Preprocessing vertex-deletion problems

Journal of Computer and System Sciences 126 (2022) 59–79

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Preprocessing vertex-deletion problems: Characterizing graph properties by low-rank adjacencies ☆,☆☆

Bart M.P. Jansen, Jari J.H. de Kroon *

Eindhoven University of Technology, the Netherlands

ARTICLE INFO

Article history:
Received 13 January 2021
Received in revised form 28 July 2021
Accepted 20 December 2021
Available online 3 January 2022

Keywords:
Kernelization
Vertex-deletion
Graph modification
Structural parameterization

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 an induced subgraph. 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. (2014) [18].

1. Introduction

Background

This paper continues a long line of investigation [2,3,18,21,26,35], 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

☆ *Related version:* An extended abstract of this work appeared under the same title in the Proceedings of the 17th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2020.

☆☆ *Funding:* This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 803421, ReduceSearch).

* Corresponding author.
E-mail addresses: b.m.p.jansen@tue.nl (B.M.P. Jansen), j.j.h.d.kroon@tue.nl (J.J.H. de Kroon).

https://doi.org/10.1016/j.jcss.2021.12.003

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 [1,6,7,11,19,24,28,30]. However, some graph modification problems such as Perfect Vertex Deletion [23] and Wheel-free Vertex Deletion [32] 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 [4,18,21,26,34] 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 [25] and P_3-subgraph-free deletion [13] not to admit polynomial kernels, under the standard assumption coNP $\not\subseteq$ NP/poly. This work therefore focuses on the following class of NP-hard [31] parameterized problems, where Π is a fixed (possibly infinite) set of graphs:

Π-free Deletion Parameterized by Vertex Cover **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 an 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. Throughout the article we use the term *parameterized by vertex cover* to refer to the variant of a graph modification problem where a vertex cover X is given in the input and its size is used as the parameter. We refer to the discussion by Fellows et al. [16, §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. [18] have investigated characteristics of Π-free Deletion problems that admit a polynomial kernel parameterized by 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 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 since any vertex cover is a solution. 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. [18] using the linear-algebraic techniques inspired by the kernel for Interval Deletion [1].

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. [18] 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 kernelization for Interval Deletion parameterized by vertex cover given by Agrawal et al. [1, p.1719] as part of their extensive work on the natural parameterization. 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 [29] from the graph), Wheel-free Deletion, and also Perfect Deletion which was an explicit open question of Fomin et al. [18, §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 coNP $\subseteq$ NP/poly.

Table 1
Kernels obtained by our framework for problems parameterized by vertex cover.

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	$\mathcal{O}(\|X\|^9)$
(Co-)Comparability Deletion	$\mathcal{O}(\|X\|^9)$
Permutation Deletion	$\mathcal{O}(\|X\|^9)$
Wheel-free Deletion	$\mathcal{O}(\|X\|^5)$

Related work

Even though the vertex cover number 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 vertex cover, unless coNP $\subseteq$ NP/poly. Similarly, it is known that (under the same assumption) Cutwidth [10], Short Secluded s-t Path [36], and 2-Club [22, Thm. 4] parameterized by vertex cover do not admit polynomial kernels.

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 contains no 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. Furthermore, we show that Almost Wheel-free Deletion does not fit in the framework. Finally we conclude in Section 5.

2. Preliminaries

Notation

For $i \in \mathbb{N}$, we denote the set $\{1, \ldots, i\}$ 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 which has vertex set S and edge set $\binom{S}{2} \cap E(G)$. 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. This operator is left associative, so $G - S - T$ should be interpreted as $(G - S) - T$. 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 $(v_1, \ldots, v_n)$ 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*. A hole is odd if n is odd. An *anti-hole* is the edge complement of a hole. For $n \geq 3$, the wheel W_n of size n is the graph on vertices $\{c, v_1, \ldots, v_n\}$ such that $(v_1, \ldots, v_n)$ 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 [9,14] is a language $Q \subseteq \Sigma^* \times \mathbb{N}$, where Σ is a finite alphabet. The notion of kernelization is formalized as follows.

Definition 2.1. Let $Q \subseteq \Sigma^* \times \mathbb{N}$ be a parameterized problem and let $f : \mathbb{N} \to \mathbb{N}$ 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. [18] that forms the basis of this work. A graph property Π is a (possibly infinite) set of graphs.

Definition 2.2 *(Definition 3, [18]).* 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. [18].

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 2.3. *Let Π' be the set of all graphs that contain a graph from a finite set Π as an 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 2.4 *(Theorem 2, [18]). 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 $|V(G)| \leq p(\text{vc}(G))$,*

then Π-free Deletion Parameterized by Vertex Cover admits a polynomial kernel with $\mathcal{O}((|X| + p(|X|)) \cdot |X|^{c_\Pi})$ 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 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. See Fig. 1 for an illustration of the following concept.

Definition 3.1 *(c-incidence vector).* Let G be a graph with a 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* $\text{inc}^{c,(Q',R')}_{(G,X)}(u)$ 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:

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

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 $\text{inc}^c_{(G,X)}$ to those for supersets of Q', R'. Informally speaking, the full vector $\text{inc}^c_{(G,X)}(u)$ encodes which subsets of up to c vertices appear in the neighborhood and non-neighborhood of vertex u. Note that since $u \notin X$ and X is a vertex cover we have $N_G(u) \subseteq X$, so that it suffices to consider subsets of X.

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.

Next we show that 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.

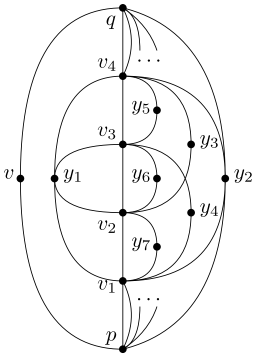

	y_1	y_2	y_3	y_4	y_5	y_6	y_7	v
$[Q=\emptyset, R=\{p,q,v_1\}]$	1	1	0	1	0	0	1	0
$[Q=\emptyset, R=\{p,q,v_1,v_2\}]$	1	0	0	0	0	0	1	0
$[Q=\emptyset, R=\{p,q,v_2,v_3\}]$	1	0	0	0	0	1	0	0
$[Q=\emptyset, R=\{p,q,v_3,v_4\}]$	1	0	0	0	1	0	0	0
$[Q=\emptyset, R=\{p,q,v_4\}]$	1	1	1	0	1	0	0	0
$[Q=\{v_1\}, R=\{p,q\}]$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

Fig. 1. Graph G with vertex cover $X=\{p,q,v_1,\ldots,v_4\}$, containing an odd hole $H=(v,p,v_1,\ldots,v_4,q)$, 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 $\mathrm{inc}^{4,(\emptyset,\{p,q\})}_{(G,X)}(u\notin X)$.

Proposition 3.2. *Let G be a graph with a 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 $c'\le c$, 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. [18] (Theorem 2.4).

