12 Graph Minors

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*, 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 graph 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 concept that is not only central to graph minor theory but has found algorithmic applications too, and *tangles*, a radically new notion of high connectivity somewhere inside a given graph.

Section 12.1 gives an introduction to *well-quasi-ordering*, a concept central to the graph minor theorem. In Section 12.2 we apply this to prove the graph minor theorem for trees. We study tree-decompositions in Sections 12.3–12.4, and tangles in Section 12.5. In Section 12.6 we look at *forbidden-minor theorems*: results in the spirit of Kuratowski's theorem (4.4.6) or Wagner's theorem (7.3.4), which describe the structure of the graphs not containing some specified graph or graphs as a minor.

In Section 12.7 we give a direct proof of the 'generalized Kuratowski' theorem that embeddability in any fixed surface can be characterized by forbidding finitely many minors. We conclude with an overview of the proof and implications of the graph minor theorem itself.

12.1 Well-quasi-ordering

$\leqslant, <$ *well-quasi-ordering* *good pair* *good/bad sequence*

A reflexive and transitive binary relation $\leqslant$ on a set X is a *quasi-ordering*. We write $x < y$ for '$x \leqslant y$ and $y \nleqslant x$'. 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 there are indices $i < j$ such that $x_i \leqslant x_j$. Then (x_i, x_j) is a *good pair* of this sequence. A sequence containing a good pair is a *good sequence*; thus, a quasi-ordering on X is a well-quasi-ordering if and only if every infinite sequence in X is good. An infinite sequence is *bad* if it is not good.

Proposition 12.1.1. *A quasi-ordering $\leqslant$ on X is a well-quasi-ordering if and only if X contains neither an infinite antichain nor an infinite strictly decreasing sequence $x_0 > x_1 > \ldots$.*

(9.1.2) *Proof.* The forward implication is trivial. Conversely, let $x_0, x_1, \ldots$ be any infinite sequence in X. Let K be the complete graph on $\mathbb{N} = \{0, 1, \ldots\}$. Colour the edges ij $(i < j)$ of K with three colours: green if $x_i \leqslant x_j$, red if $x_i > x_j$, and amber if x_i, x_j are incomparable. By Ramsey's theorem (9.1.2), K has an infinite induced subgraph whose edges all have the same colour. If there is neither an infinite antichain nor an infinite strictly decreasing sequence in X, then this colour must be green, i.e. $x_0, x_1, \ldots$ has an infinite subsequence in which every pair is good. In particular, the sequence $x_0, x_1, \ldots$ is good. $\square$

In the proof of Proposition 12.1.1, we showed more than was needed: rather than finding a single good pair in $x_0, x_1, \ldots$, we found an infinite increasing subsequence. We have thus shown the following:

Corollary 12.1.2. *If X is well-quasi-ordered, then every infinite sequence in X has an infinite increasing subsequence.* $\square$

The following lemma, and the idea of its proof, are fundamental to the theory of well-quasi-ordering. Let $\leqslant$ be a quasi-ordering on a set X.

$\leqslant$ $[X]^{<\omega}$

For finite subsets $A, B \subseteq X$, write $A \leqslant B$ if there is an injective mapping $f \colon A \to B$ such that $a \leqslant f(a)$ for all $a \in A$. This naturally extends $\leqslant$ to a quasi-ordering on $[X]^{<\omega}$, the set of all finite subsets of X.

[12.2.1] [12.7.1]

Lemma 12.1.3. *If X is well-quasi-ordered by $\leqslant$, then so is $[X]^{<\omega}$.*

Proof. Suppose that $\leqslant$ is a well-quasi-ordering on X but not on $[X]^{<\omega}$. We start by constructing a bad sequence $(A_n)_{n \in \mathbb{N}}$ in $[X]^{<\omega}$, as follows. Given $n \in \mathbb{N}$, assume inductively that A_i has been defined for every $i < n$, and that there exists a bad sequence in $[X]^{<\omega}$ starting with $A_0, \ldots, A_{n-1}$. (This is clearly true for $n = 0$: by assumption, $[X]^{<\omega}$ contains a bad sequence, and this has the empty sequence as an initial segment.) Choose $A_n \in [X]^{<\omega}$ so that some bad sequence in $[X]^{<\omega}$ starts with $A_0, \dots, A_n$ and $|A_n|$ is as small as possible.

Clearly, $(A_n)_{n \in \mathbb{N}}$ is a bad sequence in $[X]^{<\omega}$; in particular, $A_n \neq \emptyset$ for all n. For each n pick an element $a_n \in A_n$ and set $B_n := A_n \setminus \{a_n\}$.

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

$$A_0, \dots, A_{n_0-1}, B_{n_0}, B_{n_1}, B_{n_2}, \dots$$

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 graph 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) [12.7.1]
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. $T \leqslant T'$

Proof of Theorem 12.2.1. We show that the rooted trees are well-quasi-ordered by the relation $\leqslant$ defined above; this clearly implies the theorem. (12.1.3)

Suppose not. To derive a contradiction, we proceed as in the proof of Lemma 12.1.3. Given $n \in \mathbb{N}$, assume inductively that we have chosen a sequence $T_0, \dots, T_{n-1}$ of rooted trees such that some bad sequence of rooted trees starts with this sequence. Choose as T_n a minimum-order T_n

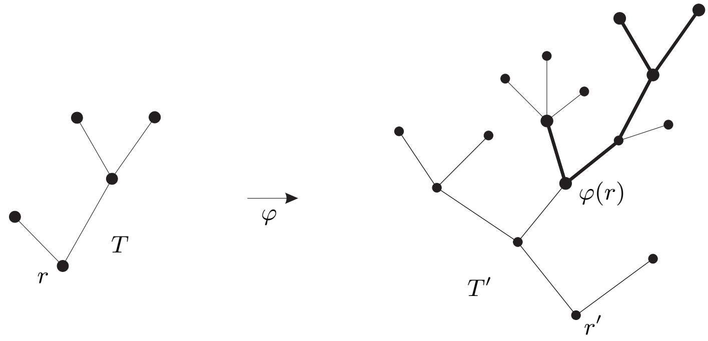

Fig. 12.2.1. An embedding of T in T' showing that $T \leqslant T'$

rooted tree such that some bad sequence starts with $T_0, \dots, T_n$. For each $n \in \mathbb{N}$, denote the root of T_n by r_n. (r_n)

(A_n) Clearly, $(T_n)_{n \in \mathbb{N}}$ is a bad sequence. For each n, let A_n denote the set of components of $T_n - r_n$, made into rooted trees by choosing the neighbours of r_n as their roots. Note that the tree-order of these trees is that induced by T_n. Let us prove that the set $A := \bigcup_{n \in \mathbb{N}} A_n$ of all these trees is well-quasi-ordered. (A)

(T^k, $n(k)$) Let $(T^k)_{k \in \mathbb{N}}$ be any sequence of trees in A. For every $k \in \mathbb{N}$ choose an $n = n(k)$ such that $T^k \in A_n$. Pick a k with smallest $n(k)$. Then

$$T_0, \dots, T_{n(k)-1}, T^k, T^{k+1}, \dots$$

is a good sequence, by the minimal choice of $T_{n(k)}$ and $T^k \subsetneq T_{n(k)}$. Let (T, T') be a good pair of this sequence. Since $(T_n)_{n \in \mathbb{N}}$ is bad, T cannot be among the first $n(k)$ members $T_0, \dots, T_{n(k)-1}$ of our sequence: then T' would be some T^i with $i \geqslant k$, i.e.

$$T \leqslant T' = T^i \leqslant T_{n(i)}\,;$$

since $n(k) \leqslant n(i)$ by the choice of k, this would make $(T, T_{n(i)})$ a good pair in the bad sequence $(T_n)_{n \in \mathbb{N}}$. Hence (T, T') is a good pair also in $(T^k)_{k \in \mathbb{N}}$, completing the proof that A is well-quasi-ordered.

(i, j) By Lemma 12.1.3,[1] the sequence $(A_n)_{n \in \mathbb{N}}$ in $[A]^{<\omega}$ has a good pair (A_i, A_j); let $f\colon A_i \to A_j$ be injective with $T \leqslant f(T)$ for all $T \in A_i$. Now extend the union of the embeddings $T \to f(T)$ to a map φ from $V(T_i)$ to $V(T_j)$ by letting $\varphi(r_i) := r_j$. This map φ preserves the tree-order of T_i, and it defines an embedding to show that $T_i \leqslant T_j$, since the edges $r_i r \in T_i$ map naturally to the paths $r_j T_j \varphi(r)$. Hence (T_i, T_j) is a good pair in our original bad sequence of rooted trees, a contradiction. □

[1] Any readers worried that we might need the lemma for sequences or multisets rather than just sets here, note that isomorphic elements of A_n are not identified: we always have $|A_n| = d(r_n)$.

12.3 Tree-decompositions

Trees are graphs with some very distinctive and fundamental properties; consider Theorem 1.5.1, 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 *bags* $V_t \subseteq V(G)$ indexed by the nodes t of T. The pair $(T, \mathcal{V})$ is called a *tree-decomposition* of G if it satisfies the following three conditions: *bags* *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 the *parts* of $(T, \mathcal{V})$, and say that $(T, \mathcal{V})$ is a tree-decomposition of G *into* the bags V_t or the parts $G[V_t]$. Condition (T3) implies that the bags of $(T, \mathcal{V})$ are organized roughly like a tree (Fig. 12.3.1). *parts*

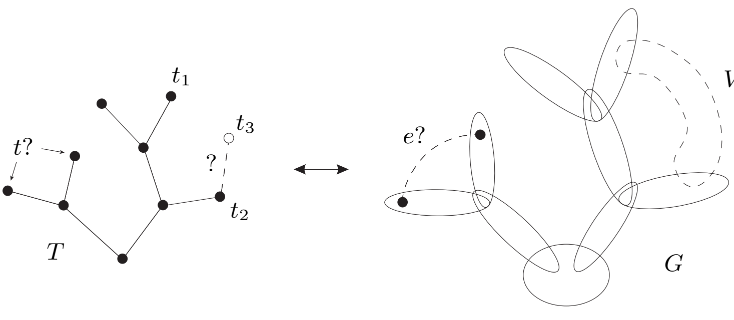

Fig. 12.3.1. Edges and bags ruled out by (T2) and (T3)

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.

[12.4.3]
[12.5]
[12.6.5]

Lemma 12.3.1. *Let t_1t_2 be any edge of T and let T_1, T_2 be the components of $T - t_1t_2$, with $t_1 \in T_1$ and $t_2 \in T_2$. Then $V_{t_1} \cap V_{t_2}$ separates $U_1 := \bigcup_{t \in T_1} V_t$ from $U_2 := \bigcup_{t \in T_2} V_t$ in G (Fig. 12.3.2).*

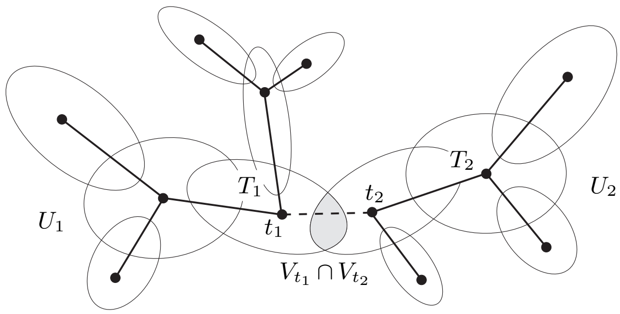

Fig. 12.3.2. $V_{t_1} \cap V_{t_2}$ separates U_1 from U_2 in G

Proof. Consider a shortest U_1–U_2 path. By (T1) it has length at most 1, so by (T2) it has length 0. By (T3), its vertex lies in $V_{t_1} \cap V_{t_2}$. □

induced separation

We shall say that the separation $\{U_1, U_2\}$ of G in Lemma 12.3.1 is *induced* by the edge t_1t_2 of T, or more generally by $(T, \mathcal{V})$. Its separator $U_1 \cap U_2 = V_{t_1} \cap V_{t_2}$ is the *adhesion set* of V_{t_1} and V_{t_2}.

adhesion

torsos

The *adhesion* of a tree-decomposition is the maximum size of its adhesion sets. (If T is trivial, we let it be zero.) The *torsos* of a tree-decomposition are the supergraphs of its parts $G[V_t]$ obtained by making their adhesion sets complete: by adding to $G[V_t]$ any edges not in G whose ends lie in a common adhesion set $V_t \cap V_{t'}$ with $tt' \in E(T)$.

Tree-decompositions are passed on to subgraphs, indeed to minors:

[12.4.1]
[12.4.4]
[12.4.3]
[12.6.2]

Lemma 12.3.2. *For every $H \subseteq G$, the pair $\big(T, (V_t \cap V(H))_{t \in T}\big)$ is a tree-decomposition of H.* □

[12.4.1]
[12.4.3]

Lemma 12.3.3. *Suppose that $G = IH$ with branch sets U_h, $h \in V(H)$. Let $f\colon V(G) \to V(H)$ be the map assigning to each vertex of G the index of the branch set containing it. For each $t \in T$ let $W_t := \{\, f(v) \mid v \in V_t \,\}$, and put $\mathcal{W} := (W_t)_{t \in T}$. Then $(T, \mathcal{W})$ is a tree-decomposition of H.*

Proof. The assertions (T1) and (T2) for $(T, \mathcal{W})$ follow immediately from the corresponding assertions for $(T, \mathcal{V})$. Now let $t_1, t_2, t_3 \in T$ be as in (T3), and consider a vertex $h \in W_{t_1} \cap W_{t_3}$ of H; we show that $h \in W_{t_2}$. By definition of W_{t_1} and W_{t_3}, the set U_h meets both V_{t_1} and V_{t_3}. As U_h is connected in G, this implies by Lemma 12.3.1 that U_h also meets V_{t_2}. Hence $h \in W_{t_2}$, by definition of W_{t_2}. □

Here is another useful consequence of Lemma 12.3.1:

Lemma 12.3.4. *Any set of vertices not contained in a bag of $(T, \mathcal{V})$ contains two vertices that are separated by an adhesion set of $(T, \mathcal{V})$.* [12.4.3]

Proof. Given $W \subseteq V(G)$, 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. Unless $V_{t_1} \cap V_{t_2}$ separates two vertices from W, 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 node of a maximal directed path in T; then all the edges of T at t are oriented towards 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:

Corollary 12.3.5. *Every complete subgraph of G is contained in some part of $(T, \mathcal{V})$.* □ [12.6.2]

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 (Exercise 19). The same applies to the property of not containing a K^r minor for some fixed r. *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 7.3.4, any graph without a K^5 minor has a supergraph with a simplicial tree-decomposition into plane triangulations and copies of the Wagner graph W, and similarly for graphs without K^4 minors (see Proposition 12.6.2).

Tree-decompositions may thus lead to intuitive structural characterizations of graph properties. A particularly simple example is the following characterization of chordal graphs:

Proposition 12.3.6. *G is chordal if and only if G has a tree-decomposition into complete parts.* [12.4.4] [12.6.2]

Proof. We apply induction on $|G|$. We first assume that G has a tree-decomposition $(T, \mathcal{V})$ such that $G[V_t]$ is complete for every $t \in T$; let us choose $(T, \mathcal{V})$ with $|T|$ minimum. If $|T| \leqslant 1$, then G is complete and hence chordal. So let $t_1t_2 \in T$ be an edge, and for $i = 1, 2$ define T_i and $G_i := G[U_i]$ as in Lemma 12.3.1. Then $G = G_1 \cup G_2$ by (T1) and (T2), and $V(G_1 \cap G_2) = V_{t_1} \cap V_{t_2}$ by the lemma; thus, $G_1 \cap G_2$ is complete. Since $(T_i, (V_t)_{t \in T_i})$ is a tree-decomposition of G_i into complete (5.5.1) parts, both G_i are chordal by the induction hypothesis. (By the choice of $(T, \mathcal{V})$, neither G_i is a subgraph of $G[V_{t_1} \cap V_{t_2}] = G_1 \cap G_2$, so both G_i are indeed smaller than G.) Since $G_1 \cap G_2$ is complete, any induced cycle in G lies in G_1 or in G_2 and hence has a chord, so G too is chordal.

Conversely, assume that G is chordal. If G is complete, there is nothing to show. If not then, by Proposition 5.5.1, G is the union of smaller chordal graphs G_1, G_2 with $G_1 \cap G_2$ complete. By the induction hypothesis, G_1 and G_2 have tree-decompositions $(T_1, \mathcal{V}_1)$ and $(T_2, \mathcal{V}_2)$ into complete parts. By Corollary 12.3.5, $G_1 \cap G_2$ lies inside one of those parts in each case, say with indices $t_1 \in T_1$ and $t_2 \in T_2$. As one easily checks, $((T_1 \cup T_2) + t_1 t_2, \mathcal{V}_1 \cup \mathcal{V}_2)$ is a tree-decomposition of G into complete parts. □

Let us wind up this section with an application of tree-decompositions to connectivity that generalizes the idea of the block-cutvertex tree from Lemma 3.1.4. A set $U \subseteq V(G)$ of at least k vertices is $(< k)$-*inseparable* in G if no two vertices from U can be separated in G by fewer than k other vertices (which may or may not lie in U).

k-block A maximal $(< k)$-inseparable set of vertices is a k-*block*. Thus, the 1-blocks of a graph are its components; its 2-blocks are the non-singleton vertex sets spanning a block as defined in Chapter 3. In general, a k-block need not induce a highly connected subgraph: the many paths between its vertices that are needed to make it $(< k)$-inseparable can all lie outside it. Its 'connectivity' is thus measured in the ambient graph; its vertices themselves may even be independent.

Theorem 12.3.7. *For every integer $k \geqslant 1$, every graph G has a tree-decomposition $(T, \mathcal{V})$ with the following properties:*

(i) *$(T, \mathcal{V})$ has adhesion $< k$.*

(ii) *Distinct k-blocks lie in different bags. Moreover, every two blocks are separated by an adhesion set that is no larger than the smallest set of vertices that separates them in G.*

(iii) *Every automorphism of G acts on the set of bags of $(T, \mathcal{V})$, and the action on $V(T)$ which this induces is an automorphism of T.*

Note that, by (i) and Lemma 12.3.4, every k-block is contained in some bag. This is unique by (i) and (T3), so the bags in (ii) are well defined. Assertion (iii) means that, for every automorphism φ of G and every $t \in T$, the set $\varphi(V_t)$ is another bag $V_{t'}$ (possibly V_t), and this map $\varphi \colon t \mapsto t'$ is an automorphism of T.[3] Such tree-decompositions are called *canonical*.

[3] This can be important both theoretically and in applications. For example, when a computer constructs a tree-decomposition, then this may depend on how we present the graph to it as input. But if it finds a canonical decomposition, then this will always be the same up to automorphism, in the sense just defined.

12.4 Tree-width

As indicated by Figure 12.3.1, the bags 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 bags, 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* $\mathrm{tw}(G)$ of G is the least width of any tree-decomposition of G. As one easily checks,[4] trees themselves have tree-width 1. *tree-width* $\mathrm{tw}(G)$

By Lemmas 12.3.2 and 12.3.3, the tree-width of a graph will never be increased by deletion or contraction: (12.3.2) (12.3.3)

Lemma 12.4.1. *If* $H \preccurlyeq G$ *then* $\mathrm{tw}(H) \leqslant \mathrm{tw}(G)$. □ [12.6.2]

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 $\mathrm{tw}(G)$ times. This takes us a step further towards a proof of the graph minor theorem:

Theorem 12.4.2. (Robertson & Seymour 1990) [12.7.1] [12.7.3]
For every integer $k > 0$*, the graphs of tree-width* $< k$ *are well-quasi-ordered by the minor relation.*

In order to make use of Theorem 12.4.2 for a proof of the full graph minor theorem, we should be able to say something about the graphs which it does not cover, i.e., to deduce some information about a graph from the assumption that its tree-width is large. Our next result, the *tree-width duality theorem*, achieves just that: it identifies a canonical obstruction to small tree-width, a structural phenomenon that occurs in a graph if and only if its tree-width is large.

