Linearity of Grid Minors in Treewidth with Applications through Bidimensionality*

Erik D. Demaine[†] MohammadTaghi Hajiaghayi[†]

Abstract

We prove that any H-minor-free graph, for a fixed graph H, of treewidth w has an $\Omega(w) \times \Omega(w)$ grid graph as a minor. Thus grid minors suffice to certify that H-minor-free graphs have large treewidth, up to constant factors. This strong relationship was previously known for the special cases of planar graphs and bounded-genus graphs, and is known not to hold for general graphs. The approach of this paper can be viewed more generally as a framework for extending combinatorial results on planar graphs to hold on H-minor-free graphs for any fixed H. Our result has many combinatorial consequences on bidimensionality theory, parameter-treewidth bounds, separator theorems, and bounded local treewidth; each of these combinatorial results has several algorithmic consequences including subexponential fixed-parameter algorithms and approximation algorithms.

1 Introduction

The $r \times r$ grid graph[1] is the canonical planar graph of treewidth $\Theta(r)$. In particular, an important result of Robertson, Seymour, and Thomas [38] is that every planar graph of treewidth w has an $\Omega(w) \times \Omega(w)$ grid graph as a minor. Thus every planar graph of large treewidth has a grid minor certifying that its treewidth is almost as large (up to constant factors).

In their Graph Minor Theory, Robertson and Seymour [36] have generalized this result in some sense to any graph excluding a fixed minor: for every graph H and integer $r > 0$, there is an integer $w > 0$ such that every H-minor-free graph with treewidth at least w has an $r \times r$ grid graph as a minor. This result has been re-proved by Robertson, Seymour, and Thomas [38], Reed [34], and Diestel, Jensen, Gorbunov, and Thomassen [22]. The best known bound on w in terms of r is as follows:

Theorem 1 ([38, Theorem 5.8]) *Every H-minor-free graph of treewidth larger than $20^{5|V(H)|^3 r}$ has an $r \times r$ grid as a minor.*

While the existence of such a relationship between treewidth and grid minors is interesting, this bound of $w = 2^{O(r)}$ is much weaker than the bound of $w = O(r)$ attainable for the special case of planar graphs. In particular, the grid they obtain from this theorem can have treewidth logarithmic in the treewidth of the original graph, which does not serve as much of a certificate of

*A preliminary version of this paper appeared in the ACM-SIAM Symposium on Discrete Algorithms (SODA 2005) [16].

[†]MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar Street, Cambridge, MA 02139, U.S.A., {edemaine, hajiagha}@mit.edu

[1]The $r \times r$ *grid* is the planar graph with r^2 vertices arranged on a square grid and with edges connecting horizontally and vertically adjacent vertices. Refer to Section 2 for other (standard) definitions and graph terminology.

large treewidth as we have for planar graphs. The main result of this paper is the following much tighter bound:

Theorem 2 *For any fixed graph H, every H-minor-free graph of treewidth w has an $\Omega(w) \times \Omega(w)$ grid as a minor.*

Thus the $r \times r$ grid is the canonical H-minor-free graph of treewidth $\Theta(r)$ for any fixed graph H. This result is best possible up to constant factors. Section 5 discusses the dependence of the constant factor in the Ω notation on the fixed graph H.

Our result cannot be generalized to arbitrary graphs: Robertson, Seymour, and Thomas [38] proved that some graphs have treewidth $\Omega(r^2 \lg r)$ but have grid minors only of size $O(r) \times O(r)$. The best known relation for general graphs is that having treewidth more than 20^{2r^5} implies the existence of an $r \times r$ grid minor [38]. The best possible bound is believed to be closer to $\Theta(r^2 \lg r)$ than $2^{\Theta(r^5)}$, perhaps even equal to $\Theta(r^2 \lg r)$ [38].

Our approach in the proof of Theorem 2 can be viewed more generally as a framework for extending combinatorial results on planar graphs to hold on H-minor-free graphs for any fixed H. The framework follows three main steps: extension from planar graphs to bounded-genus graphs, further extension to "almost-embeddable graphs", and further extension to clique sums of almost-embeddable graphs. Almost-embeddable graphs are bounded-genus graphs except for a bounded number of "local areas of non-planarity", called vortices, and for a bounded number of "apex" vertices, which can have any number of incident edges that are not properly embedded. The underpinnings of this framework is the structural characterization of H-minor-free graphs in the Robertson-Seymour Graph Minor Theory [37]. Recently this framework has been used to generalize many efficient algorithms from planar graphs to H-minor-free graphs [12, 27]. Our work shows how the framework can be applied to combinatorial results.

In addition to giving a tight bound on this basic combinatorial problem relating treewidth and grids, our result has many combinatorial consequences, each with several algorithmic consequences. To describe these consequences we first need to introduce the concept of bidimensionality.

Bidimensionality. The genesis of bidimensionality is the notion of a parameter-treewidth bound. A *parameter* $P = P(G)$ is a function mapping graphs to nonnegative integers. A *parameter-treewidth bound* is an upper bound $f(k)$ on the treewidth of a graph with parameter value k. In many cases, $f(k)$ can even be shown to be sublinear in k, often $O(\sqrt{k})$. Parameter-treewidth bounds have been established for many parameters; see e.g. [1, 30, 25, 4, 8, 32, 28, 11, 18, 20, 21, 10, 14, 12]. Essentially all of these bounds can be obtained from the general theory of bidimensional parameters, which has been introduced in a series of papers [21, 11, 12, 10]. Thus bidimensionality is the most powerful method so far for establishing parameter-treewidth bounds, encompassing all such previous results for H-minor-free graphs.

