Preprocessing Vertex-Deletion Problems: Characterizing Graph Properties by Low-Rank Adjacencies

Bart M. P. Jansen ⬤

Eindhoven University of Technology, The Netherlands
b.m.p.jansen@tue.nl

Jari J. H. de Kroon ⬤

Eindhoven University of Technology, The Netherlands
j.j.h.d.kroon@tue.nl

──── **Abstract** ────

We consider the Π-free Deletion problem parameterized by the size of a vertex cover, for a range of graph properties Π. Given an input graph G, this problem asks whether there is a subset of at most k vertices whose removal ensures the resulting graph does not contain a graph from Π as induced subgraph. Many vertex-deletion problems such as Perfect Deletion, Wheel-free Deletion, and Interval Deletion fit into this framework. We introduce the concept of *characterizing a graph property* Π *by low-rank adjacencies*, and use it as the cornerstone of a general kernelization theorem for Π-Free Deletion parameterized by the size of a vertex cover. The resulting framework captures problems such as AT-Free Deletion, Wheel-free Deletion, and Interval Deletion. Moreover, our new framework shows that the vertex-deletion problem to perfect graphs has a polynomial kernel when parameterized by vertex cover, thereby resolving an open question by Fomin et al. [JCSS 2014]. Our main technical contribution shows how linear-algebraic dependence of suitably defined vectors over $\mathbb{F}_2$ implies graph-theoretic statements about the presence of forbidden induced subgraphs.

1 Introduction

Background. This paper continues a long line of investigation [2, 3, 13, 15, 19, 27], aimed at answering the following question: how and when can an efficient preprocessing algorithm reduce the size of inputs to NP-hard problems, without changing their answers? This question can be framed and answered using the notion of kernelization, which originated in parameterized complexity theory.

In parameterized complexity theory, the complexity analysis is done not only in the size of the input, but also in terms of another complexity measure related to the input. This complexity measure is called the *parameter*. For graph problems, typical parameters are

the size of a solution, the treewidth of the graph, or the size of a minimum vertex cover (the *vertex cover number*). The latter two are often called structural parameterizations. A kernelization is a polynomial-time preprocessing algorithm with a performance guarantee. It reduces an instance (x, k) of a parameterized problem to an instance (x', k') that has an equivalent YES/NO answer, such that $|x'|$ and k' are bounded by $f(k)$ for some computable function f, called the size of the kernel. If f is a polynomial function, the parameterized problem is said to admit a polynomial kernel. Polynomial kernels are highly sought after, as they allow problem instances to be reduced to a relatively small size.

We investigate polynomial kernels for the class of *graph modification* problems, in an attempt to develop a widely applicable and generic kernelization framework. In graph modification problems, the goal is to make a small number of changes to an input graph to make it satisfy a certain property. Possible modifications are vertex deletions, edge deletions, and edge additions. In this work, we consider the problem of deleting a bounded-size set of vertices such that the resulting graph does not contain certain graphs as an induced subgraph.

The study of kernelization for graph modification problems parameterized by solution size has an interesting and rich history [1, 6, 7, 10, 14, 17, 20, 23]. However, some graph modification problems such as PERFECT VERTEX DELETION [16] and WHEEL-FREE VERTEX DELETION [25] are W[2]-hard parameterized by the solution size and therefore do not admit any kernels unless FPT = W[2]. Together with the intrinsic interest in obtaining generic kernelization theorems that apply to a large class of problems with a single parameter, this has triggered research into polynomial kernelization for graph problems under structural parameterizations [4, 13, 15, 19, 26] such as the vertex cover number. The latter parameter is often used for its mathematical elegance, and due to the fact that slightly less restrictive parameters such as the feedback vertex number already cause simple problems such as 3-COLORING not to admit polynomial kernels [18], under the standard assumption NP $\not\subseteq$ coN-P/poly. This work therefore focuses on the following class of NP-hard [24] parameterized problems, where Π is a fixed (possibly infinite) set of graphs:

Π-FREE DELETION   **Parameter:** $|X|$
Input: A graph G, a vertex cover X of G, and an integer k.
Question: Does there exist a set $S \subseteq V(G)$ of size at most k such that $G - S$ does not contain any graph from Π as induced subgraph?

The assumption that a vertex cover X is given in the input is for technical reasons. If the problem would be parameterized by an upper-bound on the vertex cover number of the graph, without giving such a vertex cover, then the kernelization algorithm would have to verify that this is indeed a correct upper bound; an NP-hard problem. Instead, in this setting we just want to allow the kernelization algorithm to exploit the structural restriction guaranteed by having a small vertex cover in the graph. We refer to the discussion by Fellows et al. [12, §2.2] for more background. To apply the kernelization algorithms for problems defined in this way, one may simply use a 2-approximate vertex cover as X.

Fomin et al. [13] have investigated characteristics of Π-FREE DELETION problems that admit a polynomial kernel parameterized by the size of a vertex cover. They introduced a generic framework that poses three conditions on the graph property Π, which are sufficient to reach a polynomial kernel for Π-FREE DELETION parameterized by vertex cover. Examples of graph properties that fit in their framework are for instance "having a chordless cycle of length at least 4" or "having an odd cycle". This results in polynomial kernels for CHORDAL DELETION and ODD CYCLE TRANSVERSAL respectively. INTERVAL DELETION does not fit

◼ **Table 1** Kernels obtained by our framework for problems parameterized by a vertex cover X.

Problem	Vertices in kernel
PERFECT DELETION	$\mathcal{O}(|X|^5)$
EVEN-HOLE-FREE DELETION (★)	$\mathcal{O}(|X|^4)$
AT-FREE DELETION	$\mathcal{O}(|X|^9)$
INTERVAL DELETION	$\mathcal{O}(|X|^9)$
WHEEL-FREE DELETION	$\mathcal{O}(|X|^5)$

in this framework, even though interval graphs are hereditary. Agrawal et al. [1] show that it admits a polynomial kernel parameterized by solution size, and therefore also by vertex cover size. They introduced a linear-algebraic technique, which assigns a vector over $\mathbb{F}_2$ to each vertex, to find an induced subgraph that preserves the size of an optimal solution by combining several disjoint bases of systems of such vectors. This formed the inspiration for our work, in which we improve the generic kernelization framework of Fomin et al. [13] using linear-algebraic techniques inspired by the kernel [1] for INTERVAL DELETION.

Results. We introduce the notion of *characterizing a graph property* Π *by low-rank adjacencies*, and use it to generalize the kernelization framework by Fomin et al. [13] significantly. The resulting kernelization algorithms consist of a single, conceptually simple reduction rule for Π-FREE DELETION, whose property-specific correctness proofs show how the linear dependence of suitably defined vectors implies certain graph-theoretic properties. This results in a simpler kernelization for INTERVAL DELETION parameterized by vertex cover compared to the one by Agrawal et al. [1]. More importantly, several vertex-deletion problems whose kernelization complexity was previously open can be covered by the framework. These include AT-FREE DELETION (eliminate all asteroidal triples [22] from the graph), WHEEL-FREE DELETION, and also PERFECT DELETION which was an explicit open question of Fomin et al. [13, §5]. An overview is given in Table 1. Moreover, we give evidence that the distinguishing property of our framework (being able to characterize Π by low-rank adjacencies) is the right one to capture kernelization complexity. While the WHEEL-FREE DELETION problem fits into our framework and therefore has a polynomial kernel, the situation is very different for the related problem ALMOST WHEEL-FREE DELETION (ensure the resulting graph does not contain any wheel, except possibly W_4). We prove the latter problem does not fit into our framework, and that it does not admit a polynomial kernel parameterized by vertex cover, unless NP $\subseteq$ coNP/poly.

