


Polynomial Kernels for Hitting Forbidden Minors under Structural Parameterizations

Bart M. P. Jansen

Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands


b.m.p.jansen@tue.nl

 <https://orcid.org/0000-0001-8204-1268>

Astrid Pieterse

Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

a.pieterse@tue.nl

 <https://orcid.org/0000-0003-3721-6721>

Abstract

We investigate polynomial-time preprocessing for the problem of hitting forbidden minors in a graph, using the framework of kernelization. For a fixed finite set of graphs \mathcal{F} , the \mathcal{F} -DELETION problem is the following: given a graph G and integer k , is it possible to delete k vertices from G to ensure the resulting graph does not contain any graph from \mathcal{F} as a minor? Earlier work by Fomin, Lokshtanov, Misra, and Saurabh [FOCS'12] showed that when \mathcal{F} contains a planar graph, an instance (G, k) can be reduced in polynomial time to an equivalent one of size $k^{\mathcal{O}(1)}$. In this work we focus on structural measures of the complexity of an instance, with the aim of giving nontrivial preprocessing guarantees for instances whose solutions are large. Motivated by several impossibility results, we parameterize the \mathcal{F} -DELETION problem by the size of a vertex modulator whose removal results in a graph of constant treedepth η .

We prove that for each set \mathcal{F} of connected graphs and constant η , the \mathcal{F} -DELETION problem parameterized by the size of a treedepth- η modulator has a polynomial kernel. Our kernelization is fully explicit and does not depend on protrusion reduction or well-quasi-ordering, which are sources of algorithmic non-constructivity in earlier works on \mathcal{F} -DELETION. Our main technical contribution is to analyze how models of a forbidden minor in a graph G with modulator X , interact with the various connected components of $G - X$. Using the language of labeled minors, we analyze the fragments of potential forbidden minor models that can remain after removing an optimal \mathcal{F} -DELETION solution from a single connected component of $G - X$. By bounding the number of different types of behavior that can occur by a polynomial in $|X|$, we obtain a polynomial kernel using a recursive preprocessing strategy. Our results extend earlier work for specific instances of \mathcal{F} -DELETION such as VERTEX COVER and FEEDBACK VERTEX SET. It also generalizes earlier preprocessing results for \mathcal{F} -DELETION parameterized by a vertex cover, which is a treedepth-one modulator.

2012 ACM Subject Classification Theory of computation \rightarrow Graph algorithms analysis, Theory of computation \rightarrow Parameterized complexity and exact algorithms

Keywords and phrases Kernelization, \mathcal{F} -minor free deletion, Treedepth modulator, Structural parameterization

Digital Object Identifier 10.4230/LIPIcs.ESA.2018.48

Related Version A full version is available at [26], <https://arxiv.org/abs/1804.08885>.

Funding This work was supported by NWO Veni grant “Frontiers in Parameterized Preprocessing” and NWO Gravitation grant “Networks”.



© Bart M. P. Jansen and Astrid Pieterse;
licensed under Creative Commons License CC-BY
26th Annual European Symposium on Algorithms (ESA 2018).

Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 48; pp. 48:1–48:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

How, and under which circumstances, can a polynomial-time algorithm prune the easy parts of an NP-hard problem input, without changing its answer? This question can rigorously be answered using the notion of kernelization [1, 23, 29] which originated in parameterized complexity theory [8, 12] where it can be naturally framed. After choosing a *complexity parameter* for the NP-hard problem of interest, which associates to every input $x \in \Sigma^*$ an integer $k \in \mathbb{N}$ that expresses its difficulty under the chosen type of measurement, the theory postulates that a good preprocessing algorithm can be captured by the notion of a *polynomial kernelization*: a polynomial-time algorithm that, given a parameterized instance $(x, k) \in \Sigma^* \times \mathbb{N}$, outputs an instance (x', k') with the same answer whose size is bounded polynomially in k . Not all parameterized problems admit polynomial kernelizations, and one can find meaningful ways to preprocess an NP-hard problem by studying those parameterizations for which it does. The study of kernelization has blossomed over the last decade, resulting in a myriad of interesting techniques for obtaining polynomial kernelizations [3, 15, 24, 31, 34], as well as frameworks for proving the non-existence of polynomial kernelizations under complexity-theoretic assumptions [1, 2, 11, 13, 20].

Originally, the study of kernelization focused on the *natural parameterizations* of (the decision variants of) search problems, where the complexity parameter k measures the size of the solution. A classic example [7, 35] is that an instance (G, k) of the k -VERTEX COVER problem, which asks whether an undirected graph G has a vertex cover of size k , can efficiently be reduced to an equivalent instance with at most $2k$ vertices. This guarantees that efficient pruning can be done on large inputs that have small vertex covers. However, such guarantees are meaningless when the smallest vertex cover contains more than half the vertices. By choosing a parameter that measures the structure of the input graph, rather than the size of the desired solution, one can hope to develop provably good preprocessing procedures even for inputs whose solutions are large. An early example of this approach was given by Jansen and Bodlaender [25], who showed that an instance of the VERTEX COVER problem can efficiently be reduced to size $\mathcal{O}(\ell^3)$, where ℓ is the size of a smallest *feedback vertex set* in G : VERTEX COVER parameterized by the size of a feedback vertex set has a cubic-vertex kernel. The result effectively conveys that large instances of VERTEX COVER that are ℓ vertex-deletions away from being acyclic, can be shrunk to size $\mathcal{O}(\ell^3)$ in polynomial time.

Problem statement. To understand the power of polynomial-time preprocessing algorithms over inputs to NP-hard problems that exhibit some structural regularities, but whose solutions are generally large, we set out to answer the following question:

For which structural parameterizations of NP-hard graph problems is it possible to obtain polynomial kernelizations?

Our goal is to answer this question for a rich class of problems, in terms of a rich class of structural parameterizations. Existing lower bounds show that, in general graphs, it is unlikely that a logical characterization exists of the problems admitting polynomial kernelizations for structural parameterizations (cf. [16, §1]), even though meta-theorems in terms of logical definability or finite integer index are possible when dealing with inputs from sparse graph families [3, 21]. We therefore target the class of \mathcal{F} -MINOR-FREE DELETION problems, henceforth abbreviated as \mathcal{F} -DELETION problems, to capture a wide class of NP-hard graph problems. Such a problem is instantiated by specifying a finite set \mathcal{F} of forbidden minors. An input then consists of a graph G and integer k , and asks whether

it is possible to find a set $Y \subseteq V(G)$ of size k such that $G - Y$ contains no graph from \mathcal{F} as a minor. This is a rich class of problems: by choosing $\mathcal{F} = \{K_2\}$ we obtain VERTEX COVER, for $\mathcal{F} = \{K_3\}$ we have FEEDBACK VERTEX SET, and for $\mathcal{F} = \{K_5, K_{3,3}\}$ we obtain the problem of making a graph planar by vertex deletions. The kernelization complexity of the solution-size parameterization of \mathcal{F} -DELETION has been the subject of intensive research [17, 18, 22, 28, 40]. In this work we attempt to find the widest class of structural parameterizations for which \mathcal{F} -DELETION admits polynomial kernels, continuing a long line of investigation into structural parameterizations for VERTEX COVER [4, 19, 25, 30, 31, 33], FEEDBACK VERTEX SET [27, 32], and other \mathcal{F} -DELETION problems [16, 21].

