A Simpler Algorithm and Shorter Proof for the Graph Minor Decomposition

[Extended Abstract] *

Ken-ichi Kawarabayashi†
National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-ku
Tokyo 101-8430, Japan
k_keniti@nii.ac.jp

Paul Wollan
University of Rome, La Sapienza
Via Salaria 113
Rome, 00198 Italy
wollan@di.uniroma1.it

ABSTRACT

At the core of the Robertson-Seymour theory of graph minors lies a powerful decomposition theorem which captures, for any fixed graph H, the common structural features of all the graphs which do not contain H as a minor. Robertson and Seymour used this result to prove Wagner's Conjecture that finite graphs are well-quasi-ordered under the graph minor relation, as well as give a polynomial time algorithm for the disjoint paths problem when the number of the terminals is fixed. The theorem has since found numerous applications, both in graph theory and theoretical computer science. The original proof runs more than 400 pages and the techniques used are highly non-trivial.

In this paper, we give a simplified algorithm for finding the decomposition based on a new constructive proof of the decomposition theorem for graphs excluding a fixed minor H. The new proof is both dramatically simpler and shorter, making these results and techniques more accessible. The algorithm runs in time $O(n^3)$, as does the original algorithm of Robertson and Seymour. Moreover, our proof gives an explicit bound on the constants in the O notation, whereas the original proof of Robertson and Seymour does not.

Categories and Subject Descriptors

G.2.2 [**Discrete Mathematics**]: Graph Theory—*graph algorithms, path and circuit problems*

General Terms

Theory

Keywords

Disjoint paths problem, graph minors, graph algorithms

*A detailed version of this paper is available at `http://research.nii.ac.jp/~k_keniti/easystruct.pdf`

†Research partly supported by Japan Society for the Promotion of Science, Grant-in-Aid for Scientific Research, by C & C Foundation, by Kayamori Foundation and by Inoue Research Award for Young Scientists.

STOC'11, June 6–8, 2011, San Jose, California, USA.

1. INTRODUCTION

A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges. The theory of graph minors was developed by Robertson and Seymour in a series of 23 papers published over more than twenty-five years. The aim of the series of papers is to prove a single result: the *graph minor theorem*, which says that in any infinite collection of finite graphs there is one that is a minor of another. As with other deep results in mathematics, the body of theory developed for the proof of the graph minor theorem has also found applications elsewhere, both within graph theory and computer science. Yet many of these applications rely not only on the general techniques developed by Robertson and Seymour to handle graph minors, but also on one particular auxiliary result which is central to the proof of the graph minor theorem: a result which approximately describes the structure of all graphs G which do not contain some fixed graph H as a minor. At a high level, the theorem says that every such a graph can be decomposed into a collection of graphs each of which can be "almost" embedded into a bounded-genus surface; the pieces can be assembled in a tree structure to obtain the original graph. This decomposition theorem is used to verify Wagner's Conjecture [30], which can be stated as follows: every minor-closed graph property (i.e. a property preserved under taking of minors) is characterized by a finite set of forbidden minors. It is also used in the proof of the correctness for the seminal graph minor algorithm, a polynomial-time algorithm for testing the presence of a fixed minor [27]. When combined together, the graph minor algorithm and the proof of Wagner's Conjecture immediately imply the existence of a polynomial-time algorithm for deciding membership in any minor-closed class of graphs. There have been numerous applications of the decomposition theorem; we discuss this more thoroughly below.

The proof of this decomposition theorem is extremely long and technical, and utilizes much of the theory of graph minors which Robertson and Seymour developed. The proof occupies the first 16 graph minor papers and is at least 400 pages long. Moreover, some of the bounds given in the proof are not explicit. The combination of the usefulness of the graph minor decomposition theorem and the difficulty of the proof have motivated the search for a more accessible proof. To quote a survey article of Lovász [19], "It would be quite important to have simpler proofs with more explicit bounds. Warning: many of us have tried, but only a few successes can be reported."

1.1 Our contribution

Building on the methods we have developed in [16], we give a much shorter proof of the graph minor decomposition theorem. We begin by giving the necessary notation to state the decomposition theorem more precisely.

A *path decomposition* of a graph G consists of linearly ordered subsets $\mathcal{B}_i$, $1 \leq i \leq k$ with $\mathcal{B}_i \subseteq V(G)$ such that for every edge $uv \in E(G)$, there exists an index i such that $u, v \in \mathcal{B}_i$ and furthermore, for all vertices x, if $x \in \mathcal{B}_i$ and $x \in \mathcal{B}_j$, then $x \in \mathcal{B}_l$ for all $i \leq l \leq j$. The *width* of the decomposition is the maximal size of a $\mathcal{B}_i$, and the *path-width* of a graph G, is the minimum width over all possible path decompositions of G. The *adhesion* of the decomposition is the $\max_{1 \leq i < j \leq k} \mathcal{B}_i \cap \mathcal{B}_j$.

In this paper, an *embedding* refers to a *2-cell embedding*, i.e. a drawing of the vertices and edges of the graph as points and arcs in a surface such that every face (region outlined by edges) is homeomorphic to a disc.

We now make explicit what we mean by "almost" embedding in a surface. Let G be a graph and Σ be a general surface. Let k be a positive integer. A *k-near embedding in Σ* consists of edge disjoint subgraphs $H_0, H_1, \ldots, H_m$ for some positive integer m satisfying the following conditions.

