GRAPH MINORS — A SURVEY

Neil Robertson*

Dept. of Mathematics, Ohio State University,
231 West 18th Ave., Columbus, Ohio 43210, U.S.A.

P. D. Seymour**

Bell Communications Research, Inc.,
600 Mountain Ave., Murray Hill, New Jersey 07974, U.S.A.

Abstract. We survey a number of results about minors of graphs which we have recently obtained. They are basically of three types:

(i) results concerning the structure of the graphs with no minor isomorphic to a fixed graph

(ii) results concerning a conjecture of K. Wagner, that for any infinite set of graphs one of its members is isomorphic to a minor of another, and

(iii) algorithmic results concerning the DISJOINT CONNECTING PATHS problem.

1. INTRODUCTION

There are two fundamental questions which motivate the work we report on here.

(A) *(K. Wagner's well-quasi-ordering conjecture). Is it true that for every infinite sequence $G_1, G_2, \ldots$ of graphs, there exist i, j with $i < j$ such that G_i is isomorphic to a minor of G_j?*

[Graphs in this paper are finite and may have loops or multiple edges. A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges.]

(B) *(The DISJOINT CONNECTING PATHS problem). If $k \geqslant 0$, is there a polynomially-bounded algorithm to decide, given a graph G and vertices $s_1, \ldots, s_k, t_1, \ldots, t_k$ of G, whether there are k mutually disjoint paths $P_1, \ldots, P_k$ of G where P_i has ends s_i, t_i $(1 \leqslant i \leqslant k)$?*

[Two paths are *disjoint* if they have no common vertices.] Some of the background to these questions is discussed in sections 2 and 3. We have obtained a number of results related to (A) and (B) (for example, both have affirmative answers if the graphs concerned are constrained to lie on any fixed surface), and in this paper we survey them.

It is an awkward time to write a survey paper such as this, because we believe that we have completely answered both (A) and (B) affirmatively, but the proofs have not yet been carefully written out and checked. We hope that this process will be completed by the time of the conference, and we will then be able to confirm or deny the tentative hopes that we proffer here.

---

* Part of this work was performed under a consulting agreement with Bell Communications Research, Inc.

** Research partially supported by NSF grants number MCS8103440 and MCS8302266.

Towards the end of this paper we sketch an outline of how these proofs should go. These proofs will in any case be reductions to our earlier theorems, and so even if they are correct they do not make the remainder of the theorems we mention in this paper redundant.

All our results are consequences of "structure theorems," that sufficiently general graphs contain minors or other sub-objects that are useful to us — or equivalently, that graphs that do not contain a useful minor have a certain restricted structure. Proving such a theorem is usually the most difficult step in proving a special case of (A) or (B), and most of this paper is devoted to structure theorems.

Incidentally, we have found that whenever we get a better structure theorem, it enables us to prove both another case of (A) *and* another case of (B). But we know of no theoretical connection between (A) and (B). Is it just coincidence that the proofs of their special cases are so closely tied?

Rather than simply present a mass of results which we have accumulated over the past three years, we have tried to indicate the order in which we found these theorems, and why we moved from each topic to the next. We hope that this gives the reader a little more insight into the tangled interrelationships that will emerge.

The paper is divided into three parts. The first three sections are introductory, and discuss the background to (A) and (B). Sections 4, 5, and 6 contain structure theorems, while sections 7 and 8 concern the application of these theorems to (A) and (B).

2. WELL-QUASI-ORDERING

The Kuratowski-Wagner theorem, that K_5 and $K_{3,3}$ are the only graphs which are minor-minimal non-planar, is probably the most basic and undoubtably the most well-known result about graph minors. Nevertheless, the corresponding problem for higher surfaces is extremely difficult, and remained completely open until 1980, when Archdeacon [1] and Glover, Huneke and Wang [7] solved the case of the projective plane. They proved that there are 35 minor-minimal "non-projective-planar" graphs.

It is possible to show that for higher genus the length of the corresponding list grows quite rapidly — Huneke claims that the list for the torus has over 800 members, for example — and it is not clear even that it must be finite. This finiteness question was raised by Erdös in the 1930's, but again remained completely open until recently. Archdeacon and Huneke [2] settled half the problem in 1981 — they proved that the list for any non-orientable surface is finite. We have shown (in 1983) that it is finite for any surface, as we shall explain later.

The list of minor-minimal graphs which cannot be drawn in some given surface is an *antichain*, that is, no member of it is isomorphic to a minor of another. In the early 1960's, Wagner proposed the conjecture that *every antichain is finite*. This evidently would settle Erdös' question — indeed, it would imply that for any property of graphs, the list of minor-minimal graphs without the property is finite. Wagner's conjecture has remained open until the present.

A *quasi-order* is a class with a reflexive transitive relation. For example, the relation "H is isomorphic to a minor of G" defines a quasi-order on the class of all graphs. A quasi-order $(Q, \leqslant)$ is a *well-quasi-order* (wqo) if for every countable sequence $q_1, q_2, \ldots$ of members of Q there exist $j > i \geqslant 1$ such that $q_i \leqslant q_j$. An equivalent requirement is that

(i) $(Q, \leqslant)$ has no infinite descending chain, and

(ii) $(Q, \leqslant)$ has no infinite antichain.

An equivalent form of Wagner's conjecture is thus that the "minor" quasi-order of the class of all graphs is a wqo, for condition (i) is trivially satisfied in this case. (We recall that all graphs considered here are finite.)