Related work. Even though the vertex cover is generally not small compared to the size of the input graph, it is not always the case that a polynomial kernel parameterized by vertex cover number exists. This was shown by Bodlaender et al. [3]. They showed that for instance the CLIQUE problem that asks whether a graph contains a clique of k vertices, does not admit a polynomial kernel parameterized by the vertex cover size, unless coNP $\subseteq$ NP/poly.

A graph is perfect if for every induced subgraph H, the chromatic number of H is equal to the size of the largest clique of H. Conjectured by Berge in 1961 and proven in the beginning of this century by Chudnovsky et al. [8], the strong perfect graph theorem states that a graph is perfect if and only if it is Berge. The forbidden induced subgraphs of Berge graphs (and hence of perfect graphs) are C_{2k+1} and $\overline{C}_{2k+1}$ for $k \geq 2$, that is, induced cycles and their edge complements of odd length at least 5. A survey of forbidden subgraph characterizations of some other hereditary graph classes is given in [5, Chapter 7].

Organization.   In Section 2 we give preliminaries and definitions used throughout this work. In Section 3 we introduce the framework. In Section 4 we show that several problems such as PERFECT DELETION and INTERVAL DELETION fit in this framework. Finally we conclude in Section 5. For statements marked ★, the proof is deferred to the full version [21].

2   Preliminaries

Notation.   For $i \in \mathbb{N}$, we denote the set $\{1, ..., i\}$ by $[i]$. For a set S, we denote the set of subsets of size at most k by $\binom{S}{\leq k} = \{S' \subseteq S \mid |S'| \leq k\}$. Similarly, $\binom{S}{k}$ denotes the set of subsets of size exactly k. We consider simple graphs that are unweighted and undirected without self-loops. A graph G has vertex and edge sets $V(G)$ and $E(G)$ respectively. An edge between vertices $u, v \in V(G)$ is an unordered pair $\{u, v\}$. For a set of vertices $S \subseteq V(G)$, by $G[S]$ we denote the graph induced by S. For $v \in V(G)$ and $S \subseteq V(G)$, by $G - v$ and $G - S$ we mean the graphs $G[V(G) \setminus \{v\}]$ and $G[V(G) \setminus S]$ respectively. We denote the open neighborhood of $v \in V(G)$ by $N_G(v) = \{u \mid \{u, v\} \in E(G)\}$. When clear from context, we sometimes omit the subscript G. For a graph G, let $\overline{G}$ be the edge complement graph of G on the same vertex set, such that for distinct $u, v \in V(G)$ we have $\{u, v\} \in E(\overline{G})$ if and only if $\{u, v\} \notin E(G)$. The path graph on n vertices $(v_1, ..., v_n)$ is denoted by P_n. Similarly, the n-vertex cycle for $n \geq 3$ is denoted by C_n. When $n \geq 4$, the graph C_n is often called a *hole*. For $n \geq 3$, the wheel W_n of size n is the graph on vertices $\{c, v_1, ..., v_n\}$ such that $(v_1, ..., v_n)$ is a cycle and c is adjacent to v_i for all $i \in [n]$. An asteroidal triple (AT) in a graph G consists of three vertices such that every pair is connected by a path that avoids the neighborhood of the third. A vertex cover in a graph G is a set of vertices that contains at least one endpoint of every edge. The minimum size of a vertex cover in a graph G is denoted by $\mathrm{vc}(G)$.

Parameterized complexity.   A parameterized problem [9, 11] is a language $Q \subseteq \Sigma^* \times \mathbb{N}$, where Σ is a finite alphabet. The notion of kernelization is formalized as follows.

▶ **Definition 1.** *Let $Q \subseteq \Sigma^* \times \mathbb{N}$ be a parameterized problem and let $f \colon \mathbb{N} \to \mathbb{N}$ be a computable function. A* kernelization *for Q of size f is an algorithm that, given an instance $(x, k) \in \Sigma^* \times \mathbb{N}$, outputs in time polynomial in $|x| + k$ an instance (x', k') (known as the kernel) such that $(x, k) \in Q$ if and only if $(x', k') \in Q$ and such that $|x'|, k' \leq f(k)$. If f is a polynomial function, then the algorithm is a* polynomial kernelization.

Previous kernelization framework.   We state some of the results from the kernelization framework by Fomin et al. [13] that forms the basis of this work. A graph property Π is a (possibly infinite) set of graphs.

▶ **Definition 2** (Definition 3, [13]). *A graph property Π is characterized by $c_\Pi \in \mathbb{N}$ adjacencies if for all graphs $G \in \Pi$, for every vertex $v \in V(G)$, there is a set $D \subseteq V(G) \setminus \{v\}$ of size at most c_Π such that all graphs G' which are obtained from G by adding or removing edges between v and vertices in $V(G) \setminus D$, are also contained in Π.*

As an example, the graph property "having a chordless cycle of length at least 4" is characterized by 3 adjacencies. The graph property "not being an interval graph" is not characterized by a finite number of adjacencies. Other examples are given by Fomin et al. [13].

Any finite graph property Π is trivially characterized by $\max_{G \in \Pi} |V(G)| - 1$ adjacencies. We state the following easily verified fact without proof.

▶ **Proposition 3.** *Let* Π' *be the set of all graphs that contain a graph from a finite set* Π *as induced subgraph. Then* Π' *is characterized by* $\max_{G \in \Pi} |V(G)| - 1$ *adjacencies.*

A graph G is vertex-minimal with respect to Π if $G \in \Pi$ and for all $S \subsetneq V(G)$ the graph $G[S]$ is not contained in Π. The following framework can be used to get polynomial kernels for the Π-FREE DELETION problem parameterized by vertex cover.

▶ **Theorem 4** (Theorem 2, [13]). *If* Π *is a graph property such that:*
   (i) Π *is characterized by* c_Π *adjacencies,*
  (ii) *every graph in* Π *contains at least one edge, and*
 (iii) *there is a non-decreasing polynomial* $p \colon \mathbb{N} \to \mathbb{N}$ *such that all graphs* G *that are vertex-minimal with respect to* Π *satisfy* $|V(G)| \leq p(\mathrm{VC}(G))$,
then Π-FREE DELETION *parameterized by the vertex cover size* x *admits a polynomial kernel with* $\mathcal{O}((x + p(x))x^{c_\Pi})$ *vertices.*

3 Framework based on low-rank adjacencies

3.1 Incidence vectors and characterizations

As a first step towards our kernelization framework for Π-FREE DELETION, we introduce an incidence vector definition (inc) that characterizes the neighborhood of a given vertex. Compared to the vector encoding used by Agrawal et al. [1] for INTERVAL DELETION, our vector definition differs because it supports arbitrarily large subsets (they consider subsets of size at most two), and because an entry of a vector simultaneously prescribes which neighbors should be present, and which neighbors should *not* be present.

