



Introducing lop-Kernels: A Framework for Kernelization Lower Bounds

Júlio Araújo, Marin Bougeret, Victor Campos, Ignasi Sau

► To cite this version:

Júlio Araújo, Marin Bougeret, Victor Campos, Ignasi Sau. Introducing lop-Kernels: A Framework for Kernelization Lower Bounds. *Algorithmica*, 2022, 84 (11), pp.3365-3406. 10.1007/s00453-022-00979-z . lirmm-03991247

HAL Id: lirmm-03991247

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-03991247>

Submitted on 17 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Introducing **lop**-kernels: a framework for kernelization lower bounds

Júlio Araújo 

Departamento de Matemática, Universidade Federal do Ceará, Fortaleza, Brazil
julio@mat.ufc.br

Marin Bougeret 

LIRMM, Université de Montpellier, CNRS, Montpellier, France
marin.bougeret@lirmm.fr

Victor Campos 

Departamento de Computação, Universidade Federal do Ceará, Fortaleza, Brazil
victoitor@ufc.br

Ignasi Sau 

LIRMM, Université de Montpellier, CNRS, Montpellier, France
ignasi.sau@lirmm.fr

Abstract

In the MAXIMUM MINIMAL VERTEX COVER (MMVC) problem, we are given a graph G and a positive integer k , and the objective is to decide whether G contains a minimal vertex cover of size at least k . Motivated by the kernelization of MMVC with parameter k , our main contribution is to introduce a simple general framework to obtain kernelization lower bounds for a certain type of kernels for optimization problems, which we call **lop-kernels**. Informally, this type of kernels is required to preserve large optimal solutions in the reduced instance, and captures the vast majority of existing kernels in the literature.

As a consequence of this framework, we show that the trivial quadratic kernel for MMVC is essentially optimal, answering a question of Boria et al. [Discret. Appl. Math. 2015], and that the known cubic kernel for MAXIMUM MINIMAL FEEDBACK VERTEX SET is also essentially optimal. We present further applications for TREE DELETION SET and for MAXIMUM INDEPENDENT SET on K_t -free graphs.

Back to the MMVC problem, given the (plausible) non-existence of subquadratic kernels for MMVC on general graphs, we provide subquadratic kernels on H -free graphs for several graphs H , such as the bull, the paw, or the complete graphs, by making use of the Erdős-Hajnal property. Finally, we prove that MMVC does not admit polynomial kernels parameterized by the size of a minimum vertex cover of the input graph, even on bipartite graphs, unless $\text{NP} \subseteq \text{coNP}/\text{poly}$.

2012 ACM Subject Classification Design and analysis of algorithms $\rightarrow$ Fixed parameter tractability.

Keywords and phrases Parameterized complexity, polynomial kernel, kernelization lower bound, maximum minimal vertex cover, Erdős-Hajnal property, induced subgraphs.

Related Version Some of the results of this article appeared in the preliminary version published in the proceedings of IPEC 2021 [4]. This article is available at <https://arxiv.org/abs/2102.02484>.

Funding *Júlio Araújo*: CNPq-Pq 304478/2018-0, CAPES-PrInt 88887.466468/2019-00 and CAPES-STIC-AmSud 88881.569474/2020-01.

Victor Campos: FUNCAP - PNE-011200061.01.00/16.

Ignasi Sau: DEMOGRAPH (ANR-16-CE40-0028), ESIGMA (ANR-17-CE23-0010), ELIT (ANR-20-CE48-0008-01), and French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027).

1 Introduction

A *vertex cover* in a graph G is a subset of vertices containing at least one endpoint of every edge. In the associated optimization problem, called MINIMUM VERTEX COVER, the objective is to find, given an input graph G , a vertex cover in G of minimum size. This problem has been one of the leitmotifs of the area of parameterized complexity [25,30], serving as a test bed for many of the most fundamental techniques. An instance of a *parameterized problem* is of the form (x, k) , where x is the total input (typically, a graph) and k is a positive integer called the *parameter*. The crucial notion is that of *fixed-parameter tractable* algorithm, FPT for short, which is an algorithm deciding whether (x, k) is a positive instance in time $f(k) \cdot |x|^{\mathcal{O}(1)}$, where f is a computable function depending only on k . In the parameterized VERTEX COVER problem, we are given a graph G and an integer parameter k , and the objective is to decide whether G contains a vertex cover of size at most k . One of the main fields within parameterized complexity is *kernelization* [35], where the objective is to decide whether an instance (x, k) of a parameterized problem can be transformed in polynomial time into an equivalent instance (x', k') whose total size is bounded by a function of k ; the reduced instance is called a *kernel*, and finding kernels of small size, typically polynomial or even linear in k in the best case, is one of the most active areas of parameterized complexity.

Considering the “max-min” version of minimization problems, that is, maximizing the size of a *minimal* solution of the corresponding problem, is a natural approach that has been applied to several problems such as DOMINATING SET [7,32] (whose “max-min” version is called UPPER DOMINATION), FEEDBACK VERTEX SET [31], or HITTING SET [5,26]. The initial motivation of this article is the “max-min” version of MINIMUM VERTEX COVER, called MAXIMUM MINIMAL VERTEX COVER, or just MMVC for short.

Previous work. In his habilitation, Fernau [34] presented FPT algorithms for MMVC as well as some results about its kernelization parameterized by the solution size k . It is easy to note, as observed in [34], that the problem admits a kernel with at most k^2 vertices: if some vertex has degree at least k , we can safely answer “yes” (cf. Lemma 2 for a proof); otherwise, the maximum degree is at most $k - 1$, and it follows that every instance without isolated vertices (which may be safely removed) that has at least k^2 vertices is a yes-instance, hence we have a trivial kernel with at most k^2 vertices. Fernau [34] presented a kernel with at most $4k$ vertices for MMVC restricted to planar instances using the algorithmic version of the Four Color Theorem [53], and claimed in [34, Corollary 4.25] a kernel with at most $2k$ vertices on general graphs using spanning trees. Unfortunately, this latter kernelization algorithm is incorrect, as we discuss in Section 6.

Boria et al. [18] initiated a study of the complexity of MMVC and presented a number of results, in particular a polynomial-time approximation algorithm with ratio $n^{1/2}$ on n -vertex graphs, and showed that, unless $P = NP$, no polynomial-time approximation algorithm with ratio $n^{1/2-\varepsilon}$ exists for any $\varepsilon > 0$. They also presented FPT algorithms for MMVC for several choices of the parameters such as the treewidth, the size of a maximum matching, or the size of a minimum vertex cover of the input graph. The authors asked explicitly whether kernels of size $o(k^2)$ exist for MMVC parameterized by k .

Zehavi [57] presented tight FPT algorithms, under the Strong Exponential Time Hypothesis, for MMVC and its weighted version parameterized by the size of a minimum vertex cover. Recently, Bonnet and Paschos [16] and Bonnet et al. [15] considered the inapproximability of MMVC in subexponential time.

Note that the MMVC problem is the dual of the well-studied MINIMUM INDEPENDENT DOMINATING SET problem (to see this, note that the complement of any minimal vertex

cover is an independent dominating set), which has applications in wireless and ad-hoc networks [46]. We refer to the survey of Goddard and Henning [39].

Our results and techniques. The starting motivation of this article is the kernelization of the MMVC problem, which has been almost unexplored so far in the literature. This initial motivation has resulted in a general framework that can be applied to a broad class of optimization problems in order to derive kernelization lower bounds.

Namely, motivated by the question of Boria et al. [18] about the existence of subquadratic kernels for MMVC, we introduce a generic framework to obtain kernelization lower bounds for a “certain type” of kernels for parameterized maximization or minimization problems (in particular, for MMVC), based on a hypothesis that guarantees an inapproximability result, typically $P \neq NP$. Informally, by “certain type” we mean kernelization algorithms that, in polynomial time, either decide the instance (by answering “yes” or “no”) or produce an equivalent instance of the considered problem in which the value of an optimal solution is “preserved”, in the sense that it may drop only by the drop suffered by the parameter; see [Subsection 3.1](#) for the formal details for the case of maximization problems. We call such kernels *large optimal preserving kernels*, or **lop-kernels** for short. Even if this type of kernels may seem restrictive, in particular we are not aware of any “non-artificial” kernel for a maximization problem, such as those that have become nowadays standard [35], which is *not* a lop-kernel. (We do have such an example for a minimization problem, as discussed later.) The core idea of our approach is to show that a lop-kernel yields a polynomial-time approximation algorithm whose ratio depends on the size (and most importantly, on the degree) of the kernel, and to use known inapproximability results to obtain the desired lower bound.

We present our framework of lop-kernels separately for maximization ([Section 3](#)) and minimization ([Section 4](#)) problems. Even if both versions are similar, they are *not* totally symmetric, and a number of technical differences pop up; we discuss them in detail as they appear in [Section 4](#). Our general result is stated in [Theorem 13](#) and [Theorem 24](#) for maximization and minimization problems, respectively. In order to be able to apply our framework to an optimization problem, we need it to be “well-behaved”, a mild condition defined in [Section 3](#) and [Section 4](#) that, for instance, for vertex-optimization problems is weaker than their decision version being in NP. Also, our results distinguish the existence of constructive or non-constructive approximation algorithms. Since our framework seems to particularly fit vertex-optimization problems, we present the particular cases of [Theorem 13](#) and [Theorem 24](#) for vertex-maximization and vertex-minimization problems in [Theorem 10](#) and [Theorem 21](#), respectively. In order to ease the application of our results to concrete problems, we provide in [Corollary 11](#) and [Corollary 22](#) the “contrapositive” versions of [Theorem 10](#) and [Theorem 21](#), respectively.

Applications of our framework. Combining [Corollary 11](#) with the $\mathcal{O}(n^{\frac{1}{2}-\varepsilon})$ -inapproximability result for MMVC by Boria et al. [18] immediately rules out (cf. [Corollary 26](#)) the existence of a lop-kernel for MMVC with $\mathcal{O}(k^{2-\varepsilon})$ vertices for any $\varepsilon > 0$, unless $P = NP$. Thus, while [Corollary 26](#) does not completely rule out the existence of subquadratic kernels for MMVC, it tells that, if such a kernel exists, it should consist of “non-standard” reduction rules.

Interestingly, our framework has consequences beyond the MMVC problem. One of them concerns the MAXIMUM MINIMAL FEEDBACK VERTEX SET (MMFVS) problem, defined in the natural way. Dublois et al. [31] recently provided a cubic kernel for MMFVS parameterized by the solution size, and proved that the problem does not admit an $\mathcal{O}(n^{\frac{2}{3}-\varepsilon})$ -approximation algorithm for any $\varepsilon > 0$, unless $P = NP$. By applying [Corollary 11](#) we directly obtain ([Corollary 27](#)) that the cubic kernel of Dublois et al. [31] is “essentially” optimal.

Another application of our results deals with the TREE DELETION SET problem. In this case, the fact that this problem does not admit a polynomial-time $\mathcal{O}(n^{1-\varepsilon})$ -approximation for any $\varepsilon > 0$ unless $P \neq NP$ [56] implies, together with Corollary 22, that TREE DELETION SET parameterized by the solution size does not admit a polynomial lop-kernel, unless $P = NP$ (Corollary 28). However, TREE DELETION SET does admit a polynomial kernel with $\mathcal{O}(k^4)$ vertices [38]. Therefore, this polynomial kernel cannot be a lop-kernel, and so far it constitutes the only non-artificial example of non-lop-kernel that we are aware of.

Our last application concerns the MAXIMUM INDEPENDENT SET problem restricted to K_t -free graphs. In particular, we show (Corollary 30) that a lop-kernel with $\mathcal{O}(k^{t-1-\varepsilon})$ vertices for MAXIMUM INDEPENDENT SET on K_t -free graphs would improve the best known approximation ratio $n^{\frac{t-2}{t-1}}$ that follows from Ramsey's theorem [52]. Finally, generalizing a conjecture of Bonnet et al. [17], we conjecture that for every fixed graph H , the MAXIMUM INDEPENDENT SET problem restricted to H -free graphs admits a polynomial lop-kernel.

Comparison with other frameworks. Compared to existing frameworks to obtain lower bounds on kernelization, such as cross-compositions [10,12], weak compositions [27,28,44], polynomial parameter transformations [8,13], or techniques to obtain lower bounds on the coefficients of linear kernels [21], or that relate approximation and kernelization [1,9,41,48,50], our approach has the advantages that it is quite simple, straightforward to apply, and relies on the same hypothesis on which the corresponding inapproximability result is based, typically the standard hypothesis that $P \neq NP$. On the negative side, it has the following two drawbacks. The first one is that, in order to obtain a non-trivial lower bound on the kernel size, it can only be applied to problems which are quite hard to approximate, for example within a factor $\mathcal{O}(n^{r-\varepsilon})$ for some constant $r > 0$, as it is the case of MMVC and MMFVS. The second, and probably most important, drawback of our techniques is that they are able to rule out the existence of what we call lop-kernels of certain sizes, but smaller non-standard kernels that do not preserve the value of large optimal solutions might, a priori, still exist (as it is the case for TREE DELETION SET, as discussed above). Hence, since our framework seems to be orthogonal to existing ones, we think that it adds to the above list of techniques to obtain kernelization lower bounds.

Other results on the kernelization of MMVC. Coming back to the MMVC problem parameterized by the solution size, given the above negative result on general graphs, we identify graph classes where MMVC is still NP-hard and admits a subquadratic kernel. In particular, we deal with graph classes defined by excluding an *induced* subgraph H that satisfies the *Erdős-Hajnal property* [33], that is, for which there exists a constant $\delta > 0$ such that every H -free graph on n vertices contains either a clique or an independent set of size n^δ . In particular, we present a kernel for MMVC with $\mathcal{O}(k^{7/4})$ vertices on the well-studied class of bull-free graphs (Theorem 32), with $\mathcal{O}(k^{\frac{2t-3}{t-1}})$ vertices on K_t -free graphs for every $t \geq 3$ (Theorem 34), and with $\mathcal{O}(k^{5/3})$ vertices on paw-free graphs (Theorem 37). To the best of our knowledge, this is the first time that the Erdős-Hajnal property is used to obtain polynomial kernels (we would like to note that it was used by Kratsch et al. [49] to obtain kernelization lower bounds).

Our strategy to obtain these subquadratic kernels on H -free graphs is as follows. By the high-degree rule mentioned above, given an instance (G, k) , we may assume that the maximum degree of G is at most $k - 1$. We find greedily a minimal vertex cover X of G . If $|X| \geq k$ we are done, so we may assume that $|X| \leq k - 1$, hence the goal is to bound the size of $S := V(G) \setminus X$. Using that $G[X]$ is also H -free, the Erdős-Hajnal property implies (Lemma 31) that X can be partitioned in polynomial time into a sublinear (in k) number of independent sets and cliques. Since S is an independent set and we may assume

that G has no isolated vertices, in order to bound $|S|$ by a subquadratic function of k , it is enough to show that, for each of the sublinearly many cliques or independent sets Y that partition X , its neighborhood in S has size $\mathcal{O}(k)$. This is easy if Y is an independent set: if $|N_S(Y)| \geq k$ we can conclude that (G, k) is a yes-instance (Lemma 2), so we may assume that $|N_S(Y)| \leq k - 1$. The case where Y is a clique is more interesting, and we need ad-hoc arguments depending on each particular excluded induced subgraph H .

We also present several positive results for MMVC restricted to other particular graph classes, such as $K_{1,t}$ -free graphs (Lemma 38), graph classes with bounded chromatic number (Lemma 39), or graph classes with bounded cliquewidth (Observation 40).

Finally, we show (Theorem 42) that MMVC, parameterized by the size of a minimum vertex cover (or of a maximum matching) of the input graph, does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$, even restricted to bipartite graphs. This result complements the FPT algorithms for MMVC under these parameterizations given by Boria et al. [18] and Zehavi [57], and shows that, in what concerns the existence of polynomial kernels for MMVC, the most natural structural parameters smaller than the solution size are not large enough to yield polynomial kernels (note that the treewidth of any graph is at most one more than its vertex cover number, hence our result rules out the existence of polynomial kernels for MMVC parameterized by treewidth as well). The proof consists of a polynomial parameter transformation from MONOTONE SAT parameterized by the number of variables. In particular, our reduction yields also the NP-hardness of MMVC on bipartite graphs, which provides an alternative proof to the one of Boliac and Lozin [14] via the NP-hardness of MINIMUM INDEPENDENT DOMINATING SET on bipartite graphs.

Organization. In Section 2 we provide some basic preliminaries about graphs, the MMVC problem, parameterized complexity, and approximation algorithms. In Section 3 (resp. Section 4) we present our framework to obtain kernelization lower bounds for maximization (minimization) problems. In both sections, the contents are split into three subsections: we start with the general definitions, then we focus on the particular and relevant case of vertex-optimization problems, and then we present the general results for what we call “well-behaved” optimization problems. In Section 5 we present several applications of the framework of *lop*-kernels for concrete problems, and in Section 6 we discuss the flaw in the linear kernel for MMVC claimed by Fernau [34]. Section 7 is devoted to the subquadratic kernels for MMVC on particular graph classes, as well as to other positive results for MMVC. Our reduction to rule out the existence of polynomial kernels for MMVC parameterized by the size of a minimum vertex cover (or a maximum matching) is presented in Section 8. We conclude the article in Section 9 with a discussion and some directions for further research.

