Treewidth reduction for the parameterized Multicut problem *

Jean Daligault, Christophe Paul, Anthony Perez, Stéphan Thomassé
LIRMM – Université de Montpellier 2, CNRS,
Montpellier, France.
{daligault|paul|perez|thomasse}@lirmm.fr

November 3, 2010

Abstract

The parameterized Multicut problem consists in deciding, given a graph, a set of requests (i.e. pairs of vertices) and an integer k, whether there exists a set of k edges which disconnects the two endpoints of each request. Determining whether Multicut is Fixed-Parameter Tractable with respect to k is one of the most important open question in parameterized complexity [5]. We show that Multicut reduces to instances of treewidth bounded in k. To that aim, we establish new reduction rules that apply to arbitrary instances of Multicut. Based on graph separability properties, these rules identify an irrelevant request that can be safely removed. As a main consequence, these rules imply that the degree of the request graph of any instance is bounded by a function of k. We prove that when the input graph has a large clique minor or a large grid minor, then we can remove an irrelevant request or contract an edge.

1 Introduction

Among the classical techniques to cope with NP-hard problems, parameterized algorithms have known a considerable development in last few years. A problem is *fixed parameter tractable* (FPT) with respect to parameter k (e.g. solution size, treewidth) if, for any instance of size n, it can be solved in time $O(f(k).n^d)$ for some fixed d. The reader is invited to refer to books [6], [7] or [19]. In this paper, we are interested in the Multicut problem parameterized by the solution size, which is considered as one of the main open problems of the fixed parameterized complexity theory [5]. Given a graph G and a set R of *requests* between pairs of vertices, called *terminals* (or *endpoints*), an (edge)-*multicut*[1] is a subset F of edges of G whose removal separates the two endpoints of every request in different connected components of $G \setminus F$:

Problem Multicut:
Input: A graph $G = (V, E)$, a set of requests R, an integer k.
Parameter: k.
Output: TRUE if there is a multicut of size at most k, otherwise FALSE.

Related results. The Multicut problem is already hard when restricted to a set of requests on a tree. Indeed, Vertex Cover can be viewed as Multicut in stars, hence Multicut is NP-complete and Max-SNP hard. Multicut and its variants have raised an extensive literature. These problems play an important role in network issues, such as routing and telecommunication (see [4]).

Multicut In Trees was already a challenging problem. Garg *et. al.* [9] proved that it admits a factor 2 approximation algorithm. Guo and Niedermeier [13] proved that Multicut In Trees is FPT with respect to the solution size. And recently, Bousquet *et. al.* [1] provided a polynomial kernel. Another variant is the

*Research supported by the AGAPE project (ANR-09-BLAN-0159)
[1]As this paper never considers vertex multicut, the term *multicut* stands for edge-multicut.

Multiwaycut problem in which a set of (non-paired) terminals has to be pairwise separated. Parameterized by the solution size Multiwaycut has been proved FPT by Marx [15]. A faster $O^\star(4^k)$ algorithm is due to Chen *et. al.* [3].

On general instances, Garg *et. al.* gave an approximation algorithm for Multicut within a logarithmic factor in [8]. However Multicut has no constant factor approximation algorithm if Khot's Unique Games Conjecture holds [2]. This fact motivates the study of the fixed parameterized tractability of Multicut. Guo *et. al.* showed in [12] that Multicut is FPT when parameterized by both the treewidth of the graph and the number of requests. Gottlob and Lee in [10] proved a stronger result: Multicut is FPT when parameterized by the treewidth of the input structure, namely the input graph whose edge set is completed by the set of request pairs.

The graph minor theorem of Roberston and Seymour implies that Multicut is non-uniformly FPT when parameterized by the solution size and the number of requests. Marx proved that Multicut is (uniformly) FPT for this latter parameterization [15]. A faster algorithm running in time $O^*(8 \cdot l)^k$ was given by Guillemot [11]. Marx *et. al.* [16] obtained FPT results for more general types of constrained Multicut problems through treewidth reduction results. However their treewidth reduction techniques do not yield FPTness of Multicut when parameterized by the solution size only. Regarding this parameterization, Marx and Razgon recently obtained a factor 2 Fixed-Parameter-Approximation for Multicut in [17].

Our results. We show that Multicut parameterized by the solution size k can be reduced to graphs of treewidth bounded by a function of k. The key is to establish conditions under which we can identify an *irrelevant* request, that is a request the removal of which yields an equivalent instance. For example, it is proved that if a vertex is an endpoint of too many requests, then one of these requests is irrelevant. In particular, this implies that the request graph of any instance has degree bounded by a function of k. Likewise, if a so called *gathered* set of terminals exists (that is a set of terminals that satisfies some separability properties), then one of these terminals is incident to an irrelevant request. Both cases yield reduction rules as the irrelevant request can be identified in FPT time. We then prove that if the graph of the input instance has large treewidth (with respect to the solution size k), then either one of the first rules applies or we can identify an edge which contraction yields an equivalent instance. The cases where the input graph has a large clique minor or a large grid minor are proved separately. Therefore we prove that Multicut is FPT when parameterized by the solution size k if and only if Multicut is FPT when restricted to graphs of treewidth bounded in k. Besides we give an $O^*((2k+1)^k)$ algorithm for the Integer Weighted Multicut In Trees. This last problem, which was left open in [1] and [13], is one of the simplest subcases of the Multicut problem on bounded treewidth graphs.

2 Preliminaries

We consider undirected graphs $G = (V(G), E(G))$ together with a set of requests $R = \{(s_1, t_1), \ldots, (s_l, t_l)\}$, with $s_i, t_i \in V(G)$ for $i = 1, \ldots, l$. Given a vertex x of $V(G)$, $N_G(x)$ denotes the set of neighbours of x (we forget the subscript when the context is clear). When $X \subseteq V(G)$ we denote by $G[X]$ the subgraph of G induced by X.

Given two vertices $x, y \in V(G)$, *contracting* x and y in G means deleting x and y and adding a new vertex with neighbourhood $N(x) \cup N(y) \setminus \{x, y\}$. An (x, y)*-cut* is a set of edges of the graph G which removal disconnects x and y. We define similarly an (x, Y)-cut and an (X, Y)-cut for X, Y two sets of vertices of $V(G)$. We say that an edge xy *belongs* to a set $X \subseteq V(G)$ if $x, y \in X$. A set of edges F *touches* an induced subgraph G' of G if an edge in F belongs to $V(G')$. Given a graph J, let us denote by L_i and call *level i of J starting from* $L_0 \subseteq V(J)$ the subset of $V(J)$ containing vertices lying at distance exactly i from L_0 in J. The *level decomposition* of J starting with $L_0 \subseteq V(J)$ is the partition of $V(J)$ into levels L_i, $i \geq 0$.

Let $\mathcal{P}$ be a partition of $V(G)$, such that every part h of $\mathcal{P}$ induces a connected subgraph of G. Let H be the graph G quotiented by $\mathcal{P}$, i.e. the graph G where each part of $\mathcal{P}$ is contracted into a single vertex. If J is a subgraph of H, then J is said to be a *minor* of G. We say that the pair $(\mathcal{P}, H)$ is a J*-model* of G. In a slight abuse of notation, we write that H is a J*-model* of G (or simply a *model* of G), leaving partition

$\mathcal{P}$ implicit. We also abusively write that a part h of $\mathcal{P}$ is a *part* of H. We want to emphasize that edges of H are *all* pairs of parts of $\mathcal{P}$ between which G induces an edge. When H' is a subgraph of H, $V(H')$ naturally denotes the set of vertices of H', which are parts, and we write $V_G(H')$ to denote the set of vertices $\cup_{h \in H'}\{x \in h\} \subseteq V(G)$.

Given an instance (G, R, k) of the parameterized MULTICUT problem (or MULTICUT for short), a k-*multicut* is a multicut of size at most k. We say that a k-multicut is *optimal* when its size is minimum among all k-multicuts. We say that a request (s_i, t_i) is *irrelevant* whenever the instance $(G, R \setminus \{(s_i, t_i)\}, k)$ is equivalent to the instance (G, R, k).

Finally, we assume that the reader is familiar with the concept of treewidth. We refer the reader to Robertson and Seymour's Graph Minors Series or to [21] for definitions and results on treewidth.

3 General reduction rules for MULTICUT

Let (G, R, k) be an instance of MULTICUT. This section is devoted to show that the following set of reduction rules is correct and can be applied in FPT time.

Rule 1 *If there exist $k+1$ edge-disjoint paths between two vertices x and y, then contract x and y.*

Rule 2 *There exist an integer $g(k)$ such that if a vertex t is an endpoint of at least $g(k)$ requests, then we can find in FPT time an irrelevant request incident to t.*

By Rule 2 we may assume that the degree of the request graph is bounded by a function of k. We say that a set $T \subseteq V$ is *gathered* if for every $F \subseteq E$, $|F| \leq k$, there exists at most one connected component in $G \setminus F$ containing more than one vertex of T. We denote this connected component by C_F. Note that a subset of a gathered set is gathered.

