The Graph Minor Theorem

June 9, 2015

Contents

1 Introduction

Let $G = (V, E)$ be a simple graph. We write $u \sim v$ if $\{u, v\} \in E$. We say that G is a *minor* of a graph $G' = (V', E')$ if G can be obtained from G' by a sequence of edge and vertex deletions and edge contractions. Equivalently, there is a function $\pi : V' \to V$ such that for every $u \in G$, $\pi^{-1}(v)$ is connected and if $u \sim v$ in G, then $u' \sim v'$ in G' for some $u' \in \pi^{-1}(u)$ and $v' \in \pi^{-1}(v)$.

If G and G' are *rooted graphs*, that is, there are distinguished vertices $r \in V$ and $r' \in V'$, then G is a *rooted minor* of G' if the above conditions are satisfied and $\pi(r') = r$.

Define a relation $\leq$ on the set of finite graphs $\mathcal{G}$ by letting $G \leq G'$ if G is isomorphic to a minor of G'. Clearly this relation is reflexive (take $\pi = \mathrm{id}_{V(G)}$ in the definition above) and transitive (if $G \leq G' \leq G''$ with functions

$$V(G'') \xrightarrow{\pi_1} V(G') \xrightarrow{\pi_2} V(G)$$

as in the definition above, take $\pi = \pi_2 \circ \pi_1$). This makes $\leq$ a *quasi-order* on $\mathcal{G}$.

Furthermore, it is easy to see that if $G \leq G'$ and $G' \leq G$, then G and G' are isomorphic. Thus we obtain from $\preceq$ a partial order on the set of isomorphism classes of graphs.

Recall that $\preceq$ is a *well-quasi order* on $\mathcal{G}$ is a quasi-order such that for any sequence $x_1, x_2, \ldots$ in $\mathcal{G}$ there exist $i < j$ with $x_i \preceq x_j$. This is equivalent to the following: $\preceq$ has no infinite antichain (set of incomparable elements) and is *well-founded* (every nonempty subset of $\mathcal{G}$ has a minimal element with respect to $\preceq$). A consequence of this is that if $\mathcal{H} \subset \mathcal{G}$ is an upper set with respect to $\preceq$ (that is, if $G \in \mathcal{H}$ and $G \preceq G'$, then $G' \in \mathcal{H}$), then $\mathcal{H}$ has finitely many minimal elements $G_1, \ldots, G_n$, and $G \in \mathcal{H}$ if and only if $G_i \preceq G$ for some i.

Robertson and Seymour in a series of papers starting with [1] prove the so-called Wagner's conjecture: that the relation $\leq$ defined above is a well-quasi order.

The equivalent condition for well-quasi ordering can then be restated: if $\mathcal{H}$ is some set of finite graphs such that if $G \in \mathcal{H}$, then any minor of G is in $\mathcal{H}$, then there is a finite set of graphs $G_1, \ldots, G_n$ such that $G \in \mathcal{H}$ if and only if G does not contain a minor isomorphic to any G_i. We say that $\mathcal{H}$ is characterized by the *forbidden minors* $G_1, \ldots, G_n$.

Indeed, a number of interesting properties of graphs that are preserved under deletions and edge contractions can be characterized by forbidden minors. For example:

- A graph is a forest (has no cycles) if and only if it has no minor isomorphic to the triangle K_3. (Proof is trivial.)
- A graph is planar if and only if it has no minor isomorphic to K_5 or to $K_{3,3}$ (see Figure 1). This is Wagner's theorem.
- A graph is outerplanar (i.e., can be embedded in the plane so that all vertices are in the outer face) if and only if it has no minor isomorphic to K_4 or to $K_{2,3}$. This is not difficult to show using the fact that outerplanar graphs can be decomposed into triangulated polygons.
- A graph can be embedded in $\mathbb{RP}^2$ if and only if it has no minor isomorphic to any of a set of 35 forbidden minors ([4]).

Wagner's conjecture implies that *all* properties preserved under deletions and contractions can be characterized by forbidden minors. (Important cases are sets of graphs that can be embedded in a given surface.)

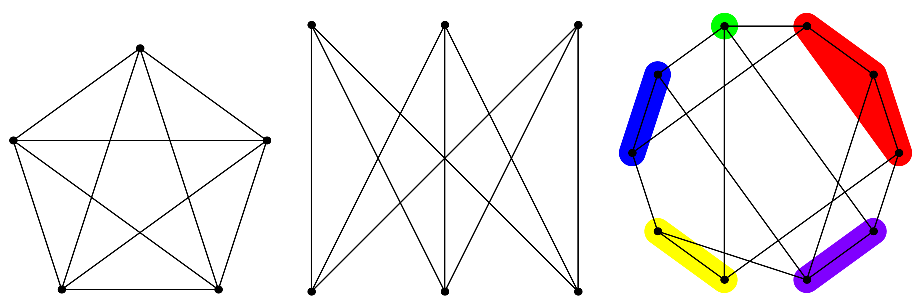

Figure 1: The two forbidden minors for planar graphs and a nonplanar graph with K_5 minor highlighted.

2 Properties excluding a planar graph

The goal of [1], [3], [4], and [5] is to show the following.

Theorem 1. *If $\mathcal{H}$ is a set of finite graphs closed under taking minors, and $\mathcal{H}$ does not contain some planar graph, then $\mathcal{H}$ can be characterized by forbidden minors.*

This is a special case of Wagner's conjecture, which omits the supposition that $\mathbb{H}$ does not contain some planar graph. We give an outline of part of Robertson and Seymour's largely self-contained proofs of Theorem 1.

2.1 Path-width and tree-width

Path-width and tree-width are invariants associated with a graph that are important to the proof of Theorem 1. We state the definitions given in [1] and [2] in clearer terms.

Let G be a graph. A *path-decomposition* of G consists of a path P (with vertices $v_1, v_2, \ldots, v_n$ connected in sequence) together with a function ρ mapping each vertex of G to a subset of the vertices of P (that is, an element of $\mathcal{P}(V(P))$, where $\mathcal{P}$ denotes the power set) such that:

(1) For each $u \in V(G)$, the subgraph of P induced by $\rho(u)$ is connected.

