

Bidimensionality and Kernels

Fedor V. Fomin^{*} Daniel Lokshantov[†] Saket Saurabh[‡]
Dimitrios M. Thilikos[§]

Abstract

Bidimensionality theory appears to be a powerful framework in the development of meta-algorithmic techniques. It was introduced by Demaine et al. [*J. ACM 2005*] as a tool to obtain sub-exponential time parameterized algorithms for bidimensional problems on H -minor free graphs. Demaine and Hajiaghayi [*SODA 2005*] extended the theory to obtain polynomial time approximation schemes (PTASs) for bidimensional problems. In this paper, we establish a third meta-algorithmic direction for bidimensionality theory by relating it to the existence of *linear kernels* for parameterized problems. In parameterized complexity, each problem instance comes with a parameter k and the parameterized problem is said to admit a *linear kernel* if there is a polynomial time algorithm, called a *kernelization* algorithm, that reduces the input instance to an equivalent instance (called *kernel*) with size linearly bounded by k . We show that “essentially” all bidimensional problems not only have sub-exponential time algorithms and PTASs but they *also* have linear kernels, affirmatively answering an open question from [*J. ACM 2005*] where the existence of linear kernels was conjectured for the first time. In particular, we prove that every minor (respectively contraction) bidimensional problem that satisfies the separation property and is of finite integer index, admits a linear kernel for classes of graphs that exclude a fixed graph (respectively an apex graph H) H as a minor. Recently, Bodlaender et al. [*FOCS 2009*] laid the foundation for obtaining meta-algorithmic results for kernelization and showed that various problems satisfying some logical and compactness properties have polynomial, even linear kernels on graphs of bounded genus. With the use of bidimensionality we are able to extend these results to minor-free and apex-minor-free graphs. Our results imply that a multitude of bidimensional problems, which include DOMINATING SET, FEEDBACK VERTEX SET, EDGE DOMINATING SET, VERTEX COVER, r -DOMINATING SET, CONNECTED DOMINATING SET, CYCLE PACK-

ING, CONNECTED VERTEX COVER, ALMOST CONSTANT TREEWIDTH, and various other vertex covering and packing problems, admit linear kernels on the corresponding graph classes. For most of these problems no polynomial kernels on H -minor-free graphs were known prior to our work.

1 Introduction

Bidimensionality was introduced as a framework for designing subexponential parameterized algorithms on sparse graphs. Alber et al. [2] obtained the first subexponential parameterized algorithm on planar graphs by solving the parameterized DOMINATING SET problem in time $2^{O(\sqrt{k})}n^{O(1)}$, where n is the input size. It appeared that not only DOMINATING SET but many other parameterized problems are solvable in subexponential time on planar graphs. The main purpose of Bidimensionality Theory was to provide a meta-algorithmic description of all these problems. The theory was developed by Demaine et al. [17, 18] and it allowed to unify and extend subexponential fixed-parameter algorithms for NP-hard graph problems to a broad range of graphs including planar graphs, map graphs, bounded-genus graphs and graphs excluding any fixed minor. This theory is built on cornerstone theorems from Graph Minors Theory of Robertson and Seymour and its initial purpose was to serve as a simple criterion of checking whether a parameterized problem is solvable in subexponential time on planar graphs, and even more generally, on graphs excluding some fixed graph as a minor. Roughly speaking, the problem is bidimensional if the solution value for the problem on a $k \times k$ -grid is $\Omega(k^2)$, and contraction/removal of edges does not increase solution value. Many natural problems are bidimensional, including DOMINATING SET, FEEDBACK VERTEX SET, EDGE DOMINATING SET, VERTEX COVER, r -DOMINATING SET, CONNECTED DOMINATING SET, CYCLE PACKING, CONNECTED VERTEX COVER, GRAPH METRIC TSP, and many others.

The second application of bidimensionality was given by Demaine and Hajiaghayi [18] who have shown that bidimensionality is a useful theory not only in the design of fast fixed-parameter algorithms

^{*}Department of Informatics, University of Bergen, Norway.

[†]Department of Informatics, University of Bergen, Norway.

[‡]The Institute of Mathematical Sciences, CIT Campus, Chennai, India.

[§]Department of Mathematics, National & Kapodistrian University of Athens, Greece. Supported by the project “Kapodistrias” (AII 02839/28.07.2008) of the National and Kapodistrian University of Athens (project code: 70/4/8757).

but also in the design of fast PTASs. They established a link between parameterized and approximation algorithms by proving that every bidimensional problem satisfying some simple separation properties has a PTAS on planar graphs and other classes of sparse graphs. We refer to the surveys [15, 20] for further information on bidimensionality and its applications.

In this work we give the third application of bidimensionality—*kernelization*. Kernelization can be seen as the strategy of analyzing preprocessing or data reduction heuristics from a parameterized complexity perspective. Parameterized complexity is basically a two-dimensional generalization of “P vs. NP” where in addition to the overall input size n , one studies the effects on computational complexity of a secondary measurement that captures additional relevant information. This additional information can be, for example, solution size or a structural restriction on the input distribution considered, such as a bound on the treewidth of an input graph. The secondary information is quantified by a positive integer k and is called the *parameter*. Parameterization can be deployed in many different ways; for general background on the theory see [21, 22, 30].