Rule 3 *If the instance is reduced under Rule 2 and there exists a gathered set of terminals T of size at least $f_1(k) = 4g(k)^3$, then we can find in FPT time an irrelevant request incident to one of these terminals.*

Lemma 1 *Rule 1 is safe and can be applied in polynomial time.*

Proof: If there exist $k+1$ edge-disjoint paths between two vertices x and y, then x and y lie in the same connected component of $G \setminus F$ for any set F of k edges. Hence contracting x and y yields an equivalent instance. Testing whether two vertices are $(k+1)$-edge-connected is polynomial by flow. □

To prove the soundness of Rule 2 and Rule 3, we study two edge-connectivity problems of independent interest. We define the central notions of this section:

- A $(zy|x)$-cut is a (z, x)-cut which is not a (z, y)-cut. Similarly, given a subset of vertices T, a $(zy|T)$-cut is a (z, T)-cut which is not a (z, y)-cut.
- A vertex $y \notin T$ is $(z|T)$-*k-linked* if there is no $(zy|T)$-cut of size at most k, *i.e.* if every (z, T)-cut of size at most k is a (z, y)-cut.

Given a graph $G = (V, E)$, $x, y, z \in V(G)$ and a positive integer k, we call TRIPLE SEPARATION the problem of finding a $(zy|x)$-cut of size at most k (if one exists) parameterized by k.

Theorem 1 *The* TRIPLE SEPARATION *problem is FPT with respect to the solution size k.*

A stronger statement has recently been proved by Marx *et. al.* [16]. Their Theorem 3.4 states that deciding whether there exists a set of k vertices separating prescribed pairs and not separating other prescribed pairs is FPT with respect to k and the number of prescribed pairs. We give a different proof of our statement for the sake of completeness.

Proof: Let $G = (V, E)$ be a graph, and $X \subseteq V(G)$. We denote by $\delta(X)$ the *border* of the set of vertices X, i.e. the set of edges having exactly one endpoint in X. Let c be the size of a minimum cut between z and x, and X a minimal set of vertices of border of size c containing x. Observe first that if $c > k$ then the answer is No. So assume that $c \leqslant k$. Notice that a minimal edge-cut separates the graph into exactly 2 connected components. Hence if $y \notin X$ then $\delta(X)$ is a $(zy|x)$-cut and the answer is YES. Otherwise, let G' be the multigraph obtained from G by contracting $V(G) \setminus X$ into a single vertex z', keeping multiple edges.

Claim 1 *There exists a $(zy|x)$-cut of size at most k in G if and only if there exists a $(z'y|x)$-cut of size at most k in G'.*

Proof. Assume first that there exists a $(z'y|x)$-cut F' of size at most k in G'. Let W be the connected component of $G' \setminus F'$ containing x. We have that $|\delta_G(W)| = |\delta_{G'}(W)|$, hence $\delta_G(W)$ is a $(zy|x)$-cut of size at most k in G.

Conversely, let F be a $(zy|x)$-cut of size at most k in G, and let Z be the connected component of $G \setminus F$ containing z and y. Let $Z' \subseteq V(G')$ be the set $Z \cap X \cup \{z'\}$. The set $\delta_{G'}(Z')$ is a $(z'y|x)$-cut in G'. By definition, we have $|\delta_{G'}(Z')| = |\delta_{G'}(Z \cap X \cup \{z'\})| = |\delta_G(Z \cup \overline{X})|$. By submodularity of the border, we know that $|\delta_G(Z)| + |\delta_G(\overline{X})| \geq |\delta_G(Z \cup \overline{X})| + |\delta_G(Z \cap \overline{X})|$. Thus, $|\delta_{G'}(Z')| \leq |\delta_G(Z)| + |\delta_G(X)| - |\delta_G(Z \cup \overline{X})|$. By minimality of X, we know that $|\delta(X)| \leq |\delta_G(Z \cup \overline{X})|$, which gives $|\delta_{G'}(Z')| \leq |\delta_G(Z)| \leq k$. $\diamond$

We are now looking for a $(z'y|x)$-cut of size at most k in G'. By minimality of X, we know that $\delta_{G'}(X)$ is the only minimum (z', x)-cut in G' (otherwise we would find a smaller set $X' \subsetneq X$ containing x of border of size c in G). We thus have to check the cuts of size l with $c + 1 \leq l \leq k$. By definition, such a cut does not contain all edges of $\delta_{G'}(X)$, because otherwise it would be a (z, y)-cut. We thus branch over the c edges adjacent to z', obtaining c new instances where the considered edge has been contracted to a single vertex $\tilde{z}$. Doing so, we strictly increase the connectivity between z' and x by minimality of X. We now have to decide if there exists a $(\tilde{z}y|x)$-cut in a graph where connectivity between $\tilde{z}$ and x is at least $c+1$. Note that if the contracted edge was $z'y$, we just need to decide whether there exists a cut of size at most k between $\tilde{z}$ and x. Since $c \leq k$ and since the connectivity between $\tilde{z}$ and x strictly increases at each step, the whole branching algorithm runs in time $O(k! \times poly(n))$. $\square$

We say that a vertex x is *k'-strongly $(z|T)$-k-linked* if and only if for every $S \subseteq T$ such that $|S| \geq |T| - k'$, x is $(z|S)$-k-linked (*i.e.* if there is no $(zx|S)$-cut of size at most k). Note that when x is k'-strongly $(z|T')$-k-linked, x is a fortiori k'-strongly $(z|T)$-k-linked in G when $T' \subseteq T$. In particular, we get the following corollary by using Theorem 1 on every subset $S \subseteq T$ such that $|S| \geq |T| - k'$, contracting set S into a single vertex:

Corollary 1 *Deciding whether a vertex x is k'-strongly $(z|T)$-k-linked and producing a witness of non-strong-linkness, i.e. a $(zx|S)$-cut for S a subset of T of size at least $|T| - k'$, is FPT in k, k' and $|T|$.*

The following Theorem is a key result of this paper:

Theorem 2 *Let $G = (V, E)$ be a graph, z a vertex of $V(G)$ and T a subset of $V(G) \setminus \{z\}$ of size at least $f(k, k')$, for a large enough $f(k, k')$. There exists a vertex $x \in T$ which is k'-strongly $(z|T \setminus \{x\})$-k-linked and which can be found in FPT time in k, k' and $|T|$.*

Proof: Assume that T has size at least $f(k, k') = k^{k+k'+1}(\frac{k'+1+k^2}{k-1} + 1) - \frac{k'+1+k^2}{k-1}$.

We initiate our algorithm with the graph $f_6 := G$ and let $T_0 := T$. We repeat the following process $k + k' + 2$ times, selecting a vertex x_i in T_i for $i = 0, \ldots, k + k' + 1$. We first test whether x_i is k'-strongly $(z|T_i \setminus \{x_i\})$-$k$-linked in G_i as in Corollary 1. If this is the case, then we are done. If this is not the case, we obtain a subset S of T_i of size at least $|T_i| - k' - 1$ and a $(z|S)$-cut C_i of size at most k which is not a (z, x_i)-cut. Denote by L_i the connected component of $G_i \setminus C_i$ containing z, and consider the connected component V_{i+1} of $G_i \setminus C_i$ not containing z and containing the largest number of vertices of S. Let G_{i+1} be the graph obtained from $G[V_{i+1} \cup L_i]$ by contracting $L_i \cup V(C_i)$ into z, where $V(C_i)$ denotes the set of vertices incident

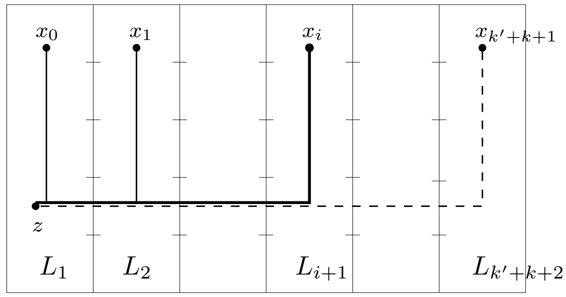

Figure 1: Vertex $x_{k'+k+1}$ is k'-strongly $(z|T \setminus \{x\})$-k-linked in the proof of Theorem 2.

to C_i. Note that $V(G_{i+1})$ contains a large subset $T_{i+1} = T_i \cap (V_{i+1} \setminus V(C_i))$ of vertices of T, and that x_i is well defined for every $i = 1, \ldots, k + k' + 1$. Indeed, $|T_i| \geq \frac{|T_{i-1}|-k'-1}{k} - k$. Equivalently, this means that $|T_i| + \frac{k'+1+k^2}{k-1} \geq \frac{1}{k}(|T_{i-1}| + \frac{k'+1+k^2}{k-1})$. Thus the sequence $(|T_i| + \frac{k'+1+k^2}{k-1})$ follows a geometric progression of factor $\frac{1}{k}$. So that $|T_{k+k'+1}| \geq 1$ we need that $|T| + \frac{k'+1+k^2}{k-1} = |T_0| + \frac{k'+1+k^2}{k-1} \geq k^{k+k'+1}(\frac{k'+1+k^2}{k-1} + |T_{k+k'+1}|) \geq k^{k+k'+1}(\frac{k'+1+k^2}{k-1} + 1)$, that is $|T| \geq k^{k+k'+1}(\frac{k'+1+k^2}{k-1} + 1) - \frac{k'+1+k^2}{k-1}$, which is the assumption of the Theorem.

If the algorithm did not stop after step $k + k' + 1$, this means that for every $i = 1, \ldots, k + k' + 1$, x_i is not k'-strongly $(z|T \setminus \{x_i\})$-k-linked in G_i.

