


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. Our main result is that a polynomial kernel for k -DOMINATING SET on graphs of twin-width at most 4 would contradict a standard complexity-theoretic assumption. The reduction is quite involved, especially to get the twin-width upper bound down to 4, and can be tweaked to work for CONNECTED k -DOMINATING SET and TOTAL k -DOMINATING SET (albeit with a worse upper bound on the twin-width). The k -INDEPENDENT SET problem admits the same lower bound by a much simpler argument, previously observed [ICALP '21], which extends to k -INDEPENDENT DOMINATING SET, k -PATH, k -INDUCED PATH, k -INDUCED MATCHING, etc.

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. Interestingly the kernel applies to graphs of Vapnik-Chervonenkis 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 optimization/decision graph 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

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 $\mathcal{P}_n, \dots, \mathcal{P}_1$ 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 $2 \leq i \leq n$, \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 $\mathcal{P}_n, \dots, \mathcal{P}_1$ of G with every part of every \mathcal{P}_i ($1 \leq i \leq n$) 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 $\varphi = \exists x_1 \dots \exists x_k \bigwedge_{1 \leq i < j \leq k} \neg(x_i = x_j \vee E(x_i, x_j))$ and k -DOMINATING SET with $\varphi = \exists x_1 \dots \exists x_k \forall x \bigvee_{1 \leq i \leq k} (x = x_i \vee E(x, x_i))$. 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-composition¹ 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 9), and (3) MAXIMUM INDEPENDENT SET is NP-hard in (subcubic) planar graphs [31]. 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 [12], 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 10), 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 SET² on “sparse”³ 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 [29, 22], bounded-genus graphs and more generally classes excluding a fixed minor [20, 26, 35], and an FPT algorithm with

¹ See Section 2 for the relevant background on how to rule out a polynomial kernel

² All the subsequent results also hold for k -INDEPENDENT SET.

³ *Sparse* is an overloaded term; here we use it as *not containing arbitrarily large bicliques as subgraphs*.

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 [33] that says that first-order model checking is FPT in any nowhere dense class.⁴ 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 [43] extend the fixed-parameter tractability of k -DOMINATING SET to any $K_{t,t}$ -free class (for a fixed t). Telle and Villanger [46] 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 [28], apex-minor-free graphs [27], but more generally in any class excluding a fixed topological minor [26]. 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) [45]. A polynomial kernel of size $O(k^{(t+1)^2})$ is obtained for $K_{t,t}$ -free graphs [43], the most general “sparse” class. Contrary to the FPT algorithm, a polynomial kernel in the parameter $k+t$ is highly unlikely [21]. More precisely, for any $\varepsilon > 0$, a kernel of size $k^{(t-1)(t-3)-\varepsilon}$ would imply that $\text{coNP} \subseteq \text{NP/poly}$ [16]. On classes of bounded expansion⁵ k -DOMINATING SET has a linear kernel, while the seemingly closely related CONNECTED k -DOMINATING SET has no polynomial kernel [24]. The latter result refines a reduction showing the same lower bound on 2-degenerate graphs [18].

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 $\text{W}[1]$ -hard (hence unlikely FPT) on circle graphs [7]. This is somewhat unexpected since DOMINATING SET is polytime solvable on permutation graphs [25], a large subclass of circle graphs. On the positive side, k -DOMINATING SET admits a polynomial kernel on so-called *c-closed graphs* [39], a far-reaching dense generalization of bounded d -degenerate graphs.

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 [43]) or are dense classes on which MINIMUM DOMINATING SET is polytime solvable, like bounded rank-width graphs [13], and (subclasses of) permutation graphs [25]. 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 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

⁴ The definition of nowhere denseness being technical and unnecessary to the current paper, we refer the interested reader to [42]. Let us just mention that bounded-degree graphs, planar graphs, and proper (topological) minor-closed classes are all nowhere dense.

⁵ 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.

overcome that obstacle and show the following.

► **Theorem 1.** *Unless $\text{coNP} \subseteq \text{NP}/\text{poly}$, 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 and CAPACITATED k -VERTEX COVER admit polynomial kernels on graphs of bounded twin-width, while such kernels are unlikely on general graphs [21]. 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 [27] via the generic framework of bidimensionality, and even for topological-minor-free classes [38]. Another powerful meta-theorem by Gajarský et al. [30] says that every problem with the so-called *finite integer index* (intuitively, that its bounded 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 [18].

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 [44].

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.*

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.*

	k -DOMINATING SET	CONNECTED k -DS	CONNECTED k -VC
general	W[2]-complete [23]	W[2]-complete [23]	FPT [14], no PK [21]
bounded expansion	LK [24]	FPT [19], no PK [24]	LK [30]
bounded biclique	PK [43]	FPT [46], no PK [18]	PK, no LK [18]
bounded degeneracy	PK [43]	FPT [32], no PK [18]	PK, no LK [18, 16]
$K_{1,3}$ -free	PK [37]	FPT, no PK [37]	LK (trivial)
$K_{1,4}$ -free	W[2]-complete [17]	W[2]-complete [17]	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	$O(k^{1.5})$-vertex kernel

■ **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 $\text{coNP} \subseteq \text{NP/poly}$. Our new results are in bold (the results without a reference nor in bold are consequences of results in bold).

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 denote by $[i, j]$ the set of integers $\{i, i + 1, \dots, j - 1, j\}$, 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.


2.1 Graph theory

Unless stated otherwise, all graphs are assumed undirected and simple, that is, they do not have parallel edges or self-loops. We denote by $V(G)$ and $E(G)$, the set of vertices and edges, respectively, of a graph G . For $S \subseteq V(G)$, we denote the *open neighborhood* (or simply *neighborhood*) of S by $N_G(S)$, i.e., the set of neighbors of S deprived of S , and the *closed neighborhood* of S by $N_G[S]$, i.e., the set $N_G(S) \cup S$. We simplify $N_G(\{v\})$ into $N_G(v)$, and $N_G[\{v\}]$ into $N_G[v]$. We may omit the subscript when G is clear from the context.

We denote by $G[S]$ the subgraph of G induced by S , and $G - S := G[V(G) \setminus S]$. An injective mapping $\eta : V(H) \rightarrow V(G)$ *witnesses* that H is a subgraph of G , if $uv \in E(H)$ implies $\eta(u)\eta(v) \in E(G)$. A bijective mapping $\eta : V(H) \rightarrow V(G)$ *witnesses* that H is a *spanning* subgraph of G , if $uv \in E(H)$ implies $\eta(u)\eta(v) \in E(G)$.

A *connected subset* (or *connected set*) $S \subseteq V(G)$ is one such that $G[S]$ is connected. A *dominating set* is a set $S \subseteq V(G)$ such that $N[S] = V(G)$. A *vertex cover* is a set $S \subseteq V(G)$ such that every edge of G has at least one of its two endpoints in S . An *independent set* is a set $S \subseteq V(G)$ such that $G[S]$ is edgeless.

For two disjoint sets $A, B \subseteq V(G)$, $E(A, B)$ denotes the set of edges in $E(G)$ with one endpoint in A and the other one in B . We also denote by $G[A, B]$ the bipartite graph



■ **Figure 1** A 5-cycle of strict half-graphs of height 6.

$(A \cup B, E(A, B))$. Two distinct vertices u, v such that $N(u) = N(v)$ are called *false twins*, and *true twins* if $N[u] = N[v]$. Two vertices are *twins* if they are false twins or true twins.

A set $S \subseteq V(G)$ is a *module* if every vertex outside S is fully adjacent to S or fully non-adjacent to S . A module is said to be *trivial* if it is a singleton or the entire vertex set, and a graph without non-trivial modules is called *prime*. The vertex set of every graph G can be partitioned into modules H_1, H_2, \dots, H_ℓ , some of which can be singletons. $\mathcal{H} = \{H_1, H_2, \dots, H_\ell\}$ is then called a *modular decomposition* of G . The *quotient graph* G/\mathcal{H} of G has vertex set \mathcal{H} and an edge between H_i and H_j whenever every vertex of H_i is adjacent to every vertex of H_j . By definition of a module, $H_i H_j$ is not an edge of G/\mathcal{H} if and only if there is no edge in G between H_i and H_j . We say that $X \subseteq V(G)$ is a *module relative to* $Y \subseteq V(G) \setminus X$ if every vertex of Y is either fully adjacent to X or fully non-adjacent to X . Hence X is a module if it is a module relative to $V(G) \setminus X$.

The *strict half-graph of height t* is (up to isomorphism) the graph with vertex set $\{a_1, \dots, a_t, b_1, \dots, b_t\}$ and edge set $\{a_i b_j : i < j, i \in [t], j \in [t]\}$. One can see $\{a_1, \dots, a_t\}$ *oriented toward* $\{b_1, \dots, b_t\}$ 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 $\{a_1^p, \dots, a_t^p : p \in [0, \ell - 1]\}$ and edge set $\{a_i^p a_j^{p+1 \bmod \ell} : i < j, i \in [t], j \in [t], p \in [0, \ell - 1]\}$. 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 1 for an example of a 5-cycle of strict half-graphs of height 6. 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 $X \subseteq V(H)$, $|\{X \cap e : e \in E(H)\}| \leq c \cdot |X|$. By extension, we may say that a graph class \mathcal{C} has VC density at most 1 if the neighborhood hypergraph of every graph in \mathcal{C} has VC density at most 1.

