Uniform Polynomial Kernel for Deletion to $K_{2,p}$ Minor-Free Graphs

William Lochet
LIRMM, Université de Montpellier, CNRS, Montpellier, France

Roohani Sharma
Department of Informatics, University of Bergen, Norway

Abstract

In the $\mathcal{F}$-Deletion problem, where $\mathcal{F}$ is a fixed finite family of graphs, the input is a graph G and an integer k, and the goal is to determine if there exists a set of at most k vertices whose deletion results in a graph that does not contain any graph of $\mathcal{F}$ as a minor. The $\mathcal{F}$-Deletion problem encapsulates a large class of natural and interesting graph problems like Vertex Cover, Feedback Vertex Set, Treewidth-η Deletion, Treedepth-η Deletion, Pathwidth-η Deletion, Outerplanar Deletion, Vertex Planarization and many more. We study the $\mathcal{F}$-Deletion problem from the kernelization perspective. In a seminal work, Fomin et al. [FOCS 2012] gave a polynomial kernel for this problem when the family $\mathcal{F}$ contains at least one planar graph. The asymptotic growth of the size of the kernel is not uniform with respect to the family $\mathcal{F}$: that is, the size of the kernel is $k^{f(\mathcal{F})}$, for some function f that depends only on $\mathcal{F}$. Later Giannopoulou et al. [TALG 2017] showed that the non-uniformity in the kernel size bound is unavoidable as Treewidth-η Deletion cannot admit a kernel of size $\mathcal{O}(k^{\frac{\eta+1}{2}-\epsilon})$, for any $\epsilon > 0$, unless NP $\subseteq$ coNP/poly. On the other hand it was also shown that Treedepth-η Deletion admits a uniform kernel of size $f(\mathcal{F}) \cdot k^6$ depicting that there are subclasses of $\mathcal{F}$ where the asymptotic kernel sizes do not grow as a function of the family $\mathcal{F}$. This work led to the question of determining classes of $\mathcal{F}$ where the problem admits uniform polynomial kernels.

In this paper, we show that if all the graphs in $\mathcal{F}$ are connected and $\mathcal{F}$ contains $K_{2,p}$ (a bipartite graph with 2 vertices on one side and p vertices on the other), then the problem admits a uniform kernel of size $f(\mathcal{F}) \cdot k^{10}$. The graph $K_{2,p}$ is one natural extension of the graph θ_p, where θ_p is a graph on two vertices and p parallel edges. The case when $\mathcal{F}$ contains θ_p has been studied earlier and serves as (the only) other example where the problem admits a uniform polynomial kernel.

2012 ACM Subject Classification Theory of computation → Fixed parameter tractability

Keywords and phrases Uniform polynomial kernel, $\mathcal{F}$-minor-free deletion, complete bipartite minor-free graphs, $K_{2,p}$, protrusions

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2024.46

1 Introduction

For any fixed finite family of (multi-)graphs $\mathcal{F}$, in the $\mathcal{F}$-Deletion problem, given as input a graph G and a positive integer k, the task is to determine whether the deletion of a set of at most k vertices results in a graph that does not contain any graph of $\mathcal{F}$ as a minor. The $\mathcal{F}$-Deletion problem encompasses various natural and interesting problems such as Vertex Cover, Feedback Vertex Set, Treewidth-η Deletion, Treedepth-η Deletion, Pathwidth-η Deletion, Outerplanar Deletion, Vertex Planarization and much more. As a result of the seminal work of Lewis and Yannakakis [18] the problem is known to be NP-complete. By a celebrated result of Robertson and Seymour [22] every $\mathcal{F}$-Deletion problem is non-uniformly FPT, that is, for every integer k, there exists an algorithm that solves the problem in $f(k) \cdot n^3$ time, where n is the number of vertices in the input graph. However, when the family $\mathcal{F}$ is given explicitly, the problem is uniformly FPT because the excluded minors for the graphs that are YES instances of the problem can be computed

35th International Symposium on Algorithms and Computation (ISAAC 2024).
Editors: Julián Mestre and Anthony Wirth; Article No. 46; pp. 46:1–46:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

explicitly from the result by Adler et al. [1]. Another breakthrough result by Fomin et al. [12] shows that the problem admits an algorithm with running time $2^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$ when all the graphs in $\mathcal{F}$ are connected and $\mathcal{F}$ contains a planar graph. The class considered by Fomin et al. seems a little restrictive at first, but it already encapsulates the classical problems mentioned above except for Vertex Planarization. In fact, the class of problems considered by Fomin et al. [12] are essentially about deleting k vertices to get to a graph of constant treewidth, since graphs that exclude a planar graph H as a minor have treewidth at most $|V(H)|^{\mathcal{O}(1)}$ [3].

One of the major highlights of the result by Fomin et al. [12] is also a polynomial (in k) kernel for the $\mathcal{F}$-Deletion problem when $\mathcal{F}$ contains a planar graph. A noteworthy feature of their kernelization algorithm is that the size of the kernel is $f(\mathcal{F}) \cdot k^{g(\mathcal{F})}$, for some functions f, g that depend only on the family $\mathcal{F}$. In particular, the exponent of the size of the kernel depends on the family $\mathcal{F}$. Such kernels are called *non-uniform* kernels as the asymptotic size of the kernel varies with the family $\mathcal{F}$. Such a result opens up questions about the existence of a kernel of size $f(\mathcal{F}) \cdot k^{\mathcal{O}(1)}$ for $\mathcal{F}$-Deletion. Such a kernel is called a *uniform polynomial kernel.*

Soon after the result of Fomin et al., Giannopoulou et al. [14] showed that the size of the polynomial kernel by Fomin et al. is essentially tight in the sense that $\mathcal{F}$-Deletion cannot admit a kernel of size $f(\mathcal{F}) \cdot k^{\mathcal{O}(1)}$, under reasonable complexity assumptions. In particular, they showed that Treewidth-η Deletion, a special case of $\mathcal{F}$-Deletion where $\mathcal{F}$ contains a planar graph, cannot admit a uniform polynomial kernel unless NP $\subseteq$ coNP/poly, even when the parameter is the vertex cover of the graph. More specifically they showed that Treewidth-η Deletion cannot admit a kernel on $\mathcal{O}(x^{\frac{\eta+1}{2}-\epsilon})$ vertices, for any $\epsilon > 0$, where x is the size of the vertex cover in the input graph. They also complemented this result by showing that Treedepth-η Deletion, another special case of $\mathcal{F}$-Deletion when $\mathcal{F}$ contains a path, admits a (uniform) kernel of size $f(\eta) \cdot k^6$ for some function f.

Other than Treedepth-η Deletion, the only other family $\mathcal{F}$ for which $\mathcal{F}$-Deletion is known to admit a uniform polynomial kernel (of size $f(\mathcal{F}) \cdot k^2 \log^{3/2} k)$) is when the graph $\theta_p \in \mathcal{F}$ [11]. Here θ_p is a graph with two distinct vertices and p parallel edges between them.

