

Parameterized Complexity Dichotomy for STEINER MULTICUT*

Karl Bringmann[†] Danny Hermelin[‡] Matthias Mnich[§] Erik Jan van Leeuwen[¶]

June 22, 2015

Abstract

We consider the STEINER MULTICUT problem, which asks, given an undirected graph G , a collection $\mathcal{T} = \{T_1, \dots, T_t\}$, $T_i \subseteq V(G)$, of terminal sets of size at most p , and an integer k , whether there is a set S of at most k edges or nodes such that of each set T_i at least one pair of terminals is in different connected components of $G \setminus S$. This problem generalizes several well-studied graph cut problems, in particular the MULTICUT problem, which corresponds to the case $p=2$. The MULTICUT problem was recently shown to be fixed-parameter tractable for the parameter k [Marx and Razgon, Bousquet et al., STOC 2011]. The question whether this result generalizes to STEINER MULTICUT motivates the present work.

We answer the question that motivated this work, and in fact provide a dichotomy of the parameterized complexity of STEINER MULTICUT on general graphs. That is, for any combination of k , t , p , and the treewidth $\text{tw}(G)$ as constant, parameter, or unbounded, and for all versions of the problem (edge deletion and node deletion with and without deletable terminals), we prove either that the problem is fixed-parameter tractable or that the problem is hard (W[1]-hard or even (para-)NP-complete). Among the many results in the paper, we highlight that:

- The edge deletion version of STEINER MULTICUT is fixed-parameter tractable for the parameter $k+t$ on general graphs (but has no polynomial kernel, even on trees). We present two independent proofs of fixed-parameter tractability. The first proof uses the randomized contractions technique of Chitnis et al. The second proof relies on several new structural lemmas, which decompose the Steiner cut into important separators and minimal s - t cuts.
- In contrast, both node deletion versions of STEINER MULTICUT are W[1]-hard for the parameter $k+t$ on general graphs.
- All versions of STEINER MULTICUT are W[1]-hard for the parameter k , even when $p=3$ and the graph is a tree plus one node. This means that the mentioned results of Marx and Razgon, and Bousquet et al. do not generalize to even the most basic instances of STEINER MULTICUT.

Since we allow k , t , p , and $\text{tw}(G)$ to be any constants, our characterization includes a dichotomy for STEINER MULTICUT on trees (for $\text{tw}(G)=1$) as well as a polynomial time versus NP-hardness dichotomy (by restricting $k, t, p, \text{tw}(G)$ to constant or unbounded).

*An extended abstract of this paper appeared in Mayr, E.W., Ollinger, N. (eds.) *32nd International Symposium on Theoretical Aspects of Computer Science (STACS 2015)*, LIPIcs volume 30, Schloss Dagstuhl, 2015, pp. 157–170.

[†]Institute of Theoretical Computer Science, ETH Zurich, Switzerland. karlb@inf.ethz.ch. Supported by the ETH Zurich Postdoctoral Fellowship Program.

[‡]Ben Gurion University of the Negev, Israel. hermelin@bgu.ac.il. Supported by the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement number 631163.11, and by the Israel Science Foundation (grant no. 551145/14).

[§]Universität Bonn, Germany. mmnich@uni-bonn.de. Supported by ERC Starting Grant 306465.

[¶]Max-Planck Institut für Informatik, Saarbrücken, Germany. erikjan@mpi-inf.mpg.de.

1 Introduction

Graph cut problems are among the most fundamental problems in algorithmic research. The classic result in this area is the polynomial-time algorithm for the s - t cut problem of Ford and Fulkerson [31] (independently proven by Elias et al. [27] and Dantzig and Fulkerson [22]). This result inspired a research program to discover the computational complexity of this problem and of more general graph cut problems. One well-studied generalization of the s - t cut problem is the MULTICUT problem, in which we want to disconnect t pairs of nodes instead of just one pair. In a recent major advance of the research program on graph cut problems, Bousquet et al. [10] and Marx and Razgon [49] showed that MULTICUT is fixed-parameter tractable in the size k of the cut only, meaning that it has an algorithm running in time $f(k) \cdot \text{poly}(|V(G)|)$ for some function f , resolving a longstanding problem in parameterized complexity (with many papers [46, 52, 35, 48] building up to this result).

In this paper, we continue the research program on generalized graph cut problems, and consider the STEINER MULTICUT problem. This problem was proposed by Klein et al. [41], and appears in several versions, depending on whether we want to delete edges or nodes, and whether we are allowed to delete terminal nodes. Formally, these versions of the STEINER MULTICUT problem are defined as follows:

{EDGE, NODE, RESTRICTED NODE} STEINER MULTICUT
Input: An undirected graph G with terminal sets $T_1, \dots, T_t \subseteq V(G)$, and an integer $k \in \mathbb{N}$.
Question: Find a set S of k {edges, nodes, non-terminal nodes} such that for $i = 1, \dots, t$ and at least one pair $u, v \in T_i$ there is no u - v path in $G \setminus S$.

Observe that MULTICUT is the special case of STEINER MULTICUT in which each terminal set has size two. In general, the terminal sets of STEINER MULTICUT can have arbitrary size, and we use p to denote $\max_i |T_i|$.

The complexity of STEINER MULTICUT has been investigated extensively, but so far only from the perspective of approximability. This line of work was initiated by Klein et al. [41], who gave an LP-based $O(\log^3(kp))$ -approximation algorithm. The approximability of several variations of the problem has also been considered [54, 33, 3]; in particular, Garg et al. [33] give an $O(\log t)$ -approximation algorithm for MULTICUT. On the hardness side, even MULTICUT is APX-hard [21, 12] and cannot be approximated within any constant factor assuming the Unique Games Conjecture [14]. We also remark that Steiner cuts (the case when $t=1$) are of interest: they are an ingredient in several LP-based approximation algorithms (for example for STEINER FOREST [1, 42]) and there is a connection to the number of edge-disjoint Steiner trees that each connect all terminals [45]. To the best of our knowledge, however, STEINER MULTICUT in its general form has not yet been considered from the perspective of parameterized complexity.

1.1 Our Contribution

In this paper, we fully chart the (parameterized) complexity landscape of STEINER MULTICUT according to k , t , p (defined as above), and the treewidth $\text{tw}(G)$. For all three versions of STEINER MULTICUT, for each possible combination of k , t , p , and $\text{tw}(G)$, where each may be either chosen as a constant, a parameter, or unbounded, we consider the complexity of STEINER MULTICUT. We show a complete dichotomy: either we provide a fixed-parameter algorithm with respect to the chosen parameters, or we prove a W[1]-hardness or (para-)NP-completeness result that rules out a fixed-parameter algorithm (unless many canonical NP-complete problems have subexponential- or polynomial-time algorithms respectively).

constants	params	EDGE STEINER MC	NODE STEINER MC	RESTR. NODE STEINER MC
k	—	poly (Sect. 2)	poly (Sect. 2)	poly (Sect. 2)
$t \leq 2$	—	poly (Sect. 2)		poly (Sect. 2)
$t=3, p=2$	—	NP-h [21]		NP-h [21]
—	k, t	*FPT (Thm. 1.1)	*W[1]-h (Thm. 1.2)	*W[1]-h (Thm. 1.2)
—	k, p, t		FPT (Sect. 2)	FPT (Sect. 2)
t	k			FPT (Sect. 2)
$p=2$	k	FPT [10, 49]	FPT [10, 49]	FPT [10, 49]
$p=3, \text{tw}=2$	k	*W[1]-h (Thm. 1.3)	*W[1]-h (Thm. 1.3)	*W[1]-h (Thm. 1.3)
—	t, tw	*FPT (Thm. 5.2)	*FPT (Thm. 5.2)	*FPT (Thm. 5.2)
$\text{tw}=1$	—		*poly (Thm. 7.1)	
$\text{tw}=1$	k	*W[2]-h (Thm. 7.2)		*W[2]-h (Thm. 7.2)
$\text{tw}=1$	k, p	*FPT (Thm. 7.4)		*FPT (Thm. 7.4)
$\text{tw}=1, p=2$	—	NP-h [12]		NP-h [12]
$\text{tw}=2, p=2$	—		NP-h [12]	

Table 1: Summary of known and new complexity results for STEINER MULTICUT, where new results are marked with *; the other entries are either known or follow easily from known results in the literature. Only maximal FPT results and minimal W[·]- or NP-hardness results are listed; empty cells are dominated by other results. E.g. EDGE STEINER MULTICUT with parameter t is hard, since it is already NP-hard for $t=3, p=2$. For NODE STEINER MULTICUT, one also has to apply the rule that $k < t$ (see Section 2) to generate a full characterization of all cases. Tree diagrams of this table are offered in Appendix B.

The dichotomy is composed of three main results, along with many smaller ones (see Table 1). These three main results are stated in the three theorems below:

Theorem 1.1. EDGE STEINER MULTICUT *is fixed-parameter tractable for the parameter $k+t$.*

Theorem 1.2. NODE STEINER MULTICUT *and* RESTR. NODE STEINER MULTICUT *are W[1]-hard for the parameter $k+t$.*

Theorem 1.3. NODE STEINER MULTICUT, EDGE STEINER MULTICUT, *and* RESTR. NODE STEINER MULTICUT *are W[1]-hard for the parameter k , even if $p=3$ and $\text{tw}(G)=2$.*

Observe the sharp gap described by Theorem 1.1 and Theorem 1.2 between the parameterized complexity of the edge deletion version versus the node deletion version; this gap does not exist for the MULTICUT problem. We also note that Theorem 1.3 implies that the fixed-parameter algorithms for MULTICUT for parameter k [10, 49] do not generalize to STEINER MULTICUT.

To obtain the fixed-parameter algorithm of Theorem 1.1, we have to avoid the brute-force choice of a pair of separated terminals of each terminal set: Although one can trivially reduce every instance of the EDGE STEINER MULTICUT problem to at most $\binom{p}{2}^t$ instances of MULTICUT parameterized by k , this only yields an $f(k) \cdot n^{O(t)}$ -time algorithm (for unbounded p). Our contribution in Theorem 1.1 is that we improve on this simple algorithm and obtain a runtime of $f(k, t) \cdot n^{O(1)}$. We give two independent proofs of Theorem 1.1:

- Our first proof uses the recent technique of Chitnis et al. [17] known as randomized contractions (even though the technique actually yields deterministic algorithms). The rough idea of the algorithm is to first determine a large subgraph G' of the input graph G , such that G' has no small cut and only has a small interface (i.e. a small number of vertices that connect the subgraph to the rest of the graph). We can then branch on the behavior of a solution on the interface to determine a set $U \subseteq E(G')$ of ‘useless’ edges, in the sense that when U is contracted in G a smallest solution (of size at most k) persists in the remaining

graph. By iterating this procedure, we can reduce the size of the graph until it is small enough to be handled by exhaustive enumeration. Our algorithm is similar to the one for EDGE MULTIWAY CUT-UNCUT in the paper by Chitnis et al. [17]; however, in contrast to that problem, there seems to be no straightforward projection of the instance onto G' in our case, and therefore more involved arguments are needed to determine the set U .

- Our second proof is based on several novel structural lemmas that show that a minimal edge Steiner cut can be decomposed into important separators and minimal s - t cuts. Using a branching strategy, we ascertain the topology of the decomposition that is promised by the structural lemmas. Since there are only few important separators of bounded size [46, 16, 49] and all relevant minimal s - t cuts lie in a graph of bounded treewidth (following the treewidth reduction techniques of Marx et al. [48]), we can then optimize over important separators and minimal s - t cuts.

The advantage of the first algorithm over the second is that it runs in single-exponential time, instead of double-exponential time. However, the second algorithm is slightly faster in terms of n . Moreover, as part of the correctness proof of the second algorithm, we present some structural lemmas that give additional insight into the properties of the cut, which may be of independent interest. Therefore, we present both algorithms.

The $W[1]$ -hardness results of Theorem 1.2 and 1.3 all rely on reductions from the MULTICOLORED CLIQUE problem [28]. For the proof of Theorem 1.3, we introduce a novel intermediate problem, NAE-INTEGGER-3-SAT, which is an integer variant of the better known NOT-ALL-EQUAL-3-SAT problem. We show that NAE-INTEGGER-3-SAT is $W[1]$ -hard parameterized by the number of variables. This is a powerful starting point for parameterized hardness reductions and should turn out to be useful to prove the hardness of other problems.

To complete our dichotomy, the second part of our paper charts the full (parameterized) complexity of STEINER MULTICUT on trees, that is, for graphs G with $\text{tw}(G)=1$. In fact, some of the hardness results that we prove for STEINER MULTICUT on general graphs even hold for trees. We also show that many of the results for trees do not carry over to graphs of bounded treewidth, the only exception being a fixed-parameter algorithm for parameters $\text{tw}(G)+t$.

We remark that our characterization induces a polynomial time vs. NP-hardness dichotomy for STEINER MULTICUT, i.e., for any choice of $k, p, t, \text{tw}(G)$ as any constants or unbounded (and all three problem variants), we either prove that STEINER MULTICUT is in P or that it is NP-hard. This characterization can be obtained from Table 1 by considering all its polynomial time and NP-hardness results as well as using the rule that any fixed-parameter algorithm induces a polynomial-time algorithm by setting all parameters to $O(1)$.

1.2 Related Work

We briefly cite the most relevant results on the parameterized complexity of graph cut problems. We already mentioned several results on the special case of STEINER MULTICUT when $p=2$ (MULTICUT) [46, 52, 35, 10, 49, 48]. MULTICUT is itself a generalization of MULTIWAY CUT, also known as MULTITERMINAL CUT, where the goal is to delete k edges or nodes to separate all terminals from each other. This problem is NP-complete even for three terminals [21] and has been extensively studied from a parameterized point of view (see, e.g., the work of Cao et al. [13] or Cygan et al. [20]). The parameterized complexity of many different other graph cut problems has also been considered in recent years [24, 17, 40, 46]. On trees, we only mention here that EDGE MULTICUT and RESTR. NODE MULTICUT remain NP-hard [12], but are fixed-parameter tractable [37, 36]. In contrast, NODE MULTICUT has a polynomial-time algorithm on trees [12].

1.3 Organization

We begin our exposition in Section 2 by giving easy results for certain parameter combinations of STEINER MULTICUT. Thereafter, we present our fixed-parameter algorithms for EDGE STEINER MULTICUT (Theorem 1.1) in Section 3 and 4. Following this, in the two subsequent sections, we present our $W[1]$ -hardness proofs: the proof of Theorem 1.3 in Section 5, and of Theorem 1.2 in Section 6. Section 7 then focuses on trees to complete our dichotomy. We conclude with some discussion and open problems in Section 8. For basic notions of parameterized complexity as well as the notion of treewidth we refer the reader to Appendix A.

2 Easy and Known Results

In this section, we collect easy and known results about the STEINER MULTICUT problem. Some of these results are scattered throughout the literature, while others are new. First, observe that whenever the cut size k is constant, we can solve the problem in polynomial time by simply guessing the desired set S of at most k edges or nodes.

Furthermore, NODE STEINER MULTICUT is trivially solvable when $t \leq k$, as in this case we may simply delete an arbitrary terminal node from each set T_i , resulting in a solution of size at most k ; thus, any instance is always a “yes”-instance in this case.

We may reduce STEINER MULTICUT to $\binom{p}{2}^t$ instances of MULTICUT by branching for each terminal set over its separated terminals. Since MULTICUT is in FPT for parameter k , we obtain a fixed-parameter algorithm for STEINER MULTICUT for parameter $k+t+p$. Also, since $\binom{p}{2}^t \leq n^{O(t)}$, STEINER MULTICUT is in FPT for parameter k and any constant t .

Now, MULTICUT on instances with $t = 1$ (i.e. instances that have only one terminal pair $|T_1| = \{s, t\}$) is polynomial-time solvable by running an $s-t$ cut algorithm. For $t=2$ a result by Yannakakis et al. [53, Lemma 1] also yields a polynomial time algorithm for MULTICUT. Again by branching over the separated terminals in both terminal sets, we obtain a polynomial time algorithm for STEINER MULTICUT for $t \leq 2$.

When there are three or more terminal sets, then STEINER MULTICUT generalizes MULTIWAY CUT and thus is NP-complete [21] even when $p=2$.

We next show that RESTR. NODE STEINER MULTICUT is as least as hard as NODE STEINER MULTICUT. Therefore, whenever NODE STEINER MULTICUT is $W[1]$ -hard (or NP-hard) for a certain combination of parameters, then so is RESTR. NODE STEINER MULTICUT.