When it comes to measuring graph complexity, a natural choice is to consider a *width measure* such as treewidth. Alas, it has long been known that even VERTEX COVER, the simplest \mathcal{F} -DELETION problem, does not admit a polynomial kernelization when parameterized by the treewidth of the input graph, assuming $\text{NP} \not\subseteq \text{coNP}/\text{poly}$.¹ Generally speaking, graph problems do not admit polynomial kernels under parameterizations that attain the maximum, rather than the sum, of the values of the connected components. We therefore use the *vertex-deletion* distance to simple graph classes \mathcal{G} as the parameter. The aforementioned result by Jansen and Bodlaender [25] shows that VERTEX COVER has a polynomial kernelization when parameterized by the vertex-deletion distance to an acyclic graph, i.e., to a graph of treewidth one. Unfortunately this formulation leaves little room for generalizations: no polynomial kernelization is possible parameterized by the distance to a graph of treewidth two [10, Theorem 11], or even pathwidth two.² We therefore cannot use the deletion distance to constant treewidth (TW) or pathwidth (PW) as our graph parameter, and use the deletion distance to constant *treedepth* (TD) instead. The parameter treedepth has recently attracted much interest [6, 14, 38], sometimes allowing better upper bounds than are possible in terms of treewidth [21, 37]. It plays an important role in the study of structural sparsity [36]. All graphs G satisfy $\text{TD}(G) \geq \text{PW}(G) \geq \text{TW}(G)$, so graphs of constant treedepth are more restricted than those of constant treewidth. We therefore study the following problem for a fixed set \mathcal{F} of connected graphs and constant $\eta \geq 1$.

\mathcal{F} -DELETION parameterized by treedepth- η modulator

Parameter: $|X|$.

Input: A graph G , integer k , and a modulator $X \subseteq V(G)$ such that $\text{TD}(G - X) \leq \eta$.

Question: Is there a set $Y \subseteq V(G)$ of size k such that $G - Y$ is \mathcal{F} -minor-free?

The restriction that \mathcal{F} contains only connected graphs is needed to ensure that a solution on a disconnected graph can be formed from solutions on its connected components, which we require in some of our proofs. The same assumption was used by Fomin et al. [18] to build a single-exponential FPT algorithm when \mathcal{F} contains a planar graph, and was later lifted in follow-up work by Kim et al. [28].

For technical reasons, we assume that a modulator X is given in the input. If no modulator is known, one can compute an approximate modulator and use it as X . For example, Gajarský et al. [21, Lemma 4.2] showed that a modulator of size at most 2^n times the optimum can be found in quadratic time. Our problem setting is related to that of Gajarský et al. [21]. They studied kernelization for a general class of graph problems that

¹ Bodlaender et al. [2, Theorem 1] show a superpolynomial kernelization lower bound for INDEPENDENT SET parameterized by treewidth. Since the parameter is not related to the solution size, this is equivalent to VERTEX COVER parameterized by treewidth. The lower bound holds under the assumption that $\text{NP} \not\subseteq \text{coNP}/\text{poly}$, which we implicitly assume when stating further lower bounds in this section.

² The lower bound is stated for distance to treewidth two, but the same proof works for pathwidth two.

includes \mathcal{F} -DELETION, parameterized by a constant-treewidth modulator, but under the additional restriction that the input graph has bounded expansion or is nowhere dense. Under this severe restriction they obtained kernelizations of linear size for a wide range of problems. This prompted Somnath Sikdar during the 2013 Workshop on Kernelization [9] to ask which types of problems admit polynomial kernelizations in *general graphs*, when parameterized by a constant-treewidth modulator; we address this question in this work.

Our results. Our main result proves the existence of polynomial kernelizations for \mathcal{F} -DELETION parameterized by a modulator whose removal leaves a graph of constant treewidth.

► **Theorem 1.** *For every fixed finite set \mathcal{F} of connected graphs and every constant η , the \mathcal{F} -DELETION problem parameterized by a treewidth- η modulator has a polynomial kernelization.*

This answers a question posed by Bougeret and Sau [4] (cf. [5]). They obtained polynomial kernels for VERTEX COVER parameterized by a constant-treewidth modulator, and asked whether their result can be extended to the FEEDBACK VERTEX SET problem. As FEEDBACK VERTEX SET is an \mathcal{F} -DELETION problem for $\mathcal{F} = \{K_3\}$, Theorem 1 shows that this is indeed the case. Theorem 1 greatly generalizes an earlier result of Fomin, Jansen, and Pilipczuk [16, Corollary 1], who proved that \mathcal{F} -DELETION parameterized by a *vertex cover* has a polynomial kernel for every fixed \mathcal{F} ; note that a vertex cover is precisely a treewidth-1 modulator.

Our kernelization is fully explicit and does not depend on protrusion replacement techniques or well-quasi-ordering, which are sources of algorithmic non-constructivity in other works [17, 18] on kernelization for \mathcal{F} -DELETION. Moreover, our general theorem allows \mathcal{F} to be any set of connected graphs, including nonplanar ones. In contrast, the kernelization for the solution-size parameterization by Fomin et al. [18] only applies when \mathcal{F} contains at least one planar graph. Hence they only capture problems where, after removing a solution, the remaining graph has constant treewidth [39]. In our case, even though the parameter value is expressed in terms of a modulator to a graph of constant treewidth and therefore constant treewidth, the graphs that result after removing an optimal solution may have unbounded treewidth. This occurs, for example, when using $\mathcal{F} = \{K_5, K_{3,3}\}$ to capture the VERTEX PLANARIZATION problem. (Whether the solution-size parameterization of VERTEX PLANARIZATION has a polynomial kernel is a notorious open problem [18].)

The degree of the polynomial in the kernel size bound grows very quickly with η . We prove that this is unavoidable, even for the simplest case of VERTEX COVER.

► **Theorem 2 (★).** *For every $\eta \geq 6$, the VERTEX COVER problem parameterized by the size of a given treewidth- η modulator X does not admit a kernelization of bitsize $\mathcal{O}(|X|^{2^{\eta-4}-\varepsilon})$ for any $\varepsilon > 0$, unless $\text{NP} \subseteq \text{coNP/poly}$.*

Techniques. To obtain a polynomial kernel for an instance (G, X, k) of \mathcal{F} -DELETION, the main challenge is to understand how the connected components \mathcal{C} of $G - X$ interact through their connections to the modulator X . Using the language of labeled minors, we analyze how minor models of a forbidden graph in \mathcal{F} may intersect the various components of $G - X$. Using these insights, we are able to characterize which components of \mathcal{C} affect the structure of optimal solutions in an essential way. On a high level, the kernelization strategy is as follows. We use the fact that a single constant-treewidth component can be analyzed efficiently, to identify a subset \mathcal{C}' of \mathcal{C} that contains $|X|^{\mathcal{O}(1)}$ essential components under our characterization. We prove that the remaining ones can be safely removed, because their interaction with the rest of the instance can be ignored. Formally speaking, we show that any optimal solution on $G' := G[X \cup \bigcup_{C \in \mathcal{C}'} C]$ can be lifted to a solution on G by

including $\Delta = \sum_{C \in \mathcal{C} \setminus \mathcal{C}' } \text{OPT}_{\mathcal{F}}(C)$ additional vertices: (G, X, k) is a YES-instance if and only if $(G', X, k - \Delta)$ is. This effectively shows that there is an optimal solution Y on G in which the non-essential components act in isolation: Y does not delete more vertices from such a component C , than would be deleted by a solution on the graph $G[C]$.

The overall kernelization follows straight-forwardly from this pruning of non-essential components by a recursive approach, similarly as in earlier work [4, 21]. The main challenge is therefore to understand which components are essential and which are not, and this is where our contribution lies. We present a stand-alone combinatorial lemma that captures our key insight in this direction. To state it, we introduce some terminology.