Let us denote $T_e := \{x_i : i = 0, \ldots, k + k'\}$. We will show that $x_{k+k'+1}$ is k'-strongly $(z|T_e)$-k-linked in G, which implies that $x_{k+k'+1}$ is k'-strongly $(z|T \setminus \{x_{k+k'+1}\})$-$k$-linked in G. Consider a path P between $x_{k+k'+1}$ and z. Note that $C_i \cap C_j = \emptyset$ for $i \neq j$, which implies that each edge of C_i has one endpoint in L_i and one endpoint in L_{i+1}. For every $i = 0, \ldots, k + k'$, P intersects L_i since P intersects C_i. For every $i = 0, \ldots, k+k'$, let P_i be a path between x_i and z in G which is included in P outside L_i. Such a path exists because L_{i+1} is connected. This implies that, for every set $S_e \subseteq T_e$ of size at least $k + k' + 1 - k' = k + 1$, and for every path P between $x_{k+k'+1}$ and z, there exist $k+1$ paths P_j for $j = 1, \ldots, k+1$ between vertices of S_e and z, such that $P_j \cap P_l \subseteq P$ for $j \neq l$. Hence, every path P between $x_{k+k'+1}$ and z must be cut by every cut of size at most k between z and S_e, which thus separates $x_{k+k'+1}$ and z. □

Lemma 2 *Rule 2 is safe and can be applied in FPT time.*

Proof: Let t be a vertex endpoint of at least $g(k) = f(k, k)$ requests, and $\tilde{T}$ be the corresponding endpoints. Since $|\tilde{T}| \geqslant f(k, k)$, Theorem 2 implies that there exists a vertex $\tilde{t} \in \tilde{T}$ which is k-strongly $(t|\tilde{T} \setminus \{\tilde{t}\})$-$k$-linked in G, so $\tilde{t}$ is $(t|\tilde{T} \setminus \{\tilde{t}\})$-k linked in G. Let F be a k-multicut of $(G, R \setminus (t, \tilde{t}), k)$. By definition, F is a $(t, \tilde{T} \setminus \{\tilde{t}\})$-cut in G, and hence a $(t, \tilde{t})$-cut in G. It follows that F is actually a k-multicut for (G, R, k) and thus that the request $(t, \tilde{t})$ is irrelevant. □

Lemma 3 *Rule 3 is safe and can be applied in FPT time.*

Proof: Let T be a gathered set of terminals of size at least $4f(k, k)^3$, and $S = \{s|(s, t) \in R, t \in T\}$. By Rule 2 each vertex $t \in T$ is the endpoint of at most $f(k, k)$ requests. Hence there exist two sets $T' = \{t'_1, \ldots, t'_p\} \subseteq T$ and $S' = \{s', \ldots, s'_p\} \subseteq S$ with $p = 4f(k, k)^2$ such that (t'_i, s'_i) is a request for every $i = 1, \ldots, p$.

We now define an auxiliary bipartite graph $B = (T' \cup S', E')$ where $t'_i s'_j \in E'$ if $i \neq j$ and there exists a cut of size at most k in G such that t'_i and s'_j lies in a same connected component different from C_F in $G \setminus F$. For every $i \neq j$ we use Theorem 1 on the graph obtained from G by contracting $T' \setminus \{t'_i\}$ into a single vertex x to test whether there exists a $(t'_i s'_j|x)$-cut of size at most k. If so, then t'_i and s'_j are in the same connected component in $G \setminus F$, which is not C_F since t'_i is cut from any vertex of $T' \setminus \{t'_i\}$, and t'_i and s'_j are adjacent in B. Hence B can be constructed in FPT time in k.

Claim 2 *Every vertex $s' \in S'$ has degree at most $f(k, k)$ in B.*

Proof. Assume there exists a vertex $s' \in S'$ with degree at least $f(k,k)$, and let x be any of its neighbours. By Theorem 2, x is k-strongly $(s'|N_B(s') \setminus \{s'\})$-$k$-linked. Since x and s' are adjacent in B, there exists a set $F \subseteq E(G)$, $|F| \leq k$ such that s' and x belongs to a same component different from C_F in $G \setminus F$. Since T is a gathered set, such a cut F must cut x from $N_B(s) \setminus \{x\}$ and is thus a cut between x and s', a contradiction. $\diamond$

Claim 3 *There exists $T'' \subseteq T'$ and $S'' \subseteq S'$ of size $f(k,k)$ such that for every set F of at most k edges, if $t'_i \in T''$ and $s'_j \in S''$ both lie in a component $C \neq C_F$ of $G \setminus F$, then $i = j$.*

Proof. By Claim 2, we have $|E'| \leqslant |S'| \cdot f(k,k)$. Hence there exists a set J of at least $\frac{|T'|}{2}$ vertices of T' with degree at most $2f(k,k)$ in B. Let $I = \{s : (t,s) \in R, t \in J\}$. The degree of $B[J \cup I]$ is at most $2f(k,k)$ implying that we can greedily find two sets $T'' = \{t''_1, \ldots, t''_{f(k,k)}\} \subseteq J$ and $S'' = \{s''_1, \ldots, s''_{f(k,k)}\} \subseteq I$ such that (t''_i, s''_i) is a request for every i and $B[T'' \cup S'']$ does not contain any edge. By construction, this means that for every set F of at most k edges if t''_i and s''_j both lie in a component $C \neq C_F$ of $G \setminus F$ then $i = j$. $\diamond$

By Theorem 2, there exists a vertex $s'' \in S''$ which is k-strongly $(z|S'' \setminus \{s''\})$-$k$-linked in $\tilde{G}$, where $\tilde{G}$ is the graph obtained from G by contracting T'' to a single vertex z. We conclude the proof by showing the following:

Claim 4 *The request (t'', s'') is irrelevant.*

Proof. Let F be a k-multicut for the instance $(G, R \setminus (t'', s''), k)$. In particular, F is a (T''_F, S''_F)-cut where $T''_F = T'' \cap C_F$ and S''_F denotes the set of terminals corresponding to $T'' \cap C_F$. Observe that F is a (z, S''_F)-cut in $\tilde{G}$: otherwise this would imply that there exists a vertex of $T'' \setminus T''_F$ lying in a component $C \neq C_F$ of $G \setminus F$ which also contains a vertex of S''_F in $G \setminus F$, which cannot be by Claim 3. Hence F is a (s'', z)-cut in $\tilde{G}$ and thus a (t'', s'')-cut in G, implying that F is actually a k-multicut for (G, R, k). $\diamond$

This concludes the proof of Lemma 3. $\square$

The rest of the paper is essentially devoted to finding a situation where Rule 3 can be applied.

4 Clique Minor

In this section, we assume that G has a large clique minor. Let H be an $f_2(k)$-clique model of G, for a large enough $f_2(k)$. Such a model can be computed in FPT time in k. Let a be a vertex of G lying in a part h_a of H. The vertex a is *special* if there exist $k+1$ vertex-disjoint paths internally in h_a between a and distinct parts of H different from h_a.

4.1 Finding a nice model

We say that a model H of a graph G is *nice* whenever it satisfies the following properties:

(1) $|V(H)| \geq f_2(k)$ and H is $(k+1)$-vertex connected.

(2) There exists at most one special vertex a in G (and we denote its part by h_a).

(3) All parts but possibly h_a have degree at most $(k+1)^2 + 1$ in H.

(4) The special vertex a separates $h_a \setminus \{a\}$ from $V(G) \setminus h_a$ in G.

Theorem 3 *Let (G, R, k) be an instance of* Multicut *reduced under Rule 1 which admits an $f_2(k)$-clique model H. There exists a nice model of G which can be computed in FPT time in k.*

Proof: Let us first show that H respects the second property of the nice model definition. Observe that any $f_2(k)$-clique model for G is in particular $(k+2)$-vertex connected since we may assume $f_2(k) \geqslant k+2$.

Claim 5 *Let (G, R, k) be an instance of* MULTICUT *reduced under Rule 1 that admits a $(k+1)$-vertex connected model H. Then G contains at most one special vertex.*

Proof. Assume by contradiction that there exist two vertices a and a' having at least $k+1$ vertex-disjoint paths $P_1, \ldots, P_{k+1}$ and $P'_1, \ldots, P'_{k+1}$ in their own parts h_a and $h_{a'}$ towards distinct parts different from h_a (resp. $h_{a'}$) in H. Since G is reduced under Rule 1, there are at most k vertex-disjoint paths between a and a' in G. Hence by Menger's theorem [18] there exists a set S of at most k vertices of G which removal disconnects a and a'. Let $\mathcal{S}$ be the set of parts of H containing vertices of S. By hypothesis, there exists at least one path P_i towards a neighbour h_i of h_a and one path P'_i towards a neighbour h'_i of h'_a such that $P_i \cup h_i$ and $P'_i \cup h'_i$ are not intersected by $\mathcal{S}$. By $(k+1)$-vertex connectivity of H there exists a path between h_i and h'_i in $H \setminus \mathcal{S}$. Since every part of $H \setminus \mathcal{S}$ is connected, such a path can be extended to a path between a and a' in $G \setminus S$, leading to a contradiction. $\diamond$

We now show a technical result:

Lemma 4 *Let $G = (V, E)$ be a graph that admits a p-vertex connected model H with $p > k$ and assume G contains at most one special vertex a. If there is a part $h \neq h_a$ with degree greater than $(p-1)(k+1)$ in H, then there exists a bipartition h_1, h_2 of h such that h_i is connected in G for $i = 1, 2$ and the partition $\mathcal{P} \setminus h \cup \{h_1, h_2\}$ of $V(G)$ is a p-vertex connected model of G, where $\mathcal{P}$ is the partition of $V(G)$ associated to H.*

