Graph Minors: Important First Steps and Generalizing Kuratowski's Theorem

R.B. Richter
University of Waterloo

October 16, 2002

Abstract

The Graph Minors Project of Neil Robertson and Paul Seymour proves that if $\{G_1, G_2, \ldots\}$ is an infinite set of graphs, then there exist i and j such that G_i is a minor of G_j. In particular, if Σ is a surface, the set $\mathcal{F}(\Sigma)$ of graphs G that do not embed in Σ but every proper minor of which does embed in Σ must be finite.

Two of the central results of the entire Graph Minors Project are: (1) if the graphs in the set $\{G_1, G_2, \ldots\}$ have bounded tree-width, then there is an i and a j such that G_i is a minor of G_j; and (2) for every integer k, there is an integer $f(k)$ such that if G has tree-width at least $f(k)$, then G has a $k \times k$ grid as a minor. We present the "simple" proofs of these two facts, the first due to Geelen, Gerards and Whittle and the second due to Diestel, Gorbunov, Jensen and Thomassen.

Finally, we give a version of an argument of Thomassen that shows how to apply (1) and (2) to prove the finiteness of $\mathcal{F}(\Sigma)$.

1 Introduction

This work is intended to be a self-contained introduction to Graph Minors. It contains a complete proof of the generalization of Kuratowski's Theorem to higher surfaces, but, more importantly, it is a major step in understanding the whole Graph Minors project. The only background a reader should need is some familiarity with connectivity issues (essentially variations of Menger's Theorem and a willingness to view cutsets from different perspectives). My own experience with the arguments presented here is that we need to be able to focus both on the big picture and on the details. There are many small points that require their own little arguments and I have attempted to provide all of those in sufficient detail to make it easier for the reader to not lose sight of the big picture.

If $\leq$ is a transitive, reflexive relation on a set S, then $\leq$ is a *well-quasi-order* if there is no infinite descending chain $s_1 > s_2 > s_3 > \dots$ and there is no infinite antichain, i.e., an infinite subset S' of S such that, for all $s, s' \in S'$, neither $s \leq s'$ nor $s' \leq s$. An equivalent formulation is that if $s_1, s_2, \dots$ is an infinite sequence of elements of S, then there exist infinitely many indices $i_1 < i_2 < \dots$ such that $s_{i_1} \leq s_{i_2} \leq \dots$.

A graph (all graphs in this work are finite) H is a *minor* of a graph G if an isomorph of H can be obtained from a subgraph of G by contracting edges. The graph H is a *proper minor* of G if H is a minor of G and H is not isomorphic to G; we write $H <_m G$ to denote that H is a proper minor of G.

Generally, it is slightly more convenient for us to work with sets rather than sequences. If $<$ is a transitive irreflexive relation on a set S (such as $<_m$ on the set of graphs) and T is an infinite subset of S, then it is an easy consequence of Ramsey's Theorem for infinite graphs (c.f., [1]) that at least one of the following happens: T contains an infinite descending chain $t_1 > t_2 > t_3 > \dots$; T contains an infinite antichain, i.e., an infinite set of pairwise incomparable elements; or T contains an infinite ascending chain $t_1 < t_2 < t_3 < \dots$. In the context of $<_m$, there are no infinite descending chains, so every infinite set of graphs either contains an infinite antichain or an infinite ascending chain.

The following theorem is the main result of the Graph Minors Project of Neil Robertson and Paul Seymour (the most recently published installment in the series is [8]; Theorem 1.1 is proved in [9]). It seems the first mention of this problem is in [14], although Diestel reports that Wagner never positively conjectured this result [2]. Nevertheless, it has become known as "Wagner's Conjecture" and the theorem is surely the deepest in graph theory today.

Theorem 1.1 *There is no infinite antichain for $<_m$ in the set of all graphs.*

Roughly speaking, the proof of Theorem 1.1 goes as follows. Let $\mathcal{G}$ be an infinite set of graphs. If the graphs in $\mathcal{G}$ have bounded tree-width (a notion to be defined and discussed in the next section), then the result is a consequence of the first important step, which we shall prove in Section 2.

Theorem 1.2 *For every positive integer k, the relation $<_m$ is a well-quasi-order on the set of graphs having tree-width at most k.*

Thus, we can assume the graphs in $\mathcal{G}$ do not have bounded tree-width. Then the second important fact shows they have large grid minors. (The k-*grid* is the graph with vertices $v_{i,j}$, $1 \leq i, j \leq k$, such that $v_{i,j}$ and $v_{i',j'}$ are adjacent if $|i - i'| + |j - j'| = 1$.)

Theorem 1.3 *For every positive integer k, there is an integer $f(k)$ such that if G has tree-width at least $f(k)$, then G has a k-grid minor.*

This theorem has been given a relatively elementary proof by Diestel, Gorbunov, Jensen and Thomassen [3]. We shall present their proof in Section 3. It shows that if the graphs in $\mathcal{G}$ do not have bounded tree-width, then they have arbitrarily large grids. Now we have a structure that we can latch on to as we try to find minors.

One example of how to use Theorem 1.3 in the proof of Theorem 1.1 is the following. Suppose H is a fixed graph and G is a graph that contains a k-grid as a minor and that furthermore, within G there are a lot (say r) of pairwise disjoint paths, whose ends are in the grid minor but otherwise disjoint from the grid minor. Furthermore, we assume that no two of the $2r$ ends of the paths are "close" together in the grid. If k and r are sufficiently large functions of H, then G contains H as a minor. Thus, in order for $\mathcal{G} = \{G_1, G_2, \ldots\}$ not to have an element G_i, $i > 1$, containing G_1 as a minor, the elements of $\mathcal{G}$ with very large grids cannot have such a set of paths. This imposes restrictions on how the rest of G_i attaches to the grid minor and this additional structure allows us to look for even more structure (this goes on about four levels deep), until eventually we deduce that $\mathcal{G}$ has two elements, one of which is a minor of another.

From these remarks, we see that Theorems 1.2 and 1.3 are fundamental building blocks of the Graph Minors Project.

Thomassen [12] made the following observation, which we shall prove in Section 4.

Theorem 1.4 *Let Σ be a fixed surface. Then there is an integer $f(\Sigma)$ such that if G is minor-minimal with the property of not embedding in Σ, then G does not contain an $f(\Sigma)$-grid as a minor.*

To see the relevance of this observation, let $\mathcal{F}(\Sigma)$ denote the set of all graphs G that minor-minimally do not embed in Σ. Note that no element of $\mathcal{F}(\Sigma)$ is a proper minor of another element of $\mathcal{F}(\Sigma)$. By Theorem 1.4 the elements of $\mathcal{F}(\Sigma)$ have bounded sized grid minors. By Theorem 1.3, the graphs in $\mathcal{F}(\Sigma)$ have bounded tree-width. Finally, by Theorem 1.2, there can only be finitely many elements in $\mathcal{F}(\Sigma)$.

Let $ir(\Sigma)$ denote the set of all graphs G with minimum degree at least 3 such that G does not embed in Σ but every proper subgraph of G does embed in Σ. The distinction here is that the graphs minimally do not embed in Σ, *up to subdivision*, rather than being *minor-minimal*.

Corollary 1.5 *Let Σ be a fixed surface. Then $ir(\Sigma)$ is finite.*

The proof is quite simple: if $G \in ir(\Sigma)$, then G does not embed in Σ and so has an element H of $\mathcal{F}(\Sigma)$ as a minor. To get from G to H we can do only contractions (since contraction and deletion commute), so that G is obtained from H by splitting vertices. But there are only a finite number of graphs with minimum degree three that can be obtained by splitting vertices in a given graph.

I would like to acknowledge contributions to this work. In 1999-2000, the Fields Institute for Research in the Mathematical Sciences held a special year in Graph Theory and Combinatorial Optimization. As part of that year, I taught a course whose purpose was to present the material contained in this article. There was also a week-long workshop in Topological Graph Theory in which Paul Seymour presented three lectures describing the main components of the Graph Minors Project. These lectures were very stimulating for me and provided both a strong base and incentive to understand the work in detail. Beginning at roughly the same time, I have had many conversations with Jim Geelen, who has been working on generalizations of Graph Minors to matroids representable over finite fields. It is fair to say that these conversations have greatly influenced the form and detail of this work, and have kept me going.

2 Decompositions of Graphs

The main point of this section is to show that the set of graphs with bounded tree-width are well-quasi-ordered with respect to minors. This section is based on [4], [5] and [2].

There are now at least two main decompositions used in this context: tree-decompositions and branch-decompositions. These have, until now, been treated as different, but related, decompositions. Tree-decompositions are more natural in the context of graphs, but branch-decompositions have been successful for simplification and generalization to matroids.

It is interesting that there is a simultaneous generalization of these two concepts which shows that they are more closely related than previously thought. We shall take the opportunity here to present this common general decomposition and to prove the basic facts about it required for well-quasi-ordering. This will also make clearer the relationship between tree- and branch-decompositions.

Subsection 2.1 introduces the notion of decomposition of a graph and describes its relation to branch- and tree-decompositions of the same graph. Subsection 2.2 proves that, given any decomposition of a graph G, there is another decomposition which is not "worse" than the original decomposition with respect to any of three parameters and such that the second decomposition satisfies the Menger-property of being linked (to be defined in that subsection). Linked decompositions are used in Subsection 2.3 to show that any set of graphs having bounded tree-width is well-quasi-ordered.

2.1 Introduction to decompositions

A *decomposition* of a graph G is a pair (T, f), where T is a tree, with set $\mathcal{L}$ of leaves, and $f : E(G) \to \mathcal{L}$ is an injection.

It is easy to understand branch-decompositions from this point of view.

A *branch-decomposition* of a graph G is a decomposition (T, f) for which T has maximum degree at most three.

Tree-decompositions are more complicated. A *tree-decomposition* of a graph G is a pair $(T, \mathcal{W})$, where T is a tree and $\mathcal{W}$ is a set of subsets W_t of $V(G)$, one for each vertex t of T, such that:

1. $\cup_{t \in V(T)} W_t = V(G)$;
2. for each edge $e = uv$ of G, there is a $t \in V(T)$ such that $u, v \in W_t$; and
3. for each $u \in V(G)$, the subtree of T induced by $\{t \mid u \in W_t\}$ is connected.

It is not so obvious how tree-decompositions relate to decompositions. Let (T, f) be a decomposition. For each non-leaf vertex $t \in V(T)$, let W_t^f denote the set of vertices u for which u is incident with edges e, e' of G such that $f(e)$ and $f(e')$ are in different components of $T - t$. Furthermore, if t is a leaf of T, then, for $e \notin f(E(G))$, W_t^f is empty, while for $t = f(e)$, W_t^f is the set of ends of e.

Lemma 2.1 *Let G be a graph without isolated vertices and let (T, f) be a decomposition of a graph G. Then $(T, \{W_t^f \mid t \in V(T)\})$ is a tree-decomposition of G.*

Proof: The definition of W_t^f implies (1) and (2). As for (3), let v be any vertex of G and suppose $t, t' \in V(T)$ are such that $v \in W_t^f \cap W_{t'}^f$. Let t_1 be any vertex on the path in T between t and t'. We show $v \in W_{t_1}^f$. Note that we may assume t, t', and t_1 are all distinct.

Let e be an edge of G incident with v such that $f(e)$ is in a component of $T - t$ not containing t' and let e' be an edge of G incident with v such that $f(e')$ is in a component of $T - t'$ not containing t. Clearly, $f(e)$ and $f(e')$ are in different components of $T - t_1$ (since t and t' are in different components of $T - t_1$). Hence $v \in W_{t_1}^f$, as required. **Q.E.D.**

The restriction to graphs without isolated vertices in Lemma 2.1 is not really important. The main aim of this section is to prove that if $\mathcal{G}$ is an infinite set of graphs, each with tree-width at most b, then there exist $G, H \in \mathcal{G}$ such that $G <_m H$. This shows that there is no infinite antichain of graphs. If this is true when the graphs in $\mathcal{G}$ have no isolated vertices, then it is also easily seen to be true when we allow isolated vertices. For suppose $\mathcal{G}$ is a set of graphs (each having at least one edge). Let $\mathcal{G}'$ denote the set of graphs obtained from the graphs in $\mathcal{G}$ by deleting all the isolated vertices. If $\mathcal{G}'$ is finite, then some member of $\mathcal{G}$ is a subgraph of another member, a contradiction. If $\mathcal{G}'$ is infinite, then, by assumption, it does not have an infinite antichain, so it must have an infinite ascending chain $G_1' <_m G_2' <_m \ldots$. Let G_i' be obtained from $G_i \in \mathcal{G}$ by deleting isolated vertices. Suppose G_i has k_i isolated vertices. Clearly, there exist $i_1 < i_2 < \ldots$ such that $k_{i_1} \leq k_{i_2} \leq \ldots$, which implies that $G_{i_1} <_m G_{i_2} <_m \ldots$, so $\mathcal{G}$ contains an infinite ascending chain, as required.

It is not true that every tree-decomposition arises from a decomposition in the manner described above. However, there is an appropriate form of a converse. In order to state such a result, we need to associate parameters with a decomposition and a tree-decomposition.

Let G be a graph and let $E \subseteq E(G)$. Set $\Lambda(E)$ to be the set of vertices u of G incident with an edge in E and with an edge in $E(G) \setminus E$. We set $\lambda(E) = |\Lambda(E)|$.

Let G be a graph and let (T, f) be a decomposition of G. For each edge τ of T, let T_1 and T_2 be the two components of $T - \tau$. Let $E_\tau \subseteq E(G)$ consist of those edges e of G such that $f(e)$ is a vertex of T_1. The sets E_τ, $E(G) \setminus E_\tau$, and $\Lambda(E_\tau)$ are all *displayed* by the decomposition (T, f). Set $\zeta(\tau) = \lambda(E_\tau)$ and $\zeta(T, f) = \max\{\zeta(\tau) \mid \tau \in E(T)\}$.

For a vertex t of T, let X_t consist of those vertices v of G incident with edges e_1 and e_2, such that $f(e_1)$ and $f(e_2)$ are in different components of $T - t$. Set $\rho(t) = |X_t|$ and $\rho(T, f) = \max\{\rho(t) \mid t \in V(T)\}$.

