Twin-Width and Polynomial Kernels

Édouard Bonnet
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Eun Jung Kim
Université Paris-Dauphine, PSL University, CNRS UMR7243, LAMSADE, Paris, France

Amadeus Reinald
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Stéphan Thomassé
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Rémi Watrigant
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Abstract

We study the existence of polynomial kernels for parameterized problems without a polynomial kernel on general graphs, when restricted to graphs of bounded twin-width. It was previously observed in [Bonnet et al., ICALP'21] that the problem k-Independent Set allows no polynomial kernel on graph of bounded twin-width by a very simple argument, which extends to several other problems such as k-Independent Dominating Set, k-Path, k-Induced Path, k-Induced Matching. In this work, we examine the k-Dominating Set and variants of k-Vertex Cover for the existence of polynomial kernels.

As a main result, we show that k-Dominating Set does not admit a polynomial kernel on graphs of twin-width at most 4 under a standard complexity-theoretic assumption. The reduction is intricate, especially due to the effort to bring the twin-width down to 4, and it can be tweaked to work for Connected k-Dominating Set and Total k-Dominating Set with a slightly worse bound on the twin-width.

On the positive side, we obtain a simple quadratic vertex kernel for Connected k-Vertex Cover and Capacitated k-Vertex Cover on graphs of bounded twin-width. These kernels rely on that graphs of bounded twin-width have Vapnik-Chervonenkis (VC) density 1, that is, for any vertex set X, the number of distinct neighborhoods in X is at most $c \cdot |X|$, where c is a constant depending only on the twin-width. Interestingly the kernel applies to any graph class of VC density 1, and does not require a witness sequence. We also present a more intricate $O(k^{1.5})$ vertex kernel for Connected k-Vertex Cover.

Finally we show that deciding if a graph has twin-width at most 1 can be done in polynomial time, and observe that most graph optimization/decision problems can be solved in polynomial time on graphs of twin-width at most 1.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis; Theory of computation → Fixed parameter tractability

Keywords and phrases Twin-width, kernelization, lower bounds, Dominating Set

Digital Object Identifier 10.4230/LIPIcs.IPEC.2021.10

Related Version *Full Version*: `https://arxiv.org/abs/2107.02882`

Funding The first, fourth, and fifth authors were supported by the ANR project DIGRAPHS (ANR-19-CE48-0013-01), the first, second, fourth, and fifth authors were supported by the ANR project TWIN-WIDTH (ANR-21-CE48-0014-01), and the second author was supported by the ANR project ASSK (ANR-18-CE40-0025-01) from French National Research Agency.

Acknowledgements We thank Noga Alon and Bart M. P. Jansen for independently asking whether k-Dominating Set admits a polynomial kernel on classes of bounded twin-width, an interesting question that led to our main result.

1 Introduction

The *twin-width* of a graph can be defined in the following way. A *partition sequence* of an n-vertex graph G, is a sequence $\mathcal{P}_n, \ldots, \mathcal{P}_1$ of partitions of its vertex set $V(G)$, such that $\mathcal{P}_n$ is the set of singletons $\{\{v\} : v \in V(G)\}$, $\mathcal{P}_1$ is the singleton set $\{V(G)\}$, and for every $2 \leqslant i \leqslant n$, $\mathcal{P}_{i-1}$ is obtained from $\mathcal{P}_i$ by merging two of its parts into one. Two parts P, P' of a same partition $\mathcal{P}$ of $V(G)$ are said *homogeneous* if either every pair of vertices $u \in P, v \in P'$ are non-adjacent, or every pair of vertices $u \in P, v \in P'$ are adjacent. Finally the twin-width of G is the least integer d such that there is partition sequence $\mathcal{P}_n, \ldots, \mathcal{P}_1$ of G with every part of every $\mathcal{P}_i$ $(1 \leqslant i \leqslant n)$ being homogeneous to every other parts of $\mathcal{P}_i$ but at most d. We call such a partition sequence a *d-sequence*.

On the one hand, a surprisingly wide variety of graphs have low twin-width. Graph classes with bounded twin-width include classes with bounded treewidth, or even rank-width, proper minor-closed classes, every hereditary proper subclass of permutation graphs, bounded-degree string graphs [6], classes with bounded queue or stack number, some expander families [4]. Furthermore on those particular classes, we can find (non necessarily optimum) $O(1)$-sequences in polynomial time. We observe that such an approximation algorithm is still missing in general graphs, but exists for *ordered* binary structures [5].

On the other hand, bounded twin-width classes have interesting algorithmic and structural properties. Remarkably, given a partition sequence witnessing that an n-vertex graph G has twin-width at most d, and a first-order sentence φ, one can decide if φ holds in G in time $f(|\varphi|, d)\, n$ for a computable, but non-elementary, function f [6]. That general framework is called *first-order model checking*, and generalizes problems like k-Independent Set with $\varphi = \exists x_1 \ldots \exists x_k \bigwedge_{1 \leqslant i < j \leqslant k} \neg(x_i = x_j \vee E(x_i, x_j))$ and k-Dominating Set with $\varphi = \exists x_1 \ldots \exists x_k \forall x \bigvee_{1 \leqslant i \leqslant k}(x = x_i \vee E(x, x_i))$. For these two particular problems, though, a much better running time of $2^{O_d(k)}\, n$ is possible [3]. In contrast, an algorithm running in time $f(k) n^{o(k)}$ for either of these problems on *general* graphs, with f being any computable function, would imply the improbable (or at least breakthrough) result that 3-SAT can be solved in subexponential time [10].

Now we know that k-Independent Set and k-Dominating Set are fixed-parameter tractable (FPT), i.e., solvable in time $f(k)\, n^{O(1)}$, on graphs of bounded twin-width given with an $O(1)$-sequence, one can then ask whether polynomial kernels exist. A *kernel* is a polytime algorithm that produces, given an instance of a parameterized problem Π, an equivalent instance of Π (i.e., the output is a YES-instance if and only if the input is a YES-instance) of size only function of the parameter. A *polynomial kernel* is a kernel for which the latter function is polynomial. Any decidable problem with a kernel is FPT, and any FPT problem admits a kernel. However not every FPT problem is believed to have a polynomial kernel. And indeed such an outcome would imply an unlikely collapse of complexity classes.

We already observed that there is a constant d such that k-Independent Set is highly unlikely to have a polynomial kernel on graphs with twin-width at most d [3]. The OR-composition is straightforward from the following facts: (1) cliques have twin-width 0 and planar graphs have bounded twin-width [6], (2) the twin-width of every graph is the maximum twin-width of its modules and quotient graph (see Lemma 8), and (3) Maximum Independent Set is NP-hard in (subcubic) planar graphs [30]. Then one can blow every vertex of a clique K_t into a distinct graph among t planar Maximum Independent Set-instances. Facts (1) and (2) imply that the constructed graph has bounded twin-width, while the correctness of the OR-composition is easy to check. Incidentally the exact same reduction rules a polynomial kernel out for k-Independent Dominating Set. Furthermore

Minimum Independent Dominating Set is NP-hard in grid graphs [11], and Maximum Independent Set is NP-hard in subdivisions of grid graphs (since these coincide with planar graphs of degree at most 4). Since these graphs have twin-width at most 4 (see Lemma 9), no polynomial kernel is likely to exist for both problems (even when a 4-sequence is given in the input). It should be noted that this simple reduction fails for k-Dominating Set: one can dominate the constructed graph by picking only two vertices (from two distinct instances).

