Subexponential Parameterized Algorithms on Bounded-Genus Graphs and H-Minor-Free Graphs

ERIK D. DEMAINE
MIT, Cambridge, Massachusetts
FEDOR V. FOMIN
University of Bergen, Bergen, Norway
MOHAMMADTAGHI HAJIAGHAYI
MIT, Cambridge, Massachusetts
and
DIMITRIOS M. THILIKOS
National and Capodistrian University of Athens

A preliminary version of this article appeared in *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms*, 2004, pp. 823–832.
F. V. Fomin is supported by Norges forskningsråd projects 160233/V30 and 160778/V30.
D. M. Thilikos is supported by EC contract IST-1999-14186: Project ALCOM-FT (Algorithms and Complexity)—Future Technologies and by the Spanish CICYT project TIC-2002-04498-C05-03 (TRACER).
Author's addresses: E. D. Demaine, M. Hajiaghayi, MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar Street, Cambridge, MA 02139, USA, e-mail: {edemaine,hajiagha}@mit.edu; F. V. Fomin, Department of Informatics, University of Bergen, N-5020 Bergen, Norway, e-mail: fomin@ii.uib.no; D. M. Thilikos, Department of Mathematics, National and Capodistrian University of Athens, Panepistimioupolis, GR-15784, Athens, Greece, e-mail: sedthilk@math.uoa.gr.

Abstract. We introduce a new framework for designing fixed-parameter algorithms with subexponential running time—$2^{O(\sqrt{k})}n^{O(1)}$. Our results apply to a broad family of graph problems, called *bidimensional problems*, which includes many domination and covering problems such as vertex cover, feedback vertex set, minimum maximal matching, dominating set, edge dominating set, disk dimension, and many others restricted to bounded-genus graphs. Furthermore, it is fairly straightforward to prove that a problem is bidimensional. In particular, our framework includes, as special cases, all previously known problems to have such subexponential algorithms. Previously, these algorithms applied to planar graphs, single-crossing-minor-free graphs, and/or map graphs; we extend these results to apply to bounded-genus graphs as well. In a parallel development of combinatorial results, we establish an upper bound on the treewidth (or branchwidth) of a bounded-genus graph that excludes some planar graph H as a minor. This bound depends linearly on the size $|V(H)|$ of the excluded graph H and the genus $g(G)$ of the graph G, and applies and extends the graph-minors work of Robertson and Seymour.

Building on these results, we develop subexponential fixed-parameter algorithms for dominating set, vertex cover, and set cover in any class of graphs excluding a fixed graph H as a minor. In particular, this general category of graphs includes planar graphs, bounded-genus graphs, single-crossing-minor-free graphs, and any class of graphs that is closed under taking minors. Specifically, the running time is $2^{O(\sqrt{k})}n^h$, where h is a constant depending only on H, which is polynomial for $k = O(\log^2 n)$. We introduce a general approach for developing algorithms on H-minor-free graphs, based on structural results about H-minor-free graphs at the heart of Robertson and Seymour's graph-minors work. We believe this approach opens the way to further development on problems in H-minor-free graphs.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*computations on discrete structures*; G.2.2 [**Mathematics of Computing**]: Discrete Mathematics—*graph algorithms, network problems*

General Terms: Algorithms, Design, Theory

Additional Key Words and Phrases: (k, r)-center, fixed-parameter algorithms, domination, planar graph, map graph

1. INTRODUCTION

Dominating set is a classic NP-complete graph optimization problem which fits into the broader class of *domination* and *covering* problems on which hundreds of papers have been written; see, for example, the survey [37]. A sample application is the problem of locating sites for emergency service facilities such as fire stations. Here we suppose that we can afford to build k fire stations to cover a city, and we require that every building is covered by at least one fire station. The problem is to find a dominating set of size k in the graph where edges represent suitable pairings of fire stations with buildings. In this application, we can afford high running time (e.g., several weeks of real time) if the resulting solution builds fewer fire stations (which are extremely expensive). Thus, we prefer exact *fixed-parameter* algorithms (which run fast provided the parameter k is small) over approximation algorithms, even if the approximation were within an additive constant. The theory of fixed-parameter algorithms and parameterized complexity has been thoroughly developed over the past few years; see, for example, [15; 27; 29; 30; 34; 4; 3].

In the last two years, several researchers have obtained exponential speedups in fixed-parameter algorithms for various problems on several classes of graphs. While most previous fixed-parameter algorithms have a running time of $2^{O(k)}n^{O(1)}$ or worse, the exponential speedups result in subexponential algorithms with running times of $2^{O(\sqrt{k})}n^{O(1)}$. For example, the first fixed-parameter algorithm for dominating set in planar graphs [2] has running time $O(8^k n)$; subsequently, a sequence of subexponential algorithms and improvements have been obtained, starting with running time $O(4^{6\sqrt{34k}}n)$ [1], then $O(2^{27\sqrt{k}}n)$ [38], and finally $O(2^{15.13\sqrt{k}}k+n^3+k^4)$ [30]. Other subexponential algorithms for other domination and covering problems on planar graphs have also been obtained [1; 4; 11; 40; 36].

However, all of these algorithms apply only to planar graphs. In another sequence of papers, these results have been generalized to wider classes of graphs: map graphs [15], which include planar graphs; $K_{3,3}$-minor-free graphs and K_5-minor-free graphs [24], which include planar graphs; and single-crossing-minor-free graphs [23; 24], which include $K_{3,3}$- and K_5-minor-free graphs. These algorithms [15; 23; 24] apply to dominating set and several other problems related to domination, covering, and logic.

Algorithms for H-minor-free graphs for a fixed graph H have been studied extensively; see, for example, [12; 35; 13; 39; 42]. In particular, it is generally believed that several algorithms for planar graphs can be generalized to H-minor-free graphs for any fixed H [35; 39; 42]. H-minor-free graphs are very general. The deep Graph-Minor Theorem of Robertson and Seymour shows that any graph class that is closed under minors is characterized by excluding a finite set of minors. In particular, any graph class that is closed under minors (other than the class of all graphs) excludes at least one minor H.

Our Results.. We introduce a framework for extending algorithms for planar graphs to apply to H-minor-free graphs for any fixed H. In particular, we design subexponential fixed-parameter algorithms for dominating set, vertex cover, and set cover (viewed as one-sided domination in a bipartite graph) for H-minor-free graphs. Our framework consists of three components, as described below. We believe that many of these components can be applied to other problems and conjectures as well.

First, we extend the algorithm for planar graphs to bounded-genus graphs. Roughly speaking, we study the structure of the solution to the problem in $k \times k$ grids, which form a representative substructure in both planar graphs and bounded-genus graphs, and capture the main difficulty of the problem for these graphs. Then using Robertson and Seymour's graph-minor theory, we repeatedly remove handles to reduce the bounded-genus graph down to a planar graph, which is essentially a grid.

Second, we extend the algorithm to *almost-embeddable* graphs that can be drawn in a bounded-genus surface except for a bounded number of "local areas of non-planarity", called vortices, and for a bounded number of "apex" vertices, which can have any number of incident edges that are not properly embedded. Because each vortex has bounded pathwidth, the number of vortices is bounded, and the number of apices is bounded, we are able to extend our approach to solve almost-embeddable graphs using our solution to bounded-genus graphs.

Third, we apply a deep theorem of Robertson and Seymour, which characterizes H-minor-free graphs as a tree structure of pieces, where each piece is an almost-embeddable graph. Using dynamic programming on such tree structures, analogous to algorithms for graphs of bounded treewidth, we are able to combine the pieces and solve the problem for H-minor-free graphs. Note that the standard bounded-treewidth methods do not suffice for general H-minor-free graphs, unlike, for example, e.g. bounded-genus graphs, because their treewidth can be arbitrarily large with respect to the parameter [14]. Our contribution is to overcome this barrier algorithmically using a two-level dynamic program in a more general tree structure called a "clique-sum decomposition".

The first step of this procedure, for bounded-genus graphs, applies to a broad class of problems called "bidimensional problems". Roughly speaking, a parameterized graph problem is *bidimensional* if the parameter is large enough (linear) in a grid and closed under contractions. Examples of bidimensional problems include vertex cover, feedback vertex set, minimum maximal matching, dominating set, edge dominating set, set cover, disk dimension, and TSP tour (in the shortest-path metric of the graph). We obtain subexponential fixed-parameter algorithms for all of these problems in bounded-genus graphs. As a special case, this generalization settles an open problem about dominating set posed by Ellis, Fan, and Fellows [28]. Along the way, we establish an upper bound on the treewidth (or branchwidth) of a bounded-genus graph that excludes some planar graph H as a minor. This bound depends linearly on the size $|V(H)|$ of the excluded graph H and the genus $g(G)$ of the graph G, and applies and extends the graph-minors work of Robertson and Seymour.

This article forms the basis of several more recent papers, for example, [17; 14; 31; 22; 16; 19; 18]. In particular, the theory of bidimensionality introduced in this article has flourished into a comprehensive body of algorithmic and combinatorial results. The consequences of this theory include tight parameter-treewidth bounds, direct seperator theorems, linearity of local treewidth, subexponential fixed-parameter algorithms, and polynomial-time approximation schemes for a broad class of problems on graphs that exclude a fixed minor. In Section 6 we describe some of these results in comparison to this article.

This article is organized as follows. First, we introduce the terminology used throughout the article, and formally define tree decompositions, treewidth, and fixed-parameter tractability in Section 2. Section 3 is devoted to graphs on surfaces. We construct a general framework for obtaining subexponential parameterized algorithms on graphs of bounded genus. First we introduce the concept of bidimensional problem, and then prove that every bidimensional problem has a subexponential parameterized algorithm on graphs of bounded genus. The proof techniques used in this section are very indirect and are based on deep theorems from Robertson and Seymour's Graph Minors XI [46] and XII [47]. As a byproduct of our results we obtain a generalization of Quickly Excluding a Planar Graph Theorem [50] for graphs of bounded genus. In Section 5 we make a further step by developing subexponential algorithms for graphs containing no fixed graph H as a minor. The proof of this result is based on combinatorial bounds from the previous section, a deep structural theorem from Graph Minors XVI [49], and complicated dynamic programming. Finally, in Section 6, we present several extensions of our results and some open problems.

2. BACKGROUND

2.1 Preliminaries

All the graphs in this article are undirected without loops. The reader is referred to standard references for appropriate background [7]. In addition, for exact definitions of various NP-hard graph-theoretic problems in this article, the reader is referred to Garey and Johnson [32].

Our graph terminology is as follows. A graph G is represented by $G = (V, E)$, where V (or $V(G)$) is the set of vertices and E (or $E(G)$) is the set of edges. We denote an edge e between u and v by $\{u, v\}$. We define n to be the number of vertices of a graph when this is clear from context. For every subset $W \subseteq V(G)$ of vertices, the subgraph of G induced by W is denoted by $G[W]$. We define the *q-neighborhood* of a vertex $v \in V(G)$, denoted by $N_G^q[v]$, to be the set of vertices of G at distance at most q from v. Notice that $v \in N_G^q[v]$. We define $N_G[v] = N_G^1[v]$ and $N_G(v) = N_G[v] - \{v\}$.

The *(disjoint) union* of two disjoint graphs G_1 and G_2, $G_1 \cup G_2$, is the graph G with merged vertex and edge sets: $V(G) = V(G_1) \cup V(G_2)$ and $E(G) = E(G_1) \cup E(G_2)$.

One way of describing classes of graphs is by using *minors*. Given an edge $e = \{u, v\}$ of a graph G, the graph G/e is obtained from G by contracting the edge e; that is, to get G/e we identify the vertices u and v and remove all loops and duplicate edges. A graph H obtained by a sequence of edge contractions 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 $H \notin \mathcal{C}$.