1. $\bigcup_{i=0}^{m} H_i = G$.
2. For all $i, j \geq 1$, $i \neq j$, $V(H_i) \cap V(H_j) \subseteq V(H_0)$.
3. For all $i > k$, $|V(H_i) \cap V(H_0)| \leq 3$.
4. There exist pairwise disjoint open discs $\Delta_1, \ldots, \Delta_m$ and an embedding $\sigma : H_0 \hookrightarrow \Sigma - \bigcup_{i=1}^{m} \Delta_i$ such that the only vertices contained in the boundary of Δ_i are exactly the vertices of $H_i \cap H_0$ for $i = 1, \ldots, m$.
5. for $1 \leq i \leq k$, let the vertices of $V(H_i) \cap V(G_0)$ be $u_1, u_2, \ldots, u_n$ for some integer n with the order given by their order on the boundary of the disc Δ_i in Δ. Then the graph H_i has a path decomposition $(\mathcal{B}_j)_{1 \leq j \leq n}$ such that $u_j \in \mathcal{B}_j$ for all $1 \leq j \leq n$.

The k-near embedding is *α-bounded* if the path decomposition of H_i in 5 has adhesion at most α for all $1 \leq i \leq k$. The k-near embedding is *totally bounded* if both $m \leq k$ and the path decomposition of H_i in 5 has width at most k for all $1 \leq i \leq k$. Finally, a graph G is *t-close* to admitting a k-near embedding if there exists a set $X \subseteq V(G)$ with $|X| \leq t$ such that $G - X$ admits a k-near embedding.

The pieces of the decomposition are combined according to "clique-sum" operations, a notion which goes back to characterizations of K_5-minor-free graphs by Wagner [33] and serves as an important tool in the theory of graph minors. Suppose G_1 and G_2 are graphs with disjoint vertex sets and let $k \geq 0$ be an integer. For $i = 1, 2$, let $W_i \subseteq V(G_i)$ form a clique of size k and let G_i' be obtained from G_i by deleting some (possibly no) edges from the induced subgraph $G_i[W_i]$. Consider a bijection $h : W_1 \to W_2$. We define a *k-sum* G of G_1 and G_2, denoted by $G = G_1 \oplus_k G_2$ or simply by $G = G_1 \oplus G_2$, to be the graph obtained from the union of G_1' and G_2' by identifying w with $h(w)$ for all $w \in W_1$. The images of the vertices of W_1 and W_2 in $G_1 \oplus_k G_2$ form the *join set*. Note that each vertex v of G has a corresponding vertex in G_1 or G_2 or both. It is also worth mentioning that $\oplus$ is not a well-defined operator: it can have a set of possible results.

Now we can state a precise form of the *Decomposition theorem*:

Theorem 1 (Theorem 1.3 [29]). *For every graph H, there exists an integer $h \geq 0$ depending only on $|V(H)|$ such that every H-minor-free graph can be obtained by at most h-sums of graphs which are h-close to admiting a totally bounded h-near embedding in some surfaces in which H cannot be embedded.*

We give a much shorter and simpler proof of Theorem 1. The proof is constructive and immediately yields an $f(|H|)n^3$ time algorithm to find such a decomposition for excluding H-minor. In addition, our proof also gives an explicit bound for the value h.

The original proof of Robertson and Seymour also gives an $f'(|H|)n^3$ algorithm to find the decomposition, if one follows all the arguments very carefully (as pointed out by Reed (private communication)). Our algorithm is an improvement, in that it is far easier and more accessible, and in that it improves the bounds on the function f and provides explicit bounds for the value h. Recently, Reed, Li and the first author announced an $O(n \log n)$ algorithm to find the decomposition theorem. The proof generally follows the argument in the graph minor theory, however some of the more technical graph minor results must be strengthened. Complete details will require more than 100 pages and are not yet fully written down.

1.2 Algorithmic applications

Algorithms for H-minor-free graphs for a fixed graph H have been studied extensively; see e.g. [3, 4, 12, 17, 20]. In particular, it is generally believed that many algorithms for planar graphs can be generalized to H-minor-free graphs for any fixed H [12, 17, 20]. The decomposition theorem provides the key insight into why this might be possible: given an algorithm for planar graphs, first extend it to handle bounded-genus graphs; then extend it further to handle graphs "almost-embeddable" into bounded-genus surfaces, and finally generalize the results to resolve the problem on repeated clique sums of graphs which are almost embedded in surfaces of bounded genus. The graph minor decomposition theorem has already been used to obtain many combinatorial results and show the existence of many efficient algorithms, despite being published only recently. Grohe [11] proves the existence of PTASs for minimum vertex cover, minimum dominating set, and maximum independent set in H-minor-free graphs. However, such an approach requires an algorithm to construct the decomposition. Our simpler proof and algorithm gives a fast and more accessible algorithm for these problems.

DeVos et al. [8] used the decomposition theorem to prove that for every integer $k \geq 1$ and every fixed graph H, every H-minor-free graph has a vertex partition into parts

$V_1, \ldots, V_k$ and edge partition $E_1, \ldots, E_k$ such that for every $i \in \{1, \ldots, k\}$, the graphs $G - V_i$ and $G - E_i$ have bounded treewidth. It follows that given an algorithm to construct the decomposition, there is a 2-approximation algorithm for graph coloring in minor-closed class of graphs, see [7]. A special case of this partitioning result restricted to planar graphs was proved by Baker [2] who used it to devise efficient approximation algorithms (and approximation schemes) for several hard approximation algorithms on planar graphs. Baker used the planar separator theorem of Lipton and Tarjan [18]. Alon, Seymour, and Thomas [1] proved a similar separator theorem for graphs excluding any fixed minor, and this result is further generalized by Kawarabayashi and Reed [15]. Eppstein [10] extended Baker's ideas to graphs in arbitrary proper minor-closed classes of graphs. Again, these methods for finding approximation algorithms for problems in minor closed classes begin with the graph minor decomposition. Thus, our proof and algorithm yields fast and more accessible algorithms for these approximation results.