2 Preliminaries

Graphs and functions. We use standard graph-theoretic notation, and we refer the reader to [29] for any undefined notation. For an integer $p \geq 1$, we let $[p]$ be the set containing all integers i with $1 \leq i \leq p$. We use $\uplus$ to denote the disjoint union. We will only consider finite undirected graphs without loops nor multiple edges, and we denote an edge between two vertices u and v by $\{u, v\}$. A subgraph H of a graph G is *induced* if H can be obtained from G by deleting a set of vertices $D = V(G) \setminus S$, and we denote $H = G[S]$. Given a graph H , a graph G is *H-free* if it does not contain any induced subgraph isomorphic to H . If $\mathcal{H}$ is a collection of graphs, a graph G is *H-free* if it is H -free for every $H \in \mathcal{H}$. For a graph G and a set $S \subseteq V(G)$, we use the notation $G \setminus S = G[V(G) \setminus S]$, and for a vertex $v \in V(G)$, we abbreviate $G \setminus \{v\}$ as $G \setminus v$. A vertex v is *complete* to a set $S \subseteq V(G)$ if v is adjacent to

every vertex in S .

The *open* (resp. *closed*) *neighborhood* of a vertex v in a graph G is denoted by $N(v)$ (resp. $N[v]$), whenever the graph G is clear from the context. For vertex sets $X, Y \subseteq V(G)$, we define $N[X] = \bigcup_{v \in X} N[v]$, $N(X) = N[X] \setminus X$, $N_Y[X] = N[X] \cap Y$, and $N_Y(X) = N_Y[X] \setminus X$. The *degree* of a vertex v in a graph G is defined as $|N(v)|$, and we denote it by $\deg_G(v)$, or just $\deg(v)$ if the graph is clear from the context. For an integer $t \geq 1$, we denote by P_t (resp. I_t , K_t) the path (resp. edgeless graph, complete graph) on t vertices. For two integers $a, b \geq 1$, we denote by $K_{a,b}$ the bipartite graph with parts of sizes a and b .

A *clique* (resp. *independent set*) in a graph G is a set of vertices that are pairwise adjacent (resp. not adjacent). A graph property is *hereditary* if whenever it holds for a graph G , it holds for all its induced subgraphs as well. Note that the properties of being an edgeless graph, a complete graph, or an independent set are hereditary. We denote by $\Delta(G)$ (resp. $\omega(G)$) the maximum vertex degree (resp. clique size) of a graph G .

A *vertex cover* of a graph G is a set of vertices containing at least one endpoint of every edge, and it is *minimal* if no proper subset of it is a vertex cover. One of the concrete problems that we study in this paper is formally stated as follows. We state it as a decision problem, since most of our results consider its parameterization by the solution size k .

MAXIMUM MINIMAL VERTEX COVER (MMVC)

Input: A graph G and a positive integer k .

Question: Does G contain a minimal vertex cover of size at least k ?

For a graph G , we denote by $\text{mmvc}(G)$ the maximum size of a minimal vertex cover of G . The following observation has been already used in previous work [18, 57].

► **Observation 1.** *Let G be a graph. A set $X \subseteq V(G)$ is a minimal vertex cover of G if and only if X is a vertex cover of G and, for every vertex $v \in X$, $N(v) \not\subseteq X$.*

The next lemma provides a useful way to conclude that we are dealing with a *yes*-instance in the kernelization algorithms presented in Section 7.

► **Lemma 2.** *Let G be a graph and let $S \subseteq V(G)$ be an independent set. There exists a minimal vertex cover of G containing $N(S)$.*

Proof. Note that, since S is an independent set, $V(G) \setminus S$ is a vertex cover of G . Hence, there exists a minimal vertex cover X of G such that $X \subseteq V(G) \setminus S$. We claim that $N(S) \subseteq X$. Suppose for the sake of contradiction that there exists a vertex $v \in N(S)$ such that $v \notin X$. Since v has a neighbor u in S and $S \cap X = \emptyset$, the edge $\{u, v\}$ would not be covered by X . ◀

Note that, in particular, Lemma 2 implies that if (G, k) is an instance of the MAXIMUM MINIMAL VERTEX COVER problem and $v \in V(G)$ is a vertex of degree at least k , then we can conclude that (G, k) is a *yes*-instance. This will allow us to assume, in our kernelization algorithms, that $\Delta(G) \leq k - 1$.

Parameterized complexity. We refer the reader to [25, 30] for basic background on parameterized complexity, and we recall here only some basic definitions used in this article. A *parameterized problem* is a language $L \subseteq \Sigma^* \times \mathbb{N}$. For an instance $I = (x, k) \in \Sigma^* \times \mathbb{N}$, k is called the *parameter*.

A parameterized problem is *fixed-parameter tractable* (FPT) if there exists an algorithm $\mathcal{A}$, a computable function f , and a constant c such that given an instance $I = (x, k)$, $\mathcal{A}$ (called an *FPT algorithm*) correctly decides whether $I \in L$ in time bounded by $f(k) \cdot |I|^c$. For instance, the VERTEX COVER problem parameterized by the size of the solution is FPT.

For an instance (x, k) of a parameterized problem Q , a *kernelization algorithm* is an algorithm $\mathcal{A}$ that, in polynomial time, generates from (x, k) an equivalent instance (x', k') of Q such that $|x'| + k' \leq f(k)$, for some computable function $f : \mathbb{N} \rightarrow \mathbb{N}$, where $|x'|$ denotes the size of x' . If $f(k)$ is bounded from above by a polynomial of the parameter, we say that Q admits a *polynomial kernel*. In particular, if $f(k)$ is bounded by a linear (resp. quadratic) function, then we say that Q admits a *linear* (resp. *quadratic*) kernel.

A *polynomial parameter transformation*, abbreviated as PPT, is an algorithm that, given an instance (x, k) of a parameterized problem A , runs in time polynomial in $|x|$ and outputs an instance (x', k') of a parameterized problem B such that k' is bounded from above by a polynomial on k and (x, k) is positive if and only if (x', k') is positive. If a parameterized problem A does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$ and there exists a PPT from A to a parameterized problem B , then B does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$ either [25].

Approximation algorithms. We refer the reader to [55] for background on approximation algorithms, and we define here only some non-standard notions used in this article.

As, when dealing with graph problems, one typically measures the size of kernels in terms of the number of vertices or edges (and not in the classical bit-size of the instance), we introduce an arbitrary notion of size as follows. Given an optimization problem Π , we say that a non-negative integer-valued function $|\cdot|$ is a *size function* if, given an instance I of Π , $|I|$ can be computed in polynomial time in the classical bit-size, and $|I|$ is upper-bounded by a polynomial in the classical bit-size.

For an optimization problem Π , an instance I of Π , and a feasible solution s of Π in I , we denote by $\text{val}_\Pi(I, s)$ the value of the objective function of Π for s . We restrict ourselves to optimization problems Π whose objective functions for feasible solutions take non-negative integer values. For a maximization (resp. minimization) problem Π and an instance I of Π , we denote by $\text{opt}_\Pi(I)$ the maximum (resp. minimum) of $\text{val}_\Pi(I, s)$ over all feasible solutions s of Π in I .

A maximization (resp. minimization) problem Π is a *vertex-maximization* (resp. *vertex-minimization*) problem if their instances consist of a graph G , and the objective is to find a vertex set $S \subseteq V(G)$ of maximum (resp. minimum) size satisfying some conditions. For instance, MAXIMUM INDEPENDENT SET and MINIMUM VERTEX COVER are typical examples of vertex-maximization and vertex-minimization problems, respectively.

For kernelization purposes, we need to consider the *decision* version of optimization problems. For a maximization (resp. minimization) problem Π whose instances are of the form I , we denote by Π_{dec} the decision problem whose instances are of the form (I, k) , where k is a non-negative integer, and where (I, k) is a *yes-instance* of Π_{dec} if $\text{opt}_\Pi(I) \geq k$ (resp. $\text{opt}_\Pi(I) \leq k$), and a *no-instance* otherwise.

Since we aim at establishing a link between the existence of certain kernels and approximation algorithms, we need to take care of *constructibility* issues. The standard definition of a kernelization algorithm [35] does not involve constructing a solution of the considered problem. On the other hand, the standard definition of an approximation algorithm [55] does take into account the construction of the corresponding solution. Hence, in order to establish such a connection, we need to consider slightly “non-standard” definitions of these objects.

Namely, we distinguish between constructive and non-constructive approximation algorithms, so that the kernelization lower bounds that we present are able to rule out constructive or non-constructive kernels (cf. the second paragraph of Section 3). To this end, we say that an algorithm for a maximization (resp. minimization) decision problem Π_{dec} *constructively decides* an instance (I, k) if, whenever it holds that $\text{opt}_\Pi(I) \geq k$ (resp.

$\text{opt}_\Pi(I) \leq k$), the algorithm outputs a feasible solution s such that $\text{val}_\Pi(I, s) \geq k$ (resp. $\text{val}_\Pi(I, s) \leq k$).

When using the term “approximation algorithm” with ratio $\rho \geq 1$ for a maximization (resp. minimization) problem Π , we assume, unless stated otherwise, that it is *constructive*, that is, that the algorithm, given an instance I of Π , outputs a feasible solution s of Π in I such that $\text{opt}_\Pi(I)/\text{val}_\Pi(I, s) \leq \rho$ (resp. $\text{val}_\Pi(I, s)/\text{opt}_\Pi(I) \leq \rho$). Note that the approximation ratio ρ is, in general, a non-negative integer-valued function that depends on I .

We define a *value-approximation algorithm* with ratio $\rho \geq 1$ for a maximization (resp. minimization) problem Π as an algorithm that, given an instance I of Π , returns a non-negative integer k such that $1 \leq \text{opt}_\Pi(I)/k \leq \rho$ (resp. $1 \leq k/\text{opt}_\Pi(I) \leq \rho$). Again, here ρ is, in general, a non-negative integer-valued function that depends on I . Note that a value-approximation algorithm is not only not required to construct a feasible solution with value k , but also not required to guarantee that such a solution exists.

3 A framework for ruling out certain polynomial kernels: the case of maximization problems

In this section we introduce our generic framework to obtain lower bounds on the size of a certain type of polynomial kernels, which we call **lop-kernels** (see Definition 5), for a broad class of maximization problems that we proceed to introduce. Informally, the framework is based on simple and self-contained arguments proving that a “small” **lop-kernel** implies the existence of a “good” approximation algorithm. Then, the contrapositive of this statement implies that inapproximability results can be turned into **lop-kernel** lower bounds.

It worth mentioning here that most of the inapproximability results in the literature hold for the *value-approximation* algorithms as defined at the end of Section 2, that is, for algorithms that are *not* required to construct in polynomial time an appropriate solution of the corresponding problem, but only to report a value within the appropriate range. In order to guarantee that it is also possible to use our framework when only the non-existence of *constructive* approximation algorithms is known, we introduce a variant of **lop-kernels**, called *constructive lop-kernels*, such that their existence implies the existence of a constructive approximation algorithm. However, in a first reading, we recommend to skip all technical details concerning constructibility.

We say that a maximization problem Π is *well-behaved* if it comes equipped with a size function (as defined in Section 2) and it satisfies the following condition, which we denote by $C^{\max}$:

There exists an algorithm that, given as input a real number c and an instance I of Π such that $\text{opt}_\Pi(I) \leq c$, runs in polynomial time for every fixed c and either decides that $\text{opt}_\Pi(I) = 0$ and provides a feasible solution s with $\text{val}_\Pi(I, s) = 0$, or provides a feasible solution s with $\text{val}_\Pi(I, s) > 0$.

Observe that most of the classical maximization problems are well-behaved, and in particular any vertex-maximization problem whose decision version belongs to NP is well-behaved, as we can enumerate all subsets of vertices of size at most c , and for each of them verify in polynomial time if it is a feasible solution. Given a well-behaved maximization problem Π , we say that a function $u : \mathbb{N} \rightarrow \mathbb{N}$ is an *upper bound function* for Π if for any instance I of Π , it holds that $\text{opt}_\Pi(I) \leq u(|I|)$, where $|\cdot|$ is the size function of Π . Throughout the paper, we assume that the notions of size used in both the size of kernels and the upper bound function are the same.

The reminder of this section is organized as follows. In [Subsection 3.1](#) we present the definition of **lop**-rules and **lop**-kernels for well-behaved maximization problems, and we prove a general technical result, namely [Lemma 8](#). In [Subsection 3.2](#) we present the connection between **lop**-kernels and approximation algorithms for vertex-maximization problems, and in [Subsection 3.3](#) we generalize it to arbitrary well-behaved maximization problems.

3.1 Definition of **lop**-rules and **lop**-kernels

► **Definition 3.** A large optimal preserving reduction rule, or **lop**-rule for short, for a well-behaved maximization problem Π , is a polynomial-time algorithm R that, given an instance (I, k) of Π_{dec} , computes another instance (I', k') of Π_{dec} with $0 \leq k' \leq k$ and such that

1. if (I, k) is a no-instance of Π_{dec} , then (I', k') is a no-instance of Π_{dec} , and
2. if (I, k) is a yes-instance of Π_{dec} , then $\text{opt}_{\Pi}(G') \geq \text{opt}_{\Pi}(G) - (k - k')$, implying that (I', k') is a yes-instance of Π_{dec} .

A **lop**-rule R is constructive if, given I and any solution s' of I' of such that $\text{val}_{\Pi}(I', s') \geq k'$, it constructs (in polynomial time) a solution s of I such that $\text{val}_{\Pi}(I, s) \geq k$.

Note that Property 2 in [Definition 3](#) is stronger than the implication “if (I, k) is a yes-instance of Π_{dec} , then (I', k') is a yes-instance of Π_{dec} ”, which would yield the definition of a classical kernelization algorithm [\[25, 30\]](#). Indeed, when we consider how this latter implication is generally proved in safeness proofs of classical kernels, one of the following scenarios often occurs:

- (a) For every solution s of I there exists a solution s' of I' with $\text{val}_{\Pi}(I', s') \geq \text{val}_{\Pi}(I, s) - (k - k')$.
- (b) For every solution s of I with $\text{val}_{\Pi}(I, s) \geq k$, there exists a solution s' of I' with $\text{val}_{\Pi}(I', s') \geq \text{val}_{\Pi}(I, s) - (k - k')$.
- (c) If there exists a solution s of I with $\text{val}_{\Pi}(I, s) \geq k$, then there exists a solution s' of I' with $\text{val}_{\Pi}(I', s') \geq k'$.

In Case (a), the rule preserves all optimal values, as it implies that $\text{opt}_{\Pi}(G') \geq \text{opt}_{\Pi}(G) - (k - k')$. In Case (b), the rule preserves only large optimal values, as it implies that if $\text{opt}_{\Pi}(G) \geq k$, then $\text{opt}_{\Pi}(G') \geq \text{opt}_{\Pi}(G) - (k - k')$, implying Property 2 above. Note that if $\text{opt}_{\Pi}(G) < k$, then $\text{opt}_{\Pi}(G')$ and $\text{opt}_{\Pi}(G)$ are not necessarily related. This justifies our choice for “large optimal preserving” rules. Case (c) corresponds to the weaker and classical implication “if (I, k) is a yes-instance of Π_{dec} , then (I', k') is a yes-instance of Π_{dec} ”.

The following observation is an immediate consequence of the definition of a **lop**-rule.

► **Observation 4.** **lop**-rules can be composed. Formally, consider two **lop**-rules R_1 and R_2 . Then, the rule R that, given an instance (I, k) of Π_{dec} , returns $R_2(R_1(I, k))$, is also a **lop**-rule. Moreover, if R_2 and R_1 are constructive, then R is also constructive.

A typical example of a **lop**-rule for a vertex-maximization problem is when we can identify a “dominant” set of vertices that can be safely included into a solution. More precisely, consider a rule that, given a graph G , finds a subset $T \subseteq V(G)$ and a graph G' such that there exists an optimal solution S^* in G such that $S^* = T \cup S'$, where S' is a solution in G' , and for every solution S' in G' , $S' \cup T$ is a solution in G . Such a rule is a (constructive) **lop**-rule, as we even fall into Case (a) described above.

Even if we are not aware of known reduction rules for vertex-maximization problems that are *not* **lop**-rules, we can artificially devise such an example. For instance, for the MMVC problem, given an instance (G, k) , if there is a vertex that has more than k neighbors of degree one, we can safely delete all but any k of them to obtain a reduced graph G' , and

leave k unchanged. Note that this rule falls into Case (c) above, since by Lemma 2 both G and G' are yes-instances of MMVC, but it does not satisfy Property 2 in Definition 3, since $\text{mmvc}(G)$ may be arbitrarily larger than $\text{mmvc}(G')$.

If we defined a **lop-kernel** as an algorithm consisting only of **lop-rules**, we would exclude from being a **lop-kernel**, for instance, a rule that detects a yes-instance as in the above paragraph. This justifies the next definition, where we also allow **lop-kernels** to decide instances.

► **Definition 5.** Let Π be a well-behaved maximization problem and let $s : \mathbb{N} \rightarrow \mathbb{N}$ be a computable function. A **lop-kernel** of size s for Π parameterized by the solution size is a polynomial-time algorithm that takes as input an instance (I, k) of Π_{dec} , and either

- decides that (I, k) is a yes-instance or a no-instance, or
- outputs a reduced instance (I', k') by applying a sequence of **lop-rules** to (I, k) , with $|I'| \leq s(k)$.

A **lop-kernel** is *constructive* if, in the first case, it constructively decides (I, k) (but in the second case it may not use constructive rules).

As it is common in kernels to exhaustively apply reduction rules, and then to either decide the reduced instance or to output it, let us introduce and discuss the following definition¹.

► **Definition 6.** Let Π be a well-behaved maximization problem and let $s : \mathbb{N} \rightarrow \mathbb{N}$ be a computable function. A **lop-kernel*** of size s for Π parameterized by the solution size is a polynomial-time algorithm that takes as input an instance (I, k) of Π_{dec} , computes an instance (I', k') by applying a (possibly empty) sequence of **lop-rules** to (I, k) , and either

- decides that (I', k') is a yes-instance or a no-instance, or
- outputs (I', k') , with $|I'| \leq s(k)$.

A **lop-kernel*** is *constructive* if, in the first case, it constructively decides (I', k') and, in the second case, it only uses constructive **lop-rules**.

