

A Parameterized Approximation Scheme for MIN k -CUT

Daniel Lokshтанov* Saket Saurabh† Vaishali Surianarayanan*

Abstract

In the MIN k -CUT problem, input is an edge weighted graph G and an integer k , and the task is to partition the vertex set into k non-empty sets, such that the total weight of the edges with endpoints in different parts is minimized. When k is part of the input, the problem is NP-complete and hard to approximate within any factor less than 2. Recently, the problem has received significant attention from the perspective of parameterized approximation. Gupta *et al.* [SODA 2018] initiated the study of FPT-approximation for the MIN k -CUT problem and gave an 1.9997-approximation algorithm running in time $2^{\mathcal{O}(k^6)}n^{\mathcal{O}(1)}$. Later, the same set of authors [FOCS 2018] designed an $(1 + \epsilon)$ -approximation algorithm that runs in time $(k/\epsilon)^{\mathcal{O}(k)}n^{k+\mathcal{O}(1)}$, and a 1.81-approximation algorithm running in time $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$. More, recently, Kawarabayashi and Lin [SODA 2020] gave a $(5/3 + \epsilon)$ -approximation for MIN k -CUT running in time $2^{\mathcal{O}(k^2 \log k)}n^{\mathcal{O}(1)}$.

In this paper we give a parameterized approximation algorithm with best possible approximation guarantee, and best possible running time dependence on said guarantee (up to Exponential Time Hypothesis (ETH) and constants in the exponent). In particular, for every $\epsilon > 0$, the algorithm obtains a $(1 + \epsilon)$ -approximate solution in time $(k/\epsilon)^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$. The main ingredients of our algorithm are: a simple sparsification procedure, a new polynomial time algorithm for decomposing a graph into highly connected parts, and a new exact algorithm with running time $s^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$ on unweighted (multi-) graphs. Here, s denotes the number of edges in a minimum k -cut. The later two are of independent interest.

*University of California, Santa Barbara, USA. Emails: daniello@ucsb.edu, vaishali@ucsb.edu

†The Institute of Mathematical Sciences, HBNI, Chennai, India. saket@imsc.res.in

1 Introduction

In this article we focus on the MIN k -CUT problem from the perspective of parameterized approximation. Here input is an edge weighted graph G and an integer k , and the task is to partition the vertex set into k non-empty sets, say \tilde{P} , such that the total weight of the edges with endpoints in different parts is minimized. We call such a partition as *optimal k -cut*, or *minimum k -cut* or simply a *k -cut*. For $k = 2$, this is just the classic GLOBAL MIN-CUT problem, which can be solved in polynomial time. In fact, for every fixed k , the problem is known to be polynomial time solvable [GH94]. However, when k is part of the input, the problem is NP-complete [GH94].

Traditionally, the problem has been extensively studied from three perspectives: (a) exact algorithms for small values of k ; (b) combinatorial upper bound on the number of minimum k -cuts; and (c) approximation algorithms. In particular, already in early 90's, Goldschmidt and Hochbaum [GH94] gave the first polynomial-time algorithm for fixed k , with running time $\mathcal{O}(n^{(0.5-o(1))k^2})$. In 1996, Karger and Stein [KS96] obtained their ingenious *contraction algorithm*, a randomized algorithm with running time $\tilde{\mathcal{O}}(n^{2(k-1)})$ ¹. The same algorithm immediately yields a proof that the number of k -cuts of minimum weight in any undirected graph on n vertices is upper bounded by $\tilde{\mathcal{O}}(n^{2(k-1)})$. In 2008, Thorup [Tho08] obtained a deterministic algorithm that essentially matched the running time of Karger and Stein [KS96]. This algorithm is based on a tree-packing theorem which allows to efficiently find a tree that crosses the optimal k -cut at most $2k - 2$ times. Then the algorithm simply enumerates over all possible sets of $2k - 2$ edges of this tree and all ways of merging the components of the tree minus these $2k - 2$ edges into k non-empty groups. From the perspective of approximation algorithms there are several $2(1 - \frac{1}{k})$ -approximation algorithms, that run in time polynomial in n and k [NR01, RS08, SV95]. This approximation ratio cannot be improved assuming the Small Set Expansion Hypothesis (SSE) [Man18].

Thus, as of early 2018, the status was as follows. The direction of polynomial time approximation algorithms was essentially settled, with factor $2(1 - \frac{1}{k})$ approximation algorithms and matching lower bounds. The fastest algorithm for small k had running time $\tilde{\mathcal{O}}(n^{2(k-1)})$, and a $f(k)n^{o(k)}$ lower bound on the running time [DEF⁺03, CFK⁺15] was known under the ETH. It remained an interesting research direction to find an algorithm with an n^{ck} running time and prove that the constant c is optimal under reasonable complexity assumptions. The best upper bound on the number of minimum k -cuts was $\tilde{\mathcal{O}}(n^{2(k-1)})$, while the best lower bound was essentially $\Omega((n/(k-1))^{k-1})$, what you obtain from the complete $(k-1)$ -partite graph. Since early 2018 the problem has seen a remarkable level of activity [CKL⁺18, GLL18a, GLL18b, GLL19, Li19, GLL20, KL20], with a new research direction of parameterized approximation being initiated, and the two remaining established research directions (exact algorithms and combinatorial bounds) essentially being completely resolved.

For exact algorithms for small values of k , Gupta et al. [GLL18a] made the first improvement in two decades, by designing an algorithm with running time $\mathcal{O}(k^{\mathcal{O}(k)}n^{(\frac{2\omega}{3}+o(1))k})$ for graphs with polynomial integer weights. Subsequently, for unweighted graphs, Li [Li19] obtained an algorithm with running time $\mathcal{O}(k^{\mathcal{O}(k)}n^{(1+o(1))k})$. For the original formulation of the problem (with general integer weights) Gupta et al. [GLL19] showed an $\mathcal{O}(n^{(1.98+o(1))k})$ -time algorithm. Their algorithm also implied a new and improved $\mathcal{O}(n^{(1.98+o(1))k})$ upper bound on the number of minimum k -cuts. Finally Gupta et al. [GLL20] fully resolved the problem by showing that for every fixed $k \geq 2$, the Karger-Stein algorithm [KS96] outputs any fixed minimum k -cut with probability at least $\hat{\mathcal{O}}(n^{-k})$, where $\hat{\mathcal{O}}(\cdot)$ hides a $2^{\mathcal{O}(\ln \ln n)^2}$ factor. This immediately gives an extremal bound of $\hat{\mathcal{O}}(n^k)$, on the number of minimum k -cuts in an n -vertex graph and an algorithm for MIN k -CUT in similar running time. Both the extremal bound and the running time of the algorithm are essentially tight (under reasonable assumptions). Indeed the extremal bound

¹ $\tilde{\mathcal{O}}$ hides the poly-logarithmic factor in the running time.

matches known lower bounds up to $\widehat{\mathcal{O}}(1)$ factors, while any further improvement to the exact algorithm would imply an improved algorithm for MAX-WEIGHT k -CLIQUE [AWW14, BT17], which has been conjectured not to exist.

The lower bound of $(2 - \epsilon)$ on polynomial time approximation algorithms [Man18], coupled with the $f(k)n^{\mathcal{O}(k)}$ running time lower bound for exact algorithms [DEF+03, CFK+15] naturally leads to the question of parameterized approximation: how good approximation ratio can we achieve if we allow the algorithm to have running time $f(k)n^{\mathcal{O}(1)}$? In 2018, Gupta et al. [GLL18b] were the first to ask this question, and showed that relaxing the running time requirement from polynomial to $f(k)n^{\mathcal{O}(1)}$ does lead to an improved approximation guarantee, by giving an 1.9997-approximation algorithm for MIN k -CUT running in time $2^{\mathcal{O}(k^6)}n^{\mathcal{O}(1)}$. Subsequently the same set of authors [GLL18a] designed an $(1 + \epsilon)$ -approximation algorithm that runs in time $(k/\epsilon)^{\mathcal{O}(k)}n^{k+\mathcal{O}(1)}$, and a 1.81-approximation algorithm running in time $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$. More recently, Kawarabayashi and Lin [KL20] gave a $(5/3 + \epsilon)$ -approximation for MIN k -CUT running in time $2^{\mathcal{O}(k^2 \log k)}n^{\mathcal{O}(1)}$. In this paper we bring the direction of parameterized approximation to its natural conclusion by giving an algorithm with best possible approximation guarantee, and best possible running time dependence on said guarantee (up to ETH and constants in the exponent). In particular, for every $\epsilon > 0$, the algorithm obtains a $(1 + \epsilon)$ -approximate solution in time $(k/\epsilon)^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$.

Theorem 1.1. *There exists a randomized algorithm that given a graph G , weight function $w : V(G) \rightarrow \mathbb{Q}_{\geq 0}$, integer k , and an $\epsilon > 0$, runs in time $(k/\epsilon)^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$ and outputs a partition \tilde{P} of $V(G)$ into k non-empty parts. With probability at least $1 - \frac{1}{n^{26}}$ the weight of \tilde{P} is at most $(1 + \epsilon)$ times the weight of an optimum k -cut of G . By weight of \tilde{P} , we mean the total weight of the edges with endpoints in different parts.*

Overview of the Algorithm. A preliminary step of our algorithm is to run the 2-approximation algorithm [NR01, RS08, SV95] to get an estimate of the value of $\text{OPT}(G, k)$ (the weight of an optimal k -cut in G). A standard rounding procedure (similar to the well-known Knapsack PTAS [KT06]) reduces the problem in a $(1 + \epsilon)$ -approximation preserving manner to unweighted multi-graphs with at most m^2/ϵ edges. From now on, *throughout this paper*, we will assume that our input graph is an *unweighted multigraph*. As long as the graph has a non-trivial 2-cut (non-trivial means that there is at least one edge crossing the cut) of weight at most $\frac{\epsilon \cdot \text{OPT}(G, k)}{k-1}$, we remove all edges of this 2-cut. This step will be repeated less than $k - 1$ times, since if it is repeated $k - 1$ times we have cut the graph into at least k connected components using at most $\epsilon \cdot \text{OPT}(G, k)$ edges. Thus the procedure stops before this, and at that point we know that (i) we have included at most $\epsilon \cdot \text{OPT}(G, k)$ edges in our solution, and (ii) the remaining graph has no non-trivial 2-cut with at most $\epsilon \cdot \text{OPT}(G, k)$ edges. We now solve the problem independently on each connected component, after guessing how many parts each component is split into (this introduces a 4^k overhead in the running time). Thus from now on we will assume that our input graph is connected and has no non-trivial 2-cut of weight at most $\frac{\epsilon \cdot \text{OPT}(G, k)}{k-1}$.

Having ensured that every non-trivial 2-cut has weight at least $\epsilon \cdot \frac{\text{OPT}(G, k)}{k-1}$ we proceed as follows. We set a probability $p = \frac{100 \log n}{\text{OPT}(G, 2) \cdot \epsilon^2}$. Then for every edge of G we flip a biased coin and keep the edge with probability p and delete it with probability $1 - p$. Call the resulting subgraph G' . A relatively simple probability computation (very similar to that of randomized cut sparsifiers [BK15, Kar94]) combining Chernoff bounds with Karger and Stein's [KS96] bound on the number of minimum 2-cuts shows that with high probability it holds that for *every* partition \tilde{P} of $V(G) = V(G')$ into k parts, the number of edges of G' crossing the partition \tilde{P} is within a $1 \pm \epsilon$ factor of p times the number of edges crossing it in G . Thus, for purposes of $(1 + \epsilon)$ -approximation one may just as well find a $(1 + \epsilon)$ -approximate solution in G' . However, in G'

we know that

$$\text{OPT}(G', k) \simeq p \cdot \text{OPT}(G, k) = \frac{100 \log n}{\text{OPT}(G, 2) \cdot \epsilon^2} \cdot \text{OPT}(G, k) \leq \frac{k \cdot 100 \log n}{\epsilon^3}.$$

In the last transition we used the assumption that $\text{OPT}(G, 2) \geq \epsilon \cdot \text{OPT}(G, k)/k - 1$. This is a pretty small number of cut edges, which naturally leads to the idea of trying to apply parameterized algorithms with parameter s , the number of edges in the optimal k -cut.

The MIN k -CUT problem is also quite well studied with $s = \text{OPT}(G, k)$ as a parameter. A brute force algorithm trying all edge sets of size s solves the problem in time $\mathcal{O}(n^{2s+\mathcal{O}(1)})$. Marx [Mar06] posed as an open problem in 2004 whether MIN k -CUT is *fixed parameter tractable* when parameterized by s , that is whether the problem admits an algorithm with running time $f(s)n^{\mathcal{O}(1)}$ for any function f . Kawarabayashi and Thorup [KT11] resolved this problem in the affirmative, and obtained an algorithm with running time $s^{s^{\mathcal{O}(s)}} n^2$. Chitnis et al. [CCH⁺16] gave an algorithm with running time $s^{\mathcal{O}(s^2)} n^{\mathcal{O}(1)}$, improving the dependence on s from double exponential to single exponential. Finally, Cygan et al. [CKL⁺18] showed that the problem is solvable in time $s^{\mathcal{O}(s)} n^{\mathcal{O}(1)}$. Unfortunately, applying the algorithm of Cygan et al. [CKL⁺18] directly on G' yields an algorithm with running time $(\frac{k \log n}{\epsilon^3})^{\mathcal{O}(\frac{k \log n}{\epsilon^3})}$ which grows super-polynomially with n , even for constant ϵ and k . Even if one obtains a parameterized algorithm for MIN k -CUT with running time $2^{\mathcal{O}(s)} n^{\mathcal{O}(1)}$ (which is an interesting open problem in itself), this would only lead to a $2^{\mathcal{O}(\frac{k \log n}{\epsilon^3})} = n^{\mathcal{O}(\frac{k}{\epsilon^3})}$ time algorithm, which is slower than the classic exact $n^{2(k-1)}$ time algorithm of Karger and Stein [KS96]. Thus, at a glance, the parameterized approach seems to come close, but hit a dead end. The key insight of our algorithm is that even though s is “small”, that k is much smaller, and that the hard instances for parameter s seem to have a value of k close to s . This leads to the natural problem of whether it is possible to design a parameterized algorithm with parameters k and s that substantially outperforms the algorithm of Cygan et al. [CKL⁺18] when $k \ll s$. Our main technical contribution is precisely such an algorithm. We state this algorithm as a separate theorem.

Theorem 1.2. *There exists an algorithm that given an unweighted multigraph G and integers k and s , determines whether G has a k -cut of weight at most s in times $s^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$.*

The algorithm of Theorem 1.2 can easily be extended using self reduction to also produce a partition \tilde{P} of G with weight at most s , if it exists. The algorithm of Theorem 1.2 is based on the $s^{\mathcal{O}(s)} n^{\mathcal{O}(1)}$ time algorithm of Cygan et al. [CKL⁺18].

Just as the algorithm of Cygan et al. [CKL⁺18], we compute a tree decomposition of the input graph G , such that the adhesions (cuts) of the decomposition are small, and the bags (the building blocks of the decomposition) are highly edge-connected. Unfortunately we are not able to use their decomposition theorem as a black box, because their running time to compute the decomposition is already $s^{\mathcal{O}(s)} n^{\mathcal{O}(1)}$. We make our own decomposition theorem with weaker properties than the one in [CKL⁺18], but computable in polynomial time. The construction of the decomposition theorem follows the template of Cygan et al. [CKL⁺18], replacing their most computationally intensive step with a polynomial time approximation algorithm.

The most technically challenging part of our algorithm is how to compute an optimal k -cut in time $s^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$, when the decomposition into highly connected pieces is provided as input. The only known way how to exploit such a decomposition is using dynamic programming. However, the dynamic programming seems to require at least $2^{\Omega(s)}$ states, even for the stronger decomposition of Cygan et al. [CKL⁺18], let alone for our weaker structural decomposition theorem. Here, the fact that k is much smaller than s comes to rescue. We prove that the dynamic programming algorithm can be “trimmed” to only populate $s^{\mathcal{O}(k)}$ states of the dynamic programming table. This trimming procedure is done by inspecting how an optimal k -cut can

interact with a tree obtained from Throup’s [Tho08] tree-packing theorem. A substantial amount of technical weight lifting is required to show that the trimmed dynamic programming table for a bag of a tree decomposition can be computed efficiently from the trimmed dynamic programming tables of its children.

