doi:10.1093/comjnl/bxm033

The Bidimensionality Theory and Its Algorithmic Applications*

Erik D. Demaine[1], MohammadTaghi Hajiaghayi[1]

[1]MIT Computer Science and Artificial Intelligence Laboratory,
32 Vassar Street, Cambridge, MA 02139, USA.
Email: {edemaine,hajiagha}@mit.edu

This paper surveys the theory of bidimensionality. This theory characterizes a broad range of graph problems ("bidimensional") that admit efficient approximate or fixed-parameter solutions in a broad range of graphs. These graph classes include planar graphs, map graphs, bounded-genus graphs, and graphs excluding any fixed minor. In particular, bidimensionality theory builds on the Graph Minor Theory of Robertson and Seymour, by extending the mathematical results and building new algorithmic tools. Here we summarize the known combinatorial and algorithmic results of bidimensionality theory, with the high-level ideas involved in their proof; we describe the previous work on which the theory is based and/or extends; and we mention several remaining open problems.

Received 26 February 2006; revised 10 August 2006

1. INTRODUCTION

The theory of bidimensionality, developed recently in a series of papers [1, 2, 3, 4, 5, 6, 7, 8, 9, 10], provides general techniques for designing efficient fixed-parameter algorithms and approximation algorithms for NP-hard graph problems in broad classes of graphs. This theory applies to graph problems that are *bidimensional* in the sense that (1) the solution value for the $k \times k$ grid graph (and similar graphs) grows with k, typically as $\Omega(k^2)$, and (2) the solution value goes down when contracting edges and optionally when deleting edges in the graph. Examples of such problems include 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, unweighted TSP tour (a walk in the graph visiting all vertices), and chordal completion (fill-in).

Bidimensional problems have many structural properties. For example, any graph in an appropriate minor-closed class has treewidth bounded above in terms of the problem's solution value, typically by the square root of that value. These properties lead to efficient—often subexponential—fixed-parameter algorithms, as well as polynomial-time approximation schemes, for bidimensional problems in many minor-closed graph classes.

*A preliminary version of this paper appeared in *Proceedings of the 12th International Symposium on Graph Drawing*, LNCS 3383, 2004, pages 517–533.

The bidimensionality theory unifies and improves several previous results. The theory is based on algorithmic and combinatorial extensions to parts of the Robertson-Seymour Graph Minor Theory, in particular initiating a parallel theory of graph contractions. The foundation of this work is the topological theory of drawings of graphs on surfaces.

This survey is organized as follows. Section 2 defines the various graph classes of increasing generality for which bidimensionality theory has been developed. Section 3 describes several structural properties of graphs in these classes, in particular from Graph Minor Theory, that form the basis of bidimensionality. Section 4 defines bidimensional problems, and gives some examples. Section 5 describes one of the main structural properties of bidimensionality, namely, that the treewidth is bounded in terms of the parameter value. Sections 6–10 describe several results in bidimensionality theory: separator theorems, bounds on local treewidth, fixed-parameter algorithms, and polynomial-time approximation schemes. Throughout, we give the high-level proof ideas, in particular how the various results fit together. Section 11 discusses the main remaining open problems in the area.

2. GRAPH CLASSES

In this section, we introduce several families of graphs, each playing an important role in both the Graph Minor Theory and the bidimensionality theory. Refer to Figure 1. All of these graph classes are generalizations of planar graphs, which are well-studied in algorithmic graph theory. Unlike planar graphs and map graphs, every other class of graphs we consider can include any particular graph G; of course, this inclusion requires a bound or excluded minor large enough depending on G. This feature distinguishes this line of research from other work considering exclusion of particular fixed minors, e.g., $K_{3,3}$, K_5, or K_6.

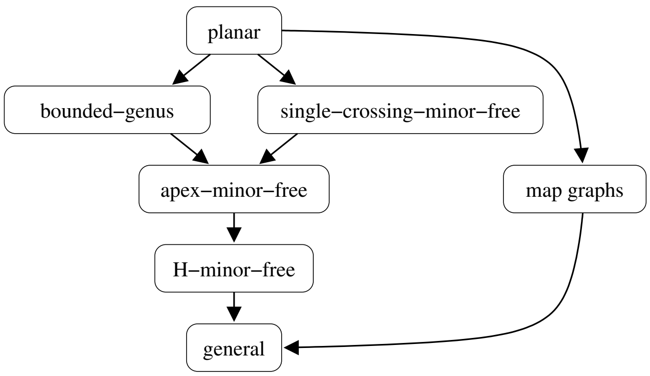

FIGURE 1. Interesting classes of graphs. Arrows point from more specific classes to more inclusive classes.

2.1. Definitions of Graph Classes

The first three classes of graphs relate to embeddings on surfaces. A graph is *planar* if it can be drawn in the plane (or the sphere) without crossings. A graph has *(Euler) genus* at most g if it can be drawn in a surface of Euler characteristic g. For orientable surfaces, the Euler genus is twice the number of handles (i.e., twice the "usual" genus), and for nonorientable surfaces, the Euler genus is the number of crosscaps. A class of graphs has *bounded genus* if every graph in the class has genus at most g for a fixed g.

Given an embedded planar graph and a two-coloring of its faces as either *nations* or *lakes*, the associated *map graph* has a vertex for each nation and an edge between two vertices corresponding to nations (faces) that share a vertex. The dual graph is defined similarly, but with adjacency requiring a shared edge instead of just a shared vertex. Map graphs were introduced by Chen, Grigni, and Papadimitriou [11] as a generalization of planar graphs that can have arbitrarily large cliques. Thorup [12] gave a polynomial-time algorithm for constructing the underlying embedded planar graph and face two-coloring for a given map graph, or determining that the given graph is not a map graph.