Demaine et al. [5] have been working on directly using the graph minor theorem in algorithmic applications. They obtained subexponential fixed-parameter algorithms for dominating set, vertex cover and set cover in any class of graphs excluding a fixed graph H as a minor. Specifically, the running time is $2^{O(\sqrt{k})}n^h$, where h is a constant depending only on H. For further applications, see the survey [6] or [7]. Our simpler proof gives not only more accessible and simpler algorithms but also faster algorithms, i.e. $h = 3$.

2. THE PROOF AND ALGORITHM

Robertson and Seymour proved several related structure theorems describing graphs which exclude a fixed minor. The decomposition theorem, Theorem 1, is the one that appears now to be best known, and which has also found the most algorithmic applications. However, Robertson and Seymour themselves [29] later dubbed it a 'red herring' in the search for the proof of the graph minor theorem. We begin this section by describing the "main" structure theorem in graph minor theory. Theorem 1 can then be relatively quickly shown from this main structure theorem.

This "main" structure theorem is designed to eliminate difficulties arising from small cutsets in a graph. Given a graph G, we recall that a *separation* of G is a pair $A, B \subseteq V(G)$ such that every edge of G has both endpoints contained either in A or in B (or in both). The order of the separation is $|A \cap B|$. Note that set $|A \cap B|$ is a cutset in the graph if $A \nsubseteq B$ and $B \nsubseteq A$, and every cutset in the graph gives rise to a separation. The structure theorem we are going to describe allows us to restrict our attention to one portion of the graph, in effect disregarding small pieces of the graph which are separated by a small cutset. We will now make this more explicit.

A particularly simple form of this structure theorem applies when the excluded minor H is planar: in that case, the said parts of G—the parts that fit together in a tree-structure and together make up all of G—have bounded size, i.e., G has bounded *treewidth*.

We will not need the explicit definition of treewidth here. Of importance to us, however, is the relationship between grid minors and the treewidth of a graph. Specifically, the $k \times k$-grid has treewidth k. Moreover, Robertson and Seymour showed [21] that there exists a function $f(k)$ such that every graph with treewidth at least $f(k)$ contains the $k \times k$-grid as a minor. Thus, grid minors offer an approximate characterization of when a given graph has large treewidth. An immediate consequence of this result is the fact mentioned above: if H is a planar graph then the graphs which do not contain H as a minor all have bounded treewidth.

If H is not planar then the set of all graphs not containing H as a minor have unbounded treewidth. Therefore this set of graphs contain arbitrarily large grids as minors. For technical reasons, it will be easier to work with *r-wall* minors. An r-wall is a grid-like graph which has maximum degree three. Every r-wall contains the $r \times r$-grid as a minor, and the $2r \times 2r$-grid contains an r-wall as a minor. Thus, the graphs with no H-minor also contain arbitrarily large walls as a minor. Since walls have maximum degree three, if G contains an r-wall as a minor, then G contains a subdivision of the r-wall as a subgraph.

Such a wall identifies, for every low-order separation of G, one side in which most of that grid or wall lies. Specifically, given a subdivision W of the r-wall in G and a separation (A, B) of order less than r, there is exactly one of A or B which contains one entire row of the wall. This is formalized by the notion of a *tangle*: the larger the treewidth of G, the larger the grid or wall, the order of the separations for which this works, and (thus) the *order* of the tangle. The advantage of working with respect to a tangle is it allows us focus on the portion of the graph containing a large wall and disregard pieces separated by small separations from the wall. We say that a α-near embedding of a graph G in a surface Σ *captures* a tangle $\mathcal{T}$ associated to a wall subdivision W in G if for every H_i, $i \geq 1$ in the definition of an α-near embedding we have that $G - H_i$ contains at least one row of the wall W. Thus, the portion of the graph, G_0, which is embedded in the surface Σ in a sense "contains" the tangle $\mathcal{T}$.

We are now ready to state the main theorem in graph minor theory describing graphs excluding a fixed H as a minor. We will refer to it as the *Structure theorem.*

THEOREM 2 (THEOREM 3.1, [29]). *For every graph R there exist integers $\theta, \alpha = \alpha(|R|) \geq 0$ such that the following holds: Let G be a graph that does not contain R as a minor and $\mathcal{T}$ be a tangle in G of order at least θ. Then there exists a subset $A \subseteq V(G)$ with $|A| \leq \alpha$ such that $G - A$ has an α-bounded α-near embedding into a surface Σ in which R cannot be drawn. Moreover, this near embedding captures $\mathcal{T} - A$.*

Assuming Theorem 2, the decomposition theorem can be proven relatively easily. We give the proof in Section 3. Thus, the majority of the work in our proof of the decomposition theorem lies in proving Theorem 2. Our main tool in proving Theorem 2 will be several new results on embedding *societies* in surfaces. We discuss this in more detail in the next subsection. We then give an outline of the proof of Theorem 2 in Subsection 2.2.

2.1 Nearly embedding a society

A *society* is a pair (G, Ω) where G is a graph and Ω is a cyclic ordering of some of the vertices of G. Societies play a key role in the graph minor series, see [23].