The algorithm of Theorem 1.2 immediately completes the proof of our $(1 + \epsilon)$ -approximation algorithm (Theorem 1.1). Specifically, applying this algorithm to the graph G' obtained by our sparsification procedure yields a $(1 + \epsilon)$ -approximation algorithm with running time

$$\begin{aligned} \left(\frac{k \cdot 100 \log n}{\epsilon^3}\right)^{\mathcal{O}(k)} n^{\mathcal{O}(1)} &= (k/\epsilon)^{\mathcal{O}(k)} (\log n)^{\mathcal{O}(k)} n^{\mathcal{O}(1)} \leq (k/\epsilon)^{\mathcal{O}(k)} (k^{\mathcal{O}(k)} + n) n^{\mathcal{O}(1)} \\ &\leq 2^{\mathcal{O}(k \log(\frac{k}{\epsilon}))} n^{\mathcal{O}(1)}. \end{aligned}$$

Setting $\epsilon = \frac{1}{n^{\mathcal{O}(1)}}$, shows that a $2^{\mathcal{O}(k \log(\frac{k}{\epsilon}))} n^{\mathcal{O}(1)}$ time algorithm would yield an exact algorithm with running time $n^{\mathcal{O}(k)}$, violating ETH.

Guide to the paper. We start by giving the notations and preliminaries that we use throughout the paper in Section 2. This section is best used as a reference, rather than being read linearly. In Section 3 we prove the $(1 + \epsilon)$ -approximation preserving reduction to instances with optimum k -cuts of size $\mathcal{O}(\frac{k \log n}{\epsilon^3})$. Since the proofs here are so similar to those for graph sparsifiers, a reader eager to get to the key ideas of the paper should look at the statement and proof of Lemma 3.1, as well as the statement of Lemma 3.2, and skip further. In Section 4, we state and prove Theorem 4.1, the polynomial time edge-unbreakable decomposition theorem. In the proof of Theorem 4.1, Subsection 4.2 contains a new approximation algorithm for finding small balanced edge-separators, while Subsections 4.1 and 4.3 are almost word-by-word identical to proofs of Cygan et al [CKL+18] and may be skipped. Section 5 contains our main technical thrust: a new exact algorithm for unweighted multi graphs with running time $s^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$, the algorithm computes and uses a decomposition such as the one provided by Theorem 4.1. Finally, in Section 6, we combine all the results obtained in the previous sections and complete the proof of our main result (Theorem 1.1). We conclude the paper with some interesting open problems in Section 7.

2 Notation and Preliminaries

In this section we give notations, and definitions that we use throughout the paper. Unless specified we will be using all general graph terminologies from the book of Diestel [Die12].

2.1 Graph and Set Theoretic Definitions and Notations

Given a graph G , we use $V(G)$ and $E(G)$ to denote the set of vertices and edges, respectively. We use $\text{cc}(G)$ to denote the number of connected components of G . In this paper our graph could be a multigraph, that is, we allow parallel edges between vertices, even though we might start with a simple graph. Given a subset X of vertices in $V(G)$, we denote by $\delta(X)$ the number of edges in G that have exactly one end point in X . By the term *smoothing* a vertex v of degree 2 in a graph G , we mean the operation of removing v from the graph and adding an edge between two of its neighbors.

For a set U , an ℓ -partition of U is a family $\tilde{P} = \{P_1, P_2, \dots, P_\ell\}$ of non-empty, pairwise disjoint sets whose union $\bigcup_{i=1}^{\ell} P_i$ is equal to U . A *partition* of U is an ℓ -partition for some positive integer ℓ . We refer to P_1, \dots, P_ℓ as the *parts* of \tilde{P} and refer to ℓ as the size of the partition and denote it by $|\tilde{P}|$. An edge uv with endpoints in U *crosses* \tilde{P} if u and v are in different parts of \tilde{P} . The number of times an edge set X crosses \tilde{P} is defined as the number of




Figure 1: (a) Example of a partition \tilde{P}' that is a refinement of a partition \tilde{P} . Partition \tilde{P}' is depicted in red and partition \tilde{P} is depicted in blue. (b) Example of a tree T that crosses a partition \tilde{P} at most $2k - 2$ times. The red edges are the edges from T that cross the partition and the green edges are the remaining edges of T .

edges in X that cross P and denoted by $\delta_X(\tilde{P})$. That is, $X' \subseteq X$ is the set containing all edges in X that crosses \tilde{P} and $\delta_X(\tilde{P}) = |X'|$. The *projection* of a partition \tilde{P} onto a subset S of U is the partition $\tilde{P}' = \{P_1 \cap S, P_2 \cap S, \dots, P_\ell \cap S\}$ with the empty sets removed. A partition \tilde{P}' is a *refinement* of \tilde{P} if every part of \tilde{P}' is a subset of some part of \tilde{P} , see Figure 1(a). We say that a partition \tilde{P} is *refined* by \tilde{P}' if \tilde{P}' is a refinement of \tilde{P} . The *combining* of a set of parts \mathcal{P} in \tilde{P} is the operation of removing all the parts in \mathcal{P} from \tilde{P} and adding their union $\bigcup_{P \in \mathcal{P}} P$ as a part in \tilde{P} .

A k -cut \tilde{S} of a graph G is a k -partition of G . Let G be an edge weighted graph G , then the *weight of the k -cut*, denoted by $w(G, \tilde{S})$ is simply the sum of the weights of edges with endpoints in different parts of \tilde{S} . When G is an unweighted multigraph, then the weight of \tilde{S} is the number of edges with endpoints in different parts of \tilde{S} (if there are q edges between a pair of vertices then it adds q to the sum). We denote this also by $w(G, \tilde{S})$. In the cases where the graph G is clear from the context, we just use $w(\tilde{S})$ instead of $w(G, \tilde{S})$. An *optimal k -cut* or a *minimum k -cut* of G is a k -cut of G that has the *minimum weight* among all k -cuts of G . We denote the weight of an optimal k -cut by $\text{OPT}(G, k)$. A *non-trivial cut* means that there is at least one edge crossing the cut. We remark that in some parts of the paper, we refer to a k -cut of a graph G as a k -partition of $V(G)$ but it will be clear from the context.

For a graph G , an *edge cut* is a pair $A, B \subseteq V(G)$ such that $A \cup B = V(G)$ and $A \cap B = \emptyset$. The order of an edge cut (A, B) is $|E(A, B)|$, that is, the number of edges with one endpoint in A and the other in B . Observe that edge cut and 2-cut are the same.

Definition 2.1 (unbreakability). *A set $X \subseteq V(G)$ is (q, s) -edge-unbreakable if every edge cut (A, B) of order at most s satisfies $|A \cap X| \leq q$ or $|B \cap X| \leq q$.*

A *separation* is a pair $A, B \subseteq V(G)$ such that $A \cup B = V(G)$ and $E(A \setminus B, B \setminus A) = \emptyset$. The order of a separation (A, B) is $|A \cap B|$.

2.2 T -tree and Treewidth

An important notion that underlies our algorithm is a notion of T -trees introduced in the seminal work of Thorup [Tho08]. For an instance (G, k) of MIN k -CUT, a T -tree of G is a spanning tree

of G such that there exists an optimal k -cut S^* of G such that T crosses S^* at most $2k - 2$ times. For an illustration of T -tree we refer to Figure 1(b).

Another important notion that we need is of tree decomposition where bags are “highly connected”. Towards this we first define tree decomposition, treewidth and associated notions and notations that we make use of. For a rooted tree τ and vertex $v \in V(\tau)$ we denote by τ_v the subtree of τ rooted at v . We refer to the vertices of τ as nodes.

A *tree decomposition* of a graph G is a pair (χ, τ) where χ is a rooted tree and τ is a function from $V(\tau)$ to $2^{V(G)}$ such that the following three conditions hold.

- (T1) $\bigcup_{t \in V(\tau)} \chi(t) = V(G)$;
- (T2) For every $uv \in E(G)$, there exists a node $t \in \tau$ such that $\chi(t)$ contains both u and v ; and
- (T3) For every vertex $u \in V(G)$, the set $T_u = \{t \in V(\tau) : u \in \chi(t)\}$, i.e., the set of nodes whose corresponding bags contain u , induces a connected subtree of τ .

For every $t \in V(\tau)$ a set $\chi(t) \subseteq V(G)$, is called a *bag*. For a tree decomposition (τ, χ) fix an edge $e = tt' \in E(\tau)$. The deletion of e from τ splits τ into two trees τ_1 and τ_2 , and naturally induces a separation (X_1, X_2) in G with $X_i := \bigcup_{t \in V(\tau_i)} \chi(t)$, which we henceforth call *the separation associated with e* . The set $A_{\tau, \chi}(e) := X_1 \cap X_2 = \chi(t) \cap \chi(t')$ is called the *adhesion* of e . We suppress the subscript if the decomposition is clear from the context.

For $s, t \in V(\tau)$ we say that s is a *descendant* of t or that t is an *ancestor* of s if t lies on the unique path from s to the root; note that a node is both an ancestor and a descendant of itself. For a node t that is not a root of τ , by $A_{\tau, \chi}(t)$ we mean the adhesion $A_{\tau, \chi}(e)$ for the edge e connecting t with its parent in τ . We extend this notation to $A_{\tau, \chi}(r) = \emptyset$ for the root r . Again, we omit the subscript if the decomposition is clear from the context. For brevity, for $t \in V(\tau)$, we use A_t to denote $A_{\tau, \chi}(t)$.

We define the following functions for convenience:

$$\begin{aligned} \gamma(t) &= \bigcup_{q \text{ descendant of } t} \chi(q), \\ \alpha(t) &= \gamma(t) \setminus A_t, \\ G_t &= G[\gamma(t)] \end{aligned}$$

We say that a rooted tree decomposition (τ, χ) of G is *compact* if for every node $t \in V(\tau)$ for which $A_t \neq \emptyset$ we have that $G[\alpha(t)]$ is connected and $N_G(\alpha(t)) = A_t$. We can extend the function χ to subsets of $V(\tau)$ in the natural way: for a subset $X \subseteq V(\tau)$, $\chi(X) = \bigcup_{x \in X} \chi(x)$. By *children*(t), we denote the set of children of t in τ . For each subset X of nodes in $V(\tau)$, we define $G_X = G[\bigcup_{t \in X} \chi(V(\tau_t))]$.

2.3 Splitters and Chernoff Bound

Let $[n]$ denote the set of integers $\{1, \dots, n\}$. We start by defining the notion of splitters.

Definition 2.2 ([NSS95]). *An (n, k, ℓ) splitter \mathcal{F} is a family of functions from $[n] \rightarrow [\ell]$ such that for all $S \subseteq [n]$, $|S| = k$, there exists a function $f \in \mathcal{F}$ that splits S evenly. That is, for all $1 \leq j, j' \leq \ell$, $|f^{-1}(j) \cap S|$ and $|f'^{-1}(j) \cap S|$ differ by at most one.*

We will need following algorithm to compute splitters with desired parameters.

Theorem 2.1 ([NSS95]). *For all $n, k \geq 1$ one can construct an (n, k, k^2) splitter family of size $k^{\mathcal{O}(1)} \log n$ in time $k^{\mathcal{O}(1)} n \log n$.*

The next lemma is a simple application of Theorem 2.1 and is used as a subroutine in our new algorithm for MIN k -CUT.

Lemma 2.1. *There exists an algorithm that takes as input a set S , two positive integers s_1 and s_2 that are less than $|S|$, and outputs a family \mathcal{S} of subsets of S having size $\mathcal{O}((s_1 + s_2)^{\mathcal{O}(s_1)} \log |S|)$ such that for any two disjoint subsets X_1 and X_2 of S of size at most s_1 and s_2 , \mathcal{S} contains a subset X that satisfies $X_1 \subseteq X$ and $X_2 \cap X = \emptyset$ in time $\mathcal{O}((s_1 + s_2)^{\mathcal{O}(s_1)} |S|^{\mathcal{O}(1)})$.*

Proof. The algorithm carries out the following steps. Order the elements in S arbitrarily, let v_j denote the j^{th} element in this ordering. Initialize $\mathcal{S} = \emptyset$. For each pair s'_1 and s'_2 such that $s'_1 \leq s_1$ and $s'_2 \leq s_2$ repeat the following procedure.

- (i) Let $s' = s'_1 + s'_2$, construct a $(|S|, s', s'^2)$ splitter \mathcal{F} .
- (ii) Construct a set \mathcal{L} of subsets of elements of S by the following procedure. For each function f in \mathcal{F} and for each set $X \subseteq [s'^2]$ having size s'_1 construct a subset C of S that contains an element $s_j \in S$ if $f(j) \in X$, that is $C = \{s_j : f(j) \in X\}$ and add it to \mathcal{L} .
- (iii) $\mathcal{S} = \mathcal{S} \cup \mathcal{L}$. Finally output \mathcal{S}

Consider two disjoint subsets X_1, X_2 of S of size at most s_1 and s_2 and the splitter \mathcal{F} constructed for the values of $s'_1 = |X_1|$ and $s'_2 = |X_2|$. By Definition 2.2, for all $X' \subseteq [|S|]$ of size $s' = s'_1 + s'_2$, there exists a function $f \in \mathcal{F}_1$ that maps each $x \in X'$ to a distinct integer in $[s'^2]$. Thus, there is a function $f \in \mathcal{F}$, that maps each j to a distinct integer in $[s'^2]$ for each $v_j \in X_1 \cup X_2$. Since the size of X_1 is s'_1 , we can infer that there exists a function $f \in \mathcal{F}$ and a set $X \subseteq [s'^2]$ of size s'_1 such that $X = \{f(j) : v_j \in X_1\}$. Therefore, by the construction of \mathcal{L} , there exists a set in C in \mathcal{L} and hence in \mathcal{S} such that $X_1 \subseteq C$ and $X_2 \cap C = \emptyset$.

For each s'_1 and s'_2 such that $s'_1 \leq s_1$ and $s'_2 \leq s_2$, it follows from Theorem 2.1 that the size of \mathcal{F} is $s'^{\mathcal{O}(1)} \log |S|$. From the construction of \mathcal{L} it is easy to see that its size is $|\mathcal{F}| \binom{s'^2}{s'_1}$.

Thus, \mathcal{S} is of size $(s_1 s_2)(s_1 + s_2)^{\mathcal{O}(1)} \log |S| \binom{(s_1 + s_2)^2}{s_1} = \mathcal{O}((s_1 + s_2)^{\mathcal{O}(s_1)} \log |S|)$ and can be computed in time $\mathcal{O}((s_1 + s_2)^{\mathcal{O}(s_1)} |S|^{\mathcal{O}(1)})$. This concludes the proof. \square

To prove the correctness of our sparsification procedure we need standard Chernoff bounds to bound the tail probability.

Lemma 2.2 (Chernoff Bound, [MU05]). *Let X_1, \dots, X_n be independent Poisson trials such that $\Pr[X_i] = p_i$. Let $X = \sum_{i=1}^n X_i$ and $\mu = E[X]$. For $0 < \delta < 1$.*

$$\Pr[|X - \mu| \geq \delta \mu] \leq 2 \cdot e^{-\mu \delta^2 / 3}.$$

3 Reduction to Instances with Logarithmic Solution Size

In this section we give a polynomial time algorithm that given an unweighted graph G , an integer k and an $\epsilon > 0$, outputs a subgraph G' of G , such that $V(G) = V(G')$ and a real number r such that with high probability it holds that for every partition \tilde{P} of $V(G) = V(G')$ into k parts, we have that $w(G, \tilde{P}) \simeq r \cdot w(G', \tilde{P})$. Furthermore, every non-trivial 2-cut of G_1 has weight at least $\frac{\epsilon \cdot \text{OPT}(G, k)}{k-1}$. Recall, that a non-trivial cut means that there is at least one edge crossing the cut. We start by a greedy procedure that will ensure that weight of each 2-cut of the sparsifier is sufficiently large.

Lemma 3.1. *There exists an algorithm that takes as input an integer $k > 0$, an unweighted graph G such that $\text{cc}(G) < k$, and a real number $0 < \epsilon \leq 1$ and outputs a subgraph G_1 of G with $V(G_1) = V(G)$, such that $|E(G)| - |E(G_1)| \leq 2\epsilon \cdot \text{OPT}(G, k)$. Further, if $\text{cc}(G_1) < k$, then each non-trivial 2-cut of G_1 has weight at least $\frac{\epsilon \cdot \text{OPT}(G, k)}{k-1}$.*

Proof. The graph G_1 is essentially obtained by a greedy algorithm that removes all 2-cuts of small weights. In particular, we obtain the graph G_1 by the following procedure. Obtain a 2-approximate k -cut of G using any of the known approximation algorithm [NR01, RS08, SV95], let w_a be the weight of this cut. Initialize $G' = G$ and as long as the graph has a non-trivial 2-cut of weight at most $\frac{\epsilon \cdot w_a}{k-1}$ and $\text{cc}(G') < k$, we remove all edges of this 2-cut from G' . Return the resulting graph G' as G_1 .