The three important parameters for a decomposition (T, f) of a graph G are: (i) the maximum degree $\Delta(T)$ of T; (ii) $\zeta(T, f)$; and (iii) $\rho(T, f)$.

We extract the following simple observations. For a vertex v in a graph G, $\delta_G(v)$ is the set of edges of G incident with v.

Lemma 2.2 *Let (T, f) be a decomposition of a graph G. Let $t \in V(T)$. Then:*

1. *if $\tau \in \delta_T(t)$, then $\Lambda(E_\tau) \subseteq X_t$;*
2. *each vertex in X_t is in at least two of the $\Lambda(E_\tau)$;*
3. *if $\tau^* \in \delta_T(t)$, then $X_t = \cup_{\tau \in \delta_T(t) \setminus \{\tau^*\}} \Lambda(E_\tau)$; and*
4. *$\zeta(T, f) \leq \rho(T, f) \leq (\Delta(T)/2)\zeta(T, f)$.*

For a graph G, $\rho(G)$ is the minimum, over all decompositions (T, f) of G, of $\rho(T, f)$.

For a tree-decomposition $(T, \mathcal{W})$, $w(T, \mathcal{W})$ is the maximum of $|W_t|$, over all vertices of T. The *tree-width* $tw(G)$ of a graph G is the minimum of $w(T, \mathcal{W})$ of any tree-decomposition $(T, \mathcal{W})$ of G. (Note that this is one more than the standard definition. Since we are only concerned here with graphs having bounded tree-width, this difference is not relevant. The main point of using $(\max |W_t|) - 1$ in the standard definition seems to be so that trees have tree-width 1.)

We can already see that the definition of a decomposition is simpler than that of a tree-decomposition. In the main proofs of this section to follow, it will therefore be helpful to work with decompositions, rather than tree-decompositions. We note that the proof of Lemma 2.1 implies $tw(G) \leq \rho(G)$. Our next result shows the reverse inequality.

Lemma 2.3 *Let G be a graph with at least one edge and let $(T, \mathcal{W})$ be a tree-decomposition of G. Then there is a decomposition (T', f) such that $\rho(T', f) \leq w(T, \mathcal{W})$.*

Proof: For each edge $e = uv$ of G, (2) of the definition of tree-decomposition implies there is a vertex t_e of T such that $\{u, v\} \subseteq W_t$. For each vertex t of T, let E_t denote the set of edges e of G such that $t_e = t$. We let T' be the tree obtained from T by adding, for each vertex t of T, $|E_t|$ new neighbours of t, all leaves of T'. We obtain the decomposition (T', f) by setting, for each edge e of G, $f(e)$ to be one of the new leaves adjacent to t, if $e \in E_t$.

Observe that if $t \in V(T') \setminus V(T)$, then $\rho(t) \leq 2$. Now if $t_0 \in V(T)$, we claim $X_{t_0} \subseteq W_{t_0}$ and therefore $\rho(t_0) \leq |W_{t_0}|$. For if $v \in X_{t_0}$, v is incident with edges e, e' such that $f(e)$ and $f(e')$ are in different components of $T' - t_0$. Let t and t' be the neighbours of $f(e)$ and $f(e')$, respectively, in T'. By the definition of f, $v \in W_t \cap W_{t'}$. If one of t and t' were equal to t_0, we would be done.

Since $f(e)$ and $f(e')$ are in different components of $T' - t_0$, it follows that t and t' are in different components of $T' - t_0$, and, therefore, in different components of $T - t_0$.

Thus, t_0 is on the path in T between t and t'. Since the subgraph of T induced by $\{s \in V(T) \mid v \in W_s\}$ is connected and $v \in W_t \cap W_{t'}$, it follows that $v \in W_{t_0}$, as required. **Q.E.D.**

As an immediate consequence, we have the following important observation.

Theorem 2.4 *For any graph G with at least one edge, $tw(G) = \rho(G)$.*

Thus, in order to prove graphs with bounded tree-width are well-quasi-ordered, it suffices to show that graphs with bounded ρ are well-quasi-ordered (and for a given bound, we are dealing with exactly the same set of graphs).

A fundamental (though very easy) property of decompositions is that they are "minor-monotone". This is summarized in the following, which we will not explicitly require in this work.

Lemma 2.5 *Let (T, f) be a decomposition of a graph G, let H be a minor of G and let $\mathcal{L}$ be the set of leaves of T. Then the decomposition (T, f') of H with $f' : E(H) \to \mathcal{L}$ defined by $f'(e) = f(e)$ is such that, for each $\tau \in E(T)$ and each $t \in V(T)$, $\zeta^H(\tau) \leq \zeta^G(\tau)$ and $\rho^H(t) \leq \rho^G(t)$.*

Proof: Any minor H of G is obtained by a sequence of the following operations: (1) delete an edge; (2) delete an isolated vertex; and (3) contract an edge. We shall show that the lemma holds in each case; the general result then follows by a trivial induction. (In the rest of this proof, we use superscripts H and G to make it clear for which graph the relevant function is being applied.)

(1) Let $H = G - e'$. For an edge τ of T, let E_τ denote the set of edges displayed by τ in the decomposition (T, f) of G, labelled so that $e' \in E_\tau$. If $v \in \Lambda^H(E_\tau \setminus \{e'\})$, then v is incident with an edge in both $E_\tau \setminus \{e'\}$ and in $E(G) \setminus E_\tau$, so $v \in \Lambda^G(E_\tau)$. Thus, $\zeta^H(\tau) \leq \zeta^G(\tau)$.

Similarly, if $t \in V(T)$ and $v \in V(H)$ is incident with edges e_1, e_2 of H such that $f'(e_1), f'(e_2)$ are in different components of $T - t$, then $f(e_1), f(e_2)$ are in different components of $T - t$. Thus, $X_t^H \subseteq X_t^G$, so $\rho^H(t) \leq \rho^G(t)$.

(2) Deleting isolated vertices has no impact on decompositions.

(3) Let H be the graph obtained from G by contracting the edge $e' = uv$ and let w be the vertex of H that replaces u and v. If $\tau \in E(T)$, then $\Lambda^H(\tau) \subseteq (\Lambda^G(\tau) \cup \{w\}) \setminus \{u, v\}$, and $w \in \Lambda^H(\tau)$ implies at least one of u and v is in $\Lambda^G(\tau)$, whence $\zeta^H(\tau) \leq \zeta^G(\tau)$.

Similarly, for $t \in V(T)$, $X_t^H \subseteq (X_t^G \cup w\}) \setminus \{u, v\}$ and if $w \in X_t^H$, then at least one of u and v is in X_t^G, so $\rho^H(t) \leq \rho^G(t)$. **Q.E.D.**

We conclude this subsection with some remarks relating the branch-width $bw(G)$ of a graph G (i.e., the smallest value of $\zeta(T, f)$ over all branch-decompositions of G) and the tree-width of G. Suppose first that (T, f) is a decomposition of G such that there is a $t \in V(T)$ of degree at least 4 in T. Let S be any tree having only degree 3 and degree 1 vertices so that S has exactly $deg_T(t)$ vertices of degree 1. Let T' be obtained from T by deleting t and identifying each neighbour of t in T with one of the leaves of S. Define $f'(e)$ to be $f(e)$.

If $t \in V(T') \setminus S$, then $\rho^{T'}(t) = \rho^T(t)$. For $s \in S$, if v is a vertex of G incident with edges e_1 and e_2 such that $f'(e_1)$ and $f'(e_2)$ are in different components of $T' - s$, then $f(e_1)$ and $f(e_2)$ are in different components of

$T - v$. Hence $X_s^{T'} \subseteq X_t^T$. That is, $\rho(T', f') \leq \rho(T, f)$, as claimed. This implies the following, which we will use in the well-quasi-ordering argument.

Lemma 2.6 *Let G be a graph. Then there is a decomposition (T, f) of G such that $\Delta(T) \leq 3$ and $\rho(T, f) = tw(G)$.*

Let (T, f) be a decomposition of G as in Lemma 2.6. Then (T, f) is also a branch-decomposition of G. Since $tw(G) = \rho(T, f) = \max\{|X_t| : t \in V(T)\}$ and, for $\tau \in \delta(t)$, $\Lambda(E_\tau) \subseteq X_t$, we see that $tw(G) \geq bw(G)$. On the other hand, let (T, f) be a branch-decomposition such that $bw(G) = \zeta(T, f)$. By Theorem 2.4 and Lemma 2.2, $tw(G) = \rho(G) \leq \rho(T, f) \leq (3/2)\zeta(T, f) = (3/2)bw(G)$.

Lemma 2.7 *For any graph G, $bw(G) \leq tw(G) \leq (3/2)bw(G)$.*

2.2 Linked decompositions

Robertson and Seymour originally proved that graphs with bounded tree-width are well-quasi-ordered by a method that has not been published in its original form. (A preprint was circulated.) In the meantime, Thomas [11] made a significant improvement on a key idea and its proof, so the published version of [7] makes use of the improvement. Thomas' proof has been simplified by Diestel [2]. The purpose of this section is to prove the analogue of Thomas' theorem for decompositions; it implies Thomas' result. It combines the techniques of [4] and [2].

In trying to prove that graphs with bounded tree-width are well-quasi-ordered, it is natural to try to "come down" the trees of tree-decompositions of the graphs, showing that at each step along the way, the graphs we have built up so far are well-quasi-ordered. As we look to go down to the next edge, we will be "gluing" together all the graphs corresponding to the edges just above. The context will be simpler if we have only two edges above — this is Lemma 2.6 from the previous subsection. It is in the "coming down" process that we require some additional information about the decomposition, and it is the purpose of this section to provide this additional information.

Let (T, f) be a decomposition of a graph G. For τ_1, τ_2 distinct elements of $E(T)$, let $A(\tau_1, \tau_2)$ denote the set of edges e of G such that $f(e)$ is in the component of $T - \tau_1$ not containing τ_2 and let $B(\tau_1, \tau_2)$ denote the set of edges e of G such that $f(e)$ is in the component of $T - \tau_2$ not containing τ_1. Let

$\lambda(\tau_1, \tau_2)$ denote the maximum number of pairwise totally disjoint paths in G from $\Lambda(A(\tau_1, \tau_2))$ to $\Lambda(B(\tau_1, \tau_2))$. We abbreviate "pairwise totally disjoint" to PTD.

On the other hand, let $P(\tau_1, \tau_2)$ denote the shortest path in T containing both τ_1 and τ_2. Let $\lambda^*(\tau_1, \tau_2) = \min\{\zeta(\tau) \mid \tau \in E(P(\tau_1, \tau_2))\}$. It follows from Menger's Theorem that $\lambda^*(\tau_1, \tau_2) \geq \lambda(\tau_1, \tau_2)$.

The decomposition (T, f) is *linked* if, for every distinct pair τ_1, τ_2 of edges of T, $\lambda^*(\tau_1, \tau_2) = \lambda(\tau_1, \tau_2)$.

We are now ready for the decomposition version of Thomas' result.

Theorem 2.8 *Let (T, f) be a decomposition of a graph G. Then there is a linked decomposition (T', f') of G such that: (i) $\Delta(T') \leq \Delta(T)$; (ii) $\zeta(T', f') \leq \zeta(T, f)$; and (iii) $\rho(T', f') \leq \rho(T, f)$.*

Proof: If (T, f) is not already linked, there exist $\tau_1, \tau_2 \in E(T)$ such that $\lambda^*(\tau_1, \tau_2) > \lambda(\tau_1, \tau_2)$. Let $A = A(\tau_1, \tau_2)$ and $B = B(\tau_1, \tau_2)$. Note that $\lambda(A) \geq \lambda^*(\tau_1, \tau_2) > \lambda(\tau_1, \tau_2) \geq 0$, so $A \neq \emptyset$. Similarly, $B \neq \emptyset$.

Two subsets $E, E' \subseteq E(G)$ *cross* if none of the four sets $E \cap E'$, $E \setminus E'$, $E' \setminus E$, and $E(G) \setminus (E \cup E')$ is empty. Notice that if E crosses E', then each of E and $E(G) \setminus E$ crosses each of E' and $E(G) \setminus E'$.

Let X be a minimal set separating $\Lambda(A)$ and $\Lambda(B)$. If u is incident with an edge of A and v is incident with an edge of B, then any (u, v)-path P must contain a vertex of $\Lambda(A)$ and a vertex of $\Lambda(B)$. (The first vertex of P incident with an edge not in A is in $\Lambda(A)$; the last incident with an edge not in B is in $\Lambda(B)$.) Hence it contains a vertex of X. This implies that no bridge of X in G contains an edge of A and an edge of B.

(Recall that, for X a subset of $V(G)$, a *bridge* of X is a subgraph of G that is either an edge of G and its ends, both of which are in X, or, a component K of $G - X$, together with all the edges from K to vertices in X, plus the ends of these edges.)

Let E denote the union of the edge sets of the bridges of X that contain edges in A. Clearly, $\Lambda(E)$ separates $\Lambda(A)$ from $\Lambda(B)$ and $\Lambda(E) \subseteq X$. The minimality of X implies $X = \Lambda(E)$. Futhermore, $A \subseteq E$.

Let E^* be chosen so that: (i) $A \subseteq E^*$; (ii) $\Lambda(E^*)$ is a smallest set separating $\Lambda(A)$ from $\Lambda(B)$; and (iii) subject to (i) and (ii), E^* crosses as few as possible of the sets E_τ displayed by edges of T.

There are three components in the graph $T - \{\tau_1, \tau_2\}$: the one containing $f(A)$ is T_A; the one containing $f(B)$ is T_B; and the third is T_M (for "middle").

Thus, $T_{A,M} = (T_A \cup T_M) + \tau_1$ is a component of $T - \tau_2$. Similarly, $T_{M,B} = (T_M \cup T_B) + \tau_2$ is a component of $T - \tau_1$.

