[8.2.3] [8.6.8] **Lemma 1.5.5.** *Let T be a normal tree in G.*

(i) *Any two vertices $x, y \in T$ are separated in G by the set $\lceil x \rceil \cap \lceil y \rceil$.*

(ii) *If $S \subseteq V(T) = V(G)$ and S is down-closed, then the components of $G - S$ are spanned by the sets $\lfloor x \rfloor$ with x minimal in $T - S$.*

Proof. (i) Let P be any x–y path in G; we show that P meets $\lceil x \rceil \cap \lceil y \rceil$. Let $t_1, \dots, t_n$ be a minimal sequence of vertices in $P \cap T$ such that $t_1 = x$ and $t_n = y$ and t_i and t_{i+1} are comparable in the tree-order of T for all i. (Such a sequence exists: the set of *all* vertices in $P \cap T$, in their natural order as they occur on P, has this property because T is normal and every segment $t_i P t_{i+1}$ is either an edge of T or a T-path.) In our minimal sequence we cannot have $t_{i-1} < t_i > t_{i+1}$ for any i, since t_{i-1} and t_{i+1} would then be comparable, and deleting t_i would yield a smaller such sequence. Thus, our sequence has the form

$$x = t_1 > \ldots > t_k < \ldots < t_n = y$$

for some $k \in \{1, \dots, n\}$. As $t_k \in \lceil x \rceil \cap \lceil y \rceil \cap V(P)$, our proof is complete.

(ii) Consider a component C of $G - S$, and let x be a minimal element of its vertex set. Then $V(C)$ has no other minimal element x': as x and x' would be incomparable, any x–x' path in C would by (i) contain a vertex below both, contradicting their minimality in $V(C)$. Hence as every vertex of C lies above some minimal element of $V(C)$, it lies above x. Conversely, every vertex $y \in \lfloor x \rfloor$ lies in C, for since S is down-closed, the ascending path xTy lies in $T - S$. Thus, $V(C) = \lfloor x \rfloor$.

Let us show that x is minimal not only in $V(C)$ but also in $T - S$. The vertices below x form a chain $\lceil t \rceil$ in T. As t is a neighbour of x, the maximality of C as a component of $G - S$ implies that $t \in S$, giving $\lceil t \rceil \subseteq S$ since S is down-closed. This completes the proof that every component of $G - S$ is spanned by a set $\lfloor x \rfloor$ with x minimal in $T - S$.

Conversely, if x is any minimal element of $T - S$, it is clearly also minimal in the component C of $G - S$ to which it belongs. Then $V(C) = \lfloor x \rfloor$ as before, i.e., $\lfloor x \rfloor$ spans this component. $\square$

Normal spanning trees are also called *depth-first search trees*, because of the way they arise in computer searches on graphs (Exercise 26). This fact is often used to prove their existence, which can also be shown by a very short and clever induction (Exercise 25). The following constructive proof, however, illuminates better how normal trees capture the structure of their host graphs.

[6.5.3] [8.2.4] **Proposition 1.5.6.** *Every connected graph contains a normal spanning tree, with any specified vertex as its root.*

1.7 Contraction and minors

In Section 1.1 we saw two fundamental containment relations between graphs: the 'subgraph' relation, and the 'induced subgraph' relation. In this section we meet two more: the 'minor' relation, and the 'topological minor' relation. Let X be a fixed graph.

A *subdivision* of X is, informally, any graph obtained from X by 'subdividing' some or all of its edges by drawing new vertices on those edges. In other words, we replace some edges of X with new paths between their ends, so that none of these paths has an inner vertex in $V(X)$ or on another new path. When G is a subdivision of X, we also say that G *is a* TX.[7] The original vertices of X are the *branch vertices of the* TX; its new vertices are called *subdividing vertices*. Note that subdividing vertices have degree 2, while branch vertices retain their degree from X.

subdivision TX *of* X

branch vertices

If a graph Y contains a TX as a subgraph, then X is a *topological minor* of Y (Fig. 1.7.1).

topological minor

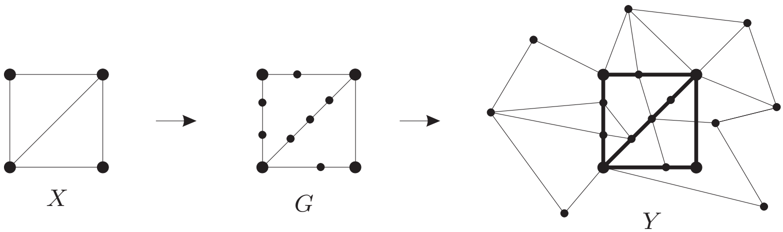

Fig. 1.7.1. The graph G is a TX, a *subdivision* of X.
As $G \subseteq Y$, this makes X a *topological minor* of Y.