Definition 3.3 *(rank-c adjacencies).* Let $c\in\mathbb{N}$ be a natural number. A 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 rank-c adjacencies for some $c\in\mathcal{O}(1)$ include ‘containing a cycle’ and ‘containing an induced wheel’. On the other hand, we will show in Theorem 4.26 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) \ X.
2: for i ← 1 to ℓ do
3:    Let V_i = {inc^c_(G,X)(y) | y ∈ Y_i} and compute a basis B_i of V_i over F_2.
4:    For each v ∈ B_i, choose a unique vertex y_v ∈ Y_i such that v = inc^c_(G,X)(y_v).
5:    Let A_i := {y_v | v ∈ B_i} and Y_{i+1} = Y_i \ A_i.
6: end for
7: return G[X ∪ ⋃_{i=1}^ℓ A_i]
```

Proposition 3.4. *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 3.5. *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 Vertex Cover *admits a polynomial kernel with $\mathcal{O}((|X| + p(|X|)) \cdot |X|^c)$ 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 3.4 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 $A_1, \ldots, A_\ell$ 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 3.2 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 3.6. *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 a 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 $\text{inc}^c_{(H,X)}(v) = \sum_{u \in D} \text{inc}^c_{(H,X)}(u)$. Let $X' = V(H) \setminus (D \cup \{v\})$.

Claim 3.7. We have $\text{inc}^c_{(H,X')}(v) = \sum_{u \in D} \text{inc}^c_{(H,X')}(u)$.

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 $\text{inc}^c_{(H,X')}(v)[Q,R] = \text{inc}^c_{(H,X')}(u)[Q,R] = 0$ for all $u \in D$ by definition, while for $R \cap (X' \setminus X) = \emptyset$ we have $\text{inc}^c_{(H,X')}(u)[Q,R] = \text{inc}^c_{(H,X)}(u)[Q \cap X, R]$ 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 3.8. We have $\text{inc}^c_{(H',X')}(v) = \sum_{u \in D} \text{inc}^c_{(H',X')}(u)$.

Proof. Immediate from Claim 3.7 since $\text{inc}^c_{(H',X')}(u)[Q,R] = \text{inc}^c_{(H,X')}(u)[R,Q]$ 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 3.9 proves closure under taking the union of two characterized properties.

Lemma 3.9. *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 a 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 $\text{inc}^{\max(c_\Pi,c_{\Pi'})}_{(H,X)}(v) = \sum_{u \in D} \text{inc}^{\max(c_\Pi,c_{\Pi'})}_{(H,X)}(u)$. By Proposition 3.2, we have $\text{inc}^{c_\Pi}_{(H,X)}(v) = \sum_{u \in D} \text{inc}^{c_\Pi}_{(H,X)}(u)$. 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 [18, Proposition 4], we do not believe the same holds for low-rank adjacencies: we do not have an analog of Lemma 3.9 for intersections. The difficulty in proving such an analog for intersections lies in the fact that Definition 3.3 does not give any control over the size of the set D' by which v is replaced to obtain a new graph in Π: when both Π and Π' are characterized by rank-c adjacencies and we consider a graph H with set $D \subseteq V(H)$ such that $H - D \in \Pi \cap \Pi'$, then when applying the definition for a vertex $v \notin D$ we obtain two sets D'_1, D'_2 such that $H - v - (D \setminus D'_1) \in \Pi$ and $H - v - (D \setminus D'_2) \in \Pi'$. However, the definition allows $D'_1 \neq D'_2$ so that we do not necessarily obtain a graph in $\Pi \cap \Pi'$. When strengthening Definition 3.3 with the additional requirement that $D' = D$ (so that v is replaced by the entire set D to obtain a new graph in Π), the intersection of two graph properties with a characterization by low-rank adjacencies has such a characterization itself. This strengthened definition would be sufficient for most of our positive results. However, as this strengthening decreases the flexibility of the framework, while the intersection variant of Lemma 3.9 is not useful for our applications, we have chosen to work with the current definition.

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 3.10. *Let G be a graph with a 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. Since X is a vertex cover we have $N_G(u) \subseteq X$ for all $u \in D \cup \{v\}$. Therefore a vertex $u \in D$ shares adjacencies with v to S if and only if u shares adjacencies with v to $S \cap X$, so that we may assume that $S \subseteq X$.

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 2.4 by replacing characterization by c adjacencies by rank-c adjacencies. From the following statement we can conclude that our framework extends Theorem 2.4.

Lemma 3.11. *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 a 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 $B \subseteq V(G) \setminus D$ of size at most c such that all graphs obtained from $G - D$ by changing adjacencies between v and $V(G) \setminus (B \cup D)$ are also contained in Π. By Lemma 3.10 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, (Co-)Comparability Deletion, Permutation 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. [18], 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 3.5. 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 4.1. *Let G be a graph, $P = (v_1, ..., v_n)$ with an even $n \geq 4$ 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 an 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 4.1 would result in a graph that contains an odd hole. Note that in Fig. 1, 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.

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

Proof. Consider some graph H with a 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 $\text{inc}^4_{(H,X)}(v) = \sum_{u \in D} \text{inc}^4_{(H,X)}(u)$. 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 $\text{inc}^{4,(\emptyset,\{p,q\})}_{(H,X)}(v) = \sum_{u \in D'} \text{inc}^{4,(\emptyset,\{p,q\})}_{(H,X)}(u)$. Such set exists by Lemma 3.10. Since C is an odd hole, $|V(P)|$ is even. Hence, by Lemma 4.1, $G[V(P) \cup \{u\}]$ 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 $D'_{\{a,b\}} \subseteq D'$ be the set of vertices that see edge $\{a, b\} \in E(P)$. Note that $D'_{\{a,b\}} = \emptyset$ in case $\{a, b\} \not\subseteq X$. Otherwise, in order to satisfy $\sum_{u \in D'} \text{inc}^{(\emptyset,\{p,q\})}(u)[\emptyset, \{p, q, a, b\}] = \text{inc}^{(\emptyset,\{p,q\})}(v)[\emptyset, \{p, q, a, b\}] = 0$ over $\mathbb{F}_2$, for $\{a, b\} \in E(P)$, we require $|D'_{\{a,b\}}|$ 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 3.6 and Lemma 3.9 we get the following.

Corollary 4.3. *The graph properties Π_{OH}, Π_{OAH}, and Π_P are characterized by rank-4 adjacencies with singleton replacements.*

Theorem 4.4. Perfect Deletion Parameterized by Vertex Cover *admits a polynomial kernel with $\mathcal{O}(|X|^5)$ vertices.*

Proof. By Corollary 4.3 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 $|V(H)| \leq 2 \cdot \text{vc}(H)$. Therefore by Theorem 3.5 it follows that Π_P-free Deletion and hence Perfect Deletion Parameterized by Vertex Cover admits a polynomial kernel with $\mathcal{O}(|X|^5)$ vertices. □

A variation of Theorem 4.2 presented in the next section 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 Vertex Cover of $\mathcal{O}(|X|^4)$ vertices.

4.2. Even-hole-free deletion

An even hole is a cycle C_n, where $n \geq 4$ is an even number. Let Π_{EH} be the set of graphs that contain an even hole as an induced subgraph. Then the Even-hole-free Deletion problem corresponds to the Π_{EH}-free Deletion problem.

Lemma 4.5. *Let G be a graph, $P = (v_1, ..., v_n)$ with an odd $n \geq 3$ 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 vertices of P. Then $G[V(P) \cup \{y\}]$ contains an even hole as an induced subgraph.*

Proof. We prove the claim by induction on n. Consider the case that $n = 3$. Since y is adjacent to both v_1 and v_3, the only way for it to see an even number of vertices of P is for y not to be adjacent to v_2. Therefore $G[V(P) \cup \{y\}]$ induces an even hole and the statement holds.

In the remainder, assume that the claim holds for $n' < n$, where $n \geq 5$ and both n' and n are odd. If y is not adjacent to v_i for any $1 < i < n$, then $G[V(P) \cup \{y\}]$ induces an even hole and the statement holds. Otherwise, since y sees an even number of path vertices it sees at least two more. Let $j > 1$ be the smallest index such that y is adjacent to v_j and let $j' < n$ be the largest index such that y is adjacent to $v_{j'}$. If j is odd, then $G[\{v_1, \ldots, v_j, y\}]$ induces an even hole and the statement holds. If j' is odd, then $G[\{v_{j'}, \ldots, v_n, y\}]$ induces an odd hole and the statement holds. Otherwise, both j and j' are even. But then $P' = (v_j, \ldots, v_{j'})$ is a path of odd length at least 3, where y is adjacent to both endpoints and sees an even number of vertices. So by the induction hypothesis, the statement holds. □

Theorem 4.6. *Π_{EH} is characterized by rank-3 adjacencies with singleton replacements.*

Proof. Consider some graph H with a vertex cover X and let $D \subseteq V(H) \setminus X$ such that $H - D \in \Pi_{EH}$. Let v be an arbitrary vertex in $V(H) \setminus (D \cup X)$ such that $\mathrm{inc}^3_{(H,X)}(v) = \sum_{u \in D} \mathrm{inc}^3_{(H,X)}(u)$. We show that $H - v - (D \setminus \{v'\}) \in \Pi_{EH}$ for some $v' \in D$.

