colour classes each containing at most $\alpha(H)$ vertices, and $(*)$ follows. V, v_i, n α, ω

To prove sufficiency, we apply induction on $n = |G|$. Assume that every induced subgraph H of G satisfies $(*)$, and suppose that G is not perfect. By the induction hypothesis, every proper induced subgraph of G is perfect. Hence, every non-empty independent set $U \subseteq V$ satisfies

$$\chi(G - U) = \omega(G - U) = \omega\,. \tag{1}$$

need only contain two types of graph: the odd cycles of length $\geqslant 5$ and their complements. (Neither of these are perfect—why?) Or, rephrased slightly:

Perfect Graph Conjecture. (Berge 1966)
A graph G is perfect if and only if neither G nor $\overline{G}$ contains an odd cycle of length at least 5 as an induced subgraph.

Clearly, this conjecture implies the perfect graph theorem. In fact, that theorem had also been conjectured by Berge: until its proof, it was known as the 'weak' version of the perfect graph conjecture, the above conjecture being the 'strong' version.

Graphs G such that neither G nor $\overline{G}$ contains an induced odd cycle of length at least 5 have been called *Berge graphs*. Thus all perfect graphs are Berge graphs, and the perfect graph conjecture claims that all Berge graphs are perfect. This has been approximately verified by Prömel & Steger (1992), who proved that the proportion of perfect graphs to Berge graphs on n vertices tends to 1 as $n \to \infty$.

Exercises

1.$^-$ Show that the four colour theorem does indeed solve the map colouring problem stated in the first sentence of the chapter. Conversely, does the 4-colourability of every map imply the four colour theorem?

2.$^-$ Show that, for the map colouring problem above, it suffices to consider maps such that no point lies on the boundary of more than three countries. How does this affect the proof of the four colour theorem?

3. Try to turn the proof of the five colour theorem into one of the four colour theorem, as follows. Defining v and H as before, assume inductively that H has a 4-colouring; then proceed as before. Where does the proof fail?

4. Calculate the chromatic number of a graph in terms of the chromatic numbers of its blocks.

5.$^-$ Show that every graph G has a vertex ordering for which the greedy algorithm uses only $\chi(G)$ colours.

6. For every $n > 1$, find a bipartite graph on $2n$ vertices, ordered in such a way that the greedy algorithm uses n rather than 2 colours.

7. Consider the following approach to vertex colouring. First, find a maximal independent set of vertices and colour these with colour 1; then find a maximal independent set of vertices in the remaining graph and colour those 2, and so on. Compare this algorithm with the greedy algorithm: which is better?

19.$^-$ An $n \times n$-matrix with entries from $\{1, \dots, n\}$ is called a *Latin square* if every element of $\{1, \dots, n\}$ appears exactly once in each column and exactly once in each row. Recast the problem of constructing Latin squares as a colouring problem.

20. Without using Proposition 5.3.1, show that $\chi'(G) = k$ for every k-regular bipartite graph G.

21. Prove Proposition 5.3.1 from the statement of the previous exercise.

22.$^+$ For every $k \in \mathbb{N}$, construct a triangle-free k-chromatic graph.

23.$^-$ Without using Theorem 5.4.2, show that every plane graph is 6-list-colourable.

24. For every integer k, find a 2-chromatic graph whose choice number is at least k.

25.$^-$ Find a general upper bound for $\mathrm{ch}'(G)$ in terms of $\chi'(G)$.

26. Compare the choice number of a graph with its colouring number: which is greater? Can you prove the analogue of Theorem 5.4.1 for the colouring number?

27.$^+$ Prove that the choice number of K_2^r is r.

28. The *total chromatic number* $\chi''(G)$ of a graph $G = (V, E)$ is the least number of colours needed to colour the vertices and edges of G simultaneously so that any adjacent or incident elements of $V \cup E$ are coloured differently. The *total colouring conjecture* says that $\chi''(G) \leqslant \Delta(G) + 2$. Bound the total chromatic number from above in terms of the list-chromatic index, and use this bound to deduce a weakening of the total colouring conjecture from the list colouring conjecture.

29.$^-$ Find a directed graph that has no kernel.

30.$^+$ Prove *Richardson's theorem*: every directed graph without odd directed cycles has a kernel.

31. Show that every bipartite planar graph is 3-list-colourable.
(Hint. Apply the previous exercise and Lemma 5.4.3.)

32.$^-$ Show that perfection is closed neither under edge deletion nor under edge contraction.

33.$^-$ Deduce Theorem 5.5.5 from the perfect graph conjecture.

34. Use König's Theorem 2.1.1 to show that the complement of any bipartite graph is perfect.

35. Using the results of this chapter, find a one-line proof of the following theorem of König, the dual of Theorem 2.1.1: in any bipartite graph without isolated vertices, the minimum number of edges meeting all vertices equals the maximum number of independent vertices.

36. A graph is called a *comparability graph* if there exists a partial ordering of its vertex set such that two vertices are adjacent if and only if they are comparable. Show that every comparability graph is perfect.

first brought to the attention of a wider public when Cayley presented it to the London Mathematical Society in 1878. A year later, Kempe published an incorrect proof, which was in 1890 modified by Heawood into a proof of the five colour theorem. In 1880, Tait announced 'further proofs' of the four colour conjecture, which never materialized; see the notes for Chapter 10.

The first generally accepted proof of the four colour theorem was published by Appel and Haken in 1977. The proof builds on ideas that can be traced back as far as Kempe's paper, and were developed largely by Birkhoff and Heesch. Very roughly, the proof sets out first to show that every plane triangulation must contain at least one of 1482 certain 'unavoidable configurations'. In a second step, a computer is used to show that each of those configurations is 'reducible', i.e., that any plane triangulation containing such a configuration can be 4-coloured by piecing together 4-colourings of smaller plane triangulations. Taken together, these two steps amount to an inductive proof that all plane triangulations, and hence all planar graphs, can be 4-coloured.

Appel & Haken's proof has not been immune to criticism, not only because of their use of a computer. The authors responded with a 741 page long algorithmic version of their proof, which addresses the various criticisms and corrects a number of errors (e.g. by adding more configurations to the 'unavoidable' list): K. Appel & W. Haken, *Every Planar Map is Four Colorable*, American Mathematical Society 1989. A much shorter proof, which is based on the same ideas (and, in particular, uses a computer in the same way) but can be more readily verified both in its verbal and its computer part, has been given by N. Robertson, D. Sanders, P.D. Seymour & R. Thomas, The four-colour theorem, *J. Combin. Theory B* **70** (1997), 2–44.

A relatively short proof of Grötzsch's theorem was found by C. Thomassen, Grötzsch's 3-color theorem and its counterparts for the torus and the projective plane, *J. Combin. Theory B* **62** (1994), 268–279. Although not touched upon in this chapter, colouring problems for graphs embedded in surfaces other than the plane form a substantial and interesting part of colouring theory; see B. Mohar & C. Thomassen, *Graphs on Surfaces*, Johns Hopkins University Press, to appear.

The proof of Brooks's theorem indicated in Exercise 15, where the greedy algorithm is applied to a carefully chosen vertex ordering, is due to Lovász (1973). Lovász (1968) was also the first to *construct* graphs of arbitrarily large girth and chromatic number, graphs whose existence Erdős had proved by probabilistic methods ten years earlier.

A. Urquhart, The graph constructions of Hajós and Ore, *J. Graph Theory* **26** (1997), 211–215, showed that not only do the graphs of chromatic number at least k each *contain* a k-constructible graph (as by Hajós's theorem); they are in fact all themselves k-constructible. Algebraic tools for showing that the chromatic number of a graph is large have been developed by Kleitman & Lovász (1982), and by Alon & Tarsi (1992); see Alon's paper cited below.

List colourings were first introduced in 1976 by Vizing. Among other things, Vizing proved the list-colouring equivalent of Brooks's theorem. Voigt (1993) constructed a plane graph of order 238 that is not 4-choosable; thus, Thomassen's list version of the five colour theorem is best possible. A stimulating survey on the list-chromatic number and how it relates to the more

6 Flows

Let us view a graph as a network: its edges carry some kind of flow—of water, electricity, data or similar. How could we model this precisely?

For a start, we ought to know how much flow passes through each edge $e = xy$, and in which direction. In our model, we could assign a positive integer k to the pair (x, y) to express that a flow of k units passes through e from x to y, or assign $-k$ to (x, y) to express that k units of flow pass through e the other way, from y to x. For such an assignment $f\colon V^2 \to \mathbb{Z}$ we would thus have $f(x, y) = -f(y, x)$ whenever x and y are adjacent vertices of G.

Typically, a network will have only a few nodes where flow enters or leaves the network; at all other nodes, the total amount of flow into that node will equal the total amount of flow out of it. For our model this means that, at most nodes x, the function f will satisfy *Kirchhoff's law*

Kirchhoff's law

$$\sum_{y \in N(x)} f(x, y) = 0\,.$$

In this chapter, we call any map $f\colon V^2 \to \mathbb{Z}$ with the above two properties a 'flow' on G. Sometimes, we shall replace $\mathbb{Z}$ with another group, and as a rule we consider multigraphs rather than graphs.[1] As it turns out, the theory of those 'flows' is not only useful as a model for real flows: it blends so well with other parts of graph theory that some deep and surprising connections become visible, connections particularly with connectivity and colouring problems.

[1] For consistency, we shall phrase some of our proposition for graphs only: those whose proofs rely on assertions proved (for graphs) earlier in the book. However, all those results remain true for multigraphs.

If f satisfies (F1), then

$$f(X,X) = 0$$

for all $X \subseteq V$. If f satisfies (F2), then

$$f(X,V) = \sum_{x \in X} f(x,V) = 0\,.$$

Together, these two basic observations imply that, in a circulation, the net flow across any cut is zero:

Proposition 6.1.1. *If f is a circulation, then $f(X,\overline{X}) = 0$ for every set $X \subseteq V$.* [6.3.1] [6.5.2] [6.6.1]

Proof. $f(X,\overline{X}) = f(X,V) - f(X,X) = 0 - 0 = 0$. □

Since bridges form cuts by themselves, Proposition 6.1.1 implies that circulations are always zero on bridges:

Corollary 6.1.2. *If f is a circulation and $e = xy$ is a bridge in G, then $f(e,x,y) = 0$.* □

6.2 Flows in networks

In this section we give a brief introduction to the kind of network flow theory that is now a standard proof technique in areas such as matching and connectivity. By way of example, we shall prove a classic result of this theory, the so-called *max-flow min-cut* theorem of Ford and Fulkerson. This theorem alone implies Menger's theorem without much difficulty (Exercise 3), which indicates some of the natural power lying in this approach.

Consider the task of modelling a network with one source s and one sink t, in which the amount of flow through a given link between two nodes is subject to a certain capacity of that link. Our aim is to determine the maximum net amount of flow through the network from s to t. Somehow, this will depend both on the structure of the network and on the various capacities of its connections—how exactly, is what we wish to find out.

Let $G = (V,E)$ be a multigraph, $s,t \in V$ two fixed vertices, and $c\colon \vec{E} \to \mathbb{N}$ a map; we call c a *capacity function* on G, and the tuple $N := (G,s,t,c)$ a *network*. Note that c is defined independently for the two directions of an edge. A function $f\colon \vec{E} \to \mathbb{R}$ is a *flow* in N if it satisfies the following three conditions (Fig. 6.2.1): $G = (V,E)$ s,t,c,N network flow

Theorem 6.2.2. (Ford & Fulkerson 1956) *max-flow min-cut theorem*
In every network, the maximum total value of a flow equals the minimum capacity of a cut.

Proof. Let $N = (G, s, t, c)$ be a network, and $G =: (V, E)$. We shall define a sequence $f_0, f_1, f_2, \ldots$ of integral flows in N of strictly increasing total value, i.e. with

$$|f_0| < |f_1| < |f_2| < \ldots$$

Clearly, the total value of an integral flow is again an integer, so in fact $|f_{n+1}| \geqslant |f_n| + 1$ for all n. Since all these numbers are bounded above by the capacity of any cut in N, our sequence will terminate with some flow f_n. Corresponding to this flow, we shall find a cut of capacity $c_n = |f_n|$. Since no flow can have a total value greater than c_n, and no cut can have a capacity less than $|f_n|$, this number is simultaneously the maximum and the minimum referred to in the theorem.

For f_0, we set $f_0(\vec{e}) := 0$ for all $\vec{e} \in \vec{E}$. Having defined an integral flow f_n in N for some $n \in \mathbb{N}$, we denote by S_n the set of all vertices v such that G contains an s–v walk $x_0 e_0 \ldots e_{\ell-1} x_\ell$ with S_n

$$f_n(\vec{e_i}) < c(\vec{e_i})$$

for all $i < \ell$; here, $\vec{e_i} := (e_i, x_i, x_{i+1})$ (and, of course, $x_0 = s$ and $x_\ell = v$).

If $t \in S_n$, let $W = x_0 e_0 \ldots e_{\ell-1} x_\ell$ be the corresponding s–t walk; without loss of generality we may assume that W does not repeat any vertices. Let W

$$\epsilon := \min \{ c(\vec{e_i}) - f_n(\vec{e_i}) \mid i < \ell \} .$$ ϵ

Then $\epsilon > 0$, and since f_n (like c) is integral by assumption, ϵ is an integer. Let

$$f_{n+1} \colon \vec{e} \mapsto \begin{cases} f_n(\vec{e}) + \epsilon & \text{for } \vec{e} = \vec{e_i},\ i = 0, \ldots, \ell - 1; \\ f_n(\vec{e}) - \epsilon & \text{for } \vec{e} = \overleftarrow{e_i},\ i = 0, \ldots, \ell - 1; \\ f_n(\vec{e}) & \text{for } e \notin W. \end{cases}$$

Intuitively, f_{n+1} is obtained from f_n by sending additional flow of value ϵ along W from s to t (Fig. 6.2.2).

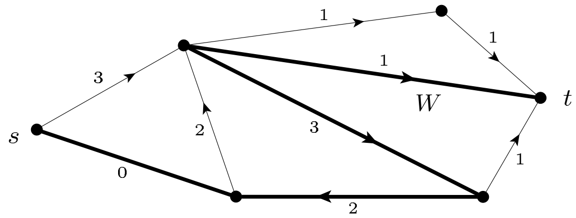

Fig. 6.2.2. An 'augmenting path' W with increment $\epsilon = 2$, for constant flow $f_n = 0$ and capacities $c = 3$

Proof. Let $G =: (V,E)$; we use induction on $m := |E|$. Let us assume first that all the edges of G are loops. Then, given any finite abelian group H, every map $\vec{E} \to H \setminus \{0\}$ is an H-flow on G. Since $|\vec{E}| = |E|$ when all edges are loops, there are $(|H|-1)^m$ such maps, and $P := x^m$ is the polynomial sought. (6.1.1)

Now assume there is an edge $e_0 = xy \in E$ that is not a loop; let $\vec{e_0} := (e_0, x, y)$ and $E' := E \setminus \{e_0\}$. We consider the multigraphs $e_0 = xy$ E'

$$G_1 := G - e_0 \quad \text{and} \quad G_2 := G/e_0 .$$

By the induction hypothesis, there are polynomials P_i for $i = 1,2$ such that, for any finite abelian group H and $k := |H|-1$, the number of H-flows on G_i is $P_i(k)$. We shall prove that the number of H-flows on G equals $P_2(k) - P_1(k)$; then $P := P_2 - P_1$ is the desired polynomial. P_1, P_2 k

Let H be given, and denote the set of all H-flows on G by F. We are trying to show that H F

$$|F| = P_2(k) - P_1(k) . \tag{1}$$

The H-flows on G_1 are precisely the restrictions to $\vec{E'}$ of those H-circulations on G that are zero on e_0 but nowhere else. Let us denote the set of these circulations on G by F_1; then F_1

$$P_1(k) = |F_1| .$$

Our aim is to show that, likewise, the H-flows on G_2 correspond bijectively to those H-circulations on G that are nowhere zero except possibly on e_0. The set F_2 of those circulations on G then satisfies F_2

$$P_2(k) = |F_2| ,$$

and F_2 is the disjoint union of F_1 and F. This will prove (1), and hence the theorem.

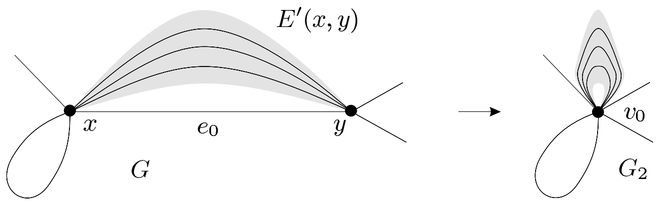

Fig. 6.3.1. Contracting the edge e_0

In G_2, let $v_0 := v_{e_0}$ be the vertex contracted from e_0 (Fig. 6.3.1; see Chapter 1.10). We are looking for a bijection $f \mapsto g$ between F_2 v_0

k — *k-flow* — *flow number* $\varphi(G)$

Let $k \geqslant 1$ be an integer and $G = (V, E)$ a multigraph. A $\mathbb{Z}$-flow f on G such that $0 < |f(\vec{e})| < k$ for all $\vec{e} \in \vec{E}$ is called a *k-flow*. Clearly, any k-flow is also an ℓ-flow for all $\ell > k$. Thus, we may ask which is the least integer k such that G admits a k-flow—assuming that such a k exists. We call this least k the *flow number* of G and denote it by $\varphi(G)$; if G has no k-flow for any k, we put $\varphi(G) := \infty$.

The task of determining flow numbers quickly leads to some of the deepest open problems in graph theory. We shall consider these later in the chapter. First, however, let us see how k-flows are related to the more general concept of H-flows.

σ_k

There is an intimate connection between k-flows and $\mathbb{Z}_k$-flows. Let σ_k denote the natural homomorphism $i \mapsto \bar{\imath}$ from $\mathbb{Z}$ to $\mathbb{Z}_k$. By composition with σ_k, every k-flow defines a $\mathbb{Z}_k$-flow. As the following theorem shows, the converse holds too: from every $\mathbb{Z}_k$-flow on G we can construct a k-flow on G. In view of Corollary 6.3.2, this means that the general question about the existence of H-flows for arbitrary groups H reduces to the corresponding question for k-flows.

[6.4.1] [6.4.2] [6.4.3] [6.4.5]

Theorem 6.3.3. (Tutte 1950)
A multigraph admits a k-flow if and only if it admits a $\mathbb{Z}_k$-flow.

g — F

Proof. Let g be a $\mathbb{Z}_k$-flow on a multigraph $G = (V, E)$; we construct a k-flow f on G. We may assume without loss of generality that G has no loops. Let F be the set of all functions $f\colon \vec{E} \to \mathbb{Z}$ that satisfy (F1), $|f(\vec{e})| < k$ for all $\vec{e} \in \vec{E}$, and $\sigma_k \circ f = g$; note that, like g, any $f \in F$ is nowhere zero.

Let us show first that $F \neq \emptyset$. Since we can express every value $g(\vec{e}) \in \mathbb{Z}_k$ as $\bar{\imath}$ with $|i| < k$ and then put $f(\vec{e}) := i$, there is clearly a map $f\colon \vec{E} \to \mathbb{Z}$ such that $|f(\vec{e})| < k$ for all $\vec{e} \in \vec{E}$ and $\sigma_k \circ f = g$. For each edge $e \in E$, let us choose one of its two directions and denote this by $\vec{e}$. We may then define $f'\colon \vec{E} \to \mathbb{Z}$ by setting $f'(\vec{e}) := f(\vec{e})$ and $f'(\overleftarrow{e}) := -f(\vec{e})$ for every $e \in E$. Then f' is a function satisfying (F1) and with values in the desired range; it remains to show that $\sigma_k \circ f'$ and g agree not only on the chosen directions $\vec{e}$ but also on their inverses $\overleftarrow{e}$. Since σ_k is a homomorphism, this is indeed so:

$$(\sigma_k \circ f')(\overleftarrow{e}) = \sigma_k(-f(\vec{e})) = -(\sigma_k \circ f)(\vec{e}) = -g(\vec{e}) = g(\overleftarrow{e}).$$

Hence $f' \in F$, so F is indeed non-empty.

f

Our aim is to find an $f \in F$ that satisfies Kirchhoff's law (F2), and is thus a k-flow. As a candidate, let us consider an $f \in F$ for which the sum

Since g is a $\mathbb{Z}_k$-flow and hence

$$\sigma_k(f(x,V)) = g(x,V) = \overline{0} \in \mathbb{Z}_k$$

and

$$\sigma_k(f(x',V)) = g(x',V) = \overline{0} \in \mathbb{Z}_k\,,$$

$f(x,V)$ and $f(x',V)$ are both multiples of k. Thus $f(x,V) \geqslant k$ and $f(x',V) \leqslant -k$, by (1) and (2). But then (4) implies that

$$|f'(x,V)| < |f(x,V)| \quad \text{and} \quad |f'(x',V)| < |f(x',V)|\,.$$

Together with (3), this gives $K(f') < K(f)$, a contradiction to the choice of f.

Therefore $K(f) = 0$ as claimed, and f is indeed a k-flow. □

Since the sum of two $\mathbb{Z}_k$-circulations is always another $\mathbb{Z}_k$-circulation, $\mathbb{Z}_k$-flows are often easier to construct (by summing over suitable partial flows) than k-flows. In this way, Theorem 6.3.3 may be of considerable help in determining whether or not some given graph has a k-flow. In the following sections we shall meet a number of examples for this.

6.4 k-Flows for small k

Trivially, a graph has a 1-flow (the empty set) if and only if it has no edges. In this section we collect a few simple examples of sufficient conditions under which a graph has a 2-, 3- or 4-flow. More examples can be found in the exercises.

Proposition 6.4.1. *A graph has a 2-flow if and only if all its degrees are even.* [6.6.1]

Proof. By Theorem 6.3.3, a graph $G = (V,E)$ has a 2-flow if and only if it has a $\mathbb{Z}_2$-flow, i.e. if and only if the constant map $\vec{E} \to \mathbb{Z}_2$ with value $\overline{1}$ satisfies (F2). This is the case if and only if all degrees are even. □ (6.3.3)

even graph

For the remainder of this chapter, let us call a graph *even* if all its vertex degrees are even.

Proposition 6.4.2. *A cubic graph has a 3-flow if and only if it is bipartite.*

Let $f_1 := \sum_{e \notin T_1} f_{1,e}$. Since each $e \notin T_1$ lies on only one cycle $C_{1,e'}$ (namely, for $e = e'$), f_1 takes only the values $\overline{1}$ and $-\overline{1}$ $(= \overline{3})$ outside T_1. Let f_1

$$F := \{\, e \in E(T_1) \mid f_1(e) = \overline{0} \,\}$$

and $f_2 := \sum_{e \in F} f_{2,e}$. As above, $f_2(e) = \overline{2} = -\overline{2}$ for all $e \in F$. Now $f := f_1 + f_2$ is the sum of $\mathbb{Z}_4$-circulations, and hence itself a $\mathbb{Z}_4$-circulation. Moreover, f is nowhere zero: on edges in F it takes the value $\overline{2}$, on edges of $T_1 - F$ it agrees with f_1 (and is hence non-zero by the choice of F), and on all edges outside T_1 it takes one of the values $\overline{1}$ or $\overline{3}$. Hence, f is a $\mathbb{Z}_4$-flow on G, and the assertion follows by Theorem 6.3.3. □ f_2 f

The following proposition describes the graphs with a 4-flow in terms of those with a 2-flow:

Proposition 6.4.5.

(i) *A graph has a 4-flow if and only if it is the union of two even subgraphs.*

(ii) *A cubic graph has a 4-flow if and only if it is 3-edge-colourable.*

Proof. Let $\mathbb{Z}_2^2 = \mathbb{Z}_2 \times \mathbb{Z}_2$ be the Klein four-group. (Thus, the elements of $\mathbb{Z}_2^2$ are the pairs (a, b) with $a, b \in \mathbb{Z}_2$, and $(a, b) + (a', b') = (a + a', b + b')$.) By Corollary 6.3.2 and Theorem 6.3.3, a graph has a 4-flow if and only if it has a $\mathbb{Z}_2^2$-flow. (6.3.2) (6.3.3)

(i) now follows directly from Proposition 6.4.1.

(ii) Let $G = (V, E)$ be a cubic graph. We assume first that G has a $\mathbb{Z}_2^2$-flow f, and define an edge colouring $E \to \mathbb{Z}_2^2 \setminus \{0\}$. As $a = -a$ for all $a \in \mathbb{Z}_2^2$, we have $f(\vec{e}) = f(\overleftarrow{e})$ for every $\vec{e} \in \vec{E}$; let us colour the edge e with this colour $f(\vec{e})$. Now if two edges with a common end v had the same colour, then these two values of f would sum to zero; by (F2), f would then assign zero to the third edge at v. As this contradicts the definition of f, our edge colouring is correct.

Conversely, since the three non-zero elements of $\mathbb{Z}_2^2$ sum to zero, every 3-edge-colouring $c\colon E \to \mathbb{Z}_2^2 \setminus \{0\}$ defines a $\mathbb{Z}_2^2$-flow on G by letting $f(\vec{e}) = f(\overleftarrow{e}) = c(e)$ for all $\vec{e} \in \vec{E}$. □

Corollary 6.4.6. *Every cubic 3-edge-colourable graph is bridgeless.* □

The proof of Lemma 6.5.1 is not entirely trivial: it is based on the so-called *orientability* of the plane, and we cannot give it here. Still, the assertion of the lemma is intuitively plausible. Indeed if we define for $e = vw$ and $e^* = xy$ the assignment $(e, v, w) \mapsto (e, v, w)^* \in \{ (e^*, x, y), (e^*, y, x) \}$ simply by turning e and its ends clockwise onto e^* (Fig. 6.5.1), then the resulting map $\vec{e} \mapsto \vec{e}^{\,*}$ satisfies the two assertions of the lemma.

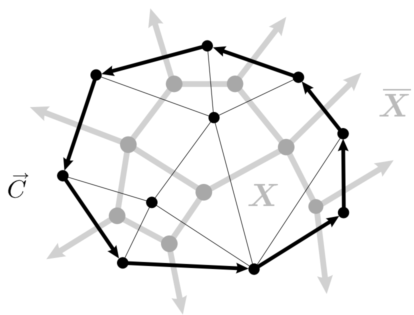

Fig. 6.5.1. Oriented cycle-cut duality

Given an abelian group H, let $f\colon \vec{E} \to H$ and $g\colon \vec{E^*} \to H$ be two maps such that f, g

$$f(\vec{e}) = g(\vec{e}^{\,*})$$

for all $\vec{e} \in \vec{E}$. For $\vec{F} \subseteq \vec{E}$, we set

$$f(\vec{F}) := \sum_{\vec{e} \in \vec{F}} f(\vec{e})\,.$$ $f(\vec{C})$ etc.

Lemma 6.5.2.

(i) *The map g satisfies (F1) if and only if f does.*

(ii) *The map g is a circulation on G^* if and only if f satisfies (F1) and $f(\vec{C}) = 0$ for every cycle $\vec{C}$ with orientation.*

Proof. Assertion (i) follows from Lemma 6.5.1 (i) and the fact that $\vec{e} \mapsto \vec{e}^{\,*}$ is bijective. (4.6.1) (6.1.1)

For the forward implication of (ii), let us assume that g is a circulation on G^*, and consider a cycle $C \subseteq G$ with some given orientation. Let $F := E(C)$. By Proposition 4.6.1, F^* is a minimal cut in G^*, i.e. $F^* = E^*(X, \overline{X})$ for some suitable $X \subseteq V^*$. By definition of f and g, Lemma 6.5.1 (ii) and Proposition 6.1.1 give

$$f(\vec{C}) = \sum_{\vec{e} \in \vec{C}} f(\vec{e}) = \sum_{\vec{d} \in \vec{E^*}(X, \overline{X})} g(\vec{d}) = g(X, \overline{X}) = 0$$

for one of the two orientations $\vec{C}$ of C. Then, by $f(\overleftarrow{C}) = -f(\vec{C})$, also

Theorem 6.5.3. (Tutte 1954)
For every dual pair G, G^* of plane multigraphs,

$$\chi(G) = \varphi(G^*) \,.$$

Proof. Let $G =: (V, E)$ and $G^* =: (V^*, E^*)$. For $|G| \in \{1, 2\}$ the (1.5.5)
assertion is easily checked; we shall assume that $|G| \geqslant 3$, and apply V, E
induction on the number of bridges in G. If $e \in G$ is a bridge then e^* V^*, E^*
is a loop, and $G^* - e^*$ is a plane dual of G/e (why?). Hence, by the induction hypothesis,

$$\chi(G) = \chi(G/e) = \varphi(G^* - e^*) = \varphi(G^*) \,;$$

for the first and the last equality we use that, by $|G| \geqslant 3$, e is not the only edge of G.

So all that remains to be checked is the induction start: let us assume that G has no bridge. If G has a loop, then G^* has a bridge, and $\chi(G) = \infty = \varphi(G^*)$ by convention. So we may also assume that G
has no loop. Then $\chi(G)$ is finite; we shall prove for given $k \geqslant 2$ that G k
is k-colourable if and only if G^* has a k-flow. As G—and hence G^*—has neither loops nor bridges, we may apply Lemmas 6.5.1 and 6.5.2 to G and G^*. Let $\vec{e} \mapsto \vec{e}^{\,*}$ be the bijection between $\vec{E}$ and $\vec{E^*}$ from Lemma 6.5.1.

We first assume that G^* has a k-flow. Then G^* also has a $\mathbb{Z}_k$-flow g. g
As before, let $f \colon \vec{E} \to \mathbb{Z}_k$ be defined by $f(\vec{e}) := g(\vec{e}^{\,*})$. We shall use f to f
define a vertex colouring $c \colon V \to \mathbb{Z}_k$ of G.

Let T be a normal spanning tree of G, with root r, say. Put $c(r) := \overline{0}$. For every other vertex $v \in V$ let $c(v) := f(\vec{P})$, where $\vec{P}$ is the $r \to v$ path in T. To check that this is a proper colouring, consider an edge $e = vw \in E$. As T is normal, we may assume that $v < w$ in the tree order of T. If e is an edge of T then $c(w) - c(v) = f(e, v, w)$ by definition of c, so $c(v) \neq c(w)$ since g (and hence f) is nowhere zero. If $e \notin T$, let $\vec{P}$ denote the $v \to w$ path in T. Then

$$c(w) - c(v) = f(\vec{P}) = -f(e, w, v) \neq \overline{0}$$

by Lemma 6.5.2 (ii).