Similarly, replacing the vertices x of X with disjoint connected graphs G_x, and the edges xy of X with non-empty sets of G_x–G_y edges, yields a graph that we shall call an IX.[8] More formally, a graph G *is an* IX if its vertex set admits a partition $\{ V_x \mid x \in V(X) \}$ into connected subsets V_x such that distinct vertices $x, y \in X$ are adjacent in X if and only if G contains a V_x–V_y edge. The sets V_x are the *branch sets of the* IX. Conversely, we say that X arises from G by *contracting* the subgraphs G_x and call it a *contraction minor* of $\cancel{Y}$.

IX

branch sets

contraction minor

If a graph Y contains an IX as a subgraph, then X is a *minor* of Y, the IX is a *model of* X in Y, and we write $X \preccurlyeq Y$ (Fig. 1.7.2).

minor, $\preccurlyeq$ *model*

[7] The 'T' stands for 'topological'. Although, formally, TX denotes a whole class of graphs, the class of all subdivisions of X, it is customary to use the expression as indicated to refer to an arbitrary member of that class.

[8] The 'I' stands for 'inflated'. As before, while IX is formally a class of graphs, those admitting a vertex partition $\{ V_x \mid x \in V(X) \}$ as described below, we use the expression as indicated to refer to an arbitrary member of that class.

Exercises

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

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

   (Hint for the circumference: induction on d.)

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

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

5. Let v_0 be a vertex in a graph G, and $D_0 := \{v_0\}$. For $n = 1, 2, \ldots$ inductively define $D_n := N_G(D_0 \cup \ldots \cup D_{n-1})$. Show that $D_n = \{v \mid d(v_0, v) = n\}$ and $D_{n+1} \subseteq N(D_n) \subseteq D_{n-1} \cup D_{n+1}$ for all $n \in \mathbb{N}$.

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

7. Prove the weakening of Theorem 1.3.4 obtained by replacing average with minimum degree. Deduce that $|G| \geqslant n_0(d/2, g)$ for every graph G as given in the theorem.

8. Show that graphs of girth at least 5 and order n have a minimum degree of $o(n)$. In other words, show that there is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that $f(n)/n \to 0$ as $n \to \infty$ and $\delta(G) \leqslant f(n)$ for all such graphs G.

9.$^+$ Show that every connected graph G contains a path or cycle of length at least $\min\{2\delta(G), |G|\}$.

10. Show that a connected graph of diameter k and minimum degree d has at least about $kd/3$ vertices but need not have substantially more.

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

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

13. Determine $\kappa(G)$ and $\lambda(G)$ for $G = P^m, C^n, K^n, K_{m,n}$ and the d-dimensional cube (Exercise 2); $d, m, n \geqslant 3$.

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

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

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

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

   Show that the statement

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

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

A. Treglown, Proof of the 1-factorization and Hamilton decomposition conjectures, *Memoirs of the AMS* (to appear).

List colourings were first introduced in 1976 by Vizing. Among other things, Vizing proved the list-colouring equivalent of Brooks's theorem. Voigt (1993) constructed a plane graph of order 238 that is not 4-choosable; thus, Thomassen's list version of the five colour theorem is best possible. A stimulating survey on the list-chromatic number and how it relates to the more classical graph invariants (including a proof of Theorem 5.4.1) is given by N. Alon, Restricted colorings of graphs, in (K. Walker, ed.) *Surveys in Combinatorics*, LMS Lecture Notes **187**, Cambridge University Press 1993. Both the list colouring conjecture and Galvin's proof of the bipartite case are originally stated for multigraphs. Kahn (1994) proved that the conjecture is asymptotically correct, as follows: given any $\epsilon > 0$, every graph G with large enough maximum degree satisfies $\mathrm{ch}'(G) \leqslant (1+\epsilon)\Delta(G)$.

The total colouring conjecture (Exercise **32**) was proposed around 1965 by Vizing and by Behzad; see Jensen & Toft for details.

