A Better-Than-2 Approximation for Weighted Tree Augmentation[*]

Vera Traub[‡]   Rico Zenklusen[§]

Abstract

We present an approximation algorithm for Weighted Tree Augmentation with approximation factor $1 + \ln 2 + \varepsilon < 1.7$. This is the first algorithm beating the longstanding factor of 2, which can be achieved through many standard techniques.

1 Introduction

The Weighted Tree Augmentation Problem (WTAP) is among the most elementary and intensively studied connectivity augmentation problems. It asks how to increase the edge-connectivity of a graph from 1 to 2 in the cheapest possible way, and is formally described as follows. An instance consists of a spanning tree $G = (V, E)$ together with a set $L \subseteq \binom{V}{2}$ of candidate edges to be added to G, which are also called *links*, and positive link weights $w \colon L \to \mathbb{R}_{>0}$. The task is to find a minimum weight subset of links $F \subseteq L$ such that $(V, E \cup F)$ is 2-edge-connected.[1] It is easy to see that even though WTAP asks to augment the edge-connectivity of a spanning tree, it does capture the problem of increasing the edge-connectivity of an arbitrary 1-edge-connected graph G to 2, because contracting all 2-edge-connected components of G leads to an equivalent WTAP instance. More generally, it is well-known that the problem of increasing the edge-connectivity of a graph from k to $k + 1$, for any odd k, reduces to WTAP (see, e.g., Cheriyan, Jordán, and Ravi [CJR99]).

Already the unweighted version of WTAP, where all links have unit weight, is NP-hard, even on trees of diameter 4, as shown by Frederickson and JáJá [FJ81]. The reduction they used was extended by Kortsarz, Krauthgamer, and Lee [KKL04] to show APX-hardness of the unweighted version of WTAP. Therefore, WTAP and special cases thereof have been heavily studied under the aspect of approximation algorithms. Prior to our work, the best approximation factor for WTAP was 2, which was first shown by Frederickson and JáJá [FJ81] in the early '80s. They reduced the problem to finding a shortest arborescence, while losing a factor of 2. Frederickson and JáJá's procedure was subsequently simplified and significantly sped up by Khuller and Thurimella [KT93]. Moreover, many classical and very versatile techniques for network design problems developed later also lead to a 2-approximation for WTAP. This includes primal-dual approaches (see Goemans, Goldberg, Plotkin, Shmoys, Tardos, and Williamson [Goe+94]), the iterative rounding technique by Jain [Jai01], and various further methods that are readily adaptable to WTAP (for example, a flow-based method by Frank and Tardos [FT89] for certain directed connectivity problems; see also discussion in [KT93] on how this relates to WTAP). However, despite extensive work on the problem and variations thereof, and some progress on special cases (see Section 1.2), the four decades old approximation factor of 2 by Frederickson and JáJá [FJ81] remained the state of the art.

1.1 Our results

In this paper, we present the first better-than-2 approximation for WTAP.

Theorem 1. *For any $\varepsilon > 0$, there is a $(1 + \ln 2 + \varepsilon)$-approximation algorithm for WTAP.*

In particular, for small enough $\varepsilon > 0$, we have $1 + \ln 2 + \varepsilon < 1.7$. Our approach significantly deviates from the numerous recent techniques introduced in the context of TAP, i.e., the unweighted version of WTAP. More precisely, we develop a relative greedy algorithm, a method introduced by Zelikovsky [Zel96] in the context of the Steiner Tree problem. Cohen and Nutov [CN13] later employed this approach to get a better-than-2 approximation for WTAP with bounded diameter, which inspired this work. Like other relative greedy algorithms, we iteratively contract well-chosen components. However, contrary to previous relative greedy approaches, including the above-mentioned ones, we rely on a (exponentially large) class of super-constant size components. Nevertheless, we can efficiently find, within our novel class of components, the best one to contract next. This circumvents a key barrier of [CN13], namely that their components must have constant size to be able to enumerate over them, which requires the underlying tree to have constant diameter. Moreover, we prove an approximate decomposition theorem for our components, which guarantees the existence of a good way to split any WTAP solution into candidate components for contraction. This

---

[1]Depending on the literature, 0-weight links may be allowed. This easily reduces to the case of strictly positive weights after including all 0-weight links in the solution in a preprocessing step and continuing on the resulting reduced WTAP instance.

theorem, applied to an optimal WTAP solution, guarantees the existence of good components to contract. We provide further details on our approach in Section 3.

1.2 Further related work

Even though significant work on tree augmentation did not improve on the canonical approximation factor of 2 for WTAP, remarkable progress has been achieved for numerous special cases. In particular, for the unweighted version, which is often simply called the *Tree Augmentation Problem (TAP)*, a long line of research [Adj18; CG18a; CG18b; Che+08; CN13; Eve+09; Fio+18; FJ81; GKZ18; KT93; KN16; KN18; Nag03; Nut17; GKZ18] led to the currently best approximation factor of 1.393 [CTZ21]. This result even applies to the more general unweighted connectivity augmentation problem, which asks to increase the edge-connectivity of an arbitrary graph G by one unit. (See also [BGJ20; Nut20] for further recent results on connectivity augmentation.) Starting with the work of Adjiashvili [Adj18] and an elegant strengthening thereof by Fiorini, Groß, Könemann, and Sanità [Fio+18], approximation factors below 2 have been obtained for WTAP in the special case where the ratio of largest to smallest link weight is bounded by a constant. Subsequent further improvements on these procedures by Grandoni, Kalaitzis, and Zenklusen [GKZ18] and Cecchetto, Traub, and Zenklusen [CTZ21] allowed for achieving approximation factors below 1.5 for this case, and an elegant application of the round-or-cut framework, first employed in this context by Nutov [Nut17], allows for obtaining better-than-2 approximations for TAP even if the ratio of largest to smallest link weight is at most logarithmic in the size of the graph. Unfortunately, all these advances very deeply exploit that the max-to-min weight ratio of the links are bounded, and it seems highly unclear whether and how they could potentially be extended to WTAP. Other special cases of WTAP, where an improvement over the approximation factor of 2 has been achieved, is when the given tree has bounded diameter (see Cohen and Nutov [CN13]), or when an optimal solution to a natural LP has no small fractional values (see Iglesias and Ravi [IR18]).

A natural generalization of WTAP is the 2-*edge-connected spanning subgraph problem* (2-ECSS). Contrary to WTAP, instead of starting with a spanning tree, one starts with an empty graph and needs to find a minimum cost set of edges leading to a 2-edge-connected graph spanning all vertices. WTAP can be cast as 2-ECSS by assigning to all tree edges a weight of zero such that they can be selected for free. Also for 2-ECSS, the best known approximation factor is 2, which can be achieved through a variety of elegant techniques [KV94], including primal-dual methods [GW95] and iterative rounding [Jai01] (see also [LRS11; WS11]). Progress beyond the factor 2 has only been achieved for unweighted 2-ECSS, where the task is to select a smallest number of edges to obtain a 2-edge-connected spanning subgraph. After the first improvements [KV94; CSS01], this led to the currently best-known approximation factor of $4/3$ [HVV19; SV14].

1.3 Organization of the paper

We start with some very brief preliminaries in Section 2, which allows us to introduce basic terminology and notation, and formalize the well-known interpretation of WTAP as a covering problem. In Section 3, we then describe our relative greedy algorithm together with the main underlying results and show how they lead to Theorem 1. In particular, we also introduce our new class of components and describe their properties needed for our relative greedy procedure. These properties are proved in the following sections. More precisely, Section 4 proves a decomposition theorem that guarantees that we can make progress by contracting a component that minimizes a natural selection function. Finally, in Section 5, we show how a component minimizing the selection function can be found efficiently through dynamic programming.

2 Preliminaries

Let $(G = (V, E), L, w)$ be a WTAP instance. WTAP is naturally described as a covering problem, where the task is to select a minimum weight set of links that cover all 1-cuts of the given tree G, which are the cuts containing a single tree edge e. Hence, the 1-cuts correspond to the edges of E. The 1-cuts that a single link $\ell \in L$ covers are described by the edge set $P_\ell \subseteq E$ of the unique path in G between the endpoints of ℓ. Because the addition of a link set $F \subseteq L$ to G makes the graph 2-edge-connected if and only if F covers all 1-cuts of G, WTAP can be formalized as the following natural covering problem.

$$\min \left\{ \sum_{\ell \in F} w(\ell) \colon F \subseteq L, \bigcup_{\ell \in F} P_\ell = E \right\} \tag{WTAP}$$

We also say that a link ℓ *covers* the edges P_ℓ. Hence, a set of links is a WTAP solution if its links cover all edges.

For a link $\ell \in L$, we denote by $V_\ell \subseteq V$ the vertices of the path with edge set P_ℓ, including its endpoints. It is often convenient to assume that the WTAP instance is *shadow-complete*, which means that for every link $\ell \in L$ and every two distinct vertices $v_1, v_2 \in V_\ell$, there is also a link $\{v_1, v_2\} \in L$ of same weight as ℓ. In this case $\{v_1, v_2\}$ is called a *shadow* of ℓ. Note that a link $\ell_1 \in L$ is a shadow of $\ell_2 \in L$ if and only if $P_{\ell_1} \subseteq P_{\ell_2}$. Clearly, any WTAP instance can be transformed into an equivalent shadow-complete one by adding, for each link $\ell \in L$, all of its shadows, each with weight $w(\ell)$. Any such added shadow of a link $\ell \in L$ that gets selected in a solution can later be replaced by ℓ without changing the weight of the solution.

3 Relative greedy algorithm for WTAP

As mentioned, our $(1 + \ln 2 + \varepsilon)$-approximation for WTAP is a relative greedy algorithm. The concept of relative greedy algorithms has been introduced by Zelikovsky [Zel96] in the context of the Steiner Tree problem, and was later leveraged by Cohen and Nutov [CN13] for WTAP with bounded diameter. The idea is to start with a weak but very well-structured approximation to the problem at hand, and then successively improve this solution by replacing parts of it. In the context of Steiner Tree, Zelikovsky [Zel96] started with a simple 2-approximation obtained by computing a minimum spanning tree over the terminals. This solution was then improved by successively finding an edge set to connect constantly many terminals, which is also called a *component*, such that the cost of the component is (significantly) cheaper than the cost of the spanning tree edges that can be removed after including the component in the solution.

For WTAP, we start with the same highly structured 2-approximation as Cohen and Nutov [CN13] did, namely one only using so-called *up-links* that are non-overlapping. Given a WTAP instance $(G = (V, E), L, w)$, we fix an arbitrary vertex $r \in V$, which we call *root* from now on; an *up-link* (with respect to r) is a link $\ell \in L$ such that one of its endpoints is on the unique path in G between the root and the other endpoint. The statement below formalizes the properties of the 2-approximate starting solution that our relative greedy procedure aims at improving. We denote by $L_{\mathrm{up}} \subseteq L$ the set of all up-links and by $\mathrm{OPT} \subseteq L$ an optimal solution to the WTAP instance.

Lemma 2 ([CN13]). [2] *Let $(G = (V, E), L, w)$ be a shadow-complete instance of WTAP. Then we can in polynomial time compute a WTAP solution $U \subseteq L_{\mathrm{up}}$ such that*

- $w(U) \leq 2 \cdot w(\mathrm{OPT})$, *and*
- *the edge sets P_u for $u \in U$ are pairwise disjoint.*

Starting with a 2-approximate up-link solution $U^* \subseteq L_{\mathrm{up}}$ as guaranteed by Lemma 2, we seek to identify a subset of the links in U^* that can be replaced by a cheaper link set $C \subseteq L$. To this end, for any set of up-links $U \subseteq L_{\mathrm{up}}$ and link set $C \subseteq L$, we denote by

$$\mathrm{Drop}_U(C) := \left\{ u \in U : P_u \subseteq \bigcup_{\ell \in C} P_\ell \right\}$$

the links of U that only cover a subset of the edges covered by links in C, and can thus safely be removed from a solution once C is added.

In a general replacement step, we have some up-links $U \subseteq U^*$ with $U \neq \emptyset$ left in our current solution, and we seek to find a link set $C \in \mathfrak{L}$, where $\mathfrak{L} \subseteq 2^L$ is a well-chosen family, such that C is a minimizer of

$$\min \left\{ \frac{w(F)}{w(\mathrm{Drop}_U(F))} : F \in \mathfrak{L} \right\} , \tag{1}$$

i.e., it has the best ratio between the weight $w(C)$ of the links to be added versus the weight of the links in $\mathrm{Drop}_U(F)$, which can safely be removed. By convention, we interpret $w(F)/w(\mathrm{Drop}_U(F))$ as being ∞ whenever $w(\mathrm{Drop}_U(F)) = 0$, i.e., $\mathrm{Drop}_U(F) = \emptyset$.

The main challenge in such a relative greedy approach lies in finding a strong family $\mathfrak{L} \subseteq 2^L$ which simultaneously fulfills the following desired properties for any set $U \subseteq L_{\mathrm{up}}$ of non-overlapping up-links:

(a) The minimization problem (1) can be solved efficiently.
(b) If $w(U)$ is significantly heavier than $w(\mathrm{OPT})$, then there is a set $C \in \mathfrak{L}$ for which $w(C)/w(\mathrm{Drop}_U(C))$ is significantly below 1.

