The Bidimensional Theory of Bounded-Genus Graphs*

Erik D. Demaine[1], MohammadTaghi Hajiaghayi[1], and Dimitrios M. Thilikos[3]

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

[2] Departament de Llenguatges i Sistemes Informàtics, Universitat Politècnica de Catalunya, Campus Nord – Mòdul C5, c/Jordi Girona Salgado 1-3, E-08034, Barcelona, Spain, sedthilk@lsi.upc.es

Abstract. Bidimensionality is a powerful tool for developing subexponential fixed-parameter algorithms for combinatorial optimization problems on graph families that exclude a minor. This paper completes the theory of bidimensionality for graphs of bounded genus (which is a minor-excluding family). Specifically we show that, for any problem whose solution value does not increase under contractions and whose solution value is large on a grid graph augmented by a bounded number of handles, the treewidth of any bounded-genus graph is at most a constant factor larger than the square root of the problem's solution value on that graph. Such bidimensional problems include vertex cover, feedback vertex set, minimum maximal matching, dominating set, edge dominating set, r-dominating set, connected dominating set, planar set cover, and diameter. This result has many algorithmic and combinatorial consequences. On the algorithmic side, by showing that an augmented grid is the prototype bounded-genus graph, we generalize and simplify many existing algorithms for such problems in graph classes excluding a minor. On the combinatorial side, our result is a step toward a theory of graph contractions analogous to the seminal theory of graph minors by Robertson and Seymour.

1 Introduction

The recent theory of fixed-parameter algorithms and parameterized complexity [13] has attracted much attention in its less than 10 years of existence. In general the goal is to understand when NP-hard problems have algorithms that are exponential only in a parameter k of the problem instead of the problem size n. Fixed-parameter algorithms whose running time is polynomial for fixed parameter values—or more precisely $f(k) \cdot n^{O(1)}$ for some (superpolynomial) function $f(k)$—make these problems efficiently solvable whenever the parameter k is reasonably small.

* The last author was supported by the EU within the 6th Framework Programme under contract 001907 (DELIS) and by the Spanish CICYT project TIC-2002-04498-C05-03 (TRACER)

In the last three 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 typical running times of $2^{O(\sqrt{k})} n^{O(1)}$. For example, the first fixed-parameter algorithm for finding a dominating set of size k in planar graphs [2] has running time $O(8^k n)$; subsequently, a sequence of subexponential algorithms and improvements have been obtained, starting with running time $O(4^{6\sqrt{34k}} n)$ [1], then $O(2^{27\sqrt{k}} n)$ [16], and finally $O(2^{15.13\sqrt{k}} k + n^3 + k^4)$ [14]. Other subexponential algorithms for other domination and covering problems on planar graphs have also been obtained [1, 3, 5, 17, 15].

All subexponential fixed-parameter algorithms developed so far are based on showing a "treewidth-parameter bound": any graph whose optimal solution has value k has treewidth at most some function $f(k)$. In many cases, $f(k)$ is sublinear in k, often $O(\sqrt{k})$. Combined with algorithms that are singly exponential in treewidth and polynomial in problem size, such a bound immediately leads to subexponential fixed-parameter algorithms.

A series of papers [7, 8, 6] introduce the notion of *bidimensionality* as a general approach for obtaining treewidth-parameter bounds and therefore subexponential algorithms. This theory captures essentially all subexponential algorithms obtained so far. Roughly speaking, a parameterized problem is *bidimensional* if the parameter is large in a "grid-like graph" (linear in the number of vertices) and either closed under contractions (*contraction-bidimensional*) or closed under minors (*minor-bidimensional*). Examples of bidimensional problems include vertex cover, feedback vertex set, minimum maximal matching, dominating set, edge dominating set, r-dominating set, connected dominating set, planar set cover, and diameter. Diameter is a simple computational problem, but its bidimensionality has important consequences as it forms the basis of locally bounded treewidth for minor-closed graph families [9].

Treewidth-parameter bounds have been established for all minor-bidimensional problems in H-minor-free graphs for any fixed graph H [8, 6]. In this case, the notion of "grid-like graph" is precisely the regular $r \times r$ square grid. However, contraction-bidimensional problems (such as dominating set) have proved substantially harder. In particular, the largest class of graphs for which a treewidth-parameter bound can be obtained is apex-minor-free graphs instead of general H-minor-free graphs [6]. ("Apex-minor-free" means "H-minor-free" where H is a graph in which the removal of one vertex leaves a planar graph.) Such a treewidth-parameter bound has been obtained for all contraction-bidimensional problems in apex-minor-free-graphs [6]. In this case, the notion of "grid-like graph" is an $r \times r$ grid augmented to have, for each vertex, $O(1)$ edges from that vertex to nonboundary vertices. (Here $O(1)$ depends on H.) Unfortunately, this treewidth-parameter bound is large: $f(k) = (\sqrt{k})^{O(\sqrt{k})}$. For a subexponential algorithm, we essentially need $f(k) = o(k)$. For apex-minor-free graphs, such a bound is known only for the special cases of dominating set and vertex cover [10, 8].