Proof: We need the following claims:

Claim 6 *Let T be a tree of maximum degree d, with a weight function ω from the nodes of T into $\{0, \ldots, M\}$, such that $\omega(T) > q(d+1)$ with $q \geq M$. There exists an edge of T separating T into two subtrees T_1 and T_2 such that $\omega(T_i) \geqslant q+1$ for $i = 1, 2$.*

Proof. Assume that T contradicts the lemma. For each edge e in T, let T_1^e and T_2^e be the connected components of $T \setminus e$. We have $\omega(T_1^e) \leq q$ or $\omega(T_2^e) \leq q$ (and these cases are mutually exclusive). Let us orient T: e gets the orientation $T_1^e \to T_2^e$ if $\omega(T_1^e) \leq q$. Note that each edge uv gets the orientation $u \to v$ whenever u is a leaf. Since this orientation is acyclic, there exists an internal node x of T such that all edges incident to x are oriented towards x. Let $T_1^x, \ldots, T_l^x$ be the connected components of $T \setminus x$. We have $l \leq d$. Thus $w(T) = \omega(T_1^x) + \ldots + \omega(T_l^x) + \omega(x) \leq l * q + M \leq q(d+1)$, a contradiction. This concludes the proof of Claim 6. $\diamond$

Claim 7 *Let H be a c-vertex-connected graph and $h \in V(H)$ be a vertex of degree at least $2c$. Let N_1, N_2 be a bipartition of $N(h)$, with $|N_i| \geq c$ for $i = 1, 2$. Let H' be the graph obtained from H by deleting h and adding two vertices h_1 and h_2, of respective neighbourhoods $N_1 \cup \{h_2\}$ and $N_2 \cup \{h_1\}$. Then H' is c-connected.*

Proof. Assume that there exists a set S of $c-1$ vertices which removal disconnects H'. If $h_i \notin S$ for $i = 1, 2$, then S disconnects H, a contradiction. If h_1 and h_2 both lie in S, then $(S \setminus \{h_1, h_2\}) \cup \{h\}$ disconnects H, a contradiction. If h_i lies in S, for $i = 1$ or $i = 2$, then $\{h_{3-i}\}$ is not a connected component of $H \setminus S$ since it has at least c neighbours besides h_i in H'. Thus $S \setminus \{h_i\} \cup \{h\}$ is a $(c-1)$-cut in H, a contradiction. This concludes the proof of Claim 7. $\diamond$

We use these two results to prove Lemma 4. Consider a part $h \neq h_a$ of degree greater than $(p-1)(k+1)$ in H. For each part h' adjacent to h in H, we distinguish one vertex $x \in h$ such that there exists an edge xx' in G with $x' \in h'$. Denote this vertex x by $v(h')$. Let T be a tree of $G[h]$ spanning all vertices x such that $v^{-1}(x) \neq \emptyset$, which is minimal by inclusion. The leaves of T are vertices such that $v^{-1}(x) \neq \emptyset$. Let ω be a weight function on h, such that $\omega(x) = |v^{-1}(x)|$.

Since G contains at most one special vertex, vertices distinct from a can have at most k neighbours in distinct other parts, implying that $\omega(x) \leqslant k$ for every vertex $x \neq a$. Moreover, the degree of any vertex in T is at most k, since a vertex with degree at least $k+1$ would have $k+1$ vertex-disjoint paths towards leaves of T and hence towards $k+1$ distinct parts different from its own part, which cannot be by hypothesis. By construction, $\omega(T)$ is equal to the degree of h in H, so $\omega(T) > (p-1)(k+1)$. By Claim 6 applied with $q = p-1$ and $d = k$, there exists a bipartition T_1, T_2 of T such that $\omega(T_i) \geqslant p$. Observe that this bipartition can be extended into a bipartition of h into two connected sets h_1, h_2 such that $|N_H(h_1)|, |N_H(h_2)| \geqslant p$, by definition of the weight function ω. By Claim 7 it follows that the partition $\mathcal{P} \setminus h \cup \{h_1, h_2\}$ defines a model H' which is p-connected. This concludes the proof of Lemma 4. □

Assume now that H does not contain a special vertex, and let h be a part of H with degree at least $(k+1)^2+1$ in H. Since H is $(k+2)$-vertex connected, it follows by Lemma 4 applied with $p = k+2$ that we can find in polynomial time a new model of G which is $(k+2)$-vertex connected, which contains at most one special vertex by Claim 5, and with size strictly greater than $V(H)$. We apply repeatedly Lemma 4 until no part has degree at least $(k+1)^2+1$ in H_i or there exists a special vertex a in a part h_a. In the former case, since the model finally obtained is $(k+2)$-vertex connected and does not contain any special vertex, it is actually a nice model, thus we are done. In the latter case, we modify the finally obtained model to obtain a new model H' such that the special vertex a cuts $h_a \setminus \{a\}$ from $V(G) \setminus h_a$ in G as follows.

For every vertex $v \neq a \in h_a$ having a neighbour in a part $h_v \neq h_a$, we denote by A_v the set of vertices disconnected from a in h_a by the removal of v. We remove the set $A_v \cup \{v\}$ from h_a and add it to h_v, repeating this process until no vertex in h_a but a is adjacent to other parts. Observe that the model H' thus obtained may no longer be $(k+2)$-vertex connected after this process: however, $H' \setminus h_a$ remains $(k+1)$-vertex connected. Since h_a has degree at least $k+1$ in H_i, it follows that H' is a $(k+1)$-vertex connected model. Observe that H' contains exactly one special vertex a by Claim 5.

Since this process may increase the degree of some parts different from h_a in H', we apply repeatedly Lemma 4 to H' to reduce its degree while preserving $(k+1)$-vertex connectivity. Once this process is over, we obtain a $(k+1)$-vertex connected model H' such that $|V(H')| \geqslant f_2(k)$, which special vertex a separates $h_a \setminus \{a\}$ from $V(G) \setminus h_a$ in G, and such that every part but possibly h_a has degree at most $(k+1)^2$. It follows that H' is a nice model, which concludes the proof of Theorem 3. □

4.2 Small and giant components

In the following we consider a nice model H of the instance (G, R, k) of the Multicut problem.

Lemma 5 *Let H be a nice model of G. Two parts not touched by a set $F \subseteq E(G)$ of at most k edges are included in the same connected component of $G \setminus F$.*

Proof: Assume that a set F of at most k edges does not touch parts h_1 and h_2. It follows that part h_1 (resp h_2) is completely included in a connected component of $G \setminus F$. By $(k+1)$-connectivity in H, there are $k+1$ disjoint paths between h_1 and h_2 in H, which cannot all be cut by a set of at most k edges. □

Corollary 2 *Let H be a nice model of G, h_1, h_2 be two parts of H and (t_1, t_2) a request with $t_1 \in h_1$ and $t_2 \in h_2$. Every k-multicut contains an edge that belongs to h_1 or an edge that belongs to h_2.*

Observe that any set of at most k edges cuts the graph G into at most $k+1$ connected components. Given a set F of edges, we call *giant component* of $G \setminus F$ and denote by C_F a connected component of $G \setminus F$ containing entirely at least $|V(H)| - k$ parts.

Lemma 6 *Let F be a set of at most k edges of a graph G admitting a nice model H. Then $G \setminus F$ has a giant component, which contains all parts not touched by F.*

Proof: Let h_1 and h_2 be two parts not touched by a set F of at most k edges. Then h_1 and h_2 belong to the same connected component in $G \setminus F$ by Lemma 5. Since there are at most k parts touched by F, all other parts are entirely included in the same connected component of $G \setminus F$. □

The connected components distinct from the giant component are called *small components*. They intersect vertex-wise altogether at most k parts.

Lemma 7 *Assume that G admits a nice model. For every set F of at most k edges, the special vertex belongs to the giant component of $G \setminus F$.*

Proof: Since a is adjacent to at least $k+1$ parts, a is adjacent to at least one vertex lying in a part not touched by F in $G \setminus F$. By Lemma 6, a cannot belong to a small component. □

Lemma 8 *Let F be an optimal k-multicut of a graph G admitting a nice model. Every small component C of $G \setminus F$ contains a vertex which is a terminal.*

Proof: Let F be an optimal k-multicut of G and e be any edge of F which touches C. If C contains no terminal then $F \setminus \{e\}$ is still a multicut, contradicting the optimality of F. □

Let $H_a = H \setminus h_a$. Adjacent vertices x and y which are far away in H_a from any terminal can be contracted as stated in the following Lemma:

Lemma 9 *Let G be a graph admitting a nice model H. Let h be a part of H_a not containing any terminal, such that every part at distance at most $k+2$ from h in H_a has no terminal. Let e be an edge of G with at least one endpoint in h. Then e does not belong to any optimal k-multicut.*

Proof: Consider an optimal k-multicut F of G, and let h' be a part adjacent to h in H. Every small component contains a terminal by Lemma 8 and intersects at most k parts by definition. Since part h (resp. h') is assumed to be too far from any terminal in H_a, part h (resp. h') can intersect a small component C of $G \setminus F$ only if $a \in C$, which contradicts Lemma 7. Hence neither part h nor part h' can intersect a small component. It follows that if e belongs to F then $F \setminus \{e\}$ is also a k-multicut, which contradicts the optimality of F. □

Hence the following reduction rule is correct:

Rule 4 *Let G be a graph admitting a nice model H and h be a part of H_a without any terminal. If every part at distance at most $k+2$ from h in H_a has no terminal, then contract an arbitrary edge e of G with at least one endpoint in h.*

When this reduction rule does no longer apply, every part is at distance at most $k+2$ in H_a from a part containing a terminal.