Analogous to prior work on relative greedy algorithms, we call the sets in the family $\mathfrak{L}$ also *components*. So far, relative greedy algorithms relied on constant-size components. In particular, for Steiner Tree, Zelikovsky's [Zel96] components were edge sets connecting constantly many terminals, and for WTAP with bounded diameter, Cohen and Nutov [CN13] considered sets of constantly many links.

Constant-size components have the obvious benefit that they allow for obtaining property (a) in a straightforward way. Moreover, their simple definition often makes it much easier to understand what one can achieve with such components, i.e., what improvements are possible by adding such a component to a solution. However, as we highlight in Figure 1, constant-size components do not allow in general for improving an up-link solution to a better-than-2 approximation. Hence, the bounded diameter restriction in [CN13] is crucial for constant-size components to work.

The components $\mathfrak{L} \subseteq 2^L$ we use are link sets with at most constant overlap on vertices. We call such link sets $O(1)$-*thin*, as formalized below.

Definition 3 (k-thin link set). *Let $(G = (V, E), L, w)$ be a WTAP instance and let $k \in \mathbb{Z}_{\geq 0}$. A link set $C \subseteq L$ is k-thin if, for each $v \in V$, we have $|\{\ell \in C : v \in V_\ell\}| \leq k$.*

---

[2]This lemma follows from the observation that an up-link only solution that is at most a factor of 2 heavier than OPT exists because one can start with OPT and replace each OPT-link ℓ by two shadows ℓ_1, ℓ_2 of ℓ that are up-links and fulfill $P_\ell = P_{\ell_1} \cup P_{\ell_2}$. Moreover, any up-link only solution can efficiently be transformed into one with only non-overlapping up-links that is no heavier by shortening up-links if necessary, i.e., replacing them by strict shadows. Finally, it remains to note that a cheapest non-overlapping up-link only solution can be found efficiently. This can be done through a dynamic program, or, alternatively, one can compute an optimal vertex solution to the canonical linear program, which is naturally integral because its constraint matrix is totally unimodular.
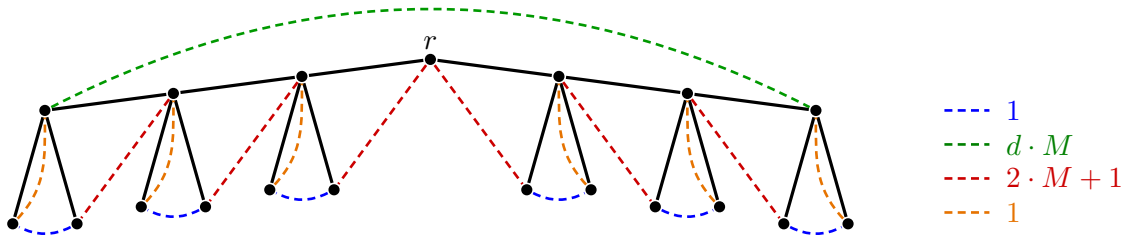
Figure 1: An instance of WTAP showing that components of constant size are not sufficient to achieve an approximation ratio below 2. The tree G is shown in black and dashed lines represent links. The link set L consists of the drawn links together with all their shadows. The weights of the links are 1 for the blue and orange links, $d \cdot M$ for the green link, and $2 \cdot M + 1$ for the red links. Here, d is the number of red/blue/orange links, i.e., $d = 6$ in this example, and M is a large constant.

Then the union of the green link and the blue links is the unique optimal solution OPT. The red and orange links form together a solution U with $w(U) = 2 \cdot w(\text{OPT})$. All red and orange links are up-links and the edge sets P_u with $u \in U$ are pairwise disjoint. (U is even a cheapest up-link only solution, i.e., our 2-approximation algorithm might indeed output this solution.) However, if we consider any set $C \subseteq L$ of at most $d/2$ links, then $w(\text{Drop}_U(C))$ is at most $w(C)$. This shows that, in general, we cannot improve the 2-approximation U by replacing a subset of U by at most k other links for some constant k.

For our algorithm, we fix the constant k in the above definition depending on the error $\varepsilon > 0$ of our $(1 + \ln 2 + \varepsilon)$-approximation, by defining our components $\mathfrak{L} \subseteq 2^L$ to be all $\lceil 2/\varepsilon \rceil$-thin link sets. Despite their super-constant size, this definition of components $\mathfrak{L}$ allows for efficiently solving (1) through a dynamic program, as formalized in the lemma below. Hence, our components $\mathfrak{L}$ fulfill property (a).

Lemma 4. *Let $k \in \mathbb{Z}_{\geq 0}$ be a constant, $(G = (V, E), L, w)$ be a WTAP instance, and $U \subseteq L_{\text{up}}$ such that the edge sets P_u for $u \in U$ are pairwise disjoint. Then we can compute in polynomial time a minimizer of*

$$\min \left\{ \frac{w(C)}{w(\text{Drop}_U(C))} : C \subseteq L \text{ is } k\text{-thin} \right\} \ .$$

We prove Lemma 4 in Section 5. Our relative greedy algorithm for WTAP is described in Algorithm 1, which, due to Lemma 4, is a polynomial-time procedure. In the algorithm, and discussion later on, we assume that $\varepsilon > 0$ is a fixed constant.

Algorithm 1: Relative greedy algorithm for WTAP

Input: A shadow-complete WTAP instance $(G = (V, E), L, w)$.
Output: A WTAP solution $F \subseteq L$ with $w(F) \leq (1 + \ln(2) + \varepsilon) \cdot w(\text{OPT})$.

1. Compute a WTAP solution $U \subseteq L_{\text{up}}$ with $w(U) \leq 2 \cdot w(\text{OPT})$ and disjoint P_u for $u \in U$.
2. Initialize $F := \emptyset$.
3. While $U \neq \emptyset$:
   - Compute a minimizer $C \in \text{argmin} \left\{ \frac{w(C)}{w(\text{Drop}_U(C))} : C \subseteq L \text{ is } \lceil 2/\varepsilon \rceil\text{-thin} \right\}$.
   - Add C to F and replace U by $U \setminus \text{Drop}_U(C)$.
4. Return F.

To make sure that our relative greedy algorithm is able to achieve approximation factors strictly below 2, it remains to show a formal version of property (b), i.e., that there is a profitable replacement step whenever $w(U)$ is significantly larger than $w(\text{OPT})$.

In order to prove that there is an improving replacement when we start, a natural reasoning, which has been used similarly in prior work, is as follows. Consider any up-link WTAP solution $U \subseteq L_{\text{up}}$ with disjoint sets P_u for $u \in U$. Ideally, we would like to find a partition of OPT into $\lceil 2/\varepsilon \rceil$-thin components $C_1, \ldots, C_q$ such that, for each $u \in U$, there is one component C_j with $u \in \text{Drop}_U(C_j)$. Notice that such a decomposition of OPT would immediately imply $\sum_{j=1}^q w(\text{Drop}_U(C_j)) \geq w(U)$. Hence, by an averaging argument we have that if $w(U) > w(\text{OPT})$, then there exists a $\lceil 2/\varepsilon \rceil$-thin component C_j with $w(C_j) < w(\text{Drop}_U(C_j))$. As shown by Cohen and Nutov [CN13], this strategy works out in the bounded diameter case when dealing with constant-size components. However, such a decomposition does not exist in general WTAP instances, even when using the significantly more general class of $O(1)$-thin components.[3] Nevertheless, as stated below, we can show that a slightly weaker statement holds, namely that such a decomposition exists if we first remove a well-chosen subset of up-links $R \subseteq U$ of small total weight. We later invoke the theorem with U being an up-link WTAP solution as guaranteed by Lemma 2 and the WTAP solution F being OPT.

Theorem 5 (decomposition theorem). *Let $(G = (V, E), L, w)$ be a WTAP instance, $F \subseteq L$ be a WTAP solution, and let $U \subseteq L_{\text{up}}$ be a set of up-links such that the sets P_u with $u \in U$ are pairwise disjoint. Then, for any $\varepsilon > 0$, there exists a partition $\mathcal{C}$ of F into $\lceil 1/\varepsilon \rceil$-thin sets and a set $R \subseteq U$ such that*
  (i) for every $u \in U \setminus R$, there exists some $C \in \mathcal{C}$ such that $P_u \subseteq \bigcup_{\ell \in C} P_\ell$, and
  (ii) $w(R) \leq \varepsilon \cdot w(U)$.

We prove the decomposition theorem in Section 4. Using Theorem 5, we readily obtain that Algorithm 1 has the desired approximation guarantee by leveraging known arguments (see, e.g., [Zel96; Grö+01; CN13]). For completeness, we provide a self-contained proof below.

Theorem 6. *For every $\varepsilon > 0$, Algorithm 1 is a $(1 + \ln 2 + \varepsilon)$-approximation algorithm for WTAP.*

Proof. Throughout the algorithm we maintain the invariant that $U \cup F$ is a WTAP solution. Hence, the returned link set F is indeed a WTAP solution and it remains to bound its weight.

Let U_0 be the link set U computed in step 1 of Algorithm 1 and let U_i denote the set U at the end of the i-th iteration of the while loop in step 3. Let C_i denote the component C chosen in the i-th iteration. We apply Theorem 5 to an optimal WTAP solution OPT and the link set U_0 to obtain a set $R \subseteq U_0$ with $w(R) \leq \frac{1}{2}\varepsilon \cdot w(U_0) \leq \varepsilon \cdot w(\text{OPT})$ and a partition $\mathcal{C}$ of OPT into $\lceil 2/\varepsilon \rceil$-thin sets.

Consider the i-th iteration of the while loop. Because there is a component $C \in \mathcal{C}$ with $P_u \subseteq \bigcup_{\ell \in C} P_\ell$ for every $u \in U_{i-1} \setminus R$, we have

$$\sum_{C \in \mathcal{C}} w\big(\text{Drop}_{U_{i-1}}(C)\big) \geq w(U_{i-1} \setminus R) \geq w(U_{i-1}) - w(R) \ .$$

This implies

$$\min_{C \in \mathcal{C}} \frac{w(C)}{w(\text{Drop}_{U_{i-1}}(C))} \leq \frac{\sum_{C \in \mathcal{C}} w(C)}{\sum_{C \in \mathcal{C}} w(\text{Drop}_{U_{i-1}}(C))} \leq \frac{w(\text{OPT})}{w(U_{i-1}) - w(R)} \ .$$

Because every $C \in \mathcal{C}$ is $\lceil 2/\varepsilon \rceil$-thin, every component $C \in \mathcal{C}$ could have been chosen by the algorithm in step 3. Thus,

$$\frac{w(C_i)}{w(U_{i-1} \setminus U_i)} = \frac{w(C_i)}{w(\text{Drop}_{U_{i-1}}(C_i))} \leq \frac{w(\text{OPT})}{w(U_{i-1}) - w(R)} \ . \tag{2}$$

---

[3]Figure 2 shows a bad example where such a decomposition does not exist. Indeed, to make sure that every link in U is covered by at least one component, all links of the WTAP solution F must be in the same component. As $v \in V_\ell$ for all $m + 1$ links $\ell \in F$, this leads to a component that is not m-thin, where m can be chosen arbitrarily large in the example.

Moreover, the algorithm could also have chosen any component consisting of a single link $u \in U_{i-1}$, which implies $w(C_i)/w(U_{i-1} \setminus U_i) \leq 1$. Combining this with (2), we obtain

$$w(C_i) \leq \min\left\{\frac{w(\mathrm{OPT})}{w(U_{i-1}) - w(R)}, 1\right\} \cdot w(U_{i-1} \setminus U_i) \leq \int_{w(U_i)}^{w(U_{i-1})} \min\left\{\frac{w(\mathrm{OPT})}{x - w(R)}, 1\right\} dx \ ,$$

where we used $w(U_i) \leq w(U_{i-1})$ and that $\min\{\frac{w(\mathrm{OPT})}{x - w(R)}, 1\}$ is monotonically decreasing in x.