Firstly, observe that a **lop-kernel*** (resp. constructive **lop-kernel***) is a **lop-kernel** (resp. constructive **lop-kernel**). Indeed, if a **lop-kernel*** decides (I', k') , then, as the definition of **lop-rules** implies that the reduced instance (I', k') is equivalent to (I, k) , it also decides (I, k) . Moreover, if a **lop-kernel*** is constructive and decides that (I', k') is a yes-instance by providing a solution s' with $\text{val}_{\Pi}(I', s') \geq k'$, then, as the rules are constructive, and according to Observation 4, we can build in polynomial time a solution s with $\text{val}_{\Pi}(I, s) \geq k$, and thus constructively decide (I, k) . Secondly, observe that a **lop-kernel** is a **lop-kernel***, but that a constructive **lop-kernel** is not necessarily a constructive **lop-kernel***. The conclusion of this discussion is that for the non-constructive versions, both definitions are equivalent, and for the constructive versions, **lop-kernels** are slightly more general. As many inapproximability results even hold for the non-constructive version of approximation, we suggest the reader to stick to the non-constructive version, and thus to choose any of the two definitions. The only case where it could make a difference would be for a problem Π for which inapproximability results are only known for ruling out constructive approximation algorithms. Then, Corollary 11 will turn this inapproximability into a kernel lower bound even for constructive **lop-kernels**, and not only for constructive **lop-kernels***. This justifies why we consider henceforth only **lop-kernels**.

¹ This definition was used in the conference version of this paper [4].

Our next objective is to prove that a **lop**-kernel yields an approximation algorithm. For this, we need the following definition, which is inspired by a similar notion introduced by Hochbaum and Shmoys [45], and referred to as *f-relaxed decision procedure* in [55].

► **Definition 7.** Let Π be a well-behaved maximization problem and let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a function. An *f*-dual-approximation algorithm for Π is a polynomial-time algorithm that, given an instance (I, k) of Π_{dec} , concludes one of the following:

- $\text{opt}_{\Pi}(I) \geq k$.
- $\text{opt}_{\Pi}(I) < f(k)$.

An *f*-dual approximation algorithm is constructive if, whenever it concludes that $\text{opt}_{\Pi}(I) \geq k$, it provides a solution s with $\text{val}_{\Pi}(I, s) \geq k$.

In the next lemma we prove that a **lop**-kernel of size s yields an *f*-dual-approximation algorithm (where f depends on s), which in turn yields a classical approximation algorithm whose ratio depends on s . To provide some insight on the statement of the next lemma, keep in mind that for vertex-maximization problems, the upper bound function u is typically the identity function.

► **Lemma 8.** Let Π be a well-behaved maximization problem with a non-decreasing upper bound function u and let $s : \mathbb{N} \rightarrow \mathbb{N}$ be a computable function. If Π admits a **lop**-kernel of size s , then Π admits an *f*-dual-approximation algorithm where $f(k) := u(s(k)) + k + 1$. Moreover, if the **lop**-kernel is constructive, then the *f*-dual-approximation algorithm is also constructive.

Proof. Let $k \in \mathbb{N}$, (I, k) be an instance of Π_{dec} , and $\mathcal{R}$ be a **lop**-kernel of size s for Π . We describe an *f*-dual approximation algorithm $\mathcal{A}$ which takes as input I and k , starts by running $\mathcal{R}$ with input (I, k) , and continues based on its possible output. If $\mathcal{R}$ decides that (I, k) is a **yes**-instance, then $\text{opt}_{\Pi}(I) \geq k$, and $\mathcal{A}$ returns $\text{opt}_{\Pi}(I) \geq k$ as well. Notice that if $\mathcal{R}$ is constructive, then it provides a solution s of I such that $\text{val}_{\Pi}(I, s) \geq k$, and $\mathcal{A}$ returns this solution as well. Otherwise, we claim that it is safe for $\mathcal{A}$ to return $\text{opt}_{\Pi}(I) \leq f(k)$. Indeed, if $\mathcal{R}$ decides that (I, k) is a **no**-instance, then $\text{opt}_{\Pi}(I) < k$, implying $\text{opt}_{\Pi}(I) < u(s(k)) + k = f(k)$. Finally, suppose that $\mathcal{R}$ outputs an equivalent instance (I', k') obtained from (I, k) using only **lop**-rules and such that $|I'| \leq s(k)$. By using [Observation 4](#) we can assume that (I', k') is obtained from (I, k) by a single **lop**-rule, and [Property 2](#) in [Definition 3](#) implies that $\text{opt}_{\Pi}(I) \leq \text{opt}_{\Pi}(I') + (k - k') \leq u(|I'|) + k \leq u(s(k)) + k < f(k)$, where we have used the fact that u is non-decreasing. ◀

Let us now turn to our main results relating the size of **lop**-kernels to the existence of approximation algorithms. To keep statements as simple as possible, we first provide in [Subsection 3.2](#) results that correspond to the specialized versions for vertex-maximization problems of the general results presented in [Subsection 3.3](#).

3.2 Connection between **lop**-kernels and approximation algorithms for vertex-maximization problems

In this subsection we deal with vertex-maximization problems. The following lemma is a folklore result [55], but as it is generally tuned for a particular function f appearing in the considered context, we need to restate it in a general form.

► **Lemma 9.** Let Π be a vertex-maximization problem whose decision version is in NP, and $f : \mathbb{N} \rightarrow \mathbb{N}$ be a computable function.

1. For every real number $c > 1$, if Π admits an f -dual-approximation algorithm with $f(k) = \mathcal{O}(k^c)$, then Π admits a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^{\frac{c-1}{c}})$ on n -vertex graphs.
2. For every real number $\beta \geq 1$, if Π admits an f -dual-approximation algorithm with $f(k) = \beta k + 1$, then Π admits a polynomial-time value-approximation algorithm with ratio $\beta + \varepsilon$ for every real number $\varepsilon > 0$.

Moreover, if the f -dual-approximation is constructive, then the corresponding approximation algorithm is also constructive.

Proof. Let A be an f -dual-approximation algorithm for Π . We proceed to construct a polynomial-time approximation algorithm for Π with the claimed ratio. We consider the two statements of the lemma separately.

Case 1: $f(k) = \mathcal{O}(k^c)$.

Given an n -vertex graph G as instance of Π , we find $k_0 \in \{0, \dots, n\}$ defined as the largest positive integer k such that algorithm A returns that $\text{opt}_\Pi(G) \geq k$. Note that k_0 can be found in polynomial time by performing at most $n + 1$ calls to algorithm A . If there is no such k_0 , or if $k_0 = 0$, then $\text{opt}_\Pi(G) < \max(f(0), f(1)) = \mathcal{O}(1)$, and since the decision version of Π is in NP, we can find an optimal solution in polynomial time by verifying all vertex subsets of size at most $\max(f(0), f(1))$. Otherwise, that is, if $k_0 \geq 1$, our approximation algorithm returns k_0 , or if is constructive it returns a solution S_0 (a subset of vertices here) such that $|S_0| \geq k_0$. Let us prove that it provides the claimed approximation ratio. We distinguish two subcases depending on the value of k_0 . Suppose first that $k_0 \geq n^{1/c}$. Since $\text{opt}_\Pi(G) \leq n$, in this case we get that

$$\frac{\text{opt}_\Pi(G)}{k_0} \leq \frac{n}{n^{1/c}} = n^{\frac{c-1}{c}}.$$

Otherwise, it holds that $k_0 < n^{1/c}$. By the definition of k_0 we have $\text{opt}_\Pi(G) < f(k_0 + 1) = \mathcal{O}((k_0 + 1)^c) = \mathcal{O}((k_0)^c)$. Thus, in this case we get that

$$\frac{\text{opt}_\Pi(G)}{k_0} = \frac{\mathcal{O}((k_0)^c)}{k_0} = \mathcal{O}((k_0)^{c-1}) = \mathcal{O}\left(n^{\frac{c-1}{c}}\right).$$

Since in both cases we have a ratio of $\mathcal{O}(n^{\frac{c-1}{c}})$, the lemma follows in Case 1.

Case 2: $f(k) = \beta k + 1$.

Let $\varepsilon > 0$ be a arbitrary real number, let $\varepsilon' = \frac{\varepsilon}{\beta}$, and let us provide a polynomial-time approximation algorithm with ratio $\beta(1 + \varepsilon') = \beta + \varepsilon$. As in Case 1, we start by finding k_0 , defined as the largest positive integer k such that algorithm A returns that $\text{opt}_\Pi(G) \geq k$. By definition of k_0 we have $\text{opt}_\Pi(G) < f(k_0 + 1) = \beta(k_0 + 1) + 1 \leq \beta(k_0 + 2)$. If $k_0 < \frac{2}{\varepsilon'}$, then $\text{opt}_\Pi(G)$ is constant, and again by enumerating all subsets of size at most $\beta(\frac{2}{\varepsilon'} + 2)$ we find an optimal solution. Otherwise, we return k_0 , or if is constructive it returns a solution S_0 (a subset of vertices here) such that $|S_0| \geq k_0$. We have $\text{opt}_\Pi(G) \leq \beta(k_0 + 2) \leq \beta(1 + \varepsilon')k_0$, concluding Case 2 of the proof. $\blacktriangleleft$

As a vertex-maximization problem whose decision version is in NP is a well-behaved problem, the hypothesis of Lemma 8 is satisfied (taking the identity function as upper bound function), and thus the following theorem is immediate by pipelining Lemma 8 and Lemma 9.

- **Theorem 10.** *Let Π be a vertex-maximization problem whose decision version is in NP.*
1. *For every real number $c > 1$, if Π admits a lop-kernel with $\mathcal{O}(k^c)$ vertices, then it admits a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^{\frac{c-1}{c}})$ on n -vertex graphs.*
 2. *For every real number $\beta \geq 1$, if Π admits a lop-kernel with βk vertices, then for any real number $\varepsilon > 0$, it admits a polynomial-time value-approximation algorithm with ratio $(\beta + 1 + \varepsilon)$.*

Moreover, if the lop-kernel is constructive, then the corresponding approximation algorithm is also constructive.

As the framework of lop-kernels is mainly defined as a tool to get lop-kernel lower bounds from inapproximability, let us explicitly formulate the contrapositive of Theorem 10. Note that, when applying it to a concrete problem Π , the inapproximability of Π will rely on some complexity assumption, typically $P \neq NP$.

- **Corollary 11.** *Let Π be a vertex-maximization problem whose decision version is in NP.*
1. *For every real number $r \in (0, 1)$, if Π does not admit a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^r)$ on n -vertex graphs, then Π parameterized by the solution size does not admit a lop-kernel with $\mathcal{O}(k^{\frac{1}{1-r}})$ vertices.*
 2. *For every real number $\beta > 1$, if Π does not admit a polynomial-time value-approximation algorithm with ratio β , then Π parameterized by the solution size does not admit a lop-kernel with $(\beta - 1 - \varepsilon)k$ vertices for any real number $\varepsilon > 0$.*

Moreover, if the non-existence of approximation algorithms only holds for constructive approximation algorithms, then the lower bound only holds for constructive lop-kernels .

3.3 Connection between lop-kernels and approximation algorithms for well-behaved maximization problems

The following lemma and theorem are the versions of Lemma 9 and Theorem 10, respectively, in the more general setting of arbitrary well-behaved maximization problems.

► **Lemma 12.** *Let Π be a well-behaved maximization problem, $a \in \mathbb{R}^+$, $u : \mathbb{N} \rightarrow \mathbb{N}$, and $f : \mathbb{N} \rightarrow \mathbb{N}$ be functions such that $u(n) = \mathcal{O}(n^a)$, u is polynomial-time computable, and f is computable. Suppose that Π has u as upper bound function and that it admits an f -dual-approximation.*

1. *If $f(k) = \mathcal{O}(k^d)$ for some real number $d > 1$, then Π admits a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^{\frac{a(d-1)}{d}})$, where n is the size of the input.*
2. *If $f(k) = \lambda k^d + k + 1$ for some real numbers $d \leq 1$ and $\lambda > 0$, then Π admits a polynomial-time value-approximation algorithm with ratio $\lambda 2^d + 3$.*

Moreover, if the f -dual-approximation algorithm is constructive, then the corresponding approximation algorithm is also constructive.

Proof. Let A be an f -dual-approximation algorithm for Π . For both cases in the statement in the lemma, we proceed to construct a polynomial-time approximation algorithm for Π with the claimed ratio.

Given an instance I of Π , we find $k_0 \in \{0, \dots, u(n)\}$ (recall that $n = |I|$) defined as the largest positive integer k such that algorithm A returns that $\text{opt}_\Pi(G) \geq k$. Note that k_0 can be found in polynomial time as $|I|$ is polynomial-time computable and its value n is polynomially upper-bounded in the classical bit-size of the instance, and that $u(n)$ can be computed in polynomial time as well. If there is no such k_0 , or if $k_0 = 0$, then $\text{opt}_\Pi(G) < \max(f(0), f(1)) = \mathcal{O}(1)$, and since Π is well-behaved, we can, given

$\max(f(0), f(1))$, decide in polynomial time if either $\text{opt}_\Pi(I) = 0$ and provide a solution s with $\text{val}_\Pi(I, s) = 0$, or provide a solution s with $\text{val}_\Pi(I, s) > 0$. In both cases we even have a constructive constant-factor approximation, and as $\text{opt}_\Pi(I) < f(1) = \lambda + 2$, we get the claimed ratio in both cases. Otherwise, that is, when $k_0 \geq 1$, our approximation algorithm returns k_0 , or if is constructive it returns a solution s_0 such that $\text{val}_\Pi(I, s_0) \geq k_0$. Let us prove that it provides the claimed approximation ratio. We now distinguish the two cases claimed in the statement of the lemma.

Case 1: $f(k) = \mathcal{O}(k^d)$.

Suppose first that $k_0 \geq n^{a/d}$. In this case we have

$$\frac{\text{opt}_\Pi(I)}{k_0} \leq \frac{u(n)}{k_0} = \frac{\mathcal{O}(n^a)}{n^{a/d}} = \mathcal{O}\left(n^{\frac{a(d-1)}{d}}\right).$$

Otherwise, it holds that $k_0 < n^{a/d}$. By the definition of k_0 we have $\text{opt}_\Pi(G) < f(k_0 + 1) = \mathcal{O}((k_0 + 1)^d) = \mathcal{O}((k_0)^d)$. Thus, in this case we get that

$$\frac{\text{opt}_\Pi(I)}{k_0} = \frac{\mathcal{O}((k_0)^d)}{k_0} = \mathcal{O}((k_0)^{d-1}) = \mathcal{O}\left(n^{\frac{a(d-1)}{d}}\right).$$

Since in both cases we have a ratio of $\mathcal{O}(n^{\frac{a(d-1)}{d}})$, the lemma follows in Case 1.

Case 2: $f(k) = \lambda k^d + k + 1$.

We have

$$\frac{\text{opt}_\Pi(I)}{k_0} \leq \frac{f(k_0 + 1)}{k_0} \leq \frac{\lambda(k_0 + 1)^d + k_0 + 2}{k_0}$$

Let

$$h(x) = \frac{\lambda(x + 1)^d + x + 2}{x}$$

and note that the approximation ratio is at most $h(k_0)$.

To bound $h(k_0)$, we proceed to show that $h(x)$ is decreasing in x when $x > 0$ and obtain the desired approximation ratio of $h(k_0) \leq h(1) = \lambda 2^d + 3$. To show that $h(x)$ is indeed decreasing in x when $x > 0$, note that

$$\frac{\partial h(x)}{\partial x} = \frac{\lambda(x + 1)^d}{x^2} \cdot \left(d \frac{x}{x + 1} - 1\right) - \frac{1}{x^2},$$

which is negative when $d \leq 1$ and $x > 0$. ◀

The next theorem follows immediately by pipelining [Lemma 8](#) and [Lemma 12](#). Namely, starting with the hypothesis of [Theorem 13](#), we first apply [Lemma 8](#) and then [Lemma 12](#) with $d = ac$ and $\lambda = \alpha\beta^a$.

► **Theorem 13.** *Let Π be a well-behaved maximization problem, $a, c \in \mathbb{R}^+$, $u : \mathbb{N} \rightarrow \mathbb{N}$, and $s : \mathbb{N} \rightarrow \mathbb{N}$ be functions such that $u(n) = \mathcal{O}(n^a)$, $s(k) = \mathcal{O}(k^c)$, u is non-decreasing, and s and u are polynomial-time computable. Suppose that Π has u as upper bound function and that it admits a log-kernel of size s , according to the same size function $|\cdot|$ associated with Π .*

1. *If $ac > 1$, then Π admits a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^{\frac{ac-1}{c}})$, where n is the size of the instance.*

2. If $ac \leq 1$ and $\alpha, \beta \in \mathbb{R}^+$ are such that $u(n) \leq \alpha n^a$ and $s(k) \leq \beta k^c$, then Π admits a polynomial-time value-approximation algorithm with ratio $\alpha\beta^a 2^{ac} + 3$.

Moreover, if the **lop**-kernel is constructive, then the corresponding approximation algorithm is also constructive.

To provide some insight on the formulas used in the statement of [Theorem 13](#), and especially on the role of the upper bound function $u(n) = \mathcal{O}(n^a)$, one can typically think of a graph problem where the output is a subset of edges, and where the size of an instance is the number of vertices of input graph. In that case, we have $a = 2$, and thus a **lop**-kernel of size (in terms of number of vertices) $\mathcal{O}(k^c)$ would only imply an $\mathcal{O}(n^{\frac{2c-1}{c}})$ -approximation algorithm, which is worse than the ratio $\mathcal{O}(n^{\frac{c-1}{c}})$ obtained in [Theorem 10](#), where $a = 1$. On the other hand, the ratio can also sometimes be slightly better, as there may exist problems with upper bound function $u(n) = \mathcal{O}(n^a)$ for some $a < 1$. Moreover, observe that for problems where $a \leq 1$, the second item covers the case of linear kernels, which corresponds to $c = 1$.