The biggest graph classes for which we know a sublinear (indeed, $O(\sqrt{k})$) treewidth-parameter bound for many contraction-bidimensional problems are single-crossing-minor-free graphs and bounded-genus graphs. ("Single-crossing-minor-free" means "H-minor-free" where H can be drawn in the plane with one crossing.) For single-crossing-minor-free graphs [12, 11] (in particular, planar graphs [7]), all contraction-bidimensional problems have a bound of $f(k) = O(\sqrt{k})$. In this case, the notion of "grid-like graph" is an $r \times r$ grid partially triangulated by additional edges that preserve planarity. For bounded-genus graphs [8], a bound of $f(k) = O(\sqrt{k})$ has been shown, for the same notion of "grid-like graphs", but only for contraction-bidimensional problems with an additional property called *α-splittability*: upon splitting a vertex, the parameter should increase by at most $\alpha = O(1)$ (or decrease).

In this paper we complete the theory of bidimensionality for bounded-genus graphs by establishing a sublinear ($f(k) = O(\sqrt{k})$) treewidth-parameter bound for general contraction-bidimensional problems in bounded-genus graphs. Our notion of "grid-like graph" is somewhat broader: a partially triangulated $r \times r$ grid (as above) with up to g additional edges ("handles"), where g is the genus of the original graph. This form of contraction-bidimensionality is more general than α-splittability,[3] and thus we generalize the results for α-splittable contraction-bidimensional problems from [8]. It is easy to construct a parameter that is contraction-bidimensional but not α-splittable, although these parameters are not "natural". So far all "natural" contraction-bidimensional parameters we have encountered are α-splittable, though we expect other interesting problems to arise that violate α-splittability.

Our results show that a partially triangulated grid with g additional edges is the prototype graph of genus g, as observed by Lovász [18]. At a high level, this property means that, to solve an (algorithmic or combinatorial) problem on a general graph of genus g, the "hardest" instance on which we should focus is the prototype graph. This property generalizes the well-known result in graph theory that the grid is the prototype planar graph. This also completes our theory of constructing such prototypes for bidimensional problems.

Further algorithmic applications of our results follow from the graph-minor theory of Robertson and Seymour (e.g., [19]) and its extensions [8, 10]. In particular, [23, 8] shows how to reduce many problems on general H-minor-free graphs to subproblems on bounded-genus graphs. Essentially, the difference between bounded-genus graphs and H-minor-free graphs are "apices" and "vortices", which are usually not an algorithmic barrier. Applying our new theory for bounded-genus graphs, we generalize the algorithmic extensions of [8, 10]. Indeed, we simplify the approaches of both [8] and [10], where it was necessary to "split" bounded-genus graphs into essentially planar graphs because of a lack of general understanding of bounded-genus graphs. Specifically, we remove the necessity of Lemmas 7.4–7.8 in [10].

[3] This statement is the contrapositive of the following property: if the parameter is k for the partially triangulated grid with g additional edges, then by α-splitting the additional edges, the parameter is at most $k + \alpha g$ on the partially triangulated grid.

Last but not least are the combinatorial aspects of our results. In a series of 23 papers (so far), Robertson and Seymour (e.g., [19]) developed the seminal theory of graphs excluding a minor, which has had many algorithmic and combinatorial applications. Our completed understanding of contraction-bidimensional parameters can be viewed as a step toward generalizing the theory of graph minors to a theory of graph contractions. Specifically, we show that any graph of genus g can be contracted to its core of a partially triangulated grid with at most g additional edges; this result generalizes an analogous result from [23] when permitting arbitrary minor operations (contractions and edge deletions). Avoiding edge deletions in this sense is particularly important for algorithmic applications because many parameters are not closed under edge deletions, while many parameters are closed under contraction.

2 Preliminaries

All the graphs in this paper are undirected without loops or multiple edges. Given a graph G, we denote by $V(G)$ the set of its vertices and by $E(G)$ the set of its edges. For any vertex $v \in V(G)$ we denote by E_v the set of edges incident to v. Moreover, we use the notation $N_G(v)$ (or simply $N(v)$) for the set of neighbors of v in G (i.e., vertices adjacent to v).