For example, a planar graph is a graph excluding both $K_{3,3}$ and K_5 as minors (Kuratowski's Theorem).

2.2 Fixed-Parameter Algorithms

Developing fast algorithms for NP-hard problems is an important issue. Downey and Fellows [27] formalized a new approach to cope with NP-hardness, called *fixed-parameter tractability*. For many NP-complete problems, the inherent combinatorial explosion can be attributed to a certain aspect of the problem, a *parameter*. The parameter is often an integer that is small in practice. The running times of simple algorithms may be exponential in the parameter but polynomial in the rest of the problem size. A problem is *fixed-parameter tractable* if it has an algorithm whose running time is $f(k)n^{O(1)}$ where n is the problem size, k is the parameter value, and f is any function (typically, $2^{\Theta(k)}$). For example, it has been shown that a vertex cover of size k can be found in $O(1.2745^k k^4 + kn)$ time [10], and hence this problem is fixed-parameter tractable.

Alber et al. [1] demonstrated a solution to finding a dominating set of size k in a planar graph in $O(4^{6\sqrt{34k}} n)$ time. This result was the first nontrivial result for the parameterized version of an NP-hard problem where the exponent of the exponential term grows sublinearly in the parameter (see also [38] and [30] for further improvements of the time bound of [1]) and it initiated the extensive study of subexponential algorithms for various parameterized problems on planar graphs. Using this result, others could obtain exponential speedup of fixed-parameter algorithms for many NP-complete problems on planar graphs (see, e.g., [11; 40; 4; 8]). (See also Cai and Juedes [9] for discussions on lower bounds of subexponential algorithms on planar graphs.) Recently, Demaine et al. [24; 23; 15] extended these results to many NP-complete problems on map graphs and graphs excluding a single-crossing-graph such as K_5 or $K_{3,3}$ as a minor. As mentioned before, we extend these results for bounded-genus graphs and more generally H-minor-free graphs for any fixed H.

2.3 Treewidth and Branchwidth

The notion of treewidth was introduced by Robertson and Seymour [43] and plays an important role in their fundamental work on graph minors. To define this notion, first we consider the representation of a graph by a tree, which is the basis of our algorithms in this article.

A *tree decomposition* of a graph G is a pair (T, χ) where T is a tree and $\chi = \{\chi_i \mid i \in V(T)\}$ is a family of subsets of $V(G)$ such that

(1) $\bigcup_{i \in V(T)} \chi_i = V(G)$;

(2) for each edge $e = \{u, v\} \in E(G)$, there is an $i \in V(T)$ such that both u and v belong to χ_i;

(3) for all $v \in V(G)$, the set of nodes $\{i \in V(T) \mid v \in \chi_i\}$ forms a connected subtree of T.

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

A *branch decomposition* of a graph G is a pair (T, τ) where T is a tree with vertices of degree 1 or 3 and τ is a bijection from the set of leaves of T to $E(G)$. The *order* of an edge e in T is the number of vertices $v \in V(G)$ such that there are leaves t_1, t_2 in T in different components of $T - e = (V(T), E(T) - e)$ with $\tau(t_1)$ and $\tau(t_2)$ both containing v as an endpoint. The *width* of (T, τ) is the maximum order over all edges of T, and the *branchwidth* of G, denoted $\mathbf{bw}(G)$, is the minimum width over all branch decompositions of G. (In the case $|E(G)| \leq 1$, we define the branchwidth to be 0; if $|E(G)| = 0$, then G has no branch decomposition; if $|E(G)| = 1$, then G has a branch decomposition consisting of a tree with one vertex, and the width of this branch decomposition is considered to be 0.)

It is known that, if H is a minor of G, then $\mathbf{tw}(H) \leq \mathbf{tw}(G)$ and $\mathbf{bw}(H) \leq \mathbf{bw}(G)$ [45]. The following connection between treewidth and branchwidth is due to Robertson and Seymour:

THEOREM 2.1. [45, Theorem 5.1] *For any connected graph G where $|E(G)| \geq 3$, $\mathbf{bw}(G) \leq \mathbf{tw}(G) + 1 \leq \frac{3}{2}\mathbf{bw}(G)$.*

3. GRAPHS ON SURFACES

3.1 Preliminaries

In this section we describe some of the machinery developed in the Graph Minors series that we use in our proofs. See also [46].

A *surface* Σ is a connected compact 2-manifold without boundary. A *line* in Σ is a subset homeomorphic to $[0, 1]$. An O-arc is a subset of Σ homeomorphic to a circle. A subset of Σ is an *open disk* if it is homeomorphic to $\{(x, y) \mid x^2 + y^2 < 1\}$, and it is a *closed disk* if it is homeomorphic to $\{(x, y) \mid x^2 + y^2 \leq 1\}$.

A 2-*cell embedding* of a graph G in a surface Σ is a drawing of the vertices as points in Σ and the edges as lines in Σ such that every region (face) bounded by edges is an open disk. To simplify notation, we do not distinguish between a vertex of G and the point of Σ used in the drawing to represent the vertex, or between an edge and the line representing it. We also consider G as the union of points corresponding to its vertices and edges. Also, a subgraph H of G can be seen as a graph H where $H \subseteq G$. A *region* of G is a connected component of $\Sigma - E(G) - V(G)$. (Every region is an open disk.) We use the notation $V(G)$, $E(G)$, and $R(G)$ for the set of the vertices, edges, and regions of G.

If $\Delta \subseteq \Sigma$, then $\overline{\Delta}$ denotes the *closure* of Δ, and the boundary of Δ is $\mathbf{bd}(\Delta) = \overline{\Delta} \cap \overline{\Sigma - \Delta}$. A vertex or an edge x is *incident* to a region r if $x \subseteq \mathbf{bd}(r)$.

A subset of Σ meeting the drawing only at vertices of G is called G-*normal*. If an O-arc is G-normal, then we call it a *noose*. The *length* of a noose is the number of vertices it meets. We say that a disk D is *bounded* by a noose N if $N = \mathbf{bd}(D)$. A graph G 2-cell embedded in a connected surface Σ is θ-*representative* if every noose of length less than θ is contractable (null-homotopic in Σ).

Tangles were introduced by Robertson and Seymour in [45]. A *separation* of a graph G is a pair (A, B) of subgraphs with $A \cup B = G$ and $E(A \cap B) = \emptyset$, and its *order* is $|V(A \cap B)|$. A *tangle of order* $\theta \geq 1$ is a set $\mathcal{T}$ of separations of G, each of order less than θ, such that

(1) for every separation (A, B) of G of order less than θ, $\mathcal{T}$ contains one of (A, B) and (B, A);
(2) if $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \mathcal{T}$, then $A_1 \cup A_2 \cup A_3 \neq G$; and
(3) if $(A, B) \in \mathcal{T}$, then $V(A) \neq V(G)$.

Let G be a graph 2-cell embedded in a connected surface Σ. A tangle $\mathcal{T}$ of order θ is *respectful* if, for every noose N in Σ of length less than θ, there is a closed disk $\Delta \subseteq \Sigma$ with $\mathbf{bd}(\Delta) = N$ such that the separation $(G \cap \Delta, G \cap \overline{\Sigma - \Delta}) \in \mathcal{T}$.

Our proofs are based on the following results from the Graph Minors series of papers by Robertson and Seymour.

THEOREM 3.1. [45, Theorem 4.3] *Let G be a graph with at least one edge. Then there is a tangle in G of order θ if and only if G has branchwidth at least θ.*

THEOREM 3.2. [46, Theorem 4.1] *Let Σ be a connected surface, not homeomorphic to a sphere; let $\theta \geq 1$; and let G be a θ-representative graph 2-cell embedded in Σ. Then there is a unique respectful tangle in G of order θ.*

Roughly speaking, a tangle of order θ assigns a notion of "inside" for each separation of order at most θ. Theorem 3.2 says that, if the surface has positive genus and the embedding is θ-representative, then every separation of order θ splits Σ into parts in such a way that exactly one part is homeomorphic to a disk, and a tangle selects the corresponding component of the graph. When the surface is the sphere, this partition is more ambiguous, and the tangle disambiguates which part is considered "inside". See [46, Section 1] for more intuition.

Our proofs also use the notion of the radial graph. Informally, the radial graph of a graph G 2-cell embedded in Σ is the bipartite graph R_G obtained by selecting a point in every region r of G and connecting it via an edge to every vertex of G incident to that region. However, a region may be incident to the same vertex "more than once", so we need a more formal definition. Precisely, R_G is a *radial graph* of a graph G 2-cell embedded in Σ if

(1) $E(G) \cap E(R_G) = V(G) \subseteq V(R_G)$;
(2) each region $r \in R(G)$ contains a unique vertex $v_r \in V(R_G)$;
(3) R_G is bipartite with a bipartition $(V(G), \{v_r : r \in R(G)\})$;
(4) if e, f are edges of R_G with the same ends $v \in V(G)$, $v_r \in V(R_G)$, then $e \cup f$ does not bound a closed disk in $r \cup \{v\}$; and
(5) R_G is maximal subject to Conditions (1)–(4).

The radial graph is unique up to isomorphism [46, Section 3].

3.2 Bounding the Representativity

Define the $(r \times r)$-*grid* to be the graph on r^2 vertices $\{(x, y) \mid 1 \le x, y \le r\}$ with edges between vertices differing by ± 1 in exactly one coordinate. A *partially triangulated* $(r \times r)$-*grid* is any planar supergraph of the $(r \times r)$-grid with the same set of vertices.

LEMMA 3.3. *Let G be a graph 2-cell embedded in a surface Σ, not homeomorphic to a sphere, of representativity at least $4r > 0$. Then G contains as a contraction a partially triangulated $(r \times r)$-grid.*

PROOF. Let $\theta = 4r$ be (a lower bound on) the representativity of G. By Theorem 3.2, G has a respectful tangle of order θ. Let $A(G)$ be the set of vertices, edges, and regions (collectively, *atoms*) of the graph G. According to [46, Section 9] (see also [47]), the existence of a respectful tangle of order θ makes it possible to define a metric d on $A(G)$ as follows:

(1) If $a = b$, then $d(a, b) = 0$.
(2) If $a \ne b$, and a and b are interior to a contractible closed walk in the radial graph R_G of length less than 2θ, then $d(a, b)$ is half the minimum length of such a walk. (Here by *interior* we mean the direction in which the walk can be contracted, and we include the boundary as part of the interior.)
(3) Otherwise, $d(a, b) = \theta$.