By taking the contrapositive of [Theorem 13](#), we obtain the following more general version of [Corollary 11](#).

► **Corollary 14.** *Let Π be a well-behaved maximization problem with a non-decreasing and polynomial-time computable upper bound function $u(n) = \mathcal{O}(n^a)$ for $a \in \mathbb{R}^+$. In what follows, the size of the instance, denoted by n , and the size of the kernel are defined according to the same size function $|\cdot|$ associated with Π .*

1. *For every real number $r \in (0, 1)$, if Π does not admit a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^r)$, then Π parameterized by the solution size does not admit a **lop**-kernel of size $\mathcal{O}(k^{\frac{1}{a-r}})$.*
2. *Suppose that $u(n) \leq \alpha n^a$ for some $\alpha \in \mathbb{R}^+$. For every real number $\beta > 1$, if Π does not admit a polynomial-time value-approximation algorithm with ratio β , then Π parameterized by the solution size does not admit a **lop**-kernel of size $\beta' k^{c'}$ for any real numbers β', c' such that $ac' \leq 1$ and $\alpha\beta'^a 2^{ac'} + 3 \leq \beta$.*

Moreover, if the non-existence of approximation algorithms only holds for constructive approximation algorithms, then the lower bound only holds for constructive **lop**-kernels.

4 A framework for ruling out certain polynomial kernels: the case of minimization problems

In this section we adapt the framework of **lop**-kernels introduced in [Section 3](#) to minimization problems. The definitions and results for minimization problems are very close to those for maximization problems, but there are a number of subtle differences that we will discuss as they appear.

We say that a minimization problem Π is *well-behaved* if it comes equipped with a size function (as defined in [Section 2](#)) and it satisfies the following condition, which we denote by $C^{\min}$:

There exists a polynomial-time algorithm that, given as input an instance I of Π , decides if $\text{opt}_{\Pi}(I) = 0$, and in this case provides a solution s where $\text{val}_{\Pi}(I, s) = 0$, or otherwise provides any solution s .

Note that condition $C^{\min}$ above is *not* the symmetric version of condition $C^{\max}$ defined at the beginning of [Section 3](#).

Given a well-behaved minimization problem Π , we say that a function $u : \mathbb{N} \rightarrow \mathbb{N}$ is an *upper bound function* for Π if for any instance I of Π and any solution s of I , we have $\text{val}_\Pi(I, s) \leq u(|I|)$, where $|\cdot|$ is the size function of Π . Note that this notion of upper bound function differs from the one given for maximization problems. Again, throughout the paper, we assume that the notions of size used in both the size of kernels and the upper bound function are the same.

The remainder of this section is organized similarly to [Section 3](#). Namely, in [Subsection 4.1](#) we present the definition of **lop-rules** and **lop-kernels** for well-behaved minimization problems, and we prove a general technical result, namely [Lemma 19](#). In [Subsection 4.2](#) we present the connection between **lop-kernels** and approximation algorithms for vertex-minimization problems, and in [Subsection 4.3](#) we generalize it to arbitrary well-behaved minimization problems.

4.1 Definition of lop-rules and lop-kernels

The following definition should be compared to [Definition 3](#).

► **Definition 15.** A large optimal preserving *reduction rule*, or **lop-rule** for short, for a well-behaved minimization problem Π , is a polynomial-time algorithm R that, given an instance (I, k) of Π_{dec} , computes another instance (I', k') of Π_{dec} with $0 \leq k' \leq k$ and such that

1. if (I, k) is a **yes-instance** of Π_{dec} , then (I', k') is a **yes-instance** of Π_{dec} , and
2. if (I, k) is a **no-instance** of Π_{dec} , then $\text{opt}_\Pi(G') \geq \text{opt}_\Pi(G) - (k - k')$, implying that (I', k') is a **no-instance** of Π_{dec} .

A **lop-rule** R is *constructive* if, for any solution s' of I' , it constructs (in polynomial time) a solution s of I such that $\text{val}_\Pi(I, s) \leq \text{val}_\Pi(I', s') + (k - k')$.

Note that Property 2 in [Definition 15](#) is stronger than the implication “if (I, k) is a **no-instance** of Π_{dec} , then (I', k') is a **no-instance** of Π_{dec} ”, which would yield the definition of a classical kernelization algorithm. Observe also that the constructibility condition implies Property 2, unlike in the maximization case. Indeed, when we consider how this latter implication is generally proved in safeness proofs of classical kernels, we generally prove its contrapositive, and one of the following scenarios often occur:

- (a) For every solution s' of I' there exists a solution s of I with $\text{val}_\Pi(I, s) \leq \text{val}_\Pi(I', s') + (k - k')$.
- (b) For every solution s' of I' with $\text{val}_\Pi(I', s') \leq k'$, there exists a solution s of I with $\text{val}_\Pi(I, s) \leq \text{val}_\Pi(I', s') + (k - k')$.
- (c) If there exists a solution s' of I' with $\text{val}_\Pi(I', s') \leq k'$, then there exists a solution s of I with $\text{val}_\Pi(I, s) \leq k$.

In Case (a), the rule preserves all optimal values, as it implies that $\text{opt}_\Pi(G') \geq \text{opt}_\Pi(G) - (k - k')$, and note that it implies Property 2 in [Definition 15](#). Compared to the maximization case, Case (c) still implies only the classical implication “if (I, k) is a **no-instance** of Π_{dec} , then (I', k') is a **no-instance** of Π_{dec} ”, but note that Case (b) no longer implies Property 2. This is one of reasons why, in our opinion, the framework of **lop-kernels** seems to be more natural when applied to maximization problems.

The following observation is again an immediate consequence of the definition of a **lop-rule**.

► **Observation 16.** **lop-rules** can be composed. Formally, consider two **lop-rules** R_1 and R_2 . Then, the rule R that, given a instance (I, k) of Π_{dec} , returns $R_2(R_1(I, k))$, is also a **lop-rule**. Moreover, if R_2 and R_1 are constructive, then R is also constructive.

A typical example of a **lop-rule** for a vertex-minimization problem is when, for some problem Π , we can identify a “dominant” set of vertices that can be safely included into a solution. More precisely, consider a rule that, given a graph G , finds a subset $T \subseteq V(G)$ and a graph G' such that there exists an optimal solution S^* in G such that $S^* = T \cup S'$, where S' is a solution in G' , and for every solution S' in G' , $S' \cup T$ is a solution in G . Such a rule is indeed a (constructive) **lop-rule**.

Even if almost all classical known reduction rules for minimization problems are **lop-rules** [25, 35], here is a simple example a non-**lop-rule**. Consider the VERTEX COVER problem, and suppose that, given an instance (G, k) , we find in G a matching M of size $k + 1$. The rule just outputs $(G', k') = (M, k)$, hence preserving the fact that (G, k) is a **no-instance**. However, this rule does not satisfy Property 2 in Definition 15, since the size of a minimum vertex cover of G may be arbitrarily large compared to k , hence the inequality $\text{opt}_\Pi(G') \geq \text{opt}_\Pi(G)$ may not hold. In Section 5 we discuss a (much more involved) reduction rule for a vertex-minimization problem, namely TREE DELETION SET, which is not a **lop-rule** either.

As in the maximization case, if we defined a **lop-kernel** as an algorithm consisting only of **lop-rules**, we would exclude from being a **lop-kernel**, for instance, the algorithm consisting of the rule that detects a **no-instance** of VERTEX COVER as in the above paragraph. This justifies the next definition, where we also allow **lop-kernels** to decide instances, and that should be compared to Definition 5.

► **Definition 17.** Let Π be a well-behaved minimization problem and let $s : \mathbb{N} \rightarrow \mathbb{N}$ be a computable function. A **lop-kernel** of size s for Π parameterized by the solution size is a polynomial-time algorithm that takes as input an instance (I, k) of Π_{dec} , and either

- decides that (I, k) is a **yes-instance** or a **no-instance**, or
- outputs a reduced instance (I', k') by applying a sequence of **lop-rules** to (I, k) , with $|I'| \leq s(k)$.

A **lop-kernel** is *constructive* if, in the first case, it constructively decides (I, k) , and, in the second case, it only uses constructive **lop-rules**.

Note that the constructibility condition in Definition 17 differs from that in Definition 5 for maximization problems. We need this stronger property in the proof of Lemma 19.

Our next objective is to prove that a **lop-kernel** yields the existence of a polynomial-time approximation algorithm. For this, we need the following definition, which is the version of Definition 7 for minimization problem.

► **Definition 18.** Let Π be a well-behaved minimization problem and let $f : \mathbb{N} \rightarrow \mathbb{N}$. An f -dual-approximation algorithm for Π is a polynomial-time algorithm that, given an instance (I, k) of Π_{dec} , concludes one of the following:

- $\text{opt}_\Pi(I) \leq f(k)$.
- $\text{opt}_\Pi(I) > k$.

An f -dual-approximation algorithm is *constructive* if, whenever it concludes that $\text{opt}_\Pi(I) \leq f(k)$, it provides a solution s with $\text{val}_\Pi(I, s) \leq f(k)$.

In the next lemma we prove that a **lop-kernel** of size s yields an f -dual-approximation algorithm (where f depends on s), which in turn yields a classical approximation algorithm whose ratio depends on s . As in the maximization case, to provide some insight on the statement of the next lemma, keep in mind that for vertex-minimization problems, the upper bound function u is typically the identity function. Note that in the next lemma, the derived function $f(k)$ differs slightly from that of Lemma 8; this is due to technical reasons motivated

by the fact that there the maximization and minimization versions of our framework are not totally symmetric.

► **Lemma 19.** *Let Π be a well-behaved minimization problem with a non-decreasing upper bound function u and let $s : \mathbb{N} \rightarrow \mathbb{N}$ be a computable function. If Π admits a **lop**-kernel of size s , then Π admits an f -dual-approximation algorithm where $f(k) := u(s(k)) + k$. Moreover, if the **lop**-kernel is constructive, then the f -dual-approximation algorithm is also constructive.*

Proof. Let $k \in \mathbb{N}$, (I, k) be an instance of Π_{dec} , and $\mathcal{R}$ be a **lop**-kernel of size s for Π . We describe an f -dual-approximation algorithm $\mathcal{A}$ which takes as input I and k , starts by running $\mathcal{R}$ with input (I, k) , and continues based on its possible output. If $\mathcal{R}$ decides that (I, k) is a **no**-instance, then $\text{opt}_{\Pi}(I) > k$, and $\mathcal{A}$ returns $\text{opt}_{\Pi}(I) > k$. If $\mathcal{R}$ decides that (I, k) is a **yes**-instance, then $\text{opt}_{\Pi}(I) \leq k$, and $\mathcal{A}$ returns $\text{opt}_{\Pi}(I) \leq k \leq f(k)$ as well. Notice that if $\mathcal{R}$ is constructive, then it provides a solution s of I such that $\text{val}_{\Pi}(I, s) \geq k$, and $\mathcal{A}$ returns this solution as well.

Finally, suppose that $\mathcal{R}$ outputs an equivalent instance (I', k') obtained from (I, k) using only **lop**-rules and such that $|I'| \leq s(k)$. By using [Observation 16](#) we can assume that (I', k') is obtained from (I, k) by a single **lop**-rule. Let us start by the non-constructive case, in which $\mathcal{A}$ returns $\text{opt}_{\Pi}(I) \leq f(k)$. If $\text{opt}_{\Pi}(I) \leq k$, then we are done as $k \leq f(k)$. If $\text{opt}_{\Pi}(I) > k$, [Property 2](#) in [Definition 15](#) implies that $\text{opt}_{\Pi}(I) \leq \text{opt}_{\Pi}(I') + (k - k') \leq u(|I'|) + k \leq u(s(k)) + k = f(k)$, where we have used that u is non-decreasing. Let us now turn to the constructive case. As Π is well-behaved and verifies $C^{\min}$, we can compute in polynomial time a solution s' (of any cost), and according to the definition of u we have $\text{val}_{\Pi}(I', s') \leq u(|I'|) \leq u(s(k))$, where we have used again that u is non-decreasing. Finally, as the rule is constructive, we can construct in polynomial time a solution s such that $\text{val}_{\Pi}(I, s) \leq \text{val}_{\Pi}(I', s') + k \leq f(k)$, and algorithm $\mathcal{A}$ returns this solution as well. ◀

Let us now turn to our main results relating the size of **lop**-kernels with the existence of approximation algorithms. As in [Section 3](#), to keep statements as simple as possible, we provide in [Subsection 4.2](#) results that correspond to the specialized versions for vertex-maximization problems of results in [Subsection 4.3](#).

4.2 Connection between **lop**-kernels and approximation algorithms for vertex-minimization problems

In this subsection we deal with vertex-minimization problems. The following lemma, which should be compared to [Lemma 9](#) Note that the hypothesis in the second item of [Lemma 9](#) is slightly different from the one below, and that the obtained approximation ratios are also slightly different.

► **Lemma 20.** *Let Π be a vertex-minimization problem whose decision version is in NP, $c > 1$ and $\beta \geq 1$ be real numbers, and $f : \mathbb{N} \rightarrow \mathbb{N}$ be a computable function.*

1. *If Π admits an f -dual-approximation algorithm where $f(k) = \mathcal{O}(k^c)$, then Π admits a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^{\frac{c-1}{c}})$ on n -vertex graphs.*
2. *If Π admits a f -dual-approximation algorithm where $f(k) = \beta k$, then Π admits a polynomial-time value-approximation algorithm with ratio β .*

Moreover, if the f -dual-approximation is constructive, then the corresponding approximation algorithm is also constructive.

Proof. Let $\mathcal{A}$ be an f -dual-approximation algorithm for Π . We proceed to construct a polynomial-time approximation algorithm for Π with the claimed ratio. We consider the two statements of the lemma separately.

Case 1: $f(k) = \mathcal{O}(k^c)$.

Given an n -vertex graph G as instance of Π , we find $k_0 \in \{0, \dots, n\}$ defined as the smallest positive integer k such that algorithm A returns that $\text{opt}_\Pi(G) \leq f(k)$. Note that k_0 can be found in polynomial time by performing at most $n + 1$ calls to algorithm A . Notice that k_0 always exists as we cannot have $\text{opt}_\Pi(G) > n$. If $k_0 = 0$, then $\text{opt}_\Pi(G) \leq f(0) = \mathcal{O}(1)$, and since the decision version of Π is in **NP**, we can find an optimal solution in polynomial time by verifying all vertex subsets of size at most $f(0)$. Otherwise, that is, if $k_0 \geq 1$, our approximation algorithm returns $f(k_0)$, or if is constructive it returns a solution S_0 (that is, a subset of vertices) such that $|S_0| \leq f(k_0)$. By definition of k_0 , we have that $\text{opt}_\Pi(G) > k_0 - 1$, or equivalently $\text{opt}_\Pi(G) \geq k_0$. Let us prove that this algorithm provides the claimed approximation ratio. We distinguish two subcases depending on the value of k_0 . Suppose first that $k_0 \geq n^{1/c}$. In this case we get that

$$\frac{f(k_0)}{\text{opt}_\Pi(G)} \leq \frac{n}{k_0} \leq \mathcal{O}(n^{\frac{c-1}{c}}).$$

Otherwise, it holds that $k_0 < n^{1/c}$. In this case we get that

$$\frac{f(k_0)}{\text{opt}_\Pi(G)} \leq \frac{f(k_0)}{k_0} < \mathcal{O}(k_0^{c-1}) = \mathcal{O}(n^{\frac{c-1}{c}}).$$

Since in both cases we have a ratio of $\mathcal{O}(n^{\frac{c-1}{c}})$, the lemma follows in Case 1.

Case 2: $f(k) = \beta k$.

As in Case 1, we start by finding k_0 defined as the smallest positive integer k such that algorithm A returns that $\text{opt}_\Pi(G) \leq f(k)$. If $k_0 = 0$ we proceed as in the first case. Otherwise, we return $f(k_0)$, or if the algorithm is constructive we return a solution S_0 (that is, a subset of vertices) such that $|S_0| \geq f(k_0)$. We have

$$\frac{f(k_0)}{\text{opt}_\Pi(G)} \leq \frac{f(k_0)}{k_0} \leq \beta,$$

and the lemma follows in Case 2. ◀

As a vertex-minimization problem whose decision version is in **NP** is a well-behaved problem, the hypothesis of [Lemma 20](#) is satisfied (by taking the identity function as upper bound function), and thus the following theorem is immediate by pipelining [Lemma 19](#) and [Lemma 20](#). The next theorem should be compared to [Theorem 10](#).

► **Theorem 21.** *Let Π be a vertex-minimization problem whose decision version is in **NP**.*

1. *For every real number $c > 1$, if Π admits a **lop-kernel** with $\mathcal{O}(k^c)$ vertices, then it admits a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^{\frac{c-1}{c}})$ on n -vertex graphs.*
2. *For every real number $c > 1$, if Π admits a **lop-kernel** with at most ck vertices, then it admits a polynomial-time value-approximation algorithm with ratio $(c + 1)$.*

*Moreover, if the **lop-kernel** is constructive, then the corresponding approximation algorithm is also constructive.*

As the framework of **lop-kernels** is mainly defined as a tool to get **lop-kernel** lower bounds from inapproximability, let us explicitly formulate the contrapositive of [Theorem 21](#). Note again that, when applying it to a concrete problem Π , the inapproximability of Π will rely on some complexity assumption, typically $\text{P} \neq \text{NP}$.