Let m denote the number of iterations of the while loop. Then $U_m = \emptyset$ and

$$w(F) = \sum_{i=1}^{m} w(C_i) \leq \sum_{i=1}^{m} \int_{w(U_i)}^{w(U_{i-1})} \min\left\{\frac{w(\mathrm{OPT})}{x - w(R)}, 1\right\} dx$$

$$= \int_{w(U_m)}^{w(U_0)} \min\left\{\frac{w(\mathrm{OPT})}{x - w(R)}, 1\right\} dx$$

$$= \int_{0}^{w(\mathrm{OPT})+w(R)} 1 \, dx + \int_{w(\mathrm{OPT})+w(R)}^{w(U_0)} \frac{w(\mathrm{OPT})}{x - w(R)} dx$$

$$= w(\mathrm{OPT}) + w(R) + \ln\left(\frac{w(U_0) - w(R)}{w(\mathrm{OPT})}\right) \cdot w(\mathrm{OPT})$$

$$\leq (1 + \varepsilon) \cdot w(\mathrm{OPT}) + \ln 2 \cdot w(\mathrm{OPT}) \ ,$$

where the last inequality follows from $w(R) \leq \varepsilon \cdot w(\mathrm{OPT})$ and $w(U_0) \leq 2 \cdot w(\mathrm{OPT})$. $\square$

As is common with relative greedy procedures, one does not need to run the while loop of Algorithm 1 until $U \neq \emptyset$, but can stop early. More precisely, as soon as a component $C \in \mathrm{argmin}\{w(C)/w(\mathrm{Drop}_U(C)) : C \subseteq L$ is $\lceil \varepsilon/2 \rceil$-thin$\}$ is computed that does not improve the solution anymore, i.e., $w(C)/w(\mathrm{Drop}_U(C)) \geq 1$, then one can return $F \cup U$ without continuing the while loop. (We recall that $w(C)/w(\mathrm{Drop}_U(C)) \geq 1$ actually implies $w(C)/w(\mathrm{Drop}_U(C)) = 1$ because C can always be chosen to be a single up-link in U.) Indeed, once such a set C is encountered, any future replacements of up-links by a component done in the while loop will also not improve the solution.

4 Proving the decomposition theorem

In this section we prove our decomposition theorem, Theorem 5. Hence, we are given a WTAP solution $F \subseteq L$ and a set $U \subseteq L_{\mathrm{up}}$ of up-links such that the edge sets P_u for $u \in U$ are pairwise disjoint. For every link $u \in U$, we will first fix a set $F_u \subseteq F$ satisfying $P_u \subseteq \bigcup_{\ell \in F_u} P_\ell$. We will then choose a set $R \subseteq U$ with $w(R) \leq \varepsilon \cdot w(U)$ and a partition $\mathcal{C}$ of F into $\lceil 1/\varepsilon \rceil$-thin sets such that, for every $u \in U \setminus R$, there is a component $C \in \mathcal{C}$ with $F_u \subseteq C$. We emphasize that, to be able to achieve a decomposition as claimed by Theorem 5, it is crucial to choose the set F_u for $u \in U$ carefully. In particular, a natural choice, used in a similar setting by Cohen and Nutov [CN13], would be to let F_u be any minimal (or any minimum cardinality) set with $P_u \subseteq \bigcup_{\ell \in F_u} P_\ell$. However, the example highlighted in Figure 2 shows that this renders it impossible to achieve the decomposition property with the strategy outlined above. Before expanding on how we choose the sets F_u, we continue the overview of our proof strategy for the decomposition theorem.

Once we fixed the sets F_u for all $u \in U$ and the set $R \subseteq U$, the definition of the partition $\mathcal{C}$ is straightforward. Two links $\ell_1, \ell_2 \in F$ are in the same component $C \in \mathcal{C}$ if and only if there is an up-link $u \in U \setminus R$ with $\ell_1, \ell_2 \in F_u$. We express this dependency through a directed graph with vertex set F. For every set F_u with $u \in U \setminus R$, this graph contains a path (F_u, A_u), which we formally define later. Hence, the connected
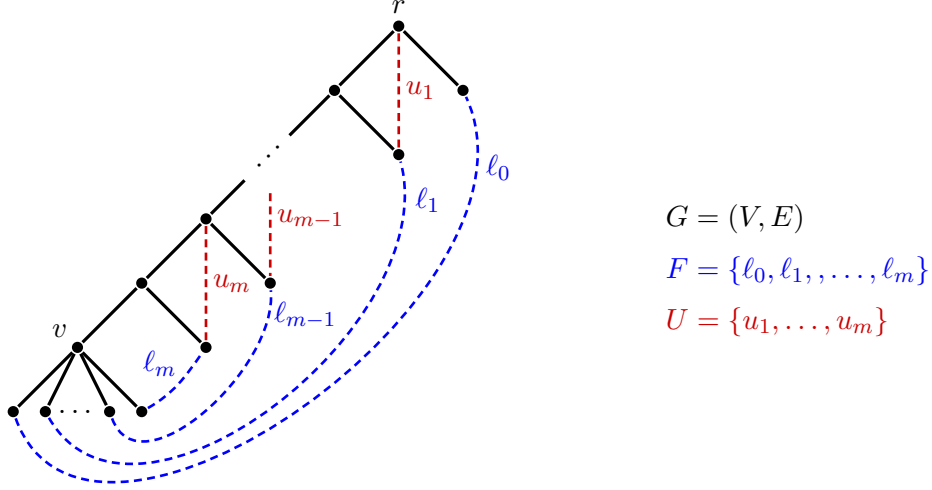
Figure 2: The figure shows an example with a tree $G = (V, E)$ (black), a WTAP solution F (blue), and a set U of up-links (red) such that the edge sets P_u with $u \in U$ are pairwise disjoint. In this example, $v \in V_\ell$ for every link $\ell \in F$. Therefore, any k-thin subset of F contains at most k links. For every link $u_i \in U$, the set $F_u = \{\ell_0, \ell_i\} \subseteq F$ is a minimal subset of F with $P_{u_i} \subseteq \bigcup_{\ell \in F_u} P_\ell$. However, with this choice of F_u, every set F_u with $u \in U$ contains the link ℓ_0. Thus, if we consider any partition $\mathcal{C}$ of F into k-thin components, the only component $C \in \mathcal{C}$ for which we can have $F_u \subseteq C$ for some $u \in U$, is the component C_0 containing ℓ_0. Because this component is k-thin, it contains at most k links and hence the number of links $u \in U$ with $F_u \subseteq C_0$ is at most $k - 1$. For $w(u) = 1$ for all $u \in U$, this shows that the total weight of the up-links $u \in U$ for which F_u is not contained in any component $C \in \mathcal{C}$ is at least $(1 - \frac{k-1}{m}) \cdot w(U)$ instead of at most $\varepsilon \cdot w(U)$ as required.

components of this dependency graph with vertex set F yield the partition $\mathcal{C}$. More precisely, $C \subseteq F$ is a part of the partition $\mathcal{C}$ if and only if the dependency graph has a connected component with vertex set C. We remark that this construction of a dependency graph has been used before by Cohen and Nutov [CN13] in the bounded diameter case with a different choice for the sets F_u.

To prove the decomposition theorem, we use such a dependency graph not only to find the partition $\mathcal{C}$, but also to choose the set $R \subseteq U$. We thus consider a dependency graph that has vertex set F and contains arcs A_u for every $u \in U$, where A_u is again the above-mentioned arc set forming a path with vertex set F_u. Then we show that there exists a set $R \subseteq U$ with $w(R) \le \varepsilon \cdot w(U)$ such that, after removing the arcs in $\bigcup_{u \in R} A_u$ from the dependency graph, every connected component (C, A) of the dependency graph fulfills that C is $\lceil 1/\varepsilon \rceil$-thin. To prove the existence of such a set R, we exploit that (with our choice of (F_u, A_u) for $u \in U$) the dependency graph has the following two properties.

(1) The dependency graph for U is a branching.
(2) Let (C, A) be a connected component of the dependency graph. If the arc set of every directed path in (C, A) has nonempty intersection with A_u for at most k up-links $u \in U$, then C is $(k + 1)$-thin.

To obtain the first of these properties, we use that the sets P_u for $u \in U$ are pairwise disjoint. Property (1) was already shown to hold for the dependency graph used by Cohen and Nutov [CN13] and holds as long as we choose F_u to be a minimal set with $P_u \subseteq \bigcup_{\ell \in F_u} P_\ell$. To prove the second property, we crucially need our particular choice of F_u as shown in Figure 3.

Once properties (1) and (2) are shown, there is a simple choice of $R \subseteq U$ to obtain the decomposition theorem, as explained in Figure 4.
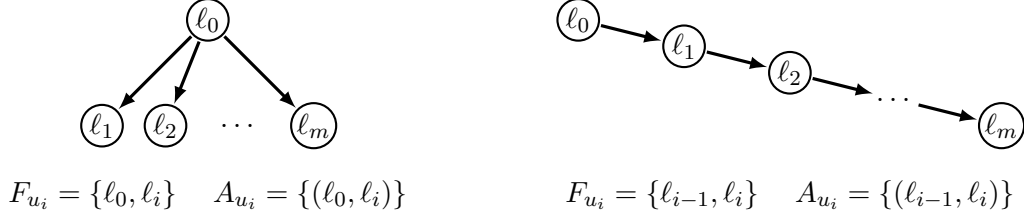
$$F_{u_i} = \{\ell_0, \ell_i\} \quad A_{u_i} = \{(\ell_0, \ell_i)\} \qquad F_{u_i} = \{\ell_{i-1}, \ell_i\} \quad A_{u_i} = \{(\ell_{i-1}, \ell_i)\}$$

Figure 3: The figure shows the dependency graph resulting from different choices of F_u for the example instance from Figure 2. The left picture shows the dependency graph that we would get if we chose $F_{u_i} = \{\ell_0, \ell_i\}$ for $i \in \{1, \dots, m\}$. We recall that we already argued in the caption of Figure 2 that this choice makes it impossible to obtain the decomposition theorem as suggested. Also, one can see that property (2) is clearly not fulfilled for the left-hand side choice. Indeed, every directed path intersects with at most one set A_u, which, if property (2) were fulfilled, should imply 2-thinness of the component; however, the thinness is $m+1$. The right picture shows the dependency graph for the choice of F_{u_i} and A_{u_i} that we will make to prove the decomposition theorem. In contrast to the left picture, the choice in the right picture fulfills property (2): The set $F = \{\ell_0, \ell_1, \dots, \ell_m\}$ is $(m+1)$-thin and the dependency graph is a path that has nonempty intersection with all m sets A_{u_i}.
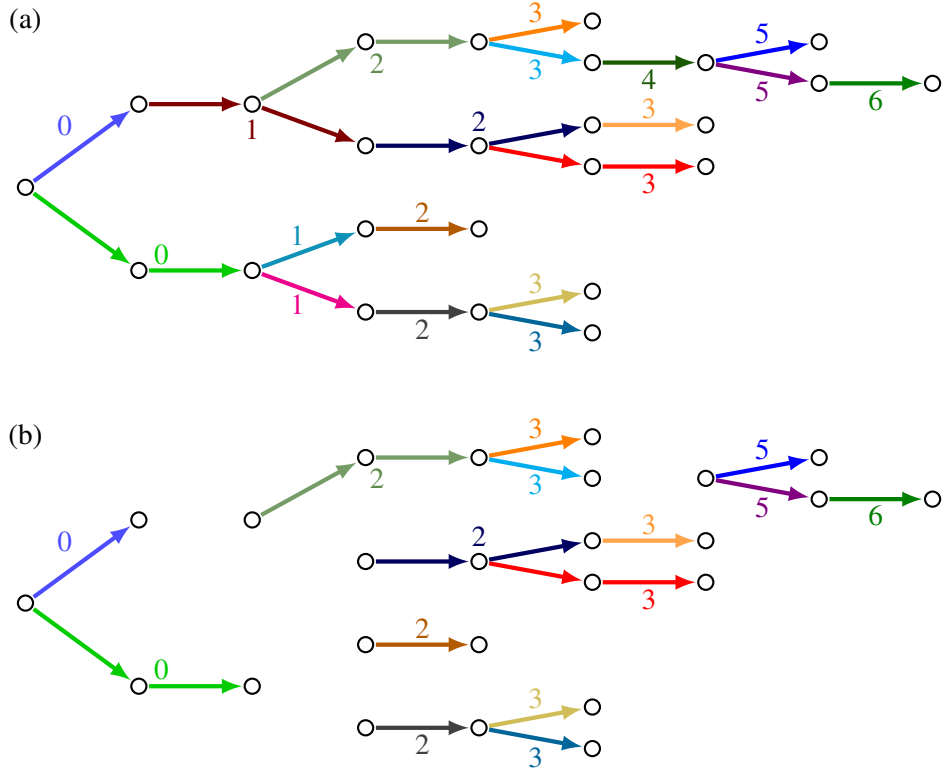

Figure 4: Picture (a) shows an example of a connected component of the dependency graph for U. The arc set of every path (F_u, A_u) is shown in a different color. In order to construct the set $R \subseteq U$, we assign a label from $\mathbb{Z}_{\geq 0}$ to each of the sets A_u that is equal to the number of different colors (or equivalently sets $A_{\overline{u}}$ for $\overline{u} \in U$) encountered on the unique path from the root of the component to the start of the path A_u. The numbers in the figure show such a labeling. Let $k = \lceil 1/\varepsilon \rceil$. The choice of the labeling guarantees that, for every $i \in \{0, \dots, k-1\}$, removing all arcs with any label $j \in \mathbb{Z}_{\geq 0}$ that satisfies $j \equiv i \pmod{k}$, leads to a dependency graph where at most $k-1$ different colors appear on each path. In other words, after removing the arcs with label $i \bmod k$, every path has nonempty intersection with at most $k-1$ different sets A_u. Then, by property (2), the connected components correspond to k-thin sets. Picture (b) shows an example for $i = 1$ and $k = 3$. We choose R to be the set of all up-links $u \in U$ for which A_u has label $i \bmod k$, where $i \in \{0, \dots, k-1\}$ is an index for which the resulting set R has minimum weight $w(R)$.