(2) If $u \sim v$ in G, then $\rho(u) \cap \rho(v) \neq \emptyset$.

(We may otherwise define the decomposition by associating a subset X_i of $V(G)$ to each vertex v_i of P and requiring that if $u \sim v$, then some X_i contains both u and v and that if $i \leq j \leq k$, then $X_i \cap X_k \subseteq X_j$. The two definitions are equivalent: given the function ρ, we may define the $X_i = \{u \in V(G) : v_i \in \rho(u)\}$. In the other direction, given the X_i, we may define $\rho(u) = \{v_i : u \in X_i\}$.)

Let the *width* of such a decomposition be $\max_i |\{u : v_i \in \rho(u)\}| - 1$. The *path-width* of G is the minimum k such that there exists a path-decomposition of G of width k.

For convenience and intuition, we will say $v \in V(P)$ *lies under* $u \in V(G)$ if $v \in \rho(u)$, as drawn in Figure 2.

Obviously, path-width is a graph invariant, as it is preserved under isomorphism. For example, one can show that a graph has path-width 0 if and only if it has no edges and that a graph has path-width 1 if and only if each of its connected components is a caterpillar (i.e., a graph obtained from a path by appending edges – "legs" – to the vertices of the path; see Figure 2). We shall return to the caterpillar example later.

For further elaboration of this definition, we give the proofs of three statements stated without proof in [1], restated under our definition of path-decomposition.

Lemma 1 ([1], 1.3-1.6)**.** *The following hold for a graph G:*

(a) If every connected component of G has path-width $\leq k$, then G has path-width $\leq k$.

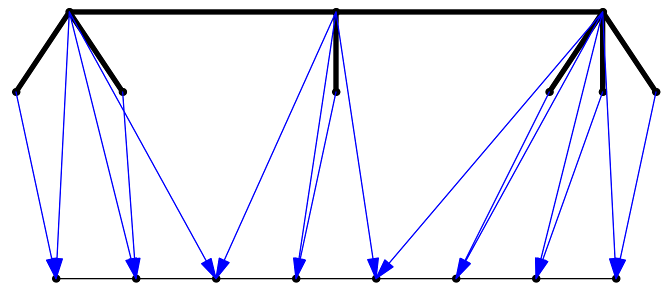

Figure 2: Path-decomposition of a caterpillar. The thick black segments are edges of the graph G, the thin black segments are edges of the path P, and blue arrows have been drawn from each vertex $u \in G$ to each vertex in $\rho(u)$.

(b) If $X \subset V(G)$ and $G \backslash X$ has path-width $\leq k$, then G has path-width $\leq k+|X|$.

(c) If $\rho : V(G) \to \mathcal{P}(V(P))$ is a path-decomposition of G, where $V(P) = \{v_1, \ldots, v_r\}$, for $1 \leq i \leq r-1$, $|\{u : \{v_i, v_{i+1}\} \subseteq \rho(u)\}| \leq k$, and for $1 \leq i \leq r$, the subgraph of G induced by $\{v \in V(G) : \rho(v) = \{i\}\}$ has path-width $\leq k'$, then G has path-width $\leq k' + 2k$.

Proof. (a) Let G_1 and G_2 be disjoint graphs. Let P_1 and P_2 be paths and $\rho_1 : V(G_1) \to \mathcal{P}(V(P_1))$ and $\rho_2 : V(G_2) \to \mathcal{P}(V(P_2))$ path-decompositions of widths k_1 and k_2, respectively. Let P be the path formed by concatenating the paths P_1 and P_2, and define $\rho : V(G_1 \cup G_2) \to V(P)$ such that $\rho|_{V(G_1)} = \rho_1$ and $\rho|_{V(G_2)} = \rho_2$. Then ρ is a path-decomposition of $G_1 \cup G_2$ of width $\max(k_1, k_2)$, because, for $v \in P$, $|\{u : v \in \rho(u)\}| = |\{u : v \in \rho_1(u)\}|$ if $v \in P_1$ and $|\{u : v \in \rho(u)\}| = |\{u : v \in \rho_2(u)\}|$ if $v \in P_2$.

This shows that the path-width of the disjoint union of two graphs is equal to the maximum of their path-widths.

(b) Let G be a graph and $X \subseteq V(G)$. Let P be a path and $\rho : V(G \setminus X) \to \mathcal{P}(V(P))$ a path-decomposition of width k. Extend ρ to a function $\rho' : V(G) \to \mathcal{P}(V(P))$ by setting $\rho'(u) = V(P)$ for $u \in X$. Then ρ' is a path-decomposition of G of width $k + |X|$, because, for $v \in P$, $\{u : v \in \rho'(u)\} = \{u : v \in \rho(u)\} \cup X$.

(c) Let $\rho : V(G) \to \mathcal{P}(V(P))$ be a path-decomposition of a graph G, where $V(P) = \{v_1, \ldots, v_r\}$, $\left|\rho^{-1}(v_i) \cap \rho^{-1}(v_{i+1})\right| \leq k$ for $1 \leq i \leq r-1$ and the subgraph of G induced by $\{v : \rho(v) = \{i\}\}$ has path-width at most k' for $1 \leq k \leq r$.

Let $P_1, \ldots, P_r$ be paths and for each i let $\rho_i : \{v : \rho(v) = \{i\}\} \to \mathcal{P}(V(P_i))$ be a path-decomposition of the subgraph of G induced by $\{v : \rho(v) = \{i\}\}$ of width at most k'.

Let Q be the path formed by concatenating the paths $P_1, \ldots, P_n$, and define $\rho' : V(G) \to \mathcal{P}(V(Q))$ by setting

$$\rho'(v) = \begin{cases} \rho_i(v) & \text{if } \rho(v) = \{i\} \text{ for some } i, \\ \bigcup_{i : v_i \in \rho(v)} V(P_i) & \text{else} \end{cases}.$$

Then ρ' is a path-decomposition of G of width at most $k' + 2k$, since if $v \in P_i$, then

$$\begin{aligned} |\{u : v \in \rho'(u)\}| &= \left|\rho_i^{-1}(v) \cup \{u : v_i \in \rho(u), |\rho(u)| > 1\}\right| \\ &\leq \left|\rho_i^{-1}(v)\right| + |\{u : \{v_i, v_{i-1}\} \subseteq \rho(u)\} \cup \{u : \{v_i, v_{i+1}\} \subseteq \rho(u)\}| \\ &\leq k' + 2k. \end{aligned}$$

(Note also that (a) follows from (c) by taking the path $P = \{v_1, v_2\}$ and letting $\rho^{-1}(v_1) = V(G_1)$ and $\rho^{-1}(v_2) = V(G_2)$.)

