

Breaching the 2-Approximation Barrier for Connectivity Augmentation: A Reduction to Steiner Tree*

Jarosław Byrka
University of Wrocław
Poland
jby@cs.uni.wroc.pl

Fabrizio Grandoni
IDSIA, USI-SUPSI
Switzerland
fabrizio@idsia.ch

Afrouz Jabal Ameli
IDSIA, USI-SUPSI
Switzerland
afrouz@idsia.ch

ABSTRACT

The basic goal of survivable network design is to build a cheap network that maintains the connectivity between given sets of nodes despite the failure of a few edges/nodes. The *Connectivity Augmentation Problem* (CAP) is arguably one of the most basic problems in this area: given a k -(edge)-connected graph G and a set of extra edges (*links*), select a minimum cardinality subset A of links such that adding A to G increases its edge connectivity to $k + 1$. Intuitively, one wants to make an *existing* network more reliable by *augmenting* it with extra edges. The best known approximation factor for this NP-hard problem is 2, and this can be achieved with multiple approaches (the first such result is in [Frederickson and Jájá'81]).

It is known [Dinitz et al.'76] that CAP can be reduced to the case $k = 1$, a.k.a. the *Tree Augmentation Problem* (TAP), for odd k , and to the case $k = 2$, a.k.a. the *Cactus Augmentation Problem* (CacAP), for even k . Several better than 2 approximation algorithms are known for TAP, culminating with a recent 1.458 approximation [Grandoni et al.'18]. However, for CacAP the best known approximation is 2.

In this paper we breach the 2 approximation barrier for CacAP, hence for CAP, by presenting a polynomial-time $2 \ln(4) - \frac{967}{1120} + \epsilon < 1.91$ approximation. From a technical point of view, our approach deviates quite substantially from the current related literature. In particular, the better-than-2 approximation algorithms for TAP either exploit greedy-style algorithms or are based on rounding carefully-designed LPs. These approaches exploit properties of TAP that do not seem to generalize to CacAP. We instead use a reduction to the Steiner tree problem which was previously used in parameterized algorithms [Basavaraju et al.'14]. This reduction is not approximation preserving, and using the current best approximation factor for Steiner tree [Byrka et al.'13] as a black-box would not be good enough to improve on 2. To achieve the latter goal, we “open the box” and exploit the specific properties of the instances of Steiner tree arising from CacAP.

*The first author is supported by the NCN grant number 2015/18/E/ST6/00456. The last 2 authors are partially supported by the SNSF Excellence Grant 200020B_182865/1.

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 '20, June 22–26, 2020, Chicago, IL, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6979-4/20/06...\$15.00
<https://doi.org/10.1145/3357713.3384301>

In our opinion this connection between approximation algorithms for survivable network design and Steiner-type problems is interesting, and it might lead to other results in the area.

CCS CONCEPTS

• **Theory of computation** → **Routing and network design problems**.

KEYWORDS

approximation algorithms, connectivity augmentation, Steiner tree.

ACM Reference Format:

Jaroslaw Byrka, Fabrizio Grandoni, and Afrouz Jabal Ameli. 2020. Breaching the 2-Approximation Barrier for Connectivity Augmentation: A Reduction to Steiner Tree. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC '20)*, June 22–26, 2020, Chicago, IL, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3357713.3384301>

1 INTRODUCTION

The basic goal of *Survivable Network Design* is to construct cheap networks that provide connectivity guarantees between pre-specified sets of nodes even after the failure of a few edges/nodes (in the following we will focus on the edge failure case). This has many applications, e.g., in transportation and telecommunication networks.

The *Connectivity Augmentation Problem* (CAP) is among the most basic survivable network design problems. Here we are given a k -(edge)-connected¹ undirected graph $G = (V, E)$ and a collection L of extra edges (*links*). The goal is to find a minimum *cardinality* subset $OPT \subseteq L$ such that $G' = (V, E \cup OPT)$ is $(k + 1)$ -connected. Intuitively, we wish to augment an existing network to make it more resilient to edge failures. Dinitz et al. [9] (see also [6, 19]) presented an approximation-preserving reduction from this problem to the case $k = 1$ for odd k , and $k = 2$ for even k . This motivates a deeper understanding of the latter two special cases.

The case $k = 1$ is also known as the *Tree Augmentation Problem* (TAP). The reason for this name is that any 2-connected component of the input graph G can be contracted, hence leading to a tree. For this problem several better than 2 approximation algorithms are known [1, 4, 5, 10, 11, 16, 21, 22, 24]. In particular, the current best approximation factor is 1.458 [16].

The case $k = 2$ is also known as the *Cactus Augmentation Problem* (CacAP), where for similar reasons we can assume that the input

¹We recall that $G = (V, E)$ is k -connected if for every subset of edges $F \subseteq E$, $|F| \leq k - 1$, the graph $G' = (V, E \setminus F)$ is connected.

graph is a cactus². Here the best-known approximation factor is still 2, and this factor can be achieved with multiple approaches [12, 15, 18, 19]. A better approximation was achieved very recently for the special case where the input cactus is a cycle [14].

Hence 2 is also the best known approximation factor for CAP in general. One might also observe that TAP can be easily reduced to CacAP by duplicating the edges of the input instance. Hence CacAP and CAP are equivalent problems in terms of approximability.

1.1 Our Results and Techniques

The main result of this paper is the first better than 2 approximation algorithm for CacAP, hence for CAP.

THEOREM 1.1. *For any constant $\varepsilon > 0$, there is a polynomial-time $2 \ln(4) - \frac{967}{1120} + \varepsilon < 1.9092 + \varepsilon$ approximation algorithm for the Cactus Augmentation problem.*

COROLLARY 1.2. *For any constant $\varepsilon > 0$, there is a polynomial-time $2 \ln(4) - \frac{967}{1120} + \varepsilon < 1.9092 + \varepsilon$ approximation algorithm for the Connectivity Augmentation problem.*

PROOF. It follows directly from Theorem 1.1 and the reduction to CacAP implied by [9]. \square

Our result is based on a reduction to the (cardinality) Steiner tree problem by Basavaraju et al. [2]. The authors use this connection to design improved parameterized algorithms (see also [23] for a related result). Recall that in the *Steiner tree* problem we are given an undirected graph $G_{ST} = (T \cup S, E_{ST})$, where T is a set of t terminals and S a set of *Steiner nodes*. Our goal is to find a tree (*Steiner tree*) $OPT_{ST} = (T \cup A, F)$ that contains all the terminals (and possibly a subset of Steiner nodes A) and has the minimum possible number of edges $|OPT_{ST}|$. Basavaraju et al. observed that, given a CacAP instance $(G = (V, E), L)$, it is possible to construct (in polynomial time) an *equivalent* Steiner tree instance $G_{ST} = (T \cup L, E_{ST})$. Here T corresponds to the nodes of degree 2 in G , L are the Steiner nodes, and the edges E_{ST} are defined properly (more details in Section 2.1). In particular, an optimal solution to G_{ST} induces an optimal solution to (G, L) and vice versa. An example of the reduction is given in Figure 1. Unfortunately, this reduction is not approximation-preserving. In particular, by working out the simple details (see also Section 2.1), one obtains that a ρ_{ST} -approximation for Steiner tree implies a $\rho \leq 3\rho_{ST} - 2$ approximation for CacAP. The current best value of ρ_{ST} is $\ln 4 + \varepsilon < 1.39$ due to Byrka, Grandoni, Rothvoss and Sanità [3]. Hence this is not good enough³ to obtain $\rho < 2$.

In order to obtain our main result we use the same algorithm as in [3], but we analyze it differently. In particular, we exploit the specific structure of the instances of Steiner tree arising from CacAP instances via the above reduction to get a substantially better approximation factor.

In more detail (see also Section 2.2), in the analysis of the algorithm in [3] one considers an optimal Steiner tree solution $OPT_{ST} = (T \cup A, F)$ rooted at some arbitrary node r , marks a random subset $F_{mar} \subseteq F$ of edges so that each Steiner node is connected to some

²We recall that a *cactus* G is a connected undirected graph in which every edge belongs to exactly one cycle. For technical reasons it is convenient to allow length-2 cycles consisting of 2 parallel edges.

³One would need $\rho_{ST} < \frac{4}{3}$ here. Notice that this is not ruled out by the current lower bounds on the approximability of Steiner tree.

terminal via marked edges, and based on F_{mar} defines a proper (random) *witness set* $W(e)$ for each $e \in F$. The cost of the approximate solution turns out to be at most $(1 + \varepsilon) \sum_{e \in F} E[H_{|W(e)|}]$, where $H_i := 1 + \frac{1}{2} + \dots + \frac{1}{i}$ is the i -th harmonic number. In particular, the authors show that $E[H_{|W(e)|}] \leq \ln 4$ for each $e \in F$, hence the claimed approximation factor.