Observe that the number of times the above procedure will be repeated is at most $k - 1$ times, since if it is repeated $k - 1$ times we have cut the graph into at least k connected components. Furthermore, in each iteration in which $E(G')$ is modified, the size of the set of edges E' removed from G' is at most $\frac{\epsilon \cdot w_a}{k-1}$. Thus $|E(G)| - |E(G_1)|$, the total number of edges removed from G to obtain G_1 is less than or equal to $\epsilon \cdot w_a$. Since $w_a \leq 2 \cdot \text{OPT}(G, k)$, it follows that $|E(G)| - |E(G_1)| \leq 2\epsilon \cdot \text{OPT}(G, k)$. Also, if $\text{cc}(G_1) < k$, then the weight of a minimum non-trivial 2-cut of G_1 is more than $\frac{\epsilon \cdot w_a}{k-1}$. Since $w_a \geq \text{OPT}(G, k)$, it follows that $\frac{\epsilon \cdot w_a}{k-1} \geq \frac{\epsilon \cdot \text{OPT}(G, k)}{k-1}$. Thus, if G_1 has less than k connected components, each non-trivial 2-cut of G_1 has weight more than $\frac{\epsilon \cdot \text{OPT}(G, k)}{k-1}$. Since, we can find a non-trivial minimum 2-cut in polynomial time, we have that the algorithm runs in polynomial time. This concludes the proof. \square

Next given an unweighted graph G , an integer k and an $\epsilon > 0$, we give a sparsification procedure that outputs a subgraph G' of G , and a number p such that with high probability it holds that for every partition \tilde{P} of $V(G) = V(G')$ into k parts, the number of edges of G' crossing the partition \tilde{P} is within a $1 \pm \epsilon$ factor of p times the number of edges crossing it in G . This procedure is very similar to that of known randomized cut sparsifiers [BK15, Kar94]).

Lemma 3.2. *There exists a polynomial time algorithm that takes as input an integer $k > 0$, an unweighted graph G with $\text{cc}(G) < k$, and a real number $\frac{1}{n} < \epsilon \leq 1$ and outputs a subgraph G_2 with $V(G_2) = V(G)$, and a real number r such that with probability at least $1 - \frac{1}{n^{26}}$, for all k -cuts \tilde{P} in G , $(1 - \epsilon) \cdot w(G, \tilde{P}) \leq w(G_2, \tilde{P}) \cdot r \leq (1 + \epsilon) \cdot w(G_2, \tilde{P})$.*

Proof. We obtain the graph G_2 by the following procedure: (i) Initialize $G' = G$. (ii) Fix $p = \frac{100 \cdot \log n}{\epsilon^2 \cdot \text{OPT}(G, 2)}$, where $\text{OPT}(G, 2)$ denotes the weight of a non-trivial minimum 2-cut of G . (iii) Keep each edge e in G' with probability p and remove with probability $1 - p$. (iv) Return G' as G_2 and $r = 1/p$.

Claim 3.1. *For every non-trivial 2-cut C in G , $|w(G_2, C) - w(G, C) \cdot p| \geq \epsilon \cdot w(G, C) \cdot p$ with probability at most $2 \cdot n^{-\frac{100 \cdot w(G, C)}{3 \cdot \text{OPT}(G, 2)}}$.*

Proof. Let us associate a random variable X_e to each edge $e \in G$, $X_e = 1$ if the edge e is in graph G_2 and 0 otherwise. The random variable $X_e = 1$ with probability p and $X_e = 0$ with probability $1 - p$. Let E' be the set of edges in the cut C in G and $X = \sum_{e \in E'} X_e$. The expectation of X (denoted by μ), is $\sum_{e \in E'} p = w(G, C) \cdot p$. Applying Chernoff bound (Lemma 2.2) on X , we get

$$\Pr[|w(G_2, C) - w(G, C) \cdot p| \geq \epsilon \cdot w(G, C) \cdot p] \leq 2 \cdot e^{-\frac{100\epsilon^2 \cdot \log n \cdot w(G, C)}{3\epsilon^2 \cdot \text{OPT}(G, 2)}} \leq 2 \cdot n^{-\frac{100 \cdot w(G, C)}{3 \cdot \text{OPT}(G, 2)}}.$$

\square

Next we would like to simultaneously apply Claim 3.1 over all non-trivial 2-cuts in G . Towards this we will first bound the number of non-trivial 2-cuts in G

Claim 3.2. *The probability that for all non-trivial 2-cuts C in G , $|w(G_2, C) - w(G, C) \cdot p| \leq \epsilon \cdot w(G, C) \cdot p$ is at least $1 - \frac{1}{n^{26}}$.*

Proof. First we show that the number of non-trivial 2-cuts in G of weight i , where $\text{OPT}(G, 2) \leq i \leq |E(G)|$ is at most $n^{2i/\text{OPT}(G,2)}$. For any $\alpha \geq 1$, a cut is called an α -*minimum cut* if its number of edges is at most α times larger than a minimum cut. In any undirected graph, and for any real number $\alpha \geq 1$, the number of α -minimum cuts is at most $n^{2\alpha}$ [Kar94]. By choosing $\alpha = i/\text{OPT}(G, 2)$, we get that the number of non-trivial 2-cuts in G of weight i is at most $n^{2i/\text{OPT}(G,2)}$.

Now we use union bound to show that the probability of all non-trivial 2-cuts C of G satisfying $|w(G_2, C) - w(G, C) \cdot p| \leq \epsilon \cdot w(G, C) \cdot p$ is at least $1 - \frac{1}{n^{26}}$. Let A denote the event that some non-trivial 2-cut C in G does not satisfy the required inequality. An upper bound on $\Pr[A]$ can be obtained by combining the Claim 3.1 along with the bound on the number of cuts of weight i as follows. However, recall that the number of edges in G , say m , is upper bounded by $m_{\text{orig}}^2/\epsilon$, where m_{orig} is the number of edges in the input graph (before we applied the Knapsack trick to make it an unweighted graph). This implies that $m \leq m_{\text{orig}}^2/\epsilon \leq n^5$.

$$\begin{aligned} \Pr[A] &\leq \sum_{i=\text{OPT}(G,2)}^m 2 \cdot n^{\frac{-100i}{3 \cdot \text{OPT}(G,2)}} \cdot n^{\frac{2i}{\text{OPT}(G,2)}} \\ &\leq \sum_{i=\text{OPT}(G,2)}^m 2 \cdot n^{-31i/\text{OPT}(G,2)} \\ &\leq \frac{1}{n^{31}} \cdot \frac{n^4}{\epsilon} \\ &\leq \frac{1}{n^{26}}. \end{aligned}$$

□

Claim 3.3. *The probability that for all k -cuts \tilde{P} in G , $(1 - \epsilon) \cdot w(G, \tilde{P}) \leq w(G_2, \tilde{P}) \cdot r \leq (1 + \epsilon) \cdot w(G, \tilde{P})$ is at least $1 - \frac{1}{n^{26}}$.*

Proof. Let $\tilde{P} = \{P_1, \dots, P_k\}$ be a k -cut in G , each part P_i is in some connected component X of G . Let \mathcal{C} contain all the non-trivial 2-cuts $C_i = (P_i, V(G_1) \setminus P_i)$ in G . The weight of \tilde{P} in G is given by,

$$\begin{aligned} w(G, \tilde{P}) &= \frac{\sum_{P_i \in \tilde{P}} w(G, (P_i, V(G_1) \setminus P_i))}{2} \\ &= \frac{\sum_{C_i \in \mathcal{C}} w(G, C_i)}{2} \end{aligned}$$

Thus, using the previous claim, with probability at least $1 - \frac{1}{n^{26}}$,

$$\begin{aligned}
|w(G_2, C_k) - w(G, C_k) \cdot p| &= \frac{\left| \sum_{C_i \in \mathcal{C}} w(G_2, C_i) - \sum_{C_i \in \mathcal{C}} w(G, C_i) \right|}{2} \\
&= \frac{\left| \sum_{C_i \in \mathcal{C}} w(G_2, C_i) - w(G, C_i) \right|}{2} \\
&= \frac{\left| \sum_{C_i \in \mathcal{C}} \epsilon \cdot w(G, C_i) \cdot p \right|}{2} \\
&= \frac{\epsilon \cdot p \cdot \left| \sum_{C_i \in \mathcal{C}} w(G, C_i) \right|}{2} \\
&= \epsilon \cdot w(G, \tilde{P}) \cdot p
\end{aligned}$$

□

The previous claim concludes the proof of the lemma. □

4 Edge Unbreakable Decomposition Theorem in Polynomial Time

In this section we give our main decomposition theorem. That is, we give a tree decomposition where each bag is edge-unbreakable. In particular we will prove the following theorem.

Theorem 4.1. *Given an n -vertex graph G and an integer s , one can in polynomial time compute a rooted compact tree decomposition (τ, χ) of G such that*

1. every adhesion of (τ, χ) is of size at most s ;
2. every bag of (τ, χ) is $((s+1)^5, s)$ -edge-unbreakable.

The proof of Theorem 4.1 closely follows the proof of Theorem 1.2 of Cygan et al. [CKL⁺18] with a few differences. The overall template of iterated refinement is *exactly* the same.

The main difference is found in Section 4.2 where we state and prove Lemma 4.2. In Cygan et al. [CKL⁺18] the corresponding lemma has vertex separators instead of edge cuts, condition 1 has $|A \cap Q| > s$ and $|A \cap Q| > s$ instead of $|A \cap Q| > (s+1)^5$ and $|A \cap Q| > (s+1)^5$ and their algorithm runs in $s^{\mathcal{O}(s)}n^{\mathcal{O}(1)}$ time instead of polynomial time. The proofs of Lemma 4.2 and Lemma 3.4 of Cygan et al. [CKL⁺18] are entirely different. The remaining parts of the proof are so similar that we **have copied their text verbatim and incorporated minor modifications to suit our need**. In particular *all of Subsections 4.1 and 4.3 are included solely for the convenience of the reader and ease of understanding*.

4.1 Improving Tree Decomposition given a Lean Witness

Our starting point for the proof of Theorem 4.1 is the definition of a lean tree decomposition of Thomas [Tho90]; we follow the formulation of [BD02].

Definition 4.1. *A tree decomposition (τ, χ) of a graph G is called lean if for every $t_1, t_2 \in V(\tau)$ and all sets $Z_1 \subseteq \chi(t_1)$ and $Z_2 \subseteq \chi(t_2)$ with $|Z_1| = |Z_2|$, either G contains $|Z_1|$ vertex-disjoint $Z_1 - Z_2$ paths, or there exists an edge $e \in E(\tau)$ on the path from t_1 to t_2 such that $|A(e)| < |Z_1|$.*

For a graph G and a tree decomposition (τ, χ) that is not lean, a quadruple (t_1, t_2, Z_1, Z_2) for which the above assertion is not true is called a *lean witness*. Note that it may happen that $t_1 = t_2$ or $Z_1 \cap Z_2 \neq \emptyset$. In particular (t, t, Z_1, Z_2) is called a *single bag lean witness*. The *order* of a lean witness is the minimum order of a separation (X_1, X_2) such that $Z_i \subseteq X_i$ for $i = 1, 2$. Observe that by Menger's theorem the following conditions are equivalent.

Claim 4.1. *For a tree decomposition (τ, χ) , a node $t \in V(\tau)$, and subsets $Z_1, Z_2 \subseteq \chi(t)$, either both or none of the following two conditions are true:*

- (t, t, Z_1, Z_2) is a single bag lean witness for (τ, χ) ,
- there is a separation (X_1, X_2) of G with $Z_i \subseteq X_i$, where $X = X_1 \cap X_2$, such that $|X| < |Z_1| = |Z_2|$, and there is a set of vertex disjoint $Z_1 - Z_2$ paths $\{P_x\}_{x \in X}$ such that $x \in V(P_x)$ for every $x \in X$.

Moreover given a single bag lean witness (t, t, Z_1, Z_2) one can find the above separation (X_1, X_2) and set of paths $\{P_x\}_{x \in X}$ in polynomial time.

A minimum order of a separation from the second point is called the *order* of the single bag lean witness (t, t, Z_1, Z_2) . To argue that the refinement process stops after a small number of steps, or that it stops at all, we define the following potential for a graph G , a tree decomposition (τ, χ) , and an integer s :

$$\Phi_{G,s}(\tau, \chi) = \sum_{t \in \tau} \max(0, |\chi(t)| - 2s - 1).$$

Lemma 4.1. *Assume we are given a graph G , an integer s , a tree decomposition (τ, χ) of G with every adhesion of size at most s , one node $q \in \tau$ with $|\chi(q)| > 2s + 1$, and a single bag lean witness (q, q, Z_1, Z_2) with $|Z_1| \leq s + 1$. Then one can in polynomial time compute a tree decomposition (τ', χ') of G with every adhesion of size at most s such that $\Phi_{G,s}(\tau, \chi) > \Phi_{G,s}(\tau', \chi')$.*

Proof. Apply Claim 4.1, yielding a separation (X_1, X_2) and a family $\{P_x\}_{x \in X}$ of vertex disjoint $Z_1 - Z_2$ paths, where $X = X_1 \cap X_2$. Note that $s + 1 \geq |Z_1| > |X|$, hence the order of (X_1, X_2) is at most s .

We construct a tree decomposition (τ', χ') as follows. First for every $i = 1, 2$, we construct a decomposition (τ^i, χ^i) of $G[X_i]$: we start with τ^i being a copy of τ , where a node $t \in V(\tau)$ corresponds to a node $t^i \in V(\tau^i)$, and $\chi^i(t^i) := \chi(t) \cap X_i$ for every $t \in V(\tau)$. Then for every $x \in X \setminus \chi(q)$ we take the node $t_x \in V(\tau)$ such that $x \in \chi(t_x)$ and t_x is closest to q among such nodes, and insert x into every bag $\chi^i(t^i)$ for t^i lying on the path between q^i (inclusive) and t_x^i (exclusive) in τ^i .

Clearly, (τ^i, χ^i) is a tree decomposition of $G[X_i]$ and $X \subseteq \chi^i(q^i)$. We construct (τ', χ') by taking τ' to be a disjoint union of τ^1 and τ^2 , with the copies of the node q connected by an edge $q^1 q^2$, and $\chi' := \chi^1 \cup \chi^2$. Since (X_1, X_2) is a separation and $X = X_1 \cap X_2$ is present in both bags $\chi^1(q^1), \chi^2(q^2)$, we infer that (τ', χ') is a tree decomposition of G .

We now argue that every adhesion of (τ', χ') is of size at most s . This is clearly true for the edge $q^1 q^2$ connecting τ^1 and τ^2 , as the adhesion there is exactly X and $|X| \leq s$.

Consider now a bag t^i in a tree (τ^i, χ^i) . The set $\chi^i(t^i) \setminus \chi(t)$ consists of some vertices of X , namely those vertices $x \in X \setminus \chi(q)$ for which t lies on the path between q (inclusive) and t_x (exclusive). However, by the properties of a tree decomposition and Menger's theorem, $\chi(t)$ contains at least one vertex of P_x that lies between x and the endpoint in Z_{3-i} , that is in $V(P_x) \cap (X_{3-i} \setminus X_i)$. This vertex is not present in $\chi^i(t^i)$ and, consequently, $|\chi^i(t^i)| \leq |\chi(t)|$. The same argumentation holds for every edge $e^i \in E(\tau^i)$ and adhesion of this edge.

We are left with analysing the potential decrease. Fix $t \in V(\tau)$. We analyse the difference between the contribution to the potential of t in (τ, χ) and the copies of t in (τ', χ') . First,

by the analysis in the previous paragraph, we have $|\chi^i(t^i)| \leq |\chi(t)|$ for $i = 1, 2$. Consequently, if $|\chi(t)| \leq 2s + 1$, then $|\chi^i(t^i)| \leq 2s + 1$ for $i = 1, 2$ and the discussed contributions are all equal to 0. Furthermore, if $|\chi^i(t^i)| \leq 2s + 1$ for some $i = 1, 2$, then $\max(|\chi(t)| - 2s - 1, 0) \geq \max(|\chi^{3-i}(t^{3-i})| - 2s - 1, 0)$ as $|\chi^{3-i}(t^{3-i})| \leq |\chi(t)|$.

Otherwise, assume that $|\chi(t)| > 2s + 1$ and $|\chi^i(t^i)| > 2s + 1$ for $i = 1, 2$. Note that $\chi^1(t^1) \cup \chi^2(t^2) \subseteq \chi(t) \cup X$ while $\chi^1(t^1) \cap \chi^2(t^2) \subseteq X$ and $|X| \leq s$. Consequently,

$$\begin{aligned} |\chi(t)| - 2s - 1 &\geq |\chi(t) \setminus X_2| + |\chi(t) \setminus X_1| - 2s - 1 \\ &\geq |(\chi(t) \setminus X_2) \cup X| + |(\chi(t) \setminus X_1) \cup X| - 4s - 1 \\ &> |\chi^1(t^1)| - 2s - 1 + |\chi^2(t^2)| - 2s - 1. \end{aligned}$$

We infer that for every bag $t \in V(\tau)$, the contribution of t to the potential $\Phi_{G,s}(\tau, \chi)$ is not smaller than the contribution of the two copies of t in $\Phi_{G,s}(\tau', \chi')$. To prove strict inequality, we show that these contributions are not equal for the bag q .

