Wagner's Conjecture*

TEH Yee Whye
009753429

January 19, 2000

Abstract

The study of graph minors is an important area of topological graph theory and graph embeddings. Central to the study of graph minors is Wagner's conjecture, which was proven by Robertson and Seymour [15]. We survey the ideas and results leading up to the proof of Wagner's conjecture and their implications.

Contents

*CSC 2410S Graph theory course project

1 Introduction

A fundamental result by Kuratowski [11] states that a graph G is planar if and only if G has no subgraphs isomorphic to any subdivision of K_5 or $K_{3,3}$. Wagner [22] showed an equivalent statement of Kuratowski's theorem that G is planar if and only if G has no K_5 nor $K_{3,3}$-minor.

König [9] then asked if a Kuratowski type characterization exists for graphs embeddable in any surface(*) S. Little progress was made until Glover et al [6] showed that there are 103 topologically irreducible graphs not embeddable in the projective plane and Archdeacon [1] showed that there are 35 minor-minimal graphs not embeddable in the projective plane. The lists of such excluded graphs for all other surfaces are not known, and evidence indicates that they might contain too many graphs to be of any practical use [20]. Archdeacon and Huneke [2] then showed that for any non-orientable surface S there exists a *finite* family of graphs F_S such that a graph G is embeddable in S if and only if G has no minor isomorphic to a member of F_S. The members of F_S are called the *excluded minors* of embeddability in S. In the same year Bodendiek and Wagner [4] showed a similar result for the orientable surfaces.

Robertson and Seymour took a less direct but more fruitful approach to the problem. In a long series of papers they proved Wagner's conjecture [15] :

Theorem 1.1 (Wagner's conjecture). *If $G_1, G_2, G_3, \ldots$ is a sequence of graphs there exists $1 \leq i < j$ such that G_i is isomorphic to a minor of G_j.*

Now consider the family F of graphs not embeddable in some given surface S. Let $F_S \subset F$ be those members in F that are minor-minimal with respect to not embeddable in S. Then every graph not embeddable in S must have a minor isomorphic to a member of F_S. By virtual of minor-minimality every graph in F_S is not a minor of any other graph in F_S. By Wagner's conjecture F_S cannot be infinite, else there would exist a sequence of graphs from F_S any two of which are incomparable.

In this paper we shall survey results related to Wagner's conjecture. In particular in section 2 we describe the preliminary definitions and background ideas. In section 3 well-quasi-orderings are introduced and Wagner's conjecture is restated as saying that the family of graphs is well-quasi-orders. Then a progression of increasingly complex objects are shown to be well-quasi-ordered : finite sequences and subsets in section 3, finite trees in section 4, graphs of bounded tree-width in section 5 and finally general graphs in section 6. We take a detour in section 7 and describe implications of Robertson and Seymour's polynomial time algorithm to solve the disjoint paths problem. Hereditary properties are described in section 8. Wagner's conjecture and the disjoint paths problem are then used to show that any hereditary property can be decided in polynomial time. Finally in section 9 we conclude with a short discussion and closing remarks.

(*)A surface is a compact 2-manifold.

2 Preliminaries, definitions and notations

We denote the sequence $x_j, x_{j+1}, \ldots, x_k$ by $(x_i)_{i=j}^k$ where k can be ∞ in which case we meant the infinite sequence $x_j, x_{j+1}, \ldots$. When there is no confusion we shall drop the indices and use (x_i). If (y_i) is a subsequence of (x_i) we denote $(y_i) \subset (x_i)$. To denote that $s_i \in S$ for all i and some S, we write $(s_i) \subset S$. Given a sequence s denote the i^{th} element of s by s_i and denote the set of elements of s by $\overline{s}$.

Given two graphs G and H we shall denote by $G \cup H$ the graph with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H)$. Similarly for $G \cap H$.

A *rooted graph* G is a graph along with a sequence $r(G)$ of *distinct* vertices from $V(G)$ called the roots of G. Let $\overline{r}(G)$ be the set of roots of G and $r_i(G)$ be the i^{th} root of G. In this survey, graphs shall mean finite undirected rooted graphs with loops and multiple edges unless otherwise specified. Although Wagner's conjecture was stated for unrooted graphs, we shall be proving rooted graphs versions of simplifications of Wagner's conjecture. The unrooted versions of the theorems clearly follow by having $r(G) = \emptyset$ for an unrooted graph G.

However trees for us shall be directed and rooted i.e. there is a unique root from which all edges radiate outwards. Let $s(T)$ be the set of components of $T - r(T)$. For $v \in V(T)$ let T_v be the maximal subtree of T rooted at v. Sometimes we may also use $T(v)$ in place of T_v where the subscripts start to clutter up. For $u, v \in V(T)$ let $[u, v]_T$ be the undirected path in T from u to v. The subscript is dropped when there is no confusion.

The vertices or edges (or both) of a graph G can be labeled. We shall denote the labels by $\phi_G(x)$ for each $x \in V(G) \cup E(G)$. Again when there is no confusion the subscripts are dropped.

Let G, H be two unrooted graphs. We say H is a *minor* of G or G has an H*-minor* if we can obtain (an isomorphic copy of) H from G via a sequence of edge/vertex deletions and edge contractions. We say G is an *expansion* of H if we obtain (an isomorphic copy of) G from H by replacing each vertex v of H by a connected subgraph G_v and replacing each edge uv originally incident on v in H by an edge uw for some $w \in V(G_v)$. We say G is a *subdivision* of H if we can obtain (an isomorphic copy of) G from H by a sequence of edge subdivisions. We say H is topologically contained in G if there is a subgraph of G that is (an isomophic to) a subdivision of H. If H is a tree then a subdivision of H is a tree G obtained from H by replacing each edge $u \rightarrow v \in E(H)$ by a directed path in the same direction as $u \rightarrow v$. If G, H are rooted graphs, we say H is a minor of G if H is a minor of G as unrooted graphs, H and G has the same number of roots and the i^{th} root of G is contracted to or corresponds with the i^{th} root of H for each i.

Our motivation for using the above definition of minor containment for rooted graphs is the following. Suppose G, G_1, G_2 are rooted graphs with $G_1 \preceq_m G_2$. Let X be a sequence of $|r(G_1)| = |r(G_2)|$ distinct vertices in G. For $i = 1, 2$, let G_i' be the graph obtained from G_i and G when we identify $r(G_i)$ with X. Then it is easy to see from the definition of minor containment that G_1' is still a minor of G_2'.

3 Well-quasi-orderings

One way of expressing Wagner's conjecture is using well-quasi-orderings.