We work with a nonstandard notion of labeled graphs. For a finite set X , an X -labeled graph is a graph in which each vertex is assigned a (possibly empty) subset of X as its labelset; we stress that multiple vertices may carry the same label on their labelset. The minor relation on graphs extends to labeled graphs in a natural way: a labeled graph H is a minor of a labeled graph G , if H can be obtained from G by repeatedly deleting an edge, deleting a vertex, deleting a label from the labelset of a vertex, or contracting an edge. When contracting an edge $\{u, v\}$ into a single vertex w , the labelset of w is formed as the union of the labelsets of u and v .

For a collection \mathcal{S} of vertex subsets of an X -labeled graph C , and a set of X -labeled graphs \mathcal{Q} , we say that all $Y \in \mathcal{S}$ leave a \mathcal{Q} -minor in C , if for all $Y \in \mathcal{S}$ the graph $C - Y$ contains some graph $H \in \mathcal{Q}$ as a labeled minor. We say that a set \mathcal{Q} of X -labeled graphs is θ -saturated for an integer θ , if for each subset $X' \subseteq X$ of size θ , the graph consisting of one vertex with labelset X' belongs to \mathcal{Q} . Our main lemma states that if all optimal solutions to \mathcal{F} -DELETION on C leave a \mathcal{Q} -minor for some suitably saturated \mathcal{Q} , then there is a small subset \mathcal{Q}^* for which the same holds.

► **Lemma 3 (Main lemma ★).** *Let \mathcal{F} be a finite set of (unlabeled) connected graphs, let X be a set of labels, let \mathcal{Q} be a $(\min_{H \in \mathcal{F}} |V(H)|)$ -saturated set of connected X -labeled graphs of at most $\max_{H \in \mathcal{F}} |E(H)| + 1$ vertices each, and let C be an X -labeled graph. If all optimal solutions to \mathcal{F} -DELETION on C leave a \mathcal{Q} -minor, then there is a subset $\mathcal{Q}^* \subseteq \mathcal{Q}$ whose size depends only on $(\mathcal{F}, \text{TD}(C))$, such that all optimal solutions leave a \mathcal{Q}^* -minor.*

In several aspects, the statement in the lemma is best-possible. In particular, we will show in Section 3 that the dependence of the size of \mathcal{Q}^* on $\text{TD}(G)$ rather than $\text{TW}(G)$ is essential and that the precondition that \mathcal{Q} is $\mathcal{O}(1)$ -saturated cannot be avoided.

Lemma 3 is the cornerstone in our understanding of which components of $G - X$ are essential. In our applications of the lemma, the graph C consists of a connected component of $G - X$ whose labels encode the adjacency of those vertices to the modulator X . The set \mathcal{Q} contains potential fragments of models of forbidden \mathcal{F} -minors, again labeled by adjacency to X , which we may be interested in destroying in C so that connections through X cannot form \mathcal{F} -minors with fragments that remain in other components of $G - X$. The lemma then essentially says that if it is not possible to select a solution that deletes a minimum number of vertices from C while simultaneously destroying all fragments in \mathcal{Q} , then there is a bounded-size subset of fragments \mathcal{Q}^* that cannot all be destroyed by such a solution. The full importance of Lemma 3 will become clear in Section 4.

Organization. Section 2 provides basic preliminaries. In Section 3, we give some of the main ideas of the proof of Lemma 3. In Section 4 we show how Theorem 1 follows from a procedure that identifies relevant components. We give the procedure and its correctness proof later in the same section, while relying on Lemma 3.

The proof of Lemma 3 is very technical and requires us to develop a framework for analyzing minor models in boundaried labeled graphs. This proof, together with the proofs of other statements marked (★), can be found in the full version [26].

2 Preliminaries

For a positive integer n we use $[n]$ as a shorthand for $\{1, \dots, n\}$. For a set S , let 2^S denote the set of all subsets of S . All graphs we consider are finite, undirected, and simple. A graph G consists of a vertex set $V(G)$ and edge set $E(G) \subseteq \binom{V(G)}{2}$. The open neighborhood of a vertex v is denoted $N_G(v)$. For a vertex set $S \subseteq V(G)$, its open neighborhood is $N_G(S) := \bigcup_{v \in S} N_G(v) \setminus S$. For an edge $\{u, v\}$ in a graph G , *contracting* $\{u, v\}$ results in the graph G' obtained from G by removing u and v , and replacing them by a new vertex w with $N_{G'}(w) = N_G(\{u, v\})$. For a vertex set $S \subseteq V(G)$, we use $G - S$ to denote the graph obtained from G by deleting all vertices in S and their incident edges. The subgraph of G induced by vertex set S is denoted $G[S]$.

► **Definition 4 (treedepth).** Treedepth is defined as follows. The trivial one-vertex graph has treedepth 1. The treedepth of a disconnected graph G with connected components C^1, \dots, C^t is $\max_{i \in [t]} \text{TD}(C^i)$. The treedepth of a connected graph G is $\min_{v \in V(G)} \text{TD}(G - \{v\}) + 1$.

► **Definition 5 (labeled graph).** Let X be a set. An X -labeled graph G is a graph G together with label function $L_G: V(G) \rightarrow 2^X$, assigning a (potentially empty) subset of labels to each vertex in G . The labeled graph G is θ -restricted if each vertex has at most θ labels.

If an edge $\{u, v\}$ is contracted in a labeled graph G to obtain a new vertex w , then the labelset of w is defined as $L_G(u) \cup L_G(v)$.

► **Definition 6 (minor model).** A *minor model* of a graph H in a graph G is a mapping $\varphi: V(H) \rightarrow 2^{V(G)}$ assigning a *branch set* $\varphi(v) \subseteq V(G)$ to each vertex $v \in V(H)$, such that:

- $G[\varphi(v)]$ is nonempty and connected for all $v \in V(H)$,
- $\varphi(v) \cap \varphi(u) = \emptyset$ for all $u \neq v \in V(H)$, and
- if $\{u, v\} \in E(H)$, then there exist $u' \in \varphi(u)$ and $v' \in \varphi(v)$ such that $\{u', v'\} \in E(G)$.

The third condition implies that one can find an *edge mapping* $\psi: E(H) \rightarrow E(G)$ such that:

- For all $\{u, v\} \in E(H)$, edge $\psi(\{u, v\})$ has one endpoint in $\varphi(u)$ and the other in $\varphi(v)$.

We will often use the existence of this edge mapping in our proofs.

For $S \subseteq V(H)$ we define $\varphi(S) := \bigcup_{v \in S} \varphi(v)$, and we define $\varphi(V(H))$ as the *range* of the minor model. A minor model φ of H in G is called *minimal* if no minor model φ' exists with $\varphi'(V(H)) \subsetneq \varphi(V(H))$.

► **Definition 7 (labeled minor model).** A *labeled minor model* of an X -labeled graph H in an X -labeled graph G is a mapping φ as in Definition 6, that additionally satisfies the following: for all $v \in V(H)$ and $\ell \in L_H(v)$ there exists $v' \in \varphi(v)$ such that $\ell \in L_G(v')$.

If G contains a (labeled) minor model of H , then we say that G contains H as a (labeled) minor and denote this as $H \preceq_m G$. Observe that G contains H as a (labeled) minor if and only if H can be obtained from G by deleting edges and vertices (and potentially labels), and contracting edges.

► **Lemma 8 (★).** Let G and H be unlabeled graphs, let $X \subseteq V(G)$, and let φ be a minimal minor model of H in G . Then $\varphi(V(H))$ intersects at most $|X| + |V(H)| + |E(H)|$ connected components of $G - X$.