Let C be an even hole in $H - D$. If $v \notin V(C)$, then for every $v' \in D$ we have $H - v - (D \setminus \{v'\}) \in \Pi_{EH}$. So suppose that $v \in C$. Let $C = (v, p, v_1, \ldots, v_{n-3}, q)$, where $|V(C)| = n$. Consider the induced path $P = (p, v_1, \ldots, v_{n-3}, q)$. 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 $\mathrm{inc}^{3,(\emptyset,\{p,q\})}_{(H,X)}(v) = \sum_{u \in D'} \mathrm{inc}^{3,(\emptyset,\{p,q\})}_{(H,X)}(u)$. Such set exists by Lemma 3.10. Since C is an even hole, $|V(P)|$ is odd. Hence, by Lemma 4.5, $G[V(P) \cup \{u\}]$ contains an odd hole if there exists some $u \in D'$ that sees an even number of vertices of P. Suppose for the sake of contradiction that every vertex in D' sees an odd number of vertices of $V(P)$. Then also every vertex in D' sees an odd number of vertices of $V(P) \setminus \{p, q\}$. Let V_u be the set of vertices in $V(P) \setminus \{p, q\}$ that are seen by $u \in D'$. Then $\sum_{u \in D'} |V_u|$ is odd as it is a sum of an odd number of odd numbers. Let $D'_a \subseteq D'$ be the set of vertices that see vertex $a \in V(P) \setminus \{p, q\}$. In order to satisfy $\sum_{u \in D'} \mathrm{inc}^{(\emptyset,\{p,q\})}(u)[\emptyset, \{p, q, a\}] = \mathrm{inc}^{(\emptyset,\{p,q\})}(v)[\emptyset, \{p, q, a\}] = 0$ over $\mathbb{F}_2$, for $a \in V(P) \setminus \{p, q\}$, we require $|D'_a|$ to be even. But then $\sum_{a \in V(P) \setminus \{p,q\}} |D'_a| = \sum_{u \in D'} |V_u|$ would also need to be an even number. This contradicts the fact that $\sum_{u \in D'} |V_u|$ is odd. Therefore there must exist some $u \in D'$ that sees an even number of vertices of P. □

Since every graph that contains an even hole has at least one edge and for every hole H, we have $|V(H)| \leq 2 \cdot \mathrm{vc}(H)$. By Theorem 3.5 we obtain:

Theorem 4.7. Even-hole-free Deletion Parameterized by Vertex Cover *admits a polynomial kernel with $\mathcal{O}(|X|^4)$ vertices.*

4.3. AT-free deletion

In his dissertation, Köhler [29] 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 large asteroidal witnesses. This completes a list given by Gallai [20], who gave a list of vertex-minimal asteroidal triples that do not contain C_5 as an induced subgraph. Köhler contributed the 5 vertex-minimal asteroidal triples that do contain C_5 as an induced subgraph, each consisting of 6 or 7 vertices. Let Π_{AT} be the set of graphs that contain an asteroidal triple. We classify graphs in Π_{AT} in several types. The most difficult cases are captured by the following definition; see Fig. 2.

Definition 4.8. The different types of *large asteroidal witnesses* (LAWs) are constructed as follows.

- $\dagger_z$-AW: A graph G such that $V(G) = \{t_l, t_r, t, c\} \cup \{b_1, \ldots, b_z\}$, where $t_l = b_0$ and $t_r = b_{z+1}$, $E(G) = \{\{t, c\}\} \cup \{\{c, b_i\} \mid i \in [z]\} \cup \{\{b_{i-1}, b_i\} \mid i \in [z+1]\}$, and $z \geq 2$.
- $\ddagger_z$-AW: A graph G such that $V(G) = \{t_l, t_r, t, c_1, c_2\} \cup \{b_1, \ldots, b_z\}$, where $t_l = b_0$ and $t_r = b_{z+1}$, $E(G) = \{\{t, c_1\}, \{t, c_2\}, \{c_1, c_2\}, \{c_1, t_l\}, \{c_2, t_r\}\} \cup \{\{c_j, b_i\} \mid i \in [z], j \in \{1, 2\}\} \cup \{\{b_{i-1}, b_i\} \mid i \in [z+1]\}$, and $z \geq 1$.
- $\diamond_z$-AW: A graph G such that $V(G) = \{t_l, t_r, t, c_1, c_2\} \cup \{b_1, \ldots, b_z\}$, where $t_l = b_0$ and $t_r = b_{z+1}$, $E(G) = \{\{t, c_1\}, \{t, c_2\}, \{c_1, t_l\}, \{c_2, t_r\}\} \cup \{\{c_j, b_i\} \mid i \in [z], j \in \{1, 2\}\} \cup \{\{b_{i-1}, b_i\} \mid i \in [z+1]\}$, and $z \geq 1$.

Here t_l, t_r, and t are called *terminal* vertices, with *top* vertex t. These three vertices form an asteroidal triple. Vertices c, c_1, and c_2 are called *center* vertices. Finally b_i for $i \in [z]$ are called *path* vertices.

Using this definition we can give the characterization of the graphs Π_{AT} containing an asteroidal triple described above.

Theorem 4.9 *([29, Cor. 1.43]). Let $\Pi_{C_{\geq 6}}$ be the set of graphs that contain a chordless cycle of length at least 6. Let Π_{LAW} be the set of graphs that contain a large asteroidal witness. There exists a set $\mathcal{S}$ of graphs on 6 or 7 vertices each, such that for Π_S the graphs which contain a graph from $\mathcal{S}$ it holds that $\Pi_{AT} = \Pi_{C_{\geq 6}} \cup \Pi_{LAW} \cup \Pi_S$.*

Fomin et al. [18, Proposition 3] show that $\Pi_{C_{\geq 6}}$ is characterized by 5 adjacencies. By Proposition 2.3 we have that Π_S is characterized by 6 adjacencies. Therefore we focus our attention on the large asteroidal witnesses.

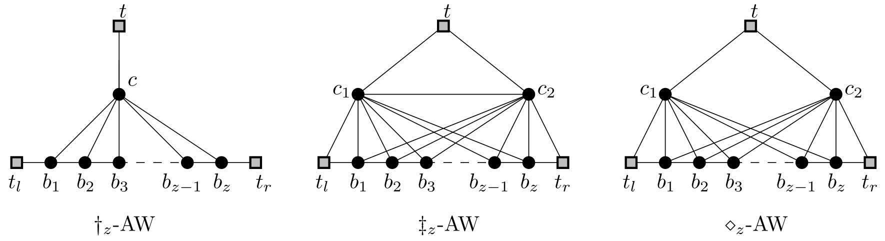

Fig. 2. Asteroidal witnesses.

The difficulty with these large asteroidal witnesses lies with the top vertex t. If t would be adjacent to a single path vertex b_i for some $i \in [z]$, then the resulting graph no longer contains an asteroidal triple. It is possible to show that, because of this difficulty, no finite adjacency characterization exists for asteroidal witnesses [12, 3.4.2]. We show that if we have a set of vertices whose vectors sum to the vector of t, then we can find a valid replacement in this set.

Theorem 4.10. *Let G be a graph with a vertex cover X, containing $O = x_z$-AW for some $x \in \{\dagger, \ddagger, \diamond\}$ according to Definition 4.8 with $t \notin X$. Let $D \subseteq V(G)$ disjoint from O and X such that $\mathrm{inc}^8_{(G,X)}(t) = \sum_{y \in D} \mathrm{inc}^8_{(G,X)}(y)$. Then $G[(V(O) \cup \{y\}) \setminus \{t\}]$ contains an asteroidal triple for some $y \in D$.*

Proof. Let $G' := G[O \cup D]$ and $X' := X \cap V(G')$. By Proposition 3.2 we have that $\mathrm{inc}^8_{(G',X')}(t) = \sum_{y \in D} \mathrm{inc}^8_{(G',X')}(y)$. Let C be the center vertices of O; then t is adjacent to the center vertices in C and not to the other terminals t_l and t_r. Let $D' \subseteq D$ be a set of vertices that share adjacencies with t to $\{t_l, t_r\} \cup C$ such that $\mathrm{inc}^{8,(T,C)}_{(G',X')}(t) = \sum_{u \in D'} \mathrm{inc}^{8,(T,C)}_{(G',X')}(u)$ where $T = \{t_l, t_r\} \cap X'$. Such D' exists by Lemma 3.10. If there are distinct $w, w' \in D'$ whose vectors are identical, then $\sum_{u \in D'} \mathrm{inc}^{8,(T,C)}_{(G',X')}(u) = \sum_{u \in D' \setminus \{w,w'\}} \mathrm{inc}^{8,(T,C)}_{(G',X')}(u)$. Therefore we can remove all duplicate vectors, and hence assume that the vectors for vertices in D' are unique. If there exists $y \in D'$ that is not adjacent to any b_i for $i \in [z]$, then $G[(V(O) \cup \{y\}) \setminus \{t\}]$ induces an x_z-AW and we are done. Otherwise, let q_y be the smallest index such that y is adjacent to b_{q_y} and let q'_y be the largest index such that y is adjacent to $b_{q'_y}$. Let $y^* \in D'$ be a vertex that maximizes $q'_{y^*} - q_{y^*}$. If $q'_{y^*} = q_{y^*}$, then y^* is adjacent to a single vertex $b_{q_{y^*}}$. However, since t is not adjacent to this vertex, there must exist $y \in D'$ that is adjacent to the same vertex $b_{q_{y^*}}$ in order to satisfy the vector sum. But then y and y^* have the same projected incidence vector, a contradiction. It remains to deal with the case that $q'_{y^*} > q_{y^*}$.

If there exists $y \in D'$ (possibly y^* itself) adjacent to $b_{q_{y^*}}$ and $b_{q'_{y^*}}$, but nonadjacent to some b_i with $q_{y^*} + 1 < i < q'_{y^*} - 1$, then t_l, t_r, and b_i form an asteroidal triple in the graph $G[(V(O) \cup \{y\}) \setminus \{t\}]$, since y provides a path between t_l and t_r avoiding $N(b_i)$. Therefore suppose there is no such y. Then any vertex $y \in D'$ adjacent to $b_{q_{y^*}}$ and $b_{q'_{y^*}}$ is also adjacent to all vertices in between, except possibly $b_{q_{y^*}+1}$ and $b_{q'_{y^*}-1}$. Let $Q = (\{t_l, t_r, b_{q_{y^*}+1}, b_{q'_{y^*}-1}\} \setminus N_G(y^*)) \cap X$ and $R = C \cup \{b_{q_{y^*}}, b_{q'_{y^*}}\} \cup (\{b_{q_{y^*}+1}, b_{q'_{y^*}-1}\} \cap N_G(y^*))$. We have $|Q| + |R| \leq 8$. The entry $\mathrm{inc}^{8,(T,C)}_{(G',X')}(y^*)[Q,R] = 1$ since we defined Q and R based on the neighborhood of y^*, but $\mathrm{inc}^{8,(T,C)}_{(G',X')}(t)[Q,R] = 0$ since t is not adjacent to the path vertices $b_{q_{y^*}}, b_{q'_{y^*}}$. So there must exist another $y' \in D'$ with $\mathrm{inc}^{8,(T,C)}_{(G',X')}(y')[Q,R] = 1$. By the maximality of the range seen by y^*, we have $q_{y^*} = q_{y'}$ and $q'_{y^*} = q'_{y'}$. But then y^* and y' have the same projected incidence vector, since they agree on the adjacency to $b_{q_{y^*}}, b_{q_{y^*}+1}, b_{q'_{y^*}-1}, b_{q'_{y^*}}$ and are adjacent to all vertices in between. This contradicts the assumption that all vectors for vertices in D' are unique. □

Theorem 4.11. *Π_{AT} is characterized by rank-8 adjacencies with singleton replacements.*

Proof. Consider some graph H with a vertex cover X and let $D \subseteq V(H) \setminus X$ such that $H - D \in \Pi_{AT}$. Let v be an arbitrary vertex in $V(H) \setminus (D \cup X)$ such that $\mathrm{inc}^8_{(H,X)}(v) = \sum_{u \in D} \mathrm{inc}^8_{(H,X)}(u)$. We show that $H - v - (D \setminus \{v'\}) \in \Pi_{AT}$ for some $v' \in D$. The statement for the cases where $H - D \in \Pi_S$ and $H - D \in C_{\geq 6}$ follow from the fact that they are characterized by 6 and 5 adjacencies respectively. In the remainder we have $H - D \in \Pi_{LAW}$.

