

Kernelization and Sparseness: the Case of Dominating Set*

Pål Grønås Drange¹, Markus Dregi², Fedor V. Fomin³,
Stephan Kreutzer⁴, Daniel Lokshtanov⁵, Marcin Pilipczuk⁶,
Michał Pilipczuk⁷, Felix Reidl⁸, Fernando Sánchez Villaamil⁹,
Saket Saurabh¹⁰, Sebastian Siebertz¹¹, and Somnath Sikdar¹²

- 1 Department of Informatics, University of Bergen, Norway
pal.drange@ii.uib.no
- 2 Department of Informatics, University of Bergen, Norway
markus.dregi@ii.uib.no
- 3 Department of Informatics, University of Bergen, Norway
fomin@ii.uib.no
- 4 Institute of Software Technology and Theoretical Computer Science,
Technische Universität Berlin, Germany
stephan.kreutzer@tu-berlin.de
- 5 Department of Informatics, University of Bergen, Norway
daniello@ii.uib.no
- 6 Institute of Informatics, University of Warsaw, Poland
marcin.pilipczuk@mimuw.edu.pl
- 7 Institute of Informatics, University of Warsaw, Poland
michal.pilipczuk@mimuw.edu.pl
- 8 Theoretical Computer Science, Department of Computer Science, RWTH
Aachen University, Germany
reidl@cs.rwth-aachen.de
- 9 Theoretical Computer Science, Department of Computer Science, RWTH
Aachen University, Germany
fernando.sanchez@cs.rwth-aachen.de
- 10 Department of Informatics, University of Bergen, Norway: and
Institute of Mathematical Sciences, India
saket.saurabh@ii.uib.no, saket@imsc.res.in
- 11 Institute of Software Technology and Theoretical Computer Science,
Technische Universität Berlin, Germany
sebastian.siebertz@tu-berlin.de
- 12 Theoretical Computer Science, Department of Computer Science, RWTH
Aachen University, Germany
sikdar@cs.rwth-aachen.de

Abstract

We prove that for every positive integer r and for every graph class \mathcal{G} of bounded expansion, the r -DOMINATING SET problem admits a linear kernel on graphs from \mathcal{G} . Moreover, in the

* The research of P. Drange, F. Fomin, M. Pilipczuk, and M. Pilipczuk leading to these results has received funding from the European Research Council (ERC) under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 267959. Mi. Pilipczuk currently holds a post-doc position at Warsaw Centre of Mathematics and Computer Science, and is supported by the Foundation for Polish Science. F. Reidl, F. Villaamil, and S. Sikdar are supported by DFG-Project RO 927/13-1 "Pragmatic Parameterized Algorithms". M. Dregi and D. Lokshtanov are supported by the BeHard grant under the recruitment programme of the of Bergen Research Foundation. S. Saurabh is supported by PARAPPROX, ERC starting grant no. 306992. S. Kreutzer has been supported by the ERC under the EU's Horizon 2020 research and innovation programme (grant agreement No 648527).



© Pål Grønås Drange, Markus Dregi, Fedor V. Fomin, Stephan Kreutzer,
Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, Felix Reidl,
Fernando Sánchez Villaamil, Saket Saurabh, Sebastian Siebertz,
and Somnath Sikdar
licensed under Creative Commons License CC-BY



SYMPOSIUM
ON THEORETICAL
ASPECTS
OF COMPUTER
SCIENCE

33rd Symposium on Theoretical Aspects of Computer Science (STACS 2016).

Editors: Nicolas Ollinger and Heribert Vollmer; Article No. 31; pp. 31:1–31:14

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

more general case when \mathcal{G} is only assumed to be nowhere dense, we give an almost linear kernel on \mathcal{G} for the classic DOMINATING SET problem, i.e., for the case $r = 1$. These results generalize a line of previous research on finding linear kernels for DOMINATING SET and r -DOMINATING SET (Alber et al., JACM 2004, Bodlaender et al., FOCS 2009, Fomin et al., SODA 2010, Fomin et al., SODA 2012, Fomin et al., STACS 2013). However, the approach taken in this work, which is based on the theory of sparse graphs, is radically different and conceptually much simpler than the previous approaches.

We complement our findings by showing that for the closely related CONNECTED DOMINATING SET problem, the existence of such kernelization algorithms is unlikely, even though the problem is known to admit a linear kernel on H -topological-minor-free graphs (Fomin et al., STACS 2013). Also, we prove that for any somewhere dense class \mathcal{G} , there is some r for which r -DOMINATING SET is $W[2]$ -hard on \mathcal{G} . Thus, our results fall short of proving a sharp dichotomy for the parameterized complexity of r -DOMINATING SET on subgraph-monotone graph classes: we conjecture that the border of tractability lies exactly between nowhere dense and somewhere dense graph classes.

1998 ACM Subject Classification F.2.2 Nonnumerical algorithms and problems

Keywords and phrases kernelization, dominating set, bounded expansion, nowhere dense

Digital Object Identifier 10.4230/LIPIcs.STACS.2016.31

1 Introduction

In the classic DOMINATING SET problem, given a graph G and an integer k , we are asked to determine the existence of a subset $D \subseteq V(G)$ of size at most k such that every vertex $u \in V(G)$ is *dominated* by D : either u belongs to D itself, or it has a neighbor that belongs to D . The r -DOMINATING SET problem, for a positive integer r , is a generalization where each vertex of D dominates all the vertices at distance at most r from it. DOMINATING SET parameterized by the target size k plays a central role in parameterized complexity as it is a predominant example of a $W[2]$ -complete problem. Recall that the main focus in parameterized complexity is on designing *fixed-parameter* algorithms, or shortly *FPT* algorithms, whose running time on an instance of size n and parameter k has to be bounded by $f(k) \cdot n^c$ for some computable function f and constant c . Downey and Fellows introduced a hierarchy of parameterized complexity classes $FPT \subseteq W[1] \subseteq W[2] \subseteq \dots$ that is believed to be strict, see [8, 13]. As DOMINATING SET is $W[2]$ -complete in general, we do not expect it to be solvable in FPT time.

However, it turns out that various restrictions on the input graph lead to robust tractability of DOMINATING SET. Out of these, a particularly fruitful line of research concerned investigation of the complexity of the problem in sparse graph classes, like planar graphs, graphs of bounded genus, or graphs excluding some fixed graph H as a minor. In these classes we can even go one step further than just showing fixed-parameter tractability: It is possible to design a linear kernel for the problem. Formally, a *kernelization algorithm* (or a *kernel*) is a polynomial-time preprocessing procedure that given an instance (I, k) of a parameterized problem outputs another instance (I', k') of the same problem which is equivalent to (I, k) , but whose total size $|I'| + k'$ is bounded by $f(k)$ for some computable function f , called the *size* of the kernel. If f is polynomial (resp. linear), then such an algorithm is called a *polynomial* (resp. *linear*) *kernel*.