4.1 The dependency graph

Next, we formally define the dependency graph of a set $U \subseteq L_{\mathrm{up}}$. Let $(G = (V, E), L, w)$ be a WTAP instance and let $F \subseteq L$ be a WTAP solution. Let $r \in V$ be the (arbitrarily chosen) root of G. The root defines a natural ancestry relationship. The *ancestors* of a vertex $v \in V$ are all vertices $z \in V$ that lie on the unique r-v path, which includes r and v (we talk about a *strict ancestor* to disallow $z = v$). Analogously, $v \in V$ is a (strict) descendant of $z \in V$ if z is a (strict) ancestor of v. For a link $\ell \in L$, we denote by $\mathrm{apex}(\ell) \in V$ the lowest common ancestor of the two endpoints of ℓ, i.e., the vertex in V_ℓ closest to the root.

For each up-link $u \in L_{\mathrm{up}}$, we choose a minimal link set $F_u \subseteq F$ with $P_u \subseteq \bigcup_{\ell \in F_u} P_\ell$ as follows. Let $u = \{t, b\}$ where t is an ancestor of b. We define v_u to be the lowest ancestor of t, i.e., the ancestor farthest away from the root r, such that P_u is covered by links in

$$B_{v_u} := \{\ell \in F : \mathrm{apex}(\ell) \text{ is a descendant of } v_u\} \ ,$$

i.e., $P_u \subseteq \bigcup_{\ell \in B_{v_u}} P_\ell$. Then we choose $F_u \subseteq B_{v_u}$ minimal such that $P_u \subseteq \bigcup_{\ell \in F_u} P_\ell$. See Figure 5 for an example of a minimal set F_u.

Due to minimality of F_u, the links $\ell \in F_u$ have a natural order. More precisely, this order is induced by how close to the root the edges of $P_{u,\ell} := P_u \setminus \bigcup_{\bar{\ell} \in F_u \setminus \{\ell\}} P_{\bar{\ell}}$ are, which are the edges in P_u for which ℓ is the only link in F_u that covers them. (See Figure 5.) To define this order formally, we start by observing that the sets $P_{u,\ell}$ are edge sets of paths.
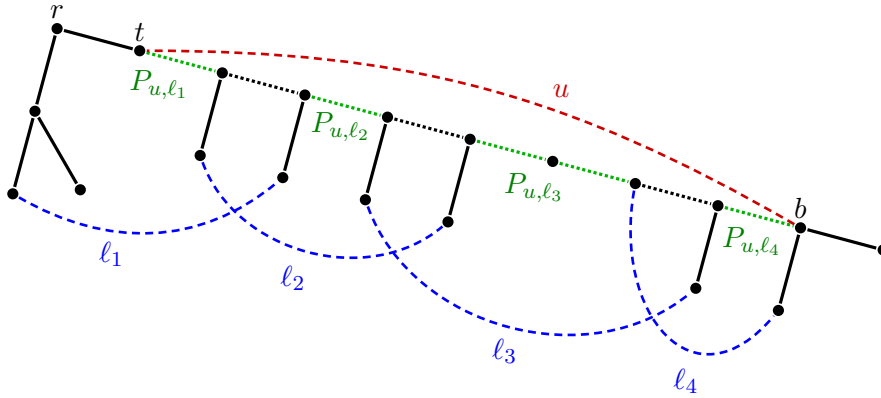

Figure 5: The picture shows the tree G (black and green), an up-link u, and a minimal set F_u of links with $P_u \subseteq \bigcup_{\ell \in F_u} P_\ell$ (blue). The edges in P_u are dotted and the edges in P_{u,ℓ_i} for $i \in \{1, 2, 3, 4\}$ are shown in green.

Lemma 7. *Let $u \in L_{\mathrm{up}}$ and $\ell \in F_u$. Then $P_{u,\ell}$ is nonempty and the edge set of a path.*

Proof. The minimality of F_u immediately implies $P_{u,\ell} \neq \emptyset$. Let $e_1, e_2, e_3 \in P_u \cap P_\ell$ be three distinct edges that appear in this order on the path (V_u, P_u). If $e_1, e_3 \in P_{u,\ell}$, then either e_2 is also contained in $P_{u,\ell}$ or there is a link $\bar{\ell} \in F_u \setminus \{\ell\}$ with $e_2 \in P_{\bar{\ell}}$. In the latter case, e_2 is contained in $P_u \cap P_{\bar{\ell}}$ which is the edge set of a subpath of (V_u, P_u). Because e_1 and e_3 are not contained in $P_{\bar{\ell}}$, we have $P_u \cap P_{\bar{\ell}} \subsetneq P_u \cap P_\ell$, contradicting the fact that $P_{u,\bar{\ell}}$ is nonempty. Hence, whenever e_1 and e_3 are contained in $P_{u,\ell}$, then all edges that appear between e_1 and e_3 on the path (V_u, P_u) are also contained in $P_{u,\ell}$. This shows that $P_{u,\ell} \subseteq P_u$ is the edge set of a path. $\square$

The edge sets $P_{u,\ell}$ with $\ell \in F_u$ are pairwise disjoint by their definition. For $\ell_1, \ell_2 \in F_u$, we define $\ell_1 \prec_u \ell_2$ if and only if the edges in P_{u,ℓ_1} appear before the edges of P_{u,ℓ_2} on the t-b path in G. By Lemma 7, this order is well-defined. An alternative characterization of the same link order is that, for $\ell_1, \ell_2 \in F_u$, we

have $\ell_1 \prec_u \ell_2$ if and only if $\mathrm{apex}(\ell_1)$ is a strict ancestor of $\mathrm{apex}(\ell_2)$. (This characterization follows from Lemma 8 (i) below.)

The dependency graph for a set $U \subseteq L_{\mathrm{up}}$ of up-links is a directed graph with vertex set F. For every up-link $u \in U$, it contains a set A_u of arcs defined as follows. For $u \in U$, let $\ell_1 \prec_u \ell_2 \prec_u \cdots \prec_u \ell_q$ be the links in F_u. Then

$$A_u := \{(\ell_i, \ell_{i+1}) \colon i \in \{1, \ldots, q-1\}\} \ .$$

The arc set of the dependency graph for U is the disjoint union of the sets A_u for all $u \in U$.

This construction immediately implies that, for every up-link $u \in U$, there is one connected component (C, A) of the dependency graph such that $F_u \subseteq C$, which implies that P_u is covered by the links in C. To prove the decomposition theorem, we will show that there exists a set $R \subseteq U$ with $w(R) \le \varepsilon \cdot w(U)$ such that, for every connected component (C, A) of the dependency graph of $U \setminus R$, the link set C is $\lceil 1/\varepsilon \rceil$-thin.

First, we show some basic properties of the dependency graph that do not rely on our particular choice of the sets F_u, except for them being a minimal link set covering P_u. We provide self-contained proofs here but remark that some of these properties, in particular property (1), have been shown already in [CN13]. Subsequently, in Section 4.2, we exploit our particular choice of the sets F_u and show how they allow for obtaining property (2).

Lemma 8. *Let $u \in L_{\mathrm{up}}$ and $(\ell_1, \ell_2) \in A_u$. Then*
  (i) $\mathrm{apex}(\ell_1)$ is a strict ancestor of $\mathrm{apex}(\ell_2)$ in the tree G, and
  (ii) P_{u,ℓ_1} is the edge set of a subpath of the $\mathrm{apex}(\ell_1)$-$\mathrm{apex}(\ell_2)$ path in G.

Proof. Let $u = \{t, b\} \in U$ be the up-link with $(\ell_1, \ell_2) \in A_u$, where t is an ancestor of b in the tree G. Let $\{a_1, b_1\} \in P_{u,\ell_1}$ and $\{a_2, b_2\} \in P_{u,\ell_2}$. Without loss of generality we may assume that a_1 is an ancestor of b_1 and a_2 is an ancestor of b_2. (Because $a_1 \ne b_1$ and $a_2 \ne b_2$, they are actually strict ancestors.) By the definition of the order $\prec_u$ and the arc set A_u, the edge $\{a_1, b_1\}$ appears before the edge $\{a_2, b_2\}$ on the t-b path in G. Because t is an ancestor of b, we conclude that a_1 is an ancestor of a_2. Moreover, because the edge $\{a_1, b_1\}$ is covered by the link ℓ_1, the vertex $\mathrm{apex}(\ell_1)$ is an ancestor of a_1. The link ℓ_2 covers $\{a_2, b_2\}$ but it does not cover $\{a_1, b_1\}$ due to the definition of P_{u,ℓ_1}. Therefore, $\mathrm{apex}(\ell_2)$ must be a descendant of b_1. We have shown that the vertex $\mathrm{apex}(\ell_1)$ is an ancestor of a_1, which in turn is a strict ancestor of b_1, which is an ancestor of $\mathrm{apex}(\ell_2)$. Hence, $\mathrm{apex}(\ell_1)$ is a strict ancestor of $\mathrm{apex}(\ell_2)$. Moreover, this shows that every edge $\{a_1, b_1\} \in P_{u,\ell_1}$ lies on the $\mathrm{apex}(\ell_1)$-$\mathrm{apex}(\ell_2)$ path in G. By Lemma 7, this implies (ii). $\square$

Lemma 9. *Let $u \in L_{\mathrm{up}}$ be an up-link and let $(\ell_1, \ell_2) \in A_u$. Let $e \in E$ be the last edge of the r-$\mathrm{apex}(\ell_2)$ path in G. Then u covers e.*

Proof. By Lemma 8, $\mathrm{apex}(\ell_1)$ is an ancestor of $\mathrm{apex}(\ell_2)$ and the edges in the nonempty set $P_{u,\ell_1} \subseteq P_u$ lie on the $\mathrm{apex}(\ell_1)$-$\mathrm{apex}(\ell_2)$ path in G. Thus, V_u contains at least one strict ancestor of $\mathrm{apex}(\ell_2)$. Because $P_u \cap P_{\ell_2}$ is nonempty, V_u contains at least one descendant of $\mathrm{apex}(\ell_2)$. Using that (V_u, P_u) is a path in the tree G, we conclude $e \in P_u$. $\square$

Lemma 10. *Let $U \subseteq L_{\mathrm{up}}$ be a set of up-links such that the sets P_u for $u \in U$ are pairwise disjoint. Then the dependency graph of U is a branching.*

Proof. By Lemma 8 (i), the dependency graph does not contain any directed cycle. Hence, it remains to show that every link has at most one incoming arc. Let $\ell \in F$ and let e be the last edge of the r-$\mathrm{apex}(\ell)$ path in G. Because the edge sets P_u with $u \in U$ are pairwise disjoint, there is at most one up-link $u \in U$ that covers e. By Lemma 9, every arc entering ℓ in the dependency graph is contained in A_u. Because A_u is the arc set of a directed path, ℓ has at most one incoming arc. $\square$

11

To prove property (2), we rely on the following lemma, which is a special case of it. We later use this special case to obtain the general statement.

Lemma 11. F_u *is 2-thin for any* $u \in L_{\mathrm{up}}$.

Proof. Let $v \in V$. With the goal of deriving a contradiction, suppose there exist distinct links $\ell_1, \ell_2, \ell_3 \in F_u$ with $v \in V_{\ell_1} \cap V_{\ell_2} \cap V_{\ell_3}$. We may assume $\ell_1 \prec_u \ell_2 \prec_u \ell_3$ without loss of generality. Let $z \in V_u$ be the vertex of V_u that is closest to v in the tree G. (In particular, if $v \in V_u$, then $z = v$.) Note that any path containing both v and a vertex of V_u must go through z. Hence, $z \in V_{\ell_i}$ for $i \in \{1, 2, 3\}$. Because $P_{\ell_i} \cap P_u$ is the edge set of a subpath of (V_u, P_u) with vertex set $V_{\ell_i} \cap V_u$ for all $i \in \{1, 2, 3\}$, and all of these paths contain the vertex z, also $(P_{\ell_1} \cap P_u) \cup (P_{\ell_3} \cap P_u)$ is the edge set of a subpath of (V_u, P_u). This subpath covers the edges in P_{u,ℓ_1} as well as the edges in P_{u,ℓ_3}. Because $\ell_1 \prec_u \ell_2 \prec_u \ell_3$, this implies that it also covers the edges in P_{u,ℓ_2}, contradicting the definition of the nonempty set P_{u,ℓ_2}. □

4.2 Thin components and the dependency graph

Most properties of the dependency graph shown so far were primarily properties of a single set A_u (or immediate consequences of these). Therefore, only minimality of the sets F_u was necessary to show them. We now move toward more global results on the connected components of the dependency graph by exploiting our particular choice of the sets F_u, with the goal to show property (2). To this end, let $U \subseteq L_{\mathrm{up}}$ be a set of up-links such that the sets P_u with $u \in U$ are pairwise disjoint. We fix a connected component (C, A) of the dependency graph of U. By Lemma 10, the connected component (C, A) is an arborescence. Even though we do not exploit this later, we note that one can show that any distinct links $\ell_1, \ell_2 \in C$ have distinct apexes, i.e., $\mathrm{apex}(\ell_1) \neq \mathrm{apex}(\ell_2)$.[4]

The next lemma is a crucial step toward property (2), as it shows that links in the same component whose paths (V_ℓ, P_ℓ) have a common vertex, must have an ancestry relationship in the dependency graph. Note that this lemma together with Lemma 11 already imply a weaker version of property (2), namely that if the arc set of every directed path in (C, A) has nonempty intersection with A_u for at most k up-links $u \in U$, then C is $2k$-thin.[5] Indeed, Lemma 12 below implies that, for any vertex $v \in V$, the links $\ell \in C$ with $v \in V_\ell$ must lie on a directed path in (C, A). Finally, Lemma 11 shows that for each of the at most k links $u \in U$ for which A_u intersects that path, there has at most 2 links $\ell \in F_u$ satisfying $v \in P_\ell$. After proving Lemma 12, we strengthen this reasoning to obtain property (2), which is tight.