We view the class of map graphs as a special case of taking fixed powers of a family of graphs. The *kth power* G^k of a graph G is the graph on the same vertex set $V(G)$ with edges connecting two vertices in G^k precisely if the distance between these vertices in G is at most k. For a bipartite graph G with bipartition $V(G) = U \cup W$, the *half-square* $G^2[U]$ is the graph on one side U of the partition, with two vertices adjacent in $G^2[U]$ precisely if the distance between these vertices in G is 2. A graph is a map graph if and only if it is the half-square of some planar bipartite graph [11]. In fact, this translation between map graphs and half-squares is constructive and takes polynomial time.

The next three classes of graphs relate to excluding minors. Given an edge $e = \{v, w\}$ in a graph G, the *contraction* of e in G is the result of identifying vertices v and w in G and removing all loops and duplicate edges. A graph H obtained by a sequence of such edge contractions starting from G is said to be a *contraction* of G. A graph H is a *minor* of G if H is a subgraph of some contraction of G. A graph class $\mathcal{C}$ is *minor-closed* 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 $H \notin \mathcal{C}$. More generally, we use the term "H-minor-free" to refer to any minor-closed graph class that excludes some fixed graph H.

A *single-crossing graph* is a minor of a graph that can be drawn in the plane with at most one pair of edges crossing. Note that a single-crossing graph may not itself be drawable with at most one crossing pair of edges; see [2]. Such graphs were first defined by Robertson and Seymour [13]. A minor-closed graph class is *single-crossing-minor-free* if it excludes a fixed single-crossing graph.

An *apex graph* is a graph in which the removal of some vertex leaves a planar graph. A graph class is *apex-minor-free* if it excludes some fixed apex graph. Such graph classes were first considered by Eppstein [14, 15], in connection to the notion of bounded local treewidth as described in Section 7.

The next section describes strong structural properties of the last three classes of graphs (minor-excluding classes) in terms of the first two classes of graphs (embeddable on surfaces) and other ingredients.

3. STRUCTURAL PROPERTIES

Graphs from single-crossing-minor-free and H-minor-free graph classes have powerful structural properties from the Graph Minor Theory. First we need to define treewidth, pathwidth, and clique sums.

3.1. Background

The notion of treewidth was introduced by Robertson and Seymour [16]. To define this notion, first we consider a representation of a graph as a tree, called a tree decomposition. 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

(i) $\bigcup_{i \in I} \chi_i = V$;

(ii) 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
(iii) 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 *width* of the tree decomposition is the maximum size of a bag in χ minus 1. The *treewidth* of a graph G, denoted $\mathbf{tw}(G)$, 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, denoted $\mathbf{pw}(G)$, is the minimum width over all possible path decompositions of G.

The notion of clique sums goes back to characterizations of $K_{3,3}$-minor-free and K_5-minor-free graphs by Wagner [17] and serves as an important tool in the Graph Minor Theory. 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 $G_i[W_i]$ 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.

3.2. Structure of Single-Crossing-Minor-Free Graphs

The structure of single-crossing-minor-free graphs can be described as follows:

Theorem 3.1. [13] *For any fixed single-crossing graph H, every H-minor-free graph can be obtained by a sequence of k-sums, $0 \leq k \leq 3$, of planar graphs and graphs of bounded treewidth, where the bound on treewidth depends on H.*

This theorem generalizes characterizations of $K_{3,3}$-minor-free and K_5-minor-free graphs [17]. A graph is $K_{3,3}$-minor-free if and only if it can be obtained by k-sums, $0 \leq k \leq 2$, of planar graphs and K_5. A graph is K_5-minor-free if and only if it can be obtained by k-sums, $0 \leq k \leq 3$, of planar graphs and of V_8 (the length-8 cycle C_8 together with four edges joining diametrically opposite vertices).

This structural property of single-crossing-minor-free graphs has since been strengthened to ensure that the summands are minors of the original graph and to provide algorithms for finding the decomposition:

Theorem 3.2. [2] *For any fixed single-crossing graph H, there is an $O(n^4)$-time algorithm to compute, given an H-minor-free graph G, a decomposition of G as a sequence of k-sums, $0 \leq k \leq 3$, of planar graphs and graphs of bounded treewidth (where the bound on treewidth depends on H), each of which is a minor of G.*

3.3. Structure of *H*-Minor-Free Graphs

The structure of H-minor-free graphs is described by a deep theorem of Robertson and Seymour [18]. This theorem is the heart of the Graph Minors series of papers, in particular forming the basis of the proof of Wagner's Conjecture in [19]. Intuitively, the 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 that 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, a graph G is *h-almost embeddable* in S if there exists a vertex set X of size at most h (the *apices*) such that $G - X$ can be written as $G_0 \cup G_1 \cup \cdots \cup G_h$, where

(i) G_0 has an embedding in S;
(ii) the graphs G_i, called *vortices*, are pairwise disjoint;
(iii) 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
(iv) the graph G_i has a path decomposition $(\mathcal{B}_u)_{u \in U_i}$ 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 *apex-free* if the set X of apices is empty.

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

Theorem 3.3. [18] *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.