Definition 3.1 (Quasi-orderings). *A relation $\preceq$ on a set S is a quasi-order if $\preceq$ is reflexive and transitive.*

Definition 3.2 (Well-quasi-orderings). *A quasi-order $\preceq$ on a set S is a well-quasi-order if for any sequence $(s_i)_{i=1}^{\infty} \subset S$ there exist $1 \leq i < j$ with $s_i \preceq s_j$.*

If $\preceq$ is a partial-order on S then $\preceq$ is a well-quasi-order if and only if there is no infinite decreasing sequence in S and there is no infinite set of incomparable elements of S.

We can define three quasi-orderings on the family of graphs. Let G and H be graphs. The first is the subgraph ordering, $H \subset G$ if H is a subgraph of G (up to isomorphism). The second is the topological ordering, $H \preceq_t G$ if H is topologically contained in G. The third ordering is the minor ordering, $H \preceq_m G$ if H is a minor of G. The subgraph ordering is not a well-quasi-order for graphs. Consider the sequence (C_i) where C_n is the n-cycle. The topological ordering is not a well-quasi-order for graphs either. For $n \geq 1$ consider two n-cycles with vertices $v_1, v_2, \ldots, v_n$ and $u_1, \ldots, u_n$ respectively. Join v_i, u_i by an edge for each $1 \leq i \leq n$ and let C'_n be the resulting graph. Then the set $\{C'_n : n \geq 1\}$ is not well-quasi-ordered by topological containment. Wagner's conjecture states that the minor ordering is a well-quasi-ordering for graphs.

The next observation shows that given any sequence from a well-quasi-ordered set, not only can we find two comparable elements we can find an infinite increasing subsequence.

Lemma 3.3. *If $\preceq$ is a well-quasi-order on S then given any sequence $(s_i)_{i=1}^{\infty} \subset S$ there exists a subsequence $(s_{k_i})_{i=1}^{\infty} \subset (s_i)_{i=1}^{\infty}$ such that $s_{k_i} \preceq s_{k_{i+1}}$ for all $i \geq 1$.*

Proof. Consider a complete countably infinite undirected graph with vertices $\{1, 2, \ldots\}$. Let $1 \leq i < j$. Color the edge ij red if $s_i \preceq s_j$. Otherwise color it blue. By Ramsey's theorem for infinite graphs there either exists $k_1 < k_2 < \cdots$ with $s_{k_i} \preceq s_{k_{i+1}}$ for all $i \geq 1$ (all edges are red) or $l_1 < l_2 < \cdots$ with $s_{l_i} \not\preceq s_{l_j}$ for all $i < j$ (all edges are blue). But $\preceq$ is a well-quasi-order, so the former holds. □

Let $\preceq$ be a well-quasi-ordering on a set S. We can extend a well-quasi-order over a set S to the family of finite sequences from S and the family of finite subsets of S. For $\mathfrak{s} = (s_i)_{i=1}^{n} \subset S$ and $\mathfrak{t} = (t_i)_{i=1}^{m} \subset S$ we say $\mathfrak{s} \preceq \mathfrak{t}$ if there is an $\alpha : \{1, 2, \ldots, n\} \to \{1, 2, \ldots, m\}$ such that $\alpha(l) < \alpha(l+1)$ and $s_l \preceq s_{\alpha(l)}$ for each $l = 1, \ldots, n$. For $\mathfrak{s}, \mathfrak{t} \subset S$ finite subsets, we say $\mathfrak{s} \preceq \mathfrak{t}$ if there is an injective map $\alpha : \mathfrak{s} \to \mathfrak{t}$ such that $s \preceq \alpha(s)$ for each $s \in \mathfrak{s}$.

Proposition 3.4 (Higman [7]). *Let $\preceq$ be a well-quasi-ordering on a set S. The family of finite sequences of S is well-quasi-ordered by $\preceq$.*

Proposition 3.5 (Higman [7]). *Let $\preceq$ be a well-quasi-ordering on S. The family of finite subsets of S is well-quasi-ordered by $\preceq$.*

The above two propositions are quite useful. However their proofs are very similar to the proof of theorem 4.1 below and we leave it to the reader as an exercise to adapt the proof to proposition 3.4 and proposition 3.5.

4 Trees

We can view finite sequences in S as directed paths with the vertices labeled by elements in S. Given labeled directed paths P and Q where $P \preceq_t Q$ as unlabeled paths, and for each vertex $v \in V(P)$ and the corresponding vertex $u \in V(Q)$ we have $\phi_P(v) \preceq \phi_Q(u)$(†), we say $P \preceq Q$. Note that this definition of $\preceq$ on labeled directed paths coincides with the definition of $\preceq$ on finite sequences in proposition 3.4. Then proposition 3.4 says that if S is well-quasi-ordered the family of directed paths labeled with elements of S is well-quasi-ordered.

We can generalize proposition 3.4 to the class of trees labeled by elements of S. Given labeled trees R and T where $R \preceq_t T$ and for each vertex $v \in V(R)$ and the corresponding vertex $u \in V(T)$ we have $\phi_R(v) \preceq \phi_T(u)$, we say $R \preceq T$. Theorem 4.1 below shows that $\preceq$ is a well-quasi-order over trees labeled with elements of S.

Theorem 4.1. *If S is well-quasi-ordered then so is the class of trees labeled by elements of S.*

Proof. Suppose the class of labeled trees is not well-quasi-ordered. Pick a minimal sequence of trees that is a counterexample inductively as follows. For $i = 1, 2, \dots$ pick a labeled tree T_i such that there is a counterexample starting with $T_1, \dots, T_i$ but there is no counterexample starting with $T_1, \dots, T_{i-1}, T_i'$ where $T_i' \in s(T_i)$.

Consider the sequence $(\phi(r(T_i)))_{i=1}^\infty$. As S is well-quasi-ordered there is a subsequence $(\phi(r(T_{k_i})))_{i=1}^\infty$ such that $\phi(r(T_{k_i})) \preceq \phi(r(T_{k_{i+1}}))$ for each $i \geq 1$.

Let $\mathfrak{T} = \cup\{s(T_{k_i}) : i \geq 1\}$. Suppose there is a sequence of trees $(T_{l_i}')_{i=1}^\infty \subset \mathfrak{T}$ where $T_{l_i}' \in s(T_{l_i})$ such that $T_{l_i}' \not\preceq T_{l_j}'$ for each $i < j$. As $\mathbb{N} = \{1, 2 \dots\}$ is well-quasi-ordered, $\preceq$ is reflexive and each $s(T_{k_i})$ is finite, we may assume that $l_i < l_{i+1}$ for each i.