A gentle introduction to the basic facts about perfect graphs and their applications is given by M.C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press 1980. A more comprehensive treatment is given in A. Schrijver, *Combinatorial optimization*, Springer 2003. Surveys on various aspects of perfect graphs are included in *Perfect Graphs* by J. Ramirez-Alfonsin & B. Reed (eds.), Wiley 2001. Our first proof of the perfect graph theorem, Theorem 5.5.4, follows Lovász's survey on perfect graphs in (L.W. Beineke and R.J. Wilson, eds.) *Selected Topics in Graph Theory 2*, Academic Press 1983. Our second proof, the proof of Theorem 5.5.6, is due to G.S. Gasparian, Minimal imperfect graphs: a simple approach, *Combinatorica* **16** (1996), 209–212.

Theorem 5.5.3 is proved in M. Chudnovsky, N. Robertson, P.D. Seymour and R. Thomas, The strong perfect graph theorem, *Ann. Math.* **164** (2006), 51–229, arXiv:math/0212070. This proof is elucidated by N. Trotignon in his 2013 survey on the arXiv:1301.5149, which also offers a short account of Lovász's proof of the (weak) perfect graph theorem. Chudnovsky, Cornuejols, Liu, Seymour and Vušković, Recognizing Berge graphs, *Combinatorica* **25** (2005), 143–186, constructed an $O(n^9)$ algorithm testing for odd holes and antiholes, and thus by the strong perfect graph theorem also for perfection.

Gýarfás's conjecture on χ-boundedness that prompted Theorem 5.5.7 is from A. Gyárfás, Problems from the world surrounding perfect graphs, Proceedings of the International Conference on Combinatorial Analysis and its Applications (Pokrzywna, 1985), *Zastos. Mat.* **19** (1987), 413–441. Part (i) of the theorem is due to A. Scott and P.D. Seymour, arXiv:1410.4118, while part (ii) was proved by M. Chudnovsky and these authors in arXiv:1506.02232.

The structure of graphs forced by forbidding some fixed induced subgraph or subgraphs, as in the strong perfect graph theorem, has been studied more widely since the proof of that theorem left a community of experts without an overriding goal. One of the central problems now studied is the *Erdős-Hajnal conjecture* that the graphs without some fixed induced subgraph ~~have linearly large~~ sets of vertices that are either independent or induce a complete subgraph. See Chapter 9.1 for a precise statement.

vertices outside $U \cup W$ that are not adjacent to u, and let V_2 contain the remaining vertices. As u is isolated in $G[V_1]$, we have $G \not\simeq G[V_1]$ and hence $G \simeq G[V_2]$. By the minimality of $|U \cup W|$, there is a vertex $v \in G[V_2] - U - W$ that is adjacent to every vertex in $U \setminus \{u\}$ and to none in W. But v is also adjacent to u, because it lies in V_2. So U, W and v satisfy $(*)$ for G, contrary to assumption.

Finally, assume that $U = \emptyset$. Then $W \neq \emptyset$. Pick $w \in W$, and consider the partition $\{V_1, V_2\}$ of V where V_1 consists of w and all its neighbours outside W. As before, $G \not\simeq G[V_1]$ and hence $G \simeq G[V_2]$. Therefore U and $W \setminus \{w\}$ satisfy $(*)$ in $G[V_2]$, with $v \in V_2 \setminus W$ say, and then U, W, v satisfy $(*)$ in G. □

Another indication of the high degree of uniformity in the structure of the Rado graph is its large automorphism group. For example, R is easily seen to be *vertex-transitive*: given any two vertices x and y, there is an automorphism of R mapping x to y.

homoge-neous

In fact, much more is true: using the back-and-forth technique, one can easily show that the Rado graph is *homogeneous*: every isomorphism between two finite induced subgraphs can be extended to an automorphism of the entire graph (Exercise 50).

R^r

Which other countable graphs are homogeneous? The complete graph $K^{\aleph_0}$ and its complement are again obvious examples. Moreover, for every integer $r \geqslant 3$ there is a homogeneous K^r-free graph R^r, constructed as follows. Let $R^r_0 := K^1$, and let R^r_{n+1} be obtained from R^r_n by joining, for every ~~subgraph~~ $H \not\supseteq K^{r-1}$ of R^r_n, a new vertex v_H to every vertex in H. Then let $R^r := \bigcup_{n \in \mathbb{N}} R^r_n$. Clearly, as the new vertices v_H of R^r_{n+1} are independent, there is no K^r in R^r_{n+1} if there was none in R^r_n, so $R^r \not\supseteq K^r$ by induction on n. Just like the Rado graph, R^r is clearly universal among the K^r-free countable graphs, and by the back-and-forth argument from the proof of Theorem 8.3.1 it is easily seen to be homogeneous.

