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 Sections 3.4 and 3.5 we look at other concepts of connectedness, more recent than the standard one but no less important: the number of H-paths in G for a subgraph H of G, and the existence of disjoint paths in G linking up specified pairs of vertices.

3.1 2-Connected graphs and subgraphs

The simplest 2-connected graphs are the cycles. All the others can be constructed inductively from a cycle by adding paths:

Proposition 3.1.1. *A graph is 2-connected if and only if it can be constructed from a cycle by successively adding H-paths to graphs H already constructed (Fig. 3.1.1).* [4.2.6]

Proof. Clearly, every graph constructed as described is 2-connected. Conversely, let a 2-connected graph G be given. Then G contains a

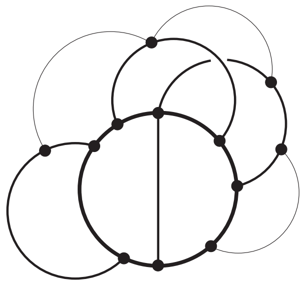

Fig. 3.1.1. The construction of 2-connected graphs

H cycle, and hence has a maximal subgraph H constructible as above. 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. □

Just as an arbitrary graph can be decomposed into its maximal connected subgraphs, or *components*, we can try to decompose a connected graph G into its maximal 2-connected subgraphs. These may not quite be disjoint, and they may not quite cover all of G. However, it is easy to weaken the notion of 'maximal 2-connected subgraph' slightly so that the subgraphs fitting the weaker notion do cover G and are still nearly disjoint. These 'blocks' fit together nicely in a tree-like fashion, which captures precisely the overall structure of G in terms of those blocks.

block Formally, a *block* is a maximal connected subgraph without a cutvertex.[1] Thus, every block 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.

Cycles and bonds are confined to a single block:

[4.6] **Lemma 3.1.2.** *Let G be any graph.*

(i) *The cycles in G are precisely the cycles in its blocks.*

(ii) *The bonds of G are precisely the bonds of its blocks.*

Proof. (i) Any cycle in G is a connected subgraph without a cutvertex, and hence lies in some maximal such subgraph. By definition, this is a block of G.

---

[1] ... of the subgraph; it may contain cutvertices of G.

(ii) The proof follows easily by repeated application of the following observation. Consider any cut in G. Let xy be one of its edges, and B the block containing it. By the maximality of B in the definition of a block, G contains no B-path. Hence every x–y path in G lies in B, so those edges of our cut that lie in B separate x from y even in G. □

As every edge lies in a unique block, belonging to a common block is an equivalence relation on the edge set of a graph. This equivalence can be expressed in two other interesting ways:

Lemma 3.1.3. *The following statements are equivalent for distinct edges e, f of a graph G:* [4.6]

(i) *The edges e, f belong to a common block of G.*

(ii) *The edges e, f belong to a common cycle in G.*

(iii) *The edges e, f belong to a common bond of G.*

Proof. (i)→(ii) It clearly suffices to prove that in a 2-connected graph any two 2-sets of vertices can be joined by two disjoint paths. This follows easily by induction based on Proposition 3.1.1.[2]

(ii)→(iii) Deleting e and f from a cycle $C \ni e, f$ leaves a partition of $V(C)$ into two connected sets. Extend this to a partitition into two connected sets of the vertex set of the component of G containing C. (How?) The edges between these sets form a bond of G containing e and f.

(iii)→(i) By Lemma 3.1.2 (ii), two edges can lie in a common bond only if they belong to the same block. □

Our last lemma on blocks shows how they fit together to form the coarse structure of G. Let A denote the set of cutvertices of G, and $\mathcal{B}$ the set of its blocks. The bipartite graph on $A \cup \mathcal{B}$ formed by the edges aB with $a \in B$ is the *block graph* of G, see Figure 3.1.2. *block graph*

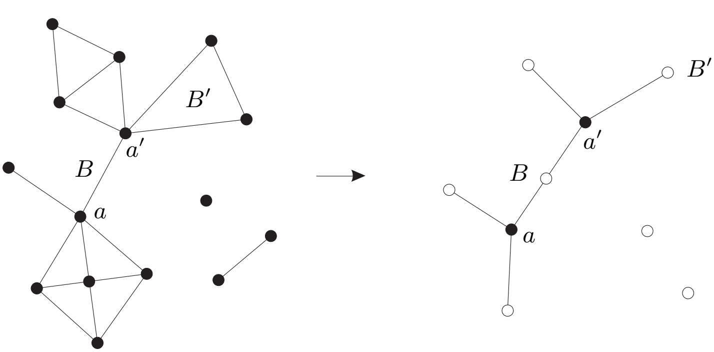

Fig. 3.1.2. A graph and its block graph

[2] See Exercise 5. Note that this is the case $k = 2$ of Menger's theorem (3.3.1).

Lemma 3.1.4. *The block graph of a connected graph is a tree.* □

Lemma 3.1.4 generalizes to graphs of higher connectivity: every $(k-1)$-connected graph has a canonical tree-like decomposition that separates all its 'k-blocks'. See Theorem 12.3.7 for the precise statement, and Exercise 17 in Chapter 12 for the case of $k = 3$.

3.2 The structure of 3-connected graphs

In this section we describe 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 theorem 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.

Proposition 3.1.1 describes how the 2-connected graphs can be constructed inductively, starting from a cycle. All the graphs constructed in the process were themselves 2-connected, so the graphs constructible in this way are precisely the 2-connected graphs. We shall now do something similar for 3-connected graphs. We shall prove that every 3-connected graph $G \neq K^4$ can be turned into a smaller 3-connected graph in two ways: by deleting an edge (and suppressing any vertices of degree 2 that may arise), and by contracting an edge. Inverting these processes will give us two independent ways of building all 3-connected graphs from a K^4.

$G \dot{-} e$ Given an edge e in a graph G, we write $G \dot{-} e$ for the *multigraph* obtained from $G - e$ by suppressing any end of e that has degree 2 in $G - e$.[3]

Lemma 3.2.1. *Let e be an edge in a graph G. If $G \dot{-} e$ is 3-connected, then so is G.*

(1.4.2) *Proof.* Thinking of G as obtained from $G \dot{-} e$ by adding e, let us call the vertices of $G \dot{-} e$ the *old* vertices of G, and any other vertex of G (which will be an end of e) a *new* vertex. Remembering that $G \dot{-} e$, being 3-connected, has no parallel edges, it is easy to see that, in G, no two vertices x_1, x_2 can separate a new vertex from all the old vertices. So it suffices to show that $\{x_1, x_2\}$ cannot separate two old vertices. If they did, then those old vertices would be separated in $G \dot{-} e$ by x'_1 and x'_2, where either $x'_i = x_i$ or, if x_i is new, x'_i is the edge of $G \dot{-} e$ subdivided by x_i. By Proposition 1.4.2, this contradicts our assumption that $G \dot{-} e$ is 3-connected. □

[3] See Chapter 1.10 for the formal definition of suppressing vertices in a multigraph. Recall also that 3-connected multigraphs cannot have multiple edges. Since parallel edges arising when a vertex is suppressed are not deleted, our assumption in Lemma 3.2.1 that the multigraph $G \dot{-} e$ is 3-connected implies that no parallel edges arise when it is formed from the graph G. Thus $G \dot{-} e$, too, is in fact a graph.

Lemma 3.2.2. *Every 3-connected graph $G \neq K^4$ has an edge e such that $G \mathbin{\dot{-}} e$ is another 3-connected graph.*

Proof. We start by showing that G contains a TK^4. Let C be a shortest cycle and $P = u \dots v$ a C-path in G. Then $\mathring{P} \neq \emptyset$ since C is induced, so $G - \{u, v\}$ contains a C–P path Q. Now $C \cup P \cup Q = TK^4$.

Thus, $TK^4 \subseteq G \neq K^4$; choose $H = TJ \subseteq G$ with some 3-connected $J \neq G$ so that $\|H\|$ is maximum. We shall find an edge e in G such that H
$G \mathbin{\dot{-}} e = J$.

As G is 3-connected but proper subdivisions of J are not, we have $H \neq G$. Let $P = u \dots v$ be an H-path in G, chosen if possible so that

u and v do not lie on the same (subdivided) edge of J. $(*)$ $P = u \dots v$

If P violates $(*)$ then $uv \in J = H$. Indeed, since G is 3-connected, the vertices subdividing an edge xy of J could be joined by an H-path P' in $G - \{x, y\}$ to a vertex not on that subdivided edge; as P' would satisfy $(*)$, this would contradict our choice of P. So $H = J$. Our assumption that P violates $(*)$ now means that $uv \in J$. Since G has no parallel u–v edges, P has an inner vertex. Now $(H - uv) \cup P$ is another TJ with more edges than H, contradicting our choice of H.