A parameterized problem with a parameter k is said to admit a *polynomial kernel* if there is a polynomial time algorithm (the degree of polynomial is independent of k), called a *kernelization* algorithm, that reduces the input instance down to an instance with size bounded by a polynomial $p(k)$ in k , while preserving the answer. Kernelization has been extensively studied in parameterized complexity, resulting in polynomial kernels for a variety of problems. Notable examples of known kernels are a $2k$ -sized vertex kernel for VERTEX COVER [13], a $355k$ kernel for DOMINATING SET on planar graphs [3], which later was improved to a $67k$ kernel [11], and an $\mathcal{O}(k^2)$ kernel for FEEDBACK VERTEX SET [32] parameterized by the solution size. One of the most intensively studied directions in kernelization is the study of problems on planar graphs and other classes of sparse graphs. This study was initiated by Alber et al. [3] who gave the first linear sized kernel for the DOMINATING SET problem on planar graphs. The work of Alber et al. [3] triggered an explosion of papers on kernelization, and kernels of linear sizes were obtained for a variety of parameterized problems on planar graphs including CONNECTED VERTEX COVER, MINIMUM EDGE DOMINATING SET, MAXIMUM TRIANGLE PACKING, EFFICIENT EDGE DOMINATING SET, INDUCED MATCHING, FULL-DEGREE SPANNING TREE, FEEDBACK VERTEX SET, CYCLE PACKING, and CONNECTED DOMINATING SET [3, 8, 9, 12, 25, 26, 27, 28, 29]. Not much was known about

kernelization on more general classes prior to this paper, except for DOMINATING SET. It was shown recently that DOMINATING SET enjoys a polynomial kernel on graphs excluding a fixed graph as a minor and on graphs of bounded degeneracy [4, 31]. We refer to the survey of Guo and Niedermeier [24] for a detailed treatment of the area of kernelization. Since most of the problems known to have polynomial kernels on planar graphs are bidimensional, the existence of links between bidimensionality and kernelization was conjectured and left as an open problem in [17].

In this work we show that every bidimensional problem with a simple separation property, which is a weaker property than the one required in the framework of Demaine and Hajiaghayi for PTASs [18] and with finite integer index (we postpone this definition till the next section) has a linear kernel on planar and even much more general classes of graphs. In this paper all the problems are parameterized by the *solution size*. Our main result is the following theorem.

THEOREM 1.1. *Every minor-bidimensional problem Π with the separation property and finite integer index has a linear kernel on graphs excluding some fixed graph as a minor. Every contraction-bidimensional problem Π with the separation property and with finite integer index has a linear kernel on graphs excluding some fixed apex graph as a minor.*

Our approach is built on recent results from Bodlaender et al. [5] who proved the first meta-theorems on kernelization. It was shown in [5] that every parameterized problem that has finite integer index and satisfies an additional surface-dependent property called quasi-compactness has a linear kernel on graphs of bounded genus.

Our kernelization algorithm uses only one reduction rule, which is based on finite integer index properties of the problem in question. This is exactly the same reduction rule as the one used in the kernelization algorithm given in [5] for obtaining kernels on planar graphs and graphs of bounded genus. The novel contribution of this paper is the way the kernel sizes for bidimensional problems are analyzed. The analysis of kernel sizes in [5] requires “topological” decompositions of the given graph, in the sense that the partitioning of the graph into regions with small border, or protrusions, strongly depends on an embedding of the graph into a surface. While such an approach works well when we have a topological embedding it seems difficult to extend it to graphs excluding some fixed graph as a minor. Instead of taking the topological approach, we apply bidimensionality and suitable variants of the Excluded Grid

Theorem [16, 19, 23]. This makes our arguments considerably simpler than the analysis in [5].

2 Definitions and Notations

In this section we give various definitions which we make use of in the paper. Let $G = (V, E)$ be a graph. A graph $G' = (V', E')$ is a *subgraph* of G if $V' \subseteq V$ and $E' \subseteq E$. The subgraph G' is called an *induced subgraph* of G if $E' = \{uv \in E \mid u, v \in V'\}$, in this case, G' is also called the subgraph *induced by* V' and denoted with $G[V']$. By $N(u)$ we denote (open) neighborhood of u that is the set of all vertices adjacent to u and by $N[u] = N(u) \cup \{u\}$. Similarly, for a subset $D \subseteq V$, we define $N[D] = \cup_{v \in D} N[v]$ and $N(D) = N[D] \setminus D$. We denote by K_h the complete graph on h vertices.

2.1 Parameterized algorithms, Kernels and Treewidth A parameterized problem Π is a subset of $\Gamma^* \times \mathbb{N}$ for some finite alphabet Γ . An instance of a parameterized problem consists of (x, k) , where k is called the parameter. We will assume that k is given in unary and hence $k \leq |x|^{\mathcal{O}(1)}$. A central notion in parameterized complexity is *fixed parameter tractability (FPT)* which means, for a given instance (x, k) , solvability in time $f(k) \cdot p(|x|)$, where f is an arbitrary function of k and p is a polynomial in the input size. The notion of *kernelization* is formally defined as follows.

DEFINITION 2.1. [Kernelization] A *kernelization algorithm*, or in short, a *kernel* for a parameterized problem $\Pi \subseteq \Gamma^* \times \mathbb{N}$ is an algorithm that given $(x, k) \in \Gamma^* \times \mathbb{N}$ outputs in time polynomial in $|x| + k$ a pair $(x', k') \in \Gamma^* \times \mathbb{N}$ such that (a) $(x, k) \in \Pi$ if and only if $(x', k') \in \Pi$ and (b) $|x'|, k' \leq g(k)$, where g is some computable function. The function g is referred to as the *size of the kernel*. If $g(k) = k^{\mathcal{O}(1)}$ or $g(k) = \mathcal{O}(k)$ then we say that Π admits a *polynomial kernel* and *linear kernel* respectively.