By the following deep theorem of Lachlan and Woodrow, the countable homogeneous graphs we have seen so far are essentially all:

Theorem 8.3.3. (Lachlan & Woodrow 1980)
Every countably infinite homogeneous graph is one of the following:

- *a disjoint union of complete graphs of the same order, or the complement of such a graph;*
- *the graph R^r or its complement, for some $r \geqslant 3$;*
- *the Rado graph R.*

To conclude this section, let us return to our original problem: for which graph properties is there a graph that is universal with this property? Most investigations into this problem have addressed it from a

Theorem 8.5.3. *Every countable rayless graph G has an unfriendly partition.*

α

Proof. To help with our formal notation, we shall think of a partition of a set V as a map $\pi\colon V \to \{0,1\}$. We apply induction on the rank of G. When this is zero then G is finite, and an unfriendly partition can be obtained by maximizing the number of edges across the partition. Suppose now that G has rank $\alpha > 0$, and assume the theorem as true for graphs of smaller rank.

U
$C_0, C_1, \ldots$
U_0, U_1, U_2

Let U be a finite set of vertices in G such that each of the components $C_0, C_1, \ldots$ of $G - U$ has rank $< \alpha$. Partition U into the set U_0 of vertices that have finite degree in G, the set U_1 of vertices that have infinitely many neighbours in some C_n, and the set U_2 of vertices that have infinite degree but only finitely many neighbours in each C_n.

$G_0, G_1, \ldots$
$n_0, n_1, \ldots$
π

For every $n \in \mathbb{N}$ let $G_n := G[U \cup V(C_0) \cup \ldots \cup V(C_n)]$. This is a graph of some rank $\alpha_n < \alpha$, so by induction it has an unfriendly partition π_n. Each of these π_n induces a partition of U. Let π_U be a partition of U induced by π_n for infinitely many n, say for $n_0 < n_1 < \ldots$. Choose n_0 large enough that G_{n_0} contains all the neighbours of vertices in U_0, and the other n_i large enough that every vertex in U_2 has more neighbours in $G_{n_i} - G_{n_{i-1}}$ than in $G_{n_{i-1}}$, for all $i > 0$. Let π be the partition of G defined by letting $\pi(v) := \pi_{n_i}(v)$ for all $v \in G_{n_i} - G_{n_{i-1}}$ and all i, where $G_{n_{-1}} := \emptyset$. Note that $\pi|_U = \pi_{n_0}|_U = \pi_U$.

Let us show that π is unfriendly. We have to check that every vertex is *happy with* π, i.e., that it has at least as many neighbours in the opposite class under π as in its own.[8] To see that a vertex $v \in G - U$ is happy with π, let i be minimal such that $v \in G_{n_i}$ and recall that v was happy with π_{n_i}. As both v and its neighbours in G lie in $U \cup V(G_{n_i} - G_{n_{i-1}})$, and π agrees with π_{n_i} on this set, v is happy also with π. Vertices in U_0 are happy with π, because they were happy with π_{n_0}, and π agrees with π_{n_0} on U_0 and all its neighbours. Vertices in U_1 are also happy. Indeed, every $u \in U_1$ has infinitely many neighbours in some C_n, and hence in some $G_{n_i} - G_{n_{i-1}}$. Then u has infinitely many opposite neighbours in $G_{n_i} - G_{n_{i-1}}$ under π_{n_i}. Since π_{n_i} agrees with π on both U and $G_{n_i} - G_{n_{i-1}}$, our vertex u has infinitely many opposite neighbours also under π. Vertices in U_2, finally, are happy with every π_{n_i}. By our choice of n_i, at least one of their opposite neighbours under π_{n_i} must lie in $G_{n_i} - G_{n_{i-1}}$. Since π_{n_i} agrees with π on both U_2 and $G_{n_i} - G_{n_{i-1}}$, this gives every $u \in U_2$ at least one opposite neighbour under π in every $G_{n_i} - G_{n_{i-1}}$. Hence u has infinitely many opposite neighbours under π, which clearly makes it happy. □

[8] It is only by tradition that such partitions are called ‘unfriendly’; our vertices love them.

30. Show that a graph G contains a $TK^{\aleph_0}$ if and only if some end of G is dominated by infinitely many vertices.

31.$^+$ Let G be a *finitely separable* graph, one in which any two vertices can be separated by finitely many edges.

    (i) Show that any two ~~rays~~ in G that cannot be separated by finitely many edges are dominated by a common vertex.