Lemma 2.1. *Any instance of NODE STEINER MULTICUT can be reduced in polynomial time to an instance of RESTR. NODE STEINER MULTICUT with the same parameter values k, t, p , and $\text{tw}(G)$.*

Proof. Take an instance (G, \mathcal{T}, k) , $\mathcal{T} = \{T_1, \dots, T_t\}$, of NODE STEINER MULTICUT and transform it to an instance of RESTR. NODE STEINER MULTICUT by adding for each terminal node $v \in T_1 \cup \dots \cup T_t$ a new pendant node v' . Then replace v by v' in every terminal set T_i . It is easy to see that the original instance admits a node cut of size k if and only if the new instance admits a node cut of size k that does not use any terminal nodes. \square

3 Tractability for Edge Deletion and Parameter $k+t$: First Proof

In this section, we prove Theorem 1.1, namely that EDGE STEINER MULTICUT parameterized by $k+t$ is fixed-parameter tractable. Here, we give the first proof, which uses the technique of randomized contractions pioneered by Chitnis et al. [17]; the second proof is in the next section.

Later we will see that this result is “maximal”, in the sense that EDGE STEINER MULTICUT is W[1]-hard parameterized by k or t alone (this follows from Theorem 7.2 and the fact that even EDGE MULTICUT is NP-hard when $t = 3$ [21] respectively), that the corresponding node deletion problem is W[1]-hard parameterized by $k + t$ (Theorem 6.1), and that there exists no polynomial kernel for EDGE STEINER MULTICUT parameterized by $k + t$ (Theorem 7.3).

We first state some notions and supporting lemmas from the paper of Chitnis et al. [17], which are needed to make our proof work. The *identification* of two vertices $v, w \in V(G)$ results in a graph G' by removing v, w , adding a new vertex vw , and if v or w is an endpoint of an edge, then we replace this endpoint by vw . Note that the identification of two vertices does not remove any edges, and generally results in a multigraph (with parallel edges and self-loops). Without confusion, we may sometimes refer to vw by its old names v or w .

The *contraction* of an edge $(v, w) \in E(G)$ results in a graph G' by removing all edges between v and w , and then identifying v and w . This is also known as contraction without removing parallel edges. Again, the result of a contraction is generally a multigraph. Given a set $F \subseteq E(G)$ of edges that induce a connected subgraph of G with $a + 1$ vertices of which v is one, after contracting all edges of F , we say that a vertices were contracted *onto* v .

Lemma 3.1 ([17]). *Given a universe \mathbb{U} and integers a', b' with $0 \leq a', b' \leq |\mathbb{U}|$, one can in time $2^{O(\min\{a', b'\} \log(a' + b'))} |\mathbb{U}| \log |\mathbb{U}|$ find a family \mathcal{F} of $2^{O(\min\{a', b'\} \log(a' + b'))} \log |\mathbb{U}|$ subsets of \mathbb{U} such that for any two disjoint sets $A, B \subseteq \mathbb{U}$ of size at most a' and b' respectively, there exists a set $S \in \mathcal{F}$ that contains all elements of A but is disjoint from B .*

Definition 3.2 ([17]). *Given two integers a, b , an (a, b) -good edge separation of a connected graph G is a partition (V_1, V_2) of $V(G)$ such that $|V_1|, |V_2| > a$, $G[V_1]$ and $G[V_2]$ are connected, and the number of edges between V_1 and V_2 is at most b .*

Lemma 3.3 ([17]). *Given a connected graph G and two integers a, b , one can decide in time $2^{O(\min\{a, b\} \log(a + b))} |V(G)|^3 \log |V(G)|$ whether G has an (a, b) -good edge separation, and if it does, find such a separation in the same time.*

We now define several notions and prove a few lemmas that are implicit in the work of Chitnis et al. [17]. We provide full proofs only for sake of completeness.

Definition 3.4. *A b -bordered subgraph of G is a connected induced subgraph G' of G such that in G at most b vertices of $V(G')$ have an edge to a vertex of $V(G) \setminus V(G')$. We call the vertices of G' that have an edge in G to a vertex of $V(G) \setminus V(G')$ the border vertices of G' .*

Lemma 3.5. *Given a connected graph G and two integers a, b (b even), one can find in time $2^{O(\min\{a, b\} \log(a + b))} |V(G)|^4 \log |V(G)|$ a b -bordered subgraph of G that does not admit an $(a, b/2)$ -good edge separation.*

Proof. Initially, let $G' = G$. We apply an iterative procedure that maintains the invariant that G' is a b -bordered subgraph of G . Note that the invariant holds for $G' = G$. Now run the algorithm of Lemma 3.3 on G' to decide the existence of an $(a, b/2)$ -good edge separation in G' . If no such separation exists, then we simply return G' . Otherwise, let (V_1, V_2) be the separation returned by the algorithm. Assume without loss of generality that V_1 contains less border vertices of G' than V_2 . Since G' is b -bordered, V_1 contains at most $b/2$ border vertices of G' . Moreover, since (V_1, V_2) is an $(a, b/2)$ -good edge separation in G' , there are at most $b/2$ edges between V_1 and V_2 , and thus at most $b/2$ vertices of V_1 have an edge to a vertex of V_2 in G' . Since G' is an induced subgraph of G , this implies at most $b/2$ vertices of V_1 have an edge to a vertex of V_2 in G . Hence, at most b vertices of V_1 have an edge to a vertex of $V(G) \setminus V_1$. Hence, $G[V_1]$ is a b -bordered subgraph of G . Set $G' = G[V_1]$ and iterate. Since V_1 and V_2 are both nonempty, this procedure terminates after $O(|V(G)|)$ steps. Combined with the running time of the algorithm of Lemma 3.3, this implies the lemma. \square

Let G be a connected graph and let a be an integer. Given a set $F \subseteq E(G)$, let G_F denote the graph obtained from G by contracting all edges of F , and then identifying into a single vertex (which we denote by h_F) all vertices onto which at least a vertices were contracted. Observe that G_F is potentially a multigraph, and that h_F might not exist.

A subset Y of the edges of a connected graph G is a *separator* if $G \setminus Y$ has more than one connected component. The set Y is a *minimal separator* if there is no $Y' \subset Y$ such that $G \setminus Y'$ has the same connected components as $G \setminus Y$.

Lemma 3.6 ([17]). *Let G be a connected graph, let a, b be two integers (b even), and let $F \subseteq E(G)$ with $|F| \leq b/2$. If G admits no $(a, b/2)$ -good edge separation, then $G \setminus F$ has at most $(b/2) + 1$ connected components, of which at most one has more than a vertices.*

Lemma 3.7. *Let G be a connected graph, let a, b be any two integers (b even) such that G does not admit an $(a, b/2)$ -good edge separation and such that $|V(G)| > a(b/2 + 1)$, and let $Y \subseteq E(G)$ with $|Y| \leq b/2$ be a minimal separator. In time $2^{O(b \log(a+b))} |E(G)| \log |E(G)|$ one can find a family \mathcal{F} of $2^{O(b \log(a+b))} \log |E(G)|$ subsets of $E(G)$ that contains a set $F_0 \subseteq E(G)$ with the following properties: (1) $F_0 \cap Y = \emptyset$, (2) h_{F_0} exists in G_{F_0} , (3) h_{F_0} is the identification of a subset of the vertices of a connected component of $G \setminus Y$, and (4) for each connected component C of $G_{F_0} \setminus \{h_{F_0}\}$, Y either contains all edges of $G_{F_0}[C \cup \{h_{F_0}\}]$ or none of these edges.*

Proof. The idea of the proof is to show that Y induces a set $F_0 \subseteq E(G)$ as in the lemma statement, which one can discover among the sets in the family \mathcal{F} returned by Lemma 3.1 for an appropriate choice of \mathbb{U} , a' , and b' .

We first describe an essential property of F_0 . Let C_0, \dots, C_ℓ be the connected components of $G \setminus Y$. By Lemma 3.6 and since Y is a separator, $1 \leq \ell \leq b/2$. Following the same lemma and the assumption that $|V(G)| > a(b/2 + 1)$, there is exactly one component of more than a vertices; without loss of generality, this is C_0 . For every connected component C_i , let T_i denote an arbitrary spanning tree of it. For each vertex v of C_0 incident to an edge of Y , let T_0^v be a subtree of T_0 that contains v and that has $a+1$ vertices; note that such a subtree exists, as T_0 has more than a vertices. Let T_0' denote the union of all of these subtrees; note that T_0' is a forest in general. Then let $A = \bigcup_{i=1}^{\ell} E(T_i) \cup E(T_0')$ and let $B = Y$. Let F_0 be an arbitrary subset of $E(G)$ that contains A but is disjoint from B .

We claim that F_0 satisfies all the properties of the lemma statement. Clearly, $F_0 \cap Y = F_0 \cap B = \emptyset$, and thus the first property holds.

Let H denote the graph obtained from G by contracting all edges of F_0 . Since $E(T_0') \subseteq A \subseteq F_0$ and T_0' contains at least one tree that is a subtree of the spanning tree T_0 and that has $a+1$ vertices and (thus) a edges, there is at least one vertex onto which at least a vertices were contracted. Hence, h_{F_0} exists, and thus the second property holds.

Observe that F_0 does not contain any edges between C_i and $G \setminus C_i$, as each such edge belongs to Y and $F_0 \cap Y = F_0 \cap B = \emptyset$. Therefore, since $\bigcup_{i=1}^{\ell} E(T_i) \subseteq F_0$, for $i = 1, \dots, \ell$, C_i gets contracted onto a single vertex of H , which we denote by c_i . Since C_i has at most a vertices, at most $a-1$ vertices are contracted onto c_i . Hence, c_1, \dots, c_ℓ exist in G_{F_0} ; in other words, they are not identified with h_{F_0} . It follows that h_{F_0} is the identification of a subset of the vertices of C_0 , a connected component of $G \setminus Y$, and thus the third property holds.

From the above observation, it follows that all edges incident to c_i in G_{F_0} for $i = 1, \dots, \ell$ belong to Y , and since Y is minimal, all edges of Y are incident to a c_i in G_{F_0} for $i = 1, \dots, \ell$. It remains to show that all edges of G_{F_0} incident to exactly one c_i for $i = 1, \dots, \ell$ have h_{F_0} as their other endpoint. Let v be any vertex of C_0 incident to an edge e of Y . By construction, v is contained in a subtree of T_0' with at least $a+1$ vertices. Since $E(T_0') \subseteq F_0$, v is contained in a connected component of $G[F_0]$ with at least $a+1$ vertices. Therefore, to obtain H , v is contracted onto a vertex w onto which at least a vertices are contracted. To obtain G_{F_0} , w is identified with

zero or more other vertices to form h_{F_0} . Since $e \notin F_0$, e is incident on h_{F_0} . Therefore, for each connected component C of $G_{F_0} \setminus \{h_{F_0}\}$, Y either contains all edges of $G_{F_0}[C \cup \{h_{F_0}\}]$ or none of these edges, and thus the fourth property holds.

Observe that since $|Y| \leq b/2$, T'_0 is built from at most $b/2$ subtrees of T_0 , each of which has $a+1$ vertices and a edges. Hence, $|A| \leq (a-1)\ell + a(b/2) \leq (2a-1)(b/2)$. Since $|B| = |Y| \leq b/2$, \mathcal{F} contains a set $F_0 \subseteq E(G)$ that contains A but is disjoint from B , where \mathcal{F} is the family returned by Lemma 3.1 for $\mathbb{U} = E(G)$, $a' = (2a-1)(b/2)$, and $b' = b/2$. The lemma follows. \square

It is important to observe that the construction of the family \mathcal{F} does not require knowledge of Y itself, beyond that it has size at most $b/2$. Moreover, note that all edges of Y are present in G_{F_0} , as $F_0 \cap Y = \emptyset$.

We are now ready to describe the algorithm for EDGE STEINER MULTICUT for the parameter $k+t$. The basic intuition is to find a part of the graph that does not have a (q,k) -good edge separation for some q , but that only has a small number of border vertices. In this part of the graph we find and contract a set of edges that are provably not part of some smallest edge Steiner multicut. We repeat this procedure until the graph is small enough to be handled by an exhaustive enumeration algorithm.

Consider an instance (G, \mathcal{T}, k) with $\mathcal{T} = \{T_1, \dots, T_t\}$ of EDGE STEINER MULTICUT. We may assume that G is connected. Let q be an integer determined later (q will depend on k and t only). We assume that $|E(G)| > q$, or we can use exhaustive enumeration to solve the problem in $tq^{O(k)}$ time.

We apply the algorithm of Lemma 3.5 to find a $2k$ -bordered subgraph G' of G that does not admit a (q,k) -good edge separation. Let B denote the set of border vertices of G' . Note that possibly $G' = G$, in which case $B = \emptyset$. The idea is now to determine a set of edges of G' that is not used by some optimal solution.

Let S be a smallest edge Steiner multicut of (G, \mathcal{T}) . Let G'' denote the graph $(B \cup (V(G) \setminus V(G')), E(G) \setminus E(G'))$. Observe that $E(G')$ and $E(G'')$ partition $E(G)$. Let $S' = S \cap E(G')$ and let $S'' = S \cap E(G'')$. We call a terminal set *active* if it is not separated in $G \setminus S''$. We call border vertices $b, b' \in B$ *paired* if there is a path between b and b' in $G'' \setminus S''$. Note that this defines an equivalence relation on B . We call an equivalence class B' of this relation *i -active* if the terminal set T_i is active and the component of $G'' \setminus S''$ that contains B' contains a terminal of T_i . Intuitively, this information suffices to compute a set $Z \subseteq E(G')$ such that $Z \cup S''$ is a smallest edge Steiner multicut of (G, \mathcal{T}) . Then we could contract (in G) all other edges of G' to get a smaller instance, and repeat this until the instance is small enough to be solved by exhaustive enumeration.

Of course, we do not know S , and thus we do not know this equivalence relation on B nor which classes are i -active for each $i = 1, \dots, t$. However, we can branch over all possibilities. In each branch, we find a small set of edges, which we mark. At the end, we contract (in G) the set of edges of G' that were not marked, and thus reduce the size of the instance. We will prove that in one of the branches, we mark a smallest set Z of edges (of size at most k) such that $(Z \cup S'')$ is a smallest edge Steiner multicut (of size at most k) of (G, \mathcal{T}) . Therefore, after contraction, a smallest edge Steiner multicut (of size at most k) of (G, \mathcal{T}) persists (if such a cut existed in the first place). In particular, we argue that we mark Z in the branch with the active classes \mathcal{T}^S , the equivalence relation \mathcal{B}^S , and i -active classes \mathcal{B}_i^S of \mathcal{B}^S for $i = 1, \dots, |\mathcal{T}^S|$ that are induced by S .

The algorithm now branches over all possibilities. Let $\mathcal{T}' = \{T'_1, \dots, T'_{t'}\}$ be an arbitrary subset of \mathcal{T} , let \mathcal{B} be an arbitrary equivalence relation on B , and let \mathcal{B}_i denote an arbitrary subset of \mathcal{B} for $i = 1, \dots, t'$. We say that we made the *right choice* if $\mathcal{T}' = \mathcal{T}^S$, $\mathcal{B} = \mathcal{B}^S$, and $\mathcal{B}_i = \mathcal{B}_i^S$ for $i = 1, \dots, t'$. The algorithm considers two cases.

Case 1: If $|V(G')| \leq q(k+1)$, then we can essentially use exhaustive enumeration. Let \tilde{G} be the graph obtained from G' by identifying two border vertices if they are in the same equivalence class of \mathcal{B} . This also makes each \mathcal{B}_i a set of vertices, which by abuse of notation, we denote by

\mathcal{B}_i as well. For $i = 1, \dots, t'$, let \tilde{T}_i be equal to $\mathcal{B}_i \cup (T_i' \cap (V(G') \setminus B))$. Then $\tilde{\mathcal{T}} = \{\tilde{T}_1, \dots, \tilde{T}_{t'}\}$. We verify that no terminal set in $\tilde{\mathcal{T}}$ is a singleton; otherwise, we can continue with the next branch.

Lemma 3.8. *Assume we made the right choice. Then S' is an edge Steiner multicut of $(\tilde{G}, \tilde{\mathcal{T}})$. Moreover, for any edge Steiner multicut X of $(\tilde{G}, \tilde{\mathcal{T}})$, $X \cup S''$ is an edge Steiner multicut of (G, \mathcal{T}) .*

Proof. Let v, w be any two vertices of \tilde{G} (possibly $v = w$) that are in the same connected component of $\tilde{G} \setminus S'$, and let P be an arbitrary v, w -path in $\tilde{G} \setminus S'$. Observe that each border vertex of \tilde{G} corresponds to a set of border vertices in G' , which are in the same connected component of $G'' \setminus S''$, since we made the right choice. Hence, P can be ‘expanded’ into a v, w -walk of $G \setminus S$, and thus v, w are in the same connected component of $G \setminus S$.