As stated in [20], the proof of this theorem in [18] in fact establishes the following stronger result. An explicit proof is also given in [21]. The additional structure provided by this theorem is useful for some mathematical and algorithmic results.

Theorem 3.4. *The clique-sum decomposition of Theorem 3.3, written as $G_1 \oplus G_2 \oplus \cdots \oplus G_k$, has the additional property that the join set of each clique-sum between $G_1 \oplus G_2 \oplus \cdots \oplus G_{i-1}$ and G_i is a subset of the apices in G_i. Furthermore, the join set of each clique-sum involving piece G_j contains at most three vertices from the bounded-genus part of G_j.*

Recently, a polynomial-time algorithm has been developed that computes the decomposition guaranteed by Theorem 3.4, for fixed H [21]. The same paper proves another algorithmic result for structural decomposition of H-minor-free graphs (an existential version of which is proved in [20]):

Theorem 3.5. [21, Theorem 3.1] *For a fixed graph H, there is a constant c_H such that, for any integer $k \geq 1$ and for every H-minor-free graph G, the vertices of G (or the edges of G) can be partitioned into $k+1$ sets such that any k of the sets induce a graph of treewidth at most $c_H k$. Furthermore, such a partition can be found in polynomial time.*

Proof idea. We use Theorem 3.4 to decompose the graph G. In each piece, we partition the bounded-genus part according to layer numbers in a breadth-first search, treated modulo $k+1$. (This idea of layerwise partition goes back to Baker's approach for PTASs in planar graphs [22].) Vortices are easy to handle because they have bounded pathwidth. Apices are easy to handle because they are bounded in number—indeed, the apices could be in all sets in the partition of a piece, and still any k sets would induce a graph of bounded treewidth. This latter fact is crucial for joining pieces across clique sums, because each join set consists entirely of apices in the child, so the child can inherit the partition of the join vertices from the parent. More care is required for edge partitions. $\square$

This theorem has several applications to approximation algorithms, as detailed in [21]. For example, we obtain a simple 2-approximation to vertex coloring of H-minor-free graphs as follows. Decompose the vertices of the H-minor-free graph into $k = 2$ bounded-treewidth graphs; then optimally color each piece via dynamic programming; and then combine the colorings by using a disjoint color set for each piece. The resulting coloring uses at most twice the optimal number of colors, because the optimal coloring of each piece is a lower bound for the original graph. A similar approach leads to constant-factor approximations for many minimization and maximization problems, in some cases even leading to polynomial-time approximation schemes (PTASs). In addition, the paper establishes a constant-factor approximation to treewidth, and approximation to half-integral multicommodity flow, in H-minor-free graphs.

3.4. Structure of Apex-Minor-Free Graphs

Apex-minor-free graph classes are an important subfamily of H-minor-free graph classes. The general structural theorem for H-minor-free graphs applies in this context as well. However, reductions developed in [6] suggest that the decomposition can be restricted to a particular form in the apex-minor-free case:

Conjecture 1. [6] For every graph H, there is 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 and whose apices are connected via edges only to vertices within vortices.

3.5. Grid Minors

The $r \times r$ grid is the canonical planar graph of treewidth $\Theta(r)$. In particular, an important result of Robertson, Seymour, and Thomas [23] 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). Recently, this result has been generalized to any H-minor-free graph class:

Theorem 3.6. [9] *For any fixed graph H, every H-minor-free graph of treewidth w has an $\Omega(w) \times \Omega(w)$ grid as a minor.*

Proof idea. We follow the decomposition provided by Theorem 3.4. Because the treewidth of a clique sum is at most the maximum treewidth of the pieces, some piece has treewidth as large as the original graph, and we can focus on that piece. We solve the bounded-genus case using results from [5]. Apices and vortices are easy to handle. Particular care is required to make the grid minor we find in the piece into a minor of the original graph, because some edges may be removed during clique sums with other pieces. $\square$

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 also best possible up to constant factors. Section 11 discusses the remaining issue of bounding the constant factor and its dependence on H.

A similar but weaker bound plays an important role in the Graph Minor Theory [24]: for any fixed 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 reproved by Robertson, Seymour, and Thomas [23], Reed [25], and Diestel, Jensen, Gorbunov, and Thomassen [26]. Among these proofs, the best known bound on w in terms of r is that every H-minor-free graph of treewidth larger than $20^{5|V(H)|^3 r}$ has an $r \times r$ grid as a minor [23]. Theorem 3.6 therefore offers an exponential (and best possible) improvement over previous results.

Theorem 3.6 cannot be generalized to arbitrary graphs: Robertson, Seymour, and Thomas [23] 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 [23]. The best possible bound is believed to be closer to $\Theta(r^2 \lg r)$ than $2^{\Theta(r^5)}$, perhaps even equal to $\Theta(r^2 \lg r)$ [23].

4. BIDIMENSIONAL PARAMETERS/PROBLEMS

Bidimensionality has been introduced and developed in a series of papers [1, 2, 3, 4, 5, 6, 7, 8, 9, 10]. Although implicitly hinted at in [1, 2, 3, 4], the first use of the term "bidimensional" was in [5].

First we define "parameters" as an alternative view on optimization problems. A *parameter* P is a function mapping graphs to nonnegative integers. The *decision problem associated with* P asks, for a given graph G and nonnegative integer k, whether $P(G) \leq k$. Many optimization problems can be phrased as such a decision problem about a graph parameter P.