Let us say that two subsets of $V(G)$ *touch* if they have a vertex in common or G contains an edge between them. A set of mutually touching connected vertex sets in G is a *bramble*. Extending our terminology of Chapter 2, we say that a subset of $V(G)$ *covers* (or is a *cover* of) a bramble $\mathcal{B}$ if it meets every element of $\mathcal{B}$. The least number of vertices covering a bramble is the *order* of that bramble. A *k-bramble* is one of order k. *touch* *bramble* *cover* *order*

[4] Indeed the '-1' in the definition of width serves no other purpose than to make this statement true.

A typical example of a bramble is the set of crosses in a grid. The
grid $k \times k$ *grid* is the graph on $\{1,\dots,k\}^2$ with the edge set

$$\{\,(i,j)(i',j') : |i-i'|+|j-j'| = 1\,\}\,.$$

The *crosses* of this grid are the k^2 sets

$$C_{ij} := \{\,(i,\ell) \mid \ell = 1,\dots,k\,\} \cup \{\,(\ell,j) \mid \ell = 1,\dots,k\,\}\,.$$

Thus, the cross C_{ij} is the union of the grid's ith row and its jth column. Clearly, the crosses of the $k \times k$ grid form a k-bramble: they are covered by any row or column, while any set of fewer than k vertices misses both a row and a column, and hence a cross.

Theorem 12.4.3. (Seymour & Thomas 1993)
Let $k \geqslant 0$ be an integer. A graph has tree-width $< k$ if and only if it contains no bramble of order $> k$.

(3.3.1) (12.3.1) (12.3.2) (12.3.3)

Proof. Let $G = (V,E)$ be a graph. For the forward implication, let $\mathcal{B}$ be any bramble in G. We show that every tree-decomposition $(T,(V_t)_{t\in T})$ of G has a bag that covers $\mathcal{B}$.

As in the proof of Lemma 12.3.4 we start by orienting the edges t_1t_2 of T. If $X := V_{t_1} \cap V_{t_2}$ covers $\mathcal{B}$, we are done. If not, then for each $B \in \mathcal{B}$ 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. 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, it has a node t all whose incident edges are oriented towards it. Then V_t covers $\mathcal{B}$ – just as in the proof of Lemma 12.3.4.

For our proof of the converse implication we need the notion of a
good *good* tree-decomposition $(T,\mathcal{V})$ of G, with $\mathcal{V} = (V_t)_{t\in T}$ say: one where $|V_t| \leqslant k$ for at least one t and $|V_t| \leqslant k$ whenever t is not a leaf of T. Then $|V_t| \leqslant k$ for the unique neighbour t of any leaf x with $|V_x| > k$.
petal The set $V_x \setminus V_t \neq \emptyset$ is the *petal* of such a leaf x with $|V_x| > k$.

Suppose now that $\mathrm{tw}(G) \geqslant k$, and let us find a bramble of order $> k$. Note that G has a good tree-decomposition – for example, with $T = K^2$ and parts $\emptyset, V$. And since $\mathrm{tw}(G) \geqslant k$, every good tree-decomposition of G has a petal. Let $\mathcal{B}$ be a minimal set of petals of good tree-decompositions satisfying

(i) $\mathcal{B}$ contains a petal of every good tree-decomposition;

(ii) $\mathcal{B}$ is closed under taking supersets among petals: if $X \subseteq X'$ are both petals of good tree-decompositions and $X \in \mathcal{B}$, then $X' \in \mathcal{B}$.

Let $\mathcal{B}' := \{\,X \in \mathcal{B} \mid G[X] \text{ is connected}\,\}$. To complete our proof we show that $\mathcal{B}'$ is a bramble of order $> k$.

Let us show first that no set Z of at most k vertices covers $\mathcal{B}'$. As $\operatorname{tw}(G) \geqslant k$ we have $V \setminus Z \neq \emptyset$, and if $C_1, \dots, C_n$ are the components of $G - Z$, then G has a tree-decomposition into Z and the sets $V(C_i) \cup Z$, whose decomposition tree is a star with Z as its central bag. This is a good tree-decomposition, so by (i) it has a petal in $\mathcal{B}$. But any such petal is one of the C_i, so it is connected and thus lies in $\mathcal{B}'$. Hence Z fails to cover $\mathcal{B}'$, as claimed.

It remains to show that every two sets in $\mathcal{B}'$ touch. In fact, let us show that every two sets in $\mathcal{B}$ touch. If not, we can find $X, Y \in \mathcal{B}$ that do not touch and are $\subseteq$-minimal in $\mathcal{B}$. Since $\mathcal{B} \setminus \{X\}$ and $\mathcal{B} \setminus \{Y\}$ still satisfy (ii), the minimality of $\mathcal{B}$ implies that they violate (i). So there exists a good tree-decomposition $(T_1, \mathcal{V}_1)$ whose only petal in $\mathcal{B}$ is X, and a good tree-decomposition $(T_2, \mathcal{V}_2)$ whose only petal in $\mathcal{B}$ is Y. Let us choose T_1 and T_2 disjoint.

Let X be the petal of $x \in T_1$, and Y the petal of $y \in T_2$. By Lemma 12.3.1, x is unique in T_1 and y is unique in T_2. Deleting any vertices outside $N(X)$ and $N(Y)$ from the adhesion sets of x and y, respectively, we may assume that their bags are exactly $X \cup N(X)$ and $Y \cup N(Y)$. (The trimmed bags still have size $> k$, since otherwise the modified tree-decomposition would have no petal in $\mathcal{B}$.) Since any petal in $(T_1, \mathcal{V}_1)$ containing Y would lie in $\mathcal{B} \setminus \{X\}$, by (ii), there is no such petal. Similarly, no petal in $(T_2, \mathcal{V}_2)$ contains X.

To complete our proof we show the following:

There is a good tree-decomposition $(T, \mathcal{V})$ all whose petals are contained in petals of $(T_1, \mathcal{V}_1)$ and $(T_2, \mathcal{V}_2)$ and in which neither X nor Y is a petal. $(*)$

Since X and Y are the only petals of $(T_1, \mathcal{V}_1)$ and $(T_2, \mathcal{V}_2)$ in $\mathcal{B}$ and have no proper subsets in $\mathcal{B}$, the $(T, \mathcal{V})$ from $(*)$ has no petal in $\mathcal{B}$, by (ii) for $\mathcal{B}$. As this contradicts (i), our proof will then be complete.

So let us prove $(*)$. As X and Y do not touch, the set $N(X)$ is disjoint from both X and Y and separates them in G. Hence G has a separation $\{A, B\}$ such that $X \subseteq A \setminus B$ and $Y \subseteq B \setminus A$. As A, B
$|N(X)| \leqslant k$ since X is a petal, choosing $\{A, B\}$ of minimum order ensures that $S := A \cap B$ has size at most k. By the minimality of S S
and Menger's Theorem 3.3.1, there is a family $\{\, P_s \mid s \in S \,\}$ of disjoint $P_s,\ Q_s$
S–$N(X)$ paths in $G[A]$ and a family $\{\, Q_s \mid s \in S \,\}$ of disjoint S–$N(Y)$ paths in $G[B]$.

Let H be the minor of G obtained by deleting $A \setminus \bigcup_{s \in S} V(P_s)$ H
and contracting each of the paths P_s. Identifying the contracted branch sets $V(P_s)$ with their representatives s, we may think of H as obtained from $G[B]$ by adding some edges on S. Let $(T_1, \mathcal{V}_1')$ be the tree-decomposition which $(T_1, \mathcal{V}_1)$ induces on H as in Lemmas 12.3.2 and 12.3.3, and think of it as a tree-decomposition of $G[B]$. Thus for any $t \in T_1$, with a

bag $V_t^1 \in \mathcal{V}_1$, say, its bag V_t in $\mathcal{V}_1'$ is

$$V_t = (V_t^1 \cap B) \cup \{\, s \in S \mid V_t^1 \cap V(P_s) \neq \emptyset \,\} \tag{1}$$

(Fig. 12.4.1). In particular, $V_x = S$, since $V_x = X \cup N(X) \subseteq A$ and $N(X)$ meets every P_s. Similarly, let J be the minor of G obtained by deleting $B \setminus \bigcup_{s \in S} V(Q_s)$ and contracting the paths Q_s, and let $(T_2, \mathcal{V}_2')$ be the tree-decomposition which $(T_2, \mathcal{V}_2)$ induces on J. As before, think of this as a tree-decomposition of $G[A]$ in which S is the bag corresponding to y.

J

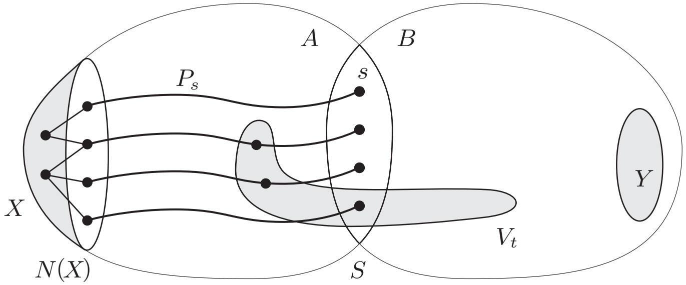

Fig. 12.4.1. To obtain V_t from $V_t^1 \cap B$, add two vertices from S

T Let T be obtained from the (disjoint) trees T_1 and T_2 by identifying x and y into a new node r. As Y and X are non-empty, x is not the only node of T_1 and y is not the only node of T_2, so r is not a leaf of T. Let $V_r := S$. For all $t \in T - r$ let V_t be the bag in $\mathcal{V}_1'$ or $\mathcal{V}_2'$ that corresponds to t there, thought of as a subset of B if $t \in T_1$, or of A if $t \in T_2$. We claim that $(T, \mathcal{V})$ with $\mathcal{V} = (V_t)_{t \in T}$ is a good tree-decomposition of G satisfying $(*)$.

V_t

$(T, \mathcal{V})$

Using that $(T_1, \mathcal{V}_1')$ and $(T_2, \mathcal{V}_2')$ are tree-decompositions of $G[B]$ and $G[A]$, it is easy to check that $(T, \mathcal{V})$ is indeed a tree-decomposition of G. The non-leaves of T are precisely those of T_1 and T_2, plus r. We have already seen that $|S| \leqslant k$. For $t \in T_1 - x$, its bag V_t in $\mathcal{V}$ is no larger than its bag V_t^1 in $\mathcal{V}_1$: by (1), there exists for every $s \in V_t \setminus V_t^1$ a vertex of P_s in $V_t^1 \setminus V_t$. Similarly, the bags V_t with $t \in T_2$ are no larger than their corresponding bags in $\mathcal{V}_2$. Thus, $(T, \mathcal{V})$ is good.

Z To show that $(T, \mathcal{V})$ satisfies $(*)$, consider a petal Z in $(T, \mathcal{V})$ of a leaf z of T. Then z is a leaf also of T_1 or T_2. Let us assume that $z \in T_1$; then $V_z \subseteq B$. By axiom (T3) for $(T, \mathcal{V})$, any vertex of $V_z \cap V_r$ lies in the adhesion set of V_z, so $Z \cap S = \emptyset$. Hence $Z \subseteq B \setminus A$. But this implies that Z lies inside the petal Z_1 of z in $(T_1, \mathcal{V}_1)$: recall that the bag V_z^1 corresponding to z in $\mathcal{V}_1$ has size at least $|V_z| > k$ (so it can have a petal) and differs from V_z only by vertices in A, and likewise for the neighbour of z in T_1 and T. Finally, $Z \neq X$ because $X \subseteq A \setminus B$, and $Z \neq Y$ since the petal $Z_1 \supseteq Z$ of z in $(T_1, \mathcal{V}_1)$ does not contain Y by assumption. This completes the proof of $(*)$, and hence of the theorem. $\square$

Often, Theorem 12.4.3 is stated in terms of the *bramble number* of a graph, the largest order of any bramble in it. The theorem then says that the tree-width of a graph is exactly one less than its bramble number.

How useful even the easy forward implication of Theorem 12.4.3 can be is exemplified once more by our example of the crosses bramble in the $k \times k$ grid: this bramble has order k, so by the theorem the $k \times k$ grid has tree-width at least $k - 1$. (Try to show this without the theorem!)

In fact, the $k \times k$ grid has tree-width k (Exercise 34). But more important than its precise value is the fact that the tree-width of grids tends to infinity with their size. For as we shall see, large grid minors pose another canonical obstruction to small tree-width: not only do large grids (and hence all graphs containing large grids as minors; cf. Lemma 12.4.1) have large tree-width, but conversely every graph of large tree-width has a large grid minor (Theorem 12.6.3).

In Section 12.5 we shall place these within the wider framework of *tangles*, another central notion in graph minor theory. Using tangles one can formulate a more general duality theory between highly connected substructures and trees, of which Theorem 12.4.3 is but a special case.

Tree-width can also be expressed as follows:

Proposition 12.4.4. $\mathrm{tw}(G) = \min\left\{\, \omega(H) - 1 \mid G \subseteq H;\ H \text{ chordal} \,\right\}$.

(12.3.2)
(12.3.5)
(12.3.6)

Proof. By Corollary 12.3.5 and Proposition 12.3.6, each of the graphs H considered for the minimum has a tree-decomposition of width $\omega(H) - 1$. Every such tree-decomposition induced one of G by Lemma 12.3.2, so $\mathrm{tw}(G) \leqslant \omega(H) - 1$ for every H.

Conversely, let us construct an H as above with $\omega(H) - 1 \leqslant \mathrm{tw}(G)$. Let $(T, \mathcal{V})$ be a tree-decomposition of G of width $\mathrm{tw}(G)$. For every $t \in T$ let K_t denote the complete graph on V_t, and put $H := \bigcup_{t \in T} K_t$. Clearly, $(T, \mathcal{V})$ is also a tree-decomposition of H. By Proposition 12.3.6, H is chordal, and by Corollary 12.3.5, $\omega(H) - 1$ is at most the width of $(T, \mathcal{V})$, i.e. at most $\mathrm{tw}(G)$. □

A tree-decomposition $(T, \mathcal{V})$ of G with $\mathcal{V} = (V_t)_{t \in T}$ is *linked*, or *lean*, if it satisfies the following condition: *linked/lean*

(T4) Given $t_1, t_2 \in T$ and vertex sets $Z_1 \subseteq V_{t_1}$ and $Z_2 \subseteq V_{t_2}$ such that $|Z_1| = |Z_2| =: k$, either G contains k disjoint Z_1–Z_2 paths or there exists an edge $tt' \in t_1 T t_2$ with $|V_t \cap V_{t'}| < k$.

The 'branches' in a lean tree-decomposition are thus stripped of any bulk not necessary to maintain their connecting qualities. Indeed, if a branch is thick (i.e. the adhesion sets $V_t \cap V_{t'}$ along a path in T are all large), then G is highly connected along this branch, and the bags themselves are no larger than their 'external connectivity' in G requires: for $t_1 = t_2$, (T4) says that two k-sets of vertices will only lie in a common bag if they cannot be separated (in G) by fewer than k vertices.

In our quest for tree-decompositions into 'small' bags, we now have two criteria to choose between: the global 'worst case' criterion of width, and the more subtle local criterion of leanness. Surprisingly, it is always possible to find a tree-decomposition that is optimal with respect to both criteria at once:

Theorem 12.4.5. (Thomas 1990)
Every graph G has a lean tree-decomposition of width $\mathrm{tw}(G)$.

Another natural feature one might ask of a tree-decomposition is that its parts, as induced subgraphs of G, be connected. Let us call such a tree-decomposition *connected*, and let the *connected tree-width* $\mathrm{ctw}(G)$ of G be the least width of a connected tree-decomposition of G.

The connected tree-width of most graphs is greater than their ordinary tree-width. For example, every cycle has tree-width 2, but $\mathrm{ctw}(C^n) = \lceil n/2 \rceil$ (Exercise 33). And unlike ordinary tree-width, the connected tree-width of a subgraph of G can be greater than that of G. (For example, let G be obtained from a long cycle by adding a chord between opposite vertices.) However, if $C \subseteq G$ is a *geodesic* cycle, i.e., if $d_C(u,v) = d_G(u,v)$ for all vertices $u, v \in C$, then $\mathrm{ctw}(C) \leqslant \mathrm{ctw}(G)$.

The presence of long geodesic cycles in a graph thus is an obstruction to small connected tree-width – as is, trivially, large ordinary tree-width. By the following theorem, however, these are the only obstructions:

Theorem 12.4.6. *There is a function $f\colon \mathbb{N}^2 \to \mathbb{N}$ such that every graph of tree-width $\leqslant w \in \mathbb{N}$ that has no geodesic cycle of length $> k \in \mathbb{N}$ has connected tree-width at most $f(w,k)$.*

12.5 Tangles

We have already in this chapter met a few types of substructures of possibly sparse graphs that are highly connected in some sense, but which are not just k-connected subgraphs (or minors) for some large k. Large grid minors are an example, the k-blocks defined at the end of Section 12.3 are another, as are brambles of high order. All these substructures have one feature in common: for every low-order separation of the graph they lie essentially, though not necessarily entirely, on one of its two sides.

$G = (V, E)$ For example, given a bramble $\mathcal{B}$ of order k in a graph $G = (V, E)$, if $\{A, B\}$ is a separation of order $< k$ then $A \setminus B$ or $B \setminus A$, but not both, contains a set from $\mathcal{B}$. This helped us prove the easy implication of the tree-width duality theorem: $\mathcal{B}$ orients every edge of the decomposition tree T of any tree-decomposition $(T, \mathcal{V})$ of adhesion $< k$ 'towards' the side of its induced separation that contains one of the sets in $\mathcal{B}$, and the edges thus oriented point to a central node t of T for which V_t covers $\mathcal{B}$.

It has turned out that, more often than not, the only feature of a highly connected substructure that mattered in the proofs of theorems about it was the information of how it 'orients' the low-order separations of G in this way. Collecting just this information together leads to a more abstract notion of a highly connected substructure, called a *tangle*. The purpose of this section is to make this precise, to prove a duality theorem for tangles in the spirit of Theorem 12.4.3, and to point out how this setting can be used to express the duality between global tree-structure of a graph and its highly connected substructures more generally.