Our analysis of the algorithm deviates from [3] for the following critical reasons:

- (1) They (i.e., the authors of [3]) mark one child edge of each Steiner node chosen uniformly at random. In our case it is convenient to *favor* child edges with one terminal endpoint (if any). The fact that this helps is not obvious in our opinion.
- (2) As mentioned above, they provide a per-edge upper bound on $E[H_{|W(e)|}]$. We rather need to average over multiple edges in order to achieve a good bound. Finding a good way to do that is not trivial in our opinion.

We remark that, from a technical point of view, our result deviates quite substantially from prior approximation algorithms for TAP. The first improvements on a 2 approximation were achieved via greedy-style algorithms and a complex case analysis [10, 21, 22, 24]. More recent approaches are based on rounding stronger and stronger LP (or SDP) relaxations for the problem [1, 4, 5, 11, 16]. We also use an LP-based rounding algorithm, which is however defined for a generic Steiner tree instance (while the properties of TAP are used only in the analysis). In our opinion the connection that we established between the approximability of survivable network design problems and Steiner-type problems might lead to other results in the future.

1.2 Related Work

One can consider a natural weighted version of CAP where each link has a positive weight and the goal is to minimize the total weight of selected links. However, in this case the best-known approximation is 2 even for TAP, and improving on this is a major open problem in the area. The techniques used in this paper seem not to generalize to the weighted case. In particular, one might use a reduction to a node-weighted version of the Steiner tree problem, however the latter problem is harder and in general allows only a logarithmic approximation [20]. Some progress on weighted TAP was made in the case of small integer weights. In particular, when the largest weight W is upper bounded by a constant, better than 2 approximation algorithms are given in [1, 11, 16]. A technique in [25] allows one to extend these results to $W = O(\log n)$. Weighted TAP also admits a $1 + \ln 2$ approximation for arbitrary weights if the input tree has constant radius [8].

A problem closely related to CAP is to build a minimum size k -edge-connected spanning subgraph of a given input graph [7, 13, 17, 27].

2 STEINER TREE AND CONNECTIVITY AUGMENTATION

In this section we present the mentioned reduction in [2] from CacAP to Steiner tree (Section 2.1). Furthermore, we describe a specific Steiner tree approximation algorithm that we will use to solve the instance arising from the above reduction (Section 2.2). We analyze the resulting approximation factor in Section 3.

combination with a different analysis that exploits the properties of the specific Steiner tree instances arising from CacAP.

We next sketch the basic properties of the algorithm and analysis in [3] that we need here. A more detailed description is given in Section C in the Appendix for the sake of completeness. The authors of [3] consider an LP relaxation DCR_k for the problem based on *directed k -components* for a proper constant parameter k depending on ε . They iteratively solve this LP, sample a directed k -component C with probability proportional to the LP values, and contract C . The process ends when all terminals are contracted into one node. This algorithm can be derandomized, and the deterministic version is good enough for our application. We do not need more details about this algorithm, other than that it runs in polynomial time.

In the analysis the authors of [3] consider any feasible Steiner tree $ST = (T \cup A, F)$, which is seen as rooted at some arbitrary node r . Then the authors define a *marking scheme* where some child edge of each internal (Steiner) node is marked. A given marking scheme defines a *witness set* $W(e)$ for each edge e : this consists of pairs of terminals $\{t', t''\}$ such that the (simple) t' - t'' path in ST contains e and precisely one unmarked edge. We let $w(e) = |W(e)|$. Notice that marked edges form a forest spanning all nodes, with exactly one terminal in each connected component. Therefore, each unmarked edge e is on a single path only, hence $w(e) = 1$. Then the authors prove the following, where $H_i := 1 + \frac{1}{2} + \dots + \frac{1}{i}$ is the i -th harmonic number.

LEMMA 2.4. [3] *For any feasible Steiner tree $ST = (T \cup A, F)$ and marking scheme, for a large enough parameter $k = O_\varepsilon(1)$, the cost of the solution computed by the above algorithm is at most $(1 + \varepsilon) \sum_{e \in F} E[H_{w(e)}]$.*

3 AN IMPROVED CacAP APPROXIMATION ALGORITHM

In this section we present our improved approximation for CacAP. The algorithm is rather simple: we just build the Steiner tree instance $G_{ST} = (T \cup L, E_{ST})$ associated with the input CacAP instance (G, L) and compute an approximate solution APX_{ST} to G via the algorithm in [3] sketched in Section 2.2. Then we derive from APX_{ST} a feasible solution APX to the input CacAP instance as described in Corollary 2.2. We let apx denote the approximation ratio of this algorithm.

In Section 3.1 we describe our alternative marking scheme and prove some of its properties. In Section 3.2 we complete the analysis of the approximation factor.

3.1 An Alternative Marking Scheme

Recall that in the analysis of the Steiner tree approximation algorithm in [3], one can focus on a specific feasible Steiner tree ST and on a specific marking scheme (so that Steiner nodes are connected to some terminal via paths of marked edges). As feasible solution ST we consider the solution $OPT_{ST} = (T \cup OPT, F)$, of cost $|OPT| + t - 1$ and with terminals being leaves, guaranteed by Lemma 2.3.

We mark edges in the following way. Let us root OPT_{ST} at some Steiner node r which is adjacent to at least one terminal. For a

Steiner node ℓ , we let $d(\ell)$, $s(\ell)$ and $t(\ell)$ be the number of its children, Steiner children, and terminal children, resp. In particular $d(\ell) = s(\ell) + t(\ell)$ and (by Remark 1) $t(\ell) \leq 2$.

For each link node ℓ , there are two options. If ℓ has at least one terminal child, we select one such child t uniformly at random, and mark edge $\{\ell, t\}$. Otherwise, we choose a child ℓ' of ℓ (ℓ' being a Steiner node) uniformly at random, and mark edge $\{\ell, \ell'\}$. Notice that this is obviously a feasible marking scheme. Observe also that in our marking we *favor* edges connecting Steiner nodes to terminals: this will be critical in our analysis. See Figure 2 for a possible marking of this type.

Let APX_{ST} be the Steiner tree computed by the algorithm. Let F_{mar} and F_{unm} be the (random) sets of marked and unmarked edges, resp., that partition F . Recall that for each $e \in F$, there exists a (random) witness set $W(e)$ of size $w(e) = |W(e)|$. Observe that each Steiner node ℓ has precisely one marked child edge $m(\ell)$. We let the *cost* $c(\ell)$ of ℓ be $E[H_{w(m(\ell))}]$. The following bound on the approximation ratio holds.

LEMMA 3.1. $apx \leq 2\varepsilon + \frac{1+\varepsilon}{|OPT|} \sum_{\ell \in OPT} c(\ell)$.

PROOF. Recall that by Lemma 2.4 the expected cost of the computed Steiner tree APX_{ST} is, modulo a factor $(1 + \varepsilon)$, at most

$$\begin{aligned} E[\sum_{e \in F} H_{w(e)}] &= E[\sum_{e \in F_{mar}} H_{w(e)} + \sum_{e \in F_{unm}} H_{w(e)}] \\ &= E[\sum_{e \in F_{mar}} H_{w(e)} + |F_{unm}|] = E[\sum_{e \in F_{mar}} H_{w(e)}] + t - 1. \end{aligned}$$

In the second-last equality above we used the fact that $w(e) = 1$ deterministically for an unmarked edge, and in the last equality above the fact that there are precisely $|OPT|$ marked edges and consequently exactly $t - 1$ unmarked ones. From APX_{ST} we derive a feasible solution APX to the input instance of cost $|APX| = |APX_{ST}| - 1 + t$ by Corollary 2.2. Hence

$$\begin{aligned} |APX| &\leq (1 + \varepsilon)(E[\sum_{e \in F_{mar}} H_{w(e)}] + t - 1) - 1 + t \\ &\leq (1 + \varepsilon)E[\sum_{e \in F_{mar}} H_{w(e)}] + 2\varepsilon|OPT|. \end{aligned}$$

In the last inequality above we used the trivial lower bound $|OPT| \geq t/2$ that we mentioned earlier. The claim follows since by definition $\sum_{e \in F_{mar}} E[H_{w(e)}] = \sum_{\ell \in OPT} c(\ell)$. \square

From the above lemma, modulo factors $(1 + \varepsilon)$, the approximation ratio of our algorithm is given by the *average cost* of Steiner nodes. The following lemma gives a generic upper bound on the cost for each non-root Steiner node based on the degree sequence of its ancestors⁴.