Suppose that $H - D$ contains $O = \dagger_z$-AW for some $z \geq 2$. We do a case distinction on v. If $v \notin V(O)$, then for any $v' \in D$ we have $H - v - (D \setminus \{v'\}) \in \Pi_{AT}$, so suppose that $v \in V(O)$. If v is top terminal t, then by Theorem 4.10 we have $H - v - (D \setminus \{v'\}) \in \Pi_{AT}$ for some $v' \in D$. If v is a center vertex c, then v is adjacent to t, b_1, and b_2 and not to t_l and t_r. Let $D' \subseteq D$ be the set of vertices that share adjacencies with v to $\{t, b_1, b_2, t_l, t_r\}$. The set D' is nonempty by Lemma 3.10. Consider any $v' \in D'$, then $\{t, t_l, t_r\}$ still is an asteroidal triple in the graph $H - v - (D \setminus \{v'\})$ and the statement holds. A similar argument based on shared adjacencies works for all cases besides the top vertex t. For each remaining possible role of v in the $\dagger_z$-AW O, we give the set of adjacencies that needs to be shared by v and the vertices in D' such that replacing v by any member of D' yields a graph in which $\{t, t_l, t_r\}$ is an asteroidal triple. For $v = t_l$, the adjacencies to be shared by D' and v are $\{b_z, c, t, t_r, b_1\}$. The case $v = t_r$ is symmetric. For $v = b_i$ for some $1 < i < z$, the adjacencies to be shared by D' and v are $\{t, t_l, t_r, b_{i-1}, b_{i+1}, c\}$. For $v = b_1$, the adjacencies to be shared by D' and v are $\{t, t_r, t_l, b_2, c\}$. Again the case for $v = b_z$ is symmetric.

Finally suppose that $H - D$ contains $O = \ddagger_z$-AW or $O = \diamond_z$-AW for some $z \geq 1$. If $v \notin V(O)$, then for any $v' \in D$ we have $H - v - (D \setminus \{v'\}) \in \Pi_{AT}$, so suppose that $v \in V(O)$. If v is top terminal t, then by Theorem 4.10 we have $H - v - (D \setminus \{v'\}) \in \Pi_{AT}$ for some $v' \in D$. For the remaining cases of v, an argument based on shared adjacencies of v and $D' \subseteq D$ to some small vertex set again suffices to prove the statement; in each case it is easy to verify that replacing v by any vertex from D' yields a graph with an asteroidal triple. If v is a center vertex c_1, then the adjacencies to be shared by D' and v are $\{t_r, t, t_l\}$. The case $v = c_2$ is symmetric. For $v = t_l$, the adjacencies to be shared by D' and v are $\{c_2, t_r, t, c_1, b_1\}$. Again the case $v = t_r$ is symmetric. For $v = b_i$ for some $1 < i < z$, the adjacencies to be shared by D' and v are $\{t, b_{i-1}, b_{i+1}\}$. For $v = b_1$, the adjacencies to be shared by D' and v are $\{t, t_l, b_{i+1}\}$. The case $v = b_z$ is again symmetric. Note that for all of these cases, the set of adjacencies S that some vertex in D needs to share with v contains at most 8 elements. This completes the proof. □

Theorem 4.12. AT-free Deletion Parameterized by Vertex Cover *admits a polynomial kernel with* $\mathcal{O}(|X|^9)$ *vertices.*

Proof. Every graph in Π_{AT} contains at least one edge. By Theorem 4.11 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 $|V(C)| \leq 2 \cdot \text{vc}(C)$. Finally each large asteroidal witness consists of an induced path with 2 or 3 additional vertices, so for all vertex-minimal graphs $G \in \Pi_{AT}$ we have $|V(G)| \in \mathcal{O}(\text{vc}(G))$. Hence, by Theorem 3.5, AT-free Deletion parameterized by the size of a vertex cover admits a polynomial kernel with $\mathcal{O}(|X|^9)$ vertices. □

4.4. Interval deletion

Interval Deletion does not fit in the framework of Fomin et al. [18], 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] as a byproduct on their work on the smaller parameterization by solution size. We show that our framework captures Interval Deletion. 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 4.3 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 4.13. Interval Deletion Parameterized by Vertex Cover *admits a polynomial kernel with* $\mathcal{O}(|X|^9)$ *vertices.*

Proof. Every graph in Π_{IV} contains at least one edge. By Theorem 4.11, Π_{AT} is characterized by rank-8 adjacencies. Furthermore, $\Pi_{C_{\geq 4}}$ is characterized by 3 adjacencies as shown by Fomin et al. [18, Proposition 3], and therefore by Lemma 3.11 also by rank-3 adjacencies. Therefore by Lemma 3.9, 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 $|V(C)| \leq 2 \cdot \text{vc}(C)$. Recall that $\Pi_{AT} = \Pi_S \cup \Pi_{C_{\geq 6}} \cup \Pi_{LAW}$. Each vertex minimal graph in Π_S contains at most 7 vertices. Finally each large asteroidal witness consists of an induced path with 2 or 3 additional vertices, so for each vertex-minimal $G \in \Pi_{IV}$ we have $|V(G)| \in \mathcal{O}(\text{vc}(G))$. Therefore by Theorem 3.5, Interval Deletion parameterized by the size of a vertex cover admits a polynomial kernel with $\mathcal{O}(|X|^9)$ vertices. □

4.5. Transitive orientability

A graph G is transitively orientable if we can assign a direction to every edge, such that for distinct vertices $a, b, c \in V(G)$ the following holds: if (a, b) and (b, c) are arcs in the orientation, then (a, c) is also an arc. Such graphs are known as *comparability graphs*. The following definition generalizes the notion of an asteroidal triple.

Definition 4.14. For $n \geq 1$, a $(2n+1)$-*asteroid* in an undirected graph G is a sequence $y_1, \ldots, y_{2n+1}$ of distinct vertices such that for each $i \in [2n+1]$ there is a $y_i y_{i+1}$-path W_i in $G - N_G[y_{i+1+n}]$, where indices are evaluated modulo $2n+1$.

An asteroidal triple is equivalent to a 3-asteroid. Gallai [20] gives the following characterization of comparability graphs G in terms of the absence of asteroids in their edge-complement $\overline{G}$.

Theorem 4.15 *(Theorem 1.20 [20]). A graph G is transitively orientable if and only if $\overline{G}$ contains no asteroid.*

Based on this characterization, Gallai gives a forbidden subgraph characterization of comparability graphs, giving the vertex-minimal of 3-asteroids (that do not contain C_5 as an induced subgraph, as it already forms a 5-asteroid) and the complements of $(2n+1)$-asteroids for $n > 1$. Recall that Π_{AT}, the set of graphs that contain an asteroidal triple, is characterized by rank-8 adjacencies. By Lemma 3.6 we have that $\overline{\Pi}_{AT}$ is also characterized by rank-8 adjacencies. Therefore we

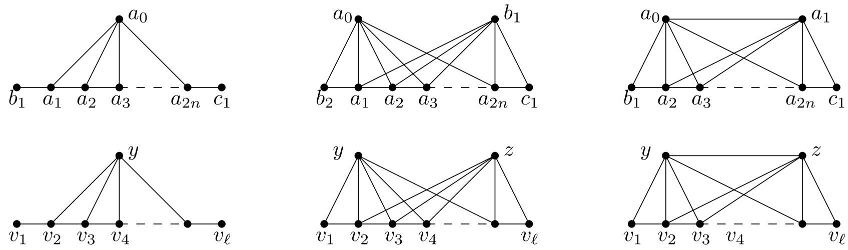

Fig. 3. Vertex minimal transitive orientability obstructions $\Gamma_2(2n+1)$, $\Gamma_3(2n+1)$, and $\Gamma_4(2n+1)$. On the top, the original labeling of the vertices as used by Gallai is shown. At the bottom, we present an alternative labeling more suitable for our argumentation.

focus on the complements of vertex-minimal $(2n+1)$-asteroids for $n > 1$. Gallai [20] introduced four graph classes; the first are the odd holes and the others are shown in Fig. 3.

Definition 4.16. For $n \geq 2$, consider the following graphs.

- $\Gamma_1(2n+1)$: The chordless cycle C_{2n+1}.
- $\Gamma_2(2n+1)$: The graph on vertices $a_0, \ldots, a_{2n}$, b_1, c_1, and with the edges $\{a_i, a_{i+1}\}$ $(0 \leq i \leq 2n$, with $a_{2n+1} = a_0)$, $\{a_0, a_i\}$ $(2 \leq i \leq 2n-1)$, $\{a_1, b_1\}$, and $\{a_{2n}, c_1\}$.
- $\Gamma_3(2n+1)$: The graph on vertices $a_0, \ldots, a_{2n}$, b_1, b_2, c_1, and with the edges $\{a_i, a_{i+1}\}$ $(0 \leq i \leq 2n$, with $a_{2n+1} = a_0)$, $\{a_0, a_i\}$ $(2 \leq i \leq 2n-1)$, $\{b_1, a_i\}$ $(1 \leq i \leq 2n)$, $\{b_1, c_1\}$, $\{b_2, a_0\}$, $\{b_2, a_1\}$, and $\{c_1, a_{2n}\}$.
- $\Gamma_4(2n+1)$: The graph on vertices $a_0, \ldots, a_{2n}$, b_1, c_1, and with the edges $\{a_i, a_{i+1}\}$ $(0 \leq i \leq 2n$, with $a_{2n+1} = a_0)$, $\{a_0, a_i\}$ $(2 \leq i \leq 2n-1)$, $\{a_1, a_j\}$ $(3 \leq j \leq 2n)$, $\{b_1, a_0\}$, $\{b_1, a_2\}$, $\{c_1, a_{2n}\}$, and $\{c_1, a_1\}$.

For $j \in [4]$, let $\Gamma_j = \bigcup_{n \geq 2} \Gamma_j(2n+1)$.

Gallai [20] describes Γ_1 in Statement 6.14, Γ_2 in Statement 6.26, Γ_3 in Statement 6.27, Γ_4 in Statement 6.32, and concludes with the following.

Theorem 4.17 *(Statement 6.33 [20]). A graph that does not contain the complement of a 3-asteroid as an induced subgraph is transitively orientable if and only if it does not contain any graph of $\bigcup_{j \in [4]} \Gamma_j$ as an induced subgraph.*

Let Π_{LAC} be the set of graphs that contain a graph from $\Gamma_2 \cup \Gamma_3 \cup \Gamma_4$. Let Π_C be the set of graphs that are not comparability. Note that $\Pi_C = \overline{\Pi}_{AT} \cup \Pi_{OH} \cup \Pi_{LAC}$. Besides comparability graphs, other related graph classes are *co-comparability* and *permutation* graphs. Co-comparability graphs are those graphs whose complement is transitively orientable, thus those that are $\overline{\Pi}_C$-free. A permutation graph corresponds to a permutation π of vertex set $[n]$, where $i < j$ are adjacent if $\pi(i) > \pi(j)$.

Theorem 4.18 *([15], cf. Theorem 4.7.1 [5]). A graph G is a permutation graph if and only if G and $\overline{G}$ are comparability graphs.*

Therefore permutation graphs are those that are $(\Pi_C \cup \overline{\Pi}_C)$-free. In order to obtain polynomial kernels for their respective deletion problems parameterized by vertex cover, the only missing ingredient is a characterization of Π_{LAC} by rank-c adjacencies. While it is not clear whether such a characterization exists, we will show how to preserve the fact that these graphs are not transitive orientable by remembering a finite number of adjacencies. Together with the following lemma, this will give us the desired kernels.

Lemma 4.19. *Let Π be a graph property such that $\Pi = \Pi_1 \cup \Pi_2$, where Π_1 is characterized by rank-c adjacencies, and for every $G \in \Pi_2$, for every $v \in V(G)$, there is a set $B \subseteq V(G) \setminus \{v\}$ of size at most d such that all graphs G' obtained from G by adding and removing edges between v and vertices in $V(G) \setminus B$ are contained in Π. Then Π is characterized by rank-$\max(c, d)$ adjacencies.*