The parameterized complexity (FPT algorithms and kernels) of k-Dominating Set[1] on "sparse"[2] classes has a rich and interesting history. Subexponential FPT algorithms with running time $2^{O(\sqrt{k})}\, n^{O(1)}$ are known in planar graphs [28, 21], bounded-genus graphs and more generally classes excluding a fixed minor [19, 25, 33], and an FPT algorithm with running time $2^{O(k)}\, n$ exists in classes excluding a fixed topological minor [2]. On these classes the mere existence of an FPT algorithm (but not the particular, enhanced running time) is subsumed by an algorithmic meta-theorem of Grohe, Kreutzer, and Siebertz [32] that says that first-order model checking is FPT in any nowhere dense class.[3] More general than nowhere dense classes are bounded-degeneracy graphs, or further, $K_{t,t}$-free classes, i.e., excluding the biclique $K_{t,t}$ as a subgraph. Alon and Gutner [2] give an FPT algorithm in d-degenerate graphs running in time $k^{O(dk)}\, n$. And Philip, Raman, and Sikdar [40] extend the fixed-parameter tractability of k-Dominating Set to any $K_{t,t}$-free class (for a fixed t). Telle and Villanger [43] further show that k-Dominating Set on $K_{t,t}$-free graphs is FPT for the combined parameter $k + t$.

In parallel to these algorithms, the existence of polynomial, or even linear, kernels have been thoroughly investigated. In 2004, Alber, Fellows, and Niedermeier [1] presented a linear kernel for k-Dominating Set on planar graphs that triggered a series of works. Linear kernels are known on planar graphs [1, 9], bounded-genus graphs [27], apex-minor-free graphs [26], but more generally in any class excluding a fixed topological minor [25]. k-Dominating Set admits a polynomial kernel on graphs of girth 5 (that is, excluding the triangle and the biclique $K_{2,2}$ as a subgraph) [42]. A polynomial kernel of size $O(k^{(t+1)^2})$ is obtained for $K_{t,t}$-free graphs [40], the most general "sparse" class. Contrary to the FPT algorithm, a polynomial kernel in the parameter $k + t$ is highly unlikely [20]. More precisely, for any $\varepsilon > 0$, a kernel of size $k^{(t-1)(t-3)-\varepsilon}$ would imply that coNP $\subseteq$ NP/poly [15]. On classes of bounded expansion[4] k-Dominating Set has a linear kernel, while the seemingly closely related Connected k-Dominating Set has no polynomial kernel [23]. The latter result refines a reduction showing the same lower bound on 2-degenerate graphs [17].

Beyond sparse classes, for which most answers turn out positive, the parameterized complexity of k-Dominating Set seems to conceal many surprises, some of which recently unraveled. We already mentioned that k-Dominating Set is FPT on bounded twin-width graphs given with an $O(1)$-sequence. Let us also mention that the same problem is actually W[1]-hard (hence unlikely FPT) on circle graphs [7]. This is somewhat unexpected since Dominating Set is polytime solvable on permutation graphs [24], a large subclass of circle graphs. On the positive side, k-Dominating Set admits a polynomial kernel on so-called *c-closed graphs* [36], a far-reaching dense generalization of bounded d-degenerate graphs.

1 All the subsequent results also hold for k-Independent Set.
2 *Sparse* is an overloaded term; here we use it as *not containing arbitrarily large bicliques as subgraphs*.
3 The definition of nowhere denseness being technical and unnecessary to the current paper, we refer the interested reader to [39]. Let us just mention that bounded-degree graphs, planar graphs, and proper (topological) minor-closed classes are all nowhere dense.
4 We will not need a definition of *expansion* here. Bounded expansion classes are more general than topological-minor-free classes and less general than nowhere dense classes.

Our results

We are back to wondering whether k-Dominating Set admits a polynomial kernel on graphs given with an $O(1)$-sequence. On the one hand, a polynomial kernel would “fit all the data points” considering that the examples of bounded twin-width classes previously given are either $K_{t,t}$-free (and one concludes with [40]) or are dense classes on which Minimum Dominating Set is polytime solvable, like bounded rank-width graphs [12], and (subclasses of) permutation graphs [24]. On the other hand, the same could be said of k-Independent Set for which we already ruled out such a kernel. Yet we will see in Section 3.1 that the above OR-composition not working for k-Dominating Set is part of a more general obstacle toward establishing its incompressibility. In the same section we lay down our plan to overcome that obstacle and show the following. Recall that the input to k-Dominating Set is a graph G and an integer k, k construed as the parameter, and the task is to decide whether G has a dominating set of size at most k.

▶ **Theorem 1.** *Unless* coNP $\subseteq$ NP*/poly, k-Dominating Set on graphs of twin-width at most 4 does not admit a polynomial kernel, even if a 4-sequence of the graph is given.*

We mentioned that the same statement holds much more directly for k-Independent Set and k-Independent Dominating Set. With analogous arguments, we can add k-Path, k-Induced Path, k-Induced Matching to the list. Local gadget modifications of the proof of Theorem 1 yield the same kernel lower bound for variants of k-Dominating Set such as Connected k-Dominating Set and Total k-Dominating Set, on graphs of bounded twin-width. More work would be necessary to get the lower bound for twin-width at most 4.

On the positive side, Connected k-Vertex Cover[5] and Capacitated k-Vertex Cover[6] admit polynomial kernels on graphs of bounded twin-width, while such kernels are unlikely on general graphs [20]. Interestingly, our kernelization algorithm does not require an $O(1)$-sequence.

▶ **Theorem 2.** *Connected k-Vertex Cover and Capacitated k-Vertex Cover admit a kernel with $O(k^2)$ vertices on any class of bounded twin-width.*

A linear kernel (in the number of vertices) is known for apex-minor-free classes [26] via the generic framework of bidimensionality, and even for topological-minor-free classes [35]. Another powerful meta-theorem by Gajarský et al. [29] says that every problem with the so-called *finite integer index* (intuitively, that its boundaried graphs provide finitely many distinct contexts) has a linear kernel on bounded expansion classes when parameterized by the vertex cover number (and more generally by the size of a smallest vertex subset whose deletion leaves the graph with bounded treedepth). In particular this yields a linear kernel for Connected k-Vertex Cover, further extending the two previous results. Besides Connected k-Vertex Cover has a polynomial kernel on $K_{t,t}$-free graphs [17].

---

[5] The problem Connected k-VC takes as input a graph G and a parameter k, and asks whether G has a vertex cover of size at most k which induces a connected subgraph of G.

[6] Given a graph G and a capacity function $c : V(G) \to N$, a *capacitated vertex cover* X of G is a vertex cover of G which admits a mapping $\rho : E(G) \to X$ assigning to each vertex $x \in X$ no more edges than its capacity, i.e., $|\rho^{-1}(x)| \leqslant c(x)$ for every $x \in X$. The goal of Capacitated k-VC is to decide, given a graph G with a capacity function $c : V(G) \to N$ and an integer k as the parameter, if G admit a capacitated vertex cover X of size at most k.

Theorem 2 is based on the following useful lemma stating that, in graphs of bounded twin-width, the number of distinct neighborhood traces inside a subset of vertices is at most linear in the size of the subset.

▶ **Lemma 3.** *There is a function f such that for every graph G of twin-width d and $X \subseteq V(G)$, the number of distinct neighborhoods in X, $|\{N(v) \cap X \ : \ v \in V(G)\}|$, is at most $f(d)|X|$.*

A more compact rewording, using the language of Vapnik-Chervonenkis parameters, is that the neighborhood set-system of graphs of bounded twin-width has VC density 1. By extension, we will say that a graph class has VC density at most 1, if its neighborhood hypergraphs do. That bounded twin-width classes have VC density 1 is an interesting property, that is shared with classes of bounded expansion. For example it implies a constant-factor approximation for Min Dominating Set (obtained in a rather different manner in [3]) via small ε-nets [8]. Lemma 3 was independently obtained by Wojciech Przybyszewski in his master thesis [41].