Now we can define bidimensionality. A parameter is $g(r)$-*bidimensional* (or just *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 ($g(r)$-*minor-bidimensional*) or contractions ($g(r)$-*contraction-bidimensional*). The exact definition of "grid-like graph" depends on the class of graphs allowed and whether we are considering minor- or contraction-bidimensionality. For minor-bidimensionality and for any H-minor-free graph class, the notion of a "grid-like graph" is defined to be the $r \times r$ *grid*, i.e., the planar graph with r^2 vertices arranged on a square grid and with edges connecting horizontally and vertically adjacent vertices. For contraction-bidimensionality, the notion of a "grid-like graph" is as follows:

(i) For planar graphs and single-crossing-minor-free graphs, a "grid-like graph" is an $r \times r$ grid partially triangulated by additional edges that preserve planarity.

(ii) For bounded-genus graphs, a "grid-like graph" is such a partially triangulated $r \times r$ grid with up to genus(G) additional edges ("handles").

(iii) For apex-minor-free graphs, a "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 the excluded apex graph.)

Contraction-bidimensionality is so far undefined for H-minor-free graphs (or general graphs).[1]

Examples of bidimensional parameters include the number of vertices, the diameter, and the size of various structures such as 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, unweighted TSP tour (a walk in the graph visiting all vertices), and chordal completion (fill-in). For example, feedback vertex set is $\Omega(r^2)$-minor-bidimensional (and thus also contraction-bidimensional) because (1) deleting or contracting an edge preserves existing feedback vertex sets, and (2) any vertex in the feedback vertex set destroys at most four squares in the $r \times r$ grid, and there are $(r-1)^2$ squares, so any feedback vertex set must have $\Omega(r^2)$ vertices. See [5, 7] for arguments of either contraction- or minor-bidimensionality for the other parameters.

We also say that the corresponding optimization problems based on these parameters, e.g., finding the minimum-size feedback vertex set, are *bidimensional*. With the exception of diameter, all of these bidimensional problems are $\Theta(r^2)$-bidimensional, which is the most common case (and in some papers used as the definition of bidimensionality). Diameter is the main exception, being only $\Theta(r)$-contraction-bidimensional for planar graphs, single-crossing-minor-free graphs, and bounded-genus graphs, and only $\Theta(\lg r)$-contraction-bidimensional for apex-minor-free graphs. See [7] for details.

5. PARAMETER-TREEWIDTH BOUNDS

The genesis of bidimensionality was in fact the notion of a parameter-treewidth bound. 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 and graph classes; see, e.g., [27, 28, 29, 30, 31, 32, 33, 3, 2, 1, 7, 6, 5]. Essentially all of these bounds can be obtained from the theory of bidimensional parameters. Thus bidimensionality is the most powerful method so far for establishing parameter-treewidth bounds, encompassing all such previous results for H-minor-free graphs.

The central result in bidimensionality that generalizes these bounds is that every bidimensional parameter has a parameter-treewidth bound, in its corresponding family of graphs as defined in Section 4 (namely, H-minor-free graphs for minor-bidimensional parameters; and

[1] For the parameters to which we have applied bidimensionality, contraction-bidimensionality does not seem to extend beyond apex-minor-free graphs, but perhaps a suitably extended definition could be found in the context of different applications or a "theory of graph contractions".

planar graphs, bounded-genus graphs, or apex-minor-free graphs for contraction-bidimensional parameters, depending on the grid-like graphs handled). More precisely, we have the following result:

Theorem 5.1. [9, 7] *If the parameter P is $g(r)$-bidimensional, then for every graph G in the family associated with the parameter P, $\mathbf{tw}(G) = O(g^{-1}(P(G)))$. In particular, if $g(r) = \Theta(r^2)$, then the bound becomes $\mathbf{tw}(G) = O(\sqrt{P(G)})$.*

Proof idea. We use Theorem 3.6 to construct an $r \times r$ grid minor R of G where $r = \Omega(\mathbf{tw}(G))$. By bidimensionality, $P(R) \geq g(r)$. If P is minor-bidimensional, then $P(G) \geq P(R) \geq g(r) \geq g(\Omega(\mathbf{tw}(G))$, which is the desired result. If P is contraction-bidimensional, we cannot delete edges to form the grid R exactly, but using results of [8, 7] we show that just performing the contractions forms the appropriate kind of "grid-like graph", serving the same role as R. □

Such a parameter-treewidth bound was originally established in the weaker form $\mathbf{tw}(G) = (g^{-1}(P(G)))^{O(g^{-1}(P(G)))}$ [7]. The stronger bound of $\mathbf{tw}(G) = O(g^{-1}(P(G)))$ was obtained first for planar graphs [3], then single-crossing-minor-free graphs [1, 2], then bounded-genus graphs [5, 8], and finally apex-minor-free graphs for contraction-bidimensional parameters and H-minor-free graphs for minor-bidimensional parameters [9] (Theorem 5.1 above) using the grid-minor bound of Theorem 3.6.

We can extend the definition of $g(r)$-minor-bidimensionality to general graphs by again defining a "grid-like graph" to be the $r \times r$ grid. Still we can obtain a parameter-treewidth bound [23, 34], but the bound is weaker: $\mathbf{tw}(G) = 2^{O(g^{-1}(k))^5}$.

6. SEPARATOR THEOREMS

If we apply the parameter-treewidth bound of Theorem 5.1 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:

Theorem 6.1. [35, Proposition 4.5], [36, Corollary 24], [9] *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 [35, Theorem 1.2]. This generalization of the classic planar separator theorem has many algorithmic applications; see e.g. [35, 37]. Also, this result shows that the structural properties of H-minor-free graphs given by Theorems 3.3 and 3.4 are powerful enough to conclude that these graphs have small separators, which we expect from such a strong theorem.