▶ **Definition 5** (c-incidence vector). *Let* G *be a graph with vertex cover* X *and let* $c \in \mathbb{N}$. *Let* $Q', R' \subseteq X$ *such that* $|Q'| + |R'| \leq c$. *We define the* c-incidence vector $\mathrm{inc}^{c,(Q',R')}_{(G,X)}(u)$ *for a vertex* $u \in V(G) \setminus X$ *as a vector over* $\mathbb{F}_2$ *that has an entry for each* $(Q, R) \in X \times X$ *with* $Q \cap R = \emptyset$ *such that* $|Q| + |R| \leq c$, $Q' \subseteq Q$ *and* $R' \subseteq R$. *It is defined as follows:*

$$\mathrm{inc}^{c,(Q',R')}_{(G,X)}(u)[Q, R] = \begin{cases} 1 & \text{if } N_G(u) \cap Q = \emptyset \text{ and } R \subseteq N_G(u), \\ 0 & \text{otherwise.} \end{cases}$$

We drop superscript (Q', R') if both Q' and R' are empty sets. The intuition behind the superscript (Q', R') is that it projects the entries of the full incidence vector $\mathrm{inc}^c_{(G,X)}$ to those for supersets of Q', R'. The c-incidence vectors can be naturally summed coordinate-wise. For ease of presentation we do not define an explicit order on the coordinates of the vector, as any arbitrary but fixed ordering suffices.

If the sum of some vectors equals some other vector with respect to a certain graph G, then this equality is preserved when decreasing c or taking induced subgraphs of G.

▶ **Proposition 6.** *Let* G *be a graph with vertex cover* X, *let* $c \in \mathbb{N}$, *and let* $D \subseteq V(G)$ *be disjoint from* X. *If* $v \in V(G) \setminus (D \cup X)$ *and* $\mathrm{inc}^c_{(G,X)}(v) = \sum_{u \in D} \mathrm{inc}^c_{(G,X)}(u)$, *then*
   ▬ $\mathrm{inc}^{c'}_{(G,X)}(v) = \sum_{u \in D} \mathrm{inc}^{c'}_{(G,X)}(u)$ *for any* $c' \leq c$, *and*
   ▬ $\mathrm{inc}^c_{(H,X \cap V(H))}(v) = \sum_{u \in D} \mathrm{inc}^c_{(H,X \cap V(H))}(u)$ *for any induced subgraph* H *of* G *that contains* D *and* v.

Proof. For the first point, observe that for any vertex $v \notin (D \cup X)$, the vector $\mathrm{inc}^{c'}_{(G,X)}(v)$ is simply a projection of $\mathrm{inc}^c_{(G,X)}(v)$ to a subset of its coordinates. Hence if the complete vector of v is equal to the sum of the complete vectors of $u \in D$, then projecting the vector of both v and of the sum to the same set of coordinates, yields identical vectors.

For the second point, observe that since X is a vertex cover of G, we have $N_G(v) \subseteq X$ for all $v \in V(G) \setminus X$. Moreover, if H is an induced subgraph of G containing D and v, then $X_H := X \cap V(H)$ is a vertex cover of H. Hence for any $u \in V(H) \setminus X_H$ the c-incidence vector $\mathrm{inc}^c_{(H, X \cap V(H))}(u)$ is well-defined. If Q, R are disjoint sets for which $\mathrm{inc}^c_{(H, X \cap V(H))}(u)[Q, R]$ is defined, then $Q, R \subseteq X_H$, so the adjacencies between u and $Q \cup R$ in the induced subgraph H are identical to those in G, which implies $\mathrm{inc}^c_{(G, X)}(u)[Q, R] = \mathrm{inc}^c_{(H, X \cap V(H))}(u)[Q, R]$. Hence when we replace a c-incidence vector with subscript (G, X) by a vector with subscript $(H, X \cap V(H))$, we essentially project the vector to a subset of its coordinates without changing any values. For the same reason as above, this preserves the fact that the vectors of D sum to that of v. ◀

We are ready to introduce the main definition, namely characterization of a graph property Π by rank-c adjacencies for some $c \in \mathbb{N}$. In our framework, this replaces characterization by c adjacencies in the framework of Fomin et al. [13] (Theorem 4).

▶ **Definition 7** (rank-c adjacencies). *Let $c \in \mathbb{N}$ be a natural number. Graph property Π is characterized by rank-c adjacencies if the following holds. For each graph H, for each vertex cover X of H, for each set $D \subseteq V(H) \setminus X$, for each $v \in V(H) \setminus (D \cup X)$, if*

- $H - D \in \Pi$, *and*
- $\mathrm{inc}^c_{(H, X)}(v) = \sum_{u \in D} \mathrm{inc}^c_{(H, X)}(u)$ *when evaluated over $\mathbb{F}_2$,*

then there exists $D' \subseteq D$ such that $H - v - (D \setminus D') \in \Pi$. If there always exists such set D' of size 1, then we say Π is characterized by rank-c adjacencies with singleton replacements.

Intuitively, the definition demands that if we have a set D such that $H - D \in \Pi$, and the c-incidence vectors of D sum to the vector of some vertex v over $\mathbb{F}_2$, then there exists $D' \subseteq D$ such that removing v from $H - D$ and adding back D' results in a graph that is still contained in Π. For example, in Section 4.1 we show that the graph property "containing an odd hole or odd-anti-hole" is characterized by rank-4 adjacencies. Using our framework, this leads to a polynomial kernel for PERFECT DELETION parameterized by vertex cover. Other examples of graph properties which are characterized by a rank-c adjacencies for some $c \in \mathcal{O}(1)$ include "containing a cycle" and "being wheel-free". On the other hand, we will show in Theorem 25 that the property "containing an induced wheel whose size is 3 or at least 5" cannot be characterized by rank-c adjacencies for any finite c.

3.2 A generic kernelization

Our kernelization framework for Π-FREE DELETION relies on a single reduction rule presented in Algorithm 1. It assigns an incidence vector to every vertex outside the vertex cover and uses linear algebra to select vertices to store in the kernel. Let us therefore recall the relevant algebraic background. A *basis* of a set S of d-dimensional vectors over a field $\mathbb{F}$ is a minimum-size subset $B \subseteq S$ such that all $\mathbf{v} \in S$ can be expressed as linear combinations of elements of B, i.e., $\mathbf{v} = \sum_{\mathbf{u} \in B} \alpha_{\mathbf{u}} \cdot \mathbf{u}$ for a suitable choice of coefficients $\alpha_{\mathbf{u}} \in \mathbb{F}$. When working over the field $\mathbb{F}_2$, the only possible coefficients are 0 and 1, which gives a basis B of S the stronger property that any vector $\mathbf{v} \in S$ can be written as $\sum_{\mathbf{u} \in B'} \mathbf{u}$, where $B' \subseteq B$ consists of those vectors which get a coefficient of 1 in the linear combination.

Our reduction algorithm repeatedly computes a basis of the incidence vectors of the remaining set of vertices, and stores the vertices corresponding to the basis in the kernel.

■ **Algorithm 1** Reduce (Graph G, vertex cover X of G, $\ell \in \mathbb{N}$, $c \in \mathbb{N}$).