The quest for small kernels for DOMINATING SET on sparse graph classes began with the groundbreaking work of Alber et al. [1], who showed the first linear kernel for the problem

on planar graphs. Another important step was the work of Alon and Gutner [2, 21], who gave an $\mathcal{O}(k^c)$ kernel for the problem on H -topological-minor free graphs, where c depends on H only. Moreover, if $H = K_{3,h}$ for some h , then the size of the kernel is actually linear. This led Alon and Gutner to pose the following excellent question: Can one characterize the families of graphs where DOMINATING SET admits a linear kernel?

The research program sketched by the works of Alber et al. [1] and Alon and Gutner [2, 21] turned out to be one of particularly fruitful directions in parameterized complexity in recent years, and eventually led to the discovery of new and deep techniques. In particular, linear kernels for DOMINATING SET have been given for bounded genus graphs [3], apex-minor-free graphs [14], H -minor-free graphs [15], and most recently H -topological-minor-free graphs [16]. In all these results, the notion of *bidimensionality* plays the central role. Using variants of the Grid Minor Theorem, it is possible to understand well the connections between the minimum possible size of a dominating set in a graph and its treewidth. The considered graph classes also admit powerful decomposition theorems that follow from the Graph Minors project of Robertson and Seymour [26], or the recent work of Grohe and Marx [20] on excluding H as a topological minor. The combination of these tools provides a robust base for a structural analysis of the input instance.

Beyond the current frontier of H -topological-minor-free graphs [16], kernelization of DOMINATING SET was studied in graphs of bounded degeneracy. Recall that a graph is called *d -degenerate* if every subgraph contains a vertex of degree at most d . Philip et al. [24] obtained a kernel of size $\mathcal{O}(k^{(d+1)^2})$ on d -degenerate graphs for constant d . However, as proved by Cygan et al. [5], the exponent of the size of the kernel needs to increase with d at least quadratically (unless $\text{NP} \subseteq \text{coNP}/\text{poly}$).

As far as r -DOMINATING SET is concerned, the current most general result gives a linear kernel for any apex-minor-free class [14], and follows from a general protrusion machinery [4]. The techniques used by Fomin et al. [15, 16] for H -(topological)-minor-free classes are tailored to the classic DOMINATING SET problem, and do not carry over to an arbitrary radius r . Therefore, up to this point the existence of linear kernels for r -DOMINATING SET on H -(topological)-minor-free classes was open.

Sparsity. The general concept of sparsity, beyond graphs with excluded (topological) minors, has been recently the subject of intensive study both from the point of view of pure graph theory and of computer science. In particular, the notions of graph classes of *bounded expansion* and *nowhere dense* graph classes have been introduced by Nešetřil and Ossona de Mendez. The main idea behind these models is to establish an abstract notion of sparsity based on known properties of well-studied sparse graph classes, e.g. H -minor-free graphs, and to develop tools for combinatorial analysis of sparse graphs based only on this abstract notion. We refer to the book of Nešetřil and Ossona de Mendez [23] for an introduction to the topic.

To measure sparsity of a graph more formally, we need a few definitions. We say that a graph H is a *minor* of G if there exists a family $(X_v)_{v \in V(H)}$ of pairwise disjoint subsets of $V(G)$ such that for every $v \in V(H)$ the graph $G[X_v]$ is connected and for every $uv \in E(H)$ there exists an edge in G between X_u and X_v . The sets X_v are called *branch sets*, and the family $(X_v)_{v \in V(H)}$ is called a *minor model* of H in G . We say that H is an *r -shallow minor* of G if there exists a minor model of H in G such that for every branch set X_v the graph $G[X_v]$ is of radius at most r . The set of all r -shallow minors of a graph G is denoted by $G \nabla r$, and for a graph class \mathcal{G} we define $\mathcal{G} \nabla r = \bigcup_{G \in \mathcal{G}} (G \nabla r)$. The essence of sparsity of a graph or a graph class is captured in the following definition.

► **Definition 1.1** (Grad and bounded expansion). For a graph G and an integer $r \geq 0$, we define the *greatest reduced average density (grad)* at depth r as

$$\nabla_r(G) = \max_{M \in \mathcal{G}_{\nabla_r}} \text{density}(M) = \max_{M \in \mathcal{G}_{\nabla_r}} |E(M)|/|V(M)|.$$

For a graph class \mathcal{G} , we let $\nabla_r(\mathcal{G}) = \sup_{G \in \mathcal{G}} \nabla_r(G)$. A graph class \mathcal{G} then has *bounded expansion* if there exists a function $f: \mathbb{N} \rightarrow \mathbb{R}$ such that for all r we have that $\nabla_r(\mathcal{G}) \leq f(r)$.

Graph classes excluding a topological minor, such as planar and bounded-degree graphs, have bounded expansion [23]. Bounded expansion implies bounded degeneracy, since the degeneracy of G lies between $\nabla_0(G)$ and $2\nabla_0(G)$. However, having bounded expansion is much stronger than just bounded degeneracy: in some sense, it requires every shallow minor to be of bounded degeneracy, with the bound on degeneracy increasing with the depth of the minors. The notion of a *nowhere dense* graph class is a further relaxation of this concept.

As far as the theory of sparsity is concerned, from the point of view of theoretical computer science, of particular importance is the program of establishing fixed-parameter tractability of model checking first order logic on sparse graphs. A long line of work resulted in FPT algorithms for model checking first order formulae on more and more general classes of sparse graphs [6, 11, 12, 17, 19, 27], similarly to the story of kernelization of DOMINATING SET. Finally, FPT algorithms for the problem have been given for graph classes of bounded expansion by Dvořák et al. [11], and very recently for nowhere dense graph classes by Grohe et al. [19]. This is the ultimate limit of this program: as proven in [22, 11], for any class \mathcal{G} that is not nowhere dense (is *somewhere dense*) and is closed under taking subgraphs, model checking first order formulae on \mathcal{G} is not fixed-parameter tractable (unless $\text{FPT} = \text{W}[1]$).

Fixed-parameter tractability of r -DOMINATING SET on nowhere dense graph classes follows immediately from the result of Grohe et al. [19], as the problem is definable in first order logic for a fixed r ; an explicit algorithm was given earlier by Dawar and Kreutzer [7].

Our results. In this work we prove that having bounded expansion or being nowhere dense is sufficient for a graph class to admit an (almost) linear kernel for DOMINATING SET. Henceforth, for a graph G , we let $\text{ds}(G)$ denote the minimum size of a dominating set of G .

► **Theorem 1.2.** *Let \mathcal{G} be a graph class of bounded expansion. There exists a polynomial-time algorithm that given a graph $G \in \mathcal{G}$ and an integer k , either correctly concludes that $\text{ds}(G) > k$ or finds a subset of vertices $Y \subseteq V(G)$ of size $\mathcal{O}(k)$ with the property that $\text{ds}(G) \leq k$ if and only if $\text{ds}(G[Y]) \leq k$.*

► **Theorem 1.3.** *Let \mathcal{G} be a nowhere dense graph class and let $\varepsilon > 0$ be a real number. There exists a polynomial-time algorithm that given a graph $G \in \mathcal{G}$ and an integer k , either correctly concludes that $\text{ds}(G) > k$ or finds a subset of vertices $Y \subseteq V(G)$ of size $\mathcal{O}(k^{1+\varepsilon})$ with the property that $\text{ds}(G) \leq k$ if and only if $\text{ds}(G[Y]) \leq k$.*

For r -DOMINATING SET, for $r > 1$, we can give a linear kernel for any graph class of bounded expansion. Unfortunately, there is a technical subtlety that does not allow us to state the kernel in a nice form as in Theorems 1.2 and 1.3. Instead, we kernelize to an annotated version of the problem, where only a given subset of vertices of G needs to be dominated. The annotated version can be reduced to the classic one by simple gadgeteering but that, unfortunately, may lead to a slight increase in the bounded expansion guarantees.