Proof. Consider a graph H with a vertex cover X and let $D \subseteq V(H) \setminus X$ such that $H - D \in \Pi$. Let $b = \max(c, d)$. Let v be an arbitrary vertex in $V(H) \setminus (D \cup X)$ such that $\text{inc}^b_{(H,X)}(v) = \sum_{u \in D} \text{inc}^b_{(H,X)}(u)$. We show that $H - v - (D \setminus D') \in \Pi$ for some $D' \subseteq D$. We do a case distinction on whether $H - D \in \Pi_1$ or $H - D \in \Pi_2$.

First suppose $H - D \in \Pi_1$. By Proposition 3.2 we have that $\text{inc}^c_{(H,X)}(v) = \sum_{u \in D} \text{inc}^c_{(H,X)}(u)$. Since Π_1 is characterized by rank-c adjacencies it follows that there exists $D' \subseteq D$ such that $H - v - (D \setminus D') \in \Pi_1$, thus $H - v - (D \setminus D') \in \Pi$.

Now suppose that $H - D \in \Pi_2$. By Proposition 3.2 we have that $\text{inc}^d_{(H,X)}(v) = \sum_{u \in D} \text{inc}^d_{(H,X)}(u)$. By the statement of the lemma, there exists a set $B \subseteq V(H) \setminus D$ of size at most d such that all graphs obtained by changing adjacencies between v and $V(H) \setminus (B \cup D)$ are contained in Π. By Lemma 3.10 there exists $w \in D$ that shares adjacencies with v to B. It follows that $H - v - (D \setminus \{w\}) \in \Pi$. $\square$

The following property of transitive orientations will be useful.

Observation 4.20. *Let G be an undirected graph and let $P = (v_1, \dots, v_\ell)$ be a simple path in G, such that $\{v_i, v_{i+2}\} \notin E(G)$ for all $i \in [\ell - 2]$. Then in any transitive orientation of G, either all edges of P are oriented towards odd-indexed vertices, or all edges of P are oriented towards even-indexed vertices.*

Using this observation we prove the main structural result concerning comparability graphs.

Lemma 4.21. *Let $G \in \Gamma_2 \cup \Gamma_3 \cup \Gamma_4$ be a graph as in Definition 4.16. Then for every $v \in V(G)$, there is a set $B \subseteq V(G) \setminus \{v\}$ of size at most 8 such that all graphs G' obtained from G by adding and removing edges between v and vertices in $V(G) \setminus B$ are not transitively orientable, and therefore contained in Π_C.*

Proof. We do a case distinction on containment of G in $\Gamma_2 \cup \Gamma_3 \cup \Gamma_4$. To simplify the presentation, we work with a different labeling of the vertices in the graphs $\Gamma_2 \cup \Gamma_3 \cup \Gamma_4$ compared to the one used by Gallai in Definition 4.16. The alternative labeling is shown in Fig. 3 and highlights the role of the horizontal path.

- First suppose $G = \Gamma_2(2n + 1)$. In the case that $v = y$, let $B = \{v_1, v_2, v_{\ell-1}, v_\ell\}$ where $\ell = 2n + 2$. Let G' be a graph obtained from G by altering adjacencies between v and $V(G) \setminus B$. Let $P' = G'[\{v_1, ..., v_\ell\}]$. Suppose that G' has a transitive orientation. As P' is an induced path in G and therefore in G', by Observation 4.20 the edges of P' need to be oriented away from each other in any transitive orientation. Without loss of generality, suppose they are directed towards the odd indexed vertices (reverse the numbering along P' otherwise). Then the edge $\{v_2, y\}$ needs to be directed towards y due to the absence of edge $\{v_1, y\}$. The edge $\{v_{\ell-1}, y\}$ needs to be directed towards $v_{\ell-1}$ due to the absence of edge $\{y, v_\ell\}$. For a transitive orientation we would then need an arc $(v_2, v_{\ell-1})$, but the edge $\{v_2, v_{\ell-1}\}$ is not in the graph since $\ell \geq 6$; a contradiction.
  In the case that $v = v_i$ for some $i \in [\ell]$, let $B = \{y, v_2, v_{\ell-1}, v_{i-1}, v_{i-2}, v_{i+1}, v_{i+2}\} \setminus \{v\}$ (ignoring indices below 1 or above ℓ). Let G'' be a graph obtained from G by altering adjacencies between v and $V(G) \setminus B$. Note that $(v_1, ..., v_\ell)$ is a simple path in G' satisfying the condition of Observation 4.20. Therefore the result follows as in the previous case, as no edges incident on y were added or removed and $\{v_2, v_{\ell-1}\} \notin E(G'')$.
- Next suppose $G = \Gamma_3(2n + 1)$. In the case that $v = y$ (the case $v = z$ is symmetric), let $B = \{z, v_1, v_2, v_{\ell-1}, v_\ell\}$ where $\ell = 2n + 2$. Let G' be a graph obtained from G by altering adjacencies between v and $V(G) \setminus B$. Let $P' = G'[\{v_1, ..., v_\ell\}]$. Suppose that G' graph has a transitive orientation. By Observation 4.20, the edges of P' need to be oriented away from each other. Suppose they are directed towards the odd indexed vertex (the even case is symmetric). Then the edge $\{v_2, z\}$ needs to be directed towards z due to $\{v_1, z\}$ being a non-edge. The edge $\{v_{\ell-1}, y\}$ needs to directed towards $v_{\ell-1}$ due to $\{y, v_\ell\}$ being a non-edge. Since y is not adjacent to z, the edge $\{v_2, y\}$ needs to be directed towards v_y. But the edge $\{v_2, v_{\ell-1}\}$ is not in the graph since $\ell \geq 6$, a contradiction.
  In the case that $v = v_i$ for some $i \in [\ell]$, let $B = \{y, z, v_2, v_{\ell-1}, v_{i-1}, v_{i-2}, v_{i+1}, v_{i+2}\} \setminus \{v\}$ (ignoring indices below 1 or above ℓ). Let G'' be a graph obtained from G by altering adjacencies between v and $V(G) \setminus B$. As before, $(v_1, ..., v_\ell)$ is a simple path in G' satisfying the condition of Observation 4.20. The result follows as in the previous case, since no edges incident on y were added or removed and $\{v_2, v_{\ell-1}\} \notin E(G'')$.
- Finally suppose $G = \Gamma_4(2n + 1)$. In the case that $v = y$ (the case $v = z$ is symmetric), let $B = \{z, v_1, v_2, v_{\ell-1}, v_\ell\}$ where $\ell = 2n + 1$. Let G' be a graph obtained from G by altering adjacencies between v and $V(G) \setminus B$. Let $P' = G'[\{v_1, ..., v_\ell\}]$. Suppose that G' graph has a transitive orientation. By Observation 4.20, the edges of P' need to be oriented away from each other. Suppose first that they are directed towards the odd indexed vertices. Then the edge $\{v_2, z\}$ needs to be directed towards z as $\{v_1, z\}$ is a non-edge, the edge $\{v_{\ell-1}, y\}$ needs to directed towards y as $\{y, v_\ell\}$ is a non-edge. Then $\{v_1, y\}$ needs to be directed towards y as $\{v_1, v_{\ell-1}\}$ is a non-edge and $\{v_\ell, z\}$ towards z as $\{v_2, v_\ell\}$ is a non-edge. But then no matter which way we orient $\{y, z\}$ we get a contradiction, since edges $\{v_1, z\}$ and $\{v_\ell, y\}$ are not in the graph. If the arcs on P' are oriented towards the even indexed vertices instead, then the argumentation is the same except that the relevant edges incident on y and z are forced to be oriented away from, instead of towards, y and z.
  In the case that $v = v_i$ for some $i \in [\ell]$, let $B = \{w, y, z, v_{i-1}, v_{i-2}, v_{i+1}, v_{i+2}\} \setminus \{v\}$ (ignoring indices below 1 or above ℓ) where $w = v_{\ell-1}$ if $i = 1$, $w = v_\ell$ if $i = 2$, $w = v_1$ if $i = \ell - 1$ and $w = v_2$ if $i = \ell$, and $w = y$ otherwise. Let G'' be a graph obtained from G by altering adjacencies between v and $V(G) \setminus B$. Our choice of B ensures that $v \notin B$, while no edges incident on $\{y, z\}$ were added or removed and $\{v_1, z\}, \{v_\ell, y\}, \{v_1, v_{\ell-1}\}, \{v_2, v_\ell\} \notin E(G'')$. Again we have that $(v_1, ..., v_\ell)$ is a simple path in G' satisfying the condition of Observation 4.20, so that the same argumentation applies as in the previous case.

Since the set B used for each case has size at most 8, this concludes the proof. □

From the previous lemmas, we can conclude the following.

Corollary 4.22. *The graph properties Π_C, $\overline{\Pi}_C$, and $\Pi_C \cup \overline{\Pi}_C$ are characterized by rank-8 adjacencies.*

Theorem 4.23. Comparability Deletion, Co-comparability Deletion, *and* Permutation Deletion Parameterized by Vertex Cover *admit a polynomial kernel with* $\mathcal{O}(|X|^9)$ *vertices.*

Proof. Every graph in Π_C and $\overline{\Pi}_C$ contains at least one edge. By Corollary 4.22 it follows that Π_C, $\overline{\Pi}_C$, and $\Pi_C \cup \overline{\Pi}_C$ are characterized by rank-8 adjacencies, where $\Pi_C \cup \overline{\Pi}_C$ is a set of forbidden induced subgraphs characterizing permutation graphs by Theorem 4.18. Finally, we claim that for any obstruction G, we have $|V(G)| \leq c \cdot \text{vc}(G)$ for some constant $c \in \mathbb{N}$. To see this, note that all the infinite families consist of at most three vertices together with a path or cycle or its complement. The claimed kernels therefore follow from Theorem 3.5. □

4.6. (Almost) wheel-free deletion

Let $\Pi_{W_{\geq 3}}$ be the set of graphs that contain a wheel of size at least 3 as an induced subgraph. Then Wheel-free Deletion corresponds to $\Pi_{W_{\geq 3}}$-free Deletion. Using a similar argument as in Theorem 4.6, we obtain a characterization by rank-4 adjacencies.

Theorem 4.24. *The graph property $\Pi_{W_{\geq 3}}$ is characterized by rank-4 adjacencies.*

Proof. Consider some graph H with a vertex cover X and set $D \subseteq V(H) \setminus X$ such that $H - D \in \Pi_{W_{\geq 3}}$. Let v be an arbitrary vertex in $V(H) \setminus (D \cup X)$ such that $\text{inc}^4_{(H,X)}(v) = \sum_{u \in D} \text{inc}^4_{(H,X)}(u)$. We show that $H - v - (D \setminus D') \in \Pi$ for some $D' \subseteq D$.