For a strictly increasing function g, a parameter is *bidimensional* if it is at least $g(r)$ in an $r \times r$ "grid-like graph" and if the parameter does not increase when taking either minors (*minor-bidimensional*) or contractions (*contraction-bidimensional*). Examples of bidimensional parameters include number of vertices, diameter, and the size of various structures, e.g., feedback vertex set, vertex cover, minimum maximal matching, face cover, a series of vertex-removal parameters, dominating set, edge dominating set, r-dominating set, connected dominating set, connected edge dominating set, connected r-dominating set, and unweighted TSP tour (a walk in the graph visiting all vertices). Parameter-treewidth bounds have been established for all minor-bidimensional parameters in H-minor-free graphs for any fixed graph H [12, 10]. In this case, the notion of "grid-like graph" is precisely the $r \times r$ grid. For contraction-bidimensional parameters, parameter-treewidth bounds have been established for apex-minor-free graphs, and this is the largest class of graphs for which such bounds can generally be obtained [10]. (An *apex-minor-free* graph family is a minor-closed graph family excluding some *apex graph*, i.e., a graph in which the removal of some vertex leaves a planar graph.) In this case, the notion of "grid-like graph" is an $r \times r$ grid *augmented* with additional edges such that each vertex is incident to $O(1)$ edges to nonboundary vertices of the grid. (Here $O(1)$ depends on H.)

Unfortunately, these parameter-treewidth bounds are large in general: $f(k) = (g^{-1}(k))^{O(g^{-1}(k))}$. For the special cases of single-crossing-minor-free graphs and bounded-genus graphs, we know tighter bounds of $f(k) = O(g^{-1}(k))$, which is the best possible bound up to constant factors. For single-crossing-minor-free graphs [21, 18] (in particular, planar graphs [11]), the notion of "grid-like graph" is an $r \times r$ grid partially triangulated by additional edges that preserve planarity. For bounded-genus graphs [19], the notion of "grid-like graph" is such a partially triangulated $r \times r$ grid with up to genus(G) additional edges ("handles"). (The same result was established for a subset of contraction-bidimensional parameters, called α-splittable parameters, previously in [12].)

Tight parameter-treewidth bounds. One consequence of our result gives the tightest possible parameter-treewidth bound for all bidimensional parameters in all possible H-minor-free graphs:

Theorem 3 *For any minor-bidimensional parameter P which is at least $g(r)$ in the $r \times r$ grid, every H-minor-free graph G has treewidth $\mathrm{tw}(G) = O(g^{-1}(P(G)))$. For any contraction-bidimensional parameter P which is at least $g(r)$ in an augmented $r \times r$ grid, every apex-minor-free graph G has treewidth $\mathrm{tw}(G) = O(g^{-1}(P(G)))$. In particular, if $g(r) = \Theta(r^2)$, then these bounds become $\mathrm{tw}(G) = O(\sqrt{P(G)})$.*

The proof of this theorem is identical to the proofs of [10, Theorem 2.3] (for minor-bidimensional parameters) and [10, Theorem 4.1] (for contraction-bidimensional parameters) except that we substitute the application of Theorem 1 with Theorem 2.

Separator theorems. If we apply the parameter-treewidth bound of Theorem 3 to the parameter of the number of vertices in the graph, which is minor-bidimensional with $g(r) = r^2$, then we immediately obtain the following (known) bound on the treewidth of an H-minor-free graph:

Corollary 1 ([5, Proposition 4.5], [27, Corollary 24]) *For any fixed graph H, every H-minor-free graph G has treewidth $O(\sqrt{|V(G)|})$.*

A consequence of this result is that every vertex-weighted H-minor-free graph G has a vertex separator of size $O(\sqrt{|V(G)|})$ whose removal splits the graph into two parts each with weight at most 2/3 of the original weight [5, Theorem 1.2]. This generalization of the classic planar separator theorem has many algorithmic applications; see e.g. [5, 3]. Also, this result shows that the Robertson-Seymour characterization of H-minor-free graphs is powerful enough to conclude that these graphs have small separators, which we expect from such a strong characterization.

Bounded local treewidth (diameter treewidth). Eppstein [23] introduced the *diameter-treewidth property* for a class of graphs, which requires that the treewidth of a graph in the class is upper bounded by a function of its diameter. He proved that a minor-closed graph family has the diameter-treewidth property precisely if the graph family excludes some apex graph. In particular, he proved that any graph in such a family with diameter D has treewidth at most $2^{2^{O(D)}}$. (A simpler proof of this result was obtained in [13].)

If we apply the parameter-treewidth bound of Theorem 3 to the diameter parameter, which is contraction-bidimensional with $g(r) = \Theta(\lg r)$ [13], then we immediately obtain the following stronger diameter-treewidth bound for apex-minor-free graphs:

Corollary 2 *For any fixed apex graph H, every H-minor-free graph of diameter D has treewidth $2^{O(D)}$.*

The diameter-treewidth property has been used extensively in a slightly modified form called the *bounded-local-treewidth property*, which requires that the treewidth of any connected subgraph of a graph in the class is upper bounded by a function of its diameter. For minor-closed graph families, which is the focus of most work in this context, these properties are identical. Graphs of bounded local treewidth have many similar properties to both planar graphs and graphs of bounded treewidth, two classes of graphs on which many problems are substantially easier. In particular, Baker's approach for polynomial-time approximation schemes (PTASs) on planar graphs [7] applies to this setting. As a result, PTASs are known for hereditary maximization problems such as maximum independent set, maximum triangle matching, maximum H-matching, and maximum tile salvage; for minimization problems such as minimum vertex cover, minimum dominating set, minimum edge-dominating set; and for subgraph isomorphism for a fixed pattern [18, 23, 29]. Graphs of bounded local treewidth also admit several efficient fixed-parameter algorithms. In particular, Frick and Grohe [26] give a general framework for deciding any property expressible in first-order logic in graphs of bounded local treewidth. Corollary 2 substantially improves the running time of these algorithms, in particular improving the running time of the PTASs from $2^{2^{2^{O(1/\varepsilon)}}} n^{O(1)}$ to $2^{2^{O(1/\varepsilon)}} n^{O(1)}$, where n is the number of vertices in the graph.