1: Let $Y_1 := V(G) \setminus X$.
2: **for** $i \leftarrow 1$ to ℓ **do**
3:      Let $V_i = \{\mathrm{inc}^c_{(G,X)}(y) \mid y \in Y_i\}$ and compute a basis B_i of V_i over $\mathbb{F}_2$.
4:      For each $\mathbf{v} \in B_i$, choose a unique vertex $y_{\mathbf{v}} \in Y_i$ such that $\mathbf{v} = \mathrm{inc}^c_{(G,X)}(y_{\mathbf{v}})$.
5:      Let $A_i := \{y_{\mathbf{v}} \mid \mathbf{v} \in B_i\}$ and $Y_{i+1} = Y_i \setminus A_i$.
6: **end for**
7: **return** $G[X \cup \bigcup_{i=1}^{\ell} A_i]$

▶ **Proposition 8.** *For a fixed $c \in \mathbb{N}$, Algorithm 1 runs in polynomial time in terms of ℓ and the size of the graph, and returns a graph on $\mathcal{O}(|X| + \ell \cdot |X|^c)$ vertices.*

Proof. Observe that for each i, the vectors in V_i have at most $2^c \cdot |\binom{X}{\leq c}| = \mathcal{O}(|X|^c)$ entries and therefore the rank of the vector space is $\mathcal{O}(|X|^c)$. Hence each computed basis contains $\mathcal{O}(|X|^c)$ vectors. For constant c, this means that each basis can be computed in polynomial time using Gaussian elimination. The remaining operations can be done in polynomial time in terms of ℓ and the size of the graph. Since $|A_i| \in \mathcal{O}(|X|^c)$ for each $i \in [\ell]$, the resulting graph has $\mathcal{O}(|X| + \ell \cdot |X|^c)$ vertices. ◀

▶ **Theorem 9.** *If Π is a graph property such that:*
(i) *Π is characterized by rank-c adjacencies,*
(ii) *every graph in Π contains at least one edge, and*
(iii) *there is a non-decreasing polynomial $p : \mathbb{N} \to \mathbb{N}$ such that all graphs G that are vertex-minimal with respect to Π satisfy $|V(G)| \leq p(\mathrm{VC}(G))$,*
then Π-FREE DELETION parameterized by the the vertex cover size x admits a polynomial kernel on $\mathcal{O}((x + p(x)) \cdot x^c)$ vertices.

Proof. Consider an instance (G, X, k) of Π-FREE DELETION. Note that if $k \geq |X|$, then we can delete the entire vertex cover to get an edgeless graph, which is Π-free by (ii), and therefore we may output a constant size YES-instance as the kernel. If $k < |X|$, let G' be the graph obtained by the procedure REDUCE(G,X,$\ell := k + 1 + p(|X|)$,c). By Proposition 8 this can be done in polynomial time and the resulting graph contains $\mathcal{O}((|X| + p(|X|)) \cdot |X|^c)$ vertices. All that is left to show is that the instance (G', X, k) is equivalent to the original instance. Since G' is an induced subgraph of G, it follows that if (G, X, k) is a YES-instance, then so is (G', X, k). In the other direction, suppose that (G', X, k) is a YES-instance with solution S. We show that S also is a solution for the original instance.

For the sake of contradiction assume that this is not the case. Then the graph $G - S$ contains an induced subgraph that belongs to Π. Let P be a minimal set of vertices of $G - S$ for which $G[P] \in \Pi$ and that minimizes $|P \setminus V(G')|$. Since S is a solution for (G', X, k), it follows that there exists a vertex $v \in P \setminus V(G')$. Moreover we have that $v \notin X$, since the graph G' returned by Algorithm 1 contains all vertices of X. The set $P \cap X$ is a vertex cover for $G[P]$, therefore by property (iii) we have that $|P| \leq p(\mathrm{VC}(G[P])) \leq p(|X|)$. Since the vertex sets $A_1, \dots, A_\ell$ computed in the REDUCE operation are disjoint, and since $|S| \leq k$, it follows that there exists an $i \in [k + 1 + p(|X|)]$ such that the set of vertices A_i corresponding to basis B_i is disjoint from both S and P.

As $v \notin V(G')$ implies $v \notin \bigcup_{i=1}^{\ell} A_i$, in each iteration of line 3 the vectors of the computed vertex set A_i span the vector of v. Hence, since we work over $\mathbb{F}_2$, there exists $D \subseteq A_i \subseteq V(G')$ such that $\mathrm{inc}^c_{(G,X)}(v) = \sum_{u \in D} \mathrm{inc}^c_{(G,X)}(u)$. Consider the graph $H := G[P \cup D]$. Since H is an induced subgraph that includes D and D is disjoint from X, by Proposition 6 it follows that $\mathrm{inc}^c_{(H, X \cap V(H))}(v) = \sum_{u \in D} \mathrm{inc}^c_{(H, X \cap V(H))}(u)$. Moreover $H - D \in \Pi$ as $H - D = G[P]$.

By the definition of rank-c adjacencies it follows that there exists $D' \subseteq D$ such that $P' = H - v - (D \setminus D') \in \Pi$. But since $|P' \setminus V(G')| < |P \setminus V(G')|$, this contradicts the minimality of P. Therefore S must be a solution for the original instance. ◄

3.3 Properties of low-rank adjacencies

In this section we present several technical lemmata dealing with low-rank adjacencies. These will be useful when applying the framework to various graph properties. The next lemma shows that if Π is characterized by low-rank adjacencies with singleton replacements, then the edge-complement graphs are as well.

▶ **Lemma 10.** *Let Π be a graph property that is characterized by rank-c adjacencies with singleton replacements. Let $\overline{\Pi}$ be the graph property such that $G \in \Pi$ if and only if $\overline{G} \in \overline{\Pi}$. Then $\overline{\Pi}$ is characterized by rank-c adjacencies with singleton replacements.*

Proof. Let H be a graph with vertex cover X. Let $D \subseteq V(H) \setminus X$ be a set such that $H - D \in \overline{\Pi}$. Consider some vertex $v \in V(H) \setminus (D \cup X)$ such that $\mathrm{inc}^c_{(H,X)}(v) = \sum_{u \in D} \mathrm{inc}^c_{(H,X)}(u)$. Let $X' = V(H) \setminus (D \cup \{v\})$.

▷ **Claim 11.** We have $\mathrm{inc}^c_{(H,X')}(v) = \sum_{u \in D} \mathrm{inc}^c_{(H,X')}(u)$.

Proof. Since vertices outside X are independent, neither v nor any vertex in D is adjacent to any vertex in $X' \setminus X$. So for any disjoint $Q, R \subseteq X'$ with $R \cap (X' \setminus X) \neq \emptyset$ we have $\mathrm{inc}^c_{(H,X')}(v)[Q, R] = \mathrm{inc}^c_{(H,X')}(u)[Q, R] = 0$ for all $u \in D$ by definition, while for $R \cap (X' \setminus X) = \emptyset$ we have $\mathrm{inc}^c_{(H,X')}(u)[Q, R] = \mathrm{inc}^c_{(H,X)}(u)[Q \cap X, R]$ for any $u \in D \cup \{v\}$. ◁

Let H' be obtained from H by (1) taking the edge complement, and then (2) turning $H'[D \cup \{v\}]$ back into an independent set (the complement made it a clique). Note that X' is a vertex cover of H'.