Since $H - D \in \Pi_{W_{\geq 3}}$, it contains a wheel W_n for some $n \geq 3$ as an induced subgraph. If $v \notin V(W_n)$, then for any $D' \subseteq D$, $H - v - (D \setminus D') \in \Pi_{W_{\geq 3}}$. So suppose $v \in V(W_n)$, then v sees three vertices of W_n that induce a complete graph if $n = 3$, or a P_3 otherwise. Let p, q, and r be these vertices. Let $D' \subseteq D$ be the set of vertices that share adjacencies with v to $\{p, q, r\}$ such that $\text{inc}^{(\emptyset,\{p,q,r\})}(v) = \sum_{u \in D'} \text{inc}^{(\emptyset,\{p,q,r\})}(u)$. Such set D' exists by Lemma 3.10. If $n = 3$, then any $d \in D'$ that is adjacent to p, q, and r is a valid replacement. Suppose that $n > 3$. If $|D'| = 1$, then this vertex must have the same adjacencies as v with respect to the wheel, and hence is a valid replacement. Finally if $|D'| > 1$, then any two vertices in D' together with p, q, and r induce a W_4, since p, q, and r induce a P_3. □

Every graph that contains a wheel contains at least one edge. For every wheel W_n, we have $|V(W_n)| \leq 2 \cdot \text{vc}(W_n)$. Therefore by Theorem 3.5 we obtain:

Theorem 4.25. Wheel-free Deletion Parameterized by Vertex Cover *admits a polynomial kernel with* $\mathcal{O}(|X|^5)$ *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 4.26. $\Pi_{W_{\neq 4}}$ *is not characterized by rank-c adjacencies for any* $c \in \mathbb{N}$.

Proof. Suppose for the sake of contradiction that $\Pi_{W_{\neq 4}}$ is characterized by rank-c adjacencies for some $c \in \mathbb{N}$. Construct a graph H as follows; see Fig. 4 for an illustration. Let $i > 2$ be the smallest integer such that $q = 2^{i-1} - 1 > c$. Vertex cover $X = (v_1, ..., v_n)$ of H is an induced cycle of length $n = 2^i - 1$. Furthermore H contains a vertex $v \notin X$ that is adjacent to all vertices v_j for $j \in [n]$. Finally H contains an independent set D disjoint from $X \cup \{v\}$. The adjacencies of D to X are defined as follows. Start with an empty set D. For each $Y \in \binom{[n]}{q}$, add a vertex d_Y to D. For $j \in [n]$ vertex d_Y is adjacent to v_j if $j \in Y$.

By Lucas' Theorem [17], we have that $\binom{m}{k}$ is divisible by prime 2, if at least one of the base 2 digits of k is larger than the respective base 2 digit of m. Since the binary representation of n and q consists only of ones, we have that $\binom{n-j}{q-j}$ is odd for all $0 \leq j \leq q$. Furthermore, for all $0 \leq j < k \leq q$, we have that $\binom{n-k}{q-j}$ is even.

For disjoint $Q, R \subseteq X$ such that $|Q| + |R| \leq c$, there are $m_{Q,R} = \binom{n-|R|-|Q|}{q-|R|}$ vertices in D adjacent to each vertex in R and to none in Q. By our choices of n and q, we have that $m_{Q,R}$ is odd if and only if $|Q| = 0$.

Clearly $H - D$ induces a wheel of size not equal to 4 with center v and cycle $(v_1, \ldots, v_n)$, as $i > 2$.

Claim 4.27. For disjoint $Q, R \subseteq X$ such that $|Q| + |R| \leq c$, we have $\text{inc}^c_{(H,X)}(v)[Q, R] = \sum_{u \in D} \text{inc}^c_{(H,X)}(u)[Q, R]$.

Proof. We distinguish two cases, depending on whether Q is empty or not. Note that by construction of H, there are exactly $m_{Q,R}$ vertices of D which are adjacent to all of R and none of Q, so there are exactly $m_{Q,R}$ vertices of $u \in D$ for which $\text{inc}^c_{(H,X)}(u)[Q, R] = 1$.

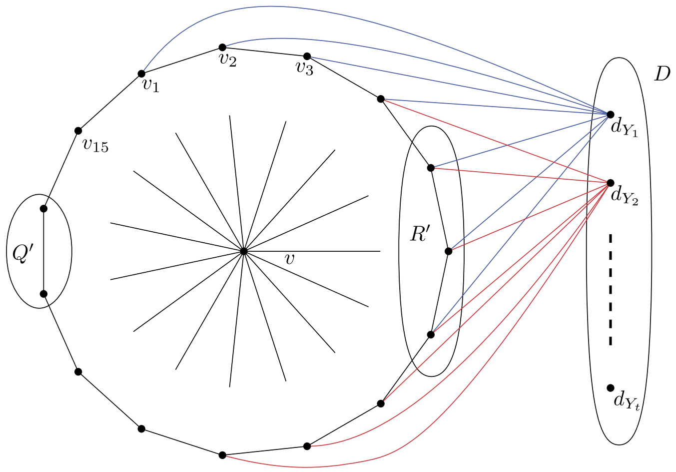

Fig. 4. Sketch of the construction in Theorem 4.26 for $i = 4$. The set D contains a vertex for each choice of 7 vertices of the cycle. $\binom{15-3}{7-3} = 495$ vertices of D are adjacent to all of R', while $\binom{15-3-2}{7-3} = 210$ vertices of D are adjacent to all of R' and none of Q'. Note that the middle vertex v_6 of R' is the center of a wheel W_4 containing $v_5, d_{Y_1}, v_7, d_{Y_2}$, but that no vertex v_i is the center of a wheel of length unequal to 4.

If $Q \neq \emptyset$, then $\text{inc}^c_{(H,X)}(v)[Q, R] = 0$ since v is adjacent to all vertices of X and therefore to a member of Q. Since $|Q| > 0$, the value $m_{Q,R}$ is even as observed above, and therefore zero over $\mathbb{F}_2$. The claim follows.

If $Q = \emptyset$, then $\text{inc}^c_{(H,X)}(v)[Q, R] = 1$ since v is adjacent to all vertices of X and therefore to all members of R and none of Q. As $|Q| = 0$, the value $m_{Q,R}$ is odd as observed above, so $\sum_{u \in D} \text{inc}^c_{(H,X)}(u)[Q, R] = 1$. $\square$

The claim above shows that $\text{inc}^c_{(H,X)}(v) = \sum_{u \in D} \text{inc}^c_{(H,X)}(u)$. From our assumption that $\Pi_{W \neq 4}$ is characterized by c adjacencies, it follows that there exists $D' \subseteq D$ such that $H - v - (D \setminus D') \in \Pi_{W \neq 4}$ and therefore contains a wheel of size unequal to 4. To derive the desired contradiction, we therefore argue that $H - v$ contains no such wheel. Observe that no vertex $v \in D$ can be used as the center of a wheel, since $N_H(v)$ is a proper subset of the induced cycle $(v_1, \ldots, v_n)$ and therefore acyclic; we use here that $n > q > c$. Finally for each $j \in [n]$, v_j can only be the center of a wheel of size 4: the graph $N_H(v_j)$ has a vertex cover $\{v_{j-1}, v_{j+1}\}$ of size two and therefore cannot contain a cycle of length five or more. Furthermore, $N_H(v_j)$ does not contain a triangle since v_{j-1}, v_{j+1} are non-adjacent and D is an independent set. Hence there exists no $D' \subseteq D$ for which $H - v - (D \setminus D') \in \Pi_{W \neq 4}$, a contradiction. $\square$

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 coNP $\subseteq$ NP/poly. This suggests that the condition of being characterized by low-rank adjacencies is the right way to capture kernelization complexity. In order to show this lower bound, we introduce some terminology.

Definition 4.28. A *polynomial compression* of a parameterized problem $Q \subseteq \Sigma^* \times \mathbb{N}$ into a language $R \subseteq \Sigma^*$ is an algorithm that takes as input an instance $(x, k) \in \Sigma^* \times \mathbb{N}$ and produces in time polynomial in $|x| + k$ a string y such that $|y| \leq p(k)$ for some polynomial p and $y \in R$ if and only if $(x, k) \in Q$.

Ruling out the existence of a polynomial compression also rules out the existence of a polynomial kernel. We use this fact to construct lower-bound proofs. For these proofs we need one more ingredient, namely polynomial parameter transformations.

Definition 4.29. Let $P, Q \subseteq \Sigma^* \times \mathbb{N}$ be two parameterized problems. An algorithm $\mathcal{A}$ is called a *polynomial parameter transformation* (PPT) from P to Q if, given an instance (x, k) of problem P, $\mathcal{A}$ produces in polynomial time an instance (x', k') of problem Q such that:

1. $(x, k) \in P$ if and only if $(x', k') \in Q$, and
2. $k' \leq p(k)$ for some fixed polynomial p.

Note that from the definition above we only require the size of the parameter in the transformed instance to be bounded in some polynomial of the original parameter. The following theorem combines this definition with polynomial compressions.

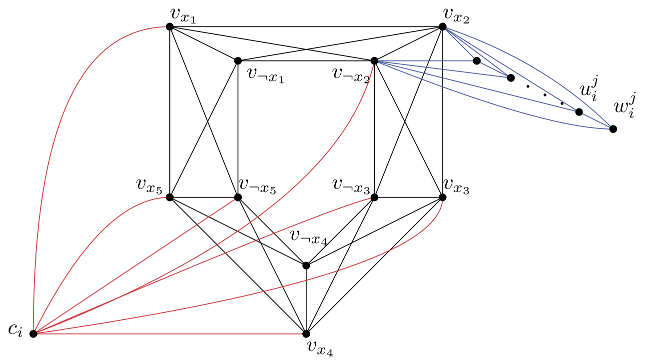

Fig. 5. Sketch of transformation for $n = 5$ variables. For simplicity the u and w vertices are only (partially) drawn for x_2. The clause vertex c_i corresponds to the clause $x_1 \vee \neg x_2 \vee x_4$.

Theorem 4.30 *(Theorem 19.2, [35]). Let P and Q be two parameterized problems such that there exists a PPT from P to Q. If Q has a polynomial compression, then P also has a polynomial compression.*

By contraposition of the theorem above, if the starting problem P does not have a polynomial compression, then Q also does not have a polynomial compression. This fact can be used to rule out polynomial kernels for parameterized problems. The starting problem for our lower bounds is CNF-SAT parameterized by the number of variables.

Theorem 4.31 *(Theorem 18.3, [35]).* CNF-SAT *parameterized by the number of variables n does not admit a polynomial compression unless* coNP $\subseteq$ NP/poly.

With the necessary definitions in place, we prove the following.

Theorem 4.32. Almost Wheel-free Deletion Parameterized by Vertex Cover *does not admit a polynomial compression unless* coNP $\subseteq$ NP/poly.

Proof. We give a polynomial parameter transformation (PPT) from the CNF-SAT problem parameterized by the number of variables. Consider an instance ϕ of the CNF-SAT problem with variables $\{x_1, ..., x_n\}$ and clauses $\{C_1, ..., C_m\}$. In the case that $n \leq 4$, try all possible truth assignments in constant time. If ϕ is satisfiable, return an empty graph with an empty vertex cover, which is obviously almost-wheel free. If ϕ is not satisfiable, return $G = W_3$ with vertex cover $X = V(G)$ and budget $k = 0$. The vertex cover is bounded by a polynomial of n and no vertex can be deleted to make the graph almost wheel-free. In the remainder we consider $n \geq 5$. We construct an instance (G, X, k) of Almost Wheel-free Deletion with a vertex cover X and solution size k as follows.

1. Starting with an empty graph G, for $i \in [n]$, add vertices v_{x_i} and $v_{\neg x_i}$ to G. Connect them with an edge.
2. For $i \in [n]$ add edges $\{v_{x_i}, v_{x_{i+1}}\}$, $\{v_{x_i}, v_{\neg x_{i+1}}\}$, $\{v_{\neg x_i}, v_{x_{i+1}}\}$ and $\{v_{\neg x_i}, v_{\neg x_{i+1}}\}$. Here $i + 1 = 1$ if $i = n$.
3. For $i \in [n]$, $j \in [n + 1]$, add vertices u_i^j and w_i^j. Add edges $\{u_i^j, w_i^j\}$, $\{u_i^j, v_{x_i}\}$, $\{u_i^j, v_{\neg x_i}\}$, $\{w_i^j, v_{x_i}\}$, and $\{w_i^j, v_{\neg x_i}\}$.