Subexponential fixed-parameter algorithms. A *fixed-parameter algorithm* is an algorithm for computing a parameter $P(G)$ of a graph G whose running time is $h(P(G))\, n^{O(1)}$ for some function h. A typical function h for many fixed-parameter algorithms is $h(k) = 2^{O(k)}$. In the last three years, several researchers have obtained *exponential speedups* in fixed-parameter algorithms in the sense that the h function reduces exponentially, e.g., to $2^{O(\sqrt{k})}$. For example, the first fixed-parameter algorithm for finding a dominating set of size k in planar graphs [2] has running time $O(8^k n)$; subsequently, a sequence of subexponential algorithms and improvements have been obtained, starting with running time $O(4^{6\sqrt{34k}} n)$ [1], then $O(2^{27\sqrt{k}} n)$ [30], and finally $O(2^{15.13\sqrt{k}} k + n^3 + k^4)$ [25]. Other subexponential algorithms for other domination and covering problems on planar graphs have also been obtained [1, 4, 8, 32, 28].

All subexponential fixed-parameter algorithms developed so far are based on showing a sublinear parameter-treewidth bound and then using an algorithm whose running time is singly exponential in treewidth and polynomial in problem size. As mentioned above, essentially all sublinear treewidth-parameter bounds proved so far can be obtained through bidimensionality. From Theorem 3 we obtain the following general result for designing subexponential fixed-parameter algorithms:

Corollary 3 *Consider a parameter P that can be computed on a graph G in $h(w)\, n^{O(1)}$ time given a tree decomposition of G of width at most w. If P is minor-bidimensional and at least $g(r)$ in the $r \times r$ grid, then there is an algorithm computing P on any H-minor-free graph G with running time $[h(O(g^{-1}(k))) + 2^{O(g^{-1}(k))}]\, n^{O(1)}$. If P is contraction-bidimensional and at least $g(r)$ in an augmented $r \times r$ grid, then there is an algorithm computing P on any apex-minor-free graph G with* *running time* $[h(O(g^{-1}(k))) + 2^{O(g^{-1}(k))}]\, n^{O(1)}$. *In particular, if* $g(r) = \Theta(r^2)$ *and* $h(w) = 2^{o(w^2)}$, *then these running times are subexponential in* k.

The proof of this corollary is identical to the proof of [10, Theorem 5.1] except that we apply the stronger parameter-treewidth bound of Theorem 3. In particular, this corollary gives subexponential fixed-parameter algorithms for many bidimensional parameters, including feedback vertex set, vertex cover, minimum maximal matching, a series of vertex-removal parameters, dominating set, edge dominating set, r-dominating set, clique-transversal set, connected dominating set, connected edge dominating set, connected r-dominating set, and unweighted TSP tour.

Approximation schemes. The bidimensionality theory has recently been extended to obtain PTASs for essentially all bidimensional parameters (including those mentioned above) in planar graphs and some generalizations [15]. These PTASs are based on techniques that generalize and in some sense unify the two main previous approaches for designing PTASs in planar graphs, namely, the Lipton-Tarjan separator approach [33] and the Baker layerwise decomposition approach [7]. However, these PTASs require a linear parameter-treewidth bound as in Theorem 3, so previously only applied to single-crossing-minor-free and bounded-genus graphs. Theorem 3 generalizes these results to all H-minor-free graphs for minor-bidimensional parameters and to all apex-minor-free graphs for contraction-bidimensional parameters. This result shows a strong connection between subexponential fixed-parameter tractability and approximation algorithms for combinatorial optimization problems on H-minor-free graphs.

Half-integral versus fractional multicommodity flow. Chekuri, Khanna, and Shephard [9] proved that, for planar graphs, the gap between the optimal half-integral multicommodity flow and the optimal fractional multicommodity flow is at most a polylogarithmic factor. Also, they gave a combinatorial proof of the result that, for planar graphs, the gap between the maximum flow and the minimum cut in product multicommodity flow (and thus uniform multicommodity flow) instances is at most a constant factor. The latter result was proved before by Klein, Plotkin, and Rao for H-minor-free graphs using primal-dual methods [31], and has many applications in embeddings of H-minor-free graphs. As mentioned by Chekuri et al. [9], our Theorem 2 can be used to generalize the half-integral/fractional gap bound and the combinatorial proof of the max-flow/min-cut gap bound to H-minor-free graphs.

Approximating treewidth and grid minors. Most recently, our linear minimax relation between treewidth and the largest grid minor has been used to develop a combinatorial "primal-dual" type algorithm, which approximates either quantity within a constant factor [17]. This algorithm also computes, in polynomial time, a tree decomposition of approximately minimum width and a grid minor of approximately maximum size.

2 Background

2.1 Preliminaries

Our graph terminology is as follows. All the graphs in this paper are undirected without loops or multiple edges. A graph G is represented by $G = (V, E)$, where V (or $V(G)$) is the set of vertices and E (or $E(G)$) is the set of edges. We denote an edge e between u and v by $\{u, v\}$. The *(disjoint)* *union* of two disjoint graphs G_1 and G_2, $G_1 \cup G_2$, is the graph G with merged vertex and edge sets: $V(G) = V(G_1) \cup V(G_2)$ and $E(G) = E(G_1) \cup E(G_2)$.