There are other wqo problems with graphs. The "subgraph" quasi-order is not a wqo, as is easily seen — a set of circuit graphs, one of each size, is an infinite antichain. (However, it is a wqo for graphs with no k-vertex path, for any fixed integer $k \geqslant 0$, as follows easily from (4.4).) A more plausible candidate for a wqo is "topological containment." We say that G *topologically contains* H if G has a subgraph which is isomorphic to a subdivision of H. (A *subdivision* of H is a graph obtained from H by replacing its edges by internally vertex-disjoint paths.) It is not a wqo in general, however — the set of graphs formed by taking a circuit graph of each size, and replacing each edge by two parallel edges, is an infinite antichain. Vázsonyi (see [11]) made two well-known conjectures:

(i) that the class of all trees is a wqo under topological containment, and

(ii) that the class of all graphs with maximum valency $\leqslant 3$ is a wqo under topological containment.

Kruskal [11] proved Vázsonyi's first conjecture, but the second has remained open. It is, however, a special case of Wagner's conjecture, for if H has maximum valency $\leqslant 3$ and G has a minor isomorphic to H then G topologically contains H.

A fourth interesting quasi-order of the class of all graphs is that from "immersion." Let us say that a *pure immersion* of a graph H is a graph obtained from H by replacing its edges by paths, where these paths are mutually edge-disjoint but may have vertices in common. We say that H is *immersed* in G if there is a subgraph of G which is isomorphic to a pure immersion of H. Nash-Williams [14] conjectures that this provides a wqo of the class of all graphs. It can be shown that this conjecture implies both of Vázsonyi's conjectures, but it seems not to be related to Wagner's conjecture.

These problems have extensions to infinite graphs, but we avoid that topic because we have nothing to say about it. We also have no new results about immersion. It is possible that "structure theorem" approach to Nash-Williams' problem might be useful, but we have not explored this.

Our approach to Wagner's conjecture is the following. Let Σ be a "structure" of graphs, for instance, being planar, or having genus $\leqslant k$, or being divisible into small pieces by small cutsets. Let us identify Σ with the class of graphs possessing this structure. Now suppose that we wish to show that Σ is well-quasi-ordered by minors, and suppose that we can prove a structure theorem within Σ of the following kind.

For every $H \in \Sigma$ there is a structure $\Sigma'(H)$ such that $G \in \Sigma'(H)$ for every graph $G \in \Sigma$ with no minor isomorphic to H.

Then it suffices to prove that for each $H \in \Sigma$, $\Sigma'(H)$ is well-quasi-ordered by minors. For if $\mathscr{C}$ is an infinite antichain with $\mathscr{C} \subseteq \Sigma$, choose $H \in \mathscr{C}$, and then $\mathscr{C}-\{H\}$ is an infinite antichain in $\Sigma'(H)$.

Let us give an example. Suppose that we wish to prove that the class of all trees is well-quasi-ordered by minors. It is convenient to work with *rooted* trees, trees with one vertex distinguished, called the *root*. We say a rooted tree is a minor of another if it can be obtained from a subgraph of the second by contraction, so that root is taken to root in the natural way. We define the *A-type* of a rooted tree as follows. A path, rooted at its end, has A-type zero; and inductively, for $i \geqslant 1$ we say that T has A-type $\leqslant i$ if there is a path P of T, with one end at the root of T, such that all its "pendant" trees (rooted at their vertex of attachment on P) have A-type $\leqslant i-1$.

Equivalently, let us take a straight-line drawing of T in the plane, and count the maximum number of turns in a path beginning at the root. This quantity, minimized over all drawings, is the A-type of T.

An easy example of a structure theorem is the following.

2.1 *Let T be a rooted tree. Then there is a number k such that every rooted tree with no minor isomorphic to G has A-type $\leqslant k$.*

Hence to prove that the class of all rooted trees is wqo by minors, it suffices to show that for every k the class of all rooted trees with A-type $\leqslant k$ is wqo by minors — and this can easily be done by induction on k, using a result of Higman that we discuss in section 7.

3. DISJOINT PATHS

The problem DISJOINT CONNECTING PATHS (DCP) is, given a graph G and vertices $s_1, ..., s_k, t_1, ..., t_k$ of G, to decide if there are k mutually disjoint paths $P_1, ..., P_k$ of G where P_i has ends s_i, t_i $(1 \leqslant i \leqslant k)$. This is superficially similar to the maximum flow problem solved by Menger's theorem, that of, given a graph G and subsets S, T of $V(G)$, deciding if there are k mutually disjoint paths between S and T. However, in DCP we are prescribing which vertex is to be paired with which, and this makes the problem much more complicated.

Indeed DCP (with k part of the input) was one of Karp's original NP-complete problems [10], and Lynch [12] showed that it remains NP-complete if G is constrained to be planar. For fixed k however it may be polynomially solvable. (An analogous situation is presented by the problem "does G have a clique of size k?" This is NP-complete, but for any fixed k it is polynomially solvable. For $k = 100$, for example, it is solvable in time $|V(G)|^{100}$.)

It is important that we are concerned with undirected graphs here. While the maximum flow problem behaves much the same for directed and for undirected graphs, DCP does not. Fortune, Hopcroft and Wyllie [6] showed that for directed graphs, DCP is NP-complete even with $k = 2$ (for directed graphs, we ask for the existence of directed paths between s_i and t_i $(i = 1, 2)$).

For $k = 2$ and undirected graphs, the problem is known to be polynomially solvable [23,24,25,26]. The algorithm is so simple that we describe it. We are given a graph G and four distinct vertices s_1, t_1, s_2, t_2 of G.

Step k (for $k = 1, 2, 3, 4$). Test if there is a separation (G_1, G_2) of G with $V(G_2) \not\subseteq V(G_1)$, $s_1, t_1, s_2, t_2 \in V(G_1)$ and $|V(G_1 \cap G_2)| = k - 1$. [$(G_1, G_2)$ is a *separation* of G if G_1, G_2 are subgraphs of G, with $G_1 \cup G_2 = G$ and $E(G_1 \cap G_2) = \emptyset$.] If we find such a separation, we construct G_1' from G_1 by adding new edges joining all pairs of vertices in $V(G_1 \cap G_2)$, set $G = G_1'$, and return to step k. If we do not find any such separation we go to step $k + 1$.