We construct a decomposition (T', f') of G as follows. Let T' be constructed from disjoint copies $T'_{A,M}$ and $T'_{M,B}$ of $T_{A,M}$ and $T_{B,M}$, respectively, with the copy of the end of τ_2 in $T'_{A,M}$ joined by a new edge τ^* to the copy of the end of τ_1 in $T'_{M,B}$. The definition of f' depends on whether $e \in E^*$ or not. If $e \in E^*$, then $f'(e)$ is the copy of the vertex $f(e)$ that is in $T'_{A,M}$, while if $e \notin E^*$, then $f'(e)$ is the copy of the vertex $f(e)$ that is in $T'_{M,B}$.

We will need to prove that: (a) $\Delta(T') \leq \Delta(T)$; (b) $\zeta(T', f') \leq \zeta(T, f)$; and (c) $\rho(T', f') \leq \rho(T, f)$. Furthermore, we will need to prove (d) (T', f') is "better" than (T, f).

Observe that (a) is a triviality.

For the rest, note that if $\tau \in E(T_A) \cup E(T_B) \cup \{\tau_1, \tau_2\}$, then it is a triviality that, for the copy τ' of τ in T', $\zeta(\tau') = \zeta(\tau)$. A similar statement holds for vertices in $V(T_A) \cup V(T_B)$.

So let $\tau \in E(T_M)$. We shall conduct all our arguments with the copy τ'_A of τ in $T'_{A,M}$, realizing that the arguments for the copy τ'_B of τ in $T'_{M,B}$ follow by symmetry. Observe that $E_{\tau'_A} = E_\tau \cap E^*$ and $E_{\tau'_B}$ is, for $\tau \in E(P(\tau_1, \tau_2))$, $E_\tau \cup E^*$, while for $\tau \notin E(P(\tau_1, \tau_2))$, $E_{\tau'_B}$ is $E_\tau \setminus E^*$.

Choose the labelling so that E_τ is the set of edges displayed by τ in (T, f) such that $E_\tau \cap B = \emptyset$. Let $\mathcal{P}$ be a set of $\lambda(E^*)$ PDT paths from $\Lambda(A)$ to $\Lambda(B)$. For each $v \in \Lambda(E^*)$, there is a unique path $P_v \in \mathcal{P}$ containing v and P_v consists of two subpaths P_v^A, from the $\Lambda(A)$ end of P_v to v, and P_v^B, from v to the $\Lambda(B)$ end of P_v.

We begin with the following observation.

Claim 1 *$\Lambda(E_\tau)$ contains $\Lambda(E_\tau \cap E^*) \setminus \Lambda(E^*)$ and, for each vertex $v \in \Lambda(E_\tau \cap E^*) \cap \Lambda(E^*)$, there is at least one vertex of P_v^B in $\Lambda(E_\tau)$.*

Proof: If $u \in \Lambda(E_\tau \cap E^*) \setminus \Lambda(E^*)$, then u is incident with an edge e_1 in $E_\tau \cap E^*$ and an edge e_2 not in $E_\tau \cap E^*$. If $e_2 \notin E^*$, then $u \in \Lambda(E^*)$, a contradiction. Hence $e_2 \notin E_\tau$, so $u \in \Lambda(E_\tau)$.

If $v \in \Lambda(E_\tau \cap E^*) \cap \Lambda(E^*)$, then v is incident with an edge in E_τ. Let $b \in \Lambda(B)$ be the end of P_v^B, so b is incident with an edge of B (which is certainly not in E_τ). Thus, P_v^B contains a vertex that is incident with an edge in E_τ and with an edge not in E_τ. **Q.E.D.**

For a vertex t of T_M, let $\tau_t \in \delta_T(t)$ denote the edge separating t from T_B and let $\delta_t = \delta_T(t) \setminus \{\tau_t\}$. By Lemma 2.2 the set X_t^T is the union, over $\tau \in \delta_t$, of the $\Lambda(E_\tau)$. For the copy t'_A of t in $T'_{A,M}$, we have the analogous set $X^{T'}_{t'_A}$. An immediate corollary to Claim 1 is:

Claim 2 *X^T_t contains $X^{T'}_{t'_A} \setminus \Lambda(E^*)$ and, for each $v \in X^{T'}_{t'_A} \cap \Lambda(E^*)$, there is at least one vertex of P^B_v in X^T_t.*

It follows immediately from Claims 1 and 2 that $\zeta(\tau'_A) \leq \zeta(\tau)$ and $\rho(t'_A) \leq \rho(t)$. This proves (b) and (c).

We now wish to prove that (T', f') is a "better" decomposition than (T, f). For a decomposition (S, g) of a graph G, and for $r \in \{0, 1, 2, \ldots, |V(G)|\}$, let $\pi(S, g, r)$ denote the number of vertices s of S for which $\rho(s) = r$. We set

$$\Pi(S, g) = (\, \pi(S, g, |V(G)|)\, ,\ \pi(S, g, |V(G)| - 1)\, ,\ \ldots\, ,\ \pi(S, g, 0)\,)$$

and we shall prove that $\Pi(T', f')$ is lexicographically smaller than $\Pi(T, f)$.

The main point is the following.

Claim 3 *If $t \in V(T_M)$ is such that $|X^T_t| = |X^{T'}_{t'_A}| > 0$, then $X^{T'}_{t'_B} \subseteq \Lambda(E^*)$.*

What Claim 3 shows is that if $r > |\Lambda(E^*)|$, then $\pi(T', f', r) \leq \pi(T, f, r)$. We shall prove below that there is at least one such r for which the inequality is strict.

The main points in the proof of Claim 3 are the following three facts.

Claim 4 *Let $A, B, C \subseteq E(G)$. If $A \cup B$ crosses C and B crosses A but not C, then A crosses C.*

Proof: We note that $C \setminus A \supset C \setminus (A \cup B) \neq \emptyset$ and $E(G) \setminus (A \cup C) \supset E(G) \setminus ((A \cup B) \cup C) \neq \emptyset$. It remains to show (a) $A \cap C \neq \emptyset$ and (b) $A \setminus C \neq \emptyset$.

We note that $E(G) \setminus (B \cup C) \supseteq E(G) \setminus ((A \cup B) \cup C) \neq \emptyset$ and $C \setminus B \supseteq C \setminus (A \cup B) \neq \emptyset$. Since B does not cross C, either $B \cap C = \emptyset$ or $B \setminus C = \emptyset$.

Case 1 $B \cap C = \emptyset$. (a) $A \cap C = (A \cup B) \cap C \neq \emptyset$. (b) $A \setminus C \supseteq A \cap B \neq \emptyset$.

Case 2 $B \setminus C = \emptyset$. (a) $A \cap C \supseteq A \cap B \neq \emptyset$. (b) $A \setminus C = (A \cup B) \setminus C \neq \emptyset$.

Q.E.D.

For a vertex t of T_M, set $E^t = E_{\tau_t}$. Note that $E^t = \cup_{\tau \in \delta_t} E_\tau$.

Claim 5 *If $|X^T_t| = |X^{T'}_{t'_A}|$, then $|\Lambda(E^* \cup E^t)| = |\Lambda(E^*)|$.*

Proof: If $u \in \Lambda(E^* \cup E^t) \setminus \Lambda(E^*)$, then u is incident with some edge e_1 not in $E^* \cup E^t$. Therefore, e_1 is not in E^* and, since $u \notin \Lambda(E^*)$, u is not incident with any edge of E^*. Since u is also incident with an edge $e_2 \in E^* \cup E^t$, it follows that $e_2 \in E^t$, whence $u \in \Lambda(E_\tau)$, for some $\tau \in \delta_t$. By Claim 2 and the fact that $|X_t^T| = |X_{t'_A}^{T'}|$, we see that u is a vertex of P_v^B for some $v \in \Lambda(E^*)$ and u is the only vertex from P_v^B in X_t^T. Thus, $|\Lambda(E^* \cup E^t)| \leq |\Lambda(E^*)|$. Since $\Lambda(E^* \cup E^t)$ separates $\Lambda(A)$ from $\Lambda(B)$, the choice of E^* implies $|\Lambda(E^* \cup E^t)| \geq |\Lambda(E^*)|$. **Q.E.D.**

Claim 6 *If $|X_t^T| = |X_{t'_A}^{T'}| > 0$ and $X_{t'_B} \not\subseteq \Lambda(E^*)$, then E^* crosses E^t.*

Proof: The set $E(G) \setminus (E^* \cup E^t)$ contains B, so it is not empty.

To see that $E^* \cap E^t$ is not empty, observe that $X_{t'_A}^{T'} = \cup_{\tau\in\delta_t}\Lambda(E_{\tau'_A}) = \cup_{\tau\in\delta_t}\Lambda(E^* \cap E_\tau)$. Since $X_{t'_A}^{T'}$ is not empty, there is a $\tau \in \delta_t$ for which $\Lambda(E^* \cap E_\tau)$ is not empty. For such a τ, $E^* \cap E_\tau$ is not empty, which in turn implies $E^* \cap E^t \neq \emptyset$.

To see that $E^* \setminus E^t \neq \emptyset$, note that $E^* \subseteq E^t$ implies $A \subseteq E^t \subseteq E(G) \setminus B$. Thus, $\Lambda(E^t)$ separates $\Lambda(A)$ from $\Lambda(B)$. But $E_* \subseteq E^t$ implies $\Lambda(E^t) = \Lambda(E^* \cup E^t)$ and, by Claim 5, $\Lambda(E^* \cup E^t)$ has the same size as $\Lambda(E^*)$. Thus, $\Lambda(E^t)$ is a smallest set separating $\Lambda(A)$ from $\Lambda(B)$ and is displayed in (T, f). But our assumption is that no edge displayed by (T, f) separating $f(A)$ from $f(B)$ has this small a label, so we have a contradiction.

Finally, to see that $E^t \setminus E^* \neq \emptyset$, let $w \in X_{t'_B}^{T'} \setminus \Lambda(E^*)$. There is some $\tau \in \delta_t$ such that w is incident with an edge $e_1 \in E_{\tau'_B}$ and with an edge $e_2 \notin E_{\tau'_B}$, where $E_{\tau'_B}$ is the subset of $E(G) \setminus B$ displayed by τ'_B in (T', f').

Since $E_{\tau'_B}$ is either $E_\tau \setminus E^*$ or $E_\tau \cup E^*$, at least one of the edges e_1, e_2 is not in E^*. Since $w \notin \Lambda(E^*)$, w cannot be incident with any edges of E^*. Thus, neither of the edges e_1, e_2 is in E^*. Thus, $e_1 \in E_\tau$. Since $E_\tau \subseteq E^t$, we have that $E^t \setminus E^* \neq \emptyset$. **Q.E.D.**

We are now ready for the proof of Claim 3.

Proof of Claim 3: If $X_{t'_B}^{T'}$ is not a subset of $\Lambda(E^*)$, then Claim 6 shows that E^* crosses E^t. If τ is any edge of T, then E^t and E_τ do not cross (since any two sets of edges displayed by (T, f) are either disjoint or one is contained in the other). It follows from Claim 4 that if $E^* \cup E^t$ crosses E_τ, then E^* crosses E_τ. Since E^* crosses E^t but $E^* \cup E^t$ does not, it follows that: (i) $E^* \cup E^t$ contains A; (ii) $\Lambda(E^* \cup E^t)$ is a smallest set separating $\Lambda(A)$ from $\Lambda(B)$; and (iii) $E^* \cup E^t$ crosses fewer of the sets displayed by the edges of T than E^* does. This contradicts the choice of E^*. **Q.E.D.**

Now we can prove that $\Pi(T', f')$ is lexicographically smaller than $\Pi(T, f)$. We suppose not. For $r > |\Lambda(E^*)|$, Lemma 3 implies $\pi(T', f', r) \le \pi(T, f, r)$, since if $|X_t^T| = r > |\Lambda(E^*)|$, at most one of $|X_{t'_A}^{T'}|$ and $|X_{t'_B}^{T'}|$ can be r. Thus, our assumption that $\Pi(T', f')$ is not lexicographically smaller than $\Pi(T, f)$ implies that if $t \in V(T_M)$ is such that $|X_t^T| = r > |\Lambda(E^*)|$, exactly one of $|X_{t'_A}^{T'}|$ and $|X_{t'_B}^{T'}|$ is r.

Let s and t be the ends of τ_1 and τ_2, respectively, in T_M. Clearly $\Lambda(A)$ is a subset of both X_s^T and $X_{s'_A}^{T'}$. Similarly, $\Lambda(B)$ is a subset of both X_t^T and $X_{t'_B}^{T'}$. The choice of A and B implies $|\Lambda(A)| > |\Lambda(E^*)|$ and $|\Lambda(B)| > |\Lambda(E^*)|$, so both $|X_{t'_A}^{T'}| = |X_t^T|$ and $|X_{t'_B}^{T'}| = |X_t^B|$.

It follows that there is an edge $\tau = pq$ in $T_M \cap P(\tau_1, \tau_2)$ such that: p is closer to s in T_M than q is; $|X_{p'_A}^{T'}| = |X_p^T|$; and $|X_{q'_B}^{T'}| = |X_q^T|$. By the choice of A and B, $|X_p^T|$ and $|X_q^T|$ are both greater than $|\Lambda(E^*)|$.

We show that E^* crosses E^p. We know $E^p = E_\tau \subseteq E^q$, $A \subseteq E^p$, $E_{\tau'_A} = E_\tau \cap E^*$, $E_{\tau'_B} = E_\tau \cup E^*$, $\Lambda(E_\tau) \subseteq X_p^T$, and $\Lambda(E_\tau) \subseteq X_q^T$.

(1) Since $A \subseteq E^* \cap E^p$, $E^* \cap E^p \neq \emptyset$.

(2) Since $B \subseteq E(G) \setminus (E^* \cup E^p)$, $E(G) \setminus (E^* \cup E^p) \neq \emptyset$.