One way of describing classes of graphs is by using *minors*, introduced as follows. *Contracting* an edge $e = \{u, v\}$ is the operation of replacing both u and v by a single vertex w whose neighbors are all vertices that were neighbors of u or v, except u and v themselves. A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges. A graph class $\mathcal{C}$ is a *minor-closed* class if any minor of any graph in $\mathcal{C}$ is also a member of $\mathcal{C}$. A minor-closed graph class $\mathcal{C}$ is H*-minor-free* if $H \notin \mathcal{C}$. For example, the class of planar graphs is both $K_{3,3}$-minor-free and K_5-minor-free.

2.2 Treewidth

The notion of treewidth was introduced by Robertson and Seymour [35] and plays an important role in their fundamental work on graph minors. To define this notion, we first consider a representation of a graph as an underlying tree, called a tree decomposition. More precisely, a *tree decomposition* of a graph $G = (V, E)$ is a pair (T, χ) in which $T = (I, F)$ is a tree and $\chi = \{\chi_i \mid i \in I\}$ is a family of subsets of $V(G)$ such that

1. $\bigcup_{i \in I} \chi_i = V$;
2. for each edge $e = \{u, v\} \in E$, there exists an $i \in I$ such that both u and v belong to χ_i; and
3. for all $v \in V$, the set of nodes $\{i \in I \mid v \in \chi_i\}$ forms a connected subtree of T.

To distinguish between vertices of the original graph G and vertices of T in the tree decomposition, we call vertices of T *nodes* and their corresponding χ_i's *bags*. The maximum size of a bag in χ minus one is called the *width* of the tree decomposition. The *treewidth* of a graph G ($\mathrm{tw}(G)$) is the minimum width over all possible tree decompositions of G. A tree decomposition is called a *path decomposition* if $T = (I, F)$ is a path. The *pathwidth* of a graph G ($\mathrm{pw}(G)$) is the minimum width over all possible path decompositions of G.

We will need the following property about how treewidth changes during small operations to faces of a graph:

Lemma 1 *Consider any graph G embedded in some surface of genus g, with $\mathrm{tw}(G) \geq 56(g+1)^2$. If G' is the result of contracting a face of G to a point, then $\mathrm{tw}(G') \leq \mathrm{tw}(G)$ and $\mathrm{tw}(G') = \Omega(\mathrm{tw}(G)/(g+1))$.*

Proof: Let f denote the face of G contracted to form G'. Because G' is a minor of G, $\mathrm{tw}(G') \leq \mathrm{tw}(G)$. Consider the graph G'' formed from graph G by adding a new vertex v in the middle of face f and adding an edge connecting v to every vertex of f. This graph G'' is embedded in the same genus-g surface as G. The treewidth of G'' is between $\mathrm{tw}(G)$ and $\mathrm{tw}(G) + 1$, because we could add v to all bags of a tree decomposition of G. By [19, Theorem 4.8], there is a sequence of contractions that brings G'' to a graph R that is a (planar) partially triangulated $r \times r$ grid augmented with at most g additional edges, where $r \geq \frac{1}{4}\mathrm{tw}(G'')/(g+1) - 12\,g - 1 \geq \frac{1}{4}\mathrm{tw}(G)/(g+1) - 12\,g - 1$. Every vertex in R can be labeled by the set of vertices in G'' that were contracted to form it. Let v_R denote the vertex in R whose label includes v. For every neighbor w of v in G, the vertex w_R in R whose label includes w has distance at most 1 from v_R in R, because contractions only decrease distances. We modify the augmented partially triangulated grid R as follows. For every neighbor w of v in G for which $w_R \neq v_R$, we delete all edges incident to w_R except $\{v_R, w_R\}$, and then we contract the edge $\{v_R, w_R\}$. The resulting graph R' is a minor of R and thus of G''. If we re-order the sequence of contractions and deletions that bring G'' to R' to start with the contractions of the edges between v and the vertices of face f (which is equivalent to contracting the face f in the original graph G), then the succeeding sequence of contractions and deletions brings G' to R'. Therefore R' is a minor of G'. By ignoring every row or column of the grid in which an edge was deleted, we obtain an $(r-g-4)\times(r-g-4)$ grid minor of G'. (There may be g such rows (resp., columns) from the g additional edges, 2 from the neighborhood of v, and 2 from the boundary of the grid.) Therefore $\mathrm{tw}(G') \geq r-g-5 \geq \frac{1}{4}\mathrm{tw}(G)/(g+1) - 13\,g - 6 = \Omega(\mathrm{tw}(G)/(g+1))$ provided $\mathrm{tw}(G) \geq 56(g+1)^2$. $\square$

2.3 Clique Sums

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 $G_i[W_i]$ with both endpoints in 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.

The following lemma shows how the treewidth changes when we apply a clique-sum operation, which plays an important role in our results.

Lemma 2 ([18, Lemma 3]) *For any two graphs G and H, $\mathrm{tw}(G \oplus H) \leq \max\{\mathrm{tw}(G), \mathrm{tw}(H)\}$.*

2.4 Clique-Sum Decompositions of H-Minor-Free Graphs

Our result uses the deep theorem of Robertson and Seymour on graphs excluding a fixed graph as a minor [37]. Intuitively, their theorem says that, for every graph H, every H-minor-free graph can be expressed as a "tree structure" of pieces, where each piece is a graph that can be drawn in a surface in which H cannot be drawn, except for a bounded number of "apex" vertices and a bounded number of "local areas of non-planarity" called *vortices*. Here the bounds depend only on H.