    (ii) Is the assumption of finite separability necessary for (i) to hold?

32. Construct a countable graph with uncountably many thick ends. Can you find a locally finite such graph?

33. Show that a locally finite connected vertex-transitive graph has exactly 0, 1, 2 or infinitely many ends.

34.$^+$ Show that the automorphisms of a graph $G = (V, E)$ act naturally on its ends, i.e., that every automorphism $\sigma\colon V \to V$ can be extended to a map $\sigma\colon \Omega(G) \to \Omega(G)$ such that $\sigma(R) \in \sigma(\omega)$ whenever R is a ray in an end ω. Prove that, if G is connected, every automorphism σ of G fixes a finite set of vertices or an end. If σ fixes no finite set of vertices, can it fix more than one end? More than two?

35.$^-$ Show that a locally finite spanning tree of a graph G contains a ray from every end of G.

36. A ray in a graph *follows* another ray if the two have infinitely many vertices in common. Show that if T is a normal spanning tree of G then every ray of G follows a unique normal ray of T.

37. Use normal spanning trees to show that a countable connected graph has either countably many or continuum many ends.

38. Show that the following assertions are equivalent for connected countable graphs G.

    (i) G has a locally finite spanning tree.

    (ii) For no finite separator $X \subseteq V(G)$ in G does $G - X$ have infinitely many components.

39. Show that every (countable) planar 3-connected graph has a locally finite spanning tree.

40. Prove the following infinite version of the Erdős-Pósa theorem: an infinite graph G either contains infinitely many disjoint cycles or it has a finite set Z of vertices such that $G - Z$ is a forest.

41. Let G be a connected graph. Call a set $U \subseteq V(G)$ *dispersed* if every ray in G can be separated from U by a finite set of vertices. (In the topology of Section 8.6, these are precisely the closed subsets of $V(G)$.)

    (i) Prove that G has a normal spanning tree if and only if $V(G)$ is a countable union of dispersed sets.

    (ii) Deduce that if G has a normal spanning tree then so does every connected minor of G.

(ii) $V_i \subseteq V_{i-1} \smallsetminus \{v_{i-1}\}$ $(i = 2, \dots, 2r-2)$;

(iii) v_{i-1} is adjacent either to all vertices in V_i or to no vertex in V_i $(i = 2, \dots, 2r-2)$.

Let $V_1 \subseteq V(G)$ be any set of 2^{2r-3} vertices, and pick $v_1 \in V_1$ arbitrarily. Then (i) holds for $i = 1$, while (ii) and (iii) hold trivially. Suppose now that V_{i-1} and $v_{i-1} \in V_{i-1}$ have been chosen so as to satisfy (i)–(iii) for $i-1$, where $1 < i \leqslant 2r-2$. Since

$$|V_{i-1} \smallsetminus \{v_{i-1}\}| = 2^{2r-1-i} - 1$$

is odd, V_{i-1} has a subset V_i satisfying (i)–(iii); we pick $v_i \in V_i$ arbitrarily.

Among the $2r-3$ vertices $v_1, \dots, v_{2r-3}$, there are $r-1$ vertices that show the same behaviour when viewed as v_{i-1} in (iii), being adjacent either to all the vertices in V_i or to none. Accordingly, these $r-1$ vertices and v_{2r-2} induce either a K^r or a $\overline{K^r}$ in G, because $v_i, \dots, v_{2r-2} \in V_i$ for all i. □

Ramsey number $R(r)$

The least integer n associated with r as in Theorem 9.1.1 is the *Ramsey number* $R(r)$ of r; our proof shows that $R(r) \leqslant 2^{2r-3}$. In Chapter 11 we shall use a simple probabilistic argument to show that $R(r)$ is bounded below by $2^{r/2}$ (Theorem 11.1.3).

Erdős-Hajnal conjecture

In other words, the largest clique or independent set of vertices that a graph of order n must contain is, asymptotically, logarithmically small in n. As soon as we forbid some fixed induced subgraph, however, it may be much bigger, ~~of size linear in n~~: The *Erdős-Hajnal conjecture* says that for every graph H there exists a constant $\delta_H > 0$ such that every graph G not containing an induced copy of H has a set of at least $|G|^{\delta_H}$ vertices that are either independent or span a complete subgraph in G.

c-colouring

$[X]^k$

monochromatic

It is customary in Ramsey theory to think of partitions as colourings: a *colouring* of (the elements of) a set X *with c colours*, or *c-colouring* for short, is simply a partition of X into c classes (indexed by the 'colours'). In particular, these colourings need not satisfy any non-adjacency requirements as in Chapter 5. Given a c-colouring of $[X]^k$, the set of all k-subsets of X, we call a set $Y \subseteq X$ *monochromatic* if all the elements of $[Y]^k$ have the same colour,[1] i.e. belong to the same of the c partition classes of $[X]^k$. Similarly, if $G = (V, E)$ is a graph and all the edges of $H \subseteq G$ have the same colour in some colouring of E, we call H a *monochromatic subgraph* of G, speak of a red (green, etc.) H in G, and so on.

