

Breaching the 2-Approximation Barrier for Connectivity Augmentation

Citation for published version (APA):

Byrka, J., Grandoni, F., & Jabal Ameli, A. (2023). Breaching the 2-Approximation Barrier for Connectivity Augmentation: A Reduction to Steiner Tree. *SIAM Journal on Computing*, 52(3), 718-739.
<https://doi.org/10.1137/21M1421143>

DOI:

[10.1137/21M1421143](https://doi.org/10.1137/21M1421143)

Document status and date:

Published: 01/06/2023

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

BREACHING THE 2-APPROXIMATION BARRIER FOR CONNECTIVITY AUGMENTATION: A REDUCTION TO STEINER TREE*

JAROSLAW BYRKA[†], FABRIZIO GRANDONI[‡], AND AFROUZ JABAL AMELI[§]

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 [G. N. Frederickson and Jájá, *SIAM J. Comput.*, 10 (1981), pp. 270–283]. It is known [E. A. Dinitz, A. V. Karzanov, and M. V. Lomonosov, *Studies in Discrete Optimization*, Nauka, Moscow, 1976, pp. 290–306] that CAP can be reduced to the case $k = 1$, also known as the *tree augmentation* problem (TAP) for odd k , and to the case $k = 2$, also known as the *cactus augmentation* problem (CacAP) for even k . Prior to the conference version of this paper [J. Byrka, F. Grandoni, and A. Jabal Ameli, *STOC'20*, ACM, New York, 2020, pp. 815–825], several better than 2 approximation algorithms were known for TAP, culminating with a recent 1.458 approximation [F. Grandoni, C. Kalaitzis, and R. Zenklusen, *STOC'18*, ACM, New York, 1918, pp. 632–645]. However, for CacAP the best known approximation was 2. In this paper we breach the 2 approximation barrier for CacAP, hence, for CAP, by presenting a polynomial-time $2\ln(4) - \frac{967}{1120} + \varepsilon < 1.91$ approximation. From a technical point of view, our approach deviates quite substantially from previous work. In particular, the better-than-2 approximation algorithms for TAP either exploit greedy-style algorithms or are based on rounding carefully designed LPs. We instead use a reduction to the Steiner tree problem which was previously used in parameterized algorithms [Basavaraju et al., *ICALP'14*, Springer, Berlin, 2014, pp. 800–811]. This reduction is not approximation preserving, and using the current best approximation factor for a Steiner tree [Byrka et al., *J. ACM*, 60 (2013), 6] 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 a Steiner tree arising from CacAP. In our opinion this connection between approximation algorithms for survivable network design and Steiner-type problems is interesting, and might lead to other results in the area.

Key words. approximation algorithms, connectivity augmentation, Steiner tree

MSC codes. 68W25, 05C85, 05C40

DOI. 10.1137/21M1421143

1. Introduction. The basic goal of *survivable network design* is to construct cheap networks that provide connectivity guarantees between prespecified 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.

*Received by the editors May 25, 2021; accepted for publication (in revised form) December 28, 2022; published electronically May 23, 2023.

<https://doi.org/10.1137/21M1421143>

Funding: The first author was supported by the NCN grant 2015/18/E/ST6/00456. The second and third authors were partially supported by SNF Excellence Grant 200020B_182865/1.

[†]University of Wrocław, Wrocław, Poland (jby@cs.uni.wroc.pl).

[‡]IDSIA, USI-SUPSI, Lugano-Viganello, Switzerland (fabrizio@idsia.ch).

[§]Eindhoven University of Technology, Eindhoven, Netherlands (jabal.farsh@gmail.com).

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, Karzanov, and Lomonosov [12] (see also [9, 24]) 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-edge-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, 7, 8, 13, 14, 19, 26, 27, 29]. Among these results the best approximation factor prior to the conference version of this paper [4] was 1.458 due to Grandoni, Kalaitzis, and Zenklusen [19] (see section 1.3 for subsequent developments).

The case $k = 2$ is also known as the *cactus augmentation problem* (CacAP), where for similar reasons we can assume that the input graph is a cactus.² Here the best-known approximation factor previous to our work [4] was 2, and this factor was achieved with multiple approaches [15, 17, 23, 24]. A better approximation was achieved recently for the special case where the input cactus is a cycle [20].