Conversely, we now assume that G has a k-colouring c. Let us define c
$f \colon \vec{E} \to \mathbb{Z}$ by

$$f(e, v, w) := c(w) - c(v) \,, \qquad f$$

and $g \colon \vec{E^*} \to \mathbb{Z}$ by $g(\vec{e}^{\,*}) := f(\vec{e})$. Clearly, f satisfies (F1) and takes g
values in $\{\pm 1, \dots, \pm(k-1)\}$, so by Lemma 6.5.2 (i) the same holds for g. By definition of f, we further have $f(\vec{C}) = 0$ for every cycle $\vec{C}$ with orientation. By Lemma 6.5.2 (ii), therefore, g is a k-flow. □

Even if true, the 4-flow conjecture will not be best possible: a K^{11}, for example, contains the Petersen graph as a minor but has a 4-flow, even a 2-flow. The conjecture appears more natural for sparser graphs, and indeed the cubic graphs form an important special case. (See the notes.)

A cubic bridgeless graph or multigraph without a 4-flow (equivalently, without a 3-edge-colouring) is called a *snark*. The 4-flow conjecture for cubic graphs says that every snark contains the Petersen graph as a minor; in this sense, the Petersen graph has thus been shown to be the smallest snark. Snarks form the hard core both of the four colour theorem and of the 5-flow conjecture: the four colour theorem is equivalent to the assertion that no snark is planar (exercise), and it is not difficult to reduce the 5-flow conjecture to the case of snarks.[5] However, although the snarks form a very special class of graphs, none of the problems mentioned seems to become much easier by this reduction.[6] *snark*

Three-Flow Conjecture. (Tutte 1972)
Every multigraph without a cut consisting of exactly one or exactly three edges has a 3-flow.

Again, the 3-flow conjecture will not be best possible: it is easy to construct graphs with three-edge cuts that have a 3-flow (exercise).

By our duality theorem (6.5.3), all three flow conjectures are true for planar graphs and thus motivated: the 3-flow conjecture translates to Grötzsch's theorem (5.1.3), the 4-flow conjecture to the four colour theorem (since the Petersen graph is not planar, it is not a minor of a planar graph), the 5-flow conjecture to the five colour theorem.

We finish this section with the main result of the chapter:

Theorem 6.6.1. (Seymour 1981)
Every bridgeless graph has a 6-flow.

(3.3.5) (6.1.1) (6.4.1)

Proof. Let $G = (V, E)$ be a bridgeless graph. Since 6-flows on the components of G will add up to a 6-flow on G, we may assume that G is connected; as G is bridgeless, it is then 2-edge-connected. Note that any two vertices in a 2-edge-connected graph lie in some common even connected subgraph—for example, in the union of two edge-disjoint paths linking these vertices by Menger's theorem (3.3.5 (ii)). We shall use this fact repeatedly.

---

[5] The same applies to another well-known conjecture, the *cycle double cover conjecture*; see Exercise 13.

[6] That snarks are elusive has been known to mathematicians for some time; cf. Lewis Carroll, *The Hunting of the Snark*, Macmillan 1876.

f_i

Assume now that $\mathbb{Z}_3$-circulations $f_n, \dots, f_i$ on G have been defined for some $i \leqslant n$, and that

$$f_i(\vec{e}) \neq \overline{0} \quad \text{for all} \quad \vec{e} \in \vec{E'} \cup \bigcup_{j>i} \vec{F_j}\,, \tag{2}$$

$\vec{F_j}$

where $\vec{F_j} := \{\, \vec{e} \in \vec{E} \mid e \in F_j \,\}$. Our aim is to define f_{i-1} in such a way that (2) also holds for $i-1$.

e

We first consider the case that $|F_i| = 1$, say $F_i = \{\, e \,\}$. We then let $f_{i-1} := f_i$, and thus have to show that f_i is non-zero on (the two directions of) e. Our assumption of $|F_i| = 1$ implies by the choice of F_i that G contains no X_i–V^{i-1} edge other than e. Since G is 2-edge-connected, it therefore has at least—and thus, by (1), exactly—one edge e' between X_i and $\overline{V^{i-1}} \setminus X_i$. We show that f_i is non-zero on e'; as $\{\, e, e' \,\}$ is a cut in G, this implies by Proposition 6.1.1 that f_i is also non-zero on e.

e'

To show that f_i is non-zero on e', we use (2): we show that $e' \in E' \cup \bigcup_{j>i} F_j$, i.e. that e' lies in no H_k and in no F_j with $j \leqslant i$. Since e' has both ends in $\overline{V^{i-1}}$, it clearly lies in no F_j with $j \leqslant i$ and in no H_k with $k < i$. But every H_k with $k \geqslant i$ is a subgraph of $G[\,\overline{V^{i-1}}\,]$. Since e' is a bridge of $G[\,\overline{V^{i-1}}\,]$ but H_k has no bridge, this means that $e' \notin H_k$. Hence, f_{i-1} does indeed satisfy (2) for $i-1$ in the case considered.

e_1, e_2
C

It remains to consider the case that $|F_i| = 2$, say $F_i = \{\, e_1, e_2 \,\}$. Since H_i and H^{i-1} are both connected, we can find a cycle C in $H^i = (H_i \cup H^{i-1}) + F_i$ that contains e_1 and e_2. If f_i is non-zero on both these edges, we again let $f_{i-1} := f_i$. Otherwise, there are directions $\vec{e_1}$ and $\vec{e_2}$ of e_1 and e_2 such that, without loss of generality, $f_i(\vec{e_1}) = \overline{0}$ and $f_i(\vec{e_2}) \in \{\, \overline{0}, \overline{1} \,\}$. Let $\vec{C}$ be the orientation of C with $\vec{e_2} \in \vec{C}$, and let g be a flow of value $\overline{1}$ around $\vec{C}$ (formally: let g be a $\mathbb{Z}_3$-circulation on G such that $g(\vec{e_2}) = \overline{1}$ and $g^{-1}(\overline{0}) = \vec{E} \setminus (\vec{C} \cup \overleftarrow{C})$). We then let $f_{i-1} := f_i + g$. By choice of the directions $\vec{e_1}$ and $\vec{e_2}$, f_{i-1} is non-zero on both edges. Since f_{i-1} agrees with f_i on all of $\vec{E'} \cup \bigcup_{j>i} \vec{F_j}$ and (2) holds for i, we again have (2) also for $i-1$.

f

Eventually, f_0 will be a $\mathbb{Z}_3$-circulation on G that is nowhere zero except possibly on edges of $H_0 \cup \ldots \cup H_n$. Composing f_0 with the map $\overline{h} \mapsto \overline{2h}$ from $\mathbb{Z}_3$ to $\mathbb{Z}_6$ ($h \in \{\, 1, 2 \,\}$), we obtain a $\mathbb{Z}_6$-circulation f on G with values in $\{\, \overline{0}, \overline{2}, \overline{4} \,\}$ for all edges lying in some H_i, and with values in $\{\, \overline{2}, \overline{4} \,\}$ for all other edges. Adding to f a 2-flow on each H_i (formally: a $\mathbb{Z}_6$-circulation on G with values in $\{\, \overline{1}, -\overline{1} \,\}$ on the edges of H_i and $\overline{0}$ otherwise; this exists by Proposition 6.4.1), we obtain a $\mathbb{Z}_6$-circulation on G that is nowhere zero. Hence, G has a 6-flow by Theorem 6.3.3. □

14.$^-$ Determine the flow number of $C^5 * K^1$, the wheel with 5 spokes.

15. Find bridgeless graphs G and $H = G - e$ such that $2 < \varphi(G) < \varphi(H)$.

16. Prove Proposition 6.4.1 without using Theorem 6.3.3.

17.$^+$ Prove *Heawood's theorem* that a plane triangulation is 3-colourable if and only if all its vertices have even degree.

18.$^-$ Find a bridgeless graph that has both a 3-flow and a cut of exactly three edges.

19. Show that the 3-flow conjecture for planar multigraphs is equivalent to Grötzsch's Theorem 5.1.3.

20. (i)$^-$ Show that the four colour theorem is equivalent to the non-existence of a planar snark, i.e. to the statement that every cubic bridgeless planar multigraph has a 4-flow.

    (ii) Can 'bridgeless' in (i) be replaced by '3-connected'?

21.$^+$ Show that a graph $G = (V, E)$ has a k-flow if and only if it admits an orientation D that directs, for every $X \subseteq V$, at least $1/k$ of the edges in $E(X, \overline{X})$ from X towards $\overline{X}$.

22.$^-$ Generalize the 6-flow Theorem 6.6.1 to multigraphs.

Notes

Network flow theory is an application of graph theory that has had a major and lasting impact on its development over decades. As is illustrated already by the fact that Menger's theorem can be deduced easily from the max-flow min-cut theorem (Exercise 3), the interaction between graphs and networks may go either way: while 'pure' results in areas such as connectivity, matching and random graphs have found applications in network flows, the intuitive power of the latter has boosted the development of proof techniques that have in turn brought about theoretic advances.

The standard reference for network flows is L.R. Ford & D.R. Fulkerson, *Flows in Networks*, Princeton University Press 1962. A more recent and comprehensive account is given by R.K. Ahuja, T.L. Magnanti & J.B. Orlin, *Network flows*, Prentice-Hall 1993. For more theoretical aspects, see A. Frank's chapter in the *Handbook of Combinatorics* (R.L. Graham, M. Grötschel & L. Lovász, eds.), North-Holland 1995. A general introduction to graph algorithms is given in A. Gibbons, *Algorithmic Graph Theory*, Cambridge University Press 1985.

If one recasts the maximum flow problem in linear programming terms, one can derive the max-flow min-cut theorem from the linear programming duality theorem; see A. Schrijver, *Theory of integer and linear programming*, Wiley 1986.

The more algebraic theory of group-valued flows and k-flows has been developed largely by Tutte; he gives a thorough account in his monograph W.T. Tutte, *Graph Theory*, Addison-Wesley 1984. Tutte's flow conjectures are

7 Substructures in Dense Graphs

In this chapter and the next, we study how global parameters of a graph, such as its edge density or chromatic number, have a bearing on the existence of certain local substructures. How many edges, for instance, do we have to give a graph on n vertices to be sure that, no matter how these edges happen to be arranged, the graph will contain a K^r subgraph for some given r? Or at least a K^r minor? Or a topological K^r minor? Will some sufficiently high average degree or chromatic number ensure that one of these substructures occurs?

Questions of this type are among the most natural ones in graph theory, and there is a host of deep and interesting results. Collectively, these are known as *extremal graph theory*.

Extremal graph problems in this sense fall neatly into two categories, as follows. If we are looking for ways to ensure by global assumptions that a graph G contains some given graph H as a *minor* (or topological minor), it will suffice to raise $\|G\|$ above the value of some linear function of $|G|$ (depending on H), i.e. to make $\varepsilon(G)$ large enough. The existence of such a function was already established in Theorem 3.6.1. The precise growth rate needed will be investigated in Chapter 8, where we study substructures of such 'sparse' graphs. Since a large enough value of ε gives rise to an H minor for any given graph H, its occurrence could be forced alternatively by raising some other global invariants (such as κ or χ) which, in turn, force up the value of ε, at least in some subgraph. This, too, will be a topic for Chapter 8.

On the other hand, if we ask what global assumptions might imply the existence of some given graph H as a *subgraph*, it will not help to raise any of the invariants ε, κ or χ, let alone any of the other invariants discussed in Chapter 1. Indeed, as mentioned in Chapter 5.2, A graph $G \not\supseteq H$ on n vertices with the largest possible number of edges is called *extremal* for n and H; its number of edges is denoted by $\text{ex}(n, H)$. Clearly, any graph G that is extremal for some n and H will also be edge-maximal with $H \not\subseteq G$. Conversely, though, edge-maximality does not imply extremality: G may well be edge-maximal with $H \not\subseteq G$ while having fewer than $\text{ex}(n, H)$ edges (Fig. 7.1.1).

extremal
$\text{ex}(n, H)$

Fig. 7.1.1. Two graphs that are edge-maximal with $P^3 \not\subseteq G$; is the right one extremal?

As a case in point, we consider our problem for $H = K^r$ (with $r > 1$). A moment's thought suggests some obvious candidates for extremality here: all complete $(r-1)$-partite graphs are edge-maximal without containing K^r. But which among these have the greatest number of edges? Clearly those whose partition sets are as equal as possible, i.e. differ in size by at most 1: if V_1, V_2 are two partition sets with $|V_1| - |V_2| \geqslant 2$, we may increase the number of edges in our complete $(r-1)$-partite graph by moving a vertex from V_1 across to V_2.

The unique complete $(r-1)$-partite graphs on $n \geqslant r-1$ vertices whose partition sets differ in size by at most 1 are called *Turán graphs*; we denote them by $T^{r-1}(n)$ and their number of edges by $t_{r-1}(n)$ (Fig. 7.1.2). For $n < r-1$ we shall formally continue to use these definitions, with the proviso that—contrary to our usual terminology—the partition sets may now be empty; then, clearly, $T^{r-1}(n) = K^n$ for all $n \leqslant r-1$.

$T^{r-1}(n)$
$t_{r-1}(n)$

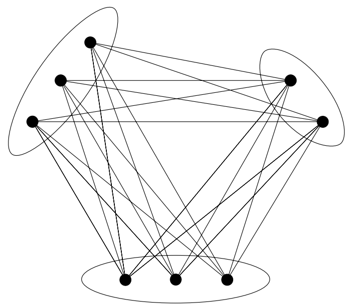

Fig. 7.1.2. The Turán graph $T^3(8)$

The following theorem tells us that $T^{r-1}(n)$ is indeed extremal for n and K^r, and as such unique; in particular, $\text{ex}(n, K^r) = t_{r-1}(n)$.

with equality whenever $r-1$ divides n (Exercise 8). It is therefore remarkable that just ϵn^2 more edges (for any fixed $\epsilon > 0$ and n large) give us not only a K^r subgraph (as does Turán's theorem) but a K^r_s for any given integer s—a graph itself teeming with K^r subgraphs:

Theorem 7.1.2. (Erdős & Stone 1946)
For all integers $r \geqslant 2$ and $s \geqslant 1$, and every $\epsilon > 0$, there exists an integer n_0 such that every graph with $n \geqslant n_0$ vertices and at least

$$t_{r-1}(n) + \epsilon n^2$$

edges contains K^r_s as a subgraph.

We shall prove this theorem in Section 7.3.

The Erdős-Stone theorem is interesting not only in its own right: it also has a most interesting corollary. In fact, it was this entirely unexpected corollary that established the theorem as a kind of meta-theorem for the extremal theory of dense graphs, and thus made it famous.

Given a graph H and an integer n, consider the number $h_n := \text{ex}(n,H)/\binom{n}{2}$: the maximum edge density that an n-vertex graph can have without containing a copy of H. Could it be that this critical density is essentially just a function of H, that h_n converges as $n \to \infty$? Theorem 7.1.2 implies this, and more: the limit of h_n is determined by a very simple function of a natural invariant of H—its chromatic number!

Corollary 7.1.3. *For every graph H with at least one edge,*

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

For the proof of Corollary 7.1.3 we need as a lemma that $t_{r-1}(n)$ never deviates much from the value it takes when $r-1$ divides n (see above), and that $t_{r-1}(n)/\binom{n}{2}$ converges accordingly. The proof of the lemma is left as an easy exercise with hint (Exercise 9).

Lemma 7.1.4. [7.1.2]

$$\lim_{n\to\infty} t_{r-1}(n)\binom{n}{2}^{-1} = \frac{r-2}{r-1}.$$

□

7.2 Szemerédi's regularity lemma

More than 20 years ago, in the course of the proof of a major result on the Ramsey properties of arithmetic progressions, Szemerédi developed a graph theoretical tool whose fundamental importance has been realized more and more in recent years: his so-called *regularity* or *uniformity lemma*. Very roughly, the lemma says that all graphs can be approximated by random graphs in the following sense: every graph can be partitioned, into a bounded number of equal parts, so that most of its edges run between different parts and the edges between any two parts are distributed fairly uniformly—just as we would expect it if they had been generated at random.

In order to state the regularity lemma precisely, we need some definitions. Let $G = (V, E)$ be a graph, and let $X, Y \subseteq V$ be disjoint. Then we denote by $\|X, Y\|$ the number of X–Y edges of G, and call $\|X, Y\|$

$$d(X, Y) := \frac{\|X, Y\|}{|X|\,|Y|}$$

$d(X, Y)$

the *density* of the pair (X, Y). (This is a real number between 0 and 1.) *density*
Given some $\epsilon > 0$, we call a pair (A, B) of disjoint sets $A, B \subseteq V$ *ϵ-regular* if all $X \subseteq A$ and $Y \subseteq B$ with *ϵ-regular pair*

$$|X| \geqslant \epsilon\,|A| \quad \text{and} \quad |Y| \geqslant \epsilon\,|B|$$

satisfy

$$\big|d(X, Y) - d(A, B)\big| \leqslant \epsilon\,.$$

The edges in an ϵ-regular pair are thus distributed fairly uniformly: the smaller ϵ, the more uniform their distribution.

Consider a partition $\{ V_0, V_1, \ldots, V_k \}$ of V in which one set V_0 has been singled out as an *exceptional set*. (This exceptional set V_0 may be empty.[3]) We call such a partition an *ϵ-regular* partition of G if it satisfies the following three conditions: *exceptional set*

ϵ-regular partition

(i) $|V_0| \leqslant \epsilon\,|V|$;
(ii) $|V_1| = \ldots = |V_k|$;
(iii) all but at most ϵk^2 of the pairs (V_i, V_j) with $1 \leqslant i < j \leqslant k$ are ϵ-regular.

The role of the exceptional set V_0 is one of pure convenience: it makes it possible to require that all the other partition sets have exactly the same size. Since condition (iii) affects only the sets $V_1, \ldots, V_k$, we

[3] So V_0 may be an exception also to our terminological rule that partition sets are not normally empty.

and for a partition $\mathcal{P} = \{ C_1, \dots, C_k \}$ of V we let

$$q(\mathcal{P}) := \sum_{i<j} q(C_i, C_j)\,.$$ $q(\mathcal{P})$

However, if $\mathcal{P} = \{ C_0, C_1, \dots, C_k \}$ is a partition of V *with exceptional set* C_0, we treat C_0 as a set of singletons and define

$$q(\mathcal{P}) := q(\tilde{\mathcal{P}})\,,$$

where $\tilde{\mathcal{P}} := \{C_1, \dots, C_k\} \cup \{\{ v \} : v \in C_0\}$. $\tilde{\mathcal{P}}$

The function $q(\mathcal{P})$ plays a pivotal role in the proof of the regularity lemma. On the one hand, it measures the uniformity of the partition $\mathcal{P}$: if $\mathcal{P}$ has too many irregular pairs (A, B), we may take the pairs (X, Y) of subsets violating the regularity of the pairs (A, B) and make those sets X and Y into partition sets of their own; as we shall prove, this refines $\mathcal{P}$ into a partition for which q is substantially greater than for $\mathcal{P}$. Here, 'substantial' means that the increase of $q(\mathcal{P})$ is bounded below by some constant depending only on ϵ. On the other hand,

$$\begin{aligned} q(\mathcal{P}) &= \sum_{i<j} q(C_i, C_j) \\ &= \sum_{i<j} \frac{|C_i|\,|C_j|}{n^2}\, d^2(C_i, C_j) \\ &\leqslant \frac{1}{n^2} \sum_{i<j} |C_i|\,|C_j| \\ &\leqslant 1\,. \end{aligned}$$

The number of times that $q(\mathcal{P})$ can be increased by a constant is thus also bounded by a constant—in other words, after some bounded number of refinements our partition will be ϵ-regular! To complete the proof of the regularity lemma, all we have to do then is to note how many sets that last partition can possibly have if we start with a partition into m sets, and to choose this number as our desired bound M.

Let us make all this precise. We begin by showing that, when we refine a partition, the value of q will not decrease:

Lemma 7.2.2.

(i) *Let $C, D \subseteq V$ be disjoint. If $\mathcal{C}$ is a partition of C and $\mathcal{D}$ is a partition of D, then $q(\mathcal{C}, \mathcal{D}) \geqslant q(C, D)$.*

(ii) *If $\mathcal{P}, \mathcal{P}'$ are partitions of V and $\mathcal{P}'$ refines $\mathcal{P}$, then $q(\mathcal{P}') \geqslant q(\mathcal{P})$.*

Let us show that $\mathcal{C}$ and $\mathcal{D}$ satisfy the conclusion of the lemma. We shall write $c_i := |C_i|$, $d_i := |D_i|$, $e_{ij} := \|C_i, D_j\|$, $c := |C|$, $d := |D|$ and $e := \|C, D\|$. As in the proof of Lemma 7.2.2, c_i, d_i, e_{ij} $c,\ d,\ e$

$$\begin{aligned} q(\mathcal{C},\mathcal{D}) &= \frac{1}{n^2}\sum_{i,j}\frac{e_{ij}^2}{c_i d_j} \\ &= \frac{1}{n^2}\left(\frac{e_{11}^2}{c_1 d_1} + \sum_{i+j>2}\frac{e_{ij}^2}{c_i d_j}\right) \\ &\underset{(1)}{\geqslant} \frac{1}{n^2}\left(\frac{e_{11}^2}{c_1 d_1} + \frac{(e-e_{11})^2}{cd - c_1 d_1}\right). \end{aligned}$$

By definition of η, we have $e_{11} = c_1 d_1 e/cd + \eta c_1 d_1$, so

$$\begin{aligned} n^2\, q(\mathcal{C},\mathcal{D}) &\geqslant \frac{1}{c_1 d_1}\left(\frac{c_1 d_1 e}{cd} + \eta c_1 d_1\right)^2 \\ &\quad + \frac{1}{cd - c_1 d_1}\left(\frac{cd - c_1 d_1}{cd} e - \eta c_1 d_1\right)^2 \\ &= \frac{c_1 d_1 e^2}{c^2 d^2} + \frac{2e\eta c_1 d_1}{cd} + \eta^2 c_1 d_1 \\ &\quad + \frac{cd - c_1 d_1}{c^2 d^2} e^2 - \frac{2e\eta c_1 d_1}{cd} + \frac{\eta^2 c_1^2 d_1^2}{cd - c_1 d_1} \\ &\geqslant \frac{e^2}{cd} + \eta^2 c_1 d_1 \\ &\underset{(2)}{\geqslant} \frac{e^2}{cd} + \epsilon^4 cd \end{aligned}$$

since $c_1 \geqslant \epsilon c$ and $d_1 \geqslant \epsilon d$ by the choice of C_1 and D_1. □

Finally, we show that if a partition has enough irregular pairs of partition sets to fall short of the definition of an ϵ-regular partition, then subpartitioning all those pairs at once results in an increase of q by a constant:

Lemma 7.2.4. *Let $0 < \epsilon \leqslant 1/4$, and let $\mathcal{P} = \{C_0, C_1, \dots, C_k\}$ be a partition of V, with exceptional set C_0 of size $|C_0| \leqslant \epsilon n$ and $|C_1| = \ldots = |C_k| =: c$. If $\mathcal{P}$ is not ϵ-regular, then there is a partition $\mathcal{P}' = \{C_0', C_1', \dots, C_\ell'\}$ of V with exceptional set C_0', where $k \leqslant \ell \leqslant k4^k$, such that $|C_0'| \leqslant |C_0| + n/2^k$, all other sets C_i' have equal size, and* c

$$q(\mathcal{P}') \geqslant q(\mathcal{P}) + \epsilon^5/2\,.$$

$C \in \mathcal{C} \smallsetminus \{ C_0 \}$, and put $C_0' := V \smallsetminus \bigcup C_i'$. Then $\mathcal{P}' = \{ C_0', C_1', \dots, C_\ell' \}$ is indeed a partition of V. Moreover, $\tilde{\mathcal{P}}'$ refines $\tilde{\mathcal{C}}$, so $\mathcal{P}'$

$$q(\mathcal{P}') \geqslant q(\mathcal{C}) \geqslant q(\mathcal{P}) + \epsilon^5/2$$

by Lemma 7.2.2 (ii). Since each set $C_i' \neq C_0'$ is also contained in one of the sets $C_1, \dots, C_k$, but no more than 4^k sets C_i' can lie inside the same C_j (by the choice of d), we also have $k \leqslant \ell \leqslant k4^k$ as required. Finally, the sets $C_1', \dots, C_\ell'$ use all but at most d vertices from each set $C \neq C_0$ of $\mathcal{C}$. Hence,

$$\begin{aligned} |C_0'| &\leqslant |C_0| + d\,|\mathcal{C}| \\ &\underset{(4)}{\leqslant} |C_0| + \frac{c}{4^k} k 2^k \\ &= |C_0| + ck/2^k \\ &\leqslant |C_0| + n/2^k. \end{aligned}$$

□

The proof of the regularity lemma now follows easily by repeated application of Lemma 7.2.4:

Proof of Lemma 7.2.1. Let $\epsilon > 0$ and $m \geqslant 1$ be given; without loss of generality, $\epsilon \leqslant 1/4$. Let $s := 2/\epsilon^5$. This number s is an upper bound on the number of iterations of Lemma 7.2.4 that can be applied to a partition of a graph before it becomes ϵ-regular; recall that $q(\mathcal{P}) \leqslant 1$ for all partitions $\mathcal{P}$. ϵ, m s

There is one formal requirement which a partition $\{ C_0, C_1, \dots, C_k \}$ with $|C_1| = \dots = |C_k|$ has to satisfy before Lemma 7.2.4 can be (re-)applied: the size $|C_0|$ of its exceptional set must not exceed ϵn. With each iteration of the lemma, however, the size of the exceptional set can grow by up to $n/2^k$. (More precisely, by up to $n/2^\ell$, where ℓ is the number of other sets in the current partition; but $\ell \geqslant k$ by the lemma, so $n/2^k$ is certainly an upper bound for the increase.) We thus want to choose k large enough that even s increments of $n/2^k$ add up to at most $\frac{1}{2}\epsilon n$, and n large enough that, for any initial value of $|C_0| < k$, we have $|C_0| \leqslant \frac{1}{2}\epsilon n$. (If we give our starting partition k non-exceptional sets $C_1, \dots, C_k$, we should allow an initial size of up to k for C_0, to be able to achieve $|C_1| = \dots = |C_k|$.)

So let $k \geqslant m$ be large enough that $2^{k-1} \geqslant s/\epsilon$. Then $s/2^k \leqslant \epsilon/2$, and hence k

$$k + \frac{s}{2^k} n \leqslant \epsilon n \tag{5}$$

whenever $k/n \leqslant \epsilon/2$, i.e. for all $n \geqslant 2k/\epsilon$.

Let us now choose M. This should be an upper bound on the number of (non-exceptional) sets in our partition after up to s iterations small, most vertices of A have about the expected number of neighbours in Y:

Lemma 7.3.1. *Let (A, B) be an ϵ-regular pair, of density d say, and let $Y \subseteq B$ have size $|Y| \geqslant \epsilon\,|B|$. Then all but at most $\epsilon\,|A|$ of the vertices in A have (each) at least $(d-\epsilon)|Y|$ neighbours in Y.*

Proof. Let $X \subseteq A$ be the set of vertices with fewer than $(d-\epsilon)|Y|$ neighbours in Y. Then $\|X, Y\| < |X|(d-\epsilon)|Y|$, so

$$d(X, Y) = \frac{\|X, Y\|}{|X|\,|Y|} < d - \epsilon = d(A, B) - \epsilon\,.$$

Since (A, B) is ϵ-regular, this implies that $|X| < \epsilon\,|A|$. □

Let G be a graph with an ϵ-regular partition $\{V_0, V_1, \dots, V_k\}$, with exceptional set V_0 and $|V_1| = \dots = |V_k| =: \ell$. Given $d \in (0, 1]$, let R be *R* the graph with vertices $V_1, \dots, V_k$ in which two vertices are adjacent if and only if they form an ϵ-regular pair in G of density $\geqslant d$. We shall call R a *regularity graph* of G with parameters ϵ, ℓ and d. Given $s \in \mathbb{N}$, let *regularity graph* us now replace every vertex V_i of R by a set V_i^s of s vertices, and every *V_i^s* edge by a complete bipartite graph between the corresponding s-sets. The resulting graph will be denoted by R_s. (For $R = K^r$, for example, *R_s* we have $R_s = K^r_s$.)

The following lemma says that subgraphs of R_s can also be found in G, provided that ϵ is small enough and the V_i are large enough. In fact, the values of ϵ and ℓ required depend only on (d and) the maximum degree of the subgraph:

Lemma 7.3.2. *For all $d \in (0, 1]$ and $\Delta \geqslant 1$ there exists an $\epsilon_0 > 0$ with* [9.2.2] *the following property: if G is any graph, H is a graph with $\Delta(H) \leqslant \Delta$, $s \in \mathbb{N}$, and R is any regularity graph of G with parameters $\epsilon \leqslant \epsilon_0$, $\ell \geqslant s/\epsilon_0$ and d, then*

$$H \subseteq R_s \;\Rightarrow\; H \subseteq G\,.$$

Proof. Given d and Δ, choose $\epsilon_0 < d$ small enough that *d, Δ*

$$\frac{\Delta + 1}{(d - \epsilon_0)^\Delta}\,\epsilon_0 \leqslant 1\,; \qquad (1)$$

ϵ_0

such a choice is possible, since $(\Delta + 1)\epsilon/(d-\epsilon)^\Delta \to 0$ as $\epsilon \to 0$. Now let *G, H, R, R_s* G, H, s and R be given as stated. Let $\{V_0, V_1, \dots, V_k\}$ be the ϵ-regular *V_i* partition of G that gave rise to R; thus, $\epsilon \leqslant \epsilon_0$, $V(R) = \{V_1, \dots, V_k\}$ *ϵ, k, ℓ* and $|V_1| = \dots = |V_k| = \ell$. Let us assume that H is actually a subgraph

We are now ready to prove the Erdős-Stone theorem.