LEMMA 3.2. *Given a non-root Steiner node ℓ , let ℓ_q be the lowest proper ancestor⁵ of ℓ with $t(\ell_q) > 0$. Let $\ell = \ell_1, \ell_2, \dots, \ell_q$, $q \geq 2$, be the simple path between ℓ and ℓ_q , and let $d_i = d(\ell_i)$. Then*

$$c(\ell) = \sum_{h=1}^{q-2} \frac{(d_{h+1} - 1)H_{d_1+\dots+d_h-h+1}}{d_2 \cdot \dots \cdot d_{h+1}} + \frac{H_{d_1+\dots+d_{q-1}-q+2}}{d_2 \cdot \dots \cdot d_{q-1}}.$$

PROOF. By definition $c(\ell) = c(\ell_1) = E[H_{w(e)}]$, where $e = m(\ell_1) = \{\ell_1, \ell_0\}$ is the marked child edge of ℓ_1 . Recall that $W(e)$ contains one entry for each path in the tree between terminals

⁴Observe that for the root r , $c(r) = H_{d(r)-1}$ deterministically.

⁵Observe that this ancestor exists since the root has this property by assumption.

that contains e and precisely one unmarked edge. In our specific case, condition on $\{\ell_0, \ell_1\}, \{\ell_1, \ell_2\}, \dots, \{\ell_{h-1}, \ell_h\}$ being a maximal sequence of consecutive marked edges. Notice that by construction $\{\ell_{q-1}, \ell_q\}$ is unmarked (since ℓ_q has a terminal child by definition), hence $h \leq q-1$. In this case $w(e) = d_1 + \dots + d_h - (h-1)$. For $h < q-1$, the mentioned event happens with probability $\frac{1}{d_2} \cdot \dots \cdot \frac{1}{d_h} \cdot \frac{d_{h+1}-1}{d_{h+1}}$. For $h = q-1$, this probability is $\frac{1}{d_2} \cdot \dots \cdot \frac{1}{d_h}$. The claim follows by computing the expectation of $H_{w(e)}$. \square

We next provide an upper bound on $c(\ell)$ as a function of $d(\ell)$ only. Let us define the following variant of H_i :

$$\hat{H}_i := \frac{1}{2}H_i + \frac{1}{4}H_{i+1} + \dots = \sum_{j \geq 0} \frac{1}{2^{j+1}} H_{i+j}.$$

One has that $\hat{H}_1 = \ln(4)$ and $\hat{H}_{j+1} = 2\hat{H}_j - H_j$. Notice that, modulo an additive ε , \hat{H}_i is precisely the approximation factor for Steiner tree achieved in [3]. The first few approximate values of \hat{H}_i are $\hat{H}_1 < 1.3863$, $\hat{H}_2 < 1.7726$, $\hat{H}_3 < 2.0452$, $\hat{H}_4 < 2.2571$, $\hat{H}_5 < 2.4308$, $\hat{H}_6 < 2.5781$, $\hat{H}_7 < 2.7062$, and $\hat{H}_8 < 2.8195$.

The proof of the following lemma, though not entirely trivial, is mostly based on algebraic manipulations and therefore we postpone it to the appendix.

LEMMA 3.3. *For any $\ell \in OPT$, $c(\ell) \leq \hat{H}_{d(\ell)}$.*

In next subsection we will see that for a carefully defined subset of Steiner nodes ℓ it is possible to obtain a better upper bound on $c(\ell)$ than the one provided by Lemma 3.3. This will be critical in our analysis since the latter bound is not strong enough.

3.2 Analysis of the Approximation Factor

In this section we upper bound the approximation factor apx as given by Lemmas 3.1 and 3.2. In order to simplify our analysis, it is convenient to focus our attention on a specific class of well-structured Steiner trees OPT_{ST} (see also Figure 2). The following lemma shows that this is (essentially) w.l.o.g.

DEFINITION 3.4. A rooted Steiner tree is well-structured if, for every Steiner node ℓ : (1) ℓ has at least 2 children and (2) ℓ has 0 or 2 terminal children.

LEMMA 3.5. *Let ρ be the supremum of $\rho(OPT_{ST}) = \frac{1}{|OPT|} \sum_{\ell \in OPT} c(\ell)$ over Steiner trees $OPT_{ST} = (T \cup OPT, F)$, and ρ_{ws} be the same quantity computed over the subset of well-structured Steiner trees OPT_{ST} of the mentioned type. Then $\rho \leq \max\{\hat{H}_1, \rho_{ws}\}$.*

PROOF. Recall that in OPT_{ST} each Steiner node ℓ has at most 2 terminal children. Consider any such tree where some Steiner node ℓ' has precisely one terminal child t . Consider the tree OPT'_{ST} which is obtained from OPT_{ST} by appending to ℓ' a second terminal child t' . Observe that the value of $c(\ell)$ does not decrease for any ℓ , and it increases for $\ell = \ell'$. Thus $\rho(OPT'_{ST}) > \rho(OPT_{ST})$. Hence ρ is equal to the supremum of $\rho(OPT_{ST})$ over the subfamily of trees that satisfies (2) in Definition 3.4.

Now consider any tree OPT_{ST} that satisfies (2), and let $o(OPT_{ST})$ be the number of its Steiner nodes with precisely one child. We prove by induction on $o(OPT_{ST})$ that $\rho(OPT_{ST}) \leq \max\{\hat{H}_1, \rho_{ws}\}$. The claim is trivially true for $o(OPT_{ST}) = 0$ since in this case OPT_{ST} is well-structured. Assume the claim is true up to $q-1 \geq 0$,

and consider $OPT_{ST} = (T \cup OPT, F)$ with $o(OPT_{ST}) = q$. Let ℓ' be any Steiner node with precisely one child ℓ'' . Observe that ℓ'' has to be a Steiner node as well by (2), and that $c(\ell') \leq \hat{H}_1$ by Lemma 3.3. Consider the tree $OPT'_{ST} = (T \cup OPT', F')$ obtained by contracting edge (ℓ', ℓ'') . We observe that OPT'_{ST} satisfies (2), $o(OPT'_{ST}) = q-1$ and $|OPT'| = |OPT| - 1$. Note also that for any Steiner node ℓ different from ℓ' and ℓ'' the value of $c(\ell)$ does not change, while for the new node $\tilde{\ell}$ resulting from the contraction one has $c(\tilde{\ell}) = c(\ell'')$. We can conclude that

$$\begin{aligned} & \frac{1}{|OPT|} \sum_{\ell \in OPT} c(\ell) \\ & \leq \frac{1}{|OPT|} (\hat{H}_1 + \sum_{\ell \in OPT \setminus \{\ell''\}} c(\ell)) \\ & = \frac{1}{|OPT|} (\hat{H}_1 + \sum_{\ell \in OPT'} c(\ell)) \leq \max\{\hat{H}_1, \frac{1}{|OPT'|} \sum_{\ell \in OPT'} c(\ell)\} \\ & \leq \max\{\hat{H}_1, \rho_{ws}\}, \end{aligned}$$

where in the last inequality we used the inductive hypothesis. \square

We next show an upper bound on ρ_{ws} which is strictly greater than \hat{H}_1 . It then follows from Lemma 3.5 that the same upper bound holds on ρ . For this goal, we next assume that OPT_{ST} is well-structured.

The upper bound on $c(\ell)$ from Lemma 3.3 is not sufficient to achieve a good approximation factor. In order to achieve a tighter bound, we consider the following classification of the Steiner nodes (see also Figure 2).

DEFINITION 3.6. A Steiner node ℓ' is a *good father* if it has at least one terminal child (hence precisely 2 such children by the above assumptions), and a *bad father* otherwise. Each Steiner child ℓ of a good father ℓ' is *good*, and all other Steiner nodes are *bad*. Let OPT_{gf} , OPT_{bf} , OPT_g and OPT_b denote the sets of good fathers, bad fathers, good nodes and bad nodes, resp.

Notice that the above classification is not affected by the random choices in the marking scheme. For good nodes, the analysis of the cost can be refined as follows.

LEMMA 3.7. *For any $\ell \in OPT_g$, $c(\ell) \leq H_{d(\ell)}$.*

PROOF. Suppose ℓ has a parent ℓ' , which is a good father by definition. This implies that the edge (ℓ', ℓ) is deterministically unmarked, hence $w(m(\ell)) = d(\ell)$ deterministically. If ℓ has no parent (i.e., it is the root r), then $w(m(\ell)) = d(\ell) - 1$. The claim follows. \square

Putting everything together, we obtain the following.