In the following, by $\text{ds}_r(G)$ we denote the minimum size of an r -DOMINATING SET in a graph G , while for $Z \subseteq V(G)$, $\text{ds}_r(G, Z)$ denotes the minimum size of a (Z, r) -dominator in G , that is, a set $D \subseteq V(G)$ that r -dominates (i.e., is at distance at most r of) every vertex of Z .

► **Theorem 1.4.** *Let \mathcal{G} be a graph class of bounded expansion, and let r be a positive integer. There exists a polynomial-time algorithm that given a graph $G \in \mathcal{G}$ and an integer k , either correctly concludes that $\text{ds}_r(G) > k$ or finds subsets of vertices $Z \subseteq W \subseteq V(G)$, where $|W| = \mathcal{O}(k)$, with the property that $\text{ds}_r(G) \leq k$ if and only if $\text{ds}_r(G[W], Z) \leq k$.*

The obtained results strongly generalize the previous results on linear kernels for DOMINATING SET on sparse graph classes [1, 3, 14, 15, 16], since all the graph classes considered in these results have bounded expansion. Moreover, by giving a linear kernel for r -DOMINATING SET on any class \mathcal{G} of bounded expansion, we obtain the same result for any H -minor-free or H -topological-minor-free class as well. The existence of such kernels was not known before.

We see the main strength of our results in that they constitute an abrupt turn in the current approach to kernelization of DOMINATING SET and r -DOMINATING SET on sparse graphs: the tools used to develop the new algorithms are radically different from all the previously applied techniques. Instead of investigating bidimensionality and treewidth, and relying on intricate decomposition theorems originating in the work on graph minors, our algorithms exploit only basic properties of bounded expansion and nowhere dense graph classes. As a result, the full version of this work presents essentially self-contained proofs of all the stated kernelization results. The only external facts that we use are basic properties of weak and centered colorings, and the constant-factor approximation algorithm for r -DOMINATING SET of Dvořák [10]. All in all, the results show that only the combinatorial sparsity of a graph class is essential for designing (almost) linear kernels for DOMINATING SET, and further topological constraints like excluding some (topological) minor are unnecessary.

Lower bounds. We complement our study by proving that for the closely related CONNECTED DOMINATING SET problem, where the sought dominating set D is additionally required to induce a connected subgraph, the existence of even polynomial kernels for bounded expansion and nowhere dense graph classes is unlikely.

► **Theorem 1.5.** *There exists a class of graphs \mathcal{G} of bounded expansion such that CONNECTED DOMINATING SET does not admit a polynomial kernel when restricted to \mathcal{G} , unless $\text{NP} \subseteq \text{coNP/poly}$. Furthermore, \mathcal{G} is closed under taking subgraphs.*

Up to this point, linear kernels for CONNECTED DOMINATING SET were given for the same family of sparse graph classes as for DOMINATING SET: a linear kernel for the problem on H -topological-minor-free graphs was obtained by Fomin et al. [16]. Hence, classes of bounded expansion constitute the point where the kernelization complexity of both problems diverge; the connectivity constraint seems to have a completely different nature, and topological properties of the graph class become necessary to handle it efficiently.

Next, we show also that nowhere dense classes form the ultimate limit of parameterized tractability of r -DOMINATING SET, as it was for model checking first order formulae.

► **Theorem 1.6.** *For every somewhere dense graph class \mathcal{G} closed under taking subgraphs, there exists an integer r such that r -DOMINATING SET is $W[2]$ -hard on graphs from \mathcal{G} .*

In this extended abstract we provide an almost complete proof of Theorems 1.2 and 1.4. The full version of the paper contains [9] full proofs of all results, as well as an extended discussion in the introduction and conclusions.

2 Neighborhoods and closure lemma

The cornerstone of our approach stems from an observation of Gajarský et al. [18] that in a graph G from a class of bounded expansion, the number of possible neighborhoods in a

given subset of vertices X is bounded linearly in $|X|$ and, moreover, the number of vertices that have many neighbors in X is also small. Let $X \subseteq V(G)$ and $v \in V(G)$. We denote by $N_X(v) = N(v) \cap X$ the neighborhood of v in X .

► **Lemma 2.1** ([18]). *Let G be a graph, $X \subseteq V(G)$ be a vertex subset, and $R = V(G) \setminus X$. Then for every integer $p \geq \nabla_1(G)$ it holds that*

1. $|\{v \in R: |N_X(v)| \geq 2p\}| \leq 2p \cdot |X|$, and
2. $|\{A \subseteq X: |A| < 2p \text{ and } \exists v \in R A = N_X(v)\}| \leq (4^p + 2p)|X|$.

Consequently, the following bound holds:

$$|\{A \subseteq X: \exists v \in R A = N_X(v)\}| \leq \left(4^{\nabla_1(G)} + 4\nabla_1(G)\right) \cdot |X|.$$

This statement is best suited for the standard DOMINATING SET problem on graphs of bounded expansion, but to extend our result to r -DOMINATING SET, we need proper generalizations. Suppose G is a graph and X is a subset of its vertices. For $u \in V(G) \setminus X$ and positive integer r , we define the r -projection of u onto X as follows: $M_r^G(u, X)$ is the set of all those vertices $w \in X$, for which there exists a path P in G that starts in u , ends in w , has length at most r , and whose all internal vertices do not belong to X . Whenever the graph is clear from the context, we omit the superscript. In the following we will use the following strengthening of Lemma 2.1(1).

Let \mathcal{G} be a class of bounded expansion, $G \in \mathcal{G}$, r a positive integer, and $X \subseteq V(G)$. A set $\text{cl}_r(X)$ is called an r -closure of X if (1) $X \subseteq \text{cl}_r(X) \subseteq V(G)$; (2) $|\text{cl}_r(X)| \leq ((r-1)\xi + 2) \cdot |X|$, where $\xi = \lceil 2\nabla_{r-1}(\mathcal{G}) \rceil$; and (3) $|M_r^G(u, \text{cl}_r(X))| \leq \xi(1 + (r-1)\xi)$ for each $u \in V(G) \setminus \text{cl}_r(X)$.