(3) Suppose $E^* \setminus E^p = \emptyset$. Then $E^* \subseteq E^p$. Because $|X_{p'_A}^{T'}| = |X_p^T| > 0$, Claim 3 implies $X_{p'_B}^{T'} \subseteq \Lambda(E^*)$. Since $\Lambda(E^* \cup E^p) = \Lambda(E_{\tau'_B}) \subseteq X_{p'_B}^{T'} \subseteq \Lambda(E^*)$, we conclude that $\Lambda(E_\tau) = \Lambda(E^p) = \Lambda(E^* \cup E^p) \subseteq \Lambda(E^*)$. Since $\Lambda(E_\tau)$ is a set separating $\Lambda(A)$ from $\Lambda(B)$, it follows that $|\Lambda(E_\tau)| \ge |\Lambda(E^*)|$, which, together with the previous containment, implies $\Lambda(E_\tau) = \Lambda(E^*)$.

But $\Lambda(E_\tau)$ is displayed in (T, f), contradicting the choice of A and B.

(4) Finally, we show that $E^p \setminus E^* \neq \emptyset$. Otherwise, $E^p \subseteq E^*$. This is the same argument, but applied to q. We have the relations $\Lambda(E_\tau) = \Lambda(E^p) = \Lambda(E^p \cap E^*) \subseteq X_{q'_A}^{T'} \subseteq \Lambda(E^*)$ (the last containment by Claim 3). Since E_τ contains A and is disjoint from B, $|\Lambda(E_\tau)| \ge |\Lambda(E^*)|$, and again we have the contradiction that $\Lambda(E_\tau) = \Lambda(E^*)$. Thus, E^* crosses E^p.

Note that $\Lambda(E^p \cup E^*) = \Lambda(E_{\tau'_B}) \subseteq X_{p'_B}^{T'} \subseteq \Lambda(E^*)$. Therefore, $\Lambda(E^p \cup E^*) = \Lambda(E^*)$. Thus, $\Lambda(E^p \cup E^*)$ is also a smallest set separating $\Lambda(A)$ from $\Lambda(B)$.

The final contradiction comes from the realization that $E^* \cup E^p$ crosses fewer of the sets displayed by (T, f) than E^* does, since Claim 4 implies everything crossed by $E^* \cup E^p$ is also crossed by E^*; but E^* crosses E^p, while $E^* \cup E^p$ does not.

To finish the proof of Theorem 2.8, choose a decomposition (T, f) so that

$\Pi(T,f)$ is lexicographically least. This implies that $\rho(T,f) = \rho(G)$ and that (T,f) is linked (since we just proved that decompositions that are not linked are not lexicographically least). **Q.E.D.**

2.3 The set of graphs with bounded tree-width is well-quasi-ordered

In this section, we shall prove that the set of graphs having bounded tree-width is well-quasi-ordered. The proof comes in two parts, the first being an argument to provide some structure in a "minimal counterexample" and the second showing the structure does not occur in graphs with bounded tree-width.

Let w be a positive integer and let $\mathcal{G}$ be an infinite set of graphs, all with tree-width at most w. For each $G \in \mathcal{G}$, we shall fix, once and for all, a linked decomposition (T,f) such that $\rho(T,f) \leq w$ and $\Delta(T) \leq 3$. We shall assume that there is a leaf r of T such that $r \notin f(E(G))$; r is the root of T. (There is no difficulty in arranging for such an r. For example, subdivide any edge of T and join a new vertex r to the subdivision vertex.) The role of r is simply to provide a direction to the tree, in order that we know how to traverse it consistently; we will work towards r. Thus, we shall think of the edges of T as being directed towards the root. We note that $X_r = \emptyset$ and if τ is the edge of T incident with r, then $E_\tau = E(G)$.

An edge τ of T is labelled with the integer $\lambda(\tau) = |\Lambda(E_\tau)|$, which is in the set $\{0, 1, \ldots, w\}$. For two edges τ, τ' of T, τ *is linked to* τ' if T has a directed path P whose first edge is τ, whose last edge is τ', $\lambda(\tau) = \lambda(\tau')$, and if τ'' is any other edge of P, then $\lambda(\tau'') \geq \lambda(\tau)$. Since T is linked, this implies that if τ is linked to τ', then in G there are $\lambda(\tau)$ PTD paths from $\Lambda(E_\tau)$ to $\Lambda(E_{\tau'})$.

More generally, a *rooted forest* is a countable collection F of rooted trees (each directed toward its root). Let F be a rooted forest and let $\lambda : E(F) \to \{0, 1, 2, \ldots, w\}$. For two edges τ, τ' of the component T of F, τ *is* λ*-linked to* τ' if T has a directed path P whose first edge is τ, whose last edge is τ', $\lambda(\tau) = \lambda(\tau')$, and if τ'' is any other edge of P, then $\lambda(\tau'') \geq \lambda(\tau)$. For a subset A of $E(F)$, the set $U_F(A)$ denotes the set of edges τ of F for which there is a directed path in F of length two whose first edge is τ and whose second edge is in A.

The lemma we need is the following version (taken from [4]) of some of the results in [7]. This is essentially Lemma 2.2 in [7], but the proof seems to be a little tidier.

Lemma 2.9 *Let F be a forest and let $\lambda : E(F) \to \{0, 1, \ldots, w\}$. Let $\preceq$ be a quasi-order on $E(F)$ such that: (i) there is no infinite descending chain; and (ii) if τ is λ-linked to τ', then $\tau' \preceq \tau$. If the edges of F are not well-quasi-ordered by $\preceq$, then there exists an infinite antichain A of edges of F such that $\preceq$ is a well-quasi-order on $U_F(A)$.*

To see how this will apply, let (T, f) and (T', f') be two of our decompositions, for G and G', respectively. For $\tau \in E(T)$ and $\tau' \in E(T')$, $\tau \preceq \tau'$ will roughly mean that the subgraph of G induced by E_τ (on the side of τ "away" from r) is a minor of the subgraph of G' induced by $E_{\tau'}$ (on the side of τ' "away" from r'). We will need the minor relation to also take into account the sets $\Lambda(E_\tau)$ and $\Lambda(E_{\tau'})$, but this is a technicality which we will ignore for the moment. The lemma says that there is an infinite antichain A of edges of F so that, looking at the tails of the edges of A, the "left" and "right" edges above the tails are well-quasi-ordered. This says that the left and right subgraphs are well-quasi-ordered, but their unions are an infinite antichain. The boundedness of the $\Lambda(E_\tau)$ will then be used in a straightforward way to obtain a contradiction.

Proof of Lemma 2.9: Suppose the lemma is false and that among all counterexamples, we have F and λ for which w is as small as possible. Let Z be the set of edges τ of F such that $\lambda(\tau) = 0$. Suppose Z is well-quasi-ordered by $\preceq$. The labelling $\lambda - 1$ on $F - Z$ gives a counterexample, but with smaller w, a contradiction. (It is very simple to see that if $\preceq$ is a well-quasi-order on both R and S, then it is a well-quasi-order on $R \cup S$.) Hence $\preceq$ is not a well-quasi-order on Z. In particular, $Z \neq \emptyset$. If $\tau, \tau' \in Z$, then τ is λ-linked to τ' if and only if there is a directed path in F starting with τ and ending with τ'.

A sequence $(\tau_1, \tau_2, \ldots)$ of edges of F is *never-increasing* if the sequence is infinite and if $i < j$ implies $\tau_i \not\preceq \tau_j$. We claim:

Claim *There is a never-increasing sequence $(\tau_1, \tau_2, \ldots)$ in Z such that, for each k, if τ_k is λ-linked to $\tau \in Z \setminus \{\tau_k\}$, then $(\tau_1, \tau_2, \ldots, \tau_{k-1}, \tau)$ is not the beginning of a never-increasing sequence in Z.*

We note that since $\preceq$ has no infinite descending chain: (i) every non-empty subset of $E(F)$ has a minimal element; and (ii) every infinite sequence either contains an infinite antichain or an infinite ascending chain.

To prove the claim, we note that the assumption that $\preceq$ is not a well-quasi-order implies there is a never-increasing sequence (e.g., any ordering of a countably infinite antichain). If we let E_1 be the set of edges τ that are the first term of some never-increasing sequence, then we can choose τ_1 to be a minimal element of E_1. Supposing we have determined $\tau_1, \tau_2, \ldots, \tau_{k-1}$, we select τ_k to be a minimal element of the set E_k of those τ such that some never-increasing sequence begins with $(\tau_1, \tau_2, \ldots, \tau_{k-1}, \tau)$.

It is easy to see that $(\tau_1, \tau_2, \ldots)$ is a never-increasing sequence. Also, the minimality of τ_k when it was selected shows that if $\tau \in Z \setminus \{\tau_k\}$ is such that $\tau \preceq \tau_k$, then $(\tau_1, \tau_2, \ldots, \tau_{k-1}, \tau)$ is not the beginning of a never-increasing sequence. This proves the claim.

Furthermore, suppose i and j are distinct indices such that $\tau_j \preceq \tau_i$. The fact that the sequence in never-increasing implies $i < j$. But then the sequence $(\tau_1, \tau_2, \ldots, \tau_{i-1}, \tau_j, \tau_{j+1}, \ldots)$ is never-increasing, contradicting the fact that τ_i is minimal in E_i. Consequently, the set $A = \{\tau_1, \tau_2, \ldots\}$ is an infinite antichain, so the assumption that the lemma fails for F implies the set $U_F(A)$ is not well-quasi-ordered. Consider the rooted subforest F' of F with roots the tails of the τ_i (keeping the directions, so the original roots are not in F'). It is another counterexample (with the same w), so the preceding argument shows that it has a never-increasing sequence $(\sigma_1, \sigma_2, \ldots)$ such that, for every i, $\lambda(\sigma_i) = 0$.

For each i, there is a unique integer $j(i)$ such that F contains a directed path from σ_i to $\tau_{j(i)}$. Let i^* be an i such that $j(i)$ is as small as possible. Consider the sequence $(\tau_1, \tau_2, \ldots, \tau_{j(i^*)-1}, \sigma_{i^*}, \sigma_{i^*+1}, \ldots)$. We claim this sequence is never-increasing. To see this, notice that the only other possibility is that $\tau_k \preceq \sigma_\ell$ for some k, ℓ such that $1 \le k \le j(i^*) - 1$ and $i^* \le \ell$. But $\sigma_\ell \preceq \tau_{j(\ell)}$ would then imply $\tau_k \preceq \tau_{j(\ell)}$ and, since $k < j(i^*) \le j(\ell)$, this contradicts the fact that $(\tau_1, \tau_2, \ldots)$ is never-increasing, thereby proving the lemma. **Q.E.D.**

In order to apply Lemma 2.9, we need to define the quasi-order $\preceq$ on the edges of the forest consisting of all the decomposition trees. This requires the concept of "rooted minor". Let G and H be graphs and let $W \subseteq V(G)$ and $X \subseteq V(H)$. The pair (G, W) is a *rooted minor* of (H, X), denoted $(G, W) \le_r (H, X)$, if there is a sequence $(H, X) = (G_1, W_1), (G_2, W_2), \ldots, (G_k, W_k)$, such that (G_k, W_k) is isomorphic (as a pair) to (G, W) (so there is an isomorphism from G to G_k that sends W to W_k) and each (G_{i+1}, W_{i+1}) is obtained from (G_i, W_i) by one of the following rules:

1. deleting an edge e, in which case $G_{i+1} = G_i - e$ and $W_{i+1} = W_i$;
2. contracting an edge $e = uv$, in which case $G_{i+1} = G_i/e$ and W_{i+1} is W_i, if $\{u, v\} \cap W_i = \emptyset$; otherwise, W_{i+1} is $(W_i \setminus \{u, v\}) \cup \{u'\}$, where u' is the vertex of contraction of e; and
3. deleting an isolated vertex u not in W_i, in which case $G_{i+1} = G_i - u$ and $W_{i+1} = W_i$.

So now let w be a positive integer and let $\mathcal{G}$ be an infinite set of graphs, all having tree-width at most w. For each $G \in \mathcal{G}$, let (T_G, f_G) be a linked decomposition such that $\rho(T_G, f_G) \leq w$ and $\Delta(T) \leq 3$. We assume that there is a leaf vertex r_G of T_G not in $f_G(E(G))$, which will be the root of T_G.

Let F be the collection of all trees T_G, for $G \in \mathcal{G}$. Define the relation $\preceq$ on the edges of F as follows. Let $\tau \in E(F)$. There is a $G_\tau \in \mathcal{G}$ such that $\tau \in E(T_{G_\tau})$. Let E_τ be that set of edges displayed by τ (in (T_{G_τ}, f_{G_τ})) chosen so that $f(E_\tau)$ is not in the component of $T_{G_\tau} - \tau$ containing r_{G_τ}. Let H_τ be the subgraph of G_τ induced by E_τ. For edges τ, τ' of F, $\tau \preceq \tau'$ if $(H_\tau, \Lambda(E_\tau)) \leq_r (H_{\tau'}, \Lambda(E_{\tau'}))$.

The following result is required in order to be able to apply Lemma 2.9. (This is where we use the fact that the decompositions are linked.) Recall that $\lambda(\tau) = |\Lambda(E_\tau)|$.

Lemma 2.10 *If τ is λ-linked in F to τ', then $\tau' \preceq \tau$.*

Proof: Since τ is λ-linked to τ', they are both in the same tree T_G, for some $G \in \mathcal{G}$. Furthermore, $\lambda(\tau) = \lambda(\tau')$, and for every edge τ'' in the shortest path in T_G containing both τ and τ', $\lambda(\tau'') \geq \lambda(\tau)$. Notice that $E_{\tau'} \subseteq E_\tau$.

The decomposition (T_G, f_G) is linked, so Menger's Theorem implies there is a set $\mathcal{P}$ of $\lambda(\tau)$ PTD paths from $\Lambda(E_\tau)$ to $\Lambda(E_{\tau'})$. To see that

$$(G_{\tau'}, \Lambda(E_{\tau'})) \leq_r (G_\tau, \Lambda(E_\tau)),$$

contract the edges in the paths in $\mathcal{P}$ and delete the remaining edges of $E_\tau \setminus E_{\tau'}$. **Q.E.D.**

We are finally prepared to prove that graphs with bounded tree-width are well-quasi-ordered.