Therefore P satisfies $(*)$. Suppressing any vertices of degree 2 in $H \cup P$ we obtain a multigraph J' such that $J' \mathbin{\dot{-}} e = J$, where e is the edge corresponding to P. By $(*)$ the edge e is not parallel to an edge of J, so J', like J, is in fact a graph. By Lemma 3.2.1, J' is 3-connected. As $TJ' = H \cup P \subseteq G$ has more edges than H, our choice of J with $\|H\|$ maximum implies $J' = G$. Thus, $G \mathbin{\dot{-}} e = J$ as claimed. □

Theorem 3.2.3. (Tutte 1966)
A graph G is 3-connected if and only if there exists a sequence $G_0, \dots, G_n$ of graphs such that

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

(ii) *G_{i+1} has an edge e such that $G_i = G_{i+1} \mathbin{\dot{-}} e$, for every $i < n$.*

Moreover, the graphs in any such sequence are all 3-connected.

Proof. If G is 3-connected, use Lemma 3.2.2 to find $G_n, \dots, G_0$ in turn. Conversely, if $G_0, \dots, G_n$ is any sequence of graphs satisfying (i) and (ii), then all these graphs, and in particular $G = G_n$, are 3-connected by Lemma 3.2.1. □

Theorem 3.2.3 enables us to construct, recursively, the entire class of 3-connected graphs. Starting from K^4, we simply add to every graph already constructed a new edge in every way compatible with (ii): between two already existing vertices, between newly inserted subdividing vertices (not on the same edge), or between one old vertex and one new subdividing vertex.

We now turn to our second method of reducing 3-connected graphs to K^4, by contracting edges. In what follows we only consider graphs, not multigraphs.

[4.4.3] **Lemma 3.2.4.** *Every 3-connected graph $G \neq K^4$ has an edge e such that G/e is again 3-connected.*

xy *Proof.* Suppose there is no such edge e. Then, for every edge $xy \in G$, the graph G/xy contains a separator 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
z and $|S| = 2$, i.e. G has a vertex $z \notin \{x, y\}$ such that $\{v_{xy}, z\}$ 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,
C every vertex in T has a neighbour in every component C of $G - T$.

We choose the edge xy, the vertex z, and the component C so that
v $|C|$ is as small as possible, and pick a neighbour v of z in C (Fig. 3.2.1). By assumption, G/zv is again not 3-connected, so again there is a vertex
w w such that $\{z, v, w\}$ separates G, and as before every vertex in $\{z, v, w\}$ has a neighbour in every component of $G - \{z, v, w\}$.

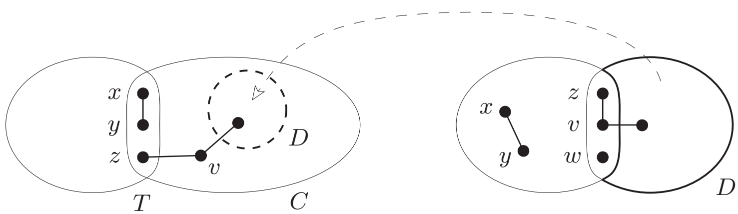

Fig. 3.2.1. Separating vertices in the proof of Lemma 3.2.4

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

Theorem 3.2.5. (Tutte 1961)
A graph G is 3-connected if and only if there exists a sequence $G_0, \ldots, G_n$ of graphs with the following two properties:

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

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

Moreover, the graphs in any such sequence are all 3-connected.

Proof. If G is 3-connected, then by Lemma 3.2.4 there is a sequence $G_n, \dots, G_0$ of 3-connected graphs satisfying (i) and (ii).

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

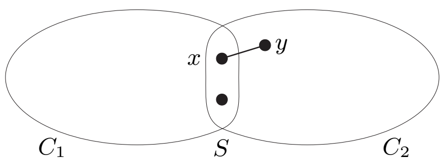

Fig. 3.2.2. The position of $xy \in G_{i+1}$ in the proof of Theorem 3.2.5

Like Theorem 3.2.3, Theorem 3.2.5 enables us to construct all 3-connected graphs inductively from K^4, by simple local alterations and without ever leaving the class of 3-connected graphs. Given a 3-connected graph already constructed, pick any vertex v and split it into two adjacent vertices v', v''; then join these to all the former neighbours of v, each to at least two. This is the essential core of a result of Tutte known as his *wheel theorem.*[4]

For larger integers k it is no longer true that in any k-connected graph we can contract an edge so as to obtain another k-connected graph. However, for every k there is a constant n_k such that in every k-connected graph we can either delete or contract an edge so that the resulting graph has no separation of order less than k in which both sides have at least n_k vertices. See the notes.

Theorem 3.2.6. (Tutte 1963) [4.5.2]
The cycle space of every 3-connected graph G is generated by its non-separating induced cycles.

Proof. Given a cycle C in G, let $b(C)$ denote the largest order of a component of $G - C$ if there is one, and put $b(C) := 0$ if $V(C) = V(G)$. Suppose the theorem is false, and consider a cycle C with $b(C)$ maximum that is not generated by non-separating induced cycles. b C

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

If $V(C) = V(G)$, then C is the sum of two cycles $C_1, C_2 \subseteq C + e$, where e a chord. As $b(C_1), b(C_2) > 0 = b(C)$, our choice of C implies that C_1 and C_2 are generated by non-separating induced cycles. But then so is their sum C, a contradiction.

B Assume now that $G - C \neq \emptyset$, and let B be a component of $G - C$ of order $|B| = b(C)$. Suppose first that

$G - B$ contains a C-path $P = u \ldots v$ such that each of the two u–v paths on C has an inner vertex in $N(B)$. $(*)$

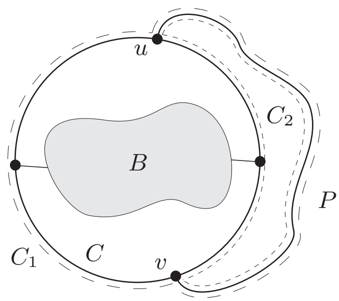

Fig. 3.2.3. C_1 and C_2 are drawn in broken lines

Then C is the sum of the two cycles $C_1, C_2 \subseteq C \cup P$ containing P, and for each of these C_i there is a component of $G - C_i$ that contains B properly (Fig. 3.2.3). Hence $b(C_i) > |B| = b(C)$, with a contradiction as earlier.

Suppose finally that $(*)$ fails. Then every vertex of C sends an edge to B. (Indeed, if not then C contains an $N(B)$-path $Q = x \ldots y$ with $\mathring{Q} \neq \emptyset$. As G is 3-connected, $C - Q \neq \emptyset$, and there is a $\mathring{Q}$–$(C - Q)$ path in $G - \{x, y\}$. Such a path P would satisfy $(*)$.) Since $V(C) = N(B)$, any chord of C would also be a path P as in $(*)$, so C has no chord. Hence unless C itself is induced and non-separating, $G - C$ has a component $B' \neq B$. Let $P = u \ldots v$ be a C-path through B', and let Q be a C–P

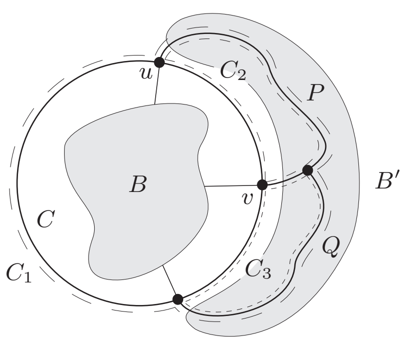

Fig. 3.2.4. Three cycles C_1, C_2, C_3 summing to C, and each missing a vertex of C that sends an edge to B

path in $G - \{u, v\}$. Note that Q too avoids B. Now $C \cup P \cup Q$ contains three cycles C_1, C_2, C_3 summing to C and each missing a vertex of C (Fig. 3.2.4). As every vertex of C sends an edge to B, we therefore have $b(C_i) > |B| = b(C)$ for every i, with the familiar contradiction. □

3.3 Menger's theorem

The following theorem is one of the cornerstones of graph theory. It is another example of the particularly attractive type of theorem disussed in the notes for Chapter 2, a discussion which equally applies here:

Theorem 3.3.1. (Menger 1927) [3.5.2] [8.2.5] [8.4.3] [12.4.3] [12.6.3]
Let $G = (V, E)$ be a graph and $A, B \subseteq V$. Then the minimum number of vertices separating A from B in G is equal to the maximum number of disjoint A–B paths in G.

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

First proof. We apply induction on $\|G\|$. If G has no edge, then $|A \cap B| = k$ and we have k trivial A–B paths. So we assume that G has an edge $e = xy$. If G has no k disjoint A–B paths, then neither does G/e; here, we count the contracted vertex v_e as an element of A (resp. B) in G/e if in G at least one of x, y lies in A (resp. B). By the induction hypothesis, G/e contains an A–B separator Y of fewer than k vertices. Among these must be the vertex v_e, since otherwise $Y \subseteq V$ would be an A–B separator in G. Then $X := (Y \setminus \{v_e\}) \cup \{x, y\}$ is an A–B separator in G of exactly k vertices.