LEMMA 3.8. *$apx \leq 2\varepsilon + \frac{1+\varepsilon}{|OPT|} \sum_{\ell \in OPT} c'(\ell)$ where*

$$c'(\ell) = \begin{cases} H_{d(\ell)} & \text{if } \ell \in OPT_g; \\ \hat{H}_{d(\ell)} & \text{if } \ell \in OPT_b. \end{cases}$$

PROOF. It follows from Lemma 3.1, by replacing $c(\ell)$ as in Lemma 3.2 with the upper bounds given by Lemmas 3.3 and 3.7. \square

We rewrite the upper bound from Lemma 3.8 as follows. Let $p \in [0, \hat{H}_2 - H_2]$ be a parameter to be fixed later. Intuitively, each good Steiner node $\ell \in OPT_g$ pays a *present* p to its (good) father

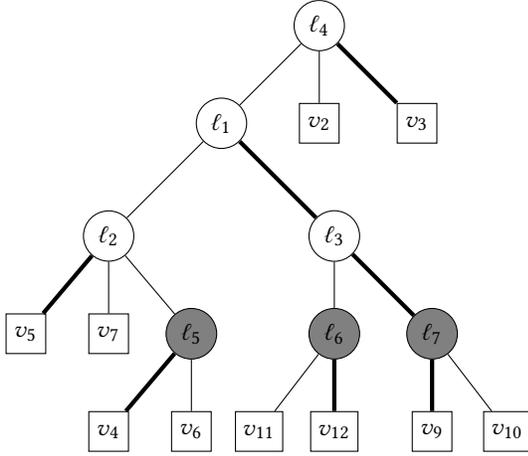


Figure 2: A feasible Steiner tree for the instance of Figure 1, which happens to be well-structured. Bold edges denote a possible marking. One has $m(\ell_3) = e := \{\ell_3, \ell_7\}$, and $W(e)$ contains $\{v_9, v_{12}\}$, $\{v_9, v_5\}$ and $\{v_9, v_3\}$. Notice that $w(e) = |W(e)| = d(\ell_3) + d(\ell_1) - 1$. Leaf-Steiner nodes are drawn in grey. Here ℓ_2 (resp., ℓ_3) is a good (resp., bad) father. Consequently ℓ_5 (resp., ℓ_6) is good (resp., bad). A feasible grouping is $g(\ell_2) = \{\ell_2\}$, $g(\ell_3) = \{\ell_3, \ell_7\}$, $g(\ell_1) = \{\ell_1, \ell_6\}$, $g(\ell_4) = \{\ell_4\}$, and $g(\ell_5) = \{\ell_5\}$.

$\ell' \in OPT_{gf}$ to thank ℓ' for making itself good. This increases the cost of ℓ by p . Symmetrically, each good father $\ell' \in OPT_{gf}$ collects presents from its (good) Steiner children and uses them to lower its own cost. Clearly by definition the total modification of the cost is zero. Let us call $c''(\ell)$ the modified costs. Then one obtains the following equality:

$$\frac{1}{|OPT|} \sum_{\ell \in OPT} c'(\ell) = \frac{1}{|OPT|} \sum_{\ell \in OPT} c''(\ell) \quad (1)$$

where

$$c''(\ell) = \begin{cases} H_{d(\ell)} + p - s(\ell)p & \text{if } \ell \in OPT_g \cap OPT_{gf}; \\ H_{d(\ell)} + p & \text{if } \ell \in OPT_g \cap OPT_{bf}; \\ \hat{H}_{d(\ell)} - s(\ell)p & \text{if } \ell \in OPT_b \cap OPT_{gf}; \\ \hat{H}_{d(\ell)} & \text{if } \ell \in OPT_b \cap OPT_{bf}. \end{cases}$$

In order to upper bound (1), we partition OPT into groups of nodes as follows (see also Figure 2).

Definition 3.9. A Steiner node ℓ is *leaf-Steiner* if it has no Steiner children (i.e., $d(\ell) = t(\ell) = 2$) and *internal-Steiner* otherwise (i.e., $s(\ell) > 0$). We let OPT_{lf} and OPT_{in} be the set of leaf-Steiner and internal-Steiner nodes, resp.

We associate to each $\ell \in OPT_{in}$ a distinct subset $OPT_{lf}(\ell)$ of precisely $s(\ell) - 1$ leaf-Steiner nodes, and let $g(\ell) = \{\ell\} \cup OPT_{lf}(\ell)$ be the *group* of ℓ . The mapping is constructed iteratively in a bottom-up fashion as follows. Initially all Steiner nodes are unprocessed. We maintain the invariant that the subtree rooted at an unprocessed leaf-Steiner node or at a processed node with unprocessed parent contains precisely one unprocessed leaf-Steiner node. Clearly the invariant holds at the beginning of the process. We consider any unprocessed internal-Steiner node ℓ whose Steiner descendants

are either processed or leaf-Steiner nodes. By the invariant, each subtree rooted at a Steiner child of ℓ (which is either an unprocessed leaf-Steiner node or a processed internal-Steiner node) contains one unprocessed leaf-Steiner node. Among this set of $s(\ell)$ unprocessed leaf-Steiner nodes, we select arbitrarily a set $OPT_{lf}(\ell)$ of size $s(\ell) - 1$ and set $g(\ell) = \{\ell\} \cup OPT_{lf}(\ell)$. All nodes in $g(\ell)$ are marked as processed. Observe that the subtree rooted at ℓ still contains an unprocessed leaf-Steiner node, hence the invariant is preserved in the following steps. At the end of the process (i.e., after processing the root r) there will be precisely one leaf-Steiner node ℓ^* which is still unprocessed, which forms a special group $g(\ell^*) = \{\ell^*\}$ on its own. Notice that the groups define a partition of OPT . In particular, $OPT = \{\ell^*\} \cup \bigcup_{\ell \in OPT_{in}} g(\ell)$. Notice also that $|g(\ell)| = s(\ell)$ for all $\ell \in OPT_{in}$ (while $|g(\ell^*)| = 1$).

Let $a(\ell)$ be the average value of $c''(\cdot)$ over the elements of $g(\ell)$. Then obviously the maximum value of $a(\ell)$ over the groups upper bounds the average value of $c''(\cdot)$:

$$\frac{1}{|OPT|} \sum_{\ell \in OPT} c''(\ell) \leq \max_{\ell \in OPT_{in} \cup \{\ell^*\}} \{a(\ell)\}. \quad (2)$$

For $\ell = \ell^*$ one has that $a(\ell^*) = c''(\ell^*) = \hat{H}_2$ if ℓ^* is bad, and $a(\ell^*) = c''(\ell^*) = H_2 + p \leq \hat{H}_2$ otherwise. For the other groups $g(\ell)$, there is always a subset of $s(\ell) - 1$ leaves whose contribution to the cost is at most \hat{H}_2 each by the same argument as above. Furthermore, we have to add the cost $c''(\ell)$. We can conclude that:

$$a(\ell) \leq \begin{cases} a_1(s(\ell)) := \frac{H_{s(\ell)+2} + p - s(\ell)p + (s(\ell)-1)\hat{H}_2}{s(\ell)} & \text{if } \ell \in OPT_g \cap OPT_{gf}; \\ a_2(s(\ell)) := \frac{H_{s(\ell)} + p + (s(\ell)-1)\hat{H}_2}{s(\ell)} & \text{if } \ell \in OPT_g \cap OPT_{bf}; \\ a_3(s(\ell)) := \frac{\hat{H}_{s(\ell)+2} - s(\ell)p + (s(\ell)-1)\hat{H}_2}{s(\ell)} & \text{if } \ell \in OPT_b \cap OPT_{gf}; \\ a_4(s(\ell)) := \frac{\hat{H}_{s(\ell)} + (s(\ell)-1)\hat{H}_2}{s(\ell)} & \text{if } \ell \in OPT_b \cap OPT_{bf}; \\ \hat{H}_2 & \text{if } \ell = \ell^*. \end{cases}$$

In the first and third case above we used the fact that $d(\ell) = s(\ell) + 2$ (ℓ is a good father, hence has 2 terminal children), while in the second and fourth case the fact that $d(\ell) = s(\ell)$ (ℓ is a bad father, hence has no terminal child).

We are now ready to prove the main result of this paper.

PROOF OF THEOREM 1.1. Consider the above algorithm. Combining Lemma 3.8 with (1) and (2) one gets

$$apx \leq 2\varepsilon + (1 + \varepsilon) \max_{i \geq 1} \{\hat{H}_2, a_1(i), a_2(i), a_3(i), a_4(i)\}. \quad (3)$$

We need the following technical result (proof in Appendix).

CLAIM 1. For any $p \in [0, \hat{H}_2 - H_2]$, the maximum of $a_1(i)$, $a_2(i)$, $a_3(i)$, and $a_4(i)$ is achieved for i at most 6, 8, 6 and 8, resp.