Step 5. We test if G can be drawn without crossings in the plane, with s_1, s_2, t_1, t_2 on the outside in that order. If so then the paths do not exist, and otherwise they do exist.

We see that the "obstruction" to the existence of the desired paths is a topological one; and it has long been our hope, and we now seem to have shown, that the same is true for the DCP problem with general k. We expand on this in section 8. There we shall see that the general (fixed k) DCP problem can be solved if and only if it can be solved for graphs which can "almost" be embedded in surfaces with genus bounded by a function of k. ("Almost" here is technical, and will be elaborated in section 6.) This is one motivation for our study of DCP for graphs which can be drawn in a fixed surface. Another is that for the DCP problem, graphs drawn on a surface are naively easier to handle than general graphs, because paths cannot cross each other.

We shall study DCP for graphs with some structure imposed. There are basically two kinds of structure involved, bounded genus (as we discussed above) and "tree structure." This last means that the graph is constructed from inherently simpler graphs by piecing them together in a tree structure, where each piece overlaps its neighbours in a subset of bounded size. This guarantees us a multiplicity of small cutsets, and "divide and conquer" methods usually serve to reduce a DCP problem on such a graph to DCP problems on the inherently simpler pieces.

There is another algorithmic question of interest here. Given a fixed graph H, how does one test if a general graph has a minor isomorphic to H? It is easy to see that if DCP is polynomially solvable for fixed k, then so is this problem. In some sense it is strictly easier — for we have often found structures within which we could solve the above minor-testing problem, and yet could not (prima facie) solve DCP; and for many types of graphs H (for instances, planar graphs) we can test if a general graph has a minor isomorphic to H in polynomial time, while the polynomial solvability of DCP in general is still not quite decided. Thus it is perhaps surprising that our proposed solution to DCP in general is based mainly on results about this minor-testing problem. This is one motivation for our separate study of the minor-testing problem; it is not just a simplification of DCP but may be the key to the entire solution.

Finally, a confession: although our algorithms are polynomially bounded, these are immensely high-order polynomials in most cases, and most of the algorithms are not practical. The algorithmic results we obtain are of interest mainly from the point of view of NP-completeness.

4. STRUCTURE THEOREMS

We shall return to these motivating problems later. For the moment we concentrate on structure theorems. The starting point for the project was Mader's use of a theorem of Erdös and Pósa. Erdös and Pósa [5] proved the following.

4.1 *For any number $k \geqslant 0$ there is a number $k' \geqslant 0$ such that for every graph G, either*

(i) *G has k vertex-disjoint circuits, or*

(ii) *there exists $X \subseteq V(G)$ with $|X| < k'$ such that $G \backslash X$ has no circuits.*

[$G \backslash X$ denotes the graph obtained from G by deleting the vertices in X.] Erdös and Pósa also showed that the best possible k' is of order $k \log k$. However, these numerical results do not concern us. Condition (ii) is not necessary and sufficient for (i) to be false, but the theorem is sharp in another sense.

Let us say a class of graphs $\mathscr{F}$ is *minor-closed* if for all $G \in \mathscr{F}$, every graph isomorphic to a minor of G is also in $\mathscr{F}$. Throughout the paper, $\mathscr{F}$ will denote an arbitrary minor-closed class of graphs, and we shall omit to say so henceforth. We can reformulate (4.1) as follows.

4.2 *For any $\mathscr{F}$, the following are equivalent:*

(i) *there exists $k \geqslant 0$ such that the graph consisting of k disjoint loops is not in $\mathscr{F}$*

(ii) *there exists $k' \geqslant 0$ such that for every $G \in \mathscr{F}$, $G \backslash X$ has no circuits for some $X \subseteq V(G)$ with $|X| < k'$.*

Mader [13] used this theorem to deduce that the class of all graphs without k vertex-disjoint circuits is well-quasi-ordered by topological containment; but we explain Mader's argument in section 7. We have found several other results which are sharp in the same way as (4.2), and we list some of them in this section.

First, an obvious analogue of (4.2) is the trivial theorem that if G has no $k+1$ matching, we can meet all its edges with $2k$ vertices. In our standard form this becomes

4.3 *For any $\mathscr{F}$, the following are equivalent:*

(i) *there exists $k \geqslant 0$ such that the graph consisting of k disjoint (non-loop) edges is not in $\mathscr{F}$*

(ii) *there exists $k' \geqslant 0$ such that for every $G \in \mathscr{F}$, $G \backslash X$ has no edges for some $X \subseteq V(G)$ with $|X| < k'$.*

We define *B-type* as follows. Graphs with $\leqslant 1$ vertex have B-type 1, and inductively for $i > 1$, G has B-type $\leqslant i$ if for some $v \in V(G)$, every component of $G \backslash \{v\}$ has B-type $\leqslant i-1$. It is rather easy to prove that a graph with no k-vertex path has B-type $< k$, and conversely that a graph with B-type $< k$ has no 2^k-vertex path. In our standard form, this becomes

4.4 *For any $\mathscr{F}$, the following are equivalent:*

(i) *there exists $k \geqslant 0$ such that $P_k \notin \mathscr{F}$*

(ii) *there exists $k' \geqslant 0$ such that every graph in $\mathscr{F}$ has B-type $< k'$.*

[P_k is the k-vertex path.] It is convenient to abbreviate this as "k-vertex path $\leftrightarrow$ B-type."