For Connected k-Vertex Cover, an improved kernel can be obtained with a more elaborate argument.

▶ **Theorem 4.** *Connected k-Vertex Cover admits a kernel with $O(k^{1.5})$ vertices on classes with VC density at most 1.*

Table 1 Kernelization results for arguably the three main problems without a polynomial kernel in general graphs, but an interesting story in sparse classes. PK stands for *polynomial kernel*, LK for *linear kernel* (in the number of vertices). The indicated lack of a kernel is under the assumption that coNP ⊆ NP/poly. Our new results are in bold (the results without a reference nor in bold are consequences of results in bold).

	k-Dominating Set	Connected k-DS	Connected k-VC
general	W[2]-complete [22]	W[2]-complete [22]	FPT [13], no PK [20]
bounded expansion	LK [23]	FPT [18], no PK [23]	LK [29]
bounded biclique	PK [40]	FPT [43], no PK [17]	PK, no LK [17]
bounded degeneracy	PK [40]	FPT [31], no PK [17]	PK, no LK [17, 15]
$K_{1,3}$-free	PK [34]	FPT, no PK [34]	LK (trivial)
$K_{1,4}$-free	W[2]-complete [16]	W[2]-complete [16]	LK (trivial)
bounded twin-width	FPT [6], no PK	FPT [6], **no PK**	$O(k^{1.5})$-vertex kernel
twin-width at most 4	FPT [6], **no PK**	FPT [6]	$O(k^{1.5})$-vertex kernel
twin-width at most 1	**in P**	**in P**	**in P**
VC density at most 1	no PK	no PK	$\mathbf{O(k^{1.5})}$**-vertex kernel**

Finally we extend cograph recognizability (cographs are exactly the graphs with twin-width 0) and prove:

▶ **Theorem 5.** *One can decide in polynomial time if a graph has twin-width at most 1.*

In case the input graph has indeed twin-width at most 1, a 1-sequence is found in polynomial time. Furthermore we observe that a wide class of graph problems is efficiently solvable on inputs of twin-width at most 1. See Table 1 for a summary of most of our results, together with the relevant pointers on other graph classes.

2 Preliminaries

We make this section light to keep the extended abstract legible; a complete preliminary section can be found in the full version. We denote by $[i, j]$ the set of integers $\{i, i+1, \ldots, j-1, j\}$, and by $[i]$ the set of integers $[1, i]$. If $\mathcal{X}$ is a set of sets, we denote by $\cup\mathcal{X}$ their union. The notation $O_d(\cdot)$ gives an asymptotic behavior when d is seen as a constant.

An injective mapping $\eta : V(H) \to V(G)$ *witnesses* that H is a subgraph of G if $uv \in E(H)$ implies $\eta(u)\eta(v) \in E(G)$. A bijective mapping $\eta : V(H) \to V(G)$ *witnesses* that H is a *spanning* subgraph of G if $uv \in E(H)$ implies $\eta(u)\eta(v) \in E(G)$.

The *strict half-graph of height t* is (up to isomorphism) the graph with vertex set $\{a_1, \ldots, a_t, b_1, \ldots, b_t\}$ and edge set $\{a_i b_j \ : \ i < j, i \in [t], j \in [t]\}$. One can see $\{a_1, \ldots, a_t\}$ *oriented toward* $\{b_1, \ldots, b_t\}$ in their realization of the relation $<$ over the indices. The *ℓ-cycle of strict half-graphs of height t* is (up to isomorphism) the graph with vertex set $\{a_1^p, \ldots, a_t^p \ : \ p \in [0, \ell-1]\}$ and edge set $\{a_i^p a_j^{p+1 \mod \ell} \ : \ i < j, i \in [t], j \in [t], p \in [0, \ell-1]\}$. Informally it is the graph obtained from an ℓ-vertex cycle by replacing every edge by a strict half-graph of height t with a consistent, say, clock-wise orientation. See Figure 2 for an example of a 5-cycle of strict half-graphs of height 6 (realized by the black edges on the rounded black boxes). A *strict half-graph* is, for some natural t, the strict half-graph of height t. A *cycle of strict half-graphs* is, for some natural ℓ, the ℓ-cycle of strict half-graphs of same height.

The *$n \times m$ grid* is the graph with vertex set $[n] \times [m]$ and edges between any pair of vertices $(x, y), (x+1, y)$ or $(x, y), (x, y+1)$. A *grid* is an $n \times m$ grid for some integer n and m. A *grid graph* is an induced subgraph of a grid. To insist that we consider a grid and not a mere grid graph, we may use the term *complete grid*.

The *neighborhood hypergraph* of a graph G has vertex set $V(G)$ and edge set $\{N(v) \ : \ v \in V(G)\}$. A family of hypergraphs $\mathcal{H}$ has *Vapnik-Chervonenkis (VC) density at most 1* if there is a constant c such that for every hypergraph $H \in \mathcal{H}$ and every $X \subseteq V(H)$, $|\{X \cap e \ : \ e \in E(H)\}| \leqslant c \cdot |X|$.

2.1 Contraction sequences and twin-width

A *trigraph* G has vertex set $V(G)$, black edge set $E(G)$, and red edge set $R(G)$, with $E(G)$ and $R(G)$ disjoint. The *total graph* of trigraph G is the graph G' with $V(G') = V(G)$ and $E(G') = E(G) \cup R(G)$. The *subtrigraph* of G induced by S is the trigraph H with $V(H) = S$, $E(H) = E(G) \cap \binom{S}{2}$, and $R(H) = R(G) \cap \binom{S}{2}$. H is then called an *induced subtrigraph* of G.

The *set of neighbors* $N_G(v)$ of a vertex v in a trigraph G consists of all the vertices adjacent to v by a black or red edge. A d-trigraph is a trigraph G such that the *red graph* $(V(G), R(G))$ has degree at most d. In that case, we also say that the trigraph has *red degree* at most d. A *contraction* or *identification* in a trigraph G consists of merging two (non-necessarily adjacent) vertices u and v into a single vertex z, and updating the edges of G in the following way. Every vertex of the symmetric difference $N_G(u) \triangle N_G(v)$ is linked to z by a red edge. Every vertex x of the intersection $N_G(u) \cap N_G(v)$ is linked to z by a black edge if both $ux \in E(G)$ and $vx \in E(G)$, and by a red edge otherwise. The rest of the edges (not incident to u or v) remain unchanged. See Figure 1 for an illustration.

A *d-sequence* (or *contraction sequence*) is a sequence of d-trigraphs $G_n, G_{n-1}, \ldots, G_1$, where $G_n = G$, $G_1 = K_1$ is the graph on a single vertex, and G_{i-1} is obtained from G_i by performing a single contraction of two (non-necessarily adjacent) vertices. We observe that G_i has precisely i vertices, for every $i \in [n]$. The twin-width of G, denoted by $tww(G)$, is the minimum integer d such that G admits a d-sequence. Note that, in what precedes, the

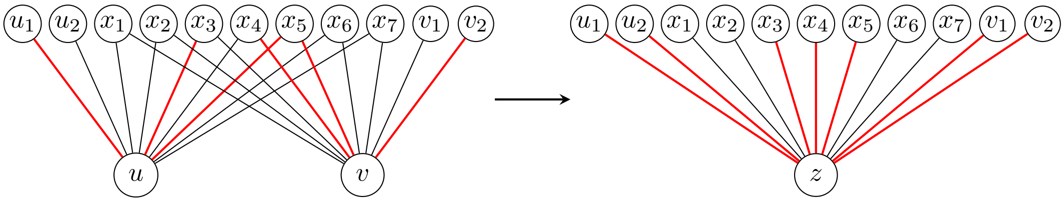

Figure 1 Contraction of vertices u and v, and how the edges of the trigraph are updated.

initial structure $G_n = G$ may be a trigraph instead of a graph. Thus we defined twin-width more generally for trigraphs. Similarly a *partial d-sequence* from a n-vertex trigraph G to an i-vertex trigraph H is a sequence of d-trigraphs $G = G_n, G_{n-1}, \dots, G_i = H$. Observe that if G has a partial d-sequence to H, and H has itself a d-sequence, then the concatenation of these sequences is a d-sequence for G.