In the context of tangles, we often denote the order $|A \cap B|$ of a separation $s = \{A, B\}$ simply as $|s|$. The *orientations* of $\{A, B\}$ are the two *oriented separations* (A, B) and (B, A). We say that (A, B) is oriented, or 'points', *towards* B and its subsets. Given oriented separations (A, B) and (C, D), we write $(A, B) \geqslant (C, D)$ if $A \subseteq C$ and $B \supseteq D$. $|s|$ *oriented separation* $\leqslant$

Two separations of G are *nested* if they have $\leqslant$-comparable orientations; otherwise they *cross*. The separation $\{E, F\}$ in Figure 12.5.1, for example, is nested with the two crossing separations $\{A, B\}$ and $\{C, D\}$, because $(E, F) \geqslant (A, B), (C, D)$. The separations induced by a tree-decomposition are clearly nested, and every nested set of separations of a graph is induced by some tree-decomposition; see after Lemma 12.3.1. *nested* *cross*

Given two separations $\{A, B\}$ and $\{C, D\}$ of G, it is easy to check that also $\{A \cup C, B \cap D\}$ is a separation of G. We call the four separations of this form the *corners* of the separations $\{A, B\}$ and $\{C, D\}$ (Fig. 12.5.1). They come in pairs of *opposite* corners: the corner opposite to $\{A \cup D, B \cap C\}$, for example, is $\{B \cup C, A \cap D\}$. Two corners that are not opposite are *adjacent*. The adjacent corners $\{A \cup C, B \cap D\}$ and $\{A \cup D, B \cap C\}$ *lie on* the same side of $\{A, B\}$, on its *B-side*. *corners* *opposite* *adjacent*

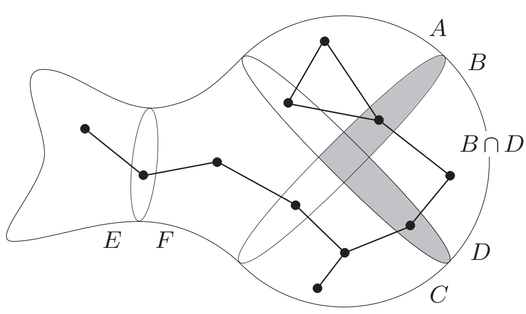

Fig. 12.5.1. The separation $\{A \cup C, B \cap D\}$ is one of the four corners of $\{A, B\}$ and $\{C, D\}$. The separation $\{E, F\}$ is nested with $\{A, B\}$ and $\{C, D\}$ and all their corners.

The sum of the orders of a pair of opposite corners always equals the sum of the orders of the original two crossing separations. For example:

$$|(A \cup D) \cap (B \cap C)| + |(B \cup C) \cap (A \cap D)| = |A \cap B| + |C \cap D|.$$

The $\leqslant$-part of this equality, the fact that opposite corners c, d of any two separations r, s satisfy

submodularity

$$|c| + |d| \leqslant |r| + |s|\,, \tag{†}$$

is sometimes referred to as *submodularity* of the order function $s \mapsto |s|$. It implies, for example, that of any two opposite corners of two separations of order $< k$ at least one also has order $< k$. More generally, (†) implies that $|c| \leqslant |r|$ or $|d| \leqslant |s|$ (or both), which is the standard way to apply it.

$\vec{S}$

Given a set S of separations, we write $\vec{S} := \{\,(A,B) \mid \{A,B\} \in S\,\}$ for the set of all their orientations. An *orientation of* S is a subset of $\vec{S}$ that contains for every element of S exactly one of its two orientations. Unless otherwise mentioned, from now on in this section S will be the set of *all* the separations of G, so $\vec{S}$ is the set of all its oriented separations.

S

consistent

A set σ of oriented separations of G is *consistent* if it does *not* contain (B,A) whenever $(A,B) \geqslant (C,D)$ with $(C,D) \in \sigma$.[5]

For example, if $(T, \mathcal{V})$ is a tree-decomposition of G with $\mathcal{V} = (V_t)_{t \in T}$, and V_t is not contained in any other bag, then orienting the separations that $(T, \mathcal{V})$ induces towards V_t orients them consistently. A k-block X of G even defines a consistent orientation of the entire set

S_k

$$S_k := \{\, s \in S : |s| < k \,\}$$

of separations of order $< k$ of G: the orientation $\{\,(A,B) \in \vec{S_k} \mid X \subseteq B\,\}$.

A bramble $\mathcal{B}$ of order $n \geqslant k$ also defines an orientation of S_k: the set $O = \{\,(A,B) \in \vec{S_k} \mid \exists X \in \mathcal{B}\colon X \subseteq B \setminus A\,\}$. Unlike in the case of k-blocks, there need not be any fixed bramble set that lies in every B with $(A,B) \in O$; indeed the intersection of all these B may be empty. (Example?) But O then shows the large order of $\mathcal{B}$ in another way, in that we cannot cover $\mathcal{B}$ by few sets A with $(A,B) \in O$: since any bramble set meeting A also meets $A \cap B$ (as it touches the bramble set in $B \setminus A$), and $|A \cap B| \leqslant k-1$, we need at least $n/(k-1)$ sets A to cover $\mathcal{B}$. The idea that this is enough not only to reflect but to constitute a kind of highly connected substructure in G has led to the following concept of a tangle.

avoids

tangle

An orientation O of a set of separations *avoids* a set $\mathcal{F}$ of subsets of $\vec{S}$ if no subset of O is an element of $\mathcal{F}$. A *tangle of order* k in G, or k-*tangle* for short, is an orientation of S_k that avoids

$\mathcal{T}$

$$\mathcal{T} := \big\{\{(A_1,B_1),(A_2,B_2),(A_3,B_3)\}\,\big|\,G[A_1] \cup G[A_2] \cup G[A_3] = G\,\big\}.$$

---

[5] Intuitively, σ is consistent if no two of its elements point away from each other. In particular, it will not contain both orientations of any given separation.

Note that the separations (A_i, B_i) in the definition of $\mathcal{T}$ need not be distinct. For example, no tangle contains a separation of the form (V, A). And if $(A,B) \geqslant (C,D)$ then no tangle contains both (B,A) and (C,D), since $G[B] \cup G[C] \supseteq G[D] \cup G[C] = G$. Thus, tangles are consistent. More generally, tangles τ clearly have the following two properties:

If τ contains the oriented separations (A,B) and (C,D), it does not *contain their oriented corner $(B \cap D, A \cup C)$.* (P)

Given separations $\{A,B\}$ and $\{C,D\}$, if τ contains (A,B) it does not *contain both $(B \cap D, A \cup C)$ and $(B \cap C, A \cup D)$.* (R)

Sets of oriented separations satifying (P) are known as *profiles*, and (P) is the *profile property* of tangles.[6] Note that it can be satisfied in two ways: either by $(A \cup C, B \cap D) \in \tau$, or by the fact that τ does not orient the corner $\{A \cup C, B \cap D\}$ at all; this happens if that corner has order k or larger, but τ is only a k-tangle. Profiles satisfying (R) are *robust*. *profile*

Like its k-blocks, compare Theorem 12.3.7, the tangles of a graph can be 'separated' by a tree-decomposition. Since the tangles in a graph see it in terms of its separations rather than its subgraphs, this is naturally expressed not in terms of the bags of the decomposition but of the separations it induces.

Let us say that a separation *distinguishes* two tangles, not necessarily of the same order, if they both orient it but do so differently. Distinct tangles of the same order are always distinguished by some separation; otherwise they would be identical. But a k-tangle cannot be distinguished from the ℓ-tangle it *induces* for $\ell < k$, the set of those of its oriented separations that have order $< \ell$. *distinguish* *induced tangle*

A separation that distinguishes two tangles does so *efficiently* if they are not distinguished by any separation of smaller order. A set T of separations *distinguishes* some *set* of tangles in a graph G *efficiently* if every two tangles in this set that are distinguished by some separation of G are distinguished efficiently by a separation in T. *efficient*

Theorem 12.5.1. (Robertson & Seymour 1991) *tree-of-tangles theorem*
Every graph G has a nested set of separations that distinguishes all the tangles in G efficiently.

It is not hard to construct from any nested set of separations, such as that in Theorem 12.5.1, a tree-decomposition that induces precisely this set of separations as defined after Lemma 12.3.1 (Exercises 16 and 51).

---

[6] If we think of tangles as orienting separations towards their 'big' side, condition (P) is reminiscent of ultrafilters: the intersection of two big sets cannot be small.

We shall give two proofs for Theorem 12.5.1. The first proof is based on a combinatorial lemma, the *splinter lemma*, which encapsulates the essence of earlier direct proofs that sought to find the required nested set of tangle-distinguishing separations directly among the corners of arbitrary efficient tangle-distinguishing separations.

Our second proof of Theorem 12.5.1 comes with an important strengthening: the nested set of tangle-distinguishing separations it finds is not constructed by a sequence of choices, as in the first proof, but is identified by an explicitly stated property of separations. As a consequence, this set is invariant under the automorphisms of the graph G. We state this result separately as Theorem 12.5.8.

In both proofs we shall need the following lemma about nested and crossing separations (see Figure 12.5.1):

Lemma 12.5.2. (Fish Lemma)
Any separation r that is nested with two crossing separations s and t is also nested with their four corners.

Proof. Since r is nested with s, the two have orientations $\vec{r} \geqslant \vec{s}$ (Ex. 40). Similarly, r and t have orientations $\vec{r}\,' \geqslant \vec{t}\,'$. If $\vec{r}\,' \neq \vec{r}$, then $\vec{r} \leqslant \vec{t}$ with $\vec{t}$ the inverse of $\vec{t}\,'$. Then $\vec{s} \leqslant \vec{r} \leqslant \vec{t}$, showing that s and t are nested. This contradicts our assumptions, so $\vec{r}\,' = \vec{r} \geqslant \vec{s}, \vec{t}\,'$.

Let $\vec{r} =: (E,F)$ and $\vec{s} =: (A,B)$ and $\vec{t}\,' =: (C,D)$ (Fig. 12.5.1). Then $(E,F) \geqslant (A \cap C, B \cup D)$ as well as, more trivially, $(E,F) \geqslant (A \cup C, B \cap D)$ and $(E,F) \geqslant (A \cup D, B \cap C)$ and $(E,F) \geqslant (C \cup B, A \cap D)$. In particular, r is nested with all four corners of s and t. $\square$

Let us say that a family $(A_1,\dots,A_n)$ of sets of separations of a graph *splinters* if, for all $i \neq j$, any crossing $a_i \in A_i \setminus A_j$ and $a_j \in A_j \setminus A_i$ have a corner in $A_i \cup A_j$.

Lemma 12.5.3. (Splinter Lemma)
Every splintering family $(A_1,\dots,A_n)$ of non-empty sets of separations has a nested family $(a_1,\dots,a_n)$ of representatives $a_i \in A_i$.

Proof. We apply induction on n, which starts trivially with $n = 1$. For
$a_1,\dots,a_{n-1}$ the induction step let $(a_1,\dots,a_{n-1})$ be a nested family of representatives $a_i \in A_i$ of $(A_1,\dots,A_{n-1})$. If any of those a_i lies in A_n, let $a_n := a_i$; then $(a_1,\dots,a_n)$ is as desired. Suppose now that $a_i \in A_i \setminus A_n$ for all
a_n $i < n$, and choose $a_n \in A_n$ so as to cross a_i for as few $i < n$ as possible.

A'_i Let us show that, for all $i < n$, the set A'_i of elements of A_i that are nested with a_n is non-empty. If $a_n \in A_i$ then $a_n \in A'_i$; we may thus assume that $a_n \in A_n \setminus A_i$ (as well as $a_i \in A_i \setminus A_n$). We may also assume that a_i and a_n cross, as otherwise $a_i \in A'_i$. Since $(A_1,\dots,A_n)$ splinters, a_i and a_n have a corner c in $A_i \cup A_n$. By the fish lemma, c is nested with all the a_j that a_n is nested with (since these a_j are also nested with a_i), and with a_i in addition. Hence $c \notin A_n$ by the choice of a_n, and therefore $c \in A_i$. Since c is nested with a_n, this yields $c \in A'_i$. This completes the proof that the A'_i are non-empty.

Let us show that the family $(A'_1, \dots, A'_{n-1})$ splinters. Consider any crossing separations $a'_i \in A'_i \smallsetminus A'_j$ and $a'_j \in A'_j \smallsetminus A'_i$ with $i \neq j$. Since a'_i is nested with a_n it cannot lie in $A_j \smallsetminus A'_j$, and similarly $a'_j \notin A_i \smallsetminus A'_i$. Thus, $a'_i \in A_i \smallsetminus A_j$ and $a'_j \in A_j \smallsetminus A_i$. As $(A_1, \dots, A_n)$ splinters, a'_i and a'_j have a corner c in $A_i \cup A_j$. We need to show that $c \in A'_i \cup A'_j$, i.e., that c is nested with a_n. But this follows from the fish lemma, since a_n is nested with both $a'_i \in A'_i$ and $a'_j \in A'_j$.

As the family $(A'_1, \dots, A'_{n-1})$ splinters, the induction hypothesis implies that it has nested representatives $a''_i \in A'_i$. As these are nested with a_n by definition of the A'_i, the family $(a''_1, \dots, a''_{n-1}, a_n)$ can serve as our desired nested family of representatives of $(A_1, \dots, A_n)$. □

Proof of Theorem 12.5.1. Pick an enumeration of the distinguishable pairs of tangles in G, and let A_i denote the set of separations of G that efficiently distinguish the ith pair. By Lemma 12.5.3 it suffices to show that the family of these A_i splinters. A_i

To prove this, consider any indices $i \neq j$ with crossing separations $s \in A_i \smallsetminus A_j$ and $t \in A_j \smallsetminus A_i$. Assume that these have orders $|s| \leqslant |t|$; that s efficiently distinguishes the ith pair of tangles, σ and σ'; and that t efficiently distinguishes the jth pair of tangles, τ and τ' (Figure 12.5.2). As $|s| \leqslant |t|$, the tangle τ orients s; we may assume it orients s as σ does. As s is not in A_j but has order $|s| \leqslant |t|$, it does not distinguish τ from τ'; so τ' too orients s as τ and σ do.

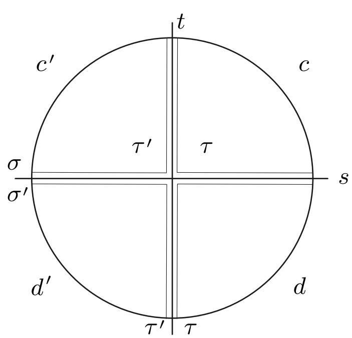

Fig. 12.5.2. Crossing separations $s \in A_i \smallsetminus A_j$ and $t \in A_j \smallsetminus A_i$ and their four corner separations c, c', d, d'

Consider the corner c of s and t at the top right in Figure 12.5.2. If c has order $|c| \leqslant |t|$, then τ orients it. Since τ has the profile property (P), and it orients t towards the right in the diagram and s upwards, it cannot orient c towards the centre. So τ orients c towards the top right. But τ' orients c towards the centre, since it orients t towards the left and is consistent. Thus, c distinguishes τ from τ', and by our assumption that $|c| \leqslant |t|$ it does so efficiently, giving $c \in A_j$ as desired. We may thus assume that $|c| > |t|$. Analogously with c' instead of c, we are done unless $|c'| > |t|$, which we therefore assume.

But now, by submodularity (†), these assumptions about the orders of c and c' imply that their opposite corners d' and d have order $< |s|$. Hence both σ and σ' orient them. As σ orients s upwards, it must orient d and d' towards the centre, by consistency. But since d and d' have order $< |s|$, they cannot distinguish σ from σ', so σ' too orients them towards the centre. As σ' orients s downwards, it thus violates (R).

This contradiction completes the proof that the family of all A_i splinters, as was our aim to show. □

The construction of a nested set of efficient tangle-distinguishing separations in our first proof of Theorem 12.5.1 makes a number of choices, those of the a_n at the start of the proof of the splinter lemma. Objects defined for a graph without making such choices, but so that they are invariant under its automorphisms,[7] are called *canonical*. Our second proof of Theorem 12.5.1 provides such a canonical nested set of tangle-distinguishing separations: the set of *all separations that efficiently distinguish some pair of tangles and cross as few other such separations as possible.* We shall make this definition precise later.

canonical

$D(\tau,\tau')$

Given a pair τ,τ' of tangles in G, let $D(\tau,\tau')$ denote the set of separations of G that distinguish τ,τ' efficiently. Our second proof of Theorem 12.5.1 rests on a property of these sets that has its own name. An *entanglement* in G is any non-empty set T of separations of G that satisfies the following condition:

entanglement

> *Whenever some $t \in T$ is crossed by a separation s of G so that two corners of s and t that lie on the same side of t have order at most $|t|$, then at least one of these corners has order exactly $|t|$ and also lies in T.* (E)

Lemma 12.5.4. *For every pair τ,τ' of distinguishable tangles in G, the separations of G that distinguish them efficiently form an entanglement.*

Proof. Let crossing separations $t \in D(\tau,\tau')$ and s be given as in condition (E), with corners c,d on the same side of t. These are as shown in Figure 12.5.2, together with the tangles τ,τ' that t distinguishes.

Since c and d have order at most $|t|$, the tangle τ orients them. As τ orients t towards the right, it cannot orient both c and d inwards, by (R). But if τ orients c outwards, say, then c distinguishes τ from τ', which orients t towards the left and hence, by consistency, orients c inwards. Since t distinguishes τ and τ' efficiently we cannot have $|c| < |t|$, so $|c| = |t|$ as required. □

[7] See the discussion following Theorem 12.3.7 for why this can be important.

We need three more general lemmas about crossing separations.

Lemma 12.5.5. *Let s, t be two crossing separations of G, with opposite corners c, d. Then every separation r of G satisfies the following:*

(i) *If r crosses c or d, it also crosses s or t.*

(ii) *If r crosses c and d, it also crosses s and t.*

Proof. (i) is equivalent to the fish lemma (12.5.2).

(ii) If $r = \{A, B\}$ and $\{C, D\} \in \{s, t\}$ are nested, they have comparable orientations; indeed we may assume that $(A, B) \geqslant (C, D)$ (Ex. 40). Since c and d are opposite corners of s and t, one of them lies on the D-side of $\{C, D\}$. This corner has a side contained in $D \subseteq B$, while its other side contains $C \supseteq A$. It is therefore nested with r. □

Given a set S of separations of G and any separation a of G, let S
$x_S(a) =: x(a)$ denote the number of separations in S that a crosses. x

Lemma 12.5.6. *Let $s, t \in S$ be two crossing separations, with opposite corners c, d. Then $x(c) + x(d) < x(s) + x(t)$.*

Proof. To evaluate x on our separations c, d, s, t, we have to count the separations $r \in S$ which they cross. Counting only separations $r \notin \{s, t\}$, we obtain $x(c) + x(d) \leqslant x(s) + x(t)$ from Lemma 12.5.5. The inequality gets strict when we allow $r \in \{s, t\}$, since s and t cross each other but not their corners c and d. □

Let $\mathcal{S}$ be any set of sets of separations of G.[8] Let us say that a separation s of G is $\mathcal{S}$-*friendly* if it lies in some $S \in \mathcal{S}$ such that no other separation in S crosses fewer separations in $\bigcup \mathcal{S}$ than s does. Note that every $S \in \mathcal{S}$ has at least one $\mathcal{S}$-friendly element. $\mathcal{S}$-friendly

Lemma 12.5.7. *Suppose that for all sets $S, S' \in \mathcal{S}$, possibly $S = S'$, and any crossing separations $s \in S$ and $s' \in S'$, one of the following two assertions holds:*

(i) *s and s' have opposite corners $c \in S$ and $c' \in S'$;*

(ii) *s and s' have pairs of opposite corners c, d in S and c', d' in S'.*

Then the $\mathcal{S}$-friendly separations of G are nested.

Proof. Let us show that crossing separations s, s' of G cannot both be $\mathcal{S}$-friendly. Suppose they are. Then $s \in S$ and $s' \in S'$ for some $S, S' \in \mathcal{S}$ witnessing this. By assumption, these satisfy (i) or (ii).