► **Lemma 2.2** (Closure lemma). *Let \mathcal{G} be a class of bounded expansion. There exists an algorithm that, given a graph $G \in \mathcal{G}$, positive integer r , and $X \subseteq V(G)$, computes the r -closure of X .*

Proof. Consider the following iterative procedure: (1) Start with $H = G$ and $Y = X$. We will maintain the invariant that $Y \subseteq V(H)$. (2) As long as there exists a vertex $u \in V(H) \setminus Y$ with $|M_r^H(u, Y)| \geq \xi$, do the following:

- Select an arbitrary subset $Z_u \subseteq M_r^H(u, Y)$ of size ξ .
 - For each $w \in Z_u$, select a path P_w that starts at u , ends at w , has length at most r , and all its internal vertices are in $V(H) \setminus Y$.
 - Modify H by contracting $\bigcup_{w \in Z_u} (V(P_w) \setminus \{w\})$ onto u , and add the obtained vertex to Y .
- Observe that in a round of the procedure above we always make a contraction of a connected subgraph of $H - Y$ of radius at most $r - 1$. Also, the resulting vertex falls into Y and hence does not participate in future contractions. Thus, at each point H is an $(r - 1)$ -shallow minor of G . For any moment of the procedure and any $u \in V(H)$, by $\tau(u)$ we denote the subset of original vertices of G that were contracted onto u during earlier rounds. Note that either $\tau(u) = \{u\}$ when u is an original vertex of G , or $|\tau(u)| \leq 1 + (r - 1)\xi$.

We claim that the presented procedure stops after at most $|X|$ rounds. Suppose otherwise, that we successfully constructed the graph H and subset Y after $|X| + 1$ rounds. Examine graph $H[Y]$. This graph has $2|X| + 1$ vertices: $|X|$ original vertices of X and $|X| + 1$ vertices that were added during the procedure. Whenever a vertex u is added to Y after contraction, then it introduces at least ξ new edges to $H[Y]$: these are edges that connect the contracted vertex with the vertices of Z_u . Hence, $H[Y]$ has at least $\xi(|X| + 1)$ edges, which means that

$$\text{density}(H[Y]) = \frac{|E(H[Y])|}{|Y|} \geq \frac{\xi(|X| + 1)}{2|X| + 1} > \nabla_{r-1}(\mathcal{G}).$$

This is a contradiction with the fact that H is an $(r - 1)$ -shallow minor of G .

Therefore, the procedure stops after at most $|X|$ rounds producing pair (H, Y) where $|M_r^H(u, Y)| < \xi$ for each $u \in V(H) \setminus Y$. Define $\text{cl}_r(X) = \tau(Y) = \bigcup_{u \in Y} \tau(u)$. Property (1) is obvious. Since $|\tau(u)| = 1$ for each original vertex of X and $|\tau(u)| \leq 1 + (r - 1)\xi$ for each u that was added during the procedure, property (2) follows. We are left with property (3).

By the construction $V(H) \setminus Y = V(G) \setminus \text{cl}_r(X)$. Take any $u \in V(H) \setminus Y$ and observe that $M_r^G(u, \text{cl}_r(X)) \subseteq \tau(M_r^H(u, Y))$. Since $|M_r^H(u, Y)| < \xi$ for each $u \in V(H) \setminus Y$ and $|\tau(u)| \leq 1 + (r - 1)\xi$ for each $u \in V(H)$, property (3) follows. \blacktriangleleft

As we can safely assume $\nabla_{r-1}(\mathcal{G}) \geq 1$ (otherwise \mathcal{G} contains only forests), we can use simplified, weaker bounds: $|\text{cl}_r(X)| \leq 3r\nabla_{r-1}(\mathcal{G}) \cdot |X|$ and $|M_r^G(u, \text{cl}_r(X))| \leq 9r\nabla_{r-1}(\mathcal{G})^2$. Observe that Lemma 2.2 is not merely a generalization of Lemma 2.1(1) to r -neighborhoods. It shows that a certain maximality property can be achieved; this property may be not true if, even for $r = 1$, we would construct $\text{cl}_r(X)$ from X by just adding all vertices with many neighbors in X .

The generalization of Lemma 2.1(2) which we will use later is the following.

► **Lemma 2.3.** *Let \mathcal{G} be a class of bounded expansion and let r be a positive integer. Let $G \in \mathcal{G}$ be a graph and $X \subseteq V(G)$. Then*

$$|\{Y : Y = M_r(u, X) \text{ for some } u \in V(G) \setminus X\}| \leq c \cdot |X|,$$

for some constant c depending only on r and the grads of \mathcal{G} .

The proof of Lemma 2.3 is essentially contained in the PhD thesis of the eighth author, Reidl [25]. For the sake of completeness, in the full version we give a self-contained proof based on the presentation of Reidl.

Finally, for the proof of Theorem 1.4 for $r > 1$ we will need the following lemma.

► **Lemma 2.4** (Short paths closure lemma). *Let \mathcal{G} be a class of bounded expansion and let r be a positive integer. Let $G \in \mathcal{G}$ be a graph and $X \subseteq V(G)$. Then there is a superset of vertices $X' \supseteq X$ with the following properties:*

1. *Whenever $\text{dist}_G(u, v) \leq r$ for some distinct $u, v \in X$, then $\text{dist}_{G[X']}(u, v) = \text{dist}_G(u, v)$.*
 2. *$|X'| \leq Q_r(\nabla_{r-1}(\mathcal{G})) \cdot |X|$ for some polynomial Q_r .*
- Moreover, X' can be computed in polynomial time.

Proof sketch. We start with $X' := X_0 := \text{cl}_r(X)$ and then, for every $u, v \in X_0$, we add to X' a path $P_{u,v}$ between u and v that is shortest among paths that do not visit X_0 apart from the endpoints, provided that it is of length at most r . The first required property of X' and polynomial time computability is immediate.

For the size bound, first note that the properties of $X_0 = \text{cl}_r(X)$ ensure that every vertex $x \in X' \setminus X_0$ can lie only on a constant number of paths $P_{u,v}$, as the endpoints of such path needs to lie in $M_r(x, X_0)$. On the other hand, if too many paths $P_{u,v}$ are pairwise disjoint, they constitute a dense r -shallow minor of G . We infer that only $\mathcal{O}(|X_0|) = \mathcal{O}(|X|)$ paths $P_{u,v}$ are added, and hence $|X'| = \mathcal{O}(|X|)$, as required. \blacktriangleleft

3 Linear kernel in graphs of bounded expansion

Let us fix a graph class \mathcal{G} that has bounded expansion, and let (G, k) be the input instance of r -DOMINATING SET, where $G \in \mathcal{G}$. We assume that \mathcal{G} is fixed, and hence so are also the values of $\nabla_i(\mathcal{G})$ for all nonnegative integers i . We start with the following pivot definition.

► **Definition 3.1** (*r*-domination core). Let G be a graph and Z be a subset of vertices. We say that Z is an *r*-domination core in G if every minimum-size (Z, r) -dominator in G is also an *r*-dominating set in G .

Clearly, the whole $V(G)$ is an *r*-domination core, but we will look for an *r*-domination core that is small in terms of k . Note that if Z is an *r*-domination core, then $\text{ds}_r(G) = \text{ds}_r(G, Z)$. Let us remark that in this definition we do not require that every (Z, r) -dominator is an *r*-dominating set in G ; there can exist (Z, r) -dominators that are not of minimum size and that do not dominate the whole graph.