4.3 Reducing the instance

Let us review the structure of the graph when none of Rule 1, Rule 2 and Rule 4 applies. Recall that we consider an instance (G, R, k) that contains a nice model H.

We consider the level decomposition of H_a starting from some part L_0 of H_a. Since H is a nice model, it follows that every part of H_a has degree at most $(k+1)^2$ in H_a. Hence every level L_i has size at most $(k+1)^2|L_{i-1}|$. Denote by d the number of non-empty levels in this decomposition. We have $f_2(k) \leq |V(H_a)| + 1 \leq (\sum_{i=0}^{d}((k+1)^2)^i) + 1 \leq (k+1)^{2(d+1)}$, so $d \geq log_{(k+1)^2}(f_2(k)-1) - 1$.

Let us show that we can find a gathered set and thus that the graph can be reduced using Rule 3. Let T be a set of maximum size of terminals in $V_G(H_a)$ belonging to different parts lying pairwise at distance at least $k+1$ in H_a. Since every part is at distance at most $k+2$ from a terminal in H_a (Rule 4), there exists at least one terminal in every $2(k+2)+1$ consecutive levels, implying that T has sufficiently large size (recall that we assumed $f_2(k)$ to be large enough).

Let us argue that T is a gathered set. Indeed, consider any set F of at most k edges. Since parts containing terminals of T are pairwise at distance at least $k+1$ in H_a, a small component of $G \setminus F$ can contain two terminals of T only if it contains the special vertex a (Lemma 6). Since a belongs to the giant component by Lemma 7, it follows that $G \setminus F$ contains at most one component containing more than one vertex of T (namely the giant component), and hence T is a gathered set. Using Rule 3, an irrelevant request can be found in FPT time. Hence we have shown in this Section that whenever G has a large clique model, we can reduce the instance in FPT time either by safely contracting an edge or by finding an irrelevant request. This concludes the large clique minor part.

5 Grid Minor

The aim of this section is to complete the treewidth reduction by showing that a graph with a large grid minor but no large clique minor can also be reduced in FPT time. Indeed, every graph with treewidth at least 2^{2m^5} contains an $(m \times m)$-grid minor [21], and a model for such a grid minor can be computed in FPT time in m.

Given a grid model $H = (V(H), E(H))$ of a graph G we denote by $\bar{H}$ a graph $(V(H), E(\bar{H}))$ where $E(\bar{H}) \subseteq E(H)$ is such that $\bar{H}$ is a grid. We call the set of vertices of degree two or three in a grid H' the *border* of the grid, and denote it by $B(H')$.

5.1 On grid minors without clique minors

In this section, we concentrate on structural properties of graphs with large grid minors but no large clique minors. Let G_k be the graph obtained from the $(k \times k)$-grid by adding the two diagonals to each internal face of the grid (see Figure 2 in Appendix).

Lemma 10 *The crossed grid G_{2k} has a k-clique minor.*

Proof: We are looking for a k-clique model inside G_{2k}, that is for k vertex-disjoint connected sets of vertices of G_{2k} such that any two of these sets are adjacent in G_{2k}. Roughly speaking, these subsets will be the union of the ith column and ith row of G_{2k} for odd i. These sets are not vertex disjoint, but this can easily be dealt with: for any crossing of two sets, we use two diagonals inside a face of the original grid to preserve connectivity while uncrossing them.

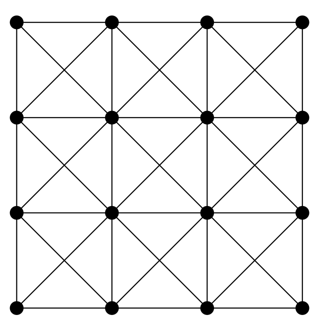

Figure 2: The crossed grid G_4.

More formally, denote by (i, j) the vertex lying on row i and column j of the crossed grid G_{2k}. Let S_i be the set of vertices defined as follows:

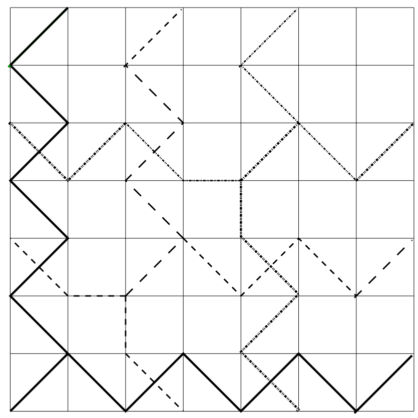

Figure 3: A clique model in the crossed grid G_8

$$S_i := \begin{cases} (i, i+l) \text{ and } (i+l, i) \text{ for even } l,\ l \geqslant 0 \\ (i+1, i+l) \text{ and } (i+l, i+1) \text{ for even } l,\ l < 0 \\ (i+1, i+l) \text{ and } (i+l, i+1) \text{ for odd } l,\ l \geqslant 0 \\ (i, i+l) \text{ and } (i+l, i) \text{ for odd } l,\ l < 0 \end{cases}$$

with $1 \leqslant i + l \leqslant 2k$. The sets S_i for odd i are a model of a K_k of G_{2k}. Indeed they are disjoint, and for every odd i, j, $i < j$ we have $(i, j) \in S_i$ and $(i + 1, j) \in S_j$ (see Figure 5.1). □

Given a grid G, an edge $e \notin E(G)$ is said to be *non-planar* if the graph $G \cup \{e\}$ is not planar. We call a planar supergraph of a grid an *augmented grid*.

We now show that a grid together with many non-planar edges has a large clique minor. In [20], Robertson and Seymour give a result (7.4) similar to what we need. To reformulate it in a form closer to our terminology:

Lemma 11 *[20] For every integer m, there exists an integer $r(m)$ such that a graph obtained from an augmented grid H by adding $\binom{m}{2}$ non-planar edges in $(V(H) \times V(H)) \setminus E(H)$, which endpoints lie at distance at least $r(m)$ from $B(H)$ in H and are pairwise lying at distance at least $r(m)$ in H, has a K_m minor.*

Note that result (7.4) from [20] concerns walls rather than grids and is stated with a different distance function, but Lemma 11 can be deduced from there.

It is important that the extra edges are non-planar in Lemma 11, and that they lie far away from the border of the grid. The fact that the extra edges are long is actually not necessary, as stated in the following stronger result.

Lemma 12 *For every integer m, there exists two integers $r(m)$ and $f_3(m)$ such that a graph obtained from an augmented grid H by adding $f_3(m)$ non-planar endpoint disjoint edges in $(V(H) \times V(H)) \setminus E(H)$, which endpoints lie at distance at least $r(m)$ from $B(H)$ in H, has a K_m minor.*

Proof: In the following we let:

$$f_3(m) = 2 \cdot 8^{2(r(m)+m)} \binom{m-1}{2} + 2 \cdot 8^{r(m)} \binom{m}{2}$$

If $2 \cdot 8^{r(m)} \binom{m}{2}$ non-planar edges have length at least m in the grid, then there exists a subset of size $\binom{m}{2}$ of these non-planar edges which endpoints lie pairwise at distance at least $r(m)$ in the grid. By Lemma 11, it follows

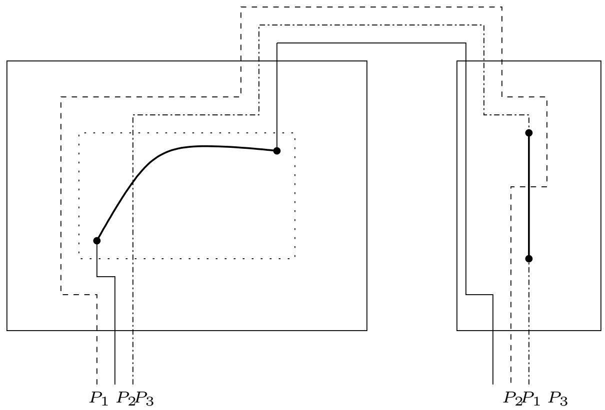

Figure 4: Illustration of the swaps made within two consecutive zones in the Proof of Lemma 12.

that H has a clique minor of size m. Otherwise, there exists a subset X of at least $f_3(m) - 2 \cdot 8^{r(m)}\binom{m}{2} \geq 2 \cdot 8^{2(r(m)+m)}\binom{m-1}{2}$ non-planar edges of length at most $r(m)$. Hence there exists a subset $S \subseteq X$ of non-planar edges which endpoints lie pairwise at distance at least $2(r(m)+m)$ of size at least $\binom{m-1}{2}$. For every edge $e = xy \in S$, let S_e be the minimal subgrid containing both x and y. The *zone* Z_e of e is the subgrid obtained from S_e by adding m layers around S_e. Note that for any two edges $e, e' \in S$ we have $Z_e \cap Z_{e'} = \varnothing$. There exists a set of m vertex disjoint paths $P_1, \ldots, P_m$ which are parallel and adjacent outside the zones and enter all zones in the grid exactly once.

A zone can be used to exchange two adjacent paths, as shown in Figure 4.