Treewidth. A *tree decomposition* of a graph $G = (V, E)$ is a pair (\mathcal{X}, T) where $T = (V_T, E_T)$ is a tree and $\mathcal{X} = \{X_i \mid i \in V_T\}$ is a collection of subsets of V such that:

1. $\bigcup_{i \in V_T} X_i = V$,
2. for each edge $xy \in E$, $\{x, y\} \subseteq X_i$ for some $i \in V_T$;
3. for each $x \in V$ the set $\{i \mid x \in X_i\}$ induces a connected subtree of T .

The *width* of the tree decomposition is $\max_{i \in V_T} |X_i| - 1$. The *treewidth* of a graph G is the minimum width over all tree decompositions of G . We denote by $\text{tw}(G)$ the treewidth of graph G .

2.2 Protrusions, t -Boundaried Graphs and Finite Integer Index Given a graph $G = (V, E)$ and

$S \subseteq V$, we define $\partial_G(S)$ as the set of vertices in S that have a neighbor in $V \setminus S$. For a set $S \subseteq V$ the neighborhood of S is $N_G(S) = \partial_G(V \setminus S)$. When it is clear from the context, we omit the subscripts. We now define the notion of a *protrusion*.

DEFINITION 2.2. [r -protrusion] Given a graph $G = (V, E)$, we say that a set $X' \subseteq V$ is an *r -protrusion* of G if $|N(X')| \leq r$ and $\text{tw}(G[X' \cup N(X')]) \leq r$.

For an r -protrusion X' , the vertex set $X = X' \cup N(X')$ is called an *extended r -protrusion*. The set X is the extended protrusion of X' and X' is the protrusion of X .

We now define the notion of *t -boundaried graphs* and various operations on them.

DEFINITION 2.3. [t -Boundaried Graphs] A *t -boundaried graph* is a graph $G = (V, E)$ with t distinguished vertices, uniquely labeled from 1 to t . The set $\partial(G)$ of labeled vertices is called the *boundary* of G . The vertices in $\partial(G)$ are referred to as *boundary vertices* or *terminals*.

For a graph $G = (V, E)$ and a vertex set $S \subseteq V$, we will sometimes consider the graph $G[S]$ as the $|\partial(S)|$ -boundaried graph with $\partial(S)$ being the boundary.

DEFINITION 2.4. [Gluing by \oplus] Let G_1 and G_2 be two t -boundaried graphs. We denote by $G_1 \oplus G_2$ the t -boundaried graph obtained by taking the disjoint union of G_1 and G_2 and identifying each vertex of $\partial(G_1)$ with the vertex of $\partial(G_2)$ with the same label; that is, we glue them together on the boundaries. In $G_1 \oplus G_2$ there is an edge between two labeled vertices if there is an edge between them in G_1 or in G_2 .

DEFINITION 2.5. [Legality] Let \mathcal{G} be a graph class, G_1 and G_2 be two t -boundaried graphs, and $G_1, G_2 \in \mathcal{G}$. We say that $G_1 \oplus G_2$ is *legal* with respect to \mathcal{G} if the unified graph $G_1 \oplus G_2 \in \mathcal{G}$. If the class \mathcal{G} is clear from the context we do not say with respect to which graph class the operation is legal.

DEFINITION 2.6. [Replacement] Let $G = (V, E)$ be a graph containing an extended r -protrusion X . Let X' be the restricted protrusion of X and let G_1 be an r -boundaried graph. The act of replacing X' with G_1 corresponds to changing G into $G[V \setminus X'] \oplus G_1$. Replacing $G[X]$ with G_1 corresponds to replacing X' with G_1 .

DEFINITION 2.7. For a parameterized problem Π on a graph class \mathcal{G} and two t -boundaried graphs G_1 and G_2 , we say that $G_1 \equiv_{\Pi} G_2$ if there exists

a constant c such that for all t -boundaried graphs G_3 and for all k : (a) $G_1 \oplus G_3$ is legal if and only if $G_2 \oplus G_3$ is legal; (b) $(G_1 \oplus G_3, k) \in \Pi$ if and only if $(G_2 \oplus G_3, k + c) \in \Pi$.

DEFINITION 2.8. [Finite Integer Index] We say that a parameterized problem Π has finite integer index in a graph class \mathcal{G} if for every t there exists a finite set \mathcal{S} of t -boundaried graphs such that $\mathcal{S} \subseteq \mathcal{G}$ and for any t -boundaried graph G_1 there exists $G_2 \in \mathcal{S}$ such that $G_2 \equiv_{\Pi} G_1$. Such a set \mathcal{S} is called a set of representatives for (Π, t) .

Note that for every t , the relation \equiv_{Π} on t -boundaried graphs is an equivalence relation. A problem Π is finite integer index, if and only if for every t , \equiv_{Π} is of finite index, that is, has a finite number of equivalence classes.

2.3 Minors and Contractions Given an edge $e = xy$ of a graph G , the graph G/e is obtained from G by contracting the edge e , that is, the endpoints x and y are replaced by a new vertex v_{xy} which is adjacent to the old neighbors of x and y (except from x and y). A graph H obtained by a sequence of edge-contractions is said to be a *contraction* of G . We denote it by $H \leq_c G$. A graph H is a *minor* of a graph G if H is the contraction of some subgraph of G and we denote it by $H \leq_m G$. We say that a graph G is *H -minor-free* when it does not contain H as a minor. We also say that a graph class \mathcal{G}_H is *H -minor-free* (or, excludes H as a minor) when all its members are H -minor-free. An *apex graph* is a graph obtained from a planar graph G by adding a vertex and making it adjacent to some of the vertices of G . A graph class \mathcal{G}_H is *apex-minor-free* if \mathcal{G}_H excludes a fixed apex graph H as a minor.