Here are useful facts about twin-width, whose proofs are trivial or can be found in the full version.

▶ **Observation 6.** *Let G be a trigraph and H be an induced subtrigraph of G. Then, $tww(H) \leqslant tww(G)$.*

▶ **Observation 7.** *Let G, G' be two trigraphs such that $V(G) = V(G')$, $R(G) \subseteq R(G')$, $E(G') \subseteq E(G)$, and $R(G) \cup E(G) \subseteq R(G') \cup E(G')$. Then $tww(G) \leqslant tww(G')$.*

▶ **Lemma 8.** *Let G be a graph and $\mathcal{H} = \{H_1, H_2, \dots, H_\ell\}$ be its modular partition. Then,*

$$tww(G) = \max\{\max_{i \in [\ell]} tww(H_i), tww(G/\mathcal{H})\}.$$

▶ **Lemma 9.** *Any trigraph whose total graph is a subdivision of a subgraph of a grid has twin-width at most 4.*

2.2 Kernels or lack thereof

For a parameterized problem $\mathcal{Q}$, a *kernel* of size bounded by a function f is a polynomial-time reduction $\rho : \Sigma^* \times N \to \Sigma^* \times N$ such that $(x, k) \in \mathcal{Q}$ if and only if $\rho(x, k) \in \mathcal{Q}$, and $|\rho(x, k)| \leqslant f(k)$. A kernel is said *linear*, *quadratic*, or *polynomial*, if the function f can be chosen linear, quadratic, or polynomial, respectively. We recall the framework of OR-cross-compositions [14], which we will rely on to show the absence of a polynomial kernel in Theorem 1.

▶ **Definition 10.** *A polynomial equivalence relation on Σ^* is an equivalence relation $\mathcal{R}$ when*
(i) *for $x, y \in \Sigma^*$, the equivalence $x\mathcal{R}y$ can be decided in time polynomial in $|x| + |y|$, and*
(ii) *$\mathcal{R}$ restricted to instances of size at most n admits polynomially many equivalence classes.*

We can now formally define an *OR-cross-composition.*

▶ **Definition 11.** *Let $\mathcal{L}$ be a language, $\mathcal{R}$ a polynomial equivalence relation on Σ^* and $\mathcal{Q}$ a parameterized problem. An* OR-cross-composition *from $\mathcal{L}$ to $\mathcal{Q}$ with respect to $\mathcal{R}$ is an algorithm taking as input t $\mathcal{R}$-equivalent instances $x_1, \dots, x_t \in \Sigma^*$, running in time polynomial in $\sum_{j=1}^{t} |x_j|$, and outputting an instance $(y, N) \in \Sigma \times N$ such that:*
(i) *N is polynomially bounded in $\max_{j \in [t]} |x_j| + \log t$,*
(ii) *$(y, N) \in \mathcal{Q}$ if and only if there exists some j such that $x_j \in \mathcal{L}$.*

We say that $\mathcal{L}$ cross-composes into $\mathcal{Q}$, and we sometimes refer to output instance (y, N) as the *composed instance*. The following result provides the lower bound under coNP $\subseteq$ NP/poly.

▶ **Theorem 12** ([14]). *If an NP-hard language $\mathcal{L}$ admits an OR-cross-composition into a parameterized problem $\mathcal{Q}$, then $\mathcal{Q}$ does not admit a polynomial kernel unless* coNP $\subseteq$ NP*/poly.*

2.3 Selected results for the short version of the paper

Due to space constraints, most proofs of the results announced in the introduction were omitted, and can be found in the long version of the paper. Section 3 is devoted to our main result: the kernel lower bound for k-Dominating Set in graphs of twin-width at most 4. We provide the intuition behind it, the definition of the particular instances we are reducing from, and finally give the construction as well as a sketch of how we can reach the twin-width bound. In Section 4, we present a sketch of proof that bounded twin-width graphs have VC density 1, and use it to get a kernel with $O(k^{1.5})$ vertices for Connected k-Vertex Cover.

3 Kernel lower bound for k-Dominating Set

3.1 Outline

Let us start explaining why we should *not* expect a simple OR-composition. A standard way to OR-compose t Dominating Set-instances is to have for each instance a "switch", that is, one vertex dominating all but one instance. Then picking the corresponding vertex in the solution, one is left with dominating one chosen instance with a given remaining budget. This is precisely what we want, but how to ensure that one does not activate two switches?

As we previously observed [3], one can use larger weights for the switches. However removing the vertex-weights cannot be done without increasing the twin-width. Another possibility is to force all the budget but one unit (for the switch) within the instances. This requires, say, k vertices called "forcers", each adjacent to a k-th fraction of each instance. Now consider the induced subgraph made by these k vertices, the t switches, and tk vertices of the instances realizing the tk possible neighborhoods toward the former $t + k$ vertices. The two neighborhoods of every pair of vertices in this graph has a large symmetric difference. Thus in particular the overall graph has unbounded twin-width. (Finally known tricks to condense the t switches into $O(\log t)$ vertices do not help, since we want the twin-width to be bounded by an absolute constant.)

So we need a more elaborate way of selecting one instance among t; one, thought primarily to keep the twin-width low. In the previous attempts, the twin-width was increasing too much because of attachments –switches and forcers– external to the instances. We will therefore have instances themselves play these roles. Say that each instance comes with a partition of its vertex set into N parts, each of which containing a vertex solely adjacent to vertices in its part. We place the t instances in a $t \times N$ two-dimensional layout, where each instance occupies a "row," while the j-th part of all the instances form the j-th "column." The switch mechanism is as follows. Every vertex in the j-th part of the i-th instance –say I_i– dominates the $j - 1$-st part of the instances with a smaller index, and the $j + 1$-st part of the instances with a larger index. In other words, we put a strict half-graph over the parts of two consecutive columns. This is done cylindrically, see Figure 2.