Recall that we assumed $|\chi(q)| > 2s + 1$. By the previous argumentation, the only chance for equal contributions of q to $\Phi_{G,s}(\tau, \chi)$ and q^1, q^2 to $\Phi_{G,s}(\tau', \chi')$ is that $|\chi^i(s^i)| \leq 2s + 1$ for some $i = 1, 2$. However, note that $\chi^{3-i}(q^{3-i}) = (\chi(q) \setminus X_i) \cup X$. Consequently, as $Z_i \subseteq \chi(q)$ and $|Z_i| > |X|$, we have $|\chi^{3-i}(q^{3-i})| < |\chi(q)|$ and hence $|\chi(q)| - 2s - 1 > \max(|\chi^{3-i}(q^{3-i})| - 2s - 1, 0)$. This finishes the proof of the lemma. \square

4.2 Finding a Lean Witness

Lemma 4.2. *There exists a polynomial time algorithm that takes as input an unweighted connected graph G , a set of terminals $Q \subseteq V(G)$ and an integer s and in polynomial time concludes either of the two statements.*

1. *There is no partition (A, B) of G of order at most s such that $|A \cap Q| \geq (s + 1)^5$ and $|B \cap Q| \geq (s + 1)^5$.*
2. *Outputs a partition (A, B) of G of order at most s such that $|A \cap Q| \geq s + 1$ and $|B \cap Q| \geq s + 1$.*

Proof. Towards the proof, we will show that if G has a partition (X, Y) of order at most s such that $|X \cap Q| \geq (s + 1)^5$ and $|Y \cap Q| \geq (s + 1)^5$, then we will always output a partition (A, B) of G of order at most s such that $|A \cap Q| \geq s + 1$ and $|B \cap Q| \geq s + 1$.

For rest of the proof we assume that we have a partition (X, Y) of order at most s such that $|X \cap Q| \geq (s + 1)^5$ and $|Y \cap Q| \geq (s + 1)^5$. Let $E(X, Y)$ be denoted by S . To design our algorithm in polynomial time we will output a polynomial size family \mathcal{F} of *connected sets* that has following properties.

1. For each set $C \in \mathcal{F}$, we have that $|C \cap Q| \geq s + 1$.
2. There exists a set $C_1 \in \mathcal{F}$ that belongs to $G[X]$ and there exists a set $C_2 \in \mathcal{F}$ that belongs to $G[Y]$.

Note that if we have \mathcal{F} of polynomial size satisfying the above properties then the algorithm just iterates over pairs (C_1, C_2) and computes a min-cut between C_1 and C_2 in polynomial time [KT06]. If this cut is of size at most s , then we can get a partition (A, B) of order at most s such that $C_1 \subseteq A$ and $C_2 \subseteq B$. Thus, we have that $|A \cap Q| \geq s + 1$ and $|B \cap Q| \geq s + 1$, concluding the proof. Next, we will show that if we do have a partition (X, Y) of order at most s such that $|X \cap Q| \geq (s + 1)^5$ and $|Y \cap Q| \geq (s + 1)^5$, then we will always find such a pair of (C_1, C_2) .

Now our job reduces to finding \mathcal{F} of polynomial size satisfying the above properties. To achieve this we start with a spanning tree T of G and use this to get the family \mathcal{F} . Let T

be rooted at some vertex. This immediately gives parent, child, ancestor, descendant relations among the vertices of the tree. We add the following connected sets to the family \mathcal{F} .

- For every vertex v of the tree, all the vertices in the rooted subtree T_v , if $|V(T_v) \cap Q| \geq s+1$.
- For every pair of vertices (u, v) in the tree where one is descendant of other, the path between them P_{uv} , if $|V(P_{uv}) \cap Q| \geq s+1$.
- Let P_{uv} be the path where one is descendant of other. Let W denote the set of children of P_{uv} . That is all those vertices whose parents are on P_{uv} . We say $w \in W$ is a *good child* if $|V(T_w) \cap Q| \geq 1$. Let W_{good} be the subset of W containing all good children. Let $|W_{\text{good}}| \geq (s+1)^2$. We partition $|W_{\text{good}}|$ into $(s+1)$ parts, say W_1, \dots, W_{s+1} such that each part contains at least $s+1$ vertices and no two sets have any element in common. Now, for all $1 \leq i \leq s+1$, add the set containing the path P_{uv} and subtree rooted at vertices in W_i . We call these sets H_1, \dots, H_{s+1} .

By our construction each set in \mathcal{F} is connected, contains at least $s+1$ terminals, and has size polynomial in n .

To prove that the \mathcal{F} has the desired property, all that remains to show is that there exists a set $C_1 \in \mathcal{F}$ that belongs to $G[X]$ and there exists a set $C_2 \in \mathcal{F}$ that belongs to $G[Y]$. We will only show that there exists a set $C_1 \in \mathcal{F}$ that belongs to $G[X]$. The proof for the other case is analogous.

Let C be a connected component in $G[X]$ (in $G-S$) with largest intersection with Q . Observe that C has at least $(s+1)^4$ terminals as deleting at most s edges can only make $s+1$ connected components. Since, T is a rooted tree, C itself is a rooted tree, let r be the root of C . For every vertex $v \in C$ incident to a cut-edge (edge in S), mark the path from v to r in C . Let M be the set of marked vertices in C . Observe that M is a *sub-tree* of C . For every $u \in M$, define the weight of u to be the number of terminals in the connected component of $C - (M - \{u\})$ that contains u . Note that the sum of weights in M is equal to number of terminal in C . Let P be the root-leaf path in M of largest weight. Since, $|S| \leq s$, we have that the weight of P ($w(P)$) is at least $\frac{w(C)}{s} \geq \frac{(s+1)^4}{s}$. Here, $w(C)$ denotes the number of terminals in C . If P contains at least $s+1$ terminals, then clearly $V(P) \in \mathcal{F}$. Thus, assume that there are at most s terminals on P . Since, $w(P)$ is at least $\frac{(s+1)^4}{s}$ we have that either there is a child u of a vertex in P , such that T_u has at least $s+1$ terminals or $|W_{\text{good}}| \geq (s+1)^2$. First let us handle the case when $|W_{\text{good}}| \geq (s+1)^2$. In this case, look at the sets H_1, \dots, H_{s+1} that we added to \mathcal{F} . We know that any edge in S can touch at most one set among H_1, \dots, H_{s+1} . This implies that there exists a set H_i that has no edge from S incident to it and thus, it is completely contained in $G[X]$. Indeed, it also contains at least $s+1$ terminals.

Now we know that $|W_{\text{good}}| < (s+1)^2$ and P has at most s terminals. Let u be the vertex of P . Then, the weight of a vertex u on P is given by the following.

$$w(u) = \begin{cases} 1 + \sum_{ux \text{ and } x \text{ is not marked}} |V(T_x) \cap Q| & \text{if } u \text{ is a terminal} \\ \sum_{ux \text{ and } x \text{ is not marked}} |V(T_x) \cap Q| & \text{otherwise} \end{cases}$$

Observe that if a vertex w is not *marked* then the subtree T_w is not incident to any edge in S because of our construction of M . Recall that, $w(P)$ is at least $\frac{(s+1)^4}{s}$. This implies that one of the edges that projected weight onto P projected at least $1/(s+1)^2$ fraction of total weight. Let ux be that edge. By the above definition of the weight of u , we have that x is unmarked and the subtree of T rooted at x has many terminals and furthermore, T_x is not incident to any edge in S and hence lies in $G[X]$. Furthermore, T_x has at least $\frac{(s+1)^4 - s}{(s+1)^2} \geq (s+1)$. This concludes the proof. \square

4.3 The Algorithm

For our proof we will use the following *cleanup procedure* on a tree decomposition (τ, χ) of a graph G : as long as there exists an edge $st \in E(\tau)$ with $\chi(s) \subseteq \chi(t)$, contract the edge st in τ , keeping the name t and the bag $\chi(t)$ at the resulting vertex. We shall say that a node s and bag $\chi(s)$ *disappears* in a cleanup step. Clearly, the final result (τ', χ') of a cleanup procedure is a tree decomposition of G and every adhesion of (τ', χ') is also an adhesion of (τ, χ) . Observe that $|E(\tau')| \leq |V(G)|$: if one roots τ' at an arbitrary vertex, going from child to parent on every edge τ' we forget at least one vertex of G , and every vertex can be forgotten only once.

It is well-known that every rooted tree decomposition can be refined to a compact one; see e.g. [BP16, Lemma 2.8]. We will need the following for our proof.

Lemma 4.3 ([CKL⁺18], Lemma 2.3). *Given a graph G and its tree decomposition (τ, χ) , one can compute in polynomial time a compact tree decomposition $(\hat{\tau}, \hat{\chi})$ of G such that every bag of $(\hat{\tau}, \hat{\chi})$ is a subset of some bag of (τ, χ) , and every adhesion of $(\hat{\tau}, \hat{\chi})$ is a subset of some adhesion of (τ, χ) .*

With Lemmas 4.2 and 4.3 in hand, we now formally prove Theorem 4.1.

Proof of Theorem 4.1. First, we can assume that G is connected, as otherwise we can compute a tree decomposition for every component separately, and then glue them up in an arbitrary fashion.

We start with a naive unrooted tree decomposition (τ, χ) that has a single bag with the entire vertex set and iteratively improve it, using Lemma 4.1, until it satisfies the conditions of Theorem 4.1, except for compactness, which we will handle in the end. We will maintain the invariant that every adhesion of (τ, χ) is of size at most s . At every step the potential $\Phi_{G,s}(\tau, \chi)$ will decrease, leading to at most $n - 2s - 1$ steps of the algorithm.

Let us now elaborate on a single step of the algorithm. There is one reason why (τ, χ) may not satisfy the conditions of Theorem 4.1: it contains a bag that is not $((s+1)^5, s)$ -edge unbreakable. Consider a bag $Q := \chi(t)$ that is not $((s+1)^5, s)$ -edge unbreakable. Now using Lemma 4.2, with graph G and terminal set Q , in polynomial time we can find a partition (A, B) of G of order at most s such that $|A \cap Q| \geq s + 1$ and $|B \cap Q| \geq s + 1$. Consequently, Lemma 4.2 allows us to find a single-bag lean witness of order at most s for the node t in polynomial time.

Thus, we uncovered a single-bag lean witness for a node $t \in V(\tau)$ satisfying $|\chi(t)| > 2s + 1$. Hence, we may refine the decomposition by applying Lemma 4.1 and proceed iteratively with the refined decomposition. As asserted by Lemma 4.1, the potential $\Phi_{G,k}(\tau, \chi)$ strictly decreases in each iteration.

We remark that between the refinement steps we need to reduce the number of edges in the decomposition to at most n using the same reasoning as in the proof of Lemma 4.3. That is, as long as there exists an edge $st \in E(\tau)$ with $\chi(s) \subseteq \chi(t)$, we can contract s onto t keeping $\chi(t)$ as the bag of the new node. A direct check shows that this operation neither increases the sizes of adhesions nor the potential $\Phi_{G,k}$ of the decomposition, while, as argued in the proof of Lemma 4.3, it bounds the number of edges of τ by n .

Observe that the potential $\Phi_{G,k}(\tau, \chi)$ is bounded polynomially in n and every iteration can be executed in polynomial time. Hence, we conclude that the refinement process finishes within polynomial time and outputs an unrooted tree decomposition (τ, χ) that satisfies all the requirements of Theorem 4.1, except for being compact. This can be remedied by applying the algorithm of Lemma 4.3. Note that neither the edge-unbreakability of bags nor the upper bound on the sizes of adhesions can deteriorate as a result of applying the algorithm of Lemma 4.3, as every bag (resp. every adhesion) of the obtained tree decomposition is a subset of a bag (resp. an adhesion) of the original one. This concludes the proof. \square

5 A $s^{\mathcal{O}(k)}$ Time Algorithm for MIN k -CUT on Unweighted Graphs

In this section we give our new algorithm for MIN k -CUT parameterized by s and k . It is a dynamic programming algorithm over an edge unbreakable tree decomposition (Theorem 4.1) and uses T -trees introduced by Thorup [Tho08] crucially.

Let G, k, s be input to MIN k -CUT. Towards our proof, in polynomial time, we first compute a rooted compact tree decomposition (τ, χ) of G , using Theorem 4.1, such that

1. every adhesion of (τ, χ) is of size at most s ;
2. every bag of (τ, χ) is $((s+1)^5, s)$ -edge-unbreakable.

Let \tilde{P}^* be an optimum k -cut of G . Secondly, we compute a family \mathcal{T} of spanning trees of G such that there exists a tree $T^* \in \mathcal{T}$ such that $E(T^*)$ crosses \tilde{P}^* at most $2k - 2$ times. Recall that T^* is said to be a T -tree of G . Such a family can be computed in polynomial time using Thorup's tree packing algorithm [Tho08].

Finally, we assume that the input graph G is connected. Else, if $\text{cc}(G) \geq k$, we return the connected components of G itself as the desired partition of G of value 0. Otherwise, if $\text{cc}(G) \leq k$, then at the expense of $k^k s^k$ in time we can guess the values of k and s for each connected component.

Thus, throughout this section, we assume that we are given (τ, χ) and a tree T from \mathcal{T} and remark that the notations $G, k, s, (\tau, \chi), T$ have fixed meanings. Also recall that in some parts, we refer to a l -cut of a graph G' as a l -partition of $V(G')$ but it will be clear from the context. We begin by defining additional terminologies and proving some results that will be helpful in defining the states of the dynamic programming.

Definition 5.1. *For a tree T' and an integer c , a partition \tilde{P} of $V(T')$ is c -admissible with respect to T' if $E(T')$ crosses \tilde{P} at most c times. For a subset X of $V(T')$ a partition \tilde{P} of X is T' -feasible if \tilde{P} is the projection onto X of a $2k - 2$ -admissible partition of $V(T')$.*

We will denote the family of all T' -feasible partitions of X by $\mathcal{F}_{T'}^X$. Further, if $X = \emptyset$ then $\mathcal{F}_{T'}^X$ will contain the partition P_\emptyset , where P_\emptyset is the partition of X that contains a single empty part.

Given a tree T' , we define the partition of $V(T')$ obtained by removing a subset E' of edges from T' as *the partition of $V(T')$ in which each part is equal to the set of vertices in some tree in the forest obtained by deleting E' from T' .*

Given a tree T' and $X \subseteq V(T')$, we define the *projection of T' onto X* as the tree $\text{proj}(T', X)$ obtained by the following procedure. Initially, set $T'' = T'$, and exhaustively apply the following modifications to T'' : (a) Delete from T'' all leaves of T'' that are not in X . (b) Smooth all degree 2 vertices v of T'' that are not in X . Once neither one of the operations (a), (b) can be applied the procedure returns T'' as the projection $\text{proj}(T', X)$. An example of $\text{proj}(T', X)$ is depicted in Figure 2.

For the projection $\text{proj}(T', X)$, we will denote its vertices and edges by $V_{T'}^X$ and $E_{T'}^X$ respectively. We encapsulate some useful properties of $\text{proj}(T', X)$ in the following observation.

Observation 5.1. *All leaves and degree 2 vertices of $\text{proj}(T', X)$ are in X . Furthermore, the number of vertices in $\text{proj}(T', X)$ is at most $2|X|$.*

The first part of the observation follows directly from the definition of the projection procedure, while the second follows from the well-known fact that in a tree, the number of vertices of degree at least three is at most the number of leaves. The following lemma shows that the set of T -feasible partitions of a set X does not change when T is projected on X . Thus, for enumerating T -feasible partitions of X it will be sufficient to work with the projection of T onto X instead, which is much more efficient.




Figure 2: (a) Tree T' is depicted by green edges and the set X by black vertices. (b) $\text{proj}(T', X)$ is depicted by red edges.

Lemma 5.1. For all $X \subseteq V(T)$, $\mathcal{F}_T^X = \mathcal{F}_{\text{proj}(T, X)}^X$. Further \mathcal{F}_T^X is of size $\mathcal{O}((4k|X|)^{2k})$ and can be computed in time proportional to its size.

Proof. For the proof, we show that every T -feasible partition of X is a $\text{proj}(T, X)$ -feasible partition of X and that every $\text{proj}(T, X)$ -feasible partition of X is a T -feasible partition of X .