Suppose first that (i) holds. Then $x(c) + x(c') < x(s) + x(s')$ by Lemma 12.5.6, where $x = x_{\hat{S}}$ for $\hat{S} = \bigcup \mathcal{S}$. So either $x(c) < x(s)$ or $x(c') < x(s')$. In the first case s is not $\mathcal{S}$-friendly, by $c \in S$ and our choice of S for s. In the second case s' is not $\mathcal{S}$-friendly, by $c' \in S'$ and our choice of S' for s'.

[8] When we apply Lemma 12.5.7 later, $\mathcal{S}$ will be a set of entanglements in G.

Suppose now that (ii) holds. We may assume that $x(s) \leqslant x(s')$. Lemma 12.5.6 now gives $x(c') + x(d') < x(s) + x(s') \leqslant 2\,x(s')$, so $x(c') < x(s')$ or $x(d') < x(s')$. In either case s' is not $\mathcal{S}$-friendly, by $c', d' \in S'$ and our choice of S' for s'. $\square$

The following general theorem about entanglements in graphs implies Theorem 12.5.1. Indeed, take as the $\mathcal{S}$ in Theorem 12.5.8 the set of all sets $D(\tau, \tau')$ of efficient distinguishers of distinguishable tangles in G; recall that, by Lemma 12.5.4, these sets $D(\tau, \tau')$ are entanglements.

Theorem 12.5.8. (Carmesin & Kurkofka, 2024)
Let $\mathcal{S}$ be any set of entanglements in a graph G. Then the $\mathcal{S}$-friendly separations of G are nested.

Proof. It suffices to show that $\mathcal{S}$ satisfies the premise of Lemma 12.5.7. So consider $S, S' \in \mathcal{S}$ with crossing separations $s \in S$ and $s' \in S'$ as in the lemma, with $|s| \leqslant |s'|$ say. Let us call a corner of s and s' *relevant* if it has order at most $|s'|$. Let us show the following:

relevant

At least three corners of s and s' are relevant. (1)

Suppose at most two corners of s and s' are relevant. By submodularity (†) and $|s| \leqslant |s'|$, at least one of any two opposite corners must be relevant, so there are adjacent relevant corners. As their opposite two corners are irrelevant and thus have order $> |s'|$, our adjacent relevant corners have order $< |s|$ by submodularity. If they lie on the same side of s, then this contradicts our assumption that S is an entanglement. If they lie on the same side of s', it contradicts our assumption that S' is an entanglement. This completes the proof of (1).

We may assume that all four corners of s and s' are relevant. (2)

To prove (2), suppose only three corners are relevant. Two of these, c' and d' say, lie on the same side of s'. One of them, c' say, then lies in S' and has order exactly $|s'|$, by (E) for S'. By submodularity, the corner c opposite c' has order at most $|s|$; in particular, c is relevant too (Fig. 12.5.3). By assumption, the remaining corner d is irrelevant and hence has order $> |s'|$. By submodularity, therefore, its opposite corner d' has order $< |s|$. Condition (E) applied to S and s now yields that $|c| = |s|$ and $c \in S$. We thus have (i) of Lemma 12.5.7, as desired. This completes the proof of (2).

By (E) for S', the separation s' has corners on both sides that have order $|s'|$ and lie in S'. Assume first that these are adjacent, say c' and d. Let c be the corner opposite c', and d' that opposite d. By submodularity, c and d' have order at most $|s|$. As they lie on the same side of s, one of them lies in S, by (E) for S. Since its opposite corner lies in S', we have (i) of Lemma 12.5.7 as desired.

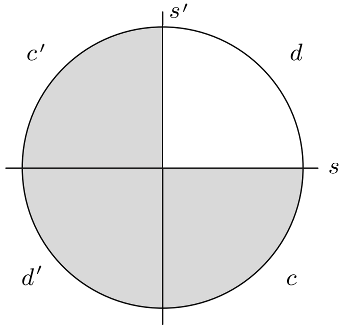

Fig. 12.5.3. Corners of s and s' in the proof of Theorem 12.5.8. Relevant corners in the proof of (2) are shaded.

Suppose now that the two corners on the different sides of s' that have order $|s'|$ and lie in S' are opposite corners, say c' and c in Figure 12.5.3. By submodularity, this means that our assumption of $|s| \leqslant |s'|$ must hold with equality: $|s| = |s'|$. Applying (E) to the side of s on which d lies, we find that d or c' lies in S. If c' does, then $c' \in S$ and $c \in S'$ satisfy (i) of Lemma 12.5.7. So we assume that $d \in S$.

Applying (E) to the other side of s, we find that c or d' lies in S. If c does, then again c and c' satisfy (i) of Lemma 12.5.7. So we assume that $d' \in S$. We now have (ii) of Lemma 12.5.7, with $d, d' \in S$ and $c, c' \in S'$. This completes our verification of the premise of Lemma 12.5.7. □

Theorem 12.5.8, like Theorem 12.5.1, tells us how all the tangles in a graph G are distinguished by a certain set N of nested separations. These are induced by a tree-decomposition of G (Exercises 16, 51). More generally, for every $k \in \mathbb{N}$ the set N_k of separations of order $< k$ in N is induced by a tree-decomposition $(T_k, \mathcal{V}_k)$ of G. Then every k-tangle τ in G orients all the separations in N_k towards V_t for some $t \in T_k$ depending on τ, where $\mathcal{V}_k = (V_t)_{t \in T_k}$ as usual. We may think of the tangle τ as 'living in' this node t of T_k, or in the bag V_t it specifies. For $i < j$, the tree T_i is obtained from T_j by contracting the edges that correspond, as in Lemma 12.3.1, to separations in $N_j \smallsetminus N_i$, while the bags of $\mathcal{V}_i$ are unions of bags of $\mathcal{V}_j$ indexed by the $t \in T_j$ that lie in a common branch set of this contraction minor $T_i \preccurlyeq T_j$.

The tree-structure which a nested set of separations imposes on G can also be made visible in a more direct way, without a detour via tree-decompositions, as follows.

The set $\vec{E}(T)$ of the oriented edges of a tree T is partially ordered by letting $\vec{e} \geqslant \vec{f}$ whenever $\vec{e} = (e, x, y)$ and $\vec{f} = (f, u, v)$ are such that the unique $\{x, y\}$–$\{u, v\}$ path in T starts in y and ends in u. Note that every two edges of T are 'nested' in our earlier sense that they have comparable orientations. Call a subset τ of $\vec{E}(T)$ *consistent* if it contains no $\overleftarrow{e}$ such $\vec{e} \geqslant \vec{f}$ *consistent* that $\vec{e} \geqslant \vec{f} \in \tau$ for some f. The consistent orientations of $E(T)$ are precisely those towards some node t of T (Ex. 50). The set of minimal elements of such an orientation, then, is the set

σ_t

$$\sigma_t := \{\, (e,s,t) \in \vec{E}(T) \mid e = st \in T \,\}$$

of all the incoming edges at t.

S-tree Given a set S of separations of G, an *S-tree* is a pair (T,α) such that T is a tree and $\alpha\colon \vec{E}(T) \to \vec{S}$ respects the orderings on these sets and commutes with inversion: $\alpha(\vec{e}) \geqslant \alpha(\vec{f})$ if $\vec{e} \geqslant \vec{f}$ (Fig. 12.5.4), and $\alpha(\overleftarrow{e}) = (B,A)$ whenever $\alpha(\vec{e}) = (A,B)$.[9] We then say that (T,α) is an
over S-tree *over* a set $\mathcal{F}$ of subsets of $\vec{S}$ if $\alpha(\sigma_t) \in \mathcal{F}$ for every node t of T. Note that we do not require α to be injective, not even on the sets σ_t.

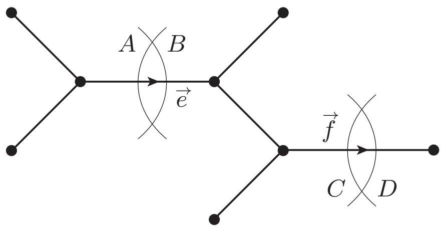

Fig. 12.5.4. An S-tree with $\alpha(\vec{e}) = (A,B) \geqslant (C,D) = \alpha(\vec{f})$

As every two edges of a tree have comparable orientations, the image of $\vec{E}(T)$ under α for an S-tree (T,α) defines a nested subset of S. Conversely, every nested set N of separations of G underlies the image of α for a suitable S-tree (T,α), with $S = N$ say. For example, a set N of efficient tangle-distinguishers as in Theorem 12.5.1 can be displayed by an N-tree (T,α) in this way. Its minors $T_1 \preccurlyeq T_2 \preccurlyeq \ldots$ defined earlier, then, display the sets N_k of efficient tangle-distinguishers of order $< k$ in N for $k = 1,2,\ldots$: those that distinguish, among others, the k-tangles in G.

The second main theorem about tangles, the *tangle–tree duality* theorem, employs S_k-trees to display the tree-structure of graphs that have *no* k-tangles for some desired k. Intuitively, the bags of the corresponding tree-decompositions will be 'too small' to be home to a tangle. To express this directly in terms of S-trees, we need another definition.

Recall that the *k-tangles* in G are the orientations of S_k that avoid $\mathcal{T}$.
star A set σ of oriented separations is a *star* if $(A,B) \geqslant (D,C)$ for all distinct
$\mathcal{T}^*$ $(A,B),(C,D) \in \sigma$ (Fig. 12.5.5). Let $\mathcal{T}^* := \{\, \sigma \in \mathcal{T} \mid \sigma \text{ is a star} \,\}$.

As we shall see in a moment, S_k-trees over $\mathcal{T}^*$ are natural certificates for the non-existence of k-tangles. The *tangle–tree duality* theorem guarantees that these certificates exist whenever a graph has no k-tangle:

---

[9] A tree-decomposition $(T,\mathcal{V})$, for example, makes T into an S-tree for the set S of separations it induces.

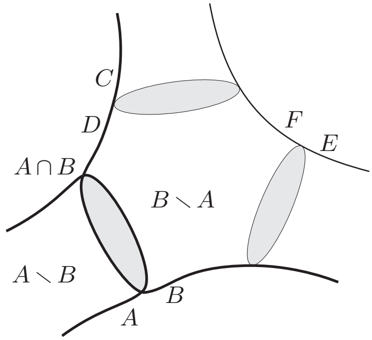

Fig. 12.5.5. The separations $(A,B),(C,D),(E,F)$ form a star

tangle–tree duality theorem

Theorem 12.5.9. (Robertson & Seymour 1991)
The following assertions are equivalent for all graphs G and integers $k > 0$:

(i) *G has no tangle of order k;*

(ii) *G has an S_k-tree over $\mathcal{T}^*$.*

Lemma 12.5.10. (Uncrossing Lemma)
Every consistent orientation of S_k that has a subset $\sigma \in \mathcal{T}$ also has a subset in $\mathcal{T}^*$.

Proof. Let $\sigma = \{\,(A_i,B_i) \mid i = 1,2,3\,\} \in \mathcal{T}$ be contained in a consistent orientation O of S_k. We show that, unless σ is a star, we can replace one of its elements by a strictly greater separation in O while keeping it in $\mathcal{T}$. In finitely many steps this will turn σ into a star: a subset of O in $\mathcal{T}^*$.

If σ is not a star we may assume that $(A_1,B_1) \not\geqslant (B_2,A_2)$. By (†) and $\sigma \subseteq \vec{S}_k$ we may further assume that $(C,D) := (A_1 \cap B_2, B_1 \cup A_2) \in \vec{S}_k$; the other case, that $(B_1 \cap A_2, A_1 \cup B_2) \in \vec{S}_k$, is analogous. Since O is consistent and $(C,D) \geqslant (A_1,B_1) \in O$, we cannot have $(D,C) \in O$. Thus, $(C,D) \in O$. But $(C,D) > (A_1,B_1)$, since either $A_1 \not\subseteq B_2$ or $B_1 \not\supseteq A_2$ by assumption. Replacing (A_1,B_1) in σ with (C,D) to obtain σ', say, gives the desired reduction: since any vertex or edge of $G[A_1]$ that does not lie in $G[C]$ lies in $G[A_2]$, and $(A_2,B_2) \in \sigma'$, we have $\sigma' \in \mathcal{T}$. $\square$

Proof of Theorem 12.5.9. (ii)→(i) Suppose G has an S_k-tree (T,α) over $\mathcal{T}^*$. Then any k-tangle τ in G defines via α an orientation of the edges of T. Let $t \in T$ be the last node of any maximal path in T whose edges are all oriented forward. Then all the edges at t are oriented towards t, and α maps these oriented edges to a star in $\mathcal{T}$. In particular, τ has a subset in $\mathcal{T}$ and thus cannot be a tangle.

(i)→(ii) To make this implication susceptible to an induction proof, our strategy is to relax both (i) and (ii) in some carefully designed way that makes it easier to prove the implication at the induction start. The relaxation, however, gets a little less in each inductive step, so that at the end it has evaporated and we have our desired implication of (i)→(ii). More precisely, we shall prove that certain sets $\mathcal{T}^+$ satisfy one of

(1) S_k has a consistent orientation that avoids $\mathcal{T}^* \cup \mathcal{T}^+$;

(2) G has an S_k-tree over $\mathcal{T}^* \cup \mathcal{T}^+$.

This yields (i)→(ii) when $\mathcal{T}^+ = \emptyset$. Indeed, if G has no k-tangle then, by Lemma 12.5.10, (1) cannot hold. So (2) must hold, but (2) is just (ii) when $\mathcal{T}^+ = \emptyset$.

$\mathcal{T}^+$

The sets $\mathcal{T}^+$ for which we prove this will depend on certain sets O of oriented separations, which we shall vary in the induction. Given O, let $\mathcal{T}^+ := \big\{\{(B,A)\} : (A,B) \in O\big\} \smallsetminus \mathcal{T}^*$. Thus, S_k-trees over $\mathcal{T}^* \cup \mathcal{T}^+$ are given more freedom compared with S_k-trees over $\mathcal{T}^*$: singleton stars associated with their leaves need not lie in $\mathcal{T}^*$, as long as they lie in $\mathcal{T}^+$. This happens when the inverse of their element lies in O but is not of the form (A,V). In the latter case, $\{(V,A)\}$ lies in $\mathcal{T}^*$, and hence not in $\mathcal{T}^+$.

The sets O we consider in our induction will be all the sets $O \subseteq \vec{S_k}$ that have the following two properties common to all k-tangles. First, they will contain all the separations (A,V) with $|A| < k$.

Second, they will be *closed upwards* in $\vec{S_k}$: they contain any $(A,B) \in \vec{S_k}$ such that $(A,B) \geqslant (C,D)$ for some $(C,D) \in O$. In particular, our induction covers the smallest such set, $O = \{\,(A,V) : |A| < k\,\}$, for which $\mathcal{T}^+ = \emptyset$.

Since all singletons $\{(B,A)\}$ in $\mathcal{T}^*$ satisfy $B = V$, by definition of $\mathcal{T}$, our sets O contain the corresponding separations (A,B). So the relationship between each O and the $\mathcal{T}^+$ it gives rise to is precisely

$$O = \big\{(A,B) \in \vec{S_k} : \{(B,A)\} \in \mathcal{T}^* \cup \mathcal{T}^+\big\}. \qquad (*)$$

S_O

If O contains a separation (X,Y) together with its inverse (Y,X), then (T,α) with $T = K^2$ and $\alpha\colon \vec{E}(T) \to \{(X,Y),(Y,X)\}$ satisfies (2), by $(*)$. We now assume that O is antisymmetric: that it contains no inverse pair of separations. Then O is a consistent orientation of some subset S_O of S_k.

Let us prove by induction on $|S_k \smallsetminus S_O|$ that $\mathcal{T}^+$ satisfies (1) or (2). At the induction start, O is an orientation of all of S_k. Hence if (1) fails then O has a subset $\sigma \in \mathcal{T}^* \cup \mathcal{T}^+$. By $(*)$ and the antisymmetry of O, we have $|\sigma| \geqslant 2$. Let T be the star $K_{1,n}$ with $n = |\sigma|$ leaves, and let α map its oriented edges (e,s,t) with s a leaf bijectively to the elements of σ. Then (T,α) satisfies (2), by definition of σ and $(*)$.

U_i, W_i

In the induction step we have $S_k \smallsetminus S_O \neq \emptyset$. Choose $\{U_1, W_1\}$ and $\{U_2, W_2\}$ in $S_k \smallsetminus S_O$ so that both (U_i, W_i) are maximal in $\vec{S_k} \smallsetminus \vec{S_O}$ and $(U_1, W_1) \geqslant (W_2, U_2)$.[10] Then the (U_i, W_i) are maximal even in $\vec{S_k} \smallsetminus O$:

[10] It is easy to see that such separations exist, just choose them in turn.

for any $(U,W) > (U_i,W_i)$ in $\vec{S_k} \smallsetminus O$ we would have $(W,U) \in O$ by the maximality of (U_i,W_i) in $\vec{S_k} \smallsetminus \vec{S_O}$, so $(W_i,U_i) > (W,U)$ would be in O (this being closed upwards in $\vec{S_k}$), contradicting the fact that $(U_i,W_i) \notin \vec{S_O}$.

Thus, the sets $O_i := O \cup \{(U_i,W_i)\}$ are again closed upwards in $\vec{S_k}$, O_i and are orientations of subsets of S_k that contain S_O properly. We may therefore apply the induction hypothesis to these O_i, to obtain (1) or (2) for $\mathcal{T}_i^+ := \big\{\{(B,A)\} : (A,B) \in O_i\big\} \smallsetminus \mathcal{T}^* \supseteq \mathcal{T}^+$ instead of $\mathcal{T}^+$. $\mathcal{T}_i^+$

Since (1) holds with $\mathcal{T}^+$ as soon as it holds with $\mathcal{T}_1^+$ or $\mathcal{T}_2^+$, we may assume that both $\mathcal{T}_i^+$ satisfy (2). Let (T_i,α_i) be the corresponding (T_i,α_i) S_k-trees over $\mathcal{T}^* \cup \mathcal{T}_i^+$. If one of these is in fact over $\mathcal{T}^* \cup \mathcal{T}^+$ we are done, so we assume not. Then each T_i has a node u_i such that $\alpha_i(e_i,w_i,u_i) = (W_i,U_i)$ for every edge $e_i = w_iu_i$ of T_i at u_i. e_i,u_i,w_i

Every such u_i must be a leaf. For otherwise $(W_i,U_i) \geqslant (U_i,W_i)$, and therefore $W_i \subseteq U_i$. But then $U_i = V$ and hence $(W_i,U_i) \in O$, contrary to the choice of (U_i,W_i). Similarly, these leaves u_i are unique. Indeed if u_i' is another leaf, with incident edge $e_i' = u_i'w_i'$ say, then $(e_i',u_i',w_i') \geqslant (e_i,w_i,u_i)$ and therefore $\alpha_i(e_i',u_i',w_i') \geqslant \alpha_i(e_i,w_i,u_i) = (W_i,U_i)$. Hence if $\alpha_i(e_i',u_i',w_i') = (U_i,W_i)$ then $(U_i,W_i) \geqslant (W_i,U_i)$, with a contradiction as earlier. Thus, our S_k-trees (T_i,α_i) are nearly over $\mathcal{T}^* \cup \mathcal{T}^+$: they are, except at their leaf u_i.

