Vertex Cover Kernelization Revisited: Upper and Lower Bounds for a Refined Parameter

Bart Jansen, Hans Bodlaender

Vertex Cover Kernelization Revisited: Upper and Lower Bounds for a Refined Parameter*

Bart M. P. Jansen[1] and Hans L. Bodlaender[1]

1 **Department of Information and Computing Sciences, Utrecht University**
P.O. Box 80.089, 3508 TB, Utrecht, The Netherlands
`{bart,hansb}@cs.uu.nl`

Abstract

Kernelization is a concept that enables the formal mathematical analysis of data reduction through the framework of parameterized complexity. Intensive research into the Vertex Cover problem has shown that there is a preprocessing algorithm which given an instance (G, k) of Vertex Cover outputs an equivalent instance (G', k') in polynomial time with the guarantee that G' has at most $2k'$ vertices (and thus $O((k')^2)$ edges) with $k' \leq k$. Using the terminology of parameterized complexity we say that k-Vertex Cover has a kernel with $2k$ vertices. There is complexity-theoretic evidence that both $2k$ vertices and $\Theta(k^2)$ edges are optimal for the kernel size. In this paper we consider the Vertex Cover problem with a different parameter, the size $\textsc{fvs}(G)$ of a minimum feedback vertex set for G. This refined parameter is structurally smaller than the parameter k associated to the vertex covering number $\textsc{vc}(G)$ since $\textsc{fvs}(G) \leq \textsc{vc}(G)$ and the difference can be arbitrarily large. We give a kernel for Vertex Cover with a number of vertices that is cubic in $\textsc{fvs}(G)$: an instance (G, X, k) of Vertex Cover, where X is a feedback vertex set for G, can be transformed in polynomial time into an equivalent instance (G', X', k') such that $k' \leq k$, $|X'| \leq |X|$ and most importantly $|V(G')| \leq 2k$ and $|V(G')| \in O(|X'|^3)$. A similar result holds when the feedback vertex set X is not given along with the input. In sharp contrast we show that the Weighted Vertex Cover problem does not have a polynomial kernel when parameterized by $\textsc{fvs}(G)$ unless the polynomial hierarchy collapses to the third level $(\mathrm{PH} = \Sigma_3^p)$. Our work is one of the first examples of research in kernelization using a non-standard parameter, and shows that this approach can yield interesting computational insights. To obtain our results we make extensive use of the combinatorial structure of independent sets in forests.

1998 ACM Subject Classification F.2.2

Keywords and phrases kernelization, lower bounds, parameterized complexity

Digital Object Identifier 10.4230/LIPIcs.STACS.2011.177

1 Introduction

The Vertex Cover problem is one of the six classic NP-complete problems discussed by Garey and Johnson in their famous work on intractability [22, GT1], and has played an important role in the development of parameterized algorithms [15, 28, 16]. A parameterized problem is a language $L \subseteq \Sigma^* \times \mathbb{N}$, and such a problem is (strongly uniform) *fixed parameter tractable* if membership of an instance (x, k) can be decided in $f(k)|x|^c$ time for some computable function f and constant c. Since the structure of Vertex Cover is so simple and elegant, it has proven to be an ideal testbed for new techniques in the context of

* This work was supported by the Netherlands Organisation for Scientific Research (NWO), project "KERNELS: Combinatorial Analysis of Data Reduction".

parameterized complexity. The problem is also highly relevant from a practical point of view because of its role in bioinformatics [1] and other problem areas.

In this work we suggest a "refined parameterization" for the Vertex Cover problem using the feedback vertex number $\text{FVS}(G)$ as the parameter, i.e. the size of a smallest vertex set whose deletion turns G into a forest. We give upper bounds on the kernel size for the unweighted version of Vertex Cover under this parameterization, and also supply a conditional superpolynomial lower bound on the kernel size for the variant of Vertex Cover where each vertex has a non-negative integral weight. But before we state our results we shall first survey the current state of the art for the parameterized analysis of Vertex Cover.

There has been an impressive series of ever-faster parameterized algorithms to solve k-Vertex Cover, which led to the current-best algorithm by Chen et al. that can decide whether a graph G has a vertex cover of size k in $O(1.2738^k + kn)$ time and polynomial space [9, 30, 8, 17]. The Vertex Cover problem has also played an important role in the development of *problem kernelization* [23]. A kernelization algorithm (or *kernel*) is a polynomial-time procedure that reduces an instance (x, k) of a parameterized decision problem to an equivalent instance (x', k') such that $|x'|, k' \leq f(k)$ for some computable function f, which is the *size* of the kernel. We also use the term kernel to refer to the reduced instance (x', k').

The k-Vertex Cover problem admits a kernel with $2k$ vertices and $O(k^2)$ edges, which has been a subject of repeated study [6, 8, 10, 2, 11] and experimentation [1, 13]. There is some complexity-theoretic evidence that the size bounds for the kernel cannot be improved. Since practically all reduction-rules found to date are approximation-preserving [28], it appears that a kernel with less than $2k$ vertices would yield a polynomial-time approximation algorithm with a performance ratio smaller than 2 which would disprove the Unique Games Conjecture [25]. A recent breakthrough result by Dell and Van Melkebeek [12] shows that there is no polynomial kernel which can be encoded into $O(k^{2-\epsilon})$ bits for any $\epsilon > 0$ unless the polynomial hierarchy collapses to the third level ($\text{PH} = \Sigma_3^p$), which suggests that the current bound of $O(k^2)$ edges is tight up to logarithmic factors.