Lemma 12. *Let* $\ell_1, \ell_2 \in C$ *with* $V_{\ell_1} \cap V_{\ell_2} \neq \emptyset$. *Then* ℓ_1 *and* ℓ_2 *have an ancestry relationship in the arborescence* (C, A), *i.e., either* ℓ_1 *is an ancestor of* ℓ_2 *or* ℓ_2 *is an ancestor of* ℓ_1.

Proof. Let $v \in V_{\ell_1} \cap V_{\ell_2}$. Then $\mathrm{apex}(\ell_1)$ and $\mathrm{apex}(\ell_2)$ are ancestors of v in the tree G. Therefore, $\mathrm{apex}(\ell_1)$ and $\mathrm{apex}(\ell_2)$ have an ancestry relation in G, say $\mathrm{apex}(\ell_1)$ is an ancestor of $\mathrm{apex}(\ell_2)$. Thus, $\mathrm{apex}(\ell_2)$ lies on the $\mathrm{apex}(\ell_1)$-v path in G. Because $v \in V_{\ell_1}$, this implies $\mathrm{apex}(\ell_2) \in V_{\ell_1}$. See the left part of Figure 6.

For the sake of deriving a contradiction, suppose that ℓ_1 and ℓ_2 have no ancestry relation in the arborescence (C, A). Then the path from the root of (C, A) to ℓ_2 does not contain ℓ_1. By Lemma 8 (i), there is an arc $a = (\ell, \bar{\ell})$ on this path such that $\mathrm{apex}(\ell)$ is a strict ancestor of $\mathrm{apex}(\ell_1)$ and $\mathrm{apex}(\bar{\ell})$ is a descendant

---

[4]This will for example follow from Lemma 12. Indeed, with the goal of deriving a contradiction, assume that there are two distinct links $\ell_1, \ell_2 \in C$ with same apex. By Lemma 12, ℓ_1 and ℓ_2 must have an ancestry relationship in (C, A) because their common apex is in $V_{\ell_1} \cap V_{\ell_2}$. This ancestry relationship is strict because $\ell_1 \neq \ell_2$. Moreover, by Lemma 8 (i), any parent-child relationship (and therefore any strict ancestry relationship) between two links in (C, A) implies that the apex of the parent is a strict ancestor of the apex of the child. This contradicts $\mathrm{apex}(\ell_1) = \mathrm{apex}(\ell_2)$.

[5]This weaker version of property (2) is already sufficient to prove our main result. However, to guarantee correctness through this weaker property, it would not suffice to consider $\lceil 2/\varepsilon \rceil$-thin components in Algorithm 1, but more general components are needed instead, for example $\lceil 4/\varepsilon \rceil$-thin ones.
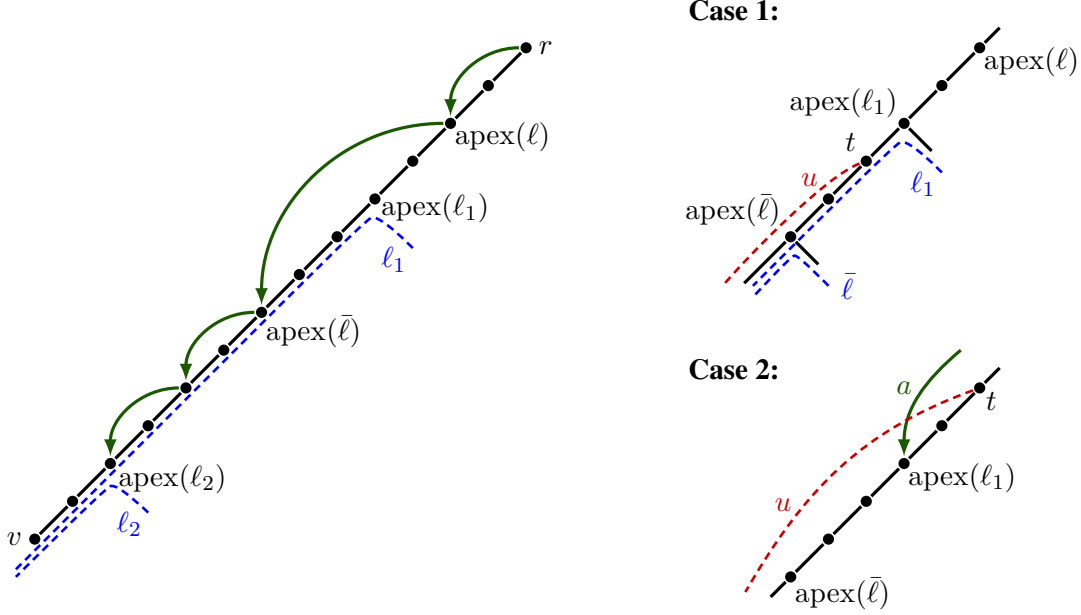
Figure 6: Illustration of the proof of Lemma 12. Here, a dark green arc $(\text{apex}(\ell), \text{apex}(\bar{\ell}))$ in the picture represents the arc $(\ell, \bar{\ell})$ in the dependency graph. Black vertices and edges show parts of the tree G and links are drawn dashed.

of $\text{apex}(\ell_1)$. Then $\text{apex}(\bar{\ell})$ lies on the $\text{apex}(\ell_1)$-$\text{apex}(\ell_2)$ path in G. This implies $\text{apex}(\bar{\ell}) \in V_{\ell_1}$ because $\text{apex}(\ell_2) \in V_{\ell_1}$.

Let $u \in U$ with $(\ell, \bar{\ell}) \in A_u$. Let t denote the endpoint of u that is closer to the root of the tree G. We distinguish two cases. In the first case, we assume that t is a descendant of $\text{apex}(\ell_1)$. (See top illustration on right-hand side of Figure 6.) Then the edges of the t-$\text{apex}(\bar{\ell})$ path are covered by ℓ_1 because $\text{apex}(\bar{\ell}) \in V_{\ell_1}$. Thus, by Lemma 8 (ii), ℓ_1 covers all edges in $P_{u,\ell}$ and hence

$$P_u \subseteq \bigcup_{\ell' \in (F_u \setminus \{\ell\}) \cup \{\ell_1\}} P_{\ell'} \ . \tag{3}$$

Because t is a descendant of $\text{apex}(\ell_1)$, it is also a descendant of $\text{apex}(\ell)$ and therefore ℓ is the first link in F_u with respect to the order $\prec_u$ by Lemma 8. We conclude that for every $\ell' \in (F_u \setminus \{\ell\}) \cup \{\ell_1\}$, the vertex $\text{apex}(\ell')$ is a strict descendant of $\text{apex}(\ell)$. Because of (3), this implies that the vertex v_u in the construction of F_u must be a strict descendant of $\text{apex}(\ell)$, contradicting $\ell \in F_u$.

Now consider the remaining second case, where t is a strict ancestor of $\text{apex}(\ell_1)$. (See bottom illustration on right-hand side of Figure 6.) Because ℓ_1 and ℓ_2 have no ancestry relation, ℓ_1 is not the root of (C, A) and hence ℓ_1 has an incoming arc $a \in A$. By Lemma 9, P_u covers the last edge of the r-$\text{apex}(\bar{\ell})$ path in G. In particular, both the descendant $\text{apex}(\bar{\ell})$ of $\text{apex}(\ell_1)$ and the strict ancestor t of $\text{apex}(\ell_1)$ are contained in V_u. This implies that u is the unique up-link in U that covers the last edge of the r-$\text{apex}(\ell_1)$ path in G. By Lemma 9, we conclude that the incoming arc a of ℓ_1 is contained in A_u. Thus, we have $\ell_1, \ell, \bar{\ell} \in F_u$. Because $\text{apex}(\ell)$ is an ancestor of $\text{apex}(\ell_1)$, which in turn is an ancestor of $\text{apex}(\bar{\ell})$, we have $\ell \prec_u \ell_1 \prec_u \bar{\ell}$ by Lemma 8 (i). This contradicts $(\ell, \bar{\ell}) \in A_u$. $\qquad \square$

Recall that in order to prove property (2) of the dependency graph, we need to give an upper bound on $|\{\ell \in C : v \in V_\ell\}|$ for all $v \in V$. The next lemma establishes this upper bound for every vertex $v \in V$ that is the apex of some link $\ell \in C$.

Lemma 13. *Let $\ell \in C$ and let H_ℓ be the arc set of the path from the root of the arborescence (C, A) to ℓ. Then*

$$\left| \left\{ \bar{\ell} \in C \setminus \{\ell\} \colon \text{apex}(\ell) \in V_{\bar{\ell}} \right\} \right| \leq |\{u \in U \colon H_\ell \cap A_u \neq \emptyset\}| \ .$$

Proof. Let $k := |\{u \in U \colon H_\ell \cap A_u \neq \emptyset\}|$. We prove the lemma by induction on k. If $k = 0$, then the link ℓ is the root of the arborescence (C, A). This implies by Lemma 8 (i) that $\text{apex}(\bar{\ell})$ is a strict descendant of $\text{apex}(\ell)$ for every $\bar{\ell} \in C \setminus \{\ell\}$. Hence, in this case $\{\bar{\ell} \in C \setminus \{\ell\} \colon \text{apex}(\ell) \in V_{\bar{\ell}}\} = \emptyset$, as desired.

Now suppose $k > 0$. By Lemma 12 (and Lemma 8 (i)), every link $\bar{\ell} \in C$ with $\text{apex}(\ell) \in V_{\bar{\ell}}$ is an ancestor of ℓ in (C, A). Because $k > 0$, the link ℓ has an incoming arc $a \in A$. Let $u \in U$ be the unique up-link with $a \in A_u$. Let $\ell_1 \in F_u$ be the first link on the directed path (F_u, A_u) in the arborescence (C, A). Then for every ancestor $\bar{\ell}$ of ℓ in the arborescence (C, A), and hence for every link $\bar{\ell} \in C$ with $\text{apex}(\ell) \in V_{\bar{\ell}}$, we have either

(i) $\bar{\ell}$ is a strict ancestor of ℓ_1 in (C, A), or

(ii) $\bar{\ell} \in F_u$.

Consider a strict ancestor $\bar{\ell}$ of ℓ_1. By Lemma 8 (i), $\text{apex}(\bar{\ell})$ is an ancestor of $\text{apex}(\ell_1)$, which in turn is an ancestor of $\text{apex}(\ell)$ in the tree G. Therefore, if $\text{apex}(\ell) \in V_{\bar{\ell}}$, then also $\text{apex}(\ell_1) \in V_{\bar{\ell}}$. By the inductive hypothesis applied to ℓ_1, this implies that there are at most $k - 1$ links in $\bar{\ell} \in C \setminus \{\ell_1\}$ that fulfill both (i) and $\text{apex}(\ell) \in V_{\bar{\ell}}$. Thus, it suffices to show that there is at most one link $\bar{\ell} \in F_u \setminus \{\ell\}$ with $\text{apex}(\ell) \in V_{\bar{\ell}}$. This holds because by Lemma 11, there are at most two links $\bar{\ell} \in F_u$ with $\text{apex}(\ell) \in V_{\bar{\ell}}$, one of which is ℓ. $\qquad\square$

Finally, we are ready to prove property (2) of the dependency graph.

Lemma 14. *Let $k \in \mathbb{Z}_{\geq 0}$. If for every path in (C, A) with arc set $H \subseteq A$, we have*

$$|\{u \in U \colon H \cap A_u \neq \emptyset\}| \leq k \ ,$$

then C is $(k + 1)$-thin.

Proof. We need to show that, for every $v \in V$, there are at most $k + 1$ links $\ell \in C$ with $v \in V_\ell$. Let $z \in V$ be the last vertex on the r-v path in G that is the apex of some link in C, and let $\ell_z \in C$ be a link with $\text{apex}(\ell_z) = z$. For every link $\ell \in C$ with $v \in V_\ell$, the vertex $\text{apex}(\ell)$ is an ancestor of v. Thus, by the choice of z, we have $z \in V_\ell$ for each such link ℓ. Hence, it suffices to show that there are at most $k + 1$ links $\ell \in C$ with $z \in V_\ell$. By Lemma 13, there are at most k links $\ell \in C \setminus \{\ell_z\}$ with $z \in V_\ell$ and hence at most $k + 1$ links $\ell \in C$ with $z \in V_\ell$. $\qquad\square$

4.3 Proof of the decomposition theorem

We are now ready to complete the proof of the decomposition theorem, which we restate here for convenience. See Figure 4 for an illustration of the proof.