Suppose that \tilde{T}_i is not separated in $\tilde{G} \setminus S'$. Consider any two terminals $s, s' \in T_i'$. We show that s and s' are in the same connected component of $G \setminus S$. Then there are several cases:

- if $s, s' \in V(G'')$, then let b and b' be border vertices reachable in $G'' \setminus S''$ from s and s' respectively (note that possibly $b = b'$). By construction, $b, b' \in \tilde{T}_i$, and since \tilde{T}_i is not separated in $\tilde{G} \setminus S'$, there is a path between b and b' in $\tilde{G} \setminus S'$. By the above observation, there is a path between b and b' in $G \setminus S$. Since b and b' are reachable in $G'' \setminus S''$ from s and s' respectively, it follows that s and s' are in the same connected component of $G \setminus S$.
- if $s, s' \in V(G')$, then $s, s' \in \tilde{T}_i$. Hence, s, s' are in the same connected component of $\tilde{G} \setminus S'$, and by the above observation, in the same connected component of $G \setminus S$.
- if $s \in V(G')$ and $s' \in V(G'')$, then by combining the reasoning of the above two cases, we can again show that s and s' are in the same connected component of $G \setminus S$.

It follows that all terminals of T_i' are in the same connected component of $G \setminus S$, contradicting that S is an edge Steiner multicut of (G, \mathcal{T}) . Therefore, S' is an edge Steiner multicut of $(\tilde{G}, \tilde{\mathcal{T}})$.

Let X be any edge Steiner multicut of $(\tilde{G}, \tilde{\mathcal{T}})$. Suppose that T_i' is not separated in $G \setminus (X \cup S'')$. Consider any two terminals $s, s' \in \tilde{T}_i$. Define $z \in V(G)$ such that $z = s$ if s is not a border vertex, or such that z is a terminal of T_i' in the connected component of $G'' \setminus S''$ that contains s otherwise. Define z' similarly. Since we made the right choice and by construction, z and z' are properly defined. Since T_i' is not separated in $G \setminus (X \cup S'')$, there is a path P in $G \setminus (X \cup S'')$ between z and z' . If s (respectively s') is a border vertex, then P contains s (respectively s'). Hence, there is a path P' between s and s' in $G \setminus (X \cup S'')$. Since we made the right choice, any subpath of P' that lies in $G'' \setminus S''$ and goes between two border vertices, can be replaced by the vertex into which these two border vertices were identified. Hence, P' can be ‘compressed’ into an s, s' -walk of $G \setminus X$, and thus s, s' are in the same connected component of $G \setminus X$. It follows that all terminals of \tilde{T}_i are in the same connected component of $\tilde{G} \setminus X$, contradicting that X is an edge Steiner multicut of $(\tilde{G}, \tilde{\mathcal{T}})$. Therefore, $X \cup S''$ is an edge Steiner multicut of (G, \mathcal{T}) . \square

Now the algorithm uses exhaustive enumeration to find a smallest edge Steiner multicut of $(\tilde{G}, \tilde{\mathcal{T}})$ of size at most k (if one exists) in $t(qk)^{O(k)}$ time. Mark this set of edges in G' .

By Lemma 3.8, if we made the right choice, we mark a set Z such that $Z \cup S''$ is a smallest edge Steiner multicut of (G, \mathcal{T}) .

Case 2: If $|V(G')| > q(k+1)$, then a more complicated approach is needed, because we cannot just use exhaustive enumeration. In fact, we cannot work with $(\tilde{G}, \tilde{\mathcal{T}})$ directly here, as \tilde{G} might have a (q, k) -good edge separation, even though G' does not.

We proceed as follows. Apply Lemma 3.7 with $a = q$ and $b = 2k$ to G' with respect to $Y := S'$, and let \mathcal{F} be the resulting family. Note that $Y = S'$ is a minimal separator, $|V(G')| > q(k+1) = a(b/2+1)$, and G' has no $(a, b/2)$ -good edge separation, and thus the lemma indeed applies. Consider an arbitrary $F \in \mathcal{F}$. We augment our definition of the *right choice* by

adding the condition that $F = F_0$, where F_0 is the family that Lemma 3.7 promises exists in \mathcal{F} . Now find G'_F . If h_F does not exist in G'_F , then we proceed to the next set F , as Lemma 3.7 promises that h_F exists if we made the right choice.

We call a set $X \subseteq E(G'_F)$ an *all-or-nothing cut* if for each connected component C' of $G'_F \setminus \{h_F\}$, X either contains all edges of $G'_F[C' \cup \{h_F\}]$ or none of these edges. Note that Lemma 3.7 promises that S' is an all-or-nothing cut in G'_F for $F = F_0 \in \mathcal{F}$.

Let \mathcal{C} denote the set of connected components of $G'_F \setminus \{h_F\}$ that contain a vertex onto which a border vertex was contracted. Let $Y \subseteq \bigcup_{C \in \mathcal{C}} E(G'_F[C \cup \{h_F\}])$ be an arbitrary set of edges that contains for each $C \in \mathcal{C}$ either all edges of $G'_F[C \cup \{h_F\}]$ or none of these edges. Note that Y is basically an all-or-nothing cut restricted to the edges induced by \mathcal{C} . The algorithm will consider all possible choices of Y . We augment our definition of the *right choice* again, by adding the condition that $Y = Y_0$, where $Y_0 := S' \cap (\bigcup_{C \in \mathcal{C}} E(G'_F[C \cup \{h_F\}]))$.

Now the algorithm deletes every edge in Y and contracts every edge in $(\bigcup_{C \in \mathcal{C}} E(G'_F[C \cup \{h_F\}])) \setminus Y$. Denote the resulting graph by H_F . Let \hat{G} be the graph obtained from H_F by identifying two border vertices if they are in the same equivalence class of \mathcal{B} . This also compresses each \mathcal{B}_i into a set of vertices, which by abuse of notation, we denote by \mathcal{B}_i as well. For $i = 1, \dots, t'$, let \hat{T}_i be equal to $\mathcal{B}_i \cup (T'_i \cap (V(\hat{G}) \setminus B))$. Then $\hat{\mathcal{T}}$ consists of all \hat{T}_i that are not already separated in H_F . We verify that no terminal set in $\hat{\mathcal{T}}$ is a singleton; otherwise, we can continue with the next branch.

Lemma 3.9. *Assume that we made the right choice. Then $S' \setminus Y$ is an edge Steiner multicut of $(\hat{G}, \hat{\mathcal{T}})$ that is an all-or-nothing cut. Moreover, for any edge Steiner multicut X of $(\hat{G}, \hat{\mathcal{T}})$, $Y \cup X \cup S''$ is an edge Steiner multicut of (G, \mathcal{T}) .*

Proof. Since Y is an all-or-nothing cut of G'_F and we made the right choice, it is immediate from Lemma 3.7 that $S' \setminus Y$ is an all-or-nothing cut of G'_F as well. Since $(\bigcup_{C \in \mathcal{C}} E(G'_F[C \cup \{h_F\}])) \setminus Y$ is also an all-or-nothing cut of G'_F , which is disjoint from S' (as we made the right choice), $S' \setminus Y$ is an all-or-nothing cut of H_F . Notice that border vertices of H_F are either isolated or equal to h_F . By the construction of \hat{G} , $S' \setminus Y$ is an all-or-nothing cut of \hat{G} .

The remainder of the proof is similar to Lemma 3.8, but more involved due to the more complex construction of \hat{G} .

Let v, w be any two vertices of \hat{G} (possibly $v = w$) that are in the same connected component of $\hat{G} \setminus (S' \setminus Y)$, and let P be an arbitrary v, w -path in $\hat{G} \setminus (S' \setminus Y)$. Observe that each border vertex of \hat{G} corresponds to a set of border vertices in H_F , which are in the same connected component of $G'' \setminus S''$, since we made the right choice. Hence, P can be ‘expanded’ into a v, w -walk W of $(G'' \setminus S'') \cup (H_F \setminus (S' \setminus Y))$. By the construction of H_F , W can be ‘expanded’ into a v, w -walk W' of $(G'' \setminus S'') \cup (G'_F \setminus S')$. Because we made the right choice, by Lemma 3.7 all vertices identified into h_F come from the same connected component of $G' \setminus S'$ and $F \cap S' = \emptyset$, and thus W' can be ‘expanded’ into a v, w -walk of $(G'' \setminus S'') \cup (G' \setminus S')$. Hence, v, w are in the same connected component of $G \setminus S$.

Suppose that \hat{T}_i is not separated in $\hat{G} \setminus (S' \setminus Y)$. Consider any two terminals $s, s' \in T'_i$. We show that s and s' are in the same connected component of $G \setminus S$. There are several cases:

- if $s, s' \in V(G'')$, then let b and b' be border vertices reachable in $G'' \setminus S''$ from s and s' respectively (note that possibly $b = b'$). By construction, $b, b' \in \hat{T}_i$, and since \hat{T}_i is not separated in $\hat{G} \setminus (S' \setminus Y)$, there is a path between b and b' in $\hat{G} \setminus (S' \setminus Y)$. By the above observation, there is a path between b and b' in $G \setminus S$. Since b and b' are reachable in $G'' \setminus S''$ from s and s' respectively, it follows that s and s' are in the same connected component of $G \setminus S$.
- if $s, s' \in V(G')$, then $s, s' \in \hat{T}_i$. Hence, s, s' are in the same connected component of $\hat{G} \setminus (S' \setminus Y)$, and by the above observation, in the same connected component of $G \setminus S$.

- if $s \in V(G')$ and $s' \in V(G'')$, then by combining the reasoning of the above two cases, we can again show that s and s' are in the same connected component of $G \setminus S$.

It follows that all terminals of T'_i are in the same connected component of $G \setminus S$, contradicting that S is an edge Steiner multicut of (G, \mathcal{T}) . Therefore, $S' \setminus Y$ is an edge Steiner multicut of $(\hat{G}, \hat{\mathcal{T}})$.

Let X be any edge Steiner multicut of $(\hat{G}, \hat{\mathcal{T}})$. Suppose that T'_i is not separated in $G \setminus (Y \cup X \cup S'')$. Consider any two terminals $s, s' \in \hat{T}_i$. Define $z \in V(G)$ such that $z = s$ if s is not a border vertex, and such that z is a terminal of T'_i in the connected component of $G'' \setminus S''$ that contains s otherwise. Define z' similarly. Since we made the right choice and by construction, z and z' are properly defined. Since T'_i is not separated in $G \setminus (Y \cup X \cup S'')$, there is a path P in $G \setminus (Y \cup X \cup S'')$ between z and z' . If s (respectively s') is a border vertex, then P contains s (respectively s'). Hence, there is a path P' between s and s' in $G \setminus (Y \cup X \cup S'')$. Observe that P' consists of several subpaths in $G' \setminus (Y \cup X)$ and several in $G'' \setminus S''$. Let Q be any maximal subpath of P' in $G' \setminus (Y \cup X)$. Since G'_F is obtained from G' by contracting edges and identifying vertices, Q corresponds to a walk W in $G'_F \setminus (Y \cup X)$ between the same vertices. By the construction of H_F , W corresponds to a walk W' in $H_F \setminus X$ between the same vertices. Finally, by the construction of \hat{G} , W' corresponds to a walk in $\hat{G} \setminus X$ between the same vertices. Since we made the right choice, any subpath of P' that lies in G'' and goes between two border vertices, can be replaced by the vertex of \hat{G} into which these two border vertices were identified. Hence, P' can be ‘compressed’ into an s, s' -walk of $\hat{G} \setminus X$, and thus s, s' are in the same connected component of $\hat{G} \setminus X$. It follows that all terminals of \hat{T}_i are in the same connected component of $\hat{G} \setminus X$, contradicting that X is an edge Steiner multicut of $(\hat{G}, \hat{\mathcal{T}})$. Therefore, $X \cup S''$ is an edge Steiner multicut of (G, \mathcal{T}) . \square

We now aim to find a smallest edge Steiner multicut X of $(\hat{G}, \hat{\mathcal{T}})$ that is an all-or-nothing cut. Let $\{C'_1, \dots, C'_u\}$ be the set of connected components of $\hat{G} \setminus \{h_F\}$. Let $\hat{\mathcal{T}}|_i$ denote the set of terminal sets in $\hat{\mathcal{T}}$ that are separated if one removes all edges of $E(\hat{G}[\{h_F\} \cup C'_i])$ from G'_F . Define $z[\mathcal{U}, i]$, where $\mathcal{U} \subseteq \hat{\mathcal{T}}$ and $1 \leq i \leq u$, as the size of the smallest all-or-nothing cut of the terminal sets in \mathcal{U} using only edges in or going out of C'_1, \dots, C'_i . Then for any $\mathcal{U} \subseteq \hat{\mathcal{T}}$,

$$z[\mathcal{U}, 1] = \begin{cases} \infty & \text{if } \mathcal{U} \not\subseteq \hat{\mathcal{T}}|_1 \\ |E(G'_F[\{h_F\} \cup C'_1])| & \text{otherwise (i.e. if } \mathcal{U} \subseteq \hat{\mathcal{T}}|_1) \end{cases}$$

and for $i > 1$,

$$z[\mathcal{U}, i] = \min \left\{ z[\mathcal{U}, i-1], |E(G'_F[\{h_F\} \cup C'_i])| + z[\mathcal{U} \setminus \hat{\mathcal{T}}|_i, i-1] \right\}.$$

Note that $z[\hat{\mathcal{T}}, u]$ holds the size of the smallest edge Steiner multicut of $(\hat{G}, \hat{\mathcal{T}})$ that is an all-or-nothing cut (if one exists). Finding the set achieving this smallest size is straightforward from the dynamic-programming table. Finally, over all choices of F and all choices of Y , mark in G' the smallest set of edges that was found if it has size at most k .

By Lemma 3.9, if we made the right choice, we mark a set Z such that $Z \cup S''$ is a smallest edge Steiner multicut of (G, \mathcal{T}) .

In both cases, let M be the set of marked edges. Now we contract all unmarked edges $E(G') \setminus M$ in G . Let \tilde{G} denote the resulting graph; note that in general \tilde{G} is a multigraph. Each time we contract an edge between two vertices u and v , we replace u and v by uv in all terminal sets in \mathcal{T} . Let $\tilde{\mathcal{T}}$ denote the resulting set of terminal sets.

Observe that if a terminal set \tilde{T}_i in $\tilde{\mathcal{T}}$ is a singleton set and T_i was not a singleton set in \mathcal{T} , then we can answer “no”. Indeed, if (G, \mathcal{T}, k) would be a “yes”-instance, then from the construction of M , for any smallest edge Steiner multicut S of (G, \mathcal{T}) there is a smallest edge Steiner multicut $Z \cup (S \setminus E(G'))$ of (G, \mathcal{T}) such that $Z \subseteq M$. In particular, since $T_i \subseteq V(G')$ for it to be contracted to a single vertex, Z (and thus also M) is an edge Steiner multicut of $(G', \{T_i\})$, which contradicts that all vertices of T_i are contracted into a single vertex.

Now it remains to prove that (G, \mathcal{T}, k) is a “yes”-instance if and only if $(\tilde{G}, \tilde{\mathcal{T}}, k)$ is. Let \tilde{S} be an edge Steiner multicut of $(\tilde{G}, \tilde{\mathcal{T}})$ of size at most k . Since \tilde{G} is obtained from G by a sequence of edge contractions (without deleting parallel edges), there is a natural mapping from $E(\tilde{G})$ to $E(G)$. Let S denote the set of edges obtained from \tilde{S} through this mapping. If S is not an edge Steiner multicut of (G, \mathcal{T}) , then there is a path between in $G \setminus S$ between any two terminals of some terminal set of \mathcal{T} . However, because of the way S was obtained from \tilde{S} , any path in $G \setminus S$ corresponds to a walk in $\tilde{G} \setminus \tilde{S}$. Hence, there is a walk in $\tilde{G} \setminus \tilde{S}$ between any two terminals of this terminal set in $\tilde{\mathcal{T}}$, a contradiction. The converse follows immediately from the construction of M , i.e. that for any smallest edge Steiner multicut S of (G, \mathcal{T}) there is a smallest edge Steiner multicut $Z \cup (S \setminus E(G'))$ of (G, \mathcal{T}) such that $Z \subseteq M$.

Finally, we analyze the running time of our algorithm. Since there are at most $2k$ border vertices and t terminal sets, there are $r = 2^{O(k t \log k)}$ different branches that we consider for \mathcal{T}' , \mathcal{B} , and \mathcal{B}_i , and in each we mark at most k edges. Choose $q = rk + 1$. Now note that $|E(G')| \geq q$:

- if $G = G'$, then this follows from the assumption that $|E(G)| > q$.
- if $G \neq G'$, then G' was obtained after considering multiple (q, k) -good separations. Hence, $|V(G')| > q$, and since G' is connected, $|E(G')| \geq q$.

Since $q = rk + 1$, at least one edge of G' was not marked and thus contracted. Therefore, $|V(G)|$ decreases by at least one, and the entire procedure finishes after at most $|V(G)|$ iterations.