Proof of Theorem 1.2: Let w, $\mathcal{G}$, (T_G, f_G) (for each $G \in \mathcal{G}$), F and $\preceq$ be as described in the paragraphs preceding Lemma 2.10. If the edges of F are well-quasi-ordered by $\preceq$, then we claim that there are $G, H \in \mathcal{G}$ such that $G <_m H$. To see this, consider an edge τ of F incident with the root vertex r_G of some T_G. We have that $E_\tau = E(G)$, so $G_\tau = G$ and $\Lambda(E_\tau) = \emptyset$. Since the infinite set of such edges τ contains two, say τ and τ', in T_G and $T_{G'}$, respectively, such that $\tau \preceq \tau'$, we see that $(G, \emptyset) \leq_r (G', \emptyset)$, which implies that $G <_m G'$, as required. Thus, we can assume the edges of F are not well-quasi-ordered by $\preceq$.

Now let τ be an edge of F incident with a leaf t of some T_G, other than the root of T_G. Then either $t \notin f_G(E(G))$ or $t = f_G(e)$, for some edge e of G. In the former case, G_τ is an empty graph and $\Lambda(E_\tau) = \emptyset$. In the latter case, G_τ is either a loop or a link and $\Lambda(E_\tau)$ has 0, 1 or 2 vertices. Thus, there are, up to isomorphism, only six possibilities for the pair $(G_\tau, \Lambda(E_\tau))$. It follows easily that these edges of F are well-quasi-ordered by $\preceq$.

Apply Lemma 2.9 to F. (Note that (i) holds trivially and (ii) holds by Lemma 2.10.) The preceding pragraph implies we can assume that the set A guaranteed in the conclusion of Lemma 2.9 does not contain any edges incident with leaves of F, other than possibly those incident with root vertices. In particular, for each $\tau \in A$, we can assume there are two edges "above" τ in $U_F(A)$. (The preceding paragraph actually only shows we can assume there is at least one edge above τ. Since $\Delta(T_G) \leq 3$, there are at most two. If there is only one, we can add another one joined to a new leaf, not the image of any edge of G.) Arbitrarily designate one of these two edges as $L(\tau)$ and the other as $R(\tau)$ (for "left" and "right").

For each $\tau \in A$, consider the sets $\Lambda(E_{L(\tau)})$. Each of these sets has at most w elements. Thus, infinitely many of them have the same size. Let us consider only these ones. We shall use w_L for their common size.

Since the edges $L(\tau)$ are well-quasi-ordered by $\preceq$ (this is part of the conclusion of Lemma 2.9), we can find an infinite ascending chain $L(\tau_1) \preceq L(\tau_2) \preceq \ldots$.

Now we repeat the above two paragraphs for the edges $R(\tau_1)$, $R(\tau_2)$, ... This gives an infinite subsequence $\tau_{i_1}, \tau_{i_2}, \ldots$ of $\tau_1, \tau_2, \ldots$ such that: (i) all $\Lambda(E_{L(\tau_{i_j})})$ have the same size w_L; (ii) all $\Lambda(E_{R(\tau_{i_j})})$ have the same size w_R; (iii) $L(\tau_{i_1}) \preceq L(\tau_{i_2}) \preceq \ldots$; and (iv) $R(\tau_{i_1}) \preceq R(\tau_{i_2}) \preceq \ldots$. For ease of notation, we relabel so that the sequence $\tau_{i_1}, \tau_{i_2}, \ldots$ is $\tau_1, \tau_2, \ldots$.

For each i, label the elements of $\Lambda(E_{L(\tau_i)})$ with distinct elements from $\{1^L, 2^L, \ldots, w_L^L\}$ and label the elements of $\Lambda(E_{R(\tau_i)})$ with distinct elements from $\{1^R, 2^R, \ldots, w_R^R\}$. Since $L(\tau_1) \preceq L(\tau_2) \preceq \ldots$, we can choose the labellings so that when we recognize G_{τ_i} as a minor of G_{τ_j} (with $i < j$), the vertex of $\Lambda(E_{L(\tau_i)})$ with label k^L corresponds to the vertex of $\Lambda(E_{L(\tau_j)})$ with label k^L, and similarly for the R-labels.

What we are trying to prove is that for some $i < j$, $(G_{\tau_i}, \Lambda(E_{\tau_i})) \leq_r (G_{\tau_j}, \Lambda(E_{\tau_j}))$, which is a contradiction of the fact that A is an antichain. The issue at hand is, how do the graphs $G_{L(\tau_i)}$ and $G_{R(\tau_i)}$ glue together to make G_{τ_i}? We need that for some $i < j$, the two pairs glue together in the same way.

Vertices in $\Lambda(E_{L(\tau_i)}) \cap \Lambda(E_{R(\tau_i)})$ get both left and right labels. Since there are only finitely many different subsets of $\{1^L, 2^L, \ldots, w_L^L\}$, infinitely often it is the same one that gives the left labels in $\Lambda(E_{L(\tau_i)}) \cap \Lambda(E_{R(\tau_i)})$. Of these, infinitely often it is the same set of right labels. Thus, we may relabel (going to this subsequence) in such a way that, for every i, the set of left labels and the set of right labels occurring in $\Lambda(E_{L(\tau_i)}) \cap \Lambda(E_{R(\tau_i)})$ is always the same.

Notice that $\Lambda(E_{\tau_i}) \subseteq \Lambda(E_{L(\tau_i)}) \cup \Lambda(E_{R(\tau_i)})$. Thus, every vertex of $\Lambda(E_{\tau_i})$ has a label. Some will have only left labels, some will have only right labels, and others will have both. Again, there are only finitely many possibilities for which labels can appear and in which combinations. Thus, we can assume that: (i) the sizes of the $\Lambda(E_{\tau_i})$ are all the same; (ii) the set of left labels that occur in the $\Lambda(E_{\tau_i})$ is always the same; and (iii) the set of right labels that occur in the $\Lambda(E_{\tau_i})$ is always the same. Note that the preceding paragraph, (ii), and (iii) combine to show that the set of left-only labels is always the same and the set of right-only labels is always the same, because if a vertex has both a left-label and a right-label, then these two labels always go together.

Now, for $i < j$,

$$(G_{L(\tau_i)}, \Lambda(E_{L(\tau_i)})) \leq_r (G_{L(\tau_j)}, \Lambda(E_{L(\tau_j)}))$$

and

$$(G_{R(\tau_i)}, \Lambda(E_{R(\tau_i)})) \leq_r (G_{R(\tau_i)}, \Lambda(E_{R(\tau_i)})).$$

Let us show that $(G_{\tau_i}, \Lambda(E_{\tau_i})) \leq_r (G_{\tau_j}, \Lambda(E_{\tau_j}))$.

The graph G_{τ_i} is obtained from $G_{L(\tau_i)}$ and $G_{R(\tau_i)}$ by identifying the vertices in $\Lambda(E_{L(\tau_i)}) \cap \Lambda(E_{R(\tau_i)})$ — this identification is the same in both G_{τ_i} and G_{τ_j}. Thus, $(G_{\tau_i}, \Lambda(E_{\tau_i})) \leq_r (G_{\tau_i}, \Lambda(E_{\tau_i}))$, as required. **Q.E.D.**

3 Graphs with large tree-width have large grid minors

In this section we shall prove:

Theorem 1.3 *Let k be a positive integer. Then there is an integer $f(k)$ such that if G is a graph with branch-width at least $f(k)$, then G has a k-grid minor.*

In order to state the two main ingredients required to prove Theorem 1.3, we need some notation.

If $A, B \subseteq V(G)$, an (A, B)-*path* is a path with one end in A, the other end in B, and otherwise disjoint from $A \cup B$.

Finally, for $A, B \subset V(G)$, $\lambda(A, B)$ denotes the maximum size of a set of PTD (A, B)-paths.

Our first ingredient is:

Theorem 3.1 *Let k be a positive integer and let G be a graph. Let A, B, X, and Y be pairwise disjoint subsets of $V(G)$. Let $\mathcal{P}$ be a set of PTD (A, B)-paths. If $|\mathcal{P}| \geq k^{2k+1}(2k^2 + 1)$ and $\lambda(X, Y) \geq k^2$, then either G contains a k-grid minor or there is a path $P \in \mathcal{P}$ and an (X, Y)-path Q such that $P \cap Q$ is empty.*

Let W, X, Y be subsets of $V(G)$ and suppose W separates X from Y (i.e., there is no (X, Y)-path in $G - W$). Let E_W be the set of edges of G that are in some path P in G such that P starts at a vertex of $X \setminus W$ and $P \cap W$ consists of at most one end of P.

We note the following easy facts:

$\Lambda(E_W) \subseteq W$;

$\Lambda(E_W) \cup (W \cap (X \cup Y))$ separates X from Y in G.

Given a subset E of $E(G)$, let $\Lambda_{(X,Y)}(E)$ denote the union of $\Lambda(E)$ with all the $x \in X$ such that x is incident with an edge of $E(G) \setminus E$ and with all $y \in Y$ such that y is incident with an edge in E. It is a routine exercise to prove that $\Lambda_{(X,Y)}(E)$ separates X from Y. Furthermore,

$\Lambda_{(X,Y)}(E_W) \subseteq W$.

Thus, all minimal sets separating X from Y are of the form $\Lambda_{(X,Y)}(E)$, for some $E \subseteq E(G)$. We comment that $\Lambda_{(X,Y)}(E)$ is not, in general, the same as $\Lambda_{(Y,X)}(E)$; they will differ as to which elements of X and Y are included.

Let G be a graph and let $E \subseteq E(G)$. Then $\Lambda(E)$ is *k-connected in $G-E$* if, for any two subsets A, B of $\Lambda(E)$ such that $|A| = |B| \leq k$, there is a set of k PTD (A, B)-paths in $G - E$, each of which is disjoint from $\Lambda(E)$, except for its ends.

Our second ingredient is:

Theorem 3.2 *Let $h \geq k \geq 2$ be positive integers and let G be a graph. Then either G has branch-width at most h or there is a subset E of $E(G)$ such that: $\lambda(E) = h$; $\Lambda(E)$ is k-connected in $G - E$; and there is a path P such that $E(P) \subseteq E$ and P contracts in G onto $\Lambda(E)$.*

Both of these theorems are essentially due to [3], but the forms we present here seem slightly simpler, with various parts suggested by Jim Geelen and Paul Seymour.

Theorem 3.1 is based on some straightforward technical results about separating sets, which we discuss in Subsection 3.1. These technical results are then used in Subsection 3.2 to obtain two collections of paths that intersect in a controlled fashion, and finally Subsection 3.3 finishes the proof of Theorem 3.1 and provides the (easy) Corollary 3.10 that is a primary ingredient in proving that a graph with large branch-width has a large grid minor.

Subsection 3.4 gives the proof of Theorem 3.2. Finally, Subsection 3.5 presents the (now easy) proof of Theorem 1.3.

3.1 Uncrossing the separating sets

In the proof of Theorem 3.1, the first, relatively easy, step is to show that we can arrange for there to be a "nice" set of (X, Y)-separating sets.

We note some basic simplifications that we may apply in proving Theorem 3.1. If $k \leq 1$, then the result is a triviality, so we shall assume $k \geq 2$. If some path P in $\mathcal{P}$ has length 1, then there is an (X, Y)-path Q that is disjoint from P.

For an edge e of G not in any path in $\mathcal{P}$, if there is a set of k^2 PTD (X, Y)-paths in $G - e$, then substitute $G - e$ for G. Repeating this step until it is no longer possible to do so, we can assume that, for every edge e of G not in any path in $\mathcal{P}$, $G - e$ does not have a set of k^2 PTD (X, Y)-paths, yet G does have such a set of paths.

If e is an edge of a path in $\mathcal{P}$ such that G/e still has a set of k^2 PTD paths, then substitute G/e for G. Repeat until this is no longer possible. Thus, we can assume that for every edge e in every path in $\mathcal{P}$, G/e does not have a set of k^2 PTD (X,Y)-paths, but G does have such a set of paths.

Notice that if $\mathcal{Q}$ is any set of k^2 PTD (X,Y)-paths, then every path in $\mathcal{Q}$ is edge-disjoint from every path in $\mathcal{P}$ (since a common edge can be contracted).

Basic Reductions *We will assume henceforth that G has the additional properties:*

1. *if e is in a path of $\mathcal{P}$, then G/e does not have a set of k^2 PTD (X,Y)-paths;*

2. *if e is not in any path of $\mathcal{P}$, then $G-e$ does not have a set of k^2 PTD (X,Y)-paths; and*

3. *$k \geq 2$.*

Let $\mathcal{Q}$ be a set of k^2 PTD paths in G and let $Q_1 = (v_0, e_1, v_1, \ldots, e_m, v_m)$ be a path in $\mathcal{Q}$. By (2), for $i = 1, 2, \ldots, m$, $G - e_i$ does not have a set of k^2 PTD paths, so there is a subset E_i of $E(G)$ such that $e_i \in E_i$ and $|\Lambda_{(X,Y)}(E_i)| = k^2$. Thus, v_i is the only vertex of Q_1 in $\Lambda_{(X,Y)}(E_i)$.

We wish for the moment to prove:

Lemma 3.3 *There is a sequence $E'_1, E'_2, \ldots, E'_m$ of subsets of $E(G)$ such that:*

1. *$E'_1 \subset E'_2 \subset \ldots \subset E'_m$;*

2. *for $i = 1, 2, \ldots, m$, $e_i \in E'_i$;*

3. *for $i = 1, 2, \ldots, m$, $|\Lambda_{(X,Y)}(E'_i)| = k^2$.*

Proof: We have a sequence that satisfies the last two items, namely $E_1, E_2, \ldots, E_m$. For any sequence $E'_1, E'_2, \ldots, E'_m$ satisfying (2) and (3), let r be the largest index such that $E'_1 \subset E'_2 \subset \ldots \subset E'_r$. If $r < m$, then let $i \leq r$ be be smallest index such that $E'_i \not\subset E'_{r+1}$. Suppose that Claim 1 fails. Then for every sequence satisfying the last two items, both $r < m$ and i exist. Let $E'_1, E'_2, \ldots, E'_m$ be a sequence for which (r, i) is lexicographically largest.