The heart of our argument lies in providing a small domination core.

► **Theorem 3.2.** *There exists a constant c_{coresize} depending on r and the grads of \mathcal{G} only and a polynomial-time algorithm that, given an instance (G, k) where $G \in \mathcal{G}$, either correctly concludes that $\text{ds}_r(G) > k$, or finds an *r*-domination core $Z \subseteq V(G)$ with $|Z| \leq c_{\text{coresize}} \cdot k$.*

We fix G and k in the following to improve readability. For the proof of Theorem 3.2 we start with $Z = V(G)$ and gradually reduce $|Z|$ by removing one vertex at a time, while maintaining the invariant that Z is an *r*-domination core. To this end, we need to prove the following lemma, from which Theorem 3.2 follows trivially as explained:

► **Lemma 3.3.** *There exists a constant c_{coresize} depending on r and the grads of \mathcal{G} only and a polynomial-time algorithm that, given an *r*-domination core $Z \subseteq V(G)$ with $|Z| > c_{\text{coresize}} \cdot k$, either correctly concludes that $\text{ds}_r(G) > k$, or finds a vertex $z \in Z$ such that $Z \setminus \{z\}$ is still an *r*-domination core.*

In Sections 3.1 and 3.2 we prove Lemma 3.3. Then, in Section 3.3 we show how Theorem 3.2 implies Theorems 1.2 and 1.4.

3.1 Iterative extraction of Z -dominators

The first phase of the algorithm of Lemma 3.3 is to build a structural decomposition of the graph G . More precisely, we try to “pull out” a small set X of vertices that *r*-dominates Z , so that after removing them, Z contains a large subset S that is $2r$ -scattered in the remaining graph.¹ Given such a structure, intuitively we can argue that in any optimal (Z, r) -dominator, vertices of X serve as “hubs” that route almost all the domination paths leading to vertices of S . This is because any vertex of $V(G) \setminus X$ can *r*-dominate only at most one vertex from S via a path that avoids X . Since S will be large compared to X , some vertices of S will be indistinguishable from the point of view of *r*-domination routed through X , and these will be precisely the vertices that can be removed from the domination core. The identification of the irrelevant dominatee will be the goal of the second phase of the algorithm, whereas the goal of this phase is to construct the pair (X, S) .

Our main tool in this part is the constant-factor approximation for *r*-DOMINATING SET proved by Dvořák [10]. We use the following convenient form; the full version of our work describes in detail how to derive it from the work of Dvořák.

► **Lemma 3.4.** *Let r be a positive integer. There is a polynomial P_r and a polynomial-time algorithm that, given a graph G , a vertex subset $Z \subseteq V(G)$ and an integer k , finds either*

- *a (Z, r) -dominator in G of size at most $P_r(\nabla_r(G)) \cdot k$, or*
- *a subset of Z of size at least $k + 1$ that is $2r$ -scattered in G .*

¹ Recall that a set $S \subseteq V(G)$ is ℓ -scattered if every two distinct vertices of S are within distance greater than ℓ ; a $2r$ -scattered set of size $k + 1$ is an obstruction for an *r*-dominating set of size k .

Let $C_{\text{dv}} = P_r(\nabla_r(\mathcal{G}))$ be the approximation ratio of the algorithm of Lemma 3.4. Given Z , we first apply the algorithm of Lemma 3.4 to G , Z , and the parameters r and k . Thus, we either find a (Z, r) -dominator Y_1 such that $|Y_1| \leq C_{\text{dv}} \cdot k$, or we find a subset $S \subseteq Z$ of size at least $k + 1$ that is $2r$ -scattered in G . In the latter case, since S is an obstruction to an r -dominating set of size at most k , we may terminate the algorithm and provide a negative answer. Hence, from now on we assume that Y_1 has been successfully constructed.

Let C_0 be a constant depending on $\nabla(\mathcal{G})$, to be defined later. Now, in search for the pair (X, S) , we inductively construct sets $X_1, Y_2, X_2, Y_3, \dots$ such that $Y_1 \subseteq X_1 \subseteq Y_2 \subseteq \dots$ using the following definitions:

- If Y_i is already defined, then set $X_i = \text{cl}_{3r}(Y_i)$.
- If X_i is already defined, then apply the algorithm of Lemma 3.4 to $G - X_i$, $Z \setminus X_i$, and the parameters r and $C_0 \cdot |X_i|$.
 1. Suppose the algorithm finds a set $S \subseteq Z \setminus X_i$ that is $2r$ -scattered in $G - X_i$ and has cardinality greater than $C_0 \cdot |X_i|$. Then we let $X = X_i$, terminate the procedure and proceed to the second phase with the pair (X, S) .
 2. Otherwise, the algorithm has found a $(Z \setminus X_i, r)$ -dominator D_{i+1} in $G - X_i$ of size at most $C_{\text{dv}} \cdot C_0 \cdot |X_i|$. Then set $Y_{i+1} = X_i \cup D_{i+1}$ and proceed.

Let $\Gamma_{\text{cl}} = 9r \nabla_{3r-1}(\mathcal{G})$ be the bound on the size blow-up in Lemma 2.2 applied to radius $3r$, and let $\Delta_{\text{cl}} = 27r \nabla_{3r-1}(\mathcal{G})^2$ be the upper bound on the sizes of $3r$ -projections given by Lemma 2.2. From Lemmas 3.4, 2.2, and a trivial induction we infer that the following bounds hold for all i for which (Y_i, X_i) were constructed:

$$|Y_i| \leq C_{\text{dv}} \Gamma_{\text{cl}}^{i-1} (1 + C_{\text{dv}} C_0)^{i-1} \cdot k \quad \text{and} \quad |X_i| \leq C_{\text{dv}} \Gamma_{\text{cl}}^i (1 + C_{\text{dv}} C_0)^{i-1} \cdot k.$$

For a nonnegative integer i , let $K_i = C_{\text{dv}} \Gamma_{\text{cl}}^i (1 + C_{\text{dv}} C_0)^{i-1}$.

In this manner, the algorithm consecutively extracts dominators D_2, D_3, D_4, \dots and performs $3r$ -closure, constructing sets X_2, X_3, X_4, \dots up to the point when case (1) is encountered. Then the computation is terminated and the sought pair (X, S) is constructed. We now claim that case (1) always happens within a constant number of iterations.

► **Lemma 3.5.** *Let $\Lambda = \sum_{i=0}^r \Delta_{\text{cl}}^i \leq (r+1) \Delta_{\text{cl}}^r$. Assuming that $|Z| > K_\Lambda \cdot k$, the construction terminates yielding some pair (X, S) before performing Λ iterations.*

Proof. For the sake of contradiction, suppose Y_Λ and X_Λ were successfully constructed. Since $|Z| > K_\Lambda \cdot k$ and $|X_\Lambda| \leq K_\Lambda \cdot k$, there is some vertex $u \in Z \setminus X_\Lambda$. For an index $1 \leq i \leq \Lambda$, we shall say that a vertex $w \in X_i \setminus X_{i-1}$ is *i-good* if there is a path P that starts at u , ends at w , has length at most r , and all its internal vertices do not belong to X_i (we denote $X_0 = \emptyset$). Vertex w is *good* if it is good for some index i .