Figure 1 Two constructions of graphs and sets \mathcal{Q} for $n = 4$, where no optimal \mathcal{F} -deletion breaks \mathcal{Q} , but for any $Q \in \mathcal{Q}$ there exists an optimal \mathcal{F} -deletion breaking $\mathcal{Q} \setminus Q$. Top: any solution breaking both \mathcal{F} and \mathcal{Q} (white vertices at the top) is larger than $\text{OPT}_{\mathcal{F}}$, but for any $Q \in \mathcal{Q}$ there is a solution of size $\text{OPT}_{\mathcal{F}}$ breaking both \mathcal{F} and $\mathcal{Q} \setminus \{Q\}$ (white vertices at the bottom).

We denote the size of an optimal \mathcal{F} -DELETION solution on G by $\text{OPT}_{\mathcal{F}}(G)$, and the set of optimal solutions by $\text{OPTSOL}_{\mathcal{F}}(G)$. In our bounds, we use the notation $\mathcal{O}_z(1)$ for some identifier(s) z to denote a constant that only depends on z .

► **Lemma 9 (★).** *Let \mathcal{F} be a fixed set of (unlabeled) graphs, let $\eta \geq 1$ be a constant, and let X be a set. For any set \mathcal{Q} of X -labeled graphs and host graph C with $\text{TD}(C) \leq \eta$, one can:*

- *compute $\text{OPT}_{\mathcal{F}}(C)$ in $\mathcal{O}_{\mathcal{F},\eta}(|V(C)|)$ time;*
- *determine whether there is a solution $Y \in \text{OPTSOL}_{\mathcal{F}}(C)$ such that $C - Y$ contains no graph from \mathcal{Q} as a labeled minor, in time $f(L, \sum_{H \in \mathcal{Q}} |V(H)|, \eta) \cdot |V(C)|$ for some function f . Here L equals the number of elements of X that appear in the labelset of at least one vertex in at least one graph of \mathcal{Q} .*

3 Overview of the main lemma

In this section we discuss Lemma 3, whose long and technical proof is deferred to the full version. The strength of the lemma comes from the fact that the bound on $|\mathcal{Q}^*|$ is *independent* of the size of the graph C and of the number of labels $|X|$ used on labelsets of vertices of C .

The statement of Lemma 3 is best-possible in several ways. First of all, the dependence of $|\mathcal{Q}^*|$ on $\text{TD}(G)$ instead of $\text{TW}(G)$ is essential. In Figure 1 (left), a construction of a graph of treewidth 2 together with a set \mathcal{Q} is shown. In this graph, no optimal $\{K_2\}$ -deletion (VERTEX COVER) breaks all graphs in \mathcal{Q} . However, for any $Q \in \mathcal{Q}$ there is an optimal vertex cover breaking $\mathcal{Q} \setminus \{Q\}$. The example in Figure 1 can easily be extended to arbitrary n , showing that there is a set \mathcal{Q} with $|\mathcal{Q}| = n$ such that no optimal vertex cover breaks \mathcal{Q} , yet there is no $\mathcal{Q}^* \subsetneq \mathcal{Q}$ such that no optimal vertex cover breaks \mathcal{Q}^* . Since $|\mathcal{Q}|$ is not bounded in terms of $\text{TW}(G) = 2$ and $\mathcal{F} = \{K_2\}$, this shows that $\text{TD}(G)$ cannot be replaced by $\text{TW}(G)$.

Secondly, the assumption that \mathcal{Q} is $(\min_{H \in \mathcal{F}} |V(H)|)$ -saturated cannot be avoided already for $\mathcal{F} = \{K_3\}$ (corresponding to FEEDBACK VERTEX SET). In Figure 1 (right) we show an example of a graph of treedepth 4 and a set \mathcal{Q} of size $2n + 2$ that consist of single vertices of two labels each, where we again cannot properly bound the size of \mathcal{Q}^* . The example is shown for $n = 4$ but can easily be generalized to arbitrary n , without increasing the treedepth. For any $\mathcal{Q}^* \subsetneq \mathcal{Q}$ there exists an optimal \mathcal{F} -deletion breaking \mathcal{Q}^* , while $|\mathcal{Q}|$ is not bounded in terms of $\text{TD}(G)$ and \mathcal{F} .

The proof of Lemma 3 follows an inductive strategy that mimics how a recursive algorithm would solve \mathcal{F} -DELETION on a bounded-treewidth graph C . We pick a vertex v whose removal decreases the treewidth, and branch on whether v is part of the solution or not. If so, we remove v and recurse on a graph of smaller treewidth; if not, then we continue looking for solutions in which v is forbidden to be removed. The process builds up a set S with the property that removing S decreases the treewidth by $|S|$, and we are only interested in solutions disjoint from S . This proceeds while $C - S$ remains connected; the branching depth is bounded since $|S| \leq \text{TD}(C)$. When $C - S$ becomes disconnected, we must take a more involved approach. We recurse on each of the connected components of $C - S$ separately and find \mathcal{F} -DELETION solutions there. But solutions for different components of $C - S$ may not combine into a solution for C , since various fragments of \mathcal{F} -minors left behind in different components of $C - S$, may be combined through their connections to S to form a forbidden minor. For this reason, when we recurse on connected components of $C - S$ we place additional restrictions on the solutions chosen there, to ensure they also break *fragments* of \mathcal{F} -minors in such a way that the solutions can be properly combined.

Our approach to bound the size of \mathcal{Q}^* is built on top of this inductive strategy. While branching over various ways to form an \mathcal{F} -DELETION solution, we additionally branch on what fragments of labeled \mathcal{Q} -minors are left behind by the solution in the various components of $C - S$. By exploiting the saturatedness of \mathcal{Q} in a crucial way, we obtain the desired bound on $|\mathcal{Q}^*|$. The formalization of these ideas requires an extensive theory of how fragments of a forbidden minor in various components of $C - S$ may combine to form a forbidden minor in C , which is developed in Appendix B of the full version of the paper.

4 Kernelization for \mathcal{F} -Deletion

In this section we describe the recursive approach to kernelize the \mathcal{F} -DELETION problem using a constant-treewidth modulator. The correctness of this strategy will crucially depend on Lemma 3. Lemma 10 identifies essential components in the input.

► **Lemma 10.** *Let \mathcal{F} be a finite set of connected graphs and let $\eta \geq 1$ be a constant. There is a polynomial-time algorithm that, given a graph G along with a modulator $X \subseteq V(G)$ such that $\text{TD}(G - X) \leq \eta$, outputs an induced subgraph G' of G together with an integer Δ such that $\text{OPT}_{\mathcal{F}}(G) = \text{OPT}_{\mathcal{F}}(G') + \Delta$ and $G' - X$ has at most $|X|^{\mathcal{O}_{\mathcal{F},\eta}(1)}$ connected components.*

Before proving this lemma, we show how it implies Theorem 1.

► **Theorem 1.** *For every fixed finite set \mathcal{F} of connected graphs and every constant η , the \mathcal{F} -DELETION problem parameterized by a treewidth- η modulator has a polynomial kernelization.*

Proof. Consider an input (G, X, k) to \mathcal{F} -DELETION. The proof is by induction on η .

($\eta = 1$) If $\text{TD}(G - X) = 1$, then $G - X$ is an independent set and any connected component of $G - X$ contains one vertex. Apply Lemma 10 to find an induced subgraph G' of G and integer Δ such that $\text{OPT}_{\mathcal{F}}(G) = \text{OPT}_{\mathcal{F}}(G') + \Delta$, which implies that (G, X, k) has answer YES if and only if $(G', X, k - \Delta)$ has answer YES. Now $G' - X$ has $|X|^{\mathcal{O}_{\mathcal{F},1}(1)}$ single-vertex connected components. It follows that $G' - X$ has at most $|X| + |X|^{\mathcal{O}_{\mathcal{F},1}(1)}$ vertices, which is polynomial in $|X|$ for fixed \mathcal{F} . Hence $(G', X, k - \Delta)$ forms a polynomial kernel.