Section 11 discusses the issue of how tight a lead constant can be obtained in such a result.

7. LOCAL TREEWIDTH

Eppstein [15] 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 is obtained in [4].)

If we apply the parameter-treewidth bound of Theorem 5.1 to the diameter parameter, which is contraction-bidimensional with $g(r) = \Theta(\lg r)$ [4], then we immediately obtain the following stronger diameter-treewidth bound for apex-minor-free graphs:

Theorem 7.1. [9] *For any fixed apex graph H, every H-minor-free graph of diameter D has treewidth $2^{O(D)}$.*

This theorem is not the best possible. In some sense, it is necessarily limited, because it still does not exploit the full structure of H-minor-free graphs from Theorem 3.3. The difficulty is that, in a grid-like graph, the $O(1)$ edges from a vertex to nonboundary vertices can accumulate to make the diameter small. However, it is possible to show that, effectively, not too many vertices can have such edges. This fact comes from the property that there are a bounded number of apices in the clique-sum decomposition of Theorem 3.3, and in an apex-minor-free graph, each apex cannot have more than a bounded number of edges to "distant" vertices. Based on this fact, a complicated proof establishes the following even stronger diameter-treewidth bound in apex-minor-free graphs:

Theorem 7.2. [6] *For any fixed apex graph H, every H-minor-free graph of diameter D has treewidth $O(D)$.*

This diameter-treewidth bound is the best possible up to constant factors. Thus this theorem establishes that, in minor-closed graph families, having any diameter-treewidth bound is equivalent to having a linear diameter-treewidth bound. As mentioned before, no minor-closed graph families beyond apex-minor-free graphs can have any diameter-treewidth bound. Theorem 7.2 is therefore the ultimate characterization of diameter-treewidth bounds in minor-closed graph families (up to constant factors).

The proof of Theorem 7.2 is the basis for Conjecture 1. In fact, Theorem 7.2 would not be hard to prove assuming Conjecture 1.

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 PTASs on planar graphs [22] 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 [2, 15, 38]. Graphs of bounded local treewidth also admit several efficient fixed-parameter algorithms. In particular, Frick and Grohe [39] give a general framework for deciding any property expressible in first-order logic in graphs of bounded local treewidth. Theorem 7.2 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 $2^{O(1/\varepsilon)} n^{O(1)}$, where n is the number of vertices in the graph.

8. 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. The exponent $O(1)$ must be independent of G; thus the exponentiality of the algorithm is bounded by the parameter $P(G)$, and the dependence on n is only polynomial. A typical function h for many fixed-parameter algorithms is $h(k) = 2^{O(k)}$. In the last five 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 [40] has running time $O(8^k n)$; subsequently, a sequence of subexponential algorithms and improvements have been obtained, starting with running time $O(4^{6\sqrt{34k}} n)$ [27], then $O(2^{27\sqrt{k}} n)$ [28], and finally $O(2^{15.13\sqrt{k}} k + n^3 + k^4)$ [29]. Other subexponential algorithms for other domination and covering problems on planar graphs have also been obtained [27, 30, 31, 32, 33].

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 in Section 5, essentially all sublinear treewidth-parameter bounds proved so far can be obtained through bidimensionality. Theorem 5.1, combined with algorithms for computing tree decompositions of approximately optimal width [41], yield the following general result for designing subexponential fixed-parameter algorithms:

THEOREM 8.1. [9, 7] *Consider a $g(r)$-bidimensional 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. Then there is an algorithm computing P on any graph G in P's corresponding graph class, 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 this running time is subexponential in k.*

In particular, this result 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, unweighted TSP tour, and chordal completion [9, 7].

For minor-bidimensional parameters, these algorithms apply to all H-minor-free graphs. The next section describes to what extent these algorithms can be extended to general graphs.

For contraction-bidimensional parameters, these algorithms apply to apex-minor-free graphs. On the other hand, subexponential fixed-parameter algorithms can be obtained for dominating set, which is contraction-bidimensional, on H-minor-free graphs [5], map graphs [3], and fixed powers of planar graphs (or even fixed powers of H-minor-free graphs) [3, 5]. These algorithms are necessarily more complicated than those produced from Theorem 8.1, because apex-minor-free graphs are precisely the minor-closed graph classes for which dominating set has a parameter-treewidth bound [7]. An intriguing open question is whether these techniques can be extended to other contraction-bidimensional problems than dominating set, for fixed powers of H-minor-free graphs and/or other classes of graphs.

9. FIXED-PARAMETER ALGORITHMS FOR GENERAL GRAPHS

As mentioned in Section 5, minor-bidimensionality can be defined for general graphs as well. In this section we show how the bidimensionality theory in this case leads to a general class of fixed-parameter algorithms.

A major result from the Graph Minor Theory (in particular [42, 19]) is that every minor-closed graph property is characterized by a finite set of forbidden minors. More precisely, for any property P on graphs such that a graph having property P implies that all its minors have property P, there is a finite set $\{H_1, H_2, \ldots, H_h\}$ of graphs such that a graph G has property P if and only if G does not have H_i as a minor for all $i = 1, 2, \ldots, h$. The algorithmic consequence of this result is that there exists an $O(n^3)$-time algorithm to decide any fixed minor-closed graph property, by finitely many calls to an $O(n^3)$-time minor test [42]. This consequence has been used to show the existence of polynomial-time algorithms for several graph problems, some of which were not previously known to be decidable [43].