This overview might suggest that there is little left to explore concerning kernelization for vertex cover, but this is far from true. All existing kernelization results for Vertex Cover use the requested size k of the vertex cover as the parameter. But there is no reason why we should not consider structurally smaller parameters, to see if we can preprocess instances of Vertex Cover such that their final size is bounded by a function of such a smaller parameter, rather than by a function of the requested set size k. We study kernelization for the Vertex Cover problem using the feedback vertex number $\text{FVS}(G)$ as the parameter. Since every vertex cover is also a feedback vertex set we find that $\text{FVS}(G) \leq \text{VC}(G)$ which shows that the feedback vertex number of a graph is a *structurally smaller* parameter than the vertex covering number: there are trees with arbitrarily large values of $\text{VC}(G)$ for which $\text{FVS}(G) = 0$. We call our parameter "refined" since it is structurally smaller than the standard parameter for the Vertex Cover problem.

Related Work. The idea of studying parameterized problems using alternative parameters is not new (see e.g. [28]), but was recently advocated by Fellows et al. [19, 20, 29] in the call to investigate the *complexity ecology* of parameters. The main idea behind this program is to determine how different parameters affect the parameterized complexity of a problem. Some recent results in this direction include FPT algorithms for graph layout problems parameterized by the vertex cover number of the graph [21] and an algorithm to decide isomorphism on graphs of bounded feedback vertex number [26]. We are aware of only two applications of this idea to give polynomial kernels using alternative parameters. Fellows et al. [20, 18] show that the problems Independent Set, Dominating Set and Hamiltonian Circuit admit linear-vertex kernels on graphs G when parameterized by the maximum number of leaves in any spanning tree of G. Very recently Uhlmann and Weller [31] gave a polynomial kernel for Two-Layer Planarization parameterized by the feedback edge set number, which is a refined structural parameter for that problem since it is smaller than the natural parameter.

Our Results. We believe that we are one of the first to present a polynomial problem kernel using a non-standard but practically relevant refined parameter. We study the following parameterized problem:

fvs-Weighted Vertex Cover
Instance: A simple undirected graph G, a weight function $w : V(G) \to \mathbb{N}^+$, a feedback vertex set $X \subseteq V(G)$ such that $G - X$ is a forest, an integer $k \geq 0$.
Parameter: The size $|X|$ of the feedback vertex set.
Question: Is there a vertex cover C of G such that $\sum_{v \in C} w(v) \leq k$?

We also consider the unweighted variant fvs-Vertex Cover in which all vertices have a weight of 1. The problems fvs-Weighted Independent Set and fvs-Independent Set are defined similarly. Throughout this work k will always represent the total size or weight of the set we are looking for; depending on the context this is either a vertex cover or an independent set.

We prove that fvs-Vertex Cover has a kernel in which the number of vertices is bounded by $\min(O(|X|^3), 2k)$. This bound is at least as small as the current-best Vertex Cover kernel, but for graphs with small feedback vertex sets our bound is significantly smaller. We also study the weighted version of the problem, and obtain a contrasting result: we show that fvs-Weighted Vertex Cover does not admit a polynomial kernel unless PH $= \Sigma_3^p$. This is very surprising since both the weighted and unweighted versions of k-Vertex Cover admit polynomial kernels and can be attacked using similar reduction rules [10]. To our knowledge we give the first example of a parameterized problem whose weighted and unweighted versions are both NP-complete and FPT, but for which the unweighted version allows a polynomial kernel but the weighted version does not.

When we present our results we will state them in terms of fvs-Independent Set and fvs-Weighted Independent Set since this simplifies the exposition. Because we are using the size of a feedback vertex set as the parameter, there are trivial parameterized reductions between these problems: an instance (G, X, k) of fvs-Vertex Cover is equivalent to an instance $(G, X, |V(G)| - k)$ of fvs-Independent Set with the same parameter value $|X|$. Hence our kernelization bounds for Independent Set carry over to Vertex Cover.

2 Preliminaries

In this work we only consider undirected, finite, simple graphs. Let G be a graph and denote its vertex set by $V(G)$ and the edge set by $E(G)$. We denote the independence number of G by $\alpha(G)$, the vertex covering number by $\textsc{vc}(G)$ and the feedback vertex number by $\textsc{fvs}(G)$. We will abbreviate maximum independent set as MIS, and feedback vertex set as FVS. For $v \in V(G)$ we denote the open and closed neighborhoods of v by $N_G(v)$ and $N_G[v]$, respectively. For a set $S \subseteq V(G)$ we have $N_G(S) := \bigcup_{v \in S} N_G(v) \setminus S$, and $N_G[S] := \bigcup_{v \in S} N_G[v]$. We write $G' \subseteq G$ if G' is a subgraph of G. The graph $G[V(G) \setminus X]$ obtained from G by deleting the vertices in X and their incident edges is denoted by $G - X$. The graph $G[E(G) \setminus Y]$ obtained

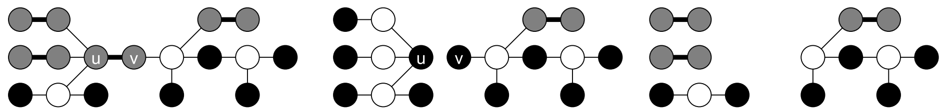

(a) Graph G and its independence decomposition. **(b)** Decomposition for graph $G/\{\{u,v\}\}$. **(c)** Decomposition for graph $G - \{u,v\}$.

Figure 1 Examples of the independence decomposition of a graph. Black vertices are in A, white vertices are in N, gray vertices are in S and the edges in M are drawn with thick lines.

from G by deleting the edges in Y but *not* their endpoints is denoted by G/Y. Carefully observe the difference between these two operators: if $\{u,v\}$ is an edge in G, then $G - \{u,v\}$ is the graph obtained from G by deleting the vertices u, v and their incident edges, whereas $G/\{\{u,v\}\}$ is the graph obtained from G by removing the edge $\{u,v\}$ while leaving the endpoints u and v intact. We note that many details had to be omitted from this extended abstract due to space restrictions; they can be found in the full version [24] of this work.

