

Improved Approximation for Tree Augmentation: Saving by Rewiring

Fabrizio Grandoni
IDSIA
Lugano, Switzerland
fabrizio@idsia.ch

Christos Kalaitzis
ETH Zurich
Zurich, Switzerland
chistos.kalaitzis@ifor.math.ethz.ch

Rico Zenklusen
ETH Zurich
Zurich, Switzerland
ricoz@math.ethz.ch

ABSTRACT

The Tree Augmentation Problem (TAP) is a fundamental network design problem in which we are given a tree and a set of additional edges, also called *links*. The task is to find a set of links, of minimum size, whose addition to the tree leads to a 2-edge-connected graph. A long line of results on TAP culminated in the previously best known approximation guarantee of 1.5 achieved by a combinatorial approach due to Kortsarz and Nutov [ACM Transactions on Algorithms 2016], and also by an SDP-based approach by Cheriyan and Gao [Algorithmica 2017]. Moreover, an elegant LP-based $(1.5 + \epsilon)$ -approximation has also been found very recently by Fiorini, Groß, Könemann, and Sanitá [SODA 2018]. In this paper, we show that an approximation factor below 1.5 can be achieved, by presenting a 1.458-approximation that is based on several new techniques.

By extending prior results of Adjiashvili [SODA 2017], we first present a black-box reduction to a very structured type of instance, which played a crucial role in recent development on the problem, and which we call k -wide. Our main contribution is a new approximation algorithm for $O(1)$ -wide tree instances with approximation guarantee strictly below 1.458, based on one of their fundamental properties: wide trees naturally decompose into smaller subtrees with a constant number of leaves. Previous approaches in similar settings rounded each subtree independently and simply combined the obtained solutions. We show that additionally, when starting with a well-chosen LP, the combined solution can be improved through a new “rewiring” technique, showing that one can replace some pairs of used links by a single link. We can rephrase the rewiring problem as a stochastic version of a matching problem, which may be of independent interest. By showing that large matchings can be obtained in this problem, we obtain that a significant number of rewirings are possible, thus leading to an approximation factor below 1.5.

CCS CONCEPTS

• **Mathematics of computing** → **Combinatorial optimization; Approximation algorithms**; • **Theory of computation** → **Routing and network design problems**;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC'18, June 25–29, 2018, Los Angeles, CA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5559-9/18/06...\$15.00

<https://doi.org/10.1145/3188745.3188898>

KEYWORDS

Tree augmentation, approximation algorithms, combinatorial optimization

ACM Reference Format:

Fabrizio Grandoni, Christos Kalaitzis, and Rico Zenklusen. 2018. Improved Approximation for Tree Augmentation: Saving by Rewiring. In *Proceedings of 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC'18)*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3188745.3188898>

1 INTRODUCTION

In the Tree Augmentation Problem (TAP) we are given an undirected tree $T = (V, E)$ and a set of links $L \subseteq \binom{V}{2}$ between pairs of vertices. The task is to find a minimum cardinality set of links $U \subseteq L$ such that $(V, E \cup U)$ is 2-edge-connected. A natural extension of TAP is its weighted version, the Weighted Tree Augmentation Problem (WTAP), in which each link $\ell \in L$ has a nonnegative cost c_ℓ , and the task is to choose a set of links $U \subseteq L$ of minimum cost such that $(V, E \cup U)$ is 2-edge-connected.

TAP is a basic Survivable Network Design problem, i.e., a problem where our goal is to enhance the connectivity properties of an input network in order to make it tolerant to edge faults. The motivation behind studying such problems is clear: networks are not static, but evolve over time, and in particular they are subject to faults that result in edges or whole sub-networks becoming unavailable. Therefore, it is desirable that we design networks in such a way that the failure of a small number of network interfaces does not disrupt its connectivity. Specifically, TAP poses one of the arguably most basic questions in this context: given a network whose connectivity is disrupted even if a single link fails, how can we introduce as few links as possible in such a way, that a single link failure will still result in a connected network? In this sense, it is a special case of a much more general problem: given a k -edge-connected network, how can we introduce new links such that the resulting network is $(k + 1)$ -edge-connected? Interestingly enough, whenever k is odd, this problem reduces to TAP (see Cheriyan, Jordán, and Ravi [5]).

Given the importance of the problem, it is not surprising that TAP, WTAP, and even special cases thereof attracted considerable interest. To begin with, WTAP was shown to be NP-hard in the 80's by Frederickson and Jájá [10], and also TAP was proven to be NP-hard later on (see Cheriyan, Jordán, and Ravi [5]). Unsurprisingly, considerable effort has been put into designing approximation algorithms.

For WTAP, the best-known approximation guarantee is 2 and was first established by Frederickson and Jájá [10]. Their algorithm was later simplified by Khuller and Thurimella [14]. A 2-approximation can also be achieved by various other techniques developed later on, including a primal-dual technique by Goemans, Goldberg, Plotkin, Shmoys, Tardos, and Williamson [11], and the iterative rounding technique by Jain [13]. For WTAP, improvements on the factor 2 have only been obtained for restricted cases, including bounded diameter trees (see Cohen and Nutov [7]), and instances where the ratio Δ between the highest and smallest cost is bounded by a constant (see Adjashvili [1], and Fiorini, Groß, Könemann, and Sanità [9]). Moreover, very recently, Nutov [20] showed that even if $\Delta = O(\log n)$, where n is the number of vertices, the factor of 2 can be beaten.

Regarding TAP, the first algorithm beating the approximation guarantee of 2 is due to Nagamochi [19], achieving an approximation factor of $1.815 + \epsilon$. This factor was subsequently improved to 1.8 by Even, Feldman, Kortsarz, and Nutov [8]. The best-known approximation guarantee is 1.5, first obtained by Kortsarz and Nutov [17]. These results are combinatorial in nature, displaying a contrast with the weighted version of the problem, where the majority of the results come from LP-based algorithms. With respect to dulling this contrast, Kortsarz and Nutov [16] provided an LP-based $7/4$ -approximation algorithm (the algorithm is combinatorial, but the analysis is LP-based). Moreover, Cheriyan and Gao [4] presented a combinatorial 1.5-approximation, whose analysis is based on an SDP obtained through Lasserre's hierarchy. More recently, a $(5/3 + \epsilon)$ -approximation algorithm that is LP-based was given by Adjashvili [1]. Fiorini, Groß, Könemann, and Sanità [9] were able to considerably strengthen Adjashvili's approach to obtain an LP-based $(1.5 + \epsilon)$ -approximation by adding to Adjashvili's LP a strong family of additional constraints, commonly known as $\{0, \frac{1}{2}\}$ -Chvátal-Gomory constraints. Hence, up to an arbitrarily small constant $\epsilon > 0$, which impacts the running time, this algorithm also matches the best factor of 1.5 achieved by Kortsarz and Nutov [17], and Cheriyan and Gao [4].

Interestingly, despite the recent LP-based and SDP-based advances, relaxations for TAP are still badly understood. In particular, even the integrality gap of the most natural LP, known as the *cut-LP*, was only very recently shown by Nutov [20] to be below 2 (namely at most $2 - \frac{2}{15}$), even though the best known lower bound is 1.5, which is due to an example of Cheriyan, Karloff, Khandekar, and Könemann [6]. Hence, it is unknown whether the stronger LPs introduced in [1] and [9] are any stronger than the canonical cut-LP. Moreover, the fact that 1.5-approximations (or $1.5 + \epsilon$, respectively) have now been obtained by three independent approaches, one combinatorial [17], one SDP-based [4], and one LP-based [9], together with the difficulty to obtain LPs or SDPs with an integrality gap strictly below 1.5 (even with a non-algorithmic proof), has led to the question whether it is possible to obtain approximation factors strictly below 1.5 for TAP (see, e.g., Könemann [15], for an example where this question was raised explicitly). In this paper, we answer this question in the affirmative.

1.1 Our Results

Our main result is the first approximation algorithm beating the factor of 1.5 for TAP.

THEOREM 1. *There exists a 1.458-approximation algorithm for TAP.*

A first step in our approach exploits key observations done by Adjashvili [1] to reduce to a very structured TAP instance, which we call k -wide, and is defined as follows.

DEFINITION 2 (k -WIDE INSTANCE, PRINCIPAL SUBTREE). *Let $k \in \mathbb{Z}_{\geq 1}$. We say that a tree $T = (V, E)$ is k -wide, if there exists a vertex $r \in V$, which we also call root, with the following property. Any subtree of T consisting of r , one of its children u , and all descendants of u has at most k leaves; moreover, these subtrees are called principal subtrees.*

Starting with the work of Adjashvili [1], several recently developed approaches [9, 20] can be interpreted as improving approximation factors and LP formulations for $O(1)$ -wide TAP instances, and propagating such improvements to general TAP instances. In a first step, we provide a black-box reduction of general TAP to $O(1)$ -wide TAP instances. In previous approaches based on Adjashvili's technique, there remained some entanglement between improvements on approximating $O(1)$ -wide TAP instances and the propagation of these improvements to general TAP instances. Our statement below avoids such entanglements and allows for a simpler presentation of our main technique.

THEOREM 3. *Let $k \in \mathbb{Z}_{\geq 1}$, $\alpha \geq 1$, and \mathcal{A} be an algorithm that is an α -approximation for TAP on k -wide instances. Then there is an $(\alpha + O(\frac{1}{\sqrt{k}}))$ -approximation algorithm for TAP, making a polynomial number of calls to \mathcal{A} and performing further operations that take time polynomially bounded in the input size. This reduction also works for randomized algorithms, where the approximation guarantees hold in expectation.*

Finally, our main contribution, which leads to Theorem 1 due to Theorem 3, is a new approximation algorithm for $O(1)$ -wide instances, whose existence is formally stated in the following theorem.

THEOREM 4. *There exists an approximation algorithm for $O(1)$ -wide TAP instances with approximation guarantee $\sqrt[3]{34/4} < 1.458$.*

For simplicity, the above results are stated only for unweighted instances. However, both Theorem 3 and Theorem 4 can be extended, in a weaker form, to WTAP instances where the ratio of the highest cost c_{\max} to the lowest cost c_{\min} is bounded by Δ . In particular, this allows for beating the approximation factor of 1.5 even for such instances by a constant that depends on Δ . Due to space constraints, we defer this discussion to the long version of the paper.

1.2 Brief Discussion of Used Techniques

The black-box reduction presented by Theorem 3 can be obtained by a careful combination of known techniques, including core reduction ideas by Adjashvili [1] and the use of the ellipsoid algorithm on a well-defined LP with a "partial" separation oracle based on an approximation algorithm for $O(1)$ -wide instances. Ellipsoid-based

techniques with such partial separation oracles have been employed successfully in prior work, including in the context of facility location [2, 3, 18], and very recently, Nutov [20] presented an elegant application of it by showing that a certain LP has integrality gap below 2 for weighted TAP with bounded ratio between highest and lowest cost.

To show our main technical result, i.e., a $\frac{\sqrt{34}}{4}$ -approximation for $O(1)$ -wide instances as claimed by Theorem 4, we introduce a new technique to improve solutions, which we call *rewiring*. More precisely, a common approach in algorithms to round solutions to relaxations of TAP is to *split* certain links into two smaller ones that cover the same edges. This approach is particularly well-suited for $O(1)$ -wide instances where one can break the TAP problem into independent ones over the principal subtrees, each of which only has $O(1)$ -many leaves, by splitting all links that go over the root. We show that, when starting with an appropriate LP, then after independently rounding the subtrees, one can further improve by identifying pairs of links across different principal subtrees that can be replaced by a single link. We show that this rewiring process can be reduced to a certain type of matching problem. We reduce the question of whether one can make a substantial improvement through rewirings, to a question about the expected cardinality of maximum matchings in a certain random graph, for which we can show that large matchings always exist in expectation.

Outline of the Paper. Section 2 proves Theorem 4 by presenting our approximation algorithm for $O(1)$ -wide TAP instances. To better highlight our main ideas, we first present a randomized approximation algorithm with an approximation factor of 1.465, and expand on the derandomization in Appendix C, and on the better approximation factor stated in Theorem 4 in Appendix D. The black-box reduction to k -wide instances is contained in Appendix B, modulo one technical result that closely follows a derivation by Adjashvili [1]. Due to space constraints, we defer a formal proof of this result to the long version of the paper. Finally, Appendix A contains some relatively standard but technical proofs that help us bound the approximation ratio of our procedure for k -wide instances.