Roughly speaking we say a graph G is *h-almost embeddable* in a surface S if there exists a set X of size at most h of vertices, called *apex vertices* or *apices*, such that $G - X$ can be obtained from a graph G_0 embedded in S by attaching at most h graphs of pathwidth at most h to G_0 within h faces in an orderly way. More precisely:

Definition 1 *A graph G is h-*almost embeddable* in S if there exists a vertex set X of size at most h called* apices *such that $G - X$ can be written as $G_0 \cup G_1 \cup \cdots \cup G_h$, where*

1. *G_0 has an embedding in S;*
2. *the graphs $G_1, \ldots, G_h$, called* vortices*, are pairwise disjoint;*
3. *there are faces $F_1, \ldots, F_h$ of G_0 in S, and there are pairwise disjoint disks $D_1, \ldots, D_h$ in S, such that for $i = 1, \ldots, h$, $D_i \subset F_i$ and $U_i := V(G_0) \cap V(G_i) = V(G_0) \cap D_i$; and*
4. *the graph G_i has a path decomposition with bags $(\mathcal{B}_u)_{u \in U_i}$ of width less than h, such that $u \in \mathcal{B}_u$ for all $u \in U_i$. The sets $\mathcal{B}_u$ are ordered by the ordering of their indices u as points along the boundary cycle of face F_i in G_0.*

An h-almost-embeddable graph is called apex-free *if there exists such an embedding in which the set X of apices is empty.*

Now, the deep result of Robertson and Seymour is as follows.

Theorem 4 ([37]) *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 that are h-almost-embeddable in some surfaces in which H cannot be embedded.*

In particular, if H is fixed, any surface in which H cannot be embedded has bounded genus. Thus, the summands in the theorem are h-almost-embeddable in bounded-genus surfaces.

3 Overview of Proof of Main Theorem

The proof of our main theorem (Theorem 2) is based on a series of reductions. Each reduction converts a given graph into a "simpler" graph whose treewidth is $\Omega(\mathrm{tw}(G))$.

The first reduction applies Theorem 4 to the original graph G, decomposing it into a clique sum of almost-embeddable graphs. By Lemma 2, at least one summand in this clique sum has treewidth at least $\mathrm{tw}(G)$. Therefore we can focus on this single summand of large treewidth. However, we note that this summand may not be a minor of G, and therefore it is not enough to prove that the summand has a large grid as a minor; we must deal with this issue later in the proof.

The second, trivial reduction is to remove the apices from the almost-embeddable graph. This reduction changes the treewidth by at most an additive constant. Now our almost-embeddable graph is apex-free.

The third reduction effectively removes the vortices from the apex-free almost-embeddable graph. This reduction uses that vortices have small pathwidth to conclude that the treewidth remains roughly the same. At this point the graph has bounded genus, because we have removed both apices and vortices.

Because the graph has bounded genus, it has a large grid as a minor. However, this grid is not useful: the graph is not necessarily a minor of the original graph G because, during the clique-sum decomposition, we may have introduced extra edges when the join set was completed into a clique. We call such edges *virtual edges*, and all other edges *actual edges*. One difficulty of Theorem 4 is that it does not guarantee that the virtual edges can be obtained by taking a minor of the original graph G, and therefore the pieces may not be minors of G. The fourth reduction overcomes this difficulty by obtaining some virtual edges by taking minors of the original graph G, and removes other virtual edges which cannot be obtained, while still preserving the treewidth up to constant factors. We call the resulting graph an *approximation graph*.

The approximation graph is both a minor of G and has bounded genus. Now we use the fact that a bounded-genus graph with treewidth w has an $\Omega(w) \times \Omega(w)$ grid as a minor. Therefore both the approximation graph and G have such a grid as a minor.

4 Proof of Main Theorem

In this section we prove Theorem 2.

First we apply Theorem 4 to the original graph G, decomposing it into a clique sum of almost-embeddable graphs.

Lemma 3 *At least one summand in the clique sum has treewidth at least* $\mathrm{tw}(G)$.

Proof: Immediate by Lemma 2. □

Let G' denote a summand in the clique sum with $\mathrm{tw}(G') \geq \mathrm{tw}(G)$. For every vertex v in G', there is a corresponding vertex $f(v)$ in G by following the definition of clique sum. Each edge $\{u, v\}$ in G' may or may not have a corresponding edge $\{f(u), f(v)\}$ in G. If the edge $\{f(u), f(v)\}$ exists in G, we say that $\{u, v\}$ is an *actual edge* in G'; otherwise, it is a *virtual edge* in G'. Virtual edges arise from removing edges from the join set during a clique sum.

Because G' is h-almost-embeddable in some bounded-genus surface, it consists of a bounded-genus graph augmented by at most h vortices and at most h apices. We remove all apices from G' to produce an apex-free h-almost-embeddable graph G''. Because adding a vertex and any collection of incident edges to a graph can increase the treewidth by at most 1, we have the following relation between the treewidths of G' and G'':

Lemma 4 $\mathrm{tw}(G'') \geq \mathrm{tw}(G') - h$.

Next we remove all vortices from G''. Let G''_0 denote the bounded-genus part of the apex-free h-almost-embeddable graph G'', and let U_i denote the set of vertices at which vortex i is attached to G''_0 (as in Definition 1). Define $G''' = G''_0 - U_1 - U_2 - \cdots - U_h$, i.e., G''' is the result of removing all vertices from vortices in G''.

Lemma 5 $\mathrm{tw}(G''') = O(\mathrm{tw}(G''))$ *for* $h = O(1)$.

Proof: Suppose G'' decomposes into $G''_0 \cup G''_1 \cup G''_2 \cup \cdots \cup G''_h$ where each G''_i, $i \geq 1$, is a vortex as in Definition 1. Define an intermediate graph $\hat{G}$ as follows. Let $U_i = \{u_i^1, u_i^2, \ldots, u_i^{m_i}\}$ be the cyclically ordered vertices of G''_0 at which vortex G''_i is attached. We obtain $\hat{G}$ by starting from G''_0 and adding edges $\{u_i^j, u_i^{j+1}\}$ where they do not already exist, and where $j + 1$ is treated modulo m_i, for each $1 \leq i \leq h$ and each $1 \leq j \leq m_i$. Because we only added a planar graph within the face corresponding to U_i, $\hat{G}$ is embeddable in the same bounded-genus surface as G''_0.