However, all of these algorithmic results (except the minor test) are nonconstructive: we are guaranteed that efficient algorithms exist, but are not told what they are. The difficulty is that we know that a finite set of forbidden minors exists, but lack "a means of identifying the elements of the set, the cardinality of the set, or even the order of the largest graph in the set" [43]. Indeed, there is a mathematical sense in which any proof of the finite-forbidden-minors theorem must be nonconstructive [44].

We can apply these graph-minor results to prove the existence of algorithms to compute parameters, provided the parameters never increase when taking a minor. For any fixed parameter and any fixed $k \geq 0$, there is an $O(n^3)$-time algorithm that decides whether a graph has parameter value $\leq k$. Unfortunately, the existence of these algorithms does not necessarily imply the existence of a single fixed-parameter algorithm that works for all $k \geq 0$, because the algorithms for individual k (in particular the set of forbidden minors) might be uncomputable. We do not even know an upper bound on the running time of these algorithms as a function of n and k, because we do not know the dependence of the size of the forbidden minors on k.

In [34], fixed-parameter algorithms are constructed for nearly all parameters that never increase when taking a minor, with explicit time bounds in terms of n and k. Essentially, by assuming a few very common properties of the parameter, we obtain a generalized form of minor-bidimensionality. The proof is then similar to that of Theorem 8.1.

Theorem 9.1. [34] *Consider a parameter P that is positive on some $g \times g$ grid, never increases when taking minors, is at least the sum over the connected components of a disconnected graph, and can be computed in $h(w)\, n^{O(1)}$ time given a width-w tree decomposition of the graph. Then there is an algorithm that decides whether P is at most k on a graph with n vertices in $\left[2^{2^{O(g\sqrt{k})^5}} + h(2^{O(g\sqrt{k})^5})\right] n^{O(1)}$ time.*

As mentioned in [34], a conjecture of Robertson, Seymour, and Thomas [23] would improve the running time to $h(O(k \lg k))\, n^{O(1)}$, which is $2^{O(k \lg k)} n^{O(1)}$ for the typical case of $h(w) = 2^{O(w)}$. This conjectured time bound almost matches the fastest known fixed-parameter algorithms for several parameters, e.g., feedback vertex set, vertex cover, and a general family of vertex-removal problems [43].

10. POLYNOMIAL-TIME APPROXIMATION SCHEMES

Recently, the bidimensionality theory has been extended to obtain polynomial-time approximation schemes (PTASs) for essentially all bidimensional parameters, including those mentioned above [10]. 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 [45] and the Baker layerwise decomposition approach [22]. The PTASs apply to H-minor-free graphs for minor-bidimensional parameters and to apex-minor-free graphs for contraction-bidimensional parameters. To achieve this level of generality, [10] uses the sublinear parameter-treewidth bound of Theorem 5.1 as well as a recent $O(1)$-approximation algorithm for treewidth in H-minor-free graphs [46, 21].

Before we can state the general theorem for constructing PTASs, we need to define a few straightforward required conditions, which are commonly satisfied by most bidimensional problems. The theorem considers families of problems in which we are given a graph and our goal is to find a minimum-size set of vertices and/or edges satisfying a certain property. Such a problem naturally defines a parameter and therefore the notion of bidimensionality. A minor-bidimensional problem has the *separation property* if it satisfies the following three conditions:

(i) If a graph G has k connected components $G_1, G_2, \ldots, G_k$, then an optimal solution for G is the union of optimal solutions for each connected component G_i.

(ii) There is a polynomial-time algorithm that, given any graph G, given any vertex cut C whose removal disconnects G into connected components $G_1, G_2, \ldots, G_k$, and given an optimal solution S_i to each connected component G_i of $G - C$, computes a solution S for G such that the number of vertices and/or edges in S within the induced subgraph $G[C \cup \cup_{i \in I} V(G_i)]$ consisting of C and some connected components of $G - C$ is $\sum_{i \in I} |S_i| \pm O(|C|)$ for any $I \subseteq \{1, 2, \ldots, k\}$. In particular, the total cost of S is at most $\mathrm{OPT}(G - C) + O(|C|)$.

(iii) Given any graph G, given any vertex cut C, and given an optimal solution OPT to G, for any union G' of some subset of connected components of $G - C$, $|\,\mathrm{OPT} \cap G'| = |\,\mathrm{OPT}(G')| \pm O(|C|)$.

For contraction-bidimensional problems, the exact requirements on the problem are slightly different but similarly straightforward. The main distinction is that the connected components are always considered together with the cut C. As a result, the merging algorithm in Condition 2 must take as input a solution to a generalized form of the problem that does not count the cost of including all vertices and edges from the cut C. We omit the exact definition of the separation property in this case in the interest of space.