This contrast in the behaviour of the asymptotic size of polynomial kernels, obtained for different specializations of $\mathcal{F}$-Deletion, leads to the question- *under what restrictions of $\mathcal{F}$, does the $\mathcal{F}$-Deletion problem admit uniform kernels?* Our study investigates this direction and exhibits an infinite collection of families $\mathcal{F}$ for which the $\mathcal{F}$-Deletion problem admits a uniform polynomial kernel.

Our Result. We show that $\mathcal{F}$-Deletion admits a kernel of size $f(\mathcal{F}) \cdot k^{10}$, when all the graphs in $\mathcal{F}$ are connected and $K_{2,p} \in \mathcal{F}$, where $K_{2,p}$ is a complete bipartite graph on 2 vertices on one side and p vertices on the other. Henceforth, for any positive integer p, let $\mathcal{F}_p$ denote an arbitrary finite family of connected graphs such that $K_{2,p} \in \mathcal{F}_p$. The $\mathcal{F}_p$-Deletion problem is formally defined as follows: given a graph G and an integer k, does there exist $X \subseteq V(G)$, $|X| \leq k$ such that $G - X$ has no graph of $\mathcal{F}_p$ as a minor?

In the remaining paper we subsume the factors depending on $\mathcal{F}_p$ in the $\mathcal{O}(\cdot)$ notation. Also, a polynomial running time refers to a running time that is polynomial in the input size where the exponent of the polynomial is an absolute constant (and hence does not depend on $\mathcal{F}_p$ or k). Thus, our kernelization algorithm runs in "purely" polynomial time.

▶ **Theorem 1.1.** *$\mathcal{F}_p$-Deletion admits a kernel of size $\mathcal{O}(k^{10})$.*

$K_{2,p}$-free graphs. The class $\mathcal{F}_p$-Deletion that we consider is, *and has to be*, more restrictive than what is considered by Fomin et al. [12], given the hardness result by Giannopoulou et al. [14]. In the following points we motivate our interests in the study of the chosen family $\mathcal{F}_p$.

(1) Generalizes outerplanarity. A graph is called *outerplanar* if there exists a planar embedding of it where all the vertices lie on the outer face. In the Outerplanar Deletion problem, given as input a graph G and an integer k the goal is to decide if the deletion of at most k vertices results in an outerplanar graph. A well-known consequence of Wagner's characterization of planar graphs implies that the Outerplanar Deletion problem is equivalent to the $\mathcal{F}$-Deletion problem where $\mathcal{F} = \{K_{2,3}, K_4\}$, where K_4 is a complete graph on 4 vertices. Clearly, $\mathcal{F}_p$-Deletion encapsulates and generalizes the Outerplanar Deletion problem.

(2) Challenge in extension from θ_p. As mentioned earlier, prior to the polynomial kernelization result of Fomin et al. [12] for general families $\mathcal{F}$ containing some planar graph, Fomin et al. [11] gave a uniform polynomial kernel for $\mathcal{F}$-Deletion when $\theta_p \in \mathcal{F}$. Such families already encapsulate classical problems like Vertex Cover, Feedback Vertex Set and Diamond Hitting [11].

Observe that $K_{2,p}$ is a natural extension of θ_p as it can be obtained from θ_p by subdividing each of its edges once. This seemingly simple extension of θ_p already poses great technical challenges, thereby disallowing to lift the kernelization techniques used in the θ_p case to the $K_{2,p}$ case. As we describe in detail later (see Section 3), the challenge in making a uniform polynomial kernel for special cases of $\mathcal{F}$-Deletion lie in what we call the *degree reduction phase* of [12]. We elaborate on this later but let us give some overview of it already here. Let S be some approximate solution to the problem of size $k^{\mathcal{O}(1)}$. Let C be some connected component of $G - S$. If one can bound the degree of each vertex of $x \in S$ in the set C by $f(\mathcal{F}) \cdot k^{\mathcal{O}(1)}$ (where the degree of k is independent of $\mathcal{F}$), then following the approach of [12], one can get a *uniform* polynomial kernel for the $\mathcal{F}$-Deletion problem.

Using the above, a uniform polynomial kernel for $\mathcal{F}$-Deletion when $\theta_p \in \mathcal{F}$ follows very easily: let S be a 1-redundant solution to the problem of size $\mathcal{O}(k^2)$, that is for each $x \in S$, $S \setminus x$ is a solution to the problem. As we will see later such sets of $\mathcal{O}(k^2)$ size can be found easily. Let C be a connected component of $G - S$. Then for any $x \in S$, the degree of x in C is at most $p - 1$, as otherwise there would be θ_p as a minor in $G[C \cup \{x\}]$, contradicting that $S \setminus \{x\}$ is a solution. Thus, in this case one can in fact, bound the degree of x in C by $\mathcal{O}(1)$.

The above simple argument for bounding the degree of x fails completely when the forbidden minor is a subdivided θ_p, that is a $K_{2,p}$. For example, consider a graph containing $n + 1$ vertices, where one vertex is adjacent to all the other n vertices and these n vertices are connected to form a path. This graph has θ_n as a minor but no $K_{2,3}$ as a minor.

(3) Interesting structural graph properties of $K_{2,p}$-free graphs. From a graph theoretic viewpoint, excluding certain classes of graphs as minors seem to give close connections to some interesting graph properties. One of the most interesting conjectures at present demonstrating this is the Hadwiger's Conjecture which states that the chromatic number of any graph that avoids K_t as a minor is at most $t - 1$. Following this line of work, graph theorists have developed a special interest in the class of graphs that exclude a *complete bipartite graph* as a minor [10, 7, 23, 5, 20, 6, 9]. Together with connectivity requirements, and possibly other assumptions, graphs with no $K_{q,p}$ as a minor can be shown to have interesting properties relating to toughness, hamiltonicity, and other traversability properties [4, 5, 21]. Particular attention has been given to the case when $q = 2$, as most of these properties appear to hold for this special case too. Note that any graph avoiding $K_{2,p}$ as a minor is at

most $p-1$-connected. Results in the literature show that 3-connected $K_{2,4}$-free graphs are Hamiltonian and the 2-connected $K_{2,4}$-free graphs have a Hamiltonian Path [10]. Also [7] shows that every 2-connected $K_{2,4}$-free graph contains two vertices whose deletion results in an outerplanar graph. Graphs on n vertices that are $K_{2,p}$-free are known to have long cycles, in particular, cycles of length at least n/p^{p-1} [5]. Another result shows that the number of edges in an n-vertex $K_{2,p}$-free graph, for $p \geq 2$, is at most $(1/2)(p+1)(n-1)$ [6]. This literature thus suggests that the class of $K_{2,p}$-free graphs exhibit interesting graph theoretic properties and hence, it could be worth to study this graph class algorithmically.