We need the following proposition on the structure of maximum independent sets in trees by Zito [32, Theorem 2], which we re-state here in terms of forests:

▶ Proposition 1. Let F be a forest. Then there is a unique partition of the vertex set $V(F)$ into subsets A, N, S such that:

1. Any MIS for F contains all vertices of A and no vertices of N.
2. For each vertex $v \in S$ there is a MIS for F containing v and a MIS for F avoiding v.
3. There is a perfect matching M in $F[S]$, and any MIS for F contains exactly one endpoint of each edge in M.
4. The matching M contains all the α-critical edges of F: for all $e \in E(F)$ it holds that $\alpha(F) < \alpha(F/\{e\}) \Leftrightarrow e \in M$.

This partition is uniquely characterized by adjacency relations. The sets A, N, S form the described partition if and only if:

I. There is a matching M on the vertices of S.
II. No vertex of A is adjacent to another vertex of A or to a vertex in S.
III. Each vertex of N is adjacent to at least two vertices of A.

We will refer to this decomposition of the vertex set of a forest F into the subsets A, N, S with the matching M as its *independence decomposition* (Figure 1).

▶ Observation 1. Let G be a graph. If G' is a vertex-induced subgraph of G then $\alpha(G) \geq \alpha(G')$, so for all $W \subseteq V(G)$ we have $\alpha(G) \geq \alpha(G - W)$. If G'' is an edge-induced subgraph of G then $\alpha(G'') \geq \alpha(G)$, so for all $Z \subseteq E(G)$ we have $\alpha(G) \leq \alpha(G/Z)$.

▶ Observation 2. If G is a graph and v is a vertex in G such that $\deg_G(v) \leq 1$ then there is a MIS for G that contains v.

3 Cubic Kernel for FVS-Independent Set

In this section we develop a cubic kernel for FVS-INDEPENDENT SET. Consider an instance (G, X, k) of the problem, which asks whether graph G with the FVS X has an independent set of size k. Throughout this section $F := G - X$ denotes the forest obtained by deleting the vertices in X. Our starting point is the current-best VERTEX COVER kernelization [8, Theorem 2.2] which exploits a theorem by Nemhauser and Trotter [27].

▶ **Theorem 1.** *There is a polynomial-time algorithm that takes an instance (G, k) of* Vertex Cover *as input, and computes in polynomial time an equivalent instance (G', k') such that G' is a vertex-induced subgraph of G with $k' \leq k$, $|V(G')| - k' \leq |V(G)| - k$ and $|V(G')| \leq 2k'$. We can ensure that G' does not contain any vertices of degree ≤ 1.*

Through the correspondence between Vertex Cover and Independent Set we can use Theorem 1 to preprocess an instance (G, X, k) of fvs-Independent Set.

▶ Reduction Rule 1. Let (G, X, k) be the current instance of fvs-Independent Set. Run the algorithm from Theorem 1 on the Vertex Cover instance $(G, |V(G)| - k)$ and let the result be $(G', |V(G')| - k')$. Obtain X' from X by deleting the vertices that were removed from G by the algorithm, and use (G', X', k') as the new instance of fvs-Independent Set.

When given an independent subset $X' \subseteq X$ of the feedback vertices we can efficiently compute the largest independent set I in G which satisfies $I \cap X = X'$: since such a set intersects X exactly in X', and since it cannot use any neighbors of X' the maximum size is $|X'| + \alpha(F - N_G(X'))$ and this is polynomial-time computable since $F - N_G(X')$ is a forest. We exploit this to assess which vertices from the FVS X might occur in a MIS of G.

▶ **Definition 2.** The number of *conflicts* $\textsc{Conf}_{F'}(X')$ induced by a subset $X' \subseteq X$ on a subforest $F' \subseteq F \subseteq G$ is defined as $\textsc{Conf}_{F'}(X') := \alpha(F') - \alpha(F' - N_G(X'))$.

This term $\textsc{Conf}_{F'}(X')$ can be interpreted as follows. Choosing vertices from X' in an independent set will prevent all their neighbors in the subforest F' from being part of the same independent set; hence if we fix some choice of vertices in X', then the number of vertices from F' we can add to this set (while maintaining independence) might be smaller than the independence number of F'. The term $\textsc{Conf}_{F'}(X')$ measures the difference between the two: informally it is the price we pay in the forest F' for choosing the vertices X' in the independent set. We can now formulate our first new reduction rules.

▶ Reduction Rule 2. If there is a vertex $v \in X$ such that $\textsc{Conf}_F(\{v\}) \geq |X|$, then delete v from the graph G and from the set X.

▶ Reduction Rule 3. If there are distinct vertices $u, v \in X$ with $\{u, v\} \notin E(G)$ for which $\textsc{Conf}_F(\{u, v\}) \geq |X|$, then add the edge $\{u, v\}$ to G.

Correctness of these two rules can be established from the following lemma.

▶ **Lemma 3.** *If $X' \subseteq X$ is a subset of feedback vertices such that $\textsc{Conf}_F(X') \geq |X|$ then there is a MIS for G that does not contain all vertices of X'.*

Proof. Assume that $I \subseteq V(G)$ is an independent set containing all vertices of X'. We will prove that there is an independent set I' which is disjoint from X' with $|I'| \geq |I|$. Since $\textsc{Conf}_F(X') \geq |X|$ it follows by definition that $\alpha(F) - \alpha(F - N_G(X')) \geq |X|$; since I cannot contain any neighbors of vertices in X' we know that $|I \cap V(F)| \leq \alpha(F - N_G(X'))$, and since $|V(G)| = |X| + |V(F)|$ we have $|I| \leq |X| + \alpha(F - N_G(X')) \leq \alpha(F)$. Hence the maximum independent set for F, which does not contain any vertices of X', is at least as large as I; this proves that for every independent set containing X' there is another independent set which is at least as large and avoids the vertices of X'. Therefore there is a MIS for G avoiding at least one vertex of X'. ◀

▶ Reduction Rule 4. If F contains a connected component T (which must be a tree) such that for all $X' \subseteq X$ with $|X'| \leq 2$ for which X' is independent in G it holds that $\textsc{Conf}_T(X') = 0$, then delete T from graph G and decrease k by $\alpha(T)$.