Given an edge $e = \{x, y\}$ of a graph G, the graph obtained from G by contracting the edge e is the graph we get if we identify the vertices x and y 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 class $\mathcal{C}$ is a *contraction-closed* class if any contraction of any graph in $\mathcal{C}$ is also a member of $\mathcal{C}$. A contraction-closed graph class $\mathcal{C}$ is *H-contraction-free* if $H \notin \mathcal{C}$. Given any graph class $\mathcal{H}$, we say that a contraction-closed graph class $\mathcal{C}$ is *$\mathcal{H}$-contraction-free* if $\mathcal{C}$ is H-contraction-free for *any* $H \in \mathcal{H}$.

2.1 Treewidth and Branchwidth

A *branch decomposition* of a graph (or a hypergraph) 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 $T(V(T), E(T) - e)$ 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, $\mathbf{bw}(G)$, is the minimum width over all branch decompositions of G. (In case where $|E(G)| \leq 1$, we define the branch-width 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 – the width of this branch decomposition is considered to be 0). The *treewidth* $\mathbf{tw}(G)$ of a graph G is a notion related to branchwidth. We need only the following relation:

Lemma 1 ([22]). *For any connected graph G where $|E(G)| \geq 3$, $\mathbf{bw}(G) \leq \mathbf{tw}(G) + 1 \leq \frac{3}{2}\mathbf{bw}(G)$.*

The main combinatorial result of this paper determines, for any k and g, a family of graphs $\mathcal{H}_{k,g}$ such that any $\mathcal{H}$-contraction-free graph G with genus g will have branchwidth $O(gk)$. To describe such a family, we will need some definitions on graph embeddings.

2.2 Graph Embeddings

A *surface* Σ is a compact 2-manifold without boundary. We will always consider connected surfaces. We denote by the $\mathbb{S}_0$ the sphere $(x, y, z \mid x^2 + y^2 + z^2 = 1)$. A *line* in Σ is subset homeomorphic to $[0, 1]$. An O-arc is a subset of Σ homeomorphic to a circle. Let G be a graph 2-cell embedded in Σ. To simplify notations 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 the points corresponding to its vertices and edges. That way, a subgraph H of G can be seen as a graph H where $H \subseteq G$. A *face* of G is a connected component of $\Sigma - E(G) - V(G)$. (Every face is an open set.) We use the notation $V(G)$, $E(G)$, and $F(G)$ for the set of the vertices, edges, and faces of G. For $\Delta \subseteq \Sigma$, $\overline{\Delta}$ is the *closure* of Δ. The boundary of Δ is $\mathbf{bd}(\Delta) = \overline{\Delta} \cap \overline{\Sigma - \Delta}$ and the interior is $\mathbf{int}(\Delta) = \overline{\Delta} - \mathbf{bd}(\Delta)$.

A subset of Σ meeting the drawing only in vertices of G is called G-*normal*. If an O-arc is G-normal then we call it *noose*. The length of a noose is the number of its vertices.

Representativity [21] is the measure of how dense is the embedding of a graph in a surface. The *representativity* (or *face-width*) $\mathbf{rep}(G)$ of a graph G embedded in surface $\Sigma \neq \mathbb{S}_0$ is the smallest length of a noncontractible noose in Σ. In other words, $\mathbf{rep}(G)$ is the smallest number k such that Σ contains a noncontractible (non null-homotopic in Σ) closed curve that intersects G in k points.

It is more convenient to work with Euler genus. The *Euler genus* $\mathbf{eg}(\Sigma)$ of a surface Σ is equal to the non-orientable genus $\tilde{g}(\Sigma)$ (or the crosscap number) if Σ is a non-orientable surface. If Σ is an orientable surface, $\mathbf{eg}(\Sigma)$ is $2g(\Sigma)$, where $g(\Sigma)$ is the orientable genus of Σ. Given a graph G its Euler genus $\mathbf{eg}(G)$ is the minimum $\mathbf{eg}(\Sigma)$ where Σ is a surface where G can be embedded.

2.3 Splitting Graphs and Surfaces

In this section we describe precisely how to cut along a noncontractible noose in order to decrease the genus of the graph until we obtain a planar graph.

Let G be a graph and let $v \in V(G)$. 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 P_v to be the graph obtained from G by (i) removing v and its incident edges; (ii) introducing two new vertices v^1, v^2; and (iii) connecting v^i with the vertices in N_i, $i = 1, 2$. If H is the result of the consecutive application of the above operation on some graph G then we say that H is a *splitting* of G. If additionally in such a splitting process we do not split vertices that are results of previous splittings then we say that H is a *fair splitting* of G.

The following lemma defines how to find a fair splitting for a given noncontractible noose. It will serve as a link between Lemma 8 and Lemma 10 in the proof of the main result of this paper.