Dirac [4] showed that every 2-connected graph with a long path also has a long circuit. Let us define the *C-type* of a graph to be the maximum of the B-types of its blocks. From Dirac's theorem we deduce

4.5 *k-vertex circuit $\leftrightarrow$ C-type.*

One would expect the "dual" problem (in the sense of planar duality), the exclusion of large "bonds," to be of the same degree of difficulty. But in fact, it was rather more complicated. The appropriate structure (*D-type*) is the following. The complete graph K_2 has D-type 1. Inductively, for $i \geqslant 2$, we say a 2-connected graph has D-type $\leqslant i$ if it can be constructed by arranging in a circle a number of 2-connected graphs, each with D-type $\leqslant i-1$, and identifying a vertex of each with a vertex of the next. For a general graph, we say its D-type is the maximum of the D-types of its blocks. Then

4.6 *$C_k^{\bullet} \leftrightarrow D$-type.*

[$C_k^{\bullet}$ is the loopless graph with 2 vertices and k edges.]

The next simplest type of graph we could think of to exclude was the k-edge star $K_{1,k}$. That is rather easy. We define

(i) if G is connected and has two adjacent vertices with valency 2, the E-type of G is the E-type of the graph obtained by contracting the edge joining these two vertices

(ii) if G is connected and has no such pair of vertices its E-type is $|V(G)|$

(iii) the E-type of a general graph is the maximum E-type of its components.

Then we have

4.7 *$K_{1,k} \leftrightarrow E$-type.*

Theorems (4.3), (4.4) and (4.7) allow us to deduce that if H is a matching, path or star then no infinite antichain contains H and there is a polynomial algorithm to test for H as a minor. It was natural to ask for a unification. This is provided by the concept of "path-width". Suppose that G can be constructed from a sequence $G_1, \ldots, G_n$ of vertex-disjoint graphs, by identifying some vertices of G_i with some of G_{i+1} $(1 \leqslant i < n)$; and that each G_i has at most $k + 1$ vertices. Then we say that G has *path-width* $\leqslant k$. The following was the first difficult theorem that we found [15].

4.8 *Binary tree of depth $k \leftrightarrow$ path-width.*

Now every forest is isomorphic to a minor of a sufficiently large binary tree; and so we deduce

4.9 *For any $\mathcal{F}$, the following are equivalent:*

(i) *some forest is not in $\mathcal{F}$*

(ii) *there exists $k' \geqslant 0$ such that every member of $\mathcal{F}$ has path-width $< k'$.*

From this we could deduce, as hoped, that for any forest H, no infinite antichain contains H, and there is a polynomial algorithm to test for H as a minor.

At this stage we switched from studying these theorems for their own sake to studying them for the sake of their applications. The object now became to find the most general type of graph H the exclusion of which would force some usable structure. We hit on a conjecture for the structure corresponding to the exclusion of a general planar graph. It remained open for some eighteen months, but was eventually proved in [19].

We need the notion of "tree-width." For path-width the "pieces" $G_1, \ldots, G_n$ are arranged in a sequence. Let us arrange them instead in a tree. Thus we have a tree T, and associated with each $t \in V(T)$ we have a graph G_t. Now suppose that G can be constructed by identifying (for each edge $\{t,t'\}$ of T) some vertices of G_t with some of $G_{t'}$. Suppose also that each G_t has at most $k + 1$ vertices. Then we say that G has *tree-width* $\leqslant k$. We proved

4.10 *$k \times k$ grid $\leftrightarrow$ tree-width.*

[The $k \times k$ *grid* is a square induced subgraph of the infinite square lattice, with k^2 vertices.] Every planar graph is isomorphic to a minor of a sufficiently large grid, and we deduce

4.11 *For any $\mathscr{F}$, the following are equivalent:*

(i) *some planar graph is not in $\mathscr{F}$*

(ii) *there exists $k' \geq 0$ such that every member of $\mathscr{F}$ has tree-width $< k'$.*

This had the desired two applications — for any planar graph H, no infinite antichain contains H, and there is a polynomial algorithm to test for H as a minor. It also had several others, one of which provides a dramatic generalization of the Erdös-Pósa theorem (4.1), as follows. For any graph H, a minimal subgraph of G with a minor isomorphic to H is called an *H-expansion* in G. Thus if H is a loop, the H-expansions in G are the circuits of G.

4.12 [19] *Let H be a planar graph. For any number $k \geq 0$ there is a number $k' \geq 0$ such that for every graph G, either*

(i) *G has k vertex-disjoint H-expansions, or*

(ii) *there exists $X \subseteq V(G)$ with $|X| < k'$ such that $G \backslash X$ has no H-expansions.*

This is best possible in that for any non-planar graph H, if we take $k = 2$ then no k' satisfies the theorem.

5. CLIQUE-SUMS

Before we continue this saga, we need to define "clique-sums." Let (G_1, G_2) be a separation of G, and let H_i be the graph obtained from G_i by adding new edges joining every pair of vertices in $V(G_1 \cap G_2)$ $(i = 1, 2)$. We say that G is the *clique-sum* of H_1 and H_2, and if $|V(G_1 \cap G_2)| \leq k$ we also say that it is the $(\leq k)$ - *sum* of H_1 and H_2. If G can be constructed by repeatedly taking $(\leq k)$-sums starting from graphs isomorphic to members of some class $\mathscr{C}$ of graphs, we write $G \in \langle \mathscr{C} \rangle_k$. We set

$$\langle \mathscr{C} \rangle = \bigcup_{k \geq 0} \langle \mathscr{C} \rangle_k .$$

It can be shown, for instance, that

(i) if $\mathscr{C} = \{K_1, K_2\}$, then $\langle \mathscr{C} \rangle$ is the class of all forests

(ii) if $\mathscr{C}$ is the class of all graphs with at most three vertices then $\langle \mathscr{C} \rangle$ is the class of "series-parallel" graphs

(iii) if $\mathscr{C}$ is the class of all graphs with at most $k+1$ vertices then $\langle \mathscr{C} \rangle$ is the class of all graphs with tree-width $\leqslant k$.