- use the $(m-1)$ first encountered zones to place P_m before P_1
- use the $(m-2)$ next zones to place P_{m-1} before P_1
- ...
- use the $(m-i)$ next zones to place P_{m-i} before P_1
- ...
- use the next zone to place P_2 before P_1

We have $\frac{m(m-1)}{2}$ zones, enough to do the $\sum_{i=1}^{m-1} i = \frac{m(m-1)}{2}$ above swaps, and a given zone is large enough to make two adjacent paths cross within it. We have reversed the original planar left-to-right order of paths P'_i, $i \in \{1, \ldots, m\}$ by making consecutive swaps, hence any two paths have crossed and hence are adjacent in the grid. Thus paths P_i, $i \in \{1, \ldots, m\}$ form a K_m minor model. □

Given a grid model H we say that $H' \subseteq H$ is a *proper subgrid* of H if it satisfies the following properties:

- $|V(H')| \geq f_4(k)$.
- H' is an induced augmented grid.
- There are at most $f_5(k)^2$ vertices in $V(G) \setminus V(H')$ with neighbours in $V(H') \setminus B(H')$. Such vertices are called *special vertices* and are denoted by W.
- Each special vertex has at least $k^2/4$ neighbours in $V(H') \setminus B(H')$.

Theorem 4 *Let G be a graph admitting an $(f_6(k) \times f_6(k))$-grid model H, with $f_6(k)$ large enough, but no $K_{f_2(k)}$ minor. There exists a proper subgrid which can be computed in polynomial time.*

Proof: We first show a technical result that will be used in the proof of Theorem 4. Given a grid model H of a graph G and an induced augmented grid H' which lies at distance at least $r(f_2(k))$ from $B(\bar{H})$ in $\bar{H}$, an edge wz where $z \in V_G(H') \setminus V_G(B(H'))$ lies at distance at least $r(f_2(k))$ of the border of H' in H' and $w \in V_G(H \setminus H')$ is called a *H'-matched edge*.

Lemma 13 *Let G be a graph admitting a grid model H with an induced augmented grid H' of size at least $f_4(k)$ which lies at distance at least $r(f_2(k)+1)$ in $\bar{H}$ from the border of $\bar{H}$. If there exist at least $f_5(k)$ endpoint disjoint H'-*matched edges *in H, with $f_5(k)$ large enough, then one of the following holds:*

(1) either G admits a grid model H_1 and an induced augmented grid $H'_1 \subseteq H_1$ such that $|H'_1| \geqslant f_4(k)$ and every set of endpoint disjoint H'_1-matched edges has size at most $f_5(k)$,

(2) or G has an $f_2(k)$ clique minor.

Proof: In the following we assume that all considered functions are large enough to prove the results. The following result follows by definition of a minor:

Observation 1 *Let G be a graph admitting a grid model H with an induced augmented grid H' such that there exist p paths internally in $G \setminus V_G(H')$ which are vertex-disjoint, between pairs of vertices $z_i, z'_i \in V_G(H')$ at distance at least $r(f_2(k))$ from the border of H' in H', $1 \leq i \leq p$. If edges $z_i z'_i$ are non-planar in H' for $i = 1, \ldots, p$, then G has as a minor a grid $\check{H}$ with p non-planar edges which endpoints lie at distance at least $r(f_2(k))$ from the border of $\check{H}$ in $\check{H}$.*

Let $W = \{w_1, \ldots, w_p\}$ be the endpoints of the H'-matched edges that belong to $V_G(H \setminus H')$, and assume $|W| \geqslant f_5(k)$. Since the graph $G \setminus V_G(H')$ is connected, let T be a tree spanning W in $G \setminus V_G(H')$. We can assume that this tree where parts have been contracted is a path, since $H \setminus V_G(H')$ has a hamiltonian path. Moreover, we can assume that any vertex of a given part has at most one neighbour in T in each of its two adjacent parts. Indeed, assume that a vertex v in a part h has more than one neighbour in T in one of its adjacent part h'. We choose any vertex $u \in N_{h'}(v)$ and remove all edges between v and $N_{h'}(v) \setminus u$ from T. We now use the connectivity of h' to find a spanning tree of all vertices in $N_{h'}(v)$, which gives a new spanning tree T' having the claimed property. Finally, we choose such a tree T with minimum maximum degree. We again distinguish two cases.

Case 1: T has degree bounded by $f_8(k)$ We use the following Lemma to exhibit a large clique minor in G:

Lemma 14 *Let T be a tree of degree at most d. Let W be a set of nodes of T. There exist $\frac{|W|-1}{d+1}$ vertex-disjoint paths in T with distinct endpoints in W.*

Proof: We root T arbitrarily. The hypothesis trivially holds for $|W| = 1$, so assume that $|W| \geq 2$. Among all pairs of vertices in W, we choose one pair (x, y) which least common ancestor u is minimum for the ancestor relation. Let T_u be the subtree of T rooted at u. As u is minimum, each subtree rooted in a son of u contains at most one vertex of W. Hence, as u has degree at most d, T_u contains at most $d+1$ vertices of W. By induction on $T \setminus T_u$, there exist $\frac{|W|-(d+1)-1}{d+1} = \frac{|W|-1}{d+1} - 1$ vertex-disjoint paths in $T \setminus T_u$ with distinct endpoints in W. We add to these paths the disjoint path $(x, y) \subseteq T_u$ to obtain the desired bound. □

Applying Lemma 14 to T and W, we find $\frac{f_5(k)-1}{f_8(k)+1}$ (which we assume to be greater than $f_3(f_2(k))$) vertex-disjoint paths between distinct vertices $w_{i,j}$. Using Lemma 1, G has as a minor a grid with $f_3(f_2(k))$ endpoint disjoint non planar edges with endpoints at distance at least $r(f_2(k))$ from the border of H'. Using Lemma 12, we deduce that G has a $K_{f_2(k)}$ minor.

Case 2: T **has a node** u **of degree more than** $f_8(k)$ Assume that u has maximum degree in T. As T where each part has been contracted to a vertex is a path, u has at least $f_8(k) - 2$ neighbours in T lying in its own part h_u. Among these vertices, we distinguish at most 4 of them, which are vertices connecting h_u to its four adjacent parts in $\bar{H}$. As T has minimum maximum degree, there must exist a set $W_u \subseteq W \cap h_u$ of size at least $f_8(k) - 6$ such that there exist vertex disjoint paths between u and vertices in W_u. For any $w = w_i \in W_u$, we use z_w to denote the corresponding vertex z_i, Z_w to denote the part of z_w, and let A_w be the set of vertices disconnected from u in T by the deletion of w. Note that $A_w \cup W_u = \emptyset$ since u has vertex disjoint paths to vertices in W_u. Let us denote by T_w the set of neighbours of $A_w \cup \{w\}$ distinct from z_w in H'.

Assume that for some vertex $w \in W_u$, there does not exist any part b_w containing a vertex in T_w such that $b_w Z_w$ would be a non-planar edge in H'. We change the partition H by removing $A_w \cup \{w\}$ from part h_u and adding it to part Z_w.

The resulting partition is still a grid model of G since all parts are still connected sets in G and the four distinguished vertices connecting h_u to its four adjacent parts in $\bar{H}$ are still in h_u. This also implies that H' still lies at distance at least $r(f_2(k))$ from the border of $\bar{H}$ in $\bar{H}$. Moreover, H' is still an induced augmented grid of H since no non-planar edge has been created in H'. We go back to the beginning of the proof of Lemma 13 with this new grid model H and augmented subgrid H' of G. We cannot loop more than n times as the number of vertices of G contained in $V_G(H')$ strictly increases. When this process terminates, if we did not find either a clique minor of G or a grid model H with an induced augmented grid H' containing at most $f_5(k)$ endpoint disjoint matched-edges towards vertices in $V_G(H \setminus H')$, then we are still in Case 2. Thus, for every vertex $w \in W_u$ there exists a vertex in $A_w \cup \{w\}$ with a neighbour z'_w in a part $Z'_w \in V(H')$ such that $Z_w Z'_w$ would be a non-planar edge in H'. Hence there exist paths P_w between z_w and z'_w for $w \in W_u$ which are vertex disjoint in h_u, as $A_w \cap A_{w'} = \emptyset$ for $w \neq w' \in W_u$.

Let $Z = \{Z_w | w \in W_u\}$ and $Z' = \{Z'_w | w \in W_u\}$. Note that by construction $Z_w \neq Z_{w'}$ for $w \neq w' \in W_u$. Consider the bipartite graph $B = (Z' \cup W_u, E)$, where $E := \{Z'_w w | w \in W_u\}$.

Claim 8 *Every part in Z' has degree at most $f_9(k)$.*

Proof. Assume by contradiction that there exists a part h in Z' of degree at least $f_9(k)$. Observe that a vertex $x \in h$ cannot have $k+1$ distinct neighbours in W_u, as there would exist $k+1$ edge-disjoint paths between u and x, which contradicts the fact that the instance is reduced under Rule 1. Let $Z'' = \{z'_w | w \in W_u\} \cap h$. Since h is connected, consider a tree T_h of minimum maximum degree spanning Z'' in h.

Assume first that T_h has minimum degree bounded by $k+1$. We apply Lemma 14 to Z'' and T_h to find $\frac{f_9(k)}{k \cdot (k+1+1)}$ (which we assume to be greater than $f_3(f_2(k)+1)$) vertex disjoint paths in T_h between vertices in Z''. This gives $f_3(f_2(k)+1)$ vertex disjoint paths in $G' := (G \setminus V_G(H') \cup \{h\}$ between vertices in Z (see Figure 5). Consider $H'' := (H' \setminus \{h\}) \cup h'$, where h' is a new part having the same adjacency with H' than h. As H'' defines an induced augmented grid, we can apply Lemma 1 to H'' and the above $f_3(f_2(k)+1)$ vertex-disjoint paths of G' to find a grid with $f_3(f_2(k)+1)$ non-planar endpoint disjoint edges with endpoints at distance at least $r(f_2(k)+1)$ from the border of H'. By Lemma 12, the graph H'' has a $K_{f_2(k)+1}$-minor, implying that G has a $K_{f_2(k)}$ minor.