Lemma 2. *Let G be a connected graph 2-cell embedded in a nonplanar surface Σ, and let N be a noncontractible noose of Σ. Then there is a fair splitting G' of G affecting the set $S = (v_1, \ldots, v_\rho)$ of the vertices of G met by N, such that (i) G' has at most two connected components, (ii) each connected component of G' can be 2-cell embedded in a surface with Euler genus strictly smaller than the Euler genus of Σ, and (iii) there are two faces f_1 and f_2, each in the 2-cell embedding of a connected component of G' (and the connected components are different for the two faces if G' is disconnected), such that the boundary of f_i, for $i \in \{1, 2\}$, contains $S_i = (v_1^i, \ldots, v_\rho^i)$ where v_j^1 and v_j^2 are the vertices created after the splitting of the vertex v_j, for $j = 1, \ldots, \rho$.*

3 Incomplete Embeddings and Their Properties

In this section we give a series of definitions and results that support the proof of the main theorem of the next section. In particular, we will need special embeddings of graphs that are incomplete, i.e., only *some* of the edges and vertices of the graph are embedded in a surface. Moreover, we will extend the definition of a contraction so that it will also consider contractions of faces for the part of the graph that is embedded.

Let Σ be a surface (orientable or not). Given a graph G, a vertex set $V \subseteq V(G)$ and an edge set $E \subseteq E(G)$ such that $\cup_{v \in V} E_v \subseteq E$, we say that G is (V, E)-embeddable in Σ if the graph G^- obtained by G if we remove from it all the vertices in V and all the edges in E, i.e., the graph $G^- = (V(G) - V, E(G) - E)$ has a 2-cell embedding in Σ. We call the graph G^- *ground* of G and we call the edges and vertices of G^- *landed*. In contrary, we call the vertices in V and E *flying*. Notice that the flying edges are partitioned into three categories: those that have both endpoints in $V(G) - V$ (we call them *bridges*), those with one endpoint in $V(G) - V$ and one endpoint in V (we call them *pillars*), and those with both endpoints in V (we call them *clouds*). From now on, whenever we refer to a graph (V, E)-embeddable in Σ we will accompany it with the corresponding 2-cell embedding of G^- in Σ.

The set of *atoms* of G with respect to some (V, E)-embedding of G in Σ is the set $A(G) = V(G) \cup E(G) \cup F(G)$ where $F(G)$ is the set of faces of the 2-cell embedding of G^- in Σ. Notice that a flying atom can only be a vertex or an edge. In this paper, we will consider the faces as open sets whose borders are cyclic sequences of edges and vertices.

3.1 Contraction Mappings

A strengthening of a graph being a contraction of another graph is for there to be a "contraction mapping" which preserves some aspects of the embedding in a

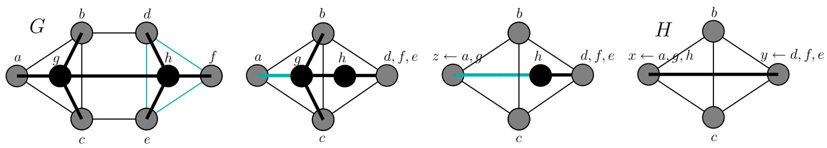

Fig. 1. An example of a contraction of a graph (V, E)-embeddable in $\mathbb{S}_0$. The contraction is shown in a three-step sequence: contracting the edges of the face $\{d, e, f\}$, then the edge $\{a, g\}$, and then edge $\{z, h\}$. A contraction mapping from G to H is defined as follows: $\phi(a)=\phi(g)=\phi(h)=\phi(\{a,g\})=\phi(\{g,h\})=x$, $\phi(b)=b$, $\phi(c)=c$, $\phi(d)=\phi(f)=\phi(e)=\phi(\{f,d\})=\phi(\{d,e\})=\phi(\{e,f\})=\phi(\{d,e,f\})=y$, $\phi(\{a,b\})=\phi(\{g,b\})=\{x,b\}$, $\phi(\{a,c\})=\phi(\{g,c\})=\{x,c\}$, $\phi(\{b,c\})=\{b,c\}$, $\phi(\{b,d\})=\{b,y\}$, $\phi(c,e)=\{c,y\}$, $\phi(\{a,b,c\})=\{x,b,c\}$, $\phi(\{b,d,e,c\})=\{b,c,y\}$, $\phi(\{h,d\})=\phi(\{h,e\})=\phi(\{h,f\})=\{x,y\}$, $\phi(\{a,b,d,f,e,c\})=\{x,b,y,c\}$.

surface during the contractions. See Fig. 1 for an example. Given two graphs G and H that are (V_G, E_G)- and (V_H, E_H)-embeddable in Σ and Σ', respectively, we say that $\phi : A(G) \to A(H)$ is a *contraction mapping* from G to H with respect to their corresponding embeddings if the following conditions are satisfied:

1. For any $v \in V(G)$, $\phi(v) \in V(H)$.
2. For any $e \in E(G)$, $\phi(e) \in E(H) \cup V(H)$.
3. For any $f \in F(G)$, $\phi(f) \in F(H) \cup E(H) \cup V(H)$.
4. For any $v \in V(H)$, $G[\phi^{-1}(v)]$ is a connected subgraph of G.
5. $\{\phi^{-1}(v) \mid v \in V(H)\}$ is a partition of $V(G)$.
6. If $\phi(\{x, y\}) = v \in V(H)$ then $\phi(x) = \phi(y) = v$.
7. If $\phi(\{x, y\}) = e \in E(H)$ then $\{\phi(x), \phi(y)\} \in E(H)$.
8. If $f \in F(G)$ and $\phi(f) = v \in V(G)$ and $f = (x_0, \ldots, x_{r-1})$ then $\phi(\{x_i, x_{i+1}\} = \phi(x_i) = v$ for any $i = 0, \ldots, r-1$ (where indices are taken modulo r).
9. If $f \in F(G)$ and if $\phi(f) = e$ (an edge of H) then there are two edges of f contained in $\phi^{-1}(e)$.
10. If $f \in F(G)$ and if $\phi(f) = g$ (a face of H) then each edge of g is *landed* and is the image of some edge in f.

Notice that, from Conditions 1, 2, and 3, the preimages of the faces of H are faces of G. The following lemma is easy.

Lemma 3. *If there exists some contraction mapping from a graph G to a graph H with respect to some embedding of G and H, then H is a contraction of G.*

3.2 Properties of Contraction Mappings

It is important that the two notions (contraction and existence of a contraction mapping) are identical in the case where G and H have no flying atoms, i.e., $V_G = V_H = E_G = E_H = \emptyset$. We choose to work with contraction mappings instead of simple contractions because they include stronger information enough to build the induction argument of Lemma 10.

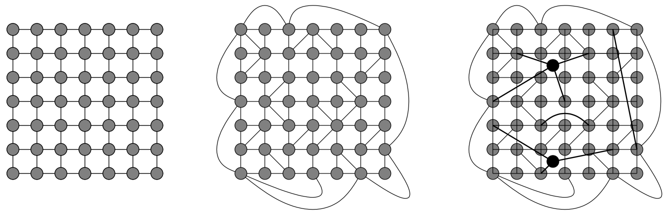

Fig. 2. A (7×7)-grid, a partially triangulated (7×7)-grid, and a $(7, 9)$-gridoid (the flying edges and vertices are the distinguished ones).

Lemma 4. *Let G be a graph (V, E)-embeddable on some surface Σ and let H be the graph occurring from G after contracting edges in $E(G^{-})$. Then $G[V] = H[V]$, H is also (V, E)-embeddable in Σ, and there exists a contraction mapping ϕ from G to H with respect to their corresponding embeddings.*

We omit the proof of this and other lemmas from this abstract.

3.3 Gridoids

A *partially triangulated* $(r \times r)$*-grid* is any graph that contains a $(r \times r)$-grid as a subgraph and is a subgraph of some triangulation of the same $(r \times r)$-grid.

We call a graph G a (r, k)*-gridoid* if it is (V, E)-embeddable in $\mathbb{S}_0$ for some pair V, E where $|E| \leq k$, $E(G[V]) = \emptyset$ (i.e., G does not have clouds), and so that G^{-} is a partial triangulated $(r' \times r')$-grid embedded on the $\mathbb{S}_0$ for some $r' \geq r$. For an example of a $(7, 9)$-gridoid and its construction; see Fig. 2.

4 Main Result

In this section we will prove that if a graph G has branchwidth more that $4k(\mathbf{eg}(G) + 1)$ then G will contain as a contraction some $(k - 12\mathbf{eg}(G), \mathbf{eg}(G))$-gridoid where $k \geq 12\mathbf{eg}(G)$.

4.1 Transformations of Gridoids

Lemma 5. *Let G be a (r, k)-gridoid $(\emptyset, E)$-embeddable in $\mathbb{S}_0$ and let $v \in V(G^{-})$. Then there exists some contraction mapping ϕ from G to some $(r - 4, k + 1)$-gridoid $(\{v\}, E \cup \{\{v, y\}\})$-embeddable in $\mathbb{S}_0$ such that $\phi(v) = v$.*

Fig. 3 illustrates the contractions in one case, where v has degree 4.