□

We also note the following.

Lemma 2. *If G is a minor of G', then the path-width of G' is not less than the path-width of G.*

Proof. It suffices, given a path-decomposition of G', to produce a path-decomposition of G of the same or lesser width.

Let $\pi : G' \to G$ be the function as in the definition of minor, and let $\rho' : G' \to \mathcal{P}(V(P))$ be a path-decomposition. We define $\rho : G \to \mathcal{P}(V(P))$ by $\rho(u) = \bigcup_{w \in \pi^{-1}(u)} \rho(w)$. It is clear that ρ is a path-decomposition of G.

Because $\left|\rho^{-1}(v)\right| = \left|\pi((\rho')^{-1}(v))\right| \leq \left|(\rho')^{-1}(v)\right|$, the width of ρ does not exceed the width of ρ', as deisred. □

A *tree-decomposition* of a graph G is defined similarly to path-decomposition, except that the path P is replaced by a tree T. Correspondingly, the *tree-width* of a graph G is the minimum k such that there exists a tree-decomposition of G of width k.

Lemma 2 holds when "path-width" is replaced with "tree-width".

2.2 Excluding a forest

An important step in the proof of Theorem 1 is the following:

Theorem 2. *Suppose F is a forest. There exists some w such that if a graph G has path-width greater than w, then G has a minor isomorphic to F.*

In this section we summarize [1]'s proof of Theorem 2.

We define several families of graphs indexed by a positive integer parameter.

The grid $\mathfrak{G}_\theta$, where $\theta > 1$, is a graph with θ^2 vertices indexed by pairs (i, j), where $1 \leq i, j \leq \theta$, and (i, j) is adjacent to $(i \pm 1, j)$ and $(i, j \pm 1)$, whenever the indices make sense.

The tree $\mathfrak{Y}_\lambda$, where $\lambda \geq 1$, is defined inductively by starting with the "Y" tree $K_{1,3}$ and performing the operation of appending two new vertices to each leaf $\lambda - 1$ times.

The tree $\mathfrak{H}_\lambda$ is obtained in the same way, but one begins with $K_{1,2}$ in place of $K_{1,3}$. Also, $\mathfrak{H}_0$ is defined to be the graph with one vertex.

The tree $\mathfrak{P}_{\gamma,\delta}$, where $\gamma, \delta \geq 1$, is formed by taking δ copies of $\mathfrak{H}_{\gamma-1}$, appending a new vertex to the "central" vertex of each copy, and joining these δ new vertices by a path. When $\mathfrak{P}_{\gamma,\delta}$ is considered as a rooted tree, its root is at one of the ends of the path.

All of these trees are illustrated in Figure 3.

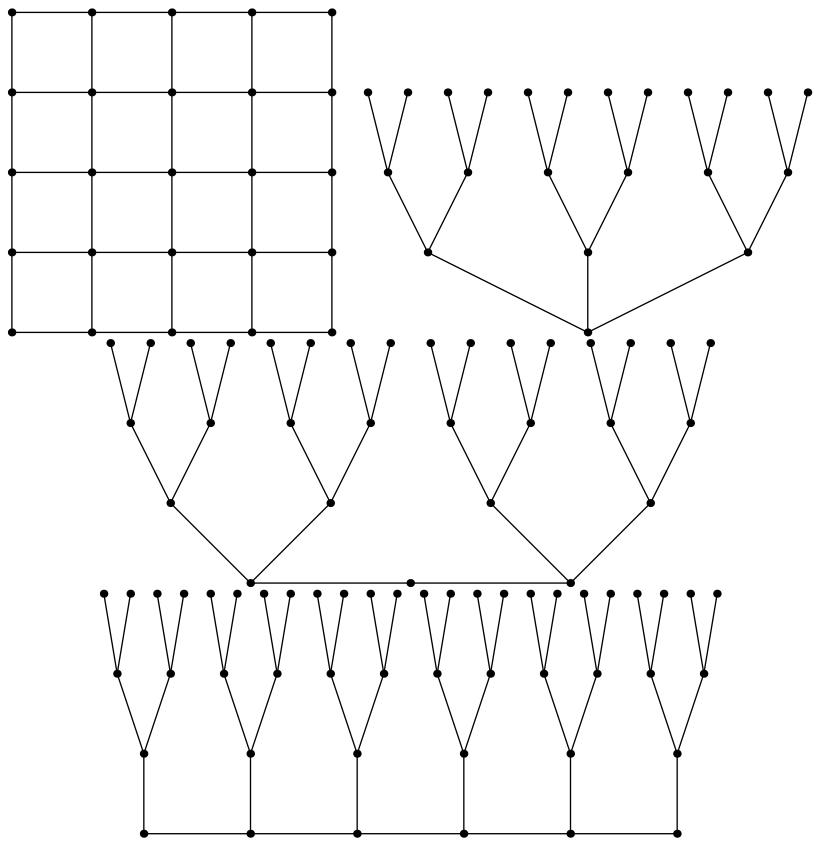

Figure 3: The graphs $\mathfrak{G}_5$, $\mathfrak{Y}_3$, $\mathfrak{H}_4$, and $\mathfrak{P}_{3,6}$.

These graphs have properties related to graph minors that will be important. The following facts are stated but not fully proven in [1]. We fill in the details.

Lemma 3 ([1],3.2-3.5)**.** *The graph $\mathfrak{Y}_\lambda$ is isomorphic to a minor of $\mathfrak{G}_\theta$, where*

$$\theta = 2^{\lceil \frac{\lambda+3}{2} \rceil}.$$

Therefore, if a graph G has no minor isomorphic to $\mathfrak{Y}_\lambda$, then G has no minor isomorphic to $\mathfrak{G}_\theta$.