Let c be any vertex in G; refer to Figure 1. For $0 \le i < \theta$, define Z_i to be the union of all atoms of distance at most i from c (where distance is measured according to the metric d). By [46, Theorem 8.10], Z_i is a nonempty simply connected set, for all i. (A subset of a surface is *simply connected* if it is connected and has no

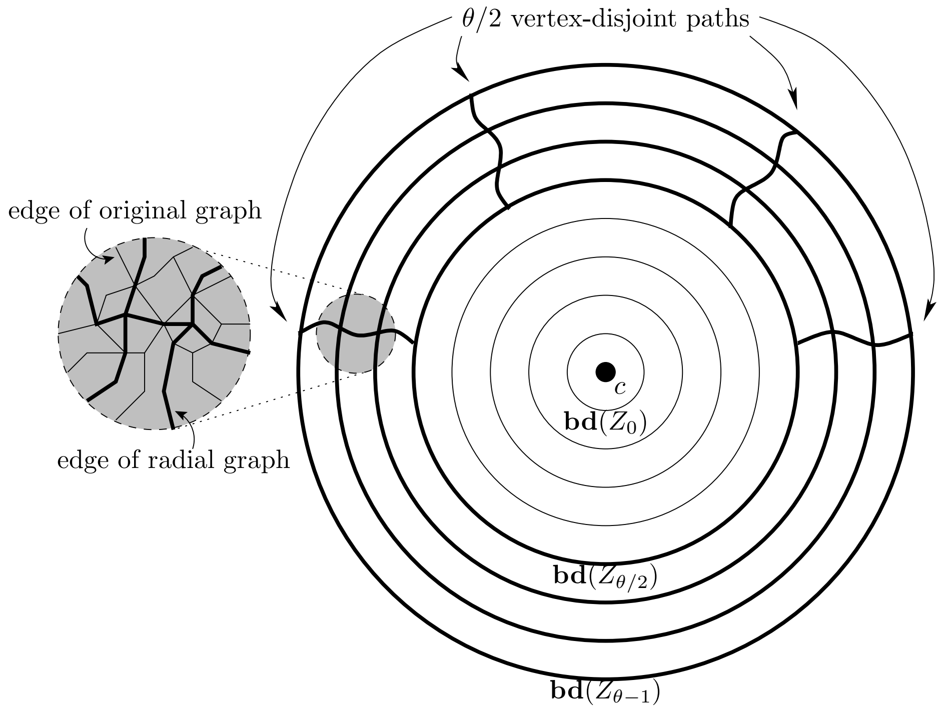

Fig. 1. The radial graph in the proof of Lemma 3.3.

noncontractible closed curves.) Thus, the boundary $\mathbf{bd}(Z_i)$ of each Z_i is a closed walk in the radial graph.

We claim that the closed walks $\mathbf{bd}(Z_i)$ and $\mathbf{bd}(Z_{i+1})$ are vertex-disjoint. Consider any vertex a of R_G on $\mathbf{bd}(Z_i)$ and an adjacent vertex b of R_G outside Z_i. The distance between a and b, measured according to d, is 1 because there is a length-2 closed walk connecting them, doubling the edge (a, b) in the radial graph. By [46, Theorem 9.1], the metric d satisfies the triangle inequality, and hence $d(c, b) \leq d(c, a) + 1 = i + 1$. In fact, this bound must hold with equality, because $b \notin Z_i$. Therefore, every vertex a of R_G on $\mathbf{bd}(Z_i)$ is surrounded on the exterior of Z_i by vertices of R_G at distance exactly $i+1$ from c, so $\mathbf{bd}(Z_i)$ is strictly enclosed by $\mathbf{bd}(Z_{i+1})$.

Consider the "annulus" $\mathcal{A} = (Z_{\theta-1} - Z_{\theta/2}) \cup \mathbf{bd}(Z_{\theta-1}) \cup \mathbf{bd}(Z_{\theta/2})$, which includes the boundary $\mathbf{bd}(\mathcal{A}) = \mathbf{bd}(Z_{\theta-1}) \cup \mathbf{bd}(Z_{\theta/2})$. We claim that there are at least $\theta/2$ vertex-disjoint paths in R_G within $\mathcal{A}$ connecting vertices of R_G in $\mathbf{bd}(Z_{\theta/2})$ to vertices of R_G in $\mathbf{bd}(Z_{\theta-1})$. By Menger's Theorem, the contrary implies the existence of a cut in $\mathcal{A}$ of size less than $\theta/2$ separating the two sets, which by simple connectedness (essentially, planarity) of $Z_{\theta-1}$ implies the existence of a cycle of length less than θ that separates the two sets, but such a cycle must be contained in $Z_{\theta/2}$.

Now we form a $(\theta/2 \times \theta/2)$-grid in the radial graph. The row lines in the grid are formed by taking, for each $i = \theta/2, \theta/2+1, \theta/2+2, \ldots, \theta-1$, the unique simple cycle that encloses c and that is a subset of the closed walk $\mathbf{bd}(Z_i)$. The column lines in the grid are formed by the $\theta/2$ vertex-disjoint paths found above. Therefore, we obtain a subdivision of the $(\theta/2 \times \theta/2)$-grid as a subgraph of the radial graph. Note that, by our construction, the rows of this grid can in fact form cycles, not just paths.

Finally, we transform this grid into a $(\theta/4 \times \theta/4)$-grid in the original graph G. Each row of the grid in the radial graph, viewed as a cycle C, corresponds in the original graph to a cyclic sequence of faces "surrounding" the row. We replace this row by the "inner half" of each face, that is, the unique simple cycle that encloses c, is enclosed by C, and whose edges are edges of these surrounding faces. In this way, each row in the radial graph maps in the original graph to a curve contained within this row line. Two adjacent mapped row lines may touch but cannot properly cross, so row lines of distance 2 or more in the grid cannot overlap when mapped to the original graph. Similarly, we can map each column of the grid in the radial graph to the original graph, trimming the ends to where they meet the second and last mapped rows (where the innermost row is considered first). Thus, by discarding the odd-numbered rows and columns, we obtain a subdivision of the $(\theta/4 \times \theta/4)$-grid in the original graph. Because $Z_{\theta-1}$ is simply connected, the grid is embedded in a simply connected subset of Σ, so if we apply contractions without deletions, we obtain a partially triangulated grid. $\square$

4. BIDIMENSIONAL PARAMETERS AND BOUNDED-GENUS GRAPHS

In this section, we define a general framework of parameterized problems for which subexponential algorithms with small constants can be obtained. Our framework is sufficiently broad that an algorithmic designer needs to check only two simple properties of any desired parameter to determine the applicability and practicality of our approach.

4.1 Definitions

Recall from Section 3.2 that a *partially triangulated* $(r \times r)$*-grid* is any planar graph obtained by adding edges between pairs of nonconsecutive vertices on a common face of a planar embedding of an $(r \times r)$-grid.

DEFINITION 4.1. A *parameter* P is any function mapping graphs to nonnegative integers. The *parameterized problem associated with* P asks, for some fixed k, whether $P(G) \le k$ for a given graph G.

DEFINITION 4.2. A parameter P is *minor bidimensional with density* δ if

(1) contracting or deleting an edge in a graph G cannot increase $P(G)$, and
(2) for the $(r \times r)$-grid R, $P(R) = (\delta r)^2 + o((\delta r)^2)$.

A parameter P is called *contraction bidimensional with density* δ if

(1) contracting an edge in a graph G cannot increase $P(G)$,
(2) for any partially triangulated $(r \times r)$-grid R, $P(R) \ge (\delta r)^2 + o((\delta r)^2)$, and
(3) δ is the smallest real number for which this inequality holds.

In either case, P is called *bidimensional*. The *density* δ of P is the minimum of the two possible densities (when both definitions are applicable). We call the sublinear function $f(x) = o(x)$ in the bound on $P(R)$ the *residual function* of P.

Notice that density assigns a positive real number, typically at most 1, to any bidimensional parameter. Interestingly, this assignment defines a total order on all such parameters.

4.2 Examples

Many parameters are bidimensional. Here we mention just a few. Examples of minor-bidimensional parameters are the following:

Vertex Cover. A *vertex cover* of a graph G is a set C of vertices such that every edge of G has at least one endpoint in C. The *vertex-cover problem* is to find a minimum-size vertex cover in a given graph G. The corresponding parameter, the size of a minimum vertex cover, is minor bidimensional with density $\delta = 1/\sqrt{2}$. (Roughly half the vertices must be in any vertex cover of the grid, and one color class in a vertex 2-coloring of the grid is a vertex cover.)

Feedback Vertex Set. A *feedback vertex set* of a graph G is a set U of vertices such that every cycle of G passes through at least one vertex of U. The size of a minimum feedback vertex size is a minor-bidimensional parameter with density $\delta \in [1/2, 1/\sqrt{2}]$. ($\delta \geq 1/2$ because there are $r^2/4 + o(r^2)$ vertex-disjoint squares in the $(r \times r)$-grid, each of which must be broken; $\delta \leq 1/\sqrt{2}$ because it suffices to remove one color class in a vertex 2-coloring of the grid.)

Minimum Maximal Matching. A *matching* in a graph G is a set E' of edges without common endpoints. A matching in G is *maximal* if it is contained by no other matching in G. The size of a minimum maximal matching is a minor-bidimensional parameter with density $\delta \in [1/\sqrt{8}, 1/\sqrt{2}]$. ($\delta \geq 1/\sqrt{8}$ because any maximal matching must include at least one edge interior to any 3×4 subgrid, and there are $r^2/8 + o(r^2)$ interior-disjoint 3×4 subgrids; $\delta \leq 1/\sqrt{2}$ because the number of edges in a matching is at most $r^2/2$.)

Examples of contraction-bidimensional parameters are

Dominating Set. A *dominating set* of a graph G is a set D of vertices of G such that each of the vertices of $V(G) - D$ is adjacent to at least one vertex of D. The size of a minimum dominating set is a contraction-bidimensional parameter with density $\delta = 1/3$. ($\delta \geq 1/3$ because every vertex dominates at most 9 vertices; $\delta \leq 1/3$ because there is a triangulation of the $(r \times r)$-grid with dominating set of size $r^2/9 + o(r^2)$.)

Edge Dominating Set. An *edge dominating set* of a graph G is a set D of edges of G such that every edge in $E(G) - D$ shares at least one endpoint with some edge in D. The size of a minimum edge domainting set is a contraction-bidimensional parameter with density $\delta = 1/\sqrt{14}$. ($\delta \geq 1/\sqrt{14}$ because every edge in a triangulated grid dominates at most 14 edges; $\delta \leq 1/\sqrt{14}$ because size-14 neighborhoods of a diagonal edge can be tiled to form a triangulated $(r \times r)$-grid requiring $r^2/14 + o(r^2)$ dominating edges.)

Many of our results can be applied not only to bidimensional parameters but also to parameters that are bounded by bidimensional parameters [24; 14]. For example, the *clique-transversal number* of a graph G is the minimum number of vertices intersecting every maximal clique of G. This parameter is not contraction-bidimensional because an edge contraction may create a new maximal clique and cause the clique-transversal number to increase. On the other hand, it is easy to see that this graph parameter always exceeds the size of a minimum dominating set. In particular, this fact can be used to obtain a parameter-treewidth bound for the clique-transversal number.

Our results can also be applied to maximization problems. For example, maximum independent set is a contraction-bidimensional parameter.

4.3 Subexponential Algorithms and Planar Graphs

Almost all known techniques for obtaining subexponential parameterized algorithms on planar graphs are based on the following "bounded-treewidth approach" [1; 30; 38]:

(I1) Prove that $\mathbf{tw}(G) \leq c\sqrt{P(G)}$ for some constant c;

(I2) Compute or approximate the treewidth (or branchwidth) of G;

(I3) Decide whether $P(G) \leq k$ as follows. If the treewidth is more than $c\sqrt{k}$, then the answer to the decision problem is NO. If treewidth is at most $c\sqrt{k}$, then run a standard dynamic program for graphs of bounded treewidth in $2^{O(\mathbf{tw}(G))}n^{O(1)} = 2^{O(\sqrt{k})}n^{O(1)}$ time.

All previously known ways of obtaining the most important step (I1) use rather complicated techniques based on separators. Next we give some hints why bidimensional parameters are important for the design of subexponential algorithms by showing how step (I1) can be performed for planar graphs. We need the following result of Robertson, Seymour, and Thomas.