Lemma 6. *Let G be a (r, k)-gridoid $(\emptyset, E)$-embeddable in $\mathbb{S}_0$ and let e be some of its flying edges. Then there exists some $(r - 4, k)$-gridoid H $(\emptyset, E')$-embeddable in $\mathbb{S}_0$ for some E' and a contraction mapping ϕ of G to H such that $\phi(e) \in V(H)$.*

Lemma 7. *Let G be a (r, k)-gridoid $(\emptyset, E)$-embeddable in $\mathbb{S}_0$ and let a be some of its atoms. Then there exists some $(r - 4, k)$-gridoid $(\emptyset, E)$-embeddable in $\mathbb{S}_0$ and a contraction mapping ϕ from G to H with respect to their corresponding embeddings such that $\phi(a) \in V(H)$.*

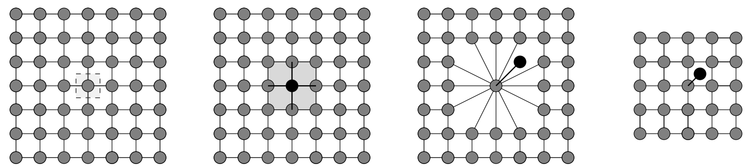

Fig. 3. An example of the first case in the proof of Lemma 5.

4.2 Excluding Gridoids as Contractions

Lemma 8. *Let G be a graph $(\emptyset,\emptyset)$-embeddable on some surface Σ. Let also H be a (r,k)-gridoid $(\emptyset,E)$-embeddable on the sphere and assume that ϕ is a contraction mapping from G to H with respect to their corresponding embeddings.*

Let now $\{v_1^i,\ldots,v_\rho^i\}, i=1,2$ be subsets of the vertices of two faces $f_i, i=1,2$ of the embedding of G where $f_1\cap f_2=\emptyset$ (we assume that the orderings of the indices in each subset respect the cyclic orderings of the vertices in $f_i, i=1,2$). Let G' be the graph obtained if we identify in G the vertex v_i^1 with the vertex v_i^2. Then, the following hold:

a. G' has some 2-cell embedding on a surface of bigger Euler genus.
b. There exist some $(r-12,k+1)$-gridoid H, $(\emptyset,E\cup\{\{e\}\})$-embeddable on the sphere such that there exists some contraction mapping from G' to H with respect their corresponding embeddings.

Proof. *a.* Let Σ be the surface where G is embedded. We define a surface Σ^- from Σ by removing the two "patches" defined by the (internal) points of the faces f_1 and f_2. Notice that G is still embeddable on Σ^- and that Σ^- is a surface with border whose connected components are the borders B_1,B_2 of the faces f_1 and f_2. We now construct a new surface from Σ^- by identifying the borders B_1 and B_2 in a way that v_i^1 is identified with v_i^2. Notice that the embedding that follows is a 2-cell embedding and that the new surface has bigger Euler genus.

b. From conditions 1, 2, and 3, $\phi(f_1)$ is either a vertex or an edge or a face of H. We apply Lemma 7 to construct a contraction mapping σ_1 from H to some $(r-4,k)$-gridoid H_1 where $\sigma_1(\phi(f_1))\in V(H_1)$. Notice again that $\sigma_1(\phi(f_2))$ is either a vertex or an edge or a face of H_1. We again use Lemma 7 to construct a contraction mapping σ_2 from H_1 to some $(r-8,k)$-gridoid H_2 where $\sigma_2(\sigma_1(\phi(f_i)))=v_i\in V(H_2), i=1,2$. We now apply Lemma 5 for v_1 and construct some contraction mapping σ_3 from H_2 to some $(r-12,k+1)$-gridoid H_3, $(\{v_1\},E\cup\{\{v_1,y\}\})$-embeddable in $\mathbb{S}_0$ such that $\sigma_3(v_1)=v_1$. Summing up we have that $\phi'=\phi\circ\sigma_1\circ\sigma_2\circ\sigma_3$ is a map from G to H_3 with respect to the $(\emptyset,\emptyset)$-embedding of G on Σ and the $(\{v_1\},E\cup\{\{v_1,y\}\})$-embeddable of H_3 in $\mathbb{S}_0$. Moreover, we have that $\phi'(f_1)=v_1$ and $\phi'(f_2)=v_2\in V(H_3)$ (to facilitate the notation we assume that $\sigma_3(v_2)=v_2$).

Notice now that if v is the result of the identification in H_3 of the vertex v_1 with the vertex v_2 we take a new graph H $(\emptyset,E\cup\{\{v,y\}\})$-embeddable in $\mathbb{S}_0$. Let A' be all the atoms of G that are not included in the faces f_1 and f_2. Notice that these atoms are not harmed while constructing G' from G and we set $\mu(a)=\phi'(a)$ for each $a\in A'$. Finally, for each atom $a\in A(G')-A$ we set $\mu(a)=v$. It now is easy to check that μ is a contraction mapping from G' to H with respect to their corresponding embeddings. As H is a $(r-12,k+1)$-gridoid we are done. $\square$