(4) Extremal limit before encapsulating planarization. In continuation of the above point, at the front of avoiding a $K_{q,p}$ as a minor, it must be noted that the case when $q = 3$ already encompasses the classical and notorious Vertex Planarization problem. In this problem, the goal is to delete at most k vertices such that the resulting graph is a planar graph. This problem is equivalent to $\mathcal{F}$-Deletion when $\mathcal{F} = \{K_{3,3}, K_5\}$ because of Wagner's characterization of planar graphs. Note that none of the graphs in $\mathcal{F}$ are planar. The first constructive FPT algorithm for this problem was given by Marx and Schlotter in 2007 [19], which was followed by an improved algorithm by Kawarabayashi [17]. This was followed by the current best known algorithm for the problem by Jansen et al. [16] in 2014 that runs in time $2^{\mathcal{O}(k \log k)} \cdot n$. Over this long period of improvements, one important open question that has intrigued the community is the question about the existence of a polynomial kernel for the problem. Recently Jansen and Wlodarczyk [15] gave a lossy polynomial kernel for this problem. But it seems for now that getting a (non-lossy) polynomial kernel for this problem may require more novel ideas. Thus, on the front of avoiding complete bipartite minors, avoiding anything beyond $K_{2,p}$ must first confront the Vertex Planarization problem.

Roadmap. In Section 2 we define basic notations and definitions. In Section 3 we describe all the (five) steps of our kernelization algorithm. In particular, we state formally the five main lemmas that we prove to give the complete proof of Theorem 1.1. In Section 4 we focus on our main technical contribution of this work which is what we call the degree reduction phase of the kernelization algorithm (step 2 of the 5 steps). We give an overview of the key ideas of this phase, followed by formal proofs. In Section 5 we conclude with some open questions. The details of the 4 other phases of the algorithm that are described in Section 3 have been omitted because of space constraints. Also the proofs of lemmas marked with $\star$ have been omitted due to space constraints.

2 Preliminaries

For standard notations and terminology that is not defined here, we refer to [8]. For the definition of kernelization and related terminology we refer to the book [13]. Throughout the paper, $h = \max_{H \in \mathcal{F}_p} |V(H)|$.

General. For positive integers $i < j$, $[i]$ denotes the set $\{1, \ldots, i\}$ and $[i, j]$ denote the set $\{i, i+1, \ldots, j\}$. Given a sequence, an *interval* is a set of consecutive entries of the sequence. The *length of the interval* is the number of entries in it.

Graphs. For $u, v \in V(G)$ a (u, v)-path in G is a path from u to v. The *internal* vertices of a path are the vertices of the path that are not its end-points. For $X \subseteq V(G)$, a path is called *X-free* if none of the vertices of X appear as internal vertices of the path. For $X, Y \subseteq V(G)$, $Z \subseteq V(G)$ is called an *(X, Y)-cut* if $G - Z$ has no path from a vertex of X to a vertex of Y. When X or Y are singletons we drop the braces around them in this notation.

Boundaried graphs and boundaried minors. For any positive integer t, a *t-boundaried graph* is a graph G together with a specially assigned vertex set of size at most t called the boundary of G. Also each vertex in the boundary is labelled with a distinct integer from $[t]$. A t-boundaried graph H is a *minor* of a t-boundaried graph G if H can be obtained from G by deleting vertices or edges or contracting edges, but never contracting edges with both endpoints being the vertices in the boundary. If we contract an edge between a boundary vertex u and a non-boundary vertex v, the resulting vertex is a boundary vertex with the same label as that of u. For a t-boundaried graph G, for any positive integer δ, the *δ-folio*(G) is the set of all t-boundaried graphs of size at most δ that can be obtained as boundaried minors in G.

3 The steps of the kernelization algorithm

The kernelization algorithm of Theorem 1.1 has five phases which we term as (1) redundant solution, (2) degree reduction, (3) component reduction, (4) protrusion decomposition, and (5) protrusion replacement. We emphasize here that the degree reduction phase is our main technical contribution of this work and our special choice of starting with a redundant solution in phase one is key for the second phase. Other than this, the overall structure of our kernelization algorithm follows the footprints of that of [12]. For any graph G, let $\mathtt{cc}(G)$ denote the set of connected components of G. Throughout the presentation, $h := \max_{F \in \mathcal{F}_p} |V(F)|$. Let (G, k) be an instance of $\mathcal{F}_p$-Deletion problem.

Redundant solution. In the *first* phase, we find a 1-redundant solution S in G of size $\mathcal{O}(k^2)$. A 1-*redundant solution* in G is a set of vertices S such that for every $x \in S$, $S \setminus \{x\}$ is a solution of $\mathcal{F}_p$-Deletion. This step is different from that of [12] where any solution modulator (for example, any approximate solution) works. Formally, we prove the following lemma.

▶ **Lemma 3.1** ($\star$, Redundant solution). *Given an instance (G, k) of $\mathcal{F}_p$-Deletion, there is a polynomial-time algorithm that either outputs $x \in V(G)$ such that $(G - \{x\}, k-1)$ is an equivalent instance of (G, k), or outputs a 1-redundant solution of (G, k) of size $\mathcal{O}(k^2)$, or concludes correctly that (G, k) is a* NO *instance of $\mathcal{F}_p$-Deletion.*

If Lemma 3.1 outputs $x \in V(G)$ then we apply a reduction rule and output the instance $(G - \{x\}, k-1)$ as an equivalent instance. If it outputs NO, then we output a trivial constant sized NO instance of $\mathcal{F}_p$-Deletion. Otherwise, we have a 1-redundant solution of size $\mathcal{O}(k^2)$, that we denote by S in all the subsequent phases.

Degree reduction. The input to this phase is the instance (G, k) together with a 1-redundant solution S. The goal is to design a reduction rule that bounds the size of the set of neighbours of x in C, for each $x \in S$ and $C \in \mathtt{cc}(G - S)$, by $k^{\mathcal{O}(1)}$. An edge uv in G is called *irrelevant* if the instance (G, k) is equivalent to the instance $(G - uv, k)$.

Let a denote the number of 2-boundaried graphs on at most $2h+6$ vertices. Let $g : [a] \to \mathbb{N}$ be a function such that $g(1) := 18 \cdot (2h+6)$ and for each $i > 1$, $g(i) := 18 \cdot (2h+6) \cdot g(i-1)$. Also let $\mathtt{degree\text{-}bound} := p^3 \cdot (|S| \cdot (p + k + 1) + 1) \cdot g(a)$. We prove the following lemma.

▶ **Lemma 3.2** (Irrelevant edge). *Let (G, k) be an instance of $\mathcal{F}_p$-Deletion and S be a 1-redundant solution in G. Let $x \in S$ and $C \in \mathtt{cc}(G - S)$. If $|N(x) \cap C| \geq \mathtt{degree\text{-}bound}$, then there exists $u \in N(x) \cap C$, such that xu is irrelevant. Moreover, such a vertex u can be found in polynomial time.*

The key insights and the proof of Lemma 3.2 are delegated to Section 4. We use Lemma 3.2 as long as there exists $x \in S$ and $C \in \texttt{cc}(G - S)$ such that $|N(x) \cap C| \geq \texttt{degree-bound}$. If u is the vertex reported by Lemma 3.2, then we apply a reduction rule and output an equivalent instance $(G - uv, k)$. Therefore at the end of this stage we can assume that for each $x \in S$ and $C \in \texttt{cc}(G - S)$, $|N(x) \cap C| < \texttt{degree-bound} = \mathcal{O}(k^3)$. It is crucial to note here that the exponent of k in $\texttt{degree-bound}$ is independent of $\mathcal{F}$.