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.*

From Theorem 1.1 and the reduction to CacAP implied by [12], we get the following.

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.*

Our result is based on a reduction to the (cardinality) Steiner tree problem by Basavaraju et al. [3]. The authors use this connection to design improved parameterized algorithms (see also [28] 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* (disjoint from T). 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. [3] observed that, given a CacAP instance $(G = (V, E), L)$, it is possible to construct (in polynomial time) an *equivalent* Steiner tree instance $G_{ST} = (T \cup L, E_{ST})$ (see also the description in [30]). Here T corresponds to the nodes of degree 2 in G , L is the set of Steiner nodes, and the edges E_{ST} are defined properly (more details in section 2.1). Intuitively, each link node $\ell \in L$ is adjacent to the terminal nodes in T which are endpoints of ℓ (if any), and two link nodes $\ell, \ell' \in L$

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.

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.

are adjacent iff the respective endpoints cannot be separated by a min-cut of G . 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 2.1. 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 a 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 et al. [5]. Hence this is not good enough³ to obtain $\rho < 2$.

In order to obtain our main result we use the same algorithm as in [5], but we analyze it differently. In particular, we exploit the specific structure of the instances of a Steiner tree arising from CacAP instances via the above reduction to get a substantially better approximation factor.

In more detail (see also section 3), in the analysis of the algorithm in [5] one considers an optimal Steiner tree solution $OPT_{ST} = (T \cup A, F)$ rooted at some arbitrary node r , marking a random subset $F_{mar} \subseteq F$ of edges so that each Steiner node is connected to some 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 [5] for the following critical reasons:

1. They (i.e., the authors of [5]) 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 [13, 26, 27, 29]. More recent approaches are based on rounding stronger and stronger LP (or semidefinite programming) relaxations for the problem [1, 7, 8, 14, 19]. 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 and previous work. One can consider a natural weighted version $WCAP$ of CAP where each link has a positive weight and the goal is to minimize the total weight of selected links. The best-known approximation for WCAP is 2. 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 [25].

Prior to the conference version of this paper, 2 was the best-known approximation factor even for the weighted version WTAP of TAP. Some progress on weighted TAP was made in the case of small integer weights. In particular, when the largest weight

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 a Steiner tree.

W is upper bounded by a constant, better than 2 approximation algorithms are given in [1, 14, 19]. A technique in [32] allows one to extend these results to $W = O(\log n)$. WTAP also admits a $1 + \ln 2$ approximation for arbitrary weights if the input tree has constant radius [11]. A better than 2 approximation can also be achieved if the fractional solution to a natural LP relaxation has nonzero entries bounded away from zero [22].

A problem closely related to CAP is to build a minimum cost k -edge-connected spanning subgraph of a given input graph [10, 16, 21, 34]. Here the best known approximation factor is $4/3$ for the unweighted case, and 2 for the weighted one.

1.3. Subsequent work. After the publication of the conference version of this paper [4], there were a few breakthroughs in the area of survivable network design. Cecchetto, Traub, and Zenklusen [6] developed an elegant and unified framework to approximate both TAP and CacAP, leading to a 1.393-approximation. Notice that this does not only greatly improves on our approximation factor for CacAP, but also on the approximation factor for TAP in [19]. Though it leads to weaker approximation factors, we believe that our connection between survivable network design and Steiner-type problems might be useful in future related work. Indeed, after our work, a similar connection was made by Nutov [31]. In more detail, he proved that an α approximation for a Steiner tree implies a $1 + \ln(4 - x) + \varepsilon$ approximation for CAP, where x is the solution to $1 + \ln(4 - x) = \alpha + (\alpha - 1)x$. This leads to an approximation factor below 2 (though worse than the one achieved in this work) using the current best approximation for Steiner tree [5]. Notice that, differently from our result, Nutov's reduction is black box, a useful feature.

Recently, Traub and Zenklusen [35] achieved a $1 + \ln 2 + \varepsilon$ approximation for WTAP, hence solving a major open problem in the area. The authors later improved their result to a $1.5 + \varepsilon$ -approximation [36]. Achieving a better than 2 approximation for WCAP remains a challenging open problem.

Grandoni, Jabal Ameli, and Traub [18] obtained the first better than 2 approximation for the forest augmentation problem, i.e., the generalization of TAP where the input graph is a forest rather than a tree.