► **Corollary 22.** *Let Π be a vertex-minimization problem whose decision version is in NP.*

1. *For every real number $r \in (0, 1)$, if Π does not admit a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^r)$ on n -vertex graphs, then Π parameterized by the solution size does not admit a lop-kernel with $\mathcal{O}(k^{\frac{1}{1-r}})$ vertices.*
2. *For every real number $\beta > 1$, if Π does not admit a polynomial-time value-approximation algorithm with ratio β , then Π parameterized by the solution size does not admit a lop-kernel with $(\beta - 1 - \varepsilon)k$ vertices for any real number $\varepsilon > 0$.*

Moreover, if the non-existence of approximation algorithms only holds for constructive approximation algorithms, then the lower bound only holds for constructive lop-kernels.

4.3 Connection between lop-kernels and approximation algorithms for well-behaved minimization problems

The following lemma and theorem are a more general version of [Lemma 20](#) and [Theorem 21](#), respectively, for well-behaved minimization problems. The next lemma should be compared to [Lemma 12](#), and note that the obtained approximation ratios in the second item of both lemmas are different.

► **Lemma 23.** *Let Π be a well-behaved minimization problem, $a \in \mathbb{R}^+$, $u : \mathbb{N} \rightarrow \mathbb{N}$, and $f : \mathbb{N} \rightarrow \mathbb{N}$ be functions such that $u(n) = \mathcal{O}(n^a)$, u is polynomial-time computable, and f is computable. Suppose that Π has u as upper bound function u and that it admits an f -dual-approximation.*

1. *If $f(k) = \mathcal{O}(k^d)$ for some real number $d > 1$, then Π admits a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^{\frac{a(d-1)}{d}})$, where n is the size of the input.*
2. *If $f(k) = \lambda k^d + k$ for some real numbers $d \leq 1$ and $\lambda > 0$, then Π admits a polynomial-time value-approximation algorithm with ratio $\lambda + 1$.*

Moreover, if the dual-approximation algorithm is constructive, then the corresponding approximation algorithm is also constructive.

Proof. Let A be an f -dual-approximation algorithm for Π . For both cases in the statement of the lemma, we proceed to construct a polynomial-time approximation algorithm for Π with the claimed ratio.

Given an instance I of Π , we find $k_0 \in \{0, \dots, u(n)\}$ (recall that $n = |I|$) defined as the smallest positive integer k such that algorithm A returns that $\text{opt}_\Pi(G) \leq f(k)$. Note that k_0 can be found in polynomial time, as $|I|$ is polynomial-time computable and its value n is polynomially upper-bounded in the classical bit-size of the instance, and that $u(n)$ can be computed in polynomial time as well. Note that k_0 always exists, as we cannot have $\text{opt}_\Pi(I) > u(n)$. If $k_0 = 0$, then as by hypothesis Π satisfies property $C^{\min}$, we can verify in polynomial time whether $\text{opt}_\Pi(I) = 0$ and, if it is the case, we provide an optimal solution s with $\text{val}_\Pi(I, s) = 0$. Otherwise, we have $\text{opt}_\Pi(I) \geq 1$, and A returns $f(k_0)$, or a solution s such that $\text{val}_\Pi(I, s) \leq f(k_0)$ if the dual-approximation algorithm is constructive. In the first case (that is, if $f(k) = \mathcal{O}(k^d)$), as $\text{opt}_\Pi(I) \geq 1$, we have a ratio $f(0)$, implying the claimed ratio. In the second case (that is, if $f(k) = \lambda k^d + k$), $f(0) = 0$, so we even have an optimal solution.

Let us now assume that $k_0 \geq 1$, and recall that $\text{opt}_\Pi(I) \geq k_0$. We distinguish the two cases claimed in the statement of the lemma.

Case 1: $f(k) = \mathcal{O}(k^d)$.

Suppose first that $k_0 \geq n^{a/d}$. In this case we have

$$\frac{f(k_0)}{\text{opt}_\Pi(I)} \leq \frac{u(n)}{k_0} \leq \frac{\mathcal{O}(n^a)}{k_0} = \mathcal{O}\left(n^{\frac{a(d-1)}{d}}\right).$$

Otherwise, it holds that $k_0 < n^{a/d}$. In this case we get that

$$\frac{f(k_0)}{\text{opt}_\Pi(I)} = \frac{\mathcal{O}((k_0)^d)}{k_0} = \mathcal{O}((k_0)^{d-1}) = \mathcal{O}\left(n^{\frac{a(d-1)}{d}}\right).$$

Since in both cases we have a ratio of $\mathcal{O}(n^{\frac{a(d-1)}{d}})$, the lemma follows in Case 1.

Case 2: $f(k) = \lambda k^d + k$.

In this case we have

$$\frac{f(k_0)}{\text{opt}_\Pi(I)} \leq \frac{\lambda k_0^d + k_0}{k_0} = \lambda(k_0)^{d-1} + 1.$$

As $d \leq 1$, the last expression in the above equation is decreasing in k_0 , and as $k_0 \geq 1$, the maximum is reached for $k_0 = 1$, and the approximation ratio claimed in Case 2 follows. ◀

The next theorem, which should be compared to [Theorem 13](#), follows immediately by pipelining [Lemma 19](#) and [Lemma 23](#). Namely, starting with the hypothesis of [Theorem 24](#), we first apply [Lemma 19](#) and then [Lemma 23](#) with $d = ac$ and $\lambda = \alpha\beta^a$.

► **Theorem 24.** *Let Π be a well-behaved minimization problem, $a, c \in \mathbb{R}^+$, $u : \mathbb{N} \rightarrow \mathbb{N}$, and $s : \mathbb{N} \rightarrow \mathbb{N}$ be functions such that $u(n) = \mathcal{O}(n^a)$, $s(k) = \mathcal{O}(k^c)$, u is non-decreasing, and s and u are polynomial-time computable. Suppose that Π has u as upper bound function and that it admits a **lop**-kernel of size s , according to the same size function $|\cdot|$ associated with Π .*

1. *If $ac > 1$, then Π admits a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^{\frac{ac-1}{c}})$.*
2. *If $ac \leq 1$ and $\alpha, \beta \in \mathbb{R}^+$ are such that $u(n) \leq \alpha n^a$ and $s(k) \leq \beta k^c$, then Π admits a polynomial-time value-approximation algorithm with ratio $\alpha\beta^a + 1$.*

*Moreover, if the **lop**-kernel is constructive, then the corresponding approximation algorithm is also constructive.*

The discussion provided right after [Theorem 13](#) also applies to the above theorem. The contrapositive of [Theorem 24](#) yields the following corollary.

► **Corollary 25.** *Let Π be a well-behaved minimization problem with a non-decreasing polynomial-time computable upper bound function $u(n) = \mathcal{O}(n^a)$ for $a \in \mathbb{R}^+$. In what follows, the size of the instance, denoted by n , and the size of the kernel are defined according to the same size function $|\cdot|$ associated with Π .*

1. *For every real number $r \in (0, 1)$, if Π does not admit a polynomial-time value-approximation algorithm with ratio $\mathcal{O}(n^r)$, then Π parameterized by the solution size does not admit a **lop**-kernel of size $\mathcal{O}(k^{\frac{1}{a-r}})$.*
2. *Suppose that $u(n) \leq \alpha n^a$ for some $\alpha \in \mathbb{R}^+$. For every real number $\beta > 1$, if Π does not admit a polynomial-time value-approximation algorithm with ratio β , then Π parameterized by the solution size does not admit a **lop**-kernel of size $\beta' k^{c'}$ for any real numbers β', c' such that $ac' \leq 1$ and $\alpha\beta'^a + 1 \leq \beta$.*

*Moreover, if the non-existence of approximation algorithms only holds for constructive approximation algorithms, then the lower bound only holds for constructive **lop**-kernels.*

5 Applications of the framework of lop-kernels

In this section we provide several applications of the framework of lop-kernels introduced in Section 3 and Section 4.

Our first application concerns the MAXIMUM MINIMAL VERTEX COVER problem, defined in Section 2. Boria et al. [18] proved that MAXIMUM MINIMAL VERTEX COVER does not admit a polynomial-time $\mathcal{O}(n^{\frac{1}{2}-\varepsilon})$ -approximation algorithm for any $\varepsilon > 0$, unless $P = NP$. Hence, by applying Corollary 11 with $r = \frac{1}{2} - \varepsilon$ we obtain the following corollary, which matches the best known kernel having $\mathcal{O}(k^2)$ vertices [34].

► **Corollary 26.** MAXIMUM MINIMAL VERTEX COVER *parameterized by the solution size does not admit a lop-kernel with $\mathcal{O}(k^{2-\varepsilon})$ vertices for any $\varepsilon > 0$, unless $P = NP$.*

Our second application is similar to the first one. In the MAXIMUM MINIMAL FEEDBACK VERTEX SET problem, given an n -vertex graph G and an integer k , the objective is to decide if there exists a minimal feedback vertex set $S \subseteq V(G)$ (i.e., a set S such that $G \setminus S$ is a forest) of size at least k . Dublois et al. [31] recently proved that the problem does not admit a polynomial-time $\mathcal{O}(n^{\frac{2}{3}-\varepsilon})$ -approximation algorithm for any $\varepsilon > 0$, unless $P = NP$. Hence, by applying Corollary 11 with $r = \frac{2}{3} - \varepsilon$ we obtain the following corollary, which matches the best known kernel with $\mathcal{O}(k^3)$ vertices also provided by Dublois et al. [31].

► **Corollary 27.** MAXIMUM MINIMAL FEEDBACK VERTEX SET *parameterized by the solution size does not admit a lop-kernel with $\mathcal{O}(k^{3-\varepsilon})$ vertices for any $\varepsilon > 0$, unless $P = NP$.*

Our third application concerns a vertex-minimization problem. In the TREE DELETION SET problem, given a graph G and an integer k , the objective is to decide whether at most k vertices can be deleted from an n -vertex graph G in order to obtain a tree. It is known that this problem does not admit a polynomial-time $\mathcal{O}(n^{1-\varepsilon})$ -approximation for any $\varepsilon > 0$ unless $P \neq NP$ [56]. Corollary 22 implies the following.

► **Corollary 28.** TREE DELETION SET *parameterized by the solution size does not admit a polynomial lop-kernel, unless $P = NP$.*

The interesting fact is that TREE DELETION SET admits a kernel with $\mathcal{O}(k^4)$ vertices [38]. This kernel is the only non-artificial example of *non-lop-kernel* that we are aware of so far. Thus, the algebraic reduction rule presented by Giannopoulou et al. [38], which is based on identifying a subset of linear equations of appropriate size that captures all solutions of size at most k , cannot be (even transformed to) a lop-rule.

Our last application deals with the MAXIMUM INDEPENDENT SET problem restricted to K_t -free graphs, for an integer $t \geq 3$. Ramsey's theorem [52] implies that, given a K_t -free graph on n vertices, it is always possible to find in polynomial time an independent set of size at least $n^{\frac{1}{t-1}}$. This directly implies a polynomial-time $n^{\frac{t-2}{t-1}}$ -approximation algorithm for MAXIMUM INDEPENDENT SET on K_t -free graphs, and a constructive lop-kernel of size k^{t-1} (indeed, if the input graph has size at least k^{t-1} , we can safely declare it a *yes-instance*). To the best of our knowledge, improving this trivial approximation factor is still open, and the only known inapproximability result is the bound $\mathcal{O}(n^{\frac{1}{4}-\varepsilon})$ on triangle-free graphs recently proved by Bonnet et al. [17], which relies on the hypothesis that $NP \not\subseteq BPP$. In the same paper [17], the authors state the following conjecture, called the “Improved Approximation Conjecture”: for every fixed graph H , there exists a constant $\varepsilon > 0$ such that MAXIMUM INDEPENDENT SET admits a (randomized) polynomial-time $n^{1-\varepsilon}$ -approximation algorithm on H -free n -vertex graphs. We state the following conjecture.

► **Conjecture 29.** *For every fixed graph H , the MAXIMUM INDEPENDENT SET problem restricted to H -free graphs admits a polynomial log-kernel.*

Corollary 11, combined with the above discussion and the inapproximability result of Bonnet et al. [17] on triangle-free graphs, imply the following.

► **Corollary 30.** *The following claims hold:*

- MAXIMUM INDEPENDENT SET parameterized by the solution size does not admit a log-kernel with $\mathcal{O}(k^{4-\varepsilon})$ vertices on triangle-free graphs for any $\varepsilon > 0$, unless $\text{NP} \subseteq \text{BPP}$.
- For every real number $\varepsilon > 0$ and every integer $t \geq 3$, a log-kernel with $\mathcal{O}(k^{t-1-\varepsilon})$ vertices for MAXIMUM INDEPENDENT SET on K_t -free graphs would improve the best known approximation ratio $n^{\frac{t-2}{t-1}}$ that follows from Ramsey's theorem [52].
- Conjecture 29 implies the Improved Approximation Conjecture of Bonnet et al. [17].

6 An attempt to obtain a linear kernel for MMVC

In this section we briefly explain the flaw in the linear kernel for MAXIMUM MINIMAL VERTEX COVER (MMVC) claimed by Fernau [34, Corollary 4.25], and that is based on joint unpublished work with Dehne, Fellows, Prieto, and Rosamond. The kernelization algorithm is a small modification of a linear kernel for the NONBLOCKER SET problem presented by Ore [51]. A set of vertices S of a graph G is a *nonblocker* if its complement is a dominating set of G , that is, for every $u \in S$ there exists $v \notin S$ with $\{u, v\} \in E(G)$. In the NONBLOCKER SET problem, we are given a graph G and an integer parameter k , and the goal is to decide whether G contains a nonblocker of size at least k . Suppose for simplicity that G is connected. The idea is to consider an arbitrary spanning tree T of G , root it arbitrarily at a vertex r , and partition $V(G) = V_0 \uplus V_1$ such that the vertices in V_0 (resp. V_1) are within even (resp. odd) distance from r in T . By construction, each of V_0 and V_1 is a nonblocker in G , so if one of them has size at least k , we can answer “yes”, and otherwise $|V(G)| \leq 2k$ and we are done.

Back to MMVC, it is observed in [34, Reduction rule 24] that a simple reduction rule allows to assume that no connected component of G is a clique (in particular, an isolated vertex). Assume again for simplicity that G is connected. It is then claimed in [34] that, using the same algorithm as for NONBLOCKER SET, the largest of V_0 and V_1 , say V_0 , can be always completed into a minimal vertex cover of G , which would immediately yield a kernel of size at most $2k$ for MMVC. Unfortunately, this claim is not true: when adding new vertices to V_0 in order to make it a vertex cover of G , we may lose the minimality property, and some vertices may need to be removed. For instance, let G be the graph obtained from a triangle on vertices u, v, w by adding $p \geq 2$ pendant vertices to each of u, v , and w . Let T be the spanning tree obtained from G by removing the edge $\{v, w\}$, and root T at vertex u . Then $|V_0| = 1 + 2p$ and $|V_1| = 2 + p$, so $|V_0| > |V_1|$, and note that the edge $\{v, w\}$ is the only edge of G not covered by V_0 . But adding either of v or w to V_0 , say v , results in a non-minimal vertex cover of G , and therefore the p pendant vertices adjacent to v have to be removed from V_0 , which yields a set of size $2 + p < \frac{|V(G)|}{2} = \frac{3+3p}{2}$, where we have used that $p \geq 2$. In fact, deciding whether a set $S \subseteq V(G)$ can be extended to a minimal vertex cover of G is an NP-complete problem [20].

7 Subquadratic kernels for MMVC on particular graph classes

In this section we present subquadratic kernels for MAXIMUM MINIMAL VERTEX COVER restricted to particular graph classes when the parameter is the solution size k . Namely, in

Subsection 7.1 we provide kernels using the Erdős-Hajnal property, and in Subsection 7.2 we provide further observations about other graph classes.

7.1 Kernels using the Erdős-Hajnal property

For a constant $\delta > 0$, a graph H is said to satisfy the *Erdős-Hajnal property with constant δ* if every H -free graph G on n vertices contains either a clique or an independent set of size n^δ . The (still open) Erdős-Hajnal conjecture [33] states that every graph H satisfies the Erdős-Hajnal property. As reported by Chudnovsky [22], the Erdős-Hajnal conjecture has been verified for only a small number of graphs, namely all graphs on at most four vertices, the *bull* (i.e., the graph obtained by adding a pendant vertex to two different vertices of a triangle), the complete graphs, and every graph that can be constructed from them using the so-called *substitution operation* [3], which we define later.

Since our goal is to use the Erdős-Hajnal property in order to obtain kernels for MAXIMUM MINIMAL VERTEX COVER, we need an algorithmic version of it. As defined by Bonnet et al. [17], for a constant $\delta > 0$, a graph H is said to satisfy the *constructive Erdős-Hajnal property with constant δ* if there exists an algorithm that takes as input an H -free graph G on n vertices, and outputs in polynomial-time a clique or an independent set of G of size at least n^δ . Fortunately for our purposes, all the graphs H shown to satisfy the Erdős-Hajnal property so far, also satisfy its constructive version [17].

In the following simple lemma we show that, if H is a graph satisfying the constructive Erdős-Hajnal property, then the vertex set of an H -free graph can be partitioned in polynomial time into “few” cliques or independent sets. This partition will then be used to obtain subquadratic kernels on H -free graphs for several graphs H .

► **Lemma 31.** *Let H be a graph satisfying the constructive Erdős-Hajnal property with constant δ . The vertex set of any H -free graph G on n vertices can be partitioned in polynomial time into a collection of cliques $\mathcal{C}$ and a collection of independent sets $\mathcal{I}$ such that $|\mathcal{C}| + |\mathcal{I}| \leq \left(\frac{1}{2^{(1-\delta)} - 1}\right) \cdot n^{1-\delta}$.*

Proof. Let G be an H -free graph on n vertices. We initialize $X_0 = V(G)$, $\mathcal{C} = \mathcal{I} = \emptyset$, and we run the following procedure as far as $|X_0| \geq 1$:

Find in polynomial time a clique or an independent set Y in $G[X_0]$ with $|Y| \geq |X_0|^\delta$. Note that this is possible since $G[X_0]$ is an H -free graph for any $X_0 \subseteq V(G)$.
Add Y to $\mathcal{C}$ or to $\mathcal{I}$ depending on whether Y is a clique or an independent set, respectively (if $|Y| = 1$, choose $\mathcal{C}$ or $\mathcal{I}$ arbitrarily). Update $X_0 \leftarrow X_0 \setminus Y$.