To prove the correctness of Rule 4 we need the following lemma.

▶ **Lemma 4.** *Let T be a connected component of F and let $X_I \subseteq X$ be an independent set in G. If $\textsc{Conf}_T(X_I) > 0$ then there is a set $X' \subseteq X_I$ with $|X'| \leq 2$ such that $\textsc{Conf}_T(X') > 0$.*

Proof. Assume the conditions stated in the lemma hold. Consider the independence decomposition of T into the sets A, N, S, and let M be a perfect matching on $T[S]$. We will try to construct a MIS I for T that does not use any vertices in $N_G(X_I)$; this must then also be a MIS for $T - N_G(X_I)$ of the same size. By the assumption that $\textsc{Conf}_T(X_I) > 0$ any independent set in T must use at least one vertex in $N_G(X_I)$ in order to be maximum, hence our construction procedure must fail somewhere; the place where it fails will provide us with a set X' as required by the statement of the lemma.

Construction of a MIS. By Proposition 1 any MIS for T must use all vertices in A, no vertices from N and exactly one endpoint of each edge in the matching M. It follows that if some vertex $v \in A$ is adjacent in G to a vertex $x \in X_I$, then $\alpha(T - \{v\}) < \alpha(T)$ and therefore $\alpha(T - N_G(x)) < \alpha(T)$ which proves that $\textsc{Conf}_T(\{x\}) > 0$; hence we can then use $X' := \{x\}$ as our desired subset to prove the claim. In the remainder of the proof we may therefore assume that no vertex of A is adjacent in G to a vertex in X_I.

We now start building our independent set I for T that avoids vertices in $N_G(X_I)$. We start by taking all vertices of A in the independent set; we do not use any vertices in $N_G(X_I)$ here since $A \cap N_G(X_I) = \emptyset$ by assumption. To augment I into a MIS for T it remains to add one endpoint of each edge in the matching M. Since the endpoints of the matching are not adjacent to vertices in A by the adjacency rules of Proposition 1, we can now restrict ourselves to the subgraph $T' := T[S]$ induced by the matched vertices since no choice of independent vertices in $T[S]$ will conflict with the choice of the vertices A. If there is a vertex v in T' such that $N_{T'}(v) = \{u\}$ and $N_G(v) \cap X_I = \emptyset$, then the edge $\{v, u\}$ must be in the matching M (since M is a perfect matching in $T[S]$). Because we must choose one of $\{u, v\}$ in a MIS for T, and by Observation 2 choosing a degree-1 vertex will never conflict with choices that are made later on, we can add v to our independent set I while respecting the invariant that no vertex in I is adjacent in G to a vertex in X_I. Since we have then chosen one endpoint of the matching edge $\{u, v\}$ in I, we can delete u, v and their incident edges to obtain a smaller graph T' (which again contains a perfect submatching of M) in which we continue the process. As long as there is a vertex with degree 1 in T' that has no neighbors in X_I then we take it into I, delete it and its neighbor, and continue. If this process ends with an empty graph, then by Proposition 1 the set I must be a MIS for T, and since it does not use any vertices adjacent to X_I it must also be a MIS for $T - N_G(X_I)$; but this proves that $\alpha(T) = \alpha(T - N_G(X_I))$ which means $\textsc{Conf}_T(X_I) = 0$, which is a contradiction to the assumption at the start of the proof. So the process must end with a non-empty graph $T' \subseteq T$ such that vertices with degree 1 in T' are adjacent in G to a vertex in X_I and for which the matching M restricted to T' is a perfect matching on T'. We use this subgraph T' to obtain a set X' as desired.

Using the subgraph to prove the claim. Consider a vertex v_0 in T' with $\deg_{T'}(v_0) = 1$, and construct a path $P = \{v_0, v_1, \ldots, v_{2p+1}\}$ by following edges of T' that are alternatingly in and out of the matching M, until arriving at a degree-1 vertex whose only neighbor was already visited. Since T' is acyclic, M restricted to T' is a perfect matching on T' and we start the process at a vertex of degree 1, it is easy to verify that there must be such a path P (there can be many; any arbitrary such path will suffice), that P must contain an even number of vertices, that the first and last vertex on P have degree-1 in T' and that the edges $\{v_{2i}, v_{2i+1}\}$ must be in M for all $0 \leq i \leq p$. Since we assumed that all degree-1 vertices in T' are adjacent in G to X_I, there exist vertices $x_1, x_2 \in X$ such that $v_0 \in N_G(x_1)$

and $v_{2p+1} \in N_G(x_2)$. We now claim that $X' := \{x_1, x_2\}$ satisfies the requirements of the statement of the lemma, i.e. that $\textsc{Conf}_T(\{x_1, x_2\}) > 0$. This fact is witnessed by considering the path P in the original tree T. Any MIS for T which avoids $N_G(\{x_1, x_2\})$ must use one endpoint of the matched edge $\{v_0, v_1\}$, and since the choice of v_0 is blocked because v_0 is a neighbor to x_1, it must use v_1. But path P shows that v_1 is adjacent in T to v_2, and hence we cannot choose v_2 in the independent set. But since $\{v_2, v_3\}$ is again a matched edge, we must use one of its endpoints; hence we must use v_3. Repeating this argument shows that we must use vertex v_{2p+1} in a MIS for T if we cannot use v_0; but the use of v_{2p+1} is also not possible if we exclude $N_G(\{x_1, x_2\})$. Hence we cannot make a MIS for T without using vertices in $N_G(\{x_1, x_2\})$ which proves that $\alpha(T) > \alpha(T - N_G(\{x_1, x_2\})$. By the definition of conflicts this proves that $\textsc{Conf}_T(X') > 0$ for $X' = \{x_1, x_2\}$, which concludes the proof. ◀

Using this lemma we can prove the correctness of Rule 4.