THEOREM 4.3. [45, Theorem 4.3], [50, Theorem 6.3] *Let $r \geq 1$ be an integer. Every planar graph with no $(r \times r)$-grid as a minor has branchwidth at most $4r - 3$.*

Using this theorem we obtain the following relation between treewidth and bidimensional parameters:

THEOREM 4.4. *Let P be a bidimensional parameter. Then for any planar graph G, $\mathbf{tw}(G) = O(\sqrt{P(G)})$.*

PROOF. First, we consider the case when P is minor-bidimensional. Suppose, for contradiction, that $\mathbf{tw}(G) > c\sqrt{P(G)}$ for a large constant c to be determined. By Theorem 2.1, $\mathbf{bw}(G) > \frac{2}{3}\mathbf{tw}(G) > \frac{2}{3}c\sqrt{P(G)}$. By Theorem 4.3, G must have an $(r \times r)$-grid R as a minor, where $r \geq \frac{1}{6}c\sqrt{P(G)}$. Let δ be the density of P. Then $|V(R)| = r^2 \leq P(R)/\delta^2 - o(r^2) \leq P(G)/\delta^2 - o(r^2)$ because P is minor-bidimensional. But $r^2 \geq \frac{1}{36}c\,P(G)$, so we get a contradiction by choosing c large enough.

If P is contraction-bidimensional, we can use the same proof with one change. After obtaining the grid R as a minor, we remove the edge deletions and take only the edge contractions that form R from G, to obtain a partially triangulated grid

R' as a contraction of G. Then, the rest of the proof uses R' instead of R; in particular, $P(R') \le P(G)$. $\square$

The class of bidimensional parameterized problems contains all parameters known from the literature to have subexponential parameterized algorithms for planar graphs [2; 1; 4; 11; 40; 36]. Recently, Cai et al. [8] defined a class of parameters, Planar TMIN_1, and proved that, for every planar graph G and parameter P in Planar TMIN_1, $\mathbf{tw}(G) = O(\sqrt{P(G)})$. Every problem in Planar TMIN_1 can be expressed as a special type of dominating-set problem on bipartite graphs. (We refer to [8] for definitions and further properties of Planar TMIN_1.) Using Theorem 4.4 it is possible to prove a similar result, establishing the bound $\mathbf{tw}(G) = O(\sqrt{P(G)})$ for most parameters P in Planar TMIN_1.

It is tempting to wonder whether every parameter admitting a $2^{O(\sqrt{k})}n^{O(1)}$-time algorithm on planar graphs is bidimensional.

4.4 Parameter-Treewidth Bound for Bounded-Genus Graphs

To extend Theorem 4.4 to graphs of bounded genus, more work needs to be done.

If P is a bidimensional parameter with density δ and residual function f, then we define the *normalization factor* of P to be the minimum number $\beta \ge 1$ such that $(\frac{\delta}{\beta}r)^2 \le (\delta r)^2 + f(\delta r)$ for all $r \ge 1$.

LEMMA 4.5. *Let P be a contraction (minor) bidimensional parameter with density δ. Then $P(G) < (\frac{\delta}{\beta}r)^2$ implies that G excludes the $(r \times r)$-grid as a minor (and all partial triangulations of the $(r \times r)$-grid as contractions).*

PROOF. If P is minor bidimensional and H is the $(r \times r)$-grid and H is a minor of G, then $P(H) \le P(G)$. Because $P(H) = (\delta r)^2 + f(\delta r)$, we have that $(\frac{\delta}{\beta}r)^2 > P(G) \ge (\delta r)^2 + f(\delta r)$, which contradicts the definition of β.

If P is contraction bidimensional and H is a partial triangulation of the $(r\times r)$-grid and H is a contraction of G, then $P(H) \le P(G)$. Because $P(H) = (\delta r)^2 + f(\delta r)$, we have that $(\frac{\delta}{\beta}r)^2 > P(G) \ge (\delta r)^2 + f(\delta r)$, which contradicts the definition of β. $\square$

Let G be a graph and let $v \in V(G)$ be a vertex. Also suppose we have a partition $\mathcal{P}_v = (N_1, N_2)$ of the set of the neighbors of v. Define the *splitting* of G with respect to v and $\mathcal{P}_v$ to be the graph obtained from G by

(1) removing v and its incident edges;

(2) introducing two new vertices v^1 and v^2; and

(3) connecting v^i with the vertices in N_i, for $i = 1, 2$.

If H is the result of consecutive application of several such operations to some graph G, then we say that H is a *splitting* of G. If, in addition, the sequence of splittings never splits a vertex that was the result of a previous splitting, then we say that H is a *fair splitting* of G. The vertices v of G involved in the splittings that make up a fair splitting are called *affected vertices*.

A parameter P is *α-splittable* if, for every graph G and for each vertex $v \in V(G)$, the result G' of splitting G with respect to v satisfies $P(G') \le P(G) + \alpha$. Many natural graph problems are α-splittable for small constants α. Examples of 1-splittable problems are dominating set, vertex cover, edge dominating set, independent set, clique-transversal set, and feedback vertex set, among many others.

For the proof of our main result on properties of bidimensional parameters, we need two technical lemmas used in induction on the genus.

It is convenient to work with *Euler genus*. The Euler genus $\mathbf{eg}(\Sigma)$ of a nonorientable surface Σ is equal to the nonorientable genus $\tilde{g}(\Sigma)$ (or the crosscap number). The Euler genus $\mathbf{eg}(\Sigma)$ of an orientable surface Σ is $2g(\Sigma)$, where $g(\Sigma)$ is the orientable genus of Σ.

The following lemma is very useful in proofs by induction on the genus. The first part of the lemma follows from [41, Lemma 4.2.4] (corresponding to a nonseparating cycle) and the second part follows from [41, Proposition 4.2.1] (corresponding to a surface-separating cycle).

LEMMA 4.6. *Let G be a connected graph 2-cell embedded in a surface Σ not homeomorphic to a sphere, and let N be a noncontractible noose on G. Then there is a fair splitting G' of G affecting the set $S = \{v_1, \ldots, v_\rho\}$ of vertices of G met by N such that one of the following holds:*

(1) *G' can be 2-cell embedded in a surface with Euler genus strictly smaller than $\mathbf{eg}(\Sigma)$; or*

(2) *each connected component G_i of G' can be 2-cell embedded in a surface with Euler genus strictly smaller than $\mathbf{eg}(\Sigma)$ and is a contraction of some graph G_i^* obtained from G after at most ρ splittings.*

The following lemma is a consequence of the definition of branchwidth.

LEMMA 4.7. *Let G be a graph and let G' be the splitting of a vertex in G. Then $\mathbf{bw}(G) \leq \mathbf{bw}(G') + 1$.*

PROOF. Consider a branch decomposition (T, τ) of G' of width $\mathbf{bw}(G')$. The same (T, τ) is also a branch decomposition of G if we replace each edge of G' with the unique correponding edge in G. The order of each edge e of T increases by at most 1 because all vertices except the split vertex have the same incident edges so are counted the same and, at worst, the split vertex is counted in G whereas its two copies in G' might not be counted (because each copy is incident to edges corresponding to leaves in only one connected component of $T - e$). □

THEOREM 4.8. *Suppose that P is an α-splittable bidimensional parameter ($\alpha \geq 0$) with density $\delta > 0$ and normalization factor $\beta \geq 1$. Then, for any (connected) graph G 2-cell embedded in a surface Σ of Euler genus $\mathbf{eg}(\Sigma)$, $\mathbf{bw}(G) \leq 4\frac{\beta}{\delta}(\mathbf{eg}(\Sigma)+1)\sqrt{P(G)+1} + 8\alpha(\frac{\beta}{\delta}(\mathbf{eg}(\Sigma)+1))^2$.*

PROOF. We induct on the Euler genus of Σ.

In the base case that $\mathbf{eg}(\Sigma) = 0$, Lemma 4.5 implies that, if $P(G) < (\frac{\delta}{\beta}r)^2$, then G excludes the $(r \times r)$-grid as a minor. This implication is precisely Lemma 4.5 when P is minor bidimensional. If P is contraction bidimensional, then the implication follows because, if a connected planar graph G can be transformed to a graph H (e.g., the $(r \times r)$-grid) via a sequence of edge contractions and/or removals, then by applying only the contractions in this sequence, we obtain a partial triangulation

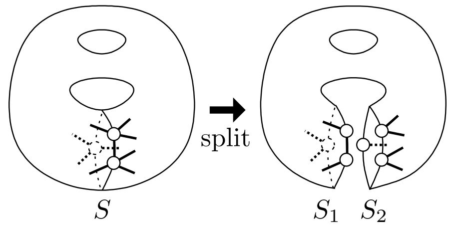

Fig. 2. Splitting a noose.

of H as a contraction of G. Now by Theorem 4.3, if $P(G) < (\frac{\delta}{\beta}r)^2$, then $\mathbf{bw}(G) \leq 4r-6$. If we set $r = \lfloor\frac{\beta}{\delta}\sqrt{P(G)}\rfloor + 1$, we have that $\mathbf{bw}(G) \leq 4\lfloor\frac{\beta}{\delta}\sqrt{P(G)}\rfloor - 2$. Because $\alpha, \beta, \delta \geq 0$, the induction base follows.

Suppose now that $\mathbf{eg}(\Sigma) \geq 1$ and that the induction hypothesis holds for any graph 2-cell embedded in a surface with Euler genus less than $\mathbf{eg}(\Sigma)$. Let G be a graph 2-cell embedded in Σ. We set $k = P(G)$ and claim that the representativity of this embedding of G is at most $4\lfloor\frac{\beta}{\delta}\sqrt{k+1}\rfloor$. Lemma 4.5 implies that, if $k < (\frac{\delta}{\beta}r)^2$, then G excludes any triangulation of the $(r \times r)$-grid as a contraction. By the contrapositive of Lemma 3.3, this implies that the representativity of G is less than $4r$. If we set $r = \lfloor\frac{\delta}{\beta}\sqrt{k+1}\rfloor + 1$, we have that the representativity of G is at most $4\lfloor\frac{\beta}{\delta}\sqrt{k+1}\rfloor$. Let N be a minimum-size noncontractible noose N on Σ meeting ρ vertices of G where $\rho \leq 4\lfloor\frac{\beta}{\delta}\sqrt{k+1}\rfloor$. By Lemma 4.6, there is a fair splitting along the vertices met by N such that either Condition 1 or Condition 2 holds; see Figure 2. Let G' be the resulting graph and let Σ' be a surface such that $\mathbf{eg}(\Sigma') \leq \mathbf{eg}(\Sigma) - 1$ and every connected component of G' is 2-cell embedable in Σ'. We claim that, given either Condition 1 or Condition 2, $\mathbf{bw}(G') \leq 4\frac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+\alpha\rho+1} + 8\alpha(\frac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2$.

Given Condition 1, we apply the induction hypothesis to G' and get that $\mathbf{bw}(G') \leq 4\frac{\beta}{\delta}(\mathbf{eg}(\Sigma')+1)\sqrt{P(G')+1} + 8\alpha(\frac{\beta}{\delta})^2(\mathbf{eg}(\Sigma')+1)^2$. Because G' is obtained from G after at most ρ splittings and P is an α-splittable parameter, we have $P(G') \leq k + \alpha\rho$. Because $\mathbf{eg}(\Sigma') \leq \mathbf{eg}(\Sigma) - 1$, we obtain $\mathbf{bw}(G') \leq 4\frac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+\alpha\rho+1} + 8\alpha(\frac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2$.

Given Condition 2, we apply the induction hypothesis to each of the connected components of G. Let G_i be such a component. We get that $\mathbf{bw}(G_i) \leq 4\frac{\beta}{\delta}(\mathbf{eg}(\Sigma')+1)\sqrt{P(G_i)+1} + 8\alpha(\frac{\beta}{\delta})^2(\mathbf{eg}(\Sigma')+1)^2$. Because G_i is a contraction of some graph G_i^* obtained from G after at most ρ splittings and P is an α-splittable parameter, we get that $P(G_i) \leq P(G_i^*) \leq k + \alpha\rho$. Again because $\mathbf{eg}(\Sigma') \leq \mathbf{eg}(\Sigma) - 1$, we have $\mathbf{bw}(G_i) \leq 4\frac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+\alpha\rho+1} + 8\alpha(\frac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2$. Because $\mathbf{bw}(G') = \max_i(\mathbf{bw}(G_i))$, we obtain $\mathbf{bw}(G') \leq 4\frac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+\alpha\rho+1} + 8\alpha(\frac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2$.