To analyze the running time, observe that the dynamic-programming procedure requires $2^{O(t)}k + O(t|E(G)|)$ time. The family \mathcal{F} has $2^{O(k^2 t \log k)} \log |E(G)|$ sets and can be constructed in $2^{O(k^2 t \log k)} |E(G)| \log |E(G)|$ time. Hence, Case 2 runs in $2^{O(k^2 t \log k)} |E(G)| \log |E(G)|$ time. Case 1 runs in $2^{O(k^2 t \log k)}$ time. Since there are $r = 2^{O(k t \log k)}$ different branches for \mathcal{T}' , \mathcal{B} , and \mathcal{B}_i that we consider, and it takes $2^{O(k^2 t \log k)} |V(G)|^4 \log |V(G)|$ time to find a $2k$ -bordered subgraph that does not admit a (q, k) -good separation, each iteration takes $2^{O(k^2 t \log k)} |V(G)|^4 \log |V(G)|$ time. Since there are at most $|V(G)|$ iterations, the total running time is $2^{O(k^2 t \log k)} |V(G)|^5 \log |V(G)|$. Actually, using the recurrence outlined by Chitnis et al. [17], one can show a bound on the running time of $2^{O(k^2 t \log k)} |V(G)|^4 \log |V(G)|$.

4 Tractability for Edge Deletion and Parameter $k + t$: Second Proof

In this section, we give the second proof of Theorem 1.1 through important separators and treewidth reduction. Our presentation of the algorithm is optimized for readability, not for the final runtime.

We first show an easy reduction to the following problem variant.

MULTIPEDAL STEINER MULTICUT

Input: A connected undirected graph G , a set $Y \subseteq V(G)$, terminal sets $\mathcal{T} = \{T_1, \dots, T_t \subseteq V(G)\}$ with $T_i \cap Y \neq \emptyset$ for all i , an integer k .

Parameter: $k + t + |Y|$

Question: Find an EDGE STEINER MULTICUT S of T_1, \dots, T_t of size at most k .

Lemma 4.1. EDGE STEINER MULTICUT for the parameter $k + t$ is fixed-parameter tractable if MULTIPEDAL STEINER MULTICUT is fixed-parameter tractable.

Proof. Note that the collection of terminal sets $\mathcal{T} = \{T_1, \dots, T_t\}$ has a hitting set Y of size at most t that can be computed in polynomial time (just pick any node $v_i \in T_i$ and set $Y := \{v_1, \dots, v_t\}$). Since $|Y| \leq t$ and t is a parameter, this is indeed a parameterized reduction to MULTIPEDAL STEINER MULTICUT. \square

In the rest of this section, we first show a fixed-parameter algorithm for the special case of MULTIPEDAL STEINER MULTICUT where $|Y|=1$ (in Section 4.1), and then generalize this algorithm to MULTIPEDAL STEINER MULTICUT (in Section 4.2). We will use the following simple fact.

Lemma 4.2. *Let $S \subseteq E(G)$, T be a terminal set and $x \in T$. Then S cuts T if and only if there is a node $v \in T$ such that S cuts v from x .*

Proof. Recall that “ S cuts T ” means that S cuts some pair of nodes in T . Thus, if there is a node $v \in T$ such that S cuts v from x then clearly S cuts T . Further, if no node $v \in T$ is cut from x by S , then T is a connected set in $G \setminus S$, so that T is not cut by S . \square

4.1 Unipedal Steiner Multicut

We first show how to solve the special case of MULTIPEDAL STEINER MULTICUT for $|Y|=1$.

UNIPEDAL STEINER MULTICUT

Input: A connected undirected graph G , a node $y \in V(G)$, terminal sets $T_1, \dots, T_t \subseteq V(G)$ with $y \in T_i$ for all i , an integer k .

Parameter: $k+t$

Question: Find an EDGE STEINER MULTICUT S of T_1, \dots, T_t of size at most k .

Theorem 4.3. UNIPEDAL STEINER MULTICUT *is fixed-parameter tractable.*

Our algorithm for UNIPEDAL STEINER MULTICUT heavily relies on the notions of important separators and closest cuts, due to Marx and Razgon [49] and Marx[47].

Let G be an undirected graph and $Y \subseteq V(G)$. A set $S \subseteq E(G)$ cuts $V(G)$ into two sets: $R^Y(S) = R_G^Y(S)$, the union of components of $G \setminus S$ that contain a node of Y (i.e., the nodes reachable from Y in $G \setminus S$), and $\overline{R}^Y(S) = \overline{R}_G^Y(S)$, the union of components of $G \setminus S$ disjoint from Y .

Definition 4.4. A cut $S \subseteq E(G)$ a Y -closest cut if it is inclusion-wise minimal (i.e., there is no $S' \subset S$ with $R^Y(S) = R^Y(S')$) and there is no $S' \subseteq E(G)$ with $|S'| \leq |S|$ and $R^Y(S') \subset R^Y(S)$.

Let $X, Y \subseteq V(G)$ be two disjoint sets. A set $S \subseteq E(G)$ is an $X-Y$ separator if no node in Y is reachable from a node in X in $G \setminus S$. For a node x we also write $x-Y$ separator instead of $\{x\}-Y$ separator and $R^x(S)$ instead of $R^{\{x\}}(S)$.

Definition 4.5 ([46]). Let $X, Y \subseteq V(G)$ be two disjoint sets and $S \subseteq E(G)$ an $X-Y$ separator. We call S an important $X-Y$ separator if it is inclusion-wise minimal and there is no $X-Y$ separator S' with $|S'| \leq |S|$ and $R^X(S') \supset R^X(S)$.

The most valuable properties of important separators with size at most ℓ are that there are not too many of them, and that we can enumerate all of them in $O^*(f(\ell))$ time. We remark that typically this is proven for the node deletion version of important separators, but a standard construction of the line graph (augmented with appropriate terminal node) shows that the same result also holds for the edge deletion variant that we consider in this section.

Theorem 4.6 ([49]). Let $X, Y \subseteq V(G)$ be two disjoint sets. The number of important $X-Y$ separators of size at most ℓ is at most 4^ℓ . Furthermore, these important separators can be enumerated in time $O(4^\ell \ell(n+m))$.

We now come to the key property of edge separators that we use for our fixed-parameter result. It shows that any Y -closest set is a *disjoint* union of important $x-Y$ separators, where x may range over all nodes in $V(G)\setminus Y$.

Fix a set $Y \subseteq V(G)$ and integer ℓ and let $x \in V(G)\setminus Y$. We denote by \mathcal{I}_x the set of important $x-Y$ separators of size at most ℓ . Further set $\mathcal{I}_\ell^Y := \bigcup_{x \in V(G)\setminus Y} \mathcal{I}_x$. We denote a *disjoint* union by \uplus .

Lemma 4.7. *Let S be a Y -closest cut. Then there are $S_1, \dots, S_m \in \mathcal{I}_\ell^Y$ with $S = \uplus_i S_i$ and $\bar{R}^Y(S) = \uplus_i \bar{R}^Y(S_i)$.*

Proof. Let S be a Y -closest cut. Let C_1, \dots, C_m be the components of $G \setminus S$ disjoint from Y . Let $S_i \subseteq S$ be the edges incident to a node in C_i and a node in $R^Y(S)$. Since C_i is not reachable from Y , the set S_i consists of all $C_i - R^Y(S)$ edges in G . Clearly, the sets S_i are pairwise disjoint. Further, if S contains an edge in $G[R^Y(S)]$ or $G[\bar{R}^Y(S)]$ then the deletion of this edge does not change $R^Y(S)$, as this would contradict S being inclusion-wise minimal. Hence, we have $S = \uplus_i S_i$. Furthermore, this shows that $\bar{R}^Y(S) = C_1 \uplus \dots \uplus C_m = \bar{R}^Y(S_1) \uplus \dots \uplus \bar{R}^Y(S_m)$.

It remains to argue that the S_i are important separators, i.e., that $S_i \in \mathcal{I}_\ell^Y$ for all i . Consider any $x \in C_i$. Since S_i consists of all $C_i - R^Y(S)$ edges in G , C_i is a connected component, and S contains no edges in $G[\bar{R}^Y(S)]$, S_i is an inclusion-wise minimal $x-Y$ separator. Assume for the sake of contradiction that S_i is not an important $x-Y$ separator. Then there exists a set $S'_i \subset E(G)$ such that $|S'_i| \leq |S_i|$ and $R^x(S'_i) \supset R^x(S_i)$. Note that $R^x(S'_i)$ contains a node z in $N(C_i) \subseteq R^Y(S_i)$ (otherwise S'_i contains S_i). Further, observe that $z \notin \bar{R}^Y(S)$ (otherwise z is in some component C_j and in $N(C_i)$, so S contains an edge in $G[\bar{R}^Y(S)]$, contradicting minimality). For notational purposes, set $S'_j := S_j$ for $j \neq i$. Consider the set $S' := \bigcup_j S'_j$. We have $|S'| \leq \sum_j |S'_j| \leq |S|$ (this uses that S is the *disjoint* union of the S_j). Further, we have $\bar{R}^Y(S') \supseteq \bigcup_j \bar{R}^Y(S'_j) \supset \bar{R}^Y(S)$, since z is in $\bar{R}^Y(S'_i) \supseteq R^x(S'_i)$ but not in $\bar{R}^Y(S)$. This contradicts S being a Y -closest cut. \square

Observe that the number of Y -closest cuts can be huge, e.g., the star with midpoint y and n outgoing edges has $\binom{n}{\ell}$ y -closest cuts of size ℓ . However, the number of important separators is much smaller: Theorem 4.6 shows that $|\mathcal{I}_\ell^Y| \leq 4^\ell n$. Since Lemma 4.7 shows that Y -closest cuts are generated by important separators, we can optimize over Y -closest cuts, although their number may be huge. We explain the details of this in the remainder of this section.

We want to emphasize that, although the above lemma resembles the Pushing Lemma [49, Lemma 3.10] at first glance, it does not hold for the node deletion variant of closest cuts and important separators (see Marx and Razgon [49] for the analogous definitions). By inspection of the following example one can see that a closest cut is in general not a *disjoint* union of important separators in the node deletion case.




Figure 1: Example showing that a closest cut is in general *not* a disjoint union of important separators in the node deletion case.

Example: Node deletion variant. *Consider the graph depicted in Figure 1. Then $\{s_1, s_2, s_3\}$ is a y -closest cut, which is the union of the important v_1-y separator $\{s_1, s_2\}$ and the important v_2-y separator $\{s_2, s_3\}$. However, $\{s_1, s_2, s_3\}$ is not a disjoint union of important separators.*

We are now ready to present the fixed-parameter algorithm for UNIPEDAL STEINER MULTICUT. Recall that we are given a graph G with a special node $y \in V(G)$ and terminal sets T_1, \dots, T_t with $y \in T_i$ for all i . We want to find an edge Steiner multicut S^* of size at most k . Recall that a set $S \subseteq E(G)$ cuts a terminal set T_i if and only if S cuts a node $v \in T_i$ from y (Lemma 4.2).

We define the *type* of an important separator $S \in \mathcal{I}_k^{\{y\}}$ as the set of all $i \in \{1, \dots, t\}$ with $T_i \cap \overline{R}^y(S) \neq \emptyset$, i.e., the set of all i such that T_i is cut by S , denoted by $\text{type}(S)$. Our algorithm is a simple dynamic program where, after a trivial initialization, we iterate over all important separators $S \in \mathcal{I}_k^{\{y\}}$ and update the table DP as follows:

```

DP[ $\emptyset$ ] := 0
for all  $\emptyset \neq I \subseteq [t]$  do
  DP[ $I$ ] :=  $\infty$ 
for all  $S \in \mathcal{I}_k^{\{y\}}$  do
  for all  $I \subseteq [t]$  do
    DP[ $I \cup \text{type}(S)$ ] :=  $\min\{\text{DP}[I \cup \text{type}(S)], \text{DP}[I] + |S|\}$ 
return DP[ $\{1, \dots, t\}$ ]

```

By Theorem 4.6 we have $|\mathcal{I}_k^{\{y\}}| \leq 4^k n$ and we can enumerate $\mathcal{I}_k^{\{y\}}$ in $O^*(4^k)$ time. Overall, the above algorithm runs in time $O^*(2^{2k+t})$. We remark that the running time of the algorithm is $O(n(n+m))$ for fixed k, t . It remains to prove correctness of the algorithm.

Lemma 4.8. *The dynamic program returns a value of at most k if and only if the optimal edge Steiner multicut S^* has size at most k (in which case the return value coincides with $|S^*|$).*

Note that determining the size of an optimal solution suffices because of self-reducibility. With Lemma 4.8, we also finish the proof of Theorem 4.3.

Proof of Lemma 4.8. Suppose that the optimal edge Steiner multicut S^* has size at most k . We can assume that S^* is a $\{y\}$ -closest cut. Otherwise, we can replace S^* by a $\{y\}$ -closest cut S' with the same (or lower) cost and $R^y(S') \subset R^y(S^*)$, and thus S' still cuts all sets T_i that are cut by S^* , see Lemma 4.2. By Lemma 4.7, there are important separators $S_1, \dots, S_m \in \mathcal{I}_k^{\{y\}}$ such that $S^* = \biguplus_i S_i$, implying that $|S^*| = \sum_i |S_i|$, and $\overline{R}^y(S^*) = \biguplus_i \overline{R}^y(S_i)$. Note that the latter implies that $\text{type}(S^*) = \bigcup_i \text{type}(S_i)$. As S^* cuts all terminal sets T_i we have $\text{type}(S^*) = \{1, \dots, t\}$, showing $\bigcup_i \text{type}(S_i) = \{1, \dots, t\}$. Hence, $\bigcup_i S_i$ appears as one term in $\text{DP}[\{1, \dots, t\}]$, so that we return a value $\text{DP}[\{1, \dots, t\}] \leq \sum_i |S_i| = |S^*| \leq k$.

For the other direction, suppose that we return $\text{DP}[\{1, \dots, t\}] = \ell$. Then there are important separators $S_1, \dots, S_m \in \mathcal{I}_k^{\{y\}}$ such that $\sum_i |S_i| = \ell$ and $\bigcup_i \text{type}(S_i) = \{1, \dots, t\}$. Consider $S := \bigcup_i S_i$. Since $\overline{R}^y(S) \supseteq \bigcup_i \overline{R}^y(S_i)$ we have $\text{type}(S) \supseteq \bigcup_i \text{type}(S_i) = \{1, \dots, t\}$, so S is a valid edge Steiner multicut. Further, $|S| \leq \sum_i |S_i|$, so we proved the existence of an edge Steiner multicut of size at most $\text{DP}[\{1, \dots, t\}]$. Together with the first direction this proves the claim. \square

4.2 Multipedal Steiner Multicut

In this section, we show a fixed-parameter algorithm for MULTIPEDAL STEINER MULTICUT in general.

We first show that any inclusion-wise minimal edge Steiner multicut S can be split into two disjoint parts, one part consisting of a union of minimal $x-y$ separators for some $x, y \in Y$, and one part where Y is not cut at all. For this, we denote by Z the union of all edges in G that appear in a minimal $x-y$ cut of size at most k for some $x, y \in Y$.

Lemma 4.9. *Let S be an inclusion-wise minimal edge Steiner multicut of G . Then we can write $S = S' \uplus S''$ where:*

- (1) $S' \subseteq Z$,

- (2) Y is contained in one component of $G \setminus S''$,
(3) $\overline{R}^y(S) = \overline{R}^y(S') \cup \overline{R}^y(S'')$ for all $y \in Y$.

Proof. In this proof we denote by $\delta(U)$ the set of edges in G that are incident to a node in U and a node in $V(G) \setminus U$, for any $U \subseteq V(G)$.

Consider an inclusion-wise minimal edge Steiner multicut S of size at most k . For $y \in Y$, let $S_y := \delta(R^y(S))$ (i.e., the edges leaving the set of nodes that are reachable from y in $G \setminus S$). Observe that $S_y \subseteq S$ and $R^y(S) = R^y(S_y)$. By the minimality of S , we have $S = \bigcup_{y \in Y} S_y$, since the deletion of any edge of S in $G[R^y(S)]$, for $y \in Y$, or in $G[\overline{R}^Y(S)]$ is redundant, as its deletion does not change $R^y(S)$ for any $y \in Y$ (and thus does not change which terminal set are disconnected).

Deleting a set S_y cuts G into some connected components. Note that if C is such a component, then $\delta(C) \subseteq S_y$ (otherwise, there is a node $v \in N(C)$ that is reachable in $G \setminus S_y$, which contradicts that C is a component) and $\delta(C)$ consists only of $C - R^y(S_y)$ edges (since S_y consists of edges incident to $R^y(S_y)$).

Let C'_1, \dots, C'_{r+1} be the components of $G \setminus S_y$ that contain some node from Y . Without loss of generality, we have $C'_{r+1} = R^y(S_y)$ (which we ignore from now on) and thus no other C'_i contains y . For such a component C'_i with $x \in Y \cap C'_i$, $1 \leq i \leq r$, the set of edges $\delta(C'_i)$ is a minimal $x - y$ cut: since $x \in C'_i, y \notin C'_i$ the set $\delta(C'_i)$ is certainly an $x - y$ cut, and since C'_i and $R^y(S_y)$ are connected and $\delta(C'_i)$ consists only of $C'_i - R^y(S_y)$ edges, it is even a minimal cut. Set $S'_y := \bigcup_{1 \leq i \leq r} \delta(C'_i) \subseteq S$ and $S' := \bigcup_{y \in Y} S'_y \subseteq S$. Note that we have $S' \subseteq Z$, since S' is a union of minimal $x - y$ cuts for some $x, y \in Y$, each of size at most $|S| \leq k$. This proves property (1).