4. Finally consider clause C_i, $i \in [m]$. Add vertex c_i to G. For every literal $\ell \in C_i$, connect c_i to v_ℓ. Here we assume no clause contains both x_q and $\neg x_q$ for $q \in [n]$. Such clauses are trivially satisfied and can be removed in polynomial time without changing the problem. For every variable x_q, $q \in [n]$, that does not correspond to a literal in C_i, connect c_i to both v_{x_q} and $v_{\neg x_q}$.

This concludes the construction of G. Fig. 5 shows a sketch of this construction for the clause $x_1 \vee \neg x_2 \vee x_4$. Assign the budget $k = n$.

Note. Vertices added in steps 1-3 form a vertex cover of size $2n + 2n(n + 1) = 2n^2 + 4n$.

Let X be the vertices added in steps 1-3 above. Since $|X| \leq p(n)$ for some polynomial $p(\cdot)$, all that is left to show is the equivalence between the original instance and the vertex-deletion problem (G, X, k). Let $\mathcal{F}$ be the set of almost wheel-free graphs. We show that there is a set $S \subseteq V(G)$ of size at most k such that $G - S$ belongs to $\mathcal{F}$ if and only if the CNF-SAT instance ϕ is satisfiable.

($\Rightarrow$) Suppose (G, X, k) is a yes-instance with solution S of size at most k such that $G - S$ belongs to $\mathcal{F}$.

Claim 4.33. For $i \in [n]$, S contains exactly one of v_{x_i} and $v_{\neg x_i}$.

Proof. Consider some $i \in [n]$. Suppose S contains neither v_{x_i} nor $v_{\neg x_i}$. For $j \in [n+1]$, the set of vertices $\{v_{x_i}, v_{\neg x_i}, u_i^j, w_i^j\}$ induce a W_3. Since $|S| \leq k$, it follows that $G - S$ would still contain a W_3 which contradicts the choice of S. It follows that S contains at least one of v_{x_i} and $v_{\neg x_i}$. Since the budget $k = n$, S must contain exactly one of v_{x_i} and $v_{\neg x_i}$. □

Consider the following truth assignment $\delta : \{x_1, ..., x_n\} \to \{0, 1\}$ such that $\delta(x_i) = 1$ if $v_{x_i} \in S$, and $\delta(x_i) = 0$ if $v_{\neg x_i} \in S$. This is well defined by the claim above.

Claim 4.34. δ satisfies ϕ.

Proof. Suppose δ does not satisfy ϕ, then there exists some $i \in [m]$ such that δ does not satisfy clause C_i. Consider the set of vertices $S' = \{v_{\neg \ell_p} \mid v_{\ell_p} \in S, p \in [n], (\ell_p = x_p \vee \ell_p = \neg x_p)\}$. These are the vertices that correspond to the inverse of the truth assignment. The vertices of S' are contained in $G - S$ and induce a cycle of length $n \geq 5$. Since C_i is not satisfied, for each literal $\ell \in C_i$ we have $v_\ell \in S'$. This means vertex c_i is adjacent to vertex v_ℓ. Now for every variable x_p, $p \in [n]$ without a literal in C_i, c_i is adjacent to both v_{x_p} and $v_{\neg x_p}$. It follows that c_i is adjacent to every vertex in S'. Hence $S' \cup \{c_i\}$ induces a wheel W_n. Since $n \geq 5$, this wheel is not of size 4 and hence forbidden, which contradicts the choice of S. It follows that δ satisfies ϕ. □

($\Leftarrow$) In the other direction, assume that there exists a truth assignment $\delta : \{x_1, ..., x_n\} \to \{0, 1\}$ that satisfies ϕ. Let $S = \{v_{\neg x_i} \mid \delta(x_i) = 0, i \in [n]\} \cup \{v_{x_i} \mid \delta(x_i) = 1, i \in [n]\}$. Obviously $|S| = n$. We show that $G - S$ belongs to $\mathcal{F}$. Let $S' = \{v_{\neg x_i} \mid \delta(x_i) = 1, i \in [n]\} \cup \{v_{x_i} \mid \delta(x_i) = 0, i \in [n]\}$, which are the vertices that correspond to the inverse of the truth assignment. Again S' forms an induced cycle. $G - S$ is a graph consisting of S' with an independent set of clause vertices c_i for $i \in [m]$ adjacent to vertices in S'. Finally each vertex $v_i \in S'$ is part of $n + 1$ edge disjoint triangles $\{v_i, u_i^j, w_i^j\}$ for $j \in [n + 1]$. Now that we know the structure of $G - S$, we do a case distinction on vertex $v \in V(G - S)$. We show that v cannot be a center vertex of a wheel of size 3 or at least 5.

- Case $v \in S'$. Since S' is an induced cycle, v has two neighbors in S', say x and y, that are not adjacent. Since the u and w vertices have degree 2 in $G - S$, they cannot be used to create an induced cycle around v. Since clause vertices c_i, for $i \in [m]$, form an independent set, they cannot appear adjacent in an induced cycle around v. Hence, the only wheel that could exist around v is $\{v, x, y, c_i, c_j\}$ for some $i, j \in [m]$ such that c_i and c_j are adjacent to v, x, and y, but this is W_4 which is not forbidden.
- Case $v = c_i$, for some $i \in [m]$. Since c_i only has neighbors in S', and S' is an induced cycle C_n, the only way c_i can be a center vertex of an induced wheel is if it is adjacent to every vertex in S'. Suppose that c_i is adjacent to every vertex in S', then no literal in clause C_i is satisfied by δ, which contradicts the choice of δ. Therefore c_i cannot be a center vertex of a wheel.
- Case $v = u_i^j$ or $v = w_i^j$ for some $i \in [n]$, $j \in [n + 1]$. Since v only has two neighbors, it cannot be a center vertex of a wheel.

Since $G - S$ does not contain a vertex that could be a center of some wheel W_p for $p = 3$ or $p \geq 5$, it follows that $G - S$ must belong to $\mathcal{F}$. Since we have shown correctness of the PPT, the result follows from Theorem 4.30. □

5. Conclusion

We have presented a framework that can be used to obtain polynomial kernels for the Π-free Deletion problem parameterized by 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. [18]. 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.

Corollary 5.1. *Let $\mathcal{F}$ be a hereditary graph class formed as the intersection of a finite nonempty subset of the following graph classes:*

1. *being wheel-free,*
2. *being odd-hole-free,*
3. *being odd-anti-hole-free,*
4. *being even-hole-free,*
5. *being AT-free,*
6. *being comparability,*
7. *being co-comparability,*
8. *being bipartite,*
9. *being $C_{\geq c}$-free for some fixed $c \geq 3$, (i.e., not containing a chordless cycle of length at least c),*
10. *being H-minor-free for some fixed graph H containing at least one edge, and*
11. *being H-free for some fixed graph H containing at least one edge.*

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.[1]

Proof. For each graph class $\mathcal{F}_i$ in the list, there exists a graph property Π_i such that:

- a graph G belongs to $\mathcal{F}_i$ if and only if it does not contain an induced subgraph belonging to Π_i,
- Π_i is characterized by rank-c_i adjacencies for some constant c_i,
- each graph in Π_i contains at least one edge, and
- there is a non-decreasing polynomial $p_i \colon \mathbb{N} \to \mathbb{N}$ such that all graphs G that are vertex-minimal with respect to Π_i satisfy $|V(G)| \leq p_i(\text{vc}(G))$.

For example, for $\mathcal{F}_i$ the class of wheel-free graphs, the set Π_i is the set of graphs containing a wheel. For the first seven graph classes, the existence of such a Π_i that is characterized by low-rank adjacencies follows from Theorem 4.24, Corollary 4.3, Theorem 4.6, Theorem 4.11, and Corollary 4.22 in this paper; the remaining conditions on Π_i are easy to verify and were derived in the proofs of Theorem 4.25, Theorem 4.4, Theorem 4.7, Theorem 4.12, and Theorem 4.23. For the remaining graph classes on the list except for the last one, the existence of such a Π_i follows from characterizations by few adjacencies due to Fomin et al. [18, Proposition 1] which imply characterizations by low-rank adjacencies via Lemma 3.11; the remaining conditions on Π_i were verified in [18, Corollary 1]. The characterization by low-rank adjacencies for the last property follows from Proposition 2.3 via Lemma 3.11. Since H is a finite graph containing at least one edge, the remaining properties on Π_i trivially hold in that case.

Now, for any hereditary graph class $\mathcal{F} = \mathcal{F}_1 \cap \ldots \cap \mathcal{F}_\ell$ that can be formed as the nonempty intersection of a finite number of classes from the corollary statement, consider the corresponding sets $\Pi_1, \ldots, \Pi_\ell$ of forbidden induced subgraphs. Since each Π_i is characterized by rank-c_i adjacencies for some $c_i \in \mathbb{N}$, by Lemma 3.9 the property $\Pi = \Pi_1 \cup \ldots \cup \Pi_\ell$ is characterized by adjacencies of rank $\max_{i \in [\ell]} c_i$. Then Π satisfies all requirements of the kernelization framework of Theorem 3.5, and since $G \in \mathcal{F}$ if and only if G has no induced subgraph belonging to Π, the result follows. □

To apply our kernelization framework for Π-free Deletion, one has to establish that Π is characterized by rank-c adjacencies for some $c \in \mathbb{N}$. We have presented several examples of Π for which this holds. Although Lemma 3.9 shows how to obtain a characterization of the union of two properties for which characterizations are known, in general it is a challenging task to determine whether a given property is characterized by low-rank adjacencies; this has to be done in a case-by-case basis. When searching for such a characterization for a graph property Π whose vertex-minimal subgraphs belong to several types, it is sometimes convenient to argue that modifying adjacencies of one type of subgraph, creates an occurrence of a different type of subgraph. In particular, this implicitly happens in the proof of Lemma 4.21. One concrete example of a hereditary graph class for which it is currently unknown whether it can be characterized by low-rank adjacencies, is the class of sun-free graphs [5, Section 7.2.2].

With respect to the size of our kernelizations, it would be interesting to see whether the exponents given by Table 1 are tight (cf. [21]).