Clearly, the above algorithm terminates in polynomial time. It remains to bound $|\mathcal{C}| + |\mathcal{I}|$, which is equal to the number of iterations of the algorithm. To this end, for a positive integer i , we say that an iteration belongs to *step i* of the algorithm if the current set X_0 at the start of the iteration satisfies $\frac{n}{2^i} < |X_0| \leq \frac{n}{2^{i-1}}$. We denote by t_i the number of iterations of the algorithm within step i . By definition, $|\mathcal{C}| + |\mathcal{I}| = \sum_{i=1}^{\infty} t_i$. Let Y be a clique or an independent set found by the algorithm within step i . Since the current set X_0 satisfies $|X_0| > \frac{n}{2^i}$, we have that $|Y| > \left(\frac{n}{2^i}\right)^\delta$. And since the sum of the sizes of the sets found before the last iteration of step i is at most $\frac{n}{2^i}$, it follows that $t_i \leq \left(\frac{n}{2^i}\right)^{1-\delta}$. Note that, in particular, $t_i = 0$ for $i > \lceil \log n \rceil$. Therefore, we conclude that

$$|\mathcal{C}| + |\mathcal{I}| = \sum_{i=1}^{\infty} t_i \leq \sum_{i=1}^{\infty} \left(\frac{n}{2^i}\right)^{1-\delta} = n^{1-\delta} \cdot \sum_{i=1}^{\infty} \left(\frac{1}{2^{1-\delta}}\right)^i = n^{1-\delta} \cdot \left(\frac{1}{2^{(1-\delta)} - 1}\right),$$

and the lemma follows. ◀

We are now ready to present the subquadratic kernel on bull-free graphs. Note that, since bipartite graphs are bull-free, MMVC restricted to bull-free graphs is NP-hard by [14] (or by Theorem 42). In the kernels presented in this section, since we can easily obtain explicit constants, we decided not to use the big-O notation.

► **Theorem 32.** *The MAXIMUM MINIMAL VERTEX COVER problem parameterized by k restricted to bull-free graphs admits a kernel with at most $c(k-1)^{7/4} + k - 1$ vertices, where $c = \frac{2}{2^{3/4}-1} < 3$.*

Proof. Let (G, k) be an instance of the MAXIMUM MINIMAL VERTEX COVER problem, where G is a bull-free graph. Recall that by Lemma 2 we can assume that the maximum degree of G is at most $k-1$. We start by finding greedily, starting from $V(G)$, a minimal vertex cover X of G . Note that X can be easily found in polynomial time by Observation 1. If $|X| \geq k$, we conclude that (G, k) is a yes-instance, so we can assume that $|X| \leq k-1$. Let $S = V(G) \setminus X$ and note that S is an independent set.

Since the bull satisfies the constructive Erdős-Hajnal property with constant $\delta = \frac{1}{4}$ [17, 23], we can apply Lemma 31 to the bull-free graph $G[X]$ and obtain in polynomial time a partition of X into a collection of cliques $\mathcal{C}$ and a collection of independent sets $\mathcal{I}$ such that $|\mathcal{C}| + |\mathcal{I}| \leq d \cdot |X|^{3/4} \leq d \cdot (k-1)^{3/4}$, where $d = \frac{1}{2^{3/4}-1} < 1.47$. Since we can assume that G has no isolated vertices, as they can be safely removed without affecting the type of the instance, it follows that

$$S = \bigcup_{C \in \mathcal{C}} N_S(C) \cup \bigcup_{I \in \mathcal{I}} N_S(I). \quad (1)$$

Hence, our objective is to bound $|N_S(Y)|$ for every $Y \in \mathcal{C} \cup \mathcal{I}$. Suppose first that $I \in \mathcal{I}$ is an independent set. From Lemma 2, if $|N_S(I)| \geq k$ we can conclude that (G, k) is a yes-instance, so we can assume henceforth that

$$\text{for every independent set } I \in \mathcal{I}, \text{ it holds } |N_S(I)| \leq k-1. \quad (2)$$


Suppose now that $C \in \mathcal{C}$ is a clique. We partition $N_S(C) = S_C^1 \uplus S_C^2$ as follows. Let S_C^1 be an inclusion-wise maximal set of vertices in $N_S(C)$ such that for any two (not necessarily distinct) vertices $x, y \in S_C^1$, $|N_C(x) \cup N_C(y)| \leq |C| - 1$. That is, S_C^1 is a maximal set in $N_S(C)$ such that the neighborhoods of its vertices pairwise do *not* cover the whole clique C . We let $S_C^2 = N_S(C) \setminus S_C^1$. The following is the crucial property of the set S_C^1 .

► **Claim 33.** The vertices in S_C^1 can be ordered $x_1, \dots, x_p$ so that $N_C(x_i) \subseteq N_C(x_j)$ whenever $i \leq j$.

Proof. In order to prove the claim, it is sufficient to prove that, for any two vertices $x, y \in S_C^1$, either $N_C(x) \subseteq N_C(y)$ or $N_C(y) \subseteq N_C(x)$. Suppose for the sake of contradiction that there exist two vertices $u \in N_C(x) \setminus N_C(y)$ and $v \in N_C(y) \setminus N_C(x)$. By definition of the set S_C^1 , there exists a vertex $w \in C \setminus (N_C(x) \cup N_C(y))$. But then the vertices x, y, u, v, w induce a bull as illustrated in Figure 1, contradicting the hypothesis that G is bull-free. ◀

Claim 33 implies in particular that, unless $S_C^1 = \emptyset$, there exists a vertex $u \in \bigcap_{x \in S_C^1} N_C(x)$. Since u has degree at most $k-1$ in G , and each vertex $x \in S_C^1$ is adjacent to u , it follows that $|S_C^1| \leq k-1$.

Let us now focus on the set S_C^2 . The definition of the set S_C^1 together with Claim 33 imply that there exists a vertex $z \in C \setminus \bigcup_{y \in S_C^1} N_C(y)$. Consider now an arbitrary vertex $x \in S_C^2$.



■ **Figure 1** Configuration considered in the proof of [Claim 33](#) and a vertex $z \in \bigcap_{x \in S_C^2} N_C(x)$.

Since x could not be added to S_C^1 , there exists a vertex $y \in S_C^1$ such that $N_C(x) \cup N_C(y) = C$. But since $z \in C \setminus \bigcup_{y \in S_C^1} N_C(y)$, necessarily $z \in N_C(x)$. It follows that $z \in \bigcap_{x \in S_C^2} N_C(x)$ (see [Figure 1](#)). Using again the fact that z has degree at most $k - 1$ in G , we obtain that $|S_C^2| \leq k - 1$. Summarizing, we have that

$$\text{for every clique } C \in \mathcal{C}, \text{ it holds } |N_S(C)| = |S_C^1| + |S_C^2| \leq 2(k - 1). \quad (3)$$

Putting all pieces together, Equations (1), (2), and (3) and the fact that $|X| \leq k - 1$ and $|\mathcal{C}| + |\mathcal{I}| \leq d \cdot |X|^{3/4}$ imply that, unless we have already concluded that (G, k) is a **yes**-instance, we have that

$$\begin{aligned} |V(G)| &= |X| + |S| = |X| + \left| \bigcup_{C \in \mathcal{C}} N_S(C) \right| + \left| \bigcup_{I \in \mathcal{I}} N_S(I) \right| \\ &\leq |X| + (|\mathcal{C}| + |\mathcal{I}|) \cdot \max_{Y \in \mathcal{C} \cup \mathcal{I}} |N_S(Y)| \leq k - 1 + d \cdot (k - 1)^{3/4} \cdot 2(k - 1) \\ &= 2d \cdot (k - 1)^{7/4} + k - 1, \end{aligned}$$

and the theorem follows. ◀

It is easy to prove that, for every integer $t \geq 2$, every K_t -free graph G on n vertices has an independent set of size $n^{\frac{1}{t-1}}$, by induction on t : for $t = 2$ the statement is trivial, and if $t \geq 3$, then either $\Delta(G) < n^{\frac{t-2}{t-1}}$, and an independent set of size $n^{\frac{1}{t-1}}$ can be found greedily by adding any vertex to it and deleting its neighborhood, or there exists a vertex $v \in V(G)$ of degree at least $n^{\frac{t-2}{t-1}}$, in which case an independent set of size $n^{\frac{1}{t-1}}$ can be found applying the inductive hypothesis to the K_{t-1} -free graph $G[N(v)]$. Clearly, this proof can be translated to a polynomial-time algorithm to find an independent set of the appropriate size. Therefore, for any integer $t \geq 2$, K_t satisfies the constructive Erdős-Hajnal property with constant $\delta = \frac{1}{t-1}$. The proof of the following theorem is a simplified version of that of [Theorem 32](#). Note that, since bipartite graphs are K_t -free for every $t \geq 3$, MMVC is NP-hard on K_t -free graphs [\[14\]](#).

► **Theorem 34.** *For every integer $t \geq 3$, the MAXIMUM MINIMAL VERTEX COVER problem parameterized by k restricted to K_t -free graphs admits a kernel with at most $c_t(k-1)^{\frac{2t-3}{t-1}} + k - 1$ vertices, where $c_t = \frac{t-1}{2^{\frac{t-2}{t-1}} - 1}$.*

Proof. As in the proof of [Theorem 32](#), given an instance (G, k) of MAXIMUM MINIMAL VERTEX COVER, where G is a K_t -free graph, we partition $V(G) = X \uplus S$, where X is a minimal vertex cover of G with $|X| \leq k - 1$, and we use [Lemma 31](#) to partition X into two collections $\mathcal{C}$ and $\mathcal{I}$ of cliques and independent sets, respectively, with $|\mathcal{C}| + |\mathcal{I}| \leq d_t \cdot |X|^{\frac{t-2}{t-1}}$, where $d_t = \frac{1}{2^{\frac{t-2}{t-1}} - 1}$. Equations (1) and (2) still hold, but now we have a much simpler version

of Equation (3): if $C \in \mathcal{C}$ is a clique then, since G is K_t -free, necessarily $|C| \leq t - 1$, which together with the fact that $\Delta(G) \leq k - 1$ yield

$$\text{for every clique } C \in \mathcal{C}, \text{ it holds } |N_S(C)| = (t - 1)(k - 1). \quad (4)$$

Combining Equations (1), (2), and (4) we get

$$|V(G)| \leq |X| + (|\mathcal{C}| + |\mathcal{I}|) \cdot \max_{Y \in \mathcal{C} \cup \mathcal{I}} |N_S(Y)| \leq k - 1 + d_t \cdot (k - 1)^{\frac{t-2}{t-1}} \cdot (t - 1)(k - 1),$$

and the theorem follows. $\blacktriangleleft$

We now extend the results of Theorem 32 and Theorem 34 to more general excluded induced graphs H , by making use of the aforementioned substitution operation. As defined by Alon et al. [3], for two graphs H_1 and H_2 on disjoint vertex sets, we say that H is *obtained from H_1 by substituting H_2 for $v \in V(H_1)$* (or just *obtained from H_1 by substituting H_2* if the vertex v in question is not important) if

- $V(H) = (V(H) \setminus \{v\}) \cup V(H_2)$,
- $H[V(H_2)] = H_2$,
- $H[V(H_1) \setminus \{v\}] = H_1 \setminus v$, and
- $u \in V(H_1)$ is adjacent in H to $w \in V(H_2)$ if and only if u is adjacent in H_1 to v .

Alon et al. [3] proved that if two graphs H_1 and H_2 satisfy Erdős-Hajnal property and H is obtained from H_1 by substituting H_2 , then H satisfies the Erdős-Hajnal property as well. More precisely, by following the details in the proof of [3, Theorem 2.1], we can derive that if H_1 and H_2 satisfy Erdős-Hajnal property with constants δ_1 and δ_2 , respectively, then H satisfies the Erdős-Hajnal property with constant $\delta = \frac{\delta_2}{\delta_1 + |V(H_1)| \cdot \delta_2}$. The same applies to the constructive version of the Erdős-Hajnal property.

For an integer $t \geq 2$, we define the t -bull as the graph obtained from K_t by adding a pendant vertex to two different vertices of the clique. Note that the 2-bull is equal to P_4 and that the 3-bull is equal to the bull. Note also that, for every $t \geq 3$, the t -bull is obtained from the bull by substituting K_{t-2} for the vertex of degree two of the bull. Therefore, by the discussion in the above paragraph, since the bull and K_{t-2} satisfy the constructive Erdős-Hajnal property with constants $\frac{1}{4}$ and $\frac{1}{t-3}$, respectively, it follows that, for every $t \geq 4$, the t -bull satisfies the constructive Erdős-Hajnal property with constant

$$\delta_t = \frac{\frac{1}{t-3}}{\frac{1}{4} + \frac{5}{t-3}} = \frac{4}{t + 17}.$$

The proof of the next theorem follows again (and generalizes) that of Theorem 32. Note that Theorem 35 corresponds to the particular case $t = 3$ of Theorem 34. Note also that bipartite graphs are t -bull-free for $t \geq 3$, hence MMVC is NP-hard on t -bull-free graphs for $t \geq 3$ [14]. On the other hand, 2-bull-free graphs are exactly P_4 -free graphs, also called cographs, which have cliquewidth at most two, hence by Observation 40 (proved later in Subsection 7.2) MMVC can be solved in polynomial time on this class.

► **Theorem 35.** *For every integer $t \geq 3$, the MAXIMUM MINIMAL VERTEX COVER problem parameterized by k restricted to t -bull-free graphs admits a kernel with at most $c_t(k - 1)^{2 - \delta_t} + k - 1$ vertices, where $\delta_3 = \frac{1}{4}$ and $\delta_t = \frac{4}{t + 17}$ for $t \geq 4$, and $c_t = \frac{t-1}{2^{(1-\delta_t)} - 1}$.*

Proof. As in the proof of Theorem 32, given an instance (G, k) of MAXIMUM MINIMAL VERTEX COVER, where G is a t -bull-free graph, we partition $V(G) = X \uplus S$, where X is a minimal vertex cover of G with $|X| \leq k - 1$, and we use Lemma 31 to partition X into two

collections $\mathcal{C}$ and $\mathcal{I}$ of cliques and independent sets, respectively, with $|\mathcal{C}| + |\mathcal{I}| \leq d_t \cdot |X|^{1-\delta_t}$, where $\delta_3 = \frac{1}{4}$ and $\delta_t = \frac{4}{t+17}$ for $t \geq 4$, and $d_t = \frac{1}{2^{(1-\delta_t)}-1}$ for every $t \geq 3$. Equations (1) and (2) still hold for every integer $t \geq 3$, but now we need slightly more involved arguments to obtain an appropriate version of Equation (3) for every $t \geq 3$.

To this end, suppose that $C \in \mathcal{C}$ is a clique. We partition $N_S(C) = S_C^1 \uplus S_C^2$ as follows. Let S_C^1 be an inclusion-wise maximal set of vertices in $N_S(C)$ such that for any two (not necessarily distinct) vertices $x, y \in S_C^1$, $|N_C(x) \cup N_C(y)| \leq |C| - (t-2)$. That is, S_C^1 is a maximal set in $N_S(C)$ such that the neighborhoods of its vertices pairwise leave at least $t-2$ uncovered vertices in the clique C . We let $S_C^2 = N_S(C) \setminus S_C^1$. The set S_C^1 satisfies exactly the same crucial property as for the case $t = 3$ (see Claim 33).

▷ **Claim 36.** For every integer $t \geq 3$, the vertices in S_C^1 can be ordered $x_1, \dots, x_p$ so that $N_C(x_i) \subseteq N_C(x_j)$ whenever $i \leq j$.

Proof. In order to prove the claim, it is sufficient to prove that, for any two vertices $x, y \in S_C^1$, either $N_C(x) \subseteq N_C(y)$ or $N_C(y) \subseteq N_C(x)$. Suppose for the sake of contradiction that there exist two vertices $u \in N_C(x) \setminus N_C(y)$ and $w \in N_C(y) \setminus N_C(x)$. By definition of the set S_C^1 , there exist $t-2$ vertices $w_1, \dots, w_{t-2} \in C \setminus (N_C(x) \cup N_C(y))$. But then the vertices $x, y, u, v, w_1, \dots, w_{t-2}$ induce a t -bull, contradicting the hypothesis that G is t -bull-free. ◀

Claim 36 implies in particular that, unless $S_C^1 = \emptyset$, there exists a vertex $u \in \bigcap_{x \in S_C^1} N_C(x)$. Since u has degree at most $k-1$ in G , and each vertex $x \in S_C^1$ is adjacent to u , it follows that $|S_C^1| \leq k-1$.