Component reduction. In this phase, we design a reduction rule whose exhaustive application guarantees that $|\texttt{cc}(G - S)| = \mathcal{O}(|S|^2 \cdot k) = \mathcal{O}(k^5)$. Formally we prove the following lemma. Let $\texttt{comp-bound}_h := (10 + p) \cdot (h + 1)^2 \cdot 2^{\binom{h+1}{2}}$. Note that $\texttt{comp-bound}_h$ depends only on h, p (and therefore only on $\mathcal{F}_p$).

▶ **Lemma 3.3** ($\star$, Component reduction). *Let (G, k) be an instance of $\mathcal{F}_p$-*Deletion *and S be some solution of $\mathcal{F}_p$-*Deletion *(not necessarily optimal). If $|\texttt{cc}(G - S)| \geq \texttt{comp-bound}_h \cdot |S|^2 \cdot k$ [1], then there exists $C \in \texttt{cc}(G - S)$ such that the instance (G, k) is equivalent to $(G - C, k)$. Moreover, such a component C can be found in polynomial time.*

Apply Lemma 3.3 on the instance (G, k) and the set S from phase 1 until $|\texttt{cc}(G - S)| = \mathcal{O}(k^5)$.

Protrusion decomposition. Note that at the end of phase three, we can bound the size of the set of neighbours of S ($|N(S)|$) by $\mathcal{O}(k^{10})$: indeed from phase three (Lemma 3.3) $|\texttt{cc}(G - S)| = \mathcal{O}(k^5)$, from phase 2, for each $x \in S$, $|N(x) \cap C| = \mathcal{O}(k^3)$ and $|S| = \mathcal{O}(k^2)$. We use this to obtain a protrusion decomposition of G with $\mathcal{O}(k^{10})$ protrusions. This is defined below.

For a graph G, let $\texttt{tw}(G)$ denote the treewidth (see [8] for the definition) of G. For a positive integer a, an *a-protrusion* in G is a set of vertices $X \subseteq V(G)$ such that $\texttt{tw}(G[X]) \leq a$ and $|N_G(X)| \leq a$. For positive integers a, b, c, an *(a, b, c)-protrusion decomposition* of G is a partition of $V(G) = V_0 \uplus V_1 \uplus \ldots \uplus V_r$ such that the following holds: (1) $|V_0| \leq a$, (2) $r \leq b$, (3) for each $i \in [c]$, $N(V_i) \subseteq V_0$, and, (4) for each $i \in [c]$, V_i is a c-protrusion in G.

In this phase, we prove the following lemma.

▶ **Lemma 3.4** ($\star$, Protrusion decomposition). *Let G be a graph and $S \subseteq V(D)$ such that $tw(G - S) \leq \eta$ and $G - S$ has at most ζ connected components. If $|N(S) \cap C| \leq \alpha$ for every connected component of $G - S$, then G admits a $(|S| + 2\alpha\eta\zeta, 6\alpha\zeta, 2\eta)$-protrusion decomposition and can be computed in polynomial time.*

We use Lemma 3.4 on G and the set S from phase 1. Since S is a solution to $\mathcal{F}_p$-Deletion, $K_{2,p} \in \mathcal{F}_p$ and $K_{2,p}$ is planar, $\texttt{tw}(G - S) = \mathcal{O}((p + 2)^9)$ [3]. Thus, in Lemma 3.4, on input G, S, $\eta = \mathcal{O}((p + 2)^9)$, from phase 3 $\zeta = \mathcal{O}(k^5)$ and, from phase 2 $\alpha = \mathcal{O}(k^3) \cdot |S| = \mathcal{O}(k^5)$. Thus, we get an $(\mathcal{O}(k^{10}), \mathcal{O}(k^{10}), \mathcal{O}(1))$-protrusion decomposition of G.

Protrusion replacement. Let $V_0 \uplus V_1 \uplus \ldots V_r$ be the $(\mathcal{O}(k^{10}), \mathcal{O}(k^{10}), \mathcal{O}(1))$-protrusion decomposition of G obtained from the previous phase. Note that in order to bound the size of the whole graph G, it remains to bound the size of protrusion V_i for each $i \in \{1, \ldots, r\}$. This is done in this final phase. To reduce the size of the protrusions, we use the fact that the $\mathcal{F}_p$-Deletion problem has Finite Integer Index. This allows one to "replace" each protrusion V_i with a vertex set whose size depends only on $\mathcal{F}_p$.

[1] The bound on $\texttt{comp-bound}_h$ has not been optimized.

Recall η, ζ, α from Lemma 3.4. The following proposition from [2, 12] implies a reduction rule if the size of any V_i, $i \in \{1, \dots, r\}$, is larger than a fixed constant q that depends only on $\mathcal{F}_p$.

▶ **Proposition 3.5** (Protrusion replacer, [2, 12])**.** *Let (G, k) be an instance of $\mathcal{F}_p$-*Deletion *and let $V_i \subseteq V(G)$ be a 2η-protrusion in G of size at least q, where q is a fixed constant that depends only on $\mathcal{F}_p$. Then there exists a polynomial-time algorithm that outputs an equivalent instance (G', k') such that $|V(G')| < |V(G)|$ and $k' \leq k$.*

Finally this implies the following corollary. Again recall η, ζ, α from Lemma 3.4 and q from Proposition 3.5.

▶ **Corollary 3.6.** *If $|V(G)| \geq |S| + 2\alpha\eta\zeta + 6q\alpha\zeta k$, then the reduction rule implied by Proposition 3.5 is applicable.*

This implies that when the reduction rule of Lemma 3.5 is not applicable, $|V(G)| = \mathcal{O}(k^{10})$. The proof of Theorem 1.1 follows from Lemmas 3.1,3.2,3.3,3.4 and Corollary 3.6.

4 The degree reduction phase

In this section we elaborate on the degree reduction phase that we described in Section 3. Recall that (G, k) is an instance of $\mathcal{F}_p$-Deletion and $S \subseteq V(G)$ is a 1-redundant solution, that is for each $x \in S$, $S \setminus \{x\}$ is a solution. Fix $x \in S$ and $C \in \mathtt{cc}(G - S)$ such that $|N(x) \cap C| \geq$ `degree-bound`. Because S is 1-redundant, $G[C \cup \{x\}]$ is $K_{2,p}$-free. Set $X := N(x) \cap C$. The goal is to design a reduction rule that bounds the size of X by $k^{\mathcal{O}(1)}$.

4.1 Overview and key insights

We give an overview of the key insights in the proof of Lemma 3.2. The degree reduction phase has two main steps. In this first step we find a subset of C that has a very nice structure and in the second step we exploit this structure to "re-build" minor models of small graphs so that they avoid an irrelevant edge.

A nicely structured set $C^\sigma \subseteq C$. As a first step, we exploit the fact that $G[C \cup \{x\}]$ is $K_{2,p}$-free to obtain a subset $C^\sigma \subseteq C$ containing $g(a)$ (recall a, g from Section 3) neighbours of x such that C^σ has a very nice "chain-like" structure. The definition of the structure is formalized in Definition 4.1 as an *x-good sequence*. Also see Figure 1. The proof of its existence is given in Lemma 4.3. Below we state informally the nice structure of C^σ that we achieve.

