Fast Algorithms for Hard Graph Problems: Bidimensionality, Minors, and Local Treewidth

Erik D. Demaine and MohammadTaghi Hajiaghayi

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

Abstract. This paper surveys the theory of bidimensional graph problems. We summarize the known combinatorial and algorithmic results of this theory, the foundational Graph Minor results on which this theory is based, and the remaining open problems.

1 Introduction

The newly developing theory of bidimensional graph problems, developed in a series of papers [DHT, DHN$^+$04, DFHT, DH04a, DFHT04b, DH04b, DFHT04a, DHT04, DH05b, DH05a], 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. 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, and unweighted TSP tour (a walk in the graph visiting all vertices).

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 many minor-closed graph classes. One type of minor-closed graph class of particular relevance has *bounded local treewidth*, in the sense that the treewidth of a graph is bounded above in terms of the diameter; indeed, such a bound is always at most linear.

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 to which bidimensionality theory applies. 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 parameters and 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 consequences of bidimensionality theory: separator theorems, bounds on local treewidth, fixed-parameter algorithms, and polynomial-time approximation schemes. Section 11 discusses the main remaining open problems in this 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 property distinguishes this line of research from other work considering exclusion of particular minors, e.g., $K_{3,3}$, K_5, or K_6.

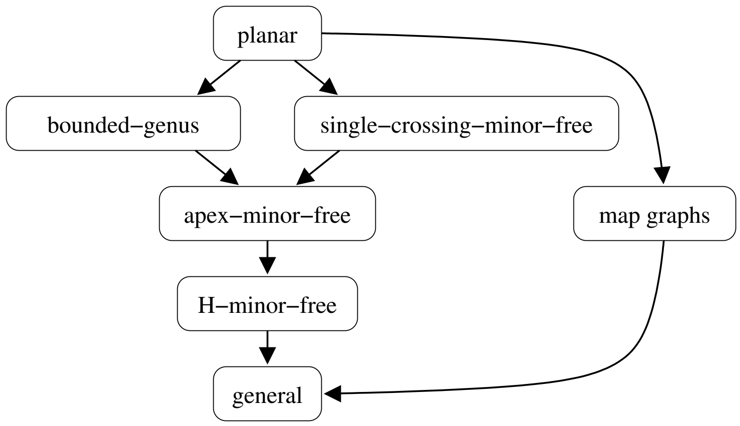

Fig. 1. Interesting classes of graphs. Arrows point from more specific classes to more inclusive classes.

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 *genus* at most g if it can be drawn in an orientable surface of genus g without crossings.[1] 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 [CGP02] as a generalization of planar graphs that can have arbitrarily large cliques. Thorup [Tho98] 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 [CGP02]. 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 [DHN+04]. Such graphs were first defined by Robertson and Seymour [RS93]. 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 [Epp95, Epp00], 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.

[1] This definition also includes graphs that can be drawn in non-orientable surfaces of low genus, because if a graph has non-orientable genus g, then it has orientable genus at most $2g$.

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 [RS86a]. 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

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 *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 [Wag37] 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 1 ([RS93]). *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 [Wag37]. A graph is $K_{3,3}$-minor-free if and only if it can be obtained by k-sums, $0 \le k \le 2$, of planar graphs and K_5. A graph is K_5-minor-free if and only if it can be obtained by k-sums, $0 \le k \le 3$, of planar graphs and V_8 (the length-8 cycle C_8 together with eight 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 2 ([DHN$^+$04]). *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, $0 \le k \le 3$, 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 [RS03]. 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 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

1. G_0 has an embedding in S;
2. the graphs G_i, 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 $(\mathcal{B}_u)_{u \in U_i}$ 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 ([RS03]). *For every graph H, there exists an integer $h \ge 0$ 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.

3.4 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 [DH04b] suggest that the decomposition can be restricted to a particular form in the apex-minor-free case:

Conjecture 1 ([DH04b]). 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 [RST94] 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 4 ([DH05b]). *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 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 [RS86b]: 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 re-proved by Robertson, Seymour, and Thomas [RST94], Reed [Ree97], and Diestel, Jensen, Gorbunov, and Thomassen [DJGT99]. 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 [RST94]. Theorem 4 therefore offers an exponential (and best possible) improvement over previous results.

Theorem 4 cannot be generalized to arbitrary graphs: Robertson, Seymour, and Thomas [RST94] 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 [RST94]. The best possible bound is believed to be closer to $\Theta(r^2 \lg r)$ than $2^{\Theta(r^5)}$, perhaps even equal to $\Theta(r^2 \lg r)$ [RST94].