2.2 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 $N_G(u) \Delta N_G(v)$ 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 2 for an illustration.



■ **Figure 2** Contraction of vertices u and v , and how the edges of the trigraph are updated.

A d -sequence (or *contraction sequence*) is a sequence of d -trigraphs G_n, G_{n-1}, \dots, 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 $tw(G)$, is the minimum integer d such that G admits a d -sequence. Note that, in what precedes, the 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 .

We may give a (partial) contraction sequence by listing the pairs of vertices to contract: $(u_1, v_1), (u_2, v_2), \dots$ where u_i, v_i are implicitly contracted to vertex $s(u_i) \cup s(v_i)$ with $s(x) = x$ if x is a set of vertices and $s(x) = \{x\}$ if x is a single vertex. Thus for instance u_2 could be equal to $\{u_1, v_1\}$.

For $u \in V(G_i)$, we denote by $u(G)$ the subset of $V(G)$ that was contracted to the single vertex u in G_n, G_{n-1}, \dots, G_i . Twin-width and d -sequences can be equivalently seen as a partition refinement process on $V(G)$. We start with the finest partition $\mathcal{P}_n = \{\{v\} : v \in V(G)\}$, and end with the coarsest partition $\mathcal{P}_1 = \{V(G)\}$. There is a *partition sequence* $\mathcal{P}_n, \mathcal{P}_{n-1}, \dots, \mathcal{P}_2, \mathcal{P}_1$ mimicking the contraction sequence, where the contraction of $u, v \in V(G_i)$ corresponds to the merge of parts $u(G_i), v(G_i) \in \mathcal{P}_i$ to form the part $u(G_i) \cup v(G_i) = z(G_{i-1}) \in \mathcal{P}_{i-1}$, while all the other parts are unchanged from \mathcal{P}_i to \mathcal{P}_{i-1} . The red degree (bounded by d) of a part $P \in \mathcal{P}_i$ now corresponds to the number of other parts $P' \in \mathcal{P}_i$ which are not fully adjacent nor fully non-adjacent to P in G . We may denote by $G_{\mathcal{P}}$ the trigraph corresponding to partition \mathcal{P} over $V(G)$. Thus $G_i = G_{\mathcal{P}_i}$.

Given a partition $\mathcal{H} = \{H_1, H_2, \dots, H_\ell\}$ of the vertex set $V(G)$ of a trigraph G , we call *quotient trigraph* G/\mathcal{H} the trigraph obtained by contracting every H_i into single vertices. Note that the resulting trigraph does not depend on the order in which the contractions are made. Thus the quotient trigraph is well-defined.

2.3 Useful twin-width bounds

It can be seen that the twin-width may only decrease when taking induced subtrigraphs.

► **Observation 6.** *Let G be a trigraph and H be an induced subtrigraph of G . Then, $\text{tw}(H) \leq \text{tw}(G)$.*

The following observation states that turning some non-edges or black edges into red edges can only increase the twin-width.

► **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 $\text{tw}(G) \leq \text{tw}(G')$.*

Proof. Indeed, any d -sequence for G' is a d -sequence for G . ◀

We will not need the next lemma in its particular form. Yet it yields some good insight on contraction sequences in general, and on the “core structure” of low twin-width used in the proof Theorem 1, in particular. Thus we give a short proof here.

► **Lemma 8.** *Cycles of strict half-graphs have twin-width at most 3.*

Proof. Consider G_t , the ℓ -cycle of strict half-graphs of height t , on vertex set $\{a_1^p, \dots, a_t^p : p \in [0, \ell - 1]\}$. By Observation 7 we may prove the stronger statement that $\text{tw}(G'_t) \leq 3$ where G'_t is obtained from G_t by adding the red edge $a_1^p a_1^{p+1 \bmod \ell}$ for every $p \in [0, \ell - 1]$. G'_t is a red cycle, and admits a 2-sequence by iteratively contracting the endpoints of any red edge. Therefore, we shall just check that the following is a partial 3-sequence from G'_t to G'_{t-1} : $(a_1^0, a_2^0), (a_1^1, a_2^1), \dots, (a_1^{\ell-1}, a_2^{\ell-1})$. Indeed $\{a_1^p, a_2^p\}$ has, in the suggested partial sequence, at most three red neighbors: $a_1^{p-1 \bmod \ell}$, $a_1^{p+1 \bmod \ell}$, and $a_2^{p+1 \bmod \ell}$; the latter two being a single vertex when $p = \ell - 1$. And all the other vertices have red degree at most 2. ◀

The next lemma was already invoked in the introduction. We will use its essence in a more general form in Section 5.1, namely that contracting vertices from a module X relative to Y do not create red edges incident to Y .

► **Lemma 9.** *Let G be a graph and $\mathcal{H} = \{H_1, H_2, \dots, H_\ell\}$ be its modular partition. Then,*

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

Proof. Consider a sequence which contracts each module H_i into a single vertex, followed by a contraction sequence of G/\mathcal{H} attaining $\text{tw}(G/\mathcal{H})$. Contracting each module H_i creates red edges only within the module H_i , thus G/\mathcal{H} has red degree 0. Hence, $\text{tw}(G)$ is upper bounded by the twin-widths of $G[H_1], G[H_2], \dots, G[H_\ell]$, and G/\mathcal{H} . The other inequality follows from Observation 6, since G/\mathcal{H} is also an induced subgraph of G . ◀

Finally we will use the following lemma to finish our 4-sequences.

► **Lemma 10.** *Any trigraph whose total graph is a subdivision of a subgraph of a grid has twin-width at most 4.*

Proof. By Observation 7, we may assume that the given trigraph is a subdivision of a complete grid consisting of red edges only. By contracting every adjacent pair of vertices one of which has degree 2, we obtain a (complete) grid. In a (red) grid on the vertex set $[n] \times [m]$, consider a matching of n leftmost horizontal edges, that is the edges connecting $(i, 1)$ and $(i, 2)$ for every $i \in [n]$. Contracting these edges from top to bottom produces intermediate red graphs of red degree at most 4, ending in a red grid graph on $[n] \times [m - 1]$. Finally a path is obtained, which can be contracted into a single vertex while keeping red degree at most 2. ◀

2.4 List of handled problems

We refer to Section 2.1 for the definitions of dominating set, vertex cover, connected set, and independent set. We will mostly deal with the following three problems, given by order of importance.

k -DOMINATING SET	Parameter: k
Input: A graph G and an integer k	
Question: Does G have a dominating set of size at most k ?	

CONNECTED k -VERTEX COVER (CONNECTED k -VC)	Parameter: k
Input: A graph G and an integer k	
Question: Does G have a vertex cover of size at most k which induces a connected subgraph of G ?	

Given a graph G and a capacity function $c : V(G) \rightarrow \mathbb{N}$, a *capacitated vertex cover* X of G is a vertex cover of G which admits a mapping $\rho : E(G) \rightarrow X$ assigning to each vertex $x \in X$ no more edges than its capacity, i.e., $|\rho^{-1}(x)| \leq c(x)$ for every $x \in X$.

CAPACITATED k -VERTEX COVER (CAPACITATED k -VC)	Parameter: k
Input: A graph G with a capacity function $c : V(G) \rightarrow \mathbb{N}$ and an integer k	
Question: Does G admit a capacitated vertex cover X of size at most k ?	

The lower bound for k -DOMINATING SET will also apply to its *connected* and *total* variants.

CONNECTED k -DOMINATING SET (CONNECTED k -DS)	Parameter: k
Input: A graph G and an integer k	
Question: Does G have a dominating set of size at most k which induces a connected subgraph of G ?	

TOTAL k -DOMINATING SET	Parameter: k
Input: A graph G and an integer k	
Question: Does G admit a set $X \subseteq V(G)$ of size at most k such that every vertex of G has a neighbor in X ?	

For completeness, we give the definition of the remaining problems that are mentioned at least once in the paper.

k -INDEPENDENT SET	Parameter: k
Input: A graph G and an integer k	
Question: Does G have an independent set of size at least k ?	

INDEPENDENT k -DOMINATING SET	Parameter: k
Input: A graph G and an integer k	
Question: Does G have an independent dominating set of size at most k ?	

k -PATH	Parameter: k
Input: A graph G and an integer k	
Question: Does G have a path of length at least k ?	

k -INDUCED PATH	Parameter: k
Input: A graph G and an integer k	
Question: Does G have an induced path of length at least k ?	

k -INDUCED MATCHINGParameter: k **Input:** A graph G and an integer k **Question:** Does G have a matching M of size at least k such that there are no edges between endpoints of M other than the edges of M ?

2.5 Kernels or lack thereof

We often identify a problem to the language made by its positive instances. For a parameterized problem \mathcal{Q} , a *kernel* of size bounded by a function f is a polynomial-time reduction $\rho: \Sigma^* \times \mathbb{N} \rightarrow \Sigma^* \times \mathbb{N}$ such that $(x, k) \in \mathcal{Q}$ if and only if $\rho(x, k) \in \mathcal{Q}$, and $|\rho(x, k)| \leq f(k)$. 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 [15], which we will rely on to show the absence of a polynomial kernel in Theorem 1. An OR-cross-composition is a polynomial reduction that takes t instances of an NP-hard problem \mathcal{L} and builds an instance of a parameterized problem \mathcal{Q} , such that the OR of the input instances is equivalent to the output instance. More precisely at least one of the t input instances is positive if and only if the output instance is positive. It furthermore allows to set some restrictions on the input NP-hard instances in the form of a polynomial equivalence relation. This relation can be useful in shaping the input instances in such a way that the composition behaves well.