The chain structure: The set C^σ contains $g(a)$ vertices of X, say ordered $(u_1, \dots, u_{g(a)})$. For each other vertex $v \in C^\sigma$, v is on some X-free (u_i, u_{i+1})-walk in C. For each $i \in [g(a)-1]$, let V_i^σ be the set of vertices on some X-free (u_i, u_{i+1})-walk in C. Then $V_i^\sigma \neq \emptyset$. For each $i, j \in [g(a) - 1]$, $i \neq j$, $V_i^\sigma \cap V_j^\sigma = \emptyset$. The set $\{u_i, u_{i+1}\} \cup V_i^\sigma$ is called a *block* of C^σ.

The boundary to C: The vertices $u_1, u_{g(a)}$ are the boundary vertices of C^σ in C. That is, no vertex of $C^\sigma \setminus \{u_1, u_r\}$ has any neighbour in $C \setminus C^\sigma$.

Neighbours in S after removing a solution: Lastly, for any $\mathcal{F}_p$-Deletion solution T of size at most $k + 1$, in $G - T$, the $N(C^\sigma) \cap S \subseteq \{x\}$.

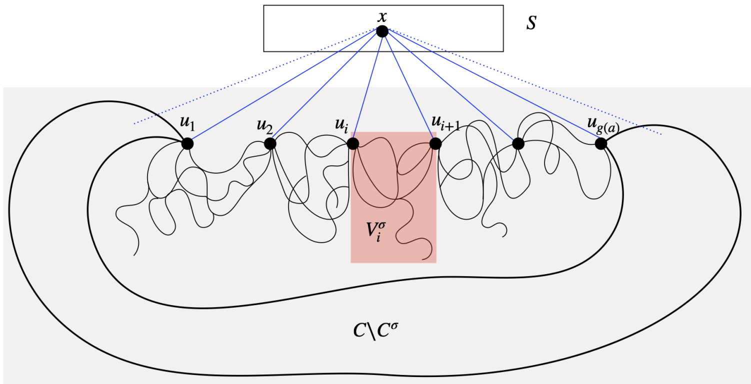

Figure 1 The structure of C^σ: The part with the grey background is a connected component C of $G - S$. The blue edges connect x to its neighbours in C.

The structure of minors of small graphs in $\mathcal{F}_p$ restricted to C^σ. In the second step, we exploit this structure of C^σ to find a vertex $u \in C^\sigma \cap X$ such that xu is an irrelevant edge. To show that xu is irrelevant, we want to show that if there is a set of at most k vertices T in $G - xu$ such that $(G - xu) - T$ is $\mathcal{F}_p$-free, then $G - T$ is also $\mathcal{F}_p$-free. In particular we will show that if $G - T$ has *any* graph H on at most h vertices as a minor, and there is a minor model of H in $G - T$ that uses the edge xu, then there is also a minor model of H in $G - T$ that does not use the edge xu (Lemmas 4.14 and 4.16).

For doing the above, the key insight is to focus on the minor model of H restricted to $C^\sigma \cup \{x\}$.

4.2 Proof of Lemma 3.2 (Irrelevant edge)

The goal of this section is to prove Lemma 3.2.

▶ **Lemma 3.2** (Irrelevant edge). *Let (G, k) be an instance of $\mathcal{F}_p$-*Deletion *and S be a 1-redundant solution in G. Let $x \in S$ and $C \in \mathtt{cc}(G - S)$. If $|N(x) \cap C| \geq$ `degree-bound`, then there exists $u \in N(x) \cap C$, such that xu is irrelevant. Moreover, such a vertex u can be found in polynomial time.*

An *x-sequence* is a sequence σ of a subset of vertices of X. If $\sigma = (u_1, \ldots, u_r)$ is an x-sequence, for each $i \in [r-1]$, the set V_i^σ contains each vertex that appears on some X-free (u_i, u_{i+1})-walk in C. An *i-block* of σ refers to the set $\{u_i, u_{i+1}\} \cup V_i^\sigma$. An *$r$-block* is simply the vertex u_r. By a *block* of σ we simply refer to some i-block of σ. By the *endpoints* of an i-block, we refer to the vertices u_i and u_{i+1}. Further $C^\sigma := \bigcup_{i=1}^{r-1} V_i^\sigma \cup \bigcup_{i=1}^{r} \{u_i\}$. The *length* of the x-good sequence σ is r.

▶ **Definition 4.1** (x-good sequence). *An x-sequence $\sigma = (u_1, \ldots, u_r)$ is called an x-good sequence if the following holds.*

1. *For every $i \in [r-1]$, there is an X-free (u_i, u_{i+1})-walk in C.*
2. *For every $i \in [2, r-1]$, $\{u_{i-1}, u_{i+1}\}$ is a (u_i, X)-cut.*
3. *For any $\mathcal{F}_p$-*Deletion *set T in G of size at most $k+1$, if $x \notin T$, $N(C^\sigma) \cap (S \setminus T) = \{x\}$.*
4. *No vertex of $C^\sigma \setminus \{u_1, u_r\}$ has a neighbour in $C \setminus C^\sigma$.*

▶ **Lemma 4.2** (⋆). *If $\sigma = (u_1, \ldots, u_r)$ is an x-good sequence then, for any minimal $\mathcal{F}_p$-Deletion solution T in G of size at most $k+1$ that does not contain x, $\{x\} \subseteq N(C^\sigma \setminus \{u_1, u_r\}) \setminus T \subseteq \{x, u_1, u_r\}$. Also, there exists a minimum $\mathcal{F}_p$-Deletion solution T^* in G such that $|T^* \cap C^\sigma| \leq 3$.*

In Section 4.2.1 we show that if X is large then there is a large x-good sequence (Lemma 4.3). In Section 4.2.2 we show how to find an irrelevant edge xu, where $u \in \sigma$, given a large x-good sequence σ.

4.2.1 Finding a large x-good sequence

The goal of this section is to prove the following lemma.

▶ **Lemma 4.3.** *If $|X| \geq p^3 \cdot (|S| \cdot (p+k+1)+1) \cdot g(a)$, then there exists an x-good sequence of length $g(a)$. In fact, such a sequence can be found in polynomial time.*

We start by taking some natural steps towards the construction of some x-sequence that has Property 1 of Definition 4.1. We later perform more steps towards proving other properties of Definition 4.1. Let us define an ordered partition of $X = (L_0, \ldots, L_l)$ inductively as follows. Fix an arbitrary vertex $x_0 \in X$ and let L_0 be $\{x_0\}$. Now suppose $L_0, \ldots, L_i$ are defined, we define L_{i+1} as the set of vertices of $X \setminus \{\bigcup_{j \in [0..i]} L_j\}$ that are reachable from L_i by X-free paths in $G[C]$. We next prove a series of claims about this partition $X = (L_0, \ldots, L_l)$. Observe that each vertex of C is on some X-free (L_i, L_{i+1})-walk.