A variant of our approach was used by Nutov [30] to obtain, among other results, the first better than 2 approximation for the node-connectivity version of TAP. The approximation factor for this problem was later improved by Angelidakis, Hyatt-Denesik, and Sanità [2], again using an approach similar to the one in this paper.

2. Steiner tree and connectivity augmentation. In this section we present the mentioned reduction in [3] 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 3). We analyze the resulting approximation factor in section 4.

2.1. A reduction to Steiner tree. Consider a CacAP instance $(G = (V, E), L)$. For a link $\ell = (v_0, v_{q+1})$, let v_1, \dots, v_q be the sequence of nodes of degree at least 4 that lie along every simple v_0 - v_{q+1} path, excluding the endpoints (that may or may not have this property). Intuitively, any simple path from v_0 to v_{q+1} will contain edges from some cycles C_0, \dots, C_q in this order. Then v_i , $i \in \{1, \dots, q\}$, is the (cut) node shared by the cycles C_{i-1} and C_i . Notice that each pair $\ell_i = \{v_i, v_{i+1}\}$, $i = 0, \dots, q$, lies along a distinct cycle C_i visited by the mentioned path. We call each such ℓ_i the *projection* of ℓ on C_i , and we let $proj(\ell)$ be the set of all projections of ℓ . Consider two links $\ell = \{x, y\}$ and $\ell' = \{x', y'\}$. Then we say that ℓ and ℓ' *cross* if one of the following two conditions holds: (1) they share one endpoint or (2) there exists a cycle

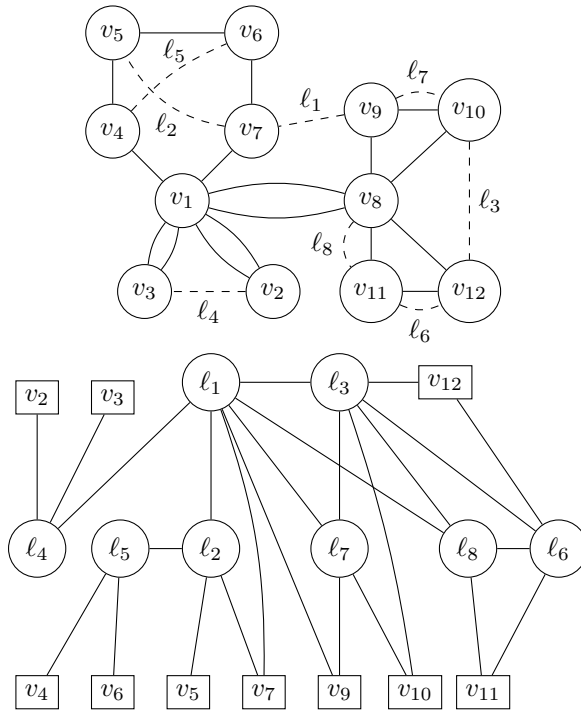


FIG. 2.1. (top) Instance of CacAP, where dashed edges denote links. The projections of l_1 are $\text{proj}(l_1) = \{\{v_7, v_1\}, \{v_1, v_8\}, \{v_8, v_9\}\}$. Link l_2 is crossing with l_1 and l_5 . (bottom) The corresponding Steiner tree instance, where square nodes denote terminals.

C such that C includes x, y, x' , and y' and taking one simple x - y path P along C , P contains exactly one node in $\{x', y'\}$ as an internal node. We say that any two links l and l' cross if there exist $l_i \in \text{proj}(l)$ and $l'_j \in \text{proj}(l')$ such that l_i and l'_j cross. See Figure 2.1 (top) for an example.

From (G, L) we construct a Steiner tree instance $G_{ST} = (T \cup S, E_{ST})$ as follows. For each one of the t nodes v of degree 2 in G , add a terminal v to T ; for each link $l \in L$, add a Steiner node l to S (i.e., $S = L$); for each $l \in L$ and endpoint $v \in T$ of l , add $\{l, v\}$ to E_{ST} ; finally, for any two links l and l' that cross, add $\{l, l'\}$ to E_{ST} . See Figure 2.1 (bottom) for an example. We observe the following simple facts.

PROPERTY 1. Each Steiner node s is adjacent to at most 2 terminals, namely, the terminals corresponding to the endpoints of the link associated with s .