(7.1.1) (7.1.4)

Proof of Theorem 7.1.2. Let $r \geqslant 2$ and $s \geqslant 1$ be given as in the statement of the theorem. For $s = 1$ the assertion follows from Turán's theorem, so we assume that $s \geqslant 2$. Let $\gamma > 0$ be given; this γ will play the role of the ϵ of the theorem. Let G be a graph with $|G| =: n$ and [margin: r, s; γ]

$$\|G\| \geqslant t_{r-1}(n) + \gamma n^2.$$ [margin: $\|G\|$]

(Thus, $\gamma < 1$.) We want to show that $K_s^r \subseteq G$ if n is large enough.

Our plan is to use the regularity lemma to show that G has a regularity graph R dense enough to contain a K^r by Turán's theorem. Then R_s contains a K_s^r, so we may hope to use Lemma 7.3.2 to deduce that $K_s^r \subseteq G$.

On input $d := \gamma$ and $\Delta := \Delta(K_s^r)$, Lemma 7.3.2 returns an $\epsilon_0 > 0$; since the lemma's assertion about ϵ_0 becomes weaker when ϵ_0 is made smaller, we may assume that [margin: d, Δ; ϵ_0]

$$\epsilon_0 < \gamma/2 < 1\,. \tag{1}$$

To apply the regularity lemma, let $m > 1/\gamma$ and choose $\epsilon > 0$ small enough that $\epsilon \leqslant \epsilon_0$ and [margin: m, ϵ]

$$\delta := 2\gamma - \epsilon^2 - 4\epsilon - d - \frac{1}{m} > 0\,;$$ [margin: δ]

this is possible, since $2\gamma - d - \frac{1}{m} > 0$. On input ϵ and m, the regularity lemma returns an integer M. Let us assume that [margin: M]

$$n \geqslant \frac{Ms}{\epsilon_0(1-\epsilon)}\,.$$ [margin: n]

Since this number is at least m, the regularity lemma provides us with an ϵ-regular partition $\{ V_0, V_1, \ldots, V_k \}$ of G, where $m \leqslant k \leqslant M$; let $|V_1| = \ldots = |V_k| =: \ell$. Then [margin: k; ℓ]

$$n \geqslant k\ell\,, \tag{2}$$

and

$$\ell = \frac{n - |V_0|}{k} \geqslant \frac{n - \epsilon n}{M} = n\,\frac{1-\epsilon}{M} \geqslant \frac{s}{\epsilon_0}$$

by the choice of n. Let R be the regularity graph of G with parameters ϵ, ℓ, d corresponding to the above partition. Since $\epsilon \leqslant \epsilon_0$ and $\ell \geqslant s/\epsilon_0$, the regularity graph R satisfies the premise of Lemma 7.3.2, and by definition of Δ we have $\Delta(K_s^r) = \Delta$. Thus in order to conclude by Lemma 7.3.2 [margin: R]

Exercises

1.$^-$ Show that $K_{1,3}$ is extremal without a P^3.

2.$^-$ Given $k > 0$, determine the extremal graphs of chromatic number at most k.

3. Determine the value of $\mathrm{ex}(n, K_{1,r})$ for all $r, n \in \mathbb{N}$.

4. Is there a graph that is edge-maximal without a K^3 minor but not extremal?

5. Show that, for every forest F, the value of $\mathrm{ex}(n, F)$ is bounded above by a linear function of n.

6.$^+$ Given $k > 0$, determine the extremal graphs without a matching of size k.

(Hint. Theorem 2.2.3 and Ex. 10, Ch. 2.)

7. Without using Turán's theorem, show that the maximum number of edges in a triangle-free graph of order $n > 1$ is $\lfloor n^2/4 \rfloor$.

8. Show that

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

with equality whenever $r-1$ divides n.

9. Show that $t_{r-1}(n)/\binom{n}{2}$ converges to $(r-2)/(r-1)$ as $n \to \infty$.

(Hint. $t_{r-1}\big((r-1)\lfloor \frac{n}{r-1} \rfloor\big) \leqslant t_{r-1}(n) \leqslant t_{r-1}\big((r-1)\lceil \frac{n}{r-1} \rceil\big)$.)

10.$^+$ Given non-adjacent vertices u, v in a graph G, denote by $G[\,u \to v\,]$ the graph obtained from G by first deleting all the edges at u and then joining u to all the neighbours of v. Show that $K^r \not\subseteq G[\,u \to v\,]$ if $K^r \not\subseteq G$. Applying this operation repeatedly to a given extremal graph for n and K^r, prove that $\mathrm{ex}(n, K^r) = t_{r-1}(n)$: in each iteration step, choose u and v so that the number of edges will not decrease, and so that eventually a complete multipartite graph is obtained.

11. Show that deleting at most $(m-s)(n-t)/s$ edges from a $K_{m,n}$ will never destroy all its $K_{s,t}$ subgraphs.

12. For $0 < s \leqslant t \leqslant n$ let $z(n,s,t)$ denote the maximum number of edges in a bipartite graph whose partition sets both have size n, and which does not contain a $K_{s,t}$. Show that $2\,\mathrm{ex}(n, K_{s,t}) \leq z(n,s,t) \leq \mathrm{ex}(2n, K_{s,t})$.

13.$^+$ Let $1 \leqslant r \leqslant n$ be integers. Let G be a bipartite graph with bipartition $\{A, B\}$, where $|A| = |B| = n$, and assume that $K_{r,r} \not\subseteq G$. Show that

$$\sum_{x \in A} \binom{d(x)}{r} \leqslant (r-1)\binom{n}{r}.$$

Using the previous exercise, deduce that $\mathrm{ex}(n, K_{r,r}) \leqslant c n^{2-1/r}$ for some constant c depending only on r.

Our version of the Erdős-Stone theorem is a slight simplification of the original. A direct proof, not using the regularity lemma, is given in L. Lovász, *Combinatorial Problems and Exercises* (2nd edn.), North-Holland 1993. Its most fundamental application, Corollary 7.1.3, was only found 20 years after the theorem, by Erdős and Simonovits (1966).

Of our two bounds on $\mathrm{ex}(n, K_{r,r})$ the upper one is thought to give the correct order of magnitude. For vastly off-diagonal complete bipartite graphs this was verified by J. Kollár, L. Rónyai & T. Szabó, Norm-graphs and bipartite Turán numbers, *Combinatorica* **16** (1996), 399–406, who proved that $\mathrm{ex}(n, K_{r,s}) \geqslant c_r n^{2-\frac{1}{r}}$ when $s > r!$.

Details about the Erdős-Sós conjecture, including an approximate solution for large k, can be found in the survey by Komlós and Simonovits cited below. The case where the tree T is a path (Exercise 18) was proved by Erdős & Gallai in 1959. It was this result, together with the easy case of stars (Exercise 17) at the other extreme, that inspired the conjecture as a possible unifying result.

The regularity lemma is proved in E. Szemerédi, Regular partitions of graphs, *Colloques Internationaux CNRS* **260** *Problèmes Combinatoires et Théorie des Graphes, Orsay* (1976), 399–401. Our rendering follows an account by Scott (personal communication). A broad survey on the regularity lemma and its applications is given by J. Komlós & M. Simonovits in (D. Miklós, V.T. Sós & T. Szőnyi, eds.) *Paul Erdős is 80*, Vol. 2, Proc. Colloq. Math. Soc. János Bolyai (1996); the concept of a regularity graph and Lemma 7.3.2 are taken from this paper. An adaptation of the regularity lemma for use with sparse graphs was developed independently by Kohayakawa and by Rödl; see Y. Kohayakawa, Szemerédi's regularity lemma for sparse graphs, in (F. Cucker & M. Shub, eds.) *Foundations of Computational Mathematics*, Selected papers of a conference held at IMPA in Rio de Janeiro, January 1997, Springer 1997.

8 Substructures in Sparse Graphs

In this chapter we study how global assumptions about a graph—on its average degree, chromatic number, or even (large) girth—can force it to contain a given graph H as a minor or topological minor. As we know already from Mader's theorem 3.6.1, there exists a function h such that an average degree of $d(G) \geqslant h(r)$ suffices to create a TK^r subgraph in G, and hence a (topological) H minor if $r \geqslant |H|$. Since a graph with n vertices and average degree d has $\frac{1}{2}dn$ edges this shows that, for every H, there is a 'constant' c (depending on H but not on n) such that a topological H minor occurs in every graph with n vertices and at least cn edges. Such graphs with a number of edges about linear[1] in their order are called *sparse*—so this is a chapter about substructures in sparse graphs. *sparse*

The first question, then, will be the analogue of Turán's theorem: given a positive integer r, what is the minimum value of the above 'constant' c for $H = K^r$, i.e. the smallest growth rate of a function $h(r)$ as in Theorem 3.6.1? This was a major open problem until very recently; we present its solution, which builds on some fascinating methods the problem has inspired over time, in Section 8.1.

If raising the average degree suffices to force the occurrence of a certain minor, then so does raising any other invariant which in turn forces up the average degree. For example, if $d(G) \geqslant c$ implies $H \preccurlyeq G$, then so will $\chi(G) \geqslant c + 1$ (by Corollary 5.2.3). However, is this best possible? Even if the value of c above is least possible for $d(G) \geqslant c$ to imply $H \preccurlyeq G$, it need not be so for $\chi(G) \geqslant c+1$ to imply $H \preccurlyeq G$. One of the most famous conjectures in graph theory, the *Hadwiger conjecture*,

[1] Compare the footnote at the beginning of Chapter 7.

existence of some large linked subset $Z \subseteq Y$. This would be the case if G were (k, ℓ)-linked for some $k \leqslant |Y|$ and $\ell \geqslant |Z|$.

As above, a large enough constant c will easily ensure that X and Y can be chosen with many vertices to spare. Another problem, however, is more serious: it will not be enough to make ℓ (and hence Z) large in absolute terms. Indeed, if k (and Y) is much larger still, it might happen that Z, although large, consists of neighbours of only a few vertices in X! We thus have to ensure that ℓ is large also relative to k. This will be the purpose of our first lemma (8.1.2): it establishes a sufficient condition for G to be $(k, \lceil k/2 \rceil)$-linked.

What is this sufficient condition? It is the assumption that G has a particularly dense minor H, one whose minimum degree exceeds $\frac{1}{2}|H|$ by a positive fraction of k. (In particular, H will be dense in the sense of Chapter 7.) In view of Theorem 3.6.2, it is not surprising that such a dense graph H is highly linked. Given sufficiently high connectivity of G (which again follows easily if c is large enough), we may then try to link up the vertices of any Y as above to distinct branch sets of H by disjoint paths in G avoiding most of the other branch sets, and thus to transfer the linking properties of H to a $\lceil k/2 \rceil$-set $Z \subseteq Y$ (Fig. 8.1.1).

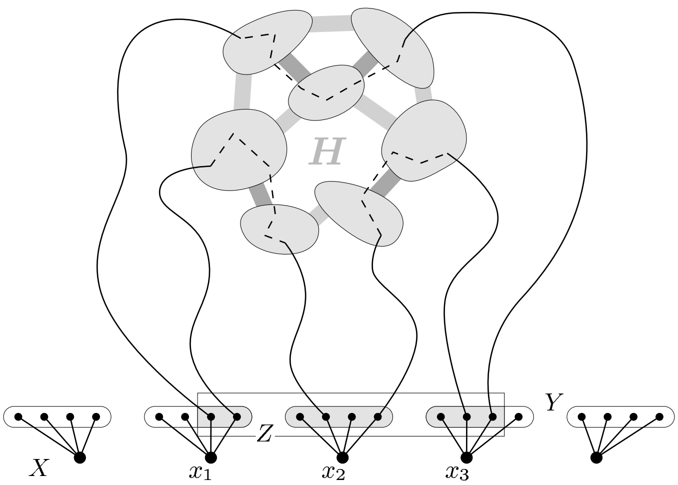

Fig. 8.1.1. Finding a TK^3 in G with branch vertices x_1, x_2, x_3

What is all the more surprising, however, is that the existence of such a dense minor H can be deduced from our assumption of $d(G) \geqslant cr^2$. This will be shown in another lemma (8.1.3); the assertion of the theorem itself will then follow easily.

Lemma 8.1.2. *If G is k-connected and has a minor H with $2\delta(H) \geqslant |H| + \frac{3}{2}k$, then G is $(k, \lceil k/2 \rceil)$-linked.*

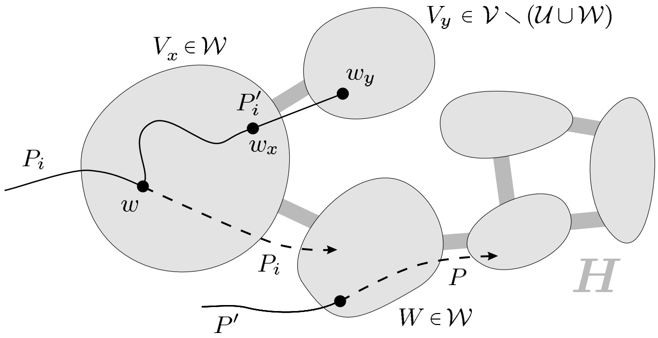

Fig. 8.1.2. If P_i does not end in V_x, we replace P_i and P by P_i' and P'

We now assume that P_i does end in V_x; then $f(P_i') = f(P_i) + 1$. As $V_x \in \mathcal{W}$, there exists a link P_j that meets V_x and leaves it again; let P_j' be the initial segment of P_j ending in V_x (Fig 8.1.3). Then $f(P_j') \leqslant f(P_j) - 1$. In fact, since replacing P_i and P_j with P_i' and P_j' in $\mathcal{P}$ yields another linkage, the choice of $\mathcal{P}$ implies that $f(P_j') = f(P_j) - 1$, so P_j ends in the branch set W it enters immediately after V_x. Then $W \in \mathcal{W}$ as before, so we may define P and P' as before. Replacing P_i, P_j and P by P_i', P_j' and P' in $\mathcal{P}$, we finally obtain a linkage that contradicts the choice of $\mathcal{P}$. This completes the proof of (1).

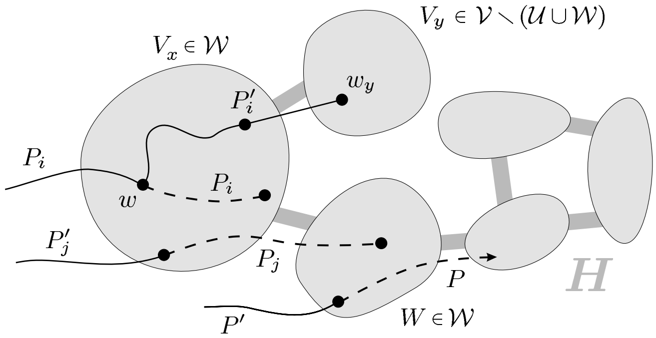

Fig. 8.1.3. If P_i ends in V_x, we replace P_i, P_j, P by P_i', P_j', P'

With the help of (1) we may define an injection $\mathcal{W} \to \mathcal{U}$ as follows: given $W \in \mathcal{W}$, choose a link that passes through W and next meets a set $U \in \mathcal{U}$, and map $W \mapsto U$. (This is indeed an injection, because different links end in different branch sets.) Thus $|\mathcal{U}| \geqslant |\mathcal{W}|$, and hence $|\mathcal{U}| \geqslant \lceil k/2 \rceil$.

Let us assume the enumeration of $v_1, \dots, v_k$ to be such that the first $u := |\mathcal{U}|$ of the links $P_1, \dots, P_k$ end in sets from $\mathcal{U}$. Since $2\delta(H) \geqslant$ u

Thus in either case we have found an integer $m \geqslant k/2$ and a graph $G_1 \preccurlyeq G$ such that m G_1

$$|G_1| \leqslant 4m \tag{1}$$

and $\delta(G_1) \geqslant 2m$, so

$$\varepsilon(G_1) \geqslant m \geqslant k/2 \geqslant 3\,. \tag{2}$$

As $2\delta(G_1) \geqslant 4m \geqslant |G_1|$, our graph G_1 is already quite a good candidate for the desired minor H of G. In order to jack up its value of 2δ by another $\frac{1}{6}k$ (as required for H), we shall reapply the above contraction process to G_1, and a little more rigorously than before: step by step, we shall contract edges as long as this results in a loss of no more than $\frac{7}{6}m$ edges per vertex. In other words, we permit a loss of edges slightly greater than maintaining $\varepsilon \geqslant m$ seems to allow. (Recall that, when we contracted G to G_0, we put this threshold at $\varepsilon(G) = k$.) *If* this second contraction process terminates with a non-empty graph H_0, then $\varepsilon(H_0)$ will be at least $\frac{7}{6}m$, higher than for G_1! The $\frac{1}{6}m$ thus gained will suffice to give the graph H_1, obtained from H_0 just as G_1 was obtained from G_0, the desired high minimum degree.

But how can we be sure that this second contraction process will indeed end with a non-empty graph? Paradoxical though it may seem, the reason is that even a permitted loss of up to $\frac{7}{6}m$ edges (and one vertex) per contraction step cannot destroy the $m\,|G_1|$ or more edges of G_1 in the $|G_1|$ steps possible: the graphs with fewer than m vertices towards the end of the process would simply be too small to be able to shed their allowance of $\frac{7}{6}m$ edges—and, by (1), these small graphs would account for about a quarter of the process!

Formally, we shall control the graphs H in the contraction process not by specifying an upper bound the number of edges to be discarded at each step, but by fixing a lower bound for $\|H\|$ in terms of $|H|$. This bound grows linearly from a value of just above $\binom{m}{2}$ for $|H| = m$ to a value of less than $4m^2$ for $|H| = 4m$. By (1) and (2), $H = G_1$ will satisfy this bound, but clearly it cannot be satisfied by any H with $|H| = m$; so the contraction process must stop somewhere earlier with $|H| > m$.

To implement this approach, let

$$f(n) := \tfrac{1}{6}m(n - m - 5)$$ f

and

$$\mathcal{H} := \left\{ H \preccurlyeq G_1 : \|H\| \geqslant m\,|H| + f(|H|) - \tbinom{m}{2} \right\}.$$ $\mathcal{H}$

By (1),

$$f(|G_1|) \leqslant f(4m) = \tfrac{1}{2}m^2 - \tfrac{5}{6}m < \tbinom{m}{2}\,,$$

so $G_1 \in \mathcal{H}$ by (2).

By the choice of x_1 and definition of H_1, therefore,

$$\begin{aligned} |H_1| - 1 &= \delta(H_0) \\ &\leqslant 2\,\varepsilon(H_0) \\ &\underset{(4)}{<} \tfrac{7}{3}m - \tfrac{4}{3}m^2/|H_0| \\ &\underset{(1)}{\leqslant} \tfrac{7}{3}m - \tfrac{1}{3}m \\ &= 2m\,, \end{aligned}$$

so $|H_1| \leqslant 2m$. Hence,

$$\begin{aligned} 2\delta(H_1) &\underset{(3)}{>} 2m + \tfrac{1}{3}m \\ &\geqslant |H_1| + \tfrac{1}{3}m \\ &\underset{(2)}{\geqslant} |H_1| + \tfrac{1}{6}k \end{aligned}$$

as claimed. □

Proof of Theorem 8.1.1. We prove the assertion for $c := 1116$. Let G be a graph with $d(G) \geqslant 1116r^2$. By Theorem 1.4.2, G has a subgraph G_0 such that (1.4.2) G_0

$$\kappa(G_0) \geqslant 279r^2 \geqslant 276r^2 + 3r\,.$$

Pick a set $X := \{\, x_1, \dots, x_{3r} \,\}$ of $3r$ vertices in G_0, and let $G_1 := G_0 - X$. For each $i = 1, \dots, 3r$ choose a set Y_i of $5r$ neighbours of x_i in G_1; let these sets Y_i be disjoint for different i. (This is possible since $\delta(G_0) \geqslant \kappa(G_0) \geqslant 15r^2 + |X|$.) X G_1, Y_i

As

$$\delta(G_1) \geqslant \kappa(G_1) \geqslant \kappa(G_0) - |X| \geqslant 276r^2,$$

we have $\varepsilon(G_1) \geqslant 138r^2$. By Lemma 8.1.3, G_1 has a minor H with $2\delta(H) \geqslant |H| + 23r^2$ and is therefore $(15r^2, 7r^2)$-linked by Lemma 8.1.2; let $Z \subseteq \bigcup_{i=1}^{3r} Y_i$ be a set of $7r^2$ vertices that is linked in G_1. Z

For all $i = 1, \dots, 3r$ let $Z_i := Z \cap Y_i$. Since Z is linked, it suffices to find r indices i with $|Z_i| \geqslant r - 1$: then the corresponding x_i will be the branch vertices of a TK^r in G_0. If r such i cannot be found, then $|Z_i| \leqslant r - 2$ for all but at most $r - 1$ indices i. But then Z_i

$$|Z| = \sum_{i=1}^{3r} |Z_i| \leqslant (r-1)\,5r + (2r+1)(r-2) < 7r^2 = |Z|\,,$$

a contradiction. □

8.2 Minors

According to Theorem 8.1.1, an average degree of cr^2 suffices to force the existence of a topological K^r minor in a given graph. If we are content with any minor, topological or not, an even smaller average degree will do: in a pioneering paper of 1968, Mader proved that every graph with an average degree of at least $cr \log r$ has a K^r minor. The following result, the analogue to Theorems 7.1.1 and 8.1.1 for general minors, determines the precise average degree needed as a function of r, up to a constant c:

Theorem 8.2.1. (Kostochka 1982; Thomason 1984)
There exists a $c \in \mathbb{R}$ such that, for every $r \in \mathbb{N}$, every graph G of average degree $d(G) \geqslant c\, r\sqrt{\log r}$ has a K^r minor. Up to the value of c, this bound is best possible as a function of r.

The easier implication of the theorem, the fact that in general an average degree of $c\, r\sqrt{\log r}$ is needed to force a K^r minor, follows from considering random graphs, to be introduced in Chapter 11. The converse implication, the fact that this average degree suffices, is proved by methods similar to those described in Section 8.1.

Rather than proving Theorem 8.2.1, we therefore devote the remainder of this section to another striking result on forcing minors. At first glance, this result is so surprising that it seems almost paradoxical: as long as we do not merely subdivide edges, we can force a K^r minor in a graph simply by raising its girth (Corollary 8.2.3)!

Theorem 8.2.2. (Thomassen 1983)
Given an integer k, every graph G with girth $g(G) \geqslant 4k-3$ and $\delta(G) \geqslant 3$ has a minor H with $\delta(H) \geqslant k$.

Proof. As $\delta(G) \geqslant 3$, every component of G contains a cycle. In particular, the assertion is trivial for $k \leqslant 2$; so let $k \geqslant 3$. Consider the vertex set V of a component of G, together with a partition $\{ V_1, \dots, V_m \}$ of V into as many connected sets V_i with at least $2k-2$ vertices each as possible. (Such a partition exists, since $|V| \geqslant g(G) > 2k-2$ and V is connected in G.) (1.5.3) V, V_i m

We first show that every $G[V_i]$ is a tree. To this end, let T_i be a spanning tree of $G[V_i]$. If $G[V_i]$ has an edge $e \notin T_i$, then $T_i + e$ contains a cycle C; by assumption, C has length at least $4k-3$. The edge (about) opposite e on C therefore separates the path $C - e$, and hence also T_i, into two components with at least $2k-2$ vertices each. Together with the sets V_j for $j \neq i$, these two components form a partition of V into $m+1$ sets that contradicts the maximality of m.

So each $G[V_i]$ is indeed a tree, i.e. $G[V_i] = T_i$. As $\delta(G) \geqslant 3$, the degrees in G of the vertices in V_i sum to at least $3|V_i|$, while the edges of T_i account for only $2|V_i| - 2$ in this sum. Hence for each i, G has T_i

8.3 Hadwiger's conjecture

As we saw in the preceding two sections, an average degree of $c\,r\sqrt{\log r}$ suffices to force an arbitrary graph to have a K^r minor, and an average degree of cr^2 forces it to have a topological K^r minor. If we replace 'average degree' above with 'chromatic number' then, with almost the same constants c, the two assertions remain true: this is because every graph with chromatic number k has a subgraph of average degree at least $k-1$ (Corollary 5.2.3).

Although both functions above, $c\,r\sqrt{\log r}$ and cr^2, are best possible (up to the constant c) for the said implications with 'average degree', the question arises whether they are still best possible with 'chromatic number'—or whether some slower-growing function would do in that case. What is lurking behind this problem about growth rates, of course, is a fundamental question about the nature of the invariant χ: can this invariant have some direct *structural* effect on a graph in terms of forcing concrete substructures, or is its effect no greater than that of the 'unstructural' property of having lots of edges somewhere, which it implies trivially?

Neither for general nor for topological minors is the answer to this question known. For general minors, however, the following conjecture of Hadwiger suggests a positive answer; the conjecture is considered by many as one of the deepest open problems in graph theory.

Conjecture. (Hadwiger 1943)
The following implication holds for every integer $r > 0$ and every graph G:

$$\chi(G) \geqslant r \;\Rightarrow\; G \succcurlyeq K^r.$$

Hadwiger's conjecture is trivial for $r \leqslant 2$, easy for $r = 3$ and $r = 4$ (exercises), and equivalent to the four colour theorem for $r = 5$ and $r = 6$. For $r \geqslant 7$, the conjecture is open. Rephrased as $G \succcurlyeq K^{\chi(G)}$, it is true for almost all graphs.[3] In general, the conjecture for $r+1$ implies it for r (exercise).

The Hadwiger conjecture for any fixed r is equivalent to the assertion that every graph without a K^r minor has an $(r-1)$-colouring. In this reformulation, the conjecture raises the question of what the graphs without a K^r minor look like: any sufficiently detailed structural description of those graphs should enable us to decide whether or not they can be $(r-1)$-coloured.

For $r = 3$, for example, the graphs without a K^r minor are precisely the forests (why?), and these are indeed 2-colourable. For $r = 4$, there

[3] See Chapter 11 for the notion of 'almost all'.

It is also possible to prove Corollary 8.3.3 by a simple direct argument (Exercise 13).

By the four colour theorem, Hadwiger's conjecture for $r = 5$ follows from the following structure theorem for the graphs without a K^5 minor, just as it follows from Proposition 8.3.1 for $r = 4$. The proof of Theorem 8.3.4 is similar to that of Proposition 8.3.1, but considerably longer. We therefore state the theorem without proof:

Theorem 8.3.4. (Wagner 1937)
Let G be an edge-maximal graph without a K^5 minor. If $|G| \geqslant 4$ then G can be constructed recursively, by pasting along triangles and K^2s, from plane triangulations and copies of the graph W (Fig. 8.3.1).

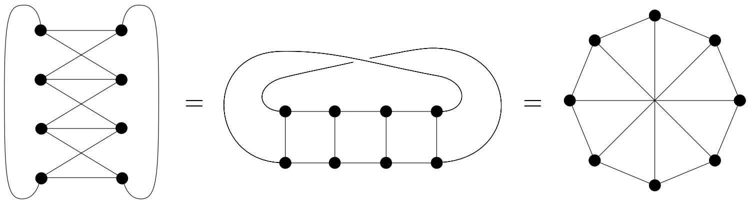

Fig. 8.3.1. Three representations of the *Wagner graph W*

Using Corollary 4.2.8, one can easily compute which of the graphs constructed as in Theorem 8.3.4 have the most edges. It turns out that these *extremal* graphs without a K^5 minor have no more edges than those that are extremal with respect to $\{ MK^5, MK_{3,3} \}$, i.e. the maximal planar graphs: 4.2.8

Corollary 8.3.5. *A graph with n vertices and no K^5 minor has at most $3n - 6$ edges.* □

Since $\chi(W) = 3$, Theorem 8.3.4 and the four colour theorem imply Hadwiger's conjecture for $r = 5$:

Corollary 8.3.6. *Hadwiger's conjecture holds for $r = 5$.* □

The Hadwiger conjecture for $r = 6$ is again substantially more difficult than the case $r = 5$, and again it relies on the four colour theorem. The proof shows (without using the four colour theorem) that any minimal-order counterexample arises from a planar graph by adding one vertex—so by the four colour theorem it is not a counterexample after all.

Theorem 8.3.7. (Robertson, Seymour & Thomas 1993)
Hadwiger's conjecture holds for $r = 6$.

9.$^+$ For which trees T is there a function $f:\mathbb{N}\to\mathbb{N}$ tending to infinity, such that every graph G with $\chi(G) < f(d(G))$ contains an induced copy of T? (In other words: can we force the chromatic number up by raising the average degree, as long as T does not occur as an induced subgraph? Or, as in Gyárfás's conjecture: will a large average degree force an induced copy of T if the chromatic number is kept small?)

10.$^-$ Derive the four colour theorem from Hadwiger's conjecture for $r = 5$.

11.$^-$ Show that Hadwiger's conjecture for $r+1$ implies the conjecture for r.

12.$^-$ Using the results from this chapter, prove the following weakening of Hadwiger's conjecture: given any $\epsilon > 0$, every graph of chromatic number at least $r^{1+\epsilon}$ has a K^r minor, provided that r is large enough.

13.$^+$ Prove Hadwiger's conjecture for $r = 4$ from first principles.

14.$^+$ Prove Hadwiger's conjecture for line graphs.

15. (i)$^-$ Show that Hadwiger's conjecture is equivalent to the statement that $G \succcurlyeq K^{\chi(G)}$ for all graphs G.

(ii) Show that any minimum-order counterexample G to Hadwiger's conjecture (as rephrased above) satisfies $K^{\chi(G)-1} \not\subseteq G$ and has a connected complement.

16. Show that any graph constructed as in Theorem 8.3.1 is edge-maximal without a K^4 minor.

17. Prove the implication $\delta(G) \geqslant 3 \Rightarrow G \supseteq TK^4$.

(Hint. Theorem 8.3.1.)

18. A multigraph is called *series-parallel* if it can be constructed recursively from a K^2 by the operations of subdividing and of doubling edges. Show that a 2-connected multigraph is series-parallel if and only if it has no (topological) K^4 minor.