▶ **Lemma 4.4** (⋆). *For every $i \in [0, l]$, $|L_i| \leq p - 1$.*

Observe that, if $|X| \geq p^3 \cdot (|S| \cdot (p+k+1)+1) \cdot g(a)$, then $l \geq p^2(|S| \cdot (p+k+1)+1)g(a)$.

The next definition and the upcoming steps help to ensure Property 2 of Definition 4.1 (this is formally proved in Lemma 4.11). For every $i \in [0, l-1]$, we say that a vertex $v \in L_i$ is *dangerous* if the set of vertices in X that are reachable from $L_i \setminus v$ by paths whose internal vertex set is disjoint from $(X \setminus L_{i+1})$, is exactly L_{i+1}. Observe that for any vertex $v \in L_i$ which is not dangerous, there is a vertex v' in L_{i+1} such that v' is not reachable from $L_i \setminus v$ by an $(X \setminus L_{i+1})$-free path. Such a vertex v' is called a *witness of a non dangerous vertex v*.

▶ **Lemma 4.5** (⋆). *The number of indices $i \in [0, l]$ such that L_i contains a dangerous vertex is at most $p - 1$.*

▶ **Lemma 4.6** (⋆). *For every $i \in [0, l-1]$, if no vertex of L_i is dangerous, then $|L_i| \leq |L_{i+1}|$.*

▶ **Lemma 4.7** (⋆). *Let $t = (|S| \cdot (p+k+1)+1)g(a)$. There exists $i \in [l-t]$ such that none of $L_i, \ldots, L_{i+t}$ contains dangerous vertices and $|L_i| = |L_{i+s}|$ for all $s \in [t]$.*

Without loss of generality, let $L_1, \ldots, L_t$ denote the interval of $(L_0, \ldots, L_l)$ from Lemma 4.7 that do not contain a dangerous vertex, where $t = (|S| \cdot (p+k+1)+1)g(a)$. Using this consecutive sequence of t sets, we will now define an x-sequence $\sigma^\star = (u_1, \ldots, u_t)$ of length t. The vertices u_j in this sequence are defined inductively as follows. Let u_1 be any vertex of L_1. Then for any $j \in [t-1]$, u_{j+1} is the witness for the non dangerous vertex u_j. Note that $\sigma^\star$ might not be an x-good sequence. In what follows, we prove some nice properties of $\sigma^\star$ and then use them to refine $\sigma^\star$ to obtain an x-good sequence. We would like to remark that this refinement procedure is required to prove Property 3 of Definition 4.1. We begin by proving a claim which will lead to the refinement. We will first show that these sets are disjoint.

▶ **Lemma 4.8.** *For every $j \in [t-1]$, (u_j, u_{j+1}) is a $(V_j^{\sigma^\star}, X)$-cut in $G[C]$.*

Proof. For the sake of contradiction, suppose that this is not the case and let P be a minimal $(V_j^{\sigma^\star}, u)$-path for some $u \in X$, in $G[C] - \{u_j, u_{j+1}\}$. If $u \in L_s$ for $s \leq j-1$, then this contradicts the fact that u_{j+1} is in L_{j+1}, and if $u \in L_s$ for $s \geq j+2$, then this contradicts the fact that u is not in L_{j+1}. If $u \in L_j$, then this contradicts the fact that u_j is the witness of u_{j-1} and if $u \in L_{j+1}$, this contradicts the fact that u_{j+1} is the witness of u_j. ◀

As a corollary of Lemma 4.8, we conclude the following.

▶ **Lemma 4.9.** *For each $i, j \in [t]$, where $i \neq j$, we have $V_i^{\sigma^\star} \cap V_j^{\sigma^\star} = \emptyset$.*

Let $S_1 \subseteq S$ be the set of all those vertices s of S such that there exists $I \subseteq [t]$ of size at least $k+p+2$ and for each $b \in I$, s is adjacent to some vertex of $V_b^{\sigma^\star} \cup u_b$ (or to u_b when $b = t$). Let $S_2 = (S \setminus S_1) \setminus x$. An index $b \in [t]$ is called *affected* if there exists $s \in S_2$ such that s is adjacent to $V_b^{\sigma^\star} \cup u_b$ (or to u_b when $b = t$). By the definition of S_2, the number of affected indices is at most $|S_2| \cdot (k+p+1) \leq |S| \cdot (k+p+1)$. Since the length of $\sigma^\star$ is $t = (|S| \cdot (p+k+1)+1)g(a)$ and the number of affected indices is at most $|S|(p+k+1)$, there exists an interval σ of $\sigma^\star$ of length $g(a)$ such that none of the indices corresponding to the subscripts of the vertices in σ are affected. Without loss of generality, let $\sigma = (u_1, \ldots, u_{g(a)})$. We will now show that σ has Property 3 of Definition 4.1.

▶ **Lemma 4.10** (Property 3). *Let T be some $\mathcal{F}_p$-DELETION set of size at most $k+1$ such that $x \notin T$. Then $N(C^\sigma) \cap (S \setminus T) = \{x\}$.*

Proof. By the definition of C^σ, $N(C^\sigma) \cap (S \setminus T)$ contains x. For the sake of contradiction, say $s \in S \setminus x$ belongs to $N(C^\sigma) \cap (S \setminus T)$. Since none of the indices corresponding to the vertices in σ are affected, we conclude that $s \in S_1 \setminus T$. Since $|T| \leq k+1$, σ is an interval of $\sigma^\star$ and from Lemma 4.9, there exists at least p indices in $[t]$ such that for each of these p indices, say b, $T \cap (V_b^{\sigma^\star} \cup u_b) = \emptyset$ (or $u_b \notin T$, if $b = t$) and s is a neighbour of each of these p sets $V_b^{\sigma^\star} \cup u_b$. Then the graph induced by $G - T$ on x, s and p of these sets contains $K_{2,p}$ as a minor in $G - T$, which is a contradiction as T is an $\mathcal{F}_p$-DELETION set. ◀

▶ **Lemma 4.11** (Property 2). *For every $j \in [2, t-1]$, $\{u_{j-1}, u_{j+1}\}$ is a (u_j, X)-cut in C.*

Proof. Suppose there is a path in C between u_j and some vertex $u \in X$ different from $\{u_{j-1}, u_{j+1}\}$. By definition of the L_i's, u belongs to either L_{j-1}, L_j or L_{j+1}. If u belongs to L_{j-1} or L_j, this contradicts the fact that u_j is the witness of u_{j-1}. If u belongs to L_{j+1}, this contradicts the fact that u is the witness of a vertex different from u_j in L_j, so we reach a contradiction. ◀

▶ **Lemma 4.12** (Property 4). *$(N(C^\sigma) \setminus \{u_1, u_{g(a)}\}) \cap C \subseteq \{u_1, u_{g(a)}\}$.*

Proof. Fix $i \in [g(a)-1]$. We will first show that no vertex of V_i^σ has a neighbour in $C \setminus C^\sigma$. For the sake of contradiction say $v \in V_i^\sigma$ is a neighbour of $w \in C \setminus C^\sigma$. Since w is on some walk between two vertices of $X \setminus \{u_1, \ldots, u_{g(a)}\}$, this implies that there is a path from v to a vertex in X that does not intersect $\{u_1, \ldots, u_{g(a)}\}$. This contradicts Lemma 4.8.