Consider the sequence $T_1, \dots, T_{l_1-1}, T_{l_1}', T_{l_2}', \dots$. For each $1 \leq i \leq l_1 - 1$ and $j \geq 1$ we have $T_{l_j}' \preceq T_{l_j}$ and $T_i \not\preceq T_{l_j}$ hence $T_i \not\preceq T_{l_j}'$ by transitivity. Also $T_i \not\preceq T_j$ for $1 \leq i < j < l_1 - 1$ and $T_{l_i}' \not\preceq T_{l_j}'$ for $i < j$. So the sequence $T_1, \dots, T_{l_1-1}, T_{l_1}', T_{l_2}', \dots$ is a counterexample too, contradicting the minimality of $(T_i)_{i=1}^\infty$. Therefore $\mathfrak{T}$ is well-quasi-ordered. By proposition 3.5 the sequence of sets of trees $(s(T_{k_i}))_{i=1}^\infty$ is well-quasi-ordered as well. So there exists $i < j$ such that $s(T_{k_i}) \preceq s(T_{k_j})$. But $\phi(r(T_{k_i})) \preceq \phi(r(T_{k_j}))$ too and a little thought demonstrates that $T_{k_i} \preceq T_{k_j}$, a contradiction and so the class of trees labeled with S is well-quasi-ordered. □

The above proof method is a useful method. The essential idea to proving that a certain set S is well-quasi-ordered is to assume it is not and pick a counterexample sequence that is minimal in some sense. Then we find another sequence based on the first that is even smaller than the first, causing a contradiction so S must be well-quasi-ordered. This "minimal bad sequence" idea was attributed by Robertson and Seymour to Nash-Williams [13].

Theorem 4.1 is a labeled version of a classical theorem by Kruskal.

Theorem 4.2 (Kruskal [10]). *The collection of finite rooted trees is well-quasi-ordered by topological containment.*

(†)Of course there might be many ways of subdividing the edges of P to obtain Q. What we mean is that there is some way of subdividing the edges of P such that we obtain Q and $\phi_P(v) \preceq \phi_Q(u)$ for each pair of corresponding vertices v and w.

5 Tree decompositions

Definition 5.1 (Tree decomposition). *A tree decomposition of a rooted graph G is a vertex and edge labeled tree T such that*

1. *for each $t \in T$, $\phi(t) \subset G$ is a rooted graph;*
2. *for each $t_1, t_2 \in T$, $E(\phi(t_1) \cap \phi(t_2)) = \emptyset$;*
3. *for each $e \in E(G)$ there is a $t \in T$ such that $e \in E(\phi(t))$;*
4. *for $t_1, t_3 \in V(T)$ and $t_2 \in V([t_1, t_3]_T)$, $\phi(t_1) \cap \phi(t_3) \subset \phi(t_2)$;*
5. *$r(\phi(r(T))) = r(G)$;*
6. *for each $u \to v \in E(T)$, $\phi(u \to v) = r(\phi(v)) = V(\phi(v) \cap \phi(u))$.*

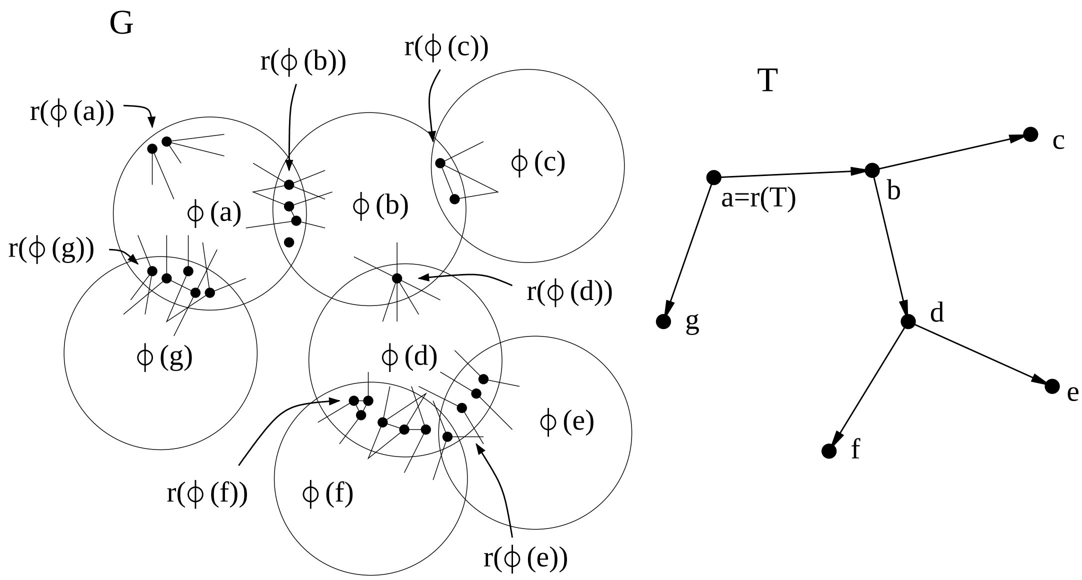

A tree decomposition T of a rooted graph G essentially breaks G down into smaller pieces (the $\phi(t)$'s for $t \in V(T)$) which we can patch together (along $r(\phi(t))$) in a tree-like fashion to get back G. The figure above depicts a tree decomposition T (on the right) and how the pieces $\phi(t)$ are patched together to form G (on the left).

For each $u \to v \in E(T)$, define the order of v to be $o(v) = |r(\phi(v))|$ and the order of $u \to v$ to be $o(u \to v) = o(v)$. For a subtree $S \subset T$ define $\phi(S) = \cup\{\phi(v) : v \in V(S)\}$ to be the underlying rooted graph which S is a tree decomposition of. The root of $\phi(S)$ is $r(\phi(S)) = r(\phi(r(S)))$.

The *width* of a tree decomposition T is $\max_{t \in V(T)} |V(\phi(t))| - 1$. The *tree-width* of a graph G is the minimum width of any tree decomposition of G. A graph has tree-width 1 if and only if it is a tree, while a graph has tree-width ≤ 2 if and only if it is a series-parallel graph. Let $w \ge 1$. In many respects graphs with tree-width $\le w$ show many characteristics of trees and can be viewed as "fat" trees whose "branches" all have width $\le w$.

Let P be a path of G connecting some vertex in $\phi(u)$ to some vertex in $\phi(v)$ where $u, v \in V(T)$. Intuitively it is clear that for each $e \in E([u, v]_T)$, P must meet some vertex in $\phi_T(e)$. We call this the *separation property* of tree decompositions.