We claim that $\mathrm{tw}(G'') \leq (h+1)^2\,(\mathrm{tw}(\hat{G})+1)$. Consider some minimum-width tree decomposition of $\hat{G}$, and consider each bag $\mathcal{B}$ of that tree decomposition. For each u_i^j that occurs in bag $\mathcal{B}$, we add to $\mathcal{B}$ the corresponding bag $\mathcal{B}_{u_i^j}$ from the path decomposition of vortex G''_i. The resulting bags form a tree decomposition of G'' because $\{u_i^1, u_i^2, \ldots, u_i^{m_i}\}$ are connected in a path in $\hat{G}$. By charging the $\leq h+1$ added vertices to the occurrence of u_i^j that triggered the addition, each bag increases in size by a factor at most $h+1$ for each of the h vortices. Thus the width of this tree decomposition of G'' is at most $(h(h+1))\,(\mathrm{tw}(\hat{G})+1) - 1$, which is stronger than the desired claim.

Let $\hat{\hat{G}}$ be the graph resulting from $\hat{G}$ by contracting the face $\{u_i^1, u_i^2, \ldots, u_i^{m_i}\}$ in $\hat{G}$ into a single vertex, for each i. Applying Lemma 1, h times, we obtain $\mathrm{tw}(\hat{\hat{G}}) = \Omega(\mathrm{tw}(\hat{G}))$ because h and the genus of the surface in which $\hat{G}$ is embedded are $O(1)$. Therefore $\mathrm{tw}(G'') = O(\mathrm{tw}(\hat{\hat{G}}))$.

Finally we delete each contracted vertex in $\hat{\hat{G}}$, which results in G'''. Thus $\mathrm{tw}(G''') \geq \mathrm{tw}(\hat{\hat{G}}) - h$, so $\mathrm{tw}(G'') = O(\mathrm{tw}(G'''))$ as desired. □

A similar technique to the proof of Lemma 5 has been used by others, e.g., [27, 12].

At this point the graph has bounded genus, because we have removed both apices and vortices. In the next step we deal with virtual edges. Intuitively, for each summand G' in the clique-sum decomposition of the original graph G, we construct a graph $\tilde{G}$ which is a minor of G and "approximately" preserves the virtual edges within G'.

For this step we need an additional property of the clique-sum decomposition obtained in the proof of Theorem 4: each clique sum involves at most three vertices from each summand other than apices and vertices in vortices of that summand. As observed by Seymour [39], this stronger form of Theorem 4 follows from exactly the same proof in [37]. Also, an explicit proof of this stronger form is given in [17, Theorem 2.2], which also proves an algorithmic version of Theorem 4.

Definition 2 *Let G' be an h-almost-embeddable graph in a clique-sum decomposition of a graph G arising from Theorem 4. The* approximation graph *of G', denoted by $\tilde{G}$, is obtained by starting from G''', removing the virtual edges, and replacing some of them as follows. In the clique-sum decomposition of G, for each clique sum involving G' with the property that the join set W has $|W \cap V(G''')| > 1$, we do the following:*

1. *If $|W \cap V(G''')| = 2$, we add an edge between these two vertices.*
2. *If $|W \cap V(G''')| = 3$ and there is more than one clique sum that contains $W \cap V(G''')$ in its join set, we add all edges between pairs of vertices in $W \cap V(G''')$.*
3. *If $|W \cap V(G''')| = 3$ and there is only one clique sum that contains $W \cap V(G''')$ in its join set, we add a new vertex v inside the triangle of $W \cap V(G''')$ on the surface and then add an edge connecting v to each vertex of $W \cap V(G''')$.*

Lemma 6 *Let G' be an h-almost-embeddable graph in a clique-sum decomposition of a graph G arising from Theorem 4. The approximation graph $\tilde{G}$ of G' is a minor of G and can be embedded in the same surface as the bounded-genus part of G'.*

Proof: First, G''' with all virtual edges removed is a minor of G, because the former graph can be constructed from G by deleting all vertices not in the summand G' and deleting all apices and vertices in vortices in G'. All that remains to show is that the edges added in Cases 1–3 of Definition 2 can also be formed as a minor of G. We use the (trivial) additional property of the clique-sum decomposition arising from the proof of Theorem 4 that each summand in the clique sum is connected even after removal of the join set. (If a summand were not connected after the removal of the join set, we could rewrite the initial clique-sum decomposition by splitting the summand into a clique sum of these pieces.) Now, for each clique sum between G' and F with the property that the join set W has $|W \cap V(G''')| > 1$, we contract F down to a single vertex v adjacent to all vertices in the join set. In Case 3, this vertex v is precisely the desired vertex v inside the triangle $W \cap V(G''')$. This triangle is guaranteed to be empty in the bounded-genus part of G' in the clique-sum decomposition arising from Theorem 4; if this were not the case, again we could rewrite the clique-sum decomposition by splitting G' into a clique sum of two pieces. Thus the resulting graph can be embedded in the same surface as the bounded-genus part of G'. In the other two cases, we contract v into a vertex of $W \cap V(G''')$—in Case 2, we contract two different v's into two different vertices of $W \cap V(G''')$—and obtain the additional edges added to $\tilde{G}$. Finally, we delete the apices and vertices in vortices in G', and delete any other summands that had $|W \cap V(G''')| \leq 1$. In the end we have contracted and deleted edges in G to obtain precisely $\tilde{G}$. $\square$

Lemma 7 $\mathrm{tw}(\tilde{G}) \geq \frac{1}{3}(\mathrm{tw}(G''') + 1) - 1$.