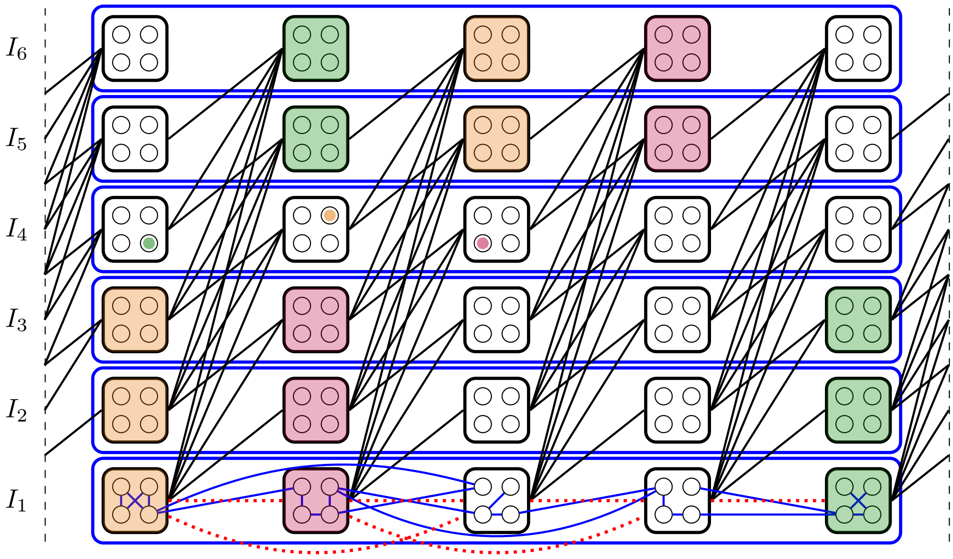

Figure 2 The overall picture. Instances $I_1, \ldots, I_t$ (here with $t = 6$) are in rows, boxed in blue, with their edge also in blue. For the sake of legibility, we only represented the edges of I_1. The red dotted edges are the red edges appearing after contracting every part (boxed in black) into a single vertex. Example of what three vertices picked in the first three parts of I_4 dominates in the other instances. Continuing consistently in I_4 would result in "switching off" all the other instances, while deviating would leave at least one part "white" and not intersected, thus one vertex not dominated.

With that mechanism, a dominating set of a fixed instance I_i (intersecting each of its parts once) is a dominating set of the overall graph. We skip here the details of the reverse direction, but the use of half-graphs and of vertices whose neighborhood in their instance is confined to their own part (the last ingredient is to have a dummy, edgeless top instance I_t) should give a feel for why no other kind of dominating sets of size N can exist.

What about the twin-width bound? Cycles of half-graphs have bounded twin-width. So a natural first step is to contract every part of every instance into a single vertex. Doing so will create some red edges within each row. To ensure that the red degree remains bounded in this first step, a part should be partially adjacent to only a bounded number of other parts. In the second step, we contract the cycle of half-graphs row by row. Thus the red edges of the different instances will progressively stack up. We need to control the accretion with the red edges of each instance mapping onto a common bounded-degree red graph. Finally in the third step, we contract the residual red graph. It should be itself of bounded twin-width, for instance by being planar.

In the next subsection, we show that Minimum Dominating Set remains NP-hard even when inputs are equipped with a vertex-partition satisfying all the properties that we came across in this outline.

3.2 Tailored NP-hardness for Dominating Set

We present a new hardness reduction for Dominating Set. The reduction is designed so as to produce carefully tamed instances, even when compared to existing NP-hardness reductions of Dominating Set (including those on planar instances of bounded degree), and this will be crucial for the subsequent OR-cross-composition, as hinted in the previous section.

The next theorem involves what we will call *snaking grids*. See Figure 3 for an illustration of the 5×10 *snaking grid*, which has $(3 \cdot (5-1)+1)(3 \cdot (10-1)+1)$ vertices. One may observe that the snaking grids are subdivisions of a wall with some extra isolated vertices. We will prefer to think of the snaking grid as a spanning subgraph of a (complete) grid, hence the particular embedding of the figure. The snaking grid is useful as it allows us to superpose a canonical hamiltonian cycle such that the maximum degree remains 3, and thus to bound the twin-width of the composed instance.

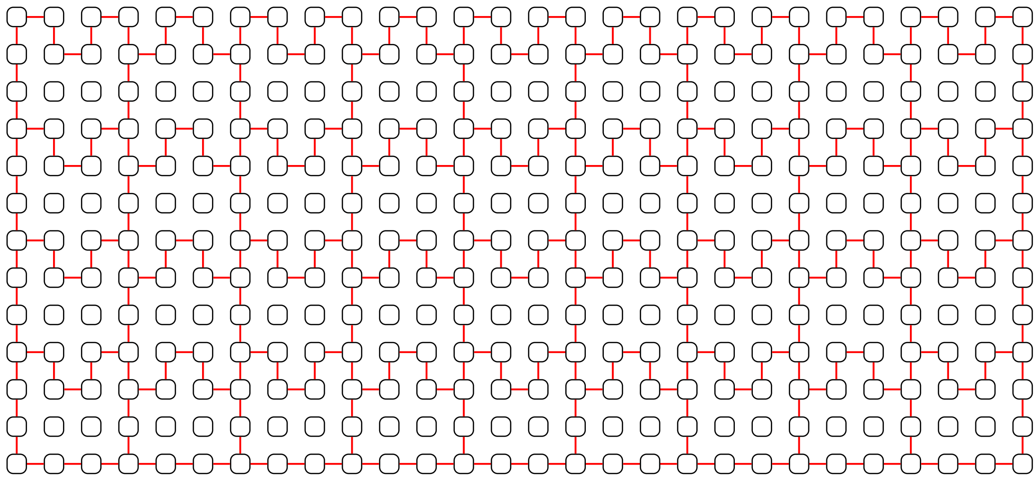

Figure 3 The 5×10 snaking grid.

The following result is obtained through a reduction from the NP-hard problem Planar 3-SAT [37].

▶ **Theorem 13.** *Dominating Set remains* NP*-hard when its input* (G, N) *comes with a vertex-partition* $\mathcal{B} = \{B_1, \dots, B_N\}$*, two positive integers* s *and* t*, and a bijective mapping* η *from* $\{B_1, \dots, B_N\}$ *to the vertex set of the* $s \times t$ *snaking grid such that:*

(i) *G has a partial 4-sequence to the quotient trigraph $G/\mathcal{B}$,*

(ii) *$G/\mathcal{B}$ is a spanning subgraph of the $s \times t$ snaking grid, with t even, witnessed by η, and*

(iii) *every dominating set of G intersects each B_i, for $i \in [N]$.*

3.3 The construction

We now describe the cross-composition from the NP-hard Dominating Set restricted as in Theorem 13 to k-Dominating Set. We use the polynomial equivalence $\mathcal{R}$ to partition all well-formed instances for the restricted k-Dominating Set (those satisfying Theorem 13) with respect to the given parameter N and dimensions (p, q) of the corresponding snaking grid. For any two well-formed instances $(I_i, N_i, \mathcal{B}_i, p_i, q_i, \eta_i), (I_\ell, N_\ell, \mathcal{B}_\ell, p_\ell, q_\ell, \eta_\ell)$, we can check in polynomial time that $N_i = N_\ell$ and $(p_i, q_i) = (p_\ell, q_\ell)$, yielding the polynomial equivalence relation. All ill-formed instances form a single class.

Consider t well-formed instances of the restricted Dominating Set which are taken from an equivalence class with respect to $\mathcal{R}$, and we may consider these instances as $(I_i, N, \mathcal{B}_i, p, q, \eta_i)_{i \in [t]}$. We will construct a k-Dominating Set instance (H, N), with the same parameter, admitting a solution if and only if at least one input instance $(I_i, N, \mathcal{B}_i, p, q, \eta_i)$ admits a solution for the restricted k-Dominating Set. Before composing the input graphs, we introduce a dummy instance in the form of graph I_{t+1} serving to ensure that any valid

(H, N) further admits a solution picking vertices in each column. I_{t+1} is an independent set of size $2N$ on which we partition $V(I_{t+1})$ through $\mathcal{B}_{i+1}$ into N classes of exactly two vertices. Note that since $I_{t+1}/\mathcal{B}_{t+1}$ is an independent set, it is a spanning subgraph of the $p \times q$ snaking grid as witnessed by any bijective η_{t+1} onto the latter.