Basic Preliminaries. Consider a TAP instance (T, L) on a tree $T = (V, E)$ with links $L \subseteq \binom{V}{2}$. For a link $\ell \in L$ we denote by $P_\ell \subseteq E$ the unique path between the endpoints of ℓ in T . Hence, TAP can be rephrased as the following covering problem

$$\min \{|U| \mid U \subseteq L, E = \cup_{\ell \in U} P_\ell\} .$$

For an edge $e \in E$, we denote by $\text{cov}(e) = \{\ell \in L \mid e \in P_\ell\}$ the set of all links covering e . More generally, for $U \subseteq E$, we denote by $\text{cov}(U) = \cup_{e \in U} \text{cov}(e)$ all links that cover at least one edge of U . The classical *cut-LP* seeks to minimize $x(L) := \sum_{\ell \in L} x_\ell$ over all x in the polytope

$$\Pi_T := \left\{ x \in [0, 1]^L \mid x(\text{cov}(e)) \geq 1 \quad \forall e \in E \right\} .$$

Moreover, for an integer $k \in \mathbb{Z}_{\geq 1}$, we use the notation $[k] := \{1, \dots, k\}$.

2 APPROXIMATING k -WIDE TAP

In this section, we prove Theorem 4 by presenting our approximation algorithm for TAP on k -wide instances. We start in Section 2.1 by presenting a novel efficiently solvable LP relaxation for k -wide

TAP instances. In Section 2.2, we present a rounding algorithm for its solution, which is based on performing rewirings. Interestingly, the number of rewirings can be lower bounded by the expected size of a matching on a particular type of random graph. In Section 2.3, we show that the latter expected size is large enough, hence implying our improved approximation factor as discussed in Section 2.4. Incidentally, this also shows that our LP relaxation has an integrality gap below 1.5 for k -wide instances. This is the first LP relaxation which is known to have this property.

Throughout this section, we assume that we are dealing with TAP instances (T, L) that are *shadow-complete*. This means that for every link $\ell \in L$, and any two vertices u, v on the path P_ℓ , we have $\{u, v\} \in L$. In this case we have $P_{\{u, v\}} \subseteq P_\ell$, and $\{u, v\}$ is called a *shadow* of ℓ . Clearly, one can assume without loss of generality that a TAP instance is shadow-complete by introducing all shadows of links. This leads to an equivalent problem because any solution using a shadow ℓ' of an original link ℓ can be modified to another solution by replacing ℓ' by ℓ , without increasing the number of links in the solution.

Moreover, we will consider different rounding algorithms whose approximation guarantees depend on the type of links involved. Given a k -wide TAP instance (T, L) with root r , we partition L into $L_{\text{cross}} \cup L_{\text{in}}$, where L_{cross} , called *cross-links*, are all the links whose endpoints are different from the root and are in different principal subtrees. The remaining set L_{in} , called *in-links*, are all the links with both endpoints in the same principal subtree. Moreover, $L_{\text{up}} \subseteq L_{\text{in}}$, called *up-links*, are all the in-links $\{u, v\} \in L_{\text{in}}$ such that one of the endpoints of $\{u, v\}$, say u , lies on the path from r to v . We sometimes also use the notions of cross-links, in-links, and up-links with respect to some vertex r for instances that are not k -wide.

2.1 A New LP Relaxation for k -Wide TAP

Recall that the cut-LP has integrality gap at least 1.5 [6].¹ We thus need to strengthen it to beat this factor. We do this by adding two sets of constraints. We start with $(0, \frac{1}{2})$ -Chvátal-Gomory cuts, which were introduced in the context of TAP by Fiorini et al. [9]. They are described by the following polytope:

$$\Pi_T^{\text{CG}} = \left\{ x \in [0, 1]^L \mid x(\pi(S)) \geq \frac{|\delta_E(S)| + 1}{2} \quad \forall S \subseteq V \right. \\ \left. \text{with } |\delta_E(S)| \text{ odd} \right\} ,$$

where $\delta_E(S)$ are all edges with precisely one endpoint in S , and $\pi(S)$ is the multiset of links whose corresponding paths intersect $\delta_E(S)$, where the multiplicity of any link ℓ is $\lceil |P_\ell \cap \delta_E(S)| / 2 \rceil$. These constraints imply the cut-LP constraints.² Moreover, one can efficiently separate over Π_T^{CG} (see [9]). A key property of Π_T^{CG} shown by Fiorini et al. [9], is that points in Π_T^{CG} can be rounded losslessly on a particular type of instance, as summarized below.

¹Notice, also for k -wide instances, the cut-LP can have an integrality gap arbitrarily close to 1.5 (for large enough k). This follows by the fact that a principal subtree can be any TAP instance with at most k leaves. Hence, for large enough k , any integrality gap instance is a k -wide TAP.

²The fact that the cut-LP constraint for any edge $e \in T$ is implied by Π_T^{CG} , and hence $\Pi_T^{\text{CG}} \subseteq \Pi_T$, follows by considering the set S corresponding to the vertex set of one of the two connected components of T when removing the edge e .

THEOREM 5 ([9]). *Let (T, L) be a TAP instance, let $x \in \Pi_T^{CG}$, and let $r \in V$. If every link $\ell \in \text{supp}(x)$ is either (i) an uplink with respect to r , or (ii) a cross-link with respect to r , then one can losslessly round x , i.e., a solution $R \subseteq L$ to the TAP problem with $|R| \leq x(L)$ can be efficiently obtained. This even holds for WTAP.*

The above theorem leads to a natural rounding algorithm for a point $x \in \Pi_T^{CG}$. This rounding procedure, which was suggested in [9], works as follows when applied to the root of a k -wide TAP instance $T = (V, E)$ with links L . We first modify x to avoid the use of in-links that are not up-links, which will bring us to the setting of Theorem 5. More precisely, for $\ell = \{u, v\} \in L_{\text{in}} \setminus L_{\text{up}}$, let $a \in V$ be the vertex closest to the root among the vertices on the path P_ℓ . Then one can modify x by increasing its value on $\{u, a\}$ and $\{v, a\}$ by $x(\ell)$, and setting the value of $x(\ell)$ to 0. Doing this for all links in $L_{\text{in}} \setminus L_{\text{up}}$, a new point $y \in \Pi_T^{CG}$ is obtained with support only on cross-links and up-links and such that $y(L_{\text{cross}}) = x(L_{\text{cross}})$ and $y(L_{\text{up}}) = 2x(L_{\text{in}}) - x(L_{\text{up}})$. Finally, by applying Theorem 5 to y , one obtains the following.

COROLLARY 6 ([9]). *Given a point $x \in \Pi_T^{CG}$, we can compute in polynomial time a solution \mathcal{A} such that*

$$|\mathcal{A}| \leq 2x(L_{\text{in}}) - x(L_{\text{up}}) + x(L_{\text{cross}}).$$

On top of the Chvátal-Gomory constraints, we introduce a second set of constraints that deviates from prior work, but can be interpreted as a strengthening of so-called *bundle constraints* introduced in [1], which were also used (with non-trivial technical adjustments) in [9, 20]. Bundle constraints make sure that when restricting an LP solution to any subtree of T with a constant number of leaves, then the value of this restricted LP solution is not lower than that of an optimal solution covering the same subtree. Being able to losslessly round an LP solution on subtrees with $k = O(1)$ leaves, by replacing the LP solution with an integral solution of no higher cost, is of particular interest in the context of k -wide TAP instances due to the following. One can first independently consider each principal subtrees of the k -wide instance and round it, and then take the union of the links used in the solutions to the different principal subtrees. As first shown in [1], this allows for obtaining a solution with at most $x(L_{\text{in}}) + 2x(L_{\text{cross}})$ links. By returning the better of this solution and the one guaranteed by Corollary 6, a 1.5-approximation is obtained for k -wide instances. Finally, by Theorem 3, this implies a $(1.5 + \epsilon)$ -approximation for any TAP instance. We highlight that this is a nice example of how Theorem 3 allows for providing a more elegant description of existing techniques, because we can restrict ourselves to k -wide instances.

We will provide techniques that improve on this rounding procedure. More precisely, our main goal now is to introduce additional constraints such that if one rounds each principal subtree independently, then there is a way to further improve the solution after we take the union of all rounded links. For this, we want to be able to round each principal subtree in a well-structured way. More precisely, we first need to be able to decompose the LP solution, when restricted to one principal subtree, into a *convex combination of integral solutions*. Moreover, we want that the links used in solutions appearing in this convex combination do not contain redundancies, which we formalize below through the notion of *shadow-minimality* for link families.

DEFINITION 7. *Consider a TAP instance (T, L) . Two links $\ell_1, \ell_2 \in L$ are shadow-minimal if there exists no shadow s of one link, say ℓ_1 , such that $P_{\ell_1} \cup P_{\ell_2} = P_s \cup P_{\ell_2}$. Moreover, a set of links $L' \subseteq L$ is shadow-minimal if all links in L' are pairwise shadow-minimal.*

The following basic observation shows that requiring shadow-minimality is not restrictive. In other words, there always exists an optimal solution that is shadow-minimal. This follows by the fact that, whenever we have a non-shadow-minimal solution, then some link can be shortened to one of its shadows without destroying feasibility.

OBSERVATION 8. *Given a TAP instance (T, L) , and given some $L_1 \subseteq L$, there exists a shadow-minimal set of links $L_2 \subseteq L$ such that $|L_1| \geq |L_2|$, and L_2 covers the same edges as L_1 . Furthermore, L_2 contains only links from L_1 , or shadows thereof.*

To describe our constraint set, which allows for rounding principal subtrees in a shadow-minimal way, we present here a compact extended formulation for simplicity. It is also possible to achieve the same goal via an exponential-size set of constraints (and no extra variables).

Fix a principal subtree $T_i = (V_i, E_i)$, $i \in [q]$. Conceptually, one could introduce a binary decision variable y_S for any set $S \subseteq \text{cov}(E_i)$ that covers all the edges of T_i and is shadow-minimal, and then enforce the constraint of picking exactly one such set S for each T_i . Clearly, there are exponentially many such sets S . We will show in the following that it suffices to focus on a polynomially small family of sets. Interestingly, even though we use shadow-minimality primarily for the rounding procedure to be introduced later, it also allows us to restrict the number of sets S we have to consider, thus obtaining a family of polynomial size. To do so, we start by observing that shadow-minimal solutions of k -wide trees contain few cross-links:

LEMMA 9. *Given a TAP instance (T, L) , where T is a k -wide tree, consider a shadow-minimal set $L' \subseteq L$. Then, for each principal subtree T_i , L' contains at most k cross-links with an endpoint inside T_i .*

PROOF. Suppose towards contradiction that there exists a principal subtree T_i such that at least $k + 1$ cross-links in L' have an endpoint inside T_i . Since T_i contains at most k leaves, we have by the pigeonhole principle that there is one leaf u such that the path from u to r in T contains the endpoints of at least two cross-links, which contradicts the shadow-minimality of L' . \square

Let $\Lambda_i \subseteq 2^{\text{cov}(E_i)}$ be the family of all shadow-minimal subsets of all cross-links with one endpoint in V_i . By Lemma 9, each set $R \in \Lambda_i$ satisfies $|R| \leq k$. Hence, Λ_i has polynomial size; more precisely, $|\Lambda_i| = O(|L|^k)$. For each $R \in \Lambda_i$, let $C(i, R) \subseteq L$ be a minimum cardinality set of links with both endpoints in V_i that satisfies

- (i) $R \cup C(i, R)$ is shadow-minimal, and
- (ii) $R \cup C(i, R)$ covers E_i .

Notice that such a set $C(i, R)$ can indeed be computed efficiently. The requirement that $R \cup C(i, R)$ covers E_i corresponds to $C(i, R)$ being a TAP solution to the residual instance consisting of the tree T_i after contracting the edges already covered by R . However, being a subtree of the k -wide tree T_i , this residual instance has no more

than k leaves, and TAP instances with constantly many leaves are known to be efficiently solvable (see, e.g., [1, 20]). To guarantee shadow-minimality of $R \cup C(i, R)$, one can first only consider links in the residual instance that do not lead to a violation of shadow-minimality with R , i.e., we delete in the residual instance any link ℓ such that $R \cup \{\ell\}$ is not shadow-minimal. Finally, any optimal solution to the residual instance that does not use any link ℓ such that $R \cup \{\ell\}$ is not shadow-minimal, can easily be transformed into a shadow-minimal one by shortening links if necessary. Let $L_i^R = R \cup C(i, R)$.

Due to Observation 8 and Lemma 9, there is an optimal solution to the k -wide TAP instance we consider that is of the form $L_1^{R_1} \cup \dots \cup L_q^{R_q}$, where $R_i \in \Lambda_i$ for $i \in [q]$. Indeed, consider some shadow-minimal optimal solution L^* ; due to Observation 8, there exists at least one. Due to Lemma 9, we know L^* contains at most k cross-links with an endpoint inside any principal subtree T_i . Let L_i^* be the cross-links of L^* that have an endpoint in T_i ; then $\bigcup_{i=1}^q (L_i^* \cup C(i, L_i^*))$ is an optimal solution of the desired form.