Let \tilde{P} be a T -feasible partition of X , we show that \tilde{P} is also a $\text{proj}(T, X)$ -feasible partition. Since \tilde{P} is T -feasible, there exists a $2k - 2$ -admissible partition $\tilde{P}_{V(T)}$ of $V(T)$ whose projection on X is \tilde{P} . Let E_1 be the set of edges from $E(T)$ that cross $\tilde{P}_{V(T)}$. We construct a set E_2 of edges from $\text{proj}(T, X)$ of size at most $2k - 2$ such that the projection of the partition obtained by removing E_2 from $\text{proj}(T, X)$ on X is \tilde{P} . For each pair $x, x' \in X$ in different parts in \tilde{P} , there is at least one edge in E_1 in the path $P(x, x')$ connecting x and x' in T , call one such arbitrary edge $e_{x, x'}$. For each $e_{x, x'}$, we will find an edge $e'_{x, x'}$ in $E(\text{proj}(T, X))$ such that $e'_{x, x'}$ lies in the path between x and x' in $\text{proj}(T, X)$ and add it to E_2 and prove that none of the edges in E_2 are in any path $P(y, y')$ in $\text{proj}(T, X)$ connecting two vertices $y, y' \in X$ that are in same component in \tilde{P} . Thus the projection of the partition obtained by removing E_2 from $\text{proj}(T, X)$ on X will be \tilde{P} and hence \tilde{P} will be $\text{proj}(T, X)$ -feasible. We now find the edge $e'_{x, x'}$. Let $e_{x, x'} = uv$, assume that $P(x, x') = P(x, u) \cup uv \cup P(v, x')$. Let u' be the last vertex in $P(x, u)$ that has not been smoothed in $\text{proj}(T, X)$ and v' be the first vertex in $P(v, x')$ that has not been smoothed in $\text{proj}(T, X)$. Note that since we never smooth out x and x' , u' and v' must exist and can be equal to x and x' respectively. We set $e'_{x, x'} = u'v'$.

We now prove that none of the edges $e'_{x, x'} \in E_2$ can lie in the path between two vertices $y, y' \in X$ that are in same components in \tilde{P} in $\text{proj}(T, X)$. Suppose not, assume some edge $e'_{x, x'} = u'v' \in E_2$ lies in the path between two vertices $y, y' \in X$ that are in the same component in \tilde{P} . Then, by the construction of $\text{proj}(T, X)$, $e_{x, x'} = uv$ must lie in the path between y, y' in T . Thus, y, y' would not have been in the same components in \tilde{P} , this is a contradiction.

For the converse, let \tilde{P} be a $\text{proj}(T, X)$ -feasible partition of X , we show that \tilde{P} is also a T -feasible partition. Since \tilde{P} is $\text{proj}(T, X)$ -feasible, there exists a $2k - 2$ -admissible partition $\tilde{P}_{V(T)}$ of $V(T)$ whose projection on X is \tilde{P} . Let E_1 be the set of edges from $E(T)$ that cross $\tilde{P}_{V(T)}$. We construct a set E_2 of edges from T of size at most $2k - 2$ such that the projection of the partition obtained by removing E_2 from T on X is \tilde{P} . For each pair $x, x' \in X$ in different parts in \tilde{P} , there is at least one edge in E_1 in the path $P(x, x')$ connecting x and x' in $\text{proj}(T, X)$, call one such arbitrary edge $e_{x, x'}$. For each $e_{x, x'}$, we will find an edge $e'_{x, x'}$ in $E(T)$ such that $e'_{x, x'}$ lies in the path between x and x' in T and add it to E_2 and prove that none of the edges in E_2

are in any path $P(y, y')$ in T connecting two vertices $y, y' \in X$ that are in same component in \tilde{P} . Thus the projection of the partition obtained by removing E_2 from T on X will be \tilde{P} and hence \tilde{P} will be a T -feasible partition of X . We now find the edge $e'_{x,x'}$. Let $e_{x,x'} = uv$, assume that $P(x, x') = P(x, u) \cup uv \cup P(v, x')$. Let v' be the vertex adjacent to vertex u in the path $P(u, v)$ between u and v in T . We set $e'_{x,x'} = uv'$. We now prove that none of the edges $e'_{x,x'} \in E_2$ can lie in the path between two vertices $y, y' \in X$ in T that are in same components in \tilde{P} . Suppose not, assume some edge $e'_{x,x'} = uv' \in E_2$ lies in the path between two vertices $y, y' \in X$ that are in same components in \tilde{P} . All the interior vertices in path $P(u, v)$ were smoothed in $\text{proj}(T, X)$. So, by construction of $\text{proj}(T, X)$, none of the interior vertices are in X nor have a path to y or y' in T that doesn't include u or v . Thus, uv was an edge in the path between y and y' in $\text{proj}(T, X)$. This implies that y, y' would not have been in the same components in \tilde{P} , this is a contradiction.

Thus, we have proved that $\mathcal{F}_T^X = \mathcal{F}_{\text{proj}(T, X)}^X$. To find the set \mathcal{F}_T^X , we do the following procedure. We initialize set $\mathcal{P} = \emptyset$. For each subset E' of edges in $\text{proj}(T, X)$ of size at most $2k - 2$, we obtain the partition \tilde{P} by removing the edges in E' from $\text{proj}(T, X)$ and add to \mathcal{P} all partitions that \tilde{P} refines. Finally we return \mathcal{P} as $\mathcal{F}_{\text{proj}(T, X)}^X$. Clearly, the set $\mathcal{F}_{\text{proj}(T, X)}^X$ is of size $\mathcal{O}((2|X|)^{2k}(2k)^{2k})$ which is $\mathcal{O}((4k|X|)^{2k})$ and the procedure takes time corresponding to its size to compute it. \square

For the rest of this section, unless stated otherwise, we will assume that the weight of any partition of a set of vertices $X \subseteq V$ is computed with respect to $G[X]$.

We now define the states of the dynamic programming algorithm over the tree τ that we will use to compute the required value. For each node $t \in V(\tau)$ we define a function $f_t: \mathcal{F}_T^{A_t} \times \{1, \dots, k\} \rightarrow \{0, \dots, s\} \cup \{\infty\}$ that our algorithm will compute. The domain of f_t consists of all pairs (\tilde{P}_{A_t}, i) where \tilde{P}_{A_t} is a T -feasible (meaning crosses T at most $2k - 2$ times) partition of A_t , where A_t is the set of vertices that are common between node t and its parent in τ , and i is a positive integer less than or equal to k . On input (\tilde{P}_{A_t}, i) , f_t returns the smallest possible weight of a T -feasible i -partition \tilde{P} of $V(G_t)$ such that \tilde{P}_{A_t} is the projection of \tilde{P} on A_t . However, if this weight is greater than s , or no such partition exists then $f_t(\tilde{P}_{A_t}, i)$ returns ∞ . The main step of the algorithm of Theorem 1.2 is an algorithm that computes $f_t(\tilde{P}_{A_t}, i)$ for every node t and pair $(\tilde{P}_{A_t}, i) \in \mathcal{F}_T^{A_t} \times \{1, \dots, k\}$ assuming that the values of $f_{t'}$ have already been computed for all children t' of t . We now state that this step can be carried out efficiently.

Lemma 5.2. *There exists an algorithm that takes as input (τ, χ) , a node $t \in V(\tau)$, a T -feasible partition \tilde{P}_{A_t} of A_t and a positive integer $i \leq k$, together with the value of $f_{t'}(\tilde{P}'_{A_{t'}}, i')$ for every child t' of t , T -feasible partition $\tilde{P}'_{A_{t'}}$ of $A_{t'}$, and positive integer $i' \leq i$, and outputs $f_t(\tilde{P}_{A_t}, i)$ in time $s^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$.*

Since a T -tree is a spanning tree of G such that there exists an optimal k -cut of G that is T -feasible, it follows that if T is a T -tree of G then $\text{OPT}(G, k)$ is given by $f_r(P_\emptyset, k)$, where r is the root of τ , recall that $A_r = \emptyset$. If T is not a T -tree of G , then $f_r(P_\emptyset, k)$ is the weight of some k -cut of G of weight at most s or ∞ and thus $f_r(P_\emptyset, k) \geq \text{OPT}(G, k)$. Therefore, we will now prove Theorem 1.2 using Lemma 5.2.

Proof of Theorem 1.2. As explained in the beginning of the section, we assume that (a) the graph is connected, (b) we are given (τ, χ) , satisfying properties mentioned in Theorem 4.1 and (c) a tree T from \mathcal{T} [Tho08]. We design an algorithm that takes all these inputs and works as follows. Starting from the leaf nodes, in a bottom up manner, for each node $t \in \tau$, the algorithm computes $\mathcal{F}_T^{A_t}$ using Lemma 5.1 and obtains the value of $f_t(\tilde{P}_{A_t}, i)$ for all pairs $(\tilde{P}_{A_t}, i) \in \mathcal{F}_T^{A_t} \times \{1, \dots, k\}$ using Lemma 5.2. Finally it returns the value $f_r(P_\emptyset, k)$, where r is the root of τ .

Since $f_r(P_\emptyset, k) \geq \text{OPT}(G, k)$ for any tree spanning tree T of G and $f_r(P_\emptyset, k) = \text{OPT}(G, k)$ if T is a T -tree, the algorithm returns the desired output.

It is easy to see that $\mathcal{F}_T^{A_t}$ is of size $\mathcal{O}((4ks)^{2k})$ which is $s^{\mathcal{O}(k)}$ and can be computed in time $s^{\mathcal{O}(k)}$ from Lemma 5.1 and the fact that $|A_t| \leq s$ for all $t \in \tau$. Thus the domain of f_t is of size $s^{\mathcal{O}(k)} \cdot k$ and each state can be computed in time $s^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$ from Lemma 5.2. Thus the algorithm runs in time $s^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. This completes the proof. \square

For proving Lemma 5.2, we will design an algorithm that takes all the inputs mentioned in Lemma 5.2 and returns a positive integer v , which is the weight of some T -feasible i -partition \tilde{P} of $V(G_t)$ having weight at most s such that \tilde{P}_{A_t} is the projection of \tilde{P} on A_t . If no such partition exists, the algorithm will return $v = \infty$. This automatically ensures that $f_t(\tilde{P}_{A_t}, i) \leq v$. The difficult part is to ensure that $f_t(\tilde{P}_{A_t}, i) \geq v$. If $f_t(\tilde{P}_{A_t}, i) = \infty$, this inequality trivially holds, and so it suffices to prove it for the cases when $f_t(\tilde{P}_{A_t}, i)$ is finite and a T -feasible i -partition \tilde{P} of $V(G_t)$ such that \tilde{P}_{A_t} is the projection of \tilde{P} on A_t and having weight at most s exists. The proof of Lemma 5.2 spans the remainder of this section and follows the route of “randomized contraction style Dynamic Programming” [CCH⁺16, CKL⁺18] but is somewhat more complicated because at every step we need to use the T -tree T to speed up computations. We begin by giving names to some of the central objects of this proof. The notation for these tools will be used throughout the section and are summarized in Figure 3.

Let $\tilde{P}(\tilde{P}_{A_t}, i)$ be a smallest possible weight T -feasible i -partition of $V(G_t)$ such that \tilde{P}_{A_t} is the projection of $\tilde{P}(\tilde{P}_{A_t}, i)$ on A_t having weight at most s . We say that $\tilde{P}(\tilde{P}_{A_t}, i)$ realises $f_t(\tilde{P}_{A_t}, i)$. We introduce the following notations assuming such a partition exists.

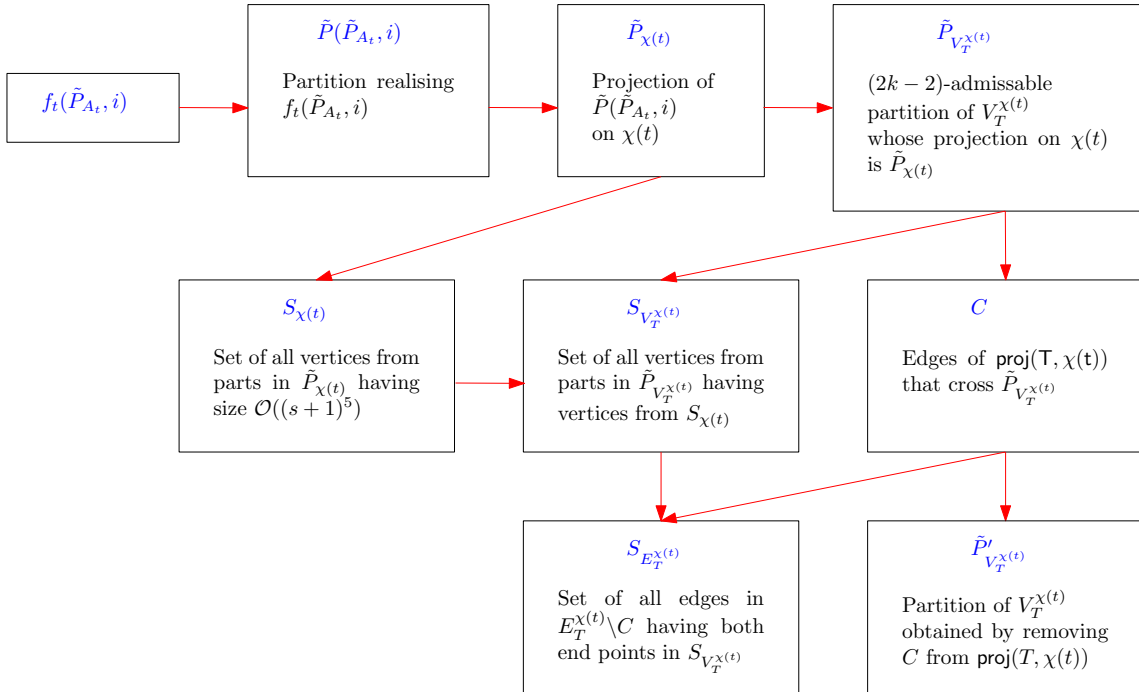


Figure 3: Various notations associated with the state $f_t(\tilde{P}_{A_t}, i)$ along with their dependencies.

Let $\tilde{P}_{\chi(t)}$ be the projection of $\tilde{P}(\tilde{P}_{A_t}, i)$ onto $\chi(t)$. Since $\tilde{P}(\tilde{P}_{A_t}, i)$ is a T -feasible i -partition of $V(G_t)$, $\tilde{P}_{\chi(t)}$ must be a T -feasible partition of $\chi(t)$. Recall that the vertices and edges of $\text{proj}(T, \chi(t))$ are denoted by $V_T^{\chi(t)}$ and $E_T^{\chi(t)}$ respectively. An example of $\chi(t)$ along with A_t and $\text{proj}(T, \chi)$ is illustrated in Figure 4. From Lemma 5.1 it follows that $\tilde{P}_{\chi(t)}$ is the projection onto




Figure 4: (a) An example of a bag $\chi(t)$ in (τ, χ) along with the adhesions (blue) associated with it, the adhesion A_t is in green. (b) An example of $\text{proj}(T, \chi(t))$ for the bag $\chi(t)$ in (a).

$\chi(t)$ of some $2k - 2$ -admissible partition $\tilde{P}_{V_T^{\chi(t)}}$ of $V_T^{\chi(t)}$. Let C be the set of edges of $\text{proj}(T, \chi(t))$ that cross $\tilde{P}_{V_T^{\chi(t)}}$, it is easy to see that $|C| \leq 2k - 2$. Let $\tilde{P}'_{V_T^{\chi(t)}}$ be the partition of $V_T^{\chi(t)}$ obtained by deleting the edges in C from $\text{proj}(T, \chi(t))$. Observe that $\tilde{P}'_{V_T^{\chi(t)}}$ is a refinement of $\tilde{P}_{V_T^{\chi(t)}}$. We remark that $\tilde{P}'_{V_T^{\chi(t)}}$ and $\tilde{P}_{V_T^{\chi(t)}}$ are different objects, and that this difference is crucial to our arguments (we apologize for the notational similarity!).

Since (τ, χ) is an $((s + 1)^5, s)$ -unbreakable tree decomposition, each part in $\tilde{P}_{\chi(t)}$ is of size $\mathcal{O}((s + 1)^5)$ except for at most one part. Let $S_{\chi(t)}$ denote the set of all vertices in parts having size $\mathcal{O}((s + 1)^5)$ in $\tilde{P}_{\chi(t)}$. The set of vertices from $\chi(t)$ in each part in $\tilde{P}_{V_T^{\chi(t)}}$ are either all from $S_{\chi(t)}$ or from $\chi(t) \setminus S_{\chi(t)}$ as $\tilde{P}_{\chi(t)}$ is the projection of $\tilde{P}_{V_T^{\chi(t)}}$ on $\chi(t)$. Let the union of the vertices in parts in $\tilde{P}_{V_T^{\chi(t)}}$ that contain vertices from $S_{\chi(t)}$ be denoted by $S_{V_T^{\chi(t)}}$, that is $S_{V_T^{\chi(t)}} = \bigcup_{P \in \tilde{P}_{V_T^{\chi(t)}}, P \cap \chi(t) \subseteq S_{\chi(t)}} P$. Observe that $S_{\chi(t)} \subseteq S_{V_T^{\chi(t)}}$. Let the set of edges in $E_T^{\chi(t)} \setminus C$ having both end points in $S_{V_T^{\chi(t)}}$ be denoted by $S_{E_T^{\chi(t)}}$. We now state a lemma that will help bound the sizes of the sets $S_{E_T^{\chi(t)}}$, and $S_{V_T^{\chi(t)}}$.