► **Definition 11.** 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 12.** 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 $(y, N) \in \Sigma \times \mathbb{N}$ 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 $\text{coNP} \subseteq \text{NP/poly}$.

► **Theorem 13.** 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 $\text{coNP} \subseteq \text{NP/poly}$.

2.6 Organization of the rest of the paper

In Section 3, we sketch the proof of our main result, Theorem 1. In Section 4, we show the NP-hardness of MINIMUM DOMINATING SET in a customized setting that will be particularly convenient for the subsequent OR-cross-composition. In Section 5 we present the OR-cross-composition, show its correctness, and prove the twin-width upper bound, thereby establishing Theorem 1. In Section 6, we present two simple $O(k^2)$ kernels for CONNECTED k -VERTEX COVER and CAPACITATED k -VERTEX COVER, and a refined $O(k^{1.5})$ kernel for the former problem. In Section 7, we show that graphs of twin-width 1 can be efficiently recognized and a 1-sequence with additional property can be produced, and argue that this extra property can be used for polynomial-time algorithm for a wide range of problems.

3 Outline of the OR-cross-composition for k -Dominating Set

Let us start explaining why we should *not* expect a simple OR-composition. After all, it is only fair to ask for a justification that the kernel lower bound for k -DOMINATING SET spans a dozen of pages when the same result for k -INDEPENDENT SET is a side note; even more so, when the former is often more intractable than the latter from a parameterized complexity standpoint.


The simplest OR-composition is the disjoint union of the input instances. By *simple* OR-composition we mean one, like for k -INDEPENDENT SET, straightforwardly based on juxtaposing the instances. 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 3.

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,



■ **Figure 3** The overall picture. Instances I_1, \dots, 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.

for instance by being planar.

In the next section, 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.


4 Tailored NP-hardness for Dominating Set

We will show in this section the following hardness result for **DOMINATING SET**. The extra properties that we get compared to existing NP-hardnesses of **DOMINATING SET** (even those on planar instances) is crucial for the subsequent **OR-cross-composition**, as hinted at in the previous section.

The statement of the next two theorems involve what we will call *snaking grids*. See Figure 4 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 motivation behind the snaking grid will become clear when designing the cross-composition and bounding its twin-width in Section 5.1: it allows to superpose a canonical hamiltonian cycle such that the maximum degree remains 3.

► **Theorem 14.** *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



■ **Figure 4** The 5×10 snaking grid.

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

We will obtain Theorem 14 as a direct consequence of the following reduction, and the fact that PLANAR 3-SAT is NP-hard [40] (see Section 4.1).

► **Theorem 15.** *There is a polynomial-time reduction from PLANAR 3-SAT to DOMINATING SET that, on n -variable m -clause PLANAR 3-SAT-instances φ produces DOMINATING SET-instances $(G, N, \mathcal{B} = \{B_1, \dots, B_N\}, \eta)$ such that:*

- (i) \mathcal{B} partitions $V(G)$, and G has a partial 4-sequence to the quotient trigraph G/\mathcal{B} ,
- (ii) G/\mathcal{B} is a spanning subgraph of the $(m+1) \times n'$ snaking grid, with n' even in $\{n, n+1\}$, witnessed by η ,
- (iii) every dominating set of G intersects each B_i , for $i \in [N]$, and
- (iv) φ is satisfiable if and only if G has a dominating set of size N .

One can observe that conditions (i), (ii), and (iii) of Theorems 14 and 15 match. Condition (iv) in Theorem 15 only insists that the polynomial-time transformation is a valid reduction.


4.1 Planar satisfiability

PLANAR 3-SAT was introduced by Lichtenstein, who showed its NP-hardness [40], as a convenient starting point to prove the intractability of planar problems. It is a restriction of 3-SAT where the variable/clause incidence graph is planar even if one adds edges between two consecutive variables for a specified ordering of the variables: x_1, x_2, \dots, x_n ; i.e., $x_i x_{i+1}$ is an edge (with index $i+1$ taken modulo n). In any PLANAR 3-SAT-instance φ one can partition the clause set \mathcal{C} into $(\mathcal{C}^+, \mathcal{C}^-)$ such that \mathcal{C}^+ and \mathcal{C}^- each admits a removal ordering, where a *removal ordering* consists of iteratively applying the two following kinds of deletions:

- removing a variable which is not present in any remaining clause, or
- removing a clause on *three consecutive variables* together with the *middle variable*,

which ends up with an empty set of clauses. *Three consecutive variables* means three variables x_i, x_j, x_k , with $i < j < k$ such that $x_{i+1}, x_{i+2}, \dots, x_{j-1}$ and $x_{j+1}, x_{j+2}, \dots, x_{k-1}$ have all been removed already. The middle variable of the clause is x_j . For an example, see Figure 5. What matters to us is that φ can be embedded in an $(|\mathcal{C}| + 1) \times n$ grid (the dashed blue lines

in Figure 5) such that every “event” (a variable “wire” turns or finishes) happens at a grid point.



■ **Figure 5** The bipartition $(\mathcal{C}^+, \mathcal{C}^-)$ of an n -variable m -clause PLANAR 3-SAT-instance, with $n = 8$ and $m = 10$. The square vertices represent the clauses. All the degree-3 vertices of the solid black (and green) graph lie at intersections of the $(m + 1) \times n$ grid represented with the dashed blue lines. The $(3m + 1) \times (3n - 2)$ grid made by the dashed blue and gray lines will embed the snaking grid. The “wire” of variable x_1 is highlighted in green.

The reduction to establish Theorem 15 will be a rather transparent substitution of local gadgets on grid points of Figure 5. We need a variable gadget (to initiate a Boolean choice), a propagation gadget (“wires” that may split), and a clause gadget (where the “wires” of three variables meet at a grid point). In addition, we will design a very simple dummy gadget, used to fill all the unoccupied grid points. We start with a description of the variable and propagation gadgets.

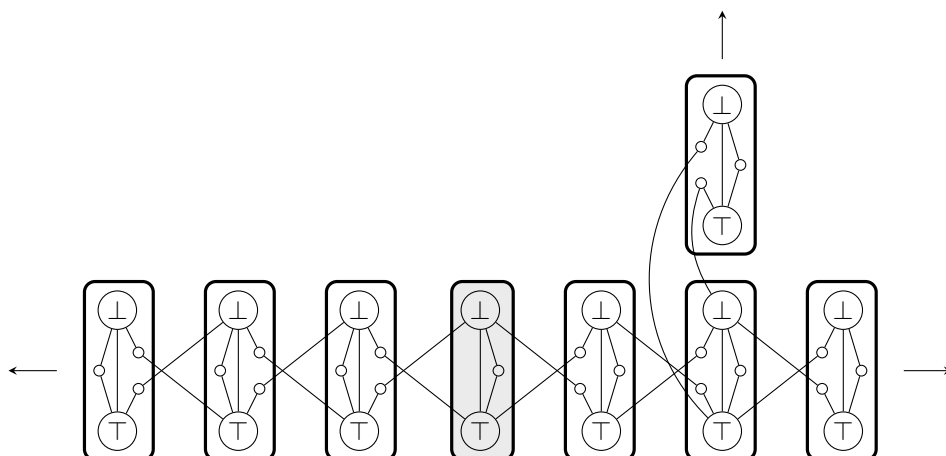
4.2 Variable and propagation gadgets

We distinguish the *initial variable gadget* and the (*regular*) *variable gadget*. The initial variable gadget is simply a triangle, one vertex of which corresponds to setting the variable to true (marked \top in the figures), a second vertex, to setting it to false (marked \perp), and a third vertex which will remain of degree 2. These three vertices form one part, say, B_ℓ of \mathcal{B} . We denote them by \top , \perp , and d , and use the subscript of the part of \mathcal{B} they belong to whenever we want to be specific (so \top_ℓ , \perp_ℓ , and d_ℓ).

The *regular variable gadget* (or *variable gadget* for short) is the *bull* graph. More precisely, it is an initial variable gadget where one adds one pendant neighbor t to \top and one pendant neighbor f to \perp . Again these five vertices form one part of \mathcal{B} , say B_ℓ , that we may denote \top_ℓ , \perp_ℓ , d_ℓ , t_ℓ , and f_ℓ . For both the initial and regular variable gadgets, the *index* of the gadget is simply ℓ .

The propagation gadget from a (possibly initial) variable gadget of index ℓ to a regular variable gadget of index ℓ' consists of the two edges $\top_\ell f_{\ell'}$ and $\perp_\ell t_{\ell'}$. The “propagation” starts at an initial variable gadget and extends in two directions. It may also “split,” that is, a same regular variable gadget can propagate to two variable gadgets, one horizontally and one vertically. See Figure 6 for a depiction of these gadgets.