▷ **Claim 12.** We have $\mathrm{inc}^c_{(H',X')}(v) = \sum_{u \in D} \mathrm{inc}^c_{(H',X')}(u)$.

Proof. Immediate from Claim 11 since $\mathrm{inc}^c_{(H',X')}(u)[Q, R] = \mathrm{inc}^c_{(H,X')}(u)[R, Q]$ for all $u \in D \cup \{v\}$. ◁

Observe that $H' - D$ is the edge-complement of $H - D$, so $H' - D \in \Pi$. Together with the previous claim, since Π is characterized by rank-c adjacencies with singleton replacements, it follows that there exists $v' \in D$ such that $G' := H' - v - (D \setminus \{v'\}) \in \Pi$. Since G' contains only a single vertex of $\{v\} \cup D$, none of its edges were edited during step (2) above, so that $G := H - v - (D \setminus \{v'\})$ is the edge-complement of G', implying $G \in \overline{\Pi}$. This shows that $\overline{\Pi}$ is characterized by rank-c adjacencies with singleton replacements. ◄

Lemma 13 proves closure under taking the union of two characterized properties.

▶ **Lemma 13.** *Let Π and Π' be graph properties characterized by rank-c_Π and rank-$c_{\Pi'}$ adjacencies (with singleton replacements), respectively. Then the property $\Pi \cup \Pi'$ is characterized by rank-$\max(c_\Pi, c_{\Pi'})$ adjacencies (with singleton replacements).*

Proof. Consider a graph H with vertex cover X and set $D \subseteq V(H) \setminus X$ such that $H - D \in \Pi \cup \Pi'$. Let $v \in V(H) \setminus (D \cup X)$ be some vertex such that $\mathrm{inc}^{\max(c_\Pi, c_{\Pi'})}_{(H,X)}(v) = \sum_{u \in D} \mathrm{inc}^{\max(c_\Pi, c_{\Pi'})}_{(H,X)}(u)$. By Proposition 6, we have $\mathrm{inc}^{c_\Pi}_{(H,X)}(v) = \sum_{u \in D} \mathrm{inc}^{c_\Pi}_{(H,X)}(u)$. If $H - D \in \Pi$, then there exists $D' \subseteq D$ such that $H - v - (D \setminus D') \in \Pi$ and hence, $H - v - (D \setminus D') \in \Pi \cup \Pi'$ (in case of singleton replacements, D' is replaced by $\{v'\}$ for some $v' \in D$). The case $H - D \in \Pi'$ is symmetric. ◄

While the *intersection* of two graph properties which are characterized by a *finite number* of adjacencies is again characterized by a finite number of adjacencies [13, Proposition 4], the same does not hold for low-rank adjacencies; there is no analog of Lemma 13 for intersections.

In a graph G, we say that vertices u and v *share adjacencies* to a set S, if $N_G(u) \cap S = N_G(v) \cap S$. The following lemma states that when we have a set D whose c-incidence vectors sum to the vector of v, then for any set S of size up to c there exists a nonempty subset $D' \subseteq D$ whose members all share adjacencies with v to S.

▶ **Lemma 14.** *Let G be a graph with vertex cover X, let $D \subseteq V(G)$ be disjoint from X, and let $c \in \mathbb{N}$. Consider a vertex $v \in V(G) \setminus (D \cup X)$. If $\mathrm{inc}^c_{(G,X)}(v) = \sum_{u \in D} \mathrm{inc}^c_{(G,X)}(u)$, then for any set $S \subseteq V(G)$ with $|S| \leq c$ there exists $D' \subseteq D$, such that:*

- $|D'| \geq 1$ *is odd,*
- *each vertex $u \in D'$ shares adjacencies with v to S, and*
- $\mathrm{inc}^{c,(Q',R')}_{(G,X)}(v) = \sum_{u \in D'} \mathrm{inc}^{c,(Q',R')}_{(G,X)}(u)$, *where $Q' = (S \setminus N_G(v)) \cap X$ and $R' = S \cap N_G(v)$.*

Proof. For any vertex $d \in D$ that does not share adjacencies with v to S, the vector $\mathrm{inc}^{c,(Q',R')}_{(G,X)}(d)$ is the vector containing only zeros. Let $D' \subseteq D$ be the set of vertices that do share adjacencies with v to S. Clearly $\mathrm{inc}^{c,(Q',R')}_{(G,X)}(v) = \sum_{u \in D'} \mathrm{inc}^{c,(Q',R')}_{(G,X)}(u)$, as removing all-zero vectors does not change the sum. Since $\mathrm{inc}^{c,(Q',R')}_{(G,X)}(v)[Q',R'] = 1$ and $\mathrm{inc}^{c,(Q',R')}_{(G,X)}(u)[Q',R'] = 1$ for all $u \in D'$, $|D'| \geq 1$ must be odd. ◀

Our framework adapts Theorem 4 by replacing characterization by c adjacencies by rank-c adjacencies. From the following statement we can conclude that our framework extends Theorem 4.

▶ **Lemma 15.** *A graph property Π characterized by c adjacencies is also characterized by rank-c adjacencies with singleton replacements.*

Proof. Let Π be a graph property characterized by c adjacencies. We show that Π is characterized by rank-c adjacencies. Let G be a graph with vertex cover X and $D \subseteq V(G) \setminus X$ be a set such that $G - D \in \Pi$. Let $v \in V(G) \setminus (D \cup X)$ be a vertex such that $\mathrm{inc}^c_{(G,X)}(v) = \sum_{u \in D} \mathrm{inc}^c_{(G,X)}(u)$.

Since Π is characterized by c adjacencies, there exists a set B of size at most c such that all graphs obtained by changing adjacencies between v and $V(G) \setminus B$ are also contained in Π. By Lemma 14 there exists $w \in D$ that shares adjacencies with v to B. Now consider the graph $G - v - (D \setminus \{w\})$. This graph is isomorphic to $G - D$ where w is matched to v and the adjacencies between v and $V(G) \setminus B$ are changed. But then by the definition of characterization by c adjacencies it follows that $G - v - (D \setminus \{w\}) \in \Pi$. ◀

4 Using the framework

In this section we give some results using our framework, which are listed in Table 1. We give polynomial kernels for Perfect Deletion, AT-free Deletion, Interval Deletion, Even-hole-free Deletion, and Wheel-free Deletion parameterized by vertex cover.

4.1 Perfect Deletion

Let Π_P be the set of graphs that contain an odd hole or an odd anti-hole. The Π_P-free Deletion problem is known as the Perfect Deletion problem. It was mentioned as an open question by Fomin et al. [13], since one can show that Π_P is not characterized by a

finite number of adjacencies. In this section we show that Π_P is characterized by rank-4 adjacencies with singleton replacements. Following this result, we show that it admits a polynomial kernel using Theorem 9. First, we give a lemma that will be helpful in the proof later on. We say that a vertex *sees an edge* if it is adjacent to both of its endpoints.