4 Bidimensional Parameters/Problems

Bidimensionality has been introduced and developed in a series of papers [DHT, DHN$^+$04, DFHT, DH04a, DFHT04b, DH04b, DFHT04a, DHT04, DH05b, DH05a]. Although implicitly hinted at in [DHT, DHN$^+$04, DFHT, DH04a], the first use of the term "bidimensional" was in [DFHT04b].

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 $P(G) \le k$. 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:

1. 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.
2. 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").
3. 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).[2]

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, and unweighted TSP tour (a walk in the graph visiting all vertices). (See [DFHT04b, DFHT04a] for arguments of either contraction- or minor-bidimensionality for the above parameters.) We also say that the corresponding optimization problems based on these parameters, e.g., finding the minimum-size dominating

[2] 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".

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.

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., [ABF⁺02, KP02, FT03, AFN04, CKL01, KLL02, GKL01, DFHT, DHN⁺04, DHT, DFHT04a, DH04b, DFHT04b]. 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. More precisely, we have the following result:

Theorem 5 ([DH05b, DFHT04a]). *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 $\mathbf{tw}(G) = O(\sqrt{P(G)})$.*

This theorem is based on the grid-minor bound from Theorem 4 and the proof of a weaker parameter-treewidth bound, $\mathbf{tw}(G) = (g^{-1}(P(G)))^{O(g^{-1}(P(G)))}$, established in [DFHT04a]. The stronger bound of $\mathbf{tw}(G) = O(g^{-1}(P(G)))$ was obtained first for planar graphs [DFHT], then single-crossing-minor-free graphs [DHT, DHN⁺04], then bounded-genus graphs [DFHT04b, DHT04], and finally apex-minor-free graphs for contraction-bidimensional parameters and H-minor-free graphs for minor-bidimensional parameters [DH05b] (Theorem 5 above).

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 [RST94, DH04c], 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 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 ([AST90, Proposition 4.5], [Gro03, Corollary 24], [DH05b]). *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 [AST90, Theorem 1.2]. This generalization of the classic planar separator theorem has many algorithmic applications; see e.g. [AST90, AFN03]. Also, this result shows that the structural properties of H-minor-free graphs given by Theorem 3 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 [Epp00] 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 [DH04a].)

If we apply the parameter-treewidth bound of Theorem 5 to the diameter parameter, which is contraction-bidimensional with $g(r) = \Theta(\lg r)$ [DH04a], then we immediately obtain the following stronger diameter-treewidth bound for apex-minor-free graphs:

Theorem 7 ([DH05b]). *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. 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, 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 8 ([DH04b]). *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 8 is therefore the ultimate characterization of diameter-treewidth bounds in minor-closed graph families (up to constant factors).

The proof of Theorem 8 is the basis for Conjecture 1. In fact, Theorem 8 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 polynomial-time approximation schemes (PTASs) on planar graphs [Bak94] 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 [DHN$^+$04, Epp00, HN02]. Graphs of bounded local treewidth also admit several efficient fixed-parameter algorithms. In particular, Frick and Grohe [FG01] give a general framework for deciding any property expressible in first-order logic in graphs of bounded local treewidth. Theorem 8 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 [AFF$^+$01] has running time $O(8^k n)$; subsequently, a sequence of subexponential algorithms and improvements have been obtained, starting with running time $O(4^{6\sqrt{34k}} n)$ [ABF$^+$02], then $O(2^{27\sqrt{k}} n)$ [KP02], and finally $O(2^{15.13\sqrt{k}} k + n^3 + k^4)$ [FT03]. Other subexponential algorithms for other domination and covering problems on planar graphs have also been obtained [ABF+02, AFN04, CKL01, KLL02, GKL01].

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 and the techniques of [DFHT04a] yield the following general result for designing subexponential fixed-parameter algorithms:

Theorem 9 ([DH05b, DFHT04a]). *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, and unweighted TSP tour.

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 [DFHT04b], map graphs [DFHT], and fixed powers of planar graphs (or even fixed powers of H-minor-free graphs) [DFHT, DFHT04b]. These algorithms are necessarily more complicated than those produced from Theorem 9, because apex-minor-free graphs are precisely the minor-closed graph classes for which domatinating set has a parameter-treewidth bound [DFHT04a]. 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 [RS95, RS]) 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 [RS95]. 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 [FL88].

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" [FL88]. Indeed, there is a mathematical sense in which any proof of the finite-forbidden-minors theorem must be nonconstructive [FRS87].

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 [DH04c], 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 the generalized form of minor-bidimensionality.