Our new LP constraints are based on this observation, and we have a variable $\lambda_i^R \geq 0$ for each principal subtree T_i and each set $R \subseteq \Lambda_i$, where $\lambda_i^R = 1$ is interpreted as including the links L_i^R in the solution. The local convex decomposition constraints we introduce ensure that x and λ are consistent, and that we only choose one set of links for each one of the q principal subtrees in T :

$$\begin{aligned} x_\ell &= \sum_{R \in \Lambda_i: \ell \in L_i^R} \lambda_i^R, & \forall i \in [q] \quad \forall \ell \in \text{cov}(E_i), \\ \sum_{R \in \Lambda_i} \lambda_i^R &= 1, & \forall i \in [q], \\ \lambda_i^R &\geq 0 & \forall i \in [q] \quad \forall R \in \Lambda_i. \end{aligned} \quad (1)$$

As discussed, there must be at least one $x \in \{0, 1\}^L$ that is the characteristic vector of an optimal solution and is feasible for the above constraints. Hence, the above constraints can be interpreted as a relaxation of the original TAP problem.

Finally, the feasible solutions of our LP consist of all tuples (x, λ) , where $x \in \mathbb{R}^L$ and λ has $\sum_{i=1}^q |\Lambda_i|$ many components, one for each $i \in [q]$ and $R \in \Lambda_i$, such that $x \in \Pi_T^{\text{CG}}$ and (x, λ) satisfies (1). Moreover, we minimize $x(L)$ over this polytope. For brevity, we call this resulting LP, the k -wide-LP. One can efficiently optimize over the k -wide-LP for any $k = O(1)$, because there is an efficient separation oracle for Π_T^{CG} , and there are polynomially many additional constraints and variables described by (1).

We summarize that the key property we gain through our k -wide-LP compared to previous linear programs, is that any solution x to the k -wide-LP, when restricted to the links $\text{cov}(E_i)$ for any $i \in [q]$, can be written as a convex combination of shadow-minimal link sets, each covering E_i .

2.2 The Rounding Algorithm

We next present our rounding algorithm based on rewirings. Let (x, λ) be an optimal fractional solution to the k -wide-LP. We return the better of two solutions obtained by two different rounding procedures. The first procedure is the one by Fiorini et al. [9], whose guarantee is stated in Corollary 6. The second rounding algorithm deviates substantially from prior work, and is one of the

main algorithmic contributions of this paper. For simplicity, we first present a randomized version of our rounding algorithm, and discuss in Appendix C how it can be derandomized.

Note that the links in x that cover the edges of each T_i are expressed as a convex combination of integral solutions L_i^R . We will interpret the coefficients λ_i^R of this convex combination as a probability distribution and sample from it, independently for each T_i . Let L_i be the (random) local solution for each T_i , and we start by considering $\bigcup_{i=1}^q L_i$, which is clearly a feasible integral solution. Notice that each in-link ℓ has both endpoints in one principal subtree T_i , and is contained in L_i with probability x_ℓ (and not contained in any other L_j). On the other hand, each cross-link ℓ has its endpoints in two different subtrees, say T_i and T_j , and is thus contained in each of L_i and L_j with probability x_ℓ (and it is not contained in any other L_h). This simple line of argumentation leads to the following guarantee.

$$\mathbb{E} \left[\left| \bigcup_{i=1}^q L_i \right| \right] \leq \mathbb{E} \left[\sum_{i=1}^q |L_i| \right] = x(L_{\text{in}}) + 2x(L_{\text{cross}}), \quad (2)$$

which is in fact the same guarantee that was also obtained by Adjiashvili [1] through a slightly different approach.

In order to improve on this, we introduce our rewiring process that will replace some pairs of cross-links appearing in different principal subtrees by a single one. To do such a replacement, we introduce the notion of *active vertices*, which will be potential endpoints of a new link that allows for removing two existing ones.

DEFINITION 10. *Let $u \in V$, and let $i \in [q]$ such that $u \in V_i$. Then u is called active if there exists a cross-link $\ell_u \in L_i \cap L_{\text{cross}}$ that has u as one of its endpoints.*

Notice that for an active vertex u , the corresponding link ℓ_u , as described in the definition, is unique; for otherwise the links L_i will not be shadow-minimal.

Our rewiring step is based on the following observation of how one can improve the solution $\bigcup_{i=1}^q L_i$. Whenever there are two active vertices u, v , say in subtrees T_i and T_j , respectively, such that $\{u, v\} \in L_{\text{cross}}$, then we can replace both the cross-link ℓ_u in L_i that has u as one of its endpoints and the cross-link ℓ_v in L_j that has v as one of its endpoints, by the single cross link $\{u, v\}$. The following proposition formalizes this operation for multiple pairs $\{u, v\}$ forming a matching M .

PROPOSITION 11. *Let $M \subseteq L_{\text{cross}}$ be a matching between active vertices. Moreover, for each active vertex u , which is part of some subtree T_i , let $\ell_u \in L_i$ be the unique cross-link in L_i with u as one of its endpoints, and we denote by $L_i(M)$ the set of all $\ell_u \in L_i$ for $u \in V_i$ being an endpoint of some edge in M . Then*

$$\mathcal{B} := \left(\bigcup_{i=1}^q (L_i \setminus L_i(M)) \right) \cup M$$

is a solution to the considered k -wide TAP instance.

PROOF. We check that all edges in each principal subtree are covered. Hence, let $i \in [q]$, and consider an edge $e \in E_i$. We recall that by construction, L_i covers all edges E_i . If e is covered by a link in $L_i \setminus L_i(M)$, then it is clearly covered by the same link in \mathcal{B} . Otherwise, if e is covered by a link $\ell \in L_i(M)$, then there is

a cross-link in M that covers all edges within E_i that ℓ covered, including the edge e . \square

Notice that the cardinality of the solution \mathcal{B} described in Proposition 11 can be bounded by

$$|\mathcal{B}| \leq \left(\sum_{i=1}^q |L_i \setminus L_i(M)| \right) + |M| = \left(\sum_{i=1}^q |L_i| \right) - |M|, \quad (3)$$

which leads to an improvement compared to (2) (after taking expectations). We call the process of replacing a pair of links ℓ_u, ℓ_v with $\{u, v\} \in L_{\text{cross}}$ by a single link $\{u, v\}$ a *rewiring*. It may happen that ℓ_u and ℓ_v are the same. Nevertheless, we still improve on the analysis presented in (2), which counts these links twice. However, it is crucial that we do not try to improve on the analysis presented in (2) by only exploiting that we improve if the same cross-link appears in two different L_i . Such events can be very unlikely, and would not be enough to improve on existing methods. Hence, we do need to do non-trivial rewirings.

Maximizing the number of rewirings h can be achieved by solving a matching problem over the graph consisting of all active vertices, and whose edges are all cross-links between active vertices. However, it turns out that we only need to consider rewirings on a subset of the active vertices, namely those of degree different from 2. This is sufficient for our analysis and simplifies the exposition. Hence, let $V_{\text{crit}} \subseteq V \setminus \{r\}$ be the set of all vertices, except for the root, with degree different from 2, and let $A \subseteq V_{\text{crit}}$ be the random set corresponding to all active vertices in V_{crit} . For brevity, we call the vertices in V_{crit} *critical vertices*. Moreover, we partition the cross-links L_{cross} into the links $L_{\text{cross}}^{\text{crit}}$ with both endpoints being critical, and the remaining links $L_{\text{cross}}^{\text{no-crit}}$. Hence, our rewiring procedure finds a maximum cardinality matching in the graph $(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2})$, say with matching edges $M \subseteq L_{\text{cross}}$, and returns the solution \mathcal{B} as described in Proposition 11. We highlight that in our description so far, our rounding algorithm is randomized due to the random choice of L_1, \dots, L_q , which also leads to a random set A of active critical vertices. To make sure that we can find a sufficiently large matching, it is crucial to observe that the distribution of A has the following properties:

$$\Pr[v \in A] = \sum_{\ell \in L_{\text{cross}}^{\text{crit}}, v \in \ell} x_\ell \quad \forall v \in V_{\text{crit}}, \text{ and} \quad (4)$$

$$\Pr[u, v \in A] = \Pr[u \in A] \cdot \Pr[v \in A] \quad \forall \{u, v\} \in L_{\text{cross}}^{\text{crit}},$$

where the first property follows from the fact that the link sets L_i are shadow-minimal, which implies that there is at most one cross-link in L_i that has v as its endpoint. Moreover, the second property follows from the fact that if $\{u, v\} \in L_{\text{cross}}^{\text{crit}}$, then u and v are in different principal subtrees, say T_i and T_j , respectively. Indeed, the independence of $\{u \in A\}$ and $\{v \in A\}$ then follows by the independence of the distributions of L_i and L_j , which holds because we round each principal subtree independently. The following lemma summarizes the key properties of the solution we obtain through our rounding algorithm based on rewiring. For any graph $G = (V, E)$, we denote by $\eta(V, E)$ the cardinality of a maximum cardinality matching in G .

LEMMA 12. *Consider a k -wide TAP instance (T, L) and a solution (x, λ) to the k -wide-LP. There exists a randomized algorithm that*

returns a solution $\mathcal{B} \subseteq L$ to the TAP instance (T, L) , such that

$$\mathbb{E}[|\mathcal{B}|] \leq x(L_{\text{in}}) + 2x(L_{\text{cross}}) - \mathbb{E}_A \left[\eta \left(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2} \right) \right],$$

where A is a random subset of V_{crit} with a distribution satisfying (4).

PROOF. The set A corresponds to the active vertices that are critical, which fulfill (4). Moreover,

$$\begin{aligned} \mathbb{E}[|\mathcal{B}|] &\leq \mathbb{E} \left[\sum_{i=1}^q |L_i| \right] - \mathbb{E}_A \left[\eta \left(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2} \right) \right] \\ &= x(L_{\text{in}}) + 2x(L_{\text{cross}}) - \mathbb{E}_A \left[\eta \left(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2} \right) \right], \end{aligned}$$

where the inequality follows from (3) and the equality from (2). \square

2.3 Lower Bounding the Matching Size

The following lemma shows that $\mathbb{E}_A[\eta(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2})]$ is large if $x(L_{\text{cross}}^{\text{crit}})$ is so.

LEMMA 13. *Let $G = (V, E)$ be an undirected graph, let $x \in \mathbb{R}_{\geq 0}^E$ such that for any $v \in V$ we have $x(\delta(v)) \leq 1$, and let A be a random subset of V with a distribution that satisfies:*

- (i) $\Pr[v \in A] = x(\delta(v)) \quad \forall v \in V$, and
- (ii) $\Pr[v, u \in A] = \Pr[u \in A] \cdot \Pr[v \in A] \quad \forall \{u, v\} \in E$.

Then there exists a matching $M \subseteq E$ in G such that

$$\mathbb{E} \left[|M \cap \binom{A}{2}| \right] \geq x(E)^2 / |V|.$$

Before we prove the lemma, observe that $M \cap \binom{A}{2}$ is clearly a matching in $(A, E \cap \binom{A}{2})$, because M is a matching. The above lemma thus implies that $\mathbb{E} \left[\eta \left(A, E \cap \binom{A}{2} \right) \right] \geq x(E)^2 / |V|$.

PROOF. We start by replacing x by a sparser point y . For this, consider the polytope

$$Q = \left\{ y \in \mathbb{R}_{\geq 0}^E \mid y(\delta(v)) = x(\delta(v)) \quad \forall v \in V \right\},$$

which is non-empty because $x \in Q$. Let y be any vertex of Q . We have $|\text{supp}(y)| \leq |V|$, which follows by the following standard sparsity argument:³ Any vertex y of Q is defined by $|E|$ -many linearly independent and tight constraints; however, Q has only $|V|$ constraints that are not nonnegativity constraints, and therefore, at least $|E| - |V|$ of the nonnegativity constraints must be tight at y , which implies $|\text{supp}(y)| \leq |V|$.

We let the matching $M \subseteq E$ be a matching obtained by the greedy algorithm for maximum weight matchings with respect to the weights y . More formally, we order the edges $E = \{e_1, \dots, e_m\}$ such that $y(e_1) \geq \dots \geq y(e_m)$. We start with $M = \emptyset$ and go through the edges in the order e_1, \dots, e_m . When considering e_i , we set $M = M \cup \{e_i\}$ if $M \cup \{e_i\}$ is a matching. We have the following (which holds for any matching M):

$$\begin{aligned} \mathbb{E} \left[|M \cap \binom{A}{2}| \right] &= \sum_{\{u, v\} \in M} x(\delta(u)) \cdot x(\delta(v)) \\ &= \sum_{\{u, v\} \in M} y(\delta(u)) \cdot y(\delta(v)), \end{aligned} \quad (5)$$

which is an immediate consequence of property (i) and (ii), and of $y \in Q$. Next, we show that the following holds:

$$\sum_{\{u, v\} \in M} y(\delta(u)) \cdot y(\delta(v)) \geq \sum_{e \in E} y(e)^2. \quad (6)$$

³Recall that the support $\text{supp}(y) \subseteq E$ of y are all edges $e \in E$ with $y(e) > 0$.