▶ **Lemma 16.** *Let G be a graph, $P = (v_1, ..., v_n)$ where $n \geq 4$ is even be an induced path in G, and let y be a vertex not on P that is adjacent to both endpoints of P and sees an even number of edges of P. Then $G[V(P) \cup \{y\}]$ contains an odd hole as induced subgraph.*

Proof. We prove the claim by induction on n. Consider the case that $n = 4$. If y would be adjacent to exactly one of v_2 or v_3, then y would see a single edge $\{v_1, v_2\}$ or $\{v_3, v_4\}$ respectively. If y would be adjacent to both v_2 and v_3, then y would see all three edges of P. Since y sees an even number of edges of P, it follows that y is only adjacent to v_1 and v_4. Then $G[V(P) \cup \{y\}]$ induces an odd hole.

In the remaining case we assume that the claim holds for $n' < n$, where $n \geq 6$ and both n' and n are even. Suppose that y sees both edges $\{v_1, v_2\}$ and $\{v_{n-1}, v_n\}$, then $P' = (v_2, ..., v_{n-1})$ is an induced path on an even number of vertices such that y is adjacent to both of its endpoints and y sees an even number edges in P'. By the induction hypothesis $G[V(P') \cup \{y\}]$ contains an odd hole, therefore $G[V(P) \cup \{y\}]$ contains an odd hole as well. If y does not see both $\{v_1, v_2\}$ and $\{v_{n-1}, v_n\}$, then assume without loss of generality that y does not see the last edge $\{v_{n-1}, v_n\}$. Let v_j for $1 \leq j < n - 1$ be the largest index before n for which y is adjacent to v_j. If j is odd, then $G[\{v_j, ..., v_n, y\}]$ induces an odd hole. Otherwise $P' = (v_1, ..., v_j)$ is an induced path on an even number of vertices, y is adjacent to both of its endpoints, and y sees an even number of edges in P'; hence the induction hypothesis applies. In all cases we get that $G[V(P) \cup \{y\}]$ contains an odd hole. ◀

Before we show the proof that Π_P is characterized by rank-4 adjacencies with singleton replacements, we give some intuition for the replacement argument. Suppose we want to replace a vertex v of some odd hole, and we have a set D where each vertex in D is adjacent to both neighbors of v in the hole. Furthermore, the 4-incidence vectors of D sum to the vector of v. Then there must exist some vertex in D that sees an even number of edges of the induced path between the neighbors of v. This together with Lemma 16 would result in a graph that contains an odd hole. Figure 1 adds to this intuition.

Let Π_{OH} be the set of graphs that contain an odd hole. We show that Π_{OH} is characterized by rank-4 adjacencies with singleton replacements.

▶ **Theorem 17.** *Π_{OH} is characterized by rank-4 adjacencies with singleton replacements.*

Proof. Consider some graph H with vertex cover X and let $D \subseteq V(H) \setminus X$ such that $H - D \in \Pi_{OH}$. Let v be an arbitrary vertex in $V(H) \setminus (D \cup X)$ such that $\text{inc}_{(H,X)}^4(v) = \sum_{u \in D} \text{inc}_{(H,X)}^4(u)$. We show that $H - v - (D \setminus \{v'\}) \in \Pi_{OH}$ for some $v' \in D$.

Let C be an odd hole in $H - D$. If $v \notin V(C)$, then for every $v' \in D$ we have $H - v - (D \setminus \{v'\}) \in \Pi_{OH}$. So suppose that $v \in C$. Let $C = (v, p, v_1, ..., v_{n-3}, q)$, where $|V(C)| = n$. Consider the induced path $P = (p, v_1, ..., v_{n-3}, q)$. We have that v is adjacent to p and q. Let $D' \subseteq D$ be a set that shares adjacencies with v to $\{p, q\}$ such that $|D'| \geq 1$ is odd and $\text{inc}_{(H,X)}^{4,(\emptyset, \{p,q\})}(v) = \sum_{u \in D'} \text{inc}_{(H,X)}^{4,(\emptyset, \{p,q\})}(u)$. Such set exists by Lemma 14. Since C is an odd hole, $|V(P)|$ is even. Hence by Lemma 16, $G[V(P) \cup \{u\}]$ contains an odd hole if there exists some $u \in D'$ that sees an even number of edges of P. Suppose for the sake of contradiction that every vertex in D' sees an odd number of edges of P. Let E_u be the set of edges in P that are seen by $u \in D'$. Then $\sum_{u \in D'} |E_u|$ is odd as it is a sum of an odd number of odd numbers. Let $D'_{\{u,w\}} \subseteq D'$ be the set of vertices that see edge $\{u, w\} \in E(P)$. In order

	y_1	y_2	y_3	y_4	y_5	y_6	y_7	v
$[Q = \emptyset, R = \{p, q, v_1\}]$	1	1	0	1	0	0	1	0
$[Q = \emptyset, R = \{p, q, v_1, v_2\}]$	1	0	0	0	0	0	1	0
$[Q = \emptyset, R = \{p, q, v_2, v_3\}]$	1	0	0	0	0	1	0	0
$[Q = \emptyset, R = \{p, q, v_3, v_4\}]$	1	0	0	0	1	0	0	0
$[Q = \emptyset, R = \{p, q, v_4\}]$	1	1	1	0	1	0	0	0
$[Q = \{v_1\}, R = \{p, q\}]$	0	0	1	0	1	1	0	1
$[Q = \{v_1\}, R = \{p, q, v_2\}]$	0	0	1	0	0	1	0	0
$[Q = \{v_2, v_4\}, R = \{p, q\}]$	0	0	0	1	0	0	0	1
$\vdots$								

Figure 1 Graph G with vertex cover $X = \{p, q, v_1, \ldots, v_4\}$, containing an odd hole $H = \{v, p, v_1, ..., v_4, q\}$, such that $\mathrm{inc}^{4,(\emptyset,\{p,q\})}_{(G,X)}(v) = \sum_{i=1}^{7} \mathrm{inc}^{4,(\emptyset,\{p,q\})}_{(G,X)}(y_i)$. All edges $\{p, y_i\}$ and $\{q, y_i\}$ for $i \in [7]$ exist, but not all are drawn. The table shows entries of the vectors $\mathrm{inc}^{4,(\emptyset,\{p,q\})}_{(G,X)}(u \notin X)$. Vertex y_2 sees an even number of edges ($\{p, v_1\}$ and $\{v_4, q\}$), and $(v_1, ..., v_4, y_2)$ is an odd hole.

to satisfy $\sum_{u \in D'} \mathrm{inc}^{(\emptyset, \{p,q\})}(u)[\emptyset, \{p, q, u, w\}] = \mathrm{inc}^{(\emptyset, \{p,q\})}(v)[\emptyset, \{p, q, u, w\}] = 0$ over $\mathbb{F}_2$, for $\{u, w\} \in E(P)$, we require $|D'_{\{u,w\}}|$ to be even. But then $\sum_{e \in E(P)} |D'_e| = \sum_{u \in D'} |E_u|$ would also need to be an even number. This contradicts the fact that $\sum_{u \in D'} |E_u|$ is odd. Therefore there must exist some $u \in D'$ that sees an even number of edges in P. ◄

Let Π_{OAH} be the set of graphs that contain an odd anti-hole. Then $\Pi_P = \Pi_{OH} \cup \Pi_{OAH}$. From applications of Lemma 10 and Lemma 13 we get the following.