Proof. First, clearly $\mathfrak{H}_{\lambda-1}$ is isomorphic to a minor of $\mathfrak{H}_\lambda$. Notice that $\mathfrak{Y}_\lambda$ is formed by joining the central vertices of $\mathfrak{H}_\lambda$ and $\mathfrak{H}_{\lambda-1}$ by an edge, while $\mathfrak{H}_{\lambda+1}$ is formed by joining the central vertices of two copies of $\mathfrak{H}_\lambda$ by two edges in series. Therefore, $\mathfrak{Y}_\lambda$ is isomorphic to a minor of $\mathfrak{H}_{\lambda+1}$.

Next, we show by induction on λ that $\mathfrak{H}_{\lambda+1}$ can be embedded in $\mathfrak{G}_{\theta-1}$, where θ is as in the statement of the lemma. It suffices to consider λ odd and to require that $\mathfrak{H}_\lambda$ correspond to the center of the grid and that no vertices of $\mathfrak{H}_{\lambda+1}$ correspond to any other vertices in the center of the grid. For $\lambda = 1$, such an embedding of $\mathfrak{H}_2$ in $\mathfrak{G}_3$ is trivial to find. Figure 4 illustrates the construction of an embedding of $\mathfrak{H}_{\lambda+2}$ in $\mathfrak{G}_{2\theta-1}$ given an embedding of $\mathfrak{H}_{\lambda+1}$ in $\mathfrak{G}_{\theta-1}$.

From this we conclude that $\mathfrak{Y}_\lambda$ is isomorphic to a minor of $\mathfrak{G}_{\theta-1}$, which is isomorphic to a minor of $\mathfrak{G}_\theta$, as desired. □

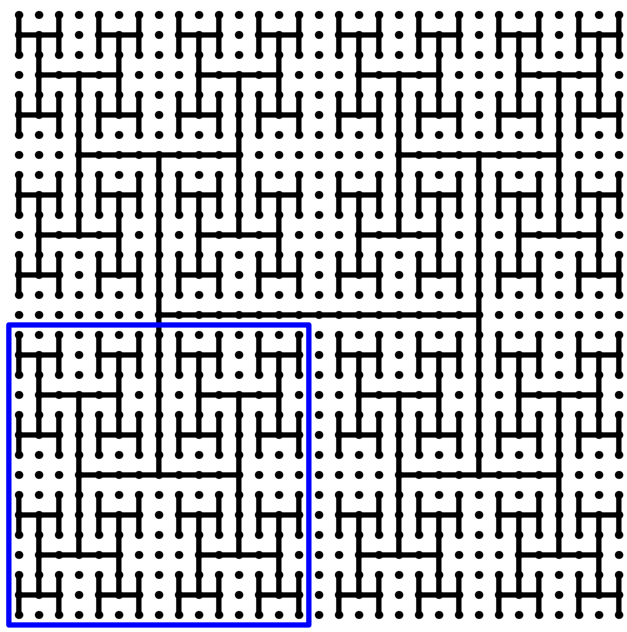

Figure 4: Inductive step in the proof of Lemma 3: embedding of $\mathfrak{H}_8$ in $\mathfrak{G}_{31}$ and (bordered) embedding of $\mathfrak{H}_6$ in $\mathfrak{G}_{15}$.

Lemma 4. *Every forest F is isomorphic to a minor of $\mathfrak{Y}_\lambda$ for some λ.*

Proof. Every forest is isomorphic to a minor of a tree because one may add edges connecting components of the forest to obtain a tree that can be transformed into the original forest by deleting the new edges.

Every tree is isomorphic to a minor of a binary tree (i.e., a tree with no vertices of degree exceeding 3) because one can repeatedly split vertices of degree higher than 3 into two vertices of lesser degree connected by an edge until one obtains a binary tree that can be transformed into the original tree by contractions of the new edges.

Finally, the graphs $\mathfrak{Y}_\lambda$ contain the full binary trees of arbitrary height $\mathfrak{H}_\lambda$. □

The critical step in the proof of Theorem 2 is motivated by some results about graphs that do not contain $\mathfrak{G}_\theta$ as a minor. The main result is the following. (Note that if $G = (V, E)$ is a graph and $A, B, W \subset V$, then W *separates* A and B for all $a \in A$ and $b \in B$, every path from a to b contains a vertex of V.)

Lemma 5 ([1],2.7)**.** *Fix θ sufficiently large ($\theta > 6$ suffices). Suppose $G = (V, E)$ is a graph and $A_1, A_2 \subseteq V$. If there exist $\frac{1}{2}\theta^2$ disjoint paths from A_1 to A_2 and there exist $\theta^{2\theta}$ disjoint connected subgraphs $B_1, \ldots, B_{\theta^{2\theta}}$ each separating A_1 and A_2, then G has a minor isomorphic to $\mathfrak{G}_\theta$.*

Sketch of proof. Step 1: If there exist disjoint paths $P_1, \ldots, P_\theta$ in G each intersecting the $\frac{1}{2}\theta \cdot \theta!$ disjoint paths $Q_1, \ldots, Q_{\frac{1}{2}\theta\cdot\theta!}$ at one vertex, in that order, then G has a minor isomorphic to $\mathfrak{G}_\theta$. The proof is a pigeonhole argument: there are $\frac{1}{2}\theta!$ orders (up to reversal) in which the P_i could intersect each Q_j, so there are $Q_{j_1}, \ldots, Q_{j_\theta}$ so that the P_i intersect the Q_{j_i} in the same order. It is trivial to extract a $\mathfrak{G}_\theta$-minor.

Step 2: If there exist disjoint paths $P_1, \ldots, P_\theta$ in G each intersecting the $\theta(\theta - 1)$ disjoint trees $R_1, \ldots, R_{\theta(\theta-1)}$ at one leaf, in that order, then G has a minor isomorphic to $\mathfrak{G}_\theta$. By performing contractions, we may reduce to the case where each R_j is a star. Figure 5 shows how to then extract a $\mathfrak{G}_\theta$-minor.