Lemma 5.3. *The size of each part P in the partition $\tilde{P}'_{V_T^{\chi(t)}}$ is at most $2(|P \cap \chi(t)| + 2k - 2)$. Furthermore the sets $S_{E_T^{\chi(t)}}$ and $S_{V_T^{\chi(t)}}$ are of size at most $2 \cdot (2k - 1)(|S_{\chi(t)}| + 2k - 2)$.*

Proof. Since $\tilde{P}'_{V_T^{\chi(t)}}$ is the partition of $V_T^{\chi(t)}$ obtained by removing the edges in C from $\text{proj}(T, \chi(t))$, each part P in it can be associated with a tree T_P in the forest obtained by removing C from $\text{proj}(T, \chi(t))$. The number of vertices from $V_T^{\chi(t)} \setminus \chi(t)$ that are a leaf or a degree two vertex in T_P is at most $2k - 2$ because from Observation 5.1 they had degree more than two in $\text{proj}(T, \chi(t))$ and thus must have been adjacent to an edge in C . The number of vertices in P having degree greater than two in T_P is bounded by the number of leaves in T_P which is bounded by $|P \cap \chi(t)| + 2k - 2$ because every leaf of T_P is either in $\chi(t)$ or not in $\chi(t)$. Thus, the number of vertices from $V_T^{\chi(t)} \setminus \chi(t)$ in P is bounded by $|P \cap \chi(t)| + 2 \cdot (2k - 2)$ and hence the size of P is at most $2 \cdot (|P \cap \chi(t)| + 2k - 2)$.

Since the size of C is at most $2k - 2$, there are at most $2k - 1$ parts in $\tilde{P}'_{V_T^{\chi(t)}}$ and thus the sizes of $S_{E_T^{\chi(t)}}$ and $S_{V_T^{\chi(t)}}$ are at most $2 \cdot (2k - 1)(|S_{\chi(t)}| + 2k - 2)$. \square

Since $S_{\chi(t)}$ is of size $\mathcal{O}(k(s+1)^5)$, the following observation directly follows from Lemma 5.3

Observation 5.2. *The set C is of size at most $2k - 2$ and the sets $S_{E_T^{\chi(t)}}$ and $S_{V_T^{\chi(t)}}$ are of size $\mathcal{O}(k^2(s+1)^5)$.*

In later parts of this section, we develop an algorithm that takes as input a subset C' of edges of $\text{proj}(T, \chi(t))$ and outputs a positive integer v greater than or equal to $f_t(\tilde{P}_{A_t}, i)$. Further if $C \subseteq C'$ and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$, then v will be equal to $f_t(\tilde{P}_{A_t}, i)$. We state this as a result below in Lemma 5.4 and will be using it to prove Lemma 5.2.

Lemma 5.4. *There exists an algorithm that takes as input (τ, χ) , a node $t \in V(\tau)$, a T -feasible partition \tilde{P}_{A_t} of A_t , a positive integer $i \leq k$, $\text{proj}(T, \chi(t))$, a set of edges $C' \subseteq E_T^{\chi(t)}$, together with the value of $f_{t'}(\tilde{P}'_{A_{t'}}, i')$ for every child t' of t , T -feasible partition $\tilde{P}'_{A_{t'}}$ of $A_{t'}$, and positive integer $i' \leq i$, and outputs an integer $v \geq f_t(\tilde{P}_{A_t}, i)$ in time $s^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$. Further if $f_t(\tilde{P}_{A_t}, i) \neq \infty$, $C \subseteq C'$, and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$, then $v \leq f_t(\tilde{P}_{A_t}, i)$.*

We will now prove that Lemma 5.2 is true assuming Lemma 5.4. In the proof, we will apply the well-known technique of splitters to obtain a family \mathcal{C} of subsets of edges in $\text{proj}(T, \chi)$ such that if $f_t(\tilde{P}_{A_t}, i) \neq \infty$, there is a set C' in the family such that $C \subseteq C'$, and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$. Then, we will apply Lemma 5.4 on each set in this family to obtain $f_t(\tilde{P}_{A_t}, i)$.

Proof of Lemma 5.2 assuming Lemma 5.4. For the proof, we propose an algorithm that takes inputs as specified in Lemma 5.2, uses Lemma 5.4 and computes $f_t(\tilde{P}_{A_t}, i)$.

Firstly, the algorithm uses Lemma 2.1 with $S = E_T^{\chi(t)}$, $s_1 = 2k - 2$, and $s_2 = \mathcal{O}(k^2(s+1)^5)$ to obtain a set \mathcal{C} of subsets of $E_T^{\chi(t)}$. For each set C' in \mathcal{C} , the algorithm calls the algorithm proposed in Lemma 5.4 with C' along with the other inputs and obtains as output a value $v_{C'}$. Finally it returns the integer $v = \min_{C' \in \mathcal{C}} v_{C'}$.

It follows from Lemma 5.4 that for each pair C' in \mathcal{C} , $v_{C'} \geq f_t(\tilde{P}_{A_t}, i)$. Thus, if $f_t(\tilde{P}_{A_t}, i) = \infty$, then the algorithm will return the value $v = \infty$. If $f_t(\tilde{P}_{A_t}, i) \neq \infty$, by Lemma 2.1 and Observation 5.2, the set \mathcal{C} contains a set C' that satisfies $C \subseteq C'$, and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$. By Lemma 5.4, for that C' , $v_{C'} = f_t(\tilde{P}_{A_t}, i)$. Thus if $f_t(\tilde{P}_{A_t}, i) \neq \infty$, the algorithm will return the integer $v = f_t(\tilde{P}_{A_t}, i)$.

The time taken by the algorithm is equal to the time taken to compute \mathcal{C} plus the time taken to compute $v_{C'}$ for each $C' \in \mathcal{C}$. From Lemma 2.1 and Lemma 5.4, it follows that the time taken by the algorithm is $s^{\mathcal{O}(k)}n^{\mathcal{O}(1)}$. This completes the proof. \square

The remainder of this section will be dedicated to proving Lemma 5.4, the crux of our result.

5.1 Nice Decompositions of $\chi(t)$

We first begin by defining *nice decompositions* of $\chi(t)$, a structure that will be helpful for proving Lemma 5.4.

Definition 5.2. *Let $\tilde{P}'_{\chi(t)}$ be a partition of $\chi(t)$, $\tilde{Q}_{\chi(t)}$ be a refinement of $\tilde{P}'_{\chi(t)}$ and O be a part in $\tilde{P}'_{\chi(t)}$ or an empty set. The triple $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ is called a nice decomposition of $\chi(t)$ if it satisfies the following properties:*

- (i) *If $O \neq \emptyset$, then the projection of $\tilde{Q}_{\chi(t)}$ on O is O , that is O is not refined by $\tilde{Q}_{\chi(t)}$ and if $O = \emptyset$, then $\tilde{P}'_{\chi(t)}$ has exactly one part.*




Figure 5: Examples of nice decomposition $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ of $\chi(t)$. (a) $O \neq \emptyset$ (b) $O = \emptyset$. The blue parts depict the parts in $\tilde{P}'_{\chi(t)}$ and the green parts depict those in $\tilde{Q}_{\chi(t)}$. In (c), (d) examples of possible edges and adhesions among the vertices of $\chi(t)$ in (a) and (b) respectively are depicted. Note that there are no adhesions and edges among different parts in $\tilde{P}'_{\chi(t)} \setminus O$.

- (ii) The size of the projection of $\tilde{Q}_{\chi(t)}$ on each part $P \in \tilde{P}'_{\chi(t)} \setminus O$ has at most $2k - 1$ parts.
- (iii) There are no cross edges between the parts in the partition $\tilde{P}'_{\chi(t)} \setminus O$.
- (iv) There are no adhesions in the set $\{A_{t'} : t' \in \text{children}(t) \text{ or } t' = t\}$ having vertices from more than one part in $\tilde{P}'_{\chi(t)} \setminus O$.

We refer to O as the center of the nice decomposition.

The structure of a nice decomposition of $\chi(t)$ is depicted by examples in Figure 5.

In order to prove Lemma 5.4, we first state and prove some results related to nice decompositions of $\chi(t)$ that will help us with the proof. Firstly, we develop an algorithm that computes a *small* family of nice decompositions of $\chi(t)$ containing a special nice decomposition in desirable cases. Then, we design another algorithm that computes a value greater than $f_t(\tilde{P}_{A_t}, i)$ when input any nice decomposition and computes $f_t(\tilde{P}_{A_t}, i)$ exactly if the input nice decomposition is special.

Lemma 5.5. *There exists an algorithm that takes as input (τ, χ) , a node $t \in V(\tau)$, $\text{proj}(\mathbb{T}, \chi(t))$, a set of edges $C' \subseteq E_T^{\chi(t)}$ and returns a set \mathcal{D} of nice decompositions of $\chi(t)$ of size $\mathcal{O}(s^{\mathcal{O}(k)} \log n)$ in time $\mathcal{O}(s^{\mathcal{O}(k)} n^{\mathcal{O}(1)})$. Furthermore, if $f_t(\tilde{P}_{A_t}, i) \neq \infty$, $C \subseteq C'$, and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$, then set \mathcal{D} contains a nice decomposition $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ such that $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}_{\chi(t)}$.*

Proof. Given $G, k, (\tau, \chi), t \in V(\tau), \text{proj}(T, \chi(t))$, and $C' \subseteq E_T^{\chi(t)}$ we design an algorithm that returns a set of nice decompositions by carrying out the following steps :

- (i) Set \tilde{R} to be the partition of $V_T^{\chi(t)}$ obtained by removing the edges in C' from $\text{proj}(T, \chi(t))$.
- (ii) If $\chi(t)$ is of size greater than $\mathcal{O}((s+1)^5)$, go to step (iii). Else, let $\tilde{P}'_{\chi(t)}$ be the partition of $\chi(t)$ having all vertices in $\chi(t)$ in one part and let $\tilde{Q}_{\chi(t)}$ be the projection of \tilde{R} on $\chi(t)$. If the number of parts in \tilde{R} is less than or equal to $2k-1$, then return the set \mathcal{D} of nice decompositions as $\{(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, \emptyset)\}$, else return $\mathcal{D} = \emptyset$.
- (iii) Let S be the set of all parts in \tilde{R} . Use Lemma 2.1 on S , $s_1 = 2k-1$ and $s_2 = (2k-1)(s^2 + 2s + k)$ and obtain a set \mathcal{S} of subsets of S .
- (iv) Construct a set \mathcal{D} of nice decompositions using the following procedure: Initialize $\mathcal{D} = \emptyset$ and for each $X \in \mathcal{S}$, repeat the following steps:
 - (a) Initialize partition $\tilde{Q} = \tilde{R}$.
 - (b) Combine all parts in \tilde{Q} that are not in X . Let this part be Q_1 .
 - (c) Form a graph G' with the parts in \tilde{Q} except Q_1 as vertices. Add an edge between two parts in G' if they have a cross edge between them or both parts have non empty intersection with some adhesion $A_{t'}, t' \in \text{children}(t) \cup t$.
 - (d) Initialize $\tilde{Q}' = \tilde{Q}$, combine all parts in connected components in G' having more than $2k-1$ parts with Q_1 in \tilde{Q}' . Call this part O' .
 - (e) Initialize $\tilde{P}' = \tilde{Q}'$. For each connected component Y in G' having at most $2k-1$ parts, combine all parts in Y in \tilde{P}' .
 - (f) Let $\tilde{Q}_{\chi(t)}$ be the projection of \tilde{Q}' on $\chi(t)$ and let $\tilde{P}'_{\chi(t)}$ be the projection of \tilde{P}' on $\chi(t)$. Let O be the projection of O' on $\chi(t)$.
 - (g) Add the pair $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ to \mathcal{D} if $O \neq \emptyset$.
- (v) Return \mathcal{D} .

Claim 5.1. *Each triple $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O) \in \mathcal{D}$ is a nice decomposition of $\chi(t)$. Further if $f_t(\tilde{P}_{A_t}, i) \neq \infty, C \subseteq C',$ and $S_{E_T^{\chi(t)}} \cap C' = \emptyset,$ then \mathcal{D} contains a nice decomposition $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ such that $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}_{\chi(t)}$.*

Proof. In the proof, will be using the notations summarized in Figure 3. Firstly, we need to prove that every triple $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ in \mathcal{D} is a nice decomposition of $\chi(t)$. For this we need to show that $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}'_{\chi(t)}$, O is either a part in $\tilde{P}'_{\chi(t)}$ or an empty set and that the triple satisfies all the properties stated in Definition 5.2. Then, to complete the proof we need to show that if $C \subseteq C',$ and $S_{E_T^{\chi(t)}} \cap C' = \emptyset,$ then \mathcal{D} contains a triple $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ such that $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}_{\chi(t)}$. We divide the proof into two cases, one where the size of $\chi(t)$ is of $\mathcal{O}((s+1)^5)$ and other where it is not.

In the first case, \mathcal{D} is either an empty set or has only one nice decomposition $\{(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)\},$ where, $\tilde{P}'_{\chi(t)}$ is the partition of $\chi(t)$ having all vertices in $\chi(t)$ in one part, $\tilde{Q}_{\chi(t)}$ is the projection of \tilde{R} on $\chi(t)$ and $O = \emptyset$. If $\mathcal{D} \neq \emptyset,$ then it is easy to see that $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}'_{\chi(t)}$ and that all the four properties are satisfied since there is only one part in $\tilde{P}'_{\chi(t)}$. Thus, \mathcal{D} is a set having zero or one nice decomposition of $\chi(t)$. In this case, since the size of $\chi(t)$ is of $\mathcal{O}((s+1)^5),$ it follows that by definition if $f_t(\tilde{P}_{A_t}, i) \neq \infty,$ then $S_{E_T^{\chi(t)}} = E_T^{\chi(t)} \setminus C$. Thus, if $C \subseteq C',$ and $S_{E_T^{\chi(t)}} \cap C' = \emptyset,$ then $C = C'$. Therefore, $\tilde{P}'_{\chi(t)}$ is the same as $\tilde{P}_{\chi(t)}$ and since $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}'_{\chi(t)},$ it follows that it is a refinement of $\tilde{P}_{\chi(t)}$ also. This completes the proof for the first case.

We now argue for the case when the size of $\chi(t)$ is not of $\mathcal{O}(s+1)^5$. By construction, every triple $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ in \mathcal{D} corresponds to some $X \in \mathcal{S}$ and has $O \neq \emptyset$.

By construction, from step (iv), \tilde{Q}' is a refinement of \tilde{P}' since \tilde{P}' is obtained by combining components in \tilde{Q}' . Since $\tilde{Q}_{\chi(t)}$ and $\tilde{P}'_{\chi(t)}$ are just projections of these partitions on $\chi(t)$, $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}'_{\chi(t)}$. Since O' is a part in \tilde{P}' and \tilde{Q}' and $O \neq \emptyset$, it is easy to see that O is part in $\tilde{P}'_{\chi(t)}$ and is not refined by $\tilde{Q}_{\chi(t)}$ as it is just the projection of O' on $\chi(t)$. In step (d), we obtain \tilde{Q}' from \tilde{Q} by combining all parts in connected components of G' that have more than $2k-1$ parts in them with Q_1 . In step (e) we construct \tilde{P}' from \tilde{Q}' by combining the parts in each connected component Y in G' having at most $2k-1$ parts into one component in \tilde{P}' . Thus, the projection of \tilde{Q}' on each part of \tilde{P}' has at most $2k-1$ parts and hence the projection of $\tilde{Q}_{\chi(t)}$ on each part of $\tilde{P}'_{\chi(t)}$ has at most $2k-1$ parts. Since in G' two parts from $\tilde{Q} \setminus Q_1$ are adjacent only if there is some cross edge between them or both have non empty intersection with some adhesion $A_{t'}$, $t' \in \text{children}(t) \cup t$, no parts in $\tilde{P}' \setminus O'$ and thus in $\tilde{P}'_{\chi(t)} \setminus O$ have any cross edge between them or have common adhesions. Therefore all the properties in Definition 5.2 are satisfied by $\tilde{P}'_{\chi(t)}$, $\tilde{Q}_{\chi(t)}$ and O and thus $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ is a nice decomposition of $\chi(t)$.

In this case, if $f_t(\tilde{P}_{A_t}, i) \neq \infty$, $C \subseteq C'$, and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$ then we need to show that there is a nice decomposition $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ in \mathcal{D} corresponding to some $X \in \mathcal{S}$ such that $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}'_{\chi(t)}$ and $O \neq \emptyset$. For this we first prove the following claim that talks about some properties of \tilde{R} which will be helpful for identifying the required X .