Let C''_1, \dots, C''_ℓ be the remaining components of $G \setminus S_y$, i.e., the components that are disjoint from Y , and set $S''_y := \bigcup_i \delta(C''_i) \subseteq S$. Further, let $S'' := \bigcup_{y \in Y} S''_y \subseteq S$. We claim that Y is contained in one component of $G \setminus S''$. Consider any $x, y \in Y$ and a $x - y$ path P in G that passes through S'' a minimum number of times (among all $x - y$ paths). Let $z \in Y$ be such that S''_z contains some edge e of P in S'' . Before e , the path P passes through S''_z (if $x \notin R^z(S)$) or it starts in $R^z(S)$ (if $x \in R^z(S)$). Similarly, after e , the path P passes through S''_z or it ends in $R^z(S)$. In any case, since $R^z(S) \subseteq R^z(S'')$ (as $S'' \subseteq S$) we can replace P by a path that avoids e by routing it through the connected set $R^z(S)$. This decreases the number of passes of P through S'' , contradicting minimality. This proves property (2).

We show that $S = S' \uplus S''$. First note that $S_y = S'_y \uplus S''_y$ by construction. Above we showed $S = \bigcup_{y \in Y} S_y$, which implies $S = S' \cup S''$. Moreover, consider any component C''_i of $G \setminus S_y$ (that is disjoint from Y , as above). Note that C''_i is disjoint from $R^x(S) = R^x(S_x)$ for any $x \in Y$, as $x \notin C''_i$ and any $C''_i - x$ path passes through $\delta(C''_i) \subseteq S$. This shows that $S''_y \cap S_x = \emptyset$ for any $x, y \in Y, x \neq y$, since any edge $e \in S''_y \cap S_x$ would be incident to all three of the disjoint sets $R^y(S_y) = R^y(S)$, $R^x(S_x) = R^x(S)$, and C''_i . Hence, S''_y is disjoint from S'_x for all $x \in Y$, so that S' and S'' are disjoint.

Finally, we show $\overline{R}^y(S) = \overline{R}^y(S') \cup \overline{R}^y(S'')$ for any $y \in Y$. With notation as above for the components of $G \setminus S_y$ we have $\overline{R}^y(S) = \overline{R}^y(S_y) = C'_1 \cup \dots \cup C'_r \cup C''_1 \cup \dots \cup C''_\ell$. By construction we have $\overline{R}^y(S'_y) = C'_1 \cup \dots \cup C'_r$ and $\overline{R}^y(S''_y) = C''_1 \cup \dots \cup C''_\ell$, implying $\overline{R}^y(S) \subseteq \overline{R}^y(S') \cup \overline{R}^y(S'')$. Since $S \supseteq S', S''$ we also have $\overline{R}^y(S) \supseteq \overline{R}^y(S') \cup \overline{R}^y(S'')$. This proves property (3). \square

We may branch for each terminal set whether it is cut by S' or by S'' , splitting \mathcal{T} into $\mathcal{T}' \uplus \mathcal{T}''$. Furthermore, we may branch over $k' + k'' = k$ such that $|S'| \leq k'$ and $|S''| \leq k''$. This branching leaves us with two subproblems:

- (1) Find an edge Steiner multicut S' of (G, \mathcal{T}', k') with $S' \subseteq Z$,
- (2) Find an edge Steiner multicut S'' of (G, \mathcal{T}'', k'') such that Y is contained in one component of $G \setminus S''$.

At first sight, these subproblems are *not independent*, since in Lemma 4.9 the sets S' and S'' are required to be *disjoint*. Nevertheless, we can solve these subproblems independently and put the solutions together to a solution of the original MULTIPEDAL STEINER MULTICUT instance.

However, first we change the above subproblems slightly to make them easier to solve. Let $Z' = Z'(G, Y)$ with $Z \subseteq Z' \subseteq V(G)$ to be fixed later. We do this because it is not known how to determine Z , but a superset Z' with nice properties can be found using ideas by Marx et al. [48], as we will explain later. Then we replace subproblem (1) by the following problem:

(1') Find an edge Steiner multicut S' of (G, \mathcal{T}', k') with $S' \subseteq Z'$.

For subproblem (2), let G_* be the graph obtained from G by adding a new node y_* and connecting y_* to every $y \in Y$ by $k+1$ parallel edges (to avoid parallel edges, one may alternatively subdivide all these parallel edges). Further, for a terminal set T , let $T_* := (T \setminus Y) \cup \{y_*\}$ and let $\mathcal{T}_* := \{T_* \mid T \in \mathcal{T}\}$. We replace subproblem (2) by the following problem. Intuitively, y_* forces Y to be connected, so that actually find an edge Steiner multicut S'' that has Y in one component of $G \setminus S''$.

(2') Find an edge Steiner multicut S'' of $(G_*, \mathcal{T}_*, k'')$.

Now we show that the original MULTIPEDAL STEINER MULTICUT instance (G, \mathcal{T}, k) has a solution if and only if for some branch both subproblems (1') and (2') have a solution, finishing a reduction to the two subproblems.

Lemma 4.10. *There is an edge Steiner multicut of (G, \mathcal{T}, k) if and only if for some branch (over $\mathcal{T}' \uplus \mathcal{T}'' = \mathcal{T}$ and $k' + k'' = k$) there is an edge Steiner multicut S' of (G, \mathcal{T}', k') with $S' \subseteq Z'$ and an edge Steiner multicut S'' of $(G_*, \mathcal{T}_*, k'')$.*

Proof. Consider a branch (over $\mathcal{T}' \uplus \mathcal{T}'' = \mathcal{T}$ and $k' + k'' = k$). Assume there is an edge Steiner multicut S' of (G, \mathcal{T}', k') with $S' \subseteq Z'$ and an edge Steiner multicut S'' of $(G_*, \mathcal{T}_*, k'')$. We can assume, without loss of generality, that S'' does not contain any edges in $E(G_*) \setminus E(G)$, since all these edges have $k+1$ parallel copies, so that picking any of these edges does not separate any nodes. Note that $S := S' \cup S'' \subseteq V(G)$ is of size $|S| \leq |S'| + |S''| \leq k' + k'' = k$. Consider a terminal set $T \in \mathcal{T}$ with $y \in T \cap Y$. If $T \in \mathcal{T}'$, then some node $v \in T$ is cut from y by S' (since if all nodes of T connect to y then T is not cut by S'). Observe that $\overline{R}^y(S) \supseteq \overline{R}^y(S') \cup \overline{R}^y(S'')$. Thus, v is also cut from y by S . A similar arguments works if $T \in \mathcal{T}''$, so that T_* is cut by S'' . Then some node $v \in T_*$ is cut from $y_* \in T_*$. Since y_* is connected to y in $G \setminus S''$, v is also cut from y by S'' . Hence, T is cut by S'' in G_* and, thus, T is cut by S in G . We have shown that S is an edge Steiner multicut of (G, \mathcal{T}, k) .

For the other direction, let S be an edge Steiner multicut of (G, \mathcal{T}, k) . Without loss of generality, we assume that S is inclusion-wise minimal. Pick sets S', S'' as in Lemma 4.9 and consider a terminal set T and $y \in T \cap Y$. Some node $v \in T$ is cut from y by S . Since we have $\overline{R}^y(S) = \overline{R}^y(S') \cup \overline{R}^y(S'')$ by property (3) in Lemma 4.9, node v is also cut from y by S' or S'' . Hence, if we let $\mathcal{T}' \subseteq \mathcal{T}$ be all terminal sets cut by S' and $\mathcal{T}'' := \mathcal{T} \setminus \mathcal{T}'$, then all terminal sets in \mathcal{T}'' are cut by S'' . Set $k' := |S'|$ and $k'' := k - k'$, then since $|S'| + |S''| = |S| \leq k$ we have $|S''| \leq k''$. Note that $\mathcal{T}' \uplus \mathcal{T}'' = \mathcal{T}$ and $k' + k'' = k$ is a valid branch. Now, S' is an edge Steiner multicut of (G, \mathcal{T}', k') with $S' \subseteq Z \subseteq Z'$ by property (1) in Lemma 4.9. Also, S'' is an edge Steiner multicut of (G, \mathcal{T}'', k'') . To see that S'' is an edge Steiner multicut of $(G_*, \mathcal{T}_*, k'')$, note that by property (2) of Lemma 4.9, Y is contained in one connected component of $G \setminus S''$. Since in the construction of G_* we only add new paths between nodes of Y , S'' is an edge Steiner multicut of $(G_*, \mathcal{T}_*, k'')$. As $y \in Y$ and y_* are connected in $G_* \setminus S''$, and $T \cap Y \neq \emptyset$ in any terminal set T , we may replace the nodes in $T \cap Y$ by y_* , so S'' still cuts T_* . Hence, S'' also cuts \mathcal{T}_* . \square

Note that $(G_*, \mathcal{T}_*, k'')$ is an instance of UNIPEDAL STEINER MULTICUT, since each terminal set contains y_* . Hence, we can solve subproblem (2') in $O^*(f(k, t))$ time by Theorem 4.3, and particularly, in $O(n(n+m))$ time for fixed k, t .

It remains to show how to solve subproblem (1'). We want to use the techniques by Marx et al. [48]. However, as their work considers the node deletion variant of cuts, we first transfer our problem to RESTR. NODE STEINER MULTICUT. Let $V := V(G)$, $E := E(G)$ and consider the following graph, which is closely related to the line graph of G :

$$H = (E \cup V, \{\{e, e'\} \mid e, e' \in E, |e \cap e'| = 1\} \cup \{\{v, e\} \mid e \in E, v \in V, v \in e\}).$$

Then EDGE STEINER MULTICUT on (G, \mathcal{T}', k') is equivalent to RESTR. NODE STEINER MULTICUT on (H, \mathcal{T}', k') , in a sense clarified in Lemma 4.12.

We make use of the following theorem, which is a variant of the Treewidth Reduction Theorem by Marx et al. [48, Theorem 2.11] (this variant follows from using only the first three lines of the proof of Marx et al. [48, Theorem 2.11]). The *torso* of a graph G with respect to a set $U \subseteq V(G)$ is the graph on node set U where $u, v \in U$ are adjacent if and only if there is a $u-v$ path in G that is internally node-disjoint from U (in particular, if $\{u, v\} \in E(G)$, then u and v are adjacent in the torso). We denote this graph by $\text{torso}(G, U)$.

Theorem 4.11 ([48]). *Let H be a graph, $Y \subseteq V(H)$, and let k be an integer. Let \tilde{Z} be the set of nodes of H participating in a minimal $x-y$ cut $\tilde{S} \subseteq V(H)$ of size at most k for some $x, y \in Y$. There is an algorithm that, in time $O^*(f(k, |Y|))$, computes a set $\tilde{Z}' \supseteq \tilde{Z} \cup Y$ such that $\text{torso}(H, \tilde{Z}')$ has treewidth at most $h(k, |Y|)$ for some function h .*

We remark that the algorithm of the above theorem runs in $O(n+m)$ time for fixed $k, |Y|$.

Now we fix $Z' := \tilde{Z}' \cap E(G)$. Note that $\text{tw}(\text{torso}(H, Z')) \leq \text{tw}(\text{torso}(H, \tilde{Z}')) \leq h(k, |Y|)$. We augment the graph $\text{torso}(H, Z')$ by the nodes $V(G)$ to a new graph H' , as follows.

For each node $v \in V(G)$, consider the connected component C of $H \setminus Z'$ that contains v . Let K be the set of neighbors of C in Z' . Since C is a connected component and by the construction of the torso, K forms a clique in $\text{torso}(H, Z')$. We add the node v to the graph $\text{torso}(H, Z')$ and make it adjacent to all elements of K . This yields a new graph H'_v . We show that H'_v also has bounded treewidth. Since K is a clique, by basic knowledge about tree decompositions, every tree decomposition of $\text{torso}(H, Z')$ has a bag containing K [4, Lemma 3.1]; let (T, \mathcal{B}) be a minimum-width tree decomposition of $\text{torso}(H, Z')$ and let $B \in \mathcal{B}$ be a bag containing K . We add a new bag $B' = B \cup \{v\}$ to (T, \mathcal{B}) and make it adjacent to B , yielding a tree decomposition of H'_v . This tree decomposition of H'_v has width at most one more than the width of (T, \mathcal{B}) , and thus the treewidth of H'_v is at most one more than the treewidth of $\text{torso}(H, Z')$. As the constructions of H'_v and $H'_{v'}$ for distinct nodes $v, v' \in V(G)$ do not interfere, in total the treewidth of H' is increased by at most one compared to the treewidth of $\text{torso}(H, Z')$. Hence, we have $\text{tw}(H') \leq h(k, |Y|) + 1$.

Lemma 4.12. *There is an edge Steiner multicut S of (G, \mathcal{T}', k') with $S \subseteq Z'$ if and only if there is a restr. node Steiner multicut S' of (H', \mathcal{T}', k') .*

Proof. Let S be a restr. node Steiner multicut of (H', \mathcal{T}', k') . Without loss of generality, we can assume that $S \subseteq Z'$, since $V(H') = Z' \cup V$ and the deletion of any non-terminal node $v \in V$ is redundant, since the neighborhood $N(v)$ is a clique (so that any path through v can be stripped of v). We show that S is also an edge Steiner multicut of (G, \mathcal{T}', k') . Assume, for the sake of contradiction, that $u, v \in T \in \mathcal{T}'$ are cut by S in H' , but there is a $u-v$ path in $G \setminus S$. Let (e_1, \dots, e_ℓ) be the edges of this path. Then $(u, e_1, \dots, e_\ell, v)$ is a $u-v$ path in $H \setminus S$ (where this time we wrote down the *nodes* of the path). Let e_i be the first and e_j be the last edge of this path that is in Z' . By basic properties of the torso [49, Proposition 3.3], e_i and e_j are also connected in $\text{torso}(H, Z') \setminus S$. Further, by the construction of H' , u is adjacent to e_i and v is adjacent to e_j in $H' \setminus S$. Hence, there is a $u-v$ path in $H' \setminus S$, which contradicts that S cuts u, v in H' .

For the other direction, let S be an edge Steiner multicut of (G, \mathcal{T}', k') with $S \subseteq Z'$. We show that S is also a restr. node Steiner multicut of (H', \mathcal{T}', k') . Assume, for the sake of contradiction, that $u, v \in T \in \mathcal{T}'$ are cut by S in G , but not in H' , so that there is a $u-v$ path in $H' \setminus S$. By the construction of $\text{torso}(H, Z')$ and H' (and since $S \subseteq Z'$), there is also a $u-v$ path in $H \setminus S$. However, any such path corresponds to a path in $G \setminus S$, which contradicts that u, v are cut by S in G . Hence, S is a node Steiner multicut of (H', \mathcal{T}', k') . Since $S \subseteq Z' \subseteq E$ and $\mathcal{T}' \subseteq V$, S is also a restr. node Steiner multicut. \square

Finally, the instance (H', \mathcal{T}', k') can be solved in $O^*(f(k, \text{tw}(G)))$ time, since the treewidth of H' is bounded in k and RESTR. NODE STEINER MULTICUT is fixed-parameter tractable for parameter $t + \text{tw}$ by Theorem 5.2. In particular, that algorithm runs in $O(n + m)$ time for fixed t, tw . Therefore, MULTIPEDAL STEINER MULTICUT is fixed-parameter tractable. The final algorithm runs in time $O(n(n + m))$ for fixed k, t ; moreover, since the treewidth of the structure returned by Theorem 4.11 is exponential, the algorithm requires double-exponential time. Combined with Lemma 4.1, this proves Theorem 1.1.

5 Steiner Multicuts for Graphs of Bounded Treewidth

In this section, we consider STEINER MULTICUT on graphs of bounded treewidth. To start the exposition, we note that EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT are NP-complete for trees and NODE STEINER MULTICUT is NP-complete on series-parallel graphs [12], which are graphs of treewidth two. This means that any efficient algorithm for STEINER MULTICUT on graphs of bounded treewidth needs an additional parameter.

We first show Theorem 1.3, namely that all variants of STEINER MULTICUT for the parameter k are W[1]-hard, even if $p=3$ and $\text{tw}(G)=2$ (but t is unbounded). The graph G is in fact a tree plus one node. We then contrast this result by showing that the problem is fixed-parameter tractable on bounded treewidth graphs when t is a parameter.