It remains to show that none of the vertices in $\{u_2, \ldots, u_{g(a)-1}\}$ have a neighbour in $C \setminus C^\sigma$. We show this in two parts. Note from the construction of the sets L_i, that for any $u_i \in L_i$, its neighbours in X are either in L_{i-1}, L_i or L_{i+1}. Because no vertex of L_i is dangerous and $|L_{i-1}| = |L_i|$, L_i is an independent set and the only potential neighbour of u_i in L_{i-1} and L_{i+1} is u_{i-1} and u_{i+1} respectively. Thus we conclude that no vertex in $\{u_2, \ldots, u_{g(a)-1}\}$ has a neighbour in $(C \setminus C^\sigma) \cap X$.

To finish the proof we need to show that no vertex of $\{u_2, \ldots, u_{g(a)-1}\}$ has a neighbour in $(C \setminus C^{\sigma}) \setminus X$. For the sake of contradiction, say u_i, for $i \in [2, g(a)-1]$, has a neighbour $w \in (C \setminus C^{\sigma}) \setminus X$. Since w is on some walk between two vertices of $X \setminus \{u_1, \ldots, u_{g(a)}\}$, this would imply a walk from a vertex of $X \setminus \{u_2, \ldots, u_{g(a)-1}\}$ to u_i, $i \in [2, g(a)-1]$. This either contradicts that $u_i \in L_i$ or that $w \notin C^{\sigma}$. ◀

From the construction of the sequence $(L_0, \ldots, L_l)$, observe that σ satisfies Property 1 of Definition 4.1. This together with Lemmas 4.10, 4.11 and 4.12 and the fact that the length of σ is $g(a)$, proves Lemma 4.3.

4.2.2 Finding an irrelevant edge

The goal of this section is to complete the proof of Lemma 3.2 using Lemma 4.3. Let $\sigma = (u_1, \ldots, u_r)$ be an x-good sequence. The *graph induced by* σ, denoted by $G[\sigma]$, is a 2-boundaried graph $G[C^{\sigma}]$ with boundary u_1, u_r. By $G[\sigma_i]$ we denote the 2-boundaried graph induced by the i-block of σ, with boundary u_i, u_{i+1}. Let $\hat{G} = G[\sigma] \cup x$ be a boundaried graph with boundary that is a subset of $\{x, u_1, u_r\}$. Let $h' = 2(h+3)$. The folio-set of $G[\sigma]$, denoted by folio-set$(G[\sigma])$, is the collection of $\{\cup_{i \in [r]}\{h'\text{-folio}(G[\sigma_i])\}\}$, where r is the length of σ.

In order to prove Lemma 3.2, it is enough to show that there is a vertex $u_{red} \in X$, such that for any $T \subseteq V(G)$ of size at most k, if there exists a graph H on h vertices that is a minor of $G - T$ and whose minor model uses the edge xu_{red}, then there exists a minor model of H in $G - T$ that does not use the edge xu. From Lemma 4.2 and Lemma 23 of [12] (stated below) this will follow from Lemma 4.14.

▶ **Proposition 4.13** (Lemma 23, [12]). *Let G_1 and G_2 be t-boundaried graphs and $G = G_1 \oplus G_2$. A graph H is a minor of G if and only if there exist $H_1 \leq_m G_1$ and $H_2 \leq_m G_2$ such that $|V(H_1)| \leq |V(H)| + t$, $|V(H_2)| \leq |V(H)| + t$ and $H \leq_m H_1 \oplus H_2$.*

▶ **Lemma 4.14.** *If the length of σ is $g(a)$, then one can find a vertex $u_{red} \in \sigma$ in polynomial time such that the following holds. Let H be a 3-boundaried graph with boundary $\{x\} \subseteq B \subseteq \{x, u_1, u_r\}$ of size at most $h+3$ that is present as a (boundaried) minor in $\hat{G} - T$, where $T \subseteq V(\hat{G}) \setminus x$ and $|T| \leq 3$. Then there exists a minor model of H in $\hat{G} - T$ that does not use the edge xu_{red}.*

▶ **Proposition 4.15** ([12]). *If ϕ is a minimal minor model of H in G, then every vertex in the minor model has degree at most $|V(H)|$ in the minor model.*

Let H be a 3-boundaried graph with boundary $\{x\} \subseteq B \subseteq \{x, u_1, u_r\}$ in G. Let ϕ be a minimal minor model of H in G. Let $\phi' = \phi \setminus x$ and H' be the (boundaried) graph witnessed by ϕ'. By Proposition 4.15, $|V(H')| \leq 2|V(H)|$. Let ϕ_1 and ϕ_2 be two minor models of some 2-boundaried graph H in $G[\sigma]$. Then ϕ_2 is said to be σ*-compatible* with ϕ_1, if for every branch set of ϕ_1 that has a vertex of σ, the corresponding branch set of ϕ_2 also has a vertex of σ. Recall that all the vertices of σ are the vertices of X and hence they are neighbours of x. With the discussion above, it is not difficult to see that to prove Lemma 4.14, it is enough to prove Lemma 4.16.

▶ **Lemma 4.16.** *If the length of σ is $g(a)$, then one can find a vertex $u_{red} \in \sigma$ in polynomial time such that the following holds. Let H be a 2-boundaried graph with boundary $B \subseteq \{u_1, u_r\}$ of size at most h' that is present as a (boundaried) minor in $\hat{G} - T$, where $T \subseteq V(G[\sigma])$ and $|T| \leq 3$. Let ϕ be a minor model of H in $\hat{G} - T$. Then there exists a minor model ϕ' of H in $\hat{G} - T$ that does not use the edge xu_{red} and is σ-compatible with ϕ.*

The following lemma will be crucially used in the arguments that follow. It depicts the structure of minors in $G[\sigma]$. Let H be a 2-boundaried minor in $G[\sigma]$. Let ϕ be a minor model of H in $G[\sigma]$. The *blocks of σ used by ϕ* refers to the collection of blocks of σ that have a non-empty intersection with the vertices of the minor model ϕ. The *crucial blocks used by ϕ* refers to those blocks which have a branch set of ϕ fully contained in it. Note that the number of crucial blocks of ϕ is at most $|V(H)|$.