From (3) and Claim 1, for any $p \in [0, \hat{H}_2 - H_2]$, one has

$$apx \leq 2\varepsilon + (1 + \varepsilon) \max\{\hat{H}_2, \max_{1 \leq i \leq 6} \{a_1(i)\}, \max_{1 \leq i \leq 8} \{a_2(i)\}, \max_{1 \leq i \leq 6} \{a_3(i)\}, \max_{1 \leq i \leq 8} \{a_4(i)\}\}. \quad (4)$$

Numerically the minimum of the right-hand side of (4) is achieved for $p \approx 0.135$, and the two largest values inside the maximum turn out to be $a_2(7)$ and $a_3(1)$. By imposing $\frac{H_7 + 6\hat{H}_2 + p}{7} = a_2(7) = a_3(1) = \hat{H}_3 - p$ one gets $p = \frac{7\hat{H}_3 - H_7 - 6\hat{H}_2}{8}$. For that value of p the

value of the maximum is precisely $\frac{H_7 + 6\hat{H}_2 + \hat{H}_3}{8} = 2 \ln 4 - \frac{967}{1120}$. The claim follows by scaling ε properly. \square

A OMITTED PROOFS FROM SECTION 3

CLAIM 2. $H_{d_1} + \sum_{j=d_1+1}^{\infty} \frac{1}{j \cdot 2^{j-d_1}} = \hat{H}_{d_1}$.

PROOF. Note that

$$\begin{aligned} \hat{H}_{d_1} &= \sum_{i=0}^{\infty} \frac{H_{d_1+i}}{2^{i+1}} = \sum_{i=0}^{\infty} \frac{H_{d_1} + \sum_{j=1}^i \frac{1}{d_1+j}}{2^{i+1}} \\ &= H_{d_1} + \sum_{j=1}^{\infty} \frac{\sum_{i=j}^{\infty} \frac{1}{2^{i+1}}}{d_1+j} = H_{d_1} + \sum_{j=1}^{\infty} \frac{1}{(d_1+j)2^j}. \end{aligned}$$

\square

PROOF OF LEMMA 3.3. The claim is trivially true if ℓ is the root since in that case $c(\ell) = H_{d(\ell)-1} < \hat{H}_{d(\ell)}$. So assume ℓ is not the root. For a generic sequence $S = (d_1, \dots, d_k)$ of positive integers, let us define

$$f(S) = \sum_{j=1}^{k-1} \frac{(d_{j+1} - 1) \cdot H_{d_1+d_2+\dots+d_j-j+1}}{d_2 \cdot d_3 \dots d_{j+1}} + \frac{H_{d_1+d_2+\dots+d_k-k+1}}{d_2 \cdot d_3 \dots d_k}.$$

Intuitively, this is the right-hand side of the equation in Lemma 3.2. For an infinite sequence $S' = (d_1, d_2, \dots)$ of positive integers, we analogously define

$$f(S') = \sum_{j=1}^{\infty} \frac{(d_{j+1} - 1) \cdot H_{d_1+d_2+\dots+d_j-j+1}}{d_2 \cdot d_3 \dots d_{j+1}}$$

Given a finite sequence $S = (d_1, \dots, d_k)$ of the above type, let $\tilde{S} = (d_1, \dots, d_k, 2, 2, \dots)$ be its infinite extension where we add an infinite sequence of 2 at the end.

CLAIM 3. $f(S) \leq f(\tilde{S})$.

PROOF. By definition

$$\begin{aligned} &f(\tilde{S}) - f(S) \\ &= \sum_{j=k}^{\infty} \frac{(d_{j+1} - 1) \cdot H_{d_1+d_2+\dots+d_j-j+1}}{d_2 \cdot d_3 \dots d_{j+1}} - \frac{H_{d_1+d_2+\dots+d_k-k+1}}{d_2 \cdot d_3 \dots d_k} \\ &\geq \sum_{j=k}^{\infty} \frac{(d_{j+1} - 1) \cdot H_{d_1+d_2+\dots+d_k-k+1}}{d_2 \cdot d_3 \dots d_{j+1}} - \frac{H_{d_1+d_2+\dots+d_k-k+1}}{d_2 \cdot d_3 \dots d_k} \\ &= \frac{H_{d_1+d_2+\dots+d_k-k+1}}{d_2 \cdot d_3 \dots d_k} \sum_{j=1}^{\infty} \frac{1}{2^j} - \frac{H_{d_1+d_2+\dots+d_k-k+1}}{d_2 \cdot d_3 \dots d_k} \\ &= \frac{H_{d_1+d_2+\dots+d_k-k+1}}{d_2 \cdot d_3 \dots d_k} - \frac{H_{d_1+d_2+\dots+d_k-k+1}}{d_2 \cdot d_3 \dots d_k} = 0. \end{aligned}$$

\square

By the above claim it is sufficient to consider infinite sequences of type \tilde{S} . We can also assume w.l.o.g. that all $d_i, i \geq 2$, in such sequences are at least 2 by the following claim.

CLAIM 4. Let $\tilde{S} = (d_1, \dots, d_k, 2, 2, \dots)$ and assume there exists $d_i = 1$ in the sequence for some $i \geq 2$. Let $\tilde{S}_i = (d_1, \dots, d_i - 1, d_i + 1, \dots, d_k, 2, 2, \dots)$ be the subsequence where the i -th entry is removed. Then $f(\tilde{S}) = f(\tilde{S}_i)$.

PROOF. Consider the entries in the sum defining $f(\tilde{S})$ and $f(\tilde{S}_i)$. The entry $j = i - 1$ in $f(\tilde{S})$ has value 0. For $j < i - 1$, the j -th entries in $f(\tilde{S})$ and $f(\tilde{S}_i)$ are identical. For $j > i - 1$, the j -th entry in $f(\tilde{S})$ is equal to the $j - 1$ -th entry in $f(\tilde{S}_i)$. \square

By the above claims we can focus on infinite sequences $S = (d_1, \dots, d_k, 2, 2, \dots)$ where $d_i \geq 2$ for $i \geq 2$. Let us prove by induction on $k \geq 2$ that $f(S) \leq \hat{H}_{d_1}$. The claim is true by definition for $k = 2$. Next consider any $k > 2$ and assume the claim is true for all values up to $k - 1$. Define $S' = (d_1 + d_2 - 1, d_3, \dots, d_k, 2, 2, \dots)$. By definition and inductive hypothesis:

$$f(S) = H_{d_1} \frac{d_2 - 1}{d_2} + \frac{f(S')}{d_2} \leq H_{d_1} \frac{d_2 - 1}{d_2} + \frac{\hat{H}_{d_1+d_2-1}}{d_2}.$$

By Claim 2,

$$\begin{aligned} &H_{d_1} \frac{d_2 - 1}{d_2} + \frac{\hat{H}_{d_1+d_2-1}}{d_2} \\ &= H_{d_1} \frac{d_2 - 1}{d_2} + \frac{1}{d_2} \left(H_{d_1+d_2-1} + \sum_{j \geq d_1+d_2} \frac{1}{j \cdot 2^{j-d_1-d_2+1}} \right) \\ &= H_{d_1} + \sum_{j=d_1+1}^{d_1+d_2-1} \frac{1}{j \cdot d_2} + \sum_{j \geq d_1+d_2} \frac{1}{j \cdot d_2 \cdot 2^{j-d_1-d_2+1}} \\ &= H_{d_1} + \sum_{j \geq d_1+1} \frac{\alpha_j}{j}, \end{aligned}$$

where

$$\alpha_j := \begin{cases} \frac{1}{d_2} & \text{for } d_1 + 1 \leq j \leq d_1 + d_2 - 1; \\ \frac{1}{j \cdot 2^{j-d_1-d_2+1}} & \text{for } j \geq d_1 + d_2. \end{cases}$$

We observe the following simple facts about the coefficients α_j .

CLAIM 5. One has:

- (1) $\sum_{j \geq d_1+1} \alpha_j = 1$.
- (2) For every $i > 1$, $\sum_{j \geq d_1+i} \alpha_j \geq \frac{1}{2^{i-1}}$.

PROOF. For (1) one has

$$\sum_{j \geq d_1+1} \alpha_j = \frac{d_2 - 1}{d_2} + \sum_{j=d_1+d_2}^{\infty} \frac{1}{d_2 \cdot 2^{j-d_1-d_2+1}} = 1 - \frac{1}{d_2} + \frac{1}{d_2}.$$

Let us prove (2). For $i \geq d_2$, one has

$$\sum_{j \geq d_1+i} \alpha_j = \sum_{j=d_1+i}^{\infty} \frac{1}{d_2 \cdot 2^{j-d_1+d_2-1}} = \frac{1}{d_2 \cdot 2^{i-d_2}} \geq \frac{1}{2^{i-1}},$$

where in the inequality we used the fact that $k \leq 2^{k-1}$ for any integer $k \geq 1$.