($\eta > 1$) For $\eta > 1$, we apply Lemma 10 on the input (G, X, k) and find G' and Δ as above. We will augment the modulator X into a superset X' to ensure that $\text{TD}(G' - X') < \eta$. To this end, we consider each connected component C of $G' - X$. If C consists of a single vertex then

its treedepth is already smaller than $\eta > 1$. Otherwise, C is a connected graph with more than one vertex, and by Definition 4 there is a vertex x_C such that $\text{TD}(C - \{x_C\}) < \text{TD}(C)$. Since the **TREEDPTH** problem parameterized by the target width is fixed-parameter tractable [38], and η is a constant, we can find such a vertex x_C by trying all options for x_C and computing the treewidth of the resulting graph in $f(\eta) \cdot n^{\mathcal{O}(1)}$ time. (Alternatively, we can compute a treedepth-decomposition of C using the algorithm of Reidl et al. [38] and take its root as x_C .) We initialize X' as X . For each component C of $G' - X$ with treedepth larger than one, we add the corresponding treedepth-decreasing vertex x_C to X' .

Since Lemma 10 guarantees that the number of connected components of $G' - X$ is polynomial in $|X|$ for fixed \mathcal{F} and η , the resulting modulator X' has size polynomial in $|X|$. Moreover, it guarantees that $\text{TD}(G' - X') < \eta$. Hence we now have an instance $(G', X', k - \Delta)$ of \mathcal{F} -DELETION parameterized by a treedepth- $(\eta - 1)$ modulator, with the same answer as (G, X, k) . We apply the kernel for the parameterization by a treedepth- $(\eta - 1)$ modulator, which outputs an instance (G^*, X^*, k^*) with the same answer as $(G', X', k - \Delta)$ and therefore as (G, X, k) . By induction, the size of G^* is bounded by some polynomial in $|X'|$, which in turn is bounded by a polynomial in $|X|$. Hence G^* has size $|X|^{\mathcal{O}_{\mathcal{F}, \eta}(1)}$ for some suitably chosen constant, and we output (G^*, X^*, k^*) as the result of the kernelization. ◀

Now we prove Lemma 10.

Proof of Lemma 10. Let \mathcal{C} be the connected components of $G - X$. To reduce their number, we have a single reduction rule stated in terms of labeled graphs. With each connected component $C \in \mathcal{C}$, we naturally associate an X -labeled graph C_L by assigning a vertex $v \in V(C)$ the labelset $N_G(v) \cap X$. We are interested in which of these labeled graphs have optimal \mathcal{F} -DELETION solutions that also hit certain fragments of potential \mathcal{F} -minor-models. We therefore define a set \mathcal{H} which is a superset of the relevant fragments. We use $\|\mathcal{F}\|$ as a shorthand for $\max_{H \in \mathcal{F}} |V(H)|$. Let \mathcal{H} consist of the connected $\|\mathcal{F}\|$ -restricted X -labeled graphs that have at most $m_{\mathcal{F}} := \max_{H \in \mathcal{F}} |E(H)|$ edges. We consider two X -labeled graphs to be identical if there is an isomorphism between them that respects the labelsets.

► **Claim 11.** $|\mathcal{H}| \in |X|^{\mathcal{O}_{\mathcal{F}}(1)}$.

Proof. Graphs in \mathcal{H} have at most $m_{\mathcal{F}} + 1$ vertices. There are less than $2^{(m_{\mathcal{F}}+1)^2}$ distinct choices for the graph structure of a member of \mathcal{H} , since there are less than 2^{n^2} different n -vertex graphs. For each vertex, there are less than $(|X| + 1)^{\|\mathcal{F}\|}$ choices for a labelset of size at most $\|\mathcal{F}\|$. Hence each graph structure H can appear with less than $((|X| + 1)^{\|\mathcal{F}\|})^{|V(H)|} \leq (|X| + 1)^{\|\mathcal{F}\| \cdot (m_{\mathcal{F}}+1)}$ different choices of labeling function, giving an overall bound $|\mathcal{H}| \leq 2^{(m_{\mathcal{F}}+1)^2} \cdot (|X| + 1)^{\|\mathcal{F}\| \cdot (m_{\mathcal{F}}+1)}$ that is polynomial in $|X|$. ◻

Choose $\gamma \in \mathcal{O}_{\mathcal{F}, \eta}(1)$ such that Lemma 3 guarantees that for this choice of \mathcal{F} and the treedepth bound η , one can always find $\mathcal{Q}^* \subseteq \mathcal{Q}$ of size at most γ . Let $\rho := |X| + \max_{H \in \mathcal{F}} (|V(H)| + |E(H)|)$, and $\tau := |X| + 1 + \gamma \cdot \rho \in \mathcal{O}_{\mathcal{F}, \eta}(|X|)$. Consider the following marking procedure.

► **Procedure 12.** For each set $\mathcal{Q} \subseteq \mathcal{H}$ of size at most γ , do the following. Let

$$\mathcal{C}_{\mathcal{Q}} := \{C \in \mathcal{C} \mid \forall Y \in \text{OPTSOL}_{\mathcal{F}}(G[C]): C_L - Y \text{ has a graph from } \mathcal{Q} \text{ as a labeled minor}\}.$$

Mark τ arbitrarily chosen components from $\mathcal{C}_{\mathcal{Q}}$, or mark all of them if there are fewer than τ .

Let $\mathcal{C}' \subseteq \mathcal{C}$ denote the marked components, $G' := G[X \cup \bigcup_{C \in \mathcal{C}'} C]$, and let $\Delta := \sum_{C \in \mathcal{C} \setminus \mathcal{C}'} \text{OPT}_{\mathcal{F}}(G[C])$. The procedure can be executed in polynomial time, using variants of Courcelle's theorem to find the sets $\mathcal{C}_{\mathcal{Q}}$. We explain how this is done in Lemma 9. Since $\gamma \in \mathcal{O}_{\mathcal{F}, \eta}(1)$, the number of subsets of \mathcal{H} over which we iterate is polynomial in $|\mathcal{H}|$ and therefore in $|X|$. Since the graphs in \mathcal{Q} are $\|\mathcal{F}\|$ -restricted, the number of labels involved is constant for fixed \mathcal{F} and η , and therefore Lemma 9 guarantees a polynomial running time.

► **Claim 13.** $|\mathcal{C}'| \leq |X|^{\mathcal{O}_{\mathcal{F}, \eta}(1)}$.

Proof. The procedure loops over $|X|^{\mathcal{O}_{\mathcal{F}, \eta}(1)}$ subsets \mathcal{Q} . For each such set, we mark at most $\tau = |X| + 1 + \gamma \cdot \rho \in \mathcal{O}_{\mathcal{F}, \eta}(|X|)$ components. \dashv

The pair (G', Δ) is the desired outcome of Lemma 10. It remains to prove that $\text{OPT}_{\mathcal{F}}(G) = \text{OPT}_{\mathcal{F}}(G') + \Delta$. This follows from Claim 14 by induction.

► **Claim 14.** For any unmarked component $C^* \in \mathcal{C} \setminus \mathcal{C}'$: $\text{OPT}_{\mathcal{F}}(G) = \text{OPT}_{\mathcal{F}}(G - V(C^*)) + \text{OPT}_{\mathcal{F}}(G[C^*])$.