Note that when expanding the left-hand side of (6) using $y(\delta(u)) = \sum_{e \in \delta(u)} y(e)$ and $y(\delta(v)) = \sum_{e \in \delta(v)} y(e)$, we obtain a sum of terms of the form $y(e) \cdot y(f)$ for some pairs $e, f \in E$. For any $e = \{u, v\} \in M$, the left-hand side of (6) has a term of the form $y(e)^2$, due to the term $y(\delta(u)) \cdot y(\delta(v))$. Now consider any non-matching edge $e = \{u, v\} \in E \setminus M$. Because e is a non-matching edge, and we constructed M by the greedy algorithm, there is at least one endpoint of e , say u , such that there is an edge $f \in M$ incident with u that satisfies $y(f) \geq y(e)$. This implies that $y(\delta(u)) \cdot y(\delta(v))$ contains a term $y(f) \cdot y(e) \geq y(e)^2$. Hence, for each $e \in E$, we identified a term on the left-hand side of (6) of value at least $y(e)^2$; moreover, all such terms on the left-hand side are different, which implies (6).

The statement of the lemma now follows by combining (5) together with (6) and the following inequality,

$$x(E)^2 = y(E)^2 \leq |\text{supp}(y)| \cdot \sum_{e \in E} y(e)^2 \leq |V| \cdot \sum_{e \in E} y(e)^2,$$

where the equality follows from $y \in Q$, the first inequality follows by the Cauchy-Schwarz inequality, and the second one by $|\text{supp}(y)| \leq |V|$. \square

Notice that the conclusion of the lemma is essentially optimal due to the following. Consider a graph $G = (V, E)$ that is a star with center $c \in V$, and let $x \in \mathbb{R}_{\geq 0}^E$ be the vector with value $1/|V|-1$ on each edge. Moreover, let A be the random set of vertices such that with probability $1/(|V|-1)$ we have $A = V$, and with probability $(|V|-2)/(|V|-1)$ we have $A = \{c\}$. In this case, $\mathbb{E}[\nu(A, E \cap \binom{A}{2})] = 1/(n-1)$ and $x(E)^2/|V| = 1/n$. Clearly, the above example can also be extended to larger values of $x(E)$ by considering disjoint stars. We highlight that this does not exclude the possibility that, in the context we use the lemma, stronger conclusions may be obtained through a different statement and analysis.

2.4 The Approximation Factor

We next upper bound the approximation factor of the considered rounding algorithm. For simplicity, and due to space constraints, we present here a slightly weaker upper bound and defer to Appendix D the refined analysis leading to the bound claimed in Theorem 4.

Consider an optimal fractional solution (x, λ) to the k -wide-LP, and let $\text{OPT}^* := x(L)$ be its value. Let $\alpha_{\text{cross}} \in [0, 1]$ be the fraction of OPT^* due to $x(L_{\text{cross}})$, i.e., $\alpha_{\text{cross}} = x(L_{\text{cross}})/\text{OPT}^*$, and define similarly $\alpha_{\text{cross}}^{\text{crit}}$, $\alpha_{\text{cross}}^{\text{no-crit}}$, α_{in} and α_{up} w.r.t. $L_{\text{cross}}^{\text{crit}}$, $L_{\text{cross}}^{\text{no-crit}}$, L_{in} and L_{up} , respectively. Observe that $\alpha_{\text{up}} \leq \alpha_{\text{in}}$, $\alpha_{\text{cross}}^{\text{crit}} + \alpha_{\text{cross}}^{\text{no-crit}} = \alpha_{\text{cross}}$, and $\alpha_{\text{cross}} + \alpha_{\text{in}} = 1$.

Previously, we established that our rewiring procedure is the stronger the larger $x(L_{\text{cross}}^{\text{crit}})$ is. Observe that if $x(L_{\text{cross}})$ is small enough, namely significantly less than $x(L)/2$, then the rounding guarantee of Lemma 12 already improves on the factor 1.5 even without any rewiring. Still it may happen that $x(L_{\text{cross}})$ is large and $x(L_{\text{cross}}^{\text{crit}})$ is small, which implies that $x(L_{\text{cross}}^{\text{no-crit}})$ is large. The next lemma shows that in this case also $x(L_{\text{up}})$ has to be large, which in turn implies that we get a stronger guarantee via Corollary 6.

LEMMA 14. *Given a k -wide TAP instance (T, L) and a solution (x, λ) to the k -wide-LP, we have that $\alpha_{\text{up}} \geq \alpha_{\text{cross}}^{\text{no-crit}}$.*

PROOF. Consider any link $\ell \in L_{\text{cross}}^{\text{no-crit}}$, and let v be any one of its non-critical endpoints. Let T_i be the principal subtree containing v ,

and, for any $S \subseteq L$, we denote by x_S be vector x restricted on the coordinates corresponding to S . From the design of the k -wide-LP, $x_{\text{cov}(T_i)}$ is a convex combination of shadow-minimal link sets L_i^R , each covering T_i . In each such solution L_i^R that contains ℓ , there must exist another link ℓ' incident on v and covering the edge right below v . The link ℓ' cannot be in L_{cross} or $L_{\text{in}} \setminus L_{\text{up}}$, because in both cases shadow-minimality of $\{\ell, \ell'\}$ would be contradicted, since one could shorten ℓ . Thus $\ell' \in L_{\text{up}}$. We can thus conclude that

$$\sum_{\ell \in L_{\text{up}}: v \in \ell} x_{\ell} \geq \sum_{\ell \in L_{\text{cross}}^{\text{no-crit}}: v \in \ell} x_{\ell}.$$

The claim follows by summing over all non-critical nodes v . \square

LEMMA 15. *There exists a randomized approximation algorithm for $O(1)$ -wide TAP instances with approximation guarantee $2(\sqrt{3}-1) < 1.465$.*

PROOF. Given a k -wide TAP instance (T, L) with $k = O(1)$, and a solution (x, λ) to the k -wide-LP, from Corollary 6 we know that we can round x to a solution \mathcal{A} such that

$$\frac{|\mathcal{A}|}{\text{OPT}^*} \leq 2\alpha_{\text{in}} - \alpha_{\text{up}} + \alpha_{\text{cross}} = 2 - \alpha_{\text{cross}} - \alpha_{\text{up}}, \quad (7)$$

where the equality follows from $\alpha_{\text{cross}} + \alpha_{\text{in}} = 1$. Furthermore, from Lemma 12 we know that we can randomly round (x, λ) to a solution \mathcal{B} such that

$$\mathbb{E}[|\mathcal{B}|] \leq x(L_{\text{in}}) + 2x(L_{\text{cross}}) - \mathbb{E}_A[\eta(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2})], \quad (8)$$

where A fulfills the conditions of Lemma 12. We proceed to upper bound $\mathbb{E}[|\mathcal{B}|]/\text{OPT}^*$. Let K be the total number of leaves. Notice that $|V_{\text{crit}}| < 2K$; for otherwise one would obtain that the average degree of T is at least 2, which contradicts that any tree has average degree strictly below 2. For each leaf node v one must have $\sum_{\ell \in L: v \in \ell} x_{\ell} \geq 1$ for the edge incident to v to be covered by the LP solution. Each link can contribute at most twice to the previous sums, and links in $L_{\text{up}} \cup L_{\text{cross}}^{\text{no-crit}}$ at most once. Thus we can conclude that

$$\begin{aligned} & 2x(L_{\text{cross}}) + 2x(L_{\text{in}}) - x(L_{\text{up}}) - x(L_{\text{cross}}^{\text{no-crit}}) \\ &= 2\text{OPT}^* - x(L_{\text{up}}) - x(L_{\text{cross}}^{\text{no-crit}}) \geq K > |V_{\text{crit}}|/2. \end{aligned} \quad (9)$$

Therefore, by Lemma 13, the expected size of the matching satisfies

$$\begin{aligned} \mathbb{E}_A \left[\eta(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2}) \right] &\geq \frac{x(L_{\text{cross}}^{\text{crit}})^2}{|V_{\text{crit}}|} \\ &\geq \frac{(\alpha_{\text{cross}}^{\text{crit}} \text{OPT}^*)^2}{(4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}}) \text{OPT}^*} \\ &= \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{(4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}})} \text{OPT}^*. \end{aligned}$$

Combining the above with (8), we conclude that

$$\begin{aligned} \frac{\mathbb{E}[|\mathcal{B}|]}{\text{OPT}^*} &\leq \alpha_{\text{in}} + 2\alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}}} \\ &= 1 + \alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}}}. \end{aligned} \quad (10)$$

Hence, by returning the better of the solutions \mathcal{A} and \mathcal{B} , we can upper bound the expected approximation factor of the algorithm

by the minimum of the right-hand sides of (7) and (10), i.e.,

$$1 + \min \left\{ 1 - \alpha_{\text{cross}} - \alpha_{\text{up}}, \alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}}} \right\}.$$

To obtain our approximation guarantee we maximize this function over $\alpha_{\text{up}}, \alpha_{\text{cross}}, \alpha_{\text{cross}}^{\text{crit}}, \alpha_{\text{cross}}^{\text{no-crit}} \in [0, 1]$ subject to the constraints: (i) $\alpha_{\text{cross}} = \alpha_{\text{cross}}^{\text{crit}} + \alpha_{\text{cross}}^{\text{no-crit}}$ (by definition), (ii) $\alpha_{\text{up}} \geq \alpha_{\text{cross}}^{\text{no-crit}}$ (by Lemma 14), and (iii) $\alpha_{\text{up}} + \alpha_{\text{cross}} \leq 1$ (because $\alpha_{\text{in}} + \alpha_{\text{cross}} = 1$ and $\alpha_{\text{up}} \leq \alpha_{\text{in}}$). In Appendix A (see Lemma 16), we show via basic calculations that this maximum, achieved for $\alpha_{\text{up}} = \alpha_{\text{cross}}^{\text{no-crit}} = 0$ and $\alpha_{\text{cross}}^{\text{crit}} = \alpha_{\text{cross}} = 4 - 2\sqrt{3}$, is $2 \cdot (\sqrt{3} - 1) < 1.465$. \square

A OMITTED PROOFS IN THE DETERMINATION OF OUR APPROXIMATION FACTOR

In this section, we show a technical statement, stated as Lemma 16 below, about the maximizer of a function used in Lemma 15, which completes the proof of Lemma 15.

LEMMA 16. *The maximum of*

$$1 + \min \left\{ 1 - \alpha_{\text{cross}} - \alpha_{\text{up}}, \alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}}} \right\}$$

over $\alpha_{\text{up}}, \alpha_{\text{cross}}, \alpha_{\text{cross}}^{\text{crit}}, \alpha_{\text{cross}}^{\text{no-crit}} \in [0, 1]$ subject to the constraints:

- (i) $\alpha_{\text{up}} \geq \alpha_{\text{cross}}^{\text{no-crit}}$,
- (ii) $\alpha_{\text{cross}} = \alpha_{\text{cross}}^{\text{crit}} + \alpha_{\text{cross}}^{\text{no-crit}}$, and
- (iii) $\alpha_{\text{up}} + \alpha_{\text{cross}} \leq 1$.

is achieved for $\alpha_{\text{up}} = \alpha_{\text{cross}}^{\text{no-crit}} = 0$ and $\alpha_{\text{cross}}^{\text{crit}} = \alpha_{\text{cross}} = 4 - 2\sqrt{3}$, leading to a value of $2(\sqrt{3} - 1)$.

PROOF. The function is non-increasing in α_{up} , and we can hence set $\alpha_{\text{up}} = \alpha_{\text{cross}}^{\text{no-crit}}$. Moreover, using $\alpha_{\text{cross}}^{\text{crit}} = \alpha_{\text{cross}} - \alpha_{\text{cross}}^{\text{no-crit}} = \alpha_{\text{cross}} - \alpha_{\text{up}}$, we have to maximize $f(\alpha_{\text{cross}}, \alpha_{\text{up}})$ defined as

$$1 + \min \left\{ 1 - \alpha_{\text{cross}} - \alpha_{\text{up}}, \alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}} - \alpha_{\text{up}})^2}{4(1 - \alpha_{\text{up}})} \right\},$$

subject to the constraints

- (a) $\alpha_{\text{up}} \leq \alpha_{\text{cross}}$ (because $\alpha_{\text{up}} = \alpha_{\text{cross}}^{\text{no-crit}} \leq \alpha_{\text{cross}}$ due to (ii)), and
- (b) $\alpha_{\text{up}} + \alpha_{\text{cross}} \leq 1$ (due to (iii)).

Let $f_1(x, y) = 1 - x - y$, $f_2(x, y) = x - \frac{(x-y)^2}{4(1-y)}$ and $f(x, y) = \min\{f_1(x, y), f_2(x, y)\}$. Using this notation, the maximization problem we consider is

$$\max_{0 \leq y \leq x \leq 1-y} f(x, y). \quad (11)$$

Let $D = \{(x, y) : 0 \leq y \leq x \leq 1 - y\}$ be the domain over which we maximize, and let $(x^*, y^*) \in D$ be a maximizer of (11). We first show that any such maximizer (x^*, y^*) must satisfy $f_1(x^*, y^*) = f_2(x^*, y^*)$.