► **Claim 3.6.** *The number of good vertices is at most $\Lambda - 1$.*

Proof of claim. Let w be any good vertex, and let P be a path certifying this. Let $q \leq r$ be the length of P , and denote the vertices of P by u_i for $0 \leq i \leq q$, where $u_0 = u$ and $u_q = w$. Observe that internal vertices of P can belong only to sets $X_j \setminus X_{j-1}$ for $j > i$, or to $V(G) \setminus X_\Lambda$. We say that a vertex u_ℓ of P is *important* if there is an index j , with $i \leq j \leq \Lambda$, such that $u_\ell \in X_j \setminus X_{j-1}$ but $u_{\ell'} \notin X_j$ for all $\ell' < \ell$. Clearly, $w = u_q$ is important. Let $\ell_1 < \ell_2 < \dots < \ell_p = q$ be the indices of important vertices on P , and let $j_1 > j_2 > \dots > j_p$ be such that $u_{\ell_i} \in X_{j_i} \setminus X_{j_i-1}$, for all $1 \leq i \leq p$. We will denote $\ell_0 = 0$, so $u_{\ell_0} = u$, and $j_0 = \Lambda + 1$ (denoting $X_{\Lambda+1} = V(G)$).

Consider any index i with $1 \leq i \leq p$. Observe that on the part between $u_{\ell_{i-1}}$ and u_{ℓ_i} , path P never entered $X_{j_{i-1}-1}$, because first such entrance would constitute an important

vertex that was not recorded. Since $u_{\ell_i} \in X_{j_{i-1}-1}$, we infer that $u_{\ell_i} \in M_r^G(u_{\ell_{i-1}}, X_{j_{i-1}-1})$. Since $X_{j_{i-1}-1} = \text{cl}_{3r}(Y_{j_{i-1}-1})$ by the construction, we infer that $|M_r^G(u_{\ell_{i-1}}, X_{j_{i-1}-1})| \leq \Delta_{\text{cl}}$. Therefore, once vertex $u_{\ell_{i-1}}$ is selected, there are at most Δ_{cl} choices for the next important vertex u_{ℓ_i} . We infer that the choice of the sequence of important vertices on P can be modeled by taking at most r decisions, each from a selection of at most Δ_{cl} options. Since w is the last important vertex, there are at most $\sum_{i=1}^r \Delta_{\text{cl}}^i = \Lambda - 1$ ways to select w . ◀

► **Claim 3.7.** *For every $1 \leq i \leq \Lambda$, there is an i -good vertex.*

Proof of claim. Recall that $D_i \subseteq X_i \setminus X_{i-1}$ is a $(Z \setminus X_{i-1}, r)$ -dominator in the graph $G - X_{i-1}$. Since $u \in Z \setminus X_{i-1}$, in $G - X_{i-1}$ there is a path P of length at most r from u to a vertex of D_i . Take w to be the first vertex of this path that belongs to $X_i \setminus X_{i-1}$. Then the prefix of P from u to w certifies that w is an i -good vertex. ◀

Claims 3.6 and 3.7 contradict each other, which finishes the proof. ◀

In Lemma 3.3 we will set $c_{\text{coresize}} = K_\Lambda$, so that Lemma 3.5 can be applied. Therefore, unless the size of Z is bounded by $K_\Lambda \cdot k$, the construction terminates within $\Lambda = \sum_{i=0}^r \Delta_{\text{cl}}^i \leq (r+1)\Delta_{\text{cl}}^r$ iterations with a pair (X, S) . By the construction of X and S , we have the following properties:

- $|X| \leq K_\Lambda \cdot k$ and $|S| > C_0 \cdot |X|$;
- X is a (Z, r) -dominator in G (because $Y_1 \subseteq X$);
- for each $u \in V(G) \setminus X$, we have $|M_{3r}^G(u, X)| \leq \Delta_{\text{cl}}$;
- $S \subseteq Z \setminus X$ and S is $2r$ -scattered in $G - X$.

3.2 Finding an irrelevant dominatee

Given G , Z , and the constructed sets X and S , we denote by $R = V(G) \setminus X$ the set of vertices outside X . Using this notation, S is $2r$ -scattered in the graph $G[R]$. Recall that for any vertex $u \in R$, we have $|M_{3r}(u, X)| \leq \Delta_{\text{cl}}$.

Define the following equivalence relation \simeq on S : for $u, v \in S$, let

$$u \simeq v \iff M_i(u, X) = M_i(v, X) \text{ for each } 1 \leq i \leq 3r.$$

Let us denote by C_{nei} the constant c given by Lemma 2.3 for class \mathcal{G} and radius $3r$. Hence, the number of different $3r$ -projections in X of vertices of R is bounded by $C_{\text{nei}} \cdot |X|$.

► **Lemma 3.8.** *The equivalence relation \simeq has at most $C_{\text{nei}} \cdot (3r)^{\Delta_{\text{cl}}} \cdot |X|$ classes.*

Proof. Observe that for each $u \in S$,

$$M_1(u, X) \subseteq M_2(u, X) \subseteq \dots \subseteq M_{3r-1}(u, X) \subseteq M_{3r}(u, X).$$

By Lemma 2.3, the number of choices for $M_{3r}(u, X)$ is at most $C_{\text{nei}} \cdot |X|$. Moreover, since $u \in R$, we have that $|M_{3r}(u, X)| \leq \Delta_{\text{cl}}$. Hence, to define the sets $M_i(u, X)$ for $1 \leq i < 3r$ it suffices, for every $w \in M_{3r}(u, X)$, to choose the smallest index j , $1 \leq j \leq 3r$, such that $w \in M_j(u, X)$. The number of such choices is at most $(3r)^{\Delta_{\text{cl}}}$, and hence the claim follows. ◀

We can finally set the constant C_0 that was introduced in the previous section. We let $C_0 = (\Delta_{\text{cl}} + 1) \cdot C_{\text{nei}} \cdot (3r)^{\Delta_{\text{cl}}}$. Since we have that $|S| > C_0 \cdot |X|$, from Lemma 3.8 and the pigeonhole principle we infer that there is a class κ of relation \simeq with $|\kappa| > \Delta_{\text{cl}} + 1$. Note that we can find such a class κ in polynomial time, by computing the classes of \simeq directly

from the definition and examining their sizes. We are ready to prove the final lemma of this section: any vertex of κ can be removed from the r -domination core Z (recall that $S \subseteq Z$), concluding the proof of Lemma 3.3.

► **Lemma 3.9.** *Let z be an arbitrary vertex of κ . Then $Z \setminus \{z\}$ is an r -domination core.*

Proof. Let $Z' = Z \setminus \{z\}$. Take any minimum-size (Z', r) -dominator D in G . If D also dominates z , then D is a minimum-size (Z, r) -dominator as well. Since Z was an r -domination core, we infer that D is an r -dominating set in G , and we are done. Hence, suppose z is not r -dominated by D . We prove that this case leads to a contradiction.