For every variable x_i of φ , we have a unique initial variable gadget, and several regular



■ **Figure 6** Initial variable gadget (slightly shaded box in the middle), seven (regular) variable gadgets, and the propagation gadgets (edges in between variable gadgets), allowing the “wire” to split. The thick boxes represent the partition \mathcal{B} , each box corresponding to an initial or regular variable gadget. The figure should be rotated 90 degrees to match Figure 5 where the propagation from the initial variable gadgets occurs vertically.

variable gadgets connected to it via a chain of propagations. We call *wire* of x_i , the graph induced by all these variable gadgets (the initial variable gadget included). The *digraph of a wire* (or, by extension, of its corresponding variable) is the directed graph whose vertices are the variable gadgets of the wire, and with an arc from a variable gadget of index ℓ to one of index ℓ' if there is a propagation gadget from the former to the latter. Following Figure 5 the digraph of every wire is an out-tree of maximum out-degree 2, rooted at the initial variable gadget. The leaves of this tree are adjacent to a grid point where we will place a clause gadget.

The following lemma ensures the expected behavior of a choice propagation. The rest of the construction (mainly the insertion of the clause gadgets) will not change the properties used in its proof, so we show it here.


▶ **Lemma 16.** *Let H be the wire of variable x , and let \vec{T} be the out-tree, digraph of H . The only two dominating sets of H of size (at most) $|V(\vec{T})|$ consists of picking \top in every gadget, or of picking \perp in every gadget.*

Proof. Every variable gadget has a vertex d only adjacent to vertices within the gadget, so a dominating set of size $|V(\vec{T})|$ has to intersect each gadget exactly once. Since the digraph of H is supposed to be an out-tree, and the vertices $t_{\ell'}$ and $f_{\ell'}$ have their neighborhood included in the gadget of index ℓ' and its unique in-neighbor (in \vec{T}), of index, say, ℓ , it suffices to observe that the only pairs of vertices dominating all the vertices in these two gadgets are $\{\top_{\ell}, \top_{\ell'}\}$ and $\{\perp_{\ell}, \perp_{\ell'}\}$. ◀

In the statement of Lemma 16, the former choice *sets x to true*, the latter, *sets x to false*.

4.3 Clause and dummy gadgets

The *clause gadget* of C_j consists of two vertices c_j, z_j , where c_j is linked to its three literals in the corresponding wire ends, and z_j is isolated. The set $\{c_j, z_j\}$ is added to \mathcal{B} . See Figure 7 for an illustration.



■ **Figure 7** Three variable wires converging to a clause gadget $C_j = x_{a_1} \vee \neg x_{a_2} \vee \neg x_{a_3}$. Again, the thick boxes represent the partition \mathcal{B} . In fact, due to the horizontal snaking, either one of x_{a_1}, x_{a_3} would appear above the clause gadget (see Figure 8).

A *dummy gadget* is simply an isolated vertex z , which makes its own singleton part $\{z\}$ in \mathcal{B} . Thus a dummy gadget is a clause gadget minus the vertex c_j . The isolated vertices z_j 's in the clause gadgets, and the dummy gadgets may seem artificial. Their purpose is to fulfill conditions (ii) and (iii). Besides isolated vertices will not remain isolated in the subsequent OR-cross-composition.

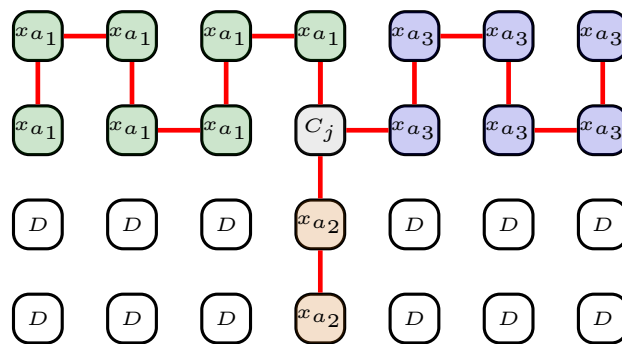
4.4 Construction

At this point, the construction is probably clear. Let n and m be the number of variables and clauses, respectively, of the PLANAR 3-SAT-instance φ . We assume that n is even (if not, we add a dummy variable not appearing in any clause). We put one initial variable gadget at each grid point (of the finer $(3m+1) \times (3n-2)$ grid) in Figure 5 with a circled vertex (n grid points in total). We place a regular variable gadget at every grid point covered by the solid black (or green) lines, and at every grid point just below any horizontal segment of a wire (to account for the snaking), except for those on the bottommost row. We orient the propagation such that the digraph of each of the n wires is an out-tree rooted at its initial variable gadget. For every clause of \mathcal{C}^- , we place a clause gadget at the corresponding grid point occupied by a square vertex in Figure 5, and link it accordingly to its literals. For every clause of \mathcal{C}^+ , the clause gadget is instead placed at the grid point just below the corresponding square vertex. This is again because of the horizontal snaking (see Figure 8). Finally we put a dummy gadget at every unoccupied grid point. This finishes the construction of the graph G , and its vertex-partition \mathcal{B} . We set N to $|\mathcal{B}|$, that is, $(3m+1) \times (3n-2)$. The embedding witness η is implicit from the construction.

Proof of Theorem 15. We now check the conditions (i) to (iv). Let us start with the correctness of the reduction, namely (iv).

Assume φ admits a satisfying assignment A , and sets each variable to the value given by A in each wire of G . Complete this solution by including all the isolated vertices of G . By design this yields a set S of size N . By Lemma 16, every vertex in a wire is dominated. The isolated vertices are dominated by themselves, and every vertex c_j in a clause gadget is dominated since A is a satisfying assignment. Thus S is a dominating set of G of size N (intersecting every part of \mathcal{B} exactly once).

Conversely, assume that G has a dominating set S of size N . As every gadget contains a vertex whose neighborhood is included in its own part of \mathcal{B} (namely d , and the isolated



■ **Figure 8** The vicinity of a clause gadget, for a clause $C_j \in \mathcal{C}^+$. Red edges materialize the edges in between two parts: propagation and edges incident to the clause gadget. Due to the horizontal snaking, the clause gadget is placed just below its position in Figure 5. Dummy gadgets are marked with a D .

vertices), S should intersect every part of \mathcal{B} exactly once. In the dummy and clause gadgets, the picked vertex should thus be the isolated vertex. By Lemma 16, in each wire a consistent choice (of taking only \top or only \perp) has to be made. Let A be the corresponding truth assignment. As S dominates each vertex c_j in the clause gadgets, it implies that A is a satisfying assignment, and ϕ is satisfiable.

Condition (i). By construction \mathcal{B} partitions $V(G)$. We want to propose a partial 4-sequence from G to G/\mathcal{B} . In other words, we wish to contract every part of \mathcal{B} into a single vertex, such that the red degree remains at most 4.

In any order, we contract the bull graph of each regular variable gadget adjacent to only two other variable gadgets in the following way: (\top, d) , $(\{\top, d\}, t)$, $(\{\top, d, t\}, \perp)$, $(\{\top, d, t, \perp\}, f)$. Internally this creates (at any moment) a single red edge. There can be up to four red edges between this part and another part of \mathcal{B} (since there are at most four leaving black edges). Observe however that these four edges can be incident to a single vertex only after the last internal contraction $(\{\top, d, t, \perp\}, f)$. Thus the red degree of these vertices remains at most 4.

Meanwhile the red degree of vertices not falling in that category is bounded by 3, in clause gadgets, and by 2, in other variable gadgets. We can then contract the initial variable gadget and the clause gadgets (there is nothing to contract in dummy gadgets). Regular variable gadgets adjacent to three parts of \mathcal{B} are now adjacent to three vertices. Indeed there is no pair of adjacent such gadgets. Thus the above suggested contraction of the bull keeps red degree at most 4.

Condition (ii). By design, G/\mathcal{B} is a subgraph of the $(m+1) \times n$ snaking grid, witnessed by η . It is spanning by introduction of the dummy gadgets.

Condition (iii). This was observed when checking the correctness of the reduction. ◀

As previously mentioned, this directly implies Theorem 14.

For CONNECTED k -DOMINATING SET and TOTAL k -DOMINATING SET on bounded twin-width graphs (with an upper bound possibly larger than 4), one does not need the snaking grid nor the dummy gadgets. In the propagation between two variable gadgets, one can add an edge between the two \top vertices, and an edge between the two \perp vertices. In the clause gadgets, one can add a neighbor y_j to z_j , adjacent to the six vertices \top and \perp in the three incident variable gadgets.

5 OR-cross-composition

We now describe the cross-composition from the NP-hard DOMINATING SET restricted as in Theorem 14 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 14) with respect to the given parameter N and dimensions (p, q) of the corresponding snaking grid. All ill-formed instances will then make a single class. Relation \mathcal{R} satisfies the conditions of Definition 11. 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 (i). Now, choosing some encoding such that all well-formed instances of $\Sigma^{\leq n}$ have at most n vertices, their parameter and snaking grid dimensions must also be bounded by n . Then, the number of equivalence classes on $\Sigma^{\leq n}$ accounting for the malformed instances is at most $n^3 + 1$, yielding (ii).