Indeed, if $f_1(x^*, y^*) < f_2(x^*, y^*)$, then (x^*, y^*) is a maximizer of $\max_{(x, y) \in D} f_1(x, y)$. However, the only maximizer of f_1 over D is clearly $x = y = 0$, for which we have $f_1(0, 0) > f_2(0, 0)$, which contradicts $f_1(x^*, y^*) < f_2(x^*, y^*)$.

On the other hand, if $f_1(x^*, y^*) > f_2(x^*, y^*)$, then (x^*, y^*) is a maximizer of $\max_{(x, y) \in D} f_2(x, y)$. Observe that

$$\frac{\partial}{\partial x} f_2(x, y) = 1 - \frac{2x - 2y}{4(1 - y)}$$

is nonnegative for all $(x, y) \in D$, which follows directly from the fact that $(x, y) \in D$ implies $y \leq \frac{1}{2}$. Therefore, if (x^*, y^*) is a maximizer of f_2 over D , it holds that $y^* = 1 - x^*$. Hence, $f_2(x^*, y^*) = x^* - \frac{(2x^* - 1)^2}{4x^*} = 1 - \frac{1}{4x^*}$, which implies that $(x^*, y^*) = (1, 0)$, which is the unique maximizer of $f_2(x, y)$ over all points in D satisfying $y = 1 - x$. However, this contradicts that $f_1(x^*, y^*) > f_2(x^*, y^*)$.

Hence, for any maximizer (x^*, y^*) of $f(x, y)$ over D we have $f_1(x^*, y^*) = f_2(x^*, y^*)$, as desired. We thus only need to consider points $(x, y) \in D$ satisfying $f_1(x, y) = f_2(x, y)$, when seeking a maximizer of (11).

By expanding $f_1(x, y) = f_2(x, y)$, we get

$$-x^2 - 6xy + 8x - 5y^2 + 8y - 4 = 0.$$

Solving the above equation for x reveals that x is in

$$\left\{ 4 - 3y - 2\sqrt{(3-y)(1-y)}, 4 - 3y + 2\sqrt{(3-y)(1-y)} \right\}.$$

Because $x \leq 1$ and $y \leq \frac{1}{2}$, we must have

$$x = 4 - 3y - 2\sqrt{(3-y)(1-y)}. \quad (12)$$

By substituting x using the above equation, we obtain

$$f_1(x, y) = f_2(x, y) = 2y + 2\sqrt{(3-y)(1-y)} - 3 =: h(y).$$

Finally, observing that the derivative of h with respect to y , which is given by

$$\frac{dh(y)}{dy} = \frac{2y - 4}{\sqrt{(3-y)(1-y)}} + 2,$$

is non-positive for $y \in [0, \frac{1}{2}]$, we have that $h(y)$ is maximized for $y^* = 0$. By (12), this leads to $x^* = 4 - 2\sqrt{3}$. Because these values for x^* and y^* satisfy $(x^*, y^*) \in D$, they are indeed a maximizer for (11), with resulting value $f_1(x^*, y^*) = f_2(x^*, y^*) = 2(\sqrt{3} - 1)$, as desired. \square

B REDUCTION TO k -WIDE INSTANCES

In this section, we show Theorem 3, i.e., that general TAP instances can be reduced to $O(1)$ -wide ones by losing only a small constant in the approximation guarantee. Let $k \in \mathbb{Z}_{\geq 1}$, and let \mathcal{A} be an algorithm that is an α -approximation for TAP on k -wide instances.

A key ingredient in our reduction is the classical cut-LP for TAP, which, we recall, is given by

$$\min \{x(L) \mid x \in \Pi_T\}.$$

To reduce general TAP to $O(1)$ -wide TAP, we start with the above cut-LP, and then use the ellipsoid method to strengthen the LP through cutting planes that we generate by using the algorithm \mathcal{A} on well-chosen sub-problems that are $O(1)$ -wide. To better illustrate the idea of our approach, we first state a reduction result that can be derived by using techniques introduced by Adjashvili [1], and which we crucially exploit later on. In short, the result shows that the problem of finding a good rounding algorithm for the cut-LP can be reduced to finding a good rounding algorithm for $O(1)$ -wide instances.

In the lemma below, a partition of $T = (V, E)$ into subtrees $T_1 = (V_1, E_1), \dots, T_q = (V_q, E_q)$ means that T_1, \dots, T_q are connected subgraphs of T (hence trees), and $\{E_1, \dots, E_q\}$ partitions E . Moreover, the notation T_i/H_i , where $H_i \subseteq E_i$ is a subset of the edges of T_i , denotes the graph obtained from T_i by contracting H_i .

LEMMA 17. *Let $k \in \mathbb{Z}_{\geq 1}$, let (T, L) be a TAP instance, and let $x \in \Pi_T$. Then one can efficiently partition T into subtrees $T_i = (V_i, E_i)$ for $i \in [q]$ and find $H_i \subseteq E_i$ for $i \in [q]$ such that:*

- (i) *For $i \in [q]$, T_i/H_i is k -wide.*
- (ii) *One can efficiently get a link set $M \subseteq L$ with $|M| = O(1/\sqrt{k}) \cdot x(L)$ that covers all edges of $\cup_{i=1}^q H_i$.*
- (iii) *$\sum_{i=1}^q x(\text{cov}(E_i)) = (1 + O(1/\sqrt{k})) \cdot x(L)$.*

Moreover, one can choose the edge sets H_i for $i \in [q]$ to form subtrees, and in the k -wide trees T_i/H_i , the contracted node that corresponds to H_i can be chosen as the root.

The above lemma follows quite directly from techniques introduced by Adjashvili [1], with some smaller changes. Due to space constraints, we defer a formal proof of this statement to the long version of this paper. In what follows, we do not strictly need the fact that the H_i can be chosen to form trees. Still, we want to highlight this additional property, which can lead to considerable simplifications when designing reductions to wide instances that do not rely on the ellipsoid method. In particular, one could use our Lemma 17 in Adjashvili's approach [1], and avoid most technicalities involving so-called compound nodes.

Notice that Lemma 17 indeed implies that a solution $x \in \Pi_T$ to the cut-LP for a TAP instance can be broken down into solutions to the cut-LPs of independent k -wide TAP instances, by increasing the objective value by a factor $1 + O(1/\sqrt{k})$. To see this, consider the problem of finding a smallest set of links that covers only the edges E_i for some $i \in [q]$, where $T_j = (V_j, E_j)$ for $j \in [q]$ are the subtrees as stated in Lemma 17. This problem can be interpreted as a TAP instance on the subtree $T_i = (V_i, E_i)$, by interpreting each link $\ell \in L$ as a link that covers $P_\ell \cap E_i$. Hence, the vector x^i obtained from x by setting to 0 all coordinates corresponding to links in $L \setminus \text{cov}(E_i)$, is a cut-LP solution for the TAP instance on T_i . Moreover, the total cost of these cut-LP solutions is

$$\sum_{i=1}^q x^i(L) = \sum_{i=1}^q x(\text{cov}(E_i)) = (1 + O(1/\sqrt{k})) \cdot x(L),$$

where the last equality follows by point (iii) of Lemma 17, and is hence only by a factor $1 + O(1/\sqrt{k})$ larger than the LP cost of x . Finally, by Lemma 17 (ii) we can cover all edges in $\cup_{i=1}^q H_i$ at small cost. Thus, for each $i \in [q]$, it remains to cover the edges $E_i \setminus H_i$, which, after contracting H_i which we already covered, reduces to a k -wide TAP instance on T_i/H_i by point (i) of Lemma 17.

The above outline highlights how one can try to reduce LP-based approaches to wide instances. In particular, it shows that if $O(1)$ -wide TAP instances have small integrality gap with respect to the cut-LP, then the cut-LP has a small integrality gap for any instance. A similar reduction approach was key in some recent progress on TAP, and also WTAP with a bounded ratio between maximal and minimal cost [1, 9]. Even though such reductions have a certain versatility, so far, they only have been carried out for specific linear programs. For example, Adjashvili [1] extended the cut-LP with

“bundle”-constraints, and Fiorini et al. [9] additionally added $\{0, 1/2\}$ -Chvátal-Gomory-constraints. In both cases the authors had to show that key properties are preserved by the decomposition. In the following, we show how, based on the decomposition guaranteed by Lemma 17, one can reduce TAP instances to $O(1)$ -wide ones without assuming a particular LP-based approach.

To obtain a black-box reduction, we show that the following holds. Given an α -approximation algorithm for k -wide TAP instances, one can efficiently compute a point $x \in \Pi_T$ and an α -approximation for each k -wide tree in the decomposition obtained by Lemma 17. This result is formalized in the following statement.

LEMMA 18. *Let $k \in \mathbb{Z}_{\geq 1}$. Given an α -approximation algorithm \mathcal{A} for TAP on k -wide instances, one can, for any TAP instance (T, L) , efficiently compute:*

- (i) *A point $x \in \Pi_T$ satisfying $x(L) \leq v^*$, where v^* is the optimal value of (T, L) ,*
- (ii) *a partition $T_i = (V_i, E_i)$ with $H_i \subseteq E_i \forall i \in [q]$ of T w.r.t. x as described by Lemma 17, and*
- (iii) *for $i \in [q]$, a link set $L_i \subseteq L$ covering $E_i \setminus H_i$ and satisfying $|L_i| \leq \alpha \cdot x(\text{cov}(E_i \setminus H_i))$.*

We first observe that Theorem 3 is an immediate consequence of the above lemma.

PROOF OF THEOREM 3. Because the partition $T_i = (V_i, E_i)$ and $H_i \subseteq E_i$ for $i \in [q]$ of T satisfy the conditions of Lemma 17, we can efficiently obtain a set of links $M \subseteq L$ by Lemma 17 such that $|M| = O(1/\sqrt{k}) \cdot x(L)$ and M covers all edges of $\cup_{i=1}^q H_i$. We return the solution $Q := M \cup \cup_{i=1}^q L_i$. Q is clearly a solution to the TAP instance (T, L) . It remains to show that its approximation guarantee is $\alpha + O(1/\sqrt{k})$, which follows due to

$$\begin{aligned} \left| M \cup \bigcup_{i=1}^q L_i \right| &\leq |M| + \sum_{i=1}^q |L_i| \\ &\leq O\left(\frac{1}{\sqrt{k}}\right) \cdot x(L) + \alpha \cdot \sum_{i=1}^q x(\text{cov}(E_i \setminus H_i)) \\ &\leq O\left(\frac{1}{\sqrt{k}}\right) \cdot x(L) + \alpha \cdot \left(1 + O\left(\frac{1}{\sqrt{k}}\right)\right) \cdot x(L) \\ &= \left(\alpha + O\left(\frac{1}{\sqrt{k}}\right)\right) \cdot x(L) \\ &= \left(\alpha + O\left(\frac{1}{\sqrt{k}}\right)\right) \cdot v^*, \end{aligned}$$

where the second inequality follows from Lemmas 17 (ii) and 18 (iii), the third inequality follows from Lemma 17 (iii), and the last equation from Lemma 18 (i). \square

It remains to show Lemma 18. To this end, let (T, L) be a TAP instance on a tree $T = (V, E)$ with links $L \subseteq \binom{V}{2}$. Consider the following polytope, parameterized by a value $v \in \mathbb{Z}_{\geq 0}$ and $k \in \mathbb{Z}_{\geq 1}$, and where we denote by $\text{OPT}(S)$, for any edge set $S \subseteq E$, the minimum number of links needed to cover all edges of S .

$$\Pi_T(v, k) := \left\{ x \in \Pi_T \mid \begin{array}{l} x(\text{cov}(U \setminus H)) \geq \text{OPT}(U \setminus H) \\ \forall U \subseteq E, H \subseteq U \text{ s.t.} \\ (V, U)/H \text{ is } k\text{-wide,} \\ x(L) \leq v \end{array} \right\}$$

To clarify, stating that $(V, U)/H$ is k -wide in the first set of constraints, implies that $U \subseteq E$ and $H \subseteq U$ must be such that $(V, U)/H$ is a tree (that is k -wide).

In words, $\Pi_T(v, k)$ describes all cut-LP solutions of value at most v with the additional requirement that on any k -wide sub-instance $(V, U)/H$ defined by edge sets $H \subseteq U \subseteq E$, the x -value of the links covering edges in $(V, U)/H$ is no lower than the optimal value for this sub-instance. Notice that if $v^* \in \mathbb{Z}_{\geq 0}$ is the optimal value of the original instance, then $\Pi_T(v^*, k)$ contains the characteristic vector of an optimal solution, and is thus indeed a relaxation of the considered TAP problem. Note that we can easily “guess” the optimal value v^* , because it lies within $\{0, \dots, |L|\}$, and thus assume that v^* is known.⁴

Using Lemma 17, it is not hard to see that the integrality gap of the linear program $\min\{x(L) \mid x \in \Pi_T(v^*, k)\}$ is bounded by $1 + O(1/\sqrt{k})$.⁵ However, optimizing (or separating) over $\Pi_T(v^*, k)$ is difficult. We therefore use a *partial separation oracle* that correctly never separates a point $y \in \mathbb{R}^L$ that is in $\Pi_T(v^*, k)$, but may sometimes not separate points $y \in \mathbb{R}^L$ that are not in $\Pi_T(v^*, k)$. The partial separation oracle for a point $y \in \mathbb{R}^L$, and an arbitrary value $v \in \mathbb{Z}_{\geq 0}$, is described by Algorithm 1.