Every vertex $s \in \kappa \setminus \{z\}$ is r -dominated by D . For each such s , let $v(s)$ be an arbitrarily chosen vertex of D that r -dominates s , and let $P(s)$ be an arbitrarily chosen path of length at most r that connects $v(s)$ with s .

► **Claim 3.10.** *For each $s \in \kappa \setminus \{z\}$, path $P(s)$ does not pass through any vertex of X (in particular $v(s) \notin X$). Consequently, vertices $v(s)$ for $s \in \kappa \setminus \{z\}$ are pairwise different.*

Proof of claim. Suppose otherwise and let w be the vertex of $V(P(s)) \cap X$ that is closest to s on $P(s)$. Then the suffix of $P(s)$ from w to s certifies that $w \in M_j(s, X)$, for j being the length of this suffix. As $s \simeq z$, we also have that $w \in M_j(z, X)$, so there is a path Q of length at most j from w to z . By concatenating the prefix of $P(s)$ from $v(s)$ to w with Q we obtain a walk of length at most r from $v(s)$ to z , a contradiction with the assumption that z is not r -dominated by D .

For the second part of the claim, suppose $v(s) = v(s')$ for some distinct $s, s' \in \kappa \setminus \{z\}$. Then the concatenation of $P(s)$ and $P(s')$ would be a path of length at most $2r$ connecting s and s' that is entirely contained in $G[R]$. This would be a contradiction with the fact that S is $2r$ -scattered in $G[R]$. ◀

Let $W = \{v(s) : s \in \kappa \setminus \{z\}\}$. From Claim 3.10 we have that $|W| = |\kappa \setminus \{z\}| \geq \Delta_{\text{cl}} + 1$. Define $D' = (D \setminus W) \cup M_{3r}(z, X)$. Since $|M_{3r}(z, X)| \leq \Delta_{\text{cl}}$, we have that $|D'| < |D|$.

► **Claim 3.11.** *D' is a (Z', r) -dominator.*

Proof of claim. For the sake of contradiction, suppose there is some $a \in Z'$ that is not r -dominated by D' . Since a was r -dominated by D and $D \setminus D' = W$, there must be a vertex $s \in \kappa \setminus \{z\}$ such that vertex $v(s)$ r -dominates a . Consequently, in G there is a path Q_0 of length at most r that leads from $v(s)$ to a . Furthermore, since X is a (Z, r) -dominator, there is a path Q_1 of length at most r that leads from a to some $x \in X$. Let Q be the concatenation of $P(s)$, Q_0 , and Q_1 ; Q is a walk of length at most $3r$ that connects s and $x \in X$.

Let x' be the first (closest to s) vertex on Q that belongs to X ; such a vertex exists as $x \in X$ is on Q . As the length of Q is at most $3r$, we have $x' \in M_{3r}(s, X)$. Since $s \simeq z$, we have $x' \in M_{3r}(z, X)$, and, consequently, $x' \in D'$. However, by Claim 3.10, x' does not lie on $P(s)$. Hence x' lies on the part of Q between $v(s)$ and x , but each vertex of this part is at distance at most r from a on Q . Thus a is r -dominated by x' , a contradiction. ◀

As $|D'| < |D|$, Claim 3.11 is a contradiction with the assumption that D is a minimum-size (Z', r) -dominator. This concludes the proof. ◀

3.3 Reducing dominators

In the rest of this section we work with arbitrary r towards the proof of Theorem 1.4. At some point we will argue that for $r = 1$, the statement of Theorem 1.2 is immediate. Having reduced the number of vertices whose domination is essential, we arrive at the situation where the vast majority of vertices serve only the role of dominators, or, when $r > 1$, they serve as connections between dominators with dominatees. Now, it is relatively easy to reduce the number of candidate dominators in one step. This immediately gives the sought kernel for $r = 1$, i.e., proves Theorem 1.2. For $r > 2$, the treatment of vertices connecting dominators and dominatees without introducing additional gadgets turns out to be problematic. Therefore, we are unable to give a kernel that is an induced subgraph of the original graph, and we resort to the statement of Theorem 1.4.

The algorithm of Theorem 1.4 works as follows. First, we apply the algorithm of Theorem 3.2 to compute a small domination core in the graph. In case the algorithm gives a negative answer, we output that $\text{ds}_r(G) > k$. Hence, from here on, we assume that we have correctly computed an r -domination core $Z_0 \subseteq V(G)$ of size $\mathcal{O}(k)$.

Compute $Z = \text{cl}_r(Z_0)$ using Lemma 2.2; then we have that $|Z| \leq 3r \nabla_{r-1}(\mathcal{G}) |Z_0| = \mathcal{O}(k)$. Observe that in any graph, any superset of an r -domination core is also an r -domination core; this follows easily from the definition. Consequently, Z is an r -domination core in G . Partition $V(G) \setminus Z$ into equivalence classes with respect to the following relation \simeq , defined similarly as in Section 3.2: For $u, v \in V(G) \setminus Z$, let:

$$u \simeq v \iff M_i(u, Z) = M_i(v, Z) \text{ for each } 1 \leq i \leq r.$$

From Lemma 2.2 we know that for each $u \in V(G) \setminus Z$, it holds that $|M_i(v, Z)| \leq 9r \nabla_{r-1}(\mathcal{G})^2$. Moreover, Lemma 2.3 implies that the number of possible different projections $M_r(u, Z)$ for $u \in V(G) \setminus Z$ is at most $c \cdot |Z|$, for some constant c depending on the grads of \mathcal{G} . Hence, using the same reasoning as in the proof of Lemma 3.8 we obtain the following. (Fully formal verification of Claims 3.12–3.15 is contained in the full version.)

► **Claim 3.12.** For $C = c \cdot r^{9r \nabla_{r-1}(\mathcal{G})^2}$, the equivalence relation \simeq has at most $C \cdot |Z|$ classes.

Construct Y as follows: start with Z and, for each equivalence class κ of relation \simeq , add an arbitrarily selected member v_κ of κ . Hence we have that $|Y| \leq (C + 1) \cdot |Z|$, so in particular $|Y| = \mathcal{O}(k)$. The following claim follows from a standard replacement argument.

► **Claim 3.13.** There exists a minimum-size r -dominating set in G that is contained in Y .

Theorem 1.2 now follows from the following immediate claim.

► **Claim 3.14.** If $r = 1$, then $\text{ds}(G) \leq k$ if and only if $\text{ds}(G[Y]) \leq k$.

For Theorem 1.4, we need not only to preserve the potential dominatees and dominators (the set Y), but also distances (up to length r) between them. To this end, we run the algorithm of Lemma 2.4 on set Y , and let $W = Y'$ be the obtained superset of Y . By Lemma 2.4 we have that $|W| = \mathcal{O}(k)$. Then Theorem 1.4 follows immediately from the following claim, which in turn follows easily from the properties of the set W promised by Lemma 2.4.

► **Claim 3.15.** $\text{ds}_r(G) \leq k$ if and only if $\text{ds}_r(G[W], Z) \leq k$.