The conditions for our composition being set, consider t instances of the restricted DOMINATING SET, equivalent with respect to \mathcal{R} . If the instances are ill-formed we output an ill-formed instance of k -DOMINATING SET. If not, their equivalence yields common parameter N and snaking grid dimensions (p, q) , letting us consider them 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}_{t+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, \dots, y_N) on the $p \times q$ snaking grid, which we now describe. From the snaking grid, we take all edges of the first row, and all edges of the first and last columns. Then, from the complete underlying grid, we take all but the first edge of each column and we complete the cycle by taking all edges between two degree one vertices in the so-built union of paths (see the darker red cycles in Figures 9 and 10). Referring to the partition of instance $i \in [t+1]$ as $\mathcal{B}_i = \{B_{i,1}, \dots, B_{i,N}\}$, 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 3. 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 $\bigcup_{i < \ell \leq t+1} B_{\ell, j+1}$ (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 10. Then, the cycle of half-graphs follows the darker red fictitious hamiltonian cycles mapping to (y_1, \dots, y_N) .

The composed k -DOMINATING SET is then (H, N) , which we can construct polynomially in $\sum_{i=1}^t |I_i|$. Indeed, N is bounded by any $|I_i|$ so the dummy instance is built linearly, then graph union and addition of the bicliques takes time polynomial in $\sum_{i=1}^t |I_i|$. Moreover, the reduction directly satisfies condition (i) of Definition 12 since the output parameter N is bounded by any $|I_i|$.

5.1 The overall construction has twin-width at most 4

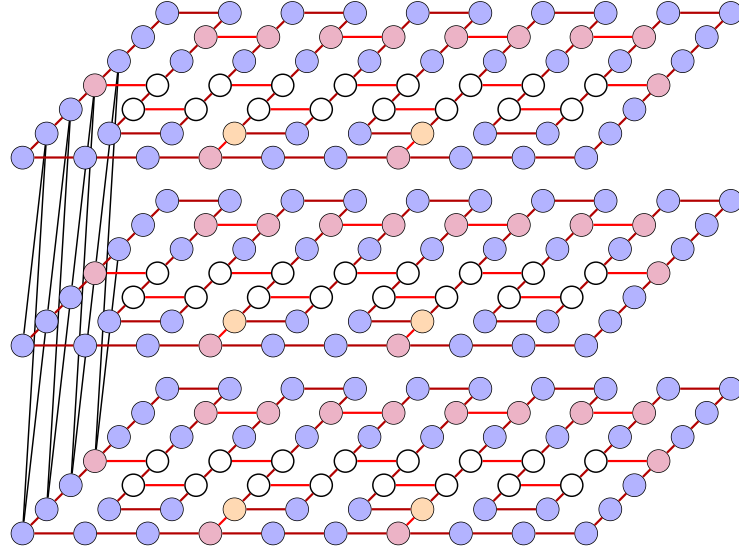
Consider the OR-cross-composition of t instances $(I_i, N, \mathcal{B}_i, p, q, \eta_i)_{i \in [t]}$ of the restricted k -DOMINATING SET described Section 5, letting (H, N) be the composed instance. Towards showing that the twin-width of H is bounded by 4, we first show that there is a partial 4-sequence contracting each $\mathcal{B}_{i,j}$ for $i \in [t+1], j \in [N]$ into a single vertex in H . After these contractions, each instance I_i/\mathcal{B}_i can be considered as a red snaking grid augmented by a hamiltonian path (Figure 9). We call *augmented snaking grid* the graph obtained by adding the (darker red) hamiltonian cycle to the snaking grid. The cycle of strict half-graphs added between the instances will then follow the hamiltonian cycles of each augmented snaking grid. Finally, we show that there is a partial 4-sequence contracting these $t+1$ grids into a single one, which is of twin-width four.



■ **Figure 9** The maximal set of red edges, with the hamiltonian cycle in darker red, and the edges of the snaking grid that are not already in the hamiltonian cycle in lighter red. A partial 4-sequence goes as follows. First contract every blue vertex with their homologous in the next layer, in any order. Then similarly contract every purple vertex, followed by every orange vertex. Finally follow the order indicated by positive integers (in increasing value). This way a vertex contracted with the next layer has at most two neighbors, or at most one non-contracted neighbor.

► **Lemma 17.** *The composed graph admits a partial 4-sequence to $H / \left(\bigcup_{i \in [t+1]} \mathcal{B}_i \right)$.*

Proof. Theorem 14 yields that each input graph $I_i, i \in [t]$ admits a partial 4-sequence to I_i/\mathcal{B}_i yielding a subgraph of the snaking grid, this is also the case for dummy instance I_{t+1} as an independent set. Our contraction sequence will consist exactly of those contractions, so



■ **Figure 10** The different layers (instances) linked by the cycle of half-graphs. Only four half-graphs are drawn for the sake of legibility.

let us see how they behave over H . Notice that for any I_i , the only contractions made above involve vertices of a single class $B_{i,j}$, which is a module with respect to $H - I_i$. Therefore, in H , the partial contraction sequence of each I_i does not create red edges towards $H - I_i$. In turn, we can contract each I_i to I_i/\mathcal{B}_i through a partial 4-sequence over H , and the only resulting red edges will belong to a single instance. ◀

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 snaking grid as depicted in Figure 9 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 $(B_{i,1}, \dots, B_{i,N})$. 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.

► **Lemma 18.** *The composed graph H has twin-width at most 4.*

Proof. By Lemma 17 and the above remarks, 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-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-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. Before proceeding, it is useful to make the following observations:

- For any grid I_i with $i > 2$, the vertices of any homologous pair $(B_{1,j}, B_{2,j})$ are each adjacent exactly to $B_{i,j+1}$ in I_i . Their contraction yields no red edges, allowing us to omit other instances and only consider the two to be contracted.
- The edges involving both snaking grids consist in the matching $(B_{1,j}, B_{2,j+1})_j$, see Figure 10, which can be assumed red while keeping red degree at most four. In our contraction sequence, this enables us to only bound the total degree at any point. Now, when contracting a pair of vertices, any vertex in the above matching adjacent to one of the vertices of our pair is also adjacent to the other because of the additional hamiltonian cycle. Therefore, the neighbors of the contracted pair will be exactly the neighbors of each vertex in its respective instance.
- Contracting a pair coarsens the partition of our vertices, so the degree of a vertex not in the pair cannot increase, thus it is enough to bound the degree of the contracted pair.

Now, consider any pair (a, b) to be contracted, such that both vertices have C already contracted, common neighbors and each P non-contracted neighbors in their respective instance. According to the second observation above, the resulting degree will be $C + 2P$. Since any vertex is of degree at most three in its respective grid $C + P \leq 3$, which would already allow us to contract the pairs in any order with degree at most six. Therefore, to bound the degree by four, the order in which we contract the pairs must be carefully chosen with respect to their position on the respective augmented snaking grids. We want to iteratively contract pairs of vertices satisfying $C + 2P \leq 4$, at each step propagating this bound to more non-contracted vertices to eventually contract all pairs. We proceed by groups of colors as depicted in Figure 9. We first contract pairs of blue vertices, those of degree two in their respective augmented snaking grid. Second, we contract pairs of purple vertices, those of degree three adjacent to two contracted vertices (necessarily blue). Third, we proceed with pairs of orange vertices, those still non contracted on the second row of the snaking grids, which are of degree three and adjacent to a blue and a purple vertex. Finally, numbered in white, we will contract the remaining degree three vertices along the corresponding snaking path.

We are now ready to describe the contraction sequence. For the blue, orange and purple pairs, the contractions can be done in any order, while we will describe the order for the last ones. For any blue pair, depending on the order of contraction we either have $(C, P) = (0, 2)$, $(C, P) = (1, 1)$ or $(C, P) = (2, 0)$. Thus, regardless of the order the contraction yields degree $C + 2P$ at most four. For any purple pair, we either have $(C, P) = (2, 1)$ or $(C, P) = (3, 0)$, again regardless of order contracting the pair yields degree at most four. The case is the same for orange pairs, and we can again contract them in any order. The graph induced by the non-contracted vertices on each instance now consists in an union of paths, with endpoints satisfying $(C, P) = (2, 1)$ and interior points such that $(C, P) = (1, 2)$. We can iteratively contract the pairs corresponding to the endpoints, satisfying $(C, P) = (2, 1)$, decreasing the length of the path by one at each step. This can be done until we are left with single vertices for each path, the corresponding pairs are then adjacent exactly to three already contracted vertices, and their contraction yields degree three.

This finishes to prove 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-sequence from our composed graph into a red augmented snaking grid. Then, as the latter is a subgraph of the red complete grid, Lemma 10 yields twin-width at most 4. ◀

5.2 Correctness

Having verified that the reduction runs in polynomial time and satisfies condition (i) of Definition 12, for the OR-cross-composition to be sound it remains to show equivalence of the instances.

► **Lemma 19.** *The composed instance is positive for k -DOMINATING SET if and only if at least one input instance is positive for the restricted k -DOMINATING SET.*

Proof. Consider the OR-cross-composition of \mathcal{R} -equivalent instances $(I_i, N)_i, i \in [t]$ of the restricted k -DOMINATING SET into k -DOMINATING SET instance (H, N) as described in Section 5. We start by showing the easier direction, the existence of a positive input instance implies that the composed instance is also positive, then we prove the converse. The reduction is such that for any set $D \in H$, D dominates H if and only if it dominates some input instance I_i .