For simplicity, our description is given for \mathcal{A} being a deterministic α -approximation algorithm for k -wide instances. The extension to randomized algorithms \mathcal{A} can be done with standard techniques. More precisely, if \mathcal{A} is a randomized procedure, one can, for any constant $\delta > 0$, obtain a $(\alpha + \delta)$ -approximation with high probability in polynomial time, by doing multiple runs of \mathcal{A} on the same instance and returning the best outcome.⁶ Hence, because we focus on polynomial-time procedures, we can assume that with high probability, \mathcal{A} returns a $(\alpha + \delta)$ -approximation whenever it is called in the following.

Notice that whenever Algorithm 1 returns a separating hyperplane for some point $y \in \mathbb{R}^L$, then this is indeed a valid one. This is clear for the first two steps of the algorithm. Moreover, if step 3 returns a hyperplane $x(\text{cov}(E_i \setminus H_i)) \geq \lceil |L_i|/\alpha \rceil$, then this hyperplane indeed separates y over $\Pi_T(v, k)$ because $|L_i| \leq \alpha \cdot \text{OPT}(E_i \setminus H_i)$, as L_i is an α -approximate solution for the TAP problem on T_i , and $\text{OPT}(E_i \setminus H_i)$ is integer. Thus, $\lceil |L_i|/\alpha \rceil \leq \text{OPT}(E_i \setminus H_i)$.

We now show that a point $x \in \Pi_T$, and links $L_i \subseteq L$ for $i \in [q]$ as guaranteed by Lemma 18, can be obtained efficiently by using any version of the ellipsoid method that efficiently returns a point in $\Pi_T(v^*, k)$; however we run the ellipsoid method with the partial

Algorithm 1: Partial separation oracle for $\Pi_T(v, k)$

- (1) If $y(L) > v$, return the separating hyperplane $x(L) \leq v$.
 - (2) If $y \notin \Pi_T$, return a separating hyperplane that separates y from Π_T .
 - (3) Decompose T into $T_i = (V_i, E_i)$ for $i \in [q]$ and find sets $H_i \subseteq E_i$ for $i \in [q]$ by using Lemma 17 with $x = y$ and parameter k . For each $i \in [q]$, use \mathcal{A} to obtain an α -approximation to the k -wide TAP instance T_i/H_i ; let $L_i \subseteq L$ be the obtained solution. If $y(\text{cov}(E_i \setminus H_i)) < \lceil |L_i|/\alpha \rceil$ for some $i \in [q]$, then return the separating hyperplane $x(\text{cov}(E_i \setminus H_i)) \geq \lceil |L_i|/\alpha \rceil$.
-

separation oracle given by Algorithm 1 instead of a true separation oracle.

PROOF OF LEMMA 18. As mentioned, we could guess the optimal value v^* of the TAP problem, because $v^* \in \{0, \dots, |L|\}$. However, we will instead use a binary search technique that only leads to an additional multiplicative factor of $\log |L|$ in the running time (instead of $|L|$). This is useful in our discussion of how our reduction can be extended to some weighted instances.

Consider a fixed $v \in \{0, \dots, |L|\}$. We run a classical ellipsoid type method to determine a point in $\Pi_T(v, k)$, if $\Pi_T(v, k)$ is non-empty. If we had a true (instead of a partial) separation oracle for $\Pi_T(v, k)$, then this could be done with a polynomial number of oracle calls and further operations taking polynomial time in the encoding lengths of the constraints. Notice that the constraints defining $\Pi_T(v, k)$ have small encoding length because all left-hand sides are 0/1-vectors, and the right-hand sides are integers within $\{0, \dots, |L|\}$. (We refer the interested reader to the excellent book of Grötschel, Lovász, and Schrijver [12] for details on the ellipsoid method.) This implies that, assuming $\Pi_T(v, k) \neq \emptyset$, if we run the ellipsoid method to find a point in $\Pi_T(v, k)$ with the partial separation oracle, then one of the following happens:

- (i) The partial separation oracle will throughout the ellipsoid method always return a separating hyperplane whenever it is called for a point $y \notin \Pi_T(v, k)$. In this case, we get a point $x \in \Pi_T(v, k)$ in polynomial time as desired.
- (ii) At some point during the ellipsoid method, the separation oracle will be invoked with a point $y \in \mathbb{R}^L$ for which no separating hyperplane is found.

By employing binary search on $v \in \{0, \dots, |L|\}$, we find a value $\bar{v} \in \{0, \dots, |L|\}$ such that running the above-described ellipsoid method with partial separation oracle on $\Pi_T(\bar{v} - 1, k)$ returns that $\Pi_T(\bar{v} - 1, k) = \emptyset$; however, for $\Pi_T(\bar{v}, k)$ it returns that the polytope is not empty. Such a \bar{v} can clearly be found through binary search over $\{0, \dots, |L|\}$ by $O(\log |L|)$ calls of the ellipsoid method with partial separation oracle. Also notice that $\bar{v} \leq v^*$, because the partial separation oracle is “weaker” than a true separation oracle, i.e., whenever it separates a point $y \in \mathbb{R}^L$, then $y \notin \Pi_T(v, k)$; however, sometimes points not in $\Pi_T(v, k)$ do not get separated.

Consider the run of the ellipsoid method with partial separation oracle on $\Pi_T(\bar{v}, k)$. If (i) happens, then we get a point $x \in \Pi_T(\bar{v}, k)$ and a corresponding partition of T into $T_i = (V_i, E_i)$ for $i \in [q]$ and

⁴As we will discuss later, we can also use a binary search technique to get an approximation of v^* that is good enough for us. This speeds up the procedure, which is interesting when extending our reduction to weighted instances with small c_{\max}/c_{\min} .

⁵Indeed, we can apply Lemma 17 to an optimal solution y to $\min\{x(L) \mid x \in \Pi_T(v^*, k)\}$, because $\Pi_T(v^*, k) \subseteq \Pi_T$. This allows for partitioning T into subtrees $T_i = (V_i, E_i)$ for $i \in [q]$ with $H_i \subseteq E_i$ that satisfy the conditions of Lemma 17 with respect to y . Point (ii) of the lemma shows that covering all edges $\cup_{i=1}^q H_i$ can be done with at most $O(1/\sqrt{k}) \cdot y(L)$ links. Moreover, Lemma 17 (iii) shows that if, for each $i \in [q]$, the edge set $E_i \setminus H_i$ can be covered by at most $y(\text{cov}(E_i))$ links, then the obtained solution uses at most $(1 + O(1/\sqrt{k}))y(L)$ links, thus leading to an integrality gap of $1 + O(1/\sqrt{k})$. This is indeed possible because $\Pi_T(v^*, k)$ contains constraints that guarantee that $E_i \setminus H_i$ can be covered by at most $y(\text{cov}(E_i \setminus H_i)) \leq y(\text{cov}(E_i))$ links, as long as $(V, E_i) \setminus H_i$ is k -wide, which holds due to Lemma 17 (i).

⁶Notice that a usual application of Markov’s inequality would guarantee that we obtain a $(1 + O(\delta)) \cdot \alpha$ -approximation with high probability. However, because we can assume $\alpha \leq 2$, as 2-approximations are well-known even for WTAP, this implies that we can obtain an $(\alpha + \delta)$ -approximation.

$H_i \subseteq E_i$ for $i \in [q]$, as stated in Lemma 17. Running for $i \in [q]$ the α -approximation \mathcal{A} on the k -wide instance $T[E_i]/H_i$, a link set $L_i \subseteq L$ is obtained such that

$$|L_i| \leq \alpha \cdot \text{OPT}(E_i \setminus H_i) \leq \alpha \cdot x(\text{cov}(E_i \setminus H_i)) ,$$

where the second inequality follows from $x \in \Pi_T(\bar{v}, k)$. Hence, the point x together with the partition of T , and the links L_i for $i \in [q]$ fulfill the conditions of Lemma 18 as desired.

Now assume that (ii) occurs, and let $x \in \mathbb{R}^L$ be a point for which the ellipsoid method called the partial separation oracle, and no separating hyperplane was returned for x . The partial separation oracle computed a partition of T into $T_i = (V_i, E_i)$ with sets $H_i \subseteq E_i$ for $i \in [q]$ using Lemma 17. Moreover, for each $i \in [q]$, a link set $L_i \subseteq L$ was determined that satisfies $x(\text{cov}(E_i \setminus H_i)) \geq \lceil |L_i|/\alpha \rceil$, because x was not separated by step 3 of the partial separation procedure. Hence, for each $i \in [q]$, we have $|L_i| \leq \alpha \cdot x(\text{cov}(E_i \setminus H_i))$. Again, the point x together with the partition of T and the links L_i for $i \in [q]$ satisfy the conditions of Lemma 18. \square

C DERANDOMIZED REWIRINGS

In this section, we show that the randomized part of our approach can be completely derandomized. More precisely, we will sketch a proof of the following variant of Lemma 15:

LEMMA 19. *There exists a deterministic approximation algorithm for $O(1)$ -wide TAP instances with approximation guarantee $2(\sqrt{3} - 1) < 1.465$.*

We stress that we choose to sketch the derandomization of the slightly weaker (in terms of approximation ratio) Lemma 15 for the sake of simplicity. Nonetheless, the derandomization of Theorem 1 works in an almost identical way.

Let us review our rounding procedure. Clearly, the only randomized part is the algorithm in Lemma 12. In particular, given a TAP instance $(T = (V, E), L)$ where T is a k -wide tree, and a k -wide-LP solution (x, λ) , this algorithm samples a local solution L_i^R for each principal subtree T_i and each $R \in \Lambda_i$ with probability λ_i^R . Let L_i be the random set of links we sample for principal subtree T_i , let A_i be the random set of active nodes (i.e., a node in some T_i that is an endpoint of a cross-link in L_i), and let A be the random variable $\cup_{i \in [q]} A_i$. Note that A_i is completely determined by L_i . The key idea behind our derandomization is the following: since L_i has a domain whose size is polynomial in $|V|$, we can apply the method of *Conditional Expectations* in order to fix a local solution for each principal subtree deterministically.

More precisely, we begin by fixing a matching M according to Lemma 13; remember that this matching is selected deterministically. Then, our rounding algorithm returns a tree augmentation whose size is at most

$$\mathbb{E} \left[\sum_{i=1}^q |L_i| - \left| M \cap \binom{A}{2} \right| \right].$$

Clearly, we can always compute $\mathbb{E}[|L_i|]$. Crucially, we observe that L_i is independent of L_j for any principal subtrees T_i and T_j (and hence the same is true for A^i and A^j). This implies that we can compute $\mathbb{E} \left[\left| M \cap \binom{A}{2} \right| \right]$, since we know $\Pr[u \in A_i, v \in A_j]$ for any principal subtrees T_i, T_j and any vertices u, v in them respectively.

Furthermore, since the domain size of any L_i is polynomial for any principal subtree T_i , we can enumerate over its domain. In particular, we can condition on $L_1 = L_1^{R_1}$, where

$$R_1 = \operatorname{argmin}_{R \in \Lambda_1} \mathbb{E} \left[\sum_{i \in [q]} |L_i| - \left| M \cap \binom{A}{2} \right| \middle| L_1 = L_1^R \right].$$

By applying this reasoning inductively, we can condition on random variable L_2, L_3 etc., until we find some

$$R_j = \operatorname{argmin}_{R \in \Lambda_j} \mathbb{E} \left[\sum_{i \in [q]} |L_i| - \left| M \cap \binom{A}{2} \right| \middle| L_t = L_t^{R_t} \forall t \leq j-1 \right],$$

for all $j \in [q]$. Since the expected cost does not increase as we condition on a variable, we found a point in the sample space of variables L_i , for all $i \in [q]$, whose cost is at most the cost of

$$\mathbb{E} \left[\sum_{i \in [q]} |L_i| - \left| M \cap \binom{A}{2} \right| \right].$$

By Lemmas 12 and 13 we thus obtain

$$\mathbb{E} \left[\sum_{i \in [q]} |L_i| - \left| M \cap \binom{A}{2} \right| \right] \leq x(L_{\text{in}}) + 2x(L_{\text{cross}}) - \frac{x^2(L_{\text{cross}}^{\text{crit}})}{|V_{\text{crit}}|} ,$$

which is the same relation obtained for the randomized procedure, and thus the analysis in Lemma 15 applies, which immediately implies Lemma 19.