It is natural to ask whether our positive kernelization results for parameterizations by vertex cover can be generalized to parameters which are less restrictive. As mentioned in the introduction, known lower bounds [13,25] rule out extending the results to parameterizations by the size of a given feedback vertex set. Since the vertex cover number equals the vertex-deletion distance to a graph of treewidth 0, while the feedback vertex number equals the vertex-deletion distance to a graph of treewidth one, these lower bounds show that relaxing the parameterization based on its relation to treewidth is infeasible. However, there is an analogous hierarchy of parameterizations based on *treedepth* (cf. [33, §6]) which may be worth investigating. The vertex cover number of a graph can also be defined as its vertex-deletion distance to a graph of treedepth 0. For each constant $\eta > 0$, one can obtain a parameterization smaller than the vertex cover number by measuring the vertex-deletion number to a graph of treedepth η. For all problems which can be formulated as hitting all minor models of a fixed set $\mathcal{F}$ of connected graphs, the parameterization by the size of a vertex-deletion set to treedepth η admits a polynomial kernel [27]. Investigating whether the existence of polynomial kernels extends to all vertex deletion properties that admit a characterization by low-rank adjacencies is an interesting direction for future work.

CRediT authorship contribution statement

Bart M.P. Jansen: Conceptualization, Formal analysis, Funding acquisition, Writing – review & editing. **Jari J.H. de Kroon:** Formal analysis, Investigation, Writing – original draft.

[1] This corrects a statement in the extended abstract of our work where 'having a Hamiltonian cycle (respectively, path)' was incorrectly included in the list. Whereas this can be characterized by low-rank adjacencies, the proof requires the set of forbidden induced subgraphs for the property to have such a characterization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We would like to thank Fedor V. Fomin for hosting Jari in Bergen (Norway).

References

[1] Akanksha Agrawal, Pranabendu Misra, Saket Saurabh, Meirav Zehavi, Interval vertex deletion admits a polynomial kernel, in: Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '19, 2019, pp. 1711–1730, http://dl.acm.org/citation.cfm?id=3310435.3310538.
[2] Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, Dimitrios M. Thilikos, (Meta) kernelization, J. ACM 63 (5) (2016) 44, https://doi.org/10.1145/2973749.
[3] H.L. Bodlaender, B.M.P. Jansen, S. Kratsch, Kernelization lower bounds by cross-composition, SIAM J. Discrete Math. 28 (1) (2014) 277–305, https://doi.org/10.1137/120880240.
[4] Marin Bougeret, Ignasi Sau, How much does a treedepth modulator help to obtain polynomial kernels beyond sparse graphs?, Algorithmica 81 (10) (2019) 4043–4068, https://doi.org/10.1007/s00453-018-0468-8.
[5] A. Brandstädt, V. Le, J. Spinrad, Graph Classes: A Survey, Society for Industrial and Applied Mathematics, 1999.
[6] Kevin Burrage, Vladimir Estivill-Castro, Michael R. Fellows, Michael A. Langston, Shev Mac, Frances A. Rosamond, The undirected feedback vertex set problem has a poly(k) kernel, in: Hans L. Bodlaender, Michael A. Langston (Eds.), Proceedings of the 2nd Second International Workshop on Parameterized and Exact Computation, IWPEC 2006, in: Lecture Notes in Computer Science, vol. 4169, Springer, 2006, pp. 192–202.
[7] Yixin Cao, Ashutosh Rai, R.B. Sandeep, Junjie Ye, A polynomial kernel for diamond-free editing, in: Yossi Azar, Hannah Bast, Grzegorz Herman (Eds.), Proceedings of the 26th Annual European Symposium on Algorithms, ESA 2018, in: LIPIcs, vol. 112, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018, pp. 10:1–10:13.
[8] Maria Chudnovsky, Neil Robertson, Paul Seymour, Robin Thomas, The strong perfect graph theorem, Ann. Math. (2) 164 (1) (2006) 51–229, https://doi.org/10.4007/annals.2006.164.51.
[9] Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, Saket Saurabh, Parameterized Algorithms, Springer, 2015.
[10] Marek Cygan, Daniel Lokshtanov, Marcin Pilipczuk, Michal Pilipczuk, Saket Saurabh, On cutwidth parameterized by vertex cover, Algorithmica 68 (4) (2014) 940–953, https://doi.org/10.1007/s00453-012-9707-6.
[11] Marek Cygan, Marcin Pilipczuk, Michal Pilipczuk, Erik Jan van Leeuwen, Marcin Wrochna, Polynomial kernelization for removing induced claws and diamonds, Theory Comput. Syst. 60 (4) (2017) 615–636, https://doi.org/10.1007/s00224-016-9689-x.
[12] Jari J.H. de Kroon, Polynomial kernels for structural parameterizations of deletion into perfect graphs and related graph classes, Master's thesis, Eindhoven University of Technology, The Netherlands, 2019, https://research.tue.nl/files/139492673/0888168_thesis_final_19_07_Jari_de_Kroon.pdf.
[13] Huib Donkers, Bart M.P. Jansen, A Turing kernelization dichotomy for structural parameterizations of F-minor-free deletion, J. Comput. Syst. Sci. 119 (2021) 164–182, https://doi.org/10.1016/j.jcss.2021.02.005.
[14] Rodney G. Downey, Michael R. Fellows, Fundamentals of Parameterized Complexity, Texts in Computer Science, Springer, 2013.
[15] Ben Dushnik, E.W. Miller, Partially ordered sets, Am. J. Math. 63 (3) (1941) 600–610, http://www.jstor.org/stable/2371374.
[16] Michael R. Fellows, Bart M.P. Jansen, Frances A. Rosamond, Towards fully multivariate algorithmics: parameter ecology and the deconstruction of computational complexity, Eur. J. Comb. 34 (3) (2013) 541–566, https://doi.org/10.1016/j.ejc.2012.04.008.
[17] N.J. Fine, Binomial coefficients modulo a prime, Am. Math. Mon. 54 (10) (1947) 589–592, http://www.jstor.org/stable/2304500.
[18] Fedor V. Fomin, Bart M.P. Jansen, Michał Pilipczuk, Preprocessing subgraph and minor problems: when does a small vertex cover help?, J. Comput. Syst. Sci. 80 (2) (2014) 468–495, https://doi.org/10.1016/j.jcss.2013.09.004.
[19] Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, 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, IEEE Computer Society, 2012, pp. 470–479.
[20] T. Gallai, Transitiv orientierbare graphen, Acta Math. Acad. Sci. Hung. 18 (1) (1967) 25–66, https://doi.org/10.1007/BF02020961.
[21] Archontia C. Giannopoulou, Bart M.P. Jansen, Daniel Lokshtanov, Saket Saurabh, Uniform kernelization complexity of hitting forbidden minors, ACM Trans. Algorithms 13 (3) (2017) 35, https://doi.org/10.1145/3029051.
[22] Sepp Hartung, Christian Komusiewicz, André Nichterlein, Parameterized algorithmics and computational experiments for finding 2-clubs, J. Graph Algorithms Appl. 19 (1) (2015) 155–190, https://doi.org/10.7155/jgaa.00352.
[23] Pinar Heggernes, Pim Van 't Hof, Bart M.P. Jansen, Stefan Kratsch, Yngve Villanger, Parameterized complexity of vertex deletion into perfect graph classes, Theor. Comput. Sci. 511 (2013) 172–180, https://doi.org/10.1016/j.tcs.2012.03.013.
[24] Yoichi Iwata, Linear-time kernelization for feedback vertex set, in: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, Anca Muscholl (Eds.), Proceedings of the 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, in: LIPIcs, vol. 80, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017, pp. 68:1–68:14.
[25] Bart M.P. Jansen, Stefan Kratsch, Data reduction for graph coloring problems, Inf. Comput. 231 (2013) 70–88, https://doi.org/10.1016/j.ic.2013.08.005.
[26] Bart M.P. Jansen, Astrid Pieterse, Optimal data reduction for graph coloring using low-degree polynomials, Algorithmica 81 (10) (2019) 3865–3889, https://doi.org/10.1007/s00453-019-00578-5.
[27] Bart M.P. Jansen, Astrid Pieterse, Polynomial kernels for hitting forbidden minors under structural parameterizations, Theor. Comput. Sci. 841 (2020) 124–166, https://doi.org/10.1016/j.tcs.2020.07.009.
[28] Bart M.P. Jansen, Marcin Pilipczuk, Approximation and kernelization for chordal vertex deletion, SIAM J. Discrete Math. 32 (3) (2018) 2258–2301, https://doi.org/10.1137/17M112035X.
[29] Ekkehard Köhler, Graphs without asteroidal triples, PhD thesis, TU, Berlin, 1999.
[30] Stefan Kratsch, Magnus Wahlström, Compression via matroids: a randomized polynomial kernel for odd cycle transversal, ACM Trans. Algorithms 10 (4) (2014) 20, https://doi.org/10.1145/2635810.
[31] John M. Lewis, Mihalis Yannakakis, The node-deletion problem for hereditary properties is NP-complete, J. Comput. Syst. Sci. 20 (2) (1980) 219–230, https://doi.org/10.1016/0022-0000(80)90060-4.
[32] Daniel Lokshtanov, Wheel-free deletion is W[2]-hard, in: Martin Grohe, Rolf Niedermeier (Eds.), Proceedings of the Third International Workshop on Parameterized and Exact Computation, IWPEC 2008, in: Lecture Notes in Computer Science, vol. 5018, Springer, 2008, pp. 141–147.
[33] Jaroslav Nešetřil, Patrice Ossona de Mendez, Sparsity - Graphs, Structures, and Algorithms, Algorithms and Combinatorics, vol. 28, Springer, 2012.

[34] Johannes Uhlmann, Mathias Weller, Two-layer planarization parameterized by feedback edge set, Theor. Comput. Sci. 494 (2013) 99–111, https://doi.org/10.1016/j.tcs.2013.01.029.

[35] Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, Meirav Zehavi, Kernelization: Theory of Parameterized Preprocessing, Cambridge University Press, 2019.

[36] René van Bevern, Till Fluschnik, Oxana Yu. Tsidulko, Parameterized algorithms and data reduction for the short secluded s-t-path problem, Networks 75 (1) (2020) 34–63, https://doi.org/10.1002/net.21904.