Choose $\{X_1,X_2\} \in S_k$ of minimum order with X_1,X_2

$$(U_1,W_1) \geqslant (X_1,X_2) \geqslant (W_2,U_2);$$

such a separation exists, because (U_1,W_1) is a candidate. We shall modify the maps α_i to maps α_i' defining S_k-trees (T_i,α_i'). These will again be over $\mathcal{T}^* \cup \mathcal{T}^+$ except at u_i, where we shall have $\alpha_i'(e_i,w_i,u_i) = (X_{3-i},X_i)$. Our plan will then be to join the newly labelled trees $T_i - u_i$ together by adding the edge w_iw_{3-i} and mapping it to their common separation (X_{3-i},X_i), to obtain our desired S_k-tree over $\mathcal{T}^* \cup \mathcal{T}^+$.

To define α_i', consider an edge e of T_i. Name its ends t,t' so that $(e_i,u_i,w_i) \geqslant (e,t,t')$. Then if $\alpha_i(e,t,t') = (A,B)$, say, let

$$\alpha_i'(e,t,t') = (A',B') := (A \cup X_i, B \cap X_{3-i})$$ α_i'

and $\alpha_i'(e,t',t) = (B',A')$. (Fig. 12.5.6).

Let us show that α_i' maps $\vec{E}(T_i)$ to $\vec{S_k}$. Given an edge $e \in E(T_i)$, let $\vec{e}$ be such that $(e_i,u_i,w_i) \geqslant \vec{e}$, as in the definition of α_i'. Then for $\alpha_i(\vec{e}) = (A,B)$ we have $\alpha_i'(\vec{e}) = (A',B') = (A \cup X_i, B \cap X_{3-i})$. By (†) we have $|A' \cap B'| \leqslant |A \cap B| < k$, as desired, if the order of

$$(Y_i,Y_{3-i}) := (A \cap X_i, B \cup X_{3-i})$$

is no less than the order of $\{X_1,X_2\}$. And it cannot be less, as that would contradict our choice of $\{X_1,X_2\}$ since $(U_1,W_1) \geqslant (Y_1,Y_2) \geqslant (W_2,U_2)$.

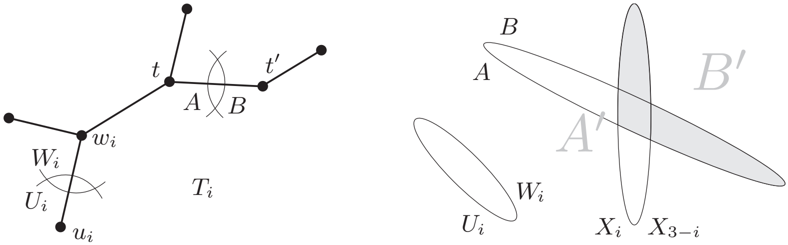

Fig. 12.5.6. Shifting (A,B) to (A',B')

Indeed, recall that $U_i \subseteq X_i$ as well as $U_i \subseteq A$, and $W_i \supseteq X_{3-i}$ as well as $W_i \supseteq B$, because $(U_i,W_i) \geqslant (X_i,X_{3-i})$ as well as $(U_i,W_i) \geqslant (A,B)$. Hence $(U_i,W_i) \geqslant (Y_i,Y_{3-i})$, and $(Y_i,Y_{3-i}) \geqslant (X_i,X_{3-i}) \geqslant (W_{3-i},U_{3-i})$.

Next, let us show that α'_i, like α_i, respects the orderings of $\vec{E}(T_i)$ and $\vec{S_k}$. Consider any $\vec{e} \geqslant \vec{f} \in \vec{E}(T_i)$, with $(A,B) = \alpha_i(\vec{e}) \geqslant \alpha_i(\vec{f}) = (C,D)$ say. Then $A \subseteq C$ and $B \supseteq D$. If $(e_i,u_i,w_i) \geqslant \vec{e}$, then $\alpha'_i(\vec{e}) =: (A',B') \geqslant (C',D') := \alpha'_i(\vec{f})$, because $A' = A \cup X_i \subseteq C \cup X_i = C'$ and $B' = B \cap X_{3-i} \supseteq D \cap X_{3-i} = D'$. On the other hand if $(e_i,u_i,w_i) \geqslant \overleftarrow{e}$ but $(e_i,u_i,w_i) \geqslant \vec{f}$, then $A' = A \cap X_{3-i} \subseteq A \subseteq C \subseteq C \cup X_i = C'$ while $B' = B \cup X_i \supseteq B \supseteq D \supseteq D \cap X_{3-i} = D'$, so again $(A',B') \geqslant (C',D')$. Since α'_i, by definition, commutes with inversions of orientations, this covers all the cases to be considered.

So (T_i,α'_i) is indeed an S_k-tree. Let us show that (T_i,α'_i), like (T_i,α_i), is an S_k-tree over $\mathcal{T}^* \cup \mathcal{T}^+$ except at u_i, where

$$\alpha'_i(e_i,w_i,u_i) = (W_i \cap X_{3-i}, U_i \cup X_i) = (X_{3-i},X_i)\,.$$

Consider a node $t \neq u_i$, and let $e = st$ be its incident edge in u_iT_it. Let $\alpha_i(e,s,t) =: (A,B)$ and $\alpha'_i(e,s,t) =: (A',B')$. As $(e_i,u_i,w_i) \geqslant (e,s,t)$,

$$(U_i,W_i) \geqslant (A,B) \geqslant (A \cup X_i, B \cap X_{3-i}) = (A',B').$$

Suppose first that $\alpha_i(\sigma_t) = \{(A,B)\}$. Let us show that t must be a leaf.

If not, then $\alpha_i(e',s',t) = (A,B)$ for another edge $e' = s't$ at t. Then $(A,B) \geqslant (B,A)$, and hence $B = V$. But then $(A,B) \in O$, and hence $(U_i,W_i) \in O$ since O is closed upwards in $\vec{S_k}$. This contradicts our choice of (U_i,W_i). Hence t is indeed a leaf of T_i.

Since (T_i,α_i) is over $\mathcal{T}^* \cup \mathcal{T}^+$ except at u_i, we thus have $(B,A) \in O$ by $(*)$. As O is closed upwards in $\vec{S_k}$, this means that $(B',A') \geqslant (B,A)$ is also in O. Thus, $\alpha'_i(\sigma_t) = \{(A',B')\} \in \mathcal{T}^* \cup \mathcal{T}^+$ by $(*)$, as desired.

Suppose now that $|\alpha_i(\sigma_t)| \geqslant 2$. Then $\alpha_i(\sigma_t) \in \mathcal{T}^*$; let us show that also $\alpha'_i(\sigma_t) \in \mathcal{T}^*$. Since α'_i respects the orderings of $\vec{E}(T_i)$ and $\vec{S_k}$, we already know that $\alpha'_i(\sigma_t)$ is a star. To show that $\alpha'_i(\sigma_t) \in \mathcal{T}$, consider any edge $e' = s't \neq st$. Let $(C,D) := \alpha_i(e',s',t)$ and $(C',D') := \alpha'_i(e',s',t)$.

Then $(e_i, u_i, w_i) \geqslant (e', t, s')$ by the choice of $e \neq e'$. Hence $C' = C \cap X_{3-i}$, but also $A' = A \cup X_i$. Thus, any vertex or edge of $G[C]$ that is not in $G[C']$ lies in $G[X_i] \subseteq G[A']$. The fact that $\alpha_i(\sigma_t) \in \mathcal{T}$ thus implies $\alpha'_i(\sigma_t) \in \mathcal{T}$, as desired.

Let T be the tree obtained from the disjoint union of $T_1 - u_1$ and $T_2 - u_2$ by joining w_1 to w_2 by a new edge e. Let $\alpha\colon \vec{E}(T) \to \vec{S_k}$ map (e, w_{3-i}, w_i) to $\alpha'_i(e_i, u_i, w_i) = (X_i, X_{3-i})$ for $i = 1, 2$, and otherwise extend the α'_i. Then α commutes with the inversion of $\vec{e}$ and of (X_1, X_2), and $\alpha(\sigma_t) = \alpha'_i(\sigma_t) \in \mathcal{T}^* \cup \mathcal{T}^+$ for all $t \in T$, in particular for $t = w_i$. Hence, (T, α) satisfies (2). □

Theorem 12.5.9, as stated above, is a special case of a more general result in which $\mathcal{T}^*$ can be replaced by other collections $\mathcal{F}$ of 'forbidden' stars of oriented separations. Given a set S of separations, an *$\mathcal{F}$-tangle* of S is a consistent orientation of S that has no subset $\sigma \in \mathcal{F}$. For *$\mathcal{F}$-tangle*

$$\mathcal{F}_k := \left\{\sigma \subseteq \vec{S} : \sigma \text{ is a star and } \left|\bigcap\{\, B : (A,B) \in \sigma\}\right| < k\right\}$$

we obtain the following more tangle-like duality theorem for tree-width:

Theorem 12.5.11. *The following two assertions are equivalent for all graphs G and integers $k > 0$:*

(i) *G has tree-width less than $k - 1$;*

(ii) *G has no $\mathcal{F}_k$-tangle of S_k.*

It is not hard to show that (i) is equivalent to the existence of an S_k-tree over $\mathcal{F}_k$ in G (Ex. 59). Theorem 12.5.11 in that form and Theorem 12.5.9 are both corollaries of a more general theorem about tangle-tree duality in so-called 'abstract separation systems'; see the notes.

Moreover, any $\mathcal{F}_k$-tangle of S_k gives rise to a k-bramble (Ex. 59). Thus, Theorem 12.5.11 implies the tree-width duality theorem (12.4.3).

12.6 Tree-decompositions and forbidden minors

If $\mathcal{H}$ is any set or class of graphs, then the class

$$\mathrm{Forb}_{\preccurlyeq}(\mathcal{H}) := \{\, G \mid G \not\succcurlyeq H \text{ for all } H \in \mathcal{H}\,\}$$ $\mathrm{Forb}_{\preccurlyeq}(\mathcal{H})$

of all graphs without a minor in $\mathcal{H}$ is a graph property, i.e. is closed under isomorphism.[11] 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*

[11] As usual, we abbreviate $\mathrm{Forb}_{\preccurlyeq}(\{H\})$ to $\mathrm{Forb}_{\preccurlyeq}(H)$.

(1.7.1) By Proposition 1.7.1, $\mathrm{Forb}_{\preccurlyeq}(\mathcal{H})$ is closed under taking minors, or *minor-closed*: if $G' \preccurlyeq G \in \mathrm{Forb}_{\preccurlyeq}(\mathcal{H})$ then $G' \in \mathrm{Forb}_{\preccurlyeq}(\mathcal{H})$. Every minor-closed property can in turn be expressed by forbidden minors:

[5.2] **Lemma 12.6.1.** *A graph property $\mathcal{P}$ can be expressed by forbidden minors if and only if it is closed under taking minors.*

$\overline{\mathcal{P}}$ *Proof.* For the 'if' part, note that $\mathcal{P} = \mathrm{Forb}_{\preccurlyeq}(\overline{\mathcal{P}})$, where $\overline{\mathcal{P}}$ is the complement of $\mathcal{P}$. □

In Section 12.7, we shall return to the general question of how a given minor-closed property is best represented by forbidden minors. In this section, we begin by looking at a particular example of such a property: bounded tree-width.

Consider the property of having tree-width less than some given integer k. By Lemmas 12.4.1 and 12.6.1, 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 12.6.2. *A graph has tree-width < 3 if and only if it has*
(7.3.1) *no K^4 minor.*

(12.3.2)
(12.3.5) *Proof.* By Corollary 12.3.5, we have $\mathrm{tw}(K^4) \geqslant 3$. By Lemma 12.4.1,
(12.3.6) therefore, a graph of tree-width < 3 cannot contain K^4 as a minor.
(12.4.1)

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 7.3.1, G' can be constructed recursively from triangles by pasting along K^2s. By induction on the number of recursion steps and Corollary 12.3.5, every graph constructible in this way has a tree-decomposition into triangles (as in the proof of Proposition 12.3.6). Such a tree-decomposition of G' has width 2, and by Lemma 12.3.2 it is also a tree-decomposition of G. □

As k grows, the list of forbidden minors characterizing the graphs of tree-width $< k$ seems to grow fast. They are known explicitly only up to $k = 4$; see the notes.

A question converse to the above is to ask for which H (other than K^3 and K^4) the tree-width of the graphs in $\mathrm{Forb}_{\preccurlyeq}(H)$ is bounded. This is the case, for example, when H is grid:

[12.7.1]
[12.7.3] **Theorem 12.6.3.** (Robertson & Seymour 1986)
For every integer r there is an integer k such that every graph of tree-width at least k has an $r \times r$ grid minor.

This *grid theorem* may, at first glance, look like just a specific and technical result. But it has a sweeping consequence:

Corollary 12.6.4. *Given a graph H, the graphs without an H minor have bounded tree-width if and only if H is planar.*

Proof. Since all grids and their minors are planar, every class $\mathrm{Forb}_{\preccurlyeq}(H)$ (4.4.6)
with a non-planar H contains all grids, which have unbounded tree-width (see after Theorem 12.4.3).

Conversely, every planar graph H is a minor of some grid: take a drawing of the graph, fatten its vertices and edges, and superimpose a sufficiently fine plane grid. Hence, by Theorem 12.6.3, the graphs in $\mathrm{Forb}_{\preccurlyeq}(H)$ have bounded tree-width as soon as H is planar. □

Theorem 12.6.3 has another interesting application. Recall that a class $\mathcal{H}$ of graphs has the *Erdős-Pósa property* if the number of vertices in a graph needed to cover all its subgraphs in $\mathcal{H}$ is bounded by a function of its maximum number of disjoint subgraphs in $\mathcal{H}$. Now let H be a H
fixed connected graph, and consider the class $\mathcal{H} = IH$ of graphs that $\mathcal{H}$
contract to a copy of H. (Thus, G has a subgraph in $\mathcal{H}$ if and only if $H \preccurlyeq G$.)

Theorem 12.6.5. (Robertson & Seymour 1986)
$\mathcal{H}$ has the Erdős-Pósa property if H is planar.

Proof. We have to find a function $f\colon \mathbb{N} \to \mathbb{N}$ such that, given $k \in \mathbb{N}$ and (12.3.1)
a graph G, either G contains k disjoint models of H or there is a set U of at most $f(k)$ vertices in G such that $H \not\preccurlyeq G - U$.

By Corollary 12.6.4, there exists for every $k \geqslant 1$ an integer w_k such that every graph of tree-width at least w_k contains the disjoint union of k copies of H (which is again planar) as a minor. Define

$$f(k) := 2f(k-1) + w_k$$

inductively, starting with $f(0) = f(1) = 0$.

To verify that f does what it should, we apply induction on k. For $k \leqslant 1$ there is nothing to show. Now let k and G be given for the induction step. If $\mathrm{tw}(G) \geqslant w_k$, we are home by definition of w_k. So assume that $\mathrm{tw}(G) < w_k$, and let $(T, (V_t)_{t \in T})$ be a tree-decomposition of G of width $< w_k$. Let us direct the edges $t_1 t_2$ of the tree T as follows. Let T_1, T_2 be the components of $T - t_1 t_2$ containing t_1 and t_2, respectively, and put

$$G_1 := G\Big[\bigcup_{t \in T_1} (V_t \setminus V_{t_2})\Big] \quad \text{and} \quad G_2 := G\Big[\bigcup_{t \in T_2} (V_t \setminus V_{t_1})\Big].$$

We direct the edge $t_1 t_2$ towards G_i if $H \preccurlyeq G_i$, thereby giving $t_1 t_2$ either one or both or neither direction.

If every edge of T receives at most one direction, we follow these to a node $t \in T$ such that no edge at t in T is directed away from t. As H is connected, this implies by Lemma 12.3.1 that V_t meets every $IH \subseteq G$. This completes the proof with $U = V_t$, since $|V_t| \leqslant w_k \leqslant f(k)$ by the choice of our tree-decomposition.

Suppose now that T has an edge t_1t_2 that received both directions. For each $i = 1, 2$ let us ask if we can cover all the models of H in G_i by at most $f(k-1)$ vertices. If we can, for both i, then by Lemma 12.3.1 the two covers combine with $V_{t_1} \cap V_{t_2}$ to the desired cover U for G. Suppose now that G_1 has no such cover. Then, by the induction hypothesis, G_1 contains $k-1$ disjoint models of H. Since t_1t_2 was also directed towards t_2, there is another such model in G_2. This gives the desired total of k disjoint models of H in G. □

Theorem 12.6.5 contains the Erdős-Pósa theorem 2.3.2 as the special case that $H = K^3$. It is best possible in that if H is non-planar, then $\mathcal{H} = IH$ does not have the Erdős-Pósa property (Exercise 62).

We conclude this section with some structure theorems for graphs not containing a given complete graph as a minor. These theorems are more difficult to prove than the results we have seen so far in this chapter, and they are not even that easy to state. But it's worth an effort: already the first of them, the excluded-K^n theorem, is both central to the proof of the graph minor theorem and can be applied elsewhere.

linear decomposition

A *linear decomposition* of G is a family $(V_i)_{i \in I}$ of vertex sets indexed by some linear order I such that $\bigcup_{i \in I} V_i = V(G)$, every edge of G has both its ends in some V_i, and $V_i \cap V_k \subseteq V_j$ whenever $i < j < k$. When G is finite, this is just a tree-decomposition whose decomposition tree is a path, and usually called a *path-decomposition*. If each V_i contains at most k vertices and k is minimal with this property, then $(V_i)_{i \in I}$ has *width* $k-1$.

$C_1, \ldots, C_k$

$S - k$

cuffs

Let S' be a subspace of a surface[12] S obtained by removing the interiors of finitely many disjoint closed discs, with boundary circles $C_1, \ldots, C_k$ say. This space is determined up to homeomorphism by S and the number k, and we denote it by $S - k$. Each C_i is the image of a continuous map $f_i \colon [0,1] \to S'$ that is injective except for $f_i(0) = f_i(1)$. We call $C_1, \ldots, C_k$ the *cuffs* of S' and the points $f_1(0), \ldots, f_k(0)$ their *roots*. The other points of each C_i are linearly ordered by f_i as images of $(0,1)$; when we use cuffs as index sets for linear decompositions below, we shall be referring to these linear orders. An *embedding* of a graph in S (or in $S - k$) is defined analogously to embeddings in the plane.

k-near embedding

Let H be a graph, S a surface, and $k \in \mathbb{N}$. We say that H is *k-nearly embeddable* in S if H has a set X of at most k vertices such that $H - X$ can be written as $H_0 \cup H_1 \cup \ldots \cup H_k$ so that

(N1) there exists an embedding $\sigma \colon H_0 \hookrightarrow S - k$ that maps only vertices to cuffs and no vertex to the root of a cuff;

[12] A compact connected 2-manifold without boundary; see Appendix B.

(N2) the graphs $H_1, \dots, H_k$ are pairwise disjoint (and may be empty), and $H_0 \cap H_i = \sigma^{-1}(\sigma(H_0) \cap C_i)$ for each i;

(N3) every H_i with $i \geqslant 1$ has a linear decomposition $(V_z^i)_{z \in C_i \cap \sigma(H_0)}$ of width $< k$ such that $\sigma^{-1}(z) \in V_z^i$ for all z.

Here, then, is the structure theorem for graphs without a K^n minor.[13] Note that, for $n = 5$, Wagner's Theorem (7.3.4) remains stronger and more precise. The case of $n = 4$ is covered by Proposition 7.3.1.

Theorem 12.6.6. (Robertson & Seymour 2003)
For every $n \geqslant 5$ there exists a $k \in \mathbb{N}$ such that every graph not containing K^n as a minor has a tree-decomposition whose torsos are k-nearly embeddable in a surface in which K^n is not embeddable.