We note the following routine exercise.

Claim *Let E, E' be subsets of $E(G)$.*

1. $\Lambda_{(X,Y)}(E \cap E') \cap \Lambda_{(X,Y)}(E \cup E') \subseteq \Lambda_{(X,Y)}(E) \cap \Lambda_{(X,Y)}(E')$;

2. $\Lambda_{(X,Y)}(E \cap E') \cup \Lambda_{(X,Y)}(E \cup E') \subseteq \Lambda_{(X,Y)}(E) \cup \Lambda_{(X,Y)}(E')$; *and*

3. $|\Lambda_{(X,Y)}(E \cap E')| + |\Lambda_{(X,Y)}(E \cup E')| \leq |\Lambda_{(X,Y)}(E)| + |\Lambda_{(X,Y)}(E')|$.

By the claim, the sequence

$$E'_1, E'_2, \ldots, E'_{i-1}, E'_i \cap E'_{r+1}, E'_{i+1}, \ldots, E'_r, E'_i \cup E'_{r+1}, E'_{r+2}, \ldots, E'_m$$

again satisfies (2) and (3) of the lemma. We note that the first r terms of this sequence still form an increasing (as subsets) sequence. By lexicographic maximality of the original sequence $E'_1, E'_2, \ldots, E'_m$, we must have $E'_r \not\subset E'_i \cup E'_{r+1}$. But notice that $E'_i \cap E'_{r+1} \subset E'_i \cup E'_{r+1}$, which contradicts maximality of i, a contradiction that proves the lemma.**Q.E.D.**

3.2 Controlling the Intersection

The point of this section is to prove the following.

Lemma 3.4 *Let ℓ and k be positive integers and let G be a graph. Let A, B, X, Y be pairwise disjoint subsets of $V(G)$. Let $\mathcal{P}$ be a set of PTD (A, B)-paths. If $|\mathcal{P}| \geq \ell(2k+1)$ and $\lambda(X, Y) \geq k$, then either there is a path $P \in \mathcal{P}$ and an (X, Y)-path Q such that $P \cap Q$ is empty or there is a subset $\mathcal{P}'$ of $\mathcal{P}$ and a set $\mathcal{Q}$ of k PTD (X, Y)-paths such that $|\mathcal{P}'| = \ell$ and every path in $\mathcal{Q}$ intersects all the paths in $\mathcal{P}'$ in the same order.*

Just to be completely clear, this last phrase means that if $P_1, P_2 \in \mathcal{P}'$ and $Q_1, Q_2 \in \mathcal{Q}$ are such that the vertices $u_1 \in V(P_1 \cap Q_1)$ and $u_2 \in V(P_2 \cap Q_1)$ occur in the order u_1, u_2 as we traverse Q_1 from U to V, then any vertex in $P_1 \cap Q_2$ occurs before any vertex in $P_2 \cap Q_2$, as we traverse Q_2 from U to V.

Proof: Let $\mathcal{Q}$ be a set of k PTD (X, Y)-paths and let $Q = (v_0, e_1, v_1, \ldots, e_m, v_m) \in \mathcal{Q}$. For $0 \leq i < j \leq m$, let $\mathcal{P}_{i,j}$ denote the set of paths P in $\mathcal{P}$ containing one of $v_i, v_{i+1}, \ldots, v_j$. Since the paths in $\mathcal{P}$ are PTD, clearly $|\mathcal{P}_{i,j+1}| \leq 1 + |\mathcal{P}_{i,j}|$. Furthermore, either $\mathcal{P}_{0,n} = \mathcal{P}$ or Q_1 is totally disjoint from some path in $\mathcal{P}$; in the latter case we are done.

Let $i_0 = 0$ and suppose $i_0, i_1, \ldots, i_{2p}$ have been determined, for some $p \geq 0$, so that: (a) $i_0 < i_1 < \ldots < i_{2p}$; and (b) for each $j = 1, 2, \ldots, t$, $i_{2j} = 1 + i_{2j-1}$; and (c) $|P_{i_{2j-2}, i_{2j-1}}| = 2k + 1$.

The index i_{2p+1} is the least $i > i_{2p}$ (if it exists) such that $|P_{i_{2p},i}| \geq 2k+1$. Note that this implies $|\mathcal{P}_{i_{2p},i_{2p+1}}| = 2k+1$ (since $|\mathcal{P}_{i,j+1}| \leq 1 + |\mathcal{P}_{i,j}|$). Then set $i_{2p+2} = 1 + i_{2p+1}$.

As long as $|\mathcal{P}| \geq (2k+1)(p+1)$, it is possible to find i_{2p+1}. Since $|\mathcal{P}| \geq (2k+1)\ell$, i_{2p+1} exists for $p = 0, 1, \ldots, \ell - 1$.

Some path P_p in $P \in \mathcal{P}_{i_{2p},i_{2p+1}}$ is disjoint from both $\Lambda_{(X,Y)}(E'_{2p})$ and $\Lambda_{(X,Y)}(E'_{i_{2p+1}})$, since there are at most k vertices in each of these sets and there are $2k+1$ PTD paths in the set $\mathcal{P}_{i_{2p},i_{2p+1}}$. Since P_p is incident with a vertex among $v_{i_{2p}}, v_{i_{2p}+1}, \ldots, v_{i_{2p+1}}$ and is disjoint from $\Lambda_{(X,Y)}(E'_{2p}) \cup \Lambda_{(X,Y)}(E'_{i_{2p+1}})$, we have $E(P_p) \subseteq E'_{i_{2p+1}} \setminus E'_{i_{2p}}$.

Let $Q \in \mathcal{Q}$ have vertex sequence $(w_0, w_1, \ldots, w_s)$, with $w_0 \in X$ and $w_s \in Y$. There are indices r, s such that $w_r \in \Lambda_{(X,Y)}(E'_{i_{2p}})$ and $w_s \in \Lambda_{(X,Y)}(E'_{i_{2p+1}})$. Since $E'_{i_{2p}} \subseteq E'_{i_{2p+1}}$, $r \leq s$.

If P_p is disjoint from Q, then the lemma is proved. Thus, we can assume P_p intersects Q at the vertex w_j. Since $E(P_p) \subseteq E'_{i_{2p+1}} \setminus E'_{i_{2p}}$ and P_p is disjoint from $\Lambda_{(X,Y)}(E'_{2p}) \cup \Lambda_{(X,Y)}(E'_{i_{2p+1}})$, we must have $r < j < s$. Thus, every vertex of $P_p \cap Q$ occurs between the intersection of Q with $\Lambda_{(X,Y)}(E'_{2p})$ and the intersection of Q with $\Lambda_{(X,Y)}(E'_{i_{2p+1}})$. That is, all the Q meet all the paths in $\mathcal{P}'$ in the same order, as required. **Q.E.D.**

3.3 Getting the grid

In this section we finish the proof of Theorem 3.1 and give several easy corollaries, the last of which is the one used in proving that graphs with large branch-width have large grid minors.

We begin with an easy fact.

Lemma 3.5 *Let $\ell \geq 2$ and $p \geq 1$ be integers. Let T be a tree with ℓ leaves and longest path of length at most p. Then $|V(T)| \leq (\ell-2)(p-1)+p+1$.*

Proof: We proceed by induction on ℓ, the base $\ell = 2$ being trivial. If $\ell > 2$, let P be a minimal path in T with one end a leaf and the other end adjacent to a vertex of degree at least 3 in T. Clearly, every vertex of P has degree at most two in T, the length of P is at most $p-2$, $T' = T - V(P)$ is a tree with $\ell - 1$ leaves, and $|V(T)| = |V(T')| + |V(P)|$. **Q.E.D.**

We are now ready for the final argument.

Lemma 3.6 *Let k and m be integers such that $k \geq 2$ and $m \geq k^{2k+1}$ and let G be a graph. Let A, B, X, Y be pairwise disjoint subsets of $V(G)$. Let $\mathcal{P}$ be a set of m PTD (A, B)-paths and let $\mathcal{Q}$ be a set of $k^2 - 3k + 5$ PTD (X, Y)-paths such that the paths in $\mathcal{P}$ are all edge-disjoint from the paths in $\mathcal{Q}$ and every path in $\mathcal{Q}$ meets the paths in $\mathcal{P}$ in the same order. Then G contains a k-grid minor.*

Proof: Let $\mathcal{P} = \{P_1, P_2, \ldots, P_m\}$ be such that every path in $\mathcal{Q}$ meets the paths in $\mathcal{P}$ in the order $P_1, P_2, \ldots, P_m$. For each $i = 1, 2, \ldots, m$, let K_i be the graph whose vertices are the paths in $\mathcal{Q}$ and $Q, Q' \in \mathcal{Q}$ are adjacent in K_i if there is a subpath of P_i that has one end in Q, one end in Q' and is otherwise disjoint from all the paths in $\mathcal{Q}$. Since P_i intersects all the paths in $\mathcal{Q}$, K_i is connected. Let T_i be a spanning tree of K_i.

Since T_i has $(k-2)(k-1-1) + (k-1) + 2$ vertices, Lemma 3.5 implies either T_i has at least k leaves or T_i has a path of length at least $k-1$. Let $(Q_{i,1}, Q_{i,2}, \ldots, Q_{i,k})$ be a sequence of vertices of T_i such that either they are all leaves of T_i or they are (in order) the vertex sequence of a path in T_i. Thus, for each $j = 2, 3, \ldots, k$, there is a subpath $P_{i,j}$ of P_i whose ends are in the paths $Q_{i,j-1}$ and $Q_{i,j}$ and is otherwise disjoint from $Q_{i,1} \cup Q_{i,2} \cup \ldots \cup Q_{i,k}$.

There are only $s = k^2 - 3k + 5$ paths in $\mathcal{Q}$, so there are only $s!/(s-k)! \leq s^k \leq k^{2k}$ different possible sequences $(Q_{i,1}, Q_{i,2}, \ldots, Q_{i,k})$. Thus, some sequence occurs at least $m/s^k \geq k$ times. Let $(Q_1, Q_2, \ldots, Q_k)$ be such a sequence, and let $i_1 < i_2 < \ldots < i_k$ be among those indices i such that

$$(Q_{i,1}, Q_{i,2}, \ldots, Q_{i,k}) = (Q_1, Q_2, \ldots, Q_k).$$

The Q_r, $r = 1, 2, \ldots, k$ are going to be the "horizontal" paths in the k-grid, while the P_{i_j}, $j = 1, 2, \ldots, k$, are going to provide the "vertical" paths in the k-grid. Specifically, the vertical paths are going to be the subpaths of P_{i_j} from Q_{r-1} to Q_r, $r = 2, 3, \ldots, k$, that are otherwise disjoint from $Q_1 \cup Q_2 \cup \ldots \cup Q_k$.

In order to see exactly the k-grid minor, we contract the portions of Q_r, $r = 2, 3, \ldots, k-1$, between the ends in Q_r of the segments of P_{i_j} joining Q_r to Q_{r-1} and Q_{r+1}. It is here that we use the fact that all the Q_r meet the P_{i_j} in the same order. **Q.E.D.**

We can now easily prove Theorem 3.1.

Proof of Theorem 3.1: Recall the Basic Reductions from Subsection 3.1. We assume that no (X, Y)-path is disjoint from any path in $\mathcal{P}$. By Lemma

3.4, there is a subset $\mathcal{P}'$ of $\mathcal{P}$ and a set $\mathcal{Q}$ of k^2 PTD (X, Y)-paths such that $|\mathcal{P}'| = k^{2k+1}$ and every path in $\mathcal{Q}$ intersects all the paths in $\mathcal{P}'$ in the same order.

Now we apply Lemma 3.6 to $\mathcal{P}'$ and $\mathcal{Q}$ to obtain the k-grid minor of G, as required. **Q.E.D.**

With the hard work done, we can readily move on through some improvements to Theorem 3.1. The functions $f_1(k) = k^{2k+1}(2k^2 + 1)$ and $f_2(k) = k^2$ from Theorem 3.1 will be used in the following corollaries. By abstracting away from the particular values of f_1 and f_2, the main points of these corollaries should be clearer. We emphasize that it is Corollary 3.10 that we are aiming for; the others are just steps along the way.

Corollary 3.7 *Let k and n be positive integers and let G be a graph. Let A, B, X, and Y be pairwise disjoint subsets of $V(G)$. Let $\mathcal{P}$ be a set of PTD (A, B)-paths. If $|\mathcal{P}| \geq f_1(k)$ and $\lambda(X, Y) \geq nf_1(k)f_2(k)$, then either G contains a k-grid minor or there is a path $P \in \mathcal{P}$ and a set $\mathcal{Q}$ of (X, Y)-paths such that $|\mathcal{Q}| \geq n$ and $\{P\} \cup \mathcal{Q}$ is a set of PTD paths.*

Proof: Let $\mathcal{Q}_0$ be a fixed set of PTD (X, Y)-paths such that $|\mathcal{Q}_0| = \lambda(X, Y)$. Let $X_1, X_2, \ldots, X_{nf_1(k)}$ be a set of pairwise disjoint subsets of X, each of size $f_2(k)$, such that each vertex x in each X_i is the origin of some path Q_x in $\mathcal{Q}_0$. Let Y_i denote the set of termini of the Q_x, $x \in X_i$. Finally, let $\mathcal{P}'$ be a subset of $\mathcal{P}$ of size $f_1(k)$.

Let G_i denote the union of the paths in $\mathcal{P}'$ and the paths in $\mathcal{Q}_0$ between X_i and Y_i. Theorem 3.1 applies to G_i, so either G has a k-grid minor or there is a path $P_i \in \mathcal{P}'$ and an (X_i, Y_i)-path Q_i such that $P_i \cap Q_i$ is empty. We are done if, for any $i = 1, 2, \ldots, n$, we have a k-grid minor of G_i.