Clique-sums seem to be intimately connected with excluded minors. The classes $\langle \mathscr{C} \rangle$ of (i) and (ii) above are the classes of graphs with no minor isomorphic to a loop or to K_4, respectively. (See [3].) There are more difficult theorems of the same type:

(iv) [8] if $\mathscr{C}$ is the class of all planar graphs together with K_5, then $\langle \mathscr{C} \rangle_2$ is the class of all graphs with no $K_{3,3}$ minor

(v) [27] if $\mathscr{C}$ is the class of all planar graphs together with the "four-rung Möbius ladder", then $\langle \mathscr{C} \rangle_3$ is the class of all graphs with no K_5 minor.

These show that (4.10) is not the most general theorem in which clique-sums are significant. Motivated by (iv) and (v) and our desire to find structure theorems for the exclusion of the most general graphs possible, we derived from (4.10) the following, which in a way generalizes (iv) and (v) above.

5.1 *For any $\mathscr{F}$ the following are equivalent:*

(i) *some graph with crossing number $\leqslant 1$ is not in $\mathscr{F}$*

(ii) *there exists $k' \geqslant 0$ such that if $\mathscr{C}_1$ denotes the class of all planar graphs and $\mathscr{C}_2$ denotes the class of all graphs with tree-width at most k', then*

$$\mathscr{F} \subseteq \langle \mathscr{C}_1 \cup \mathscr{C}_2 \rangle_3 .$$

6. MINORS AND SURFACES

(5.1) is a fairly easy extension of (4.10), and our program of producing ever more general structure theorems languished at this point for some time. Even obtaining a form of (5.1) for graphs with crossing number $\leqslant 2$ seemed too difficult to be worth the effort.

However, before proving (4.10) we had proved a special case of it where $\mathscr{F}$ is restricted to be a class of *planar* graphs, and for that the proof was easy [17]. We decided therefore to look for similar results where the members of $\mathscr{F}$ were constrained to lie on some fixed surface Σ (briefly, "$\mathscr{F}$ embeds in Σ"). A second motivation for considering what might seem an unnatural special case was Erdös' old question about the finiteness of the list of minimal graphs which cannot be drawn in a fixed surface. It is easy to see that all members of such a list have bounded genus; and so Erdös' question would be answered affirmatively if it could be shown that there is no infinite antichain all members of which can be drawn in a fixed surface. This last we hoped would follow from a structure theorem of the kind we were considering, as indeed it did.

The structure theorem we found is, surprisingly, simpler for higher surfaces than the corresponding theorem is for the sphere. Let G be a graph with a drawing Γ in a surface Σ. We define the *representativeness* of Γ to be the minimum, over all non-null-homotopic non-self-intersecting closed paths P in Σ, of the number of points of P used by Γ. Our theorem is the following.

6.1 *Let Σ be a connected surface different from the sphere. Then for any $\mathcal{F}$ which embeds in Σ, the following are equivalent:*

(i) *some graph which can be drawn in Σ is not in $\mathcal{F}$*

(ii) *there exists $k' \geqslant 0$ such that every drawing in Σ of every member of $\mathcal{F}$ has representativeness $< k'$.*

Although this has quite a simple form, we needed some elaborate machinery for its proof [21]. With the aid of this machinery we returned to the general problem, and recently were able to prove what we regard as the "ultimate" structure theorem. To state it we need some further definitions.

Let Σ be a surface with boundary (that is, it has some finite number of holes punched in it). Let us draw a graph G in Σ, allowing vertices to be drawn on the boundary of Σ. For each vertex v of G drawn on the boundary of Σ, let us take a new graph G_v with $V(G) \cap V(G_v) = \{v\}$, where the G_v's for different vertices v are vertex-disjoint. Now for all consecutive pairs of vertices v, v' on the boundary of Σ, let us identify some vertices of G_v with some of $G_{v'}$. Let G' be the graph produced from combining G and the G_v's in this way. If each G_v has at most $k+1$ vertices, we say that G' is (Σ, k) - *drawable*. Secondly, if $G \backslash X$ is (Σ, k)-drawable, where $X \subseteq V(G)$ and $|X| \leqslant w$, we say that G is (Σ, k, w)-*drawable*. Our "ultimate" structure theorem is the following. Half of its proof is in [23], while the other half is being written.

6.2 *For any $\mathcal{F}$, the following are equivalent:*

(i) *some graph is not in $\mathcal{F}$*

(ii) *there exists Σ, k, w such that $\mathcal{F} \subseteq \langle \mathcal{C} \rangle$, where $\mathcal{C}$ is the class of (Σ, k, w)-drawable graphs.*

As we mentioned in the introduction, we believe that this settles both problems (A) and (B) stated there; but the details of these applications of (6.2) still have to be checked.

7. APPLICATIONS TO WELL-QUASI-ORDERING

We turn now to the derivation from our structure theorems of results related to Wagner's conjecture.

Let $(Q, \leqslant)$ be a wqo, and let σ, σ' be finite sequences of members of Q. We say $\sigma \leqslant \sigma'$ if σ' has a subsequence σ'' of the same length as σ, such that σ'' dominates σ "componentwise." It is a fundamental theorem of Higman [9] that this defines a wqo on the set of finite sequences over Q.

Let us see how this can be used to show that the class of graphs with path-width $\leqslant w$ is wqo by minors. Now if G has path-width $\leqslant w$ then it is constructed from a sequence $G_1, ..., G_n$ of graphs of bounded size, by overlapping each with the next. A first attempt at a proof might then be to encode G as this finite sequence $\sigma(G)$ say, recording on each member of the sequence how it is to overlap its neighbours. Certainly the encodings $\sigma(G)$ lie in a wqo set, because of Higman's theorem and the fact the graphs of bounded size are wqo by the "subgraph" quasi-order, even when their vertices are labelled. The difficulty lies in the fact that it is possible that $\sigma(G) \leqslant \sigma(G')$ even if G is not isomorphic to a minor of G'; for if $\sigma(G)$ and $\sigma(G')$ have different lengths, and $\sigma(G')$ has a redundant internal entry, we have no satisfactory way of "contracting out" the corresponding part of G'.