Step 3: If there exist $\frac{1}{2}\theta^2$ disjoint paths $P_1, \ldots, P_{\frac{1}{2}\theta^2}$ in G each intersecting the $\theta^{2\theta-2}$ disjoint subgraphs $B_1, \ldots, B_{\theta^{2\theta-2}}$ in that order, then G has a minor isomorphic to $\mathfrak{G}_\theta$. This can be deduced from the previous two steps, since any connected graph with at least $\frac{1}{2}\theta^2$ vertices has a minor isomorphic to a path with θ vertices or to a star with θ leaves.

Step 4: If $A_1, A_2 \subseteq V$, there is a unique set of disjoint paths $P_1, \ldots, P_m$ from A_1 to A_2 such that every vertex of G lies on one of the P_i, and there exist $2m\theta^{2\theta-2}\binom{m}{\frac{1}{2}\theta^2}$ disjoint connected subgraphs each intersecting at least $\frac{1}{2}\theta^2$ of the P_i, then G has a minor isomorphic to $\mathfrak{G}_\theta$. To prove this, one constructs a labeling of each vertex $v \in V$ with an integer $\mu(v)$ such that the labels are strictly increasing along each path P_i and for each n, $\{v : \mu(v) < n\}$ is separated from $\{v : \mu(v) \geq n\}$ by $\bigcup_i \{$the first vertex v of P_i with $\mu(v) \geq n\}$.

The result follows easily from Step 4. □

Now we can finally relate containment of specific minors to path-width:

Theorem 3 ([1],3.6(i))**.** *Fix $\lambda > 2$, $\gamma \geq 0$, and $\delta \geq 2$, and let $\theta = 2^{\lceil \frac{\lambda+3}{2} \rceil}$. If*

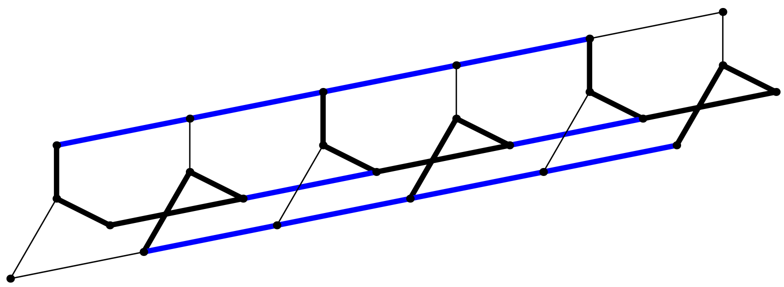

Figure 5: $\theta = 3$. Three paths intersect $\theta(\theta-1) = 6$ stars. The $\mathfrak{G}_3$ minor is highlighted in thick black and blue.

the path-width of a rooted connected graph G is greater than

$$w = \underbrace{\left(2^{\theta^6}\right)^{\left(2^{\theta^6}\right)^{\cdots \left(2^{\theta^6}\right)^{\delta}}}}_{\gamma},$$

then G has a rooted minor isomorphic to $\mathfrak{P}_{\gamma,\delta}$ or a minor isomorphic to $\mathfrak{G}_\theta$.

The proof is a long technical application of the lemmas about path-width and Lemma 5 above. Another family of trees $\mathfrak{Q}_{\gamma,\delta}$, obtained by appending two new vertices to each leaf in one of the copies of $\mathfrak{H}_{\gamma-2}$ in $\mathfrak{P}_{\gamma-1,\delta}$, is defined. Bounds on the path-width of rooted graphs having no minor isomorphic to $\mathfrak{G}_\theta$ and no rooted minor isomorphic to $\mathfrak{P}_{\gamma,\delta}$ or $\mathfrak{Q}_{\gamma,\delta}$ are established by a joint induction.

Proof of Theorem 2. By Lemma 4, choose λ such that F is isomorphic to a minor of $\mathfrak{Y}_\lambda$. Let $\theta = 2^{\lceil \frac{\lambda+3}{2} \rceil}$. By the previous theorem with $\gamma = \lambda$ and $\delta = 3$, there exists w such that if the path-width of G (therefore, of every connected component of G) is greater than w, then every connected component of G has a minor isomorphic to $\mathfrak{P}_{\lambda,3}$ or to $\mathfrak{G}_\theta$.

But $\mathfrak{P}_{\lambda,3}$ clearly contains a minor isomorphic to $\mathfrak{Y}_\lambda$, and G_θ contains such a minor by Lemma 3. Therefore, G contains a minor isomorpic to $\mathfrak{Y}_\lambda$, so G contains a minor isomorphic to F. $\square$

To summarize, any forest is isomorphic to a minor of some $\mathfrak{Y}_\lambda$ and therefore of some $\mathfrak{G}_\theta$; we introduced the graphs $\mathfrak{P}_{\gamma,\delta}$ to show that any graph not having a minor isomorphic to $\mathfrak{G}_\theta$ or $\mathfrak{P}_{\gamma,\delta}$ has bounded path-width, from which it followed that any graph not having a minor isomorphic to $\mathfrak{Y}_\lambda$ has bounded path-width.

As a corollary to Theorem 2, for any set $\mathcal{F}$ of graphs of bounded path-width there exist a forest F such that no $G \in \mathcal{F}$ has a minor isomorphic to F. For example, suppose $\mathcal{F}$ is the set of caterpillars, which, as mentioned above, have path-width bounded by 1. Indeed, no caterpillar has a minor isomoprhic to the graph obtained from $K_{1,3} = \mathfrak{Y}_1$ by replacing each edge by two edges in series.

(The converse is not true: not every graph not containing this graph as a minor has path-width 0 or 1. The graphs K_3 or even K_4 provides counterexamples with path-width 2 and 3, respectively. On the other hand, this graph is contained in $\mathfrak{Y}_2$, so $\lambda = 2$, $\theta = 8$, and any graph with no minor isomorphic to it has path-width at most

$$\left(2^{8^6}\right)^{\left(2^{8^6}\right)^3}.$$

This is less desirable – but not optimal.)

2.3 Planar graphs and sleeve unions

Theorem 4 ([3])**.** *Suppose F is a planar graph. There exists some w such that if a planar graph G has tree-width greater than w, then G has a minor isomorphic to F.*

This result is analogous to Theorem 2 in that "tree" has been replaced with "planar graph" and "path-width" with "tree-width".