Proof. Let $\widehat{G} := G - V(C^*)$. Clearly, any solution for the graph G can be partitioned into a solution for \widehat{G} and a solution for $G[C^*]$, so that $\text{OPT}_{\mathcal{F}}(G) \geq \text{OPT}_{\mathcal{F}}(\widehat{G}) + \text{OPT}_{\mathcal{F}}(G[C^*])$. We focus on proving the converse. Let $\widehat{Y} \in \text{OPTSOL}_{\mathcal{F}}(\widehat{G})$ be an optimal solution on \widehat{G} . Let $X_0 := X \setminus \widehat{Y}$ and let $\mathcal{H}_0 \subseteq \mathcal{H}$ contain those graphs for which the labelset of each vertex is contained in X_0 . Now define:

$$\mathcal{Q} := \{H \in \mathcal{H}_0 \mid \text{there are fewer than } \rho \text{ components } C \text{ of } \widehat{G} - X \text{ (1)}$$

$$\text{whose } X\text{-labeled version } C_L - \widehat{Y} \text{ contains } H \text{ as } X\text{-labeled minor}\}.$$

Intuitively, one may think of \mathcal{Q} as those labeled graphs (that represent potential fragments of forbidden \mathcal{F} -minors) that can be realized in only few ($\rho \in \mathcal{O}_{\mathcal{F}}(|X|)$) components of $\widehat{G} - X$ after removing the solution \widehat{Y} . When lifting the solution \widehat{Y} in \widehat{G} to a solution in G by adding a solution in C^* , it will be crucial to break all X -labeled minor models of \mathcal{Q} in C^* ; the fragments $\mathcal{H}_0 \setminus \mathcal{Q}$ that remain in *many* different components turn out to be irrelevant.

For a subset $X' \subseteq X_0$ of labels, let $I_{X'}$ be the labeled graph consisting of a single vertex with labelset X' . Let $n_{\mathcal{F}} := \min_{H \in \mathcal{F}} |V(H)|$ and observe that $n_{\mathcal{F}} \leq \rho$. We prove:

$$\forall X' \subseteq X_0, |X'| = n_{\mathcal{F}}: I_{X'} \in \mathcal{Q}. \quad (2)$$

Suppose $I_{X'} \notin \mathcal{Q}$ for suitable X' . Then there are $\rho \geq n_{\mathcal{F}}$ components of $\widehat{G} - X$ that have $I_{X'}$ as labeled minor after removing the solution \widehat{Y} . Take $n_{\mathcal{F}}$ such components $C_L^1, \dots, C_L^{n_{\mathcal{F}}}$, and associate each one to a distinct vertex of $X' \subseteq V(\widehat{G}) \setminus \widehat{Y}$. The fact that $I_{X'}$ is a labeled minor of $C_L^i - \widehat{Y}$ for each i , implies that in each such component there is a connected vertex subset $S_i \subseteq V(C_L^i) \setminus \widehat{Y}$ such that each label of X' appears at least once on a vertex of S_i . Considering the corresponding vertex subset in $\widehat{G} - \widehat{Y}$ and taking into account that the labeling of C_L^i represents adjacency to X in G , this implies that we can contract each S_i into a single vertex s_i that becomes adjacent to all vertices of X' . Then contract each s_i into a distinct vertex of X' : these minor operations on graph $\widehat{G} - \widehat{Y}$ turn X' into a clique of size $n_{\mathcal{F}}$. Hence any graph on $n_{\mathcal{F}}$ vertices is a minor of $\widehat{G} - \widehat{Y}$, contradicting that $\widehat{G} - \widehat{Y}$ is \mathcal{F} -minor-free since \mathcal{F} has a graph on $n_{\mathcal{F}}$ vertices. So (2) holds.

Now consider the unmarked component C^* in the statement of Claim 14, and consider its labeled version C_L^* . We say that a vertex set Y *breaks* the minor models of the X_0 -labeled




Figure 2 This figure shows how to define H_L based on H and \tilde{G} , and how to modify the minor model of H in \tilde{G} such that it uses fewer vertices of C^* , in the proof of (4) in Claim 14.

graphs \mathcal{Q} in C_L^* , or simply *breaks* \mathcal{Q} in C_L^* , if $C_L^* - Y$ does not contain any graph in \mathcal{Q} as a labeled minor. We first show the following.

$$\exists Y^* \in \text{OPTSOL}_{\mathcal{F}}(G[C^*]): Y^* \text{ breaks } \mathcal{Q} \text{ in } C_L^*. \tag{3}$$

To establish (3), assume that no solution of size $\text{OPT}_{\mathcal{F}}(G[C^*])$ in $G[C^*]$ breaks \mathcal{Q} . We will use Lemma 3, together with our marking scheme, to argue for a contradiction. Observe that (2) implies that \mathcal{Q} is an $n_{\mathcal{F}}$ -saturated set of X_0 -labeled graphs. If no optimal solution on $G[C^*]$ breaks \mathcal{Q} , then by Lemma 3 there is a set $\mathcal{Q}^* \subseteq \mathcal{Q}$ of size at most γ such that no optimal solution on $G[C^*]$ breaks \mathcal{Q}^* . Since the assumption that (3) does not hold means that the unmarked C^* was eligible to be marked for the set $\mathcal{C}_{\mathcal{Q}^*}$ in our procedure above, it has marked τ other components $C^1, \dots, C^\tau \in \mathcal{C}_{\mathcal{Q}^*}$ of $G - X$. For each $i \in [\tau]$, there is no \mathcal{F} -DELETION solution of size $\text{OPT}_{\mathcal{F}}(G[C^i])$ in $G[C^i]$ that breaks \mathcal{Q}^* in the labeled version C_L^i . Since $\mathcal{Q}^* \subseteq \mathcal{Q}$, by (1) we have for each graph $H \in \mathcal{Q}^*$ that there are fewer than ρ components C^i among C^1, \dots, C^τ for which $C_L^i - \hat{Y}$ contains H as a labeled minor. Since $|\mathcal{Q}^*| \leq \gamma$, it follows that there are at most $\gamma \cdot \rho$ indices $i \in [\tau]$ for which $C_L^i - Y$ contains some graph from \mathcal{Q}^* as a labeled minor. But since $\tau = |X| + 1 + \gamma \cdot \rho$, there are at least $|X| + 1$ components C_L^i in which all \mathcal{Q}^* -minors are broken by \hat{Y} . Since no *optimal* solution breaks \mathcal{Q}^* in the marked components, we have $|\hat{Y} \cap V(C^i)| > \text{OPT}_{\mathcal{F}}(G[C^i])$ for at least $|X| + 1$ components. But this contradicts that \hat{Y} is an optimal solution to \mathcal{F} -DELETION on \hat{G} : since \mathcal{F} consists of connected graphs, we can form a solution \hat{Y}' by taking X together with a set of size $\text{OPT}_{\mathcal{F}}(\hat{G}[C])$ from each component C of $\hat{G} - X$. Since $|\hat{Y}' \cap V(C)| \leq |\hat{Y} \cap V(C)|$ for all $C \in \mathcal{C}$, with strict inequality for at least $|X| + 1$ components, we have $|\hat{Y}'| < |\hat{Y}|$. This contradicts that \hat{Y} is an optimal solution and establishes (3).

Hence there exists a solution Y^* in C_L^* breaking \mathcal{Q} of size $\text{OPT}_{\mathcal{F}}(G[C^*])$. We prove:

$$\hat{Y} \cup Y^* \text{ is a solution to } \mathcal{F}\text{-DELETION on } G. \tag{4}$$

This will complete the proof of Claim 14, since $|\hat{Y} \cup Y^*| = \text{OPT}_{\mathcal{F}}(\hat{G}) + \text{OPT}_{\mathcal{F}}(G[C^*])$. Assume for a contradiction that $\tilde{G} := G - (\hat{Y} \cup Y^*)$ contains some graph $H \in \mathcal{F}$ as a minor. Consider a minimal minor model of H in \tilde{G} , which is given by a vertex mapping $\varphi: V(H) \rightarrow 2^{V(\tilde{G})}$, and let $\psi: E(H) \rightarrow E(\tilde{G})$ be a corresponding edge mapping.