For $2 \leq i \leq d_2 - 1$, one has:

$$\sum_{j \geq d_1+i} \alpha_j = \frac{d_2 - i}{d_2} + \frac{1}{d_2} = \frac{d_2 - i + 1}{d_2} \geq \frac{1}{i} \geq \frac{1}{2^{i-1}},$$

where in the first inequality above we used the fact that $\frac{k-j+1}{k}$ is a decreasing function of $k \geq j + 1$ and $d_2 \geq i + 1$, and in the second inequality again the fact that $k \leq 2^{k-1}$ for $k \geq 1$. \square

Intuitively, the term $A = \sum_{j=d_1+1}^{\infty} \frac{\alpha_j}{j}$ is a convex combination of terms of type $1/j$ under the constraint that the sum of the tail coefficients is large enough. An obvious upper bound on A is obtained by choosing coefficients β_j that respect the constraints on α_j given by Claim 5, and at the same time are as large as possible on the smallest terms of the sum. An easy induction shows that the best choice is $\beta_j = \frac{1}{2^{j-d_1}}$ for all $j \geq d_1 + 1$. Thus we can conclude

$$\begin{aligned} f(S) &\leq H_{d_1} + \sum_{j \geq d_1+1} \frac{\alpha_j}{j} \leq H_{d_1} + \sum_{j \geq d_1+1} \frac{\beta_j}{j} \\ &= H_{d_1} + \sum_{j=d_1+1}^{\infty} \frac{1}{j \cdot 2^{j-d_1}} = \hat{H}_{d_1}, \end{aligned}$$

where last equality comes from Claim 2. \square

PROOF OF CLAIM 1. Consider $a_1(i)$. Excluding a fixed additive term $\hat{H}_2 - p$, the value of this function is $a'_1(i) := \frac{H_{i+2}-x}{i}$, where $x = \hat{H}_2 - p \in (0, \hat{H}_2]$. Taking the discrete derivative

$$a'_1(i+1) - a'_1(i) = \frac{x + \frac{i-1}{i+3} - H_{i+3}}{i(i+1)}$$

one might observe that this is negative for $i \geq 6$ since $x + \frac{i-1}{i+3} \leq \hat{H}_2 + 1 < 2.7726 < H_9 > 2.8289$. The reader might skip the following cases that are analogous.

Consider now $a_2(i)$. Excluding a fixed additive term \hat{H}_2 , the value of this function is $a'_2(i) := \frac{H_i-x}{i}$, where $x = \hat{H}_2 - p \in (0, \hat{H}_2]$. One has

$$a'_2(i+1) - a'_2(i) = \frac{x+1-H_{i+1}}{i(i+1)},$$

which is negative for $i \geq 8$ since $x+1 \leq \hat{H}_2 + 1 < 2.7726 < H_9 > 2.8289$.

Consider next $a_3(i)$. Excluding a fixed additive term $\hat{H}_2 - p$, the value of this function is $a'_3(i) := \frac{\hat{H}_{i+2}-\hat{H}_2}{i}$. One has

$$\begin{aligned} a'_3(i+1) - a'_3(i) &= \frac{\hat{H}_2 - \hat{H}_{i+2}}{i(i+1)} + \sum_{j \geq 1} \frac{1}{2^j(i+1)(i+j+2)} \\ &\leq \frac{\hat{H}_2 + 1 - \hat{H}_{i+2}}{i(i+1)}, \end{aligned}$$

which is negative for $i \geq 6$ since $\hat{H}_2 + 1 < 2.7726 < \hat{H}_8 > 2.8194$.

It remains to consider $a_4(i)$. Excluding a fixed additive term \hat{H}_2 , the value of this function is $a'_4(i) := \frac{H_i-\hat{H}_2}{i}$. One has

$$\begin{aligned} a'_4(i+1) - a'_4(i) &= \frac{\hat{H}_2 - \hat{H}_i}{i(i+1)} + \sum_{j \geq 1} \frac{1}{2^j(i+1)(i+j)} \leq \frac{\hat{H}_2 + 1 - \hat{H}_i}{i(i+1)}, \end{aligned}$$

which is negative for $i \geq 8$ since $\hat{H}_2 + 1 < 2.7726 < \hat{H}_8 > 2.8194$. \square

B DETAILS ON THE REDUCTION TO STEINER TREE

PROOF OF LEMMA 2.1. \Leftarrow Assume by contradiction that A is not a feasible CacAP solution. Then there exists a 2-edge cut $\{e_1, e_2\}$, for two edges e_1, e_2 belonging to some cycle C of G , which is not covered by any link in A . Let $G_L = (V_L, E_L)$ and $G_R = (V_R, E_R)$ be the two (vertex disjoint) connected components identified by this

cut. Let also t_L and t_R be any two degree 2 nodes in V_L and V_R , respectively. (Observe that these nodes must exist.) By assumption there exists a (simple) path $P = t_L, \ell_1, \dots, \ell_q, t_R$ between t_L and t_R in $G_{ST}[T \cup A]$, where all ℓ_i 's are link nodes. Since $\{e_1, e_2\}$ is not covered, each such link has both endpoints either in V_L or in V_R . Furthermore, ℓ_1 and ℓ_q have one endpoint in V_L and V_R , resp. Hence there must be two consecutive links ℓ_i and ℓ_{i+1} where ℓ_i has both endpoints in V_L and ℓ_{i+1} both endpoints in V_R . These links cannot be crossing, therefore contradicting the fact that $\{\ell_i, \ell_{i+1}\}$ is an edge of G_{ST} .

\Rightarrow We first observe that, w.l.o.g., we can replace each link ℓ with its projections $proj(\ell)$. The feasibility of A is preserved. The same holds for the connected components of $G_{ST}[T \cup A]$ since the links in $proj(\ell)$ induce a path in G_{ST} . Thus for simplicity we assume that all links in A have both their endpoints in the same cycle. Let C_1, \dots, C_k be the cycles of G . For any cycle C_i of the cactus G let A_i be the set of links in A with both their endpoints in C_i . The following lemma shows that $G_{ST}[A_i]$ is connected.

LEMMA B.1. *Let $G = (V, E)$ be an input cactus of CacAP which consists of exactly one cycle and let A be a feasible solution for G . Then $G_{ST}[A]$ is connected.*

PROOF. Assume that $G_{ST}[A]$ is not connected. Then A can be partitioned in L_R and L_B , such that for any $l_R \in L_R$ and $l_B \in L_B$, l_R does not cross l_B . We call the links in L_R red links and the links in L_B blue links. We can also partition V in V_R and V_B , such that the endpoints of red links belong to V_R and the endpoints of blue links belongs to V_B . Therefore we call V_B and V_R , blue vertices and red vertices respectively.

Let V_1, V_2, \dots, V_{2k} be the partition of vertices of the cycle G into maximal consecutive blocks of vertices of the same color, so that $V_1 \cup V_3 \cup \dots \cup V_{2k-1} = V_R$ and $V_2 \cup V_4 \cup \dots \cup V_{2k} = V_B$.

We say that a link $\ell = \{u, w\} \in A$ is *nice*, if u and v belong to different blocks V_i and V_j , $i \neq j$. We say that an edge $e = \{u, v\} \in E$ is *colorful* if u is red and v is blue or vice versa. Note that G has precisely $2k$ colorful edges. If there is no nice link in A , then any pair of colorful edges of G is not covered by A , which is a contradiction.

Assume that $\ell = \{u, v\} \in A$ is a nice link, such that the distance between u and v in the cycle G is minimum. Assume that $u \in V_1$ and $v \in V_{2x+1}$ (and therefore these are red vertices) and also that the vertices of V_2 are in the shortest path from u to v . Now let e_1 and e_2 be the colorful edges such that exactly one of their endpoints is in V_2 . Now we show that the cut formed by e_1 and e_2 is not covered by A .

Assume that $\{e_1, e_2\}$ is covered, then there should be a link $\ell_1 = (w, z)$ such that $w \in V_2$ and $z \notin V_2$. Then either this link is a nice link that crosses ℓ , which is a contradiction since $\ell \in L_R$ and $\ell_1 \in L_B$, or ℓ_1 is a nice link such that the distance of w and z is less than the distance of u and v . \square

For every pair of cycles C_i and C_j that share a vertex v , there is a link $\ell_i \in A_i$ and $\ell_j \in A_j$ which are incident to v , thus ℓ_i and ℓ_j cross. We can conclude that $G_{ST}[A]$ is connected. Finally, since A is feasible, there exists at least one link $\ell \in A$ incident to each node t of degree 2 in G , which implies that the edge $\{\ell, t\}$ belongs to E_{ST} . Thus $G_{ST}[T \cup A]$ is also connected.