Theorem 5 (decomposition theorem)**.** *Let $(G = (V, E), L, w)$ be a WTAP instance, $F \subseteq L$ be a WTAP solution, and let $U \subseteq L_{\text{up}}$ be a set of up-links such that the sets P_u with $u \in U$ are pairwise disjoint. Then, for any $\varepsilon > 0$, there exists a partition $\mathcal{C}$ of F into $\lceil 1/\varepsilon \rceil$-thin sets and a set $R \subseteq U$ such that*

(i) *for every $u \in U \setminus R$, there exists some $C \in \mathcal{C}$ such that $P_u \subseteq \bigcup_{\ell \in C} P_\ell$, and*

(ii) *$w(R) \leq \varepsilon \cdot w(U)$.*

Proof. Let $k := \lceil 1/\varepsilon \rceil$. We start by defining, independently for every connected component (C, A) of the dependency graph of U, a labeling $c \colon A \to \mathbb{Z}_{\geq 0}$, where arcs contained in the same set A_u with $u \in U$ will have the same label. (Hence, this can be interpreted as a labeling of the sets A_u as we did in our brief description in Figure 4.) For all arcs a in a path (F_u, A_u) (with $u \in U$) that start at the root of the

14

arborescence (C, A), we set $c(a) := 0$. For a path (F_u, A_u) that starts at a link ℓ that is not the root of (C, A), let $j \in \mathbb{Z}_{\geq 0}$ be the label of the incoming arc of ℓ. Then we set $c(a) = j + 1$ for all $a \in A_u$. Because (C, A) is an arborescence and we can consider the paths (A_u, F_u) in an order of increasing distance of their start point from the root of (C, A), this indeed defines a labeling $c \colon A \to \mathbb{Z}_{\geq 0}$.

For $i \in \{0, \ldots, k-1\}$, let $R_i \subseteq U$ be the set of up-links in U for which the arcs in A_u have a label j with $j \equiv i \pmod{k}$. Then $\{R_0, R_1, \ldots, R_{k-1}\}$ is a partition of U. Hence, there exists some $i \in \{0, \ldots, k-1\}$ such that $w(R_i) \leq {w(U)}/{k} \leq \varepsilon \cdot w(U)$, and we set $R = R_i$.

This completes the construction of $R \subseteq U$ with $w(R) \leq \varepsilon \cdot w(U)$. We choose the partition $\mathcal{C}$ of F to be the collection of the vertex sets of the connected components of the dependency graph of $U \setminus R$. The dependency graph of $U \setminus R$ arises from the dependency graph of U by deleting the arcs of each k-th label, starting with label i. Therefore, by the construction of the labeling $c \colon A \to \mathbb{Z}_{\geq 0}$, every path in the dependency graph of $U \setminus R$ has nonempty intersection with A_u for at most $k - 1$ links in $U \setminus R$. By Lemma 14, this implies that all elements of $\mathcal{C}$ are k-thin. Finally, we observe that by the definition of $\mathcal{C}$, we have that, for every $u \in U \setminus R$, there exists some $C \in \mathcal{C}$ with $F_u \subseteq C$. This shows property (i). □

5 Finding optimal thin components

In this section we prove Lemma 4, i.e., we prove that, for any constant $k \in \mathbb{Z}_{\geq 0}$, we can efficiently find a k-thin set $C \subseteq L$ that minimizes ${w(C)}/{w(\mathrm{Drop}_U(C))}$. To this end, we compute the optimal ratio

$$\rho^* := \min \left\{ \frac{w(C)}{w(\mathrm{Drop}_U(C))} : C \subseteq L \text{ is } k\text{-thin} \right\} \tag{4}$$

and a corresponding minimizer through a binary search procedure that relies on a dynamic program to decide whether some value $\rho \in \mathbb{R}$ is larger or smaller than ρ^*. (We recall that ${w(C)}/{w(\mathrm{Drop}_U(C))}$ is interpreted as ∞ whenever $w(\mathrm{Drop}_U(C)) = 0$, which, due to strictly positive link weights, is equivalent to $C = \emptyset$.)

5.1 Reducing to slack maximization

The question of whether some value $\rho \in \mathbb{R}$ is larger or smaller than ρ^* naturally reduces to maximizing the following slack function $\mathrm{slack}_\rho(C)$ over all k-thin sets $C \subseteq L$:

$$\mathrm{slack}_\rho(C) := \rho \cdot w(\mathrm{Drop}_U(C)) - w(C) \ .$$

More precisely, we have the following simple yet very helpful equivalence, which immediately follows from the definition of slack_ρ.

Observation 15. *Let $\rho \in \mathbb{R}$ and $C \subseteq L$ with $C \neq \emptyset$. Then $\mathrm{slack}_\rho(C) \geq 0$ if and only if $\frac{w(C)}{w(\mathrm{Drop}_U(C))} \leq \rho$.*

Hence, the question how a given $\rho \in \mathbb{R}$ compares to ρ^*, which is what we need to apply binary search, reduces to maximizing $\mathrm{slack}_\rho(C)$ over nonempty k-thin sets $C \subseteq L$, as formalized below.

Observation 16. *Let $\rho \in \mathbb{R}$. Then the following two statements are equivalent:*
(i) $\rho \geq \rho^*$.
(ii) $\max\{\mathrm{slack}_\rho(C) \colon C \subseteq L \text{ is } k\text{-thin and } C \neq \emptyset\} \geq 0$.

Proof. We have $\rho \geq \rho^*$ if and only if there is a k-thin set $C \subseteq L$ with ${w(C)}/{w(\mathrm{Drop}_U(C))} \leq \rho$. Notice that C must be nonempty, for otherwise we would have $\mathrm{Drop}_U(C) = \emptyset$ and ${w(C)}/{w(\mathrm{Drop}_U(C))}$ would have been interpreted as ∞, which violates finiteness of ρ. By Observation 15, we thus obtain that $\rho \geq \rho^*$ if and only if there is a nonempty k-thin set $C \subseteq L$ with $\mathrm{slack}_\rho(C) \geq 0$, as desired. □

Hence, to compute the value of ρ^* by binary search, it suffices to have an algorithm for the maximization problem in point (ii) of Observation 16. In Section 5.2, we describe a dynamic program that solves this maximization problem (or decides that the maximum $\mathrm{slack}_\rho(C)$ is negative). The result of our dynamic program is summarized in the lemma below.

Lemma 17. *Let $k \in \mathbb{Z}_{\geq 0}$ be a constant. Given a number $\rho \in \mathbb{R}$, a WTAP instance $(G = (V, E), L, w)$, and $U \subseteq L_{\mathrm{up}}$ such that the edge sets P_u for $u \in U$ are pairwise disjoint, we can in polynomial time compute a k-thin set $\overline{C} \subseteq L$ that maximizes $\mathrm{slack}_\rho(C)$ over all k-thin sets $C \subseteq L$. Moreover, if there is a nonempty maximizer, then $\overline{C} \neq \emptyset$.*

Note that Lemma 17 indeed implies that, for any $\rho \in \mathbb{R}$, we can decide in polynomial time whether $\max\{\mathrm{slack}_\rho(C)\colon C \subseteq L \text{ is } k\text{-thin and } C \neq \emptyset\} \geq 0$, due to the following. Let $\overline{C}$ be a k-thin set as described in Lemma 17. If $\overline{C} \neq \emptyset$, then $\mathrm{slack}_\rho(\overline{C}) = \max\{\mathrm{slack}_\rho(C)\colon C \subseteq L \text{ is } k\text{-thin and } C \neq \emptyset\}$ because $\overline{C}$ maximizes $\mathrm{slack}_\rho(\overline{C})$ over all k-thin sets $C \subseteq L$. Otherwise, if $\overline{C} = \emptyset$, then the maximum value of $\mathrm{slack}_\rho(C)$ over all k-thin sets is $\mathrm{slack}_\rho(\overline{C}) = 0$, and because Lemma 17 would have returned a nonempty maximizer if there had been one, we have $\max\{\mathrm{slack}_\rho(C)\colon C \subseteq L \text{ is } k\text{-thin and } C \neq \emptyset\} < 0$.

Before expanding on our dynamic program, we observe that using Lemma 17 to perform binary search over ρ readily implies Lemma 4, which we restate below for convenience.

Lemma 4. *Let $k \in \mathbb{Z}_{\geq 0}$ be a constant, $(G = (V, E), L, w)$ be a WTAP instance, and $U \subseteq L_{\mathrm{up}}$ such that the edge sets P_u for $u \in U$ are pairwise disjoint. Then we can compute in polynomial time a minimizer of*

$$\min\left\{\frac{w(C)}{w(\mathrm{Drop}_U(C))}\colon C \subseteq L \text{ is } k\text{-thin}\right\} \ .$$

Proof. We may assume without loss of generality that $w\colon L \to \mathbb{Z}_{>0}$ is integral by scaling up the weights if necessary. Moreover, we assume $U \neq \emptyset$, as the problem is trivial otherwise. First, observe that $0 \leq \rho^* \leq 1$ because our algorithm can always choose $C = \{u\}$ for any $u \in U$. As discussed, for any $\rho \in [0, 1]$, we can use Lemma 17 together with Observation 16 to decide in polynomial time whether $\rho \geq \rho^*$ or $\rho < \rho^*$. Moreover, if $\rho \geq \rho^*$, we obtain a nonempty k-thin set $C \subseteq L$ with $\mathrm{slack}_\rho(C) \geq 0$. By Observation 15, we then have $w(C)/w(\mathrm{Drop}_U(C)) \leq \rho$. Thus, using binary search, we can in polynomial time determine an interval $[a, b]$ with $\rho^* \in [a, b]$ and $b - a < 1/w(U)^2$, together with a k-thin set $C \subseteq L$ satisfying $w(C)/w(\mathrm{Drop}_U(C)) \leq b$.

We claim that this component C is a minimizer of $w(C)/w(\mathrm{Drop}_U(C))$ among all k-thin sets, as desired, i.e., $w(C)/\mathrm{Drop}_U(C) = \rho^*$. We suppose $w(C)/\mathrm{Drop}_U(C) > \rho^*$ and derive a contradiction. Let $C^* \subseteq L$ be such that $w(C^*)/\mathrm{Drop}_U(C^*) = \rho^*$. Because w is integral, we obtain the following contradiction:

$$\frac{1}{w(U)^2} > b - a \geq \frac{w(C)}{w(\mathrm{Drop}_U(C))} - \frac{w(C^*)}{w(\mathrm{Drop}_U(C^*))} \geq \frac{1}{w(\mathrm{Drop}_U(C)) \cdot w(\mathrm{Drop}_U(C^*))} \geq \frac{1}{w(U)^2} \ ,$$

where the penultimate fraction in the above chain of inequalities has a non-zero denominator because $w(C)/w(\mathrm{Drop}_U(C)) - w(C^*)/w(\mathrm{Drop}_U(C^*)) = w(C)/w(\mathrm{Drop}_U(C)) - \rho^* > 0$ by assumption. Hence, the component C that we computed fulfills $w(C)/\mathrm{Drop}_U(C) = \rho^*$. $\square$

We remark that instead of binary search, as used above, one could also employ Megiddo's [Meg79] parametric search technique to find the value of ρ^* in strongly polynomial time. This works out because our dynamic program to obtain $\overline{C}$ as described in Lemma 17 is a strongly polynomial time algorithm, where the value of ρ appears linearly in each comparison that we perform during the algorithm.

5.2 Proving Lemma 17 by dynamic programming

We now discuss a dynamic programming algorithm that computes in polynomial time a maximizer $\overline{C} \subseteq L$ of $\mathrm{slack}_\rho(C)$ over all k-thin sets $C \subseteq L$, with the additional property that $\overline{C} \neq \emptyset$ if there is a nonempty maximizer, thus implying Lemma 17.

The dynamic program follows the canonical approach of going from leaves toward the root r to build such a k-thin maximizer $\overline{C} \subseteq L$. More precisely, It computes k-thin link sets in subtrees of G that successively get combined to eventually obtain $\overline{C}$. To formalize this approach, we use the following notation to deal with subtrees. To refer to the vertices of a subtree with root $v \in V$, we denote by $D_v \subseteq V$ the set of all descendants of v in G. Moreover, to refer to links (or edges) contained in a subtree, we write, for any set X of links (or edges), $X[D_v] \subseteq X$ to denote all links (or edges) in X with both endpoints in D_v. Additionally, $\delta_X(D_v) \subseteq X$ denotes the set of links (or edges) in X with exactly one endpoint in D_v. Recall that the up-links U have disjoint edge sets P_u for $u \in U$. Hence, for any $v \in V$, the set $\delta_U(D_v)$ contains at most one up-link, and, if it does, then this up-link covers the last edge of the r-v path in G.