Theorem 12.6.6 is true also for infinite graphs; see the notes.

Note that there are only finitely many surfaces in which K^n is not embeddable. The set of those surfaces in the statement of Theorem 12.6.6 could therefore be replaced by just two surfaces: the orientable and the non-orientable surface of maximum genus in this set.

Theorem 12.6.6 also has a converse, though only a qualitative one. A decomposition as described does not by itself preclude the presence of a K^n minor. But for every n there is an r such that no graph with such a decomposition has a K^r minor. This is because the adhesion sets of the tree-decomposition have bounded size, e.g. by $2k + n$, since they induce complete subgraphs in the torsos, and these are k-nearly embeddable in a surface that does not accommodate K^n.

For graphs without a given topological minor, there is a related structure theorem:

Theorem 12.6.7. (Grohe & Marx 2012)
For every $n \geqslant 5$ there exists a $k \in \mathbb{N}$ such that every graph not containing K^n as a topological minor has a tree-decomposition whose torsos are either k-nearly embeddable in a surface of Euler genus $\leqslant k$ or have at most k vertices of degree $> k$.

(See Appendix B for the definition of the Euler genus of a surface.)

There are also structure theorems for excluding infinite minors, and we now state two of these.

First, the structure theorem for excluding $K^{\aleph_0}$. Call a graph H *nearly planar* if H has a finite set X of vertices such that $H - X$ can be written as $H_0 \cup H_1$ so that (N1–2) hold with $S = S^2$ (the sphere) and $k = 1$, while (N3) holds with $k = |X|$. (In other words, deleting nearly planar

[13] Robertson and Seymour proved several versions of this theorem, of which Theorem 12.6.6 is the simplest. See the notes.

a bounded number of vertices makes H planar except for a subgraph of bounded linear width sewn on to the unique cuff of $S^2 - 1$.) A tree-decomposition $(T, (V_t)_{t \in T})$ of a graph G has *finite adhesion* if all its adhesion sets are finite and for every infinite path $t_1 t_2 \dots$ in T the value of $\liminf_{i \to \infty} |V_{t_i} \cap V_{t_{i+1}}|$ is finite.

finite adhesion

Unlike its counterpart for K^n, the excluded-$K^{\aleph_0}$ structure theorem has a direct converse. It thus characterizes the graphs without a $K^{\aleph_0}$ minor, as follows:

Theorem 12.6.8. *A graph G has no $K^{\aleph_0}$ minor if and only if G has a tree-decomposition of finite adhesion whose torsos are nearly planar.*

Finally, a structure theorem for excluding $K^{\aleph_0}$ as a topological minor. Let us say that G has *finite tree-width* if G admits a tree-decomposition $(T, (V_t)_{t \in T})$ into finite parts such that for every infinite path $t_1 t_2 \dots$ in T the set $\bigcup_{j \geqslant 1} \bigcap_{i \geqslant j} V_{t_i}$ is finite.

finite tree-width

Theorem 12.6.9. *The following assertions are equivalent for connected graphs G:*

(i) *G does not contain $K^{\aleph_0}$ as a topological minor;*

(ii) *G has finite tree-width;*

(iii) *G has a normal spanning tree T such that for every ray R in T there are only finitely many vertices v such that G contains an infinite v–$(R - v)$ fan.*

12.7 The graph minor theorem

Graph properties that are closed under taking minors occur frequently in graph theory. Among the most natural examples are the properties of being embeddable in some fixed surface, such as planarity.

By Kuratowski's theorem, planarity can be expressed by forbidding the minors K^5 and $K_{3,3}$. This is a *good characterization* of planarity in the following sense. Suppose we wish to persuade someone that a certain graph is planar: this is easy (at least intuitively) if we can produce a drawing of the graph. But how do we persuade someone that a graph is non-planar? By Kuratowski's theorem, there is also an easy way to do that: we just have to exhibit an IK^5 or $IK_{3,3}$ in our graph, as an easily checked 'certificate' for non-planarity. Our simple Proposition 12.6.2 is another example of a good characterization: if a graph has tree width < 3, we can prove this by exhibiting a suitable tree-decomposition; if not, we can produce an IK^4 as evidence.

Theorems that characterize a property $\mathcal{P}$ by a set of forbidden minors are doubtless among the most attractive results in graph theory. As we saw in Lemma 12.6.1, such a characterization exists whenever $\mathcal{P}$ is minor-closed: then $\mathcal{P} = \mathrm{Forb}_{\preccurlyeq}(\overline{\mathcal{P}})$, where $\overline{\mathcal{P}}$ is the complement of $\mathcal{P}$. However, one naturally seeks to make the set of forbidden minors as small as possible. And there is indeed a unique smallest such set: the set (12.6.1)

$$\mathcal{K}_{\mathcal{P}} := \{\, H \mid H \text{ is } \preccurlyeq\text{-minimal in } \overline{\mathcal{P}} \,\}$$

Kuratowski set $\mathcal{K}_{\mathcal{P}}$

satisfies $\mathcal{P} = \mathrm{Forb}_{\preccurlyeq}(\mathcal{K}_{\mathcal{P}})$ and is contained in every other set $\mathcal{H}$ such that $\mathcal{P} = \mathrm{Forb}_{\preccurlyeq}(\mathcal{H})$. We call $\mathcal{K}_{\mathcal{P}}$ the *Kuratowski set* for $\mathcal{P}$.

Clearly, the elements of $\mathcal{K}_{\mathcal{P}}$ are incomparable under the minor relation $\preccurlyeq$. Now the *graph minor theorem* of Robertson & Seymour says that any set of $\preccurlyeq$-incomparable graphs must be finite:

Theorem 12.7.1. (Robertson & Seymour 1986–2004) *graph minor theorem*
The finite graphs are well-quasi-ordered by the minor relation $\preccurlyeq$.

We shall give a sketch of the proof of the graph minor theorem at the end of this section.

Corollary 12.7.2. *The Kuratowski set for any minor-closed graph property is finite.* □

As a special case of Corollary 12.7.2 we have, at least in principle, a Kuratowski-type theorem for every surface S: the property $\mathcal{P}(S)$ of embeddability in S is characterized by the finite set $\mathcal{K}_{\mathcal{P}(S)}$ of forbidden minors. $\mathcal{P}(S)$

Corollary 12.7.3. *For every surface S there exists a finite set of graphs $H_1, \ldots, H_n$ such that a graph is embeddable in S if and only if it contains none of $H_1, \ldots, H_n$ as a minor.* □

While Corollary 12.7.3 is immediate from the graph minor theorem, it can also be proved more directly. It is our next goal to do this. The main step is to prove that the graphs in $\mathcal{K}_{\mathcal{P}(S)}$ do not contain arbitrarily large grids as minors (Lemma 12.7.4). Then their tree-width is bounded (Theorem 12.6.3), so $\mathcal{K}_{\mathcal{P}(S)}$ is well-quasi-ordered (Theorem 12.4.2) and therefore finite.

The proof of Lemma 12.7.4 gives a good impression of the interplay between graph minors and surface topology, which – by way of Theorem 12.6.6, which we could not prove here – is also one of the key ingredients of the proof of the graph minor theorem. Appendix B summarizes the necessary background on surfaces, including a lemma

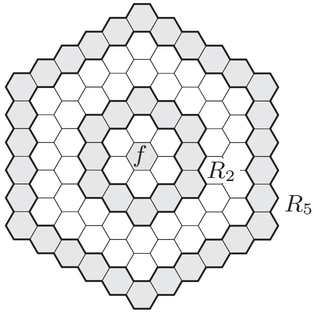

Fig. 12.7.1. The hexagonal grid H^6 with central face f and rings R_2 and R_5

used in the proof. For convenience (cf. Proposition 1.7.3 (ii)), we shall work with hexagonal rather than square grids.

H^r Denote by H^r the plane hexagonal grid whose dual has radius r (Figure 12.7.1). The face corresponding to the central vertex of its dual
faces is its *central face*. (Generally, when we speak of the *faces* of H^r, we mean its hexagonal faces, not its outer face.) Call a subgrid H^k of H^r
standard *standard* if their central faces coincide. We write S_k for the perimeter
$S_1, \dots, S_r$ cycle of the standard subgrid H^k in H^r; for example, S_1 is the hexagon
ring R_k bounding the central face of H^r. The *ring* R_k is the subgraph of H^r formed by S_k and S_{k+1} and the edges between them.

Lemma 12.7.4. *For every surface S there exists an integer r such that no graph that is minimal with the property of not being embeddable in S contains H^r as a topological minor.*

(4.1.2)
(4.2.2)
(4.3.2)
(App. B)

Proof. Let G be a graph that cannot be embedded in S and is minimal with this property. Our proof will run roughly as follows. Since G is minimally not embeddable in S, we can embed it in an only slightly larger surface S'. If G contains a very large H^r grid, then by Lemma B.6 some large H^m subgrid will be flat in S', that is, the union of its faces in S' will be a disc D'. We then pick an edge e from the middle of this H^m grid and embed $G - e$ in S. Again by Lemma B.6, one of the rings of our H^m will be flat in S. In this ring we can embed the (planar) subgraph of G which our first embedding had placed in D'; note that this subgraph contains the edge e. The rest of G can then be embedded in S outside this ring much as before, yielding an embedding of all of G in S (a contradiction).

More formally, let $\varepsilon := \varepsilon(S)$ denote the Euler genus of S. Let r
ε be large enough that H^r contains $\varepsilon + 3$ disjoint copies of H^{m+1}, where
r, m $m := 3\varepsilon + 4$. We show that G has no TH^r subgraph.

Let $e' = u'v'$ be any edge of G, and choose an embedding σ' of $G - e'$ in S. Choose a face with u' on its boundary, and another with v'

on its boundary. Cut a disc out of each face and add a handle between the two holes, to obtain a surface S' of Euler genus $\varepsilon + 2$ (Lemma B.3). Embedding e' along this handle, extend σ' to an embedding of G in S'. $\sigma' \colon G \hookrightarrow S'$

Suppose G has a subgraph $H = TH^r$. Let $f \colon H^r \to H$ map the vertices of H^r to the corresponding branch vertices of H, and its edges to the corresponding paths in H between those vertices. Let us show that H^r has a subgrid H^m (not necessarily standard) whose hexagonal face boundaries correspond (by $\sigma' \circ f$) to circles in S' that bound disjoint open discs there. H f

By the choice of r, we can find $\varepsilon + 3$ disjoint copies of H^{m+1} in H^r. The standard subgrids H^m of these H^{m+1} are not only disjoint, but sufficiently spaced out in H^r that their deletion leaves a tree $T \subseteq H^r$ that sends an edge to each of them (Figure 12.7.2). Hence whenever we pick one hexagon from each of these H^m and delete the images C of those hexagons in S', the component D_0 of the remainder of S' that contains $(\sigma' \circ f)(T)$ meets all those C in its boundary. By Lemma B.6 and $\varepsilon(S') = \varepsilon + 2$, therefore, it cannot be true that none of our circles C bounds a disc in S' that is disjoint from $(\sigma' \circ f)(T)$.

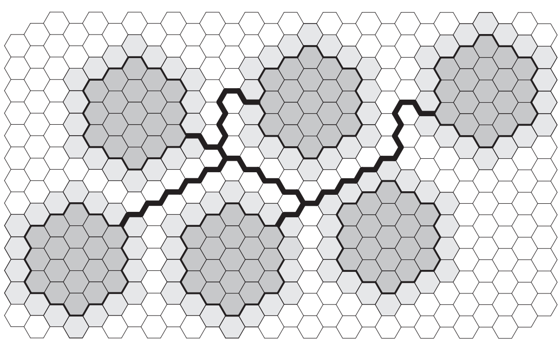

Fig. 12.7.2. Disjoint copies of H^m ($m = 3$) linked up by a tree in the rest of H^r

Hence for one of our copies of H^m in H^r, the image of every hexagon in S' bounds an open disc that is disjoint from $(\sigma' \circ f)(T)$. Let us show that these discs are disjoint. If not, then one of them, D say, contains a point x from the boundary of another such disc. But then D also contains $(\sigma' \circ f)(T)$, contrary to assumption, because we can walk from x to $(\sigma' \circ f)(T)$ in $(\sigma' \circ f)(H^r) \subseteq S'$ avoiding the boundary of D.

From now on, we shall work with this fixed H^m and will no longer consider its supergraph H^r. We write $C_i := f(S_i)$ for the images in G of the concentric cycles S_i of this H^m ($i = 1, \ldots, m$). C_i

Pick an edge $e = uv$ of C_1, and choose an embedding σ of $G - e$ in S. As before, Lemma B.6 implies that one of the $\varepsilon + 1$ disjoint rings R_{3i+2} in our H^m ($i = 0, \ldots, \varepsilon$), R_k say, has the property that its hexagons e $\sigma \colon G - e \hookrightarrow S$ k

correspond (by $\sigma \circ f$) to circles in S that bound disjoint open discs there
R (Figure 12.7.3). Let $R \supseteq (\sigma \circ f)(R_k)$ be the closure in S of the union
of those discs, which is a cylinder in S. One of its two boundary circles
C is the image under σ of the cycle $C := C_{k+1}$ in G to which f maps the
perimeter cycle S_{k+1} of our special ring $R_k \subseteq H^m$.

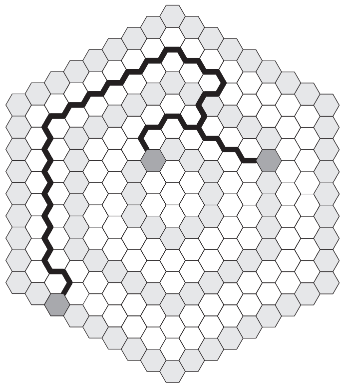

Fig. 12.7.3. A tree linking up hexagons selected from the rings $R_2, R_5, R_8 \ldots$

H' Let $H' := f(H^{k+1}) \subseteq G$, where H^{k+1} is standard in our H^m. Recall
that $\sigma' \circ f$ maps the hexagons of H^{k+1} to circles in S' bounding disjoint
open discs there. The closure in S' of the union of these discs is a disc
D' D' in S', bounded by $\sigma'(C)$. Deleting a small open disc inside D' that
R' does not meet $\sigma'(G)$, we obtain a cylinder $R' \subseteq S'$ that contains $\sigma'(H')$.

We shall now combine the embeddings $\sigma: G - e \hookrightarrow S$ and $\sigma': G \hookrightarrow S'$
σ'' to an embedding $\sigma'': G \hookrightarrow S$, which will contradict the choice of G.
Let $\varphi: \sigma'(C) \to \sigma(C)$ be a homeomorphism between the images of C in
S' and in S that commutes with these embeddings, i.e., is such that
φ $\sigma|_C = (\varphi \circ \sigma')|_C$. Then extend this to a homeomorphism $\varphi: R' \to R$.
The idea now is to define σ'' as $\varphi \circ \sigma'$ on the part of G which σ' maps
to D' (which includes the edge e on which σ is undefined), and as σ on
the rest of G (Fig. 12.7.4).

To make these two partial maps compatible, we start by defining σ'' on C as $\sigma|_C = (\varphi \circ \sigma')|_C$. Next, we define σ'' separately on the components of $G - C$. Since $\sigma'(C)$ bounds the disc D' in S', we know that σ' maps each component J of $G - C$ either entirely to D' or entirely to $S' \setminus D'$. On all the components J such that $\sigma'(J) \subseteq D'$, and on all the edges they send to G, we define σ'' as $\varphi \circ \sigma'$. Thus, σ'' embeds these components in R. Since $e \in f(H^k) = H' - C$, this includes the component of $G - C$ that contains e.

It remains to define σ'' on the components of $G - C$ which σ' maps to $S' \setminus D'$. As $\sigma'(C_k) \subseteq D'$, these do not meet C_k. Since $\sigma(C \cup C_k)$ is

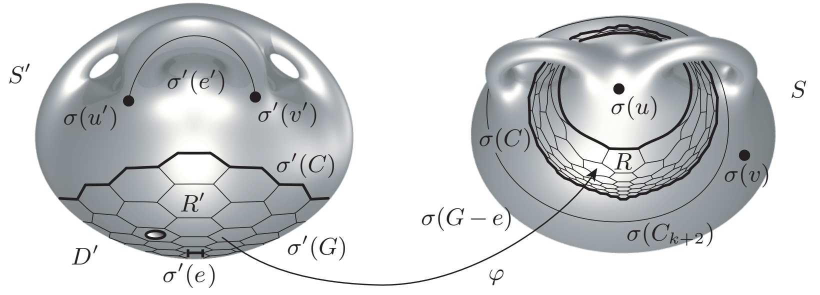

Fig. 12.7.4. Combining $\sigma'\colon G \hookrightarrow S'$ and $\sigma\colon G-e \hookrightarrow S$ to $\sigma''\colon G \hookrightarrow \mathrm{S}$

the frontier of R in S, this means that $\sigma(J) \subseteq S \setminus R$ or $\sigma(J) \subseteq R$ for every such component J.

For the component J_0 of $G-C$ that contains C_{k+2} we cannot have J_0 $\sigma(J_0) \subseteq R$: as $S_{k+2} \cap R_k = \emptyset$, this would mean that $\sigma(C_{k+2})$ lies in a disc $D \subseteq R$ corresponding to a face of R_k, which is impossible since S_{k+2} sends edges to vertices of S_{k+1} outside the boundary of that face. We thus have $\sigma(J_0) \subseteq S \setminus R$, and define σ'' as σ on J_0 and on all the J_0–C edges of G.

Next, consider any remaining component J of $G-C$ that sends no edge to C. If $\sigma(J) \subseteq S \setminus R$, we define σ'' on J as σ. If $\sigma(J) \subseteq R$, then J is planar. Since J sends no edge to C, we can have σ'' map J to any open disc in R that has not so far been used by σ''.

It remains to define σ'' on the components $J \neq J_0$ of $G-C$ which σ' maps to $S' \setminus D'$ and for which G contains a J–C edge. Let $\mathcal{J}$ be the $\mathcal{J}$ set of all those components J. We shall group them by the way they attach to C, and define σ'' for these groups in turn.