One input instance is positive implies that the composed instance is positive

Assume there exists some $\ell \neq t + 1$ such that (I_ℓ, N) is a positive instance for the restricted k -DOMINATING SET, let us consider dominating set D for I_ℓ , with $|D| = N$ and satisfying the conditions of Theorem 14. We show that D is also a dominating set for H .

Since our reduction preserves the edges of each instance, D dominates $V(I_\ell)$ over H by assumption. To show that $V(H) - V(I_\ell) = \bigcup_{i \neq \ell, j \in [N]} B_{i,j}$ is dominated, let us consider any cell $B_{i,j}$ with $i \neq \ell$. If $i < \ell$, by construction $B_{\ell, j-1}, B_{i,j}$ form a biclique in H , then D intersects $B_{\ell, j-1}$ via item (iii) of Theorem 14 so $B_{i,j}$ is dominated. Otherwise, $i > \ell$ and $B_{\ell, j+1}, B_{i,j}$ form a biclique in H , now D must intersect $B_{\ell, j+1}$ and again $B_{i,j}$ is also dominated. Therefore D is indeed a dominating set of H , yielding the backwards implication.

The composed instance is positive implies that one input instance is positive

Conversely, assume (H, N) is a positive instance for k -DOMINATING SET, considering any solution $D \subseteq V(H)$, we show that D is necessarily also a dominating set of some I_ℓ , yielding a positive instance among the input ones.

We leverage the fact that D dominates each dummy partition class $B_{t+1, j+1}$ to yield that $|D \cap C_j| \geq 1$ for any j . By construction, each $B_{t+1, j+1}$ is only adjacent to vertices in column C_j , so we already know D intersects each $C_j \cup B_{t+1, j+1}$ for $j \in [N]$, and since $|D| = N$ each such intersection consists of a single vertex d_j . Now, the two vertices of $B_{t+1, j+1}$ are adjacent exactly to C_j , so the choice of d_j as one of them would prevent our solution from dominating the other. Then $d_j \in C_j = \bigcup_{i=1}^t B_{i,j}$, letting us define the row choice r_j for column j as the unique $r_j \in [t]$ such that $d_j \in B_{r_j, j}$.

We now show the existence of some $i \in [t]$ such that $D \subseteq V(I_i)$. From the necessity to choose one vertex per column shown above, this is equivalent to showing that for some $i \in [N]$ the row choices are constant: $r_j = i$ for $j \in [N]$. We first argue that both the even choice sequence $(r_{2j})_j$ and the odd choice sequence $(r_{2j+1})_j$ must be constant (notice that if N is odd, both contain all row choices). Considering the indices as corresponding to rows modulo N , both sequences are periodic. If one is not constant we can consider a pair of indices $(j, j+2)$ such that corresponding row choices $k = r_j, k'' = r_{j+2}$ are such that $k > k''$.

With $d_j \in B_{k,j}, d_{j+2} \in B_{k'', j}$, let us consider the adjacencies between D and column C_{j+1} towards showing an absurdity. On the one hand, the edges added in our composition between $\{d_j, d_{j+2}\}$ and C_{j+1} are exactly $\{(d_j, v) : v \in \bigcup_{i > k} B_{i, j+1}\}$ and $\{(d_{j+2}, v) : v \in \bigcup_{i < k''} B_{i, j+1}\}$. On the other hand, for any $d \in D \setminus \{d_j, d_{j+2}\}$, any edge between d and C_{j+1}


necessarily belongs to some instance $I_i, i \in [t]$. Combining the two observations above: for $i \in [k'', k]$, any edge (a, d) with $a \in B_{i,j+1}, d \in D$ must be an edge of I_i . Since $k > k''$, interval $[k'', k]$ is of size at least two. We get that $B_{k,j+1}$ is dominated by $D \cap V(I_k)$ in I_k , while $B_{k'',j+1}$ is dominated by $D \cap I_{k-1}$ in I_{k-1} , see Figure 11. Now, by Theorem 14 (iii), for any $i \in [t]$, $B_{i,j+1} \not\subseteq N(V(I_i) \setminus B_{i,j+1})$, yielding that both $D \cap B_{k,j+1}$ and $D \cap B_{k'',j+1}$ are non-empty. This is absurd since D must intersect C_{j+1} in exactly one vertex, so both $(r_{2j})_j$ and $(r_{2j+1})_j$ are constant.

Now, let us call k the row choice for even columns and k' the row choice for odd ones, if N is odd, $k = k'$ and the choice is constant. If $k \neq k'$, we can choose without loss of generality three consecutive columns C_j, C_{j+1}, C_{j+2} such that $r_j = r_{j+2} = k$ and $r_{j+1} = k'$ with $k' > k$. Then, we proceed similarly to the last paragraph. Since $k' > k$, the class $B_{k,j+1}$ is not dominated by d_j, d_{j+2} , by construction of the half-graph. Therefore, it is necessarily dominated by $D \cap V(I_k)$ both in H and in I_k . Again, Theorem 14 (iii) yields that $B_{k,j+1} \not\subseteq N(V(I_k) \setminus B_{k,j+1})$ so $d_{j+1} \in B_{k,j+1}$, contradicting with $d_{j+1} \in B_{k',j+1}$. Therefore $k = k'$ and the sequence of row choices $(r_j)_j$ is constant.

The latter yields $k \in [t]$ such that $D \subseteq V(I_k)$, by assumption D dominates $V(I_k)$ on H , while no edges are added between two vertices in $V(I_k)$ by our construction, thus D is a dominating set of I_k . ◀

► **Theorem 1.** *Unless $\text{coNP} \subseteq \text{NP}/\text{poly}$, 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.*

Proof. Lemma 19 proves that the OR-cross-compositions described in the beginning of Section 5 is correct. The twin-width bound of the composed instance is given in Lemma 18. We can thus invoke Theorem 13 in order to obtain the desired result. ◀



■ **Figure 11** In the middle column C_{j+1} , pink classes are those dominated by the choice of d_j and d_{j+2} in the neighboring columns. At least two classes, highlighted in green and corresponding to the interval $[k'', k]$, would need to be dominated by a single vertex within the column, which cannot be done.

6 Polynomial kernels

We present two simple kernelization algorithms for CONNECTED k -VERTEX COVER and CAPACITATED k -VERTEX COVER with kernels on $O(k^2)$ vertices. We also provide an improved kernel on $O(k^{1.5})$ vertices for CONNECTED k -VERTEX COVER.

A folklore in the context of k -VERTEX COVER kernelization is that keeping at most $k + 1$ twin vertices is a safe reduction rule. Hence any module outside of a vertex cover X can be trimmed down to $k + 1$ vertices. Similar reduction rules are also known for CONNECTED k -VERTEX COVER and CAPACITATED k -VERTEX COVER [16]. Therefore to get a polynomial kernel, one simply needs to polynomially bound the number of modules outside X , or equivalently, the number of distinct neighborhoods in X .

This will be done by proving 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.

To state the next two theorems, we need to introduce some vocabulary on 0, 1-matrices. A *row* (resp. *column*) *partition* of a matrix M is a partition of its row set (resp. column set). A *row* (resp. *column*) *division* is a row (resp. column) partition where every part is a consecutive set of rows (resp. columns), or *interval*. A *division* of a matrix M is a pair $(\mathcal{R}, \mathcal{C})$ where \mathcal{R} is a row division and \mathcal{C} is a column division. A *cell* or *zone* of a matrix M with division $(\mathcal{R}, \mathcal{C})$ 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 *t-grid minor* of a 0, 1-matrix M is a t -division of M where every cell contains at least one 1-entry. 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. A *t-mixed minor* of a matrix M is a t -division of M where every cell is mixed. It can be observed that this is equivalent to the existence of a t -division where every cell contains a corner [6].

We first recall the Marcus-Tardos theorem, a celebrated result in combinatorics.

► **Theorem 20** ([41]). *For every integer t , there is some c_t such that every $n \times m$ 0, 1-matrix M with at least $c_t \max(n, m)$ entries 1 has a t -grid minor.*

In what follows c_t will always denote the bound in Theorem 20. The best upper bound for c_t is currently $8/3(t + 1)2^{2t}$ [11], that we take as the definition of c_t .

The following result is shown in the first paper of the series dedicated to twin-width.

► **Theorem 21** ([6]). *If, for every vertex ordering, the adjacency matrix of a graph G has a $2t + 2$ -mixed minor, then G has twin-width larger than t .*

We prove that in a 0, 1-matrix of small twin-width, the number of *distinct* columns is linearly bounded in the total number of rows.

► **Theorem 22**. *Let t be a positive integer. Let M be a 0, 1-matrix with s rows, and at least $2^{4c_{2t+2}} \cdot s$ distinct columns. Then M has a $2t + 2$ -mixed minor.*

Proof. We denote by R and C the sets of rows and columns, respectively, of M . Without loss of generality, we may assume that all the columns of M are distinct. We consider a division of C into s parts $\{C_1, \dots, C_s\}$, each C_i consisting of $2^{4c_{2t+2}}$ consecutive (distinct) columns. Note that the submatrix of M consisting of the columns C_i has rank at least $4c_{2t+2}$ in the binary field \mathbb{F}_2 , for all $i \in [s]$.