2.4 Grids and their triangulations. Let r be a positive integer, $r \geq 2$. The $(r \times r)$ -grid is the Cartesian product of two paths of lengths $r - 1$. A vertex of a grid is a *corner* if it has degree 2. Thus each $(r \times r)$ -grid has 4 corners. A vertex of a $(r \times r)$ -grid is called *internal* if it has degree 4, otherwise it is called *external*. Let Γ_r be the graph obtained from the $(r \times r)$ -grid by triangulating internal faces of the $(r \times r)$ -grid such that all internal vertices become of degree 6, all non-corner external vertices are of degree 4, and then one corner of degree two is joined by edges with all vertices of the external face. The graph Γ_6 is shown in Fig. 1.

2.5 Bidimensionality and Separation property

DEFINITION 2.9. ([17, 23]) A parameterized problem Π is *minor-bidimensional* if




Figure 1: Graph Γ_6 .

1. For any pair of graphs $H \leq_m G$ and integer k , $(G, k) \in \Pi \Rightarrow (H, k) \in \Pi$. In other words, contracting or deleting an edge in a graph G cannot increase the parameter; and
2. there is $\delta > 0$ such that for every $(r \times r)$ -grid R , $(R, k) \notin \Pi$ for every $k \leq \delta r^2$. In other words, the value of the solution on R should be at least δr^2 .

A parameterized problem Π is called *contraction-bidimensional* if

1. For any pair of graphs $H \leq_c G$ and integer k , $(G, k) \in \Pi \Rightarrow (H, k) \in \Pi$, and
2. there is $\delta > 0$ such that $(\Gamma_r, k) \notin \Pi$ for every $k \leq \delta r^2$.

In either case, Π is called *bidimensional*.

DEFINITION 2.10. For a parameterized problem Π , let $\bar{\Pi}$ denote the set of all no instances of Π . A minor (contraction) bidimensional problem Π is called a *minimization problem* if for all $(G, k) \in \Pi$, $\text{tw}(G) \leq \mathcal{O}(\sqrt{k})$ whenever G excludes a fixed (apex) graph H as a minor. A minor (contraction) bidimensional problem Π is called a *maximization problem* if for all $(G, k) \in \bar{\Pi}$, $\text{tw}(G) \leq \mathcal{O}(\sqrt{k})$ whenever G excludes a fixed (apex) graph H as a minor.

Demaine and Hajiaghayi [18] define the *separation* property for problems, and show how separability together with bidimensionality is useful to obtain PTASes on H -minor-free graphs. In our setting a slightly weaker notion of separability is sufficient. In particular the following definition is just the requirement 3 of the definition of separability in [18].

DEFINITION 2.11. A *minor-bidimensional problem* has the *separation property* if given any graph G , given any vertex cut S , and given an optimal solution OPT to G , for any union G' of some subset of connected components of $G \setminus S$, $|OPT \cap G'|$ is between $|OPT(G')| - \mathcal{O}(|S|)$ and $|OPT(G')| + \mathcal{O}(|S|)$.

For contraction-bidimensional parameters we have a slightly different definition of the separation property.

DEFINITION 2.12. A contraction-bidimensional problem has the separation property if the following condition is satisfied. Given a graph G , a vertex cut S whose removal disconnects G into connected components C_1, C_2, \dots, C_k , and a subset $I \subseteq \{1, 2, \dots, k\}$, we define G_I to be the graph obtained from G by contracting for every $j \notin I$ the component C_j into the vertex in $N(C_j)$ with the lowest index. Let OPT be an optimal solution to G . Then for any subset $I \subseteq \{1, 2, \dots, k\}$, $|OPT \cap G_I|$ is between $|OPT(G_I)| - \mathcal{O}(|S|)$ and $|OPT(G_I)| + \mathcal{O}(|S|)$.

3 Combinatorial Properties of Separable Bidimensional Problems

In this section we first show that any non-trivial instance of a separable bidimensional problem must have a $O(k)$ sized vertex set whose deletion makes the treewidth of the input graph constant. Second, we show how to exploit such a set to find a large protrusion in the input graph. We need the following well known lemma, see e.g. [6], on separators in graphs of bounded treewidth.

LEMMA 3.1. Let $G = (V, E)$ be a graph of treewidth at most t and $w : V \rightarrow \mathbb{R}^+ \cup \{0\}$ be a weight function. Then there is a set $S \subset V$ of size at most $t + 1$ such that for every connected component $G[C]$ of $G[V \setminus S]$, $w(C) \leq w(V)/2$. Furthermore the connected components C_1, \dots, C_ℓ of $G[V \setminus S]$ can be grouped into two sets \mathcal{C}_1 and \mathcal{C}_2 such that $\frac{w(V) - w(S)}{3} \leq w(\mathcal{C}_i) \leq \frac{2(w(V) - w(S))}{3}$, for $i \in \{1, 2\}$.

We also need the following result proved in [16, 19, 23].

PROPOSITION 3.1. ([16, 19, 23]) Let G be a connected graph excluding a fixed graph H as a minor. Then there exists some constant c such that if $\text{tw}(G) \geq c \cdot r^2$, then G contains the $r \times r$ -grid as a minor. Moreover, if H is an apex graph, then G contains Γ_r as a contraction.