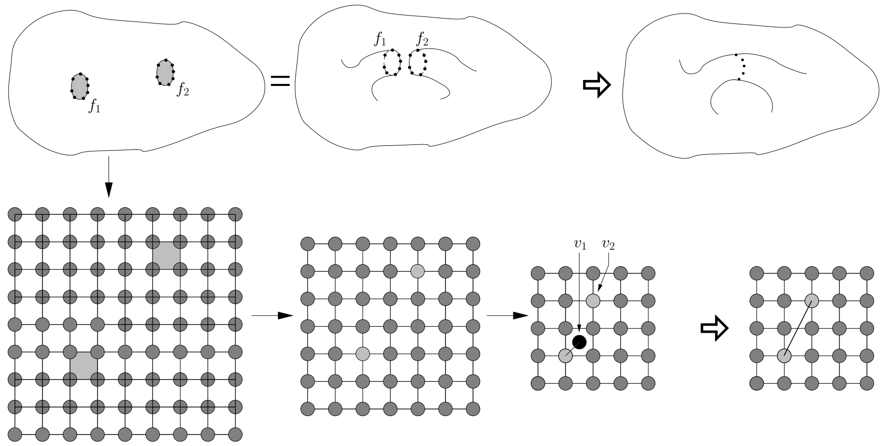

Fig. 4. An example of the transformations in the proof of Lemma 10.

The following is one of the main results in [8].

Theorem 1. *Let G be a graph 2-cell embedded in a non-planar surface Σ of representativity at least θ. Then one can contract edges in G to obtain partially triangulated $(\theta/4 \times \theta/4)$-grid.*

We also need the following easy lemma.

Lemma 9. *Let G be a graph and let H be the graph occurring from G after splitting some vertex $v \in V(G)$. Then $\mathbf{bw}(H) \leq \mathbf{bw}(G) + 1$.*

We are now ready to prove the central result of this section.

Lemma 10. *Let G be a graph $(\emptyset, \emptyset)$-embeddable on a surface Σ of Euler genus g and assume that $\mathbf{bw}(G) \geq 4(r-12g)(g+1)$. Then there exists some $(r-12g, g)$-gridoid H, $(\emptyset, E)$-embeddable in $\mathbb{S}_0$ such that there exists some contraction mapping from G to H with respect to their corresponding embeddings.*

Proof. First, if the graph G is disconnected, we discard all but one connected component C such that $\mathbf{bw}(C) = \mathbf{bw}(G)$.

We use induction on g. Clearly, if $g = 0$, G is a planar graph and after applying Lemma 3, the result follows from the planar exclusion theorem of RS. (The induction base use strongly the fact that for conventional embeddings the contraction relation is identical to our mapping.)

Suppose now that $g \geq 1$ and the theorem holds for any graph embeddable in a surface with Euler genus less than g. Refer to Fig. 4. If the representativity of G is at least $4(r - 12g)$, then by Theorem 1 we can contract edges in G to obtain a partially triangulated $((r - 12g) \times (r - 12g))$-grid (with no additional edges), and we are done. Otherwise, the representativity of G is less than $4(r - 12g)$. In this case, the smallest noncontractible noose has vertex set S of size less than $4(r-12g)$. Let G' be a splitting of G with respect to S as in Lemma 2. Recall that G' is now $(\emptyset, \emptyset)$-embeddable on a surface of Euler genus $g' \leq g - 1$.

By Lemma 9, the branchwidth of G' is at least the branchwidth of G minus $|S|$. As $|S| \leq 4(r-12g)$ we have that $\mathbf{bw}(G') \geq 4(r-12g)(g+1) - 4(r-12g) = 4(r-12g)g \geq 4(r-12g)(g'+1)$. By the induction hypothesis there exist some $(r-12g', g')$-gridoid H', $(\emptyset, E)$-embeddable in $\mathbb{S}_0$ such that there exists some contraction mapping from G' to H with respect to their corresponding embeddings. From Lemma 8, there exist some $(r-12g'-12, g'+1)$-gridoid H, $(\emptyset, E \cup \{\{e\}\})$-embeddable on the sphere such that there exists some contraction mapping from G to H with respect to their corresponding embeddings. As $r - 12g' - 12 \geq r - 12g$ and $g' + 1 \leq g$, we are done. □

And we have the conclusion of this section.

Theorem 2. *If a graph G excludes all $(k - 12\mathbf{eg}(G), \mathbf{eg}(G))$-gridoids as contractions, for some $k \geq 12\mathbf{eg}(G)$, then G has branchwidth at most $4k(\mathbf{eg}(G)+1)$.*

By Lemma 1 we can obtain a treewidth-parameter bound as desired.