In our proof, we will need to understand when a given society (G, Ω) is t-close to having a totally bounded k-near embedding in the disc Δ for given values t and k. Two

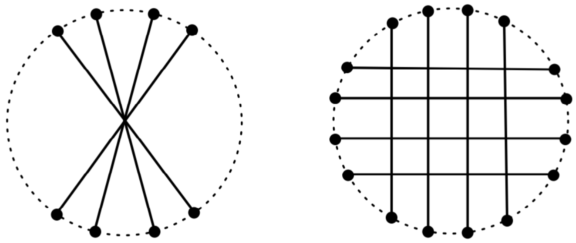

Figure 1: An example of a 4-crosscap and a 4-handle.

possible obstructions are what's known as *k-crosscaps* and *k-handles*. A k-crosscap consists of k disjoint, pairwise crossing paths with their endpoints in Ω. A k-handle consists of $2k$ pairwise disjoint paths $P_1, \dots, P_k, Q_1, \dots, Q_k$ each with their endpoints in Ω which satisfy the following. For all i, j, we have that P_i and P_j do not cross and similarly, Q_i and Q_j do not cross, and alternatively, the paths P_i and Q_j do cross. See Figure 1.

It would be convenient if these were the only obstructions. Unfortunately, we must define several *patterns* consisting of disjoint paths with their endpoints in Ω. We will not need their exact definition here, and we will simply refer to the pattern with k paths as $\mathcal{P}_k$.

THEOREM 3. *Let $k \geq 1$ be given. There exists a value $\alpha = \alpha(k)$ such that for every society (G, Ω), one of the following holds.*

1. *(G, Ω) contains a k-crosscap, a k-handle, or the pattern $\mathcal{P}_k$.*
2. *(G, Ω) is α-close to admitting a α-bounded α-near embedding in the disc Δ such that all the vertices of Ω are embedding on the boundary of Δ in the order indicated by Ω*

Theorem 3 generalizes the main theorem of [23]. Let Δ be the disc and let D_1 and D_2 be two disjoint, open discs in Δ which do not intersect the boundary of Δ. Let Σ_1 be the surface obtained by deleting D_1 and gluing a crosscap to the boundary of D_1. Let Σ_2 be the surface obtained by deleting D_1 and D_2 from Δ and gluing a handle onto the boundaries of D_1 and D_2. The existence of a k-crosscap or k-handle for large k in a society (G, Ω) implies that the natural surface in which we should attempt to embed (G, Ω) is Σ_1, or Σ_2, respectively. When we can do so is given in the following theorem.

THEOREM 4. *Let $k \geq 1$ be given. There exists a value $f = f(k)$ and $\alpha = \alpha(k)$ satisfying the following. Let (G, Ω) be a society containing a f-crosscap (f-handle) $\mathcal{Q}$ but which does not contain the pattern $\mathcal{P}_k$. Then there exist edge disjoint subgraphs H_1 and H_2 of G such that $H_1 \cup H_2 = G$ which satisfy the following.*

1. *H_1 contains a k-crosscap (k-handle) $\mathcal{Q}'$ which is a subgraph of $\mathcal{Q}$.*
2. *$\Omega \subseteq V(H_1)$ and (H_1, Ω) is α-close to admitting a 0-near embedding in Σ_1 (Σ_2).*
3. *There exists a single face of the near embedding in Σ_1 or Σ_2 which contains all the vertices of $V(H_1) \cap V(H_2)$.*

The proofs of Theorems 3 and 4 contain the main technical work in the new proof. We will see in the next subsection how these theorems come into the proof of Theorem 2.

2.2 Outline of the proof of Theorem 2

We now give an overview of our proof for the structure theorem, Theorem 2. Our starting point is so called *Weak Structure theorem* which is proved in Graph Minor XIII [27]. Roughly it says that, given any wall subdivision W of size $f(t, k)$ in a given graph G, either G has a K_t-minor, or G has a vertex set Z of order at most t^2 such that $G - Z$ has a subwall W' of W of size k which induces essentially planar embedding in $G - Z$. More explicitly:

THEOREM 5 (THEOREM 9.4 [27]). *Let G be a graph. Let W be a wall in G. Either G contains a K_t minor which cannot be separated from W by a small order separation, or there exists a subset $A \subseteq V(G)$, a sub-wall W' of W and edge disjoint subgraphs H_1 and H_2 of G satisfying the following.*

1. *$|A| \leq t^2$ and $G - A = H_1 \cup H_2$.*
2. *W' is a subgraph of H_1 and H_1 has a 0-near embedding in the disc Δ such that at least t vertices of degree 3 in W' are embedded in the disc and all the vertices of $V(H_1) \cap V(H_2)$ are embedding on the boundary of Δ.*

Note that given the structure in Theorem 5, we have a natural society given by the subgraph H_2 and the cyclic ordering of the vertices of $V(H_1) \cap V(H_2)$ given by their embedding on the boundary of the disc Δ.

We begin with a graph G not containing K_t as a minor, and now our proof adapts some ideas in [28, 29] to grow a large subgraph of G which essentially embeds on a surface with large representativity. We recall that a *non-contractable curve* in the surface is simply a curve that cannot be contracted continuously to a point on the surface. The *representativity* of an embedded graph is minimum number of times a non-contractable curve C intersects the embedded graph, with the minimum taken over all such non-contractable curves C.

We proceed inductively maintaining two subgraphs H_1 and H_2 of G and a subset of the vertices Z such that

1. $H_1 \cup H_2 = G - Z$,
2. H_1 essentially embeds in a surface Σ with big representativity (technically, H_1 has a 0-near embedding in Σ), and
3. there is a single face F of the embedding of H_1 containing the vertices of $V(H_1) \cap V(H_2)$.