▶ **Lemma 4.17** ($\star$). *Let H be a connected 2-boundaried minor in $G[\sigma]$. Let ϕ be some minor model of H in $G[\sigma]$. Then there exists a minor model ϕ' of H in $G[\sigma]$ obtained from ϕ by replacing the vertices in the non-crucial blocks used by ϕ with any arbitrary path between the endpoints of the block.*

▶ **Definition 4.18** (Chunk partition of σ). *Let σ be an x-good sequence. A* chunk *of σ is an interval of σ. A* chunk partition *of σ is a partition of the blocks of σ into intervals. Let σ' be a chunk of σ. Then folio-set(σ') is a set containing the h'- folio($G[\sigma_i]$), for each i-block in σ'. A chunk partition of σ is called* uniform *if the folio-sets of all chunks in the partition are same. In particular, for any $j \in [a]$, a chunk partition of σ is called j-*uniform *if it is uniform and the size of the folio-set of each chunk is exactly j.*

▶ **Lemma 4.19** ($\star$, Finding an i-uniform chunk partition). *Let σ be an x-good sequence and let the size of the folio-set(σ) be exactly i. If the length of σ is $g(i)$ then there exists an interval of σ, say σ', which admits a j-uniform chunk partition of length $18h'$, for some $j \in [i]$.*

▶ **Lemma 4.20** ($\star$, Replacement Lemma). *Let σ be an x-good sequence and $\chi = (\chi_1, \dots, \chi_s)$ be an i-uniform chunk partition of σ, for some $i \in [a]$. Let $s \geq h'$. Let H be a 2-boundaried minor in $G[\sigma]$ of size at most h'. Then H is present as a 2-boundaried minor in every graph that is induced on any h' sized interval of χ. Moreover, the later minor model of H is σ-compatible with the former minor model of H.*

▶ **Lemma 4.21** ($\star$). *Let σ be an x-good sequence of length $g(a)$ that admits an i-uniform chunk partition, for some $i \in [a]$, of length $18h'$. Then Lemma 4.16 holds.*

From Lemmas 4.21 and 4.19, Lemma 4.16 follows. Lemma 4.16 together with Lemma 4.3 finishes the proof of Lemma 3.2.

5 Conclusion

In this article we showed that $\mathcal{F}$-Deletion where all graphs in $\mathcal{F}$ are connected and $\mathcal{F}$ contains $K_{2,p}$ admits a uniform polynomial kernel of size $\mathcal{O}(k^{10})$. This result is the third example where $\mathcal{F}$-Deletion admits a uniform polynomial kernel; the first two being the Treedepth-η Deletion and the case when $\mathcal{F}$ contains θ_p. The most interesting aspect of our result is defining and obtaining an extremely structured set of vertices that have a small effective boundary. This structure is exploited to reduce the degree of the vertices to $k^{\mathcal{O}(1)}$.

We conclude with some intriguing open questions. Our result does not extend to the case when $\mathcal{F}$ is allowed to contain disconnected graphs. The first question is: can one obtain a uniform polynomial kernel when $\mathcal{F}$ contains $K_{2,p}$ and other possibly disconnected graphs? In fact, handling disconnected graphs in the kernelization algorithm of [12] is one point which introduces non-uniform bounds. Can this step of the kernelization algorithm of [12] be made to work without introducing non-uniformity? Or even more specifically, can we find some non-trivial families $\mathcal{F}$ which contain disconnected graphs but admit uniform polynomial kernels? Lastly, can we characterize the families $\mathcal{F}$ that admit a uniform polynomial kernel?

As long as we do not resolve the last question completely, one would be interested in finding more and more non-trivial families for which the problem admits a uniform polynomial kernel.

References

1 Isolde Adler, Martin Grohe, and Stephan Kreutzer. Computing excluded minors. In Shang-Hua Teng, editor, *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 641–650. SIAM, 2008. URL: `http://dl.acm.org/citation.cfm?id=1347082.1347153`.

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

3 Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM (JACM)*, 63(5):1–65, 2016. `doi:10.1145/2820609`.

4 Guantao Chen, Yoshimi Egawa, Ken-ichi Kawarabayashi, Bojan Mohar, and Katsuhiro Ota. Toughness of-minor-free graphs. *The Electronic Journal of Combinatorics [electronic only]*, 18(1):Research–Paper, 2011.

5 Guantao Chen, Laura Sheppardson, Xingxing Yu, and Wenan Zang. The circumference of a graph with no $K_{3,t}$-minor. *Journal of Combinatorial Theory, Series B*, 96(6):822–845, 2006. `doi:10.1016/J.JCTB.2006.02.006`.

6 Maria Chudnovsky, Bruce Reed, and Paul Seymour. The edge-density for $K_{2,t}$ minors. *Journal of Combinatorial Theory, Series B*, 101(1):18–46, 2011. `doi:10.1016/J.JCTB.2010.09.001`.

7 Youssou Dieng. *Décomposition arborescente des graphes planaires et routage compact*. PhD thesis, Bordeaux 1, 2009.

8 Reinhard Diestel. *Graph theory*. Springer Publishing Company, Incorporated, 2018.

9 Guoli Ding. Graphs without large $K_{2,n}$-minors. *arXiv preprint*, 2017. `arXiv:1702.01355`.

10 Mark N Ellingham, Emily A Marshall, Kenta Ozeki, and Shoichi Tsuchiya. A characterization of $K_{2,4}$-minor-free graphs. *SIAM Journal on Discrete Mathematics*, 30(2):955–975, 2016. `doi:10.1137/140986517`.

11 Fedor V Fomin, Daniel Lokshtanov, Neeldhara Misra, Geevarghese Philip, and Saket Saurabh. Hitting forbidden minors: Approximation and kernelization. *SIAM Journal on Discrete Mathematics*, 30(1):383–410, 2016. `doi:10.1137/140997889`.

12 Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Planar f-deletion: Approximation, kernelization and optimal FPT algorithms. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 470–479. IEEE Computer Society, 2012. See `http://www.ii.uib.no/~daniello/papers/PFDFullV1.pdf` for the full version. `doi:10.1109/FOCS.2012.62`.

13 Fedor V Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: theory of parameterized preprocessing*. Cambridge University Press, 2019.

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

15 Bart M. P. Jansen and Michal Wlodarczyk. Lossy planarization: a constant-factor approximate kernelization for planar vertex deletion. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 900–913. ACM, 2022. `doi:10.1145/3519935.3520021`.

16 Bart MP Jansen, Daniel Lokshtanov, and Saket Saurabh. A near-optimal planarization algorithm. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1802–1811. SIAM, 2014.

17 Ken-ichi Kawarabayashi. Planarity allowing few error vertices in linear time. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 639–648. IEEE, 2009. doi:10.1109/FOCS.2009.45.

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

19 Dániel Marx and Ildikó Schlotter. Obtaining a planar graph by vertex deletion. *Algorithmica*, 62(3-4):807–822, 2012. doi:10.1007/S00453-010-9484-Z.

20 Joseph Samuel Myers. The extremal function for unbalanced bipartite minors. *Discrete mathematics*, 271(1-3):209–222, 2003. doi:10.1016/S0012-365X(03)00051-7.

21 Katsuhiro Ota and Kenta Ozeki. Spanning trees in 3-connected $K_{3,t}$-minor-free graphs. *Journal of Combinatorial Theory, Series B*, 102(5):1179–1188, 2012. doi:10.1016/J.JCTB.2012.07.002.

22 Neil Robertson and PD Seymour. Graph minors. Xlll. The disjoint paths problem. *J. Combin. Theory Ser. B*, 63:65–110, 1995.

23 Noah Streib and Stephen J Young. Dimension and structure for a poset of graph minors. *University of Louisville Department of Mathematics*, 2010.