Suppose that $S \subset T$ are trees. Then clearly $S \preceq_t T$. The following variation of this property for tree decompositions turns out to be very useful. Under certain conditions we want $\phi(S) \preceq_m \phi(T)$ if $S \subset T$ are tree decompositions. Clearly $\phi(S) \subset \phi(T)$ so the only tricky part is to perform edge contractions such that $r(\phi(T))$ is contracted to $r(\phi(S))$. A necessary and sufficient condition for this is that $|r(\phi(S))| = |r(\phi(T))| = k$ and there are k disjoint paths connecting $r_i(\phi(T))$ to $r_i(\phi(S))$ for each $1 \leq i \leq k$. Consider the subgraph $G' = \phi(S) \cup \{$ the k disjoint paths $\} \subset \phi(T)$ with $r(G') = r(T)$. Contracting the k disjoint paths of G' to $r(\phi(S))$, we see that $\phi(S) \preceq_m G' \preceq_m \phi(T)$. As a result of the separation property and Menger's theorem a necessary condition is that for each edge e on the path from $r(T)$ to $r(S)$ we have $o(e) \geq k$. Lemma 5.2 below, due in part to Thomas [21] and in part to Robertson and Seymour [18] shows that there are always tree decompositions in which the necessary condition is also sufficient. We call such tree decompositions *linked.*

Lemma 5.2. *If G has tree-width $\leq w$ there exists a tree decomposition T of G with width $\leq w$ such that for every directed path P from u to v in T with $o(u) = o(v) \leq o(e)$ for all $e \in E(P)$, $\phi(T_v)$ is a minor of $\phi(T_u)$.*

Idea of proof. Start with a tree decomposition T of G with width $\leq w$. If T is linked we are done. Otherwise there is a directed path P from u to v with $o(u) = o(v) \leq o(e)$ for all $e \in E(P)$ but $\phi(T_v)$ is not a minor of $\phi(T_u)$. Suppose there are $k = o(u)$ disjoint paths from $\overline{r}(\phi(u))$ to $\overline{r}(\phi(v))$. Then we can reorder the vertices in $r(\phi(u))$ such that the k paths match $r_i(\phi(u))$ to $r_i(\phi(v))$ for all $1 \leq i \leq k$. If there are $< k$ disjoint paths, by Menger's theorem there is a separation S with $|S| < k$ separating $r(\phi(u))$ from $r(\phi(v))$.

The intuitive idea now is to transform the tree decomposition such that there is an edge $e \in E(P)$ with $\phi(e) = S$. This way the antecedent becomes false and the lemma is trivially satisfied for the path P. To do this we essentially "pinch" the tree decomposition at S.

The figure below illustrates the idea for a simple case. On the left we have a path $P = u \to a \to v$. for each $t \in V(P)$, $o(t) = 4$ but there is a separation S with $|S| = 3$. So we "pinch" the tree decomposition at S, separating $\phi(a)$ into two parts : $\phi(a')$ on the left of S and $\phi(a'')$ on the right, with $r(\phi(a')) = r(\phi(a))$ and $\overline{r}(\phi(a'')) = S$. This is shown on the right. The situation is more complicated when S spans a few pieces of the tree decomposition, but the idea is generally the same. We repeat this procedure until we obtain a linked tree decomposition.

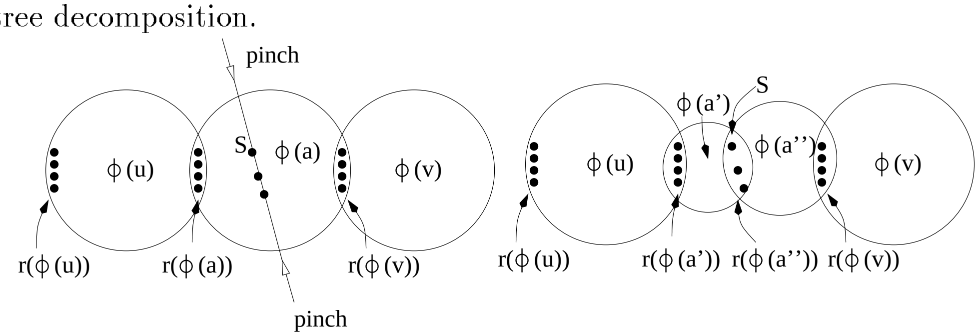

□

A lemma by Robertson and Seymour [18] illustrates why linked tree decompositions are important. They allow us to find "minimal bad sequences" if graphs of bounded tree-widths are not well-quasi-ordered.

Lemma 5.3. *Let $w \geq 1$ and $(T_i)_{i=1}^{\infty}$ be a sequence of linked tree decompositions each with width $\leq w$. Suppose that $\phi(T_i) \not\preceq_m \phi(T_j)$ for each $1 \leq i < j$. Then there exists $k_1 < k_2 < \cdots$ and $t_i \in V(T_{k_i})$ such that $\phi(T_{k_i}(t_i)) \not\preceq_m \phi(T_{k_j}(t_j))$ but $\cup\{\phi(T') : T' \in s(T_{k_i}(t_i)), i \geq 1\}$ is well-quasi-ordered with respect to $\preceq_m$.*

Theorem 5.4 (Robertson and Seymour [18]). *For each $w \geq 0$ the collection of graphs of tree-width $\leq w$ is well-quasi-ordered by minor containment.*

Proof. Suppose the collection is not well-quasi-ordered. Let $(G_i)_{i=1}^{\infty}$ be a sequence of graphs with tree-width $\leq w$ such that $G_i \not\preceq_m G_j$ for every $i < j$. By lemma 5.2 G_i has a linked tree decomposition T_i with width $\leq w$ for each i. By lemma 5.3 there exists a sequence $(T_i^{(1)})_{i=1}^{\infty}$ of subtrees of some T_i's such that $\phi(T_i^{(1)}) \not\preceq_m \phi(T_j^{(1)})$ for each $i < j$ but $\mathfrak{G} = \cup\{\phi(T') : T' \in s(T_i^{(1)}), i \geq 1\}$ is well-quasi-ordered with respect to $\preceq_m$.

Now $|V(\phi(r(T_i^{(1)})))| \leq w + 1$ for all i so there are only finitely many possible choices for $\phi(r(T_i^{(1)}))$ up to isomorphism. Pick a subsequence $(T_i^{(2)}) \subset (T_i^{(1)})$ such that $\phi(r(T_i^{(2)}))$ is isomorphic to $\phi(r(T_j^{(2)}))$ and $r(\phi(r(T_i^{(2)})))$ is mapped to $r(\phi(r(T_j^{(2)})))$ under the isomorphism for all $i, j \geq 1$. For simplicity we shall assume that $\phi(r(T_i^{(2)})) = G_0$ for all $i \geq 1$.