To build up some intuition for the dynamic program, consider a vertex $v \in V$ and the subtree below this vertex, i.e., the one with vertices D_v. To better understand how partial solutions for this subtree can get extended to larger subtrees, let us first consider a k-thin set of links $Q \subseteq L$ such that each link $\ell \in Q$ interacts with the subtree below v, i.e., ℓ has at least one endpoint in D_v. To later extend Q to a bigger subtree (i.e., to use Q in the *propagation step* of a dynamic program), there are only few things we need to know about Q. More precisely, let us partition Q into the links $C = Q[D_v]$ with both endpoints in D_v and the links $Y = \delta_U(D_v)$ with a single endpoint in D_v. See Figure 7. The crucial characteristics of Q that we need to know for propagation are:

- the set Y, which satisfies $|Y| \leq k$ because Q is k-thin;
- the slack when considering only those covered up-links that are contained in $U[D_v]$ and only accounting for the cost of links in C; we denote this slack by

$$\mathrm{slack}_\rho(C, Y, v) := \rho \cdot w\Big(\mathrm{Drop}_{U[D_v]}(C \cup Y)\Big) - w(C) \ ;$$

- if there is a link $u \in \delta_U(D_v)$, then we need to know whether the edges $P_u[D_v]$ are covered by $Q = C \cup Y$. This information is needed to decide whether u can later be dropped if Q becomes part of a larger link set that covers the edges in $P_u \setminus P_u[D_v]$.

Our dynamic program will therefore construct link sets C for triples (v, Y, x) consisting of
  (a) a vertex $v \in V$,
  (b) a set $Y \subseteq \delta_L(D_v)$ with $|Y| \leq k$, and
  (c) $x \in \{+, -\}$.
The value of x being $+$ indicates that there is an up-link $u \in \delta_U(D_v)$ and the triple (v, Y, x) represents a solution $C \cup Y$ that covers all edges of $P_u[D_v]$; otherwise, x should equal $-$. We denote by $\mathcal{T} \subseteq V \times 2^L \times \{+, -\}$ all triples (v, Y, x) fulfilling (a)–(c). We now define formally, when a link set $C \subseteq L[D_v]$ corresponds to the triple (v, Y, x), i.e., it complies with the above-mentioned interpretation of a triple.

Definition 18 (link set corresponding to (v, Y, x)). *Let $(v, Y, x) \in \mathcal{T}$. A link set $C \subseteq L[D_v]$ corresponds to the triple (v, Y, x) if $C \cup Y$ is k-thin and the following holds. If x equals $+$, then for C to correspond to (v, Y, x) we require that there exists an up-link $u \in \delta_U(D_v)$ and that $P_u[D_v] \subseteq \cup_{\ell \in C \cup Y} P_\ell$.*

We call a triple $(v, Y, x) \in \mathcal{T}$ *feasible* if there exists a link set $C \subseteq L[D_v]$ that corresponds to (v, Y, x). Otherwise, we call (v, Y, x) *infeasible*. By the above definition, this can be rephrased as follows.

Observation 19. *A triple $(v, Y, x) \in \mathcal{T}$ is infeasible if x equals $+$ and either*
  - $\delta_U(D_v) = \emptyset$, *or*

$Y \subseteq \delta_L(D_v)$

$C \subseteq L[D_v]$

$\mathrm{Drop}_{U[D_v]}(C \cup Y) \subseteq U[D_v]$

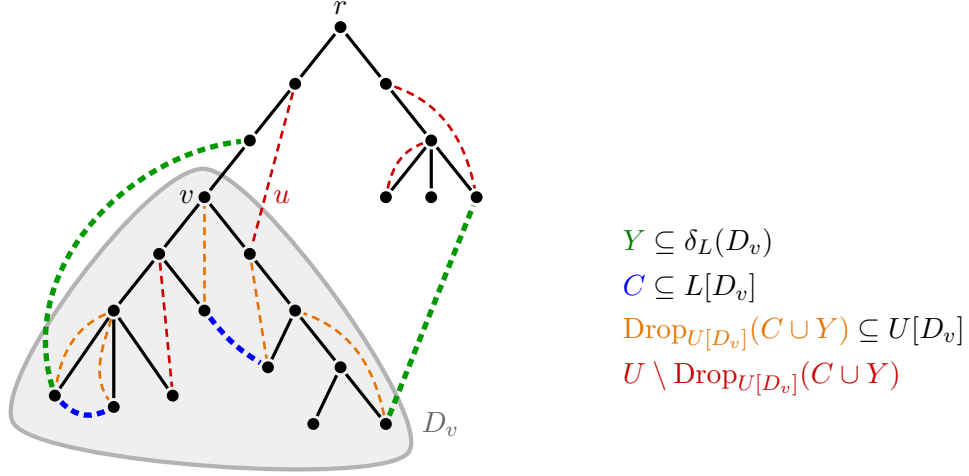$U \setminus \mathrm{Drop}_{U[D_v]}(C \cup Y)$

Figure 7: Illustration of the partial solutions we compute in the dynamic program. The blue links are those in C and the green ones those in Y. Up-links in U are drawn in orange and red, where the orange ones are those in $\mathrm{Drop}_{U[D_v]}(C \cup Y)$. In this example there is a link $u \in \delta_U(D_v)$. Moreover, the edges in $P_u[D_v]$ are covered by the links in $C \cup Y$. Thus, in this example, the set C corresponds to the triple $(v, Y, +)$, assuming $k \geq 3$.

- *there is a (single) up-link $u \in \delta_U(D_v)$, but there is no link set $C \subseteq L[D_v]$ such that $C \cup Y$ is k-thin and covers $P_u[D_v]$.*

Otherwise, the triple (v, Y, x) is called feasible.

For each triple $(v, Y, x) \in \mathcal{T}$, our dynamic program decides whether it is infeasible and, if not, computes a set $C \subseteq L[D_v]$ with the following properties:

(i) $C \cup Y$ is k-thin;

(ii) if x equals $+$, in which case there exists a link $u \in \delta_U(D_v)$ because the triple (v, Y, x) is feasible, we have $P_u[D_v] \subseteq \bigcup_{\ell \in C \cup Y} P_\ell$;

(iii) C maximizes $\mathrm{slack}_\rho(C, Y, v)$ among all link sets $C \subseteq L[D_v]$ satisfying (i) and (ii). Moreover, if there is a nonempty maximizer, then $C \neq \emptyset$.

See Figure 7. We denote the set C that we compute for a feasible triple $(v, Y, x) \in \mathcal{T}$ by $C(v, Y, x)$. Following standard terminology for dynamic programs, the set $C(v, Y, x)$ is called the *table entry* for the feasible triple (v, Y, x). If $(v, Y, x) \in \mathcal{T}$ is infeasible, then we simply store in the table entry that this triple is infeasible.

Note that being able to efficiently compute, for all feasible triples $(v, Y, x) \in \mathcal{T}$, a link set $C(v, Y, x)$ satisfying the properties (i)–(iii), implies Lemma 17. Indeed, because $U[D_r] = U$ and $L[D_v] = L$, the set $C = C(r, \emptyset, -)$ maximizes $\mathrm{slack}_\rho(C) = \mathrm{slack}_\rho(C, \emptyset, r)$ among all k-thin sets $C \subseteq L$. Moreover, if there is a nonempty maximizer, then, due to property (iii), the computed maximizer $C(r, \emptyset, -)$ is nonempty.

We now discuss how to compute the table entries by starting from the leaves and propagating them up to the root. The table entries for leaves are trivial to compute and hence we focus on the propagation step of the dynamic program. More precisely, we discuss how to compute a table entry for a triple $(v, Y, x) \in \mathcal{T}$, assuming that we already computed the table entries for all triples $(v', Y', x') \in \mathcal{T}$, where v' is a child of v in G.

Let $(v, Y, x) \in \mathcal{T}$ and let $v_1, \ldots, v_m$ be the children of v in G. First, if x equals $+$ and $\delta_U(D_v) = \emptyset$, then the triple (v, Y, x) is infeasible and we save this information as the table entry. Hence, in what follows, assume that there is a (single) link in $\delta_U(D_v)$ if x equals $+$. To compute the table entry $C(v, Y, x)$, we use the following observation about how any link set $C \subseteq L[D_v]$ that corresponds to the triple (v, Y, x) naturally

decomposes into link sets contained in the subtrees of the children of v and constantly many further links. (Think of C as a maximizer $C(v, Y, x)$ we try to find.) More precisely, $C \cup Y$ can be partitioned into the sets

- $C_i := C \cap L[D_{v_i}]$ for $i \in \{1, \ldots, m\}$, and
- $\overline{Y} := Y \cup \{\ell \in C : v \in V_\ell\}$.

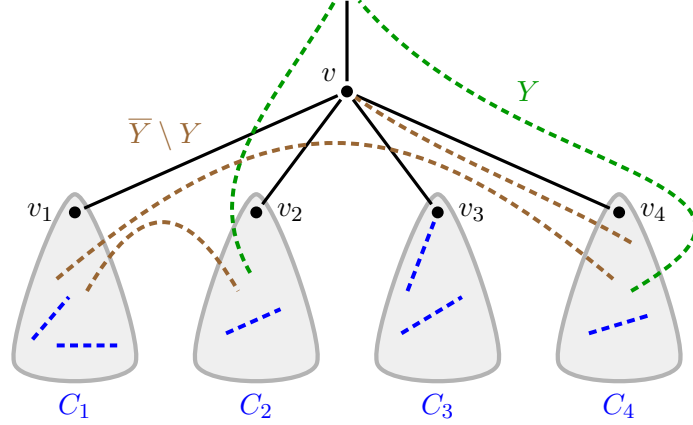Then $C = (\overline{Y} \setminus Y) \cup \bigcup_{i=1}^m C_i$. See Figure 8.

Figure 8: Illustration of the sets C_i, Y, and $\overline{Y}$.

Note that $|\overline{Y}| \le k$ because $\overline{Y}$ is a subset of the k-thin set $C \cup Y$ and all links $\ell \in \overline{Y}$ fulfill $v \in V_\ell$. Thus, $\overline{Y}$ has the following properties:

- $\overline{Y} \subseteq \{\ell \in L : v \in V_\ell\}$;
- $\overline{Y} \cap \delta_L(D_v) = Y$;
- $|\overline{Y}| \le k$;
- if x equals $+$ and if the link $u \in \delta_U(D_v)$ satisfies $u \notin \delta_U(v)$, then $u \in \delta_U(D_{v_i})$ for some child v_i of v in G and we have $\overline{Y} \cap \delta_L(D_{v_i}) \ne \emptyset$.

Let $\mathcal{Y}$ denote the family of all sets $\overline{Y} \subseteq L$ with these four properties. Because $|\overline{Y}| \le k$ for all $\overline{Y} \in \mathcal{Y}$ and k is constant, the family $\mathcal{Y}$ has only polynomial size.

So far we have shown that for any set C corresponding to the triple (v, Y, x), the set $C \cup Y$ can be partitioned into sets $\overline{Y} \in \mathcal{Y}$ and sets $C_i \subseteq L[D_{v_i}]$ for $i = 1, \ldots, m$. Thus, to maximize $\text{slack}_\rho(C, Y, v)$ over links sets $C \subseteq L[D_v]$ corresponding to the triple (v, Y, x), we can proceed as follows. First we enumerate over $\overline{Y}$, which can be done efficiently because $\mathcal{Y}$ has only polynomially many elements. Then we find, for each $\overline{Y} \in \mathcal{Y}$, sets $C_i \subseteq L[D_{v_i}]$ for $i \in \{1, \ldots, m\}$ that maximize $\text{slack}_\rho(C, Y, v)$ for $C := (\overline{Y} \setminus Y) \cup \bigcup_{i=1}^m C_i$.

We now discuss how we can find the sets C_i for a fixed $\overline{Y} \in \mathcal{Y}$. More precisely, we show how to efficiently find sets $C_i \subseteq L[D_{v_i}]$ for $i \in \{1, \ldots, m\}$ such that

$$C_{\overline{Y}} := (\overline{Y} \setminus Y) \cup \bigcup_{i=1}^m C_i \tag{5}$$

is k-thin, corresponds to the triple (v, Y, x), and maximizes $\text{slack}_\rho(C_{\overline{Y}}, Y, v)$ among all such sets $C_{\overline{Y}}$. If there is a non-empty maximizer, then the set $C_{\overline{Y}}$ we compute is nonempty. Moreover, if there are no sets C_i such that the resulting set $C_{\overline{Y}}$ as defined in (5) is k-thin and corresponds to the triple (v, Y, x), then we will detect this.

This is all that remains to be done, because if (v, Y, x) is feasible, then any set $C_{\overline{Y}}$ that maximizes $\text{slack}_\rho(C_{\overline{Y}}, Y, v)$ among all $\overline{Y} \in \mathcal{Y}$ is an optimal entry for the triple $(v, Y, x) \in \mathcal{T}$ (where we choose

$C_{\overline{Y}} \neq \emptyset$ if a non-empty maximizer exists). Otherwise, the triple (v, Y, x) is infeasible. We detect this because we cannot find a set $C_{\overline{Y}}$ corresponding to (v, Y, x) for any $\overline{Y} \in \mathcal{Y}$.