19. Prove Corollary 8.3.5.

20. Characterize the graphs with n vertices and more than $3n-6$ edges that contain no $TK_{3,3}$. In particular, determine $\mathrm{ex}(n, TK_{3,3})$.

(Hint. By a theorem of Wagner, every edge-maximal graph without a $K_{3,3}$ minor can be constructed recursively from maximal planar graphs and copies of K^5 by pasting along K^2s.)

21. By a theorem of Pelikán, every graph of minimum degree at least 4 contains a subdivision of K^5_-, a K^5 minus an edge. Using this theorem, prove Thomassen's 1974 result that every graph with $n \geqslant 5$ vertices and at least $4n-10$ edges contains a TK^5.

(Hint. Show by induction on $|G|$ that if $\|G\| \geqslant 4n-10$ then for every vertex $x \in G$ there is a $TK^5 \subseteq G$ in which x is not a branch vertex.)

general H, Mader improved Theorem 8.1.4 by weakening the requirement of $\delta(G) \geqslant d$ to $d(G) \geqslant d-1+\epsilon$ for arbitrary $\epsilon > 0$ (where now the girth k required to force a TH in such graphs G depends on ϵ as well as on H); see W. Mader, Subdivisions of a graph of maximal degree $n+1$ in graphs of average degree $n+\epsilon$ and large girth, manuscript 1999.

Theorem 8.1.5 is due to A.D. Scott, Induced trees in graphs of large chromatic number, *J. Graph Theory* **24** (1997), 297–311. Theorem 8.2.1 was proved independently by Kostochka (1982; English translation: A.V. Kostochka, Lower bounds of the Hadwiger number of graphs by their average degree, *Combinatorica* **4** (1984), 307–316) and by A.G. Thomason, An extremal function for contractions of graphs, *Math. Proc. Camb. Phil. Soc.* **95** (1984), 261–265. Theorem 8.2.2 was taken from Thomassen's survey, Paths, Circuits and Subdivisions, in (L.W. Beineke & R.J. Wilson, eds.) *Selected Topics in Graph Theory 3*, Academic Press 1988.

The proof of Hadwiger's conjecture for $r = 4$, hinted at in Exercise 13, is given by Hadwiger himself in the 1943 paper containing his conjecture. For a while, there was a counterpart to Hadwiger's conjecture for topological minors, the conjecture of Hajós that $\chi(G) \geqslant r$ even implies $G \supseteq TK^r$. A counterexample to this conjecture was found in 1979 by Catlin; a little later, Erdős and Fajtlowicz even proved that Hajós's conjecture is false for *almost all* graphs (see Chapter 11).

Mader's Theorem 8.3.8 that $3n-5$ edges force a topological K^5 minor had been conjectured by Dirac in 1964. Its proof comprises two papers: W. Mader, $3n-5$ edges do force a subdivision of K_5, *Combinatorica* **18** (1998), 569–595; and W. Mader, An extremal problem for subdivisions of K_5^-, *J. Graph Theory* **30** (1999), 261–276. His proof of Theorem 8.3.9 has not been published yet. Dirac's conjecture has been extended by Seymour, who conjectures that every 5-connected non-planar graph should contain a TK^5 (unpublished).

9 Ramsey Theory for Graphs

In this chapter we set out from a type of problem which, on the face of it, appears to be similar to the theme of the last two chapters: what kind of substructures are necessarily present in every large enough graph?

The regularity lemma of Chapter 7.2 provides one possible answer to this question, saying as it does that every (large) graph G contains large random-like bipartite subgraphs. If we are looking for more definite substructures, however, such as subgraphs isomorphic to some given graphs H, then these H will have to be sufficiently complementary in kind to cater for the variety allowed for G. For example: given an integer r, does every large enough graph contain either a K^r or an induced $\overline{K^r}$? Does every large enough connected graph contain either a K^r or else a large induced path or star?

Despite its similarity to extremal problems in that we are looking for local implications of global assumptions, the above type of question leads to a kind of mathematics with a distinctive flavour of its own. Indeed, the theorems and proofs in this chapter have more in common with similar results in algebra or geometry, say, than with most other areas of graph theory. The study of their underlying methods, therefore, is generally regarded as a combinatorial subject in its own right: the discipline of *Ramsey theory*.

In line with the subject of this book, we shall focus on results that are naturally expressed in terms of graphs. Even from the viewpoint of general Ramsey theory, however, this is not as much of a limitation as it might seem: graphs are a natural setting for Ramsey problems, and the material in this chapter brings out a sufficient variety of ideas and methods to convey some of the fascination of the theory as a whole.

(ii) $V_i \subseteq V_{i-1} \setminus \{ v_{i-1} \}$ $(i = 2, \dots, 2r-2)$;

(iii) v_{i-1} is adjacent either to all vertices in V_i or to no vertex in V_i $(i = 2, \dots, 2r-2)$.

Let $V_1 \subseteq V(G)$ be any set of 2^{2r-3} vertices, and pick $v_1 \in V_1$ arbitrarily. Then (i) holds for $i = 1$, while (ii) and (iii) hold trivially. Suppose now that V_{i-1} and $v_{i-1} \in V_{i-1}$ have been chosen so as to satisfy (i)–(iii) for $i-1$, where $1 < i \leqslant 2r-2$. Since

$$|V_{i-1} \setminus \{ v_{i-1} \}| = 2^{2r-1-i} - 1$$

is odd, V_{i-1} has a subset V_i satisfying (i)–(iii); we pick $v_i \in V_i$ arbitrarily.

Among the $2r-3$ vertices $v_1, \dots, v_{2r-3}$, there are $r-1$ vertices that show the same behaviour when viewed as v_{i-1} in (iii), being adjacent either to all the vertices in V_i or to none. Accordingly, these $r-1$ vertices and v_{2r-2} induce either a K^r or a $\overline{K^r}$ in G, because $v_i, \dots, v_{2r-2} \in V_i$ for all i. □

The least integer n associated with r as in Theorem 9.1.1 is the *Ramsey number* $R(r)$ of r; our proof shows that $R(r) \leqslant 2^{2r-3}$. In Chapter 11 we shall use a simple probabilistic argument to show that $R(r)$ is bounded below by $2^{r/2}$ (Theorem 11.1.3). *Ramsey number $R(r)$*

It is customary in Ramsey theory to think of partitions as colourings: a *colouring* of (the elements of) a set X *with c colours*, or *c-colouring* for short, is simply a partition of X into c classes (indexed by the 'colours'). In particular, these colourings need not satisfy any non-adjacency requirements as in Chapter 5. Given a c-colouring of $[X]^k$, the set of all k-subsets of X, we call a set $Y \subseteq X$ *monochromatic* if all the elements of $[Y]^k$ have the same colour,[1] i.e. belong to the same of the c partition classes of $[X]^k$. Similarly, if $G = (V, E)$ is a graph and and all the edges of $H \subseteq G$ have the same colour in some colouring of E, we call H a *monochromatic subgraph* of G, speak of a red (green, etc.) H in G, and so on. *c-colouring* *$[X]^k$* *monochromatic*

In the above terminology, Ramsey's theorem can be expressed as follows: for every r there exists an n such that, given any n-set X, every 2-colouring of $[X]^2$ yields a monochromatic r-set $Y \subseteq X$. Interestingly, this assertion remains true for c-colourings of $[X]^k$ with arbitrary c and k—with almost exactly the same proof!

To avoid repetition, we shall use this opportunity to demonstrate a common alternative proof technique: we first prove an infinite version of the general Ramsey theorem (which is easier, because we need not worry about numbers), and then deduce the finite version by a so-called *compactness argument*.

[1] Note that Y is called monochromatic, but it is the elements of $[Y]^k$, not of Y, that are (equally) coloured.

paths, infinitely many agree even on their vertex v_2 in V_2—and so on. Although the set of paths considered decreases from step to step, it is still infinite after any finite number of steps, so v_n gets defined for every $n \in \mathbb{N}$. By definition, each vertex v_n is adjacent to v_{n-1} on one of those paths, so $v_0 v_1 \dots$ is indeed an infinite path. □

Theorem 9.1.4. *For all $k, c, r \geqslant 1$ there exists an $n \geqslant k$ such that every n-set X has a monochromatic r-subset with respect to any c-colouring of $[X]^k$.* [9.3.3]

Proof. As is customary in set theory, we denote by $n \in \mathbb{N}$ (also) the set $\{0, \dots, n-1\}$. Suppose the assertion fails for some k, c, r. Then for every $n \geqslant k$ there exist an n-set, without loss of generality the set n, and a c-colouring $[n]^k \to c$ such that n contains no monochromatic r-set. Let us call such colourings *bad*; we are thus assuming that for every $n \geqslant k$ there exists a bad colouring of $[n]^k$. Our aim is to combine these into a bad colouring of $[\mathbb{N}]^k$, which will contradict Theorem 9.1.2.

k, c, r

bad colouring

For every $n \geqslant k$ let $V_n \neq \emptyset$ be the set of bad colourings of $[n]^k$. For $n > k$, the restriction $f(g)$ of any $g \in V_n$ to $[n-1]^k$ is still bad, and hence lies in V_{n-1}. By the infinity lemma, there is an infinite sequence $g_k, g_{k+1}, \dots$ of bad colourings $g_n \in V_n$ such that $f(g_n) = g_{n-1}$ for all $n > k$. For every $m \geqslant k$, all colourings g_n with $n \geqslant m$ agree on $[m]^k$, so for each $Y \in [\mathbb{N}]^k$ the value of $g_n(Y)$ coincides for all $n > \max Y$. Let us define $g(Y)$ as this common value $g_n(Y)$. Then g is a bad colouring of $[\mathbb{N}]^k$: every r-set $S \subseteq \mathbb{N}$ is contained in some sufficiently large n, so S cannot be monochromatic since g coincides on $[n]^k$ with the bad colouring g_n. □

The least integer n associated with k, c, r as in Theorem 9.1.4 is the *Ramsey number* for these parameters; we denote it by $R(k, c, r)$.

Ramsey number $R(k,c,r)$

9.2 Ramsey numbers

Ramsey's theorem may be rephrased as follows: if $H = K^r$ and G is a graph with sufficiently many vertices, then either G itself or its complement $\overline{G}$ contains a copy of H as a subgraph. Clearly, the same is true for any graph H, simply because $H \subseteq K^h$ for $h := |H|$.

However, if we ask for the *least* n such that every graph G with n vertices has the above property—this is the *Ramsey number* $R(H)$ of H—then the above question makes sense: if H has only few edges, it should embed more easily in G or $\overline{G}$, and we would expect $R(H)$ to be smaller than the Ramsey number $R(h) = R(K^h)$.

Ramsey number $R(H)$

A little more generally, let $R(H_1, H_2)$ denote the least $n \in \mathbb{N}$ such that $H_1 \subseteq G$ or $H_2 \subseteq \overline{G}$ for every graph G of order n. For most graphs

$R(H_1, H_2)$

or R''_s. Since $\chi(H) \leqslant \Delta(H)+1 \leqslant \Delta+1$, this will be the case if $s \geqslant \alpha(H)$ and we can find a $K^{\Delta+1}$ in R' or in R''. But that is easy to ensure: we just need that $K^r \subseteq R$, where r is the Ramsey number of $\Delta+1$, which will follow from Turán's theorem because R is dense.

For the formal proof let now $\Delta \geqslant 1$ be given. On input $d := 1/2$ and Δ, Lemma 7.3.2 returns an ϵ_0; since the lemma's assertion about ϵ_0 becomes weaker if ϵ_0 is made smaller, we may assume that $\epsilon_0 < 1$. Let $m := R(\Delta+1)$ be the Ramsey number of $\Delta+1$. Let $\epsilon \leqslant \epsilon_0$ be positive but small enough that, for $k = m$ (and hence for all $k \geqslant m$), Δ, d ϵ_0 m, ϵ

$$2\epsilon < \frac{1}{m-1} - \frac{1}{k}\,. \tag{1}$$

Finally, let M be the integer returned by the regularity lemma (7.2.1) on input ϵ and m. M

All the quantities defined so far depend only on Δ. We shall prove the theorem with

$$c := \frac{M}{\epsilon_0\,(1-\epsilon)}\,.$$ c

So let H with $\Delta(H) \leqslant \Delta$ be given, and let $s := |H|$. Let G be an arbitrary graph of order $n \geqslant c\,|H|$; we show that $H \subseteq G$ or $H \subseteq \overline{G}$. s G, n

By Lemma 7.2.1, G has an ϵ-regular partition $\{\, V_0, V_1, \dots, V_k \,\}$ with exceptional set V_0 and $|V_1| = \ldots = |V_k| =: \ell$, where $m \leqslant k \leqslant M$. Then k ℓ

$$\ell = \frac{n-|V_0|}{k} \geqslant \frac{n-\epsilon n}{M} = n\,\frac{1-\epsilon}{M} \geqslant cs\,\frac{1-\epsilon}{M} = \frac{s}{\epsilon_0}\,. \tag{2}$$

Let R be the regularity graph with parameters $\epsilon, \ell, 0$ corresponding to this partition. By definition, R has k vertices and R

$$\begin{aligned}
\|R\| &\geqslant \binom{k}{2} - \epsilon k^2 \\
&= \tfrac12 k^2 \Big(1 - \frac{1}{k} - 2\epsilon\Big) \\
&\underset{(1)}{>} \tfrac12 k^2 \Big(1 - \frac{1}{k} - \frac{1}{m-1} + \frac{1}{k}\Big) \\
&= \tfrac12 k^2\,\frac{m-2}{m-1} \\
&\geqslant t_{m-1}(k)
\end{aligned}$$

edges. By Theorem 7.1.1, therefore, R has a subgraph $K = K^m$. K

We now colour the edges of R with two colours: red if the edge corresponds to a pair (V_i, V_j) of density at least 1/2, and green otherwise. Let R' be the spanning subgraph of R formed by the red edges, and R''

9.3 Induced Ramsey theorems

Ramsey's theorem can be rephrased as follows. For every graph $H = K^r$ there *exists* a graph G such that every 2-colouring of the edges of G yields a monochromatic $H \subseteq G$; as it turns out, this is witnessed by any large enough complete graph as G. Let us now change the problem slightly and ask for a graph G in which every 2-edge-colouring yields a monochromatic *induced* $H \subseteq G$, where H is now an arbitrary given graph.

This slight modification changes the character of the problem dramatically. What is needed now is no longer a simple proof that G is 'big enough' (as for Theorem 9.1.1), but a careful construction: the construction of a graph that, however we bipartition its edges, contains an induced copy of H with all edges in one partition class. We shall call such a graph a *Ramsey graph* for H. *Ramsey graph*

The fact that such a Ramsey graph exists for every choice of H is one of the fundamental results of graph Ramsey theory. It was proved around 1973, independently by Deuber, by Erdős, Hajnal & Pósa, and by Rödl.

Theorem 9.3.1. *Every graph has a Ramsey graph. In other words, for every graph H there exists a graph G that, for every partition $\{ E_1, E_2 \}$ of $E(G)$, has an induced subgraph H with $E(H) \subseteq E_1$ or $E(H) \subseteq E_2$.*

We give two proofs. Each of these is highly individual, yet each offers a glimpse of true Ramsey theory: the graphs involved are used as hardly more than bricks in the construction, but the edifice is impressive.

First proof. In our construction of the desired Ramsey graph we shall repeatedly replace vertices of a graph $G = (V, E)$ already constructed by copies of another graph H. For a vertex set $U \subseteq V$ let $G[U \to H]$ denote the graph obtained from G by replacing the vertices $u \in U$ with copies $H(u)$ of H and joining each $H(u)$ completely to all $H(u')$ with $uu' \in E$ and to all vertices $v \in V \setminus U$ with $uv \in E$ (Fig. 9.3.1). Formally, $G[U \to H]$ $H(u)$

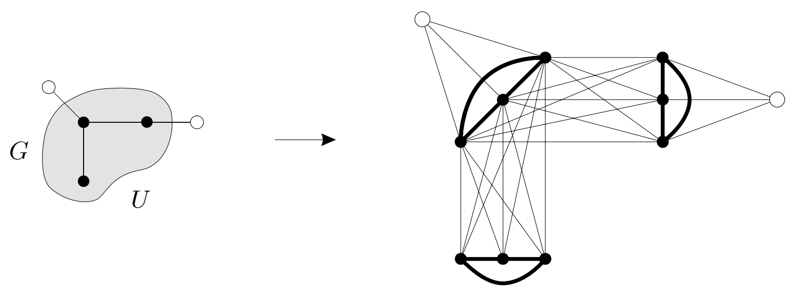

Fig. 9.3.1. A graph $G[U \to H]$ with $H = K^3$

the set of all vertices in V^i with origin w, and the map f links vertices with the same origin across the various G^i.

By the induction hypothesis, there are Ramsey graphs

$$G_1 := G(H_1, H_2') \quad \text{and} \quad G_2 := G(H_1', H_2)\,.$$ G_1, G_2

Let G^0 be a copy of G_1, and set $V^0 := V(G^0)$. Let $W_0', \dots, W_{n-1}'$ be the subsets of V^0 spanning an H_2' in G^0. Thus, n is defined as the number of induced copies of H_2' in G^0, and we shall construct a graph G^i for every set W_{i-1}', $i = 1, \dots, n$. Since H_1 has an edge, $n \geqslant 1$: otherwise G^0 could not be a $G(H_1, H_2')$. For $i = 0, \dots, n-1$, let W_i'' be the image of $V(H_2'')$ under some isomorphism $H_2' \to G^0 [\, W_i' \,]$. G^0, V^0 W_i' n W_i''

Assume now that $G^0, \dots, G^{i-1}$ and $V^0, \dots, V^{i-1}$ have been defined for some $i \geqslant 1$, and that f has been defined on $V^1 \cup \dots \cup V^{i-1}$ and satisfies (1) for all $j \leqslant i$. We expand G^{i-1} to G^i in two steps. For the first step, consider the set U^{i-1} of all the vertices $v \in V^{i-1}$ whose origin $f^{i-1}(v)$ lies in W_{i-1}''. (For $i = 1$, this gives $U^0 = W_0''$.) Expand G^{i-1} to a graph $\tilde{G}^{i-1}$ by replacing every vertex $u \in U^{i-1}$ with a copy $G_2(u)$ of G_2, i.e. let U^{i-1} $G_2(u)$

$$\tilde{G}^{i-1} := G^{i-1} [\, U^{i-1} \to G_2 \,]$$ $\tilde{G}^{i-1}$

(see Figures 9.3.2 and 9.3.3). Set $f(u') := u$ for all $u \in U^{i-1}$ and

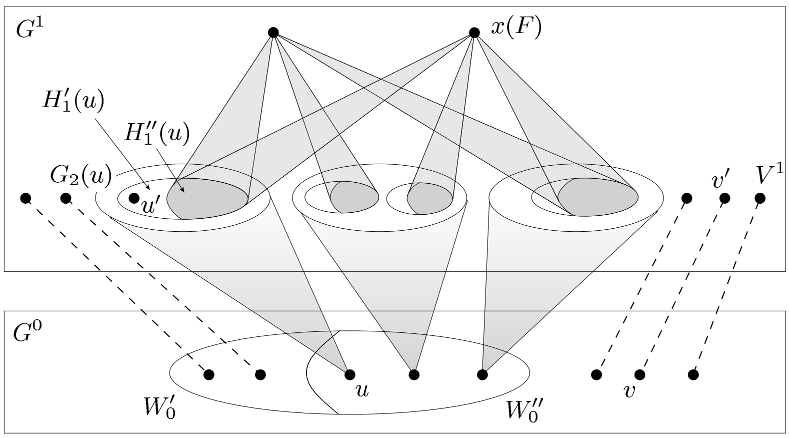

Fig. 9.3.2. The construction of G^1

$u' \in G_2(u)$, and $f(v') := v$ for all $v' = (v, \emptyset)$ with $v \in V^{i-1} \setminus U^{i-1}$. (Recall that $(v, \emptyset)$ is simply the unexpanded copy of a vertex $v \in G^{i-1}$ in $\tilde{G}^{i-1}$.) Let V^i be the set of those vertices v' or u' of $\tilde{G}^{i-1}$ for which f has thus been defined, i.e. the vertices that either correspond directly to a vertex v in V^{i-1} or else belong to an expansion $G_2(u)$ of such a vertex u. Then (1) holds for i. Also, if we assume (2) inductively for V^i

to G^{i-1}: just map every y_u to u and every $(v, \emptyset)$ to v. Our edge colouring of G^i thus induces an edge colouring of G^{i-1}. If this colouring yields an induced $H_1 \subseteq G^{i-1}$ coloured 1 or an induced $H_2 \subseteq G^{i-1}$ coloured 2, we have these also in $\hat{G}^{i-1} \subseteq G^i$ and are again home.

By (∗∗) for $i-1$ we may therefore assume that G^{i-1} has an induced subgraph H' coloured 2, with $V(H') \subseteq V^{i-1}$, and such that the restriction of f^{i-1} to $V(H')$ is an isomorphism from H' to $G^0[W'_k] \simeq H'_2$ for some $k \in \{i-1, \dots, n-1\}$. Let $\hat{H}'$ be the corresponding induced subgraph of $\hat{G}^{i-1} \subseteq G^i$ (also coloured 2); then $V(\hat{H}') \subseteq V^i$, H' $\hat{H}'$

$$f^i(V(\hat{H}')) = f^{i-1}(V(H')) = W'_k,$$

and $f^i: \hat{H}' \to G^0[W'_k]$ is an isomorphism.

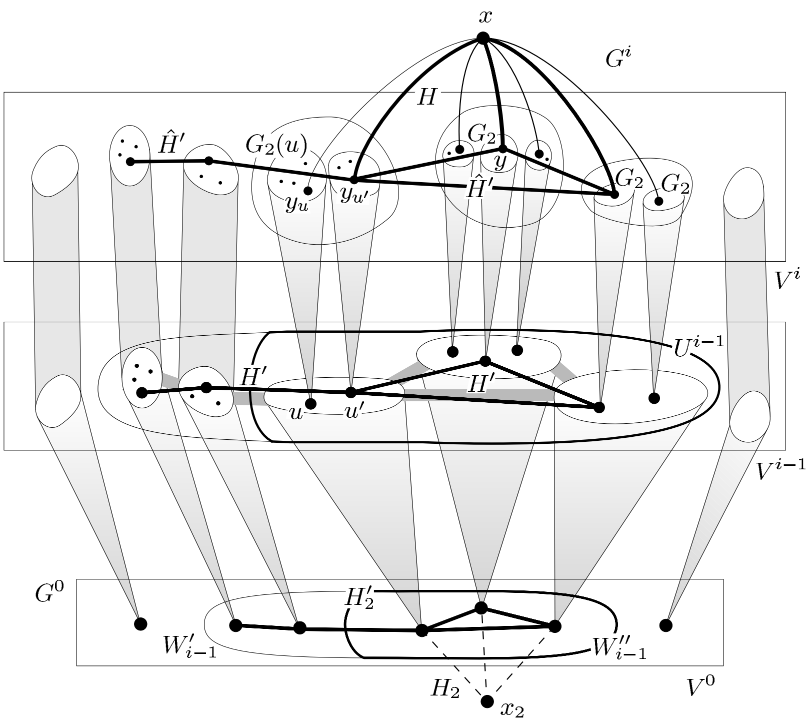

Fig. 9.3.3. A monochromatic copy of H_2 in G^i

If $k \geqslant i$, this completes the proof of (∗∗) with $H := \hat{H}'$; we therefore assume that $k < i$, and hence $k = i-1$ (Fig. 9.3.3). By definition of U^{i-1} and $\hat{G}^{i-1}$, the inverse image of W''_{i-1} under the isomorphism $f^i: \hat{H}' \to G^0[W'_{i-1}]$ is a subset of $\hat{U}^{i-1}$. Since x is joined to precisely those vertices of $\hat{H}'$ that lie in $\hat{U}^{i-1}$, and all these edges xy_u have colour 2, the graph $\hat{H}'$ and x together induce in G^i a copy of H_2 coloured 2, and the proof of (∗∗) is complete. □

Our second lemma already covers the bipartite case of the theorem: it says that every bipartite graph has a Ramsey graph—even a bipartite one.

Lemma 9.3.3. *For every bipartite graph P there exists a bipartite graph P' such that for every 2-colouring of the edges of P' there is an embedding $\varphi: P \to P'$ for which all the edges of $\varphi(P)$ have the same colour.*

Proof. We may assume by Lemma 9.3.2 that P has the form $(X, [X]^k, E)$ (9.1.4)
with $E = \{ xY \mid x \in Y \}$. We show the assertion for the graph $P' :=$ P, X, k, E
$(X', [X']^{k'}, E')$, where $k' := 2k-1$, X' is any set of cardinality P', X', k'

$$|X'| = R\left(k',\, 2\tbinom{k'}{k},\, k\,|X| + k - 1\right),$$

(this is the Ramsey number defined after Theorem 9.1.4), and

$$E' := \{ x'Y' \mid x' \in Y' \}.$$ E'

Let us then colour the edges of P' with two colours α and β. Of the α, β
$|Y'| = 2k-1$ edges incident with a vertex $Y' \in [X']^{k'}$, at least k must have the same colour. For each Y' we may therefore choose a fixed k-set
$Z' \subseteq Y'$ such that all the edges $x'Y'$ with $x' \in Z'$ have the same colour; Z'
we shall call this colour *associated* with Y'. *associated*

The sets Z' can lie within their supersets Y' in $\binom{k'}{k}$ ways, as follows. Let X' be linearly ordered. Then for every $Y' \in [X']^{k'}$ there is a unique
order-preserving bijection $\sigma_{Y'}: Y' \to \{1, \dots, k'\}$, which maps Z' to one $\sigma_{Y'}$
of $\binom{k'}{k}$ possible images.

We now colour $[X']^{k'}$ with the $2\binom{k'}{k}$ elements of the set

$$[\{1, \dots, k'\}]^k \times \{\alpha, \beta\}$$

as colours, giving each $Y' \in [X']^{k'}$ as its colour the pair $(\sigma_{Y'}(Z'), \gamma)$, where γ is the colour α or β associated with Y'. Since $|X'|$ was chosen as the Ramsey number with parameters k', $2\binom{k'}{k}$ and $k\,|X| + k - 1$, we
know that X' has a monochromatic subset W of cardinality $k\,|X| + k - 1$. W
All Z' with $Y' \subseteq W$ thus lie within their Y' in the same way, i.e. there exists an $S \in [\{1, \dots, k'\}]^k$ such that $\sigma_{Y'}(Z') = S$ for all $Y' \in [W]^{k'}$,
and all $Y' \in [W]^{k'}$ are associated with the same colour, say with α. α

We now construct the desired embedding φ of P in P'. We first $\varphi|_X$
define φ on $X =: \{x_1, \dots, x_n\}$, choosing images $\varphi(x_i) =: w_i \in W$ so x_i, w_i, n
that $w_i < w_j$ in our ordering of X' whenever $i < j$. Moreover, we choose the w_i so that exactly $k-1$ elements of W are smaller than w_1, exactly $k-1$ lie between w_i and w_{i+1} for $i = 1, \dots, n-1$, and exactly $k-1$ are bigger than w_n. Since $|W| = kn + k - 1$, this can indeed be done (Fig. 9.3.4).

Second proof of Theorem 9.3.1. Let H be given as in the theorem, and let $n := R(r)$ be the Ramsey number of $r := |H|$. Then, for every 2-colouring of its edges, the graph $K = K^n$ contains a monochromatic copy of H—although not necessarily induced. r, n K

We start by constructing a graph G^0, as follows. Imagine the vertices of K to be arranged in a column, and replace every vertex by a row of $\binom{n}{r}$ vertices. Then each of the $\binom{n}{r}$ columns arising can be associated with one of the $\binom{n}{r}$ ways of embedding $V(H)$ in $V(K)$; let us furnish this column with the edges of such a copy of H. The graph G^0 thus arising consists of $\binom{n}{r}$ disjoint copies of H and $(n-r)\binom{n}{r}$ isolated vertices (Fig. 9.3.5).

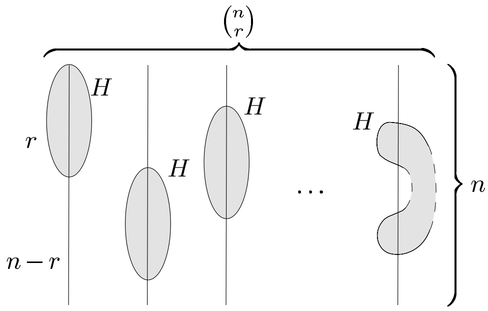

Fig. 9.3.5. The graph G^0

In order to define G^0 formally, we assume that $V(K) = \{1, \dots, n\}$ and choose copies $H_1, \dots, H_{\binom{n}{r}}$ of H in K with pairwise distinct vertex sets. (Thus, on each r-set in $V(K)$ we have one fixed copy H_j of H.) We then define

$$V(G^0) := \left\{(i,j) \mid i = 1, \dots, n;\ j = 1, \dots, \tbinom{n}{r}\right\}$$

$$E(G^0) := \bigcup_{j=1}^{\binom{n}{r}} \left\{(i,j)(i',j) \mid ii' \in E(H_j)\right\}.$$ G^0

The idea of the proof now is as follows. Our aim is to reduce the general case of the theorem to the bipartite case dealt with in Lemma 9.3.3. Applying the lemma iteratively to all the pairs of rows of G^0, we construct a very large graph G such that for every edge colouring of G there is an induced copy of G^0 in G that is monochromatic on all the bipartite subgraphs induced by its pairs of rows, i.e. in which edges between the same two rows always have the same colour. The projection of this $G^0 \subseteq G$ to $\{1, \dots, n\}$ (by contracting its rows) then defines an edge colouring of K. By the choice of $|K|$, one of the $H_j \subseteq K$ will be monochromatic. But this H_j occurs with the same colouring in the jth column of our G^0, where it is an induced subgraph of G^0, and hence of G.