□

C SOME DETAILS ABOUT THE STEINER TREE APPROXIMATION ALGORITHM IN [3]

We will briefly discuss the $\ln(4) + \epsilon$ approximation algorithm from [3] for the Steiner tree problem. For a complete presentation of the Steiner tree algorithm we refer to the original paper [3]. The algorithm is based on the Directed Component Relaxation (DCR) of the Steiner tree problem.

$$\min \sum_{C \in \mathcal{C}} c(C)x_C \quad (\text{DCR}) \quad (5)$$

$$\text{s.t.} \quad \sum_{C \in \delta_C^+(U)} x_C \geq 1 \quad \forall \emptyset \neq U \subseteq T \setminus \{r\} \quad (6)$$

$$x_C \geq 0 \quad \forall C \in \mathcal{C}. \quad (7)$$

Here \mathcal{C} is a set of directed components, where each directed component C is a minimum-cost Steiner tree (of cost $c(C)$) over a subset of terminals. Furthermore, the leaves of C are precisely its terminals, and C is directed towards a specific terminal: the latter node is the sink of C , and the remaining terminals are the sources of C . Intuitively, our goal is to buy a minimum-cost subset of directed components so that they induce a directed path from each terminal to the root. In more detail, for any cut U that separates some non-root terminal from the root, let $\delta_C^+(U)$ be the set of components with some source in U and the sink not in U . Then every feasible solution has to buy some component in $\delta_C^+(U)$. The DCR relaxation follows naturally.

After restricting DCR to solutions that only use components with at most k terminals we obtain DCR_k . For constant k , DCR_k has a polynomial number of variables. Furthermore, the separation problem can be solved in polynomial time via a reduction to minimum cut. Therefore DCR_k can be solved in polynomial time. Moreover, the value of DCR_k is known to be a $(1 + \epsilon)$ -approximation of the value of DCR for large enough $k = O_\epsilon(1)$.

The iterative randomised rounding algorithm from [3], until all terminals are connected to the root, in iterations $t = 1, 2, 3 \dots$, does the following:

- solve DCR_k for the current instance of the Steiner tree problem to get x^t ;
- sample a component C^t from \mathcal{C}_k with probability proportional to x_C^t ;
- contract the sampled component C^t .

For the ease of the analysis, by adding dummy components w.l.o.g. one may assume that the total number of components in the fractional solution remains constant across the iterations of the algorithm, i.e., $\sum_{C \in \mathcal{C}} x_C^t = M$ for a proper M for all $t = 1, 2, \dots$. It is argued that after t iterations of the algorithm, having bought the first t sampled components, the residual instance of the problem is expected to be less costly. To this end a reference solution S^t is constructed such that $S^t \cup \bigcup_{t'=1}^{t-1} C^{t'}$ connects all the terminals. The initial reference solution $S^1 = \text{OPT}_{ST}$ is an optimal solution to the Steiner tree instance of cost opt . Consecutive reference solutions S^2, S^3, \dots are obtained by gradually deleting edges that are no

longer necessary due to the connectivity provided by the already sampled components.

Key to estimate the expected cost of the final solution is to bound the number of iterations until a particular edge $e \in S^1$ can be removed. Define $D(e) = \max\{t | e \in S^t\}$. In [3] (proof of Theorem 21) it is shown that there exist a randomised process of constructing reference solutions S^1, S^2, \dots such that $E[D(e)] \leq \ln(4) \cdot M$, which allows one to bound the total expected cost of sampled components as $E\left[\sum_{t \geq 1} c(C^t)\right] \leq (\ln(4) + \epsilon) \cdot \text{opt}$. Note that the above *per-edge* guaranty allows for easily handling arbitrary costs of individual edges. In our application to (unweighted) CacAP, we need to average over multiple edges to achieve a good enough bound.

C.1 Witness Tree and Witness Sets

We next slightly abuse notation and sometimes denote in the same way a tree and its set of edges. The construction of reference solutions S^1, S^2, \dots is not trivial. It involves:

- construction of a terminal spanning tree W , called the *witness tree*, based on randomised marking (selection) of a subset of edges of S^1 . Each edge e of S^1 is associated with a proper subset $W(e) \subseteq W$ (witness set of e);
- randomised deletion of a proper subset of W in response to selecting a particular component C^t in iteration t ;
- removing an edge e from S^t when all edges $W(e)$ have already been deleted.

In the following we discuss the main idea behind and the key properties of each of the three above mentioned processes. We also pin-point the element of the analysis that can be modified in order to utilise the specific properties of the instance we obtain from the reduction from CacAP.

Construction of the witness tree. The high level idea behind the witness tree is that we need to always satisfy the condition that $S^t \cup \bigcup_{t'=1}^{t-1} C^{t'}$ connects all the terminals, which is that the remaining fragments of the initial reference solution S^1 together with the already sampled components must provide sufficient connectivity. To this end a simpler object providing connectivity is constructed. It is an auxiliary tree W whose node set is the terminals of the instance (edges of W are independent of the edges of the input graph). It will be easier to delete edges from W in response to sampling components rather than deleting them directly from S^t .

We will now discuss methods to construct W . Intuitively, removing edges from a Steiner tree (in response to receiving connectivity from a component) is directly possible for only a subset of edges of the Steiner tree. In particular it appears more difficult to remove a Steiner vertex (and hence a path connecting a Steiner vertex to a terminal). This is related to the concept of *Loss* and *Loss contracting algorithms* (see, e.g., [26]), where one accepts that the cost of the system of paths connecting Steiner nodes to terminals is not removable.

Consider the following procedure: For each component⁶ S' of the Steiner tree S^1 select a single Steiner vertex s and draw the component as a tree rooted in s . For every Steiner vertex of S' select and mark a single edge going down (away from s). Note that

⁶Recall that a full component is a maximal subtree whose terminals are exactly its leaves.

for each Steiner vertex v the marked edges will form a unique path towards a leaf containing terminal $t(v)$. Note also that connected components formed by the marked edges will all have a single terminal node. Construct $W(S')$ by adding to $E(W(S'))$ an edge $\{t(u), t(v)\}$ for each unmarked edge $\{u, v\}$ of S' .⁷ Observe that the above constructed graph $W(S')$ is a tree spanning the terminals of S' . By repeating this procedure for all full components of S^1 we obtain tree W spanning all terminals of the Steiner tree instance.

So far we did not specify how to select the edge below Steiner node $v \in S'$ to be marked. In [3] the tree was assumed to be binary, and the edge would be selected at random by tossing a fair coin. In the current paper we use a different marking strategy as discussed in Section 3.1.

Marking edges of the witness tree. When edges of the witness tree W become unnecessary, we mark them. We keep the invariant that the unmarked edges of W together with the already collected components are sufficient to connect all terminals. Still, given a fixed collection of the already sampled components, the choice of which edges of W to mark is not obvious. In [3] a randomised marking scheme was considered. It was shown (Lemma 19 in [3]) that there exists a random process marking edges in W in response to sampled components, such that for every edge $e \in W$ not marked until iteration t the probability that it is marked in iteration t is at least $1/M$. In the current work we continue using the mentioned “uniform” witness tree marking process, and utilise the following lemma.

LEMMA C.1 (LEMMA 20 IN [3]). *Let $\tilde{W} \subseteq W$. Then the expected number of iterations until all edges in \tilde{W} are marked is at most $H_{|\tilde{W}|} \cdot M$.*

Removing edges of the reference tree S^t . Which edges of the reference tree can be removed? Clearly it suffices if S^t provides the same terminal connectivity as the unmarked edges of the witness tree W . Note that a single edge $e \in W$ corresponds to a single path $p(e)$ in S^1 . It then suffices to keep the edges of S^1 that occur in a path $p(e)$ of at least one unmarked edge $e \in W$.

We introduce the following notation: for an edge f in S^1 let $W(f) = \{e \in W \mid f \in p(e)\}$, we call $W(f)$ to be the *witness set* of f . Therefore, at iteration t , the reference solution S^t contains the edges from S^1 whose witness sets are not fully marked until iteration $t - 1$.

Observe that the expected number of iterations an edge f from the reference solution survives (until being removed) $E[D(f)]$ can be expressed using only the size of its witness set $W(f)$.

COROLLARY C.2. *Let $f \in S^1$, then $E[D(f)] \leq H_{|W(f)|} \cdot M$.*

Following the argument from the proof of Theorem 21 in [3], we also get

COROLLARY C.3. *For $k = O_\epsilon(1)$ large enough, the total cost of components bought by the algorithm is at most*

$$\frac{1 + \epsilon}{M} \sum_{f \in S^1} E[D(f)] \cdot c(f) \leq (1 + \epsilon) \cdot \sum_{f \in S^1} H_{|W(f)|} \cdot c(f)$$

⁷Note that in [3] the role of marked and unmarked edges was reversed. It was irrelevant for the analysis in [3] as it was assumed that the tree S' is binary. In this paper however we will exploit the high degree of Steiner nodes in S' and hence prefer to mark the “Loss” edges.