PROPERTY 2. The neighbors of each terminal t are Steiner nodes which form a clique: this clique corresponds to the links that share the node corresponding to t as a common endpoint.

We will critically exploit the following lemma sketched in [3, Lemma 1]. For the sake of completeness we give a (more detailed) proof of it in Appendix B.

LEMMA 2.1 ([3]). $A \subseteq L$ is a feasible solution to a CacAP instance (G, L) iff, in the corresponding Steiner tree instance $G_{ST} = (T \cup L, E_{ST})$, $G_{ST}[T \cup A]$ is connected.

Notice that the above reduction is not approximation preserving. Still, we can state the following.

COROLLARY 2.2. *Any optimum solution OPT to the input CacAP instance, induces a solution OPT_{ST} of cost $|OPT_{ST}| = |OPT| + t - 1$ for the associated Steiner tree instance, where t is the number of terminals in the latter instance. Vice versa, given a solution APX_{ST} to the Steiner tree instance, one can construct in polynomial time a solution APX to the input CacAP instance with $|APX| = |APX_{ST}| - t + 1$.*

Proof. Both claims follow directly from Lemma 2.1. For the first claim, it is sufficient to observe that a spanning tree of $G_{ST}[T \cup OPT]$ contains $t + |OPT| - 1$ edges. For the second claim, observe that the Steiner nodes in APX_{ST} induce a feasible solution to CacAP. The claim follows since $|APX_{ST}| = s + t - 1$, where s is the number of Steiner nodes in APX_{ST} . \square

We will exploit also the following simple fact.

LEMMA 2.3. *Given some optimal solution OPT to a CacAP instance, there is a feasible solution OPT_{ST} to the associated Steiner tree instance with $|OPT_{ST}| = |OPT| + t - 1$, where terminals have degree exactly 1 (namely, all the terminals are leaves).*

Proof. Given any feasible solution ST to the problem, we can transform it into a solution ST' of the same cost, where some terminal v of degree $d(v) \geq 2$ in ST has degree $d(v) - 1$ in ST' . In order to do that, consider any terminal v adjacent to two Steiner nodes ℓ and ℓ' in ST . By Property 2, ℓ and ℓ' are adjacent. Hence $ST' := ST \cup \{\ell, \ell'\} \setminus \{v, \ell'\}$ is a feasible Steiner tree of the same cost and with the desired property.

By iteratively applying the above process to the solution OPT_{ST} guaranteed by Corollary 2.2 one obtains the desired solution. \square

As mentioned earlier, a ρ_{ST} approximation for a Steiner tree (used as a *black box*) provides a $3\rho_{ST} - 2$ approximation for CacAP by the above construction. Indeed, the Steiner tree instance has cost at most $|OPT| + t - 1$ by Corollary 2.2, hence an approximate solution APX_{ST} would cost at most $\rho_{ST}(|OPT| + t - 1)$. By the same corollary, we can convert this into a solution APX to CacAP of a cost at most $\rho_{ST}(|OPT| + t - 1) - t + 1$. Next observe that $|OPT| \geq t/2$. Indeed, any node of degree 2 in the CacAP instance needs to have at least one link incident to it in a feasible solution, and a link can be incident to at most 2 such nodes. Thus $|APX| \leq 3\rho_{ST}|OPT| - 2|OPT|$. In order to improve on this simple bound, we will have to *open the box*.

Let us remark that the size $|APX|$ of the approximate solution for CacAP that we will compute is precisely the number of Steiner nodes involved in the solution APX_{ST} that we compute for the corresponding Steiner tree instance. In our 1.91 approximation we implicitly address all the Steiner tree instances satisfying Properties 1 and 2. Therefore, we implicitly achieve the same approximation factor for the latter instances of a Steiner tree, where the objective function is to minimize the number of Steiner nodes in the computed Steiner tree.

3. Steiner tree via iterative randomized rounding. As we mentioned in the introduction, the current best $(\ln 4 + \varepsilon)$ -approximate Steiner tree algorithm from [5], used as a black box, is not good enough to break the 2-approximation barrier for CacAP. However, it turns out that the same algorithm achieves this goal in combination with a different analysis that exploits the properties of the specific Steiner tree instances arising from CacAP.