Theorem 10 ([DH04c]). *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 [DH04c], a conjecture of Robertson, Seymour, and Thomas [RST94] 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 [FL88].

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 [DH05a]. 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 [LT80] and the Baker layerwise decomposition approach [Bak94]. 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, [DH05a] uses the sublinear parameter-treewidth bound of Theorem 5 as well as a recent $O(1)$-approximation algorithm for treewidth in H-minor-free graphs [FHL04].

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:

1. 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.
2. 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 $\mathrm{OPT}(G - C) + O(|C|)$.
3. 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 11 ([DH05a]). *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.*

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 [RST94] 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. If this conjecture is correct, we would obtain nearly as good parameter-treewidth bounds for minor-bidimensional parameters as in the H-minor-free case. A similar generalization of parameter-treewidth bounds beyond apex-minor-free graphs is not possible for all contraction-bidimensional parameters, e.g., dominating set [DFHT04a], 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.

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 4. 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 [AST90] has a lead factor of $|V(H)|^{3/2}$. In fact, Alon, Seymour, and Thomas [AST90] 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 4, which would imply the corresponding bound for separators.

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 [AGK+98] and subset feedback vertex set [ENZ00]. 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.

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

References

[ABF+02] 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(4):461–493, 2002.

[AFF+01] 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 *Proc. 26th Internat. Symp. Math. Found. Comput. Sci.*, LNCS 2136, pp. 111–122, 2001.

[AFN03] J. Alber, H. Fernau, and R. Niedermeier. Graph separators: a parameterized view. *J. Comput. System Sci.*, 67(4):808–832, Dec. 2003.

[AFN04] J. Alber, H. Fernau, and R. Niedermeier. Parameterized complexity: exponential speed-up for planar graph problems. *J. Algorithms*, 52(1):26–56, 2004.

[AGK+98] S. Arora, M. Grigni, D. Karger, P. Klein, and A. Woloszyn. A polynomial-time approximation scheme for weighted planar graph TSP. In *Proc. 9th ACM-SIAM Symp. Discr. Alg.*, pp. 33–41, 1998.

[AST90] N. Alon, P. Seymour, and R. Thomas. A separator theorem for nonplanar graphs. *J. Amer. Math. Soc.*, 3(4):801–808, 1990.

[Bak94] B. S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. Assoc. Comput. Mach.*, 41(1):153–180, 1994.

[CGP02] Z.-Z. Chen, M. Grigni, and C. H. Papadimitriou. Map graphs. *J. ACM*, 49(2):127–138, 2002.

[CKL01] M.-S. Chang, T. Kloks, and C.-M. Lee. Maximum clique transversals. In *Proc. 27th Internat. Workshop Graph-Theor. Concepts in Comput. Sci.*, LNCS 2204, pp. 32–43, 2001.

[DFHT] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Fixed-parameter algorithms for the (k, r)-center in planar graphs and map graphs. *ACM Trans. Alg.*. To appear. A preliminary version appears in *Proc. 30th Internat. Colloq. Automata Lang. Prog.*, LNCS 2719, 2003, pp. 829–844.

[DFHT04a] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Bidimensional parameters and local treewidth. *SIAM J. Discrete Math.*, 2004. To appear. A preliminary version appears in *Proc. 6th Latin Amer. Symp. Theoret. Inf.*, LNCS 2976, 2004, pp. 109–118.

[DFHT04b] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs. In *Proc. 15th ACM-SIAM Symp. Discr. Alg.*, pp. 823–832, 2004.

[DH04a] E. D. Demaine and M. Hajiaghayi. Diameter and treewidth in minor-closed graph families, revisited. *Algorithmica*, 40(3):211–215, Aug. 2004.

[DH04b] E. D. Demaine and M. Hajiaghayi. Equivalence of local treewidth and linear local treewidth and its algorithmic applications. In *Proc. 15th ACM-SIAM Symp. Discr. Alg.*, pp. 833–842, 2004.

[DH04c] E. D. Demaine and M. Hajiaghayi. Quickly deciding minor-closed parameters in general graphs. Manuscript, 2004.

[DH05a] E. D. Demaine and M. Hajiaghayi. Bidimensionality: New connections between FPT algorithms and PTASs. In *Proc. 16th ACM-SIAM Symp. Discr. Alg.*, 2005. To appear.

[DH05b] E. D. Demaine and M. Hajiaghayi. Graphs excluding a fixed minor have grids as large as treewidth, with combinatorial and algorithmic applications through bidimensionality. In *Proc. 16th ACM-SIAM Symp. Discr. Alg.*, 2005. To appear.