LEMMA 3.2. Let Π be a minor-bidimensional problem with the separation property. For every $(G = (V, E), k) \in \Pi$, where G does not contain a fixed graph H as a minor, there is a subset of vertices $S \subseteq V$ and a constant t , such that $|S| = \mathcal{O}(k)$, and $\text{tw}(G[V \setminus S]) \leq t$.

Proof. The definition of minor-bidimensionality together with Proposition 3.1 imply that for every

$(G, k) \in \Pi$ there is a constant d' such that $\text{tw}(G) \leq d'\sqrt{k}$. Let us fix a solution Z of size k for G and a weight function $w : V \rightarrow \{0, 1\}$ which assigns 1 to every vertex in Z and 0 otherwise. By Lemma 3.1, there exists a separator X of G of size at most $d'\sqrt{k} + 1$ such that the connected components of $G[V \setminus X]$ can be grouped into two parts \mathcal{C}_1 and \mathcal{C}_2 such that $\frac{k - w(X)}{3} \leq w(\mathcal{C}_i) \leq \frac{2(k - w(X))}{3}$, $i \in \{1, 2\}$.

We want to construct a set S such that the treewidth of $G[V \setminus S]$ is at most t . We start the construction by putting $S := X$. Let Z_1 and Z_2 be an optimum solution to Π on graphs $G_1 = G[\mathcal{C}_1]$ and $G_2 = G[\mathcal{C}_2]$ respectively. Since Π is separable, we have that $|Z \cap \mathcal{C}_i| = |Z_i| \pm \mathcal{O}(|X|)$ for $i \in \{1, 2\}$. We grow S by recursively applying Lemma 3.1 to find balanced separators in G_1 and G_2 respectively and adding them to S . Since Π is minor-bidimensional problem with the separation property, we have that in recursive step for a graph G' with solution of size ℓ we find a separator of size $\mathcal{O}(\sqrt{\ell})$. Let $\mu(G, Z, k)$ denote the size of the set S , we are looking for. Then we get the following recurrence for $\mu(G, k)$,

$$\begin{aligned} \mu(G, Z, k) &\leq \mu\left(G_1, Z_1, \frac{k}{3} + (d\sqrt{k})\right) \\ &\quad + \mu\left(G_2, Z_2, \frac{2k}{3} + (d\sqrt{k})\right) \\ &\quad + (d'\sqrt{k} + 1) \end{aligned}$$

where d and d' are constants. We stop recursing when the solution in the current graph is at most some fixed constant t' . Again, by making use of Proposition 3.1 and minor-bidimensionality of Π , we conclude that the treewidth of the graphs at the leaves of the recursion tree is at most some fixed constant t . Hence we have found a set S such that the treewidth of $G[V \setminus S]$ is at most t . By applying the generalization of Master's Theorem due to Akra and Bazzi [1] to the recurrence above we have that the size of S is $\mathcal{O}(k)$.

The proof of the next lemma for contraction-bidimensional problems Π satisfying the separation property goes along the proof of Lemma 3.2, with the only difference that instead of using the first part of Proposition 3.1, we use the part on contractions.

LEMMA 3.3. Let Π be a contraction-bidimensional problem with the separation property. For every $(G = (V, E), k) \in \Pi$, where G is an apex-minor-free graph, there is a subset of vertices $S \subseteq V$ and constant t , such that $|S| = \mathcal{O}(k)$ and $\text{tw}(G[V \setminus S]) \leq t$.

Proof. The proof of this lemma is almost identical to the proof of Lemma 3.2. The only difference is in the definition of the graphs G_1 and G_2 .

Let C_1, C_2, \dots, C_ℓ be the connected components of $G[V \setminus S]$ and let $I_i \subseteq \{1, 2, \dots, \ell\}$ such that $C_i = \cup_{j \in I_i} C_j$. We set $G_1 := G_{I_1}$ and $G_2 := G_{I_2}$ as in Definition 2.12.

We proceed to show that any H -minor-free n -vertex graph G with a vertex set S of size k such that $\text{tw}(G \setminus S)$ is constant must have a protrusion of size at least cn/k for some fixed c . This result is crucial for our analysis of the kernel sizes.

LEMMA 3.4. *For every fixed graph H and constant t there are constants ζ and r that satisfy the following. For any n -vertex graph G which excludes H as a minor and has a vertex set S of size k such that $\text{tw}(G \setminus S) \leq t$, G has an extended r -protrusion of size at least $\zeta n/k$.*

Proof. For a fixed graph H and constants t , p and q define $\mu(k, x)$ to be the maximum number of vertices in a graph G excluding H as a minor, such that G has a vertex set S of size k such that $\text{tw}(G \setminus S) \leq t$ and there is no set Z of size at least x such that $|\partial(Z)| \leq p$ and $|Z \cap S| \leq q$. We argue that for p and q chosen sufficiently large compared to t and H (but chosen independently of k and x), $\mu(k, x) = O(kx)$.

First, consider a graph G excluding H as a minor, such that G has a vertex set S of size k such that $\text{tw}(G \setminus S) \leq t$. We prove that $\text{tw}(G) = O(\sqrt{k})$. By Proposition 3.1, there is a constant d' such that G contains a $(d'\text{tw}(G) \times d'\text{tw}(G))$ -grid as a minor. Thus G contains $(d'\text{tw}(G)/(t+1))^2$ disjoint $(t+1 \times t+1)$ -grids as minors. Since the treewidth of a $(t+1 \times t+1)$ -grid is $t+1$, S must contain at least one vertex from each of the $(t+1 \times t+1)$ -grid minors. Thus $|S| = k \geq (d'\text{tw}(G)/(t+1))^2$ and there is a constant d such that $\text{tw}(G) \leq d\sqrt{k}$.