For the reduction to prove Theorem 1.3 we introduce an intermediate problem, that we call (MONOTONE) NAE-INTEGGER-3-SAT. In this problem, we are given variables x_1, \dots, x_k that each take a value in $\{1, \dots, n\}$ and clauses C_1, \dots, C_m of the form

$$\text{NAE}(x_{i_1} \leq a_1, x_{i_2} \leq a_2, x_{i_3} \leq a_3),$$

where $a_1, a_2, a_3 \in \{1, \dots, n\}$, and such a clause is satisfied if not all three inequalities are true and not all are false (i.e., they are “not all equal”). The goal is to find an assignment of the variables that satisfies all given clauses. We remark that NAE-INTEGGER-3-SAT generalizes MONOTONE NAE-3-SAT (by restriction to $n=2$), and that NAE-INTEGGER-3-SAT can be solved in time $O(m \cdot n^k)$, by enumerating all assignments. We complement this with a W[1]-hardness result for parameter k .

To prove that NAE-INTEGGER-3-SAT is W[1]-hard parameterized by k , we reduce from MULTICOLORED CLIQUE. In that problem, which is known to be W[1]-hard [28], we are given a graph G and a (proper) coloring of G using k colors, and the goal is to decide if G has a clique that contains at least one node with each of the k colors. We use V_i to denote the set of nodes of color i , set $n_i = |V_i|$, and $E_{i,j}$ to denote the set of edges with one endpoint in V_i and the other in V_j .

Lemma 5.1. *NAE-INTEGGER-3-SAT is W[1]-hard for parameter k .*

Proof. Let (G, k) be an instance of MULTICOLORED CLIQUE. We create an instance of NAE-INTEGGER-3-SAT on variables x_i , one for each color $1 \leq i \leq k$, and y_{ij} , one for each pair of colors $1 \leq i < j \leq k$. We identify the nodes V_i with the integers $\{1, \dots, n_i\}$ in an arbitrary way. We restrict x_i to $\{1, \dots, n_i\}$ using the clause $\text{NAE}(x_i \leq 0, x_i \leq 0, x_i \leq n_i)$, and write $x_i = u$ if the number x_i corresponds to node u . Analogously, we can identify the edges $uv \in E_{i,j}$ with numbers in

$\{1, \dots, |E_{i,j}|\}$ and write $y_{ij} = uv$ if we pick the number corresponding to edge uv . Consider the following two constraints (for any edge uv , with u of color i and v of color j):

$$y_{ij} = uv \Rightarrow x_i = u, \quad \text{and} \quad y_{ij} = uv \Rightarrow x_j = v .$$

If we can encode these constraints with NAE-clauses, then any satisfying assignment of the constructed NAE-INTEGER-3-SAT instance corresponds to a clique in G , as all chosen pairs y_{ij} correspond to edges, and edges sharing a color i picked the same node x_i . We focus on the first constraint; the second constrained is handled similarly. Note that the first constraint is equivalent to

$$y_{ij} = uv \Rightarrow x_i \geq u, \quad \text{and} \quad y_{ij} = uv \Rightarrow x_i \leq u .$$

Again, without loss of generality, we focus on the first of these constraints. It is equivalent to

$$y_{ij} < uv \vee y_{ij} > uv \vee x_i \geq u,$$

which in turn can be written as

$$\text{NAE}(y_{ij} < uv, y_{ij} > uv, x_i \geq u),$$

since $y_{ij} < uv, y_{ij} > uv$ cannot both be true. Note that we can replace any inequality $x < a$ by $x \leq a - 1$ (and similarly for $x > a$). Hence, we can encode all desired constraints if we may use “ \leq ” and “ \geq ” inequalities, not only “ \leq ” inequalities, as is the case in the definition of NAE-INTEGER-3-SAT.

In the remainder of this proof, we reduce NAE-INTEGER-3-SAT with “ \leq ” and “ \geq ” inequalities to the original variant with only “ \leq ” inequalities. Given any instance of NAE-INTEGER-3-SAT with both types of inequalities, for any variable x we introduce a new variable \bar{x} . For any $1 \leq v \leq n$, we add the constraint $\text{NAE}(x \leq v, x \geq v, \bar{x} \leq n - v)$. This enforces $\bar{x} = n + 1 - x$. Finally, we replace any inequality $x \geq v$ by $\bar{x} \leq n + 1 - v$. This yields an equivalent NAE-INTEGER-3-SAT instance with only “ \leq ” inequalities. \square

We are now ready to prove Theorem 1.3.

Proof of Theorem 1.3. We first give a reduction from NAE-INTEGER-3-SAT to EDGE STEINER MULTICUT (satisfying $p=3$ and $\text{tw}(G)=2$). Then we show how to generalize the reduction to RESTRICTED NODE STEINER MULTICUT and NODE STEINER MULTICUT.

Consider an instance of NAE-INTEGER-3-SAT on variables x_1, \dots, x_k taking values in $\{1, \dots, n\}$ with clauses C_1, \dots, C_m . Take k paths consisting of n edges and identify their start nodes (to a common node s) and end nodes (to a common node t), respectively. The resulting graph G has $\text{tw}(G)=2$, since it is not a tree, but becomes a tree after deleting s (or t). Let v_j^i be the j -th node on the i -th path from s to t , so that $v_0^i = s$ and $v_n^i = t$. For each clause $\text{NAE}(x_{i_1} \leq a_1, x_{i_2} \leq a_2, x_{i_3} \leq a_3)$ we introduce a terminal set $\{v_{a_1}^{i_1}, v_{a_2}^{i_2}, v_{a_3}^{i_3}\}$ (note that we can assume $0 \leq a_j \leq n$ without loss of generality). Further, we let $\{s, t\}$ be a terminal set and set the cut size to k , i.e., we allow to delete k edges. This finishes the construction. In order to separate s from t we need to cut at least one edge of each of the k paths that connect s and t , and because the cut size is k we have to delete *exactly one* edge per path. Say we delete the x_i -th edge on the i -th path. This splits G into two components, one containing s and the other containing t . Note that we separate nodes v_j^i and $v_{j'}^{i'}$ by cutting at $x_i \leq j$ and $x_{i'} > j'$ (or with both inequalities the other way round), since then v_j^i is in the t -component and $v_{j'}^{i'}$ in the s -component. Hence, the following are equivalent:

- the terminal set $\{v_{a_1}^{i_1}, v_{a_2}^{i_2}, v_{a_3}^{i_3}\}$ is disconnected;
- some pair of nodes in this set is disconnected;
- among the inequalities $x_{i_j} \leq a_j$, $j=1,2,3$, one is true and one is false;
- the clause $\text{NAE}(x_{i_1} \leq a_1, x_{i_2} \leq a_2, x_{i_3} \leq a_3)$ is satisfied.

Therefore, the given NAE-INTEGER-3-SAT instance is equivalent to the constructed EDGE STEINER MULTICUT instance.

To prove hardness for NODE STEINER MULTICUT, we adapt the construction for EDGE STEINER MULTICUT. We take the same graph, but let each path contain $2n-1$ internal nodes.

As before, denote by v_j^i the j -th node on the i -th path, so that now $v_0^i = s$ and $v_{2n}^i = t$. We introduce terminal sets $\{v_1^i, v_{2n-1}^i\}$ for all $1 \leq i \leq k$. To separate these sets we have to delete at least one inner node of each of the k paths (i.e., a node among v_1^i, \dots, v_{2n-1}^i). By setting the cut size to k we make sure that we delete exactly one inner node of every path. Say we delete the x'_i -th node $v_{x'_i}^i$ of the i -th path. We call the node v_{2a-2}^i and v_{2a-1}^i the *representatives* of “ $x_i = a$ ”. Thus, for each clause $\text{NAE}(x_{i_1} \leq a_1, x_{i_2} \leq a_2, x_{i_3} \leq a_3)$ we introduce the terminal sets $\{v_{2a_1-c_1}^{i_1}, v_{2a_2-c_2}^{i_2}, v_{2a_3-c_3}^{i_3}\}$ for all $c_1, c_2, c_3 \in \{1, 2\}$, meaning that for all representatives we need to separate some pair. Note that after deleting the x'_i -th node, at least one representative of “ $x_i = a$ ” survives, and the one or two surviving representatives lie in the same component of $G \setminus \{x'_1, \dots, x'_k\}$. We let x_i be the minimal value a^* such that a surviving representative of “ $x_i = a^*$ ” lies in the t -component. Then the surviving representatives of “ $x_i = a$ ” are contained in the t -component if and only if $x_i \leq a$. Hence, taking a terminal set with surviving representatives $\{v_{2a_1-c_1}^{i_1}, v_{2a_2-c_2}^{i_2}, v_{2a_3-c_3}^{i_3}\}$, the terminal set is satisfied if and only if the clause $\text{NAE}(x_{i_1} \leq a_1, x_{i_2} \leq a_2, x_{i_3} \leq a_3)$ is satisfied, completing the proof.

For **RESTR. NODE STEINER MULTICUT**, hardness now follows from Lemma 2.1. \square

We now contrast the above theorem by showing that **STEINER MULTICUT** is fixed-parameter tractable for the parameter t if the graph has bounded treewidth.

Theorem 5.2. **NODE STEINER MULTICUT**, **EDGE STEINER MULTICUT**, and **RESTR. NODE STEINER MULTICUT** are fixed-parameter tractable for parameter $t + \text{tw}(G)$.

Proof. We first present an MSOL formula for **NODE STEINER MULTICUT**, extending the work of Gottlob and Lee [34] and Marx et al. [48] for **NODE MULTICUT**. Gottlob and Lee [34] construct an MSOL formula $\text{connects}(R, x, y)$, equal to

$$R(x) \wedge R(y) \wedge \forall P((P(x) \wedge \neg P(y)) \rightarrow (\exists v, w(R(v) \wedge R(w) \wedge P(v) \wedge \neg P(w) \wedge \text{adj}(v, w))))),$$

which expresses that x and y are connected in the subgraph of G induced by R . We then extend the formula of Marx et al. [48] to $p > 2$. We construct the following MSOL formula:

$$\bigwedge_{i=1}^t \exists x, y (T_i(x) \wedge T_i(y) \wedge \neg(x = y) \wedge \forall R(\text{connects}(R, x, y) \rightarrow \exists v (X(v) \wedge R(v)))),$$

where X is the potential cut set and $T_i(x)$ is true if and only if x is in T_i . A slight modification yields MSOL-formulas for **EDGE STEINER MULTICUT** and **RESTR. NODE STEINER MULTICUT**. Note that X is a free set variable in the formulae.

To apply these formulae, we use Bodlaender’s algorithm [5] to find a tree decomposition of G of width $\text{tw}(G)$ in time $f(\text{tw}(G)) \cdot \text{poly}(|V(G)|)$, for some computable function f . Then we input the necessary MSOL formula ϕ and the tree decomposition into the algorithm of Arnborg et al. [2], which runs in time $g(|\phi|, \text{tw}(G))$ for some function g and finds a smallest set X that satisfies ϕ . Since $|\phi|$ is polynomial in t , the theorem follows. \square

We remark that one can develop an algorithm that has a single-exponential running time of $O^*(2^{O(\text{tw}(t + \log \text{tw}))})$ via dynamic programming. Because such an algorithm is a straightforward application of known techniques, we only provide a sketch by defining the entries of the dynamic programming table. First we compute an approximately optimal tree decomposition in time $O^*(2^{O(\text{tw})})$ [8]. For each bag B in the tree decomposition we have a dynamic programming table with entries as follows. Let S be any set of k' edges, $0 \leq k' \leq k$, in the graph G_B induced by the subtree below B . Let $\mathcal{T}' \subset \mathcal{T}$ be the terminal sets that are cut by S even after making B into a clique; these terminal sets are cut no matter what we do above B . Let C_1, \dots, C_r be the connected components of $G_B \setminus S$ and let P_1, \dots, P_ℓ be the induced partitioning of the vertices in B . For any $T \in \mathcal{T} \setminus \mathcal{T}'$ and any part P_i we store whether a vertex in T is connected to a vertex in P_i in G_B . This reachability information together with k' , \mathcal{T}' , and P_1, \dots, P_ℓ can be stored using $O(\text{tw}(t + \log \text{tw}))$ bits. Thus, the dynamic programming table has $2^{O(\text{tw}(t + \log \text{tw}))}$ possible entries,

and for each possible entry we store whether it is realized by some set of edges S . At the root, from this information we can extract whether there exists an edge Steiner multicut.

6 Hardness for Cutsizes k and Number of Terminal Sets t

In this section, we consider the STEINER MULTICUT problem on general graphs parameterized by $k+t$. We show that both node deletion versions of the problem, NODE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT, are $W[1]$ -hard for this parameter.

Theorem 6.1. *NODE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT are $W[1]$ -hard for the parameter $k+t$.*

Proof. We present a parameterized reduction from the MULTICOLORED CLIQUE problem [28] to NODE STEINER MULTICUT. Let (H, k) be an instance of MULTICOLORED CLIQUE, and let V_i and $E_{i,j}$ be as in the definition of MULTICOLORED CLIQUE. We then create the following instance of NODE STEINER MULTICUT. First, we subdivide each edge of H , and let $N_{i,j}$ denote the set of nodes that were created when subdividing the edges of $E_{i,j}$. Then, add a complete graph C with $2k$ nodes, where we denote the nodes of C by c_1, \dots, c_{2k} , and make all nodes of V_i adjacent to c_{2i-1} and c_{2i} for each $i = 1, \dots, k$. Let G denote the resulting graph. Observe that $G[V(H)]$ and $G[\bigcup_{i,j} N_{i,j}]$ are both independent sets of G . We then create terminal sets $T_i = V_i \cup \{c_{2i-1}\}$ and $T'_i = V_i \cup \{c_{2i}\}$, and terminal sets $T_{i,j} = N_{i,j} \cup V(C)$. Let $\mathcal{T} = \{T_i, T'_i \mid i = 1, \dots, k\} \cup \{T_{i,j} \mid i \neq j, i, j = 1, \dots, k\}$. Then the created instance is (G, \mathcal{T}, k) .

Suppose that (H, k) is a “yes”-instance of MULTICOLORED CLIQUE, and let K denote a clique of H such that $V(K) \cap V_i \neq \emptyset$ for each i . Pick a node $v_i \in V(K) \cap V_i$ and let $S = \{v_i \mid i = 1, \dots, k\}$. Observe that v_i disconnects terminal sets T_i and T'_i . Further, if we let $n_{i,j}$ denote the subdivision node of the edge $(v_i, v_j) \in E(H)$, then v_i and v_j disconnect $n_{i,j}$ from the rest of $T_{i,j}$. Finally, $|S| = k$. Therefore, (G, \mathcal{T}, k) is a “yes”-instance of NODE STEINER MULTICUT.

Suppose that (G, \mathcal{T}, k) is a “yes”-instance of NODE STEINER MULTICUT, and let $S \subseteq V(G)$ denote a node Steiner multicut of G with respect to terminal sets \mathcal{T} such that $|S| \leq k$. We claim that $H[S]$ is a multicolored clique of H . First, observe that to disconnect T_i and T'_i , we need that $S \cap T_i \neq \emptyset$ and $S \cap T'_i \neq \emptyset$. Since $|S| \leq k$, we know that $S \cap T_i \cap T'_i \neq \emptyset$ for each $i = 1, \dots, k$. This implies that $|S| = k$, that $S \subseteq V(H)$, and that $S \cap V_i \neq \emptyset$ for each $i = 1, \dots, k$. It remains to show that $H[S]$ is a clique. Let v_i denote the node in $S \cap V_i$. Suppose that $(v_i, v_j) \notin E(H)$ for some i, j . Then consider the terminal set $T_{i,j}$, and observe that for any node $n \in N_{i,j}$ at least one endpoint of the edge corresponding to n is not in S . Since $S \cap V(C) = \emptyset$, this implies that $T_{i,j}$ is not disconnected by S , a contradiction. It follows that $H[S]$ is a clique, and thus (H, k) is a “yes”-instance of MULTICOLORED CLIQUE.

For RESTR. NODE STEINER MULTICUT hardness follows from the statement for NODE STEINER MULTICUT and Lemma 2.1. \square

7 Steiner Multicuts in Trees

In this section, we consider the complexity of STEINER MULTICUT on trees and completely characterize it with respect to various parameters. Throughout, we consider instances (G, \mathcal{T}, k) of STEINER MULTICUT (in either the edge, node, or restricted-node variant) where G is a tree. Here we assume that G is rooted at an arbitrary node $r \in V(G)$. For each terminal set $T \in \mathcal{T}$, let G_T denote the subtree of G induced by the terminals of T . That is, G_T is the union of the shortest paths between each pair of terminals of T . Since G is rooted, G_T also has a root, denoted by $r(G_T)$ and called a *terminal root*.

7.1 Polynomial-time Algorithm for Node Steiner Multicut

It is known that NODE MULTICUT can be solved in polynomial time on trees [12]. Here, we extend this algorithm to NODE STEINER MULTICUT.