To prove this theorem, we define another family of graphs $\mathfrak{C}_{r,s}$, the cylinders. The graph $\mathfrak{C}_{r,s}$ is the product of the cycle with r vertices and the path with s vertices. The s copies of the cycle with r vertices in $\mathfrak{C}_{r,s}$ are labeled $C_1, \ldots, C_s$ in the obvious order, where C_s is the "inner" cycle. (See Figure 6.) We also let $\mathfrak{N}_{\lambda,s}$ be the graph formed by identifying the leaves of $\mathfrak{Y}_\lambda$ with the vertices on the inner cycle of $\mathfrak{C}_{3\cdot 2^{\lambda-1},s}$. (Notice that $\mathfrak{Y}_\lambda$ indeed has $3 \cdot 2^{\lambda-1}$ leaves.)

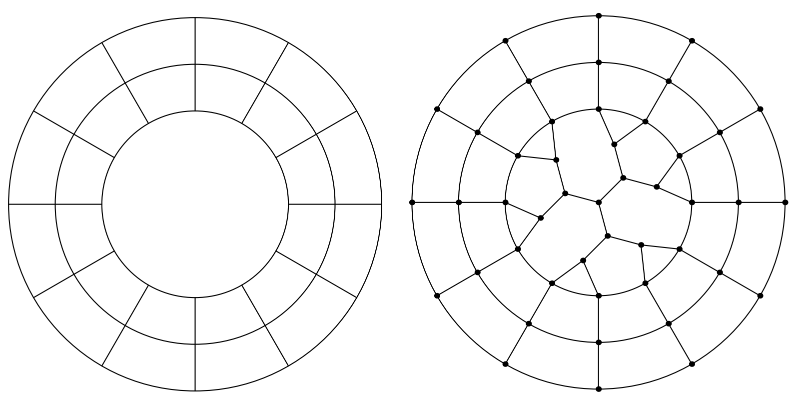

Figure 6: The graphs $\mathfrak{C}_{12,3}$ and $\mathfrak{N}_{3,3}$.

We need several simple results about tree-width as it relates to cylinders. They will be used to define an operation on graphs that, in some sense, respects tree-width.

Lemma 6 ([3],3.8)**.** *Let H be a graph, and let G be formed from H by identifying some r vertices $v_1, \ldots, v_r$ of H with the r vertices of the inner cycle C_s of $\mathfrak{C}_{r,s}$.*

Suppose G has tree-width w. Then H has a tree-decomposition (T, ρ) of width not exceeding w such that $\bigcap_{i=1}^{r} \rho(v_i)$ is nonempty.

That is, given a tree-decomposition of H with a cylinder glued to it, we may produce a tree-decomposition of H of lesser width such that some vertex of the tree lies under all the vertices of H that were glued to the cylinder.

The proof of Lemma 6 is technical and we omit it. The main intermediate step concerns tree-decompositions of cylinders: if (T, ρ) is a tree-decomposition of $\mathfrak{C}_{r,s}$, then there exists some $t \in T$ such that there are r vertex-disjoint paths from $\{u : t \in \rho(u)\}$ to C_s. This is used to produce a tree-decomposition of H given a tree-decomposition of G.

This motivates the notion of the sleeve union of graphs. Let H_1 and H_2 be graphs that each contain $\mathfrak{C}_{r,s}$ as a minor. Let G_1 be formed from H_1 by identifying some r vertices $u_1, \ldots, u_r$ of H_1 with the inner cycle of a copy of $\mathfrak{C}_{r,s}$ to H_1. Let G_2 similarly be formed from H_2 by attaching a copy of $\mathfrak{C}_{r,s}$ along the vertices $v_1, \ldots, v_r$. Let G be the graph formed from H_1 and H_2 by identifying u_i with v_i. Then G is the *sleeve union* of G_1 and G_2.

As illustrated in Figure 7, one can imagine G_1 and G_2 sliding into each other as hands into sleeves. The graphs G_1 and G_2 are both minors of G: G_1 is formed by replacing H_1 by its minor $\mathfrak{C}_{r,s}$ in G, and G_2 is similarly formed by deletions and contractions in the H_2 portion of G.

This operation behaves well with tree-width.

Lemma 7 ([3],3.10)**.** *If G is the sleeve union of G_1 and G_2, then the tree-width of G is the maximum of the tree-widths of G_1 and G_2.*

Proof. Let w be this maximum.

Let H_1 and H_2 be subgraphs of G as in the definition of sleeve union. Because G_1 and G_2 both have tree-width w, by Lemma 6, one can find tree-decompositions (T_1, ρ_1) of H_1 and (T_2, ρ_2) of H_2, both of width at most w, such that there are $t_1 \in T_1$ and $t_2 \in T_2$ with $t_1 \in \rho(u_i)$ and $t_2 \in \rho(v_i)$ for all i.

Let T be the tree formed from the disjoint union of T_1 and T_2 by adding an edge from t_1 to t_2, and define a tree-decomposition (T, ρ) of G by

$$\rho(v) = \begin{cases} \rho_1(v) & v \in H_1 \setminus H_1, \\ \rho_2(v) & v \in H_2 \setminus H_1, \\ \rho_1(v) \cup \rho_2(v) & v \in H_1 \cap H_2 \end{cases}.$$

Clearly this is a tree-decomposition of G of width w.

Conversely, the tree-width of G is at least w because G_1 and G_2 are minors of G, by Lemma 2. □

So far, we have not used planarity of any of the graphs involved. We now consider planar graphs together with their embeddings in the plane.