Let r be a finite sequence of distinct vertices in $V(G_0)$ and $S_i^r = \{\phi(T') : T' \in s(T_i^{(2)})$ and $r(\phi(T')) = r\}$ for each $i \geq 1$. Since $\mathfrak{G}$ is well-quasi-ordered, by proposition 3.5 and lemma 3.3 there is a subsequence $(S_{k_i}^r)_{i=1}^{\infty} \subset (S_i^r)_{i=1}^{\infty}$ such that for each $i < j$ there is an injective map $\alpha : S_{k_i}^r \to S_{k_j}^r$ with $G' \preceq_m \alpha(G')$ for each $G' \in S_{k_i}^r$. Apply this for each r (there are only finitely many of them), and we obtain a subsequence $(T_i^{(3)}) \subset (T_i^{(2)})$ such that for each $i < j$ there is an injective map $\alpha : s(T_i^{(3)}) \to s(T_j^{(3)})$ with $\phi(T') \preceq_m \alpha(\phi(T'))$ and $r(\phi(T')) = r(\alpha(\phi(T')))$ for each $T' \in s(T_i^{(3)})$. But $\phi(r(T_i^{(3)})) = \phi(r(T_j^{(3)}))$ too and it is not hard to see that $\phi(T_i^{(3)}) \preceq_m \phi(T_j^{(3)})$. This is a contradiction, so the collection of graphs of tree-width $\leq w$ is well-quasi-ordered. □

An application of theorem 5.4 is another simplification of Wagner's conjecture.

Theorem 5.5. *If H is a planar graph the set of graphs with no H-minor is well-quasi-ordered by minor containment.*

Corollary 5.6. *If $(G_i)_{i=1}^{\infty}$ is a sequence of graphs with G_1 planar then there are $1 \leq i < j$ such that G_i is a minor of G_j.*

Proof. If there is an $i \geq 2$ such that G_1 is a minor of G_i we are done. Otherwise $\{G_i : i \geq 2\}$ is well-quasi-ordered by theorem 5.5 so there are $2 \leq i < j$ such that G_i is a minor of G_j. □

Theorem 5.5 easily follows from theorem 5.4 and the following due to Robertson and Seymour [17].

Theorem 5.7. *If H is a planar graph then the set of graphs with no H-minor has bounded tree-width.*

Sketch proof. For $n \geq 1$ let the n-grid be a graph with vertex set $\{(i,j) : 1 \leq i,j \leq n\}$ and two vertices (i,j) and (i',j') are connected by an edge if and only if $|i-i'|+|j-j'|=1$. It can be shown that the n-grid has tree-width n [16]. The tree-width theorem [17] shows that for each n there exists an $m(n) \geq 1$ such that every graph with tree-width $\geq m(n)$ has an n-grid-minor. Hence a graph has large tree-width if and only if it has a large grid-minor.

Let H be a planar graph. There is a planar graph H' with maximum degree 3 such that H is a minor of H' (consider expanding every vertex of H of degree ≥ 3 into a connected subgraph with each vertex having degree 1 or 2 while preserving planarity). Now embed H' on the plane such that each edge is a line segment and transform each line segment into a polygonal arc consisting of vertical and horizontal segments. Now it is easy to see that H' is a subdivision of some large n-grid, hence H is a minor of the n-grid. Now every graph with tree-width $\geq m(n)$ has an H-minor. □

Note that theorem 5.7 is best possible in the sense that if H is not a planar graph there are graphs of arbitrarily large tree-width with no H-minor (consider the n-grids).

Because of their tree-like structure, many NP-complete problems can be efficiently solved for graphs of bounded tree-widths. For example Arnborg and Proskurowski [3] have shown that determining whether a graph can be k-colored and whether a graph has a Hamiltonian cycle can be done in linear time for graphs of bounded tree-widths. Coupled with Bodlaender's linear time algorithm to determine whether a graph has tree-width $\leq w$ [5], we obtain algorithms running in linear time that either returns that a graph has tree-width $> w$, or is k-colorable (respectively has a Hamiltonian cycle), or not k-colorable (respectively has no Hamiltonian cycle) for each fixed w and k.

6 Wagner's conjecture

In this section we describe Robertson and Seymour's proof of Wagner's conjecture [15]. The material presented here is based largely upon a survey by Thomas [20], because Robertson and Seymour's paper [15] has not been published yet(‡).

The proof is similar to theorem 5.5 and corollary 5.6. Given a graph G let $\mathbb{G} - G = \{G' : G \text{ is not a minor of } G'\}$. Take a sequence of graphs $(G_i)_{i=1}^{\infty}$. If there is an $i > 1$ such that $G_1 \preceq_m G_i$ we are done. Otherwise $G_i \in \mathbb{G} - G_1$ for each $i > 1$ so if we can show that $\mathbb{G} - G_1$ is well-quasi-ordered for each G_1 then we can find $1 < i < j$ with $G_i \preceq_m G_j$ and we are done. Theorem 5.7 showed that if G_1 is planar then graphs in $\mathbb{G} - G_1$ has bounded tree-width (i.e. each piece in the decomposition has a bounded number of vertices) hence is well-quasi-ordered. Robertson and Seymour [14] showed that for a general graph G_1, each graph in $\mathbb{G} - G_1$ also has a certain tree decomposition in which each piece is bounded in some sense and this can be exploited to show that $\mathbb{G} - G_1$ is well-quasi-ordered along the same lines as theorem 5.4. We shall describe Robertson and Seymour's theorem concerning the structure of graphs in $\mathbb{G} - G_1$.

(‡)I suspect that the paper has been submitted by this time, since the manuscript was already written as early as 1996 and the result was widely accepted. Unfortunately it has not been published yet. I am appalled to see the time it took for papers to get published in academic journals. Judging from the papers I used for this survey it seems the average time required is at least three to five years!

Definition 6.1 (Vortex decomposition). *Let G be a graph and U be a cyclic ordering of a subset of $V(G)$. For each $u \in U$ let $X_u \subset V(G)$. We say $(X_u)_{u\in U}$ is a vortex decomposition of (G, U) if*

1. *for each $u \in U$, $u \in X_u$;*
2. *$\cup_{u\in U} X_u = V(G)$;*
3. *for each $vw \in E(G)$ there is a $u \in U$ with $v, w \in X_u$;*
4. *if u_1, u_2, u_3, u_4 occur in U in that order then $X_{u_1} \cap X_{u_3} \subset X_{u_2} \cup X_{u_4}$.*

The last condition means that if $v \in V(G)$ then $\{u \in U : v \in X_u\}$ is either empty, a contiguous segment of U or U itself. The width of a vortex decomposition is $\max_{u\in U} |X_u|$.