Since there are $n|\mathcal{P}'|$ paths P_i, for some $P \in \mathcal{P}'$ there are at least n values of i for which $P = P_i$. Therefore, we may set $\mathcal{Q}$ to be the set $\{Q_i \mid P = P_i\}$. **Q.E.D.**

Corollary 3.8 *Let k, m, n be positive integers and let G be a graph. Let A, B, X, Y be pairwise disjoint subsets of $V(G)$. Let $\mathcal{P}$ be a set of PTD (A, B)-paths. If $|\mathcal{P}| \geq f_1(k) + m - 1$ and $\lambda(X, Y) \geq n(f_1(k)f_2(k))^m$, then either G contains a k-grid minor or there exist $\mathcal{P}' \subseteq \mathcal{P}$ and a set $\mathcal{Q}$ of (X, Y)-paths such that: $|\mathcal{P}'| \geq m$; $|\mathcal{Q}| \geq n$; and $\mathcal{P}' \cup \mathcal{Q}$ is a set of PTD paths.*

Proof: We proceed by induction on m, the case $m = 1$ being Corollary 3.7. For the inductive step, apply Corollary 3.7 to find either a k-grid minor of G or $P_1 \in \mathcal{P}$ and a set $\mathcal{Q}_1$ of (X, Y)-paths such that $|\mathcal{Q}_1| \geq n(f_1(k)f_2(k))^{m-1}$ and $\{P_1\} \cup \mathcal{Q}_1$ is a set of PTD paths.

Now apply the inductive assumption to the subgraph of G consisting of the union of the paths in $(\mathcal{P} \setminus \{P_1\}) \cup \mathcal{Q}_1$. **Q.E.D.**

Corollary 3.9 *Let k, m_i $(i = 1, 2, \ldots, r)$, n be positive integers and let G be a graph. Let A_i, B_i $(i = 1, 2, \ldots, r)$, X, Y be pairwise disjoint subsets of $V(G)$. For $i = 1, 2, \ldots, r$, let $\mathcal{P}_i$ be a set of PTD (A_i, B_i)-paths, so that $\cup_{i=1}^r \mathcal{P}_i$ is a set of PTD paths. If, for $i = 1, 2, \ldots, r$, $|\mathcal{P}_i| \geq f_1(k) + m_i - 1$ and $\lambda(X, Y) \geq n(f_1(k)f_2(k))^{(\sum_{i=1}^r m_i)}$, then either G contains a k-grid minor or there exist, for $i = 1, 2, \ldots, r$, $\mathcal{P}'_i \subseteq \mathcal{P}_i$ and a set $\mathcal{Q}$ of (X, Y)-paths such that: for $i = 1, 2, \ldots, r$, $|\mathcal{P}'_i| \geq m_i$; $|\mathcal{Q}| \geq n$; and $(\cup_{i=1}^r \mathcal{P}'_i) \cup \mathcal{Q}$ is a set of PTD paths.*

Proof: We proceed by induction on r, the case $r = 1$ being Corollary 3.8. If $r > 1$, by Corollary 3.8 either G contains a k-grid minor or there is a set $\mathcal{Q}_1$ of (X, Y)-paths and a subset $\mathcal{P}'_1 \subseteq \mathcal{P}_1$ such that: $|\mathcal{P}'_1| \geq m_1$; $|\mathcal{Q}_1| \geq n(f_1(k)f_2(k))^{\sum_{i=2}^r m_i}$; and $\mathcal{P}'_i \cup \mathcal{Q}_1$ is a set of PTD paths. The result now follows easily by the inductive assumption. **Q.E.D.**

Corollary 3.10 *Let k, m_i $(i = 1, 2, \ldots, r)$, n_j $(j = 1, 2, \ldots, s)$ be positive integers and let G be a graph. Let (A_i, B_i), $i = 1, 2, \ldots, r$, (X_j, Y_j), $j = 1, 2, \ldots, s$ be pairwise disjoint subsets of $V(G)$. For $i = 1, 2, \ldots, r$, let $\mathcal{P}_i$ be a set of (A_i, B_i)-paths so that $\cup_{i=1}^r \mathcal{P}_i$ is a set of PTD paths. Set $M = \sum_{i=1}^r m_i + \sum_{j=1}^s n_j$. If, for $i = 1, 2, \ldots, r$, $|\mathcal{P}_i| \geq sf_1(k) + m_i - s$ and, for each $j = 1, 2, \ldots, s$,*

$$\lambda(X_j, Y_j) \geq (sf_1(k) + n_i)(f_1(k)f_2(k))^{(s(r+s)f_1(k)+M)},$$

then either G contains a k-grid minor or there exist, for $i = 1, 2, \ldots, r$, $\mathcal{P}'_i \subseteq \mathcal{P}_i$ and, for $j = 1, 2, \ldots, s$, a set $\mathcal{Q}_j$ of (X_j, Y_j)-paths such that: for $i = 1, 2, \ldots, r$, $|\mathcal{P}'_i| \geq m_i$; for $j = 1, 2, \ldots, s$, $|\mathcal{Q}_j| \geq n_j$; and $(\cup_{i=1}^r \mathcal{P}'_i) \cup (\cup_{j=1}^s \mathcal{Q}_j)$ is a set of PTD paths.

Proof: We proceed by induction on s. If $s = 1$, then the result follows from Corollary 3.9. For $s > 1$, apply Corollary 3.9 to the (A_i, B_i) and (X_1, Y_1), to find either a k-grid minor or $\mathcal{P}'_i \subseteq \mathcal{P}_i$, $i = 1, 2, \ldots, r$ and a set $\mathcal{Q}_1$ of (X_1, Y_1)-paths such that: $|\mathcal{P}'_i| \geq (s-1)f_1(k) + m_i - 1$; $|\mathcal{Q}_1| \geq (s-1)f_1(k) + n_1 - 1$; and $(\cup_{i=1}^{r} \mathcal{P}'_i) \cup \mathcal{Q}_1$ is a set of PTD paths. The result now follows easily by induction. **Q.E.D.**

We note that in Corollary 3.10 we can allow $r = 0$ (in which case we set $(A_1, B_1) = (X_1, Y_1)$ and let $\mathcal{P}_1$ be any set of $\lambda(X_1, Y_1)$ PTD paths) and we can allow $s = 0$ (a triviality). We will actually use this result when $r = 0$.

3.4 Branch-width and connected sets

We are now ready to prove Theorem 3.2, which we restate here for easy reference.

Theorem 3.2 *Let $h \geq k \geq 2$ be positive integers and let G be a graph. Then either G has branch-width at most h or there is a subset E of $E(G)$ such that: $\lambda(E) = h$; $\Lambda(E)$ is k-connected in $G - E$; and there is a path P such that $E(P) \subseteq E$ and P contracts in G onto $\Lambda(E)$.*

The correct image for the path P is the following: $P = (v_0, v_1, \ldots, v_r)$ has its ends $v_0, v_r \in \Lambda(E)$ and there are $0 < i_1 < \ldots < i_s < r$ and a set $\{P_1, P_2, \ldots, P_s\}$ of $s = |\Lambda(E)| - 2$ PTD paths such that P_j joins v_{i_j} to a vertex of $\Lambda(E)$, and each P_j meets P only in v_{i_j}.

Proof: We suppose the branch-width of G is more than h. Let E be a maximal subset of $E(G)$ such that:

1. $\lambda(E) \leq h$;

2. if (T, f) is a branch-decomposition such that $E = E_\tau$ for some edge τ of T, then some edge τ' in the component of $T - \tau$ not containing $f(E)$ has $\zeta(\tau') > h$; and

3. there is a path P in G such that $E(P) \subseteq E$ and P contracts in G onto $\Lambda(E)$.

To see that E exists, it suffices to observe that $\emptyset$ satisfies all three conditions. Moreover, (2) implies $E \neq E(G)$.

Part 1: $\lambda(E) = h$.

For if it is less than h, let $v \in \Lambda(E)$ be an end of P. By definition of $\Lambda(E)$, there is an edge $e \in E(G) \setminus E$ incident with v. Clearly, $\lambda(E \cup \{e\}) \leq 1 + \lambda(E) \leq h$, and there is a path P' whose edges are all in $E \cup \{e\}$ that contracts onto $\Lambda(E \cup \{e\})$.

The contradiction comes from (2). By maximality of E, there is a branch-decomposition (T, f) of G such that $E \cup \{e\} = E_\tau$, for some edge τ of T, and every edge τ' in the component of $T - \tau$ not containing $f(e)$ has $\zeta(\tau') \leq h$.

Let T_1 and T_2 be the components of $T - \tau$ so that $f(e)$ is a vertex of T_1. Let (T', f') be the branch-decomposition of G obtained as follows. The vertices of T' are those of T, with two new vertices t_1, t_2, which are adjacent in T'. In addition, t_2 is adjacent to the two vertices of T incident with τ. All other adjacencies in T' are in complete agreement with adjacency in $T - \tau$.

We let $f'(e) = t_1$, and otherwise $f'(e') = f(e')$. Clearly: the label $\zeta(t_1 t_2)$ is at most 2; the edge joining T_1 to t_2 has label $\lambda(E) < h$; the edge joining t_2 to T_2 has label $\lambda(E \cup \{e\}) \leq h$; and every edge in T_2 has the same label as in (T, f), which is to say these labels are all $\leq h$, contradicting (2). Hence, $\lambda(E) = h$, as claimed.

Part 2: $\Lambda(E)$ is k-connected in $G - E$.

If not, there are subsets A, B of $\Lambda(E)$ such that $|A| = |B| = s \leq k$ and there is no set of s PTD (A, B)-paths in $G - E$, each path disjoint from $\Lambda(E)$ except for its ends. Choose A and B so that s is as small as possible.

Let H be the subgraph of G induced by the edges in $E(G) \setminus E$, with the vertices of $\Lambda(E) \setminus (A \cup B)$ deleted. (Since we need to work in both G and H, for the next few paragraphs, we will include, as a superscript, the name of the graph (either G or H) in which we are determining $\Lambda(F)$, for some set F of edges.)

Let $E' \subseteq E(H)$ be such that $S = \Lambda^H_{(A,B)}(E')$is a smallest set that separates A from B in H. Clearly, $|S| = s - 1$ (by the choice of s). By Menger's Theorem, there is a set $\mathcal{Q}$ of $s - 1$ PTD (A, S)-paths in H.

Claim 1 $\Lambda^G(E \cup E') \subseteq (\Lambda^G(E) \setminus A) \cup S$.

Proof: Suppose $z \in \Lambda^G(E \cup E')$. Then z is incident with an edge $\bar{e} \notin E \cup E'$ and an edge $e \in E \cup E'$. If $e \in E$, then $z \in \Lambda^G(E)$, while if $e \in E'$, then $z \in \Lambda^H(E') \subseteq S$. Thus, $\Lambda^G(E \cup E') \subseteq \Lambda^G(E) \cup S$.

Let $u \in A \setminus S$. Then u is not incident with an edge of $E(H) \setminus E'$, so every edge in G incident with u is in $E \cup E'$. Hence $u \notin \Lambda^G(E \cup E')$. **Q.E.D.**

It is an immediate consequence of Claim 1 that $\lambda_G(E \cup E') < h$ (since $A \subseteq \Lambda^G(E)$ and $|S| < |A|$). The set $\mathcal{Q}$ of PTD (A, S)-paths shows there is a path P' using only edges in $E \cup E'$ that contracts to $\Lambda^(E \cup E')$. The maximality of E shows that there is a branch-decomposition (T_A, f_A) of G such that, for some edge τ_A of T_A, $f_A(E \cup E')$ is contained in the set of leaves of the component T_A^1 of $T_A - \tau_A$ and $f_A(E(G) \setminus (E \cup E'))$ is contained in the set of leaves of the other component T_A^2 of $T_A - \tau_A$ so that, for $\tau \in E(T_A^2)$, $\zeta(\tau) \leq h$.

In exactly the same way, we can proceed with B in place of A. For ease of notation, let $E'' = E(H) \setminus E'$. There is a branch-decomposition (T_B, f_B) of G and an edge τ_B of T_B such that $f_B(E \cup E'')$ is contained in the leaves of the component T_B^1 of $T_B - \tau_B$ and $f_B(E(G) \setminus (E \cup E''))$ is contained in the leaves of the other component T_B^2 of $T_B - \tau_B$ so that, for $\tau \in E(T_B^2)$, $\zeta(\tau) \leq h$.

Now we describe a third branch-decomposition (T, f) of G. Let T be any tree with maximum degree at most three containing a vertex t such that $T - t$ has three components, one isomorphic to T_A^2, one isomorphic to T_B^2, and the third, T_3, having at least $|E|$ leaves that are leaves of T. Moreover, the edge joining T_A^2 to t is incident with the end of τ_A in T_A^2, and similarly for the edge joining T_B^2 to t. Define f by $f(a) = f_A(a)$, for $a \in E''$, $f(b) = f_B(b)$, for $b \in E'$, and, for $e \in E$, $f(e)$ is any leaf of T_3 that is a leaf of T, chosen so that f is an injection. It is easy to see that if τ is the edge of T joining T_3 to t, then $E = E_\tau$, while if τ' is any edge of $T_A^2 \cup T_B^2$, then $\zeta(\tau') \leq h$, contradicting the choice of E.

Therefore, $\Lambda(E)$ is k-connected in E. **Q.E.D.**

3.5 The final details

We are now ready to move on to the main result of this section. We restate it here for ease of reference.

Theorem 1.3 *Let k be a positive integer. If G is a graph with branch-width at least $2k^2(2k^2 f_1(k) - 2k^2 + 2)(f_1(k)f_2(k))^{(4k^4(f_1(k)-1)+2k^2)} + 1$, then G has a k-grid minor.*

As a brief comment, recall $f_1(k) = k^{2k+1}(2k^2+1)$ and $f_2(k) = k^2$. Thus, the branch-width we require for a k-grid minor is doubly exponential in k, i.e., roughly $k^{(8k^{2k+5})}$. The best result to date is by Robertson, Seymour and

Thomas [10], who proved that tree-width of order 2^{9k^5} suffices to ensure the existence of a k-grid minor. Not surprisingly, this requires more sophisticated methods than those used here.