We now consider the graph $G - e$. Since $x, y \in X$, every A–X separator in $G - e$ is also an A–B separator in G and hence contains at least k vertices. So by induction there are k disjoint A–X paths in $G - e$, and similarly there are k disjoint X–B paths in $G - e$. As X separates A from B, these two path systems do not meet outside X, and can thus be combined to k disjoint A–B paths. □

Let $\mathcal{P}$ be a set of disjoint A–B paths, and let $\mathcal{Q}$ be another such set. We say that $\mathcal{Q}$ *exceeds* $\mathcal{P}$ if the set of vertices in A that lie on a path in $\mathcal{P}$ is a proper subset of the set of vertices in A that lie on a path in $\mathcal{Q}$, and likewise for B. Then, in particular, $|\mathcal{Q}| \geqslant |\mathcal{P}| + 1$. *exceeds*

Second proof. We prove the following stronger statement:

If $\mathcal{P}$ is any set of fewer than $k = k\,(G, A, B)$ disjoint A–B paths in G, then G contains a set of $|\mathcal{P}|+1$ disjoint A–B paths exceeding $\mathcal{P}$.

Keeping G and A fixed, we let B vary and apply induction on $|\bigcup \mathcal{P}|$. Let R be an A–B path that avoids the (fewer than k) vertices of B that lie on a path in $\mathcal{P}$. If R avoids all the paths in $\mathcal{P}$, then $\mathcal{P} \cup \{R\}$ exceeds $\mathcal{P}$, as desired. (This will happen when $\mathcal{P} = \emptyset$, so the induction starts.) If not, let x be the last vertex of R that lies on some $P \in \mathcal{P}$ (Fig. 3.3.1).

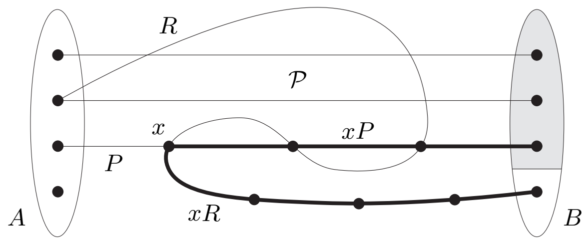

Fig. 3.3.1. Paths in the second proof of Menger's theorem

Put

$$B' := B \cup V(xP \cup xR) \quad \text{and} \quad \mathcal{P}' := \big(\mathcal{P} \smallsetminus \{P\}\big) \cup \{Px\}\,.$$

Then $|\mathcal{P}'| = |\mathcal{P}|$ but $|\bigcup \mathcal{P}'| < |\bigcup \mathcal{P}|$, and $k(G, A, B') \geqslant k(G, A, B)$, so by the induction hypothesis there is a set $\mathcal{Q}'$ of $|\mathcal{P}'|+1$ disjoint A–B' paths exceeding $\mathcal{P}'$. Then $\mathcal{Q}'$ contains a path Q ending at x, and a unique path Q' whose last vertex y is not among the last vertices of the paths in $\mathcal{P}'$.

If $y \notin xP$, we let $\mathcal{Q}$ be obtained from $\mathcal{Q}'$ by appending xP to Q, and appending yR to Q' if $y \notin B$. Otherwise $y \in \mathring{x}P$, and we let $\mathcal{Q}$ be obtained from $\mathcal{Q}'$ by appending xR to Q and yP to Q'. In all cases $\mathcal{Q}$ exceeds $\mathcal{P}$, as desired. □

Applied to a bipartite graph, Menger's theorem specializes to the assertion of König's theorem (2.1.1). For our third proof, we shall adapt the alternating path proof of König's theorem to the more general set-up of Theorem 3.3.1. Let again G, A, B be given, and let $\mathcal{P}$ be a set of disjoint A–B paths in G. Let us say that an A–B separator $X \subseteq V$ lies *on* $\mathcal{P}$ if it consists of a choice of exactly one vertex from each path in $\mathcal{P}$. If we can find such a separator X, then clearly $k \leqslant |X| = |\mathcal{P}|$, and Menger's theorem will be proved.

$\mathcal{P}$

on

Put

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

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

Let a walk $W = x_0 e_0 x_1 e_1 \dots e_{n-1} x_n$ in G with $e_i \neq e_j$ for $i \neq j$ be said to *alternate* with respect to $\mathcal{P}$ (Fig. 3.3.2) if it starts in $A \setminus V[\mathcal{P}]$ and the following three conditions hold for all $i < n$ (with $e_{-1} := e_0$ in (iii)): W, x_i, e_i *alternating walk*

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

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

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

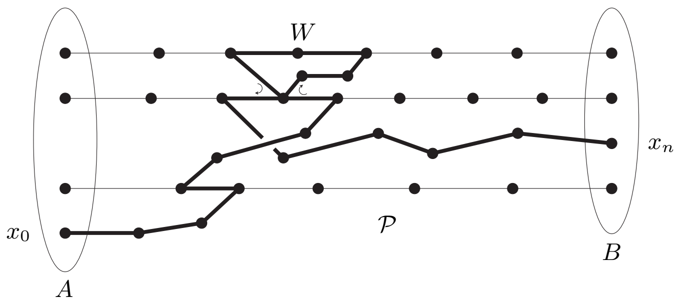

Fig. 3.3.2. An alternating walk from A to B

Note that, by (ii), any vertex outside $V[\mathcal{P}]$ occurs at most once on W. And since the edges e_i of W are all distinct, (iii) implies that any vertex $v \in V[\mathcal{P}]$ occurs at most twice on W. For $v \neq x_n$, this can happen in exactly the following two ways. If $x_i = x_j$ with $0 < i < j < n$, then

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

Unless otherwise stated, any use of the word 'alternate' below will refer to our fixed path system $\mathcal{P}$.

The next two lemmas together make up our third proof of Menger's theorem. We state and prove them in a way that makes them reusable in Chapter 8, when we prove Menger's theorem for infinite graphs.

Lemma 3.3.2. *If an alternating walk W as above ends in $B \setminus V[\mathcal{P}]$, then G contains a set of disjoint A–B paths exceeding $\mathcal{P}$.* [8.4.7]

Proof. We may assume that W has only its first vertex in $A \setminus V[\mathcal{P}]$ and only its last vertex in $B \setminus V[\mathcal{P}]$. Let H be the graph on $V(G)$ 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 vertex $a \in (A \cap V[\mathcal{P}]) \cup \{x_0\}$, 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 is an A–B path. (When G is infinite, this last conclusion uses the fact that W meets only finitely many paths in $\mathcal{P}$, and hence every component of H is finite.)

Similarly, for every $b \in (B \cap V[\mathcal{P}]) \cup \{x_n\}$ there is an A–B path in H that ends in b. The set of A–B paths in H therefore exceeds $\mathcal{P}$. $\square$

[8.4.7] **Lemma 3.3.3.** *If no alternating walk W as above ends in $B \setminus V[\mathcal{P}]$, then G contains an A–B separator on $\mathcal{P}$.*

Proof. Let

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

and

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

x_P For every path $P \in \mathcal{P}$, let x_P be the last vertex of P that lies on some alternating walk; if no such vertex exists, let x_P be the first vertex of P. Our aim is to show that

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

meets every A–B path in G; then X is an A–B separator on $\mathcal{P}$.

Q Suppose there is an A–B path Q that avoids X. We know that Q meets $V[\mathcal{P}]$, as otherwise it would be an alternating walk ending in B_2. Now the A–$V[\mathcal{P}]$ path in Q is either an alternating walk or consists only of the first vertex of some path in $\mathcal{P}$. Therefore Q also meets the vertex set $V[\mathcal{P}']$ of

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

y, P, x Let y be the last vertex of Q in $V[\mathcal{P}']$, say $y \in P \in \mathcal{P}$, and let $x := x_P$. As Q avoids X and hence x, we have $y \in P\mathring{x}$. In particular, $x = x_P$ is
W not the first vertex of P, and so there is an alternating walk W ending at x. Then $W \cup xPyQ$ is a walk from A_2 to B (Fig. 3.3.3). If this walk

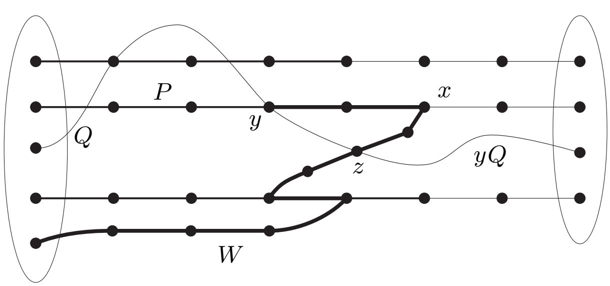

Fig. 3.3.3. Alternating walks in the proof of Lemma 3.3.3.

alternates and ends in B_2, we have our desired contradiction.