Let us now focus on the set S_C^2 . The definition of the set S_C^1 together with Claim 36 imply that there exist at least $t-2$ vertices $z_1, \dots, z_{t-2} \in C \setminus \bigcup_{y \in S_C^1} N_C(y)$. Consider now an arbitrary vertex $x \in S_C^2$. Since x could not be added to S_C^1 , there exists a vertex $y \in S_C^1$ such that $|N_C(x) \cup N_C(y)| \geq |C| - (t-1)$. But since $z_1, \dots, z_{t-2} \in C \setminus \bigcup_{y \in S_C^1} N_C(y)$, there exists an index $j \in [t-2]$ such that $z_j \in N_C(x)$. That is, every vertex $x \in S_C^2$ is adjacent to at least one of the vertices $z_1, \dots, z_{t-2}$. Using again the fact that each of the vertices $z_1, \dots, z_{t-2}$ has degree at most $k-1$ in G , we obtain that $|S_C^2| \leq (t-2)(k-1)$. Summarizing, we have that

$$\text{for every clique } C \in \mathcal{C}, \text{ it holds } |N_S(C)| = |S_C^1| + |S_C^2| \leq (t-1)(k-1). \quad (5)$$

Putting all pieces together, Equations (1), (2), and (5) and the fact that $|X| \leq k-1$ and $|\mathcal{C}| + |\mathcal{I}| \leq d_t \cdot |X|^{1-\delta_t}$ imply that, unless we have already concluded that (G, k) is a yes-instance, we have that

$$|V(G)| \leq |X| + (|\mathcal{C}| + |\mathcal{I}|) \cdot \max_{Y \in \mathcal{C} \cup \mathcal{I}} |N_S(Y)| \leq k-1 + d_t \cdot (k-1)^{1-\delta_t} \cdot (t-1)(k-1),$$

and the theorem follows. ◀

Let the *paw* be the graph obtained from a triangle by adding a pendant edge. Gyárfás [42] showed that the paw satisfies the constructive Erdős-Hajnal property with constant $\delta = \frac{1}{3}$. Note that bipartite graphs are paw-free, hence MMVC is NP-hard on paw-free graphs [14].

► **Theorem 37.** *The MAXIMUM MINIMAL VERTEX COVER problem parameterized by k restricted to paw-free graphs admits a kernel with at most $c(k-1)^{5/3} + k-1$ vertices, where $c = \frac{2}{2^{2/3}-1} < 3.41$.*

Proof. Given an instance (G, k) of MAXIMUM MINIMAL VERTEX COVER, where G is a paw-free graph, we again partition $V(G) = X \uplus S$, where X is a minimal vertex cover of G with $|X| \leq k - 1$, and we use Lemma 31 to partition X into two collections $\mathcal{C}$ and $\mathcal{I}$ of cliques and independent sets, respectively, with $|\mathcal{C}| + |\mathcal{I}| \leq d \cdot |X|^{2/3}$, where $d = \frac{1}{2^{2/3}-1}$. Equations (1) and (2) still hold, and we can again obtain in a simpler way an appropriate version of Equation (3). Indeed, let $C \in \mathcal{C}$ be a clique, and our goal is to bound $|N_S(C)|$. If $|C| = 1$ then by the fact that $\Delta(G) \leq k - 1$ we get that $|N_S(C)| \leq k - 1$, so assume that $|C| \geq 2$. Suppose for the sake of contradiction that there exists a vertex $v \in N_S(C)$ such that $|N_C(v)| \leq |C| - 2$. Let $w \in N_C(v)$ and let z_1, z_2 be two vertices in $C \setminus N_C(v)$. Then the vertices v, w, z_1, z_2 induce a paw, contradicting the hypothesis that G is paw-free. Therefore, for every vertex $v \in N_S(C)$ it holds that $|N_C(v)| \geq |C| - 1$. Hence, the number of edges in G between C and $N_S(C)$ is at least $|N_S(C)| \cdot (|C| - 1)$ and, since $\Delta(G) \leq k - 1$, at most $|C| \cdot (k - 1)$. Using that $|C| \geq 2$, it follows that

$$\text{for every clique } C \in \mathcal{C}, \text{ it holds } |N_S(C)| \leq \frac{|C|}{|C| - 1} \cdot (k - 1) \leq 2(k - 1). \quad (6)$$

Putting all pieces together, Equations (1), (2), and (6) and the fact that $|X| \leq k - 1$ and $|\mathcal{C}| + |\mathcal{I}| \leq d \cdot |X|^{2/3}$ imply that, unless we have already concluded that (G, k) is a yes-instance, we have that

$$|V(G)| \leq |X| + (|\mathcal{C}| + |\mathcal{I}|) \cdot \max_{Y \in \mathcal{C} \cup \mathcal{I}} |N_S(Y)| \leq k - 1 + d \cdot (k - 1)^{2/3} \cdot 2(k - 1),$$

and the theorem follows. $\blacktriangleleft$

7.2 Remarks on other graph classes

In this subsection we provide additional observations about the complexity of the MAXIMUM MINIMAL VERTEX COVER problem restricted to special graph classes.

► **Lemma 38.** *For every integer $t \geq 1$, the MAXIMUM MINIMAL VERTEX COVER problem parameterized by k restricted to $K_{1,t}$ -free graphs admits a kernel with at most $t(k - 1)$ vertices.*

Proof. Given an instance (G, k) of MAXIMUM MINIMAL VERTEX COVER, where G is a $K_{1,t}$ -free graph, we again partition $V(G) = X \uplus S$, where X is a minimal vertex cover of G with $|X| \leq k - 1$. Since G is $K_{1,t}$ -free and S is an independent set, it holds that for every $v \in X$, $|N_S(v)| \leq t - 1$, and since we can assume that G contains no isolated vertex, we obtain that $|V(G)| = |X| + |\bigcup_{v \in X} N_S(v)| \leq k - 1 + (t - 1)(k - 1) = t(k - 1)$. $\blacktriangleleft$

Let $\mathcal{C}$ be a graph class such that there exists a polynomial-time algorithm that, given a graph $G \in \mathcal{C}$, outputs a proper coloring of the vertices of G using at most c colors, for some integer $c \geq 1$. We say that such a graph class $\mathcal{C}$ is **poly- χ - c -bounded**. Examples of poly- χ - c -bounded classes are planar graphs, minor-free graphs, or, more generally, graphs of bounded expansion. We note that Fernau [34, Corollary 4.14] provides a similar observation for the particular case of planar graphs.

► **Lemma 39.** *For every integer $c \geq 1$, the MAXIMUM MINIMAL VERTEX COVER problem parameterized by k restricted to the class of poly- χ - c -bounded graphs admits a kernel with at most $c(k - 1)$ vertices.*

Proof. Given an instance (G, k) of MMVC, where G belongs to a poly- χ - c -bounded class, we first compute in polynomial time a proper vertex coloring of G using at most c colors. We may clearly assume that G has no isolated vertices, as such vertices can be safely removed.

Let $V(G) = S_1 \uplus \dots \uplus S_c$ be the corresponding partition of $V(G)$ into independent sets. By Lemma 2, for every $i \in [c]$ there exists a minimal vertex cover of G that contains $N(S_i)$. Hence, if for some $i \in [c]$ we have that $|N(S_i)| \geq k$, we can safely answer “yes”, so we may assume that, for every $i \in [c]$, $|N(S_i)| \leq k - 1$. Since G has no isolated vertices and every set S_i is an independent set, it follows that $V(G) = \bigcup_{i \in [c]} N(S_i)$, so we have that $|V(G)| \leq \sum_{i \in [c]} |N(S_i)| \leq c(k - 1)$. $\blacktriangleleft$

Another graph class $\mathcal{K}$ that allows for linear kernels is defined such that, for every graph $G \in \mathcal{K}$, the minimum size of a dominating set of G is equal to the size of a minimum *independent* dominating set of G . We furthermore ask $\mathcal{K}$ to be hereditary. Such graphs have been studied, for instance, in [2, 54], and include in particular $K_{1,3}$ -free graphs (note that a generalization to $K_{1,t}$ -graphs is given in Lemma 38). Let us see why the class $\mathcal{K}$ allows for a linear kernel. As discussed at the end of Section 8, the complement of a dominating set is called a nonblocker, and the NONBLOCKER SET problem admits a linear kernel [34]. On the other hand, the complement of an independent dominating set is a minimal vertex cover. Hence, if $G \in \mathcal{K}$, an instance (G, k) of NONBLOCKER SET is positive if and only if (G, k) is a positive instance of MMVC. Note the linear kernel for the NONBLOCKER SET problem discussed at the end of Section 8 outputs a subgraph G' of G , and we have that $G' \in \mathcal{K}$ since $\mathcal{K}$ is hereditary. Hence, the equivalence between NONBLOCKER SET and MMVC also holds for G' , and it follows that this kernel is also a linear kernel for MMVC restricted to graphs in $\mathcal{K}$.

Our last contribution in this section concerns graph classes of bounded *cliquewidth*. Cliquewidth, which we do not need to define here, is a graph parameter that is “smaller” than treewidth in the sense that graph classes of bounded treewidth have also bounded cliquewidth (the opposite is not true, as cliques have cliquewidth one but unbounded treewidth); see [24] for the formal definition.

The variation of *monadic second order* logic of graphs called MSO_1 is defined by a syntax that includes the logical connectives $\vee, \wedge, \neg$, variables for vertices, edges, sets of vertices (but *not* sets of edges), the quantifiers $\forall, \exists$ that can be applied to these variables, and the binary relations expressing whether a vertex belongs to a set, whether an edge is incident to vertex, whether two vertices are adjacent, and whether two sets are equal. It is well-known that finding a minimum or maximum weight vertex set that satisfies a given graph property expressed in MSO_1 can be solved in linear time on graphs of cliquewidth bounded by a constant [6, 24].

► **Observation 40.** *The MAXIMUM MINIMAL VERTEX COVER problem can be expressed in MSO_1 , and therefore it can be solved in linear time when restricted to any graph class of cliquewidth bounded by a constant.*

Proof. Given a graph G , we can express the property of a vertex set S being a minimal vertex cover of G in the syntax of MSO_1 as follows: for every pair of vertices u, v such that u is adjacent to v , $u \in S$ or $v \in S$ (this guarantees that S is a vertex cover of G), and for every vertex $v \in V(G)$, $v \notin S$ or there exists a vertex u adjacent to v such that $u \notin S$ (this guarantees, by Observation 1, that S is minimal). $\blacktriangleleft$

Let the *diamond* be the graph obtained from K_4 by removing an edge. Since Brandstadt [19] proved that $\{P_5, \text{diamond}\}$ -free graphs have bounded cliquewidth, from Observation 40 we immediately get the following corollary.

► **Corollary 41.** *The MAXIMUM MINIMAL VERTEX COVER problem restricted to $\{P_5, \text{diamond}\}$ -free graphs can be solved in linear time.*


8 Ruling out polynomial kernels for MMVC for smaller parameters

In this section we rule out, assuming that $\text{NP} \not\subseteq \text{coNP}/\text{poly}$, the existence of polynomial kernels for MMVC parameterized by the size of a minimum vertex cover of the input graph. As mentioned in the introduction, the reduction given in [Theorem 42](#) also provides an alternative proof of the NP-completeness of MMVC on bipartite graphs, which also follows from [\[14\]](#). We note that the existing NP-hardness reductions for MMVC, such as the one in [\[14\]](#), do not seem to be easily modifiable so to yield the non-existence of polynomial kernels.

► **Theorem 42.** *The MAXIMUM MINIMAL VERTEX COVER problem parameterized by the size of a minimum vertex cover (or of a maximum matching) of the input graph does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP}/\text{poly}$, even restricted to bipartite graphs.*

Proof. We present a PPT from MONOTONE SAT parameterized by the number of variables, which is also an NP-completeness reduction. The MONOTONE SAT problem is the restriction of the SAT problem to formulas in which the literals in each clause are either all positive or all negative. This problem is well-known to be NP-complete [\[37\]](#), and it is easy to see that, when parameterized by the number of variables, it does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP}/\text{poly}$. Indeed, Fortnow and Santhanam [\[36\]](#) proved that the SAT problem parameterized by the number of variables does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP}/\text{poly}$, and the classical reduction from SAT to MONOTONE SAT that replaces each variable with a “positive” and a “negative” variable and adds extra clauses appropriately [\[37\]](#) is in fact a PPT when the parameter is the number of variables.

Given an instance ϕ of MONOTONE SAT, where the formula ϕ contains n variables and m clauses, we construct in polynomial time an instance (G, k) of MAXIMUM MINIMAL VERTEX COVER as follows. For each variable x_i of ϕ , $i \in [n]$, we add to G four vertices $\ell_i, x_i^+, x_i^-, r_i$ and three edges $\{\ell_i, x_i^+\}, \{x_i^+, x_i^-\}, \{x_i^-, r_i\}$, hence inducing a P_4 . We call the vertex x_i^+ (resp. x_i^-) a *positive* (resp. a *negative*) vertex of G . For each clause C_j of ϕ , $j \in [m]$, we add to G a vertex c_j , which we connect to the positive or negative vertices corresponding to the literals contained in C_j . This concludes the construction of G , which is illustrated in [Figure 2\(a\)](#). Note that, since ϕ is a monotone formula, G is a bipartite graph. Note also that the set of vertices $\{x_i^+, x_i^- \mid i \in [n]\}$ is a minimum vertex cover of G of size $2n$, and that the set of edges $\{\{\ell_i, x_i^+\}, \{x_i^-, r_i\} \mid i \in [n]\}$ is a maximum matching of G of size $2n$. We claim that ϕ is satisfiable if and only if G contains a minimal vertex cover of size $k := 2n + m$.



■ **Figure 2** (a) Illustration of the graph G built from the formula ϕ in the proof of [Theorem 42](#). (b) A minimal vertex cover X of G is shown with larger red vertices.

Suppose first that ϕ is satisfiable, and let σ be an assignment of the variables that satisfies

all the clauses in ϕ . We proceed to define a minimal vertex cover X of G of size k . First, add to X all the clause vertices $\{c_j \mid j \in [m]\}$. For every $i \in [n]$, if $\sigma(x_i) = \text{true}$ (resp. $\sigma(x_i) = \text{false}$), add to X vertices x_i^- and ℓ_i (resp. x_i^+ and r_i). See Figure 2(b) for an illustration, where the set X is shown with larger red vertices. Clearly, X is a vertex cover of G . To see that it is minimal, by Observation 1 it is enough to verify that, for every vertex $v \in X$, $N[v] \not\subseteq X$. This condition holds easily for all vertices in X that are in the P_4 's, since for each P_4 its vertices in X are not adjacent. Let c_j be a clause vertex. Since σ is a satisfying assignment of the variables, there exists a variable x_i such that if $\sigma(x_i) = \text{true}$ (resp. $\sigma(x_i) = \text{false}$) then $x_i \in C_j$ (resp. $\bar{x}_i \in C_j$). By definition of X , if $\sigma(x_i) = \text{true}$ (resp. $\sigma(x_i) = \text{false}$) then $x_i^+ \notin X$ (resp. $x_i^- \notin X$), and by construction of G we have that $x_i^+ \in N(c_j)$ (resp. $x_i^- \in N(c_j)$), so in both cases $N[c_j] \not\subseteq X$.

Conversely, suppose that G contains a minimal vertex cover X of size k , and we proceed to define a variable assignment σ as follows. For $i \in [n]$, as $\{x_i^+, x_i^-\} \in E(G)$ we have that X contains one or two vertices in the set $\{x_i^+, x_i^-\}$. If $x_i^+ \notin X$ (resp. $x_i^- \notin X$) we set $\sigma(x_i) = \text{true}$ (resp. $\sigma(x_i) = \text{false}$), and if both x_i^+ and x_i^- belong to X we set $\sigma(x_i)$ to **true** or to **false** arbitrarily. We claim that σ satisfies all the clauses in ϕ . For $i \in [n]$, let P^i be the P_4 of G induced by the vertices $\ell_i, x_i^+, x_i^-, r_i$. Since X is a vertex cover, clearly $|X \cap V(P^i)| \geq 2$. We claim that $|X \cap V(P^i)| = 2$. Indeed, if $|X \cap V(P^i)| \geq 3$, then $\{\ell_i, x_i^+\} \subseteq X$ or $\{x_i^-, r_i\} \subseteq X$ (or both). But then $N[\ell_i] \subseteq X$ or $N[r_i] \subseteq X$ (or both), contradicting Observation 1. Thus, $|X \cap V(P^i)| = 2$, which implies that $|X \cap \bigcup_{i \in [n]} V(P^i)| = 2n$, hence necessarily X contains the whole set $\{c_j \mid j \in [m]\}$ of clause vertices. Consider an arbitrary clause vertex c_j . Since X is minimal and $c_j \in X$, by Observation 1 there exists a neighbor of c_j in G that is *not* in X , and by definition of σ it follows that the literal corresponding to that neighbor of c_j satisfies clause C_j . Thus, σ is a satisfying assignment and the proof is complete.

Finally, note that the above reduction is also an NP-completeness reduction from MONOTONE SAT to MAXIMUM MINIMAL VERTEX COVER on bipartite graphs. ◀

9 Conclusions and further research

Motivated by the existence of subquadratic kernels for the MAXIMUM MINIMAL VERTEX COVER problem parameterized by the solution size, we introduced a general framework to rule out certain types of kernels, which we called **lop-kernels**, for optimization problems. This “lop” assumption does not seem to be very restrictive, as the vast majority of known kernels are in fact **lop-kernels** [35]. For instance, the classical kernels for VERTEX COVER, such as those using the high-degree rule, the crown decomposition rule, or the Nemhauser-Trotter rule [35], are **lop-kernels**. More involved kernels, such as those based on protrusion replacement [11], are also **lop-kernels**. However, we discussed in Section 5 an example of a polynomial kernel for a vertex-minimization problem, namely TREE DELETION SET [38], which is *not* a **lop-kernel**. We still do not know of a similar example that is a vertex-maximization problem.

For several technical reasons, we think that the framework of **lop-kernels** seems to be more suited for maximization problems. In this direction, we showed that a direct application of our general result for vertex-maximization problems (Corollary 11) yields kernelization lower bounds for MMVC (Corollary 26) and MMFVS (Corollary 27), matching the sizes of the best known kernels for these problems. We also presented consequences of our results for the MAXIMUM INDEPENDENT SET problem restricted to K_t -free graphs (Corollary 30) and conjectured (Conjecture 29) that, for every fixed graph H , the MAXIMUM INDEPENDENT SET problem restricted to H -free graphs admits a polynomial **lop-kernel**. For (vertex-)minimization problems, the only application that we were able to find concerns the TREE DELETION SET

problem (Corollary 28).