Out of all possible minimal minor models of H in \tilde{G} , select a model (φ, ψ) that minimizes the quantity $|\varphi(V(H)) \cap V(C^*)|$. Observe that if $\varphi(V(H)) \cap V(C^*) = \emptyset$, then φ is also a valid model in $\hat{G} - \hat{Y}$, contradicting that \hat{Y} is a solution to \mathcal{F} -DELETION on \hat{G} . So in the remainder we consider the case that the minor model contains at least one vertex of C^* . We will build a minimal minor model of H in G using strictly fewer vertices of C^* , thereby contradicting the choice of (φ, ψ) . Consider the X_0 -labeled subgraph H'_L of \tilde{G} obtained by the following procedure, which is illustrated in Figure 2:

1. Start from the X_0 -labeled subgraph of \tilde{G} induced by $\bigcup_{v \in V(H)} \varphi(v) \cap V(C^*)$, where each vertex u has labelset $N_G(u) \cap X_0$. As observed above, this subgraph is not empty.

2. Remove all edges from this subgraph, except those in the range of ψ and those that connect two vertices that belong to a common branch set under φ .
3. Contract every edge between two vertices that belong to a common branch set of φ , obtaining an X_0 -labeled graph H'_L . (Recall that labelsets merge during edge contraction.) Observe that H'_L has at most $|E(H)|$ edges, since each edge remaining in H'_L corresponds to an edge in the range of ψ . We claim that H'_L is an $n_{\mathcal{F}}$ -restricted graph: the labelset of each vertex has size less than $n_{\mathcal{F}}$. To see this, observe that if some vertex of H'_L has a labelset $X' \subseteq X_0$ of size at least $n_{\mathcal{F}}$, then the pre-image of this vertex corresponds to a connected vertex subset A of $\varphi(V(H)) \cap V(C^*)$ such that $|N_G(A) \cap X_0| \geq n_{\mathcal{F}}$. Since (φ, ψ) is a minor model in $\tilde{G} = G - (\hat{Y} \cup Y^*)$, this would imply that $C_L^* - Y^*$ has the one-vertex graph $I_{X'}$ with labelset X' as a labeled minor. But $I_{X'} \in \mathcal{Q}$ by (2), while Y^* breaks all labeled \mathcal{Q} -minors in C_L^* by definition; a contradiction. Hence H'_L is indeed $n_{\mathcal{F}}$ -restricted.

Let H_L be an arbitrary connected component of H'_L . Since H_L is connected, $n_{\mathcal{F}}$ -restricted, and contains at most $|E(H)|$ edges, we have $H_L \in \mathcal{H}_0$. As H_L clearly occurs as a labeled minor of $C_L^* - Y^*$, while Y^* breaks \mathcal{Q} in C_L^* , we have $H_L \notin \mathcal{Q}$. By definition of \mathcal{Q} , this implies there are at least ρ connected components C^1, \dots, C^ρ of $\tilde{G} - X$ such that $C_L^i - \hat{Y}$ contains H_L as X_0 -labeled minor for each $i \in [\rho]$. By Lemma 8, the minimal model (φ, ψ) in \tilde{G} intersects at most $|X| + |V(H)| + |E(H)| \leq \rho$ components of $\tilde{G} - X$ and therefore of $G - X$. Since $\varphi(V(H))$ also intersects $C^* \setminus \{C^1, \dots, C^\rho\}$, it follows that some C^i is disjoint from the range of (φ, ψ) .

To finish the argument, fix C^i such that $\varphi(V(H)) \cap V(C^i) = \emptyset$ and $C_L^i - \hat{Y}$ contains H_L as X_0 -labeled minor. Let T denote the vertices of $\varphi(V(H)) \cap V(C^*)$ whose contraction in the process above resulted in the connected component H_L of H'_L . Then it is straightforward to verify that $G[(\varphi(V(H)) \setminus T) \cup (C^i - \hat{Y})]$ contains H as a minor. The role that vertices of T played in the minor model (φ, ψ) can be replaced by the vertices of $C_L^i - \hat{Y}$: each edge of ψ that was realized between vertices of T yielded an edge of H_L which is realized by a labeled H_L -minor in $C_L^i - \hat{Y}$; each fragment of a branch set that was realized within C^* yielded a vertex of H_L that is realized in the H_L -minor in $C_L^i - \hat{Y}$; and finally the connectivity of the branch sets is ensured because the labeling ensures that for all fragments of branch sets in T that were adjacent to vertices of $X - \hat{Y} = X_0$, the branch set of the H_L -minor in $C^i - \hat{Y}$ realizing that fragment is also adjacent to all those vertices of X_0 . Hence there is a *minimal* H -minor in \tilde{G} whose range is a subset of $(\varphi(V(H)) \setminus T) \cup (C^i - \hat{Y})$. Since $T \subseteq C^*$ is not empty, this contradicts our choice of (φ, ψ) as a minimal H -model minimizing the intersection with C^* . \lrcorner

This concludes the proof of Lemma 10. \blacktriangleleft

5 Conclusion

Our goal in this paper was to obtain polynomial kernelizations for a wide range of graph problems, in terms of a rich class of structural parameterizations. We obtained polynomial kernelizations for \mathcal{F} -DELETION problems parameterized by a constant-treewidth modulator. The kernelization algorithm as presented here is only of theoretical interest. While the kernel size is polynomial for fixed \mathcal{F} and η , the degree of the polynomial grows very quickly with \mathcal{F} and η . It would be desirable to have a *uniformly polynomial* kernel size, of the form $f(\mathcal{F}, \eta)|X|^c$ for some constant c and function f . Unfortunately, Theorem 2 shows that even for the simplest choice of \mathcal{F} , corresponding to the VERTEX COVER problem, the degree of the polynomial must depend exponentially on η and no uniformly polynomial kernelization exists. The bad news also extends in the other direction: when taking the

simplest choice for η and working with a treedepth-one modulator (a vertex cover), the degree of the polynomial in the kernel size for \mathcal{F} -DELETION must depend on \mathcal{F} [22, Theorem 1.1] and a uniformly-polynomial kernel does not exist.