Now, fix a weight function $w : V \rightarrow \mathbb{R}^+$ such that $w(v) = 1$ if $v \in S$ and $w(v) = 0$ otherwise. By Lemma 3.1, there is a partitioning of V into a separator X of size at most $d\sqrt{k} + 1$ and two sets \mathcal{C}_1 and \mathcal{C}_2 such that $N(\mathcal{C}_i) \subseteq X$ and $|\mathcal{C}_i \cap S| \leq 2k/3$. Define $S'_1 = S \cap (\mathcal{C}_1 \cup X)$ and $S'_2 = S \cap (\mathcal{C}_2 \cup X)$. Now, for each $i \in \{1, 2\}$, a subset Z of $\mathcal{C}_i \cup X$ satisfying $|\partial(Z)| \leq p$ (in $G[\mathcal{C}_i \cup X]$) and $|Z \cap S'_i| \leq q$ satisfies $|\partial(Z)| \leq p$ (in G) and $|Z \cap S_i| \leq q$. Hence

$$\begin{aligned} \mu(k, x) &\leq \mu\left(\frac{k}{3} + (d\sqrt{k}) + 1, x\right) \\ &\quad + \mu\left(\frac{2k}{3} + (d\sqrt{k}) + 1, x\right) \end{aligned}$$

However, if $k + t \leq p$ and $k \leq q$, then $\mu(k, x) < x$. Thus, choosing p and q sufficiently large compared to t and d and applying the generalization of Master's Theorem due to Akra and Bazzi [1] to the recurrence above we have that $\mu(k, x) = O(kx)$.

Notice now that a set Z such that $|\delta(Z)| \leq p$ and $|(Z \cap S)| \leq q$ has treewidth at most $t + q$ and hence is an extended $(p+t+q)$ -protrusion. Choosing $r = p + t + q$ and ζ to be a constant such that $\zeta \mu(k, x) \leq kx$ completes the proof of the lemma.

4 Kernelization for Separable Bidimensional Problems

In this section we give our kernelization algorithm. First we give the reduction rule which is applied exhaustively on the input to reduce its size. This rule is from [5] and it has been just modified here to work for graphs excluding a fixed graph H as a minor and satisfying finite integer index. Finally, in the last subsection, we prove Theorem 1.1 using the auxiliary results we have proved so far.

4.1 Reduction Rule The only reduction rule we apply has the following form:

If there is a constant size separator such that after its removal we obtain a connected component of unbounded size and of constant treewidth, then we replace this component with a subgraph of constant size.

We need the following generalization of lemma from [5] to H -minor free graphs. The proof is almost identical and we provide it here for completeness.

LEMMA 4.1. ([5]) *Let Π has finite integer index. Then there exists a constant $c(\Pi, \gamma)$ and an algorithm that given a graph $G = (V, E) \in \mathcal{G}_H$, an integer k and an extended γ -protrusion X in G with $|X| > c(\Pi, \gamma)$, runs in time $O(|X|)$ and returns a graph $G^* = (V^*, E^*) \in \mathcal{G}_H$ and an integer k^* such that $|V^*| < |V|$, $k^* \leq k$, and $(G^*, k^*) \in \Pi$ if and only if $(G, k) \in \Pi$.*

Proof. Let \mathcal{S} be a set of representatives for (Π, r) and let $c = \max_{Y \in \mathcal{S}} |Y|$. Similarly, let \mathcal{S}' be a set of representatives for $(\Pi, 2r)$ and let $c' = \max_{Y \in \mathcal{S}'} |Y|$. If $|X| > 3c'$ we find an extended $2r$ -protrusion $X' \subseteq X$ such that $c' < |X'| \leq 3c'$ and work on X' instead of X . This can be done in time $O(|X|)$ since $G[X]$ has treewidth at most r . From now on, we assume that $|X| \leq 3c'$. This initial step is the only step of the algorithm that does not work with constant size structures, and hence the running time of the algorithm is upper bounded by $O(|X|)$. The algorithm proceeds as follows.

Because Π has finite integer index there is a graph $H = (V_H, E_H) \in \mathcal{S}$ such that $H \equiv_{\Pi} G[X]$. We show how to compute H from X . Since k is given in unary, we have that $k \leq |V|^p$. Let $k_{max} = (6c')^p$. For every $G_1 = (V_1, E_1) \in \mathcal{S}$,

$G_2 = (V_2, E_2) \in \mathcal{S}$ and $k' \leq k_{max}$ we compute whether $(G_1 \oplus G_2, k') \in \Pi$. For each such triple the computation can be done in time $\mathcal{O}((|V_1| + |V_2|)^p)$ since Π has finite integer index [10, 14]. Now, for every $G_1 \in \mathcal{S}$ and $k' \leq (|X| + |V_1|)^p$ we compute whether $(G[X] \oplus G_1, k') \in \Pi$. When all these computations are done, the results are stored in a table.

It is not hard to see that $H \equiv_{\Pi} G[X]$ if and only if there exists a constant c such that for all $G_2 \in \mathcal{S}$ and $k' \leq k_{max}$, $(H \oplus G_2, k') \in \Pi \Leftrightarrow (G[X] \oplus G_2, k' + c) \in \Pi$. Also, c is the constant such that for all r -boundaried graphs G_2 and integers k' , $(H \oplus G_2, k') \in \Pi \Leftrightarrow (G[X] \oplus G_2, k' + c) \in \Pi$. For each $H \in \mathcal{C}$ we can check whether $H \equiv_{\Pi} G[X]$ using this condition and the pre-computed table, and if $H \equiv_{\Pi} G[X]$, find the constant c .