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

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

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

If W' does indeed meet $\mathring{y}Q$, we let z be the first vertex of W' on $\mathring{y}Q$ z and set $W'' := W'zQ$. Otherwise we set $W'' := W' \cup yQ$. In both W'' cases W'' alternates with respect to $\mathcal{P}'$, because W' does and $\mathring{y}Q$ avoids $V[\mathcal{P}']$. (W'' satisfies condition (iii) at y in the second case even if y occurs twice on W', because W'' then contains the entire walk W' and not just its initial segment $W'y$.) By definition of $\mathcal{P}'$, therefore, W'' avoids $V[\mathcal{P}] \smallsetminus V[\mathcal{P}']$. Thus W'' also alternates with respect to $\mathcal{P}$ and ends in B_2, contrary to our assumptions. □

Third proof of Menger's theorem. Let $\mathcal{P}$ contain as many disjoint A–B paths in G as possible. Then by Lemma 3.3.2, no alternating walk ends in $B \smallsetminus V[\mathcal{P}]$. By Lemma 3.3.3, this implies that G has an A–B separator X on $\mathcal{P}$, giving $k \leqslant |X| = |\mathcal{P}|$ as desired. □

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

Corollary 3.3.4. *For $B \subseteq V$ and $a \in V \smallsetminus B$, the minimum number of vertices separating a from B in G is equal to the maximum number of paths forming an a–B fan in G.*

Proof. Apply Theorem 3.3.1 to $G - a$ with $A := N_G(a)$. □

Corollary 3.3.5. *Let a and b be two distinct vertices of G.*

(i) *If $ab \notin E$, then the minimum number of vertices 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 to $G - \{a, b\}$, with $A := N_G(a)$ and $B := N_G(b)$.

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

[4.2.7] [6.6.1] [9.4.2]

Theorem 3.3.6. (Global Version of Menger's Theorem)

(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 3.3.5 (i), a and b are adjacent; let $G' := G - ab$. Then G' contains at most $k-2$ independent a–b paths. By Corollary 3.3.5 (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.

(ii) follows straight from Corollary 3.3.5 (ii). □

3.4 *A*-paths and Mader's theorem

In analogy to Menger's theorem we may consider the following questions. Given a set A of vertices in a graph, how many disjoint, edge-disjoint, or independent A-paths can we find in that graph? Is the largest number of such paths related to the least size of a vertex or edge cover of all its A-paths?

An early theorem of Gallai goes in this direction. In Erdős-Pósa spirit – see Chapter 2.3 – it says that if a graph fails to contain many disjoint A-paths it has a small set of vertices that covers all its A-paths, and thus 'totally disconnects' A:

Theorem 3.4.1. (Gallai 1961)
Let A be a set of vertices in a graph G, and $k \geqslant 1$ an integer. Then G either contains k disjoint A-paths or has a set of at most $2k-2$ vertices that meets every A-path.

The graph K^{2k-1}, with A its entire vertex set, shows that $2k-2$ is lowest possible as a general bound, one that depends only on k.

Theorem 3.4.1, as well as its analogue for edge-disjoint paths (with the same bound of $2k-2$), will be an easy consequence of Mader's theorem below, via its Corollary 3.4.3. See Exercises 25 and 30.

Mader's theorem solves our initial problem for independent H-paths rather than A-paths. Recall that, for $H \subseteq G$, an *H-path* in G is an A-path for $A = V(H)$ that is not just a single edge of H. If H is induced in G, as it will be for Mader's theorem, the H-paths in G are thus precisely its A-paths of length at least 2, for $A = V(H)$.

In addition to implying Gallai's theorem and its edge-analogue, Mader's theorem also implies a similar statement about independent A-paths: just subdivide any A-paths of length 1, which turns them into H-paths (with $V(H) = A$) without changing the number of A-paths.

In the fashion of Menger's theorem, Mader's theorem observes an upper bound on the number of independent H-paths that arises naturally from the size of certain separators, and then states that this bound is always attained by some set of paths.

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

In contrast to Menger's theorem, this bound can still be improved. The minimality of $|X \cup F|$ implies that no edge in F has an end in X: otherwise this edge would not be needed in the separator. Let $Y := V(G-H) \setminus X$, and denote by $\mathcal{C}_F$ the set of components of the $\mathcal{C}_F$
graph (Y, F). Since every H-path avoiding X contains an edge from F, it has at least two vertices in ∂C for some $C \in \mathcal{C}_F$, where ∂C denotes ∂C
the set of vertices in C that send an edge of G to $G - X - C$ (Fig. 3.4.1).

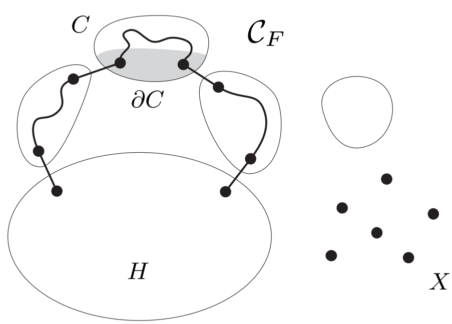

Fig. 3.4.1. An H-path in $G - X$

The number of independent H-paths in G is therefore bounded above by

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

where the minimum is taken over all X and F as described above:

X $X \subseteq V(G-H)$ and $F \subseteq E(G-H-X)$ such that every H-path in G has a vertex or an edge in $X \cup F$.

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

Theorem 3.4.2. (Mader 1978)
Given a graph G with an induced subgraph H, there are always $M_H(G)$ independent H-paths in G.

In order to obtain more immediate 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 $\kappa_H(G)$ the least cardinality of a vertex set $X \subseteq V(G-H)$ that meets every H-path in G. Similarly, we let $\lambda_H(G)$ denote the least cardinality of an edge set $F \subseteq E(G)$ that meets every H-path in G.

$\kappa_H(G)$

$\lambda_H(G)$

Corollary 3.4.3. *Given a graph G with an induced subgraph H, there are at least $\frac{1}{2}\kappa_H(G)$ independent H-paths and at least $\frac{1}{2}\lambda_H(G)$ edge-disjoint H-paths in G.*

k *Proof.* To prove the first assertion, let k be the maximum number of independent H-paths in G. By Theorem 3.4.2, there are sets $X \subseteq V(G-H)$ and $F \subseteq E(G-H-X)$ with

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

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

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

as claimed.

The second assertion follows from the first by considering the line graph of G (Exercise 27). □

It may come as a surprise to see that the bounds in Corollary 3.4.3 are best possible (as general bounds): one can find examples for G and H where G contains no more than $\frac{1}{2}\kappa_H(G)$ independent H-paths or no more than $\frac{1}{2}\lambda_H(G)$ edge-disjoint H-paths (Exercises 28 and 29).

3.5 Linking pairs of vertices

Let G be a graph, and let $X \subseteq V(G)$ be a set of vertices. We call X *linked* in G if whenever we pick distinct vertices $s_1, \dots, s_\ell$, $t_1, \dots, t_\ell$ in X we can find disjoint paths $P_1, \dots, P_\ell$ in G such that each P_i links s_i to t_i and has no inner vertex in X. Thus, unlike in Menger's theorem, we are not merely asking for disjoint paths between two *sets* of vertices: we insist that each of these paths shall link a specified pair of endvertices. *linked*

If $|G| \geqslant 2k$ and every set X of at most $2k$ vertices is linked in G, then G is k-*linked*. Clearly, this is equivalent to requiring merely that $|G| \geqslant 2k$ and disjoint paths $P_i = s_i \dots t_i$ exist for every choice of *exactly* $2k$ distinct vertices $s_1, \dots, s_k$, $t_1, \dots, t_k$: just add dummy vertices to X to bring it up to size $2k$. In practice, the latter is easier to prove, because we need not worry about inner vertices in X. *k-linked*

Clearly, every k-linked graph is k-connected. The converse, however, seems far from true: being k-linked is clearly a much stronger property than k-connectedness. Still, we shall prove in this section that we can force a graph to be k-linked by assuming that it is $f(k)$-connected, for some function $f\colon \mathbb{N} \to \mathbb{N}$. We first borrow a lemma from Chapter 7 to give a nice and simple proof that such a function f exists at all. In the remainder of the section we then prove that f can even be chosen linear.

The basic idea in the simple proof is as follows. If we can prove that G contains a subdivision K of a large complete graph, we can use Menger's theorem to link the vertices of X disjointly to branch vertices of K, and then hope to pair them up as desired through the subdivided edges of K. This requires, of course, that our paths do not hit too many of the subdivided edges before reaching the branch vertices of K.

The lemma saying that large enough connectivity does indeed force the existence of such a complete topological minor K will be proved in Chapter 7.2, where we consider several results of this type. By Theorem 1.4.3 it suffices to assume that G has large average degree:

Lemma 3.5.1. *There is a function $h\colon \mathbb{N} \to \mathbb{N}$ such that, for every $r \in \mathbb{N}$, every graph of average degree at least $h(r)$ contains K^r as a topological minor.*

Theorem 3.5.2. (Jung 1970; Larman & Mani 1970)
There is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that every $f(k)$-connected graph is k-linked, for all $k \in \mathbb{N}$.