▶ **Lemma 5.** *Rule 4 is correct: if T is a connected component in F such that for all $X' \subseteq X$ which are independent in G and satisfy $|X'| \leq 2$ it holds that $\textsc{Conf}_T(X') = 0$, then $\alpha(G) = \alpha(G - T) + \alpha(T)$.*

Proof. Assume the conditions in the statement of the lemma hold. It is trivial to see that $\alpha(G) \leq \alpha(G - T) + \alpha(T)$. To establish the lemma we only need to prove that $\alpha(G) \geq \alpha(G - T) + \alpha(T)$, which we will do by showing that any independent set I_{G-T} in $G - T$ can be transformed to an independent set of size at least $|I_{G-T}| + \alpha(T)$ in G. So consider such an independent set I_{G-T}, and let $X_I := I_{G-T} \cap X$ be the set of vertices which belong to both I_{G-T} and the feedback vertex set X. Suppose that $\alpha(T) > \alpha(T - N_G(X_I))$. Then by Lemma 4 there is a subset $X' \subseteq X_I$ with $|X'| \leq 2$ such that $\textsc{Conf}_T(X') > 0$. Since X_I is an independent set, such a subset X' would also be independent; but by the preconditions to this lemma such a set X' does not exist and therefore we must have $\alpha(T) = \alpha(T - N_G(X_I))$.

Now we show how to transform I_{G-T} into an independent set for G of the requested size. Let I_T be a MIS in $T - N_G(X_I)$, which has size $\alpha(T - N_G(X_I)) = \alpha(T)$. It is easy to verify that $I_{G-T} \cup I_T$ must be an independent set in G because vertices of T are only adjacent to vertices of $G - T$ which are contained in X. Hence the set $I_{G-T} \cup I_T$ is independent in G and it has size $|I_{G-T}| + \alpha(T)$. Since this argument applies to any independent set I_{G-T} in graph $G - T$ it holds in particular for a MIS in $G - T$, which proves that $\alpha(G) \geq \alpha(G - T) + \alpha(T)$ which proves the claim. ◀

We introduce the concept of blockability for the statement of the last two reduction rules.

▶ **Definition 6.** We say that the pair $x, y \in V(G) \setminus X$ is X-*blockable* if G contains an independent set $X' \subseteq X$ of size $|X'| \leq 2$ such that $\{x, y\} \subseteq N_G(X')$.

This can be interpreted as follows: any independent set in G that contains X' cannot contain x or y, so the pair x, y is blocked from being in an independent set by choosing X'. It follows directly from the definition that if x, y is not X-blockable, then for any combination of $u \in N_G(x) \cap X$ and $v \in N_G(y) \cap X$ we must have $\{u, v\} \in E(G)$.

▶ Reduction Rule 5. If there are distinct vertices $u, v \in V(G) \setminus X$ which are adjacent in G and are not X-blockable such that $\deg_F(u), \deg_F(v) \leq 2$ then reduce the graph as follows. Delete vertices u, v and decrease k by 1. If u has a neighbor t in F which is not v, then make all vertices of $N_G(v) \cap X$ adjacent to t. If v has a neighbor w in F which is not u, then make all vertices of $N_G(u) \cap X$ adjacent to w. If the vertices t, w exist then they must be unique; add the edge $\{t, w\}$ to the graph.

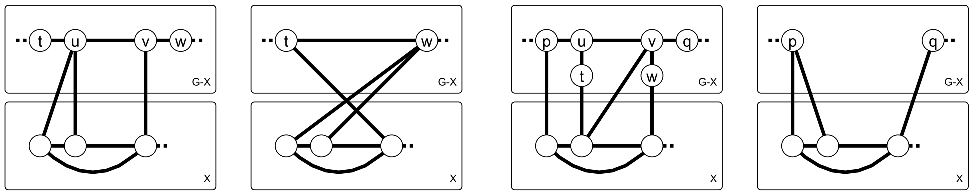

(a) Rule 5: Shrinking unblockable degree-2 paths in trees. ($k' := k - 1$)

(b) Rule 6: Removing unblockable pendants in trees. ($k' := k - 2$)

Figure 2 Illustrations of two reduction rules. The original structure is shown on the left, and the image on the right shows the structure after the reduction. Feedback vertices X are drawn in the bottom container, whereas the forest $G - X$ is visualized in the top container.

▶ Reduction Rule 6. If there are distinct vertices t, u, v, w in $V(G) \setminus X$ such that $\deg_F(u) = \deg_F(v) = 3$, $N_F(t) = \{u\}$, $N_F(w) = \{v\}$ and $\{u, v\} \in E(G)$ such that none of the pairs $\{u, t\}$, $\{v, w\}$, $\{t, w\}$ are X-blockable, then reduce as follows. Let $\{p\} = N_F(u) \setminus \{t, v\}$ and let $\{q\} = N_F(v) \setminus \{w, u\}$. Delete $\{t, u, v, w\}$ and their incident edges from G, decrease k by 2, make p adjacent to all vertices of $N_G(t) \cap X$ and make q adjacent to all vertices of $N_G(w) \cap X$.