9.4 Ramsey properties and connectivity

According to Ramsey's theorem, every large enough graph G has a very dense or a very sparse induced subgraph of given order, a K^r or $\overline{K^r}$. If we assume that G is connected, we can say a little more:

Proposition 9.4.1. *For every $r \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that every connected graph of order at least n contains K^r, $K_{1,r}$ or P^r as an induced subgraph.*

Proof. Let $d+1$ be the Ramsey number of r, let $n > 1 + rd^r$, and let G be a graph of order at least n. If G has a vertex v of degree at least $d+1$ then, by Theorem 9.1.1 and the choice of d, either $N(v)$ induces a K^r in G or $\{v\} \cup N(v)$ induces a $K_{1,r}$. On the other hand, if $\Delta(G) \leqslant d$, then by Proposition 1.3.3 G has radius $> r$, and hence contains two vertices at a distance $\geqslant r$. Any shortest path in G between these two vertices contains a P^r. □

The collection of 'typical' induced subgraphs in Proposition 9.4.1 is smallest possible in the following sense. If $\mathcal{G}$ is *any* set of connected graphs with the same property, i.e. such that, given $r \in \mathbb{N}$, every large enough connected graph G contains an induced copy of a graph of order $\geqslant r$ from $\mathcal{G}$, then $\mathcal{G}$ contains arbitrarily large complete graphs, stars and paths. (Note that if we take a complete graph, a star or a path as G, and then all its subgraphs are again of that type.) But Proposition 9.4.1 tells us that we need no more than these.

In principle, we could look for a set like $\mathcal{G}$ for any assumed connectivity k. We could try to find a 'minimal' set (in the above sense) of typical k-connected graphs, one such that every large k-connected graph has a large subgraph in this set. Unfortunately, $\mathcal{G}$ seems to grow very quickly with k: already for $k = 2$ it becomes thoroughly messy if (as for $k = 1$) we insist that those subgraphs be induced. By relaxing our specification of containment from 'induced subgraph' to 'topological minor' and on to 'minor', however, we can give some neat characterizations up to $k = 4$.

Proposition 9.4.2. *For every $r \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that every 2-connected graph of order at least n contains C^r or $K_{2,r}$ as a topological minor.*

(1.3.3)
(3.3.5)

Proof. Let d be the n associated with r in Proposition 9.4.1, and let G be a 2-connected graph with more than $1 + rd^r$ vertices. By Proposition 1.3.3, either G has a vertex of degree $> d$ or $\text{diam}(G) \geqslant \text{rad}(G) > r$.

In the latter case let $a, b \in G$ be two vertices at distance $> r$. By Menger's theorem (3.3.5), G contains two independent a–b paths. These form a cycle of length $> r$.

5. Sketch a proof of the following theorem of Erdős and Szekeres: for every $k \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that among any n points in the plane, no three of them collinear, there are k points spanning a convex k-gon, i.e. such that none of them lies in the convex hull of the others.

6. Prove the following result of Schur: for every $k \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that, for every partition of $\{1, \dots, n\}$ into k sets, at least one of the subsets contains numbers x, y, z such that $x + y = z$.

7. Let $(X, \leqslant)$ be a totally ordered set, and let $G = (V, E)$ be the graph on $V := [X]^2$ with $E := \{(x, y)(x', y') \mid x < y = x' < y'\}$.
   (i) Show that G contains no triangle.
   (ii) Show that $\chi(G)$ will get arbitrarily large if $|X|$ is chosen large enough.

8. A family of sets is called a Δ-*system* if every two of the sets have the same intersection. Show that every infinite family of sets of the same finite cardinality contains an infinite Δ-system.

9. Prove the following weakening of Scott's Theorem 8.1.5: for every $r \in \mathbb{N}$ and every tree T there exists a $k \in \mathbb{N}$ such that every graph G with $\chi(G) \geqslant k$ and $\omega(G) < r$ contains a subdivision of T in which no two branch vertices are adjacent in G (unless they are adjacent in T).

10. Use the infinity lemma to show that, given $k \in \mathbb{N}$, a countably infinite graph is k-colourable (in the sense of Chapter 5) if all its finite subgraphs are k-colourable.

11. Let $m, n \in \mathbb{N}$, and assume that $m - 1$ divides $n - 1$. Show that every tree T of order m satisfies $R(T, K_{1,n}) = m + n - 1$.

12. Prove that $2^c < R(2, c, 3) \leqslant 3c!$ for every $c \in \mathbb{N}$.
    (Hint. Induction on c.)

13.$^-$ Derive the statement $(*)$ in the first proof of Theorem 9.3.1 from the theorem itself, i.e. show that $(*)$ is only formally stronger than the theorem.

14. Show that the Ramsey graph G for H constructed in the second proof of Theorem 9.3.1 does indeed satisfy $\omega(G) = \omega(H)$.

15. Show that, given any two graphs H_1 and H_2, there exists a graph $G = G(H_1, H_2)$ such that, for every *vertex*-colouring of G with colours 1 and 2, there is either an induced copy of H_1 coloured 1 or an induced copy of H_2 coloured 2 in G.
    (Hint. Apply induction as in the first proof of Theorem 9.3.1.)

16. Show that every infinite connected graph contains an infinite path or an infinite star.

17.$^-$ The K^r from Ramsey's theorem, last sighted in Proposition 9.4.1, conspicuously fails to make an appearance from Proposition 9.4.2 onwards. Can it be excused?

Our first proof of Theorem 9.3.1 is based on W. Deuber, A generalization of Ramsey's theorem, in (A. Hajnal, R. Rado & V.T. Sós, eds.) *Infinite and finite sets*, North-Holland 1975. The same volume contains the alternative proof of this theorem by Erdős, Hajnal and Pósa. Rödl proved the same result in his MSc thesis at the Charles University, Prague, in 1973. Our second proof of Theorem 9.3.1, which preserves the clique number of H for G, is due to J. Nešetřil & V. Rödl, A short proof of the existence of restricted Ramsey graphs by means of a partite construction, *Combinatorica* **1** (1981), 199–202.

The two theorems in Section 9.4 are due to B. Oporowski, J. Oxley & R. Thomas, Typical subgraphs of 3- and 4-connected graphs, *J. Combin. Theory B* **57** (1993), 239–257.

10 Hamilton Cycles

In Chapter 1.8 we briefly discussed the problem of when a graph contains an Euler tour, a closed walk traversing every edge exactly once. The simple Theorem 1.8.1 solved that problem quite satisfactorily. Let us now ask the analogous question for vertices: when does a graph G contain a closed walk that contains every vertex of G exactly once? If $|G| \geqslant 3$, then any such walk is a cycle: a *Hamilton cycle* of G. If G has a Hamilton cycle, it is called *hamiltonian*. Similarly, a path in G containing every vertex of G is a *Hamilton path*.

Hamilton cycle

Hamilton path

To determine whether or not a given graph has a Hamilton cycle is much harder than deciding whether it is Eulerian, and no good characterization[1] is known of the graphs that do. We shall begin this chapter by presenting the standard sufficient conditions for the existence of a Hamilton cycle (Sections 10.1 and 10.2). The rest of the chapter is then devoted to the beautiful theorem of Fleischner that the 'square' of every 2-connected graph has a Hamilton cycle. This is one of the main results in the field of Hamilton cycles. The simple proof we present (due to Říha) is still a little longer than other proofs in this book, but not difficult.

10.1 Simple sufficient conditions

What kind of condition might be sufficient for the existence of a Hamilton cycle in a graph G? Purely global assumptions, like high edge density, will not be enough: we cannot do without the local property that every vertex has at least two neighbours. But neither is any large (but constant) minimum degree sufficient: it is easy to find graphs without a Hamilton cycle whose minimum degree exceeds any given constant bound.

[1] The notion of a 'good characterization' can be made precise; see the introduction to Chapter 12.5 and the notes for Chapter 12.

Proposition 10.1.2. *Every graph G with $|G| \geqslant 3$ and $\kappa(G) \geqslant \alpha(G)$ has a Hamilton cycle.*

Proof. Put $\kappa(G) =: k$, and let C be a longest cycle in G. Enumerate the vertices of C cyclically, say as $V(C) = \{ v_i \mid i \in \mathbb{Z}_n \}$ with $v_i v_{i+1} \in E(C)$ for all $i \in \mathbb{Z}_n$. If C is not a Hamilton cycle, pick a vertex $v \in G - C$ and a v–C fan $\mathcal{F} = \{ P_i \mid i \in I \}$ in G, where $I \subseteq \mathbb{Z}_n$ and each P_i ends in v_i. Let $\mathcal{F}$ be chosen with maximum cardinality; then $vv_j \notin E(G)$ for any $j \notin I$, and (3.3.3) k

$$|\mathcal{F}| \geqslant \min\{ k, |C| \} \tag{1}$$

by Menger's theorem (3.3.3).

For every $i \in I$, we have $i+1 \notin I$: otherwise, $(C \cup P_i \cup P_{i+1}) - v_i v_{i+1}$ would be a cycle longer than C (Fig. 10.1.2, left). Thus $|\mathcal{F}| < |C|$, and hence $|I| = |\mathcal{F}| \geqslant k$ by (1). Furthermore, $v_{i+1} v_{j+1} \notin E(G)$ for all $i, j \in I$, as otherwise $(C \cup P_i \cup P_j) + v_{i+1} v_{j+1} - v_i v_{i+1} - v_j v_{j+1}$ would be a cycle longer than C (Fig. 10.1.2, right). Hence $\{ v_{i+1} \mid i \in I \} \cup \{ v \}$ is a set of $k+1$ or more independent vertices in G, contradicting $\alpha(G) \leqslant k$. □

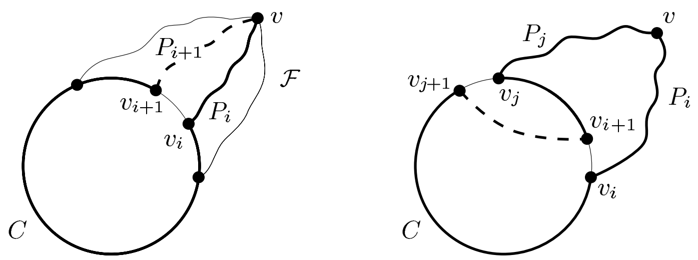

Fig. 10.1.2. Two cycles longer than C

It may come as a surprise to learn that hamiltonicity for planar graphs is related to the four colour problem. As we noted in Chapter 6.6, the four colour theorem is equivalent to the non-existence of a planar snark, i.e. to the assertion that every bridgeless planar cubic graph has a 4-flow. It is easily checked that 'bridgeless' can be replaced with '3-connected' in this assertion, and that every hamiltonian graph has a 4-flow (Ex. 12, Ch. 6). For a proof of the four colour theorem, therefore, it would suffice to show that every 3-connected planar cubic graph has a Hamilton cycle!

Unfortunately, this is not the case: the first counterexample was found by Tutte in 1946. Ten years later, Tutte proved the following deep theorem as a best possible weakening:

Theorem 10.1.3. (Tutte 1956)
Every 4-connected planar graph has a Hamilton cycle.

Then $I \cup J \subseteq \{1, \ldots, n-1\}$, and $I \cap J = \emptyset$ because G has no Hamilton cycle. Hence

$$d(x) + d(y) = |I| + |J| < n\,, \tag{3}$$

so $h := d(x) < n/2$ by the choice of x. h

Since $x_i y \notin E$ for all $i \in I$, all these x_i were candidates for the choice of x (together with y). Our choice of $\{x, y\}$ with $d(x) + d(y)$ maximum thus implies that $d(x_i) \leqslant d(x)$ for all $i \in I$. Hence G has at least $|I| = h$ vertices of degree at most h, so $d_h \leqslant h$. By (2), this implies that $d_{n-h} \geqslant n - h$, i.e. the $h+1$ vertices $v_{n-h}, \ldots, v_n$ all have degree at least $n - h$. Since $d(x) = h$, one of these vertices, z say, is not adjacent to x. Since z

$$d(x) + d(z) \geqslant h + (n-h) = n\,,$$

this contradicts the choice of x and y by (3).

Let us now show that, conversely, for every sequence $(a_1, \ldots, a_n)$ of the theorem with

$$a_h \leqslant h \quad \text{and} \quad a_{n-h} \leqslant n - h - 1$$

for some $h < n/2$, there exists a graph that has a pointwise greater degree sequence than $(a_1, \ldots, a_n)$ but no Hamilton cycle. Clearly it suffices, given h, to show this for the greatest such sequence $(a_1, \ldots, a_n)$, the sequence h

$$(\underbrace{h, \ldots, h}_{h \text{ times}}, \underbrace{n-h-1, \ldots, n-h-1}_{n-2h \text{ times}}, \underbrace{n-1, \ldots, n-1}_{h \text{ times}})\,. \tag{4}$$

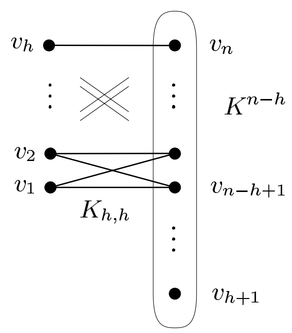

Fig. 10.2.1. Any cycle containing $v_1, \ldots, v_h$ misses v_{h+1}

As Figure 10.2.1 shows, there is indeed a graph with degree sequence (4) but no Hamilton cycle: the graph with vertices $v_1, \ldots, v_n$ and edge set

$$\{\, v_i v_j \mid i, j > h \,\} \cup \{\, v_i v_j \mid i \leqslant h;\ j > n-h \,\}\,,$$

Proof. (i) If k is even, let $Q := v_0v_2 \dots v_{k-2}v_kv_{k-1}v_{k-3} \dots v_3v_1$. If k is odd, let $Q := v_0v_2 \dots v_{k-1}v_kv_{k-2} \dots v_3v_1$.

(ii) If k is even, let $Q := v_0v_2 \dots v_{k-2}v_k$; if k is odd, let $Q := v_0v_1v_3 \dots v_{k-2}v_k$. In both cases, let Q' be the u–w path on the remaining vertices of G^2. □

Lemma 10.3.3. *Let $G = (V, E)$ be a cubic multigraph with a Hamilton cycle C. Let $e \in E(C)$ and $f \in E \setminus E(C)$ be edges with a common end v (Fig. 10.3.2). Then there exists a closed walk in G that traverses e once, every other edge of C once or twice, and every edge in $E \setminus E(C)$ once. This walk can be chosen to contain the triple (e, v, f), that is, it traverses e in the direction of v and then leaves v by the edge f.*

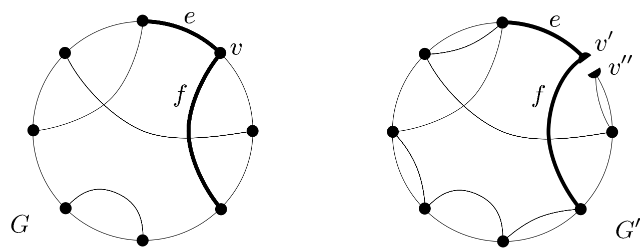

Fig. 10.3.2. The multigraphs G and G' in Lemma 10.3.3

Proof. By Proposition 1.2.1, C has even length. Replace every other edge of C by a double edge, in such a way that e does not get replaced. In the arising 4-regular multigraph G', split v into two vertices v', v'', making v' incident with e and f, and v'' incident with the other two edges at v (Fig. 10.3.2). By Theorem 1.8.1 this multigraph has an Euler tour, which induces the desired walk in G. □ (1.2.1) (1.8.1)

Lemma 10.3.4. *For every 2-connected graph G and $x \in V(G)$, there is a cycle $C \subseteq G$ that contains x as well as a vertex $y \neq x$ with $N_G(y) \subseteq V(C)$.*

Proof. If G has a Hamilton cycle, there is nothing more to show. If not, let $C' \subseteq G$ be any cycle containing x; such a cycle exists, since G is 2-connected. Let D be a component of $G - C'$. Assume that C' and D are chosen so that $|D|$ is minimal. Since G is 2-connected, D has at least two neighbours on C'. Then C' contains a path P between two such neighbours u and v, whose interior $\mathring{P}$ does not contain x and has no neighbour in D (Fig. 10.3.3). Replacing P in C' by a u–v path through D, we obtain a cycle C that contains x and a vertex $y \in D$. If y had a neighbour z in $G - C$, then z would lie in a component $D' \subsetneq D$ of $G - C$, contradicting the choice of C' and D. Hence all the neighbours of y lie on C, and C satisfies the assertion of the lemma. □

> *The elements of $\mathcal{P}$ are pairwise disjoint paths in G^2 avoiding C, and $V(G) = V(C) \cup \bigcup_{P \in \mathcal{P}} V(P)$. Every end y of a path $P \in \mathcal{P}$ has a neighbour on C in G;* we choose such a neighbour and call it the *foot* of P at y. (1)

foot

If $P \in \mathcal{P}$ is trivial, then P has exactly one foot. If P is non-trivial, then P has a foot at each of its ends. These two feet need not be distinct, however; so any non-trivial P has either one or two feet.

We shall now modify $\mathcal{P}$ a little, preserving the properties summarized under (1); no properties of $\mathcal{P}$ other than those will be used later in the proof. If a vertex of C is a foot of two distinct paths $P, P' \in \mathcal{P}$, say at $y \in P$ and at $y' \in P'$, then yy' is an edge and $Pyy'P'$ is a path in G^2; we replace P and P' in $\mathcal{P}$ by this path. We repeat this modification of $\mathcal{P}$ until the following holds:

No vertex of C is a foot of two distinct paths in $\mathcal{P}$. (2)

For $i = 1, 2$ let $\mathcal{P}_i \subseteq \mathcal{P}$ denote the set of all paths in $\mathcal{P}$ with exactly i feet, and let $X_i \subseteq V(C)$ denote the set of all feet of paths in $\mathcal{P}_i$. Then $X_1 \cap X_2 = \emptyset$ by (2), and $y^* \notin X_1 \cup X_2$.

$\mathcal{P}_1, \mathcal{P}_2$
X_1, X_2

Let us also simplify G a little; again, these changes will affect neither the paths in $\mathcal{P}$ nor the validity of (1) and (2). First, we shall assume from now on that all elements of $\mathcal{P}$ are paths in G itself, not just in G^2. This assumption may give us some additional edges for G^2, but we shall not use these in our construction of the desired Hamilton cycle H. (Indeed, H will contain all the paths from $\mathcal{P}$ whole, as subpaths.) Thus if H lies in G^2 and satisfies $(*)$ for the modified version of G, it will do so also for the original. For every $P \in \mathcal{P}$, we further delete all P–C edges in G except those between the ends of P and its corresponding feet. Finally, we delete all chords of C in G. We are thus assuming without loss of generality:

> *The only edges of G between C and a path $P \in \mathcal{P}$ are the two edges between the ends of P and its corresponding feet. (If $|P| = 1$, these two edges coincide.) The only edges of G with both ends on C are the edges of C itself.* (3)

Our goal is to construct the desired Hamilton cycle H of G^2 from the paths in $\mathcal{P}$ and suitable paths in C^2. As a first approximation, we shall construct a closed walk W in the graph

$$\tilde{G} := G - \bigcup \mathcal{P}_1 ,$$

$\tilde{G}$

a walk that will already satisfy a $(*)$-type condition and traverse every path in $\mathcal{P}_2$ exactly once. Later, we shall modify W so that it passes through every vertex of C exactly once and, finally, so as to include the

Type 3: W traverses I twice, on separate occasions (i.e., there is no triple as above).

By definition of W, the interval I^* is of type 1. The vertex x in the definition of a type 2 interval will be called the *dead end* of that interval. Finally, since Q^* is a subpath of W and W traverses both I^* and P^* only once, we have:

dead end

The interval to the right of I^* is of type 2 and has its dead end on the left. (4)

Consider a fixed interval $I = [x_1, x_2]$. Let y_1 be the neighbour of x_1, and y_2 the neighbour of x_2 on a path in $\mathcal{P}_2$. Let I^- denote the interval to the left of I.

I, x_1, x_2
y_1, y_2
I^-

Suppose first that I is of type 1. We then leave W unchanged on I. If $I \neq I^*$ we choose as $e(v)$, for each $v \in \mathring{I}$, the edge to the left of v. As $I^- \neq I^*$ by (4), and hence $x_1 \neq x^*$, these choices of $e(v)$ satisfy $(**)$. If $I = I^*$, we define $e(v)$ as the edge left of v if $v \in (x_1, x^*] \cap \mathring{I}$, and as the edge right of v if $v \in (x^*, x_2)$. These choices of $e(v)$ are again compatible with $(**)$.

Suppose now that I is of type 2. Assume first that x_2 is the dead end of I. Then W contains the walk $y_1 x_1 I x_2 I x_1 I^-$ (possibly in reverse order). We now apply Lemma 10.3.2 (i) with $P := y_1 x_1 I \mathring{x}_2$, and replace in W the subwalk $y_1 x_1 I x_2 I x_1$ by the y_1–x_1 path $Q \subseteq G^2$ of the lemma (Fig. 10.3.5). Then $V(\mathring{Q}) = V(P) \setminus \{ y_1, x_1 \} = V(\mathring{I})$. The vertices

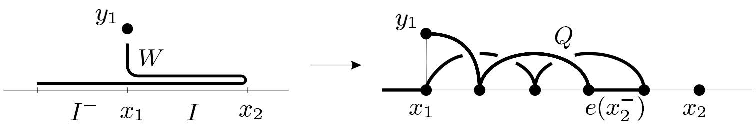

Fig. 10.3.5. How to modify W on an interval of type 2

$v \in (x_1, x_2^-)$ are each bridged by an edge of Q, which we choose as $e(v)$. As $e(x_2^-)$ we choose the edge to the left of x_2^- (unless $x_2^- = x_1$). This edge, too, lies on Q, by the lemma. Moreover, by (4) it is not incident with x^* (since x_2 is the dead end of I, by assumption) and hence satisfies $(**)$. The case that x_1 is the dead end of I can be treated in the same way: using Lemma 10.3.2 (i), we replace in W the subwalk $y_2 x_2 I x_1 I x_2$ by a y_2–x_2 path $Q \subseteq G^2$ with $V(\mathring{Q}) = V(\mathring{I})$, choose as $e(v)$ for $v \in (x_1^+, x_2)$ an edge of Q bridging v, and define $e(x_1^+)$ as the edge to the right of x_1^+ (unless $x_1^+ = x_2$).

Suppose finally that I is of type 3. Since W traverses the edge $y_1 x_1$ only once and the interval I^- no more than twice, W contains $y_1 x_1 I$ and $I^- \cup I$ as subpaths, and I^- is of type 1. By (4), however, $I^- \neq I^*$. Hence, when $e(v)$ was defined for the vertices $v \in \mathring{I}^-$, the rightmost edge $x_1^- x_1$ of I^- was not chosen as $e(v)$ for any v, so we may now replace this every $v \in C - y^*$, set

$$e(v) := \begin{cases} vv^+ & \text{if } v \in [x^*, y^*) \\ vv^- & \text{if } v \in (y^*, x^*). \end{cases}$$

(Here, $[x^*, y^*)$ and (y^*, x^*) denote the obvious paths in C defined analogously to intervals.) As before, this map $v \mapsto e(v)$ is injective, satisfies $(**)$, and is defined on a superset of X_1; recall that y^* cannot lie in X_1 by definition.

Let $P \in \mathcal{P}_1$ be a path to be incorporated into $\tilde{H}$, say with foot $v \in X_1$ and ends y_1, y_2. (If $|P| = 1$, then $y_1 = y_2$.) Our aim is to replace the edge $e := e(v)$ in $\tilde{H}$ by P; we thus have to show that the ends of P are joined to those of e by suitable edges of G^2. P, v y_1, y_2 e

By (2) and (3), v has only two neighbours in $\tilde{G}$, its neighbours x_1, x_2 on C. If v is incident with e, i.e. if $e = vx_i$ with $i \in \{1, 2\}$, we replace e by the path $vy_1Py_2x_i \subseteq G^2$ (Fig. 10.3.7). If v is not incident

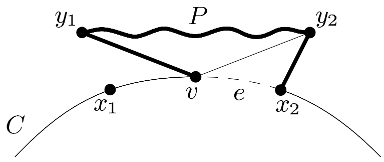

Fig. 10.3.7. Replacing the edge e in $\tilde{H}$

with e then e bridges v, by $(**)$. Then $e = x_1x_2$, and we replace e by the path $x_1y_1Py_2x_2 \subseteq G^2$ (Fig. 10.3.8). Since $v \mapsto e(v)$ is injective on X_1, assertion (2) implies that all these modifications of $\tilde{H}$ (one for every $P \in \mathcal{P}_1$) can be performed independently, and hence produce a Hamilton cycle H of G^2. H

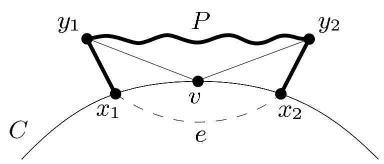

Fig. 10.3.8. Replacing the edge e in $\tilde{H}$

Let us finally check that H satisfies $(*)$, i.e. that both edges of H at x^* lie in G. Since $(*)$ holds for $\tilde{H}$, it suffices to show that any edge $e = x^*z$ of $\tilde{H}$ that is not in H (and hence has the form $e = e(v)$ for some $v \in X_1$) was replaced by an x^*–z path whose first edge lies in G. e, z v

Where can the vertex v lie? Let us show that v must be incident with e. If not then $\mathcal{P}_2 \neq \emptyset$, and e bridges v. Now $\mathcal{P}_2 \neq \emptyset$ and $v \in X_1$ together imply that $|C - y^*| \geqslant |X_1| + 2\,|\mathcal{P}_2| \geqslant 3$, so $|C| \geqslant 4$. As $e \in G$ (by $(*)$ for $\tilde{H}$), the fact that e bridges v thus contradicts (3).

7.$^-$ Let G be a graph with fewer than i vertices of degree at most i, for every $i < |G|/2$. Use Chvátal's theorem to show that G is hamiltonian. (Thus in particular, Chvátal's theorem implies Dirac's theorem.)

8. Find a connected graph G whose square G^2 has no Hamilton cycle.

9.$^+$ Show by induction on $|G|$ that the third power G^3 of a connected graph G contains a Hamilton path between any two vertices. Deduce that G^3 is hamiltonian.

10. Show that the square of a 2-connected graph contains a Hamilton path between any two vertices.

11. An oriented complete graph is called a *tournament*. Show that every tournament contains a (directed) Hamilton path.

12.$^+$ Let G be a graph in which every vertex has odd degree. Show that every edge of G lies on an even number of Hamilton cycles.

(Hint. Let $xy \in E(G)$ be given. The Hamilton cycles through xy correspond to the Hamilton paths in $G - xy$ from x to y. Consider the set $\mathcal{H}$ of *all* Hamilton paths in $G - xy$ starting at x, and show that an even number of these end in y. To show this, define a graph on $\mathcal{H}$ so that the desired assertion follows from Proposition 1.2.1.)

Notes

The problem of finding a Hamilton cycle in a graph has the same kind of origin as its Euler tour counterpart and the four colour problem: all three problems come from mathematical puzzles older than graph theory itself. What began as a game invented by W.R. Hamilton in 1857—in which 'Hamilton cycles' had to be found on the graph of the dodecahedron—reemerged over a hundred years later as a combinatorial optimization problem of prime importance: the *travelling salesman problem*. Here, a salesman has to visit a number of customers, and his problem is to arrange these in a suitable circular route. (For reasons not included in the mathematical brief, the route has to be such that after visiting a customer the salesman does not pass through that town again.) Much of the motivation for considering Hamilton cycles comes from variations of this algorithmic problem.

A detailed discussion of the various degree conditions for hamiltonicity referred to at the beginning of Section 10.2 can be found in R. Halin, *Graphentheorie*, Wissenschaftliche Buchgesellschaft 1980. All the relevant references for Sections 10.1 and 10.2 can be found there, or in B. Bollobás, *Extremal Graph Theory*, Academic Press 1978.

The 'proof' of the four colour theorem indicated at the end of Section 10.1, which is based on the (false) premise that every 3-connected cubic planar graph is hamiltonian, is usually attributed to the Scottish mathematician P.G. Tait. Following Kempe's flawed proof of 1879 (see the notes for Chapter 5), it seems that Tait believed to be in possession of at least one 'new proof of Kempe's theorem'. However, when he addressed the Edinburgh Mathematical Society on

11 Random Graphs

At various points in this book, we already encountered the following fundamental theorem of Erdős: *for every integer k there is a graph G with $g(G) > k$ and $\chi(G) > k$.* In plain English: there exist graphs combining arbitrarily large girth with arbitrarily high chromatic number.

How could one prove such a theorem? The standard approach would be to construct a graph with those two properties, possibly in steps by induction on k. However, this is anything but straightforward: the global nature of the second property forced by the first, namely, that the graph should have high chromatic number 'overall' but be acyclic (and hence 2-colourable) locally, flies in the face of any attempt to build it up, constructively, from smaller pieces that have the same or similar properties.

In his pioneering paper of 1959, Erdős took a radically different approach: for each n he defined a probability space on the set of graphs with n vertices, and showed that, for some carefully chosen probability measures, the probability that an n-vertex graph has both of the above properties is positive for all large enough n.

This approach, now called the *probabilistic method*, has since unfolded into a sophisticated and versatile proof technique, in graph theory as much as in other branches of discrete mathematics. The theory of *random graphs* is now a subject in its own right. The aim of this chapter is to offer an elementary but rigorous introduction to random graphs: no more than is necessary to understand its basic concepts, ideas and techniques, but enough to give an inkling of the power and elegance hidden behind the calculations.