We first sketch the basic property of the algorithm and analysis in [5] that we need here and express it in the form of Lemma 3.1. For the sake of completeness, we include a more detailed description and a sketch of the proof of the lemma in the following section 3.1. The authors of [5] 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 (more details in sections 3.1 and 3.2) the authors of [5] consider any feasible Steiner tree $ST = (T \cup A, F)$. They interpret each full component⁴ S' of ST as a tree rooted at some Steiner node r of S' (if there is no such node, it can be created by splitting the single edge in S'). Then the authors define a *marking scheme* where some child edge of each internal (Steiner) node of S' is marked. Notice that the marked edges induce a collection of disjoint paths in each full component S' : such paths span the nodes of S' and each such path contains precisely one terminal (as an endpoint). A given marking scheme defines a witness set $W(e)$ for each edge e in S' : this consists of all the pairs of terminals $\{t', t''\}$ in S' such that the t' - t'' path in S' contains e and precisely one unmarked edge. We let $w(e) = |W(e)|$. Notice that for each unmarked edge e there exists exactly one such t' - t'' path, hence $w(e) = 1$ (we will later use this property in Lemma 4.1). Then the authors prove the following, where $H_i := 1 + \frac{1}{2} + \dots + \frac{1}{i}$ is the i th harmonic number.

LEMMA 3.1 ([5]). *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 $(1 + \varepsilon) \sum_{e \in F} E[H_{w(e)}]$.*

3.1. Some details about the Steiner tree approximation algorithm in [5]. For a complete presentation of the Steiner tree algorithm we refer to the original paper [5]. Here we sketch the main ideas. The algorithm is based on the following directed component relaxation (DCR) of the Steiner tree problem:

$$(3.1) \quad \min \sum_{C \in \mathcal{C}} c(C)x_C \quad (\text{DCR})$$

$$(3.2) \quad \text{s.t.} \quad \sum_{C \in \delta_{\mathcal{C}}^+(U)} x_C \geq 1 \quad \forall \emptyset \neq U \subseteq T \setminus \{r\},$$

$$(3.3) \quad x_C \geq 0 \quad \forall C \in \mathcal{C}.$$

Here \mathcal{C} is the 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 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 nonroot terminal from the root, let $\delta_{\mathcal{C}}^+(U)$ 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 $\delta_{\mathcal{C}}^+(U)$. The DCR relaxation follows naturally.

⁴Recall that a full component is a maximal subtree whose terminals are exactly its leaves.

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 randomized rounding algorithm from [5], 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 without loss of generality (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 estimating 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 $D(e) = \max\{t \mid e \in S^t\}$. In [5, proof of Theorem 21] it is shown that there exist a randomized 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 $E[\sum_{t \geq 1} c(C^t)] \leq (\ln(4) + \epsilon) \cdot \text{opt}$. 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.

3.2. 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 randomized 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);
- randomized 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 ideas behind our approach and the key properties of each of the three above-mentioned processes.

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 (while the edges of W are not necessarily 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 node (and hence a path connecting a Steiner node to a terminal). This is related to the concept of *loss* and *loss contracting algorithms* (see, e.g., [33]), 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 full component S' of the Steiner tree S^1 select a single Steiner node r and interpret S' as a tree rooted at r . For every Steiner node s of S' , mark one edge between s and one of its children. Note that for each Steiner node s the marked edges will form a unique path towards a leaf containing terminal $t(s)$. Note also that connected components formed by the marked edges will all contain 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' .⁵ 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 a tree W spanning all terminals of the Steiner tree instance.

So far we did not specify how to select the edge below the Steiner node $v \in S'$ to be marked. In [5] 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 4.1.

Removing edges of the witness tree. When edges of the witness tree W become unnecessary, we remove them. We keep the invariant so that the (not removed) 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 remove is not obvious. In [5] a randomized scheme was considered. It was shown (Lemma 19 in [5]) that there exists a random process removing edges from W in response to sampled components, such that for every edge $e \in W$ not removed before iteration t , the probability that it is removed in iteration t is at least $1/M$. In the current work we continue using the mentioned “uniform” witness tree edge removing process, and utilize the following lemma.

LEMMA 3.2 (Lemma 20 in [5]). *Let $\tilde{W} \subseteq W$. Then the expected number of iterations until all edges in \tilde{W} are removed 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 not removed 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 (still not removed) 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 $W(f)$ 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 removed until iteration $t - 1$.