We believe that our results could be applied to other vertex-maximization problems, in particular to the “max-min” version of other vertex-minimization problems, as they seem to be quite hard to approximate. It would be interesting to find other examples of vertex-minimization problems, other than TREE DELETION SET, where our results could be applied. Here, the natural candidates seem to be the “min-max” version of vertex-maximization problems, which seem to have been almost unexplored so far.

Our general results for maximization (Theorem 13) and minimization (Theorem 24) problems take into account an upper bound function $u(n) = \mathcal{O}(n^a)$ that upper-bounds the size of optimal solutions of the considered problems. All our applications presented in Section 5 correspond to vertex problems, that is, to the case $a = 1$. We leave for further research to find applications of our results for problems with superlinear upper bound functions, such as edge problems, for which $a = 2$.

Our results are also able to derive lower bounds on the multiplicative constants of the existing kernels (cf. for instance the second item of Corollary 11 and Corollary 22). We still do not have any relevant application of this type for a concrete problem. For instance, if we apply Corollary 22 to the VERTEX COVER problem, relying on the non-existence of a $(2 - \varepsilon)$ -approximation under the Unique Games Conjecture [47], we rule out the existence of a $\text{lo}p$ -kernel with $(1 - \varepsilon)k$ vertices, which is not particularly interesting.

We presented (Section 7) subquadratic kernels for MAXIMUM MINIMAL VERTEX COVER on H -free graphs for some graphs H satisfying the (constructive) Erdős-Hajnal property, such as the bull, the complete graphs, or the paw. It would be interesting to obtain subquadratic kernels for other graphs H satisfying the Erdős-Hajnal property, such as C_4 , the diamond, P_5 , or C_5 . Note that, from [42], C_4 and the diamond satisfy the constructive Erdős-Hajnal property with constant $\delta \geq 1/3$. Note also that the graphs constructed in the reduction of Theorem 42 are $\{C_5, \text{diamond}\}$ -free, as they are bipartite, hence MMVC is NP-hard on this class, in contrast to the fact (Corollary 41) that MMVC can be solved in linear time on $\{P_5, \text{diamond}\}$ -free graphs. To the best of our knowledge, the complexity on P_5 -free graphs is open, as well as on $K_{1,t}$ graphs for $t \geq 3$ (see Lemma 38). It is worth mentioning that P_5 -free graphs have unbounded cliquewidth, because co-bipartite graphs, which are P_5 -free, have unbounded cliquewidth.

As defined in Section 2, for a graph G let $\text{mmvc}(G)$ be the maximum size of a minimal vertex cover of G . Boria et al. [18] proved that if G is an n -vertex graph without isolated vertices, then $\text{mmvc}(G) \geq \lfloor n^{1/2} \rfloor$. Note that this immediately yields a quadratic kernel for MMVC: if $k \leq \lfloor n^{1/2} \rfloor$ we answer “yes”, otherwise $n \leq k^2$. By the same argument, if $\mathcal{C}$ is a graph class such that every n -vertex graph $G \in \mathcal{C}$ without isolated vertices satisfies $\text{mmvc}(G) \geq n^{1/2+\varepsilon}$, for some $\varepsilon > 0$, then MMVC restricted to $\mathcal{C}$ admits a (subquadratic) kernel with at most $k^{\frac{2}{1+2\varepsilon}}$ vertices. It might be possible that this is the case for some of the H -free graph classes for which we provided subquadratic kernels in Section 7: we were not able to find any counterexample, that is, a family of n -vertex H -free graphs G for which $\text{mmvc}(G) = \Theta(n^{1/2})$. In particular, the case of triangle-free graphs seems particularly interesting. Haviland [43] and Goddard and Lyle [40] established upper bounds on the size of a minimum independent dominating set (that is, the complement of a minimal vertex cover) of triangle-free graphs. It follows from their results [40, 43] that there exist n -vertex triangle-free graphs G with $\text{mmvc}(G) = \Theta(n^{2/3} \cdot \log n)$, hence if such a constant $\varepsilon > 0$ as discussed above exists for triangle-free graphs, necessarily $\varepsilon \leq \frac{2}{3} - \frac{1}{2} = \frac{1}{6}$. Therefore, the smallest kernel that we may obtain in this way on triangle-free graphs would have $k^{\frac{2}{1+2\varepsilon}} \leq k^{3/2}$ vertices, which matches the size of the kernel that we obtained in Theorem 34 for the particular case

$t = 3$, disregarding lower-order terms and multiplicative constants. Finding such a constant $\varepsilon > 0$ on H -free graphs for small graphs H , in particular on triangle-free graphs, looks like a challenging problem, having interesting connections with the Ramsey numbers [40, 43].

Acknowledgement. We would like to thank Michael Lampis (resp. Magnus Wahlström, Venkatesh Raman) for pointing us to reference [31] (resp. reference [38], references [9, 49]).

References

- 1 Faisal N. Abu-Khzam, Cristina Bazgan, Morgan Chopin, and Henning Fernau. Approximation algorithms inspired by kernelization methods. In *Proc. of the 25th International Symposium on Algorithms and Computation (ISAAC)*, volume 8889 of *LNCS*, pages 479–490, 2014. doi:10.1007/978-3-319-13075-0_38.
- 2 Robert B. Allan and Renu C. Laskar. On domination and independent domination numbers of a graph. *Discrete Mathematics*, 23(2):73–76, 1978. doi:10.1016/0012-365X(78)90105-X.
- 3 Noga Alon, János Pach, and József Solymosi. Ramsey-type theorems with forbidden subgraphs. *Combinatorica*, 21(2):155–170, 2001. doi:10.1007/s004930100016.
- 4 Júlio Araújo, Marin Bougeret, Victor Campos, and Ignasi Sau. A New Framework for Kernelization Lower Bounds: The Case of Maximum Minimal Vertex Cover. In *Proc. of the 16th International Symposium on Parameterized and Exact Computation (IPEC)*, volume 214 of *LIPICs*, pages 4:1–4:19, 2021. doi:10.4230/LIPICs.IPEC.2021.4.
- 5 Júlio Araújo, Marin Bougeret, Victor A. Campos, and Ignasi Sau. Parameterized complexity of computing maximum minimal blocking and hitting sets. Manuscript submitted for publication, 2021. arXiv:2102.03404.
- 6 Stefan Arnborg, Jens Lagergren, and Detlef Seese. Easy problems for tree-decomposable graphs. *Journal of Algorithms*, 12(2):308–340, 1991. doi:10.1016/0196-6774(91)90006-K.
- 7 Cristina Bazgan, Ljiljana Brankovic, Katrin Casel, Henning Fernau, Klaus Jansen, Kim-Manuel Klein, Michael Lampis, Mathieu Liedloff, Jérôme Monnot, and Vangelis Th. Paschos. The many facets of upper domination. *Theoretical Computer Science*, 717:2–25, 2018. doi:10.1016/j.tcs.2017.05.042.
- 8 Daniel Binkele-Raible, Henning Fernau, Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Yngve Villanger. Kernel(s) for problems with no kernel: On out-trees with many leaves. *ACM Transactions on Algorithms*, 8(4):38:1–38:19, 2012. doi:10.1145/2344422.2344428.
- 9 Arindam Biswas, Venkatesh Raman, and Saket Saurabh. Approximation in (poly-) logarithmic space. In *Proc. of the 45th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 170 of *LIPICs*, pages 16:1–16:15, 2020. doi:10.4230/LIPICs.MFCS.2020.16.
- 10 Hans L. Bodlaender, Rodney G. Downey, Michael R. Fellows, and Danny Hermelin. On problems without polynomial kernels. *Journal of Computer and System Sciences*, 75(8):423–434, 2009. doi:10.1016/j.jcss.2009.04.001.
- 11 Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. (Meta) Kernelization. *Journal of the ACM*, 63(5):44:1–44:69, 2016.
- 12 Hans L. Bodlaender, Bart M. P. Jansen, and Stefan Kratsch. Kernelization lower bounds by cross-composition. *SIAM Journal on Discrete Mathematics*, 28(1):277–305, 2014. doi:10.1137/120880240.
- 13 Hans L. Bodlaender, Stéphan Thomassé, and Anders Yeo. Kernel bounds for disjoint cycles and disjoint paths. *Theoretical Computer Science*, 412(35):4570–4578, 2011. doi:10.1016/j.tcs.2011.04.039.
- 14 Rodica Boliac and Vadim V. Lozin. Independent domination in finitely defined classes of graphs. *Theoretical Computer Science*, 301(1-3):271–284, 2003. doi:10.1016/S0304-3975(02)00586-8.
- 15 Édouard Bonnet, Michael Lampis, and Vangelis Th. Paschos. Time-approximation trade-offs for inapproximable problems. *Journal of Computer and System Sciences*, 92:171–180, 2018. doi:10.1016/j.jcss.2017.09.009.

- 16 Édouard Bonnet and Vangelis Th. Paschos. Sparsification and subexponential approximation. *Acta Informatica*, 55(1):1–15, 2018. doi:10.1007/s00236-016-0281-2.
- 17 Édouard Bonnet, Stéphan Thomassé, Xuan Thang Tran, and Rémi Watrigant. An algorithmic weakening of the Erdős-Hajnal conjecture. In *Proc. of the 28th Annual European Symposium on Algorithms (ESA)*, volume 173 of *LIPIcs*, pages 23:1–23:18, 2020. doi:10.4230/LIPIcs.ESA.2020.23.
- 18 Nicolas Boria, Federico Della Croce, and Vangelis Th. Paschos. On the max min vertex cover problem. *Discrete Applied Mathematics*, 196:62–71, 2015. doi:10.1016/j.dam.2014.06.001.
- 19 Andreas Brandstädt. $(P_5, \text{diamond})$ -free graphs revisited: structure and linear time optimization. *Discrete Applied Mathematics*, 138(1-2):13–27, 2004. doi:10.1016/S0166-218X(03)00266-X.
- 20 Katrin Casel, Henning Fernau, Mehdi Khosravian Ghadikolaei, Jérôme Monnot, and Florian Sikora. Extension of Vertex Cover and Independent Set in Some Classes of Graphs. In *Proc. of the 11th International Conference on Algorithms and Complexity (CIAC)*, volume 11485 of *LNCS*, pages 124–136, 2019. doi:10.1007/978-3-030-17402-6_11.
- 21 Jianer Chen, Henning Fernau, Iyad A. Kanj, and Ge Xia. Parametric duality and kernelization: Lower bounds and upper bounds on kernel size. *SIAM Journal on Computing*, 37(4):1077–1106, 2007. doi:10.1137/050646354.
- 22 Maria Chudnovsky. The Erdős-Hajnal Conjecture - A Survey. *Journal of Graph Theory*, 75(2):178–190, 2014. doi:10.1002/jgt.21730.
- 23 Maria Chudnovsky and Shmuel Safra. The Erdős-Hajnal conjecture for bull-free graphs. *Journal of Combinatorial Theory, Series B*, 98(6):1301–1310, 2008. doi:10.1016/j.jctb.2008.02.005.
- 24 Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems*, 33(2):125–150, 2000. doi:10.1007/s002249910009.
- 25 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.
- 26 Peter Damaschke. Parameterized algorithms for double hypergraph dualization with rank limitation and maximum minimal vertex cover. *Discrete Optimization*, 8(1):18–24, 2011. doi:10.1016/j.disopt.2010.02.006.
- 27 Holger Dell and Dániel Marx. Kernelization of packing problems. In *Proc. of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 68–81, 2012. doi:10.1137/1.9781611973099.6.
- 28 Holger Dell and Dieter van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. *Journal of the ACM*, 61(4):23:1–23:27, 2014. doi:10.1145/2629620.
- 29 Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012. URL: <https://dblp.org/rec/books/daglib/0030488.bib>.
- 30 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.
- 31 Louis Dublois, Tesshu Hanaka, Mehdi Khosravian Ghadikolaei, Michael Lampis, and Nikolaos Melissinos. (In)approximability of Maximum Minimal FVS. In *Proc. of the 31st International Symposium on Algorithms and Computation (ISAAC)*, volume 181 of *LIPIcs*, pages 3:1–3:14, 2020. The cubic kernel appears in the full version, available at <https://arxiv.org/abs/2009.09971>. doi:10.4230/LIPIcs.ISAAC.2020.3.
- 32 Louis Dublois, Michael Lampis, and Vangelis Th. Paschos. Upper Dominating Set: Tight Algorithms for Pathwidth and Sub-Exponential Approximation. *CoRR*, abs/2101.07550, 2021. arXiv:2101.07550.
- 33 Paul Erdős and András Hajnal. Ramsey-type theorems. *Discrete Applied Mathematics*, 25(1-2):37–52, 1989. doi:10.1016/0166-218X(89)90045-0.

- 34 Henning Fernau. *Parameterized algorithms: a graph-theoretic approach*. Habilitationsschrift, Universität Tübingen, 2005. URL: <http://www.informatik.uni-trier.de/~fernau/papers/habil.pdf>.
- 35 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization. Theory of Parameterized Preprocessing*. Cambridge University Press, 2019. doi:10.1017/9781107415157.
- 36 Lance Fortnow and Rahul Santhanam. Infeasibility of instance compression and succinct PCPs for NP. *Journal of Computer and System Sciences*, 77(1):91–106, 2011. doi:10.1016/j.jcss.2010.06.007.
- 37 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979. doi:<https://dl.acm.org/doi/book/10.5555/574848>.
- 38 Archontia C. Giannopoulou, Daniel Lokshtanov, Saket Saurabh, and Ondrej Suchý. Tree Deletion Set has a polynomial kernel but no $\text{OPT}^{O(1)}$ approximation. *SIAM Journal on Discrete Mathematics*, 30(3):1371–1384, 2016. doi:10.1137/15M1038876.
- 39 Wayne Goddard and Michael A. Henning. Independent domination in graphs: A survey and recent results. *Discrete Mathematics*, 313(7):839–854, 2013. doi:10.1016/j.disc.2012.11.031.
- 40 Wayne Goddard and Jeremy Lyle. Independent dominating sets in triangle-free graphs. *Journal of Combinatorial Optimization*, 23(1):9–20, 2012. doi:10.1007/s10878-010-9336-4.
- 41 Jiong Guo, Iyad A. Kanj, and Stefan Kratsch. Safe approximation and its relation to kernelization. In *Proc. of the 6th International Symposium on Parameterized and Exact Computation (IPEC)*, volume 7112 of *LNCS*, pages 169–180, 2011. doi:10.1007/978-3-642-28050-4_14.
- 42 András Gyárfás. Reflections on a Problem of Erdős and Hajnal. In Ronald L. Graham, Jaroslav Nešetřil, and Steve Butler, editors, *The Mathematics of Paul Erdős II*, pages 135–141. Springer, 2013. doi:10.1007/978-1-4614-7254-4_11.
- 43 Julie Haviland. Independent domination in triangle-free graphs. *Discrete Mathematics*, 308(16):3545–3550, 2008. doi:10.1016/j.disc.2007.07.010.
- 44 Danny Hermelin and Xi Wu. Weak compositions and their applications to polynomial lower bounds for kernelization. In *Proc. of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 104–113, 2012. doi:10.1137/1.9781611973099.9.
- 45 Dorit S. Hochbaum and David B. Shmoys. Using dual approximation algorithms for scheduling problems theoretical and practical results. *Journal of the ACM*, 34(1):144–162, 1987. doi:10.1145/7531.7535.
- 46 Johann L. Hurink and Tim Nieberg. Approximating minimum independent dominating sets in wireless networks. *Information Processing Letters*, 109(2):155–160, 2008. doi:10.1016/j.ipl.2008.09.021.
- 47 Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. *Journal of Computer and System Sciences*, 74(3):335–349, 2008. doi:10.1016/j.jcss.2007.06.019.
- 48 Stefan Kratsch. Polynomial kernelizations for $\text{MIN } F^+II_1$ and MAX NP . *Algorithmica*, 63(1-2):532–550, 2012. doi:10.1007/s00453-011-9559-5.
- 49 Stefan Kratsch, Marcin Pilipczuk, Ashutosh Rai, and Venkatesh Raman. Kernel lower bounds using co-nondeterminism: Finding induced hereditary subgraphs. *ACM Transactions on Computation Theory*, 7(1):4:1–4:18, 2014. doi:10.1145/2691321.
- 50 Daniel Lokshtanov, Fahad Panolan, M. S. Ramanujan, and Saket Saurabh. Lossy kernelization. In *Proc. of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, pages 224–237, 2017. doi:10.1145/3055399.3055456.
- 51 Oystein Ore. *Theory of Graphs*. American Mathematical Society Colloquium Publications, volume 38, 1962. URL: <https://bookstore.ams.org/coll-38>.
- 52 Frank P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, s2-30(1):264–286, 1930. doi:10.1112/plms/s2-30.1.264.

- 53 Neil Robertson, Daniel P. Sanders, Paul D. Seymour, and Robin Thomas. The four-colour theorem. *Journal of Combinatorial Theory, Series B*, 70(1):2–44, 1997. doi:10.1006/jctb.1997.1750.
- 54 Jerzy Topp and Lutz Volkmann. On graphs with equal domination and independent domination numbers. *Discrete Mathematics*, 96(1):75–80, 1991. doi:10.1016/0012-365X(91)90473-F.
- 55 David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011. URL: http://www.cambridge.org/de/knowledge/isbn/item5759340/?site_locale=de_DE.
- 56 Mihalis Yannakakis. The effect of a connectivity requirement on the complexity of maximum subgraph problems. *Journal of the ACM*, 26(4):618–630, 1979. doi:10.1145/322154.322157.
- 57 Meirav Zehavi. Maximum Minimal Vertex Cover Parameterized by Vertex Cover. *SIAM Journal on Discrete Mathematics*, 31(4):2440–2456, 2017. doi:10.1137/16M109017X.