Suppose G is a connected planar graph together with an embedding M in the plane. Let G' be its planar dual, with the dual embedding M' in the plane. The *radius* of M is the maximum distance from the vertex of G' corresponding

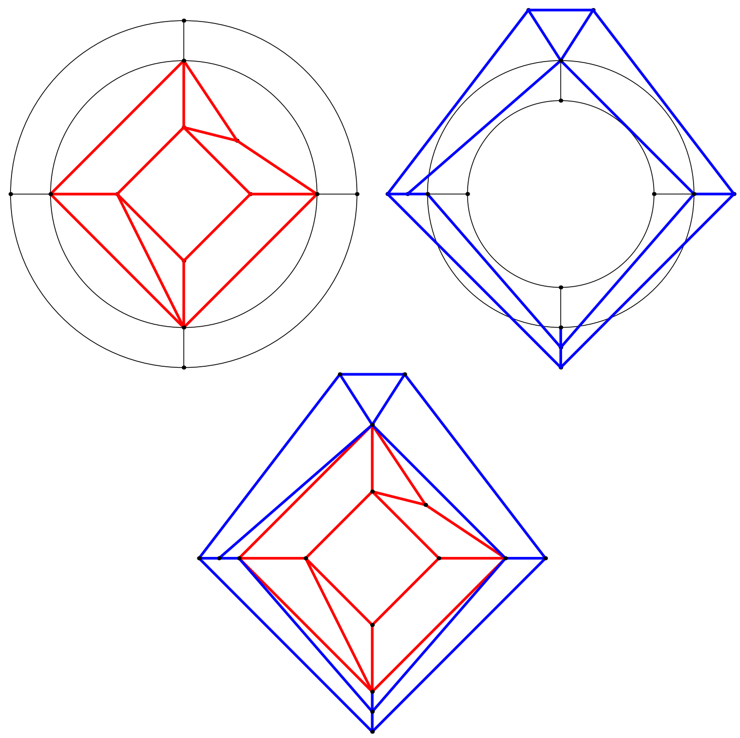

Figure 7: Sleeve union of two graphs, each containing $\mathfrak{C}_{4,2}$ as a minor.

to the infinite face of G to any other vertex of G'. The *radius* of a graph G is the minimum radius of an embedding of G. For example, the embeddings of $\mathfrak{C}_{12,3}$ and $\mathfrak{N}_{3,3}$ shown above in Figure 6 both have radius 3; in fact, 3 is the radius of each of these graphs, since they can be shown to have no embedding of radius 1 or 2.

Now we shall use the graphs $\mathfrak{N}_{\lambda,s}$ to relate radius to tree-width.

One can show the following:

Lemma 8 ([3],2.4)**.** *If G is a planar graph of radius s, then G is isomorphic to a minor of $\mathfrak{N}_{\lambda,s}$ for some λ.*

The proof proceeds as follows: we find a sequence of s circuits, each one contained in the region bounded by the previous one, in an embedding of G; removing these circuits yields a graph with radius 0, that is, a forest. Finally, observe that any forest is isomorphic to a minor of some $\mathfrak{Y}_\lambda$ (as shown in the previous section) and that if $\lambda < \lambda'$, then $\mathfrak{N}_{\lambda,s}$ is isomorphic to a minor of $\mathfrak{N}_{\lambda',s}$.

Lemma 9 ([3],2.5)**.** *The tree-width of $\mathfrak{N}_{\lambda,s}$ does not exceed $3s+1$.*

This is shown by explicit construction of a tree-decomposition of $\mathfrak{N}_{\lambda,s}$, where the tree is isomorphic to $\mathfrak{Y}_{\lambda-1}$. A "proof by colourful picture" is shown in Figure 8.

Figure 8: Tree-decomposition of $\mathfrak{N}_{3,3}$. The tree $\mathfrak{Y}_2$ is thought of as the copy of $\mathfrak{Y}_3$ at the center with its leaves removed, shown in thick black. Each vertex v is coloured with $\{u : v \in \rho(u)\}$.

Corollary 1. *A planar graph of radius s has tree-width not exceeding $3s + 1$.*

Proof. From the previous two lemmas. □

Finally, we relate sleeve unions and cylinders to planarity and radius.

An embedding of a graph G in the plane is said to *major* $\mathfrak{C}_{r,s}$ if there are circuits $C_1, \ldots, C_s$ in G, where each C_{i+1} is contained in the region bounded by C_i, and paths $P_1, \ldots, P_r$ from C_1 to C_s such that the intersection of P_i with each C_j is connected. (That is, G has a "topological minor" isomorphic to $\mathfrak{C}_{r,s}$: there are nested circuits in G corresponding to the nested circuits of $\mathfrak{C}_{r,s}$.) In particular, if G majors $\mathfrak{C}_{r,s}$, then G has radius at least s.

One can show the following fact about decompositions of graphs into sleeve unions:

Lemma 10 ([3],4.3)**.** *Suppose G can not be expressed as the sleeve union of two graphs.*

(a) If G majors $\mathfrak{C}_{r,s}$, where r is even and $r < s$, then G majors $\mathfrak{C}_{r+1,s-r}$.

(b) If G majors $\mathfrak{C}_{r,s}$, where r is odd and $r < s - 1$, then G majors $\mathfrak{C}_{r+1,s-r-1}$.

From this it follows by induction that if G cannot be expressed as the sleeve union of two graphs and does not major $\mathfrak{C}_{r,s}$, then G has radius at most $s + \left\lfloor \frac{1}{2}(r^2 - 1) \right\rfloor$ and, by Corollary 1, tree-width at most $3\left(s + \left\lfloor \frac{1}{2}(r^2 - 1) \right\rfloor\right) - 2$.

This decomposition theorem, combined with Lemma 7, allows us to prove the theorem.

Proof of Theorem 4. Let H be a planar graph. There exists r such that $\mathfrak{C}_{r,r}$ has a minor isomorphic to H. Let $w = 3\left(r + \left\lfloor \frac{1}{2}(r^2 - 1) \right\rfloor\right) - 2$.

Suppose G is a planar graph having tree-width greater than w.

If G can be expressed as a sleeve union of two of its minors G_1 and G_2, then by Lemma 7 at least one of G_1 and G_2 (assume G_1) has tree-width greater than w, and we reduce to a smaller case by replacing G by G_1.

So suppose G cannot be expressed as a sleeve union of two of its minors. But, by the previous lemma, G majors $\mathfrak{C}_{r,r}$, so G has a minor isomorphic to H, as desired. □

2.4 Planar graphs and grids

Theorem 5 ([5])**.** *Suppose F is a planar graph. There exists some w such that if a graph G has tree-width greater than w, then G has a minor isomorphic to F.*

This result is analogous to Theorem 2 in that "tree" has been replaced with "graph" and "path-width" with "tree-width". It strengthens Theorem 4.