Because G' is the result of at most ρ consecutive vertex splittings on G, Lemma 4.7 yields that $\mathbf{bw}(G) \leq \mathbf{bw}(G') + \rho$. Therefore,

$$\begin{aligned}\mathbf{bw}(G) &\leq 4\tfrac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+\alpha\rho+1} + 8\alpha(\tfrac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2 + \rho \\ &\leq 4\tfrac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+\alpha(4\tfrac{\beta}{\delta}\sqrt{k+1})+1} + 8\alpha(\tfrac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2 + 4\tfrac{\beta}{\delta}\sqrt{k+1}\end{aligned}$$

$$
\begin{aligned}
&= 4\tfrac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{(\sqrt{k+1})(\sqrt{k+1}+4\alpha\tfrac{\beta}{\delta})+8\alpha(\tfrac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2}+4\tfrac{\beta}{\delta}\sqrt{k+1} \\
&\leq 4\tfrac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{(\sqrt{k+1}+4\alpha\tfrac{\beta}{\delta})(\sqrt{k+1}+4\alpha\tfrac{\beta}{\delta})+8\alpha(\tfrac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2}+4\tfrac{\beta}{\delta}\sqrt{k+1}, \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{because } \alpha,\beta,\delta \geq 0 \\
&= 4\tfrac{\beta}{\delta}\mathbf{eg}(\Sigma)(\sqrt{k+1}+4\alpha\tfrac{\beta}{\delta})+8\alpha(\tfrac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2+4\tfrac{\beta}{\delta}\sqrt{k+1} \\
&= 4\tfrac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+1}+16\alpha(\tfrac{\beta}{\delta})^2\mathbf{eg}(\Sigma)+8\alpha(\tfrac{\beta}{\delta})^2(\mathbf{eg}(\Sigma))^2+4\tfrac{\beta}{\delta}\sqrt{k+1} \\
&= 4\tfrac{\beta}{\delta}(\mathbf{eg}(\Sigma)+1)\sqrt{k+1}+8\alpha(\tfrac{\beta}{\delta})^2(\mathbf{eg}(\Sigma)^2+2\mathbf{eg}(\Sigma)) \\
&\leq 4\tfrac{\beta}{\delta}(\mathbf{eg}(\Sigma)+1)\sqrt{k+1}+8\alpha(\tfrac{\beta}{\delta})^2(\mathbf{eg}(\Sigma)^2+2\mathbf{eg}(\Sigma)+1), \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{because } \alpha,\beta,\delta \geq 0 \\
&= 4\tfrac{\beta}{\delta}(\mathbf{eg}(\Sigma)+1)\sqrt{k+1}+8\alpha(\tfrac{\beta}{\delta}(\mathbf{eg}(\Sigma)+1))^2.
\end{aligned}
$$

□

Theorem 4.8 is a general theorem that applies to any α-splittable bidimensional parameter. For minor-bidimensional parameters, the bound for branchwidth can be further improved.

THEOREM 4.9. *Suppose that P is a minor-bidimensional parameter with density $\delta \leq 1$ and normalization factor $\beta \geq 1$. Then, for any graph G 2-cell embedded in a surface Σ of Euler genus $\mathbf{eg}(\Sigma)$, $\mathbf{bw}(G) \leq 4\frac{\beta}{\delta}(\mathbf{eg}(\Sigma)+1)\sqrt{P(G)+1}$.*

PROOF. The proof is similar to the proof of Theorem 4.8. The only difference is that, instead of a fair splitting along the vertices of a minimum-size noncontractible noose, we just remove vertices of the noose from the graph. Because the parameter is minor bidimensional, the parameter cannot increase by this operation. The rest of the proof proceeds as before. Let G be a graph 2-cell embedded in a surface Σ of Euler genus $\mathbf{eg}(\Sigma)$, and let $k = P(G)$. We have the following substantially simpler inequality than the one in Theorem 4.8:

$$
\begin{aligned}
\mathbf{bw}(G) &\leq 4\tfrac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+1}+\rho \leq 4\tfrac{\beta}{\delta}\mathbf{eg}(\Sigma)\sqrt{k+1}+4\tfrac{\beta}{\delta}\sqrt{k+1} \\
&= 4\tfrac{\beta}{\delta}(\mathbf{eg}(\Sigma)+1)\sqrt{k+1}.
\end{aligned}
$$

□

4.5 Combinatorial Results and Further Improvements

As a consequence of Theorem 4.9, we establish an upper bound on the treewidth (or branchwidth) of a bounded-genus graph that excludes some planar graph H as a minor.

As part of their seminal Graph Minors series, Robertson and Seymour proved the following:

THEOREM 4.10. [44] *If G excludes a planar graph H as a minor, then the branchwidth of G is at most b_H and the treewidth of G is at most t_H, where b_H and t_H are constants depending only on H.*

The current best estimate of these constants is the exponential upper bound $t_H \leq 20^{2(2|V(H)|+4|E(H)|)^5}$ [50]. However, it is known that planar graphs can be excluded "quickly" from planar graphs. More precisely, the following result says that, for planar graphs, the constants depend only linearly on the size of H:

THEOREM 4.11. [50] *If G is planar and excludes a planar graph H as a minor, then the branchwidth of G is at most* $4(2|V(H)| + 4|E(H)|) - 3$.

This theorem follows from combining Theorem 4.3 with Theorem 1.5 of [50] that every planar graph H is a minor of an $(r \times r)$-grid where $r = 2|V(H)| + 4|E(H)|$.

Essentially the same proofs of Theorems 4.8 and 4.9 yield the following generalization of Theorem 4.3 for graphs of bounded genus. In fact, though, we can prove the following result directly from Theorem 4.9.

THEOREM 4.12. *If G is a graph of Euler genus $\mathbf{eg}(G)$ with branchwidth more than $4r(\mathbf{eg}(G) + 1)$, then G has the $(r \times r)$-grid as a minor.*

PROOF. Consider the parameter $\xi(G) = \max\{r^2 \mid G$ has an $(r \times r)$-grid as a minor$\}$. This parameter never increases when taking minors, and has value r^2 on the $(r \times r)$-grid, so is minor bidimensional with density 1 and normalization factor 1. If G excludes the $(r \times r)$-grid as a minor, then $\xi(G) < r^2$, so $\xi(G) \leq r^2 - 1$. By Theorem 4.9, we have that $\mathbf{bw}(G) \leq 4(\mathbf{eg}(G) + 1)\sqrt{\xi(G) + 1} \leq 4(\mathbf{eg}(G) + 1)r$, proving the contrapositive of the theorem. □

As above, by combining Theorem 4.12 with [50, Theorem 1.5], we obtain the following generalization of Theorem 4.11:

COROLLARY 4.13. *If G is a graph of Euler genus $\mathbf{eg}(G)$ that excludes a planar graph H as a minor, then its branchwidth is at most* $4(2|V(H)|+4|E(H)|)(\mathbf{eg}(G)+1)$.

4.6 Algorithmic Consequences

As we already discussed, the combinatorial upper bounds for branchwidth/treewidth are used for constructing subexponential parameterized algorithms as follows. Let G be a graph and P be a parameterized problem we need to solve on G. First one constructs a branch/tree decomposition of G that is optimal or "almost" optimal. A $(\theta, \gamma, \lambda)$-*approximation scheme* for branchwidth/treewidth consists of, for every w, an $O(2^{\gamma w} n^{\lambda})$-time algorithm that, given a graph G, either reports that G has branchwidth/treewidth at least w or produces a branch/tree decomposition of G with width at most θw. For example, the current best schemes are a $(3 + 2/3, 3.698, 3 + \epsilon)$-approximation scheme for treewidth [5] and a $(3, \lg 27, 2)$-approximation scheme for branchwidth [48].

If the branchwidth/treewidth of a graph is "large", then combinatorial upper bounds come into play and we conclude that P has no solution on G. Otherwise we run a dynamic program on graphs of bounded branchwidth/treewidth and compute $P(G)$.

Thus, we conclude with the main algorithmic result of this section:

THEOREM 4.14. *Let P be a bidimensional parameter with density δ and normalization factor β. Suppose P is either minor bidimensional, in which case we set $\mu = 0$, or P is contraction bidimensional and α-splittable, in which case we set $\mu = 2$. Suppose that there is an algorithm for the associated parameterized problem that runs in $O(2^{aw} n^b)$ time given a tree/branch decomposition of the graph G with width w. Suppose also that we have a $(\theta, \gamma, \lambda)$-approximation scheme for treewidth/branchwidth. Set $\tau = 1$ in the case of branchwidth and $\tau = 1.5$ in the* *case of treewidth. Then the parameterized problem asking whether $P(G) \leq k$ can be solved in $O(2^{\max\{a\theta,\gamma\}\tau 4\frac{\beta}{\delta}(g(G)+1)(\sqrt{k+1}+\mu\alpha\frac{\beta}{\delta}(g(G)+1))}n^{\max\{b,\lambda\}})$ time.*

The existence of an $O(2^{aw}n^b)$-time algorithm for treewidth/branchwidth w holds for many examples of bidimensional parameters with small values of a and b; see [1; 4; 11; 24; 30; 40; 31]. Observe that the correctness of our algorithms is simply based on Theorems 4.8 and 4.9, despite their nonalgorithmic natures, and $(\theta, \gamma, \lambda)$-approximation scheme for branch/tree decomposition. We note that the time bounds we provide do not contain any hidden constants, and the constants are reasonably low for a broad collection of problems covering all the problems for which $2^{O(\sqrt{k})}n^{O(1)}$-time algorithms already exist.

5. H-MINOR-FREE GRAPHS

In this section, we show how the results on graphs of bounded genus can be generalized on graphs with excluded minors.

5.1 Clique Sums

Suppose G_1 and G_2 are graphs with disjoint vertex sets and let $k \geq 0$ be an integer. For $i = 1, 2$, let $W_i \subseteq V(G_i)$ form a clique of size k and let G'_i be obtained from G_i by deleting some (possibly no) edges from $G_i[W_i]$ with both endpoints in W_i. Consider a bijection $h : W_1 \to W_2$. We define a *k-sum* G of G_1 and G_2, denoted by $G = G_1 \oplus_k G_2$ or simply by $G = G_1 \oplus G_2$, to be the graph obtained from the union of G'_1 and G'_2 by identifying w with $h(w)$ for all $w \in W_1$. See Figure 3. The images of the vertices of W_1 and W_2 in $G_1 \oplus_k G_2$ form the *join set*.

Note that each vertex v of G has a corresponding vertex in G_1 or G_2 or both. It is also worth mentioning that $\oplus$ is not a well-defined operator: it can have a set of possible results.

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

LEMMA 5.1. [21, Lemma 3] *For any two graphs G and H, $\mathbf{tw}(G \oplus H) \leq \max\{\mathbf{tw}(G), \mathbf{tw}(H)\}$.*

5.2 Characterizations of H-Minor-Free Graphs

Our result uses the deep theorem of Robertson and Seymour [49] on graphs excluding a non-planar graph as a minor. 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 nonplanarity" called *vortices*. Here, the bounds depend only on H.