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

⁵Note that in [5] the role of marked and unmarked edges was reversed. It was irrelevant for the analysis in [5] 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.

COROLLARY 3.3. *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 [5], we also get the following.

COROLLARY 3.4. *For $k = O_\varepsilon(1)$ large enough, the total expected cost of the components bought by the algorithm is at most*

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

Therefore, it suffices to analyze how the marking scheme used in the construction of the witness tree affects the 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 nodes in the initial optimal solution S^1 , and the fact that all the edges of S^1 have the same cost.

4. 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 [5] sketched in section 3. 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 4.1 we describe our alternative marking scheme and prove some of its properties. In section 4.2 we complete the analysis of the approximation factor.

4.1. An alternative marking scheme. Recall that in the analysis of the Steiner tree approximation algorithm in [5], 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). Let OPT be some optimal solution to the considered CacAP instance. As a 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. Consider each full component S' of OPT_{ST} . W.l.o.g, S' contains at least one Steiner node (otherwise, we can create it by splitting one edge). Let us root S' at some Steiner node r which is adjacent to at least one terminal (notice that such an r must exist). 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, respectively. In particular $d(\ell) = s(\ell) + t(\ell)$ and (by Property 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 a feasible marking scheme, namely, for each Steiner node we mark exactly one child edge. Observe also that in our marking we favor edges connecting Steiner nodes to terminals: this will be critical in our analysis.⁶ See Figure 4.1 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, respectively, that partition F . Recall that for each $e \in F$, there exists a (random) witness set $W(e)$ of size

⁶While we are able to show that our marking scheme leads to a better than 2 approximation, we are not able to show that the same cannot be achieved with the original marking scheme in [5].




FIG. 4.1. A feasible Steiner tree for the instance of Figure 2.1, 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 gray. 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 $g(\ell_5) = \{\ell_5\}$.

$w(e) = |W(e)|$. 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 4.1. $apx \leq 2\varepsilon + \frac{1+\varepsilon}{|OPT|} \sum_{\ell \in OPT} c(\ell)$.

Proof. Recall that by Lemma 3.1 the expected cost of the computed Steiner tree APX_{ST} is, modulo a factor $(1 + \varepsilon)$, at most

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

In the second to 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 $|APX| = |APX_{ST}| - t + 1$ by Corollary 2.2. Hence

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

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 nonroot Steiner node based on the degree sequence of its ancestors.⁷

LEMMA 4.2. *Given a nonroot Steiner node ℓ , let ℓ_q be the lowest proper ancestor⁸ 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 $d_i = d(\ell_i)$. Then⁹*

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

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 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 $\frac{1}{d_2} \times \dots \times \frac{1}{d_h} \times \frac{d_{h+1}-1}{d_{h+1}}$. For $h = q - 1$, this probability is $\frac{1}{d_2} \times \dots \times \frac{1}{d_h}$. 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}_1 is precisely the approximation factor for the Steiner tree achieved in [5]. 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 Appendix A.

LEMMA 4.3. *For any $\ell \in OPT$, $c(\ell) \leq \hat{H}_{d(\ell)}$.*

In the 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 4.3. This will be critical in our analysis since the latter bound is not strong enough.

4.2. Analysis of the approximation factor. In this section we upper bound the approximation factor *apx* as given by Lemmas 4.1 and 4.2. 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 4.1). The following lemma shows that this is (essentially) w.l.o.g.

DEFINITION 4.4. *A rooted Steiner tree is well-structured if, for every Steiner node ℓ , the following two conditions hold:*

⁷Observe that for the root r , $c(r) = H_{d(r)-1}$ deterministically.

⁸Observe that this ancestor exists since the root has this property by assumption.

⁹The value of a product of type $a_i \cdot a_{i+1} \cdots a_j$ for $j < i$ is assumed to be 1 by definition.

1. ℓ has at least 2 children and
2. ℓ has 0 or 2 terminal children.

LEMMA 4.5. 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 by Property 1 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 4.4.