We then consider the society given by H_2 and the cyclic order of the vertices of $V(H_1) \cap V(H_2)$ given by the boundary of the face F (denote as Ω this cyclic order of $V(H_1) \cap V(H_2)$). If there existed a set $Z_2 \subseteq V(H_2)$ with $|Z_2| \leq \alpha$ such that $(H_2 - Z_2, \Omega - Z_2)$ admitted an α-bounded α-near embedding of (G, Ω) in the disc, then we could "glue" it onto the embedding of H_1 in Σ and find the desired α-bounded α-near embedding of $G - (Z \cup Z_2)$.

We apply Theorem 3 to attempt to find such a nice embedding of (H_2, Ω). We can (by picking α sufficiently large), assume that we find either a large crosscap or handle, or a large pattern $\mathcal{P}_k$. However, it is straightforward to show that if k is large, then H_1 along with the pattern $\mathcal{P}_k$ yields a large clique minor. Thus, we may assume that we always find a large crosscap or handle. We apply Theorem 4 to the society (H_2, Ω) along with the large crosscap or handle. We find subgraphs J_1 and J_2 and subset X of vertices such that (J_1, Ω) embeds in the disc plus a crosscap or handle.

Returning to our graph original graph G, we see that $G - (Z \cup X) = H_1 \cup J_1 \cup J_2$ and that $H_1 \cup J_1$ has a 0-near embedding in the surface Σ' obtained by adding a single crosscap or handle to Σ. Moreover, by using the fact that (J_1, Ω) contains a large handle or crosscap, we can ensure that this embedding of $H_1 \cup J_1$ embeds with large representativity (although it will be somewhat smaller than the representativity of the embedding of H_1 in Σ). Finally, there exists a single face of the embedding in Σ' which contains all the vertices of $V(H_1 \cup J_1) \cap V(J_2) = V(J_1) \cap V(J_2)$.

Thus, at each inductive step we grow the genus of the surface maintaining large representativity. The base case in the induction is the weak structure theorem with the large graph essentially embedded in the sphere. The whole process must eventually stop since there is a theorem [22] which states that for a sufficiently large genus surface and a sufficiently large amount of representativity, any graph embedded with such representativity must contain K_t as a minor.

2.3 Improvements over the original

Let us clarify why our proof is much shorter than the original proof by Robertson and Seymour.

1. We only need to introduce three parameters in our proof (the genus of the surface into which we embed, the representativity of the embedding, and the size of the set of vertices to be deleted). On the other hand, Robertson and Seymour consider seven parameters, see Graph Minor XVI [29]. This leads to a particularly technical and sensitive induction hypothesis in the Robertson-Seymour proof.

2. Our starting point is the "weak structure theorem" as above, while Robertson and Seymour begins with the grid theorem. Our initial setup immediately allows us to focus on a single "special" face and a single "society" because the rest of the graph is essentially embedded into a disc. On the other hand, Robertson and Seymour need to analyze how the rest of the graph attaches to the grid. This requires a lot of work, as in Graph Minors XIV, XV and XVI [25, 28, 29].

3. We maintain one "special" face and one society, as above, at each inductive step. Alternatively, Robertson and Seymour have to deal with many societies simultaneously. Thus our proof makes the genus addition step much easier. This issue is discussed in Graph Minor XVI [29].

4. In the inductive step, we maintain a subgraph which is essentially embedded in the surface with large representativity. Robertson and Seymour only maintain a subdivision of G embedded in the surface. This difference allows us to deal with the "connectivity" issue more easily. This issue is mainly discussed in Graph Minor XV [28].

5. In addition to a subgraph embedded in the surface, Robertson and Seymour maintain a set of "long" jumps attaching to distant faces of a graph in the surface (see Graph Minor XVI [29]). However, in our case, we simply maintain an embedding on a surface with large representativity.

6. The biggest advantage for our proof is that we do not have to worry about "distance" on a surface. More precisely, Robertson and Seymour have to maintain a subdivision embedded on a surface, but in each inductive step they delete large portions of the subgraph embedded in the surface. In order to maintain the "long" jumps mentioned above in 5, they rely on a technical distance measure for the graph embedded in the surface. This issue is actually quite troublesome in the Graph Minors Series; both Graph Minors XI and XII [24, 26] are devoted to this distance measure. Alternatively, because we do not have to maintain such "long" jumps, we do not rely on this distance measure.

2.4 Extracting an algorithm

Our proof is constructive and can be converted into a polynomial time algorithm (in fact an $O(n^2)$ time algorithm for fixed H) to obtain the structure given in the main structure theorem.

As subroutines, we only need the following.

1. In the arguments, we often need to find, for some two vertices s, t and fixed constant k, either k disjoint paths between s and t, or a vertex set of order at most k that separates s and t.

2. We also need to find a 0-near embedding in the disc for a given society (G, Ω).

Concerning the first point, we can use the result by Nagamochi and Ibaraki [13] to find one of the desired outcomes in $O(n)$ time. Concerning the second point, there is now an $O(n)$ time algorithm to find such a near embedding by Kapadia, Li and Reed [14]. This improves the previous best known result by Tholey [32] who gives $O(m\alpha(m, n))$ time algorithm, where the function $\alpha(m, n)$ is the inverse of the Ackermann function (see Tarjan [31]).