Erdős's theorem asserts the existence of a graph with certain properties: it is a perfectly ordinary assertion showing no trace of the randomness employed in its proof. There are also results in random graphs that are generically random even in their statement: these are theorems about *almost all* graphs, a notion we shall meet in Section 11.3. In the

Thus, formally, an element of Ω is a map ω assigning to every $e \in [V]^2$ either 0_e or 1_e, and the probability measure P on Ω is the product measure of all the measures P_e. In practice, of course, we identify ω with the graph G on V whose edge set is *P*

$$E(G) = \{ e \mid \omega(e) = 1_e \},$$

and call G a *random graph* on V with edge probability p. *random graph*

Following standard probabilistic terminology, we may now call any set of graphs on V an *event* in $\mathcal{G}(n,p)$. In particular, for every $e \in [V]^2$ the set *event*

$$A_e := \{ \omega \mid \omega(e) = 1_e \}$$ A_e

of all graphs G on V with $e \in E(G)$ is an event: the event that e is an edge of G. For these events, we can now prove formally what had been our guiding intuition all along:

Proposition 11.1.1. *The events A_e are independent and occur with probability p.*

Proof. By definition,

$$A_e = \{ 1_e \} \times \prod_{e' \neq e} \Omega_{e'}.$$

Since P is the product measure of all the measures P_e, this implies

$$P(A_e) = p \cdot \prod_{e' \neq e} 1 = p.$$

Similarly, if $\{ e_1, \dots, e_k \}$ is any subset of $[V]^2$, then

$$\begin{aligned} P(A_{e_1} \cap \ldots \cap A_{e_k}) &= P\Big(\{ 1_{e_1} \} \times \ldots \times \{ 1_{e_k} \} \times \prod_{e \notin \{ e_1, \dots, e_k \}} \Omega_e \Big) \\ &= p^k \\ &= P(A_{e_1}) \cdots P(A_{e_k}). \end{aligned}$$

□

As noted before, P is determined uniquely by the value of p and our assumption that the events A_e are independent. In order to calculate probabilities in $\mathcal{G}(n,p)$, it therefore generally suffices to work with these two assumptions: our concrete model for $\mathcal{G}(n,p)$ has served its purpose and will not be needed again.

As a simple example of such a calculation, consider the event that G contains some fixed graph H on a subset of V as a subgraph; let $|H| =: k$ and $\|H\| =: \ell$. The probability of this event $H \subseteq G$ is the product of the probabilities A_e over all the edges $e \in H$, so $P[H \subseteq G] = p^\ell$. In k ℓ

$$\begin{aligned} P[\alpha(G) \geqslant k],\ P[\omega(G) \geqslant k] &\leqslant \binom{n}{k} \left(\tfrac{1}{2}\right)^{\binom{k}{2}} \\ &< \left(n^k/2^k\right) 2^{-\frac{1}{2}k(k-1)} \\ &\leqslant \left(2^{k^2/2}/2^k\right) 2^{-\frac{1}{2}k(k-1)} \\ &= 2^{-k/2} \\ &< \tfrac{1}{2}\,. \end{aligned}$$

□

In the context of random graphs, each of the familiar graph invariants (like average degree, connectivity, girth, chromatic number, and so on) may be interpreted as a non-negative *random variable* on $\mathcal{G}(n,p)$, a function *random variable*

$$X : \mathcal{G}(n,p) \to [0,\infty)\,.$$

The *mean* or *expected* value of X is the number *mean* *expectation*

$$E(X) := \sum_{G \in \mathcal{G}(n,p)} P(\{G\}) \cdot X(G)\,.$$

$E(X)$

Note that the operator E, the *expectation*, is linear: we have $E(X+Y) = E(X)+E(Y)$ and $E(\lambda X) = \lambda E(X)$ for any two random variables X, Y on $\mathcal{G}(n,p)$ and $\lambda \in \mathbb{R}$.

Computing the mean of a random variable X can be a simple and effective way to establish the existence of a graph G such that $X(G) < a$ for some fixed $a > 0$ and, moreover, G has some desired property $\mathcal{P}$. Indeed, if the expected value of X is small, then $X(G)$ cannot be large for more than a few graphs in $\mathcal{G}(n,p)$, because $X(G) \geqslant 0$ for all $G \in \mathcal{G}(n,p)$. Hence X must be small for many graphs in $\mathcal{G}(n,p)$, and it is reasonable to expect that among these we may find one with the desired property $\mathcal{P}$.

This simple idea lies at the heart of countless non-constructive existence proofs using random graphs, including the proof of Erdős's theorem presented in the next section. Quantified, it takes the form of the following lemma, whose proof follows at once from the definition of the expectation and the additivity of P:

Lemma 11.1.4. (Markov's Inequality) [11.2.2]
Let $X \geqslant 0$ be a random variable on $\mathcal{G}(n,p)$ and $a > 0$. Then [11.4.1] [11.4.3]

$$P[X \geqslant a] \leqslant E(X)/a\,.$$

Proof.

$$E(X) = \sum_{G \in \mathcal{G}(n,p)} P(\{G\}) \cdot X(G)$$

How many such cycles $C = v_0 \dots v_{k-1} v_0$ are there? There are $(n)_k$ sequences $v_0 \dots v_{k-1}$ of distinct vertices in V, and each cycle is identified by $2k$ of those sequences—so there are exactly $(n)_k/2k$ such cycles.

Our random variable X assigns to every graph G its number of k-cycles. Clearly, this is the sum of all the values $X_C(G)$, where C varies over the $(n)_k/2k$ cycles of length k with vertices in V:

$$X = \sum_C X_C \,.$$

Since the expectation is linear, (1) thus implies

$$E(X) = E\Big(\sum_C X_C\Big) = \sum_C E(X_C) = \frac{(n)_k}{2k}\, p^k$$

as claimed. □

11.2 The probabilistic method

Very roughly, the *probabilistic method* in discrete mathematics has developed from the following idea. In order to prove the existence of an object with some desired property, one defines a probability space on some larger—and certainly non-empty—class of objects, and then shows that an element of this space has the desired property with positive probability. The 'objects' inhabiting this probability space may be of any kind: partitions or orderings of the vertices of some fixed graph arise as naturally as mappings, embeddings and, of course, graphs themselves. In this section, we illustrate the probabilistic method by giving a detailed account of one of its earliest results: of Erdős's classic theorem on large girth and chromatic number.

Erdős's theorem says that, given any positive integer k, there is a graph G with girth $g(G) > k$ and chromatic number $\chi(G) > k$. Let us call cycles of length at most k *short*, and sets of $|G|/k$ or more vertices *big*. For a proof of Erdős's theorem, it suffices to find a graph G without short cycles and without big independent sets of vertices: then the colour classes in any vertex colouring of G are *small* (not big), so we need more than k colours to colour G. *short* *big/small*

How can we find such a graph G? If we choose p small enough, then a random graph in $\mathcal{G}(n,p)$ is unlikely to contain any (short) cycles. If we choose p large enough, then G is unlikely to have big independent vertex sets. So the question is: do these two ranges of p overlap, that is, can we choose p so that, for some n, it is both small enough to give $P[\,g \leqslant k\,] < \frac{1}{2}$ and large enough for $P[\,\alpha \geqslant n/k\,] < \frac{1}{2}$? If so, then

Since $p \geqslant (6k \ln n) n^{-1}$ for n large, we thus obtain for $r := \lceil \frac{1}{2} n/k \rceil$

$$\lim_{n\to\infty} P\,[\,\alpha \geqslant \tfrac{1}{2} n/k\,] = \lim_{n\to\infty} P\,[\,\alpha \geqslant r\,] = 0\,,$$

as claimed. □

Theorem 11.2.2. (Erdős 1959)
For every integer k there exists a graph H with girth $g(H) > k$ and chromatic number $\chi(H) > k$. [9.2.3]

(11.1.4)
(11.1.5)

Proof. Assume that $k \geqslant 3$, fix ϵ with $0 < \epsilon < 1/k$, and let $p := n^{\epsilon - 1}$. Let $X(G)$ denote the number of short cycles in a random graph $G \in \mathcal{G}(n,p)$, i.e. its number of cycles of length at most k. p, ϵ, X

By Lemma 11.1.5, we have

$$E(X) = \sum_{i=3}^{k} \frac{(n)_i}{2i} p^i \leqslant \tfrac{1}{2} \sum_{i=3}^{k} n^i p^i \leqslant \tfrac{1}{2}(k-2)\, n^k p^k\,;$$

note that $(np)^i \leqslant (np)^k$, because $np = n^{\epsilon} \geqslant 1$. By Lemma 11.1.4,

$$\begin{aligned} P\,[\,X \geqslant n/2\,] &\leqslant E(X)/(n/2) \\ &\leqslant (k-2)\, n^{k-1} p^k \\ &= (k-2)\, n^{k-1} n^{(\epsilon-1)k} \\ &= (k-2)\, n^{k\epsilon - 1}. \end{aligned}$$

As $k\epsilon - 1 < 0$ by our choice of ϵ, this implies that

$$\lim_{n\to\infty} P\,[\,X \geqslant n/2\,] = 0\,.$$

Let n be large enough that $P\,[\,X \geqslant n/2\,] < \frac{1}{2}$ and $P\,[\,\alpha \geqslant \frac{1}{2} n/k\,] < \frac{1}{2}$; n
the latter is possible by our choice of p and Lemma 11.2.1. Then there is a graph $G \in \mathcal{G}(n,p)$ with fewer than $n/2$ short cycles and $\alpha(G) < \frac{1}{2} n/k$. From each of those cycles delete a vertex, and let H be the graph obtained. Then $|H| \geqslant n/2$ and H has no short cycles, so $g(H) > k$. By definition of G,

$$\chi(H) \geqslant \frac{|H|}{\alpha(H)} \geqslant \frac{n/2}{\alpha(G)} > k\,.$$

□

Corollary 11.2.3. *There are graphs with arbitrarily large girth and arbitrarily large values of the invariants κ, ε and δ.*

Proof. Apply Corollary 5.2.3 and Theorem 1.4.2. □ (1.4.2) (5.2.3)

Proof. For fixed U, W and $v \in G - (U \cup W)$, the probability that v is adjacent to all the vertices in U but to none in W, is

$$p^{|U|} q^{|W|} \geqslant p^i q^j .$$

Hence, the probability that no suitable v exists for these U and W, is

$$(1 - p^{|U|} q^{|W|})^{n - |U| - |W|} \leqslant (1 - p^i q^j)^{n-i-j}$$

(for $n \geqslant i + j$), since the corresponding events are independent for different v. As there are no more than n^{i+j} pairs of such sets U, W in $V(G)$ (encode sets U of fewer than i points as non-injective maps $\{0, \dots, i-1\} \to \{0, \dots, n-1\}$, etc.), the probability that some such pair has no suitable v is at most

$$n^{i+j} (1 - p^i q^j)^{n-i-j} ,$$

which tends to zero as $n \to \infty$ since $1 - p^i q^j < 1$. □

Corollary 11.3.3. *For every constant $p \in (0,1)$ and $k \in \mathbb{N}$, almost every graph in $\mathcal{G}(n,p)$ is k-connected.*

Proof. By Lemma 11.3.2, it is enough to show that every graph in $\mathcal{P}_{2,k-1}$ is k-connected. But this is easy: any graph in $\mathcal{P}_{2,k-1}$ has order at least $k+2$, and if W is a set of fewer than k vertices, then by definition of $\mathcal{P}_{2,k-1}$ any other two vertices x, y have a common neighbour $v \notin W$; in particular, W does not separate x from y. □

In the proof of Corollary 11.3.3, we showed substantially more than was asked for: rather than finding, for any two vertices $x, y \notin W$, some x–y path avoiding W, we showed that x and y have a common neighbour outside W; thus, all the paths needed to establish the desired connectivity could in fact be chosen of length 2. What seemed like a clever trick in this particular proof is in fact indicative of a more fundamental phenomenon for constant edge probabilities: by an easy result in logic, any statement about graphs expressed by quantifying over vertices only (rather than over sets or sequences of vertices)[4] is either almost surely true or almost surely false. All such statements, or their negations, are in fact immediate consequences of an assertion that the graph has property $\mathcal{P}_{i,j}$, for some suitable i, j.

As a last example of an 'almost all' result we now show that almost every graph has a surprisingly high chromatic number:

[4] In the terminology of logic: any first order sentence in the language of graph theory

Let us then see what happens if p is allowed to vary with n. Almost immediately, a fascinating picture unfolds. For edge probabilities p whose order of magnitude lies below n^{-2}, a random graph $G \in \mathcal{G}(n,p)$ almost surely has no edges at all. As p grows, G acquires more and more structure: from about $p = \sqrt{n}\,n^{-2}$ onwards, it almost surely has a component with more than two vertices, these components grow into trees, and around $p = n^{-1}$ the first cycles are born. Soon, some of these will have several crossing chords, making the graph non-planar. At the same time, one component outgrows the others, until it devours them around $p = (\log n)n^{-1}$, making the graph connected. Hardly later, at $p = (1+\epsilon)(\log n)n^{-1}$, our graph almost surely has a Hamilton cycle!

It has become customary to compare this development of random graphs as p grows to the evolution of an organism: for each $p = p(n)$, one thinks of the properties shared by almost all graphs in $\mathcal{G}(n,p)$ as properties of 'the' typical random graph $G \in \mathcal{G}(n,p)$, and studies how G changes its features with the growth rate of p. As with other species, the evolution of random graphs happens in relatively sudden jumps: the critical edge probabilities mentioned above are thresholds below which almost no graph and above which almost every graph has the property considered. More precisely, we call a real function $t = t(n)$ with $t(n) \neq 0$ for all n a *threshold function* for a graph property $\mathcal{P}$ if the following holds for all $p = p(n)$, and $G \in \mathcal{G}(n,p)$:

threshold function

$$\lim_{n\to\infty} P\,[\,G \in \mathcal{P}\,] = \begin{cases} 0 & \text{if } p/t \to 0 \text{ as } n \to \infty \\ 1 & \text{if } p/t \to \infty \text{ as } n \to \infty. \end{cases}$$

If $\mathcal{P}$ has a threshold function t, then clearly any positive multiple ct of t is also a threshold function for $\mathcal{P}$; thus, threshold functions in the above sense are only ever unique up to a multiplicative constant.[5]

Which graph properties have threshold functions? Natural candidates for such properties are *increasing* ones, properties closed under the addition of edges. (Graph properties of the form $\{\,G \mid G \supseteq H\,\}$, with H fixed, are common increasing properties; connectedness is another.) And indeed, Bollobás & Thomason (1987) have shown that all increasing properties, trivial exceptions aside, have threshold functions.

In the next section we shall study a general method to compute threshold functions.

---

[5] Our notion of threshold reflects only the crudest interesting level of screening: for some properties, such as connectedness, one can define sharper thresholds where the constant factor is crucial. Note also the role of the constant factor in our comparison of connectedness with hamiltonicity in the previous paragraph.

Note that μ and σ^2 always refer to a random variable on some fixed probability space. In our setting, where we consider the spaces $\mathcal{G}(n,p)$, both quantities are functions of n.

The following lemma says exactly what we need: that X cannot deviate a lot from its mean too often.

Lemma 11.4.1. (Chebyshev's Inequality)
For all real $\lambda > 0$,

$$P\big[\,|X-\mu| \geqslant \lambda\,\big] \;\leqslant\; \sigma^2/\lambda^2.$$

Proof. By Lemma 11.1.4 and definition of σ^2, (11.1.4)

$$P\,[\,|X-\mu| \geqslant \lambda\,] = P\,[\,(X-\mu)^2 \geqslant \lambda^2\,] \leqslant \sigma^2/\lambda^2.$$

□

For a proof that $X(G) > 0$ for almost all $G \in \mathcal{G}(n,p)$, Chebyshev's inequality can be used as follows:

Lemma 11.4.2. *If $\mu > 0$ for n large, and $\sigma^2/\mu^2 \to 0$ as $n \to \infty$, then $X(G) > 0$ for almost all $G \in \mathcal{G}(n,p)$.*

Proof. Any graph G with $X(G) = 0$ satisfies $|X(G) - \mu| = \mu$. Hence Lemma 11.4.1 implies with $\lambda := \mu$ that

$$P\,[\,X=0\,] \;\leqslant\; P\big[\,|X-\mu| \geqslant \mu\,\big] \;\leqslant\; \sigma^2/\mu^2 \underset{n\to\infty}{\longrightarrow} 0\,.$$

Since $X \geqslant 0$, this means that $X > 0$ almost surely, i.e. that $X(G) > 0$ for almost all $G \in \mathcal{G}(n,p)$. □

As the main result of this section, we now prove a theorem that will at once give us threshold functions for a number of natural properties. Given a graph H, we denote by $\mathcal{P}_H$ the graph property of containing a copy of H as a subgraph. We shall call H *balanced* if $\varepsilon(H') \leqslant \varepsilon(H)$ for all subgraphs H' of H. $\mathcal{P}_H$ *balanced*

Theorem 11.4.3. (Erdős & Rényi 1960)
If H is a balanced graph with k vertices and $\ell \geqslant 1$ edges, then $t(n) := n^{-k/\ell}$ is a threshold function for $\mathcal{P}_H$. k, ℓ t

We now come to the second part of the proof: we show that almost all $G \in \mathcal{G}(n,p)$ lie in $\mathcal{P}_H$ if $\gamma \to \infty$ as $n \to \infty$. Note first that, for $n \geqslant k$,

$$\begin{aligned}\binom{n}{k} n^{-k} &= \frac{1}{k!}\left(\frac{n}{n}\cdots\frac{n-k+1}{n}\right)\\ &\geqslant \frac{1}{k!}\left(\frac{n-k+1}{n}\right)^k\\ &\geqslant \frac{1}{k!}\left(1-\frac{k-1}{k}\right)^k; \qquad (6)\end{aligned}$$

thus, n^k exceeds $\binom{n}{k}$ by no more than a factor independent of n.

Our goal is to apply Lemma 11.4.2, and hence to bound $\sigma^2/\mu^2 = \big(E(X^2)-\mu^2\big)/\mu^2$ from above. As in (3) we have

$$E(X^2) = \sum_{(H',H'')\in\mathcal{H}^2} P\,[\,H'\cup H'' \subseteq G\,]\,. \qquad (7)$$

Let us then calculate these probabilities $P\,[\,H'\cup H'' \subseteq G\,]$. Given $H', H'' \in \mathcal{H}$, we have

$$P\,[\,H'\cup H''\subseteq G\,] \;=\; p^{2\ell-\|H'\cap H''\|}.$$

Since H is balanced, $\varepsilon(H'\cap H'') \leqslant \varepsilon(H) = \ell/k$. With $|H'\cap H''| =: i$ i
this yields $\|H'\cap H''\| \leqslant i\ell/k$, so by $0 \leqslant p \leqslant 1$,

$$P\,[\,H'\cup H''\subseteq G\,] \;\leqslant\; p^{2\ell - i\ell/k}. \qquad (8)$$

We have now estimated the individual summands in (7); what does this imply for the sum as a whole? Since (8) depends on the parameter $i = |H'\cap H''|$, we partition the range $\mathcal{H}^2$ of the sum in (7) into the subsets

$$\mathcal{H}_i^2 := \big\{\,(H',H'')\in\mathcal{H}^2 : |H'\cap H''| = i\,\big\}\,, \qquad i = 0,\dots,k,$$ $\mathcal{H}_i^2$

and calculate for each $\mathcal{H}_i^2$ the corresponding sum

$$A_i := \sum\nolimits_i P\,[\,H'\cup H''\subseteq G\,]$$ A_i

by itself. (Here, as below, we use $\sum_i$ to denote sums over all pairs $\sum_i$
$(H',H'')\in\mathcal{H}_i^2$.)

If $i = 0$ then H' and H'' are disjoint, so the events $H' \subseteq G$ and $H'' \subseteq G$ are independent. Hence,

$$A_0 \;=\; \sum\nolimits_0 P\,[\,H'\cup H''\subseteq G\,]$$

By Lemma 11.4.2, therefore, $X > 0$ almost surely, i.e. almost all $G \in \mathcal{G}(n,p)$ have a subgraph isomorphic to H and hence lie in $\mathcal{P}_H$. □

Theorem 11.4.3 allows us to read off threshold functions for a number of natural graph properties.

Corollary 11.4.4. *If $k \geqslant 3$, then $t(n) = n^{-1}$ is a threshold function for the property of containing a k-cycle.* □

Interestingly, the threshold function in Corollary 11.4.4 is independent of the cycle length k considered: in the evolution of random graphs, cycles of all (constant) lengths appear at about the same time!

There is a similar phenomenon for trees. Here, the threshold function does depend on the order of the tree considered, but not on its shape:

Corollary 11.4.5. *If T is a tree of order $k \geqslant 2$, then $t(n) = n^{-k/(k-1)}$ is a threshold function for the property of containing a copy of T.*

We finally have the following result for complete subgraphs:

Corollary 11.4.6. *If $k \geqslant 2$, then $t(n) = n^{-2/(k-1)}$ is a threshold function for the property of containing a K^k.*

Proof. K^k is balanced, because $\varepsilon(K^i) = \frac{1}{2}(i-1) < \frac{1}{2}(k-1) = \varepsilon(K^k)$ for $i < k$. With $\ell := \|K^k\| = \frac{1}{2}k(k-1)$, we obtain $n^{-k/\ell} = n^{-2/(k-1)}$. □

It is not difficult to adapt the proof of Theorem 11.4.3 to the case that H is unbalanced. The threshold then becomes $t(n) = n^{-1/\varepsilon'(H)}$, where $\varepsilon'(H) := \max\{\, \varepsilon(F) \mid F \subseteq H \,\}$; see Exercise 22.

Exercises

1.$^-$ What is the probability that a random graph in $\mathcal{G}(n,p)$ has exactly m edges, for $0 \leqslant m \leqslant \binom{n}{2}$ fixed?

2. What is the expected number of edges in $G \in \mathcal{G}(n,p)$?

3. What is the expected number of K^r-subgraphs in $G \in \mathcal{G}(n,p)$?

4. Characterize the graphs that occur as a subgraph in every graph of sufficiently large average degree.

5. In the usual terminology of measure spaces (and in particular, of probability spaces), the phrase 'almost all' is used to refer to a set of points whose complement has measure zero. Rather than considering a limit of probabilities in $\mathcal{G}(n,p)$ as $n \to \infty$, would it not be more natural to define a probability space on the set of *all* finite graphs (one copy of each) and to investigate properties of 'almost all' graphs in this space, in the sense above?

21.$^+$ Given a graph H, let $\mathcal{P}$ be the property of containing an induced copy of H. If H is complete then, by Corollary 11.4.6, $\mathcal{P}$ has a threshold function. Show that $\mathcal{P}$ has no threshold function if H is not complete.

22.$^+$ Prove the following version of Theorem 11.4.3 for unbalanced subgraphs. Let H be any graph with at least one edge, and put $\varepsilon'(H) := \max\{\varepsilon(F) \mid \emptyset \neq F \subseteq H\}$. Then the threshold function for $\mathcal{P}_H$ is $t(n) = n^{-1/\varepsilon'(H)}$.

(Hint. Imitate the proof of Theorem 11.4.3. Instead of the sets $\mathcal{H}_i$, consider the sets $\mathcal{H}_F^2 := \{(H', H'') \in \mathcal{H}^2 \mid H' \cap H'' = F\}$. Replace the distinction between the cases of $i = 0$ and $i > 0$ by the distinction between the cases of $\|F\| = 0$ and $\|F\| > 0$.)

Notes

There are a number of monographs and texts on the subject of random graphs. The most comprehensive of these is B. Bollobás, *Random Graphs*, Academic Press 1985. Another advanced but very readable monograph is S. Janson, T. Łuczak & A. Ruciński, *Topics in Random Graphs*, in preparation; this concentrates on areas developed since *Random Graphs* was published. E.M. Palmer, *Graphical Evolution*, Wiley 1985, covers material similar to parts of *Random Graphs* but is written in a more elementary way. Compact introductions going beyond what is covered in this chapter are given by B. Bollobás, *Graph Theory*, Springer GTM 63, 1979, and by M. Karoński, *Handbook of Combinatorics* (R.L. Graham, M. Grötschel & L. Lovász, eds.), North-Holland 1995.

A stimulating advanced introduction to the use of random techniques in discrete mathematics more generally is given by N. Alon & J.H. Spencer, *The Probabilistic Method*, Wiley 1992. One of the attractions of this book lies in the way it shows probabilistic methods to be relevant in proofs of entirely deterministic theorems, where nobody would suspect it. Another example for this phenomenon is Alon's proof of Theorem 5.4.1; see the notes for Chapter 5.

The probabilistic method had its first origins in the 1940s, one of its earliest results being Erdős's probabilistic lower bound for Ramsey numbers (Theorem 11.1.3). Lemma 11.3.2 about the properties $\mathcal{P}_{i,j}$ is taken from Bollobás's Springer text cited above. A very readable rendering of the proof that, for constant p, every first order sentence about graphs is either almost surely true or almost surely false, is given by P. Winkler, Random structures and zero-one laws, in (N.W. Sauer et al., eds.) *Finite and Infinite Combinatorics in Sets and Logic* (NATO ASI Series C **411**), Kluwer 1993.

The seminal paper on graph evolution is P. Erdős & A. Rényi, On the evolution of random graphs, *Publ. Math. Inst. Hungar. Acad. Sci.* **5** (1960), 17–61. This paper also includes Theorem 11.4.3 and its proof. The generalization of this theorem to unbalanced subgraphs was first proved by Bollobás in 1981, using advanced methods; a simple adaptation of the original Erdős-Renyi proof was found by Ruciński & Vince (1986), and is presented in Karoński's Handbook chapter.

12 Minors, Trees, and WQO

Our goal in this last chapter is a single theorem, one which dwarfs any other result in graph theory and may doubtless be counted among the deepest theorems that mathematics has to offer: *in every infinite set of graphs there are two such that one is a minor of the other.* This *graph minor theorem* (or *minor theorem* for short), inconspicuous though it may look at first glance, has made a fundamental impact both outside graph theory and within. Its proof, due to Neil Robertson and Paul Seymour, takes well over 500 pages.

So we have to be modest: of the actual proof of the minor theorem, this chapter will convey only a very rough impression. However, as with most truly fundamental results, the proof has sparked off the development of methods of quite independent interest and potential. This is true particularly for the use of *tree-decompositions*, a technique we shall meet in Sections 12.3 and 12.4. Section 12.1 gives an introduction to *well-quasi-ordering*, a concept central to the minor theorem. In Section 12.2 we apply this concept to prove the minor theorem for trees. The chapter finishes with an overview in Section 12.5 of the proof of the general graph minor theorem, and of some of its immediate consequences.

12.1 Well-quasi-ordering

well-quasi-ordering

A reflexive and transitive relation is called a *quasi-ordering*. A quasi-ordering $\leqslant$ on X is a *well-quasi-ordering*, and the elements of X are *well-quasi-ordered* by $\leqslant$, if for every infinite sequence $x_0, x_1, \ldots$ in X

By Corollary 12.1.2, the sequence $(a_n)_{n\in\mathbb{N}}$ has an infinite increasing subsequence $(a_{n_i})_{i\in\mathbb{N}}$. By the minimal choice of A_{n_0}, the sequence

$$A_0, \ldots, A_{n_0-1}, B_{n_0}, B_{n_1}, B_{n_2}, \ldots$$

is good; consider a good pair. Since $(A_n)_{n\in\mathbb{N}}$ is bad, this pair cannot have the form (A_i, A_j) or (A_i, B_j), as $B_j \leqslant A_j$. So it has the form (B_i, B_j). Extending the injection $B_i \to B_j$ by $a_i \mapsto a_j$, we deduce again that (A_i, A_j) is good, a contradiction. □

12.2 The graph minor theorem for trees

The minor theorem can be expressed by saying that the finite graphs are well-quasi-ordered by the minor relation $\preccurlyeq$. Indeed, by Proposition 12.1.1 and the obvious fact that no strictly descending sequence of minors can be infinite, being well-quasi-ordered is equivalent to the non-existence of an infinite antichain, the formulation used earlier.

In this section, we prove a strong version of the graph minor theorem for trees:

Theorem 12.2.1. (Kruskal 1960)
The finite trees are well-quasi-ordered by the topological minor relation.

We shall base the proof of Theorem 12.2.1 on the following notion of an embedding between rooted trees, which strengthens the usual embedding as a topological minor. Consider two trees T and T', with roots r and r' say. Let us write $T \leqslant T'$ if there exists an isomorphism φ, from some subdivision of T to a subtree T'' of T', that preserves the tree-order on $V(T)$ associated with T and r. (Thus if $x < y$ in T then $\varphi(x) < \varphi(y)$ in T'; see Fig. 12.2.1.) As one easily checks, this is a quasi-ordering on the class of all rooted trees.

$T \leqslant T'$

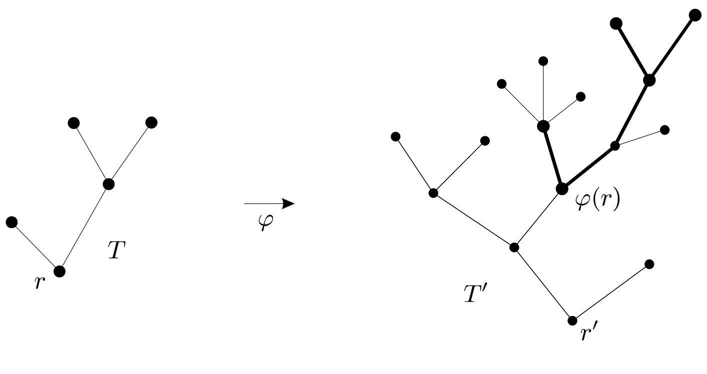

Fig. 12.2.1. An embedding of T in T' showing that $T \leqslant T'$

12.3 Tree-decompositions

Trees are graphs with some very distinctive and fundamental properties; consider Theorem 1.5.1 and Corollary 1.5.2, or the more sophisticated example of Kruskal's theorem. It is therefore legitimate to ask to what degree those properties can be transferred to more general graphs, graphs that are not themselves trees but tree-like in some sense.[2] In this section, we study a concept of tree-likeness that permits generalizations of all the tree properties referred to above (including Kruskal's theorem), and which plays a crucial role in the proof of the graph minor theorem.

Let G be a graph, T a tree, and let $\mathcal{V} = (V_t)_{t \in T}$ be a family of vertex sets $V_t \subseteq V(G)$ indexed by the vertices t of T. The pair $(T, \mathcal{V})$ is called a *tree-decomposition* of G if it satisfies the following three conditions:

tree-decomposition

(T1) $V(G) = \bigcup_{t \in T} V_t$;

(T2) for every edge $e \in G$ there exists a $t \in T$ such that both ends of e lie in V_t;

(T3) $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever $t_1, t_2, t_3 \in T$ satisfy $t_2 \in t_1 T t_3$.