Therefore there is a row division of $M[R, C_i]$ into at least $2c_{2t+2}$ (row) parts, each zone of which has rank at least 2, thus is mixed, and, by an observation in [6], contains a corner. These $2c_{2t+2}$ corners in $M[R, C_i]$ are on pairwise disjoint pairs of consecutive rows. 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, \dots, C_s have at least c_{2t+2} mixed zones when divided by, say, \mathcal{R}^1 . Consider the column division $\{C'_1, \dots, C'_{s'}\}$ with $s' \geq \lceil s/2 \rceil$, coarsening of $\{C_1, \dots, C_s\}$ such that each part C'_j contains exactly one column part C_i with the property of the previous sentence.

Let M' be a $|\mathcal{R}^1| \times s'$ 0, 1-matrix with a 1-entry at positions where the cell of $(\mathcal{R}^1, \{C'_1, \dots, C'_{s'}\})$ is mixed, and a 0-entry otherwise. Note that M' has at most $\lfloor s/2 \rfloor$ rows and $s' \geq \lceil s/2 \rceil$ columns. Moreover each column of M' contains by design at least c_{2t+2} 1-entries. By Theorem 20, M' admits a $(2t+2)$ -grid minor. Thus M has a $(2t+2)$ -mixed minor, which implies that M has twin-width larger than t by Theorem 21. ◀

We conclude the following.

► **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 $2^{4c_{2t+2}}|X|$.*

Proof. 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 . By Theorem 22, it implies that M has a $2t+2$ -mixed minor. By Theorem 21, this in turn implies that G has twin-width more than t . ◀

6.1 Quadratic vertex kernels

For completeness, we state and prove the folklore reduction rule for k -VERTEX COVER variants.

► **Reduction Rule 1** (Reduction Rule for CONNECTED k -VERTEX COVER). Let X be a vertex cover of G . If there is a set $S \subseteq V(G) \setminus X$ with the same neighborhood in X and $|S| > k$, delete a vertex of S .

► **Lemma 23.** *Let G be a graph and G' be a graph obtained by applying Reduction Rule 1. Then (G, k) is a YES-instance to CONNECTED k -VERTEX COVER if and only if (G', k) is.*

Proof. Let $s \in S$ be the vertex such that $G' := G - s$. Suppose that T is a connected vertex cover of G of size at most k . If $s \notin T$, then T is a connected vertex cover of G' . If $s \in T$, there is at least one vertex s' in $S \setminus T$. Now $T' := (T \setminus \{s\}) \cup \{s'\}$ is a connected vertex cover of G' of size $|T|$.

Conversely let T be a connected vertex cover of G' of size at most k . We claim that T is also a connected vertex cover in G . As $G[T]$ is connected and $|T| \leq k$, we shall just check that all the edges of G are covered by T . This is because $S \setminus s$ cannot be totally included in T ; if so, T is not connected because $S \setminus s$ consists of (at least) k pairwise independent vertices. Therefore, the common neighborhood $N(S)$ should be included in T . ◀

► **Proposition 24.** *CONNECTED k -VERTEX COVER admits a kernel on $O_t(k^2)$ vertices when the input graphs have twin-width at most t .*

Proof. Let (G, k) be an instance of CONNECTED k -VERTEX COVER and let X be a vertex cover obtained using a 2-approximation algorithm for MIN VERTEX COVER (given by any maximal matching). If $|X| \geq 2k + 1$, an optimal vertex cover and thus an optimal connected vertex cover has size at least $k + 1$. Then one can correctly output a trivial NO-instance. Henceforth we assume that $|X| \leq 2k$. Apply Reduction Rule 1 exhaustively, and let (G', k) be the resulting instance. By Lemma 3 and by the construction of G' , $V(G') \setminus X$ can be partitioned into at most $2^{4c_2+2} \cdot |X|$ modules and each module consists of at most $k + 1$ vertices. Therefore, $V(G')$ has at most $c_t \cdot (k + 1) + 2k$ vertices. That (G', k) is an equivalent to (G, k) is implied by Lemma 23. \blacktriangleleft

► **Reduction Rule 2** (Reduction Rule for CAPACITATED k -VERTEX COVER). Let $(G, k, c : V(G) \rightarrow \mathbb{N})$ be an instance to CAPACITATED k -VERTEX COVER and X be a vertex cover of G . If there is a set $S \subseteq V(G) \setminus X$ with the same neighborhood in X and $|S| > k + 1$, delete a vertex of $s \in S$ with the minimum capacity and decrease the capacity of each neighbor of s by one.

► **Lemma 25.** *Let (G, k, c) be an instance to CAPACITATED k -VERTEX COVER and (G', k, c') be an instance obtained by applying Reduction Rule 2. Then (G, k, c) is a YES-instance if and only if (G', k, c') is.*

Proof. Again let s be such that $G' := G - s$. We first check the forward direction. A c -capacitated vertex cover of G is not necessarily a c' -capacitated vertex cover of G' as we decreased the capacities of neighbors of s . However we claim that if G admits a c -capacitated vertex cover T of size at most k , then it also admits a c -capacitated vertex cover of G of size at most k not containing s . Indeed, if $s \in T$, then there exists a vertex $s' \in S \setminus T$ and not that all neighbors of s' (thus s) are in T . Therefore, $(T \setminus \{s\}) \cup \{s'\}$ is a vertex cover of G of size at most k . It is further c -capacitated because $c(s') \geq c(s)$. Now we may assume that T does not contain s . Set T is a c' -capacitated vertex cover of G' because the decrease of capacity of each vertex $v \in N(s)$ by one unit is canceled out by the absence of edge sv in G' .

Conversely, if there is a c' -capacitated vertex cover T of G' of size at most k , note that all neighbors of $S \setminus \{s\}$ must be in T . Therefore, T is a c -capacitated vertex cover of G where each edge incident with s is covered by the residual capacity of the other endpoint. \blacktriangleleft

Therefore, CAPACITATED k -VERTEX COVER admits a quadratic vertex kernel, which follows the proof of Proposition 24 verbatim.

► **Proposition 26.** *CAPACITATED k -VERTEX COVER admits a kernel on $O_t(k^2)$ vertices when the input graphs have twin-width at most t .*

Theorem 2 is a direct consequence of Propositions 24 and 26.

6.2 Improved kernel for Connected k -Vertex Cover

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, and a simple linear kernel when the input class is further restricted to be sparse, that is, $K_{s,s}$ -free.

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, \dots, 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 .

► **Reduction Rule 3.** If there is $i \in [q]$ with $X_i \neq \emptyset$ and $|Y_i| \geq |X_i| + 2$, then delete a vertex of Y_i .

► **Lemma 27.** *Let G be a graph and G' be a graph obtained by applying Reduction Rule 3. Then (G, k) is a YES-instance to CONNECTED k -VERTEX COVER if and only if (G', k) is.*

Proof. Let $Y_i \subseteq V(G) \setminus X$ be a module within which Reduction Rule 3 is applied, and $y \in Y_i$ be the deleted vertex, that is, such that $G' := G - y$. Set $Y'_i := Y_i \setminus \{y\}$. We first observe that the vertices of X^b are mandatory for any feasible solution to (G', k) . Let T be an arbitrary connected vertex cover of G' of size at most k .

▷ Claim 28. $X^b \subseteq T$.

PROOF OF THE CLAIM: Suppose some vertex $w \in X^b$ is not contained in T and thus all its neighbors belong to T . Note that w has at least k neighbors outside X in G' , even when the deleted vertex y was a neighbor of w in G . We conclude that $T = N_{G'}(w) \setminus X$, which contradicts that $G'[T]$ is connected. ◊

Because there are all possible edges between Y'_i and X_i , T must contain at least one of Y'_i and X_i entirely. The next claim says that T can be assumed to fully contain X_i .

▷ Claim 29. Let y' be a vertex of Y'_i . If $Y'_i \subseteq T$, then $T' := (T \setminus Y'_i) \cup (X_i \cup \{y'\})$ is a connected vertex cover of G' of size at most k .

PROOF OF THE CLAIM: Notice that $X^b \cup X_i$ is the neighborhood of Y'_i in G' . By Claim 28, T' is clearly a vertex cover of G' and y' provides any connection between a pair of vertices in T' that Y'_i used to provide. Finally, that Y'_i is obtained from Y_i after Reduction Rule 3 means that $|Y'_i| \geq |X_i| + 1 = |X_i \cup \{y'\}|$, and thus $|T'| \leq |T|$. ◊

Due to Claim 29, we may assume that T contains X_i . Now for any edge xy incident with the deleted vertex y , x is either in X^b or X_i , and thus the edge xy is covered by T by Claim 28 and the assumption $X_i \subseteq T$. It follows that T is a feasible solution to (G, k) .