Theorem 10.1. [10] *Consider a bidimensional problem satisfying the separation property. Suppose that the problem can be solved on a graph G with n vertices in $f(n, \mathbf{tw}(G))$ time. Suppose also that the problem can be approximated within a factor of α in $g(n)$ time. For contraction-bidimensional problems, suppose further that both of these algorithms also apply to the generalized form of the problem. Then there is a $(1+\varepsilon)$-approximation algorithm whose running time is $O(nf(n, O(\alpha^2/\varepsilon)) + n^3 g(n))$ for the corresponding graph class of the bidimensional problem.*

Proof idea. The central idea is to use a separator theorem like Theorem 6.1, but to bound the separator size by $O(\sqrt{\mathrm{OPT}})$ instead of $O(\sqrt{|V(G)|})$. Without such a bound, the overhead of a single separator might incur a cost far greater than OPT, and this has been the limiting factor of the separator approach in the past. From Theorem 5.1 and bidimensionality, we obtain that the treewidth of G is $O(\sqrt{\mathrm{OPT}})$, and this leads to a separator that bisects the number of vertices. However, for an effective recursion, we need to find a separator that bisects OPT. This challenge is particularly difficult because we do not know OPT and its distribution in the tree decomposition. □

This result shows a strong connection between subexponential fixed-parameter tractability and approximation algorithms for combinatorial optimization problems on H-minor-free graphs. In particular, this result yields a PTAS for the following minor-bidimensional problems in H-minor-free graphs: feedback vertex set, face cover (defined just for planar graphs), vertex cover, minimum maximal matching, and a series of vertex-removal problems. Furthermore, the result yields a PTAS for the following contraction-bidimensional problems in apex-minor-free graphs: dominating set, edge dominating set, R-dominating set, connected dominating set, connected edge dominating set, connected R-dominating set, and clique-transversal set.

11. OPEN PROBLEMS

Several combinatorial and algorithmic open problems remain in the theory of bidimensionality and related concepts.

One interesting direction is to generalize bidimensionality to handle general graphs, not just H-minor-free graph classes. As mentioned in Section 5, the natural generalization of minor-bidimensionality still yields a parameter-treewidth bound, but it is very large. This direction essentially asks for the size of the largest grid minor guaranteed to exist in any graph of treewidth w. Robertson, Seymour, and Thomas [23] proved that every graph of treewidth larger than 20^{2r^5} has an $r \times r$ grid as a minor, but that some graphs of treewidth $\Omega(r^2 \lg r)$ have no grid larger than $O(r) \times O(r)$, conjecturing that the right requirement on treewidth for an $r \times r$ grid is closer to the $\Theta(r^2 \lg r)$ lower bound. Other work leads to a conjecture of $\Theta(r^3)$ being the right bound [47]. If either of these conjectures are correct, we would obtain nearly as good parameter-treewidth bounds for minor-bidimensional parameters as in the H-minor-free case. Polynomial upper bounds have recently been established for classes of graphs that do not exclude a fixed minor, namely, map graphs and power graphs [47].

A similar generalization of parameter-treewidth bounds beyond apex-minor-free graphs is not possible for all contraction-bidimensional parameters, e.g., dominating set [7], but it would still be quite interesting to explore an analogous "theory of graph contractions" paralleling the Graph Minor Theory. Such a theory would be an interesting and powerful tool for handling problems that are closed under contractions but not minors, and therefore deserves more focus. Unfortunately, the contraction analog of Wagner's Conjecture has recently been disproved for series-parallel graphs [47] (although it holds for trees, triangulated planar graphs, and 2-connected outerplanar graphs). As a consequence, general techniques for deciding contraction-closed properties in general graphs would need to use a different approach than that of Graph Minor Theory. Nonetheless, it would be interesting to generalize the subexponential fixed-parameter algorithms and polynomial-time approximation schemes for contraction-bidimensional problems from apex-minor-free graphs to H-minor-free graphs. So far such results are known only for dominating set [5, 36], but there are many other natural contraction-bidimensional problems to tackle, such as unweighted TSP. Another challenging problem for approximation is minimum chordal completion (or minimum fill-in), which is contraction-bidimensional yet no constant-factor approximation algorithm is known even for planar graphs.

Another interesting direction is to obtain the best constant factors in terms of the fixed excluded minor H. These constants are particularly important in the context of the exponent in the running time of a fixed-parameter algorithm. At the heart of all such constant factors is the lead constant in Theorem 3.6. This 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 [35] has a lead factor of $|V(H)|^{3/2}$. In fact, Alon, Seymour, and Thomas [35] 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 3.6, which would imply the corresponding bound for separators. Explicit bounds have recently been constructed for the case of $H = K_{3,k}$ using different techniques [47].

A third interesting direction is to generalize the polynomial-time approximation schemes that come out of bidimensionality to more general algorithmic problems that do not correspond directly to bidimensional parameters. One general family of such problems arises when adding weights to vertices and/or edges, and the goal is e.g. to find the minimum-weight dominating set. It is difficult to define bidimensionality of the corresponding weighted parameter because its value is no longer well-defined on an $r \times r$ grid: the parameter value now depends on the weights of vertices in such a grid. Another family of such problems arises when placing constraints (e.g., on coverage or domination) only on subsets of vertices and/or edges. Examples of such problems include Steiner tree [48] and subset feedback vertex set [49]. Again it is difficult to define bidimensionality in such cases because the value of the parameter on a grid depends on which vertices and/or edges of the grid are in the subset.

Finally, we mention some broader directions in which bidimensionality theory may be generalized. Instead of basing bidimensionality on treewidth, it might be possible to use cliquewidth, a parameter that also leads to efficient algorithms when it is small, and which is bounded whenever treewidth is bounded; see, e.g., [50]. We already obtain a parameter-cliquewidth bound from the existing theory, but an alternative theory may allow us to bound the cliquewidth in more general graphs than H-minor-free graphs. Another direction might be to generalize to more general families of matroids that have excluded-minor characterizations; see, e.g., [51]. It would also be interesting to capture problems such as graph coloring in H-minor-free graphs, for which the best possible approximation factor is between 4/3 and 2 [21], while it still may be possible to obtain an additive $O(1)$ approximation.