However, suppose instead that each G_i has two subsets A_i and B_i, both of cardinality k, and there are k mutually disjoint paths of G between A_i and B_i; and that G is constructed from $G_1, ..., G_n$ by identifying A_i with B_{i+1} $(1 \leqslant i < n)$. We say that such a graph G is *k-uniform* over any class of graphs containing $G_1, ..., G_n$. Now suppose G, G' are both k-uniform over some class, and we encode them as before as $\sigma(G), \sigma(G')$. Now it *is* true that $\sigma(G) \leqslant \sigma(G')$ implies that G is isomorphic to a minor of G', for we can remove any unwanted section of G' by contracting the edges of the k paths which run through it. (We need a little care to ensure that the permutation induced by these paths is consistent with the way we record on each G_i the sets A_i and B_i in the encoding process, but that is easy.)

Now it is not really necessary for the above argument that the G_i's be of bounded size; it is enough that we know a suitable well-quasi-ordering result for the triples (G_i, A_i, B_i). Let us combine this with a theorem that every graph of path-width $\leqslant w$ is 0-uniform over the class of 1-uniform graphs over the class of 2-uniform graphs over ... over the class of w-uniform graphs over the class of graphs with at most $w + 1$ vertices. Then, by applying our result $w + 1$ times, we deduce that the class of graphs with path-width $\leqslant w$ is wqo by minors.

We deduce that no infinite antichain contains a forest; for if it did, then its other members would constitute an infinite antichain with all members of bounded path-width, by (4.9), contrary to what we have just proved.

As for tree-width, basically the same approach works. We may think of Higman's theorem as a result about paths with vertices labelled from some wqo. Kruskal [11] proved the generalization to trees with vertices labelled from a wqo. We hope then to encode a graph of bounded tree-width as a labelled tree, where the labels are better-connected graphs with tree-width no larger, and to apply Kruskal's theorem repeatedly. There is a serious difficulty, however. With path-width, only two subsets of G_i (A_i and B_i) needed to be encoded. With tree-width this number is unbounded, and the analogous encoding procedure does not work. To overcome this, we needed to work with hypergraphs, labelling edges of the hypergraph rather than subsets of vertices of the graphs. It is complicated technically, and we do not discuss it further here. This is the theory of "patchworks" of [18].

We deduce from this and (4.11) that no infinite antichain contains a planar graph.

The second technique we want to mention is that used by Mader [13] to deduce from the Erdös-Pósa theorem that the class of graphs without k vertex-disjoint circuits is wqo by topological containment. Let G be a graph without k disjoint circuits. We define an encoding of G as follows. Choose k' as in (4.1), and choose $X \subseteq V(G)$ with $|X| < k'$ such that $G\backslash X$ has no circuits. Take an arbitrary numbering $x_1, ..., x_r$ of the elements of X, and for each $v \in V(G) - X$, define $n(v) = (n_1, ..., n_r)$, where n_i is the number of edges between x_i and v. The set of all such finite sequences, ordered by Higman's order, is a wqo. Let G^+ be $G\backslash X$ with all vertices v labelled with $n(v)$, and let $\sigma(G)$ be (G^+, G^-) where G^- is the restriction of G to X. It is easy to see that if G, G' are two such graphs, and $\sigma(G) \leqslant \sigma(G')$ with the natural ordering, then G is topologically contained in G'. Moreover, the natural ordering of all such pairs (G^+, G^-) is a wqo by Kruskal's theorem; and Mader's theorem follows.

More generally let $\mathscr{C}$ be a class of graphs which forms a well-quasi-order under "labelled" topological containment, when its members are vertex-labelled from any wqo (such as the class of all forests, by Kruskal's theorem). For any integer k', the class of all graphs G such that $G\backslash X \in \mathscr{C}$ for some $X \subseteq V(G)$ with $|X| < k'$ is wqo by topological containment. (The proof is the same.) An analogous result holds for minor containment.

A third technique worth mentioning is our method of proving that the class of all graphs which can be drawn on a fixed surface is wqo by minors [22]. That runs as follows. Suppose the result is false, and choose a surface Σ as simple as possible which embeds an infinite antichain $\mathscr{A}$ say. Now Σ is not the sphere, because no infinite antichain contains even one planar graph. Choose $H \in \mathscr{A}$. Then no $G \in \mathscr{A} - \{H\}$ has a minor isomorphic to H, and so by (6.1), there exists $k' \geqslant 0$ such that every member of $\mathscr{A} - \{H\}$ has a drawing in Σ with representativeness $< k'$.

Now there are only finitely many different non-self-intersecting closed paths in Σ up to homeomorphism; and so we may choose one, P say, such that there is an infinite antichain $\mathscr{A}_0$ of graphs which all have drawings in Σ using only k' points of P, and these points are used only to represent vertices of the graphs (not edges). Let Σ' be the surface obtained from Σ by cutting along P, and for each $G \in \mathscr{A}_0$, let G' be the graph obtained from G by "splitting" the vertices of G drawn on P in the natural way. We thus obtain an infinite set of graphs drawn in Σ'. If this were an infinite antichain, we would have a contradiction to our choice of Σ. Unfortunately it need not be. The difficulty lies in keeping track of the special vertices which have been split. But there is only a bounded number of these. Let us go back and begin again, now working with graphs drawn in Σ with a bounded number of distinguished vertices, drawn in $bd(\Sigma)$. Again, choose Σ as simple as possible embedding an infinite antichain of such objects. The argument given before again (essentially) applies, and we do now obtain a contradiction to the choice of Σ.