Claim 5.2. *If $C \subseteq C'$, and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$, then \tilde{R} will satisfy the following properties.*

(P1) \tilde{R} is a refinement of $\tilde{P}'_{V_T^{\chi(t)}}$.

(P2) Projection of \tilde{R} on $S_{V_T^{\chi(t)}}$ will be the same as that of $\tilde{P}'_{V_T^{\chi(t)}}$ on $S_{V_T^{\chi(t)}}$.

(P3) Each part R in \tilde{R} has all vertices from either $S_{V_T^{\chi(t)}}$ or from $V_T^{\chi(t)} \setminus S_{V_T^{\chi(t)}}$, but not both.

Proof. Since $C \subseteq C'$ and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$, \tilde{R} is a refinement of $\tilde{P}'_{V_T^{\chi(t)}}$ and the projection of \tilde{R} on $S_{V_T^{\chi(t)}}$ will be the same as that of $\tilde{P}'_{V_T^{\chi(t)}}$ on $S_{V_T^{\chi(t)}}$. Recall that the partition $\tilde{P}'_{V_T^{\chi(t)}}$ is the partition of $V_T^{\chi(t)}$ obtained by removing C from $\text{proj}(T, \chi(t))$. Since, each part in $\tilde{P}'_{V_T^{\chi(t)}}$ has vertices from either $S_{V_T^{\chi(t)}}$ or from $V_T^{\chi(t)} \setminus S_{V_T^{\chi(t)}}$, but not both, the same holds for the parts in \tilde{R} . \square

Let S' be the parts in \tilde{R} having vertices from $S_{V_T^{\chi(t)}}$. By (P2) and (P3), S' is a subset of S which are the parts in \tilde{R} and S' is of size at most $2k-1$. Also, let N be the parts in \tilde{R} which are not in S' and have cross edges to some part in S' or have non empty intersection with an adhesion $A_{t'}$, $t' \in \text{children}(t) \cup t$ that also has non empty intersection with some part in S' . That is N are the parts in $S \setminus S'$ that are adjacent to the parts in S' in G' . We now prove that every part P in S' has at most $s^2 + 2s + k$ neighbours in G' .

Each adhesion $A_{t'}$, $t' \in \text{children}(t)$ that intersects different parts in S can be associated with at least one unique cross edge in $\tilde{P}(\tilde{P}_{A_t}, i)$ from $E(G_{t'}) \setminus E(G[A_{t'}])$ since (τ, χ) is a compact tree decomposition and thus the graph obtained by removing the edges between any two vertices in $A_{t'}$ from $G_{t'}$ is connected. Since the number of vertices in an adhesion is at most s , the number of parts that P can share an adhesion with is at most s^2 . The number of cross edges in $\tilde{P}(\tilde{P}_{A_t}, i)$ are at most s , thus the number of parts that P can share an edge in $E(G_t)$ with are at most s . In addition, the adhesion A_t cannot be associated with a unique cross edge as it is the bag t for

which we are computing the state. Thus, P can be adjacent to the parts that have vertices from A_t and there are at most s of these parts. Finally, since the small part that P was contained in $\tilde{P}_{V_T^{\chi(t)}}$ is broken down into at most k parts in $\tilde{P}'_{V_T^{\chi(t)}}$, at most k edges can be saved so P can have at most k extra neighbours. Thus, P has at most $s^2 + 2s + k$ neighbours in G' . Therefore, since $|S'| \leq 2k - 1$, $N \leq (2k - 1)(s^2 + 2s + k)$.

Thus from Lemma 2.1, there exists $X \in \mathcal{S}$ such that $S' \subseteq X$ and $S' \cap N = \emptyset$. In \tilde{Q} , all parts that are not in X are contained in Q_1 because of step (b). Therefore, all parts in N are contained in Q_1 and the connected components in G' either have all its parts from S' or no parts from S' . As $|S'| \leq 2k - 1$, all connected components in \tilde{Q} having more than $2k - 1$ components contain no parts from S' . Thus \tilde{Q}' has all the parts from S' in it, and therefore $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}_{\chi(t)}$.

To complete the proof, we need to prove that $O \neq \emptyset$. Observe that each part in $\tilde{P}'_{\chi(t)}$ has all vertices from $S_{\chi(t)}$ or all from $\chi(t) \setminus S_{\chi(t)}$. Since we are in the case when the size of $\chi(t)$ is not of $\mathcal{O}((s + 1)^5)$, $\chi(t) \setminus S_{\chi(t)} \neq \emptyset$. Thus the number of parts in $\tilde{P}'_{\chi(t)}$ is more than one. If $O = \emptyset$, then since there are no edges or common adhesions between different parts in $\tilde{P}'_{\chi(t)}$, it means that G_t is disconnected. But because (τ, χ) is a compact tree decomposition, G_t is connected, and therefore O cannot be an empty set. This completes the proof of the claim. \square

Claim 5.3. \mathcal{D} is of size $\mathcal{O}(s^{\mathcal{O}(k)} \log n)$ and is computed in time $\mathcal{O}(s^{\mathcal{O}(k)} n^{\mathcal{O}(1)})$.

Proof. In the case when the size of $\chi(t)$ is of $\mathcal{O}((s + 1)^5)$, \mathcal{D} is of size at most one and is computed in polynomial time. In the other case, \mathcal{S} is of size $\mathcal{O}(s^{\mathcal{O}(k)} \log n)$ from Lemma 2.1. Since for each $X \in \mathcal{S}$ we add at most one nice decomposition to set \mathcal{D} in step (vi), \mathcal{D} is of size $\mathcal{O}(ks)^{\mathcal{O}(k)} \log n$. Time taken to compute \mathcal{S} in step (v) is $\mathcal{O}(s^{\mathcal{O}(k)} n^{\mathcal{O}(1)})$ from Lemma 2.1. All other steps take polynomial time, in particular, for each set $X \in \mathcal{S}$ it takes polynomial time to compute $\tilde{P}'_{\chi(t)}$, $\tilde{Q}_{\chi(t)}$ and O and add it to \mathcal{D} . Thus, the time taken to compute \mathcal{D} in this case is $\mathcal{O}(s^{\mathcal{O}(k)} n^{\mathcal{O}(1)})$. \square

Claims 5.2 and 5.3 complete the proof of Lemma 5.5. \square

With this, we start the second part of our section, an algorithm that takes as input a nice decomposition of $\chi(t)$ along with other inputs and outputs a value greater than or equal to $f_t(\tilde{P}_{A_t}, i)$. Further if the partition $\tilde{Q}_{\chi(t)}$ in the nice decomposition satisfies a special property, then the algorithm will output $f_t(\tilde{P}_{A_t}, i)$. We now state this as a Lemma.

Lemma 5.6. *There exists an algorithm that takes as input as input (τ, χ) , a node $t \in V(\tau)$, a T -feasible partition \tilde{P}_{A_t} of A_t , a positive integer $i \leq k$, a nice decomposition $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ of $\chi(t)$, together with the value of $f_{t'}(\tilde{P}'_{A_{t'}}, i')$ for every child t' of t , T -feasible partition $\tilde{P}'_{A_{t'}}$ of $A_{t'}$, and positive integer $i' \leq i$, and returns a positive integer v such that $f_t(\tilde{P}_{A_t}, i) \leq v$ or $v = \infty$ in time $k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. Furthermore, if $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}_{\chi(t)}$, then $v \leq f_t(\tilde{P}_{A_t}, i)$.*

To prove Lemma 5.6, we design a dynamic programming algorithm on a given nice decomposition $(\tilde{P}'_{\chi(t)}, \tilde{Q}_{\chi(t)}, O)$ of $\chi(t)$. Let p denote the number of parts in $\tilde{P}'_{\chi(t)} \setminus O$, it is easy to see that by the definition of a nice decomposition, if $O = \emptyset$, then $p = 1$ and if not then $p = |\tilde{P}'_{\chi(t)}| - 1$. We arbitrarily order all the parts in $\tilde{P}'_{\chi(t)}$ except O and denote by P_l the l^{th} part in this order, where $l \geq 1$ and denote by $P_{\leq l} = \bigcup_{x \leq l} P_x$, the union of all parts P_x , $x \leq l$.

Given a non negative integer l that is less than or equal to p , we define the set $\mathcal{A}(l)$ to be the set of all adhesions in the set $\{A_{t'} : t' \in \text{children}(t)\}$ that have vertices only from $P_l \cup O$ and have non empty intersection with P_l . Also, we denote by $\mathcal{A}_{\leq}(l) = \bigcup_{0 \leq l' \leq l} \mathcal{A}(l')$, the union

of all sets $\mathcal{A}(l')$ where $0 \leq l' \leq l$. Further, we define the graph $G(l) = G[O \cup P_l] \cup \bigcup_{A_{t'} \in \mathcal{A}(l)} G_{t'}$, to be the subgraph of G induced by all the vertices in $O \cup P_l \cup \bigcup_{A_{t'} \in \mathcal{A}(l)} V(G_{t'})$ and define the graph $G_{\leq}(l) = G[O \cup P_{\leq l}] \cup \bigcup_{A_{t'} \in \mathcal{A}_{\leq}(l)} G_{t'}$, to be the subgraph of G induced by all the vertices in $O \cup P_{\leq l} \cup \bigcup_{A_{t'} \in \mathcal{A}_{\leq}(l)} V(G_{t'})$.

We define two function $h, g : \{0, \dots, p\} \times \{1, \dots, i\} \longrightarrow \{0, \dots, s\} \cup \{\infty\}$ that our algorithm will compute. The domain of h and g consists of all pairs (l, j) where l is a non negative integer less than or equal to p and j is a positive integer less than or equal to the input integer i . On input (l, j) , h returns the smallest possible weight of a j -partition \tilde{P} of $G(l)$ such that the projection of \tilde{P} on $O \cup P_l$ is refined by the projection of $\tilde{Q}_{\chi(t)}$ on $O \cup P_l$ and if $A_t \subseteq O \cup P_l$ then \tilde{P}_{A_t} is the projection of \tilde{P} on A_t and g returns the smallest possible weight of a j -partition \tilde{P} of $G_{\leq}(l)$ such that the projection of \tilde{P} on $O \cup P_{\leq l}$ is refined by the projection of $\tilde{Q}_{\chi(t)}$ on $O \cup P_{\leq l}$ and if $A_t \subseteq O \cup P_{\leq l}$ then \tilde{P}_{A_t} is the projection of \tilde{P} on A_t . However, if this weight is greater than s , or no such partition exists then both $h(l, j)$ and $g(l, j)$ return ∞ .

The main step of an algorithm for Lemma 5.6 is an algorithm that computes $g(l, j)$ for every pair $(l, j) \in \{0, \dots, p\} \times \{1, \dots, i\}$ assuming that the value of $g(l', j')$ has been computed for every pair $(l', j') \in \{0, \dots, l-1\} \times \{1, \dots, i\}$. We now state that this step can be carried out.

Lemma 5.7. *There exists an algorithm that takes all the inputs of Lemma 5.6, along with a non negative integer $l \leq p$, a positive integer $j \leq i$ and if $l \geq 1$ the value of $g(l-1, j')$ for every positive integer $j' \leq j$ and returns a positive integer $v = g(l, j)$ in time $k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$.*

Observe that $g(p, i)$ is the smallest possible weight of an i -partition \tilde{P} of G_t such that $\tilde{Q}_{\chi(t)}$ is a refinement of the projection of \tilde{P} on $\chi(t)$ and the projection of \tilde{P} on A_t is \tilde{P}_{A_t} . By definition, $f_t(\tilde{P}_{A_t}, i)$ is the weight of the smallest possible i -partition whose projection on A_t is \tilde{P}_{A_t} , thus it follows that $g(p, i) \geq f_t(\tilde{P}_{A_t}, i)$. If $f_t(\tilde{P}_{A_t}, i) \neq \infty$ and $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}_{\chi(t)}$ then $g(p, i) \leq f_t(\tilde{P}_{A_t}, i)$ since by definition, $\tilde{P}(\tilde{P}_{A_t}, i)$ is a i -partition of G_t that has weight $f_t(\tilde{P}_{A_t}, i)$, whose projection on $\chi(t)$ is $\tilde{P}_{\chi(t)}$, and whose projection on A_t is \tilde{P}_{A_t} . We state these relations between $g(p, i)$ and $f_t(\tilde{P}_{A_t}, i)$ in the following observation.

Observation 5.3. *$g(p, i)$ is always greater than or equal to $f_t(\tilde{P}_{A_t}, i)$ and if $f_t(\tilde{P}_{A_t}, i) \neq \infty$ and $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}_{\chi(t)}$ then $g(p, i)$ is less than or equal to $f_t(\tilde{P}_{A_t}, i)$.*

We are now ready to prove Lemma 5.6 using Lemma 5.7 and Observation 5.3.

Proof of Lemma 5.6 assuming Lemma 5.7. For proving Lemma 5.6 we propose the following algorithm. The algorithm will take inputs as specified in the Lemma, compute the function g using Lemma 5.7 and finally returns an integer $v = g(p, i)$. From Observation 5.3, it follows that the algorithm outputs a value $v \geq f_t(\tilde{P}_{A_t}, i)$. Further if $f_t(\tilde{P}_{A_t}, i) \neq \infty$ and $\tilde{Q}_{\chi(t)}$ is a refinement of $\tilde{P}_{\chi(t)}$ then $v \leq f_t(\tilde{P}_{A_t}, i)$. Since the domain of g has $p \cdot i$ values, $p \leq n$ and $i \leq k$, from Lemma 5.7 it is easy to see that the algorithm runs in time $k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. \square

Since, $G_{\leq}(l) = G_{\leq}(l-1) \cup G(l)$ and all common edges of $G_{\leq}(l-1)$ and $G(l)$ are in $G(0)$, it follows that for any pair $(l, j) \in \{0, \dots, p\} \times \{1, \dots, i\}$ the following equation holds if $O \neq \emptyset$.

$$g(l, j) = \begin{cases} \min_{1 \leq j' \leq j} g(l-1, j') + h(l, j-j'+1) & \text{if } l \geq 1 \\ h(l, j) & \text{if } l = 0 \end{cases} \quad (1)$$

If $O = \emptyset$, then there is only one part in $\tilde{P}'_{\chi(t)}$ of the nice decomposition and thus we can compute $g(1, j) = h(1, j)$, for all $j \leq i$ directly.

We are now ready to prove Lemma 5.7, the proof will have an algorithm that computes $h(l, j')$ efficiently for all $j' \leq j$.

Proof of Lemma 5.7. Let $\mathcal{R}(l, j)$ be the set of partitions \tilde{R} of $O \cup P_l$ having at most j components such that the projection of $\tilde{Q}_{\chi(t)}$ on $O \cup P_l$ is a refinement of \tilde{R} and if $A_t \subseteq P_{\leq l}$, then \tilde{P}_{A_t} is the projection of \tilde{R} on A_t .

We define the function $h' : \{1, \dots, j\} \times \mathcal{R}(l, j) \rightarrow \{0, \dots, s\} \cup \{\infty\}$ that our algorithm will compute. On input (j', \tilde{R}) , where $1 \leq j' \leq j$ and $\tilde{R} \in \mathcal{R}(l, j)$, h' returns the smallest possible weight of a j' -partition \tilde{P} of $G(i)$ such that the projection of \tilde{P} on $O \cup P_l$ is \tilde{R} .

Clearly, for all positive integers $j' \leq j$, $h(l, j')$ satisfies the following equation.

$$h(l, j') = \min_{\tilde{R} \in \mathcal{R}(l, j)} h'(j', \tilde{R}) \quad (2)$$

Thus, to complete the proof, it is sufficient to compute $h'(j', \tilde{R})$ for all $j' \leq j$ given \tilde{R} . For this, we design a knapsack style dynamic programming algorithm on the set $\mathcal{A}(l)$ of adhesions associated with P_l . Arbitrarily order the adhesions in the set $\mathcal{A}(l)$ and let a be a positive integer less than or equal to $|\mathcal{A}(l)|$, let t_a denote the child of t such that A_{t_a} is the a^{th} element in $\mathcal{A}(l)$. Let $G(l, a)$ denote the graph obtained by removing the edges among the vertices in A_{t_a} from the graph $G[O \cup P_l] \cup G_{t_a}$ and let $G(l, 0)$ denote the graph $G[O \cup P_l]$. Also, let $G_{\leq}(l, a) = \bigcup_{0 \leq a' \leq a} G(l, a')$.

We define two function $\zeta, \nu : \{0, \dots, |\mathcal{A}(l)|\} \times \{1, \dots, j\} \rightarrow \{0, \dots, s\} \cup \{\infty\}$ that our algorithm will compute. The domain of ζ and ν consists of all pairs (a, r) where a is a non negative integer less than or equal to $|\mathcal{A}(l)|$ and r is a positive integer less than or equal to the input integer j . On input (a, r) , ζ returns the smallest possible weight of a r -partition \tilde{P} of $G(l, a)$ such that the projection of \tilde{P} on $O \cup P_l$ is \tilde{R} and ν returns the smallest possible weight of a r -partition \tilde{P} of $G_{\leq}(l, a)$ such that the projection of \tilde{P} on $O \cup P_l$ is \tilde{R} . However, if this weight is greater than s , or no such partition exists then both $\zeta(a, r)$ and $\nu(a, r)$ return ∞ .