Conditions (T1) and (T2) together say that G is the union of the subgraphs $G[V_t]$; we call these subgraphs and the sets V_t themselves the *parts* of $(T, \mathcal{V})$ and say that $(T, \mathcal{V})$ is a tree-decomposition of G *into* these parts. Condition (T3) implies that the parts of $(T, \mathcal{V})$ are organized roughly like a tree (Fig. 12.3.1).

parts *into*

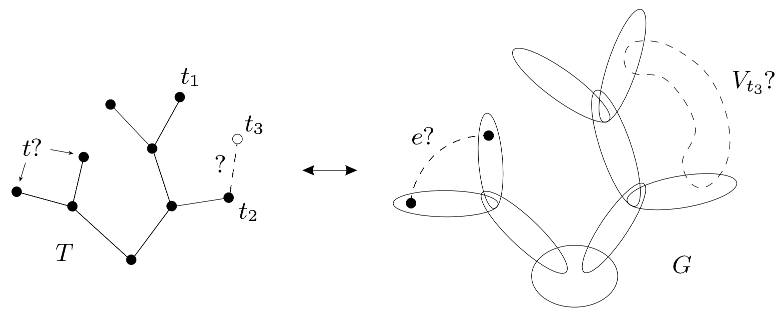

Fig. 12.3.1. Edges and parts ruled out by (T2) and (T3)

Before we discuss the role that tree-decompositions play in the proof of the minor theorem, let us note some of their basic properties. Consider a fixed tree-decomposition $(T, \mathcal{V})$ of G, with $\mathcal{V} = (V_t)_{t \in T}$ as above. $(T, \mathcal{V})$, V_t

Perhaps the most important feature of a tree-decomposition is that it transfers the separation properties of its tree to the graph decomposed:

[2] What exactly this 'sense' should be will depend both on the property considered and on its intended application.

Lemma 12.3.4. *Given a set $W \subseteq V(G)$, there is either a $t \in T$ such that $W \subseteq V_t$, or there are vertices $w_1, w_2 \in W$ and an edge $t_1t_2 \in T$ such that w_1, w_2 lie outside the set $V_{t_1} \cap V_{t_2}$ and are separated by it in G.*

Proof. Let us orient the edges of T as follows. For each edge $t_1t_2 \in T$, define U_1, U_2 as in Lemma 12.3.1; then $V_{t_1} \cap V_{t_2}$ separates U_1 from U_2. If $V_{t_1} \cap V_{t_2}$ does not separate any two vertices of W that lie outside it, we can find an $i \in \{1, 2\}$ such that $W \subseteq U_i$, and orient t_1t_2 towards t_i.

Let t be the last vertex of a maximal directed path in T; we claim that $W \subseteq V_t$. Given $w \in W$, let $t' \in T$ be such that $w \in V_{t'}$. If $t' \neq t$, then the edge e at t that separates t' from t in T is directed towards t, so w also lies in $V_{t''}$ for some t'' in the component of $T - e$ containing t. Therefore $w \in V_t$ by (T3). □

The following special case of Lemma 12.3.4 is used particularly often:

Lemma 12.3.5. *Any complete subgraph of G is contained in some part of $(T, \mathcal{V})$.* □ [12.4.2]

As indicated by Figure 12.3.1, the parts of $(T, \mathcal{V})$ reflect the structure of the tree T, so in this sense the graph G decomposed resembles a tree. However, this is valuable only inasmuch as the structure of G within each part is negligible: the smaller the parts, the closer the resemblance.

This observation motivates the following definition. The *width* of $(T, \mathcal{V})$ is the number *width*

$$\max \left\{ |V_t| - 1 : t \in T \right\},$$

and the *tree-width* $\text{tw}(G)$ of G is the least width of any tree-decomposition of G. As one easily checks,[3] trees themselves have tree-width 1. *tree-width* $\text{tw}(G)$

By Lemmas 12.3.2 and 12.3.3, the tree-width of a graph will never be increased by deletion or contraction:

Proposition 12.3.6. *If $H \preccurlyeq G$ then* $\text{tw}(H) \leqslant \text{tw}(G)$. □

Graphs of bounded tree-width are sufficiently similar to trees that it becomes possible to adapt the proof of Kruskal's theorem to the class of these graphs; very roughly, one has to iterate the 'minimal bad sequence' argument from the proof of Lemma 12.1.3 $\text{tw}(G)$ times. This takes us a step further towards a proof of the graph minor theorem:

Theorem 12.3.7. (Robertson & Seymour 1990)
For every integer $k > 0$, the graphs of tree-width $< k$ are well-quasi-ordered by the minor relation.

---

[3] Indeed the '-1' in the definition of width serves no other purpose than to make this statement true.

for each B disjoint from X there is an $i \in \{1, 2\}$ such that $B \subseteq U_i \setminus X$ (defined as in Lemma 12.3.1); recall that B is connected. Moreover, this i is the same for all such B, because they touch. We now orient the edge t_1t_2 towards t_i.

If every edge of T is oriented in this way and t is the last vertex of a maximal directed path in T, then V_t meets every set in $\mathcal{B}$—just as in the proof of Lemma 12.3.4.

To prove the forward direction, we now assume that G contains no bramble of order $> k$. We show that for every bramble $\mathcal{B}$ in G there is a $\mathcal{B}$-*admissible* tree-decomposition of G, one in which any part of order $> k$ fails to cover $\mathcal{B}$. For $\mathcal{B} = \emptyset$ this implies that $\text{tw}(G) < k$, because every set covers the empty bramble. *$\mathcal{B}$-admissible*

Let $\mathcal{B}$ be given, and assume inductively that for every bramble $\mathcal{B}'$ with more sets than $\mathcal{B}$ there is a $\mathcal{B}'$-admissible tree-decomposition of G. (The induction starts, since no bramble in G has more than $2^{|G|}$ sets.) Let $X \subseteq V(G)$ be a cover of $\mathcal{B}$ with as few vertices as possible; then $\ell := |X| \leqslant k$ is the order of $\mathcal{B}$. Our aim is to show the following: $\mathcal{B}$ X ℓ

For every component C of $G - X$ there exists a $\mathcal{B}$-admissible tree-decomposition of $G[X \cup V(C)]$ with X as a part. $(*)$

Then these tree-decompositions can be combined to a $\mathcal{B}$-admissible tree-decomposition of G by identifying their nodes corresponding to X. (If $X = V(G)$, then the tree-decomposition with X as its only part is $\mathcal{B}$-admissible.)

So let C be a fixed component of $G - X$, write $H := G[X \cup V(C)]$, and put $\mathcal{B}' := \mathcal{B} \cup \{C\}$. If $\mathcal{B}'$ is not a bramble then C fails to touch some element of $\mathcal{B}$, and hence $Y := V(C) \cup N(C)$ does not cover $\mathcal{B}$. Then the tree-decomposition of H consisting of the two parts X and Y satisfies $(*)$. C, H $\mathcal{B}'$

So we may assume that $\mathcal{B}'$ is a bramble. Since X covers $\mathcal{B}$ but not $\mathcal{B}'$, we have $|\mathcal{B}'| > |\mathcal{B}|$. Our induction hypothesis therefore ensures that G has a $\mathcal{B}'$-admissible tree-decomposition $(T, (V_t)_{t \in T})$. If this decomposition is also $\mathcal{B}$-admissible, there is nothing more to show. If not, then one of its parts of order $> k$, V_s say, covers $\mathcal{B}$. Since no set of fewer than ℓ vertices covers $\mathcal{B}$, Lemma 12.3.8 implies with Menger's theorem (3.3.1) that V_s and X are linked by ℓ disjoint paths $P_1, \dots, P_\ell$. As V_s fails to cover $\mathcal{B}'$ and hence lies in $G - C$, the paths P_i meet H only in their ends $x_i \in X$. $T, (V_t)_{t \in T}$ s P_i x_i

For each $i = 1, \dots, \ell$ pick a $t_i \in T$ with $x_i \in V_{t_i}$, and let t_i

$$W_t := (V_t \cap (X \cup V(C))) \cup \{x_i \mid t \in sTt_i\};$$

for all $t \in T$ (Fig. 12.3.3). Then $(T, (W_t)_{t \in T})$ is the tree-decomposition which $(T, (V_t)_{t \in T})$ induces on H (cf. Lemma 12.3.2), except that a few

Given any two vertices $t_1, t_2 \in T$, Lemma 12.3.1 implies that every V_t with $t \in t_1Tt_2$ separates V_{t_1} from V_{t_2} in G. Let us call our tree-decomposition $(T, \mathcal{V})$ of G *linked*, or *lean*,[4] if it satisfies the following condition: *linked/lean*

(T4) given any $s \in \mathbb{N}$ and $t_1, t_2 \in T$, either G contains s disjoint V_{t_1}–V_{t_2} paths or there exists a $t \in t_1Tt_2$ such that $|V_t| < s$.

The 'branches' in a lean tree-decomposition are thus stripped of any bulk not necessary to maintain their connecting qualities: if a branch is thick (the parts along a path in T large), then G is highly connected along this branch.

In our quest for tree-decompositions into 'small' parts, we now have two criteria to choose between: the global 'worst case' criterion of width, which ensures that T is nontrivial (unless G is complete) but says nothing about the tree-likeness of G among parts other than the largest, and the more subtle local criterion of leanness, which ensures tree-likeness everywhere along T but might be difficult to achieve except with trivial or near-trivial T. Surprisingly, though, it is always possible to find a tree-decomposition that is optimal with respect to both criteria at once:

Theorem 12.3.10. (Thomas 1990)
Every graph G has a lean tree-decomposition of width $\mathrm{tw}(G)$.

The proof of Theorem 12.3.10 is not too long but technical, and we shall not present it. The fact that this theorem gives us a very useful property of minimum-width tree-decompositions 'for free' has made it a valuable tool wherever tree-decompositions are applied.

The tree-decomposition $(T, \mathcal{V})$ of G is called *simplicial* if all the separators $V_{t_1} \cap V_{t_2}$ induce complete subgraphs in G. This assumption can enable us to lift assertions about the parts of the decomposition to G itself. For example, if all the parts in a simplicial tree-decomposition of G are k-colourable, then so is G (proof?). The same applies to the property of not containing a K^r minor for some fixed r. Algorithmically, it is easy to obtain a simplicial tree-decomposition of a given graph into irreducible parts. Indeed, all we have to do is split the graph recursively along complete separators; if these are always chosen minimal, then the set of parts obtained will even be unique (Exercise 22). *simplicial*

Conversely, if G can be constructed recursively from a set $\mathcal{H}$ of graphs by pasting along complete subgraphs, then G has a simplicial tree-decomposition into elements of $\mathcal{H}$. For example, by Wagner's Theorem 8.3.4, any graph without a K^5 minor has a supergraph with a simplicial tree-decomposition into plane triangulations and copies of the

[4] depending on which of the two dual aspects of (T4) we wish to emphasize

12.4 Tree-width and forbidden minors

If $\mathcal{H}$ is any set or class of graphs, then the class

$$\text{Forb}_{\preccurlyeq}(\mathcal{H}) := \{\, G \mid G \not\succcurlyeq H \text{ for all } H \in \mathcal{H} \,\}$$ $\text{Forb}_{\preccurlyeq}(\mathcal{H})$

of all graphs without a minor in $\mathcal{H}$ is a graph property, i.e. is closed under isomorphism.[5] When it is written as above, we say that this property is expressed by specifying the graphs $H \in \mathcal{H}$ as *forbidden* (or *excluded*) *minors*. *forbidden minors*

By Proposition 1.7.3, $\text{Forb}_{\preccurlyeq}(\mathcal{H})$ is closed under taking minors: if $G' \preccurlyeq G \in \text{Forb}_{\preccurlyeq}(\mathcal{H})$ then $G' \in \text{Forb}_{\preccurlyeq}(\mathcal{H})$. Graph properties that are closed under taking minors will be called *hereditary* in this chapter. Every hereditary property can in turn be expressed by forbidden minors: (1.7.3) *hereditary*

Proposition 12.4.1. *A graph property $\mathcal{P}$ can be expressed by forbidden minors if and only if it is hereditary.* [12.5.1]

Proof. For the 'if' part, note that $\mathcal{P} = \text{Forb}_{\preccurlyeq}(\overline{\mathcal{P}})$, where $\overline{\mathcal{P}}$ is the complement of $\mathcal{P}$. □ $\overline{\mathcal{P}}$

In Section 12.5, we shall return to the general question of how a given hereditary property is best represented by forbidden minors. In this section, we are interested in one particular type of hereditary property: bounded tree-width.

Thus, let us consider the property of having tree-width less than some given integer k. By Propositions 12.3.6 and 12.4.1, this property can be expressed by forbidden minors. Choosing their set $\mathcal{H}$ as small as possible, we find that $\mathcal{H} = \{\, K^3 \,\}$ for $k = 2$: the graphs of tree-width < 2 are precisely the forests. For $k = 3$, we have $\mathcal{H} = \{\, K^4 \,\}$:

Proposition 12.4.2. *A graph has tree-width < 3 if and only if it has no K^4 minor.*

(8.3.1) (12.3.2) (12.3.5) (12.3.11)

Proof. By Lemma 12.3.5, we have $\text{tw}(K^4) \geqslant 3$. By Proposition 12.3.6, therefore, a graph of tree-width < 3 cannot contain K^4 as a minor.

Conversely, let G be a graph without a K^4 minor; we assume that $|G| \geqslant 3$. Add edges to G until the graph G' obtained is edge-maximal without a K^4 minor. By Proposition 8.3.1, G' can be constructed recursively from triangles by pasting along K^2s. By induction on the number of recursion steps and Lemma 12.3.5, every graph constructible in this way has a tree-decomposition into triangles (as in the proof of Proposition 12.3.11). Such a tree-decomposition of G' has width 2, and by Lemma 12.3.2 it is also a tree-decomposition of G. □

[5] As usual, we abbreviate $\text{Forb}_{\preccurlyeq}(\{\, H \,\})$ to $\text{Forb}_{\preccurlyeq}(H)$.

Y–Z paths in G that have no inner vertex or edge in $G[X]$. Note that the vertex set of a k-connected subgraph of G need not be externally k-connected in G. On the other hand, any horizontal path in the $r \times r$ grid is externally k-connected in that grid for every $k \leqslant r$. (How?)

One of the first things we shall prove below is that any graph of large enough tree-width—not just grids—contains a large externally k-connected set of vertices (Lemma 12.4.5). Conversely, it is easy to show that large externally k-connected sets (with k large) can exist only in graphs of large tree-width (Exercise 30). So, like large grid minors, these sets form a canonical obstruction to small tree-width: they can be found in a graph if and only if its tree-width is large.

An ordered pair (A, B) of subgraphs of G will be called a *premesh* in G if $G = A \cup B$ and A contains a tree T such that *premesh*

(i) T has maximum degree ≤ 3;

(ii) every vertex of $A \cap B$ lies in T and has degree ≤ 2 in T;

(iii) T has a leaf in $A \cap B$, or $|T| = 1$ and $T \subseteq A \cap B$.

The *order* of such a premesh is the number $|A \cap B|$, and if $V(A \cap B)$ is externally k-connected in B then this premesh is a k-*mesh* in G. *order* *k-mesh*

Lemma 12.4.5. *Let G be a graph and let $h \geq k \geq 1$ be integers. If G contains no k-mesh of order h then G has tree-width $< h + k - 1$.*

Proof. We may assume that G is connected. Let $U \subseteq V(G)$ be maximal such that $G[U]$ has a tree-decomposition $\mathcal{D}$ of width $< h + k - 1$ with the additional property that, for every component C of $G - U$, the neighbours of C in U lie in one part of $\mathcal{D}$ and $(G - C, \tilde{C})$ is a premesh of order $\leq h$, where $\tilde{C} := G[V(C) \cup N(C)]$. Clearly, $U \neq \emptyset$. U $\mathcal{D}$ $\tilde{C}$

We claim that $U = V(G)$. Suppose not. Let C be a component of $G - U$, put $X := N(C)$, and let T be a tree associated with the premesh $(G - C, \tilde{C})$. C X T

By assumption, $|X| \leq h$; let us show that equality holds here. If not, let $u \in X$ be a leaf of T (respectively $\{u\} := V(T)$) as in (iii), and let $v \in C$ be a neighbour of u. Put $U' := U \cup \{v\}$ and $X' := X \cup \{v\}$, let T' be the tree obtained from T by joining v to u, and let $\mathcal{D}'$ be the tree-decomposition of $G[U']$ obtained from $\mathcal{D}$ by adding X' as a new part (joined to a part of $\mathcal{D}$ containing X, which exists by our choice of U; see Fig. 12.4.1). Clearly $\mathcal{D}'$ still has width $< h + k - 1$. Consider a component C' of $G - U'$. If $C' \cap C = \emptyset$ then C' is also a component of $G - U$, so $N(C')$ lies inside a part of $\mathcal{D}$ (and hence of $\mathcal{D}'$), and $(G - C', \tilde{C}')$ is a premesh of order $\leq h$ by assumption. If $C' \cap C \neq \emptyset$, then $C' \subseteq C$ and $N(C') \subseteq X'$. Moreover, $v \in N(C')$: otherwise $N(C') \subseteq X$ would separate C' from v, contradicting the fact that C' and v lie in the same component C of $G - X$. Since v is a leaf of T', it is straightforward to and $Z \setminus S$; we assume it has none in $Y \setminus S$. Let T' be the union of T and all the Y–S subpaths of paths P_s with $s \in N(C') \cap C$; since these subpaths start in $Y \setminus S$ and have no inner vertices in X', they cannot meet C'. Therefore $(G - C', \tilde{C}')$ is a premesh with tree T' and leaf v; the degree conditions on T' are easily checked. Its order is $|N(C')| \leq |X| - |Y| + |S| = h - |Y| + k' < h$, a contradiction to the maximality of U. □

Lemma 12.4.6. *Let $k \geq 2$ be an integer. Let T be a tree of maximum degree $\leqslant 3$ and $X \subseteq V(T)$. Then T has a set F of edges such that every component of $T - F$ has between k and $2k - 1$ vertices in X, except that one such component may have fewer vertices in X.*

Proof. We apply induction on $|X|$. If $|X| \leq 2k - 1$ we put $F = \emptyset$. So assume that $|X| \geq 2k$. Let e be an edge of T such that some component T' of $T - e$ has at least k vertices in X and $|T'|$ is as small as possible. As $\Delta(T) \leq 3$, the end of e in T' has degree at most two in T', so the minimality of T' implies that $|X \cap V(T')| \leq 2k - 1$. Applying the induction hypothesis to $T - T'$ we complete the proof. □

Lemma 12.4.7. *Let G be a bipartite graph with bipartition (A, B), $|A| = a$, $|B| = b$, and let $c \leq a$ and $d \leq b$ be positive integers. Assume that G has at most $(a - c)(b - d)/d$ edges. Then there exist $C \subseteq A$ and $D \subseteq B$ such that $|C| = c$ and $|D| = d$ and $C \cup D$ is independent in G.*

Proof. As $||G|| \leq (a - c)(b - d)/d$, fewer than $b - d$ vertices in B have more than $(a - c)/d$ neighbours in A. Choose $D \subseteq B$ so that $|D| = d$ and each vertex in D has at most $(a - c)/d$ neighbours in A. Then D sends a total of at most $a - c$ edges to A, so A has a subset C of c vertices without a neighbour in D. □

Given a tree T, call an r-tuple $(x_1, \ldots, x_r)$ of distinct vertices of T *good* if, for every $j = 1, \ldots, r - 1$, the x_j–x_{j+1} path in T contains none of the other vertices in this r-tuple. *good r-tuple*

Lemma 12.4.8. *Every tree T of order at least $r(r - 1)$ contains a good r-tuple of vertices.*

Proof. Pick a vertex $x \in T$. Then T is the union of its subpaths xTy, where y ranges over its leaves. Hence unless one of these paths has at least r vertices, T has at least $|T|/(r - 1) \geqslant r$ leaves. Since any path of r vertices and any set of r leaves gives rise to a good r-tuple in T, this proves the assertion. □

Our next lemma shows how to obtain a grid from two large systems of paths that intersect in a particularly orderly way.

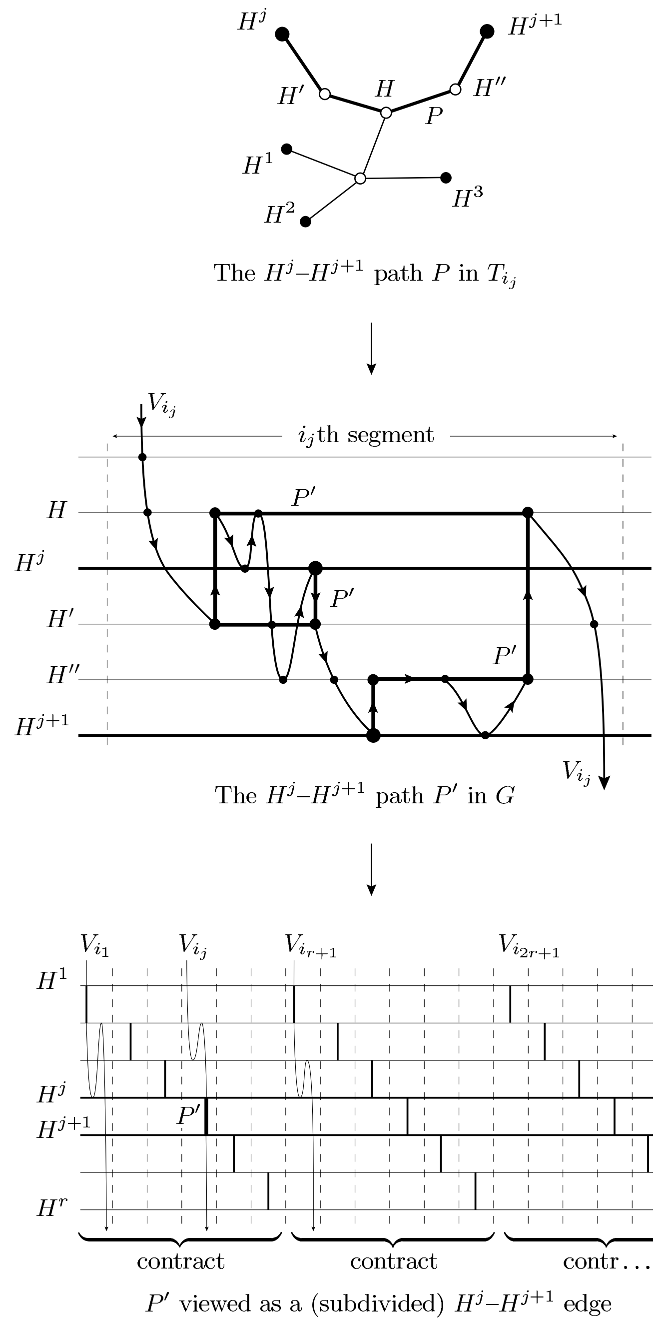

Fig. 12.4.4. An H^j–H^{j+1} path in T_{i_j} inducing segments of V_{i_j} for the jth edge of the grid's first vertical path

X $A_1, \ldots, A_m$ $\mathcal{P}_{ij}$

premesh (A, B); then $X := V(A \cap B) \subseteq V(T)$. By Lemma 12.4.6, T has $|X|/(2k-1) - 1 = m$ disjoint subtrees each containing at least k vertices of X; let $A_1, \ldots, A_m$ be the vertex sets of these trees. By definition of a k-mesh, B contains for all $1 \leq i < j \leq m$ a set $\mathcal{P}_{ij}$ of k disjoint A_i–A_j paths that have no inner vertices in A. These sets $\mathcal{P}_{ij}$ will shrink a little and be otherwise modified later in the proof, but they will always consist of 'many' disjoint A_i–A_j paths.

One option in our proof will be to find single paths $P_{ij} \in \mathcal{P}_{ij}$ that are disjoint for different pairs ij and thus link up the sets A_i to form a K^m minor of G. If this fails, we shall instead exhibit two specific sets $\mathcal{P}_{ij}$ and $\mathcal{P}_{pq}$ such that many paths of $\mathcal{P}_{ij}$ meet many paths of $\mathcal{P}_{pq}$, forming an $r \times r$ grid between them by Lemma 12.4.9.

σ

Let us impose a linear ordering on the index pairs ij by fixing an arbitrary bijection $\sigma : \{\, ij \mid 1 \leq i < j \leq m \,\} \to \{\, 0, 1, \ldots, \binom{m}{2} - 1 \,\}$. For $\ell = 0, 1, \ldots$ in turn, we shall consider the pair pq with $\sigma(pq) = \ell$ and choose an A_p–A_q path P_{pq} that is disjoint from all previously selected such paths, i.e. from the paths P_{st} with $\sigma(st) < \ell$. At the same time, we shall replace all the 'later' sets $\mathcal{P}_{ij}$—or what has become of them—by smaller sets containing only paths that are disjoint from P_{pq}. Thus for each pair ij, we shall define a sequence $\mathcal{P}_{ij} = \mathcal{P}^0_{ij}, \mathcal{P}^1_{ij}, \ldots$ of smaller and smaller sets of paths, which eventually collapses to $\mathcal{P}^\ell_{ij} = \{\, P_{ij} \,\}$ when ℓ has risen to $\ell = \sigma(ij)$.

ℓ^*

More formally, let $\ell^* \leq \binom{m}{2}$ be the greatest integer such that, for all $0 \leq \ell < \ell^*$ and all $1 \leqslant i < j \leqslant m$, there exist sets $\mathcal{P}^\ell_{ij}$ satisfying the following five conditions:

(i) $\mathcal{P}^\ell_{ij}$ is a non-empty set of disjoint A_i–A_j paths in B that meet A only in their endpoints.

H^ℓ_{ij}

Whenever a set $\mathcal{P}^\ell_{ij}$ is defined, we shall write $H^\ell_{ij} := \bigcup \mathcal{P}^\ell_{ij}$ for the union of its paths.

P_{ij}

(ii) If $\sigma(ij) < \ell$ then $\mathcal{P}^\ell_{ij}$ has exactly one element P_{ij}, and P_{ij} does not meet any path belonging to a set $\mathcal{P}^\ell_{st}$ with $ij \neq st$.

(iii) If $\sigma(ij) = \ell$, then $|\mathcal{P}^\ell_{ij}| = k/c^{2\ell}$.

(iv) If $\sigma(ij) > \ell$, then $|\mathcal{P}^\ell_{ij}| = k/c^{2\ell+1}$.

(v) If $\ell = \sigma(pq) < \sigma(ij)$, then for every $e \in E(H^\ell_{ij}) \setminus E(H^\ell_{pq})$ there are no $k/c^{2\ell+1}$ disjoint A_i–A_j paths in the graph $(H^\ell_{pq} \cup H^\ell_{ij}) - e$.

Note that, by (iv), the paths considered in (v) do exist in H^ℓ_{ij}. The purpose of (v) is to force those paths to reuse edges from H^ℓ_{pq} whenever possible, using new edges $e \notin H^\ell_{pq}$ only if necessary. Note further that since $\sigma(ij) < \binom{m}{2}$ by definition of σ, conditions (iii) and (iv) give $|\mathcal{P}^\ell_{ij}| \geq c^2$ whenever $\sigma(ij) \geq \ell$.

Clearly if $\ell^* = \binom{m}{2}$ then by (i) and (ii) we have a (subdivided) K^m minor with branch sets $A_1, \ldots, A_m$ in G. Suppose then that $\ell^* < \binom{m}{2}$.

appropriate segments, we shall first pick a path $Q \in \mathcal{H}$ to serve as a yardstick: we shall divide Q into segments each meeting lots of paths from $\mathcal{V}$, select a 'non-crossing' subset $V_1, \ldots, V_d$ of these vertical paths, one from each segment (which is the most delicate task; we shall need condition (v) from the definition of the sets $\mathcal{P}_{ij}^{\ell}$ here), and finally divide the other horizontal paths into the 'induced' segments, accommodating one V_n each.

So let us pick a path $Q \in \mathcal{H}$, and put Q

$$d := \lfloor \sqrt{c}/m \rfloor = \lfloor r^{2r+4}/m \rfloor \geq r^{2r+2}.$$ d

Note that $|\mathcal{V}| \geqslant (c/m^2)|\mathcal{P}_{ij}^{\ell}| \geqslant d^2|\mathcal{P}_{ij}^{\ell}|$.