Proof: To prove that $\mathrm{tw}(G''') \leq 3(\mathrm{tw}(\tilde{G}) + 1) - 1$, we start from a minimum-width tree decomposition of $\tilde{G}$ and convert it into a tree decomposition of G'''. We need only consider Case 3 in Definition 2 because otherwise $\tilde{G}$ is identical to G'''. For each occurrence of an added vertex v from Case 3 in a bag $\mathcal{B}$ in the tree decomposition of $\tilde{G}$, we replace v in $\mathcal{B}$ with all three vertices from $W \cap V(G''')$. The result is a tree decomposition of G''' where each bag has increased in size by at most a factor of 3. $\square$

By Lemma 6, the approximation graph $\tilde{G}$ is both a minor of G and has bounded genus. By [12, Theorem 3.5], every bounded-genus graph with treewidth $\Omega(r)$ has an $r \times r$ grid as a minor. By Lemmas 3, 4, 5, and 7, $\mathrm{tw}(\tilde{G}) = \Omega(\mathrm{tw}(G))$. Therefore $\tilde{G}$ and thus G have an $\Omega(\mathrm{tw}(G)) \times \Omega(\mathrm{tw}(G))$ grid as a minor. This concludes the proof of Theorem 2.

5 Conclusion and Further Remarks

We have shown that every graph excluding a fixed minor has a grid minor whose treewidth is within a constant factor of the graph's treewidth. Such a tight connection has many combinatorial and algorithmic applications through the theory of bidimensionality. These applications suggest two directions for improvement and generalization.

First, the constant factor we obtain is likely not the best possible. The dependence of the factor on H is of particular interest because it can severely affect the running time of algorithms based on this result. The factor must be $\Omega(\sqrt{|V(H)|} \lg |V(H)|)$, because otherwise such a bound would contradict the lower bound for general graphs. An upper bound near this lower bound (in particular, polynomial in $|V(H)|$) is not out of the question: the bound on the size of separators in [5] has a lead factor of $|V(H)|^{3/2}$. In fact, Alon, Seymour, and Thomas [5] suspect that the correct factor for separators is $\Theta(|V(H)|)$, which holds e.g. in bounded-genus graphs. We also suspect that the same bound holds for the factor in Theorem 2, which would imply the corresponding bound for separators.

Second, it would be interesting to determine the tightest possible relation between treewidth and grid minors in general graphs. This problem was posed by Robertson, Seymour, and Thomas [38]; the answer is that the treewidth must be somewhere between $\Theta(r^2 \lg r)$ and $2^{\Theta(r^5)}$. A bound closer to $\Theta(r^2 \lg r)$ might result in efficient algorithms for computing minor-bidimensional parameters in general graphs.

Finally, it would be interesting to obtain a constant-factor (polynomial-time) approximation algorithm for treewidth in H-minor-free graphs for a fixed H, through the framework used in this paper. Constant-factor approximation algorithms for treewidth are known for planar graphs [40], single-crossing-minor-free graphs [18]. Using a different approach, a constant-factor approximation was recently obtained for H-minor-free graphs for a fixed H [24]. For general graphs, the best known approximation ratio was $O(\lg \mathrm{tw}(G))$ [6], until the recent $O(\sqrt{\lg \mathrm{tw}(G)})$-approximation of [24]. These approximation algorithms have recently been used in fixed-parameter and approximation algorithms; see e.g. [18, 10, 15].

Acknowledgments

The authors thank Paul D. Seymour for many helpful discussions and for providing a portal into the Graph Minor Theory and revealing some of its hidden structure that we use in this paper.

References

[1] J. Alber, H. L. Bodlaender, H. Fernau, T. Kloks, and R. Niedermeier, *Fixed parameter algorithms for dominating set and related problems on planar graphs*, Algorithmica, 33 (2002), pp. 461–493.

[2] J. Alber, H. Fan, M. R. Fellows, H. Fernau, R. Niedermeier, F. A. Rosamond, and U. Stege, *Refined search tree technique for DOMINATING SET on planar graphs*, in

Proceedings of the 26th International Symposium on Mathematical Foundations of Computer Science, vol. 2136 of Lecture Notes in Computer Science, 2001, pp. 111–122.

[3] J. Alber, H. Fernau, and R. Niedermeier, *Graph separators: a parameterized view*, Journal of Computer and System Sciences, 67 (2003), pp. 808–832.

[4] J. Alber, H. Fernau, and R. Niedermeier, *Parameterized complexity: exponential speed-up for planar graph problems*, Journal of Algorithms, 52 (2004), pp. 26–56.

[5] N. Alon, P. Seymour, and R. Thomas, *A separator theorem for nonplanar graphs*, Journal of the American Mathematical Society, 3 (1990), pp. 801–808.

[6] E. Amir, *Efficient approximation for triangulation of minimum treewidth*, in Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence (UAI-2001), San Francisco, CA, 2001, Morgan Kaufmann Publishers, pp. 7–15.

[7] B. S. Baker, *Approximation algorithms for NP-complete problems on planar graphs*, Journal of the Association for Computing Machinery, 41 (1994), pp. 153–180.

[8] M.-S. Chang, T. Kloks, and C.-M. Lee, *Maximum clique transversals*, in Proceedings of the 27th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2001), vol. 2204 of Lecture Notes in Computer Science, 2001, pp. 32–43.

[9] C. Chekuri, S. Khanna, and F. B. Shepherd, *Edge-disjoint paths in planar graphs*, in Proceedings of the 45th Symposium on Foundations of Computer Science (FOCS 2004), 2004, pp. 71–80.

[10] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos, *Bidimensional parameters and local treewidth*, SIAM Journal on Discrete Mathematics, 18 (2004), pp. 501–511.