Proof: Since G does not have branch-width

$$F(k) = 2k^2(2k^2 f_1(k) - 2k^2 + 2)(f_1(k) f_2(k))^{(4k^4(f_1(k)-1)+2k^2)},$$

by Theorem 3.2, there is a subset E of G such that: $\lambda(E) = F(k)$; $\Lambda(E)$ is $(F(k)/k^2)$-connected in $E(G) \setminus E$; and there is a path P such that $E(P) \subseteq E$ and P contracts in G onto $\Lambda(E)$. Let $(v_1, v_2, \ldots, v_{F(k)})$ be the vertices of $\Lambda(E)$ in the order they appear in the contraction of P. Let $V_1, V_2, \ldots, V_{k^2}$ be the partition of $\Lambda(E)$ obtained by putting $F(k)/k^2$ consecutive vertices into each of these sets.

Thinking of the path P as giving us a Hamilton path in the k-grid minor that starts and ends in different corners of the grid, we see that we need only to connect in pairs some of the sets V_i in order to make a k-grid minor. We will attempt to join these pairs of V_i using only edges in $E(G) \setminus E$, in order not to interfere with the Hamilton path. Each V_i is to be joined to at most two others.

For each i, let j and j' be the indices such that we are trying to join V_i to both V_j and $V_{j'}$. (Some i have only one such index j.) Arbitrarily partition V_i into two sets $V_{i,j}$ and $V_{i,j'}$ of size $F(k)/2k^2$.

Now we consider the pairs $(V_{i,j}, V_{j,i})$. Note that, for each pair i, j,

$$\lambda(V_{i,j}, V_{j,i}) \geq (2k^2 f_1(k) - 2k^2 + 2)(f_1(k) f_2(k))^{(4k^4(f_1(k)-1)+2k^2)}.$$

Thus, we can apply Corollary 3.10 with $r = 0$ and $s \leq 2k^2$; we deduce that either there is a k-grid minor in G or we can find disjoint paths in $E(G) \setminus E$ between the required pairs so that we complete a k-grid minor of G. In either case, G has a k-grid minor. **Q.E.D.**

4 Embedding large grids in surfaces

In this section we shall prove Theorem 1.4. The presentation here is based on Geelen, Richter, and Salazar [6].

A *surface* is a compact connected 2-manifold with boundary. It is well-known that every surface is homemomorphic either to a *sphere with handles* or a *sphere with crosscaps*. (For a comprehensible proof of this classification of surfaces, see [13].) We take as known the fact that if a surface Σ is a sphere with h handles and k crosscaps, then either $k = 0$ and Σ is a sphere with h handles or $k > 0$ and Σ is homeomorphic to the sphere with $2h + k$ crosscaps. Thus, the *Euler genus* $\varepsilon(\Sigma) = 2h + k$ is an invariant of Σ. (It is related to the Euler characteristic $\chi(\Sigma)$ by $\varepsilon(\Sigma) = 2 - \chi(\Sigma)$. The torus and Klein bottle both have $\varepsilon = 2$.)

The main topological component we need is the following. (Here, and for the rest of this work, if $\mathcal{C}$ is a set, then $\cup\mathcal{C} = \cup_{\gamma\in\mathcal{C}}\gamma$.)

Lemma 4.1 *Let Σ be a surface and let $\mathcal{C}$ be a set of pairwise disjoint simple closed noncontractible curves in Σ such that some component of $\Sigma \setminus (\cup\mathcal{C})$ has all the curves in $\mathcal{C}$ in its boundary. Then $|\mathcal{C}| \leq \varepsilon(\Sigma)$.*

Proof: We may assume that $|\mathcal{C}|$ is finite. Let Γ be a component of $\Sigma\setminus(\cup\mathcal{C})$ such that all the curves of $\mathcal{C}$ are contained in the boundary of Γ. Suppose that $\mathcal{C}$ contains a separating curve C, and let Σ_1 and Σ_2 be the components of $\Sigma \setminus C$, where Γ is contained in Σ_1. Since C is noncontractible, Σ_2 is not a disc, and, hence, contains a nonseparating closed curve C'. Replacing C with C' in $\mathcal{C}$ reduces the number of separating curves in $\mathcal{C}$. Thus, we may assume that $\mathcal{C}$ contains only nonseparating curves.

Let n be the number of components of $\Sigma \setminus (\cup\mathcal{C})$, let h be the number of 2–sided curves in $\mathcal{C}$, and let k be the number of 1–sided curves in $\mathcal{C}$. Since every curve is incident with Γ and every 1–sided curve is incident only with Γ, some collection of $h - (n - 1)$ 2–sided curves in $\mathcal{C}$ does not separate Σ. Each curve in this collection contributes a handle to Σ. Since the 1–sided curves are pairwise disjoint, each contributes a cross-cap to Σ. Thus, $\varepsilon(\Sigma) \geq 2(h - (n - 1)) + k$.

If $n = 1$, then $\varepsilon(\Sigma) \geq 2h + k \geq h + k = |\mathcal{C}|$, as required. Thus, we may assume that $n > 1$.

As $\mathcal{C}$ does not contain a separating curve, every component of $\Sigma \setminus (\cup\mathcal{C})$ other than Γ has at least two curves in its boundary. Each curve in the boundary of such a component is 2-sided and has Γ on the other side, so $h \geq 2(n-1)$. Therefore, $\varepsilon(\Sigma) \geq 2(h-(n-1))+k \geq h+k = |\mathcal{C}|$, as required. **Q.E.D.**

We need to turn Lemma 4.1 into a form that is useful for graphs. A subgraph H of a graph G *separates* G if H has more than one bridge in G.

Lemma 4.2 *Let G be a connected graph embedded in a surface Σ and let $\mathcal{C}$ be a set of pairwise disjoint cycles in G such that $\cup\mathcal{C}$ does not separate G. If every cycle in $\mathcal{C}$ is noncontractible in Σ, then $|\mathcal{C}| \leq \varepsilon(\Sigma)$.*

Proof: Let $H = \cup\mathcal{C}$. Since H does not separate G, there is one component of $\Sigma \setminus H$ in which the rest of G is embedded. Since G is connected, each cycle in H is in the boundary of this component. By Lemma 4.1, $|\mathcal{C}| \leq \varepsilon(\Sigma)$. **Q.E.D.**

More specifically, we focus on embeddings of grids. We use Lemma 4.2 to show that if a large grid is embedded in a surface, then some large subgrid is embedded in a disc (Lemma 4.3) and some cylinder that "goes around the centre" is also embedded in a disc (Lemma 4.4).

The *boundary cycle* of a k-grid is the cycle that bounds the infinite face in the usual planar embedding of the grid (that is, consists of the subgraph induced by the $v_{i,j}$ such that $\{i, j\} \cap \{1, k\} \neq \emptyset$).

Lemma 4.3 *Let t, k, n be positive integers such that $n \geq t(k + 1)$ and let G be an n-grid. If G is embedded in a surface Σ of Euler genus at most $t^2 - 1$, then some k-subgrid of G is embedded in a closed disc in Σ such that the boundary cycle of the k-grid is the boundary of the disc.*

Proof: Clearly G contains t^2 pairwise disjoint k-subgrids and no two vertices from distinct subgrids are adjacent in the grid. By Lemma 4.2, not all of these subgrids can have noncontractible 4-cycles, so one has only contractible 4-cycles. Each of these bounds a closed disc and the union of these closed discs is the required closed disc. **Q.E.D.**

For the cylinder, let G be a k-grid. Let $1 \leq t \leq k/2$ be a given integer. A *t-collar* is a subgraph of G induced by, for some positive integer $i \leq k/2-t+1$, the vertices of G at distance at least $i - 1$ and at most $i + t - 1$ from the boundary cycle of the grid. The *exterior cycle* and *interior cycle* of this t-collar are the cycles induced by the vertices at distance $i - 1$ and distance $i + t - 1$, respectively. We note that the exterior cycle is the boundary cycle of a $(k - 2i + 2)$-subgrid.

We will only use 2-collars.

Lemma 4.4 *Let Σ be a surface of Euler genus ε and let t be a positive integer. Let $k \geq 2(\varepsilon + 1)(t + 2)$. Let G be a k-grid embedded in Σ. Then G contains a t-collar embedded in a cylinder in Σ.*

Proof: For each $i = 1, 2, \ldots, \lfloor k/2 \rfloor - t + 1$, let C_i be the t-collar consisting of the vertices at distance at least $i-1$ and at most $i+t-1$ from the boundary cycle of G.

By Lemma 4.2, one of the t-collars C_i, $i \in \{1, 1+(t+2), 1+2(t+2), \ldots, 1+\varepsilon(t+2)\}$, contains only contractible 4-cycles. For one in which all 4-cycles are contractible, the 4-cycles bound discs and the union of these discs is a cylinder. **Q.E.D.**

We are now prepared for the proof of Theorem 1.4. More precisely, we prove the following.

Theorem 4.5 *Let Σ be a surface of Euler genus ε. Let*

$$n \geq \left\lceil \sqrt{\varepsilon + 3} \right\rceil (8\varepsilon + 9)$$

and let G be a graph containing an n-grid as a minor. If, for every $e \in E(G)$, $G - e$ embeds in Σ, then G embeds in Σ.

An edge e of a k-grid is *central* if either (1) $k = 2m - 1$ is odd and e is incident with $v_{m,m}$ or (2) $k = 2m$ is even and e is in the 4-cycle induced by $v_{m,m}$, $v_{m,m+1}$, $v_{m+1,m}$, and $v_{m+1,m+1}$.

Proof: Let e be any edge of G. An embedding of $G - e$ in Σ can be used to obtain an embedding of G in Σ', the surface obtained from Σ by adding a handle.

Let M be a minimal subgraph of G that contracts to the n-grid G_n. The embedding of M in Σ' contracts (in Σ') to an embedding of G_n.

Since $\varepsilon(\Sigma') = \varepsilon(\Sigma) + 2 = \varepsilon + 2$, Lemma 4.3 implies that some $(8\varepsilon + 9)$-subgrid G' of G_n is embedded in a disc in Σ'. Let K be a minimal subgraph of M that contracts to G'. Then K is also embedded in a disc.

Let e be a central edge of K. Let K' be a minimal subgraph of $K - e$ that contracts to an $(8\varepsilon + 8)$-grid.

Embed $G - e$ in Σ. By Lemma 4.4, there is a 2-collar J of K' that is embedded in a cylinder. Let I be the interior cycle of J, let E be the exterior cycle of J, and let C be the cycle in $J - V(I \cup E)$.

Let B_E be the bridge of C in G that contains E, let B_I be the bridge of C in G that contains I, and let $\mathcal{B}$ denote the set of all other bridges of C. We claim that we can arrange the embeddings of G in Σ' and $G - e$ in Σ so that if $B \in \mathcal{B}$, then B is in the cylinder bounded by $C \cup I$ either in both embeddings or in neither embedding.

Let $\mathcal{B}'$ denote the subset of $\mathcal{B}$ consisting of those bridges that are in the cylinder bounded by $C \cup I$ in one embedding but not the other. Let $\mathcal{O}$ denote the "overlap diagram" for the bridges in $\mathcal{B}'$: its vertices are the bridges in $\mathcal{B}'$ and two bridges are adjacent in $\mathcal{O}$ if they cannot be simultaneously embedded on the same side of C.

Since every bridge in $\mathcal{B}$ must be embedded in the cylinder bounded by $E \cup I$, $\mathcal{O}$ is bipartite. One side of the bipartition corresponds to those bridges that are inside the cylinder bounded by $C \cup I$ in one embedding and the other side of the bipartition corresponds to the bridges that are outside in the same embedding. Because we are in the cylinder, we can simply switch the embeddings of the bridges in $\mathcal{B}'$ in one of the two embeddings, so that each one is either in both cylinders or in neither cylinder.

Thus, the subgraphs of G that are contained in the cylinders bounded by $C \cup I$ are the same in both embeddings.

In order to obtain an embedding of G in Σ, let Δ_1 be the disc in Σ' bounded by C and let Δ_2 be a small closed disc in Δ_1, disjoint from G. Let Γ be the cylinder in Σ' obtained by deleting the interior of Δ_2 from Δ_1. It is bounded by C and some other simple closed curve.

Identifying the two copies of C, replace the cylinder in Σ bounded by $C \cup I$ with the cylinder Γ. This, together with the cylinder bounded by $C \cup E$ and the embedding of B_E in Σ yields an embedding of G in Σ. **Q.E.D.**

References

[1] R. Diestel, Graph Theory, Second edition, Springer-Verlag, New York, 2000.

[2] R. Diestel, *Tree-width duality*, preprint.

[3] R. Diestel, Y. Gorbunov, T. Jensen, and C. Thomassen, *Highly connected sets and the excluded grid theorem*, J. Combin. Theory Ser. B, **75** (1999), 61–73.

[4] J. Geelen, B. Gerards, and G. Whittle, *Branch-width in matroids*, preprint.

[5] J. Geelen, B. Gerards, and G. Whittle, *Tree-like decompositions*, in preparation.

[6] J. Geelen, R.B. Richter, and G. Salazar, *Embedding grids in surfaces*, preprint.

[7] N. Robertson and P.D. Seymour, *Graph minors IV. Tree-width and well-quasi-ordering*, J. Combin. Theory Ser. B **48** (1990), 227-254.

[8] N. Roberston and P.D. Seymour, *Graph Minors XVII. Taming a vortex*, J. Combin. Theory Ser. B **77** (1999), 162-210.

[9] N. Robertson and P.D. Seymour, *Graph Minors XX. Wagner's Conjecture*, manuscript (January 1996).

[10] N. Robertson, P.D. Seymour, and R. Thomas, *Quickly excluding a planar graph*, J. Combin. Theory Ser. B **62** (1994), 323-348.

[11] R. Thomas, *A Menger-like property of tree-width: the finite case*, J. Combin. Theory Ser. B **48** (1990), 67-76.

[12] C. Thomassen, *A simpler proof of the excluded minor theorem for higher surfaces*, J. Combin. Theory Ser. B **70** (1997), 306-311.

[13] C. Thomassen, *The Jordan-Schönflies theorem and the classification of surfaces*, Amer. Math. Monthly **99** (1992), 116-130.

[14] K. Wagner, *Graphentheory*, BJ Hochsultashenbucher 248/248a, Mannheim, 1970, 61.