We first show how to order the partition classes of each instance in the same way with respect to their mapping onto the snaking grid. This ordering will follow a fictitious hamiltonian cycle $(y_1, \ldots, y_N)$ on the $p \times q$ snaking grid, in the way depicted as the darker red cycles in Figure 4. Referring to the partition of instance $i \in [t+1]$ as $\mathcal{B}_i = \{B_{i,1}, \ldots, B_{i,N}\}$, we can assume up to the reordering above that $\eta_i(B_{i,j}) = y_j$.

Now, considering all instances over H, a representation of the construction that follows is given in Figure 2. It will be useful to consider the instances in a grid such that $B_{i,j}$ is the cell in the i-th row and j-th column, and we will use the term partition class or cell interchangeably. We can then see instance I_i as row i, and define regular instance columns, omitting the dummy instance, as $C_j = \bigcup_{i \in [t]} B_{i,j}$ for $j \in [N]$.

Construction. We start building our composed graph H as the union of all instances $(I_i)_{i \in [t+1]}$, that is, $V(H) = \bigcup_{i \in [t+1]} V(I_i)$ and $E(I_i) \subseteq E(H)$ for $i \in [t+1]$. Then, our cross-composition proceeds by adding a cycle of strict half-graphs over columns $(C_j)_{j \in [N+1]}$: for $i \in [t+1], j \in [N]$, $B_{i,j}$ forms a biclique with $\bigcup_{i < \ell \leqslant t+1} B_{\ell,j+1}$ (accounting for indices j modulo N). Notice then that the only edges added above lie between columns $C_j, C_{j'}$ with $|j' - j| = 1$, so any edge between two columns differing by at least two indices is an edge of I_i. Each instance class $B_{t+1,j}$ is then adjacent exactly to C_{j-1}. Having ordered the classes of each instance in the same way with respect to their mapping onto the snaking grid, column C_j consists of homologous vertices, all in the same position on their respective grids, see Figure 4. Then, the cycle of half-graphs follows the darker red fictitious hamiltonian cycles mapping to $(y_1, \ldots, y_N)$. The following lemma allows to conclude the proof of Theorem 1.

▶ **Lemma 14.** *The composed instance has twin-width at most 4, and is positive if and only if one of the input instance is positive.*

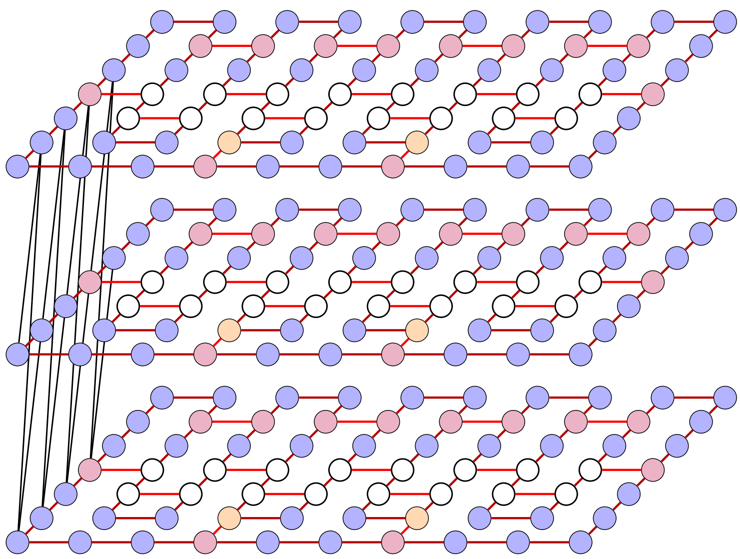

Figure 4 The different layers (instances) linked by the cycle of half-graphs. Only four half-graphs are drawn for the sake of legibility.

How to get twin-width 4, sketch. Let us sketch how to bound the twin-width of the composed graph H. One first contract each $B_{i,j}$ and obtain a trigraph as in Figure 4. It can be shown that the composed graph H admits a partial 4-contraction sequence to $H/\left(\bigcup_{i\in[t+1]} \mathcal{B}_i\right)$. Now, for the purpose of bounding the twin-width of the composed graph, it is useful to note that Observation 7 allows us to add red edges. At this point, the only vertices of large degree stem from the strict half-graphs. We keep those edges black, and we now turn each instance into a red augmented (i.e. additional edges) snaking grid as follows. Since each instance is a spanning subgraph of the $p \times q$ snaking grid, we can first assume that it is a (fully) red snaking grid. Then, the red augmented snaking grid is built by further adding red cycle $(B_{i,1}, ..., B_{i,N})$. By our choice of ordering in the composition, this cycle is the same on every instance with respect to their mapping on the $p \times q$ snaking grid.

We can now describe the contraction of $t + 1$ red augmented snaking grids $(I_i)_{i\in[t+1]}$, abusing notation for the now quotiented instances, with the black edges of our composition. We will exhibit a partial 4-contraction sequence eventually contracting every column, now consisting of $t + 1$ homologous vertices, that is, all vertices at the same position on their respective snaking grid into a single one. The proof will proceed by induction on the number of augmented snaking grids, our hypothesis at step t being that there exists a partial 4-contraction sequence from t augmented snaking grids to a single one, accounting for the black edges added in the composition. This being true for $t = 1$, assume the result holds for some t and let us consider case $t + 1$. We will deal with the two bottommost augmented snaking grids I_1, I_2 in the half-graphs, contracting pairs of homologous vertices, corresponding to the quotiented $(B_{1,j}, B_{2,j})$ thanks to the ordering chosen in the composition.

We argue that the contraction of the first two red snaking grids can be done while bounding the red degree by four. Since the only contracted pairs were homologous, this results in a red augmented snaking grid with no red edges towards grids $i > 2$. The remaining edges of the strict half-graph cycle still form one of height t, which is exactly the induction case for t and achieves to prove the induction. Therefore, there is a partial 4-contraction sequence from our composed graph into a red augmented snaking grid. Then, as the latter is a subgraph of the red complete grid, Lemma 9 yields twin-width at most 4.

4 Polynomial kernels

Let us prove Lemma 3 which, we repeat, is equivalent to saying that the neighborhood hypergraphs of graphs of bounded twin-width have VC density 1. This feature is shared with classes of bounded expansion. Lemma 3 is of independent interest as it opens the door to a common algorithmic treatment for classes of bounded twin-width and of bounded expansion.

We need to introduce some vocabulary on 0, 1-matrices. A *row* (resp. *column*) *division* is a row (resp. column) partition where every part is a consecutive set of rows (resp. columns). A *cell* or *zone* of a matrix M with row and columns divisions $(\mathcal{R}, \mathcal{C})$ is a submatrix $M[R_i, C_j]$ with $R_i \in \mathcal{R}$ and $C_j \in \mathcal{C}$. A *t-division* is a division $(\mathcal{R}, \mathcal{C})$ with $|\mathcal{R}| = |\mathcal{C}| = t$. A matrix is *mixed* if it has at least two distinct rows and at least two distinct columns. A *corner* is a 2×2 contiguous submatrix which is mixed.