▶ **Corollary 18.** *Graph properties Π_{OH}, Π_{OAH}, and Π_P are characterized by rank-4 adjacencies with singleton replacements.*

▶ **Theorem 19.** PERFECT DELETION *parameterized by the size of a vertex cover admits a polynomial kernel on $\mathcal{O}(|X|^5)$ vertices.*

Proof. By Corollary 18 we have that Π_P is characterized by rank-4 adjacencies with singleton replacements. Each graph in Π_P contains at least one edge. For each odd hole or odd anti-hole H, we have $|V(H)| \leq 2 \cdot \mathrm{vc}(H)$. Therefore by Theorem 9 it follows that Π_P-FREE DELETION and hence PERFECT DELETION parameterized by vertex cover admits a polynomial kernel on $\mathcal{O}(|X|^5)$ vertices. ◄

A variation of Theorem 17 presented in the full version [21] shows that the set Π_{EH} of graphs containing an even hole are characterized by rank-3 adjacencies, which leads to a kernel for EVEN-HOLE-FREE DELETION parameterized by the size of a vertex cover of $\mathcal{O}(|X|^4)$ vertices.

4.2 AT-free Deletion

In his dissertation, Köhler [22] gives a forbidden subgraph characterization of graphs without asteroidal triples. This forbidden subgraph characterization consists of 15 small graphs on 6 or 7 vertices each, chordless cycles of length at least 6, and three infinite families often called asteroidal witnesses. Let Π_{AT} be the set of graphs that contain an asteroidal triple. A technical case analysis leads to the following results.

▶ **Theorem 20 (★).** Π_{AT} *is characterized by rank-8 adjacencies with singleton replacements.*

▶ **Theorem 21.** *AT-FREE DELETION parameterized by the size of a vertex cover admits a polynomial kernel on $\mathcal{O}(|X|^9)$ vertices.*

Proof. Every graph in Π_{AT} contains at least one edge. By Theorem 20 it follows that Π_{AT} is characterized by rank-8 adjacencies with singleton replacements. Each small graph has at most 7 vertices. For each cycle C, we have $|V(C)| \leq 2 \cdot \mathrm{vc}(C)$. Finally each asteroidal witness consists of an induced path with 2 or 3 additional vertices, hence there exists $c \in \mathbb{N}$ such that for $G \in \Pi_{AT}$, $|V(G)| \leq c \cdot \mathrm{vc}(G)$. Hence by Theorem 9, AT-FREE DELETION parameterized by the size of a vertex cover admits a polynomial kernel on $\mathcal{O}(|X|^9)$ vertices. ◀

4.3 Interval Deletion

INTERVAL DELETION does not fit in the framework of Fomin et al. [13], since one can show that its forbidden subgraph characterization is not characterized by a finite number of adjacencies. It was shown to admit a polynomial kernel by Agrawal et al. [1]. We show that our framework captures this result. Consider the graph property $\Pi_{IV} = \Pi_{AT} \cup \Pi_{C_{\geq 4}}$, where Π_{AT} is the set of graphs that contain an asteroidal triple as in Section 4.2 and $\Pi_{C_{\geq 4}}$ is the set of graphs that contain an induced cycle of length at least 4. Making a graph Π_{IV}-free makes it chordal and AT-free, therefore Π_{IV}-FREE DELETION corresponds to INTERVAL DELETION.

▶ **Theorem 22.** *INTERVAL DELETION parameterized by the size of a vertex cover admits a polynomial kernel on $\mathcal{O}(|X|^9)$ vertices.*

Proof. Every graph in Π_{IV} contains at least one edge. By Theorem 20, Π_{AT} is characterized by rank-8 adjacencies. Furthermore, $\Pi_{C_{\geq 4}}$ is characterized by 3 adjacencies as shown by Fomin et al. [13, Proposition 3], and therefore by Lemma 15 also by rank-3 adjacencies. Therefore by Lemma 13, it follows that Π_{IV} is also characterized by rank-8 adjacencies. Each vertex minimal graph in $\Pi_{C_{\geq 4}}$ is a cycle C, for which we have $|V(C)| \leq 2 \cdot \mathrm{vc}(C)$. Recall that $\Pi_{AT} = \Pi_S \cup \Pi_{C_{\geq 6}} \cup \Pi_{AW}$. Each vertex minimal graph in Π_S contains at most 7 vertices. Finally each asteroidal witness consists of an induced path with 2 or 3 additional vertices, hence there exists $c \in \mathbb{N}$ such that for $G \in \Pi_{IV}$, $|V(G)| \leq c \cdot \mathrm{vc}(G)$. Therefore by Theorem 9, INTERVAL DELETION parameterized by the size of a vertex cover admits a polynomial kernel on $\mathcal{O}(|X|^9)$ vertices. ◀

4.4 (Almost) Wheel-free Deletion

Let $\Pi_{W_{\geq 3}}$ be the set of graphs that contain a wheel of size at least 3 as induced subgraph. Then WHEEL-FREE DELETION corresponds to $\Pi_{W_{\geq 3}}$-FREE DELETION. We present a characterization by rank-4 adjacencies.

▶ **Theorem 23 (★).** $\Pi_{W_{\geq 3}}$ *is characterized by rank-4 adjacencies.*

Every graph that contains a wheel contains at least one edge. For every wheel W_n, we have $|V(W_n)| \leq 2 \cdot \mathrm{vc}(W_n)$. Therefore by Theorem 9 we obtain:

▶ **Theorem 24.** *WHEEL-FREE DELETION parameterized by the size of a vertex cover admits a polynomial kernel on $\mathcal{O}(|X|^5)$ vertices.*

It turns out that this good algorithmic behavior is very fragile. Let $\Pi_{W_{\neq 4}}$ be the set of graphs that contain a wheel of size 3, or at least 5. Then $\Pi_{W_{\neq 4}}$-FREE DELETION corresponds to ALMOST WHEEL-FREE DELETION. While $\Pi_{W_{\geq 3}}$ can be characterized by rank-4 adjacencies, the following shows that $\Pi_{W_{\neq 4}}$ is not characterized by adjacencies of any finite rank, and therefore does not fall within the scope of our kernelization framework.

▶ **Theorem 25 (★).** $\Pi_{W_{\neq 4}}$ *is not characterized by rank-c adjacencies for any* $c \in \mathbb{N}$.

This is not a deficiency of our framework; we prove that the problem does not have any polynomial compression, and therefore no polynomial kernel, unless NP $\subseteq$ coNP/poly.

▶ **Theorem 26 (★).** ALMOST WHEEL-FREE DELETION *parameterized by vertex cover does not admit a polynomial compression unless* coNP $\subseteq$ NP/poly.

This suggests that the condition of being characterized by low-rank adjacencies is the right way to capture kernelization complexity.

5 Conclusion

We have presented a framework that can be used to obtain polynomial kernels for the Π-FREE DELETION problem parameterized by the size of a vertex cover, based on the novel concept of characterizations by low-rank adjacencies. Our framework significantly extends the scope of the earlier framework of Fomin et al. [13]. In addition to the examples given in Table 1, the framework can be applied to obtain kernels for a wide range of vertex-deletion problems. Using the fact that graph properties characterized by low-rank adjacencies are closed under taking a union (Lemma 13), together with the characterizations by low-rank adjacencies developed here, and characterizations by few adjacencies by Fomin et al. [13, Table 1], we obtain the following.