Proof. We prove the assertion for $f(k) = h(3k) + 2k$, where h is a function as in Lemma 3.5.1. Let G be an $f(k)$-connected graph. Then $d(G) \geqslant \delta(G) \geqslant \kappa(G) \geqslant h(3k)$; let K be a TK^{3k} in G as provided by Lemma 3.5.1, and let U denote its set of branch vertices. (3.3.1) G K U

For the proof that G is k-linked, let distinct vertices $s_1, \dots, s_k$ and $t_1, \dots, t_k$ be given. By definition of $f(k)$, we have $\kappa(G) \geqslant 2k$. s_i, t_i

Hence by Menger's theorem (3.3.1), G contains disjoint paths $P_1, \dots, P_k$, P_i, Q_i $Q_1, \dots, Q_k$, such that each P_i starts in s_i, each Q_i starts in t_i, and all $\mathcal{P}$ these paths end in U but have no inner vertices in U. Let the set $\mathcal{P}$ of these paths be chosen so that their total number of edges outside $E(K)$ is as small as possible.

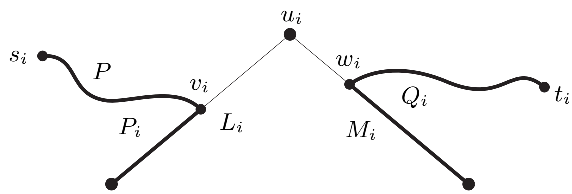

Fig. 3.5.1. Constructing an s_i–t_i path via u_i

Let $u_1, \dots, u_k$ be those k vertices in U that are not an end of a path in $\mathcal{P}$. For each $i = 1, \dots, k$, let L_i be the U-path in K (i.e., the subdivided edge of the K^{3k}) from u_i to the end of P_i in U, and let v_i be the first vertex of L_i on any path $P \in \mathcal{P}$. By definition of $\mathcal{P}$, P has no more edges outside $E(K)$ than $P v_i L_i u_i$ does, so $v_i P = v_i L_i$ and hence $P = P_i$ (Fig. 3.5.1). Similarly, if M_i denotes the U-path in K from u_i to the end of Q_i in U, and w_i denotes the first vertex of M_i on any path in $\mathcal{P}$, then this path is Q_i. Then the paths $s_i P_i v_i L_i u_i M_i w_i Q_i t_i$ are disjoint for different i and show that G is k-linked. □

The function h of Lemma 3.5.1 which our proof in Chapter 7.2 will yield will be exponential in r, and will therefore give only an exponential upper bound for the function $f(k)$ in Theorem 3.5.2. As $2\varepsilon(G) \geqslant \delta(G) \geqslant \kappa(G)$, the following result implies the linear bound of $f(k) = 16k$:

[7.2.3] **Theorem 3.5.3.** (Thomas & Wollan 2005)
Let G be a graph and $k \in \mathbb{N}$. If G is $2k$-connected and $\varepsilon(G) \geqslant 8k$, then G is k-linked.

We begin our proof of Theorem 3.5.3 with a lemma.

Lemma 3.5.4. *Any graph H with $\delta(H) \geqslant 8k \geqslant |H|/2$ has a k-linked subgraph.*

Proof. If H itself is k-linked there is nothing to show, so suppose not. X Then we can find a set X of $2k$ vertices $s_1, \dots, s_k, t_1, \dots, t_k$ that cannot $\mathcal{P}$ be linked in H by disjoint paths $P_i = s_i \dots t_i$. Let $\mathcal{P}$ be a set of as many such paths as possible, without inner vertices in X and all of length at most 7. If there are several such sets $\mathcal{P}$, we choose one with $|\bigcup \mathcal{P}|$ $s_1,\ t_1$ minimum. We may assume that $\mathcal{P}$ contains no path from s_1 to t_1. Let

J be the subgraph of H induced by X and all the vertices on the paths in $\mathcal{P}$, and let $K := H - J$. J, K

Note that each vertex $v \in K$ has at most three neighbours on any given $P_i \in \mathcal{P}$: if it had four, then replacing the segment uP_iw between its first and its last neighbour on P_i by the path uvw would reduce $|\bigcup \mathcal{P}|$ and thus contradict our choice of $\mathcal{P}$. So v has at most 3 neighbours in J for every $i = 1, \ldots, k$, at most $3k$ in total. As $\delta(H) \geqslant 8k$ by assumption, as well as $|H| \leqslant 16k$ and $|X| = 2k$, we deduce that

$$\delta(K) \geqslant 5k \quad \text{and} \quad |K| \leqslant 14k\,. \tag{1}$$

Our next aim is to show that K is disconnected. Since each of the paths in $\mathcal{P}$ has at most eight vertices, we have $|J - \{s_1, t_1\}| \leqslant 8(k-1)$. Therefore s_1 has a neighbour s in K, and t_1 has a neighbour t in K. Put $S := \{\, s' \in K \mid d_K(s,s') \leqslant 2 \,\}$ and $T := \{\, t' \in K \mid d_K(t,t') \leqslant 2 \,\}$. Since $H - \bigcup \mathcal{P}$ contains no s_1–t_1 path of length at most 7, we have $S \cap T = \emptyset$, and there is no S–T edge in K. To prove that K is disconnected, it thus suffices to show that $V(K) = S \cup T$. But for any vertex $v \in K - (S \cup T)$ the sets $N_K(s)$, $N_K(t)$ and $N_K(v)$ are disjoint and each have size at least $5k$, contradicting (1).

So K is disconnected; let C be its smaller component. By (1),

$$2\delta(C) \geqslant 2\delta(K) \geqslant 7k + 3k \geqslant \tfrac{1}{2}\,|K| + 3k \geqslant |C| + 3k\,. \tag{2}$$

We complete the proof by showing that C is k-linked. As $\delta(C) \geqslant 5k$, we have $|C| \geqslant 2k$. Let Y be a set of at most $2k$ vertices in C. By (2), every two vertices in Y have at least $3k$ common neighbours, at least k of which lie outside Y. We can therefore link any desired $\ell \leqslant k$ pairs of vertices in Y inductively by paths of length 2 whose inner vertex lies outside Y. □

Before we launch into the proof of Theorem 3.5.3, let us look at its main ideas. To prove that G is k-linked, we have to consider a given set X of up to $2k$ vertices and show that X is linked in G. Ideally, we would like to use Lemma 3.5.4 to find a linked subgraph L somewhere in G, and then use our assumption of $\kappa(G) \geqslant 2k$ to obtain a set of $|X|$ disjoint X–L paths by Menger's theorem (3.3.1). Then X could be linked via these paths and L, completing the proof.

Unfortunately, we cannot expect to find a subgraph H such that $\delta(H) \geqslant 8k$ and $|H| \leqslant 16k$ (in which L could be found by Lemma 3.5.4); cf. Corollary 11.2.3. However, it is not too difficult to find a minor $H \preccurlyeq G$ that has such a subgraph (Ex. 21, Ch. 7), even so that the vertices of X come to lie in distinct branch sets of H. We may then regard X as a subset of $V(H)$, and Lemma 3.5.4 provides us with a linked subgraph L of H. The only problem now is that H need no longer be $2k$-connected, that is, our assumption of $\kappa(G) \geqslant 2k$ will not ensure that we can link X to L by $|X|$ disjoint paths in H.

And here comes the clever bit of the proof: it relaxes the assumption of $\kappa \geqslant 2k$ to a weaker assumption that does get passed on to H. This weaker assumption is that *if* we can separate X from another part of G (or H) by fewer than $|X|$ vertices, then this other part must be 'light': roughly, its own value of ε must not exceed $8k$. If X then fails to link to L by $|X|$ disjoint paths, and hence H has a separation $\{A, B\}$ with $X \subseteq A$ and $L \subseteq B$ and $|A \cap B| < |X|$, we know that ε is still at least $8k$ on $H[A]$, because the B-part of H was light.

The idea now is to continue the proof inside $H' := H[A]$ by induction. This still needs some ingenuity, since it is not enough that ε is large on H': we also need that for every low-order separation (A', B') of H' with $X \subseteq A'$ the B'-part is light. That need not be true. But when it fails, we shall be able to use induction on $H'[B']$ to show that $A' \cap B'$ is linked in $H'[B']$, and use this for our proof that X is linked in H.