Given a surface S and $k \geq 0$ let $S - k$ be the topological space obtained by removing k disjoint disks from S. Note that each connected component of the boundary of S is homeomorphic to the circle.

Definition 6.2 ((S, k)-nearly drawable). *Given a surface S, $k \geq 0$ and $C_1, \ldots, C_k$ the components of the boundaries of $S - k$, and a graph G we say G is (S, k)-nearly drawable if there exists $X \subset V(G)$ with $|X| \leq k$ such that $G - X$ can be written as $G_0 \cup G_1 \cup \cdots \cup G_k$ where*

1. *G_0 is embeddable in $S - k$;*
2. *G_i for $i \geq 1$ are mutually disjoint;*
3. *for each i, $U_i = V(G_0) \cap V(G_i) = V(G_0) \cap C_i$;*
4. *for each i, (G_i, U_i) has a vortex decomposition with width $\leq k$ where we order U_i according to the cyclic order of elements of U_i on C_i.*

Note that G is (S, k)-nearly drawable means that G can mostly be embedded on S except for the k "small" regions $G_1, \ldots, G_k$ and X.

Definition 6.3 ($\mathfrak{G}$-tree decomposition). *Let T be a tree decomposition and $\mathfrak{G}$ a family of graphs. For each $e \in E(G)$ let K_e be the clique with vertices $V(\phi(e))$. For $t \in V(T)$ let $\psi(t) = (\cup\{K_e : e \text{ is incident on } t\}) \cup \phi(t)$. We say T is a $\mathfrak{G}$-tree decomposition if $\psi(t) \in \mathfrak{G}$ for each $t \in V(T)$.*

This generalizes the notion of tree-widths. If $\mathfrak{G}_w$ is the set of graphs with at most $w + 1$ vertices, then tree decompositions of width $\leq w$ are exactly $\mathfrak{G}_w$-tree decompositions.

Theorem 6.4 (Robertson and Seymour [14]). *Given a graph G there exists $k \geq 0$ such that every graph in $\mathbb{G} - G$ has a $\mathfrak{G}$-tree decomposition, where $\mathfrak{G}$ is the set of graphs that can be (S, k)-nearly drawn where S is a surface on which G is not embeddable.*

Note that the set of surfaces on which G cannot be embedded is finite. Hence $\mathfrak{G}$ as in theorem 6.4 is bounded in the sense that they must all be almost embeddable on some surface from a finite set of surfaces, that there can be at most k regions in them that cannot be embedded, and that these obstructions must be vortex decompositions with width $\leq k$. Robertson and Seymour then showed that graphs with structures as described in theorem 6.4 are well-quasi-ordered, completing the proof of Wagner's conjecture.

7 The disjoint paths problem

The discussion preceding the definition of a linked tree decomposition in section 5 hints at the usefulness of determining whether there are disjoint paths joining correspond vertices from two sequences. The disjoint paths problem is as follows : given a graph G and vertices $s_1, \ldots, s_k$ and $t_1, \ldots, t_k$ are there disjoint paths joining s_i to t_i for each $1 \le i \le k$? This is not the same as the network flow problem, which imposes no ordering on the way the vertices can be connected up. In fact the disjoint paths problem is NP-complete if we allow k to vary [8]. A result of Robertson and Seymour shows that the disjoint paths problem can be solved in polynomial time if k is fixed [19]. This is an important result because it implies polynomial time algorithms for deciding any hereditary property (section 8).

Algorithm 7.1 (Disjoint paths). *For all $k \ge 1$ there exists a polynomial time algorithm which solves the disjoint paths problem for k fixed.*

The polynomial time algorithm for the disjoint paths problem implies a polynomial time algorithm for whether a given graph contains a fixed graph as a minor. We first show that there is a polynomial time algorithm for topological containment.

Algorithm 7.2 (Topological containment). *For any graph H there exists a polynomial time algorithm which when given a graph G determines whether G topologically contains H.*

Proof. Let $\mathfrak{v} = |V(H)|$ and $\mathfrak{e} = |E(H)|$. Consider a graph G with n vertices. Note $H \preceq_t G$ if and only if there are injective maps $f : V(H) \to V(G)$ and $g : E(H) \to \{P : P \text{ is a path in } G\}$ such that for $uv \in E(H)$, $g(uv)$ is a path in G connecting $f(u)$ to $f(v)$ and the paths $g(E(H))$ are mutually vertex-disjoint except possibly at the endpoints.

To determine whether $H \preceq_t G$, we first iterate over all possibilities for f. There are only $\binom{n}{\mathfrak{v}} \in O(n^{\mathfrak{v}})$ possible f's. For each $v \in V(H)$ we replace $f(v) \in V(G)$ with a clique K_v of size $\deg_H v$. Replace each $uf(v) \in E(G)$ with $\deg_H v$ edges uw for each $w \in V(K_v)$. Now for each $e = uv \in E(H)$ we assign $s_e \in V(K_u)$ and $t_e \in V(K_v)$ such that if $e, f \in E(H)$ and $e \ne f$, we have $s_e \ne s_f$ and $t_e \ne t_f$. Invoking algorithm 7.1, if there are disjoint paths joining s_e to t_e for each $e \in E(H)$ then we can construct g as in the previous paragraph and so $H \preceq_t G$. ☐

Lemma 7.3. *For all graphs H there exists graphs $H_1, \ldots, H_k$ such that any graph G has a H-minor if and only G topologically contains H_i for some $1 \le i \le k$.*

Proof. Note that G has a H-minor if and only if we can obtain a subgraph of G from H by a sequence of vertex expansions into an edge (along with its endpoints). If each vertex of H has degree at most 3, then every G with a H-minor topologically contains H. This is because every vertex expansion of H into an edge is an elementary subdivision, and every vertex of the resulting graph still has degree at most 3.

Let $C(H) = \sum_{v \in V(H)} \max\{0, \deg v - 3\}$ and consider expanding a vertex $v \in V(H)$ with $\deg v > 3$ into an edge $v_1 v_2$ such that both v_1 and v_2 has degree at least 3. Such an expansion always decrease $C(H)$ by at least 1, hence we can perform at most $C(H)$ such expansions before the resulting graph H' has $C(H') = 0$, i.e. every vertex of H' has degree at most 3.