▶ **Corollary 27.** *Let* $\mathcal{F}$ *be a hereditary graph class defined by an arbitrary combination of the following properties: being wheel-free, being odd-hole-free, being odd-anti-hole-free, being even-hole-free, being AT-free, being bipartite, being* $C_{\geq c}$*-free for some fixed* $c \in \mathbb{N}$*, being H-minor-free for some fixed graph H, being H-free for some fixed graph H containing at least one edge, and having a Hamiltonian cycle (respectively, path). Then the problem of testing whether an input graph G can be turned into a member of* $\mathcal{F}$ *by removing at most k vertices, has a polynomial kernel parameterized by vertex cover.*

It would be interesting to see whether the exponents given by Table 1 are tight (cf. [15]).

───── **References** ─────

1   Akanksha Agrawal, Pranabendu Misra, Saket Saurabh, and Meirav Zehavi. Interval vertex deletion admits a polynomial kernel. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pages 1711–1730, 2019. URL: `http://dl.acm.org/citation.cfm?id=3310435.3310538`.

2   Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. (Meta) Kernelization. *J. ACM*, 63(5):44:1–44:69, 2016. `doi:10.1145/2973749`.

3   H.L. Bodlaender, B.M.P. Jansen, and S. Kratsch. Kernelization lower bounds by cross-composition. *SIAM Journal on Discrete Mathematics*, 28(1):277–305, 2014. `doi:10.1137/120880240`.

4   Marin Bougeret and Ignasi Sau. How much does a treedepth modulator help to obtain polynomial kernels beyond sparse graphs? *Algorithmica*, 81(10):4043–4068, 2019. `doi:10.1007/s00453-018-0468-8`.

5   A. Brandstädt, V. Le, and J. Spinrad. *Graph Classes: A Survey*. Society for Industrial and Applied Mathematics, 1999. `doi:10.1137/1.9780898719796`.

6   Kevin Burrage, Vladimir Estivill-Castro, Michael R. Fellows, Michael A. Langston, Shev Mac, and Frances A. Rosamond. The undirected feedback vertex set problem has a poly(k)

kernel. In Hans L. Bodlaender and Michael A. Langston, editors, *Proceedings of the 2nd Second International Workshop on Parameterized and Exact Computation, IWPEC 2006*, volume 4169 of *Lecture Notes in Computer Science*, pages 192–202. Springer, 2006. `doi:10.1007/11847250_18`.

7 Yixin Cao, Ashutosh Rai, R. B. Sandeep, and Junjie Ye. A polynomial kernel for diamond-free editing. In Yossi Azar, Hannah Bast, and Grzegorz Herman, editors, *Proceedings of the 26th Annual European Symposium on Algorithms, ESA 2018*, volume 112 of *LIPIcs*, pages 10:1–10:13. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. `doi:10.4230/LIPIcs.ESA.2018.10`.

8 Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The strong perfect graph theorem. *Ann. Math. (2)*, 164(1):51–229, 2006. `doi:10.4007/annals.2006.164.51`.

9 Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.

10 Marek Cygan, Marcin Pilipczuk, Michal Pilipczuk, Erik Jan van Leeuwen, and Marcin Wrochna. Polynomial kernelization for removing induced claws and diamonds. *Theory Comput. Syst.*, 60(4):615–636, 2017. `doi:10.1007/s00224-016-9689-x`.

11 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`.

12 Michael R. Fellows, Bart M. P. Jansen, and Frances A. Rosamond. Towards fully multivariate algorithmics: Parameter ecology and the deconstruction of computational complexity. *Eur. J. Comb.*, 34(3):541–566, 2013. `doi:10.1016/j.ejc.2012.04.008`.

13 Fedor V. Fomin, Bart M.P. Jansen, and Michał Pilipczuk. Preprocessing subgraph and minor problems: When does a small vertex cover help? *Journal of Computer and System Sciences*, 80(2):468–495, 2014. `doi:10.1016/j.jcss.2013.09.004`.

14 Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Planar $\mathcal{F}$-deletion: Approximation, kernelization and optimal FPT algorithms. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012*, pages 470–479. IEEE Computer Society, 2012. `doi:10.1109/FOCS.2012.62`.

15 Archontia C. Giannopoulou, Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. Uniform kernelization complexity of hitting forbidden minors. *ACM Trans. Algorithms*, 13(3):35:1–35:35, 2017. `doi:10.1145/3029051`.

16 Pinar Heggernes, Pim Van 't Hof, Bart M. P. Jansen, Stefan Kratsch, and Yngve Villanger. Parameterized complexity of vertex deletion into perfect graph classes. *Theor. Comput. Sci.*, 511:172–180, 2013. `doi:10.1016/j.tcs.2012.03.013`.

17 Yoichi Iwata. Linear-time kernelization for feedback vertex set. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017*, volume 80 of *LIPIcs*, pages 68:1–68:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. `doi:10.4230/LIPIcs.ICALP.2017.68`.

18 Bart M. P. Jansen and Stefan Kratsch. Data reduction for graph coloring problems. *Inf. Comput.*, 231:70–88, 2013. `doi:10.1016/j.ic.2013.08.005`.

19 Bart M. P. Jansen and Astrid Pieterse. Optimal data reduction for graph coloring using low-degree polynomials. *Algorithmica*, 81(10):3865–3889, 2019. `doi:10.1007/s00453-019-00578-5`.

20 Bart M. P. Jansen and Marcin Pilipczuk. Approximation and kernelization for chordal vertex deletion. *SIAM J. Discrete Math.*, 32(3):2258–2301, 2018. `doi:10.1137/17M112035X`.

21 Bart M.P. Jansen and Jari J.H. de Kroon. Preprocessing vertex-deletion problems: Characterizing graph properties by low-rank adjacencies. *CoRR*, abs/2004.08818, 2020. `arXiv:2004.08818`.

22 Ekkehard Köhler. *Graphs without asteroidal triples*. PhD thesis, TU Berlin, 1999.

23 Stefan Kratsch and Magnus Wahlström. Compression via matroids: A randomized polynomial kernel for odd cycle transversal. *ACM Trans. Algorithms*, 10(4):20:1–20:15, 2014. `doi:10.1145/2635810`.

24   John M. Lewis and Mihalis Yannakakis. The node-deletion problem for hereditary properties
         is NP-complete. *Journal of Computer and System Sciences*, 20(2):219–230, 1980. `doi:`
         `10.1016/0022-0000(80)90060-4`.

25   Daniel Lokshtanov. Wheel-free deletion is W[2]-hard. In Martin Grohe and Rolf Niedermeier,
         editors, *Proceedings of the Third International Workshop on Parameterized and Exact Com-
         putation, IWPEC 2008*, volume 5018 of *Lecture Notes in Computer Science*, pages 141–147.
         Springer, 2008. `doi:10.1007/978-3-540-79723-4_14`.

26   Johannes Uhlmann and Mathias Weller. Two-layer planarization parameterized by feedback
         edge set. *Theor. Comput. Sci.*, 494:99–111, 2013. `doi:10.1016/j.tcs.2013.01.029`.

27   Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization:
         Theory of Parameterized Preprocessing*. Cambridge University Press, 2019. `doi:10.1017/`
         `9781107415157`.