X-separation

Given $k \in \mathbb{N}$, a graph G, and $A, B, X \subseteq V(G)$, call the ordered pair (A, B) an X-*separation* of G if $\{A, B\}$ is a proper separation of G of order at most $|X|$ and $X \subseteq A$. An X-separation (A, B) is *small* if $|A \cap B| < |X|$, and *linked* if $A \cap B$ is linked in $G[B]$.

small/linked

$\| \ \|^+$
light
heavy

Call a set $U \subseteq V(G)$ *light* in G if $\|U\|^+ \leqslant 8k\,|U|$, where $\|U\|^+$ denotes the number of edges of G with at least one end in U. A set of vertices is *heavy* if it is not light.

k

Proof of Theorem 3.5.3. We shall prove the following, for fixed $k \in \mathbb{N}$:

$G = (V, E)$
X

Let $G = (V, E)$ be a graph and $X \subseteq V$ a set of at most $2k$ vertices. If $V \setminus X$ is heavy and for every small X-separation (A, B) the set $B \setminus A$ is light, then X is linked in G. $(*)$

To see that $(*)$ implies the theorem, assume that $\kappa(G) \geqslant 2k$ and $\varepsilon(G) \geqslant 8k$, and let X be a set of exactly $2k$ vertices. Then G has no small X-separation. And $V \setminus X$ is heavy, since

$$\|V \setminus X\|^+ \geqslant \|G\| - \binom{2k}{2} > 8k\,|V| - 16k^2 = 8k\,|V \setminus X|\,.$$

By $(*)$, X is linked in G, completing the proof that G is k-linked.

We prove $(*)$ by induction on $|G|$, and for each value of $|G|$ by induction on $\|V \setminus X\|^+$. If $|G| = 1$ then X is linked in G. For the induction step, let G and X be given as in $(*)$. We first prove the following:

We may assume that G has no linked X-separation. (1)

(A, B)
S

For our proof of (1), suppose that G has a linked X-separation (A, B). Let us choose one with A minimal, and put $S := A \cap B$.

We first consider the case that $|S| = |X|$. If $G[A]$ contains $|X|$ disjoint X–S paths, then X is linked in G because (A,B) is linked, completing the proof of $(*)$. If not, then by Menger's theorem (3.3.1) $G[A]$ has a small X-separation (A',B') such that $B' \supseteq S$. If we choose this of minimum order, i.e. with $|A' \cap B'|$ minimum, we can link $A' \cap B'$ to S in $G[B']$ by $|A' \cap B'|$ disjoint paths, again by Menger's theorem. But then $(A', B' \cup B)$ is a linked X-separation of G that contradicts the choice of (A,B).

So $|S| < |X|$. Let G' be obtained from $G[A]$ by adding any missing edges on S, so that $G'[S]$ is a complete subgraph of G'. As (A,B) is now a small X-separation, our assumption in $(*)$ says that $B \setminus A$ is light in G. Thus, G' arises from G by deleting $|B \setminus A|$ vertices outside X and at most $8k\,|B \setminus A|$ edges, and possibly adding some edges. As $V \setminus X$ is heavy in G, this implies that G'

$$A \setminus X \textit{ is heavy in } G'.$$

In order to be able to apply the induction hypothesis to G', let us show next that for every small X-separation (A',B') of G' the set $B' \setminus A'$ is light in G'. Suppose not, and choose a counterexample (A',B') with B' minimal. As $G'[S]$ is complete, we have $S \subseteq A'$ or $S \subseteq B'$. (A',B')

If $S \subseteq A'$ then $B \cap B' \subseteq S \subseteq A'$, so $(A' \cup B, B')$ is a small X-separation of G. Moreover,

$$B' \setminus (A' \cup B) = B' \setminus A',$$

and no edge of $G' - E$ on S is incident with this set (Fig 3.5.2). Our assumption that this set is heavy in G', by the choice of (A',B'), therefore implies that it is heavy also in G. As $(A' \cup B, B')$ is a small X-separation of G, this contradicts our assumptions in $(*)$.

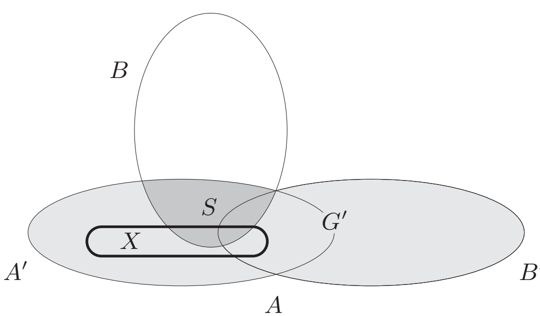

Fig. 3.5.2. If $S \subseteq A'$, then $(A' \cup B, B')$ is an X-separation of G

Hence $S \subseteq B'$. By our choice of (A',B'), the graph $G'' := G'[B']$ satisfies the premise of $(*)$ for $X'' := A' \cap B'$. Indeed, $B' \setminus X'' = B' \setminus A'$

is heavy, and by the minimality of B' any small X''-separation (A'', B'') of G'' will be such that $B'' \setminus A''$ is light, because $(A' \cup A'', B'')$ will be a small X-separation of G', and $B'' \setminus A'' = B'' \setminus (A' \cup A'')$.

By the induction hypothesis, therefore, X'' is linked in G''. But then X'' is also linked in $G[B' \cup B]$: as S was linked in $G[B]$, we simply replace any edges added on S in the definition of G' by disjoint paths through B (Fig. 3.5.3). But now $(A', B' \cup B)$ is a linked X-separation of G that violates the minimality of A in the choice of (A, B).

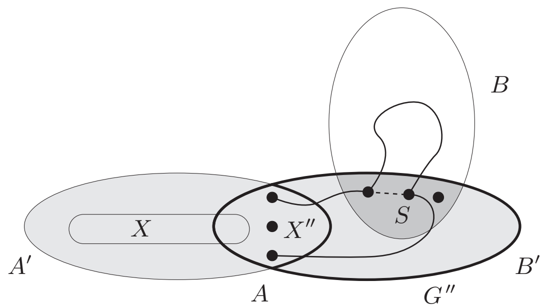

Fig. 3.5.3. If $S \subseteq B'$, then $(A', B' \cup B)$ is linked in G

We have thus shown that G' satisfies the premise of $(*)$ with respect to X. Since $\{A, B\}$ is a proper separation, G' has fewer vertices than G. By the induction hypothesis, therefore, X is linked in G'. Replacing edges of $G' - E$ on S by paths through B as before, we can turn any linkage of X in G' into one in G, completing the proof of $(*)$. This completes the proof of (1).

Our next goal is to show that, by the induction hypothesis, we may assume that G has not only large average degree but even large minimum degree. For our proof that X is linked in G, let $s_1, \dots, s_\ell, t_1, \dots, t_\ell$ be the distinct vertices in X which we wish to link by disjoint paths $P_i = s_i \dots t_i$. Let us add to G any missing edges on X except those of

$G[X]$ the form $s_i t_i$; as the paths P_i are not allowed to have inner vertices in X, these new edges affect neither the premise nor the conclusion in $(*)$.

After this modification, we can now prove the following:

We may assume that any two adjacent vertices u, v which do not both lie in X have at least $8k-1$ common neighbours. (2)

$e = uv$ To prove (2), let $e = uv$ be such an edge, let n denote the number of
n common neighbours of u and v, and let $G' := G/e$ be the graph obtained
G' by contracting e. Since u, v are not both in X we may view X as a subset
V' also of $V' := V(G')$, replacing u or v in X with the contracted vertex v_e if $X \cap \{u, v\} \neq \emptyset$. Our aim is to show that unless $n \geqslant 8k - 1$ as desired in (2), G' satisfies the premise of $(*)$. Then X will be linked in G' by

the induction hypothesis, so the desired paths $P_1, \dots, P_\ell$ exist in G'. If one of them contains v_e, replacing v_e by u or v or uv turns it into a path in G, completing the proof of $(*)$.

In order to show that G' satisfies the premise of $(*)$ with respect to X, let us show first that $V' \smallsetminus X$ is heavy. Since $V \smallsetminus X$ was heavy and $|V' \smallsetminus X| = |V \smallsetminus X| - 1$, it suffices to show that the contraction of e resulted in the loss of at most $8k$ edges incident with a vertex outside X. If u and v are both outside X, then the number of such edges lost is only $n+1$: one edge at every common neighbour of u and v, as well as e. But if $u \in X$, then $v \notin X$, and we lost all the X–v edges xv of G with $x \neq u$, too: while xv counted towards $\|V \smallsetminus X\|^+$, the edge xv_e lies in $G'[X]$ and does not count towards $\|V' \smallsetminus X\|^+$. If x is not a common neighbour of u and v, then this is an additional loss. But u is adjacent to every $x \in X \smallsetminus \{u\}$ except at most one (by our assumption about $G[X]$), so every such x except at most one is in fact a common neighbour of u and v. Thus in total, we lost at most $n+2$ edges. Unless $n \geqslant 8k-1$ (which would prove (2) directly for u and v), this means that we lost at most $8k$ edges, as desired for our proof that $V' \smallsetminus X$ is heavy.