Theorem 7.1. NODE STEINER MULTICUT *can be decided in linear time on trees.*

Proof. Let (G, \mathcal{T}, k) be an instance of NODE STEINER MULTICUT where G is a tree. Let r and the terminal roots of G be defined as before. A crucial observation is that, for any $T \in \mathcal{T}$, if the subtree $G' \supseteq G_T$ of G rooted at $r(G_T)$ contains no terminal roots except $r(G_T)$, then there is an optimal solution that contains $r(G_T)$. To see this, note that any solution must contain at least one node of G' in order to disconnect T . Further, since G' has no terminal roots except $r(G_T)$, the set of terminal sets disconnected when removing $r(G_T)$ is a superset of the set of terminal sets disconnected when removing any node of G' .

Using the above crucial observation, a greedy strategy becomes apparent. First, if $k=0$ and $\mathcal{T} \neq \emptyset$, then return “no”; if $k \geq 0$ and $\mathcal{T} = \emptyset$, then return “yes”. Otherwise, compute the terminal root of each terminal set. Then, find a terminal root c that is deepest in G . Add c to the cut, and recurse on the instance $(G', \mathcal{T}', k-1)$, where G' is G minus the subtree rooted at c , and \mathcal{T}' is \mathcal{T} minus every terminal set that is disconnected by c .

The above algorithm can be implemented in linear time, as follows. First, we find all terminal roots. To this end, index the nodes of G according to a post-order traversal. Then, for each terminal set $T \in \mathcal{T}$, its terminal root is the lowest common ancestor of the lowest and highest indexed terminal in T . We can use the lowest common ancestor data structure of Harel and Tarjan [39] to find all these lowest common ancestors, and thus all terminal roots, in linear time. Second, we perform the greedy algorithm. For each node of G , we maintain a queue with a set of terminals on that node. Now apply an inverse breadth-first search; that is, number the node first visited by a breadth-first search (i.e., r) by $|V(G)|$, and the last one by 1, and then visit the nodes starting with number 1 up to $|V(G)|$. Suppose that we visit a node v . If v is not a terminal root or v is a terminal root for a terminal set that has already been disconnected, then merge the set of terminals on v with the set of terminals on the parent of v . If v is a terminal root for a terminal set that has not been cut yet, then add v to the cut and mark all terminal sets that have a terminal on v as cut. Observe that the instance must at least store (say, in a queue) on each node which terminal sets contain that node as a terminal (or, equivalently, for each terminal set all nodes that it contains). Therefore, the above algorithm indeed runs in linear time in terms of the input size $n+m+\sum_{T \in \mathcal{T}} |T|$. \square

7.2 Parameterized Analysis for Edge Steiner Multicut and Restr. Node Steiner Multicut

The situation for EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT is completely different than for NODE STEINER MULTICUT. Călinescu et al. [12] proved that both EDGE MULTICUT and RESTR. NODE MULTICUT are NP-complete on trees. This implies that EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT are para-NP-complete on trees for parameter p . In this subsection, we explore all possible other parameterizations for these problems on trees.

7.2.1 Parameter k

We briefly recall the definition of the HITTING SET problem, which is known to be W[2]-complete [25]. Recall that HITTING SET is the problem of, given a universe U , a family \mathcal{F} of subsets of U , and an integer k , to decide whether there is a set $H \subseteq U$ with $|H| \leq k$ such that $F \cap H \neq \emptyset$ for each $F \in \mathcal{F}$ (i.e., H is a hitting set of size at most k).

Theorem 7.2. *EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT are $W[2]$ -hard on trees for the parameter k , even if the terminal sets are pairwise disjoint.*

Proof. We reduce from HITTING SET. Let (U, \mathcal{F}, k) be an instance of this problem, and let $\mathcal{F}(u) = \{F \in \mathcal{F} \mid u \in F\}$. We now build the following graph. For each $u \in U$, we build a path P_u on $|\mathcal{F}(u)|$ nodes. Further, let $\sigma_u: \mathcal{F}(u) \rightarrow V(P_u)$ be an arbitrary bijection. We then add a root r , and for each $u \in U$, we connect r to an end of P_u through a new edge e_u . This is the tree G . We now build terminal sets. Let $T_F = \{r\} \cup \{\sigma_u(F) \mid \forall u \in F\}$ for any $F \in \mathcal{F}$, and let $\mathcal{T} = \{T_F \mid F \in \mathcal{F}\}$. The final instance of EDGE STEINER MULTICUT is (G, \mathcal{T}, k) .

Suppose that \mathcal{F} has a hitting set H of size at most k . Let $S = \{e_u \mid u \in H\}$. Note that if some terminal set T_F is not disconnected in $G \setminus S$, then this contradicts that $H \cap F \neq \emptyset$. Further, $|S| \leq k$, and thus S is a solution for (G, \mathcal{T}, k) .

Suppose that (G, \mathcal{T}, k) has a solution, and let S be a solution such that $S \cap \{e_u \mid u \in V(G)\}$ is maximum. Since e_u disconnects at least the terminal sets disconnected by any edge of P_u , $S \setminus \{e_u \mid u \in U\} = \emptyset$. Let $H = \{u \mid e_u \in S\}$. Note that if $H \cap F = \emptyset$ for some $F \in \mathcal{F}$, then T_F is not disconnected by S . Further, $|H| \leq k$, and thus H is solution for (U, \mathcal{F}, k) .

A modification of the reduction yields the result for RESTR. NODE STEINER MULTICUT. Let (U, \mathcal{F}, k) be an instance of HITTING SET. We again build the paths P_u and the terminal sets T_F (which include the root r). However, we do not connect r to an end of P_u directly; instead, for each $u \in U$, we add a node n_u , and connect n_u to both r and an end of P_u . Let G denote the resulting tree, and let $\mathcal{T} = \{T_F \mid F \in \mathcal{F}\}$. The final instance of RESTR. NODE STEINER MULTICUT is (G, \mathcal{T}, k) .

Suppose that \mathcal{F} has a hitting set H of size at most k . Let $S = \{n_u \mid u \in H\}$. Note that if some terminal set T_F is not disconnected in $G \setminus S$, then this contradicts that $H \cap F \neq \emptyset$. Further, $|S| \leq k$, and thus S is a solution for (G, \mathcal{T}, k) .

Suppose that (G, \mathcal{T}, k) has a solution S . Since the only non-terminal nodes are the n_u , we let $H = \{u \mid n_u \in S\}$ and note that $|H| = |S| \leq k$. Further, if $H \cap F = \emptyset$ for some $F \in \mathcal{F}$, then T_F is not disconnected by S . Hence, H is solution for (U, \mathcal{F}, k) . \square

Note that in the constructions of both theorems, we can obtain a slightly different tree structure by replacing the paths P_u with stars. Further, if we do not insist that the terminal sets are pairwise disjoint, then both constructions can be simplified to a star by contracting each path P_u into a single node.

7.2.2 Kernelization for t and $k+t$

We now consider whether EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT have a polynomial kernel on trees for the parameter t . We prove the following stronger result.

Theorem 7.3. *EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT have no polynomial kernel on trees for the parameter $k+t$, even if the terminal sets are pairwise disjoint, unless the polynomial hierarchy collapses to the third level.*

Note that this theorem indeed implies that EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT have no polynomial kernel on trees for the parameter t .

We need the observation that HITTING SET has no polynomial kernel for the parameter solution size plus the number of sets, unless the polynomial hierarchy collapses to the third level. To see this, Dom et al. [23, Theorem 2] prove that SET COVER has no polynomial kernel for the parameter solution size plus the size of the universe, unless the polynomial hierarchy collapses to the third level. Consider an instance (E, \mathcal{S}, k) of SET COVER. For each $S \in \mathcal{S}$, create an element u_S , and let U be the set of all these elements. For each $e \in E$, create a set $F_e = \{u_S \mid e \in S\}$, and let \mathcal{F} be the family of all these sets. Then, (E, \mathcal{S}, k) is a “yes”-instance

of SET COVER if and only if (U, \mathcal{F}, k) is a “yes”-instance of HITTING SET. Note that $|\mathcal{F}| = |E|$. This completes a polynomial parameter transformation. Since both HITTING SET and SET COVER are NP-complete, the observation follows from Bodlaender et al. [9].

Proof of Theorem 7.3. Note that the reductions of Theorem 7.2 and Theorem 7.2 give polynomial parameter transformations for $k+t$. Further, EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT on trees and HITTING SET are NP-complete. The theorem then follows from Bodlaender et al. [9]. \square

7.2.3 Parameter $k+p$

This result generalizes the known algorithms for EDGE MULTICUT and RESTR. NODE STEINER MULTICUT on trees for the parameter k [37, 36].

Theorem 7.4. *EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT are fixed-parameter tractable on trees for parameter $k+p$.*

Proof. Let (G, \mathcal{T}, k) be an instance of EDGE STEINER MULTICUT, where G is a tree. Let r and the terminal roots of G be defined as before. We use the following crucial observation, which is similar to the crucial observation made in Theorem 7.1: for any $T \in \mathcal{T}$, if the subtree $G' \supseteq G_T$ of G rooted at $r(G_T)$ contains no terminal roots except $r(G_T)$, then there is an optimal solution that contains an edge of G_T incident to $r(G_T)$. To see this, note that any solution must contain at least one edge of G_T in order to disconnect T . Further, since G' has no terminal roots except $r(G_T)$, the set of terminal sets disconnected when removing an edge $e \in E(G_T)$ is a subset of the set of terminal sets disconnected when removing the edge incident to $r(G_T)$ on the path in G from e to $r(G_T)$.

Using the above crucial observation, a branching strategy becomes apparent. First, if $k=0$ and $\mathcal{T} \neq \emptyset$, then return “no”; if $k \geq 0$ and $\mathcal{T} = \emptyset$, then return “yes”. Otherwise, compute the terminal roots of each terminal set. Then, find a terminal set T for which $r(G_T)$ is deepest in G . Since $|T| \leq p$, G_T has at most p edges incident on $r(G_T)$. Branch on all such edges e . Then, add e to the cut, and recurse on the instance $(G', \mathcal{T}', k-1)$, where G' is the connected component of $G \setminus e$ that contains $r(G_T)$, and \mathcal{T}' is \mathcal{T} minus every terminal set that is disconnected by e . If any of the branches results in a “yes”, then return “yes”; otherwise, return “no”.

Correctness of the algorithm is immediate from the crucial observation. Further, observe that the algorithm has at most p^k branches.

To prove the theorem for RESTR. NODE STEINER MULTICUT, we essentially combine the crucial observation of Theorem 7.1 with the above crucial observation for the edge version. Let (G, \mathcal{T}, k) be an instance of RESTR. NODE STEINER MULTICUT, where G is a tree. As a pre-processing step, we contract any edge of which both endpoints are a terminal. Since we are not allowed to delete terminals, this is safe. However, if such a contraction makes a terminal set become a singleton set, then we may return “no”. By abuse of notation, we call the resulting tree G and the resulting family of terminal sets \mathcal{T} . Root G at a node r , and define G_T and terminal roots $r(G_T)$ as before.

We follow a branching strategy. First, if $k=0$ and $\mathcal{T} \neq \emptyset$, then return “no”; if $k \geq 0$ and $\mathcal{T} = \emptyset$, then return “yes”. Otherwise, compute the terminal roots of each terminal set. Then, find a terminal set T for which $r(G_T)$ is deepest in G . If $r(G_T)$ is not a terminal node, then following the crucial observation of Theorem 7.1, we add $r(G_T)$ to the cut, and recurse on the instance $(G', \mathcal{T}', k-1)$, where G' is G minus the subtree rooted at $r(G_T)$, and \mathcal{T}' is \mathcal{T} minus every terminal set that is disconnected by $r(G_T)$. If $r(G_T)$ is a terminal node, then none of the children of $r(G_T)$ are a terminal. Since $|T| \leq p$, $r(G_T)$ has at most p neighbors in G_T . Branch on all such neighbors c . Then, add c to the cut, and recurse on the instance $(G', \mathcal{T}', k-1)$, where G'

is G minus the subtree rooted at c , and \mathcal{T}' is \mathcal{T} minus every terminal set that is disconnected by c . If any of the branches results in a “yes”, then return “yes”; otherwise, return “no”.

Correctness of the algorithm is immediate from the crucial observation of Theorem 7.1 and the above crucial observation for the edge version. Further, observe that the algorithm has at most p^k branches. The theorem follows. \square

Since EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT are both fixed-parameter tractable on trees for the parameter $k+p$, it is natural to ask whether these problems admit a polynomial kernel. We answer this question negatively.

Theorem 7.5. *EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT do not admit polynomial kernels on trees for parameter $k+p$, even if the terminal sets are pairwise disjoint, unless the polynomial hierarchy collapses to the third level.*

Proof. It is known that HITTING SET has no polynomial kernel when parameterized by the solution size and the maximum size of any set, unless the polynomial hierarchy collapses to the third level [23, Theorem 6]. Moreover, EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT on trees and HITTING SET are NP-complete. Since the reductions given in Theorem 7.2 and Theorem 7.2 are polynomial parameter transformations from HITTING SET for the parameter solution size plus maximum set size to EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT, respectively, for parameter $k+p$, the theorem follows from Bodlaender et al. [9]. \square

7.2.4 Parameter t

We note that the result of the following theorem is actually dominated by Theorem 5.2. However, since the theorem is much simpler to prove for trees and does not rely on MSOL, we give its proof for completeness.

Theorem 7.6. *EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT are fixed-parameter tractable on trees for parameter t .*

Proof. Let (G, \mathcal{T}, k) be an instance of EDGE STEINER MULTICUT, where G is a tree. For each $T \in \mathcal{T}$, let G_T again denote the subtree of G induced by the terminals in T . Let $\mathcal{F} = \{E(G_T) \mid T \in \mathcal{T}\}$. Observe that $(E(G), \mathcal{F}, k)$ is a “yes”-instance of HITTING SET if and only if (G, \mathcal{T}, k) is a “yes”-instance of EDGE STEINER MULTICUT. Recall that HITTING SET is fixed-parameter tractable parameterized by the number of sets [30]. Since $|\mathcal{F}| = t$, the theorem follows.

This algorithm can be easily adapted for RESTR. NODE STEINER MULTICUT: the universe of the HITTING SET instance becomes $V(G)$, and the family of sets becomes $\{V(G_T) - X \mid T \in \mathcal{T}\}$, where X is the set of terminal nodes. \square

To contrast Theorem 7.3, we prove that both EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT admit a polynomial kernel on trees for parameter $t+p$.

Theorem 7.7. *EDGE STEINER MULTICUT admits a polynomial kernel on trees for parameter $t+p$.*

Proof. Let (G, \mathcal{T}, k) be an instance of EDGE STEINER MULTICUT, where G is a tree. We describe two reduction rules.

The first reduction rule concerns leaves. Let v be a leaf of G and suppose that v is not a terminal. Let u be the neighbor of v in G . Then no minimal solution contains the edge (u, v) . Hence, we may remove v without changing the feasibility of the instance. Apply this reduction rule exhaustively, until each leaf of G is a terminal.

The second reduction rule concerns nodes of degree two. Consider any triple of nodes $u, v, w \in V(G)$ such that $N(v) = \{u, w\}$ and v is not a terminal. Then no minimal solution selects both (u, v) and (v, w) . Hence, we may remove v and connect u, w without changing the feasibility of the instance. Apply this reduction rule exhaustively, until each internal node of G is either a terminal or of degree at least three.

Observe that after applying the first reduction rule exhaustively, G has $O(tp)$ leafs. After applying the second reduction rule, G has $O(tp)$ internal nodes that are terminals. Further, any non-terminal internal node has degree at least three. Since the number of non-terminal internal nodes cannot exceed the number of leafs, G has $O(tp)$ nodes in total. \square

Theorem 7.8. *RESTR. NODE STEINER MULTICUT has a polynomial kernel on trees for parameter $t+p$.*

Proof. Let (G, \mathcal{T}, k) be an instance of *RESTR. NODE STEINER MULTICUT*, where G is a tree. We describe two reduction rules. The first reduction rule is the same as in Theorem 7.7: remove any leaf u that is not a terminal. This rule is safe, as no minimal solution contains u . Apply this reduction rule exhaustively, until each leaf of G is a terminal.

The second reduction rule is a modified version of the second reduction rule of Theorem 7.7. Consider any triple of distinct nodes $(u, v, w) \in V(G)^3$ such that $N(v) = \{u, w\}$ and u and v are not a terminal. Then no minimal solution selects both u and v . Hence, we may remove v and connect u, w without changing the feasibility of the instance. Apply this reduction rule exhaustively, until each internal node of G either is a terminal, or is of degree at least three, or is of degree two and has only terminal neighbors.

Observe that after applying the first reduction rule exhaustively, G has $O(tp)$ leafs. After applying the second reduction rule, G has $O(tp)$ internal nodes that are terminals. Further, any non-terminal internal node of degree two has a terminal as child, and thus their number is bounded by $O(tp)$ as well. Finally, the number of non-terminal internal nodes of degree at least three cannot exceed the number of leafs. Hence, G has $O(tp)$ nodes in total. \square

8 Discussion

We provided a comprehensive computational complexity analysis of the *STEINER MULTICUT* problem with respect to fundamental parameters, culminating in either a fixed-parameter algorithm or a $W[1]$ -hardness result for *every* combination of parameters. This way, we generalize known tractability results for special cases of *STEINER MULTICUT*, and chart the boundary of tractability for other cases. See Table 1 for a complete overview.