Conversely, let T be a connected vertex cover of (G, k) of size at most k . Following the same arguments as the above claims, one can easily check that T must contain both X^b and X_i , and contains at most one vertex of Y_i . As we can modify T so that it does not contain y , T is a feasible solution to (G', k) . ◀

► **Proposition 30.** *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 every connected component of G contains at least one edge. If G contains more than one connected component, then clearly (G, k) is a NO-instance and we output the 4-vertex graph with two isolated edges. Therefore, we can assume that G is connected. With a 2-approximation algorithm for VERTEX COVER, one can find a vertex cover X of G and assume that $|X| \leq 2k$. 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 3 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 3 with the vertex cover X at hand. We classify X into X^b and X^s as before, and Y_1, \dots, 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 $q \leq 2 \cdot 2^{4c_{2t+2}k}$. 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

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

Suppose that $|Y| - q > 2 \cdot 2^{2c_{2t+2}} \cdot k^{1.5}$. Now,

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

and hence there are more than $2k^2$ edges between X^s and Y . With $|X^s| \leq 2k$, 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, G' has at most

$$|X| + |Y| \leq 2k + 2 \cdot 2^{2c_{2t+2}} \cdot k^{1.5} + q \leq 2k + 2 \cdot 2^{2c_{2t+2}} \cdot k^{1.5} + 2 \cdot 2^{4c_{2t+2}} \cdot k = O_t(k^{1.5})$$

vertices as claimed. \blacktriangleleft

Note that the proof of Proposition 30 only uses the fact that the input graphs have VC density at most 1, so we in fact established Theorem 4.

7 Graphs of twin-width 1

Cycles on at least five vertices and their complements have twin-width 2. Thus graphs of twin-width 1 do not have induced cycles of length at least 5, nor their complements. In particular, they are perfect. Hence computing the independence, clique, and coloring numbers can be done in polynomial-time [34].

Graphs of twin-width 0 are cographs, and can be recognized in linear time [36]. It turns out that twin-width 1 graphs are also efficiently recognizable. For that, we need the following technical lemma.

► Lemma 31. *Let G be a prime graph of twin-width 1, and let $G = G_n, \dots, G_2, G_1$ be a 1-contraction sequence. Then the trigraphs G_i have exactly one red edge, except for G_n and G_1 which have no red edge.*

Proof. The idea of the proof is to start with the 3-vertex trigraph G_3 and to rewind the contraction sequence back to $G_n = G$. Let a, b, c be the three vertices of G_3 . Since G_3 is a 1-trigraph, there is at most one red edge in G_3 . Since G is a prime graph, there has to be at least one red edge in G_3 , say between b and c . Furthermore, a has to be a vertex in G , otherwise $a(G)$ is a module of size at least two in G . Moreover, a is adjacent to, say, b , but non-adjacent to the other vertex, c since otherwise $b(G) \cup c(G)$ is a module.

We show by induction the following property (P) for $i \in [2, n-1]$:

Property P: in G_i , there is exactly one red edge, say, uv (u and v can depend on i) and all the vertices that are not u nor v are original vertices of G .

We already observed that Property P holds for G_3 . It also holds for G_2 because G contains no non-trivial module. Assuming that G_i satisfies Property P for some $i \in [3, n-2]$, let us prove that G_{i+1} also satisfies Property P. Let uv be the red edge in G_i . By assumption, all the other vertices are vertices in G . Hence, the graph G_{i+1} is obtained from G_i by splitting u or v , say u , into two vertices u_1 and u_2 . Now, there is at most one red edge in $G_{i+1}[\{u_1, u_2, v\}]$

since otherwise a vertex shall be incident with two red edges. Recall that at least one vertex of G_{i+1} , thus one of $\{u_1, u_2, v\}$ is not an original vertex of G . Such a vertex must be incident with a red edge because it is not a module. Therefore, $G_{i+1}[\{u_1, u_2, v\}]$ has a unique red edge. Without loss of generality, assume that the red edge is u_2v . Again, u_1 should be an original vertex of G , i.e., $|u_1(G)| = 1$, because it has no incident red edge and therefore $u_1(G)$ forms a module. So Property P holds for G_{i+1} . ◀

We can efficiently detect graphs of twin-width 1 using the previous lemma and an inductive scheme.

► **Theorem 32.** *Twin-width 1 graphs can be recognized in polynomial time. Moreover, a 1-contraction sequence where each trigraph has at most one red edge can be constructed in polynomial time.*

Proof. Let G be a given graph. We show by induction on n , the number of vertices of G , how to decide whether or not it has twin-width 1. If G has one vertex, it has twin-width 1 (even 0). If G is not a prime graph, then each maximal module of G and the quotient graph of G by its (pairwise disjoint) maximal modules all have at most $n - 1$ vertices. So Lemma 9 and the inductive step allow us to conclude. We can therefore assume that G is a prime graph.

As a technicality, it is more convenient to assume that G is a trigraph with exactly one red edge. We can in quadratic time guess the first 1-contraction (which creates exactly one red edge since the contractions of twins are safe by Observation 6). Recall that G has exactly one red edge, say uv . Let us establish the next 1-contraction. By Lemma 31, we can limit ourselves to contractions which keep the number of red edges to at most one. Therefore, the next contraction cannot be between two vertices in $V(G) \setminus \{u, v\}$. Indeed, that would either add at least one new red edge to the trigraph or the two contracted vertices would be twins, and as such, in a module of size at least two.

Any contraction between $w \in V(G) \setminus \{u, v\}$ and u (or v) that creates a new red edge can be disregarded again by Lemma 31. The only way that such a contraction does not increase the number of red edges is if the contraction actually results in the induced subgraph $G - \{w\}$. By Observation 6, this contraction can be safely performed. When no such safe contraction is available, we know that only the current red edge can possibly be contracted. Hence, we can safely make this contraction, and we obtain a new graph with $n - 1$ vertices. Observe that in this process, we never branch on two possible contractions.

The second statement immediately follows by concatenating the 1-contraction sequence of modules and the quotient graph, each of twin-width at most 1, as in Lemma 9. ◀

We remark that polynomial time algorithms are not hard to obtain for, say, CONNECTED k -VERTEX COVER, k -DOMINATING SET or CONNECTED k -DS, when a 1-contraction sequence as in Theorem 32 is given. Consider CONNECTED k -VERTEX COVER for instance. For each red component, we "trace" a partial solution X , i.e., a vertex set covering all original edges whose endpoints belong to the said red component, by recording for each vertex u of the component whether $u(G)$ is fully contained in X , intersecting (but not fully contained in) X , or not intersecting at all. Among all partial solutions leaving the same trace over the red component, we remember the minimum size of a solution. One can readily check that the optimal value of a partial solution per trace can be updated in constant time for each newly created red component.

Readers may further notice that this dynamic programming works not only when there is at most one red edge, but more generally when a red component size is bounded by a

constant. In a manuscript under preparation, the authors show that MSOL_1 -expressible problems can be solved in $f(d) \cdot n^{O(1)}$ -time if the input graph admits a contraction sequence in which each red component has size at most d .

Acknowledgments

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.

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. URL: <https://arxiv.org/abs/2007.14161>, 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. URL: <https://arxiv.org/abs/2102.03117>, 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 Köneemann, 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 Josef Cibulka and Jan Kyncl. Füredi-Hajnal limits are typically subexponential. *CoRR*, abs/1607.07491, 2016. URL: <http://arxiv.org/abs/1607.07491>, arXiv:1607.07491.
- 12 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)90358-0.
- 13 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.
- 14 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.
- 15 Marek Cygan, Fedor V Fomin, Lukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 4. Springer, 2015.
 - 16 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.
 - 17 Marek Cygan, Geevarghese Philip, Marcin Pilipczuk, Michał 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.
 - 18 Marek Cygan, Marcin Pilipczuk, Michał 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.
 - 19 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.
 - 20 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.
 - 21 Michael Dom, Daniel Lokshantov, and Saket Saurabh. Kernelization Lower Bounds Through Colors and IDs. *ACM Trans. Algorithms*, 11(2):13:1–13:20, 2014. doi:10.1145/2650261.
 - 22 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.
 - 23 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.
 - 24 Pål Grønås Drange, Markus Sortland Dregi, Fedor V. Fomin, Stephan Kreutzer, Daniel Lokshantov, Marcin Pilipczuk, Michał 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.
 - 25 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.
 - 26 Fedor V. Fomin, Daniel Lokshantov, 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.
 - 27 Fedor V. Fomin, Daniel Lokshantov, Saket Saurabh, and Dimitrios M. Thilikos. Bidimensionality and kernels. *SIAM J. Comput.*, 49(6):1397–1422, 2020. doi:10.1137/16M1080264.
 - 28 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.

- 29 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.
- 30 Jakub Gajarský, Petr Hliněný, Jan Obdržálek, 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.
- 31 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.
- 32 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.
- 33 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.
- 34 Martin Grötschel, László Lovász, and Alexander Schrijver. Corrigendum to our paper "the ellipsoid method and its consequences in combinatorial optimization". *Comb.*, 4(4):291–295, 1984. doi:10.1007/BF02579139.
- 35 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.
- 36 Michel Habib and Christophe Paul. A simple linear time algorithm for cograph recognition. *Discret. Appl. Math.*, 145(2):183–197, 2005. doi:10.1016/j.dam.2004.01.011.
- 37 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.
- 38 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.
- 39 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.
- 40 David Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.*, 11(2):329–343, 1982. doi:10.1137/0211025.
- 41 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.
- 42 Jaroslav Nešetřil 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.
- 43 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.
- 44 Wojciech Przybyszewski and Szymon Toruńczyk. personal communication, 2021.
- 45 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.

- 46 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.