Roughly speaking we say a graph G is *h-almost-embeddable* in a surface Σ 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 Σ by attaching at most h graphs of pathwidth at most h to G_0 within h faces in an orderly way. More precisely:

DEFINITION 5.2. A graph G is *h-almost-embeddable* in a surface Σ if there exists a vertex set X of size at most h called *apices* such that $G - X$ can be written as

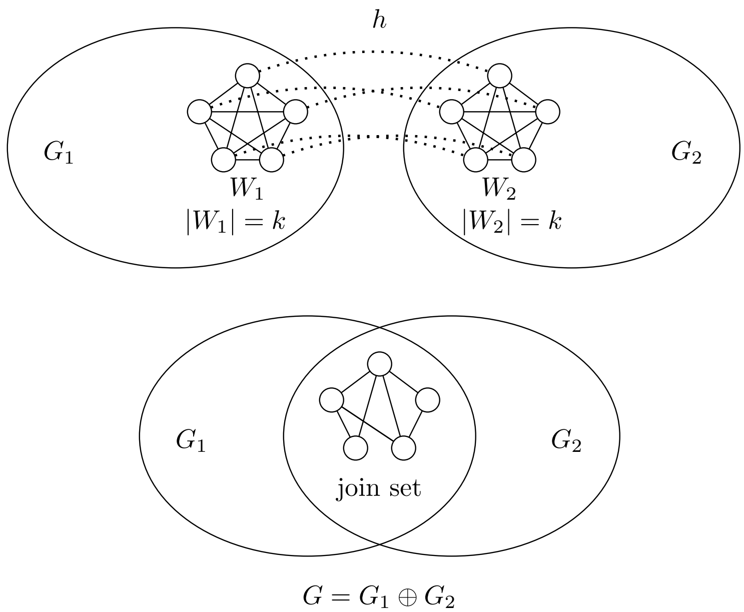

Fig. 3. A k-sum of two graphs G_1 and G_2.

$G_0 \cup G_1 \cup \cdots \cup G_h$, where

(1) G_0 has an embedding in Σ;
(2) the graphs G_i, called *vortices*, are pairwise disjoint;
(3) there are faces $F_1, \ldots, F_h$ of G_0 in Σ, and there are pairwise disjoint disks $D_1, \ldots, D_h$ in Σ, 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 called *apex free* if the set X of apices is empty.

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

THEOREM 5.3. [49] *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 graphs 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 graphs in bounded-genus surfaces. This structural theorem plays an important role in obtaining the rest of the results of this article.

Another way to view Theorem 5.3 is that every H-minor-free graph G has a tree decomposition (T, χ) such that, for each node $i \in V(T)$, the induced subgraph $G[\chi_i]$ augmented with additional edges to form a clique on the vertices that overlap with the parent's bag, and a clique on the vertices that overlap with each child's bag, is h-almost-embeddable in a bounded-genus surface. (This augmented graph is called the *torso* $[\chi_i]$ in, e.g., [33; 26].) The intersections between bag χ_i and its parent's bag, and with each child's bag, correspond to the join sets in the clique-sum decomposition. Our development primarily follows the original clique-sum viewpoint of Robertson and Seymour, but we will also occasionally view the sums as being organized into the tree T.

Theorem 5.3 is very general and appeared in print only recently. However, several nice applications (see, e.g., [6; 33; 25]) are already known.

In [20] the following algorithmic version of Theorem 5.3 is shown:

THEOREM 5.4. [20] *For any graph H, there is an algorithm with running time $n^{O(1)}$ that either computes a clique-sum decomposition as in Theorem 5.3 for any given H-minor-free graph G, or outputs that G is not H-minor-free. Here n is the number of vertices in G, and the exponent in the running time depends on H.*

In this article, we show that, given the tree decompositions computed by Theorem 5.4, we can obtain efficient algorithms for problems on H-minor-free graphs. Although our main development is in terms of dominating set, our approach can be viewed as a guideline for solving other problems on H-minor-free graphs. Some further results in this direction are described in Section 6.

5.3 Almost-Embeddable Graphs and r-Dominating Set

In order to treat each term separately in the clique-sum decomposition of an H-minor-free graph, we need to solve a more general problem than dominating set. This r-dominating set problem, which also arises in facility location, is also contraction-bidimensional. This property enables us to obtain a parameter-treewidth bound for this problem as well.

DEFINITION 5.5. Let G be a graph. A subset $D \subseteq V(G)$ of vertices *r-dominates* another subset $S \subseteq V(G)$ of vertices if each vertex in S is at distance at most r from a vertex in D. We say that D is an *r-dominating set* if it r-dominates $V(G)$.

We need the following result proved in [15].

LEMMA 5.6. [15] *Let $\rho, k, r \geq 1$ be integers and G be a planar graph having an r-dominating set of size k and containing a $(\rho \times \rho)$-grid as a minor. Then $k \geq (\frac{\rho-2r}{2r+1})^2$.*

In other words, Lemma 5.6 says that, for any fixed r, r-dominating set is a bidimensional parameter. It is also easy to see that it is 1-splittable. Thus, Theorem 4.8 yields the following lemma.

LEMMA 5.7. *For any constant r, if a graph G of genus g has an r-dominating set of size at most k, then the treewidth of G is $O(g\sqrt{k} + g^2)$.*

Now we extend this result to apex-free h-almost-embeddable graphs. Before expressing this result, we need the following slight modification of [33, Lemma 2].

LEMMA 5.8. *Let $G = G_0 \cup G_1 \cup \cdots \cup G_h$ be an apex-free h-almost-embeddable graph. For $1 \leq i \leq h$, let $(\mathcal{B}_u)_{u \in U_i}$ be the path decomposition of vortex G_i of width* *at most h. Suppose that, for each $1 \le i \le h$, the vertices $U_i = \{u_i^1, u_i^2, \ldots, u_i^{m_i}\}$ form a path in G_0. Then, $\mathbf{tw}(G) \le (h^2+1)(\mathbf{tw}(G_0)+1) - 1$.*

PROOF. Let $\mathcal{B}$ be a bag of a tree decomposition of G_0 of minimum width $\mathbf{tw}(G_0)$. For each index $1 \le i \le h$, and for each vertex $u \in \mathcal{B} \cap U_i$, we add to $\mathcal{B}$ the corresponding bag $\mathcal{B}_u$ of the path decomposition of G_i. The size of each $\mathcal{B}_u$ is at most h, and the original size of $\mathcal{B}$ is at most $\mathbf{tw}(G_0)+1$. Thus such additions increase the size of $\mathcal{B}$ by at most $h^2(\mathbf{tw}(G_0)+1)$. Performing these additions for all bags $\mathcal{B}$ of a tree decomposition increases the maximum bag size from $\mathbf{tw}(G_0)+1$ to $(h^2+1)(\mathbf{tw}(G_0)+1)$. It can be easily seen that the resulting set of bags $\mathcal{B}$ form a tree decomposition of G, because each U_i forms a path in G_0. □

LEMMA 5.9. *Let r be a constant and let $G = G_0 \cup G_1 \cup \cdots \cup G_h$ be an apex-free h-almost-embeddable graph on a surface Σ of genus g. Let k be the size of a set $D \subseteq V(G)$ that r-dominates $V(G_0)$. Then $\mathbf{tw}(G) = O(h^2(g\sqrt{k+h} + g^2))$. In particular, for fixed g and h, $\mathbf{tw}(G) = O(\sqrt{k})$.*

PROOF. For each $1 \le i \le h$, let $(\mathcal{B}_u)_{u \in U_i}$ be the path decomposition of vertex G_i, where $U_i = \{u_i^1, u_i^2, \ldots, u_i^{m_i}\}$. Let G_0' be the graph obtained from G_0 by adding new vertices $C = \{c_1, c_2, \cdots, c_h\}$ and edges (c_i, u_i^j) and (u_i^j, u_i^{j+1}) (where $j+1$ is treated modulo m_i) for all $1 \le i \le h$ and $1 \le j \le m_i$. Because G_0 is embeddable in Σ, G_0' is also embeddable in Σ. G_0' has an r-dominating set of size at most $k+h$, namely, $(D \cap V(G_0)) \cup C$. By Lemma 5.7, $\mathbf{tw}(G_0') = O(g\sqrt{k+h}+g^2)$. The subgraph $G_0'' = G_0' - C$ of G_0' satisfies the same treewidth bound: $\mathbf{tw}(G_0'') = O(g\sqrt{k+h}+g^2)$. Also, in G_0'', the vertices U_i, $1 \le i \le h$, form a path. By Lemma 5.8, the treewidth of $G'' = G_0'' \cup G_1 \cup \cdots \cup G_h$ is $O(h^2(g\sqrt{k+h}+g^2))$. Finally, because G is a subgraph of G'', $\mathbf{tw}(G) \le \mathbf{tw}(G'')$. □

5.4 H-Minor-Free Graphs and Dominating Set

Now that we have an understanding of r-dominating set in apex-free almost-embeddable graphs, we return to the original problem of dominating set in the more general setting of H-minor-free graphs. For this section we use the notation G^* for the entire H-minor-free graph so that the primary object of interest, an almost-embeddable piece of G^*, can be referred to as G. The main result of this section is the following algorithmic result.

THEOREM 5.10. *One can test whether an H-minor-free graph G^* has a dominating set of size at most k in time $2^{O(\sqrt{k})} n^{O(1)}$, where the constants in the exponents depend on H.*

The main intuition behind the proof of Theorem 5.10 is as follows. The algorithm consists of two levels of dynamic programming. The top-level dynamic program is over the clique-sum decomposition of G^*. Within each subproblem, we can focus on a single almost-embeddable graph G. If G is apex free, then we have a parameter-treewidth bound on G by Lemma 5.9. However, a single apex vertex in G can dominate many vertices; hence, in general, we cannot bound the treewidth of G. Therefore, the algorithm guesses which apex vertices are in the dominating set, and removes the vertices of G that become "irrelavant" to our problem. (Roughly speaking, a vertex is irrelevant if it is already dominated, and it cannot be used to dominate anyone else; however, the precise definition is more complicated because of clique-sums.) If we remove the apex vertices in this way, then we show how to obtain a parameter-treewidth bound for the remaining graph in Theorem 5.12. Once we have a parameter-treewidth bound, the bottom-level dynamic program solves (a generalized form of) the problem on this graph and thus G.

Before detailing the proof, we need more precise definitions.

DEFINITION 5.11. Consider a clique-sum decomposition of an H-minor-free graph G^* in accordance with Theorem 5.3, organized into a tree structure (T, χ) as described in Section 5.2. Let G be one term in the clique-sum decomposition of G^* that is h-almost embeddable on a surface of genus g, with apex set X. If we remove from T the node of T corresponding to term G, we obtain a forest T' of p subtrees; let $G_1, G_2, \ldots, G_p$ denote the clique-sums of the terms corresponding to the nodes in each connected component of T'. We say that G is *clique-summed* with each G_i, $1 \le i \le p$, with *join set* $W_i = V(G) \cap V(G_i)$. Because the clique-sums are at most h-sums, $|W_i| \le h$. A clique W_i is called *fully dominated* by a subset $S \subseteq V(G)$ of vertices in G if $V(G_i) - X \subseteq N_{G^*}(S)$; otherwise, clique W_i is called *partially dominated* by S. A vertex v of G is *fully dominated* by a set S if $N_{G^*[V(G)-X]}(v) \subseteq N_{G^*}(S)$.

Note that the only edges that can appear in G but not in G^* are the edges among vertices of W_i, $1 \le i \le p$.