See Figure 2 for an illustration of the final two reduction rules, which are meant to reduce the sizes of the trees in the forest F. The correctness of these rules can be proven by an exchange argument. Whereas Rule 4 deletes a tree T from the forest F when we can derive that for every independent set in $G - T$ we can obtain an independent set in G which is $\alpha(T)$ vertices larger, these last reduction rules act *locally* within one tree, but according to the same principle. Instead of working on an entire connected component of F, they reduce subtrees $T' \subseteq F$ in situations where we can derive that every independent set in X can be augmented with $\alpha(T')$ vertices from T'. In Rule 5 we reduce the subtree on vertices $\{u, v\}$ which has independence number 1, and in Rule 6 we reduce the subtree on vertices $\{u, v, t, w\}$ with independence number 2. Connections between the vertices adjacent to the reduced subtree are made to enforce that removal of the subtree does not affect the types of interactions between the neighboring vertices.

When no reduction rules can be applied to an instance, we call it *reduced*. In reduced instances the number of vertices in F must be bounded by a function of $|X|$, which can be proven using the following notion.

▶ **Definition 7.** Let F' be a subforest of F, and define the number of *active conflicts* induced on F' by the feedback set X as follows: $\textsc{Active}_{F'}(X) := \sum_{X' \in \mathcal{X}} \textsc{Conf}_{F'}(X')$ using the abbreviation $\mathcal{X} := \{X' \mid X' \subseteq X \wedge |X'| \leq 2 \wedge X' \text{ is independent in } G\}$.

The number of active conflicts induced on F in a reduced instance is polynomially bounded in $|X|$. For every $v \in X$ we have $\textsc{Conf}_F(\{v\}) < |X|$ by Rule 2, and every pair of distinct non-adjacent vertices $\{u, v\} \subseteq X$ satisfies $\textsc{Conf}_F(\{u, v\}) < |X|$ by Rule 3. Hence for every reduced instance we have $\textsc{Active}_F(X) \leq |X|^2 + \binom{|X|}{2}|X|$. A technical proof shows that in a reduced instance the number of active conflicts induced on the forest F is at least $\frac{1}{83}|V(F)|$. By combining this with the bound on the number of active conflicts, we can bound the size of reduced instances and obtain a kernelization algorithm. The algorithm exhaustively applies the six reduction rules, and the analysis then shows that the instance must be small when no more reduction rules can be applied. Using the duality of Vertex Cover and Independent Set we also obtain a kernel for fvs-Vertex Cover as a corollary.

▶ **Theorem 8.** FVS-INDEPENDENT SET *has a kernel with a cubic number of vertices: there is a polynomial-time algorithm that transforms an instance* (G, X, k) *into an equivalent instance* (G', X', k') *such that* $|X'| \leq |X|$, $k' \leq k$, $|V(G')| - k' \leq |V(G)| - k$, $|V(G')| \leq 2(|V(G')| - k')$ *and* $|V(G')| \leq |X| + 83|X|^3$.

▶ **Corollary 9.** FVS-VERTEX COVER *has a kernel with* $\min(2k, |X| + 83|X|^3)$ *vertices.*

4 No Polynomial Kernel for FVS-Weighted Independent Set

In this section we show that the introduction of vertex weights makes the parameterized INDEPENDENT SET problem harder to kernelize, by proving that FVS-WEIGHTED INDEPENDENT SET does not have a polynomial kernel unless PH $= \Sigma_3^p$. To establish this result, we introduce a new parameterized problem called t-PAIRED VECTOR AGREEMENT and show that it does not have a polynomial kernel unless PH $= \Sigma_3^p$. We then finish the proof by giving a polynomial-parameter transformation [5, 14] to FVS-WEIGHTED INDEPENDENT SET.

t-PAIRED VECTOR AGREEMENT
Instance: A list L consisting of t pairs of vectors (a^i, b^i) for $1 \leq i \leq t$ such that each vector is an element of $\{0, 1, \#, ?\}^m$, and an integer $k \geq 0$.
Parameter: The number of pairs t.
Question: Is it possible to choose one vector from each pair, such that the chosen vectors S induce at most k conflict positions? A position $1 \leq j \leq m$ in a vector is a *conflict* position if some chosen vector $v \in S$ has $v_j = \#$, or if we have chosen vectors $u, v \in S$ such that $u_j = 0$ and $v_j = 1$.

The framework for proving that a parameterized problem does not have a polynomial kernel unless PH $= \Sigma_3^p$ requires us to establish that the corresponding classical problem is NP-complete. A reduction from VERTEX COVER shows that the classic problem PAIRED VECTOR AGREEMENT is NP-complete. By exploiting the fact that t-PAIRED VECTOR AGREEMENT can be solved in $O(2^t p(m))$ time for some polynomial p (by trying all possible combinations of vectors), we can build an OR-composition algorithm for the paired agreement problem using a bitmask selection strategy; the techniques we use here are similar to those employed by Dom et al. [14]. These two facts prove that t-PAIRED VECTOR AGREEMENT does not have a polynomial kernel unless PH $= \Sigma_3^p$. To relate these results to FVS-WEIGHTED INDEPENDENT SET we use the following transformation.

▶ **Lemma 10.** *There is a polynomial-parameter reduction from* t-PAIRED VECTOR AGREEMENT *to* FVS-WEIGHTED INDEPENDENT SET*.*

Proof. Let (L, t, m, k) be an instance of t-PAIRED VECTOR AGREEMENT. We may assume that $k < m$ otherwise the answer to the instance is trivially YES. We show how to build an equivalent instance (G', w', X', k') of FVS-INDEPENDENT SET in polynomial time such that $|X'| = 2t$, which implies the existence of a polynomial-parameter reduction.

The graph G' has $2(t + m)$ vertices, and is defined as follows. For each index $1 \leq i \leq t$ there is a pair of vertices v_i^a, v_i^b which are connected by an edge, and have weight $2(t + m)$. For each vector position $1 \leq j \leq m$ there are vertices p_j^0, p_j^1 which are connected by an edge, and have weight 1. The vertices v_i^a and v_i^b correspond to the vectors a^i, b^i of the t-PAIRED VECTOR AGREEMENT instance, and are connected to the position-vertices as follows. Let v be a vertex v_i^a or v_i^b corresponding to the vector VEC(v) which is a_i or b_i, respectively. For $1 \leq i \leq t$ vertex v is adjacent in G' to all p_j^0 for which vector VEC(v) has a 0 at position j; it is also adjacent to all p_j^1 for which vector $\textsc{vec}(v)$ has a 1 at position j, and finally vertex v is adjacent to all $\{p_j^0, p_j^1\}$ for which vector $\textsc{vec}(v)$ has a # at position j. This concludes the definition of the structure of graph G'.

