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

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

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

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

... whether distinct vertices $s_1, \dots, s_k$ and $t_1, \dots, t_k$ in a graph can be linked by disjoint paths $P_i = s_i \dots t_i$. (This yields an $O(n^{k+3})$ algorithm to decide 'k-linked'.)

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

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

In each of e and e^*, the unique point of $\mathring{e}^* \cap \mathring{e}$ should be an inner point of a straight line segment.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$$AB = J.$$

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

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

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

6.6 Tutte's flow conjectures

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

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

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

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

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

In the 2nd edition, this is no longer a formal exercise.

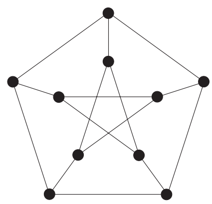

Fig. 6.6.1. The Petersen graph

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

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

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

Thus in either case we have found an integer $m \geqslant k/2$ and a graph $G_1 \preccurlyeq G$ such that

$$|G_1| \leqslant 4m \tag{1}$$

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.

vigorously

Replace $\varepsilon(H_0)$ by $\delta(H_0)$; 'higher' after replacing m with k, not in absolute terms.

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 on 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

$$f(n) := \tfrac{1}{6}m(n - m - 5)$$

and

$$\mathcal{H} := \left\{ H \preccurlyeq G_1 : \|H\| \geqslant m\,|H| + f(|H|) - \tbinom{m}{2} \right\}.$$

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).

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* $H_1, \dots, H_n$ *such that* $\mathrm{Forb}_{\preccurlyeq}(H_1, \dots, H_n)$ *contains precisely the graphs not embeddable in* S. *delete 'not'*

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

$$G_0, G_1, G_2, \dots$$

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 $G_1, G_2, \dots$ 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

Symbol Index

Unfortunately, the entire symbol index printed in the second edition is old. This page and the next show the corrected version.

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.

The page numbers on this page are shown as they should be, not as they are in the printed edition.