After we have found a $H \in \mathcal{S}$ and the corresponding constant c , such that $H \equiv_{\Pi} G[X]$, we make G^* from G by replacing the extended r -protrusion X with H . Also, we set $k^* = k - c$. Since $|X| > c$ and H has at most c vertices, $|V^*| < |V|$. By the choice of H and c , $(G^*, k^*) \in \Pi$ if and only if $(G, k) \in \Pi$. This concludes the proof.

4.2 Proof of Theorem 1.1

Proof. [Proof of Theorem 1.1] We give the proof only for minor-bidimensional problems. The proof of the theorem for contraction-bidimensional problems goes along the same line with the only difference that instead of Lemma 3.2 we need to use Lemma 3.3.

We first give a proof for a minor bidimensional problem Π that is also a minimization problem. Let Π be a bidimensional problem with the separation property and finite integer index. By Lemma 3.2, for $(G = (V, E), k) \in \Pi$, there is a set $S \subseteq V$ and a constant t , such that $|S| \leq t \cdot k$ and the treewidth of $G_S = G \setminus S = (V_S, E_S)$ is at most t . By Lemma 3.4, G contains an r -protrusion of size at least $\zeta|V|/tk$. The reduction algorithm exhaustively applies Lemma 4.1, with $\gamma = r$. Since an irreducible instance contains no r -protrusion of size at least $c(\Pi, \gamma)$ it follows that an irreducible instance $(G = (V, E), k) \in \Pi$ must satisfy $\zeta|V|/tk < c(\Pi, \gamma)$. Thus $|V|$ is at most $k \cdot tc(\Pi, \gamma)/\zeta = O(k)$. For the case that Π is a minor bidimensional maximization problem an identical argument shows that for any irreducible instance $(G, k) \in \bar{\Pi}$, $|V|$ is at most $O(k)$.

Now we show that our kernelization procedure runs in polynomial time. Observe that we can find a protrusion by guessing the boundary which is of constant size. Once given a protrusion X we can replace it with an equivalent instance in $O(|X|)$ time using the Lemma 4.1. This concludes that the

kernelization algorithm runs in polynomial time.

Theorem 1.1 has the following corollary. The list in the corollary is not exhaustive, and contains most of the well-known parameterized problems for which our theorem gives linear kernels.

COROLLARY 4.1. FEEDBACK VERTEX SET, VERTEX COVER, CONNECTED VERTEX COVER, EDGE DOMINATING SET, MINIMUM MAXIMUM MATCHING, ALMOST CONSTANT TREewidth, CYCLE PACKING and INDEPENDENT SET admit a linear kernel on H -minor free graphs. DOMINATING SET, r -DOMINATING SET, EDGE DOMINATING SET, CONNECTED DOMINATING SET TRIANGLE PACKING, and INDUCED MATCHING admit a linear kernel on an apex minor free graphs.

5 Conclusion

In this paper we related the meta-algorithmic theory of bidimensionality and the field of kernelization and affirmatively resolved an open question raised in [17]. In particular, we showed that every minor (contraction) bidimensional problems satisfying a separation property admit a linear kernel on graphs excluding a fixed graph (apex graph) H as a minor. Our results have two advantages over the linear sized kernels obtained in [5]. First, our results apply to much larger classes of sparse graphs, second our kernels are obtained using relatively simple combinatorial arguments instead of an approach based on topological decomposition. The requirement of separability in addition to bidimensionality is unavoidable, as k -PATH, the problem of finding a k -sized path in the given graph, is bidimensional but it is known not to admit polynomial kernel even on planar graphs unless the polynomial hierarchy collapses to the third level [7].

We conclude the article with an open problem - do all "reasonable" contraction-bidimensional parameters have linear kernels on H -minor-free graphs? In particular, does DOMINATING SET and CONNECTED DOMINATING SET?

References

- [1] M. AKRA AND L. BAZZI, *On the solution of linear recurrence equations*, Computational Optimization and Applications, 10 (1998), pp. 195–210.
- [2] J. ALBER, H. L. BODLAENDER, H. FERNAU, T. KLOKS, AND R. NIEDERMEIER, *Fixed parameter algorithms for dominating set and related problems on planar graphs*, Algorithmica, 33 (2002), pp. 461–493.
- [3] J. ALBER, M. R. FELLOWS, AND R. NIEDERMEIER, *Polynomial-time data reduction for dominating sets*, J. ACM, 51 (2004), pp. 363–384.