We leave several interesting questions for future research. A first possible extension of our work is to consider generalizations of *STEINER MULTICUT*, for example to *REQUIREMENT MULTICUT*, where each terminal set T_i must be cut into $r_i \geq 2$ components. The approximability of this problem has been considered in several papers [51, 50, 38]. From a parameterized viewpoint, we note that all hardness results of this paper carry over to *REQUIREMENT MULTICUT*. Yet, it is an intriguing open question which of our fixed-parameter algorithms generalize to *REQUIREMENT CUT*.

A second possible extension is to consider directed graphs. Already *MULTICUT* is $W[1]$ -hard in this case [49] for parameter cut size k , even on acyclic directed graphs [43]. On the other hand, *MULTICUT* is fixed-parameter tractable for the parameter $k+t$ in directed acyclic graphs [43]. It would be interesting whether this result generalizes to *STEINER MULTICUT*.

A third possible extension is to investigate which problems admit polynomial kernels. While we have resolved many kernelization questions in this paper, several open problems remain (see Table 2), in particular whether there is a polynomial kernel for the parameters $k+t+p$ on general

constants	params	poly kernel?
$\text{tw}(G)=1$	t,p	✓ (Thm. 7.7, 7.8)
$\text{tw}(G)=1$	k,p	X (Thm. 7.5)
$\text{tw}(G)=1$	k,t	X (Thm. 7.3)
$\text{tw}(G)=1, p \geq 3$	k	open
$\text{tw}(G)=1, p=2$	k	✓ [11, 15]
$p=2, t$	k	✓ [44]
$p=2$	k	X [19]
	k, t, p	open
$p=2$	k, t	open
$p=3, t=3$	k	open
t	k	open

Table 2: A summary of results about the kernelization of EDGE STEINER MULTICUT; the top three lines also hold for RESTR. NODE STEINER MULTICUT. We note that the kernelization complexity of EDGE STEINER MULTICUT and RESTR. NODE STEINER MULTICUT are completely characterized for trees, except that the case that $p \geq 3$ is a constant and k is a parameter is still open. We have listed all minimal and maximal open questions (among all cases with $\text{tw}(G)=1$ or unbounded).

graphs. Answers in this research direction might shed new light on some long-standing open questions [18] on the existence of polynomial kernels for MULTICUT for parameter $k+t$ (currently, only a kernel of size $k^{O(\sqrt{t})}$ is known [44], and there is no kernel of size polynomial in k only [19]).

Acknowledgements. We thank Magnus Wahlström for an insight that helped in proving Lemma 5.1.

References

- [1] A. Agrawal, P. Klein, and R. Ravi. When trees collide: an approximation algorithm for the generalized Steiner problem on networks. *SIAM J. Comput.*, 24(3):440–456, 1995.
- [2] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *J. Algorithms*, 12(2):308 – 340, 1991.
- [3] A. Avidor and M. Langberg. The multi-multiway cut problem. *Theoret. Comput. Sci.*, 377(1-3):35–42, 2007.
- [4] H. Bodlaender and R. Möhring. The pathwidth and treewidth of cographs. *SIAM Journal on Discrete Mathematics*, 6(2):181–188, 1993.
- [5] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.
- [6] H. L. Bodlaender. A partial k -arboretum of graphs of bounded treewidth. *Theoret. Comput. Sci.*, 209:1–45, 209.
- [7] H. L. Bodlaender, R. G. Downey, M. R. Fellows, and D. Hermelin. On problems without polynomial kernels. *J. Comput. Syst. Sci.*, 75(8):423–434, 2009.

- [8] H. L. Bodlaender, P. G. Drange, M. S. Dregi, F. V. Fomin, D. Lokshtanov, and M. Pilipczuk. An $o(c^k n)$ 5-approximation algorithm for treewidth. In *Proc. FOCS'13*, pages 499–508, 2013.
- [9] H. L. Bodlaender, S. Thomassé, and A. Yeo. Kernel bounds for disjoint cycles and disjoint paths. *Theoret. Comput. Sci.*, 412(35):4570–4578, 2011.
- [10] N. Bousquet, J. Daligault, and S. Thomassé. Multicut is FPT. In *Proc. STOC 2011*, pages 459–468, 2011.
- [11] N. Bousquet, J. Daligault, S. Thomassé, and A. Yeo. A polynomial kernel for multicut in trees. In *Proc. STACS 2009*, volume 3 of *Leibniz Int. Proc. Inform.*, pages 183–194. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2009.
- [12] G. Călinescu, C. G. Fernandes, and B. Reed. Multicuts in unweighted graphs and digraphs with bounded degree and bounded tree-width. *J. Algorithms*, 48(2):333–359, 2003.
- [13] Y. Cao, J. Chen, and J.-H. Fan. An $O^*(1.84^k)$ parameterized algorithm for the multiterminal cut problem. *Information Processing Letters*, 114(4):167–173, 2014.
- [14] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Comput. Complexity*, 15(2):94–114, 2006.
- [15] J. Chen, J.-H. Fan, I. Kanj, Y. Liu, and F. Zhang. Multicut in trees viewed through the eyes of vertex cover. *J. Comput. System Sci.*, 78(5):1637–1650, 2012.
- [16] J. Chen, Y. Liu, and S. Lu. An improved parameterized algorithm for the minimum node multiway cut problem. *Algorithmica*, 55(1):1–13, 2009.
- [17] R. Chitnis, M. Cygan, M. Hajiaghayi, M. Pilipczuk, and M. Pilipczuk. Designing FPT algorithms for cut problems using randomized contractions. In *Proc. FOCS 2012*, pages 460–469, 2012.
- [18] M. Cygan, Ł. Kowalik, and M. Pilipczuk. Open problems from workshop on kernels, April 2013. <http://worker2013.mimuw.edu.pl/slides/worker-opl.pdf>.
- [19] M. Cygan, S. Kratsch, M. Pilipczuk, M. Pilipczuk, and M. Wahlström. Clique cover and graph separation: new incompressibility results. *ACM Transactions on Computation Theory (TOCT)*, 6(2), 2014.
- [20] M. Cygan, M. Pilipczuk, M. Pilipczuk, and J. O. Wojtaszczyk. On multiway cut parameterized above lower bounds. *ACM Transactions on Computation Theory (TOCT)*, 5(1), 2013.
- [21] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM J. Comput.*, 23(4):864–894, 1994.
- [22] G. B. Dantzig and D. R. Fulkerson. On the max-flow min-cut theorem of networks. In *Linear inequalities and related systems*, Annals of Mathematics Studies, no. 38, pages 215–221. Princeton University Press, Princeton, N. J., 1956.
- [23] M. Dom, D. Lokshtanov, and S. Saurabh. Incompressibility through colors and IDs. In *Proc. ICALP 2009*, volume 5555 of *Lecture Notes Comput. Sci.*, pages 378–389. Springer, Berlin, 2009.

- [24] R. G. Downey, V. Estivill-Castro, M. Fellows, E. Prieto, and F. A. Rosamond. Cutting up is hard to do: The parameterised complexity of k -cut and related problems. *Electr. Notes Theor. Comput. Sci.*, 78(0):209 – 222, 2003.
- [25] R. G. Downey and M. R. Fellows. *Parameterized complexity*. Monographs in Computer Science. Springer-Verlag, New York, 1999.
- [26] A. Drucker. New limits to classical and quantum instance compression. In *Proc. FOCS 2012*, pages 609–618. IEEE Computer Society, 2012.
- [27] P. Elias, A. Feinstein, and C. Shannon. A note on the maximum flow through a network. *Information Theory, IRE Transactions on*, 2(4):117–119, 1956.
- [28] M. R. Fellows, D. Hermelin, F. Rosamond, and S. Vialette. On the parameterized complexity of multiple-interval graph problems. *Theoret. Comput. Sci.*, 410(1):53–61, 2009.
- [29] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
- [30] F. V. Fomin, D. Kratsch, and G. J. Woeginger. Exact (exponential) algorithms for the dominating set problem. In *Proc. WG 2004*, volume 3353 of *Lecture Notes Comput. Sci.*, pages 245–256. Springer, Berlin, 2004.
- [31] L. R. Ford, Jr. and D. R. Fulkerson. Maximal flow through a network. *Canad. J. Math.*, 8:399–404, 1956.
- [32] L. Fortnow and R. Santhanam. Infeasibility of instance compression and succinct PCPs for NP. *J. Comput. Syst. Sci.*, 77(1):91–106, 2011.
- [33] N. Garg, V. V. Vazirani, and M. Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM J. Comput.*, 25(2):235–251, 1996.
- [34] G. Gottlob and S. T. Lee. A logical approach to multicut problems. *Inform. Process. Lett.*, 103(4):136–141, 2007.
- [35] S. Guillemot. FPT algorithms for path-transversal and cycle-transversal problems. *Discrete Optim.*, 8(1):61–71, 2011.
- [36] J. Guo, F. Hüffner, E. Kenar, R. Niedermeier, and J. Uhlmann. Complexity and exact algorithms for multicut. In *Proc. SOFSEM 2006*, volume 3831 of *Lecture Notes Comput. Sci.*, pages 303–312, 2006.
- [37] J. Guo and R. Niedermeier. Fixed-parameter tractability and data reduction for multicut in trees. *Networks*, 46(3):124–135, 2005.
- [38] A. Gupta, V. Nagarajan, and R. Ravi. An improved approximation algorithm for requirement cut. *Oper. Res. Lett.*, 38(4):322–325, 2010.
- [39] D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984.
- [40] K.-i. Kawarabayashi and M. Thorup. The minimum k -way cut of bounded size is fixed-parameter tractable. In *Proc. FOCS 2011*, pages 160–169. IEEE Computer Soc., 2011.
- [41] P. N. Klein, S. A. Plotkin, S. Rao, and É. Tardos. Approximation algorithms for Steiner and directed multicuts. *J. Algorithms*, 22(2):241–269, 1997.

- [42] J. Könemann, S. Leonardi, G. Schäfer, and S. H. M. van Zwam. A group-strategyproof cost sharing mechanism for the Steiner Forest game. *SIAM J. Comput.*, 37(5):1319–1341, 2008.
- [43] S. Kratsch, M. Pilipczuk, M. Pilipczuk, and M. Wahlström. Fixed-parameter tractability of multicut in directed acyclic graphs. In *Proc. ICALP 2012*, volume 7391 of *Lecture Notes Comput. Sci.*, pages 581–593. Springer, Heidelberg, 2012.
- [44] S. Kratsch and M. Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. In *Proc. FOCS 2012*, pages 450–459, 2012.
- [45] L. C. Lau. An approximate max-Steiner-tree-packing min-Steiner-cut theorem. *Combinatorica*, 27(1):71–90, 2007.
- [46] D. Marx. Parameterized graph separation problems. *Theoret. Comput. Sci.*, 351(3):394–406, 2006.
- [47] D. Marx. Important separators and parameterized algorithms. In *Proc. WG 2011*, volume 6986 of *Lecture Notes Comput. Sci.*, pages 5–10, 2011.
- [48] D. Marx, B. O’Sullivan, and I. Razgon. Finding small separators in linear time via treewidth reduction. *ACM Trans. Algorithms*, 9(4):30:1–30:35, Oct. 2013.
- [49] D. Marx and I. Razgon. Fixed-parameter tractability of multicut parameterized by the size of the cutset. In *Proc. STOC 2011*, pages 469–478, 2011.
- [50] A. Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *Proc. FOCS 2009*, pages 3–12. IEEE Computer Soc., 2009.
- [51] V. Nagarajan and R. Ravi. Approximation algorithms for requirement cut on graphs. *Algorithmica*, 56(2):198–213, 2010.
- [52] M. Xiao. Simple and improved parameterized algorithms for multiterminal cuts. *Theory Comput. Syst.*, 46(4):723–736, 2010.
- [53] M. Yannakakis, P. C. Kanellakis, S. S. Cosmadakis, and C. H. Papadimitriou. Cutting and partitioning a graph after a fixed pattern (extended abstract). In *Proc. ICALP 1983*, volume 154 of *Lecture Notes Comput. Sci.*, pages 712–722, 1983.
- [54] B. Yu and J. Cheriyan. Approximation algorithms for feasible cut and multicut problems. In *Proc. ESA 1995*, volume 979 of *Lecture Notes Comput. Sci.*, pages 394–408. Springer, Berlin, 1995.

A Basic Definitions

In this section, we give formal definitions of several core notions that are used in this paper. First and foremost, a *parameter* of a problem instance I of problem Π is an integer $k(I)$. Usually, this parameter is part of the problem definition, so we commonly write it as just k , and it is often equal to the size of the optimum or desired solution to the problem instance. A problem is *fixed-parameter tractable* if it admits a *fixed-parameter algorithm*, which decides instances I of Π in time $f(k(I)) \cdot |I|^{O(1)}$ -time for some computable function f . The class of fixed-parameter tractable problems is denoted by FPT. As evidence that a problem Π is unlikely to be fixed-parameter tractable, one can show that Π is W[1]- or W[2]-hard. Informally, W[1] is the class of problems that are as hard as CLIQUE parameterized by the maximum clique size, and W[2] is the class of problems that are as hard as HITTING SET parameterized by the minimum hitting set size.

To prove hardness of Π , one can give an *parameterized reduction* from a $W[\cdot]$ -hard problem Π' to Π that maps every instance I' of Π' with parameter k' to an instance I of Π with parameter $k \leq g(k')$ for some computable function g such that I' can be computed in time $f(k') \cdot |I|^{O(1)}$ for some computable function f , and I is a “yes”-instance if and only if I' is a “yes”-instance. In case f and g are polynomials, such a reduction is called a *polynomial parameter transformation*. A second way to prove hardness is by showing that the problem is NP-complete even if the parameter k is a constant; in that case, the parameterized problem is called *para-NP-complete*.

An equivalent notion to fixed-parameter tractability is that of a “kernel”. Given an instance I with parameter k of a problem Π , a *kernelization algorithm* computes in polynomial time an instance I' (a *kernel*) of Π with size $|I'| \leq f(k)$, for some computable function f , such that I is a “yes”-instance if and only if I' is. Any decidable fixed-parameter tractable problem for a parameter k with running time $f(k) \cdot n^{O(1)}$ admits a kernel with the same function f . Conversely, no $W[\cdot]$ -hard problem admits a kernel, unless $FPT = W[\cdot]$. For NP-hard problems, a fixed-parameter algorithm can generally only yield kernels with an superpolynomial function f , unless $P = NP$. However, using different techniques it is sometimes possible to obtain a *polynomial kernel*, for which the function f is a polynomial. For yet other problems, the existence of a polynomial kernel is unlikely [26, 7, 32], for example because the polynomial hierarchy would collapse to the third level. In this paper, we use that for a parameterized problem Π' , for which a polynomial kernel is known to be unlikely, and a polynomial parameter transformation from Π' to a parameterized problem Π , it is also unlikely that Π has a polynomial kernel [9].

A good introduction into parameterized complexity can be found in the books by Downey and Fellows [25] and Flum and Grohe [29].

The final notion that we define is that of “treewidth”. Given an undirected graph G , a *tree decomposition* of G is a tree T and a collection of *bags* $\mathcal{B} = \{B_t \subseteq V(G) \mid t \in V(T)\}$ such that

- for each node $v \in V(G)$, the set $\{t \mid v \in B_t\}$ is a nonempty set that induces a connected subtree of T ;
- for each edge $\{u, v\} \in E(G)$, there is a bag B_t such that $u, v \in B_t$.

The *width* of a tree decomposition (T, \mathcal{B}) is $\max_{B \in \mathcal{B}} \{|B| - 1\}$, and the *treewidth* $\text{tw}(G)$ of G is the minimum width of all tree decompositions of G . We refer to Bodlaender [6] for some basic properties of tree decompositions.

B Tree Diagrams

We provide tree diagrams of Table 1 for each of the three versions of STEINER MULTICUT. The diagrams exhibit both the FPT versus $W[1]$ -hardness/para-NP-completeness dichotomy and the polynomial-time versus NP-completeness dichotomy. To see the latter, remove all branches marked only ‘param’ from each of the diagrams.

In the diagrams, “T.x” refers to Theorem x, “S.x” refers to Section x, “const” means that the preceding variable is taken to be a fixed constant, “param” means that the preceding variable is taken to be a parameter, and “unb” (unbounded) means that we make no assumptions on the preceding variable.




Figure 2: Tree diagram for RESTRICTED NODE STEINER MULTICUT.




Figure 3: Tree diagram for EDGE STEINER MULTICUT.




Figure 4: Tree diagram for NODE STEINER MULTICUT. For the cases marked with a †, one has to realize that an instance with $k \geq t$ is always a “yes”-instance, and thus we may always assume that $k < t$.