It is easy to see that $h'(j', \tilde{R})$ satisfies the following equation just by the definition of ν .

$$h'(j', \tilde{R}) = \nu(|\mathcal{A}(l)|, j') \quad (3)$$

Since $G_{\leq}(l, a) = G_{\leq}(l, a-1) \cup G(l, a)$ and there are no common edges between $G_{\leq}(l, a-1)$ and $G(l, a)$, for any pair $(a, r) \in \{0, \dots, |\mathcal{A}(l)|\} \times \{1, \dots, j\}$, the following equation holds, where $\tilde{R}_{A_{t_a}}$ denotes the projection of \tilde{R} on the adhesion A_{t_a} .

$$\nu(a, r) = \begin{cases} \min_{0 \leq r' \leq r} \nu(a-1, r') + \zeta(a, r-r' + |\tilde{R}_{A_{t_a}}|) & \text{if } a \geq 1 \\ w(\tilde{R}) & \text{if } a = 0 \text{ and } |\tilde{R}| = r \\ \infty & \text{otherwise} \end{cases} \quad (4)$$

Observe that $\zeta(a, r-r' + |\tilde{R}_{A_{t_a}}|)$ can be computed by the following equation just by the definition of ζ and f ,

$$\zeta(a, r-r' + |\tilde{R}_{A_{t_a}}|) = f_{t_a}(\tilde{R}_{A_{t_a}}, r-r' + |\tilde{R}_{A_{t_a}}|) - w(\tilde{R}_{A_{t_a}}) \quad (5)$$

Combining equations (3), (4) and (5), it is easy to see that given $j' \leq j$ and $\tilde{R} \in \mathcal{R}(l, j)$, $h'(j', \tilde{R})$ can be computed in $\mathcal{O}(nk)$ time. The size of $\mathcal{R}(l, j)$ is at most j^{2k} since from properties (i) and (ii) of Definition 5.2, it follows that the projection of $\tilde{Q}_{\chi(t)}$ on $P_l \cup O$ has at most $2k$ parts. Thus, from equation (2), it follows that $h(l, j')$ can be computed in time $k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. Combining this with equation (1), it is clear that $g(l, j)$ can be computed in time $k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. This completes the proof. \square

We now are ready to complete the proof of Lemma 5.4 for which we obtained results throughout this section.

Proof of Lemma 5.4. For the proof, we propose the following algorithm. Firstly, the algorithm will take inputs as specified in Lemma 5.4 and obtain a set of nice decompositions \mathcal{D} from Lemma 5.5. Then for each nice decomposition $D \in \mathcal{D}$, the algorithm obtains a value v_D from Lemma 5.6. Finally it returns the positive integer $v = \min_{D \in \mathcal{D}} v_D$. If $\mathcal{D} = \emptyset$, then return $v = \infty$. Combining Lemma 5.5 and Lemma 5.6, it is easy to see that the algorithm will return $v \geq f_t(\tilde{P}_{A_t}, i)$. Further if $f_t(\tilde{P}_{A_t}, i) \neq \infty$, $C \subseteq C'$, and $S_{E_T^{\chi(t)}} \cap C' = \emptyset$, then it will return $v = f_t(\tilde{P}_{A_t}, i)$. The time taken by the algorithm is $\mathcal{O}((s)^{\mathcal{O}(k)} n^{\mathcal{O}(1)})$ since \mathcal{D} has at most $\mathcal{O}((s)^{\mathcal{O}(k)} \log n)$ sets and each set can be computed in $k^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$ time from Lemma 5.5 and Lemma 5.6. This proves Lemma 5.4 \square

6 Combining all the Pieces: Proof of Theorem 1.1

In this section we conjure all the pieces we obtained so far and give a proof of our main result (Theorem 1.1).

Proof of Theorem 1.1. Let G, k, ϵ be the input to MIN k -CUT. Also let $|V(G)| = n$. We will output a partition \tilde{P} with weight v such that $v \leq (1 + \epsilon)\text{OPT}(G, k)$ with probability at least $1 - \frac{1}{n^{26}}$. We first check if $\text{cc}(G) \geq k$. If this is true then we know that G has optimum k -cut of value 0 and we can return G itself. If $\epsilon < \frac{1}{n}$ then we run the best known exact algorithm on (G, k) running in time $n^{\mathcal{O}(k)}$ and return the exact answer [KS96, Tho08, GLL20]. Clearly, $n^{\mathcal{O}(k)}$ is $(k/\epsilon)^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$ in this case. So from now onwards we assume that $\text{cc}(G) \leq k$ and $\epsilon \geq \frac{1}{n}$. Also, we set $\epsilon' = \frac{\epsilon}{10}$.

Now we apply a standard rounding procedure (similar to the well-known Knapsack PTAS [KT06]) that reduces the problem in a $(1 + \epsilon')$ -approximation preserving manner to an unweighted multi-graph with at most m^2/ϵ edges. This implies that now the graph has at most n^5 edges (counting multiplicities). Let this graph be G^* . Next we apply Lemma 3.1 and obtain a subgraph G_1 of G^* with $V(G_1) = V(G^*)$, such that $q = |E(G^*)| - |E(G_1)| \leq 2\epsilon' \cdot \text{OPT}(G^*, k)$. Further, if $\text{cc}(G_1) < k$, then each non-trivial 2-cut of G_1 has weight at least $\frac{\epsilon' \cdot \text{OPT}(G^*, k)}{k-1}$. If $\text{cc}(G_1) = k$, then we return the connected components as a partition \tilde{P} . The cost of returned solution is

$$v \leq 2\epsilon' \cdot \text{OPT}(G^*, k) \leq 2\epsilon' \cdot (1 + \epsilon')\text{OPT}(G, k) \leq (1 + \epsilon)\text{OPT}(G, k).$$

Thus, we assume that $\text{cc}(G_1) < k$ and each non-trivial 2-cut of G_1 has weight at least $\frac{\epsilon' \cdot \text{OPT}(G^*, k)}{k-1}$. Now we apply the sparsification procedure described in Lemma 3.2 and obtain a subgraph G_2 with $V(G_2) = V(G_1)$, and a real number r such that with probability at least $1 - \frac{1}{n^{26}}$, for all k -cuts \tilde{P} in G_1 , $(1 - \epsilon') \cdot w(G_1, \tilde{P}) \leq w(G_2, \tilde{P}) \cdot r \leq (1 + \epsilon') \cdot w(G_1, \tilde{P})$. Here, $r = 1/p$, where $p = \frac{100 \cdot \log n}{\epsilon'^2 \cdot \text{OPT}(G_1, 2)}$. However, in G_2 we know that

$$\begin{aligned} \text{OPT}(G_2, k) &\leq (1 + \epsilon') \cdot p \cdot \text{OPT}(G_1, k) \\ &= \frac{(1 + \epsilon') 100 \log n}{\text{OPT}(G_1, 2) \cdot \epsilon'^2} \cdot \text{OPT}(G_1, k) \\ &\leq \frac{(1 + \epsilon') 100 \log n}{\text{OPT}(G_1, 2) \cdot \epsilon'^2} \cdot \text{OPT}(G^*, k) \\ &\leq \frac{(1 + \epsilon') \cdot k \cdot 100 \log n}{\epsilon'^3} \\ &\leq \frac{\beta \cdot k \cdot 100 \log n}{\epsilon^3}. \end{aligned}$$

Here, β is a fixed constant. Further, in the second last transition we used the assumption that $\text{OPT}(G_1, 2) \geq \epsilon' \cdot \text{OPT}(G^*, k)/k - 1$. Now we apply the Theorem 1.2 on G_2 and for each

$s \in \left\{1, \dots, \frac{\beta \cdot k \cdot 100 \log n}{\epsilon^3}\right\}$ and solve the problem exactly in time $s^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. Thus the running time of the algorithm is upper bounded by

$$\begin{aligned} \left(\frac{\beta \cdot k \cdot 100 \log n}{\epsilon^3}\right)^{\mathcal{O}(k)} n^{\mathcal{O}(1)} &= (k/\epsilon)^{\mathcal{O}(k)} (\log n)^{\mathcal{O}(k)} n^{\mathcal{O}(1)} \leq (k/\epsilon)^{\mathcal{O}(k)} (k^{\mathcal{O}(k)} + n) n^{\mathcal{O}(1)} \\ &\leq 2^{\mathcal{O}(k \log(\frac{k}{\epsilon}))} n^{\mathcal{O}(1)}. \end{aligned}$$

This completes the running time analysis. All that remains is to show that what we get is an $(1 + \epsilon)$ -approximation algorithm. Let s be the minimum value for which the algorithm described in Theorem 1.2 returns yes. This implies that $s = \text{OPT}(G_2, k)$. Let \tilde{P} be the corresponding partition. We return $s \cdot r + q$ as value of the cut and \tilde{P} as a solution. Here, the value is the sum of edges deleted when we obtained G_1 from G^* and the value returned by Theorem 1.2 when ran on (G_2, k, s) . Now we have that

$$\begin{aligned} \text{OPT}(G_2, k) \cdot r + q &\leq (1 + \epsilon') \cdot \text{OPT}(G_1, k) + 2\epsilon' \cdot \text{OPT}(G^*, k) \\ &\leq (1 + \epsilon') \cdot \text{OPT}(G^*, k) + 2\epsilon' \cdot \text{OPT}(G^*, k) \\ &\leq (1 + 3\epsilon') \cdot \text{OPT}(G^*, k) \\ &\leq (1 + 3\epsilon') \cdot (1 + \epsilon') \cdot \text{OPT}(G, k) \\ &\leq (1 + \epsilon) \cdot \text{OPT}(G, k) \end{aligned}$$

The correctness of the algorithm follows from Lemmas 3.1, 3.2 and Theorem 1.2. This concludes the proof. \square

7 Conclusion

In this paper we gave a parameterized approximation algorithm with best possible approximation guarantee, and best possible running time dependence on said guarantee (upto ETH and constants in the exponent) for MIN k -CUT. In particular, for every $\epsilon > 0$, the algorithm computes a $(1 + \epsilon)$ -approximate solution in time $(k/\epsilon)^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$. Along the way we also obtained a new exact algorithm with running time $s^{\mathcal{O}(k)} n^{\mathcal{O}(1)}$ on unweighted (multi-) graphs, where s denotes the number of edges in a minimum k -cut. For an even more complete understanding of the parameterized approximation complexity of MIN k -CUT one could explore the possibility of lossy kernels of polynomial size in s or k , and ratios between $2 - \epsilon$ and $1 + \epsilon$. Finally, an intriguing open problem about the parameterized complexity of MIN k -CUT is whether the problem admits an algorithm with running time $2^{\mathcal{O}(s)} n^{\mathcal{O}(1)}$.

References

- [AWW14] Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. Consequences of faster alignment of sequences. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 39–51. Springer, 2014. 2
- [BD02] Patrick Bellenbaum and Reinhard Diestel. Two short proofs concerning tree-decompositions. *Combinatorics, Probability & Computing*, 11(6):541–547, 2002. 10

- [BK15] András A. Benczúr and David R. Karger. Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM J. Comput.*, 44(2):290–319, 2015. [2](#), [8](#)
- [BP16] Mikołaj Bojańczyk and Michal Pilipczuk. Definability equals recognizability for graphs of bounded treewidth. *CoRR*, abs/1605.03045, 2016. Extended abstract appeared in the proceedings of LICS 2016. [14](#)
- [BT17] Arturs Backurs and Christos Tzamos. Improving viterbi is hard: Better runtimes imply faster clique algorithms. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 311–321. PMLR, 2017. [2](#)
- [CCH⁺16] Rajesh Chitnis, Marek Cygan, MohammadTaghi Hajiaghayi, Marcin Pilipczuk, and Michał Pilipczuk. Designing FPT algorithms for cut problems using randomized contractions. *SIAM J. Comput.*, 45(4):1171–1229, 2016. [3](#), [18](#)
- [CFK⁺15] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. [1](#), [2](#)
- [CKL⁺18] Marek Cygan, Pawel Komosa, Daniel Lokshtanov, Michal Pilipczuk, Marcin Pilipczuk, Saket Saurabh, and Magnus Wahlstrom. Randomized contractions meet lean decompositions. *CoRR*, abs/1810.06864, 2018. [1](#), [3](#), [4](#), [10](#), [14](#), [18](#)
- [DEF⁺03] Rodney G. Downey, Vladimir Estivill-Castro, Michael R. Fellows, Elena Prieto-Rodriguez, and Frances A. Rosamond. Cutting up is hard to do: the parameterized complexity of k -cut and related problems. *Electron. Notes Theor. Comput. Sci.*, 78:209–222, 2003. [1](#), [2](#)
- [Die12] Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012. [4](#)
- [GH94] Olivier Goldschmidt and Dorit S. Hochbaum. A polynomial algorithm for the k -cut problem for fixed k . *Math. Oper. Res.*, 19(1):24–37, 1994. [1](#)
- [GLL18a] Anupam Gupta, Euiwoong Lee, and Jason Li. Faster exact and approximate algorithms for k -cut. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 113–123. IEEE Computer Society, 2018. [1](#), [2](#)
- [GLL18b] Anupam Gupta, Euiwoong Lee, and Jason Li. An FPT algorithm beating 2-approximation for k -cut. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2821–2837. SIAM, 2018. [1](#), [2](#)
- [GLL19] Anupam Gupta, Euiwoong Lee, and Jason Li. The number of minimum k -cuts: improving the karger-stein bound. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 229–240. ACM, 2019. [1](#)
- [GLL20] Anupam Gupta, Euiwoong Lee, and Jason Li. The number of minimum k -cuts: Improving the karger-stein bound. *To appear in STOC 2020*, abs/1906.00417, 2020. [1](#), [27](#)

- [Kar94] David R. Karger. Using randomized sparsification to approximate minimum cuts. In Daniel Dominic Sleator, editor, *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, 23-25 January 1994, Arlington, Virginia, USA*, pages 424–432. ACM/SIAM, 1994. [2](#), [8](#), [9](#)
- [KL20] Ken-ichi Kawarabayashi and Bingkai Lin. A nearly $5/3$ -approximation FPT algorithm for min- k -cut. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 990–999. SIAM, 2020. [1](#), [2](#)
- [KS96] David R. Karger and Clifford Stein. A new approach to the minimum cut problem. *J. ACM*, 43(4):601–640, 1996. [1](#), [2](#), [3](#), [27](#)
- [KT06] Jon M. Kleinberg and Éva Tardos. *Algorithm design*. Addison-Wesley, 2006. [2](#), [12](#), [27](#)
- [KT11] Ken-ichi Kawarabayashi and Mikkell Thorup. The minimum k -way cut of bounded size is fixed-parameter tractable. In *52nd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 160–169. IEEE Computer Society, 2011. [3](#)
- [Li19] Jason Li. Faster minimum k -cut of a simple graph. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1056–1077. IEEE Computer Society, 2019. [1](#)
- [Man18] Pasin Manurangsi. Inapproximability of maximum biclique problems, minimum k -cut and densest at-least- k -subgraph from the small set expansion hypothesis. *Algorithms*, 11(1):10, 2018. [1](#), [2](#)
- [Mar06] Dániel Marx. Parameterized graph separation problems. *Theor. Comput. Sci.*, 351(3):394–406, 2006. [3](#)
- [MU05] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005. [7](#)
- [NR01] Joseph Naor and Yuval Rabani. Tree packing and approximating k -cuts. In S. Rao Kosaraju, editor, *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, January 7-9, 2001, Washington, DC, USA*, pages 26–27. ACM/SIAM, 2001. [1](#), [2](#), [8](#)
- [NSS95] Moni Naor, Leonard J. Schulman, and Aravind Srinivasan. Splitters and near-optimal derandomization. In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23-25 October 1995*, pages 182–191. IEEE Computer Society, 1995. [6](#)
- [RS08] R Ravi and Amitabh Sinha. Approximating k -cuts using network strength as a lagrangean relaxation. *European Journal of Operational Research*, 186(1):77–90, 2008. [1](#), [2](#), [8](#)
- [SV95] Huzur Saran and Vijay V. Vazirani. Finding k cuts within twice the optimal. *SIAM J. Comput.*, 24(1):101–108, 1995. [1](#), [2](#), [8](#)
- [Tho90] Robin Thomas. A Menger-like property of tree-width: The finite case. *J. Comb. Theory, Ser. B*, 48(1):67–76, 1990. [10](#)

- [Tho08] Mikkel Thorup. Minimum k -way cuts via deterministic greedy tree packing. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 159–166. ACM, 2008. [1](#), [4](#), [5](#), [15](#), [17](#), [27](#)