4 Conclusions

We have shown that, for each $r \geq 1$, r -DOMINATING SET admits a linear kernel on any graph class of bounded expansion. Before this work, the most general family of graph classes where

such a kernelization result was known were apex-minor-free graphs [14], whereas in the case of the classic DOMINATING SET, linear kernels were shown also for general H -minor-free [15] and H -topological-minor-free classes [16]. Moreover, for $r = 1$, i.e., the DOMINATING SET problem, we can also give a kernel on any nowhere dense class of graphs, at the cost of increasing the size bound to almost linear, i.e., $\mathcal{O}(k^{1+\varepsilon})$ for any $\varepsilon > 0$. These results vastly and broadly extend the current frontier of kernelization results for domination problems on sparse graph classes.

The most important question left is understanding the kernelization complexity of r -DOMINATING SET on nowhere dense graph classes. So far we know that this problem admits a linear kernel on any class of bounded expansion, for each r , whereas on any somewhere dense class closed under taking subgraphs, for some r it is W[2]-hard. Our approach for bounded expansion graph classes fails to generalize to nowhere dense classes mostly because of technical reasons. We believe that, in fact, for any nowhere dense class \mathcal{G} and any positive integer r , r -DOMINATING SET has an almost linear kernel on \mathcal{G} . Together with the lower bound of Theorem 1.6, this would confirm the following dichotomy conjecture that we pose:

- **Conjecture 1.** Let \mathcal{G} be a graph class closed under taking subgraphs and $r \in \mathbb{N}$. Then:
- If \mathcal{G} is nowhere dense, then for every $r \geq 1$ and real $\varepsilon > 0$, r -DOMINATING SET admits an $\mathcal{O}(k^{1+\varepsilon})$ kernel on \mathcal{G} .
 - If \mathcal{G} is somewhere dense, then r -DOMINATING SET is W[2]-hard on \mathcal{G} for some $r \geq 1$.

References

- 1 Jochen Alber, Michael R. Fellows, and Rolf Niedermeier. Polynomial-time data reduction for Dominating Set. *J. ACM*, 51(3):363–384, 2004.
- 2 Noga Alon and Shai Gutner. Kernels for the Dominating Set problem on graphs with an excluded minor. *Electronic Colloquium on Comp. Complexity (ECCC)*, 15(066), 2008.
- 3 H. Bodlaender, F. V. Fomin, D. Lokshtanov, E. Penninkx, S. Saurabh, and D. M. Thilikos. (Meta) Kernelization. In *Proceedings of the 50th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 629–638. IEEE, 2009.
- 4 Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. (meta) Kernelization. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pages 629–638. IEEE Computer Society, 2009.
- 5 Marek Cygan, Fabrizio Grandoni, and Danny Hermelin. Tight kernel bounds for problems on graphs with small degeneracy. In *Proc. of the 21st Annual European Symp. on Algorithms (ESA)*, volume 8125 of *LNCS*, pages 361–372. Springer, 2013.
- 6 Anuj Dawar, Martin Grohe, and Stephan Kreutzer. Locally excluding a minor. In *Proceedings of the 22nd IEEE Symposium on Logic in Computer Science (LICS)*, pages 270–279. IEEE Computer Society, 2007.
- 7 Anuj Dawar and Stephan Kreutzer. Domination problems in nowhere-dense classes. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 4 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 157–168. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2009.
- 8 Rodney G. Downey and Michael R. Fellows. *Parameterized complexity*. Springer-Verlag, New York, 1999.
- 9 Pål Grønås Drange, Markus Dregi, Fedor V. Fomin, Stephan Kreutzer, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, Felix Reidl, Saket Saurabh, Fernando Sánchez Villamil, and Somnath Sikdar. Kernelization and Sparseness: the case of Dominating Set. *CoRR*, abs/1309.4022v1, 2014.
- 10 Z. Dvořák. Constant-factor approximation of the domination number in sparse graphs. *Eur. J. Comb.*, 34(5):833–840, 2013.

- 11 Zdenek Dvořák, Daniel Král', and Robin Thomas. Testing first-order properties for subclasses of sparse graphs. *J. ACM*, 60(5):36, 2013.
- 12 Jörg Flum and Martin Grohe. Fixed-parameter tractability, definability, and model-checking. *SIAM J. Comput.*, 31(1):113–145, 2001.
- 13 Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, 2006.
- 14 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Bidimensionality and kernels. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 503–510. SIAM, 2010.
- 15 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Linear kernels for (Connected) Dominating Set on H -minor-free graphs. In *Proc. of the 22nd Annual ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 82–93. SIAM, 2012.
- 16 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Linear kernels for (Connected) Dominating Set on graphs with excluded topological subgraphs. In *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 20 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 92–103. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2013.
- 17 Markus Frick and Martin Grohe. Deciding first-order properties of locally tree-decomposable structures. *J. ACM*, 48(6):1184–1206, 2001.
- 18 J. Gajarský, P. Hliněný, J. Obdržálek, S. Ordyniak, F. Reidl, P. Rossmanith, F. Sánchez Villaamil, and S. Sikdar. Kernelization using structural parameters on sparse graph classes. In *Proceedings of the 21st Annual European Symposium on Algorithms (ESA)*, volume 8125 of *Lecture Notes in Comput. Sci.*, pages 529–540. Springer, 2013. Full version available at <http://arxiv.org/abs/1302.6863>.
- 19 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 89–98. ACM, 2014.
- 20 Martin Grohe and Dániel Marx. Structure theorem and isomorphism test for graphs with excluded topological subgraphs. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, pages 173–192. ACM, 2012.
- 21 Shai Gutner. Polynomial kernels and faster algorithms for the Dominating Set problem on graphs with an excluded minor. In *Proceedings of the 4th International Workshop on Parameterized and Exact Computation (IWPEC)*, volume 5917 of *Lecture Notes in Comput. Sci.*, pages 246–257. Springer, 2009.
- 22 Stephan Kreutzer. Algorithmic meta-theorems. In J. Esparza, C. Michaux, and C. Steinhorn, editors, *Finite and Algorithmic Model Theory*, chapter 5, pages 177–270. Cambridge University Press, 2011.
- 23 J. Nešetřil and P. Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and Combinatorics*. Springer, 2012.
- 24 Geevarghese Philip, Venkatesh Raman, and Somnath Sikdar. Polynomial kernels for Dominating Set in graphs of bounded degeneracy and beyond. *ACM Transactions on Algorithms*, 9(1):11, 2012. doi:10.1145/2390176.2390187.
- 25 Felix Reidl. *Structural sparseness and complex networks*. Dr., Aachen, Techn. Hochsch., Aachen, 2015. Aachen, Techn. Hochsch., Diss., 2015. URL: <https://publications.rwth-aachen.de/record/565064>.
- 26 Neil Robertson and Paul D. Seymour. Graph Minors. XVI. Excluding a non-planar graph. *J. Comb. Theory, Ser. B*, 89(1):43–76, 2003.
- 27 Detlef Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 6(6):505–526, 1996.