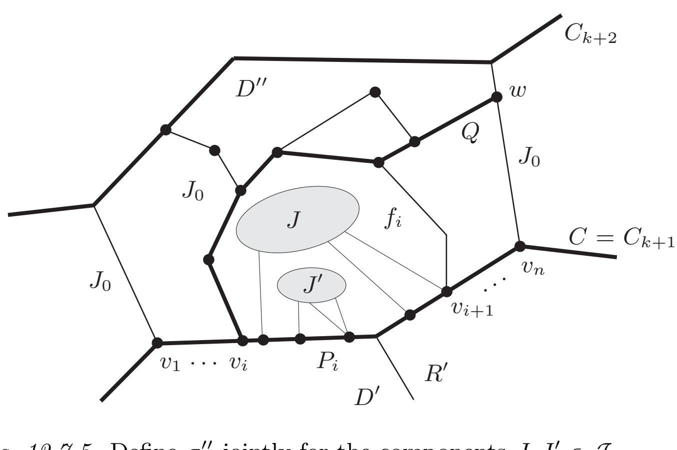

Fig. 12.7.5. Define σ'' jointly for the components $J, J' \in \mathcal{J}$ that attach to the same $P_i \subseteq C$

Since $m \geqslant k+2$, the disc D' lies inside a larger disc in S', which is D'' the union of D' and closed discs D'' bounded by the images under $\sigma' \circ f$ of the hexagons in R_{k+1}. By definition of $\mathcal{J}$, the embedding σ' maps every $J \in \mathcal{J}$ to such a disc D'' (Fig. 12.7.5). On the path P in C such that $\sigma'(P) = \sigma'(C) \cap D''$ (which is the image under f of one or two consecutive edges on S_{k+1}), let $v_1, \ldots, v_n$ be the vertices with a neighbour P_i in J_0, in their natural order along P, and write P_i for the segment of P from v_i to v_{i+1}. For any v_i with $1 < i < n$, pick a v_i–J_0 edge and extend it through J_0 to a path Q from v_i to C_{k+2} (which exists by definition of J_0); let w be its first vertex that σ' maps to the boundary circle of D''. By Lemma 4.1.2 applied to $\sigma'(v_i Q w)$ and the two arcs joining $\sigma'(v_i)$ to $\sigma'(w)$ along the boundary circle of D'', there is no arc through D'' that links $\sigma'(P_{i-1})$ to $\sigma'(P_i)$ but avoids $\sigma'(v_i Q w)$. Hence, every $J \in \mathcal{J}$ with $\sigma'(J) \subseteq D''$ has all its neighbours on C in the same P_i, and σ' maps f_i J to the face f_i of the plane graph $\sigma'(G[J_0 \cup C]) \cap D''$ whose boundary contains P_i. We shall define σ'' jointly on all those $J \in \mathcal{J}$ which σ' maps to this f_i, for $i = 1, \ldots, n-1$ in turn.

To do so, we choose an open disc D_i in $S \setminus R$ that has a boundary circle containing $\sigma(P_i)$ and avoids the image of σ'' as defined until now. Such D_i exists in a strip neighbourhood of $\sigma(C)$ in S, because components $J' \in \mathcal{J}$ attaching to a segment $P_j \neq P_i$ of C send no edge to $\mathring{P}_i$. Choose a homeomorphism φ_i from the boundary circle of f_i to that of D_i so that $\sigma|_{P_i} = (\varphi_i \circ \sigma')|_{P_i}$, and extend this to a homeomorphism φ_i from the closure of f_i in S' to the closure of D_i in S. For every $J \in \mathcal{J}$ with $\sigma'(J) \subseteq f_i$, and for all J–C edges of G, define σ'' as $\varphi_i \circ \sigma'$. □

(1.7.3)
(12.4.2)
(12.6.3)

Proof of Corollary 12.7.3. By their minimality, the graphs in $\mathcal{K}_{\mathcal{P}(S)}$ are incomparable under the minor-relation. If their tree-width is bounded, then $\mathcal{K}_{\mathcal{P}(S)}$ is well-quasi-ordered by the minor relation (Theorem 12.4.2), and hence must be finite. So assume their tree-width is unbounded, and let r be as in Lemma 12.7.4. By Theorem 12.6.3, some $H \in \mathcal{K}_{\mathcal{P}(S)}$ has a grid minor large enough to contain H^r. By Proposition 1.7.3, H^r is a topological minor of H, contrary to the choice of r. □

We finally come to the proof of the graph minor theorem itself. The complete proof would still fill a book or two, but we are well equipped now to get a good understanding of its main ideas and overall structure. For background on surfaces, we once more refer to Appendix B.

(12.1.3)
(12.2.1)
(12.4.2)
(12.6.3)

Proof of the graph minor theorem (sketch). 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 (Corollary 12.6.4) and are therefore well-quasi-ordered by Theorem 12.4.2. 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. But as in Wagner's theorem (7.3.4) for $n = 5$, the parts in these tree-decompositions are no longer constrained in terms of order now but in more subtle structural terms. Roughly speaking, for every n there exists a finite set $\mathcal{S}$ of surfaces such that every graph without a K^n minor has a tree-decomposition into parts each 'nearly' embeddable in one of the surfaces $S \in \mathcal{S}$; see Theorem 12.6.6. By a generalization of Theorem 12.4.2 – and hence of Kruskal's theorem – it now suffices, essentially, to prove that the set of all the parts in these tree-decompositions is well-quasi-ordered: then the graphs decomposing into these parts are well-quasi-ordered, too. Since $\mathcal{S}$ is finite, every infinite sequence of such parts has an infinite subsequence whose members are all (nearly) embeddable in the same surface $S \in \mathcal{S}$. Thus all we have to show is that, given any surface S, all the graphs embeddable in S are well-quasi-ordered by the minor relation.

This is shown by induction on the Euler genus of S, using the same approach as before: if $H_0, H_1, H_2, \ldots$ is an infinite sequence of graphs embeddable in S, we may assume that none of the graphs $H_1, H_2, \ldots$ contains H_0 as a minor. If $S = S^2$ we are back in the case that H_0 is planar, so the induction starts. For the induction step we now assume that $S \neq S^2$. Again, the exclusion of H_0 as a minor constrains the structure of the graphs $H_1, H_2, \ldots$, this time topologically: each H_i with $i \geqslant 1$ has an embedding in S which meets some circle $C_i \subseteq S$ that does not bound a disc in S in no more than a bounded number of vertices (and no edges), say in $X_i \subseteq V(H_i)$. (The bound on $|X_i|$ depends on H_0, but not on H_i.) Cutting along C_i and capping the hole(s), we obtain one or two new surfaces of smaller Euler genus. If the cut produces only one new surface S_i, then our embedding of $H_i - X_i$ still counts as a near-embedding of H_i in S_i (since X_i is small). If this happens for infinitely many i, then infinitely many of the surfaces S_i are also the same, and the induction hypothesis gives us a good pair among the corresponding graphs H_i. On the other hand, if we get two surfaces S_i' and S_i'' for infinitely many i (without loss of generality the same two surfaces), then H_i decomposes accordingly into subgraphs H_i' and H_i'' embedded in these surfaces, with $V(H_i' \cap H_i'') = X_i$. The set of all these subgraphs taken together is again well-quasi-ordered by the induction hypothesis, and hence so are the pairs (H_i', H_i'') by Lemma 12.1.3. Using a sharpening of the lemma that takes into account not only the graphs H_i' and H_i'' themselves but also how X_i lies inside them, we finally obtain indices i, j not only with $H_i' \preccurlyeq H_j'$ and $H_i'' \preccurlyeq H_j''$, but also such that these minor embeddings extend to the desired minor embedding of H_i in H_j – completing the proof of the graph minor theorem.

The graph minor theorem does not extend to graphs of arbitrary cardinality, but it might extend to countable graphs. Whether or not it does appears to be a difficult problem. It may be related to the following intriguing conjecture, which easily implies the graph minor theorem for finite graphs (Exercise 68). Call a graph H a *proper minor* of G if there is a contraction from a subgraph of G onto H that is not an isomorphism from G to H.

Self-minor conjecture. (Seymour 1980s)
Every countably infinite graph is a proper minor of itself.

In addition to its impact on 'pure' graph theory, the graph minor theorem has had far-reaching algorithmic consequences. Using their structure theorem for the graphs in $\mathrm{Forb}_{\preccurlyeq}(K^n)$, Theorem 12.6.6, Robertson and Seymour have shown that testing for any fixed minor is 'fast': for every graph H there is a polynomial-time algorithm[14] that decides whether or not the input graph contains H as a minor. By the minor theorem, then, every minor-closed graph property $\mathcal{P}$ can be decided in polynomial (even cubic) time: if $\mathcal{K}_{\mathcal{P}} = \{H_1, \dots, H_k\}$ is the corresponding set of 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 closed under taking minors: 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

[14] indeed a cubic one – although with an enormous constant depending on H

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

$$\text{Forb}_{\leqslant}(X) := \{\, a \in A \mid a \ngeqslant x \text{ for all } x \in X \,\}.$$

Show that $\leqslant$ is a well-quasi-ordering on A if and only if every subset $B \subseteq A$ that is closed under $\leqslant$ (i.e. such that $x \leqslant y \in B \Rightarrow x \in B$) can be written as $B = \text{Forb}_{\leqslant}(X)$ with finite X.

3. Prove Proposition 12.1.1 and Corollary 12.1.2 directly, without using Ramsey's theorem.

4.$^-$ Show that the relation $\leqslant$ between rooted trees defined in the text is indeed a quasi-ordering.

5. Show that the finite trees are not well-quasi-ordered by the subgraph relation.

6. The last step of the proof of Kruskal's theorem considers a 'topological' embedding of T_m in T_n that maps the root of T_m to the root of T_n. Suppose we assume inductively that the trees of A_m are embedded in the trees of A_n in the same way, with roots mapped to roots. We thus seem to obtain a proof that the finite rooted trees are well-quasi-ordered by the subgraph relation, even with roots mapped to roots. Where is the error?

7. Extend Kruskal's theorem to trees whose vertices are labelled from a well-quasi-ordered set. The tree embedding is defined as before but in addition respects the ordering of the labels.

8. Are the connected finite graphs well-quasi-ordered by contraction alone (i.e. by taking minors without deleting edges or vertices)?

9.$^+$ Relax the minor relation by not insisting that branch sets be connected. Show that the finite graphs are well-quasi-ordered by this relation.

10.$^+$ Show that the finite graphs are not well-quasi-ordered by the topological minor relation.

11.$^+$ Given $k \in \mathbb{N}$, is the class $\{\, G \mid G \not\supseteq P^k \,\}$ well-quasi-ordered by the subgraph relation?

12.$^-$ Let G be a graph, T a tree, and $\mathcal{V} = (V_t)_{t \in T}$ a family of subsets of $V(G)$. Show that $(T, \mathcal{V})$ is a tree-decomposition of G if and only if

(i) for every $v \in V(G)$ the set $T_v := \{ t \mid v \in V_t \}$ is connected in T;

(ii) $T_u \cap T_v \neq \emptyset$ for every edge uv of G.

13.$^-$ Consider a tree-decomposition of a graph G in which some parts contain other parts. Modify it into a tree-decomposition whose parts are the $\subseteq$-maximal parts of the first decomposition. How does the new tree arise from the old?

14. Let G be a graph, T a set, and $(V_t)_{t \in T}$ a family of subsets of $V(G)$ satisfying (T1) and (T2) from the definition of a tree-decomposition. Show that there exists a tree on T that makes (T3) true if and only if there exists an enumeration $t_1, \ldots, t_n$ of T such that for every $k = 2, \ldots, n$ there is a $j < k$ satisfying $V_{t_k} \cap \bigcup_{i<k} V_{t_i} \subseteq V_{t_j}$.

(The new condition tends to be more convenient to check than (T3). It can help, for example, with the construction of a tree-decomposition into a given set of parts.)

15. Prove the following converse of Lemma 12.3.1: if $(T, \mathcal{V})$ satisfies condition (T1) and the statement of the lemma, then $(T, \mathcal{V})$ is a tree-decomposition of G.

16. Recall that two separations $\{U_1, U_2\}$ and $\{W_1, W_2\}$ of G are *nested* if we can choose $i, j \in \{1, 2\}$ so that $U_i \subseteq W_j$ and $U_{3-i} \supseteq W_{3-j}$.

(i) Show that the separations $S_e := \{U_1, U_2\}$ in Lemma 12.3.1 are pairwise nested (for different choices of the edge $e = t_1t_2 \in T$).

(ii)$^+$ Conversely, show that given a set N of nested separations of G there is a tree-decomposition $(T, \mathcal{V})$ of G such that $N = \{ S_e \mid e \in E(T) \}$.

17.$^+$ Prove Theorem 12.3.7 for $k = 3$. Specifically, prove Tutte's theorem that every 2-connected graph has a tree-decomposition of adhesion 2 whose torsos are each either 3-connected or a cycle. Conversely, show that every graph with such a tree-decomposition is 2-connected.

(Hint. Try the tree-decomposition defined, as in Exercise 16 (ii), by the set of all separations of order 2 that are nested with all other such separations.)

18. Describe the tree-decomposition of a contraction minor H of G which a given tree-decomposition of G induces as in Lemma 12.3.3, in terms subtrees of T (as in Exercise 12).

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. Use the previous exercise to 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 (Figure 7.3.1).

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. Show that for $n \geqslant 3$ the graphs K^n, C^n, an arbitrary tree of order n, and the $n \times n$ grid have tree-decompositions of widths $n-1$, 2, 1, and n, respectively. For K^n and C^n show that these values are best possible.

25. Can the tree-width of a subdivision of a graph G be smaller than $\mathrm{tw}(G)$? Can it be larger?

26. Show that the tree-width of a finite graph is at least its minimum degree. Is this still true for infinite graphs?

27.$^+$ Show that if a graph has circumference $k \neq 0$, then its tree-width is at most $k-1$.

28.$^+$ A graph is called *outerplanar* if it has a drawing in which every vertex lies on the boundary of the outer face. Show that outerplanar graphs can have arbitrarily large tree-width, or find the best upper bound.

A tree-decomposition whose tree is a path is a *path-decomposition*. The *path-width* $\mathrm{pw}(G)$ of G is the least width of a path-decomposition of G.

29. 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 Exercise 44, Chapter 5.)

30. (continued)

Prove the following analogue of Proposition 12.4.4 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.

31.$^+$ Do trees have unbounded path-width?

A *transaction* of a sequence $(v_1, \dots, v_n)$ of vertices is a set of disjoint paths from an initial segment $\{v_1, \dots, v_i\}$ to the rest, $\{v_{i+1}, \dots, v_n\}$.

32.$^+$ Given $k \in \mathbb{N}$ and a sequence $v_1, \dots, v_n$ of vertices in a graph G, show that G has a path-decomposition $(V_1, \dots, V_n)$ of adhesion $\leqslant k$, with $v_i \in V_i$ for all i, if and only if G contains no transaction $\mathcal{P}$ of $(v_1, \dots, v_n)$ of order $|\mathcal{P}| > k$.

33. Show that the cycle C^n has connected tree-width $\lceil n/2 \rceil$.

34. Show that the $n \times n$ grid has tree-width n.

35.$^-$ 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 $\mathcal{B}$.

36.$^-$ Let $\mathcal{P}$ be a minor-closed graph property. Show that strengthening the notion of a minor (for example, to that of topological minor) increases the set of forbidden minors required to characterize $\mathcal{P}$.

37. Deduce from the graph minor theorem that every minor-closed property can be expressed by forbidding finitely many topological minors. Is the same true for every property that is closed under taking topological minors?

Call a set $X \subseteq V(G)$ of vertices *k-connected in G* if $|X| \geqslant k$ and for all subsets $Y, Z \subseteq X$ with $|Y| = |Z| \leqslant k$ there are $|Y|$ disjoint Y–Z paths in G.

38.$^+$ Show that the tree-width of a graph G is large if and only if it contains a large set of vertices that is k-connected in G for some large k. For example, show that graphs of tree-width $< k$ contain no $(k+1)$-connected set of $3k$ vertices, and that graphs containing no $(k+1)$-connected set of $3k$ vertices have tree-width $< 4k$.

39. (continued)

(i)$^+$ Find an $\mathbb{N} \to \mathbb{N}^2$ function $k \mapsto (h, \ell)$ such that every graph with an ℓ-connected set of h vertices contains a bramble of order $> k$.

(ii)$^-$ Using the last exercise, deduce the following weakening of the difficult implication of Theorem 12.4.3: given k, every graph of large enough tree-width $f(k)$ contains a bramble of order $> k$.

40.$^-$ Show that if separations r, s are nested, they have orientations $\vec{r} \geqslant \vec{s}$.

41. Two cuts are *nested* if a side of one is contained in a side of the other. Show that the k-cuts in a k-edge-connected graph are nested if k is odd.

42. Characterize the 2-tangles in a graph.

43. Find a 4-tangle in the 3-dimensional cube.

When τ is a tangle and $(A, B) \in \tau$, we call A the *small side* of $\{A, B\}$ in τ and B its *big side*.

44. Let G be a graph with a tangle τ of order k.

(i) Show that every graph $G' \succcurlyeq G$ also has a k-tangle.

(ii) Justify the notion of a 'small side' by showing that if $(A, B) \in \tau$ and $\{A', B'\}$ is a separation of order $< k$ with $A' \subseteq A$ or $B' \supseteq B$, then $(A', B') \in \tau$.

(iii) Deduce from the profile property of τ that the intersection X of the big sides of any star in τ contains at least k vertices.

(iv) Deduce that, for every set X of fewer than k vertices, exactly one of the components C of $G - X$ is 'big', in the sense that $(V(G-C), X \cup V(C)) \in \tau$.

45.[+] Is (iv) of the previous exercise true also in infinite graphs (with k finite)?

A *decider set* for a tangle τ in G is a set $X \subseteq V(G)$ such that $|X \cap A| < |X \cap B|$ for every $(A, B) \in \tau$.

46.[+] Show that a k-tangle induced by a $2k$-tangle always has a decider set.

47. Show that separations $\{A, B\}$ and $\{C, D\}$ cross if and only if both A and B meet both C and D outside $A \cap B \cap C \cap D$.

48. Does every set of separations in a graph admit a consistent orientation?

49. Show that for every consistent orientation O of the set N of the separations $S_e := \{U_1, U_2\}$ in Lemma 12.3.1, one for every edge e of T, there exists a node t of T such that O orients every separation in N towards V_t.

50.[−] Show that orienting the edges of a tree $T = (V, E)$ towards some fixed node t is consistent for the partial ordering on $\vec{E}$ defined in Section 12.5. Is this a bijection between V and the consistent orientations of E?

51. (continued)

Show that S-trees are just a formal way to display, by a concrete tree, the 'tree-like' structure of nested sets of separations:

(i) Given an S-tree (T, α), show that $\alpha(\vec{E}(T)) \subseteq \vec{S}$ is nested. (A set $\sigma \subseteq \vec{S}$ is *nested* if $\{\, s \mid \vec{s} \in \sigma \,\}$ is nested.)

(ii)[+] Given a nested set S of separations, construct an S-tree (T, α) with surjective $\alpha \colon \vec{E}(T) \to \vec{S}$.

(iii) Derive Exercise 16 (ii).

52. (continued)

Prove the assertions about trees of tangles and tree-decompositions made in the long paragraph following the proof of Theorem 12.5.8.

53.[+] Show that in any graph of order n there are at most n tangles.

54. Recall from the proof of Theorem 12.4.3 how a k-bramble orients S_k. Is this orientation of S_k a tangle? Is it a profile?

55. Show the following implications for a graph G:

(i) G contains a k-block $\Rightarrow$ G has a bramble of order k.

(ii) G has a tangle of order k $\Rightarrow$ G has a bramble of order k.

(iii) G has a bramble of order $3k$ $\Rightarrow$ G has a tangle of order k.

(iv) G contains a k-block $\Rightarrow$ G has a k-tangle or $|B| \leqslant \frac{3}{2}(k-1)$.

Is there a function $f \colon \mathbb{N} \to \mathbb{N}$ such that, for every $k \in \mathbb{N}$, if G has a tangle of order at least $f(k)$ then it contains a k-block?

56. Show that k-blocks of size at least $3k/2$ define robust profiles of S_k. Strengthen Theorem 12.3.7 for such blocks, to find a canonical tree-decomposition which, for variable k, efficiently distinguishes all pairs of k-blocks of size at least $3k/2$ such that neither includes the other.

57.$^+$ Define as the *order* of a bipartition $\{A, B\}$ of the vertex set of a graph G the number $\|A, B\|$ of its A–B edges. Define *orientations* of such bipartitions in the obvious way. Call an orientation of the set of all the vertex bipartitions of order $< k$ an *edge-tangle of order k* if it has no three elements $(A_1, B_1), (A_2, B_2), (A_3, B_3)$ such that $A_1 \cup A_2 \cup A_3 = V(G)$. Show that the proofs of Theorems 12.5.1 and 12.5.8 yield tree-of-tangle theorems for edge tangles, the second of them canonical.