▶ **Lemma 3.** *For every graph G of twin-width t and $X \subseteq V(G)$, the number of distinct neighborhoods in X, $|\{N(v) \cap X \ : \ v \in V(G)\}|$, is at most $2^{4c_{2t+2}}|X|$ for some constant c_t depending only on t.*

Proof sketch. We assume for the sake of contradiction that $|\{N(v) \cap X \ : \ v \in V(G)\}| > 2^{4c_{2t+2}}|X|$. For every vertex ordering of G, its adjacency matrix M along this order contains an $|X| \times 2^{4c_{2t+2}}$ submatrix without two equal columns; namely the submatrix of the adjacencies between X and $2^{4c_{2t+2}}$ vertices with a pairwise distinct neighborhood in X. It can be shown that M has a $2t+2$-mixed minor in the following way. Let $s := |X|$ and create a column division $C_1, \ldots, C_s$ into s column parts, each consisting of $2^{4c_{2t+2}}$ consecutive columns. Note that the submatrix of M consisting of the part C_i has rank at least $4c_{2t+2}$ in the binary field F_2, for all $i \in [s]$. Therefore the rows of $M[R, C_i]$ allow a row division into at least $2c_{2t+2}$ parts so that each cell has rank at least 2, thus is mixed, and, by an observation in [6], contains a corner. Let us consider $\mathcal{R}^1$ the row division grouping each pair of rows with indices $2i-1, 2i$, and $\mathcal{R}^2$, grouping each pair of rows $2i, 2i+1$, for $i \in [\lceil s/2 \rceil]$. Observe that one of the two divisions $(\mathcal{R}^1, \{C_i\})$, $(\mathcal{R}^2, \{C_i\})$ of $M[R, C_i]$ contains at least c_{2t+2} zones with a corner, hence mixed.

Without loss of generality, we may assume that at least $\lceil s/2 \rceil$ column parts among $C_1, \ldots, C_s$ have at least c_{2t+2} mixed zones when divided by, say, $\mathcal{R}^1$. Consider the column division $\mathcal{C}' = \{C'_1, \ldots, C'_{s'}\}$ with $s' \geqslant \lceil s/2 \rceil$, coarsening of $\{C_1, \ldots, C_s\}$ such that each part C'_j contains exactly one column part C_i with the property of the previous sentence. Now the number of mixed zones in the division $(\mathcal{R}^1, \mathcal{C}')$ is at least $c_{2t+2} \cdot s'$, and with the correct choice of the constant c_{2t+2}, the celebrated Marcus-Tardos theorem [38] concludes that the division can be further coarsened into $2t+2$-division, where each cell contains a corner. By Grid Minor Theorem of [6], this in turn implies that G has twin-width more than t. ◀

Using the previous lemma, we present here a kernelization algorithm for Connected k-Vertex Cover on bounded twin-width graphs which leads to an instance on $O(k^{1.5})$ vertices. Let X be a vertex cover of G, and let X^b (resp. X^s) be the subsets of X containing all vertices of X with at least $k+1$, respectively at most k, neighbors in $V(G) \setminus X$. Let $Y_1, \ldots, Y_q$ be the partition of $V(G) \setminus X$ into maximal modules. For each $i \in [q]$, let X_i be the neighbors of Y_i in X^s. We use one reduction rule, for which the proof of safeness can be found in the full version.

▶ **Reduction Rule 1.** If there is $i \in [q]$ with $X_i \neq \emptyset$ and $|Y_i| \geqslant |X_i| + 2$, delete a vertex of Y_i.

▶ **Proposition 15.** *Connected k-Vertex Cover admits a kernel on $O_t(k^{1.5})$ vertices when the input graphs have twin-width at most t.*

Proof. Let (G, k) be the input instance of Connected k-Vertex Cover. We can safely remove any isolated vertex, and assume that G is connected (otherwise it is a NO-instance). With a 2-approximation algorithm for Vertex Cover, one can find a vertex cover X of G and assume that $|X| \leqslant 2k$. Indeed if this is not the case, we can correctly output a trivial no-instance because G does not admit a connected vertex cover of size at most k.

Note that Reduction Rule 1 does not disconnect the given graph as we remove a vertex only when it has a twin. Let (G', k) be an instance obtained by exhaustively applying Reduction Rule 1 with the vertex cover X at hand. We classify X into X^b and X^s as before, and $Y_1, \ldots, Y_q$ denote the partition of $Y := V(G') \setminus X$ into maximal modules. For each $i \in [q]$, X_i is the neighbors of Y_i in X^s. By Lemma 3, we have $q \leqslant f(t) \cdot 2k$. Because the edge set between X^s and Y is decomposed into the edge sets of complete bipartite graphs on (Y_i, X_i) over $i \in [q]$, the number of edges between X^s and Y is at least

$$\sum_{i=1}^{q} |Y_i| \cdot |X_i| \geqslant \sum_{i=1}^{q} (|Y_i| - 1)^2 \geqslant \frac{1}{q} \cdot \left(\sum_{i=1}^{q} (|Y_i| - 1) \right)^2 \geqslant \frac{1}{q} \cdot (|Y| - q)^2.$$

Suppose that $|Y| - q > 2 \cdot f(t)^{0.5} \cdot k^{1.5}$. Now,

$$\frac{1}{q} \cdot (|Y| - q)^2 > \frac{4 \cdot f(t) \cdot k^3}{2 \cdot f(t) \cdot k} = 2k^2,$$

and hence there are more than $2k^2$ edges between X^s and Y. With $|X^s| \leqslant 2k$, this implies that there exists a vertex in X^s which has more than k neighbors in Y, contradicting the definition of X^s. To conclude, the number of vertices of G' is at most

$$|X| + |Y| \leqslant 2k + 2f(t)k^{1.5} + q \leqslant 2k + 2f(t)k^{1.5} + 2f(t)k = O_t(k^{1.5})$$ ◀

Note that the proof of Proposition 15 only uses the fact that the input graphs have VC density at most 1, so we in fact established Theorem 4.

References

1 Jochen Alber, Michael R. Fellows, and Rolf Niedermeier. Polynomial-time data reduction for dominating set. *J. ACM*, 51(3):363–384, 2004. doi:10.1145/990308.990309.

2 Noga Alon and Shai Gutner. Linear time algorithms for finding a dominating set of fixed size in degenerated graphs. *Algorithmica*, 54(4):544–556, 2009. doi:10.1007/s00453-008-9204-0.

3 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: Max Independent Set and Coloring. *CoRR*, abs/2007.14161, 2020. arXiv:2007.14161.

4 Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1977–1996, 2021. doi:10.1137/1.9781611976465.118.

5 Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. Twin-width IV: ordered graphs and matrices. *CoRR*, abs/2102.03117, 2021. arXiv:2102.03117.

6 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 601–612. IEEE, 2020. doi:10.1109/FOCS46700.2020.00062.

7 Nicolas Bousquet, Daniel Gonçalves, George B. Mertzios, Christophe Paul, Ignasi Sau, and Stéphan Thomassé. Parameterized domination in circle graphs. *Theory Comput. Syst.*, 54(1):45–72, 2014. doi:10.1007/s00224-013-9478-8.

8 Timothy M. Chan, Elyot Grant, Jochen Könemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1576–1585. SIAM, 2012. doi:10.1137/1.9781611973099.125.

9 Jianer Chen, Henning Fernau, Iyad A. Kanj, and Ge Xia. Parametric duality and kernelization: Lower bounds and upper bounds on kernel size. *SIAM J. Comput.*, 37(4):1077–1106, 2007. doi:10.1137/050646354.

10 Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.*, 72(8):1346–1367, 2006. doi:10.1016/j.jcss.2006.04.007.

11 Brent N. Clark, Charles J. Colbourn, and David S. Johnson. Unit disk graphs. *Discret. Math.*, 86(1-3):165–177, 1990. doi:10.1016/0012-365X(90)90358-O.

12 Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000. doi:10.1007/s002249910009.

13 Marek Cygan. Deterministic parameterized connected vertex cover. In Fedor V. Fomin and Petteri Kaski, editors, *Algorithm Theory - SWAT 2012 - 13th Scandinavian Symposium and Workshops, Helsinki, Finland, July 4-6, 2012. Proceedings*, volume 7357 of *Lecture Notes in Computer Science*, pages 95–106. Springer, 2012. doi:10.1007/978-3-642-31155-0_9.