At each point we have a choice from among at most $|V(H)| + C(H)$ vertices to expand and a choice from among $2^{|V(H)|+C(H)}$ possible partitions of the neighbors of the vertex. Hence there are at most $((|V(H)| + C(H))2^{|V(H)|+C(H)})^{C(H)}$ graphs $\mathcal{H}$ such that G has a H-minor if and only if G topologically contains some $H' \in \mathcal{H}$. This is because every graph obtained from H using such expansion has to be a subdivision of some $H' \in \mathcal{H}$. □

Algorithm 7.4 (Minor containment). *For all graphs H there exists a polynomial time algorithm which when given a graph G determines whether there exists a H-minor in G.*

Proof. By lemma 7.3 there are graphs $H_1, \ldots, H_k$ such that G has a H-minor if and only if G has a subgraph that is a subdivision of some H_i, $1 \leq i \leq k$. So to determine whether G has a subgraph with H as a minor we only need to apply algorithm 7.2 to G and each H_i. □

Robertson and Seymour [19] actually showed that algorithm 7.1 and algorithm 7.4 can be implemented in $O(|V(G)|^3)$ time. They described an algorithm that generalized both algorithm 7.1 and algorithm 7.4. Let $\delta \geq 0$ and G be a rooted graph. Define the *δ-folio* of G to be the set of all minors H of G such that $|V(H) - r(H)| \leq \delta$ and $|E(H)| \leq \delta$.

Algorithm 7.5 (Folio). *For each $\zeta, \delta \geq 0$ there is an algorithm with run time $O(|V(G)|^3)$ which when given a rooted graph G with $|r(G)| \leq \zeta$ determines the δ-folio of G.*

Algorithm 7.1 can be solved by computing the 0-folio of G with $r(G) = \{s_i, t_i : 1 \leq i \leq k\}$ and determining whether the relevant disjoint paths are present in the folio; while algorithm 7.4 can be solved by determining the $\max\{|V(H)|, |E(H)|\}$-folio of G with $r(G) = \emptyset$.

Algorithm 7.5 is very complicated and beyond the scope of this survey. The reader is referred to [19] for the actual algorithm. We shall only give a very rough outline. The branch-width of a graph is a notion related to the tree-width. In fact it is within a constant multiplicative factor of the tree-width. Given a graph G and $w \geq 1$ we can determine whether G has branch-width $\leq 3w$ or G has branch-width $\geq w$. If G has bounded branch-width (i.e. $\leq 3w$) then there is a simple algorithm to determine the required folio. This is not surprising given that branch-width is closely related to tree-width, and many NP-complete problems are solvable in polynomial time for graphs of bounded tree-width. Otherwise if G has large branch-width (i.e. $\geq w$) then due to the high connectivity of G we can find a vertex that is irrelevant to the required folio (i.e. the folio is the same whether the vertex is present or removed). We remove the irrelevant vertex and repeat. As each repetition takes $O(|V(G)|^2)$ time the algorithm has a run-time of $O(|V(G)|^3)$.

8 Hereditary properties

Definition 8.1 (Hereditary property). *A property P is hereditary if for all G with property P and all minors $H \preceq_m G$, H has property P.*

Examples of hereditary properties are "embeddable in a given surface S" and "has tree-width $\leq w$". The following is a direct consequence of the proof of Wagner's conjecture.

Theorem 8.2. *For any hereditary property P there exists a finite set $\mathcal{H}$ of graphs such that a graph G has property P if and only if G has no H-minor for all $H \in \mathcal{H}$.*

Proof. Let $\mathcal{H}$ be the set of all graphs H such that H does not have property P yet for every non-trivial minor H' of H, H' has property P. Then every graph G without property P must have a minor in $\mathcal{H}$. Suppose $\mathcal{H}$ is infinite. Pick a sequence $(H_i)_{i=1}^{\infty}$ of distinct graphs from $\mathcal{H}$. As each H_i is minor minimal $H_i \not\preceq_m H_j$ for each $i < j$. This contradicts Wagner's conjecture so $\mathcal{H}$ must be finite. □

Theorem 8.2 is one of the most important results in recent graph theory. A powerful consequence of theorem 8.2 is a Kuratowski type characterization for graphs embeddable in any surface.

Theorem 8.3. *For any surface S there exists graphs $H_1, H_2, \ldots, H_k$ such that a graph G is embeddable in S if and only if G has no H_i-minor for all $1 \leq i \leq k$.*

As a consequence of the $O(|V(G)|^3)$ time complexity of the disjoint paths problem, given any hereditary property P we can determine whether a graph G has property P in $O(|V(G)|^3)$ time too.

Algorithm 8.4 (Hereditary). *For all hereditary property P there exists an $O(|V(G)|^3)$ time algorithm which when given a graph G determines whether G has property P.*

Proof. By Wagner's conjecture there are graphs $H_1, \ldots, H_k$ such that G has property P if and only if G does not have any subgraph with H_i as a minor for all $1 \leq i \leq k$. To determine whether G has any subgraph with H_i as a minor apply algorithm 7.4. □

A disadvantage of the approach is that because theorem 8.2 uses a non-constructive proof we do not know how to construct the excluded minors or the polynomial time algorithm given any hereditary property.

As corollary to algorithm 8.4 and theorem 8.3 we have the following

Algorithm 8.5 (Embeddability). *For all surfaces S there exists a polynomial time algorithm which when given a graph G determines whether G can be embedded in S.*

Note that if H is a minor of some subgraph of G, and G has tree-width $\leq w$ then H has tree-width $\leq w$. So applying algorithm 8.4

Algorithm 8.6 (Tree-width). *For all $w \geq 1$ there exists a polynomial time algorithm which when given a graph G determines if G has tree-width $\leq w$.*

It is interesting to note that both algorithm 8.5 and algorithm 8.6 can actually be solved in linear time. Bodlaender [5] constructed an algorithm that determines whether a graph has tree-width $\leq w$ and if so returns a tree decomposition of the graph, all in linear time. Mohar [12] constructed an algorithm that determines whether a graph can be embedded in a given surface S in linear time as well. If it can, the algorithm returns an embedding of the graph in S and if it cannot, it returns a smallest subgraph not embeddable in S. The method used by Mohar directly tries to find an embedding of the graph and is independent of Robertson and Seymour's approach. In fact it leads to a constructive proof of theorem 8.2 and gives an upper bound on the number of edges that the excluded minors can have.

9 Discussion and closing remarks

In this survey we have described material surrounding Wagner's conjecture. Besides the important direct implications of Wagner's conjecture like theorem 8.2 and theorem 8.3, the ideas useful in understanding and proving Wagner's conjecture have found uses in other areas of graph theory. For example, tree-widths and related notions of path-widths and branch-widths as well as the disjoint paths algorithm have many algorithmic applications. The enormous theoretical impact and the many useful algorithmic applications of this line of research shows that this has been a very important area of graph theory in recent years.