58. Discuss the following three shortcuts in the proof of Theorem 12.5.11, all designed to avoid shifting. Does any of them work?

   (i) In the induction step, pick any separation $s \in S_k \setminus S_O$ of your choice. (If you think it helps, choose it of minimum order.) Add each of its two orientations $\vec{s}, \overleftarrow{s}$ to O to obtain S_k-trees that are over $\mathcal{T}^* \cup \mathcal{T}^+$ except at one leaf, whose label is $\{\overleftarrow{s}\}$ or $\{\vec{s}\}$, respectively. Merge the trees at these leaves as in the text.

   Does it worry you that the extensions of O are no longer closed upwards in $\vec{S_k}$? If so, why? Then try the following adaptations of (i):

   (ii) Dispense with the requirement that O be closed upwards in $\vec{S_k}$.

   (iii) Add to O not just $\vec{s}$ or $\overleftarrow{s}$, but their entire up-closure in $\vec{S_k}$.

59. To illustrate the equivalence of Theorems 12.4.3 and 12.5.11, prove the following assertions for all graphs G:

   (i) G has tree-width $< k-1$ if and only if it has an S_k-tree over $\mathcal{F}_k$.

   (ii) G has a k-bramble if and only if it admits an $\mathcal{F}_k$-tangle of S_k.

60.$^+$ Modify the proof of Theorem 12.5.9 to obtain a proof of Theorem 12.5.11.

61. Extend Theorem 12.6.5 as follows. Let H be a connected planar graph, let $\mathcal{X}$ be any set of connected graphs including H, and let $\mathcal{H} := \{IX \mid X \in \mathcal{X}\}$. Show that $\mathcal{H}$ has the Erdős-Pósa property, witnessed by the same function f as defined in the proof of Theorem 12.6.5. Explain how it is possible that f depends on H but not on any of the other graphs in $\mathcal{X}$.

62.$^+$ Show that, for every non-planar graph H, the class IH fails to have the Erdős-Pósa property.

   (Hint. Embed H in a surface.)

63.$^+$ Let $\mathcal{H}$ be a class of connected graphs, and $k \geqslant 1$ an integer. Without using any theorems from this chapter, show that $\mathcal{H}$ has the Erdős-Pósa property for graphs without a k-tangle.

64.$^+$ Show that the four ingredients to the structure of the graphs in $\mathrm{Forb}_{\preccurlyeq}(K^n)$ as described in Theorem 12.6.6 – tree-decomposition, an *apex set* X, arbitrary surfaces $S \not\hookleftarrow K^n$, and *vortices* $H_1, \ldots, H_k$ – are all needed to capture all the graphs in $\mathrm{Forb}_{\preccurlyeq}(K^n)$. More precisely, find examples of graphs in $\mathrm{Forb}_{\preccurlyeq}(K^n)$ showing that Theorem 12.6.6 becomes false if we require in addition that the tree-decomposition has only one part, or that X is always empty, or that S is always the sphere, or that $H_1, \ldots, H_k$ are always empty. No exact proofs are required.

65.$^+$ (continued)

Show that, unlike in Theorem 12.6.6, the surfaces used in Theorem 12.6.7 cannot be limited to those in which K^n cannot be drawn. (As before, no exact proofs are required.)

66. Without using the graph minor theorem, show that the chromatic number of the graphs in any $\preccurlyeq$-antichain is bounded.

67. Let S_g denote the orientable surface obtained from the sphere by adding g handles. Find a lower bound for $|\mathcal{K}_{\mathcal{P}(S)}|$ in terms of g.

(Hint. The smallest g such that a given graph can be embedded in S_g is its *orientable genus*. Use the theorem that the orientable genus of a graph is equal to the sum of the genera of its blocks.)

68. Deduce the graph minor theorem from the self-minor conjecture.

69. Prove Theorem 12.6.9, assuming that G has a normal spanning tree.

70. Let G be a locally finite graph obtained from the $\mathbb{Z} \times \mathbb{Z}$ grid H by adding an infinite set of edges xy with $d_H(x, y)$ unbounded. Show that $G \succcurlyeq K^{\aleph_0}$. Can you do the same if the distances $d_H(x, y)$ are bounded (but at least 3)?

71. Is the infinite $\mathbb{Z} \times \mathbb{Z}$ grid a minor of the $\mathbb{Z} \times \mathbb{N}$ grid? Is the latter a minor of the $\mathbb{N} \times \mathbb{N}$ grid?

72.$^+$ Extend Proposition 12.3.6 to infinite graphs not containing an infinite complete subgraph.

73. Using the previous exercise, prove that if every finite subgraph of G has tree-width less than $k \in \mathbb{N}$ then so does G.

74. Show that no assumption of large finite connectivity can ensure that a countable graph has a K^r minor when $r \geqslant 5$. However, using the previous exercise show that sufficiently large finite connectivity forces an infinite graph to contain any given planar minor.

Notes

Robertson & Seymour have traditionally referred to the graph minor theorem as *Wagner's conjecture*. Wagner did indeed discuss this problem in the 1960s with his then students, Halin and Mader, and it seems that Mader conjectured a positive solution. Wagner himself always insisted that he did not – even after the graph minor theorem had been proved.

Robertson & Seymour's proof of the graph minor theorem is given in the numbers IV–VII, IX–XII and XIV–XXII of their series of over 20 papers under the common title of *Graph Minors*, most of which appeared in the *Journal of Combinatorial Theory, Series B*, between 1983 and 2012. Of their theorems cited in this chapter, Theorem 12.4.2 is from Graph Minors IV, Theorems 12.5.9 and 12.5.1 from Graph Minors X, Theorems 12.6.3 and 12.6.5 from Graph Minors V, and Theorem 12.6.6 from Graph Minors XVI.

Kruskal's theorem on the well-quasi-ordering of finite trees was first published in J.B. Kruskal, Well-quasi ordering, the tree theorem, and Vászonyi's conjecture, *Trans. Amer. Math. Soc.* **95** (1960), 210–225. Our proof is due to Nash-Williams, who introduced the versatile proof technique of choosing a 'minimal bad sequence'. This technique was also used in our proof of Higman's Lemma 12.1.3.

Nash-Williams generalized Kruskal's theorem to infinite graphs. This extension is much more difficult than the finite case. Its proof introduces as a tool the notion of *better-quasi-ordering*, a concept that has profoundly influenced well-quasi-ordering theory. The graph minor theorem is false for uncountable graphs; this was shown by R. Thomas, A counterexample to 'Wagner's conjecture' for infinite graphs, *Math. Proc. Camb. Phil. Soc.* **103** (1988), 55–57. Whether or not the countable graphs are well-quasi-ordered as minors, and whether the finite (or the countable) graphs are better-quasi-ordered as minors, are related questions that remain wide open. Both are related also to the self-minor conjecture. This, too, was originally intended to include graphs of arbitrary cardinality, but was disproved for uncountable graphs by B. Oporowski, A counterexample to Seymour's self-minor conjecture, *J. Graph Theory* **14** (1990), 521–524.

Doubling all the edges of an cycle, and then subdividing every new edge once, yields a graph which is incomparable under the topological minor relation to any graph obtained in this way from a cycle of any other length. This shows that the finite graphs are not well-quasi-ordered by the topological minor relation; cf. Exercise 10. C.-H. Liu, *Graph Structures and Well-Quasi-Ordering*, PhD thesis, Georgia Institute of Technology (2014), showed that this example is essentially the only obstruction to well-quasi-ordering by the topological minor relation: given any integer k, the finite graphs that do not contain, as a topological minor, a graph obtained from a path P^k in this way are well-quasi-ordered as topological minors.

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 seminal paper that introduced simplicial tree-decompositions to prove Theorem 7.3.4; cf. Exercise 21.) Simplicial tree-decompositions are treated in depth in R. Diestel, *Graph Decompositions*, Oxford University Press 1990.

An instructive introductory survey on tree-width, brambles and tangles is given by B.A. Reed in (R.A. Bailey, ed) *Surveys in Combinatorics 1997*, Cambridge University Press 1997, 87–162. Reed introduced the term 'bramble'; in Seymour & Thomas's original paper they are called 'screens'.

Theorem 12.3.7 is extracted from J. Carmesin, R. Diestel, F. Hundertmark & M. Stein, Connectivity and tree structure in finite graphs, *Combinatorica* **34** (2014), 1–35, arXiv:1105.1611. Under mild additional assumptions one can show that the tree-decompositions constructed for the proof of Theorem 12.3.7 refine each other as k grows: the decomposition for $k+1$ induces tree-decompositions of the torsos of the decomposition for k and is therefore compatible with that decomposition. Just as in Theorems 12.5.1 and 12.5.8, one thus obtains one overall tree-decomposition whose induced separations separate every two blocks that can be separated at all, i.e., that are not just some k-block contained in a larger ℓ-block (for $\ell < k$).

The tree-width duality theorem, Theorem 12.4.3, is due to P.D. Seymour and R. Thomas, Graph searching and a min-max theorem for tree-width, *J. Comb. Theory, Ser. B* **58** (1993), 22–33. A short version of this proof was included in earlier editions of this book and can be found in P. Bellenbaum & R. Diestel, Two short proofs concerning tree-decompositions, *Comb. Probab. Comput.* **11** (2002), 541–547 (which also offers a short proof of Theorem 12.4.5). The proof presented in the text follows an idea of F. Mazoit, personal communication 2013. The simplest proof, perhaps – and the only one not using Menger's theorem – is via Theorem 12.5.11; see Exercises 59–60 and their hints.

Historically, tree-width duality evolved with a few quirks. As Robertson and Seymour developed the theory of tree-decompositions, they simultaneously looked for witnesses to large tree-width, as a way to proceed with the proof of the graph minor theorem when the graphs in question have unbounded tree-width. The result of this search was the notion of a tangle – with hindsight, perhaps the deepest single innovation for graph theory stemming from this proof. Numerically, however, the duality did not exactly fit: while large tree-width implies the existence of a large-order tangle and vice versa, one loses a small constant factor in the conversion. Instead of adjusting the notion of a tangle to repair this, however (e.g., as in Theorem 12.5.11), Robertson and Seymour simply changed the notion of a tree-decomposition to a new concept called *branch-decompositions*, which are exactly dual to tangles (except for very small k). To tie up the loose ends, Seymour and Thomas later introduced brambles and Theorem 12.4.3 to provide exact duality for tree-width too; but brambles, though interesting, never assumed the significance of tangles.

Theorem 12.4.5 is from R. Thomas, A Menger-like property of tree-width; the finite case, *J. Comb. Theory, Ser. B* **48** (1990), 67–76. Theorem 12.4.6 is from R. Diestel & M. Müller, Connected tree-width, *Combinatorica* **38** (2018), 381–398, arXiv:1211.7353. This paper also includes a proof that $\mathrm{ctw}(C) \leqslant \mathrm{ctw}(G)$ if C is a geodesic cycle in G.

In older papers on tangles, as well as in the 5th edition of this book, the partial ordering between oriented separations is the reverse of that defined in

Section 12.5. Our new definition seems more natural, since $\leqslant$ implies $\subseteq$ for the sides to which the separations point: if $(A,B) \leqslant (C,D)$ then $B \subseteq C$. This is also better compatible with the tangle theory of set bipartitions, where it is customary to refer to an oriented partition (A,B) simply as B (since A is determined as $A = V \setminus B$); see the book reference below for more on such tangles and their applications.

Profiles more general than tangles are studied in R. Diestel, F. Hundertmark & S. Lemanczyk, Profiles of separations: in graphs, matroids, and beyond, *Combinatorica* **39,** 37–75. This paper gave the first canonical proof of the tree-of-tangles theorem, Theorem 12.5.1. The tree-of-tangles theorem it proves for profiles of so-called abstract separation systems also implies Theorem 12.3.7 and Exercise 57, since blocks and edge-tangles induce profiles. Indeed this is how they came by their name: as the 'profiles' of blocks visible on the screen of the low-order separations of a graph, which they orient.

Our first proof of the tree-of-tangles theorem, and in particular the splinter lemma on which it is based, are due to C. Elbracht, J. Kneip & M. Teegen, Trees of tangles in abstract separation systems, *J. Comb. Theory, Ser. A* **180** (2021), arXiv:1909.09030. Its canonical strengthening, Theorem 12.5.8, is due to J. Carmesin & J. Kurkofka, Entanglements, *J. Comb. Theory, Ser. B* **164** (2024), 17–28, arXiv:2205.11488. This paper also give examples of entanglements that are not of the form $D(\tau, \tau')$. Thus, Theorem 12.5.8 is also more general than Theorem 12.5.1, not only stronger.

Our proof of Theorem 12.5.9 is adapted from R. Diestel & S. Oum, Tangle-tree duality in abstract separation systems, *Adv. Math.* **377** (2021), 107470; arXiv:1701.02509. In this paper, a duality theory is developed for tangles in abstract separation systems, not necessarily of graphs. Its main result contains Theorems 12.5.9 and 12.5.11 as special cases.

The theory of tangles in graphs, including its main two theorems, has been extended to more general combinatorial structures such as matroids or set partitions. In this general form it can be applied outside mathematics, in areas as diverse as clustering in data analysis, finding mindsets in political science or psychology, or consumer behaviour in economics. This is explored in R. Diestel, *Tangles: a structural approach to artificial intelligence in the empirical sciences*, Cambridge University Press 2024. Excerpts, an electronic edition, and open-source tangle software are available from `tangles-book.com`.

The Kuratowski set for the graphs of tree-width < 4 have been determined by S. Arnborg, D.G. Corneil and A. Proskurowski, Forbidden minors characterization of partial 3-trees, *Discrete Math.* **80** (1990), 1–19. They are: K^5, the octahedron $K_{2,2,2}$, the 5-prism $C^5 \times K^2$, and the Wagner graph W. The Kuratowski set $\mathcal{K}_{\mathcal{P}(S)}$ for a given surface S has been determined explicitly for only one surface other than the sphere, the projective plane. It consists of 35 forbidden minors; see D. Archdeacon, A Kuratowski theorem for the projective plane, *J. Graph Theory* **5** (1981), 243–246. It is not difficult to show that $|\mathcal{K}_{\mathcal{P}(S)}|$ grows rapidly with the genus of S (Exercise 67).

A survey of finite forbidden minor theorems is given in Chapter 6.1 of R. Diestel, *Graph Decompositions*, Oxford University Press 1990. More recent developments are surveyed in R. Thomas, Recent excluded minor theorems, in (J.D. Lamb & D.A. Preece, eds) *Surveys in Combinatorics 1999*, Cambridge University Press 1999, 201–222. A survey of infinite forbidden minor theorems was given by N. Robertson, P.D. Seymour & R. Thomas, Excluding infinite minors, *Discrete Math.* **95** (1991), 303–319.

The first short proof of the grid theorem, Theorem 12.6.3, was given by R. Diestel, K.Yu. Gorbunov, T.R. Jensen & C. Thomassen, Highly connected sets and the excluded grid theorem, *J. Comb. Theory, Ser. B* **75** (1999), 61–73. This proof was included in editions 2–4 of this book. It was further simplified by A. Leaf and P.D. Seymour, Treewidth and planar minors, *J. Comb. Theory, Ser. B* **111** (2015) 38–53. The first proof with polynomial bound was obtained by C. Chekuri and J. Chuzhoy, Polynomial bounds for the grid-minor theorem, *J. ACM* **63** (2016), 1–65; arXiv:1602.02629.

As a forerunner to the grid theorem, 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, *Comb. Probab. Comput.* **4** (1995), 27–30. It also follows from the abstract tangle duality theorem of Diestel and Oum cited earlier.

Theorem 12.6.6 is the earliest version of Robertson and Seymour's structure theorem for the graphs without a K^n minor. It has become known as the 'red herring' version – a phrase coined by Robertson and Seymour themselves, referring to its role in their proof of the graph minor theorem. It nonetheless remains the most-often applied version of the structure theorem, especially in algorithmic contexts. The strongest version so far, designed with future applications in mind, is given in R. Diestel, K. Kawarabayashi, Th. Müller & P. Wollan, On the excluded minor structure theorem for graphs of large tree-width, *J. Comb. Theory, Ser. B* **102** (2012), 1189–1210, arXiv:0910.0946. Its proof is based on Theorem 12.6.6. A short proof of Theorem 12.6.6 itself was recently given by K. Kawarabayashi, R. Thomas & P. Wollan, arXiv:1207.6927 and arXiv:2010.12397.

The structure Theorem 12.6.7 for excluding topological minors is due to M. Grohe and D. Marx, Structure theorem and isomorphism test for graphs with excluded topological subgraphs, *Proc. 44th ann. ACM symp. theory of computing* (STOC 2012), 173–192, arXiv:1111.1109.

The existence of normal spanning trees for graphs with no topological $K^{\aleph_0}$ minor was proved by R. Halin, Simplicial decompositions of infinite graphs, in: (B. Bollobás, ed.) *Advances in Graph Theory, Annals of Discrete Mathematics* **3**, North-Holland 1978. Its strengthening, part (iii) of Theorem 12.6.9, was observed in R. Diestel, The depth-first search tree structure of $TK_{\aleph_0}$-free graphs, *J. Comb. Theory, Ser. B* **61** (1994), 260–262. Part (iii) easily implies part (ii), which had been proved independently by N. Robertson, P.D. Seymour & R. Thomas, Excluding infinite clique subdivisions, *Trans. Amer. Math. Soc.* **332** (1992), 211–223. Theorem 12.6.8 and the infinite case of Theorem 12.6.6 were proved in R. Diestel & R. Thomas, Excluding a countable clique, *J. Comb. Theory, Ser. B* **76** (1999), 41–67. The proof of Theorem 12.6.8 builds on the main result of N. Robertson, P.D. Seymour & R. Thomas, Excluding infinite clique minors, *Mem. Amer. Math. Soc.* **118** (1995).

Our proof of the 'generalized Kuratowski theorem', Corollary 12.7.3, was inspired by J. Geelen, B. Richter & G. Salazar, Embedding grids in surfaces, *Eur. J. Comb.* **25** (2004), 785–792. An alternative proof, which bypasses Theorem 12.4.2 by proving directly that the graphs in $\mathcal{K}_{\mathcal{P}(S)}$ have bounded order, is given by B. Mohar & C. Thomassen, *Graphs on Surfaces*, Johns Hopkins University Press 2001. Mohar (see there) also developed a set of algorithms, one for each surface, that decide embeddability in that surface in linear time. As a corollary, he obtains an independent and constructive proof of Corollary 12.7.3.

For every graph X, Graph Minors XIII gives an explicit algorithm that decides in cubic time for every input graph G whether $X \preccurlyeq G$. The constants in the cubic polynomials bounding the running time of these algorithms depend on X but are constructively bounded from above.

The concept of a 'good characterization' of a graph property was first suggested by J. Edmonds, Minimum partition of a matroid into independent subsets, *J. Research of the National Bureau of Standards (B)* **69** (1965) 67–72. In the language of complexity theory, a characterization is *good* if it specifies two assertions about a graph such that, given any graph G, the first assertion holds for G if and only if the second fails, and such that each assertion, if true for G, provides a certificate for its truth that can be checked in polynomial time. Thus every good characterization has the corollary that the decision problem corresponding to the property it characterizes lies in NP $\cap$ co-NP.