D A REFINED APPROXIMATION FACTOR

In order to establish a stronger approximation guarantee than that of Lemma 15, we first provide a strengthening of Lemma 13, described by the lemma below. Its proof is almost identical to the proof of Lemma 13, except for a more careful application of the Cauchy-Schwarz inequality:

LEMMA 20. *Let $G = (V, E)$ be an undirected graph, let $x \in \mathbb{R}_{\geq 0}^E$ such that for any $v \in V$ we have $x(\delta(v)) \leq 1$, and let A be a random subset of V with a distribution that satisfies:*

- (i) $\Pr[v \in A] = x(\delta(v)) \quad \forall v \in V$, and
- (ii) $\Pr[v \in A \text{ and } u \in A] = \Pr[u \in A] \cdot \Pr[v \in A] \quad \forall \{u, v\} \in E$.

Moreover, let $z \in \mathbb{R}_{\geq 0}^E$ by a vertex of the polytope

$$Q = \left\{ y \in \mathbb{R}_{\geq 0}^E \mid y(\delta(v)) = x(\delta(v)) \quad \forall v \in V \right\} .$$

Then,

- $|\operatorname{supp}(z)| \leq |V|$, and
- there exists a matching $M \subseteq E$ in G such that

$$\mathbb{E} \left[\left| M \cap \binom{A}{2} \right| \right] \geq \sum_{e \in E} z_e^2 \geq \frac{(\sum_{e \in E_1} z_e)^2}{|E_1|} + \frac{(\sum_{e \in E_2} z_e)^2}{|E_2|} ,$$

for any $E_1, E_2 \subseteq \operatorname{supp}(z)$ with $E_1 \cap E_2 = \emptyset$.

PROOF. Clearly, $x \in Q$. Moreover, for any vertex z of Q we indeed have $|\operatorname{supp}(z)| \leq |V|$. This follows through the following standard sparsity argument: Any vertex z of Q is defined by $|E|$ -many linearly independent and tight constraints; however, Q has only $|V|$ constraints that are not nonnegativity constraints, and

therefore, at least $|E| - |V|$ of the nonnegativity constraints must be tight at z , which implies $|\text{supp}(z)| \leq |V|$.

We let the matching $M \subseteq E$ be a matching obtained by the greedy algorithm for maximum weight matchings with respect to the weights z . More formally, we order the edges $E = \{e_1, \dots, e_m\}$ such that $z(e_1) \geq \dots \geq z(e_m)$. We start with $M = \emptyset$ and go through the edges in the order e_1, \dots, e_m . When considering e_i , we set $M = M \cup \{e_i\}$ if $M \cup \{e_i\}$ is a matching.

We have the following (which holds for any matching M):

$$\begin{aligned} \mathbb{E} \left[\left| M \cap \binom{A}{2} \right| \right] &= \sum_{\{u,v\} \in M} x(\delta(u)) \cdot x(\delta(v)) \\ &= \sum_{\{u,v\} \in M} z(\delta(u)) \cdot z(\delta(v)) , \end{aligned} \quad (13)$$

which is an immediate consequence of property (i) and (ii), and of $z \in Q$. Next, we show that the following holds:

$$\sum_{\{u,v\} \in M} z(\delta(u)) \cdot z(\delta(v)) \geq \sum_{e \in E} z(e)^2 . \quad (14)$$

Note that expanding the left-hand side of (14) by using $z(\delta(u)) = \sum_{e \in \delta(u)} z(e)$ and $z(\delta(v)) = \sum_{e \in \delta(v)} z(e)$, leads to a sum of terms of the form $z(e) \cdot z(f)$ for some pairs $e, f \in E$. For any $e = \{u, v\} \in M$, the left-hand side of (14) has a term of the form $z(e)^2$, due to the term $z(\delta(u)) \cdot z(\delta(v))$. Now consider any non-matching edge $e = \{u, v\} \in E \setminus M$. Because e is a non-matching edge, and we constructed M by the greedy algorithm, there is at least one endpoint of e , say u , such that there is an edge $f \in M$ incident with u that satisfies $z(f) \geq z(e)$. This implies that $z(\delta(u)) \cdot z(\delta(v))$ contains a term $z(f) \cdot z(v) \geq z(e) \cdot z(v) \geq z(e)^2$. Hence, for each $e \in E$, we identified a term on the left-hand side of (14) of value at least $z(e)^2$; moreover, all such terms on the left-hand side are different, which implies (14).

The statement of the lemma now follows by combining (13) and (14) with the following inequality, which holds for any $E_1, E_2 \subseteq \text{supp}(z)$ such that $E_1 \cap E_2 = \emptyset$, and follows by the Cauchy-Schwarz inequality:

$$\sum_{e \in E} z(e)^2 \geq \sum_{e \in E_1} z(e)^2 + \sum_{e \in E_2} z(e)^2 \geq \frac{z(E_1)^2}{|E_1|} + \frac{z(E_2)^2}{|E_2|} .$$

□

Given the above lemma, we are now ready to prove the randomized version of Theorem 1:

PROOF OF THEOREM 1. Given a k -wide TAP instance (T, L) with $k = O(1)$, and a solution (x, λ) to the k -wide-LP, by Lemma 12 we can randomly round (x, λ) to a solution \mathcal{B} such that

$$\mathbb{E}[\|\mathcal{B}\|] \leq x(L_{\text{in}}) + 2x(L_{\text{cross}}) - \mathbb{E}_A \left[\eta \left(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2} \right) \right] , \quad (15)$$

where A fulfills the conditions of Lemma 12.

We apply Lemma 20 to the graph $H = (V_{\text{crit}}, L_{\text{cross}}^{\text{crit}})$ and the restriction of x to $L_{\text{cross}}^{\text{crit}}$, where, as usual, $A \subseteq V_{\text{crit}}$ will be the random set of active nodes. To apply Lemma 20, we define $z \in \mathbb{R}_{\geq 0}^{L_{\text{cross}}^{\text{crit}}}$ to be any vertex of the polytope

$$Q := \left\{ y \in \mathbb{R}_{\geq 0}^{L_{\text{cross}}^{\text{crit}}} \mid y(\delta_H(v)) = x(\delta_H(v)) \quad \forall v \in V_{\text{crit}} \right\} .$$

Lemma 20 then implies that

$$|\text{supp}(z)| \leq |V_{\text{crit}}| , \quad (16)$$

and the expected size of a maximum cardinality matching in graph $(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2})$ is lower bounded by

$$\mathbb{E}_A \left[\eta \left(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2} \right) \right] \geq \frac{z(E_1)^2}{|E_1|} + \frac{z(E_2)^2}{|E_2|} ,$$

for any $E_1, E_2 \subseteq \text{supp}(z)$ with $E_1 \cap E_2 = \emptyset$.

Let L_{2l} be the links in $L_{\text{cross}}^{\text{crit}} \cap \text{supp}(z)$ whose endpoints are both leaves, and $L_{0/1l}$ be the remaining links in $L_{\text{cross}}^{\text{crit}} \cap \text{supp}(z)$. In what follows, we set $E_1 = L_{0/1l}$ and $E_2 = L_{2l}$, and will use inequality (17) for these two sets, thus leading to

$$\mathbb{E}_A \left[\eta \left(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2} \right) \right] \geq \frac{z(L_{2l})^2}{|L_{2l}|} + \frac{z(L_{0/1l})^2}{|L_{0/1l}|} . \quad (17)$$

Moreover, because z is a vertex of Q , we have

$$|L_{2l}| \leq K , \quad (18)$$

where K is the number of leaves of T . This follows by the fact that, except for nonnegativity constraints, there are only K constraints in Q that involve links with both endpoints being leaves.

Consider $z' \in \mathbb{R}^L$ defined by

$$z'_\ell = \begin{cases} z_\ell & \text{for } \ell \in L_{\text{cross}}^{\text{crit}} , \\ x_\ell & \text{otherwise} . \end{cases}$$

Observe that $z' \in \Pi_T$. We define $\phi \in [0, 1]$ so that $z'(L_{0/1l}) = \phi \cdot z'(L_{\text{cross}}^{\text{crit}}) = \phi \cdot x(L_{\text{cross}}^{\text{crit}})$. (Hence $z'(L_{2l}) = (1 - \phi) \cdot z'(L_{\text{cross}}^{\text{crit}}) = (1 - \phi) \cdot x(L_{\text{cross}}^{\text{crit}})$.)

We proceed by deriving an upper bound for $\mathbb{E}[\|\mathcal{B}\|]/\text{OPT}^*$, where $\text{OPT}^* = x(L) = z'(L)$ is the optimal value of the k -wide-LP, as usual. Notice that $|V_{\text{crit}}| < 2K$, which readily follows from the fact that the average degree of T must be strictly below 2, because T is a tree. Moreover, we must have

$$\sum_{\ell \in L: v \in \ell} z'_\ell \geq 1 \quad \text{for each leaf node } v \in V$$

for the edge incident with v to be covered by the LP solution $z' \in \Pi_T$. Each link can contribute to at most two of the sums above, for two different leaf nodes, and the links in $L_{\text{up}} \cup L_{\text{cross}}^{\text{no-crit}} \cup L_{0/1l}$ to at most one of these sums. Thus we can conclude that

$$\begin{aligned} 2z'(L_{\text{cross}}) + 2z'(L_{\text{in}}) - z'(L_{\text{up}}) - z'(L_{\text{cross}}^{\text{no-crit}}) - z'(L_{0/1l}) \\ \geq K \geq \frac{|V_{\text{crit}}|}{2} . \end{aligned} \quad (19)$$

Let $\alpha_{\text{in}} = x(L_{\text{in}})/\text{OPT}^* = z'(L_{\text{in}})/\text{OPT}^*$ be the fraction of OPT^* that corresponds to links in L_{in} . We define analogously $\alpha_{\text{up}}, \alpha_{\text{cross}}, \alpha_{\text{cross}}^{\text{crit}}, \alpha_{\text{cross}}^{\text{no-crit}}, \alpha_{2l}$, and $\alpha_{0/1l}$. By Lemma 13, the expected size of the maximum cardinality matching in $(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2})$ satisfies

$$\begin{aligned} \mathbb{E}_A \left[\eta \left(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2} \right) \right] &\geq \frac{x(L_{\text{cross}}^{\text{crit}})^2}{|V_{\text{crit}}|} = \frac{z'(L_{\text{cross}}^{\text{crit}})^2}{|V_{\text{crit}}|} \\ &\geq \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}} - 2\alpha_{0/1l}} \cdot \text{OPT}^* , \end{aligned} \quad (20)$$

where the last inequality above follows from (19). Combining (15) and (20), we conclude that

$$\begin{aligned} \frac{\mathbb{E}[|\mathcal{B}|]}{\text{OPT}^*} &\leq \alpha_{\text{in}} + 2\alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}} - 2\alpha_{0/11}} \\ &= 1 + \alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}} - 2\phi\alpha_{\text{cross}}^{\text{crit}}}. \end{aligned} \quad (21)$$

Observe that $|L_{2l}| + |L_{0/11}| = |\text{supp}(z)| \leq |V_{\text{crit}}| \leq 2K$, where the first inequality follows from 16. For simplicity, we define $q := K - |L_{2l}|$, and notice that $q \geq 0$ due to (18). Hence, $|L_{2l}| = K + q$ and $|L_{0/11}| \leq 2K - |L_{2l}| = K + q$. We thus obtain by (17)

$$\begin{aligned} \mathbb{E}_A \left[\eta(A, L_{\text{cross}}^{\text{crit}} \cap \binom{A}{2}) \right] &\geq \frac{z'(L_{2l})^2}{K - q} + \frac{z'(L_{0/11})^2}{K + q} \\ &= \frac{((1 - \phi)z'(L_{\text{cross}}^{\text{crit}}))^2}{K - q} + \frac{(\phi z'(L_{\text{cross}}^{\text{crit}}))^2}{K + q}. \end{aligned} \quad (22)$$

For $\phi \leq 0.5$, the right-hand side of (22) is growing in q , hence the worst matching size cost is achieved for $q = 0$ in this case, and is at least

$$\frac{1}{K} \cdot \left(((1 - \phi)z'(L_{\text{cross}}^{\text{crit}}))^2 + (\phi z'(L_{\text{cross}}^{\text{crit}}))^2 \right),$$

which implies together with (15) the following inequality for any $\phi \leq 0.5$:

$$\frac{\mathbb{E}[|\mathcal{B}|]}{\text{OPT}^*} \leq 1 + \alpha_{\text{cross}} - \frac{((1 - \phi)\alpha_{\text{cross}}^{\text{crit}})^2 + (\phi\alpha_{\text{cross}}^{\text{crit}})^2}{2 - \alpha_{\text{up}} - \alpha_{\text{cross}}^{\text{no-crit}} - \phi\alpha_{\text{cross}}^{\text{crit}}}. \quad (23)$$

Finally, from Corollary 6 we know that we can round x to a solution \mathcal{A} such that

$$\frac{|\mathcal{A}|}{\text{OPT}^*} \leq 2\alpha_{\text{in}} - \alpha_{\text{up}} + \alpha_{\text{cross}} = 2 - \alpha_{\text{cross}} - \alpha_{\text{up}}, \quad (24)$$

where the equality follows from $\alpha_{\text{cross}} + \alpha_{\text{in}} = 1$.