Assume now that T_h has a vertex u_h of degree at least $k+1$. Assume that u_h has maximum degree in T_h. In particular, this means that there exists $k+1$ vertices of Z' with vertex disjoint paths to u_h in T_h. As we may assume $\frac{f_9(k)}{k} \geq k+1$, there exist $k+1$ edge-disjoint paths between u and u_h, which contradicts the fact that the instance is reduced under Rule 1.

This completes the proof of Claim 8. ◇

By Claim 8 we can greedily find $f_{10}(k) = \frac{|W_u|}{f_9(k)}$ vertex-disjoint paths in $V(G) \setminus H'$ between distinct vertices of H' of the form $z_i w_i z'_i$. By Observation 1, G admits as a minor a grid $\bar{H}$ with at least $\frac{f_8(k)-6}{f_9(k)}$ (assumed to be greater than $f_3(f_2(k))$) endpoint disjoint non-planar edges lying at distance at least $r(f_2(k))$ from the border of $\bar{H}$ in $\bar{H}$. Together with Lemma 12, this implies that G admits a $K_{f_2(k)}$-minor. This completes the proof of Lemma 13.

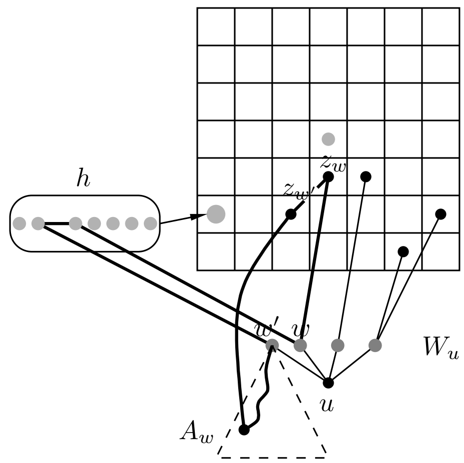

Figure 5: Illustration of the vertex disjoint paths found in the proof of Claim 8.

□

Proof of Theorem 4. Once again we assume that the considered functions are large enough for our purpose. We need the following technical result:

Lemma 15 *Let H be an $(m \times m)$-grid and Z be a set of i^2 vertices of H. Then H contains an $(\lfloor \frac{m}{i+1} \rfloor \times \lfloor \frac{m}{i+1} \rfloor)$-subgrid that does not contain any vertex from Z.*

Proof: We refine H by considering $(i+1)^2$ vertex-disjoint subgrids of size $\lfloor \frac{m}{i+1} \rfloor \cdot \lfloor \frac{m}{i+1} \rfloor$ each. As there are more grids than elements of Z there exists one such subgrid containing no vertex from Z, which is of size $\lfloor \frac{m}{i+1} \rfloor \cdot \lfloor \frac{m}{i+1} \rfloor$. □

Consider the refinement of the subgrid obtained from H by removing its $r(f_2(k))$ first layers in $\bar{H}$, in vertex disjoint subgrids $H_{i,j}$, for $1 \leq i, j \leq 2f_2(k)$ of size $\frac{f_6(k)}{2f_2(k)} \cdot \frac{f_6(k)}{2f_2(k)}$ each.

Claim 9 *If every subgrid $H_{i,j}$ is a non-planar graph, then G has the $(2f_2(k) \times 2f_2(k))$-crossed grid as a minor.*

Together with Lemma 10, Claim 9 implies that there exist $1 \leq i, j \leq 2f_2(k)$ such that $H_{i,j}$ is an induced augmented grid. In the following, we use H' to denote $H_{i,j}$.

Claim 10 *If every set of H'-matched edges has size at most $f_5(k)$ then G admits a proper subgrid model.*

Proof. Recall that an H'-matched edge is an edge wz where $z \in V_G(H') \setminus V_G(B(H'))$ lies at distance at least $r(f_2(k))$ of the border of H' in H' and $w \in V_G(H \setminus H')$. Let $Out = V(G) \setminus V_G(H')$ denote the set of vertices of G outside the grid H', and In the set of parts of $V(H')$ lying at distance at least $r(f_2(k))$ from the border of H' in H'. Let B be the bipartite graph with vertex bipartition (Out, In), and where a part $h \in In$ is adjacent to a vertex $x \in Out$ if there exists a vertex $y \in h$ such that xy is an edge of G. Denote by A the set of parts in In with degree at least $f_5(k)$ in B. Since every set of H'-matched edges has size at most $f_5(k)$, we have that $|A| < f_5(k)$. Hence there exists a subgrid $H'' \subseteq In$ of size at least $\frac{f_6(k)}{\sqrt{|f_5(k)|}}$ which contains no part of A by Lemma 15. Let C be the set of parts in H'' with degree at least 1 in B. If $|C| \leq f_5(k)$, Lemma 15 gives a subgrid H''' of H'' of size at least $\frac{f_6(k)}{f_5(k)}$ (which we assume to be greater than $f_4(k)$) which contains no part in C. Parts in H''' have no neighbours in $V(G) \setminus V_G(H')$, hence H''' is a proper subgrid. Otherwise if $|C| > f_5(k)$, let D be the set of vertices in Out with at least one neighbour in H''. We have by hypothesis that $|D| \leqslant f_5(k)^2$ (otherwise we would find a large set of H'-matched edges).

Hence there exists a grid model with a subgrid H'' that respects all properties of the definition of a proper subgrid but the last one. We now show how to deal with the special vertices. Let $H_0 = H''$. For $i \geq 0$, if there exists a vertex in Out adjacent in B to a set P_i of parts in H_i of size less than $k^2/4$, then we let H_{i+1} be the subgrid of H_i obtained by Lemma 15 which contains no vertex in P_i. Otherwise H_i is a proper subgrid and we stop. The above process can be repeated at most $D < f_5(k)^2$ times. Hence the proper subgrid found by this process has size at least $\frac{f_6(k)}{f_5(k)^{f_5(k)}}$ (assumed to be greater $f_4(k)$). $\diamond$

To complete the proof of Theorem 4, it remains to show that we can find a subgrid H_m of sufficiently large size with at most $f_5(k)$ H_m-matched edges. Assume there are at least $f_5(k)$ H'-matched edges. Using Lemma 13, this implies either that G has a large clique minor or that G admits a grid model H with an induced augmented grid H_m having at most $f_5(k)$ H_m-matched edges. Since the former case is impossible by assumption, it follows by Claim 10 that G admits a proper subgrid model. This completes the proof of Theorem 4.

□

5.2 Reducing the instance

Assume now that G has an $(f_6(k) \times f_6(k))$ grid model H but no $f_2(k)$-clique model. By Theorem 4 we can moreover assume that G has a proper subgrid H'. Given a set F of edges, a *giant component* denotes in this case a connected component of $G \setminus F$ containing entirely at least $|V(H)| - k^2/4$ parts. We do not give detailed proofs of our claim when they are similar to claims in the previous section.

Lemma 16 *Let F be a set of at most k edges of a graph G admitting a grid model. Then $G \setminus F$ has a* giant component, *denoted by C_F.*

This is analogous to Lemma 6. The way to cut the most vertices in a grid with k edges is to cut off a corner of size $(k/2 \times k/2)$. We call the other connected components the *small components* (and they intersect altogether at most $k^2/4$ parts).

Lemma 17 *For every set F of at most k edges the special vertices belong to the giant component of $G \setminus F$.*

This is analogous to Lemma 7.

Lemma 18 *Let F be an optimal k-multicut of a graph G admitting a grid model H. Every small component C of $G \setminus F$ contains a vertex which is a terminal.*

This is analogous to Lemma 8.

Let us now consider the refinement of the subgrid H' into vertex-disjoint subgrids $H'_{i,j}$ of size $f_7(k) \cdot f_7(k)$ each, where $f_7(k)$ is large enough.

5.2.1 There exists a subgrid without terminal

Let $1 \leqslant i, j \leqslant f_7(k)$ be such that $H'_{i,j}$ does not contain any terminal, and $h_{i,j}$ be a part of $H'_{i,j}$ at maximum distance in $H'_{i,j}$ of the border of the grid $H'_{i,j}$. We define $x_{i,j}$ to be a vertex of $h_{i,j}$ and $y_{i,j}$ to be any neighbour of $x_{i,j}$ in G.

Claim 11 *For every k-multicut F, $x_{i,j}$ and $y_{i,j}$ both belong to the giant component of $G \setminus F$.*

Proof. Recall that every small component contains a terminal and intersects at most $k^2/4$ parts, by Lemma 16 and Lemma 18. If part $h_{i,j}$ lies at distance more than $k^2/4+1$ in H of a part containing a terminal, then $x_{i,j}$ has no path in G to a terminal intersecting at most $k^2/4+1$ parts, and so $x_{i,j}$ and $y_{i,j}$ must both belong to the giant component of any k-multicut. Assume that there exists a path in G between $x_{i,j}$ and a terminal intersecting at most $k^2/4+1$ parts. This path uses an edge $z_{i,j}w_{i,j}$ with $z_{i,j} \in H'_{i,j}$ for some special vertex $w_{i,j}$. Indeed $H'_{i,j}$ has size greater than $2(k^2/4+2)$, so $h_{i,j}$ is at distance more than $k^2/4+1$ from the border of $H'_{i,j}$, and $H'_{i,j}$ contains no terminal. Since the special vertices belong to the giant component of $G \setminus F$ for every set F of at most k edges by Lemma 17, it follows that $x_{i,j}$ and $y_{i,j}$ belong to the giant component of $G \setminus F$. $\diamond$

Since $x_{i,j}$ and $y_{i,j}$ belong to the giant component of $G \setminus F$ for every k-multicut F, we can safely contract them.