THEOREM 5.12. *Let G be an h-almost embeddable on a surface of genus g in a clique-sum decomposition of a graph G^*. Suppose G is clique-summed with graphs $G_1, \ldots, G_p$ via join sets $W_1, \ldots, W_p$, where $|W_i| \le h$, $1 \le i \le p$. Suppose G^* has a dominating set D of size at most k. Then there is a subset $S \subseteq D$ of size at most h such that, if we form the graph $\hat{G}$ by removing all vertices fully dominated by S that are not included in any partially dominated clique W_i from G, then $\mathbf{tw}(\hat{G}) = O(h^2 g\sqrt{k+h} + g^2) = O(\sqrt{k})$.*

PROOF. Suppose X is the set of apices in G, so that $G - X$ is an apex-free h-almost embeddable graph. Let $S = X \cap D$. We claim that S is our desired set. The rest of the proof is as follows: we construct a set $\hat{D}$ of size at most k for $\hat{G} - X$ which 2-dominates every vertex v of $\hat{G} - X$ which is not included in any vortex. Then since $\hat{G} - X$ is apex-free h-almost-embeddable on a surface of genus g with a 2-dominating set of size at most k desired by Lemma 5.9, it has treewidth at most $O(h^2 g\sqrt{k+h} + g^2)$. Then we can add vertices of X to all bags and still have a tree decomposition of width $O(h^2 g\sqrt{k+h} + g^2)$, as desired. We construct $\hat{D}$ from D as follows. First, we set $\hat{D} = D \cap V(G)$. For each $1 \le i \le p$, if $D \cap (V(G_i) - W_i) \ne \emptyset$ and $W_i \not\subseteq X$, we add an arbitrary vertex $w \in W_i - X$ to $\hat{D}$. Here we say a vertex v of D is *mapped* to a vertex w of $\hat{D}$ if $v = w$ or if $v \in D \cap (V(G_i) - W_i)$ and vertex $w \in W_i - X$ is the one that we have added to $\hat{D}$. One can easily observe that since each new vertex in $\hat{D}$ is in fact accounted by a unique vertex in D, $|\hat{D}| \le k$. It only remains to show that D is a 2-dominating set for $\hat{G} - X$. If a vertex $v \in V(\hat{G}) - X$ is not fully dominated, then there exists a vertex $w \in N_G(v)$ which is not dominated by S and thus not dominated by X (since $S = D \cap X$). This means that v is 2-dominated by a vertex u of $\hat{G} - X$ which dominates w (we note that u can be originally a vertex u' in $(V(G_i) - W_i) \cap D$ which is mapped to u in $\hat{D}$). Also, we note that for each clique W_i in which there is a mapped vertex of D, this vertex dominates all vertices of $W_i - X$ in $\hat{G} - X$ and thus we keep the whole clique $W_i - X$ in G. It only remains to show that every vertex of a partially dominated clique W_i is 2-dominated by a vertex of $\hat{G} - X$. We consider two cases: if $W_i \cap S = \emptyset$, since $V(G_i) - W_i \neq \emptyset$, there must exists a (mapped) vertex of $\hat{D}$ in $W_i - X$ and we are done. Now assume $W_i \cap S \neq \emptyset$. If $W_i \subset X$ then $W_i \cap (V(\hat{G}) - X) = \emptyset$ and we are done (since there is no clique in $\hat{G} - X$ at all). Otherwise, there exists a vertex $W_i - X$. If $(V(G_i) - W_i) \subseteq N_{G^*}(S) \neq \emptyset$, then $V(G_i) \cap D \neq \emptyset$. Thus there exists a mapped vertex $w \in W_i - X$ and we have 1-dominated vertices of $W_i - X$. As mentioned before if $D \cap (W_i - X) \neq \emptyset$, vertices $W_i - X$ are 1-dominated and we are done. The only remaining case is the case in which there exists a vertex $w \in W_i - X$ which is dominated by a vertex $x \in V(G)$ and by assumption $w \notin N_{G^*}(S)$ (we note that in this case, there is no dominating vertex in $V(G_i) - W_i$ for any i for which $w \in W_i$.) This means that vertex x is not fully dominated and thus it remains in $\hat{G}$. In addition, vertex x 2-dominates all vertices of $W_i - X$, since W_i is a clique in G and thus all vertices of $W_i - X$ are 2-dominated. This completes the proof of the theorem. ☐

We are now ready to prove Theorem 5.10.

Proof of Theorem 5.10. First, we use the $n^{O(1)}$-time algorithm of Theorem 5.4 to obtain the clique-sum decomposition of graph G^*. As mentioned before, this clique-sum decomposition can be considered as a generalized tree decomposition of G^*.

More precisely, we consider the clique-sum decomposition as a rooted tree. We try to find a dominating set of size at most k in this graph using a two-level dynamic program. Suppose a graph G is an h-almost-embeddable graph on a surface of genus g in a clique-sum decomposition of a graph G^*. Assume G is clique-summed with graphs $G_0, G_1, \ldots, G_p$ via join sets $W_0, W_1, \ldots, W_p$, where $|W_i| \leq h$, $0 \leq i \leq p$. Also assume that G_0 is considered as the parent of G and $G_1, \ldots, G_p$ are considered as children of G.

Colorings.. The subproblems in our first-level dynamic program are defined by a coloring of the vertices in W_i. Each vertex will be assigned one of 3 colors, labelled 0, $\uparrow 1$, and $\downarrow 1$. The meaning of the coloring of a vertex v is as follows. Color 0 represents that vertex v belongs to the chosen dominating set. Colors $\downarrow 1$ and $\uparrow 1$ represent that the vertex v is not in the chosen dominating set. Such a vertex v must have a neighbor w in the dominating set (i.e., colored 0); we say that vertex w *resolves* vertex v. Color $\downarrow 1$ for vertex v represents that the dominating vertex w is in the subtree of the clique-sum decomposition rooted at the current graph G, whereas $\uparrow 1$ represents that the dominating vertex w is elsewhere in the clique-sum decomposition. Intuitively, the vertices colored $\downarrow 1$ have already been resolved, whereas the vertices colored $\uparrow 1$ still need to be assigned to a dominating vertex.

Locally Valid Colorings.. A coloring of the vertices of W_i is called *locally valid* with respect to sets $S_1, S_2 \subseteq V(G)$ if the following properties hold:

—for any two adjacent vertices v and w in W_i, if v is colored 0, w is colored $\downarrow 1$; and

—if $v \in S_1 \cap W_i$, then v is colored 0; and
—if $v \in S_2 \cap W_i$, then v is not colored 0.

Our colorings are similar to that of previous work (e.g., [1]), but we use them in a new dynamic-programming framework that acts over clique-sum decompositions instead of tree decompositions.

Dynamic Program Subproblems.. Our first-level dynamic program has one subproblem for each graph G in the clique-sum decomposition and for each coloring c of the vertices in W_0. Because each join set has at most h vertices, the number of subproblems is $O(n \cdot 3^h)$. We define $D(G, c)$ to be the size of the minimum "semi"-dominating set of the vertices in subtree rooted at G subject to the following restrictions:

(1) Vertices colored $\downarrow 1$ are adjacent to at least one vertex in the dominating set. (Vertices colored $\uparrow 1$ are dominated "for free".)
(2) Vertices colored 0 are precisely the vertices in the dominating set.
(3) Vertices in W_0 are colored according to c.

If we solve every such subproblem, then in particular, we solve the subproblems involving the root node of the clique-sum decomposition and in which every vertex is colored 0 or $\downarrow 1$. The final dominating set of size k is given by the best solution to these subproblems.

Induction Step.. Suppose for each coloring c of W_i, $1 \le i \le p$, we know $D(G_i, c)$. If the graph G is of size at most h, then we can try all colorings in $O(3^h \cdot h^2) = O(1)$ time (where the factor of h^2 is for checking validity). Thus, we focus on almost-embeddable graphs G. First, we guess a subset X of size at most h. Then for each subset S of X, we put the vertices of S in the dominating set and forbid vertices of $X - S$ from being in the dominating set. Now we remove from G all fully dominated vertices of $G - X$ that are not included in any partially dominated clique W_i. Call the resulting graph $\hat{G}$. By Theorem 5.12, $\mathbf{tw}(\hat{G}) = O(\sqrt{k})$, or else we can ignore this subset S. We can obtain such a tree decomposition of width $3 + 2/3$ times optimum (or determine that $\mathbf{tw}(\hat{G})$ is too large), in $2^{O(\sqrt{k})} n^{3+\epsilon}$ time by a result of Amir [5]. All vertices absent from this tree decomposition are fully dominated and thus, in any minimum dominating set that includes S, they will not appear except the following case. It is possible that up to $|X - S| = O(h)$ vertices, which are either fully dominated or belong to $V(G_i) - W_i$ where W_i is fully dominated, appear in the dominating set to dominate vertices of $X - S$. Call the set of such vertices S'. We can guess this set S' by choosing at most h vertices among the discarded vertices that have at least one neighbor in $X - S$, and then add S' to the dominating set. On the other hand, for any partially dominated clique W_i, we know that all of its vertices are present in the tree decomposition; because they form a clique, there is a bag α_i in any tree decomposition that contains all vertices of W_i. We find α_i in our tree decomposition and map W_i and G_i to this bag. We also assume W_0 is contained in all bags, because its size is at most h. Now, for each coloring c of W_0, we run the dynamic program of Alber et al. [1] on the tree decomposition, with the restriction that the colorings of the bags are locally valid with respect to $S_1 := S \cup S'$ and $S_2 := X - S$, and are consistent with the coloring

c of W_0. For each bag α_i to which we mapped G_i, we add to the cost of the bag the value $D(G_i, c')$ for the current coloring c' of W_i. Using this dynamic program, we can obtain $D(G, c)$ for each coloring c of W_0.

Running Time.. The running time for each coloring c of W_0 and each choice of S is $2^{O(\sqrt{k})}n$ according to [1]. We have 3^h choices for c, $O(n^{h+1})$ choices for X, $O(2^h)$ choices for S, and $O(n^{h+1})$ choices for S'. Thus the running time for this inductive step is $6^h n^{2h+2} 2^{O(\sqrt{k})}$. There are $O(n)$ graphs in the clique-sum decomposition of G. Therefore, the total running time of the algorithm is $O(6^h n^{2h+3} 2^{O(\sqrt{k})}) + n^{O(1)}$ (the latter term for creating the clique-sum decomposition), which is $2^{O(\sqrt{k})} n^{O(1)}$ as desired. □

6. CONCLUSIONS AND FURTHER WORK

We have shown how to obtain subexponential fixed-parameter algorithms for the broad class of bidimensional problems on bounded-genus graphs, and for dominating set on general H-minor-free graphs for any fixed H. Our approach can also be used to obtain subexponential algorithms for other problems on H-minor-free graphs. We now demonstrate some examples of such problems.

The first example is vertex cover, where we use the following reduction. For a graph G, let G' be the graph obtained from G by adding a path of length two between any pair of adjacent vertices. The following lemma is obvious.

LEMMA 6.1. *For any K_h-minor-free graph G, $h \geq 4$, and integer $k \geq 1$,*

—*G' is K_h-minor-free, and*

—*G has a vertex cover of size $\leq k$ if and only if G' has a dominating set of size $\leq k$.*

Combining Lemma 6.1 with Theorem 5.10, we conclude that parameterized vertex cover can be solved in subexponential time on graphs with an excluded minor.