Many things have been omitted from this survey, the most glaring of which is an in depth discussion on the proof of Wagner's conjecture. I felt that a more detailed description of the disjoint paths algorithm is also justified but the algorithm is very complicated and beyond the scope (and length) of this survey. The reader is referred to [19]. Completely missing are many other excluded minor theorems out there which make up much of graph structure theory. For a survey of these theorems the reader is referred to Thomas [20].

References

[1] D. Archdeacon. *A Kurotowski theorem for the projective plane.* PhD thesis, Ohio State University, 1980.

[2] D. Archdeacon and J. P. Huneke. A Kurotowski theorem for nonorientable surfaces. *Journal of Combinatorial Theory B*, 46:173–231, 1989.

[3] S. Arnborg and A. Proskurowski. Linear time algorithms for NP-hard problems on graphs embedded in k-trees. *Discrete Applied Mathematics*, 23:11–24, 1989.

[4] R. Bodendiek and K. Wagner. Solution to König's graph embedding problem. *Mathematische Nachrichten*, 140:251–272, 1989.

[5] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions. *SIAM Journal on Computing*, 25(6):1305–1317, 1996.

[6] H. H. Glover, J. P. Huneke, and C. S. Wang. 103 graphs that are irreducible for the projective plane. *Journal of Combinatorial Theory B*, 27:332–370, 1979.

[7] G. Higman. Ordering by divisibility in abstract algebras. *Proceedings of the London Mathematical Society*, 2:326–336, 1952.

[8] R. M. Karp. On the complexity of combinatorial problems. *Networks*, 5:45–68, 1975.

[9] D. König. Theorie des endlichen und unendlichen graphen. Technical report, Akademische Verlagsgesellschaft M.B.H. Leipzig, 1936. English translation published by Birkhäuser 1990.

[10] J. B. Kruskal. Well-quasi-ordering, the tree theorem, and vázsonyi's conjecture. *Transactions of the American Mathematical Society*, 95:210–225, 1960.

[11] K. Kuratowski. Sur le problème des courbes gauches en topologie. *Fund. Math*, 15:271–282, 1930.

[12] B. Mohar. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM Journal on Discrete Mathematics*, 12(1):6–26, 1999.

[13] C. St. J. A. Nash-Williams. On well-quasi-ordering finite trees. *Proceedings of the Cambridge Philosophical Society*, 59:833–835, 1963.

[14] N. Robertson and P. D. Seymour. Graph minors XVI : Excluding a non-planar graph. Submitted.

[15] N. Robertson and P. D. Seymour. Graph minors XX : Wagner's conjecture. Manuscript.

[16] N. Robertson and P. D. Seymour. Graph minors II : algorithmic aspects of tree-width. *Journal of Algorithms*, 7:309–322, 1986.

[17] N. Robertson and P. D. Seymour. Graph minors V : excluding a planar graph. *Journal of Combinatorial Theory B*, 41:92–114, 1986.

[18] N. Robertson and P. D. Seymour. Graph minors IV : tree-width and well-quasi-ordering. *Journal of Combinatorial Theory B*, 48:227–254, 1990.

[19] N. Robertson and P. D. Seymour. Graph minors XIII : the disjoint paths problem. *Journal of Combinatorial Theory B*, 63:65–110, 1995.

[20] R. Thomas. Recent excluded minor theorems. A survey for the British Combinatorial Conference.

[21] R. Thomas. A Menger-like property of tree-width : the finite case. *Journal of Combinatorial Theory B*, 48:67–76, 1990.

[22] K. Wagner. Über eine eigenschaft der ebenen komplexe. *Mathematische Annalen*, 114:570–590, 1937.

A The graph minors series

Robertson and Seymour authored a long series of papers dealing with the various aspects of Wagner's conjecture. We list here the papers in the series known to the author.

1. N. Robertson and P. D. Seymour. Graph minors I : excluding a forest. *Journal of Combinatorial Theory B*, 35:39-61, 1983.
2. N. Robertson and P. D. Seymour. Graph minors II : algorithmic aspects of tree-width. *Journal of Algorithms*, 7:309-322, 1986.
3. N. Robertson and P. D. Seymour. Graph minors III : planar tree-width. *Journal of Combinatorial Theory B*, 36:49-64, 1984.
4. N. Robertson and P. D. Seymour. Graph minors IV : tree-width and well-quasi-ordering. *Journal of Combinatorial Theory B*, 48:227-254, 1990.
5. N. Robertson and P. D. Seymour. Graph minors V : excluding a planar graph. *Journal of Combinatorial Theory B*, 41:92-114, 1986.
6. N. Robertson and P. D. Seymour. Graph minors VI : disjoint paths across a disc. *Journal of Combinatorial Theory B*, 41:115-138, 1986.
7. N. Robertson and P. D. Seymour. Graph minors VII : disjoint paths on a surface. *Journal of Combinatorial Theory B*, 45:212-254, 1988.
8. N. Robertson and P. D. Seymour. Graph minors VIII : a Kuratowski theorem for general surfaces. *Journal of Combinatorial Theory B*, 48:255-288, 1990.
9. N. Robertson and P. D. Seymour. Graph minors IX : disjoint crossing paths. *Journal of Combinatorial Theory B*, 49:40-77, 1990.
10. N. Robertson and P. D. Seymour. Graph minors X : obstructions to tree-decomposition. *Journal of Combinatorial Theory B*, 52:153-190, 1991.
11. N. Robertson and P. D. Seymour. Graph minors XI : circuits on a surface. *Journal of Combinatorial Theory B*, 60:72-106, 1994.
12. N. Robertson and P. D. Seymour. Graph minors XII : distance on a surface. *Journal of Combinatorial Theory B*, 64:240-272, 1995.
13. N. Robertson and P. D. Seymour. Graph minors XIII : the disjoint paths problem. *Journal of Combinatorial Theory B*, 63:65-110, 1995.
14. N. Robertson and P. D. Seymour. Graph minors XIV : extending an embedding. *Journal of Combinatorial Theory B*, 65:23-50, 1995.
15. N. Robertson and P. D. Seymour. Graph minors XV : giant steps. *Journal of Combinatorial Theory B*, 68:112-148, 1996.
16. N. Robertson and P. D. Seymour. Graph minors XVI : excluding a non-planar graph. Submitted.
17. N. Robertson and P. D. Seymour. Graph minors XVII : taming a vortex. Submitted.
18. N. Robertson and P. D. Seymour. Graph minors XX : Wagner's conjecture. Manuscript.
19. N. Robertson and P. D. Seymour. Graph minors XXII : irrelevant vertices in linkage problems. In preparation.