In the above terminology, Ramsey's theorem can be expressed as follows: for every r there exists an n such that, given any n-set X,

---

[1] Note that Y is called monochromatic, but it is the elements of $[Y]^k$, not of Y, that are (equally) coloured.

Indeed, 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 $|N(X)| \leqslant k$ 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 and Menger's Theorem 3.3.1, there is a family $\{ P_s \mid s \in S \}$ of disjoint S–$N(X)$ paths in $G[A]$ and a family $\{ Q_s \mid s \in S \}$ of disjoint S–$N(Y)$ paths in $G[B]$. A, B S $P_s,\ Q_s$

Let H be the minor of G obtained by deleting $A \setminus \bigcup_{s \in S} V(P_s)$ 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 the part $V_t^1 \in \mathcal{V}_1$ say, its part V_t in $\mathcal{V}_1'$ is H

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

(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 part corresponding to y. J

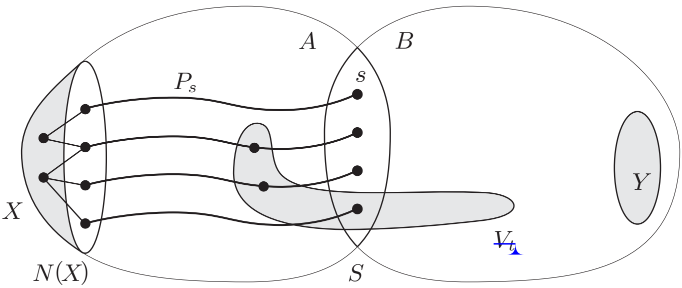

Fig. 12.4.1. To obtain V_t from $V_t \cap B$, add two vertices from S

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 part 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 $(*)$. T 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 part V_t in $\mathcal{V}$ is no tree 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, sometimes, the only feature of a highly connected substructure that we really care about is 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 tree-structure and highly connected substructures more generally.

The *orientations* of a separation $\{A, B\}$ of G are the two *oriented separations* (A, B) and (B, A). We say that (A, B) is *oriented*, or *pointing, towards* B and its subsets. Given a set S of separations, we write $\vec{S} := \big\{(A, B) \mid \{A, B\} \in S\big\}$ for the set of all their orientations. An *orientation of* S is a subset O of $\vec{S}$ that contains for every element of S exactly one of its two orientations. We say that O *avoids* a collection $\mathcal{F}$ of sets of oriented separations if no subset of O lies in $\mathcal{F}$.

oriented separation

$\vec{S}$

avoids

Given oriented separations (A, B) and (C, D) of G, let us write $(A, B) \leqslant (C, D)$ if $A \subseteq C$ and $B \supseteq D$. A set σ of oriented separations of G is *consistent* if it does *not* contain (B, A) whenever $(A, B) \leqslant (C, D)$ with $(C, D) \in \sigma$.[4] And σ is a *star* of oriented separations if $(A, B) \leqslant (D, C)$ for all distinct $(A, B), (C, D) \in \sigma$ (Fig. 12.5.1).

$\leqslant$

consistent

star

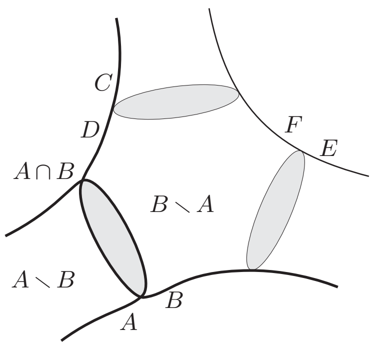

Fig. 12.5.1. The separations $(A, B), (C, D), (E, F)$ form a star

For example, if $(T, \mathcal{V})$ is a tree-decomposition of G with $\mathcal{V} = (V_t)_{t \in T}$, then orienting the separations induced by the edges of T towards some fixed V_t orients them consistently. And the separations corresponding to

[4] 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.