References

- 1 Hans L. Bodlaender. Kernelization: New upper and lower bound techniques. In *Proc. 4th IWPEC*, pages 17–37, 2009. doi:10.1007/978-3-642-11269-0_2.
- 2 Hans L. Bodlaender, Rodney G. Downey, Michael R. Fellows, and Danny Hermelin. On problems without polynomial kernels. *J. Comput. Syst. Sci.*, 75(8):423–434, 2009. doi:10.1016/j.jcss.2009.04.001.
- 3 Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. (Meta) Kernelization. *J. ACM*, 63(5):44:1–44:69, 2016. doi:10.1145/2973749.
- 4 Marin Bougeret and Ignasi Sau. How much does a treedepth modulator help to obtain polynomial kernels beyond sparse graphs? *CoRR*, abs/1609.08095, 2016. URL: <http://arxiv.org/abs/1609.08095>.
- 5 Marin Bougeret and Ignasi Sau. How much does a treedepth modulator help to obtain polynomial kernels beyond sparse graphs? In *Proc. 12th IPEC (2017)*, volume 89, pages 10:1–10:13, 2018. doi:10.4230/LIPIcs.IPEC.2017.10.
- 6 Hubie Chen and Moritz Müller. One hierarchy spawns another: Graph deconstructions and the complexity classification of conjunctive queries. In *Proc. CSL-LICS 2014*, pages 32:1–32:10. ACM, 2014. doi:10.1145/2603088.2603107.
- 7 Jianer Chen, Iyad A. Kanj, and Weijia Jia. Vertex cover: Further observations and further improvements. *J. Algorithms*, 41(2):280–301, 2001. doi:10.1006/jagm.2001.1186.
- 8 Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.
- 9 Marek Cygan, Łukasz Kowalik, and Marcin Pilipczuk. Open problems from worker 2013, the workshop on kernels, April 2013. URL: <http://worker2013.mimuw.edu.pl/slides/worker-opl.pdf>.
- 10 Marek Cygan, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. On the hardness of losing width. *Theory Comput. Syst.*, 54(1):73–82, 2014. doi:10.1007/s00224-013-9480-1.
- 11 Holger Dell and Dieter van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. *J. ACM*, 61(4):23:1–23:27, 2014. doi:10.1145/2629620.
- 12 Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013. doi:10.1007/978-1-4471-5559-1.
- 13 Andrew Drucker. New limits to classical and quantum instance compression. In *Proc. 53rd FOCS*, pages 609–618, 2012. doi:10.1109/FOCS.2012.71.
- 14 Michael Elberfeld, Martin Grohe, and Till Tantau. Where first-order and monadic second-order logic coincide. *ACM Trans. Comput. Log.*, 17(4):25:1–25:18, 2016. doi:10.1145/2946799.
- 15 Henning Fernau. Kernelization, Turing kernels. In *Encyclopedia of Algorithms*, pages 1043–1045. Springer, 2016. doi:10.1007/978-1-4939-2864-4_528.
- 16 Fedor V. Fomin, Bart M. P. Jansen, and Michał Pilipczuk. Preprocessing subgraph and minor problems: When does a small vertex cover help? *J. Comput. Syst. Sci.*, 80(2):468–495, 2014. doi:10.1016/j.jcss.2013.09.004.

- 17 Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, Geevarghese Philip, and Saket Saurabh. Hitting forbidden minors: Approximation and kernelization. In *Proc. 28th STACS*, pages 189–200, 2011. doi:10.4230/LIPIcs.STACS.2011.189.
- 18 Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Planar \mathcal{F} -Deletion: Approximation, kernelization and optimal FPT algorithms. In *Proc. 53rd FOCS*, pages 470–479, 2012. doi:10.1109/FOCS.2012.62.
- 19 Fedor V. Fomin and Torstein J. F. Strømme. Vertex cover structural parameterization revisited. In *Proc. 42nd WG*, volume 9941 of *LNCS*, pages 171–182, 2016. doi:10.1007/978-3-662-53536-3_15.
- 20 Lance Fortnow and Rahul Santhanam. Infeasibility of instance compression and succinct PCPs for NP. *J. Comput. Syst. Sci.*, 77(1):91–106, 2011. doi:10.1016/j.jcss.2010.06.007.
- 21 Jakub Gajarský, Petr Hliněný, Jan Obdržálek, Sebastian Ordyniak, Felix Reidl, Peter Rossmanith, Fernando Sánchez Villaamil, and Somnath Sikdar. Kernelization using structural parameters on sparse graph classes. *J. Comput. Syst. Sci.*, 84:219–242, 2017. doi:10.1016/j.jcss.2016.09.002.
- 22 Archontia C. Giannopoulou, Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. Uniform kernelization complexity of hitting forbidden minors. *ACM Trans. Algorithms*, 13(3):35:1–35:35, 2017. doi:10.1145/3029051.
- 23 Jiong Guo and Rolf Niedermeier. Invitation to data reduction and problem kernelization. *SIGACT News*, 38(1):31–45, 2007. doi:10.1145/1233481.1233493.
- 24 Gregory Gutin. Kernelization, constraint satisfaction problems parameterized above average. In *Encyclopedia of Algorithms*, pages 1011–1013. Springer, 2016. doi:10.1007/978-1-4939-2864-4_524.
- 25 Bart M. P. Jansen and Hans L. Bodlaender. Vertex cover kernelization revisited - Upper and lower bounds for a refined parameter. *Theory Comput. Syst.*, 53(2):263–299, 2013. doi:10.1007/s00224-012-9393-4.
- 26 Bart M. P. Jansen and Astrid Pieterse. Polynomial kernels for hitting forbidden minors under structural parameterizations. *CoRR*, abs/1804.08885, 2018. arXiv:1804.08885.
- 27 Bart M. P. Jansen, Venkatesh Raman, and Martin Vatshelle. Parameter ecology for feedback vertex set. *Tsinghua Science and Technology*, 19(4):387–409, 2014. doi:10.1109/TST.2014.6867520.
- 28 Eun Jung Kim, Alexander Langer, Christophe Paul, Felix Reidl, Peter Rossmanith, Ignasi Sau, and Somnath Sikdar. Linear kernels and single-exponential algorithms via protrusion decompositions. *ACM Trans. Algorithms*, 12(2):21:1–21:41, 2016. doi:10.1145/2797140.
- 29 Stefan Kratsch. Recent developments in kernelization: A survey. *Bulletin of the EATCS*, 113:58–97, 2014.
- 30 Stefan Kratsch. A randomized polynomial kernelization for vertex cover with a smaller parameter. In *Proc. 24th ESA*, volume 57 of *LIPIcs*, pages 59:1–59:17, 2016. doi:10.4230/LIPIcs.ESA.2016.59.
- 31 Stefan Kratsch and Magnus Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. In *Proc. 53rd FOCS*, pages 450–459, 2012. doi:10.1109/FOCS.2012.46.
- 32 Diptapriyo Majumdar. Structural parameterizations of feedback vertex set. In *Proc. 11th IPEC*, volume 63 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:16, 2017. doi:10.4230/LIPIcs.IPEC.2016.21.
- 33 Diptapriyo Majumdar, Venkatesh Raman, and Saket Saurabh. Kernels for structural parameterizations of vertex cover - case of small degree modulators. In *Proc. 10th IPEC*, volume 43 of *LIPIcs*, pages 331–342, 2015. doi:10.4230/LIPIcs.IPEC.2015.331.

- 34 Neeldhara Misra. Kernelization, planar \mathcal{F} -Deletion. In *Encyclopedia of Algorithms*, pages 1033–1036. Springer, 2016. doi:10.1007/978-1-4939-2864-4_527.
- 35 G.L. Nemhauser and L.E. Trotter (jr.). Vertex packings: structural properties and algorithms. *Math. Program.*, 8:232–248, 1975. doi:10.1007/BF01580444.
- 36 J. Nešetřil and P. Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and Combinatorics*. Springer, 2012. doi:10.1007/978-3-642-27875-4.
- 37 Michał Pilipczuk and Marcin Wrochna. On space efficiency of algorithms working on structural decompositions of graphs. In *Proc. 33rd STACS*, volume 47 of *LIPIcs*, pages 57:1–57:15, 2016. doi:10.4230/LIPIcs.STACS.2016.57.
- 38 Felix Reidl, Peter Rossmanith, Fernando Sanchez Villaamil, and Somnath Sikdar. A faster parameterized algorithm for treedepth. In *Proc. 41st ICALP*, pages 931–942, 2014. doi:10.1007/978-3-662-43948-7_77.
- 39 Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986. doi:10.1016/0095-8956(86)90030-4.
- 40 Stéphan Thomassé. A $4k^2$ kernel for feedback vertex set. *ACM Trans. Algorithms*, 6(2), 2010. doi:10.1145/1721837.1721848.