[11] ——, *Fixed-parameter algorithms for (k, r)-center in planar graphs and map graphs*, ACM Transactions on Algorithms, 1 (2005), pp. 33–47.

[12] ——, *Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs*, Journal of the ACM, 52 (2005), pp. 866–893.

[13] E. D. Demaine and M. Hajiaghayi, *Diameter and treewidth in minor-closed graph families, revisited*, Algorithmica, 40 (2004), pp. 211–215.

[14] ——, *Equivalence of local treewidth and linear local treewidth and its algorithmic applications*, in Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA'04), January 2004, pp. 833–842.

[15] E. D. Demaine and M. Hajiaghayi, *Bidimensionality: New connections between FPT algorithms and PTASs*, in Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005), Vancouver, January 2005, pp. 590–601.

[16] ——, *Graphs excluding a fixed minor have grids as large as treewidth, with combinatorial and algorithmic applications through bidimensionality*, in Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005), Vancouver, January 2005, pp. 682–689.

[17] 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, Pittsburgh, PA, October 2005, pp. 637–646. http://theory.csail.mit.edu/~edemaine/papers/Decomposition_FOCS2005/

[18] E. D. Demaine, M. Hajiaghayi, N. Nishimura, P. Ragde, and D. M. Thilikos, *Approximation algorithms for classes of graphs excluding single-crossing graphs as minors*, Journal of Computer and System Sciences, 69 (2004), pp. 166–195.

[19] E. D. Demaine, M. Hajiaghayi, and D. M. Thilikos, *The bidimensional theory of bounded-genus graphs*, SIAM Journal on Discrete Mathematics. To appear.

[20] ——, *A 1.5-approximation for treewidth of graphs excluding a graph with one crossing*, in Proceedings of the 5th International Workshop on Approximation Algorithms for Combinatorial Optimization (Italy, APPROX 2002), vol. 2462 of Lecture Notes in Computer Science, 2002, pp. 67–80.

[21] E. D. Demaine, M. Hajiaghayi, and D. M. Thilikos, *Exponential speedup of fixed-parameter algorithms for classes of graphs excluding single-crossing graphs as minors*, Algorithmica, 41 (2005), pp. 245–267.

[22] R. Diestel, T. R. Jensen, K. Y. Gorbunov, and C. Thomassen, *Highly connected sets and the excluded grid theorem*, Journal of Combinatorial Theory, Series B, 75 (1999), pp. 61–73.

[23] D. Eppstein, *Diameter and treewidth in minor-closed graph families*, Algorithmica, 27 (2000), pp. 275–291.

[24] U. Feige, M. Hajiaghayi, and J. R. Lee, *Improved approximation algorithms for minimum-weight vertex separators*, in Proceedings of the 37th ACM Symposium on Theory of Computing (STOC 2005), Baltimore, May 2005, pp. 563–572.

[25] F. V. Fomin and D. M. Thilikos, *Dominating sets in planar graphs: Branch-width and exponential speed-up*, in Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2003), 2003, pp. 168–177.

[26] M. Frick and M. Grohe, *Deciding first-order properties of locally tree-decomposable structures*, Journal of the ACM, 48 (2001), pp. 1184–1206.

[27] M. Grohe, *Local tree-width, excluded minors, and approximation algorithms*, Combinatorica, 23 (2003), pp. 613–632.

[28] G. Gutin, T. Kloks, and C. M. Lee, *Kernels in planar digraphs*, in Optimization Online, Mathematical Programming Society, Philadelphia, 2001.

[29] M. Hajiaghayi and N. Nishimura, *Subgraph isomorphism, log-bounded fragmentation and graphs of (locally) bounded treewidth*, in Proc. the 27th International Symposium on Mathematical Foundations of Computer Science (Poland, 2002), 2002, pp. 305–318.

[30] I. Kanj and L. Perković, *Improved parameterized algorithms for planar dominating set*, in Proceedings of the 27th International Symposium on Mathematical Foundations of Computer Science (MFCS 2002), vol. 2420 of Lecture Notes in Computer Science, 2002, pp. 399–410.

[31] 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), 1993, pp. 682–690.

[32] T. Kloks, C. M. Lee, and J. Liu, *New algorithms for k-face cover, k-feedback vertex set, and k-disjoint set on plane and planar graphs*, in Proceedings of the 28th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2002), vol. 2573 of Lecture Notes in Computer Science, 2002, pp. 282–295.

[33] R. J. Lipton and R. E. Tarjan, *Applications of a planar separator theorem*, SIAM Journal on Computing, 9 (1980), pp. 615–627.

[34] B. A. Reed, *Tree width and tangles: a new connectivity measure and some applications*, in Surveys in combinatorics, 1997 (London), vol. 241 of London Math. Soc. Lecture Note Ser., Cambridge Univ. Press, Cambridge, 1997, pp. 87–162.

[35] N. Robertson and P. D. Seymour, *Graph minors. II. Algorithmic aspects of tree-width*, Journal of Algorithms, 7 (1986), pp. 309–322.

[36] N. Robertson and P. D. Seymour, *Graph minors. V. Excluding a planar graph*, Journal of Combinatorial Theory, Series B, 41 (1986), pp. 92–114.

[37] ——, *Graph minors. XVI. Excluding a non-planar graph*, Journal of Combinatorial Theory, Series B, 89 (2003), pp. 43–76.

[38] N. Robertson, P. D. Seymour, and R. Thomas, *Quickly excluding a planar graph*, Journal of Combinatorial Theory, Series B, 62 (1994), pp. 323–348.

[39] P. D. Seymour, *Personal communication*, May 2004.

[40] P. D. Seymour and R. Thomas, *Call routing and the ratcatcher*, Combinatorica, 14 (1994), pp. 217–241.