At each step of the proof of the main structure theorem, we may proceed by repeatedly calling the above two operations. Therefore, we obtain an $O(n^2)$ time algorithm to find the structure ensured by the main structure theorem. For the decomposition theorem, Theorem 1, in each iteration, we apply the algorithm of the structure theorem and then recurse on a smaller graph. Thus we obtain an $O(n^3)$ time algorithm to construct the structure in Theorem 1. How we extract the $O(n^3)$ time algorithm for the graph minor decomposition from the algorithm for the graph minor structure theorem is treated in more detail in the next section.

3. PROOF OF THEOREM 1

In this section, we prove Theorem 1 assuming Theorem 2. For the definition of the *torso, tangle, tree decomposition, vortices*, we refer the reader to Diestel's book [9]. We recall that a *separation* of a graph G is a pair (A, B) of subsets of vertices such that $G[A] \cup G[B] = G$.

THEOREM 6. *For every graph R there exist integers α and θ such that for every graph G that does not contain R as a minor and every $Z \subseteq V(G)$ with $|Z| \leq 3\theta - 2$ there is a rooted tree decomposition $\{V_t \mid t \in T\}$ of G with root r such* *that for every $t \in V(T)$, there is a surface Σ_t in which R cannot be embedded, and a subset $A_t \subseteq V(G_t)$ of the torso G_t of V_t with $|A_t| \leq \alpha$ such that $G_t - A_t$ has an α-bounded α-near embedding into Σ_t with the following properties:*

1. *There are at most α vortices.*
2. *All vortices have path-width at most α.*
3. *For every $t' \in V(T)$ with $tt' \in E(T)$ and $t \in rTt'$ there is a vertex set X which is either*
   - *(a) two consecutive parts of a vortex decomposition in G_t or*
   - *(b) a subset of $V(G_t)$ and induces in G_t a K_1, a K_2 or a triangle face in Σ_t.*

   such that $V_t \cap V_{t'} \subseteq X \cup A_t$.
4. *For every $t' \in T$ with $tt' \in E(T)$ and $t' \in rTt$ the overlap $V_t \cap V_{t'}$ is contained in $A_{t'}$.*

Furthermore $Z \subseteq A_r$. We say that the part V_r accommodates *Z.*

Proof. Applying Theorem 2 with the given graph R yields two constants $\hat{\alpha}$ and $\hat{\theta}$. Let $\theta := \max(\hat{\theta}, 3\hat{\alpha}+1)$ and $\alpha := 4\theta - 2$.

The proof proceeds by induction on $|V(G)|$. We may assume that $|Z| = 3\theta - 2$, since if it is smaller we may arbitrarily add vertices to Z. Note, we may assume that such vertices exist, as the theorem is trivialy truel for $|V(G)| < \alpha$.

Claim 1. *We may assume that there is no separation (A, B) of order at most θ such that both $|Z \setminus A|$ and $|Z \setminus B|$ are of size at least $|A \cap B|$.*

Proof. Otherwise, let $Z_A := (A \cap Z) \cup (A \cap B)$. By assumption, $|A \cap B| \leq |Z \setminus A|$ and therefore, $|Z_A| \leq |Z|$. We apply our theorem inductively to $G[A]$ and Z_A, which yields a tree decomposition of $G[A]$ with one part G_A such that the apex set of the embedding of its torso contains Z_A. Similarly, we apply the theorem to $G[B]$ and $Z_B := (B \cap Z) \cup (A \cap B)$. We combine these two tree decompositions by joining a new part $Z \cup (A \cap B)$ to both G_A and G_B and obtain a tree decomposition of G with the desired properties of the theorem: The new part contains at most $|Z| + |A \cap B| \leq 4\theta - 2$ vertices, so all these can be put into the apex set of an α-near embedding. Further, the new part contains Z. This proves the claim. □

Let $\mathcal{T}$ be the set of separations (A, B) of G of order less than θ such that $|Z \cap B| > |Z \cap A|$. With this definition, we make the following claim.

Claim 2. *$\mathcal{T}$ is a tangle of G of order θ.*

Proof. For every separation (A, B) of G of order less than θ, one of the sets $Z \setminus B$ and $Z \setminus A$ contains at least θ vertices, as $|Z| = 3\theta - 2$, but not both by (1). Therefore, property (i) of the definition of a tangle holds. We deduce further, that for every $(A, B) \in \mathcal{T}$, the small side A contains less than θ vertices from Z. Hence, the union of three small sides cannot be $V(G)$ as it contains at most $3\theta - 3$ vertices from Z, which shows property (ii) and proves the claim. □

From Claim 1 and the definition of $\mathcal{T}$, we conclude the following claim.

Claim 3. *$|(A - B) \cap Z| < |A \cap B|$ for every $(A, B) \in \mathcal{T}$.*

Theorem 2 implies that there exists a subset $\hat{A} \subseteq V(G)$ with $|\hat{A}| \leq \alpha$ such that there exists an $\hat{\alpha}$-bounded $\hat{\alpha}$-near embedding of G in some surface Σ that captures $\mathcal{T}$. At a high level, our plan is now to split up G at separators consisting of apex vertices, society vertices $\Omega(H_i)$ for $i > \hat{\alpha}$ and vertices of single parts of vortex decompositions of a vortex H_i for $i \leq \hat{\alpha}$. We obtain a part that contains H_0 and which we know how to embed α-nearly; this part is going to be one part of a new tree decomposition. We find tree decompositions for all subgraphs of G that we split off inductively and eventually combine these tree decompositions to a new one that satisfies our theorem.