One may verify that a position vertex p_j^x is adjacent to exactly 1 other position vertex p_j^{1-x}, which implies that the graph induced by the position vertices $p_j^{0,1}$ has maximum degree 1 and is therefore a forest; this shows that the vector-vertices $v_i^{a/b}$ form a feedback vertex set for G' and thus we define the feedback vertex set for our instance as $X' := \{v_i^a, v_i^b \mid 1 \leq i \leq t\}$ which has size exactly $2t$. We now ask for an independent set of total weight at least $k' := 2t(t+m) + (m-k)$, which completes the description of instance (G', w', X', k'). It is easy to see that this instance can be computed in polynomial time from the instance (L, t, m, k). The proof that these two instances are equivalent is not difficult, and has been deferred to the full version of this paper. ◀

By standard kernelization lower-bound techniques (see [5, 14]) Lemma 10 implies:

▶ **Theorem 11.** *The problems* FVS-WEIGHTED INDEPENDENT SET *and* FVS-WEIGHTED VERTEX COVER *do not admit polynomial kernels unless* $PH = \Sigma_3^p$.

It is interesting to note that an instance (G', w', X', k') of FVS-INDEPENDENT SET resulting from the polynomial-parameter transformation of Lemma 10 has a very restricted graph structure: every connected component of the forest $G' - X'$ is a path on two vertices. Hence our proof shows that even using the parameter "number of vertex deletions needed to turn the graph into a disjoint union of P_2's" (a structurally *larger* parameter than the FVS size) there is no polynomial kernel unless $\mathrm{PH} = \Sigma_3^p$.

5 Conclusion

We have given upper and lower bounds on the size of kernels for the VERTEX COVER and INDEPENDENT SET problems using the parameter $\textsc{fvs}(G)$. It would be very interesting to perform experiments with our new reduction rules to see whether they offer significant benefits over the existing VERTEX COVER kernel on real-world instances. This result is one of the first examples of a polynomial kernel using a "refined" parameter which is structurally smaller than the standard parameterization. The contrasting result on the weighted problem shows that there is a rich structure waiting to be uncovered when studying kernelization using non-standard parameters. The kernel we have presented for FVS-VERTEX COVER contains $O(|X|^3)$ vertices and can therefore be encoded in $O(|X|^6)$ bits using an adjacency matrix. The results of Dell and Van Melkebeek [12] imply that it is unlikely that there exists a kernel which can be encoded in $O(|X|^{2-\epsilon})$ bits for any $\epsilon > 0$. It might be possible to improve the size of the kernel to a quadratic or even a linear number of vertices, by employing new reduction rules. The current reduction rules can be seen as analogs of the traditional "high degree" rule for the VERTEX COVER problem, and it would be interesting to see whether it is possible to find analogs of crown reduction rules when using $\textsc{fvs}(G)$ as the parameter.

Although we have assumed throughout the paper that a feedback vertex set is supplied with the input, we can drop this restriction by applying the known polynomial-time 2-approximation algorithm for FVS [3]. Observe that the reduction algorithm does not require that the supplied set X is a *minimum* feedback vertex set; the kernelization algorithm works if X is *any* feedback vertex set, and the size of the output instance depends on the size of the FVS that is supplied. Hence if we compute a 2-approximate FVS and supply it to the kernelization algorithm, the bound on the number of vertices in the output instance is only a factor 2 worse than when running the kernelization using a *minimum* FVS.

This paper has focused on the decision version of the Vertex Cover problem, but the data reduction rules given here can also be translated to the optimization version to obtain the following result: given a graph G there is a polynomial-time algorithm that computes a graph G' and a non-negative integer c such that $\textsc{vc}(G) = \textsc{vc}(G') + c$ with $|V(G')| \leq 2\,\textsc{vc}(G)$ and $|V(G')| \in O(\textsc{fvs}(G)^3)$; and a vertex cover S' for G' can be transformed back into a vertex cover of G of size $|S'| + c$ in polynomial time.

The approach of studying Vertex Cover parameterized by $\textsc{fvs}(G)$ fits into the broad context of "parameterizing away from triviality" [28, 7], since the parameter $\textsc{fvs}(G)$ measures how many vertex-deletions are needed to reduce G to a forest in which Vertex Cover can be solved trivially in polynomial time. As there is a wide variety of restricted graph classes for which Vertex Cover is in P, this opens up a multitude of possibilities for non-standard parameterizations. For every graph class $\mathcal{G}$ which is closed under vertex deletion and for which the Vertex Cover problem is in P, the Vertex Cover problem is in FPT when parameterized by the size of a set X such that $G - X \in \mathcal{G}$, assuming that X is given as part of the input. Recent work [4] into this direction shows that whenever $\mathcal{G}$ contains all cliques the resulting parameterized problem does not have a polynomial kernel unless PH $= \Sigma_3^p$. Examples of such classes $\mathcal{G}$ are chordal graphs, interval graphs and other types of perfect graphs. We conclude with two specific open problems. Is there a polynomial kernel using the deletion distance from a bipartite graph as the parameter? Does the Vertex Cover problem parameterized by the size of a minimum set X such that $\textsc{treewidth}(G - X) \leq i$ have a polynomial kernel for every fixed i, or is there some value of i for which this problem does not have a polynomial kernel? The classic Vertex Cover kernelizations can be interpreted as the case $i = 0$ whereas this paper supplies the result for $i = 1$. It appears that many interesting insights are waiting to be discovered in this direction.