To find optimal sets C_i for a fixed $\overline{Y} \in \mathcal{Y}$, we first observe that, for any sets $C_i \subseteq L[D_{v_i}]$ for $i \in \{1, \ldots, m\}$, we have

$$\text{slack}_\rho(C_{\overline{Y}}, Y, v) = \sum_{i=1}^m \text{slack}_\rho\Big(C_i, \ \overline{Y} \cap \delta_L(D_{v_i}), \ v_i\Big) + \rho \cdot \sum_{\substack{i \in I^+: \\ u_i \in \text{Drop}_U(C_i \cup \overline{Y})}} w(u_i) - w(\overline{Y} \setminus Y) \ , \tag{6}$$

where $I^+ \subseteq \{1, \ldots, m\}$ is the set of indices $i \in \{1, \ldots, m\}$ for which there is an up-link $u_i \in \delta_U(v) \cap \delta_U(D_{v_i})$. Moreover, $C_{\overline{Y}} \cup Y$ is k-thin if and only if the sets C_i are chosen such that $C_i \cup (\overline{Y} \cap \delta_L(D_{v_i}))$ is k-thin for all $i \in \{1, \ldots, m\}$.

For each $i \in \{1, \ldots, m\}$ let $Y_i := (\overline{Y} \cap \delta_L(D_{v_i}))$. Note that $u_i \in \text{Drop}_U(C_i \cup \overline{Y})$ if and only if $u_i \in \text{Drop}_U(C_i \cup Y_i)$. Because of (6), finding optimal sets C_i reduces to finding, for each $i \in \{1, \ldots, m\}$, a set $C_i \subseteq L[D_{v_i}]$ that maximizes

$$\begin{cases} \text{slack}_\rho(C_i, Y_i, v_i) + \rho \cdot w(u_i) & \text{if } i \in I^+ \text{ and } u_i \in \text{Drop}_U(C_i \cup Y_i), \\ \text{slack}_\rho(C_i, Y_i, v_i) & \text{otherwise} \end{cases}$$

among all sets $C_i \subseteq L[D_{v_i}]$ for which $C_i \cup Y_i$ is k-thin. For each $i \in \{1, \ldots, m\}$, such a maximizer C_i is obtained as follows. Let $C_i^- := C(v_i, Y_i, -)$. Moreover, if the triple $(v_i, Y_i, +)$ is feasible, then we also define $C_i^+ := C(v_i, Y_i, +)$. We choose C_i to be either C_i^- or, if $(v_i, Y_i, +)$ is feasible, possibly C_i^+, as described in the below case distinction. (See Figure 9 for an illustration of the cases.)
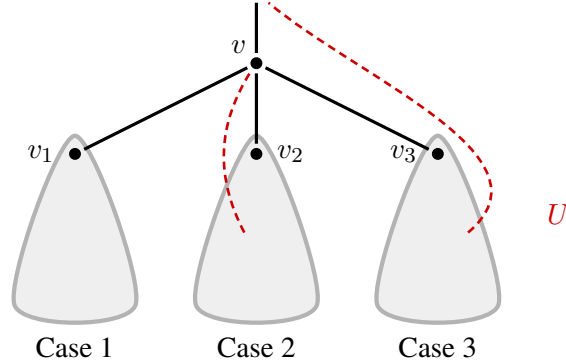

Figure 9: Illustration of the three cases for choosing $C_i \in \{C_i^-, C_i^+\}$.

Case 1: The set $\delta_U(D_{v_i})$ is empty.

In this case we set $C_i := C_i^-$.

Case 2: An up-link $u_i \in \delta_U(D_{v_i}) \cap \delta_U(v)$ exists. (Equivalently, $i \in I^+$.)

If $Y_i := \overline{Y} \cap \delta_L(D_{v_i}) = \emptyset$ or $(v_i, Y_i, +)$ is infeasible, we set $C_i := C_i^-$. Otherwise,

$$C_i := \begin{cases} C_i^+ & \text{if } \text{slack}_\rho\Big(C_i^+, \ Y_i, \ v_i\Big) + \rho \cdot w(u_i) \geq \text{slack}_\rho\Big(C_i^-, \ Y_i, \ v_i\Big) \ , \\ C_i^- & \text{otherwise} \ . \end{cases}$$

Case 3: An up-link $u_i \in \delta_U(D_{v_i}) \cap \delta_U(D_v)$ exists.

If x equals $+$ and the tuple $(v_i, Y_i, +)$ is infeasible, then there are no sets $C_i \subseteq L[D_{v_i}]$ such that $C_{\overline{Y}}$ as defined in (5) corresponds to the triple (v, Y, x). Otherwise, we set $C_i := C_i^x$.

20

One can easily check that, for a feasible triple (v, Y, x) and fixed set $\overline{Y} \in \mathcal{Y}$, the above choice of C_i leads to a k-thin set $C_{\overline{Y}}$ as defined in (5) that corresponds to (v, Y, x) and, among all possible choices for the sets C_i, maximizes $\mathrm{slack}_\rho(C_{\overline{Y}}, Y, v)$. Moreover, if there is no choice of the C_i that leads to a set $C_{\overline{Y}}$ that corresponds to the triple (v, Y, x), then this will be correctly detected in the third case above. Finally, the set $C_{\overline{Y}}$ is nonempty whenever there is a nonempty maximizer, because it is composed of sets C_i that are nonempty whenever possible.

It remains to analyze the running time of the algorithm. The number of triples (v, Y, x) we consider is no more than $|\mathcal{T}| \leq |V| \cdot |V|^{2k} \cdot 2$. (Note that $|V|^{2k}$ is an upper bound on the number of subset Y of $L \subseteq \binom{V}{2}$ of up to k links.) For each of these, the number of sets $\overline{Y}$ we enumerate can be bounded by $|\mathcal{Y}| \leq |V|^{2k}$ because $|\overline{Y}| \leq k$ for all $\overline{Y} \in \mathcal{Y}$. Because the number of children of a vertex v is at most $|V|$, computing $C_{\overline{Y}}$ for a fixed set $\overline{Y}$ (and a fixed triple $(v, Y, x) \in \mathcal{T}$) takes time polynomially bounded in $|V|$. Thus, the algorithm takes $|V|^{O(k)}$ time altogether, which is polynomial for constant k. This completes the proof of Lemma 17.

References

[Adj18]   D. Adjiashvili. "Beating Approximation Factor Two for Weighted Tree Augmentation with Bounded Costs". In: *ACM Transactions on Algorithms* 15.2 (2018), 19:1–19:26.

[BGJ20]   J. Byrka, F. Grandoni, and A. Jabal Ameli. "Breaching the 2-Approximation Barrier for Connectivity Augmentation: a Reduction to Steiner Tree". In: *Proceedings of 52nd ACM Symposium on Theory of Computing (STOC)*. 2020.

[CG18a]   J. Cheriyan and Z. Gao. "Approximating (Unweighted) Tree Augmentation via Lift-and-Project, Part I: Stemless TAP". In: *Algorithmica* 80 (2018), pp. 530–559.

[CG18b]   J. Cheriyan and Z. Gao. "Approximating (Unweighted) Tree Augmentation via Lift-and-Project, Part II". In: *Algorithmica* 80.2 (2018), pp. 608–651.

[Che+08]   J. Cheriyan, H. Karloff, R. Khandekar, and J. Könemann. "On the Integrality Ratio for Tree Augmentation". In: *Operations Research Letters* 36.4 (2008), pp. 399–401.

[CJR99]   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)*. 1999, pp. 510–520.

[CN13]   N. Cohen and Z. Nutov. "A $(1 + \ln 2)$-Approximation Algorithm for Minimum-Cost 2-Edge-Connectivity Augmentation of Trees with Constant Radius". In: *Theoretical Computer Science* 489 (2013), pp. 67–74.

[CSS01]   J. Cheriyan, A. Sebő, and Z. Szigeti. "Improving on the 1.5-Approximation of a Smallest 2-Edge Connected Spanning Subgraph". In: *SIAM Journal on Discrete Mathematics* 14 (2001), pp. 170–180.

[CTZ21]   F. Cecchetto, V. Traub, and R. Zenklusen. "Bridging the Gap Between Tree and Connectivity Augmentation: Unified and Stronger Approaches". In: *Proceedings of 53rd Annual ACM Symposium on Theory of Computing (STOC)*. To appear. See also https://arxiv.org/abs/2012.00086. 2021.

[Eve+09]   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". In: *ACM Transactions on Algorithms* 5.2 (2009), 21:1–21:17.

[Fio+18]   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)*. 2018, pp. 817–831.

[FJ81]   G. N. Frederickson and J. JáJá. "Approximation Algorithms for Several Graph Augmentation Problems". In: *SIAM Journal on Computing* 10.2 (1981), pp. 270–283.

[FT89]   A. Frank and É. Tardos. "An Application of Submodular Flows". In: *Linear Algebra and its Applications* 114–115 (1989), pp. 329–348.

[GKZ18]   F. Grandoni, C. Kalaitzis, and R. Zenklusen. "Improved Approximation for Tree Augmentation: Saving by Rewiring". In: *Proceedings of 50th ACM Symposium on Theory of Computing (STOC)*. 2018, pp. 632–645.

[Goe+94]   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)*. 1994, pp. 223–232.

[Grö+01]   C. Gröpl, S. Hougardy, T. Nierhoff, and H. J. Prömel. "Approximation Algorithms for the Steiner Tree Problem in Graphs". In: *Steiner Trees in Industry*. Ed. by Xiu Zhen Cheng and Ding-Zhu Du. Springer, 2001, pp. 235–279. ISBN: 978-1-4613-0255-1. URL: https://doi.org/10.1007/978-1-4613-0255-1_7.

[GW95]   M. X. Goemans and D. P. Williamson. "A General Approximation Technique for Constrained Forest Problems". In: *SIAM Journal on Computing* 24.2 (1995), pp. 296–317.

[HVV19]   C. Hunkenschröder, S. Vempala, and A. Vetta. "A $4/3$-Approximation Algorithm for the Minimum 2-Edge Connected Subgraph Problem". In: *ACM Transactions on Algorithms* 15.4 (2019), 55:1–55:28.

[IR18]   J. Iglesias and R. Ravi. *Coloring Down: $3/2$-approximation for special cases of the weighted tree augmentation problem*. https://arxiv.org/abs/1707.05240. 2018.

[Jai01]   K. Jain. "A Factor 2 Approximation Algorithm for the Generalized Steiner Network Problem". In: *Combinatorica* 21 (2001), pp. 39–60.

[KKL04]   G. Kortsarz, R. Krauthgamer, and J. R. Lee. "Hardness of Approximation for Vertex-Connectivity Network Design Problems". In: *SIAM Journal on Computing* 33.3 (2004), pp. 704–720.

[KN16]   G. Kortsarz and Z. Nutov. "A Simplified 1.5-Approximation Algorithm for Augmenting Edge-Connectivity of a Graph from 1 to 2". In: *ACM Transactions on Algorithms* 12.2 (2016), 23:1–23:20.

[KN18]   G. Kortsarz and Z. Nutov. "LP-Relaxations for Tree Augmentation". In: *Discrete Applied Mathematics* 239 (2018), pp. 94–105.

[KT93]   S. Khuller and R. Thurimella. "Approximation algorithms for graph augmentation". In: *Journal of Algorithms* 14.2 (1993), pp. 214–225.

[KV94]   S. Khuller and U. Vishkin. "Biconnectivity Approximations and Graph Carvings". In: *Journal of the ACM* 41.2 (1994), pp. 214–235.

[LRS11]   L. C. Lau, R. Ravi, and M. Singh. *Iterative Methods in Combinatorial Optimization*. 1st. New York, NY, USA: Cambridge University Press, 2011. ISBN: 978-0-521-18943-9.

[Meg79]   N. Megiddo. "Combinatorial Optimization with Rational Objective Functions". In: *Mathematics of Operations Research* 4.4 (1979), pp. 414–424.

[Nag03]   H. Nagamochi. "An Approximation for Finding a Smallest 2-Edge-Connected Subgraph Containing a Specified Spanning Tree". In: *Discrete Applied Mathematics* 126.1 (2003), pp. 83–113.

[Nut17]   Z. Nutov. "On the Tree Augmentation Problem". In: *Proceedings of 25th Annual Symposium on Algorithms (ESA)*. 2017, 61:1–61:14.

[Nut20]   Z. Nutov. *Approximation Algorithms for Connectivity Augmentation Problems*. https://arxiv.org/abs/2009.13257. 2020.

[SV14]   A. Sebő and J. Vygen. "Shorter tours by nicer ears: 7/5-Approximation for the graph-TSP, 3/2 for the path version, and 4/3 for two-edge-connected subgraphs". In: *Combinatorica* 34.5 (2014), pp. 597–629. DOI: 10.1007/s00493-014-2960-3.

[WS11]   D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011. ISBN: 978-0521195270.

[Zel96]   A. Zelikovsky. *Better approximation bounds for the network and Euclidean Steiner tree problems*. Tech. rep. CS-96-06. University of Virginia, 1996.