Let us consider a small vortex (H_i, Ω_i) for $i \in \{\hat{\alpha}+1, \ldots, m\}$. Our embedding captures $\mathcal{T}$, therefore the separation $(V(H_i) \cup \hat{A}, V(G - (V(H_i) - V(H_0))) \cup \hat{A})$, whose order is smaller than $3 + |\hat{A}| \leq \theta$, lies in $\mathcal{T}$. By (3), H_i contains less than θ vertices of Z. Thus, $Z' := \Omega_i \cup \hat{A} \cup (Z \cap V(H_i))$ contains at most $3 + \hat{\alpha} + \theta \leq 3\theta - 1$ vertices. We apply our theorem inductively to the smaller graph $G[V(H_i) \cup \hat{A}]$ with Z'. Let H^i be a part of the resulting tree decomposition $(T^i, \mathcal{H}^i)$ that accommodates Z'. Let $\mathcal{W} = \{H_{\hat{\alpha}+1}, \ldots, H_m\}$.

For every vortex (H_i, Ω_i) with $\Omega_i = \{w_1^i, \ldots, w_{n(i)}^i\}$ for $i = 1, \ldots, \hat{\alpha}$, let us choose a decomposition $(\hat{X}_1^i, \ldots, \hat{X}_{n(i)}^i)$ of depth at most $\hat{\alpha}$. We define

$$X_j^i := \begin{cases} \big(\hat{X}_1^i \cap \hat{X}_2^i\big) \cup \{w_1^i\} & \text{for } j = 1 \\ \big(\hat{X}_j^i \cap (\hat{X}_{j-1}^i \cup \hat{X}_{j+1}^i)\big) \cup \{w_j^i\} & \text{for } 1 < j < n(i) \\ \big(\hat{X}_{n(i)}^i \cap \hat{X}_{n(i)-1}^i\big) \cup \{w_{n(i)}^i\} & \text{for } j = n(i) \end{cases}$$

By H_i^- we denote the graph on $X_1^i \cup \ldots \cup X_{n(i)}^i$ where every X_j^i induces a complete graph but no further edges are present. Now, as the depth of (H_i, Ω_i) is at most $\hat{\alpha}$, every X_j^i contains at most $2\hat{\alpha}+1$ vertices and thus, $(X_1^i, \ldots, X_{n(i)}^i)$ is a decomposition of the vortex $V_i^- := (H_i^-, \Omega_i)$ of width at most $2\hat{\alpha}+1 \leq \alpha$. Let $\mathcal{V}$ denote the set of these new vortices.

For every $j = 1, \ldots, n(i)$, the pair

$$\big(\hat{X}_j^i \cup \hat{A}, (V(G) - (\hat{X}_j^i - X_j^i)) \cup \hat{A}\big)$$

is a separation of order at most $|X_j^i \cup \hat{A}| \leq 2\hat{\alpha}+1+\hat{\alpha} \leq \theta$. As before, our embedding captures $\mathcal{T}$ and thus, the separation lies in $\mathcal{T}$. By (3), at most $\theta - 1$ vertices from Z lie in $\hat{X}_j^i$. Let $Z' := X_j^i \cup \hat{A} \cup (Z \cap \hat{X}_j^i)$. This set contains at most $3\theta - 1$ vertices and, similar to before, we can apply our theorem inductively to the smaller graph $G[\hat{X}_j^i \cup \hat{A}]$ with Z'. We obtain a tree decomposition $(T_j^i, \mathcal{H}_j^i)$ of this graph, with one part H_j^i accommodating Z'.

Now, with $V_0 := V(H_0) \cup \hat{A}$, we can write

$$G = G[H_0] \cup \big(\bigcup \mathcal{W}\big) \cup \big(\bigcup \{G[\hat{X}_j^i] : V_i \in \mathcal{V}, 1 \leq j \leq n(i)\}\big).$$

By induction, we obtained tree decompositions for all vortices in $\mathcal{W}$ and all the graphs $G[\hat{X}_j^i]$ with the required properties. We can now construct a tree decomposition of G: We just add a new vertex v_0 representing V_0 to the union of all the trees T^i and T_j^i and add edges from v_0 to every vertex representing an H^i or an H_j^i we found in our proof.

We still have to check that the torso of the new part V_0 can be α-nearly embedded as desired. But this is easy: Let H_0' be the graph resulting from H_0 if we add an edge xy for every two nonadjacent vertices x and y that lie in a common vortex $V \in \mathcal{W}$. We can extend the embedding $\sigma : G_0 \hookrightarrow \Sigma$ to an embedding $\sigma' : G_0' \hookrightarrow \Sigma$ by mapping the new edges disjointly to the discs $D(V)$. Then, $G' := H_0' \cup \bigcup H_i^- \cup \hat{A}$ is the torso of V_0 in our new tree decomposition and we have an α-bounded α-near embedding of $G' - \hat{A}$ in Σ. This completes the proof of Theorem 6.

Observe that the above proof of Theorem 6 is constructive and consequently, it gives rise to an $O(n^3)$ time algorithm to construct the decomposition as in Theorem 6.

To extract the algorithm, the first step is to define the tangle $\mathcal{T}$ of order Θ. This can be done in $O(n)$ time given that $|Z| \leq 3\Theta - 2$ is an absolute constant and we just need the standard max-flow, minimum-cut algorithm. Then we apply the algorithm to give the structure in Theorem 2 with respect to the tangle $\mathcal{T}$, as in the previous section. Finally, we recursively apply this algorithm to all the graphs in $\mathcal{W}$ and all the graphs $G[\hat{X}_j^i]$ with the required properties. We can put these decompositions together, as in the proof of Theorem 6, to obtain a tree decomposition of G by adding a new vertex v_0 representing $V_0 := V(H_0) \cup \bigcup H_i^- \cup \hat{A}$ (as in the above proof) to the union of all the trees T^i and T_j^i and add edges from v_0 to every vertex representing an H^i or an H_j^i we found in our proof.