[DHN+04] 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. *J. Comput. System Sci.*, 69(2):166–195, 2004.

[DHT] 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*. To appear. A preliminary version appears in *Proc. 13th Internat. Symp. Alg. Comput.*, LNCS 2518, 2002, pp. 262–273.

[DHT04] E. D. Demaine, M. Hajiaghayi, and D. M. Thilikos. The bidimensional theory of bounded-genus graphs. In *Proc. 29th Internat. Symp. Math. Found. Comput. Sci.*, pp. 191–203, 2004.

[DJGT99] R. Diestel, T. R. Jensen, K. Y. Gorbunov, and C. Thomassen. Highly connected sets and the excluded grid theorem. *J. Combin. Theory Ser. B*, 75(1):61–73, 1999.

[ENZ00] G. Even, J. Naor, and L. Zosin. An 8-approximation algorithm for the subset feedback vertex set problem. *SIAM J. Comput.*, 30(4):1231–1252, 2000.

[Epp95] D. Eppstein. Subgraph isomorphism in planar graphs and related problems. In *Proc. 6th ACM-SIAM Symp. Discr. Alg.*, pp. 632–640, 1995.

[Epp00] D. Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27(3-4):275–291, 2000.

[FG01] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48(6):1184–1206, 2001.

[FHL04] U. Feige, M. Hajiaghayi, and J. R. Lee. On improving approximate vertex separator and its algorithmic applications. Manuscript, 2004.

[FL88] M. R. Fellows and M. A. Langston. Nonconstructive tools for proving polynomial-time decidability. *Journal of the ACM*, 35(3):727–739, 1988.

[FRS87] H. Friedman, N. Robertson, and P. Seymour. The metamathematics of the graph minor theorem. In *Logic and combinatorics*, vol. 65 of *Contemp. Math.*, pp. 229–261. Amer. Math. Soc., Providence, RI, 1987.

[FT03] F. V. Fomin and D. M. Thilikos. Dominating sets in planar graphs: Branch-width and exponential speed-up. In *Proc. 14th ACM-SIAM Sympos. Discrete Algorithms*, pp. 168–177, 2003.

[GKL01] G. Gutin, T. Kloks, and C. M. Lee. Kernels in planar digraphs. In *Optimization Online*, 2001.

[Gro03] M. Grohe. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003.

[HN02] M. Hajiaghayi and N. Nishimura. Subgraph isomorphism, log-bounded fragmentation and graphs of (locally) bounded treewidth. In *Proc. 27th Internat. Symp. Math. Found. Comput. Sci.*, LNCS 2420, pp. 305–318, 2002.

[KLL02] 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 *Proc. 28th Internat. Workshop Graph-Theor. Concepts Comput. Sci.*, LNCS 2573, pp. 282–295, 2002.

[KP02] I. Kanj and L. Perković. Improved parameterized algorithms for planar dominating set. In *Proc. 27th Internat. Symp. Math. Found. Comput. Sci.*, LNCS 2420, pp. 399–410, 2002.

[LT80] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.

[Ree97] B. A. Reed. Tree width and tangles: a new connectivity measure and some applications. In *Surveys in combinatorics*, vol. 241 of *London Math. Soc. Lecture Note Ser.*, pp. 87–162. Cambridge Univ. Press, 1997.

[RS] N. Robertson and P. D. Seymour. Graph minors. XX. Wagner's conjecture. *J. Combin. Theory Ser. B.* To appear.

[RS86a] N. Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.

[RS86b] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *J. Combin. Theory Ser. B*, 41(1):92–114, 1986.

[RS93] N. Robertson and P. Seymour. Excluding a graph with one crossing. In *Graph structure theory*, pp. 669–675. Amer. Math. Soc., 1993.

[RS95] N. Robertson and P. D. Seymour. Graph minors. XII. Distance on a surface. *J. Combin. Theory Ser. B*, 64(2):240–272, 1995.

[RS03] N. Robertson and P. D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Combin. Theory Ser. B*, 89(1):43–76, 2003.

[RST94] N. Robertson, P. D. Seymour, and R. Thomas. Quickly excluding a planar graph. *J. Combin. Theory Ser. B*, 62(2):323–348, 1994.

[Tho98] M. Thorup. Map graphs in polynomial time. In *Proc. 39th IEEE Sympos. Found. Comput. Sci.*, pp. 396–407, 1998.

[Wag37] K. Wagner. Über eine Eigenschaft der eben Komplexe. *Deutsche Math.*, 2:280–285, 1937.