- [4] N. ALON AND S. GUTNER, *Kernels for the dominating set problem on graphs with an excluded minor*, Tech. Rep. TR08-066, Electronic Colloquium on Computational Complexity (ECCC), 2008.
- [5] H. BODLAENDER, F. V. FOMIN, D. LOKSHTANOV, E. PENNINKX, S. SAURABH, AND D. M. THILIKOS, *(Meta) kernelization*, in FOCS'09, IEEE Comput. Soc. Press, 2009, p. to appear.
- [6] H. L. BODLAENDER, *A partial k -arboretum of graphs with bounded treewidth*, Theor. Comp. Sc., 209 (1998), pp. 1–45.
- [7] H. L. BODLAENDER, R. G. DOWNEY, M. R. FELLOWS, AND D. HERMELIN, *On problems without polynomial kernels*, in Proc. 35th ICALP, vol. 5125 of LNCS, Springer, 2008, pp. 563–574.
- [8] H. L. BODLAENDER AND E. PENNINKX, *A linear kernel for planar feedback vertex set*, in Proceedings Third International Workshop on Parameterized and Exact Computation, IWPEC 2008, M. Grohe and R. Niedermeier, eds., Springer Verlag, Lecture Notes in Computer Science, vol. 5018, 2008, pp. 160–171.
- [9] H. L. BODLAENDER, E. PENNINKX, AND R. B. TAN, *A linear kernel for the k -disjoint cycle problem on planar graphs*, in Proceedings of the 19th International Symposium on Algorithms and Computation (ISAAC), vol. 5369 of LNCS, Springer, 2008, pp. 306–317.
- [10] H. L. BODLAENDER AND B. VAN ANTWERPENDE FLUITER, *Reduction algorithms for graphs of small treewidth*, Information and Computation, 167 (2001), pp. 86–119.
- [11] J. CHEN, H. FERNAU, I. A. KANJ, AND G. XIA, *Parametric duality and kernelization: Lower bounds and upper bounds on kernel size*, in Proc. 22nd STACS, vol. 3404 of LNCS, Springer, 2005, pp. 269–280.
- [12] ———, *Parametric duality and kernelization: Lower bounds and upper bounds on kernel size*, SIAM J. Comput., 37 (2007), pp. 1077–1106.
- [13] J. CHEN, I. A. KANJ, AND W. JIA, *Vertex Cover: Further observations and further improvements*, J. Algorithms, 41 (2001), pp. 280–301.
- [14] B. DE FLUITER, *Algorithms for Graphs of Small Treewidth*, PhD thesis, Utrecht University, 1997.
- [15] E. DEMAINE AND M. HAJIAGHAYI, *The bidimensionality theory and its algorithmic applications*, The computer Journal, (2007), pp. 332–337.
- [16] E. D. DEMAINE, F. V. FOMIN, M. HAJIAGHAYI, AND D. M. THILIKOS, *Bidimensional parameters and local treewidth*, SIAM J. Discrete Math., 18 (2004/05), pp. 501–511.
- [17] E. D. DEMAINE, F. V. FOMIN, M. HAJIAGHAYI, AND D. M. THILIKOS, *Subexponential parameterized algorithms on bounded-genus graphs and H -minor-free graphs*, J. ACM, 52 (2005), pp. 866–893.
- [18] E. D. DEMAINE AND M. HAJIAGHAYI, *Bidimensionality: New connections between FPT algorithms and PTASs*, in Proc. 16th SODA, ACM/SIAM, 2005, pp. 590–601.
- [19] E. D. DEMAINE AND M. HAJIAGHAYI, *Linearity of grid minors in treewidth with applications through bidimensionality*, Combinatorica, 28 (2008), pp. 19–36.
- [20] F. DORN, F. V. FOMIN, AND D. M. THILIKOS, *Subexponential parameterized algorithms*, Comp. Sc. Rev., 2 (2008), pp. 29–39.
- [21] R. G. DOWNEY AND M. R. FELLOWS, *Parameterized Complexity*, Springer, 1999.
- [22] J. FLUM AND M. GROHE, *Parameterized Complexity Theory*, Texts in Theoretical Computer Science. An EATCS Series, Springer-Verlag, Berlin, 2006.
- [23] F. V. FOMIN, P. A. GOLOVACH, AND D. M. THILIKOS, *Contraction bidimensionality: the accurate picture*, in ESA'09, LNCS, 2009, p. to appear.
- [24] J. GUO AND R. NIEDERMEIER, *Invitation to data reduction and problem kernelization*, SIGACT News, 38 (2007), pp. 31–45.
- [25] ———, *Linear problem kernels for NP-hard problems on planar graphs*, in Proc. 34th ICALP, vol. 4596 of LNCS, Springer, 2007, pp. 375–386.
- [26] J. GUO, R. NIEDERMEIER, AND S. WERNICKE, *Fixed-parameter tractability results for full-degree spanning tree and its dual*, in Proceedings of the Second International Workshop on Parameterized and Exact Computation (IWPEC), vol. 4169 of LNCS, Springer, 2006, pp. 203–214.
- [27] I. A. KANJ, M. J. PELSMAJER, G. XIA, AND M. SCHAEFER, *On the induced matching problem*, in Proceedings of the 25th Annual Symposium on Theoretical Aspects of Computer Science (STACS 2008), vol. 08001, Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, Berlin, 2008, pp. 397–408.
- [28] D. LOKSHTANOV, M. MNICH, AND SAURABH, *Linear kernel for planar connected dominating set*, in Proceedings of Theory and Applications of Models of Computation, (TAMC 2009), Lecture Notes in Comput. Sci., Springer, 2009.
- [29] H. MOSER AND S. SIKDAR, *The parameterized complexity of the induced matching problem in planar graphs*, in Proceedings First Annual International Workshop Frontiers in Algorithmics (FAW), vol. 4613 of LNCS, Springer, 2007, pp. 325–336.
- [30] R. NIEDERMEIER, *Invitation to fixed-parameter algorithms*, vol. 31 of Oxford Lecture Series in Mathematics and its Applications, Oxford University Press, Oxford, 2006.
- [31] G. PHILIP, V. RAMAN, AND S. SIKDAR, *Polynomial kernels for dominating set in $K_{i,j}$ -free and d -degenerate graphs*, manuscript, 2009.
- [32] S. THOMASSÉ, *A quadratic kernel for feedback vertex set*, in Proc. 20th SODA, ACM/SIAM, 2009, pp. 115–119.