For $n = 1, 2, \ldots, d-1$ let e_n be the first edge of Q (on its way from A_i to A_j) such that the initial component Q_n of $Q - e_n$ meets at least $nd\,|\mathcal{P}_{ij}^{\ell}|$ different paths from $\mathcal{V}$, and such that e_n is not an edge of H_{pq}^{ℓ}. As any two vertices of Q that lie on different paths from $\mathcal{V}$ are separated in Q by an edge not in H_{pq}^{ℓ}, each of these Q_n meets exactly $nd\,|\mathcal{P}_{ij}^{\ell}|$ paths from $\mathcal{V}$. Put $Q_0 := \emptyset$ and $Q_d := Q$. Since $|\mathcal{V}| \geq d^2|\mathcal{P}_{ij}^{\ell}|$, we have thus divided Q into d consecutive disjoint segments $Q'_n := Q_n - Q_{n-1}$ $(n = 1, \ldots, d)$ each meeting at least $d\,|\mathcal{P}_{ij}^{\ell}|$ paths from $\mathcal{V}$. e_n Q_n $Q'_1, \ldots, Q'_d$

For each $n = 1, \ldots, d-1$, Menger's theorem (3.3.1) and conditions (iv) and (v) imply that $H_{pq}^{\ell} \cup H_{ij}^{\ell}$ has a set S_n of $|\mathcal{P}_{ij}^{\ell}| - 1$ vertices such that $(H_{pq}^{\ell} \cup H_{ij}^{\ell}) - e_n - S_n$ contains no path from A_i to A_j. Let S denote the union of all these sets S_n. Then $|S| < d\,|\mathcal{P}_{ij}^{\ell}|$, so each Q'_n meets at least one path $V_n \in \mathcal{V}$ that avoids S (Fig. 12.4.5). S_n S V_n

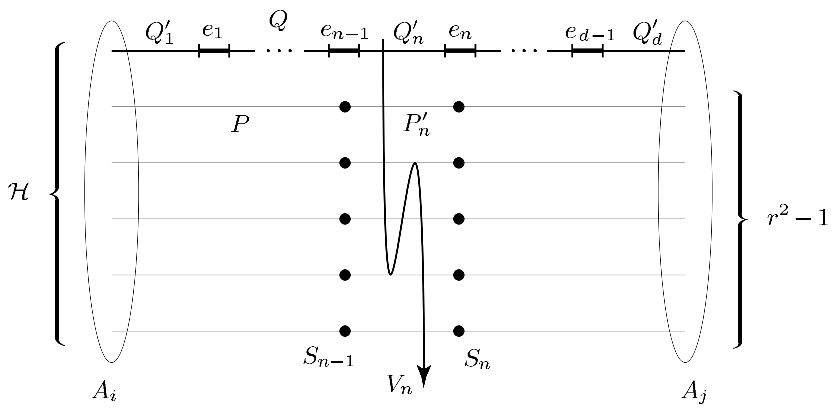

Fig. 12.4.5. V_n meets every horizontal path but avoids S

Clearly, each S_n consists of a choice of exactly one vertex x from every path $P \in \mathcal{P}_{ij}^{\ell} \setminus \{ Q \}$. Denote the initial component of $P - x$ by P_n, put $P_0 := \emptyset$ and $P_d := P$, and let $P'_n := P_n - P_{n-1}$ for $n = 1, \ldots, d$. The separation properties of the sets S_n now imply that $V_n \cap P \subseteq P'_n$ for $P'_1, \ldots, P'_d$

Theorem 12.5.2. (Graph Minor Theorem; Robertson & Seymour) *The finite graphs are well-quasi-ordered by the minor relation $\preccurlyeq$.*

So every $\mathcal{H}_\mathcal{P}$ is finite, i.e. every hereditary graph property can be represented by finitely many forbidden minors:

Corollary 12.5.3. *Every graph property that is closed under taking minors can be expressed as* $\mathrm{Forb}_\preccurlyeq(\mathcal{H})$ *with finite* $\mathcal{H}$. □

As a special case of Corollary 12.5.3 we have, at least in principle, a Kuratowski-type theorem for every surface:

Corollary 12.5.4. *For every surface S there exists a finite set of graphs $H_1, \ldots, H_n$ such that* $\mathrm{Forb}_\preccurlyeq(H_1, \ldots, H_n)$ *contains precisely the graphs not embeddable in S.*

The minimal set of forbidden minors has been determined explicitly for only one surface other than the sphere: for the projective plane it is known to consist of 35 forbidden minors. It is not difficult to show that the number of forbidden minors grows rapidly with the genus of the surface (Exercise 34).

The complete proof of the graph minor theorem would fill a book or two. For all its complexity in detail, however, its basic idea is easy to grasp. We have to show that every infinite sequence

$$G_0, G_1, G_2, \ldots$$

of finite graphs contains a good pair: two graphs $G_i \preccurlyeq G_j$ with $i < j$. We may assume that $G_0 \not\preccurlyeq G_i$ for all $i \geqslant 1$, since G_0 forms a good pair with any graph G_i of which it is a minor. Thus all the graphs $G_1, G_2, \ldots$ lie in $\mathrm{Forb}_\preccurlyeq(G_0)$, and we may use the structure common to these graphs in our search for a good pair.

We have already seen how this works when G_0 is planar: then the graphs in $\mathrm{Forb}_\preccurlyeq(G_0)$ have bounded tree-width (Theorem 12.4.3) and are therefore well-quasi-ordered by Theorem 12.3.7. In general, we need only consider the cases of $G_0 = K^n$: since $G_0 \preccurlyeq K^n$ for $n := |G_0|$, we may assume that $K^n \not\preccurlyeq G_i$ for all $i \geqslant 1$.

The proof now follows the same lines as above: again the graphs in $\mathrm{Forb}_\preccurlyeq(K^n)$ can be characterized by their tree-decompositions, and again their tree structure helps, as in Kruskal's theorem, with the proof that they are well-quasi-ordered. The parts in these tree-decompositions are no longer restricted in terms of order now, but they are constrained in more subtle structural terms. Roughly speaking, for every n there exists a finite set $\mathcal{S}$ of closed surfaces such that every graph without a K^n minor has a simplicial tree-decomposition into parts each 'nearly' embedding in there is a polynomial-time algorithm[7] that decides whether or not the input graph contains H as a minor. By the minor theorem, then, every hereditary graph property $\mathcal{P}$ can be decided in polynomial (even cubic) time: if $H_1, \ldots, H_k$ are the corresponding minimal forbidden minors, then testing a graph G for membership in $\mathcal{P}$ reduces to testing the k assertions $H_i \preccurlyeq G$!

The following example gives an indication of how deeply this algorithmic corollary affects the complexity theory of graph algorithms. Let us call a graph *knotless* if it can be embedded in $\mathbb{R}^3$ so that none of its cycles forms a non-trivial knot. Before the graph minor theorem, it was an open problem whether knotlessness is decidable, that is, whether *any* algorithm exists (no matter how slow) that decides for any given graph whether or not that graph is knotless. To this day, no such algorithm is known. The property of knotlessness, however, is easily 'seen' to be hereditary: contracting an edge of a graph embedded in 3-space will not create a knot where none had been before. Hence, by the minor theorem, there *exists* an algorithm that decides knotlessness—even in polynomial (cubic) time!

However spectacular such unexpected solutions to long-standing problems may be, viewing the graph minor theorem merely in terms of its corollaries will not do it justice. At least as important are the techniques developed for its proof, the various ways in which minors are handled or constructed. Most of these have not even been touched upon here, yet they seem set to influence the development of graph theory for many years to come.

Exercises

1.$^-$ Let $\leqslant$ be a quasi-ordering on a set X. Call two elements $x, y \in X$ *equivalent* if both $x \leqslant y$ and $y \leqslant x$. Show that this is indeed an equivalence relation on X, and that $\leqslant$ induces a partial ordering on the set of equivalence classes.

2. Let $(A, \leqslant)$ be a quasi-ordering. For subsets $X \subseteq A$ write

$$\mathrm{Forb}_{\leqslant}(X) := \{\, a \in A \mid a \not\geqslant x \text{ for all } x \in X \,\}.$$

Show that $\leqslant$ is a well-quasi-ordering on A if and only if every subset $B \subseteq A$ that is closed under $\geqslant$ (i.e. such that $x \leqslant y \in B \Rightarrow x \in B$) can be written as $B = \mathrm{Forb}_{\leqslant}(X)$ with finite X.

3. Prove Proposition 12.1.1 and Corollary 12.1.2 directly, without using Ramsey's theorem.

[7] indeed a cubic one—although with a typically enormous constant depending on H

17. Let $\mathcal{B}$ be a maximum-order bramble in a graph G. Show that every minimum-width tree-decomposition of G has a unique part covering $\mathcal{B}$.

18.[+] In the second half of the proof of Theorem 12.3.9, let H' be the union of H and the paths $P_1, \dots, P_\ell$, let H'' be the graph obtained from H' by contracting each P_i, and let $(T, (W_t'')_{t \in T})$ be the tree-decomposition induced on H'' (as in Lemma 12.3.3) by the decomposition that $(T, (V_t)_{t \in T})$ induces on H'. Is this, after the obvious identification of H'' with H, the same decomposition as the one used in the proof, i.e. is $W_t'' = W_t$ for all $t \in T$?

19. Show that any graph with a simplicial tree-decomposition into k-colourable parts is itself k-colourable.

20. Let $\mathcal{H}$ be a set of graphs, and let G be constructed recursively from elements of $\mathcal{H}$ by pasting along complete subgraphs. Show that G has a simplicial tree-decomposition into elements of $\mathcal{H}$.

21. Given a tree-decomposition $(T, (V_t)_{t \in T})$ of G and $t \in T$, let H_t denote the graph obtained from $G[V_t]$ by adding all the edges xy such that $x, y \in V_t \cap V_{t'}$ for some neighbour t' of t in T; the graphs H_t are called the *torsos* of this tree-decomposition. Show that G has no K^5 minor if and only if G has a tree-decomposition in which every torso is either planar or a copy of the Wagner graph W (Fig. 8.3.1).

22.[+] Call a graph *irreducible* if it is not separated by any complete subgraph. Every (finite) graph G can be decomposed into irreducible induced subgraphs, as follows. If G has a separating complete subgraph S, then decompose G into proper induced subgraphs G' and G'' with $G = G' \cup G''$ and $G' \cap G'' = S$. Then decompose G' and G'' in the same way, and so on, until all the graphs obtained are irreducible. By Exercise 20, G has a simplicial tree-decomposition into these irreducible subgraphs. Show that they are uniquely determined if the complete separators were all chosen minimal.

23.[+] If $\mathcal{F}$ is a family of sets, then the graph G on $\mathcal{F}$ with $XY \in E(G) \Leftrightarrow X \cap Y \neq \emptyset$ is called the *intersection graph* of $\mathcal{F}$. Show that a graph is chordal if and only if it is isomorphic to the intersection graph of a family of (vertex sets of) subtrees of a tree.

24. A tree-decomposition of a graph is called a *path-decomposition* if its decomposition tree is a path. Show that a graph has a path-decomposition into complete graphs if and only if it is isomorphic to an interval graph. (Interval graphs are defined in Ex. 37, Ch. 5.)

25. (continued)

The *path-width* $\mathrm{pw}(G)$ of a graph G is the least width of a path-decomposition of G. Prove the following analogue of Corollary 12.3.12 for path-width: every graph G satisfies $\mathrm{pw}(G) = \min \omega(H) - 1$, where the minimum is taken over all interval graphs H containing G.

26.[+] Do trees have unbounded path-width?

(1988), 55–57. Whether or not the minor theorem extends to countable graphs remains an open problem.

The notions of tree-decomposition and tree-width were first introduced (under different names) by R. Halin, S-functions for graphs, *J. Geometry* **8** (1976), 171–186. Among other things, Halin showed that grids can have arbitrarily large tree-width. Robertson & Seymour reintroduced the two concepts, apparently unaware of Halin's paper, with direct reference to K. Wagner, Über eine Eigenschaft der ebenen Komplexe, *Math. Ann.* **114** (1937), 570–590. (This is the classic paper that introduced simplicial tree-decompositions to prove Theorem 8.3.4; cf. Exercise 21.) Simplicial tree-decompositions are treated in depth in R. Diestel, *Graph Decompositions*, Oxford University Press 1990.

Robertson & Seymour themselves usually refer to the graph minor theorem as *Wagner's conjecture*. It seems that Wagner did indeed discuss this problem in the 1960s with his then students Halin and Mader. However, Wagner apparently never conjectured a positive solution; he certainly rejected any credit for the 'conjecture' when it had been proved.

Robertson & Seymour's proof of the graph minor theorem is given in the numbers IV–VII, IX–XII and XIV–XX of their series of over 20 papers under the common title of *Graph Minors*, which has been appearing in the *Journal of Combinatorial Theory, Series B*, since 1983. Of their theorems cited in this chapter, Theorem 12.3.7 is from Graph Minors IV, while Theorems 12.4.3 and 12.4.4 are from Graph Minors V. Our short proof of these latter theorems is from R. Diestel, K.Yu. Gorbunov, T.R. Jensen & C. Thomassen, Highly connected sets and the excluded grid theorem, *J. Combin. Theory B* **75** (1999), 61–73.

Theorem 12.3.9 is due to P.D. Seymour & R. Thomas, Graph searching and a min-max theorem for tree-width, *J. Combin. Theory B* **58** (1993), 22–33. Our proof is a simplification of the original proof. B.A. Reed gives an instructive introductory survey on tree-width and graph minors, including some algorithmic aspects, in (R.A. Bailey, ed) *Surveys in Combinatorics 1997*, Cambridge University Press 1997, 87–162. Reed also introduced the term 'bramble'; in Seymour & Thomas's paper, brambles are called 'screens'.

The obstructions to small tree-width actually used in the proof of the graph minor theorem are not brambles but so-called *tangles*. Tangles are more powerful than brambles and well worth studying. See Graph Minors X or Reed's survey for an introduction to tangles and their relation to brambles and tree-decompositions.

Theorem 12.3.10 is due to R. Thomas, A Menger-like property of tree-width; the finite case, *J. Combin. Theory B* **48** (1990), 67–76.

As a forerunner to Theorem 12.4.3, Robertson & Seymour proved its following analogue for path-width (Graph Minors I): excluding a graph H as a minor bounds the path-width of a graph if and only if H is a forest. A short proof of this result, with optimal bounds, can be found in the first edition of this book, or in R. Diestel, Graph Minors I: a short proof of the path width theorem, *Combinatorics, Probability and Computing* **4** (1995), 27–30.

The 35 minimal forbidden minors for graphs to be embedded in the projective plane were determined by D. Archdeacon, A Kuratowski theorem for the projective plane, *J. Graph Theory* **5** (1981), 243–246. An upper bound for

Hints for all the Exercises

Caveat. These hints are intended to set on the right track anyone who has already spent some time over an exercise but somehow failed to make much progress. They are not designed to be particularly intelligible without such an initial attempt, and they will rarely spoil the fun by giving away the key idea. They may, however, narrow ones mind by focusing on what is just one of several possible ways to think about a problem...

Hints for Chapter 1

1.$^-$ How many edges are there at each vertex?

2. Average degree and edges: consider the vertex degrees. Diameter: how do you determine the distance between two vertices from the corresponding 0–1 sequences? Girth: does the graph have a cycle of length 3? Circumference: partition the d-dimensional cube into cubes of lower dimension and use induction.

3. Consider how the path intersects C. Where can you see cycles, and can they all be short?

4.$^-$ Can you find graphs for which Proposition 1.3.2 holds with equality?

5. Estimate the distances within G as seen from a central vertex.

6.$^+$ Consider the cases $d = 2$ and $d > 2$ separately. For $d > 2$, give a sharper bound on $|D_i|$ for $i > 0$ than the one used in the proof of Proposition 1.3.3.

7.$^-$ Assume the contrary, and derive a contradiction.

8.$^-$ Find two vertices that are linked by two independent paths.

Hints for Chapter 2

1. Compare the given matching with a matching of maximum cardinality.
2. Augmenting paths.
3. If you have $S \subsetneqq S' \subseteq A$ with $|S| = |N(S)|$ in the finite case, the marriage condition ensures that $N(S) \subsetneqq N(S')$: increasing S makes more neighbours available. Use the fact that this fails when S is infinite.
4. Apply the marriage theorem.
5. Construct a bipartite graph in which A is one side, and the other side consists of a suitable number of copies of the sets A_i. Define the edge set of the graph so that the desired result can be derived from the marriage theorem.
6.$^+$ Construct chains in the power set lattice of X as follows. For each $k < n/2$, use the marriage theorem to find a 1–1 map φ from the set A of k-subsets to the set B of $(k+1)$-subsets of X such that $Y \subseteq \varphi(Y)$ for all $Y \in A$.
7. Decide where the leaves should go: in factor-critical components or in S?
8. Distinguish between the cases of $|S| \leq 1$ and $|S| \geq 2$.
9. The case $S = \emptyset$ is easy. In the other case, look for a vertex that meets every maximum-cardinality matching. What are the consequences of this for the other vertices?
10. For the 'if' direction consider the graph suggested in the hint: does it have a 1-factor? If not, then consider the set of vertices supplied by Tutte's 1-factor theorem. For an alternative solution, apply the remarks on maximum-cardinality matchings which follow Theorem 2.2.3.
11.$^-$ Corollary 2.2.2.
12. Let G be a bipartite graph that satisfies the marriage condition, with bipartition (A, B) say. Reduce the problem to the case of $|A| = |B|$. To verify the premise of Tutte's theorem for a given set $S \subseteq V(G)$, bound $|S|$ from below in terms of the number of components of $G - S$ that contain more vertices from A than from B and vice versa.
13.$^-$ Consider any smallest path cover.
14. Direct all the edges from A to B.
15.$^-$ Dilworth.
16. Start with the set of minimal elements in P.
17. Think of the elements of A as being smaller than their neighbours in B.
18.$^+$ Let $(x, y) \leqslant (x', y')$ if and only if $x \leqslant x'$ and $y \leqslant y'$.

Hints for Chapter 4

1. Embed the vertices inductively. Where should you *not* put the new vertex?
2.$^-$ Figure 1.6.2.
3.$^-$ Make the given graph connected.
4. This is a generalization of Corollary 4.2.8.
5. Theorem 3.5.4.
6. Imitate the proof of Corollary 4.2.8.
7. Proposition 4.2.10.
8.$^-$ Express the difference between the two drawings as a formal statement about vertices, faces, and the incidences between them.
9. Combinatorially: use the definition. Topologically: express the relative position of the short (respectively, the long) branches of G' formally as a property of G' which any topological ismorphism would preserve but G lacks.
10.$^-$ Reflexivity, symmetry, transitivity.
11. Look for a graph whose drawings all look the same, but which admits an automorphism that does not extend to a homeomorphism of the plane. Interpret this automorphism as $\sigma_2 \circ \sigma_1^{-1}$.
12.$^+$ Star-shape: every inner face contains a point that sees the entire face boundary.
13. Work with plane rather than planar graphs.
14. (i) The set $\mathcal{X}$ may be infinite.
    (ii) If Y is a TX, then every TY is also a TX.
15.$^-$ By the next exercise, any counterexample can be disconnected by at most two vertices.
16. Incorporate the extra condition into the induction hypothesis of the proof. It may help to disallow polygons with 180 degree angles.
17. Number of edges.
18. Use that maximal planar graphs are 3-connected, and that the neighbours of each vertex induce a cycle.
19. If $G = G_1 \cup G_2$ with $G_1 \cap G_2 = \overline{K^2}$, we have a problem. This will go away if we embed a little more than necessary.
20. Use a suitable modification of the given graph G to simulate outerplanarity.
21. Use the fact that $\mathcal{C}(G)$ is the direct sum of $\mathcal{C}(G_1)$ and $\mathcal{C}(G_2)$.
22.$^+$ Euler.
23. The face boundaries generate $\mathcal{C}(G)$.
24.$^-$ Which are the faces that e^* (viewed as a polygonal arc) can meet?
25.$^-$ How many vertices does it have?
26.$^-$ Join two given vertices of the dual by a straight line, and use this to find a path between them in the dual graph.

15.$^+$ (i) How will v_1 and v_2 be coloured, and how v_n?

(ii) Consider the subgraph induced by the neighbours of v_n.

16. For the induction start, explicitly calculate $P_G(k)$ for $|G| = n$ and $\|G\| = 0$.

17.$^+$ Derive from the polynomial the number of edges and the number of components of G; see the previous exercise.

18. Imitate the proof of Theorem 5.2.5.

19.$^-$ $K_{n,n}$.

20. How are edge colourings related to matchings?

21. Construct a bipartite $\Delta(G)$-regular graph that contains G as subgraph. It may be necessary to add some vertices.

22.$^+$ Induction on k. In the induction step $k \to k+1$, consider using several copies of the graph you found for k.

23.$^-$ Vertex degrees.

24. $K_{n,n}$. To choose n so that $K_{n,n}$ is not even k-choosable, consider lists of k-subsets of a k^2-set.

25.$^-$ Vizing's theorem.

26. All you need are the definitions, Proposition 5.2.2, and a standard argument from Chapter 1.2.

27.$^+$ Try induction on r. In the induction step, you would like to to delete one pair of vertices and only one colour from the other vertices' lists. What can you say about the lists if this is impossible? This information alone will enable you to find a colouring directly, without even looking at the graph again.

28. Show that $\chi''(G) \leqslant \text{ch}'(G) + 2$, and use this to deduce $\chi''(G) \leqslant \Delta(G) + 3$ from the list colouring conjecture.

29.$^-$ Do bipartite graphs have a kernel?

30.$^+$ Call a set S of vertices in a directed graph D a *core* if D contains a directed v–S path for every vertex $v \in D - S$. If, in addition, D contains no directed path between any two vertices of S, call S a *strong core*. Show first that every core contains a strong core. Next, define inductively a partition of $V(D)$ into 'levels' $L_0, \dots, L_n$ such that, for even i, L_i is a suitable strong core in $D_i := D - (L_0 \cup \dots \cup L_{i-1})$, while for odd i, L_i consists of the vertices of D_i that send an edge to L_{i-1}. Show that, if D has no directed odd cycle, the even levels together form a kernel of D.

31. Construct the orientation needed for Lemma 5.4.3 in steps: if, in the current orientation, there are still vertices v with $d^+(v) \geqslant 3$, reverse the directions of an edge at v and take care of the knock-on effect of this change. If you need to bound the average degree of a bipartite planar graph, remember Euler's formula.

32.$^-$ Start with a non-perfect graph.

33.$^-$ Do odd cycles or their complements satisfy $(*)$?

34. Exercise 12, Chapter 3.

4.$^-$ H-flows are nowhere zero, by definition.

5.$^-$ Use the definition and Proposition 6.1.1.

6.$^-$ Do subgraphs also count as minors?

7.$^-$ Try $k = 2, 3, \dots$ in turn. In searching for a k-flow, tentatively fix the flow value through an edge and investigate which consequences this has for the adjacent edges.

8. To establish uniqueness, consider cuts of a special type.

9. Express G as the union of cycles.

10. Combine $\mathbb{Z}_2$-flows on suitable subgraphs to a flow on G.

11.$^+$ Begin by sending a small amount of flow through every edge outside T.

12. View G as the union of suitably chosen cycles.

13. Corollary 6.3.2 and Proposition 6.4.1.

14.$^-$ Duality.

15. Take as H your favourite graph of large flow number. Can you decrease its flow number by adding edges?

16. Euler.

17.$^+$ Theorem 6.5.3.

18.$^-$ Search among small cubic graphs.

19. Theorem 6.5.3.

20. (i) Theorem 6.5.3.

(ii) Yes it can. Show, by considering a smallest counterexample, that if every 3-connected cubic planar multigraph is 3-edge-colourable (and hence has a 4-flow), then so is every bridgeless cubic planar multigraph.

21.$^+$ For the 'only if' implication apply Proposition 6.1.1. Conversely, consider a circulation f on G, with values in $\{0, \pm 1, \dots, \pm(k-1)\}$, that respects the given orientation (i.e. is positive or zero on the edge directions assigned by D) and is zero on as few edges as possible. Then show that f is nowhere zero, as follows. If f is zero on $e = st \in E$ and D directs e from t to s, define a network $N = (G, s, t, c)$ such that any flow in N of positive total value contradicts the choice of f, but any cut in N of zero capacity contradicts the property assumed for D.

22.$^-$ Convert the given multigraph into a graph with the same flow properties.

Hints for Chapter 7

1.$^-$ Straightforward from the definition.

2.$^-$ When constructing the graphs, start by fixing the colour classes.

3. It is not difficult to determine an upper bound for $\text{ex}(n, K_{1,r})$. What remains to be proved is that this bound can be achieved for all r and n.

4. Proposition 1.7.2 (ii).

5. Proposition 1.2.2 and Corollary 1.5.4.

Hints for Chapter 8

1. For the induction step, partition the vertex set of the given graph G into two sets V_1 and V_2 so that colourings of $G[V_1]$ and $G[V_2]$ can be combined to a colouring of G.

2. Imitate the start of the proof of Lemma 8.1.3.

3.$^-$ Does a large chromatic number force up the average degree? If in doubt, consult Chapter 5.

4.$^+$ Try parallel paths in the grid as branch sets.

5.$^+$ How can we best make a TK^{2r} fit into a $K_{s,s}$ when we want to keep s small?

6. Split the argument into the cases of $k = 0$ and $k \geqslant 1$.

7. How are the two lemmas used in the proof of the theorem?

8. Study the motivational chat preceding the definition of f in the proof.

9.$^+$ Consider your favourite graphs with high average degree and low chromatic number. Which trees do they contain induced? Is there some reason to expect that exactly these trees may always be found induced in graphs of large average degree and small chromatic number?

10.$^-$ What does planarity have to do with minors?

11.$^-$ Consider a suitable supergraph.

12.$^-$ Average degree.

13.$^+$ Show by induction on $|G|$ that any 3-colouring of an induced cycle in $G \not\succcurlyeq K^4$ extends to all of G.

14.$^+$ Reduce the statement to critical k-chromatic graphs and apply Vizing's theorem.

15. (i) is easy. In the first part of (ii), distinguish between the cases that the graph is or is not separated by a $K^{\chi(G)-1}$. Show the second part by induction on the chromatic number. In the induction step split the vertex set of the graph into two subsets.

16. Induction on the number of construction steps.

17. Induction on $|G|$.

18. Note the previous exercise.

19. Which of the graphs constructed as in Theorem 8.3.4 have the largest average degree?

20. Which of the graphs constructed as in the hint have the largest average degree?

21. Consider the subgraph of G induced by the neighbours of x.

4. Figure 10.1.1.
5. Induction on k with n fixed; for the induction step consider $\overline{G}$.
6.$^-$ Recall the definition of a hamiltonian sequence.
7.$^-$ On which kind of vertices does the Chvátal condition come to bear? To check the validity of the condition for G, first find such a vertex.
8. How does an arbitrary connected graph differ from the kind of graph whose square contains a Hamilton cycle by Fleischner's theorem? How could this difference obstruct the existence of a Hamilton cycle?
9.$^+$ In the induction step consider a minimal cut.
10. Condition (∗) in the proof of Fleischner's theorem.
11. Induction.
12.$^+$ How can a Hamilton path $P \in \mathcal{H}$ be modified into another? In how many ways? What has this got to do with the degree in G of the last vertex of P?

Hints for Chapter 11

1.$^-$ Consider a fixed choice of m edges on $\{0, 1, \ldots, n\}$. What is the probability that $G \in \mathcal{G}(n, p)$ has precisely this edge set?
2. Consider the appropriate indicator random variables, as in the proof of Lemma 11.1.5.
3. Consider the appropriate indicator random variables.
4. Erdős.
5. What would be the measure of the set $\{G\}$ for a fixed G?
6. Consider the complementary properties.
7.$^-$ $\mathcal{P}_{2,1}$.
8. Apply Lemma 11.3.2.
9. Induction on $|H|$ with the aid of Exercise 6.
10.$^+$ (i) Given a pair U, U', calculate the probability that every other vertex is joined incorrectly to $U \cup U'$. What, then, is the probability that this happens for *some* pair U, U'?

    (ii) Enumerate the vertices of G and G' jointly, and construct an isomorphism $G \to G'$ inductively.
11. Imitate the proof of Lemma 11.2.1.
12. Imitate the proof of Proposition 11.3.1. To bound the probabilities involved, use the inequality $1 - x \leqslant e^{-x}$ as in the proof of Lemma 11.2.1.
13.$^+$ (i) Calculate the expected number of isolated vertices, and apply Lemma 11.4.2 as in the proof of Theorem 11.4.3.

    (ii) Linearity.
14.$^+$ Chapter 8.2, the proof of Erdős's theorem, and a bit of Chebyshev.

11. For the forward implication, apply Corollary 1.5.2. For the converse, use induction on k.
12. To prove (T2), consider the edge e of Figure 12.3.1. Checking (T3) is easy.
13. For the first question, recall Proposition 12.3.6. For the second, try to modify a tree-decomposition of G into one of the TG without increasing its width.
14. Lemma 12.3.1 relates the separation properties of a graph G to those of its decomposition tree T. This exercise illuminates this relationship from the dual viewpoint of connectedness: how are the connected subgraphs of G related to those of T?
15.$^-$ This is just a reformulation of Theorem 12.3.9.
16. Modify the proof given in the text that the $k \times k$ grid has tree-width at least $k-1$.
17. Existence was shown in Theorem 12.3.9; the task is to show uniqueness.
18.$^+$ Work out an explicit description of the sets W'_t similar to the definition of the W_t, and compare the two.
19. Induction.
20. Induction.
21. Use the previous exercise and a result from Chapter 8.3. And don't despair at a subgraph of W!
22.$^+$ Show that the parts are precisely the maximal irreducible induced subgraphs of G.
23.$^+$ Exercise 14.
24. For the forward implication, interpret the subpaths of the decomposition path as intervals. Which subpath corresponds naturally to a given vertex of G?
25. Follow the proof of Corollary 12.3.12.
26.$^+$ They do. To prove it, show first that every connected graph G contains a path whose deletion decreases the path-width of G. Then apply induction on a suitable set of trees, deleting a suitable path in the induction step.
27. Consider minimal sets such as $\mathcal{X}_\mathcal{P}$ in Proposition 12.5.1.
28. To answer the first part, construct for each forbidden minor X a finite set of graphs whose exclusion as topological minors is equivalent to forbidding X as a minor. For the second part recall Exercise 8.
29. Find the required paths one by one.
30.$^+$ One direction is just a weakening of Lemma 12.4.5. For the other, imitate the proof of Lemma 12.3.4.
31.$^+$ Let X be an externally ℓ-connected set of h vertices in a graph G, where h and ℓ are large. Consider a small separator S in G: clearly, most of X will lie in the same component of $G-S$. Try to make these 'large' components, perhaps together with their separators S, into the desired connected vertex sets.

Index

Page numbers in italics refer to definitions; in the case of author names, they refer to theorems due to that author. The alphabetical order ignores letters that stand as variables; for example, 'k-colouring' is listed under the letter c.

Symbol Index

The entries in this index are divided into two groups. Entries involving only mathematical symbols (i.e. no letters except variables) are listed on the first page, grouped loosely by logical function. The entry '[ ]', for example, refers to the definition of induced subgraphs $H[U]$ on page 4 as well as to the definition of face boundaries $G[f]$ on page 72.

Entries involving fixed letters as constituent parts are listed on the second page, in typographical groups ordered alphabetically by those letters. Letters standing as variables are ignored in the ordering.

Reinhard Diestel received a PhD from the University of Cambridge, following research 1983–86 as a scholar of Trinity College under Béla Bollobás. He was a Fellow of St. John's College, Cambridge, from 1986 to 1990. Research appointments and scholarships have taken him to Bielefeld (Germany), Oxford and the US. He became a professor in Chemnitz in 1994 and has held a chair at Hamburg since 1999.

Reinhard Diestel's main area of research is graph theory, including infinite graph theory. He has published numerous papers and a research monograph, *Graph Decompositions* (Oxford 1990).