Since we only need to apply the algorithm to construct the structure in Theorem 2 recursively, we obtain an $g(t)n^3$ time algorithm to construct the structure as in Theorem 6 for some function $g(t)$, as claimed.

4. REFERENCES

[1] N. Alon, P. D. Seymour, and R. Thomas. A separator theorem for non-planar graphs. *J. Amer. Math. Soc.*, 3:801–809, 1990.

[2] B. S. Baker. Approximation algorithms for np-complete problems on planar graphs. *J. Assoc. Comput. Mach.*, 41:153–180, 1994.

[3] M. Charikar and A. Sahai. Dimension reduction in the l_1 norm. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)*, pages 551–560. IEEE, 2002.

[4] C. Chekuri, A. Gupta, I. Newman, Y. Rabinovich, and S. Alistair. Embedding k-outerplanar graphs into ℓ_1. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'03)*, pages 527–536. ACM-SIAM, 2003.

[5] E. D. Demaine, F. Fomin, M. Hajiaghayi, and D. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and h-minor-free graphs. *J. ACM*, 52:1–29, 2005.

[6] E. D. Demaine and M. Hajiaghayi. Fast algorithms for hard graph problems: Bidimensionality, minors and local treewidth. *Lect. Note Comput. Sci.*, 3383:517–533, 2004.

[7] E. D. Demaine, M. Hajiaghayi, and K. Kawarabayashi. Algorithmic graph minor theory: Decomposition, approximation and coloring. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 637–646. IEEE, 2005.

[8] M. DeVos, G. Ding, B. Oporowski, D. Sanders, B. Reed, P. D. Seymour, and D. Vertigan. Excluding any graph as a minor allows a low tree-width 2-coloring. *J. Combin. Theory Ser. B*, 91:25–41, 2004.

[9] R. Diestel. *Graph Theory*. Springer-Verlag, Berlin, 2005.

[10] D. Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27:275–291, 2000.

[11] M. Grohe. Local tree-width, excluded minors and approximation algorithms. *Combinatorica*, 23:613–632, 2003.

[12] A. Gupta, I. Newman, Y. Rabinovich, and S. Alistair. Cuts, trees and ℓ_1-embeddings of graphs. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS'99)*, pages 399–409. IEEE, 1999.

[13] T. Ibaraki and H. Nagamochi. A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph. *Algorithmica*, 7:583–596, 1992.

[14] R. Kapadia, Z. Li, and B. Reed. A linear time algorithm to test the 2-path problem. *manuscript*.

[15] K. Kawarabayashi and B. Reed. A separator theorem in minor-closed classes. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS'10)*, pages 153–162. IEEE, 2010.

[16] K. Kawarabayashi and P. Wollan. A shorter proof of the graph minors algorithm- the unique linkage theorem. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing(STOC'10)*, pages 687–694. ACM, 2010.

[17] P. N. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC'93)*, pages 682–690. ACM, 1993.

[18] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9:615–627, 1980.

[19] L. Lovasz. Graph minor theory. *B. Amer. Math. Soc.*, 43:75–86, 2005.

[20] S. A. Plotkin, S. Rao, and W. D. Smith. Shallow excluded minors and improved graph decompositions. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'94)*, pages 462–470. ACM-SIAM, 1994.

[21] N. Robertson and P. D. Seymour. Graph minors. v: Excluding a planar graph. *J. Combin. Theory Ser. B*, 41:92–114, 1986.

[22] N. Robertson and P. D. Seymour. Graph minors. vii: Disjoint paths on a surface. *J. Combin. Theory Ser. B*, 45:212–254, 1988.

[23] N. Robertson and P. D. Seymour. Graph minors. ix: Disjoint crossed paths. *J. Combin. Theory Ser. B*, 49:40–77, 1990.

[24] N. Robertson and P. D. Seymour. Graph minors. xi: Circuits on a surface. *J. Combin. Theory Ser. B*, 60:72–106, 1994.

[25] N. Robertson and P. D. Seymour. Graph minors. ivx: Extending an embedding. *J. Combin. Theory Ser. B*, 65:23–50, 1995.

[26] N. Robertson and P. D. Seymour. Graph minors. xii: Distance on a surface. *J. Combin. Theory Ser. B*, 64:240–272, 1995.

[27] N. Robertson and P. D. Seymour. Graph minors. xiii:

The disjoint paths problem. *J. Combin. Theory Ser. B*, 63:65–110, 1995.

[28] N. Robertson and P. D. Seymour. Graph minors. xv: Giant steps. *J. Combin. Theory Ser. B*, 68:112–148, 1996.

[29] N. Robertson and P. D. Seymour. Graph minors. xvi: Excluding a non-planar graph. *J. Combin. Theory Ser. B*, 89:43–76, 2003.

[30] N. Robertson and P. D. Seymour. Graph minors. xx: Wagner's conjecture. *J. Combin. Theory Ser. B*, 92:325–357, 2004.

[31] R. E. Tarjan. *Data structures and network algorithmsl*. SIAM, Philadelphia, Pennsylvania, 1983.

[32] T. Tholey. Solving the 2-disjoint paths problem in nearly linear time. *Theory Comput Syst*, 39:51–78, 2004.

[33] K. Wagner. Uber eine eigenschaft der ebenen komplexe. *Math. Ann.*, 114:570–590, 1937.