It remains to show that for every small X-separation (A', B') of G' the set $B' \smallsetminus A'$ is light. Let (A', B') be a counterexample, chosen with B' minimal. Then $G'[B']$, as in the proof of (1), satisfies the premise of $(*)$ with respect to $X' := A' \cap B'$. Hence X' is linked in $G'[B']$ by induction. Let A and B be obtained from A' and B' by replacing v_e, where applicable, with both u and v, and put $X'' := A \cap B$. We shall prove that the separation (A, B) of G contradicts our assumption (1).

(A', B')

X'

(A, B), X''

Let us consider two possible positions of v_e in turn. If v_e lies in $B' \smallsetminus A'$, then $u, v \in B \smallsetminus A$. Then $X'' = X'$ is linked in $G[B]$, because it is linked in $G'[B']$: if v_e occurs on one of the linking paths for X', just replace it by u or v or uv as earlier. This contradicts (1). The other case is that v_e lies in A', possibly in X'. We show that $G[B]$ satisfies the premise of $(*)$ with respect to X''; then X'' will be linked in $G[B]$ by induction, again contradicting (1). Since (A', B') is a small X-separation, we have

$$|X''| \leqslant |X'| + 1 \leqslant |X| \leqslant 2k\,,$$

even if v_e lies in X'. Moreover, $B \smallsetminus X'' = B' \smallsetminus A'$ is heavy in G, because it is heavy in G' by the choice of (A', B'). Now consider a small X''-separation (A'', B'') of $G[B]$. Then $(A \cup A'', B'')$ is a small X-separation of G, because $|X''| \leqslant |X|$. Therefore $B'' \smallsetminus A'' = B'' \smallsetminus (A \cup A'')$ is light by the assumption in $(*)$. Hence $G[B]$ does satisfy the premise of $(*)$ for X'', completing the proof of (2).

Using induction by contracting an edge, we have just shown that the vertices in $V \smallsetminus X$ may be assumed to have large degree. Using induction by deleting an edge, we now show that their degrees cannot be too large.

d^* Since $(*)$ holds if $V = X$, we may assume that $V \setminus X \neq \emptyset$; let d^* denote the smallest degree in G of a vertex in $V \setminus X$. Let us prove the following:

$$\text{We may assume that } 8k \leqslant d^* \leqslant 16k - 1. \tag{3}$$

The lower bound in (3) follows from (2) if we assume that G has no isolated vertex outside X, which we may clearly assume by induction. $e = uv$ For the upper bound, let us see what happens if we delete an edge e whose ends u, v are not both in X. If $G - e$ satisfies the premise of $(*)$ with respect to X, then X is linked in $G - e$ by induction, and hence in G. If not, then either $V \setminus X$ is light in $G - e$, or $G - e$ has a small X-separation (A, B) such that $B \setminus A$ is heavy. If the latter happens then e must be an $(A \setminus B)$–$(B \setminus A)$ edge: otherwise, (A, B) would be a small X-separation also of G, and $B \setminus A$ would be heavy also in G, in contradiction to our assumptions in $(*)$. But if e is such an edge then any common neighbours of u and v lie in $A \cap B$, so there are fewer than $|X| \leqslant 2k$ such neighbours. This contradicts (2).

So $V \setminus X$ must be light in $G - e$. For G, this yields

$$\|V \setminus X\|^+ \leqslant 8k\,|V \setminus X| + 1\,. \tag{4}$$

In order to show that this implies the desired upper bound for d^*, let us $f(x)$ estimate the number $f(x)$ of edges that a vertex $x \in X$ sends to $V \setminus X$. There must be at least one such edge, xy say, as otherwise $(X, V \setminus \{x\})$ would be a small X-separation of G that contradicts our assumptions in $(*)$. But then, by (2), x and y have at least $8k - 1$ common neighbours, at most $2k - 1$ of which lie in X. Hence $f(x) \geqslant 6k$. As

$$2\,\|V \setminus X\|^+ = \sum_{v \in V \setminus X} d_G(v) + \sum_{x \in X} f(x)\,,$$

an assumption of $d^* \geqslant 16k$ would thus imply that

$$2\,(8k\,|V \setminus X| + 1) \underset{(4)}{\geqslant} 2\,\|V \setminus X\|^+ \geqslant 16k\,|V \setminus X| + 6k\,|X|\,,$$

yielding the contradiction of $2 \geqslant 6k\,|X|$. This completes the proof of (3).

To complete our proof of $(*)$, pick a vertex $v_0 \in V \setminus X$ of degree d^*, and consider the subgraph H induced in G by v_0 and its neighbours. By (2) we have $\delta(H) \geqslant 8k$, and by (3) and the choice of v_0 we have $|H| \leqslant 16k$. By Lemma 3.5.4, then, H has a k-linked subgraph; let L be its vertex set. By definition of 'k-linked', we have $|L| \geqslant 2k \geqslant |X|$. If G contains $|X|$ disjoint X–L paths, then X is linked in G, as desired. If not, then G has a small X-separation (A, B) with $L \subseteq B$. If we choose (A, B) of minimum order, then $G[B]$ contains $|A \cap B|$ disjoint $(A \cap B)$–L paths by Menger's theorem (3.3.1). But then (A, B) is a linked X-separation that contradicts (1). □

Exercises

For the first three exercises let G be a graph with vertices a and b, and let $X \subseteq V(G) \setminus \{a, b\}$ be an a–b separator in G.

1.$^{-}$ Show that X is minimal as an a–b separator 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.$^{-}$ (continued)

Let $X' \subseteq V(G) \setminus \{a, b\}$ be another a–b separator, 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).

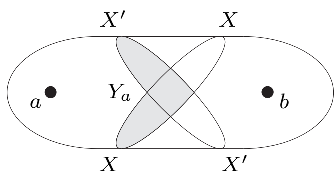

3. (continued)

Are Y_a and Y_b minimal a–b separators if X and X' are? Are $|Y_a|$ and $|Y_b|$ minimum for a–b separators if $|X|$ and $|X'|$ are?

4. Let X and X' be minimal separators in G. Show that if X meets at least two components of $G - X'$, then X' meets all the components of $G - X$ and X meets *all* the components of $G - X'$.

5.$^{-}$ Deduce the $k = 2$ case of Menger's theorem (3.3.1) from Proposition 3.1.1.

6.$^{-}$ Prove the elementary properties of blocks mentioned after their formal definition.

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

8.$^{-}$ Let G be a k-connected graph, and let xy be an edge of G. Show that G/xy is k-connected if and only if $G - \{x, y\}$ is $(k-1)$-connected.

9. (i) Let e be an edge in a 2-connected graph $G \neq K^3$. Show that either $G - e$ or G/e is again 2-connected.

(ii) Does every 2-connected graph $G \neq K^3$ have an edge e such that G/e is still 2-connected?

10.$^{+}$ Let e be an edge in a 3-connected graph $G \neq K^4$. Show that either $G \mathbin{\dot{-}} e$ or G/e is again 3-connected.

11. Show without using Theorem 3.2.6 that every edge of a 3-connected graph lies on some non-separating induced cycle.

12. Give an inductive proof of Theorem 3.2.6 based on Lemma 3.2.2. You may use the previous exercise.

13.$^+$ Give an inductive proof of Theorem 3.2.6 based on Lemma 3.2.4.

14.$^+$ Show that every transitive graph G with $\kappa(G) = 2$ is a cycle.

15.$^-$ At which point does the first proof of Menger's theorem fail if we assign the contracted vertex v_e to A in G/e only if A in G contains both ends of e, and similarly for B?

16. When one tries to prove an unknown implication $a \Rightarrow b$, it can be dangerous to attempt to prove $a \Rightarrow c \Rightarrow b$ for some assertion c that clearly implies b: if c is too strong, then $a \Rightarrow c$ may fail even if $a \Rightarrow b$ is true. But the first proof of Menger's theorem appears to be doing just that: it proves the seemingly very strong assertion that, given G, A, B and *any* edge e of G, we can contract or delete e without decreasing $k(G, A, B)$ – from which the existence of k disjoint A–B paths follows easily by induction. Can one see already at the outset of the proof that this route will in fact not be dangerous?

17. (i) Find the error in the following 'simple proof' of Menger's theorem (3.3.1). Let X be an A–B separator of minimum size. Denote by G_A the subgraph of G induced by X and all the components of $G - X$ that meet A, and define G_B correspondingly. By the minimality of X, there can be no A–X separator in G_A with fewer than $|X|$ vertices, so G_A contains k disjoint A–X paths by induction. Similarly, G_B contains k disjoint X–B paths. Together, all these paths form the desired A–B paths in G.

    (ii) Fill the gap in the proof of (i) by considering, if possible, a vertex or edge outside A and B.

18. Prove Menger's theorem by induction on $\|G\|$, as follows. Given an edge $e = xy$, consider a smallest A–B separator S in $G - e$. Show that the induction hypothesis implies a solution for G unless $S \cup \{x\}$ and $S \cup \{y\}$ are smallest A–B separators in G. Then show that if choosing neither of these separators as X in the previous exercise gives a valid proof, there is only one easy case left to do.