Another example is the set-cover problem. Given a collection $C = \{C_1, C_2, \ldots, C_m\}$ of subsets of a finite set $S = \{s_1, s_2, \ldots, s_n\}$, a *set cover* is a subcollection $C' \subseteq C$ such that $\bigcup_{C_i \in C'} C_i = S$. The *minimum set cover* problem is to find a cover of minimum size. For an instance (C, S) of minimum set cover, its graph G_S is a bipartite graph with bipartition (C, S). Vertices s_i and C_j are adjacent in G_S if and only if $s_i \in C_j$. Theorem 5.10 can be used to prove that minimum set cover can be solved in subexponential time when G_S is H-minor free for some fixed graph H. Specifically, for a given graph G_S, we construct an auxiliary graph A_S by adding new vertices v, u, w and making v adjacent to $\{u, w, C_1, C_2, \ldots, C_m\}$. Then

—(C, S) has a set cover of size $\leq k$ if and only if A_S has a dominating set of size $\leq k + 1$, and

—if G_S is K_h-minor-free, then A_S is K_{h+1}-minor-free.

It is reasonable to believe that Theorem 5.10 generalizes to obtain a subexponential fixed-parameter algorithm for the (k, r)-*center problem* on H-minor-free graphs. The (k, r)-center problem is a generalization of the dominating-set problem in which the goal is to determine whether an input graph G has at most k vertices (called *centers*) such that every vertex of G is within distance at most r

from some center. Demaine et al. [15] consider this problem for planar graphs and map graphs, and present a generalization of dynamic programming mentioned in the proof of Theorem 5.10 to solve the (k, r)-center problem for graphs of bounded treewidth/branchwidth. This dynamic program and Theorem 5.12 can be generalized to establish the desired result for H-minor-free graphs. A consequence is that we can solve the dominating-set problem in constant powers of H-minor-free graphs, which is the most general class of graphs so far for which one can obtain the exponential speedup.

It is an open and tempting question whether our technique can be generalized to solve in subexponential time on H-minor-free graphs every problem that can be solved in subexponential time on bounded-genus graphs. Recent positive progress on this question has been made [19]. Based on our results, they obtain subexponential algorithms for any minor-bidimensional problem on H-minor-free graphs, and for any contraction-bidimensional problem on apex-minor-free graphs. (A graph is *apex-minor-free* if it excludes a fixed apex graph; an *apex graph* is a graph in which the removal of a vertex leaves a planar graph.) Note that these results, while general, cannot be applied directly to dominating set on H-minor-free graphs. In particular, it remains open to extend the algorithmic approaches of Section 5 for H-minor-free graphs to all bidimensional parameters.

We also suspect that there is a strong connection between bidimensional parameters and the existence of linear-size kernels for the corresponding parameterized problems in bounded-genus graphs. Such a linear kernel has recently been obtained for dominating set [31].

Another question asked in the conference version of this article is whether the upper bounds of Theorems 4.8 and 4.9 can be extended to larger graph classes. The first steps in this direction were obtained in [14] for minor-closed graph families. A graph family $\mathcal{F}$ has the *domination-treewidth property* if there is some function $f(d)$ such that, for every graph $G \in \mathcal{F}$ with dominating set of size $\leq k$, $\mathbf{tw}(G) \leq f(k)$. In [14] it is shown that a minor-closed graph family has the domination-treewidth property if and only if the family has bounded local treewidth. In [17] it is shown further that, for any minor-closed graph family $\mathcal{F}$ of bounded local treewidth, $\mathbf{tw}(G) = O(\sqrt{P(G)})$ for any $G \in \mathcal{F}$, where P is the dominating-set parameter. More recently the same result has been established for any bidimensional parameter P [19].

The theory of bidimensionality can also be applied to obtain fixed-parameter algorithms and polynomial-time approximation schemes for most bidimensional problems on planar graphs and more generally H-minor-free graphs. We refer the reader to [18; 19] for details.

Finally, we point out that all papers cited in this section were based on the results of this article.

Acknowledgments

The authors are indebted to Paul D. Seymour for many discussions that led to combinatorial results of this article and for providing a portal into the Graph Minor Theory. We also thank Naomi Nishimura and Prabhakar Ragde for encouragement, helpful discussions, and advice. Finally, we thank the three anonymous referees for

their helpful comments.

REFERENCES

1. 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,* 4, 461–493.
2. 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. In *Proceedings of the 26th International Symposium on Mathematical Foundations of Computer Science (MFCS 2001)*. Lecture Notes in Computer Science, vol. 2136. 111–122.
3. Alber, J., Fellows, M. R., and Niedermeier, R. 2004. Polynomial-time data reduction for dominating set. *Journal of the ACM 51,* 3, 363–384.
4. Alber, J., Fernau, H., and Niedermeier, R. 2004. Parameterized complexity: exponential speed-up for planar graph problems. *Journal of Algorithms 52,* 1, 26–56.
5. Amir, E. 2001. Efficient approximation for triangulation of minimum treewidth. In *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence (UAI-2001)*. Morgan Kaufmann Publishers, San Francisco, CA, 7–15.
6. Böhme, T., Maharry, J., and Mohar, B. 2002. $K_{a,k}$ minors in graphs of bounded treewidth. *Journal of Combinatorial Theory, Series B 86,* 1, 133–147.
7. Bondy, J. A. and Murty, U. S. R. 1976. *Graph Theory with Applications*. American Elsevier Publishing Co., Inc., New York.
8. Cai, L., Fellows, M., Juedes, D., and Rosamond, F. 2001. On efficient polynomial-time approximation schemes for problems on planar structures. Manuscript.
9. Cai, L. and Juedes, D. 2003. On the existence of subexponential parameterized algorithms. *Journal of Computer and System Sciences 67,* 4, 789–807.
10. Chandran, L. S. and Grandoni, F. 2005. Refined memorization for vertex cover. *Information Processing Letters 93,* 3, 123–131.
11. Chang, M.-S., Kloks, T., and Lee, C.-M. 2001. Maximum clique transversals. In *Proceedings of the 27th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2001)*. Lecture Notes in Computer Science, vol. 2204. 32–43.
12. Charikar, M. and Sahai, A. 2002. Dimension reduction in the l_1 norm. In *Proceedings of the 43th Annual Symposium on Foundations of Computer Science (FOCS'02)*. 551–560.
13. Chekuri, C., Gupta, A., Newman, I., Rabinovich, Y., and Alistair, S. 2003. Embedding k-outerplanar graphs into ℓ_1. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'03)*. 527–536.
14. Demaine, E. D., Fomin, F. V., Hajiaghayi, M., and Thilikos, D. M. 2004. Bidimensional parameters and local treewidth. *SIAM Journal on Discrete Mathematics 18,* 3 (December), 501–511.
15. 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,* 1, 33–47.
16. Demaine, E. D. and Hajiaghayi, M. 2004a. Diameter and treewidth in minor-closed graph families, revisited. *Algorithmica 40,* 3 (August), 211–215.
17. Demaine, E. D. and Hajiaghayi, M. 2004b. Equivalence of local treewidth and linear local treewidth and its algorithmic applications. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA'04)*. 833–842.
18. Demaine, E. D. and Hajiaghayi, M. 2005a. Bidimensionality: New connections between FPT algorithms and PTASs. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*. Vancouver, 590–601.
19. Demaine, E. D. and Hajiaghayi, M. 2005b. Graphs excluding a fixed minor have grids as large as treewidth, with combinatorial and algorithmic applications through bidimensionality. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*. Vancouver, 682–689.
20. Demaine, E. D., Hajiaghayi, M., and Kawarabayashi, K. 2005. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *Proceedings of the 46th Annual IEEE*

Symposium on Foundations of Computer Science. Pittsburgh, PA. To appear.

21. 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,* 2 (September), 166–195.
22. Demaine, E. D., Hajiaghayi, M., and Thilikos, D. M. The bidimensional theory of bounded-genus graphs. *SIAM Journal on Discrete Mathematics*. To appear.
23. Demaine, E. D., Hajiaghayi, M., and Thilikos, D. M. 2002. A 1.5-approximation for treewidth of graphs excluding a graph with one crossing. In *Proceedings of the 5th International Workshop on Approximation Algorithms for Combinatorial Optimization (Italy, APPROX 2002)*. Lecture Notes in Computer Science, vol. 2462. 67–80.
24. 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,* 4 (February), 245–267.
25. 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,* 1, 25–41.
26. Diestel, R. and Thomas, R. 1999. Excluding a countable clique. *Journal of Combinatorial Theory, Series B 76,* 1, 41–67.
27. Downey, R. G. and Fellows, M. R. 1999. *Parameterized Complexity*. Springer-Verlag, New York.
28. Ellis, J., Fan, H., and Fellows, M. 2004. The dominating set problem is fixed parameter tractable for graphs of bounded genus. *Journal of Algorithms 52,* 2, 152–168.
29. Fellows, M. R. 2001. Parameterized complexity: The main ideas and some research frontiers. In *Proceedings of the 12th Annual International Symposium on Algorithms and Computation (ISAAC 2001)*. Lecture Notes in Computer Science, vol. 2223. 291–307.
30. Fomin, F. V. and Thilikos, D. M. 2003. Dominating sets in planar graphs: Branch-width and exponential speed-up. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2003)*. 168–177.
31. Fomin, F. V. and Thilikos, D. M. 2004. Fast parameterized algorithms for graphs on surfaces: Linear kernel and exponential speed-up. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming (ICALP 2004)*. Turku, Finland, 581–592.
32. Garey, M. R. and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman and Co., San Francisco, Calif.
33. Grohe, M. 2003. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica 23,* 4, 613–632.
34. Grohe, M. and Flum, J. 2002. The parameterized complexity of counting problems. In *Proceedings of the 43rd IEEE Symposium on Foundations of Comupter Science (FOCS 2002)*. 538–547.
35. Gupta, A., Newman, I., Rabinovich, Y., and Alistair, S. 1999. Cuts, trees and ℓ_1-embeddings of graphs. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS'99)*. 399–409.
36. Gutin, G., Kloks, T., and Lee, C. M. 2001. Kernels in planar digraphs. In *Optimization Online*. Mathematical Programming Society, Philadelphia.
37. Haynes, T. W., Hedetniemi, S. T., and Slater, P. J. 1998. *Fundamentals of Domination in Graphs*. Marcel Dekker Inc.
38. Kanj, I. and Perković, L. 2002. Improved parameterized algorithms for planar dominating set. In *Proceedings of the 27th International Symposium on Mathematical Foundations of Computer Science*. Lecture Notes in Computer Science, vol. 2420. 399–410.
39. Klein, P. N., Plotkin, S. A., and Rao, S. 1993. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC'93)*. 682–690.
40. 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. In *Proceedings of the 28th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2002)*. Lecture Notes in Computer Science, vol. 2573. 282–295.

41. Mohar, B. and Thomassen, C. 2001. *Graphs on surfaces.* Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD.
42. Plotkin, S. A., Rao, S., and Smith, W. D. 1994. Shallow excluded minors and improved graph decompositions. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'94)*. 462–470.
43. Robertson, N. and Seymour, P. D. 1986a. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms 7,* 3, 309–322.
44. Robertson, N. and Seymour, P. D. 1986b. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B 41,* 1, 92–114.
45. Robertson, N. and Seymour, P. D. 1991. Graph minors. X. Obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B 52,* 2, 153–190.
46. Robertson, N. and Seymour, P. D. 1994. Graph minors. XI. Circuits on a surface. *Journal of Combinatorial Theory, Series B 60,* 1, 72–106.
47. Robertson, N. and Seymour, P. D. 1995a. Graph minors. XII. Distance on a surface. *Journal of Combinatorial Theory, Series B 64,* 2, 240–272.
48. Robertson, N. and Seymour, P. D. 1995b. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B 63,* 1, 65–110.
49. Robertson, N. and Seymour, P. D. 2003. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B 89,* 1, 43–76.
50. Robertson, N., Seymour, P. D., and Thomas, R. 1994. Quickly excluding a planar graph. *Journal of Combinatorial Theory, Series B 62,* 2, 323–348.

received october 2004; reviewed july 2005; accepted july 2005