The method of proof will be similar to that of Theorem 2. The following is analogous to Lemma 3.

Lemma 11 ([5],2.1)**.** *For all sufficiently large θ, there exists w such that if a graph G has tree-width greater than w, then G has a minor isomorphic to $\mathfrak{G}_\theta$.*

The constant w is even more impractical than in Lemma 3, as will be seen in the proof sketch below.

The idea of the proof is to show that if G has no minor isomorphic to $\mathfrak{G}_\theta$, then G does not contain certain structures that involve two pairwise intersecting families of subgraphs.

Namely, an (m, n)-*web* in G is a set of paths $P_1, \ldots, P_m$, $Q_1, \ldots, Q_n$ in G such that the P_i are disjoint, the Q_i are disjoint, and each P_i intersects each Q_j. We also require that the P_i have no edges in common with the Q_j.

An (m, n)-*mesh* in G is defined in the same way, but with "paths" replaced by "connected subgraphs" and permitting that the P_i and Q_j not be edge-disjoint.

Sketch of proof of Lemma 11. We first define a (huge) constant w_1 depending on θ. Let α be a function recursively defined by $\alpha(2, n) = n + 1$ and $\alpha(k, n) = 2^{n\theta^4} + \alpha\left(k - 1, 2^{n\theta^4} + n + 1\right)$ for $n \geq 3$. Let $\theta' = 2\alpha\left(\frac{1}{2}\theta^2, \frac{1}{2}\theta^2\right)$. Let $w_1 = \phi_0 + 2\phi_1 + \cdots + 2\phi_{\theta'-1} + \phi_{\theta'}$, where $\phi_{\theta'} = \frac{1}{2}\theta^2$ and ϕ_i is defined recursively from $\phi_{\theta'}$ by $\phi_i = \phi_{i+1} 2^{\phi_{i+1}\theta^2}$.

The motivation for this bizarre constant is that one can show, using results about grids mentioned two sections previously, that if a graph G contains a (w_1, w_1)-web, then G contains a minor isomorphic to $\mathfrak{G}_\theta$.

Now, let $w_2' = \left(\frac{1}{2}\theta^2\right)^{w_1 - 1}$ and $w_3' = w_1\binom{w_2'}{w_1} + \frac{1}{2}\theta^2\binom{w_2'}{\frac{1}{2}\theta^2}$; let $w_2 = \left(\frac{1}{2}\theta^2\right)^{w_3' - 1}$ and $w_3 = w_2\binom{w_2}{w_3'} + \frac{1}{2}\theta^2\binom{w_2}{\frac{1}{2}\theta^2}$. An involved pigeonhole argument shows that if a graph G contains a (w_2, w_3)-mesh, then G contains a (w_2', w_3')-mesh where each of the connected subgraphs is a path, and that then G contains a (w_1, w_1)-web and therefore a minor isomorphic to $\mathfrak{G}_\theta$.

Finally, choose $w > \alpha(w_2, w_3)(1 + \frac{3}{4}(w_2(3^{w_2} - 1)))$. One can show by an argument similar to the proof of Theorem 7 that if a graph G has tree-width greater than w, then G has a minor isomorphic to $\mathfrak{G}_\theta$. The reduction operation in this case is not decomposition of G as a sleeve union, but writing G as the union of two subgraphs with relatively "few" vertices in common and of roughly equal size. One can show that if G cannot be written as such a union, then G has a minor isomorphic to $\mathfrak{G}_\theta$ by showing it has a (w_2, w_3)-mesh.

Specifically, "few" means fewer than $\alpha(w_2, w_3)$ and "roughly equal" means each subgraph has no more than $1 - \frac{1}{\frac{3}{4}(w_2(3^{w_2}-1))}$ of the number of vertices in G – this explains this factors in the definition of w. □

This lemma implies the theorem just as Lemma 3 implied Theorem 2.

Proof of theorem 2. Suppose F is a planar graph. Let θ be sufficiently large ($\theta \geq 6$) such that F is isomorphic to a minor of $\mathfrak{G}_\theta$. Choose w as in the previous lemma. Then, if a graph G has tree-width greater than w, then G has a minor isomorphic to $\mathfrak{G}_\theta$, and hence has a minor isomorphic to F. □

Let us see how impractical this is. For example, the trees are the connected graphs not having the planar graph K_3 as a minor. The results above required $\theta \geq 6$, and $\mathfrak{G}_6$ contains K_3 as a minor. So, let $\theta = 6$. But even the constant $\theta' = 2\alpha(18, 18)$ defined in the proof of Lemma 11 would not fit on this page! On the other hand, trees are precisely the connected graphs with tree-width not exceeding the (considerably smaller) constant 1.

2.5 Families with bounded tree-width

Theorem 1 is implied by Theorem 5 and the following:

Theorem 6. *Fix an integer k. The relation $\preceq$ is a well-quasi order on the set of graphs with tree-width not greater than k.*

The previous two theorems imply the main result.

Proof of Theorem 1. Suppose $\mathcal{H}$ is a set of finite graphs closed under taking minors and $\mathcal{H}$ does not contain some planar graph F. By Theorem 5, $\mathcal{H}$ has bounded tree-width. By Theorem 6, $\preceq$ restricts to a well-quasi order on $\mathcal{H}$, so $\mathcal{H}$ has finitely many minimal elements. □

References

[1] N. Robertson and P.D. Seymour, Graph Minors. I. Excluding a Forest, Journal of Combinatorial Theory 35 (1983).

[2] N. Robertson and P.D. Seymour, Graph Minors. II. Algorithmic Aspects of Tree-Width, Journal of Algorithms 7 (1986).

[3] N. Robertson and P.D. Seymour, Graph Minors. III. Planar Tree-Width, Journal of Combinatorial Theory 36 (1984).

[4] N. Robertson and P.D. Seymour, Graph Minors. IV. Tree-Width and Well Quasi-Ordering, Journal of Combinatorial Theory 48 (1990).

[5] N. Robertson and P.D. Seymour, Graph Minors. V. Excluding a Planar Graph, Journal of Combinatorial Theory 41 (1986).