ACKNOWLEDGMENTS

We thank Fedor Fomin, Naomi Nishimura, Prabhakar Ragde, Paul Seymour, and Dimitrios Thilikos for their fruitful collaboration and discussions in this area.

REFERENCES

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

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

[3] Demaine, E. D., Fomin, F. V., Hajiaghayi, M., and Thilikos, D. M. (2005) Fixed-parameter algorithms for (k, r)-center in planar graphs and map graphs. *ACM Transactions on Algorithms*, **1**, 33–47.

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

[5] Demaine, E. D., Fomin, F. V., Hajiaghayi, M., and Thilikos, D. M. (2005) Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs. *Journal of the ACM*, **52**, 866–893.

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

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

[8] Demaine, E. D., Hajiaghayi, M., and Thilikos, D. M. (2006) The bidimensional theory of bounded-genus graphs. *SIAM Journal on Discrete Mathematics*, **20**, 357–371.

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

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

[11] Chen, Z.-Z., Grigni, M., and Papadimitriou, C. H. (2002) Map graphs. *Journal of the ACM*, **49**, 127–138.

[12] Thorup, M. (1998) Map graphs in polynomial time. *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, pp. 396–407.

[13] Robertson, N. and Seymour, P. (1993) Excluding a graph with one crossing. *Graph structure theory (Seattle, 1991)*, pp. 669–675. Amer. Math. Soc.

[14] Eppstein, D. (1995) Subgraph isomorphism in planar graphs and related problems. *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco, CA, 1995)*, pp. 632–640.

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

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

[17] Wagner, K. (1937) Über eine Eigenschaft der eben Komplexe. *Mathematische Annalen*, **114**, 570–590.

[18] Robertson, N. and Seymour, P. D. (2003) Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, **89**, 43–76.

[19] Robertson, N. and Seymour, P. D. (2004) Graph minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B*, **92**, 325–357.

[20] DeVos, M., Ding, G., Oporowski, B., Sanders, D. P., Reed, B., Seymour, P., and Vertigan, D. (2004) Excluding any graph as a minor allows a low tree-width 2-coloring. *Journal of Combinatorial Theory, Series B*, **91**, 25–41.

[21] Demaine, E. D., Hajiaghayi, M., and Kawarabayashi, K. (2005) Algorithmic graph minor theory: Decomposition, approximation, and coloring. *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, Pittsburgh, PA, October, pp. 637–646.

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

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

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

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

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

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

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

[29] Fomin, F. V. and Thilikos, D. M. (2006) Dominating sets in planar graphs: branch-width and exponential speed-up. *SIAM Journal on Computing*, **36**, 281–309.

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

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

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

[33] Gutin, G., Kloks, T., Lee, C. M., and Yao, A. (2005) Kernels in planar digraphs. *Journal of Computer and System Sciences*, **71**, 174–184.

[34] Demaine, E. D. and Hajiaghayi, M. (2007) Quickly deciding minor-closed parameters in general graphs. *European Journal of Combinatorics*, **28**, 311–314.

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

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

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

[38] Hajiaghayi, M. and Nishimura, N. (2007) Subgraph isomorphism, log-bounded fragmentation and graphs of (locally) bounded treewidth. *Journal of Computer and System Sciences*, **73**, 755–768.

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

[40] Alber, J., Fan, H., Fellows, M. R., Fernau, H., Niedermeier, R., Rosamond, F. A., and Stege, U. (2001) Refined search tree technique for DOMINATING SET on planar graphs. *Proceedings of the 26th International Symposium on Mathematical Foundations of Computer Science*, Lecture Notes in Computer Science, **2136**, pp. 111–122.

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

[42] Robertson, N. and Seymour, P. D. (1995) Graph minors. XII. Distance on a surface. *Journal of Combinatorial Theory, Series B*, **64**, 240–272.

[43] Fellows, M. R. and Langston, M. A. (1988) Nonconstructive tools for proving polynomial-time decidability. *Journal of the ACM*, **35**, 727–739.

[44] Friedman, H., Robertson, N., and Seymour, P. (1987) The metamathematics of the graph minor theorem. *Logic and combinatorics (Arcata, Calif., 1985)*, Contemp. Math., **65**, pp. 229–261. Amer. Math. Soc.

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

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

[47] Demaine, E. D., Hajiaghayi, M., and Kawarabayashi, K. (2006) Algorithmic graph minor theory: Improved grid minor bounds and Wagner's contraction. *Proceedings of the 17th Annual International Symposium on Algorithms and Computation*, Calcutta, India, December, Lecture Notes in Computer Science, **4288**, pp. 3–15.

[48] Arora, S., Grigni, M., Karger, D., Klein, P., and Woloszyn, A. (1998) A polynomial-time approximation scheme for weighted planar graph TSP. *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 33–41.

[49] Even, G., Naor, J., and Zosin, L. (2000) An 8-approximation algorithm for the subset feedback vertex set problem. *SIAM Journal on Computing*, **30**, 1231–1252.

[50] Corneil, D. G. and Rotics, U. (2005) On the relationship between clique-width and treewidth. *SIAM Journal on Computing*, **34**, 825–847.

[51] Geelen, J. F., Gerards, A. M. H., Robertson, N., and Whittle, G. P. (2003) On the excluded minors for the

matroids of branch-width k. *J. Combin. Theory Ser. B*, **88**, 261–265.