19. Work out the details of the proof of Corollary 3.3.5 (ii).

20.$^-$ Show that the least number of edges separating two disjoint sets A, B of vertices in a graph G equals the maximum number of edge-disjoint A–B paths in G.

21. Is there a function $f : \mathbb{N} \to \mathbb{N}$ such that, for every $k \in \mathbb{N}$, every $f(k)$-edge-connected graph contains between any two vertices k edge-disjoint paths every two of which traverse their common vertices in the same order?

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

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

24. Find a subset D of the plane and two infinite subsets $A, B \subseteq D$ such that for every finite set $X \subseteq D$ there is an A–B arc in $D \setminus X$ but D contains no infinite set of disjoint A–B arcs.

25.$^+$ Prove the following weakening of Gallai's theorem without the help of Mader's theorem. Find a constant c such that, for every k, G and A, either G contains k disjoint A-paths or G has a set of at most ck vertices that meets all its A-paths.

26.$^+$ Given a collection $\mathcal{S}$ of disjoint vertex sets in a graph $G = (V, E)$, a path in G is an *$\mathcal{S}$-path* if it joins distinct sets in $\mathcal{S}$ and has no inner vertices in $S := \bigcup \mathcal{S}$. Show that the following version of Mader's theorem is equivalent to Theorem 3.4.2: The maximum number of disjoint $\mathcal{S}$-paths in G is always equal to the minimum value of $|V_0| + \sum_{i=1}^{r} \lfloor \frac{1}{2} |\partial V_i| \rfloor$ taken over all partitions $\{V_0, \dots, V_r\}$ of V such that every $\mathcal{S}$-path in $G - V_0$ has an edge spanned by some V_i, where ∂V_i is the set of vertices in V_i that lie in S or have a neighbour outside $V_0 \cup V_i$.

27. Derive the edge part of Corollary 3.4.3 from the vertex part.

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

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

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

30.$^-$ Deduce Gallai's theorem, as well as an edge analogue with the same bound of $2k - 2$, from Mader's theorem via Corollary 3.4.3.

31. Derive a suitable version of Menger's theorem from Mader's theorem.

32.$^+$ 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'.)

33.$^-$ Show that $2k$-edge-connected graphs are *k-edge-linked* in the sense that for all distinct vertices $s_1, \dots, s_k, t_1, \dots, t_k$ there are edge-disjoint paths $P_i = s_i \dots t_i$ for $i = 1, \dots, k$.

34.$^-$ Show that k-linked graphs are $(2k-1)$-connected. Are they even $2k$-connected?

35. For every $k \in \mathbb{N}$ find an $\ell = \ell(k)$, as large as possible, such that not every ℓ-connected graph is k-linked.

36. Show that if G is k-linked and $s_1, \dots, s_k, t_1, \dots, t_k$ are not necessarily distinct vertices such that $s_i \neq t_i$ for all i, then G contains independent paths $P_i = s_i \dots t_i$ for $i = 1, \dots, k$.

37. Go through the proof of Theorem 3.5.3 monitoring the use of $\|V \smallsetminus X\|^+$. How would the proof fail if $\|G[V \smallsetminus X]\|$ was used instead? Which arguments would become simpler?

38. In the informal discussion preceding the proof of Theorem 3.5.3 we noted that, by Corollary 11.2.3, we cannot expect to find in G a subgraph H that satisfies the premise of Lemma 3.5.4. But then the proof of $(*)$ does find such a subgraph H. Can you explain this?

39. Use Theorem 3.5.3 to show that the function h in Lemma 3.5.1 can be chosen as $h(r) = cr^2$, for some $c \in \mathbb{N}$.

Notes

Although connectivity theorems are doubtless among the most natural, and also the most applicable, results in graph theory, there is still no monograph on this subject. The most comprehensive sources to date are A. Schrijver, *Combinatorial optimization*, Springer 2003, and A. Frank, *Connections in combinatorial optimization*, Oxford University Press 2011. 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.

Theorem 3.2.3 is often attributed to Barnette and Grünbaum (1969). It can also be extracted from W.T. Tutte, *Connectivity in graphs*, Oxford University Press 1966. Tutte's *wheel theorem*, proved in W.T. Tutte, A theory of 3-connected graphs, *Nederl. Akad. Wet. Proc. Ser. A* **64** (1961), 441–455, differs from our Theorem 3.2.5 as follows. As an alternative to the contraction of an edge in the reduction step, the wheel theorem also allows its deletion. Of edges to be contracted, however, it requires that they do not lie in any triangle. The starting set for the construction of all 3-connected graphs therefore consists of all wheels, not only K^4.

Tutte's wheel theorem has been extended to 3-connected graphs H other than K^4: from any 3-connected graph $G \succcurlyeq H$ that is not a wheel we can obtain H by contracting or deleting edges step by step, remaining 3-connected at every step. (As in Tutte's theorem, one is not allowed to contract edges that lie in a triangle.) This was proved by S. Negami, A characterization of 3-connected graphs containing a given graph, *J. Comb. Theory, Ser. B* **32** (1982), 69–74. It also follows from an earlier theorem of Seymour on matroid decompositions, and is sometimes called Seymour's *splitter theorem* for 3-connected graphs.

The fact, mentioned after the proof of Theorem 3.2.5, that in k-connected graphs we can either delete or contract an edge so that the resulting graphs have no separations of order $< k$ with unboundedly large sides follows from Lemma 3.1 of J. Geelen, B. Gerards, N. Robertson and G. Whittle, On the excluded minors for the matroids of branch-width k, *J. Comb. Theory, Ser. B* **88** (2003), 261–265.

Our proof of Theorem 3.2.6 is the original from W.T. Tutte, How to draw a graph, *Proc. Lond. Math. Soc.* **13** (1963), 743–767. Alternative proofs are indicated in Exercises 12 and 13.

An approach to the study of connectivity not touched upon in this chapter is the investigation of edge-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. Exercise 17 in Chapter 1 says that every edge-minimal k-*edge*-connected graph has a vertex of degree exactly k. Halin (1969) proved that this holds even for all edge-minimal k-connected graphs.

The existence of a vertex of small degree can be 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.

Menger's theorem goes back to his paper, Zur allgemeinen Kurventheorie, *Fundamenta Math.* **10** (1927), 96–115. It is probably the most-used classical result in graph theory. Our first proof is due to Halin; he published it only in his book, from where our proof is extracted. The second is due to T. Böhme, F. Göring and J. Harant, Menger's theorem, *J. Graph Theory* **37** (2001), 35–36, the third to T. Grünwald (later Gallai), Ein neuer Beweis eines Mengerschen Satzes, *J. Lond. Math. Soc.* **13** (1938), 188–192. A fourth proof is sketched in Exercise 18, and in Chapter 6 we shall obtain a fifth proof as an application of a theorem about network flows (Ch. 6, Ex. 3.) The global version of Menger's theorem (Theorem 3.3.6) was first stated and proved by Whitney (1932). Topological generalizations of Menger's theorem have been known since the 1930s; see C. Thomassen and A. Vella, Graph-like continua and Menger's theorem, *Combinatorica* **28** (2008), 595–623.

Theorem 3.4.1 is due to T. Gallai, Maximum-Minimum Sätze und verallgemeinerte Faktoren von Graphen, *Acta Math. Hungar.* **12** (1961), 131–173.

Mader's Theorem 3.4.2 is taken from W. Mader, Über die Maximalzahl kreuzungsfreier H-Wege, *Arch. Math.* **31** (1978), 387–402; our formulation is easily seen to be equivalent to the original. The shortest proof known to me is given by Schrijver in his book. The theorem may be viewed as a common generalization of Menger's theorem and Tutte's 1-factor theorem (Exercise 32).

Theorem 3.5.3 is due to R. Thomas and P. Wollan, An improved linear bound for graph linkages, *Eur. J. Comb.* **26** (2005), 309–324. Using a more involved version of Lemma 3.5.4, they prove that $2k$-connected graphs even with only $\varepsilon \geqslant 5k$ must be k-linked. And for graphs of large enough girth the condition on ε can be dropped altogether: as shown by W. Mader, Topological subgraphs in graphs of large girth, *Combinatorica* **18** (1998), 405–412, such graphs are k-linked as soon as they are $2k$-connected, which is best possible.

(Mader assumes a lower bound on the girth that depends on k, but this is not necessary; see D. Kühn & D. Osthus, Topological minors in graphs of large girth, *J. Comb. Theory, Ser. B* **86** (2002), 364–380.) In fact, for every $s \in \mathbb{N}$ there exists a k_s such that if $G \not\supseteq K_{s,s}$ and $\kappa(G) \geqslant 2k \geqslant k_s$ then G is k-linked; see D. Kühn & D. Osthus, Complete minors in $K_{s,s}$-free graphs, *Combinatorica* **25** (2005) 49–64.