We have explained three techniques here. It is encouraging that these would seem to be exactly those techniques needed to exploit (6.2). (The fourth feature of (6.2), the "rings" of little graphs attached onto the boundary of the surfaces, is easily handled — we delete each such little graph, and label the vertex where it was attached to the surface with a description of the deleted graph.)

8. APPLICATIONS TO DISJOINT PATH PROBLEMS

Central to this is the idea of "splitting" a vertex. Let us take an instance of DCP — G is a graph, and $s_1, \ldots, t_k$ are vertices of G. (We say that k is the *index* of the instance.) Suppose $v \in V(G)$, and the edges of G incident with v have a natural partition into (say) two sets E_1, E_2. Let G' be the graph obtained by "splitting" v into two vertices v_1, v_2, where the edges now incident with v_i are those in $E_i (i = 1,2)$. It is easy to see that the original DCP problem can be solved if we can solve $2k + 2$ DCP problems in G' each with index $k + 1$.

As a reduction method this is of course terrible. Not only is one problem replaced by $2k + 2$, but also the index increases, and even the size of the graph increases. However, it can sometimes be useful. Suppose for example that G has a small cutset, dividing $V(G)$ into two roughly equal pieces, and the vertices $s_1, \ldots, t_k$ are also roughly evenly distributed between these two sides. If we split all the vertices of the cutset in the natural way, the graph becomes disconnected, and we can solve the derived DCP problems by working with the components separately. The DCP problems on each component are much simpler than the original on G; and if only this reduction process could be applied to the components as well, we would have a polynomial algorithm.

Let us see how to use this idea to solve DCP with bounded index k for graphs with bounded tree-width $\leqslant w$ [16]. Fix c with $k, 3w \leqslant c$. Now for any subset $X \subseteq V(G)$, there is a cutset of cardinality $\leqslant w$, dividing X into two parts, neither of which is bigger than twice the other; and we can find such a cutset in polynomial time. (Simply try all subsets of $V(G)$ with cardinality $\leqslant w$.) Let us choose such a cutset for $X = V(G)$, and another for $X = \{s_1, ..., t_k\}$; and split in the natural way all vertices in both cutsets. We find that our original problem is reduced to a number (bounded by some function of c) of DCP problems on graphs which still have tree-width $\leqslant w$, which have at most about two-thirds of the original number of vertices; and it can be checked that they all have index still at most c. (This was the reason for the choice of $c \geqslant 3w, k$.) This then yields a polynomial algorithm.

We can use the same algorithm on graphs not known to have tree-width $\leqslant w$ if we permit three possible outcomes:

(i) the paths exist

(ii) the paths do not exist

(iii) G has tree-width $> w$.

For we simply search for the small cutsets, as before; and if we do not find them, we conclude that (iii) holds.

This permits us, for a fixed planar graph H, to test in polynomial time if an input graph G has a minor isomorphic to H. For we choose w so that, by (4.11), every graph with tree-width $> w$ has a minor isomorphic to H. Then we can use the algorithm above to solve DCP and hence to test for H as a minor; for if it ever outputs (iii) we can stop, since H is present.

Now secondly, let us discuss DCP when G is restricted to embed in some connected surface Σ. Our approach is once again to use a structure theorem, that for sufficiently general graphs the paths exist; and if our input graph is not "sufficiently general" we exploit its failure, to reduce to a number of simpler problems. We can prove for example that if Σ is not the sphere, and a drawing of G has sufficiently high representativeness and the vertices $s_1, ..., t_k$ are sufficiently far apart, then the required paths exist. But this structure theorem is not good enough, because we cannot exploit every case in which its hypotheses fail to hold. Certainly if the representativeness is too small, we can exploit that — we split all the vertices in the offending non-null-homotopic closed path (choosing such a path which meets the drawing only in vertices of G) and reduce to a number of problems on a simpler surface. But what do we do if some of the s_i's and t_i's are too close together?

Our problem is solved by using the following better structure theorem.

8.1 [21] *Let Σ be a connected surface with boundary, not the sphere with one, two, or three holes. For any $k \geqslant 0$ there exists N so large that the following is true. Let G be a graph with a drawing in Σ, and let $s_1, ..., t_k$ be the vertices of G which are drawn in $bd(\Sigma)$. Suppose that*

(i) *in Σ (not necessarily in G) there are k disjoint paths linking s_i and t_i $(1 \leqslant i \leqslant k)$*

(ii) *for every non-null-homotopic non-self-intersecting closed path P which simply surrounds one hole in Σ, the number of points of P used by the drawing is at least as large as the number of points of the boundary of the corresponding hole used by the drawing*

(iii) *for every non-null-homotopic non-self-intersecting closed path P of any other kind, the number of points of P used by the drawing is at least N.*
Then the required paths exist.

This can easily be used in an algorithm. If (i) fails then certainly the paths do not exist in G. If some P fails to satisfy (ii) or (iii) we cut along it, splitting the vertices of G on it; and we find that we have reduced to a number of problems either on the same surface with smaller index, or on to a simpler surface. If Σ is a sphere with three holes a slight variation of (8.1) can be used. If Σ is a disc or cylinder we need a different technique, but it exists — it is even practical [20].

We have also found (8.1) to be very useful theoretically. It is this, for instance, that is the machinery behind the proof of (6.1) and (6.2).

Lastly, let us sketch a way in which (6.2) may yield a polynomial algorithm for DCP in general. We are given a graph G, and vertices $s_1, ..., t_k$ of G (which we may assume distinct). Put $Z = \{s_1, ..., t_k\}$.