References

1. Faisal N. Abu-Khzam, Rebecca L. Collins, Michael R. Fellows, Michael A. Langston, W. Henry Suters, and Christopher T. Symons. Kernelization algorithms for the vertex cover problem: Theory and experiments. In *Proc. 6th ALENEX/ANALC*, pages 62–69, 2004.
2. Faisal N. Abu-Khzam, Michael R. Fellows, Michael A. Langston, and W. Henry Suters. Crown structures for vertex cover kernelization. *Theory Comput. Syst.*, 41(3):411–430, 2007.
3. Vineet Bafna, Piotr Berman, and Toshihiro Fujito. A 2-approximation algorithm for the undirected feedback vertex set problem. *SIAM Journal on Discrete Mathematics*, 12(3):289–297, 1999.
4. Hans L. Bodlaender, Bart M. P. Jansen, and Stefan Kratsch. Cross-composition: A new technique for kernelization lower bounds. *CoRR*, abs/1011.4224, 2010.
5. Hans L. Bodlaender, Stéphan Thomassé, and Anders Yeo. Kernel bounds for disjoint cycles and disjoint paths. In *Proc. 17th ESA*, pages 635–646, 2009.
6. Jonathan F. Buss and Judy Goldsmith. Nondeterminism within P. *SIAM J. Comput.*, 22(3):560–572, 1993.
7. Leizhen Cai. Parameterized complexity of vertex colouring. *Discrete Applied Mathematics*, 127(3):415–429, 2003.
8. Jianer Chen, Iyad A. Kanj, and Weijia Jia. Vertex cover: Further observations and further improvements. *J. Algorithms*, 41(2):280–301, 2001.
9. Jianer Chen, Iyad A. Kanj, and Ge Xia. Improved parameterized upper bounds for vertex cover. In *Proc. 31st MFCS*, pages 238–249, 2006.

10 Miroslav Chlebík and Janka Chlebíková. Crown reductions for the minimum weighted vertex cover problem. *Discrete Applied Mathematics*, 156(3):292–312, 2008.

11 Benny Chor, Mike Fellows, and David W. Juedes. Linear kernels in linear time, or how to save k colors in $O(n^2)$ steps. In *Proc. 30th WG*, pages 257–269, 2004.

12 Holger Dell and Dieter van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. In *Proc. 42nd STOC*, pages 251–260, 2010.

13 Josep Díaz, Jordi Petit, and Dimitrios M. Thilikos. Kernels for the vertex cover problem on the preferred attachment model. In *Proc. 5th WEA*, pages 231–240, 2006.

14 Michael Dom, Daniel Lokshtanov, and Saket Saurabh. Incompressibility through colors and IDs. In *Proc. 36th ICALP*, pages 378–389, 2009.

15 Rod Downey and Michael R. Fellows. *Parameterized Complexity.* Monographs in Computer Science. Springer, New York, 1999.

16 Rodney G. Downey, Michael R. Fellows, and Michael A. Langston, editors. *The Computer Journal: Special Issue on Parameterized Complexity*, volume 51, 2008.

17 Rodney G. Downey, Michael R. Fellows, and Ulrike Stege. Parameterized complexity: A framework for systematically confronting computational intractability. In *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 49–99, 1997.

18 Vladimir Estivill-Castro, Michael Fellows, Michael Langston, and Frances Rosamond. FPT is P-time extremal structure I. In *Proc. 1st ACiD*, pages 1–41, 2005.

19 Michael R. Fellows. Towards fully multivariate algorithmics: Some new results and directions in parameter ecology. In *Proc. 20th IWOCA*, pages 2–10, 2009.

20 Michael R. Fellows, Daniel Lokshtanov, Neeldhara Misra, Matthias Mnich, Frances A. Rosamond, and Saket Saurabh. The complexity ecology of parameters: An illustration using bounded max leaf number. *Theory Comput. Syst.*, 45(4):822–848, 2009.

21 Michael R. Fellows, Daniel Lokshtanov, Neeldhara Misra, Frances A. Rosamond, and Saket Saurabh. Graph layout problems parameterized by vertex cover. In *Proc. 19th ISAAC*, pages 294–305, 2008.

22 Michael R. Garey and David S. Johnson. *Computers and Intractability, A Guide to the Theory of NP-Completeness.* W.H. Freeman and Company, New York, 1979.

23 Jiong Guo and Rolf Niedermeier. Invitation to data reduction and problem kernelization. *SIGACT News*, 38(1):31–45, 2007.

24 Bart M. P. Jansen and Hans L. Bodlaender. Vertex cover kernelization revisited: Upper and lower bounds for a refined parameter. *CoRR*, abs/1012.4701, 2010.

25 Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.

26 Stefan Kratsch and Pascal Schweitzer. Isomorphism for graphs of bounded feedback vertex set number. In *Proc. 12th SWAT*, pages 81–92, 2010.

27 G.L. Nemhauser and L.E.jun. Trotter. Vertex packings: structural properties and algorithms. *Math. Program.*, 8:232–248, 1975.

28 Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms.* Oxford University Press, 2006.

29 Rolf Niedermeier. Reflections on multivariate algorithmics and problem parameterization. In *Proc. 27th STACS*, pages 17–32, 2010.

30 Rolf Niedermeier and Peter Rossmanith. On efficient fixed-parameter algorithms for weighted vertex cover. *J. Algorithms*, 47(2):63–77, 2003.

31 Johannes Uhlmann and Mathias Weller. Two-layer planarization parameterized by feedback edge set. In *Proc. 7th TAMC*, pages 431–442, 2010.

32 Jennifer Zito. The structure and maximum number of maximum independent sets in trees. *J. Graph Theory*, 15(2):207–221, 1991.