5 Algorithmic Consequences

Define the *parameter* corresponding to an optimization problem to be the function mapping graphs to the solution value of the optimization problem. In particular, *deciding* a parameter corresponds to computing whether the solution value is at most a specified value k. A parameter is *contraction-bidimensional* if (1) its value does not increase under taking of contractions and (2) its value on a $(r, O(1))$-gridoid is $\Omega(r^2)$.[4]

Theorem 3. *Consider a contraction-bidimensional parameter P such that, given a tree decomposition of width at most w for a graph G, the parameter can be decided in $h(w) \cdot n^{O(1)}$ time. Then we can decide parameter P on a bounded-genus graph G in $h(O(\sqrt{k})) \cdot n^{O(1)} + 2^{O(\sqrt{k})} n^{3+\epsilon}$ time.*

Corollary 1. *Vertex cover, minimum maximal matching, dominating set, edge dominating set, r-dominating set (for fixed r), and clique-transversal set can be solved on bounded-genus graphs in $2^{O(\sqrt{k})} n^{3+\epsilon}$ time, where k is the size of the optimal solution. Feedback vertex set and connected dominating set can be solved on bounded-genus graphs in $2^{O(\sqrt{k}\log k)} n^{3+\epsilon}$ time.*

Acknowledgments. Thanks go to Fedor Fomin for early collaboration on this project, and to the anonymous referees for their helpful feedback on the paper.

References

1. J. Alber, H. L. Bodlaender, H. Fernau, T. Kloks, and R. Niedermeier. Fixed parameter algorithms for dominating set and related problems on planar graphs. *Algorithmica*, 33(4):461–493, 2002.

[4] The requirement of $\Omega(r^2)$ can be weakened to allow any function $g(r)$, as in [8, 6]; the only consequence is that $\sqrt{k}$ gets replaced by $g^{-1}(r)$.

2. J. Alber, H. Fan, M. R. Fellows, H. Fernau, R. Niedermeier, F. A. Rosamond, and U. Stege. Refined search tree technique for DOMINATING SET on planar graphs. In *MFCS 2001*, LNCS 2136, pp. 111–122.
3. J. Alber, H. Fernau, and R. Niedermeier. Parameterized complexity: Exponential speed-up for planar graph problems. *ECCC*, 8(23), 2001.
4. E. Amir. Efficient approximation for triangulation of minimum treewidth. In *UAI 2001*, pp. 7–15.
5. M.-S. Chang, T. Kloks, and C.-M. Lee. Maximum clique transversals. In *WG 2001*, pp. 300–310.
6. E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Bidimensional parameters and local treewidth. *SIAM J. Discr. Math.* To appear. A preliminary version appears in *LATIN 2004*, LNCS 2976, 2004, pp. 109–118.
7. E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Fixed-parameter algorithms for the (k, r)-center in planar graphs and map graphs. In *ICALP 2003*, pp. 829–844.
8. E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs. In *SODA 2004*, pp. 823–832.
9. E. D. Demaine and M. Hajiaghayi. Diameter and treewidth in minor-closed graph families, revisited. *Algorithmica*. To appear.
10. E. D. Demaine and M. Hajiaghayi. Equivalence of local treewidth and linear local treewidth and its algorithmic applications. In *SODA 2004*, pp. 833–842.
11. E. D. Demaine, M. Hajiaghayi, N. Nishimura, P. Ragde, and D. M. Thilikos. Approximation algorithms for classes of graphs excluding single-crossing graphs as minors. *J. Comput. System Sci.*. To appear.
12. E. D. Demaine, M. Hajiaghayi, and D. M. Thilikos. Exponential speedup of fixed parameter algorithms on $K_{3,3}$-minor-free or K_5-minor-free graphs. In *ISAAC 2002*, LNCS 2518, pp. 262–273.
13. R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer-Verlag, 1999.
14. F. V. Fomin and D. M. Thilikos. Dominating sets in planar graphs: Branch-width and exponential speed-up. In *SODA 2003*, pp. 168–177.
15. G. Gutin, T. Kloks, and C. Lee. Kernels in planar digraphs. In *Optimization Online*. Mathematical Programming Society, 2001.
16. I. Kanj and L. Perković. Improved parameterized algorithms for planar dominating set. In *MFCS 2002*, LNCS 2420, pp. 399–410.
17. T. Kloks, C. M. Lee, and J. Liu. New algorithms for k-face cover, k-feedback vertex set, and k-disjoint set on plane and planar graphs. In *WG 2002*, LNCS 2573, pp. 282–296.
18. L. Lovász. Personal communication, Dec. 2003.
19. N. Robertson and P. D. Seymour. Graph minors — a survey. In *Surveys in Combinatorics*, pp. 153–171. Cambridge Univ. Press, 1985.
20. N. Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.
21. N. Robertson and P. D. Seymour. Graph minors. VII. Disjoint paths on a surface. *J. Combin. Theory Ser. B*, 45(2):212–254, 1988.
22. N. Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, 52(2):153–190, 1991.
23. N. Robertson and P. D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Combin. Theory Ser. B*, 89(1):43–76, 2003.