Step 1. Test if there is a separation (G_1, G_2) of G with $Z \subseteq V(G_1)$, $|V(G_1 \cap G_2)| < 2k$ and $|V(G_2)|$ large. If there is one, we list all those pairings of $V(G_1 \cap G_2)$ which correspond to disjoint collections of paths in G_2, each path beginning and ending in $V(G_1 \cap G_2)$. We can do this because each candidate pairing corresponds to a DCP problem of index $\leqslant k - 1$ which we can therefore solve, by induction on k. Having constructed this list, we find a small graph with the same list, and substitute it for G_2. (The preparation of these small graphs is preprocessing.) We continue these substitutions until no such separation remains.

Step 2. Let Σ be an connected orientable surface of genus k, and let H be a graph which has a drawing in Σ with high representativity. Now we can follow the proof of (6.2) with a polynomial algorithm, and hence construct either

(i) a minor H' of G isomorphic to H, or

(ii) a presentation of G as a member of $\langle \mathscr{C} \rangle$, where $\mathscr{C}$ is the class of (Σ', k', w')-drawable graphs, and Σ', k', w' are independent of G.

Step 3. If (i) above occurred, then the required paths exist. For there are $2k$ disjoint paths of G linking Z to $V(H')$, in such a way that their ends in $V(H')$ do not "block each other in" — this can be deduced from the absence of separations eliminated in step 1. But since the genus of Σ is so high, the ends of these paths can be joined in any required pairing by disjoint paths of H' — this follows from (8.1). Hence as claimed, the required paths exist. We assume then that (ii) occurs.

Step 4. We use "divide and conquer" methods to reduce to a number of DCP problems on (Σ', k', w')-drawable graphs. For each such problem, we can eliminate the $\leqslant w'$ special vertices, and reduce to a number of DCP problems on (Σ', k')-drawable graphs.

There remains the question of how to solve DCP problems on (Σ', k')-drawable graphs. If $k' = 0$ this is what we previously discussed. If $k' \neq 0$, it seems that we need a variation on (8.1), which we think we have found, but it is too complex to describe here.

REFERENCES

[1] D. Archdeacon, *A Kuratowski Theorem for the Projective Plane*, Thesis, The Ohio State University (1980).

[2] D. Archdeacon and P. Huneke, "On irreducible graphs for non-orientable surfaces," preprint (1981).

[3] G. Dirac, "A property of 4-chromatic graphs and remarks on critical graphs," *J. London Math. Soc.* **27** (1952), 85-92.

[4] G. Dirac, "Some theorems on abstract graphs," *Proc. London Math. Soc.* (*Ser.* 3), **2** (1952), 69-81.

[5] P. Erdös and L. Pósa, "On the independent circuits contained in a graph," *Canad. J. Math.* **17** (1965), 347-352.

[6] S. Fortune, J. Hopcroft and J. Wyllie, "The directed subgraph homeomorphism problem," *J. Theoret. Comput. Sci.* **10** (1980), 111-121.

[7] H. Glover, P. Huneke and C. S. Wang, "103 graphs that are irreducible for the projective plane," *J. Combinatorial Theory* (*Ser. B*) **27** (1979), 332-370.

[8] D. W. Hall, "A note on primitive skew curves," *Bull. Amer. Math. Soc.* **49** (1943), 935-937.

[9] G. Higman, "Ordering by divisibility in abstract algebras," *Proc. London Math. Soc.* **2** (1952), 326-336.

[10] R. M. Karp, "On the complexity of combinatorial problems," *Networks* **5** (1975), 45-68.

[11] J. Kruskal, "Well-quasi-ordering, the tree theorem, and Vázsonyi's conjecture," *Trans. Amer. Math. Soc.* **95** (1960), 210-225.

[12] J. F. Lynch, "The equivalence of theorem proving and the interconnection problem," *ACM SIGDA Newsletter* **5:3** (1976).

[13] W. Mader, "Wohlquasigeordnete Klassen endlicher Graphen," *J. Combinatorial Theory* (*Ser. B*), **12** (1972), 105-122.

[14] C.St.J.A. Nash-Williams, "On well-quasi-ordering infinite trees," *Proc. Cambridge Phil. Soc.* **61** (1965), 697-720.

[15] N. Robertson and P. D. Seymour, "Graph minors. I. Excluding a forest," *J. Combinatorial Theory* (*Ser. B*), 39-61.

[16] N. Robertson and P. D. Seymour, "Graph minors. II. Algorithmic aspects of tree-width," *J. Algorithms*, to appear.

[17] N. Robertson and P. D. Seymour, "Graph minors. III. Planar tree-width," *J. Combinatorial Theory* (*Ser. B*), 49-64.

[18] N. Robertson and P. D. Seymour, "Graph minors. IV. Tree-width and well-quasi-ordering," submitted.

[19] N. Robertson and P. D. Seymour, "Graph minors. V. Excluding a planar graph," submitted.

[20] N. Robertson and P. D. Seymour, "Graph minors. VI. Disjoint paths across a disc," submitted.

[21] N. Robertson and P. D. Seymour, "Graph minors. VII. Disjoint paths on a surface," submitted.

[22] N. Robertson and P. D. Seymour, "Graph minors. VIII. A Kuratowski theorem for general surfaces," submitted.

[23] N. Robertson and P. D. Seymour, "Graph minors. IX. Disjoint crossing paths," submitted.

[24] P. D. Seymour, "Disjoint paths in graphs," *Discrete Math.* **29** (1980), 293-309.

[25] Y. Shiloach, "A polynomial solution to the undirected two paths problem," *J. Assoc. Comput. Mach.* **27** (1980), 445-456.

[26] C. Thomassen, "2-linked graphs," *European J. Combinatorics* **1** (1980), 371-378.

[27] K. Wagner, "Uber eine Eigenshaft der ebenen Komplexe," *Math. Ann.* **114** (1937), 570-590.