5.2.2 All subgrids contain a terminal

We find a gathered set using arguments similar to the ones used in the previous section. Consider the level decomposition of $G \setminus W$ starting from $L_0 := (H \setminus (H' \cup W)) \cup B(H')$. Since H' is a proper subgrid and hence contains no non-planar edges, if follows that there are at least $\frac{f_4(k)}{2}$ levels in the level decomposition of $G \setminus W$ starting from L_0. Let T be a set of maximum size of terminals in $V_G(L_1 \cup \ldots \cup L_d)$ belonging to different parts lying at distance at least $k^2/4$ in $H_{>1} := V_H(L_1 \cup \ldots \cup L_d)$. Since every subgrid $H'_{i,j}$ in H' contains a terminal and has size $f_7(k) \cdot f_7(k)$ there is a terminal in every $3f_7(k)$ consecutive levels, implying that T has sufficiently large size to apply Rule 3.

We now show that T is a gathered set. Consider any set F of at most k edges. Since parts containing terminals of T are pairwise at distance at least $k^2/4$ in $H_{>1}$ and since small components intersect at most $k^2/4$ parts by Lemma 16, a small component of $G \setminus F$ can contain two terminals of T only if it contains a special vertex $w \in W$. Since w belongs to the giant component by Lemma 17, it follows that $G \setminus F$ contains at most one component (the giant component) containing more than one vertex of T and hence T is a gathered set. Using Rule 3, there exists an irrelevant request adjacent to a terminal in T which can be found in FPT time.

This completes the reduction of this section, showing that we can reduce in FPT time in k a graph with a large grid minor but with no large clique minor.

6 FPT algorithm for Integer Weighted Multicut In Trees

Multicut is known to be FPT on trees [13], and even to have a polynomial kernel [1]. One of the simplest subcases of Multicut left open is that of trees with multi-edges [1, 13]. In this section we give a fixed-parameter algorithm for the integer weighted case of Multicut In Trees, defined as follows. Given a tree $T = (V, E)$, a set of requests P, a weight function $\omega : E \to \mathbb{N}^+$ and an integer k, decide if there is a multicut S such that $\omega(S) \leqslant k$, where the weight of a set of edges is defined to be the sum of the weights of its edges.

In [14], Guo *et. al.* gave an algorithm solving this problem in time $O(3^d.mn^2)$, where the parameter d is the maximum number of requests going through a node or an edge of the given tree. To our knowledge, no fixed-parameter algorithm for such a problem when parameterized by the size of the multicut is known. We provide such an algorithm, using branching and reduction rules, running in time $O^\star((2k+1)^k)$.

Rule 5 *If there exists a request (x, y) of length at most $2k+1$, then branch on every edge e in the path $P_{x,y}$ between x and y in T, adding e to the multicut and letting $k = k - \omega(e)$.*

We exhaustively apply branching Rule 5, answering FALSE if $k < 0$. Note that this branching rule can be applied at most k times since every application strictly decreases the parameter. If the resulting instance has no request and $k \geq 0$ we answer TRUE. If there exists an edge of weight at least $k+1$ we contract it.

Then we use the following reduction rule:

Rule 6 *Let $P = \{u_1, \dots, u_p\}$ be a path of maximum length in T and let $e_i = u_iu_{i+1}$ for all $1 \leqslant i \leqslant k+1$. If there exists i such that $1 \leqslant i \leqslant k+1$ and $\omega(e_i) \geqslant \omega(e_{i+1})$, then we can contract e_i.*

Proof: Observe that P has length at least $2k+2$ since T is reduced under Rule 5. First of all, we show that if there is a request (x, y) with $x = u_i$ for $1 \leqslant i \leqslant k+1$, then $P_{x,y}$ contains $\{u_i, \dots, u_{k+2}\}$. Assume by contradiction that there exists a request (u_i, y) with $1 \leqslant i \leqslant k+1$ that does not go through u_{k+2}. As we assumed that all requests have length at least $2k+1$, we know that the distance between u_i and y in T must be at least $k+1$, that is $|P_{y,u_i}| > k$. But then the path $P_{y,u_i} \cup (P \setminus \{u_1, \dots, u_{i-1}\})$ is a path of length larger than P, which is a contradiction. Thus any request than contains e_i also contains e_{i+1}. As $w(e_i) \geq w(e_{i+1})$, this implies that there always exists an optimal k-multicut not containing e_i, and thus it is safe to contract e_i. □

When none of Rule 5 and Rule 6 apply, if there still exists a request, then the longest path of T has length more than $2k+1$. Denote by $(e_i)_{1 \leq i \leq i+2}$ the first $i+2$ edges of this path. We have $w(e_{i+1}) > w(e_i)$ for $1 \leq i \leq i+1$, hence $w(e_{k+2}) \geq k+1$, and it is safe to contract e_{k+2}. Hence when this branching/reduction process stops after at most $O^\star((2k+1)^k)$ steps.

Conclusion

Deciding whether the Multicut problem parameterized by the solution size k is fixed parameter tractable is one of the most important open question in parameterized complexity [5]. In this paper we have shown that this problem can be reduced in FPT time to graphs of treewidth bounded by a function of k. Hence the main open question that remains is whether it is fixed parameter tractable to decide the Multicut problem in graphs of treewidth bounded by a function of k. As a first step, we proposed an algorithm for the Integer Weighted Multicut In Trees problem, which is a particular subcase of Multicut for graphs of bounded treewidth. We believe that Multicut with this parameterization should be fixed parameter tractable on graphs of treewidth bounded in k. We thus conclude our paper with the following:

Conjecture 1 *The Multicut problem parameterized by the solution size k is fixed-parameter tractable on graphs of treewidth bounded in k.*

Aknowledgements. We thank Gwenaël Joret for pointing out the proof of Lemma 10.

References

[1] N. Bousquet, J. Daligault, S. Thomassé, and A. Yeo. A polynomial kernel for multicut in trees. In *STACS*, volume 09001 of *Dagstuhl Seminar Proceedings*, pages 183–194, 2009.

[2] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Comput. Complex.*, 15(2):94–114, 2006.

[3] J. Chen, Y. Liu, and S. Lu. An improved parameterized algorithm for the minimum node multiway cut problem. *Algorithmica*, 55(1):1–13, 2009.

[4] M.-C. Costa, L. Létocart, and F. Roupin. Minimal multicut and maximal integer multiflow: A survey. *European Journal of Operational Research*, 162(1):55–69, 2005.

[5] E. D. Demaine, M. Hajiaghayi, and D. Marx. 09511 open problems – parameterized complexity and approximation algorithms. In *Parameterized complexity and approximation algorithms*, number 09511 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, 2010.

[6] R. G. Downey and M. R. Fellows. *Parameterized Complexity.* Springer, 1999.

[7] J. Flum and M. Grohe. *Parameterized complexity theory.* Springer-Verlag New York Inc, 2006.

[8] N. Garg, V. V. Vazirani, and M. Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. In *STOC*, pages 698–707, 1993.

[9] N. Garg, V. V. Vazirani, and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, 1997.

[10] G. Gottlob and S. Tien Lee. A logical approach to multicut problems. *Inf. Process. Lett*, 103(4):136–141, 2007.

[11] S. Guillemot. FPT algorithms for path-transversals and cycle-transversals problems in graphs. In *IWPEC*, pages 129–140, 2008.

[12] J. Guo, F. Hüffner, E. Kenar, R. Niedermeier, and J. Uhlmann. Complexity and exact algorithms for multicut. In *SOFSEM*, pages 303–312, 2006.

[13] J. Guo and R. Niedermeier. Fixed-parameter tractability and data reduction for multicut in trees. *Networks*, 46(3):124–135, 2005.

[14] J. Guo and R. Niedermeier. Exact algorithms and applications for tree-like weighted set cover. *J. Discrete Algorithms*, 4(4):608–622, 2006.

[15] D. Marx. Parameterized graph separation problems. *Theor. Comput. Sci.*, 351(3):394–406, 2006.

[16] D. Marx, B. O'Sullivan, and I. Razgon. Treewidth reduction for constrained separation and bipartization problems. In *STACS*, 2010.

[17] D. Marx and I. Razgon. Constant ratio fixed-parameter approximation of the edge multicut problem. In *ESA*, volume 5757 of *Lecture Notes in Computer Science*, pages 647–658. Springer, 2009.

[18] K. Menger. Zur allgemeinen kurventheorie. *Fund. Math.*, 10:96–115, 1927.

[19] R. Niedermeier. *Invitation to Fixed Parameter Algorithms (Oxford Lecture Series in Mathematics and Its Applications).* Oxford University Press, USA, March 2006.

[20] N. Robertson and P.D. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory, Ser. B*, 63(1):65–110, 1995.

[21] N. Robertson, P.D. Seymour, and R. Thomas. Quickly excluding a planar graph. *J. Combin. Theory Ser. B*, 62(2):323–348, 1994.