14 Marek Cygan, Fedor V Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 4. Springer, 2015.

15 Marek Cygan, Fabrizio Grandoni, and Danny Hermelin. Tight kernel bounds for problems on graphs with small degeneracy. *ACM Trans. Algorithms*, 13(3):43:1–43:22, 2017. doi:10.1145/3108239.

16 Marek Cygan, Geevarghese Philip, Marcin Pilipczuk, Michal Pilipczuk, and Jakub Onufry Wojtaszczyk. Dominating set is fixed parameter tractable in claw-free graphs. *Theor. Comput. Sci.*, 412(50):6982–7000, 2011. doi:10.1016/j.tcs.2011.09.010.

17 Marek Cygan, Marcin Pilipczuk, Michal Pilipczuk, and Jakub Onufry Wojtaszczyk. Kernelization hardness of connectivity problems in d-degenerate graphs. *Discret. Appl. Math.*, 160(15):2131–2141, 2012. doi:10.1016/j.dam.2012.05.016.

18 Anuj Dawar and Stephan Kreutzer. Domination problems in nowhere-dense classes. In Ravi Kannan and K. Narayan Kumar, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2009, December 15-17, 2009, IIT Kanpur, India*, volume 4 of *LIPIcs*, pages 157–168. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2009. doi:10.4230/LIPIcs.FSTTCS.2009.2315.

19 Erik D. Demaine, Fedor V. Fomin, Mohammad Taghi Hajiaghayi, and Dimitrios M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and H-minor-free graphs. *J. ACM*, 52(6):866–893, 2005. doi:10.1145/1101821.1101823.

20 Michael Dom, Daniel Lokshtanov, and Saket Saurabh. Kernelization Lower Bounds Through Colors and IDs. *ACM Trans. Algorithms*, 11(2):13:1–13:20, 2014. doi:10.1145/2650261.

21 Frederic Dorn. Dynamic programming and fast matrix multiplication. In Yossi Azar and Thomas Erlebach, editors, *Algorithms - ESA 2006, 14th Annual European Symposium, Zurich, Switzerland, September 11-13, 2006, Proceedings*, volume 4168 of *Lecture Notes in Computer Science*, pages 280–291. Springer, 2006. doi:10.1007/11841036_27.

22 Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013. doi:10.1007/978-1-4471-5559-1.

23 Pål Grønås Drange, Markus Sortland Dregi, Fedor V. Fomin, Stephan Kreutzer, Daniel Lokshtanov, Marcin Pilipczuk, Michal Pilipczuk, Felix Reidl, Fernando Sánchez Villaamil, Saket Saurabh, Sebastian Siebertz, and Somnath Sikdar. Kernelization and sparseness: the case of dominating set. In Nicolas Ollinger and Heribert Vollmer, editors, *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016, February 17-20, 2016, Orléans, France*, volume 47 of *LIPIcs*, pages 31:1–31:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.STACS.2016.31.

24 Martin Farber and J. Mark Keil. Domination in permutation graphs. *J. Algorithms*, 6(3):309–321, 1985. doi:10.1016/0196-6774(85)90001-X.

25 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Kernels for (connected) dominating set on graphs with excluded topological minors. *ACM Trans. Algorithms*, 14(1):6:1–6:31, 2018. doi:10.1145/3155298.

26 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Bidimensionality and kernels. *SIAM J. Comput.*, 49(6):1397–1422, 2020. doi:10.1137/16M1080264.

27 Fedor V. Fomin and Dimitrios M. Thilikos. Fast parameterized algorithms for graphs on surfaces: Linear kernel and exponential speed-up. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings*, volume 3142 of *Lecture Notes in Computer Science*, pages 581–592. Springer, 2004. doi:10.1007/978-3-540-27836-8_50.

28 Fedor V. Fomin and Dimitrios M. Thilikos. Dominating sets in planar graphs: Branch-width and exponential speed-up. *SIAM J. Comput.*, 36(2):281–309, 2006. doi:10.1137/S0097539702419649.

29 Jakub Gajarský, Petr Hlinený, Jan Obdrzálek, Sebastian Ordyniak, Felix Reidl, Peter Rossmanith, Fernando Sánchez Villaamil, and Somnath Sikdar. Kernelization using structural parameters on sparse graph classes. *J. Comput. Syst. Sci.*, 84:219–242, 2017. doi:10.1016/j.jcss.2016.09.002.

30 M. R. Garey, David S. Johnson, and Larry J. Stockmeyer. Some simplified np-complete graph problems. *Theor. Comput. Sci.*, 1(3):237–267, 1976. doi:10.1016/0304-3975(76)90059-1.

31 Petr A. Golovach and Yngve Villanger. Parameterized complexity for domination problems on degenerate graphs. In Hajo Broersma, Thomas Erlebach, Tom Friedetzky, and Daniël Paulusma, editors, *Graph-Theoretic Concepts in Computer Science, 34th International Workshop, WG 2008, Durham, UK, June 30 - July 2, 2008. Revised Papers*, volume 5344 of *Lecture Notes in Computer Science*, pages 195–205, 2008. doi:10.1007/978-3-540-92248-3_18.

32 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017. doi:10.1145/3051095.

33 Shai Gutner. Polynomial kernels and faster algorithms for the dominating set problem on graphs with an excluded minor. In Jianer Chen and Fedor V. Fomin, editors, *Parameterized and Exact Computation, 4th International Workshop, IWPEC 2009, Copenhagen, Denmark, September 10-11, 2009, Revised Selected Papers*, volume 5917 of *Lecture Notes in Computer Science*, pages 246–257. Springer, 2009. doi:10.1007/978-3-642-11269-0_20.

34 Danny Hermelin, Matthias Mnich, Erik Jan van Leeuwen, and Gerhard J. Woeginger. Domination when the stars are out. *ACM Trans. Algorithms*, 15(2):25:1–25:90, 2019. doi:10.1145/3301445.

35 Eun Jung Kim, Alexander Langer, Christophe Paul, Felix Reidl, Peter Rossmanith, Ignasi Sau, and Somnath Sikdar. Linear kernels and single-exponential algorithms via protrusion decompositions. *ACM Trans. Algorithms*, 12(2):21:1–21:41, 2016. doi:10.1145/2797140.

36 Tomohiro Koana, Christian Komusiewicz, and Frank Sommer. Exploiting c-closure in kernelization algorithms for graph problems. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, volume 173 of *LIPIcs*, pages 65:1–65:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ESA.2020.65.

37 David Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.*, 11(2):329–343, 1982. doi:10.1137/0211025.

38 Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Comb. Theory, Ser. A*, 107(1):153–160, 2004. doi:10.1016/j.jcta.2004.04.002.

39 Jaroslav Nesetril and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012. doi:10.1007/978-3-642-27875-4.

40 Geevarghese Philip, Venkatesh Raman, and Somnath Sikdar. Polynomial kernels for dominating set in graphs of bounded degeneracy and beyond. *ACM Trans. Algorithms*, 9(1):11:1–11:23, 2012. doi:10.1145/2390176.2390187.

41 Wojciech Przybyszewski and Szymon Toruńczyk. personal communication, 2021.

42 Venkatesh Raman and Saket Saurabh. Short Cycles Make W -hard Problems Hard: FPT Algorithms for W-hard Problems in Graphs with no Short Cycles. *Algorithmica*, 52(2):203–225, 2008. doi:10.1007/s00453-007-9148-9.

43 Jan Arne Telle and Yngve Villanger. FPT algorithms for domination in sparse graphs and beyond. *Theor. Comput. Sci.*, 770:62–68, 2019. doi:10.1016/j.tcs.2018.10.030.