Therefore, it suffices to analyse how the marking scheme used in the construction of the witness tree affects distributions of the sizes of the witness sets for the individual edges of S^1 . To this end we will exploit two properties of our instances: the high degree of the Steiner vertices in the optimal solutions, and the fact that all edges of S^1 have the same cost.

ACKNOWLEDGMENTS

This work is highly in debt with Saket Saurabh. During a visit of the second author to Bergen University a few years ago, Saket mentioned the possibility to use the reduction to Steiner tree to approximate connectivity augmentation problems, possibly with an ad-hoc analysis. The result in this paper follows precisely that path (though finding a good enough analysis was not easy). The second author is also grateful to M. S. Ramanujan and L. Vegh for several helpful discussions on this topic.

REFERENCES

- [1] David Adjiashvili. 2017. Beating Approximation Factor Two for Weighted Tree Augmentation with Bounded Costs. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, January 16-19*. SIAM, Barcelona, Spain, Hotel Porta Fira, 2384–2399. DOI: <http://dx.doi.org/10.1137/1.9781611974782.157>
- [2] Manu Basavaraju, Fedor V. Fomin, Petr A. Golovach, Pranabendu Misra, M. S. Ramanujan, and Saket Saurabh. 2014. Parameterized Algorithms to Preserve Connectivity. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, July 8-11, 2014, Proceedings, Part I*. Springer, Copenhagen, Denmark, 800–811. DOI: http://dx.doi.org/10.1007/978-3-662-43948-7_66
- [3] Jaroslav Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. 2013. Steiner Tree Approximation via Iterative Randomized Rounding. *J. ACM* 60, 1 (2013), 6:1–6:33.
- [4] Joseph Cheriyan and Zhihan Gao. 2018. Approximating (Unweighted) Tree Augmentation via Lift-and-Project, Part I: Stemless TAP. *Algorithmica* 80, 2 (2018), 530–559. DOI: <http://dx.doi.org/10.1007/s00453-016-0270-4>
- [5] Joseph Cheriyan and Zhihan Gao. 2018. Approximating (Unweighted) Tree Augmentation via Lift-and-Project, Part II. *Algorithmica* 80, 2 (2018), 608–651. DOI: <http://dx.doi.org/10.1007/s00453-017-0275-7>
- [6] Joseph Cheriyan, Tibor Jordán, and R. Ravi. 1999. On 2-Coverings and 2-Packings of Laminar Families. In *Algorithms - ESA '99, 7th Annual European Symposium, July 16-18, 1999, Proceedings*. Springer, Prague, Czech Republic, 510–520. DOI: http://dx.doi.org/10.1007/3-540-48481-7_44
- [7] Joseph Cheriyan and Ramakrishna Thurimella. 2000. Approximating Minimum-Size k -Connected Spanning Subgraphs via Matching. *SIAM J. Comput.* 30, 2 (2000), 528–560. DOI: <http://dx.doi.org/10.1137/S009753979833920X>
- [8] Nachshon Cohen and Zeev Nutov. 2013. A $(1+\ln 2)$ -approximation algorithm for minimum-cost 2-edge-connectivity augmentation of trees with constant radius. *Theor. Comput. Sci.* 489-490 (2013), 67–74. DOI: <http://dx.doi.org/10.1016/j.tcs.2013.04.004>
- [9] E. A. Dinits, A. V. Karzanov, and M. V. Lomonosov. 1976. On the structure of a family of minimal weighted cuts in a graph. *Studies in Discrete Optimization* (1976), 290–306.
- [10] G. Even, J. Feldman, G. Kortsarz, and Z. Nutov. 2009. A 1.8 Approximation Algorithm for Augmenting Edge-connectivity of a Graph from 1 to 2. *ACM Transactions on Algorithms* 5, 2 (2009), 21:1–21:17.
- [11] Samuel Fiorini, Martin Groß, Jochen Köneemann, and Laura Sanità. 2018. Approximating Weighted Tree Augmentation via Chvátal-Gomory Cuts. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, January 7-10, 2018, Artur Czumaj (Ed.)*. SIAM, New Orleans, LA, USA, 817–831. DOI: <http://dx.doi.org/10.1137/1.9781611975031.53>
- [12] G. N. Frederickson and J. JáJá. 1981. Approximation Algorithms for Several Graph Augmentation Problems. *SIAM J. Comput.* 10, 2 (1981), 270–283.
- [13] Harold N. Gabow and Suzanne Gallagher. 2012. Iterated Rounding Algorithms for the Smallest k -Edge Connected Spanning Subgraph. *SIAM J. Comput.* 41, 1 (2012), 61–103. DOI: <http://dx.doi.org/10.1137/080732572>
- [14] Waldo Gálvez, Fabrizio Grandoni, Afrouz Jabal Ameli, and Krzysztof Sornat. 2019. On the Cycle Augmentation Problem: Hardness and Approximation Algorithms. In *Approximation and Online Algorithms - 17th International Workshop, WAOA 2019, September 12-13, 2019, Revised Selected Papers*. Springer, Munich, Germany, 138–153. DOI: http://dx.doi.org/10.1007/978-3-030-39479-0_10

- [15] Michel X. Goemans, Andrew V. Goldberg, Serge A. Plotkin, David B. Shmoys, Éva Tardos, and David P. Williamson. 1994. Improved Approximation Algorithms for Network Design Problems. In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms. 23-25 January 1994*. ACM/SIAM, Arlington, Virginia, USA, 223–232. <http://dl.acm.org/citation.cfm?id=314464.314497>
- [16] Fabrizio Grandoni, Christos Kalaitzis, and Rico Zenklusen. 2018. Improved approximation for tree augmentation: saving by rewiring. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, June 25-29, 2018*. ACM, Los Angeles, CA, USA, 632–645. DOI : <http://dx.doi.org/10.1145/3188745.3188898>
- [17] Christoph Hunkenschroder, Santosh S. Vempala, and Adrian Vetta. 2019. A $4/3$ -Approximation Algorithm for the Minimum 2-Edge Connected Subgraph Problem. *ACM Trans. Algorithms* 15, 4 (2019), 55:1–55:28. DOI : <http://dx.doi.org/10.1145/3341599>
- [18] Kamal Jain. 2001. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica* 21, 1 (2001), 39–60.
- [19] S. Khuller and R. Thurimella. 1993. Approximation algorithms for graph augmentation. *Journal of Algorithms* 14, 2 (1993), 214–225.
- [20] Philip Klein and R Ravi. 1995. A nearly best-possible approximation algorithm for node-weighted Steiner trees. *Journal of Algorithms* 19, 1 (1995), 104–115.
- [21] Guy Kortsarz and Zeev Nutov. 2016. LP-Relaxations for Tree Augmentation. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2016, September 7-9, 2016*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Paris, France, 13:1–13:16. DOI : <http://dx.doi.org/10.4230/LIPICs.APPROX-RANDOM.2016.13>
- [22] G. Kortsarz and Z. Nutov. 2016. A Simplified 1.5-Approximation Algorithm for Augmenting Edge-Connectivity of a Graph from 1 to 2. *ACM Transactions on Algorithms (TALG)* 12, 2 (2016), 23:1–23:20.
- [23] Dániel Marx and László A. Végh. 2015. Fixed-Parameter Algorithms for Minimum-Cost Edge-Connectivity Augmentation. *ACM Trans. Algorithms* 11, 4 (2015), 27:1–27:24. DOI : <http://dx.doi.org/10.1145/2700210>
- [24] H. Nagamochi. 2003. An Approximation for Finding a Smallest 2-Edge-Connected Subgraph Containing a Specified Spanning Tree. *Discrete Applied Mathematics* 126, 1 (2003), 83–113.
- [25] Zeev Nutov. 2017. On the Tree Augmentation Problem. In *25th Annual European Symposium on Algorithms, ESA 2017, September 4-6, 2017*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Vienna, Austria, 61:1–61:14. DOI : <http://dx.doi.org/10.4230/LIPICs.ESA.2017.61>
- [26] Gabriel Robins and Alexander Zelikovsky. 2005. Tighter bounds for graph Steiner tree approximation. *SIAM Journal on Discrete Mathematics* 19, 1 (2005), 122–134.
- [27] András Sebő and Jens Vygen. 2014. Shorter Tours by Nicer Ears: $7/5$ -approximation for graphic TSP, $3/2$ for the path version, and $4/3$ for two-edge-connected subgraphs. *Combinatorica* 34(5) (2014), 597–629. arXiv:1201.1870 <http://arxiv.org/abs/1201.1870>