Overall, from (24), (21), and (23), we get that if $\phi \leq 0.5$, the approximation ratio is at most

$$1 + \min \left\{ 1 - \alpha_{\text{cross}} - \alpha_{\text{up}}, \alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}} - 2\phi\alpha_{\text{cross}}^{\text{crit}}}, \alpha_{\text{cross}} - \frac{((1 - \phi)\alpha_{\text{cross}}^{\text{crit}})^2 + (\phi\alpha_{\text{cross}}^{\text{crit}})^2}{2 - \alpha_{\text{up}} - \alpha_{\text{cross}}^{\text{no-crit}} - \phi\alpha_{\text{cross}}^{\text{crit}}} \right\}, \quad (25)$$

and if $\phi \geq 0.5$ it is at most

$$1 + \min \left\{ 1 - \alpha_{\text{cross}} - \alpha_{\text{up}}, \alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}}^{\text{crit}})^2}{4 - 2\alpha_{\text{up}} - 2\alpha_{\text{cross}}^{\text{no-crit}} - 2\phi\alpha_{\text{cross}}^{\text{crit}}} \right\}. \quad (26)$$

First, observe that all terms inside the minima in (25) and (26) are non-increasing in α_{up} and $\alpha_{\text{cross}}^{\text{no-crit}}$. Therefore, the maxima of both (25) and (26) over all values $\alpha_{\text{cross}}, \alpha_{\text{up}}, \alpha_{\text{cross}}^{\text{no-crit}}, \alpha_{\text{cross}}^{\text{crit}} \in [0, 1]$ have $\alpha_{\text{up}} = \alpha_{\text{cross}}^{\text{no-crit}} = 0$. Hence, we focus on finding the maxima of (25) and (26) for $\alpha_{\text{up}} = \alpha_{\text{cross}}^{\text{no-crit}} = 0$. Furthermore, if $\phi \geq 0.5$, then (26) is non-increasing in ϕ , which means we can also assume $\phi \leq 0.5$.

Therefore, we conclude the ratio is at most

$$1 + \min_{\alpha_{\text{cross}} \in [0, 1], \phi \in [0, 1/2]} \left\{ 1 - \alpha_{\text{cross}}, \alpha_{\text{cross}} - \frac{(\alpha_{\text{cross}})^2}{4 - 2\phi\alpha_{\text{cross}}}, \alpha_{\text{cross}} - \frac{((1 - \phi)\alpha_{\text{cross}})^2 + (\phi\alpha_{\text{cross}})^2}{2 - \phi\alpha_{\text{cross}}} \right\}.$$

Finally, observe that in this range, the third term is always dominated by the second one because $1 - 2\phi + 2\phi^2 \geq 0.5$. Therefore, the approximation ratio is at most

$$1 + \min_{\alpha_{\text{cross}} \in [0, 1], \phi \in [0, 1/2]} \left\{ 1 - \alpha_{\text{cross}}, \alpha_{\text{cross}} - \frac{((1 - \phi)\alpha_{\text{cross}})^2 + (\phi\alpha_{\text{cross}})^2}{2 - \phi\alpha_{\text{cross}}} \right\}.$$

We finish the proof by showing that the maximum ratio attainable under these constraints is $\sqrt{34}/4 < 1.458$.

Let $f_1(x, y) := 1 - x$, $f_2(x, y) := x - x^2 \cdot \frac{1 - 2y + 2y^2}{2 - xy}$, and $f(x, y) := \min\{f_1(x, y), f_2(x, y)\}$. In the sequel, we will only consider $(x, y) \in [0, 1] \times [0, \frac{1}{2}]$, since our goal is to find

$$\max_{(x, y) \in [0, 1] \times [0, \frac{1}{2}]} f(x, y).$$

Let $D = [0, 1] \times [0, \frac{1}{2}]$. Consider any maximizer (x^*, y^*) of f over D . We will show that (x^*, y^*) satisfies $f_1(x^*, y^*) = f_2(x^*, y^*)$.

To this end, observe that if $f_1(x^*, y^*) < f_2(x^*, y^*)$, then (x^*, y^*) is a local maximizer of f_1 . This directly implies that $x^* = 0$, which holds for any local maximizer of f_1 in D . However, $f_1(0, y) = 1 > f_2(0, y)$ for any $y \in [0, \frac{1}{2}]$, which is a contradiction.

Conversely, consider the case $f_1(x^*, y^*) > f_2(x^*, y^*)$. First of all, observe that

$$\frac{\partial}{\partial x} f_2(x, y) = 1 - (1 - 2y + 2y^2) \cdot \frac{2x(2 - xy) + yx^2}{(2 - xy)^2}.$$

Now, observe that $f(x^*, y^*) \geq f_2(\frac{1}{2}, \frac{1}{2}) \geq 0.4$. Since $f_1(x^*, y^*) > f_2(x^*, y^*) \geq 0.4$, it follows that $x^* \leq 0.6$. Furthermore, for $x \leq 0.6$, we have

$$\begin{aligned} \frac{\partial}{\partial x} f_2(x, y) &= 1 - (1 - 2y + 2y^2) \cdot \frac{4x - x^2y}{(2 - xy)^2} \\ &\geq 1 - \frac{\frac{12}{5} - \frac{9}{25}y}{1.7^2} \geq 1 - \frac{\frac{12}{5}}{1.7^2} > 0. \end{aligned}$$

Hence, $f_2(x, y)$ is strictly increasing in x for $x \leq 0.6$, which implies $x^* \geq 0.6$. Moreover, because $\frac{\partial}{\partial x} f_2(x, y)$ and $f_1(x, y)$ are continuous, we even get $x^* > 0.6$. However, this contradicts $x^* \leq 0.6$. Therefore, we conclude that for any maximizer (x^*, y^*) of f , $f_1(x^*, y^*) = f_2(x^*, y^*)$.

Now, we wish to solve $f_1(x, y) = f_2(x, y)$ for x . We have:

$$1 - x = x - x^2 \cdot \frac{1 - 2y + 2y^2}{2 - xy},$$

which implies

$$0 = x^2(1 + 2y^2) - x(4 + y) + 2,$$

and thus leads to

$$x = \frac{4 + y \pm \sqrt{(4 + y)^2 - 8(1 + 2y^2)}}{2(1 + 2y^2)} .$$

Since $x \in [0, 1]$, we need to pick

$$x = \frac{4 + y - \sqrt{(4 + y)^2 - 8(1 + 2y^2)}}{2(1 + 2y^2)} . \quad (27)$$

Finally, we want to maximize $f_1(x, y) = f_2(x, y) = 1 - x$, which is equal to the following:

$$\begin{aligned} & 1 - \frac{4 + y - \sqrt{(4 + y)^2 - 8(1 + 2y^2)}}{2(1 + 2y^2)} \\ = & 1 - \frac{1}{2(1 + 2y^2)} \frac{(4 + y)^2 - \left(\sqrt{(4 + y)^2 - 8(1 + 2y^2)}\right)^2}{4 + y + \sqrt{(4 + y)^2 - 8(1 + 2y^2)}} \\ = & 1 - \frac{8(1 + 2y^2)}{2(1 + 2y^2)(4 + y + \sqrt{(4 + y)^2 - 8(1 + 2y^2)})} \\ = & 1 - \frac{4}{4 + y + \sqrt{(4 + y)^2 - 8(1 + 2y^2)}} . \end{aligned}$$

Equivalently, we want to maximize

$$h(y) := 4 + y + \sqrt{(4 + y)^2 - 8(1 + 2y^2)} .$$

We have

$$\frac{\partial}{\partial y} h(y) = 1 + \frac{1}{2} \frac{8 - 30y}{\sqrt{(4 + y)^2 - 8(1 + 2y^2)}} .$$

Setting $\frac{\partial}{\partial y} h(y)$ to 0, we finally obtain

$$30y^2 - 16y + 1 = 0 ,$$

and thus

$$y = \frac{1}{30} \cdot (8 \pm \sqrt{34}) .$$

One can conclude that the value of y must be set to $y = (8 + \sqrt{34})/30$, because by setting $\frac{\partial}{\partial y} h(y)$ to 0, we equated $8 - 30y$ (which is negative for $y = (8 - \sqrt{34})/30$) to a square root of a positive real number. By checking that $h(0)$ and $h(1)$ are smaller than $h((8 + \sqrt{34})/30)$, we determine that $h(y)$ is indeed maximized for $y^* = (8 + \sqrt{34})/30$. By determining x^* from y^* through (27), we obtain $x^* = 2 - \sqrt{34}/4$, and the resulting approximation guarantee is therefore

$$1 + f(x^*, y^*) = 1 + f_1(x^*, y^*) = 2 - x^* = \frac{\sqrt{34}}{4} < 1.458 ,$$

as desired. \square

Acknowledgments

The first author was partially supported by Swiss National Science Foundation grant 200021_159697. The second and third author were supported by Swiss National Science Foundation grant 200021_165866. The authors are grateful to the anonymous reviewers for various helpful comments.

REFERENCES

- [1] D. Adjiashvili. Beating approximation factor two for weighted tree augmentation with bounded costs. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2384–2399, 2017.
- [2] H.-C. An, M. Singh, and O. Svensson. LP-based algorithms for capacitated facility location. *SIAM Journal on Computing*, 46(1):272–306, 2017.
- [3] R. D. Carr, L. K. Fleischer, V. J. Leung, and C. A. Phillips. Strengthening integrality gaps for capacitated network design and covering problems. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 106–115, 2000.
- [4] J. Cheriyan and Z. Gao. Approximating (unweighted) tree augmentation via lift-and-project, part II. *Algorithmica*, 2017.
- [5] J. Cheriyan, T. Jordán, and R. Ravi. On 2-coverings and 2-packings of laminar families. In *Proceedings of 7th Annual European Symposium on Algorithms (ESA)*, pages 510–520, 1999.
- [6] J. Cheriyan, H. Karloff, R. Khandekar, and J. Könemann. On the integrality ratio for tree augmentation. *Operations Research Letters*, 36(4):399–401, 2008.
- [7] N. Cohen and Z. Nutov. A $(1 + \ln 2)$ -approximation algorithm for minimum-cost 2-edge-connectivity augmentation of trees with constant radius. *Theoretical Computer Science*, 489:67–74, 2013.
- [8] G. Even, J. Feldman, G. Kortsarz, and Z. Nutov. A 1.8 approximation algorithm for augmenting edge-connectivity of a graph from 1 to 2. *ACM Transactions on Algorithms*, 5(2):21:1–21:17, 2009.
- [9] S. Fiorini, M. Groß, J. Könemann, and L. Sanità. Approximating weighted tree augmentation via Chvátal-Gomory Cuts. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 817–831, 2018.
- [10] G. N. Frederickson and J. JáJá. Approximation algorithms for several graph augmentation problems. *SIAM Journal on Computing*, 10(2):270–283, 1981.
- [11] M. X. Goemans, A. V. Goldberg, S. Plotkin, D. B. Shmoys, É. Tardos, and D. P. Williamson. Improved approximation algorithms for network design problems. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 223–232, 1994.
- [12] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer, second corrected edition, 1993.
- [13] K. Jain. A factor 2 approximation algorithm for the generalized Steiner Network Problem. *Combinatorica*, 21:39–60, 2001.
- [14] S. Khuller and R. Thurimella. Approximation algorithms for graph augmentation. *Journal of Algorithms*, 14(2):214–225, 1993.
- [15] J. Könemann. Improved approximation for tree augmentation via Chvátal-Gomory Cuts. Invited talk at MAPSP, 2017. Slides available at <http://www.mapsp2017.ma.tum.de/MAPSP2017-Koenemann.pdf>.
- [16] G. Kortsarz and Z. Nutov. LP-relaxations for tree augmentation. In *Proceedings of 19th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 13:1–13:16, 2016.
- [17] G. Kortsarz and Z. Nutov. A simplified 1.5-approximation algorithm for augmenting edge-connectivity of a graph from 1 to 2. *ACM Transactions on Algorithms*, 12(2):23:1–23:20, 2016.
- [18] R. Levi, A. Lodi, and M. Sviridenko. Approximation algorithms for the capacitated multi-item lot-sizing problem via flow-cover inequalities. *Mathematics of Operations Research*, 33(2):461–474, 2008.
- [19] H. Nagamochi. An approximation for finding a smallest 2-edge-connected subgraph containing a specified spanning tree. *Discrete Applied Mathematics*, 126(1):83–113, 2003.
- [20] Z. Nutov. On the tree augmentation problem. In *Proceedings of 25th Annual Symposium on Algorithms (ESA)*, pages 61:1–61:14, 2017.