Now consider any tree OPT_{ST} that satisfies (2), and let $oc(OPT_{ST})$ be the number of its Steiner nodes with precisely one child. We prove by induction on $oc(OPT_{ST})$ that $\rho(OPT_{ST}) \leq \max\{\hat{H}_1, \rho_{ws}\}$. The claim is trivially true for $oc(OPT_{ST}) = 0$ 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 $oc(OPT_{ST}) = q$. 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 4.3. Consider the tree $OPT'_{ST} = (T \cup OPT', F')$ obtained by contracting the edge (ℓ', ℓ'') . We observe that OPT'_{ST} satisfies (2), $oc(OPT'_{ST}) = q - 1$, 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

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

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 4.5 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 4.3 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 4.1).

DEFINITION 4.6. 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 Property 1), 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, respectively.

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 4.7. *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 4.8. *$apx \leq 2\varepsilon + \frac{1+\varepsilon}{|OPT|} \sum_{\ell \in OPT} c'(\ell)$, 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. The proof follows from Lemma 4.1, by replacing $c(\ell)$ as in Lemma 4.2 with the upper bounds given by Lemmas 4.3 and 4.7. \square

We rewrite the upper bound from Lemma 4.8 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 $\ell' \in OPT_{gf}$ 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,

$$(4.1) \quad \frac{1}{|OPT|} \sum_{\ell \in OPT} c'(\ell) = \frac{1}{|OPT|} \sum_{\ell \in OPT} c''(\ell),$$

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 (4.1), we partition OPT into groups of nodes as follows (see also Figure 4.1).

DEFINITION 4.9. *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, respectively.*

We associate with 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 $c''(\cdot)$:

$$(4.2) \quad \frac{1}{|OPT|} \sum_{\ell \in OPT} c''(\ell) \leq \max_{\ell \in OPT_{in} \cup \{\ell^*\}} \{a(\ell)\}.$$

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 4.8 with (4.1) and (4.2) one gets

$$(4.3) \quad apx \leq 2\varepsilon + (1 + \varepsilon) \max_{i \geq 1} \{\hat{H}_2, a_1(i), a_2(i), a_3(i), a_4(i)\}.$$

Notice that the above approximation factor is a function of the parameter $p \in [0, \hat{H}_2 - H_2]$ which still needs to be fixed. In order to choose a convenient p , we need the following technical result (proof in Appendix A).

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, respectively.*

From (4.3) and Claim 1, for any $p \in [0, \hat{H}_2 - H_2]$, one has

$$(4.4) \quad apx \leq 2\varepsilon + (1 + \varepsilon) \max\{\hat{H}_2, \max_{1 \leq i \leq 6} \{a_1(i)\}, \max_{1 \leq i \leq 8} \{a_2(i)\}, \max_{1 \leq i \leq 6} \{a_3(i)\}, \max_{1 \leq i \leq 8} \{a_4(i)\}\}.$$

Numerically the minimum of the right-hand side of (4.4) is achieved for $p \simeq 0.135$, 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.

Appendix A. Omitted proofs from section 4.

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

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

Proof of Lemma 4.3. 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 we next assume that ℓ is not the root. For a generic sequence $S = (d_1, \dots, d_k)$ of positive integers, let us define

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

By Lemma 4.2, in order to prove the claim it is sufficient to show that $f(S) \leq \hat{H}_{d_1}$. For an infinite sequence $S' = (d_1, d_2, \dots)$ of positive integers, we analogously define

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

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

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

Proof. By definition

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

By Claim 3 it is sufficient to consider infinite sequences of type \bar{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 $\bar{S} = (d_1, \dots, d_k, 2, 2, \dots)$ and assume there exists $d_i = 1$ in the sequence for some $i \geq 2$. Let $\bar{S}_i = (d_1, \dots, d_{i-1}, d_{i+1}, \dots, d_k, 2, 2, \dots)$ be the subsequence where the i th entry is removed. Then $f(\bar{S}) = f(\bar{S}_i)$.

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

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

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

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

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

where the last equality comes from Claim 2.

Summarizing, given a nonroot Steiner node ℓ and the associated values $S = (d_1, \dots, d_{q-1})$, $q \geq 2$, as defined in Lemma 4.2, we have that, for $\bar{S} = (d_1, \dots, d_{q-1}, 2, 2, \dots)$,

$$c(\ell) \stackrel{\text{Lem. 4.2}}{=} f(S) \stackrel{\text{Claim 3}}{\leq} f(\bar{S}) \leq \hat{H}_{d_1} = \hat{H}_{d(\ell)}. \quad \square$$

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 $x + \frac{i+1}{i+3} \leq \hat{H}_2 + 1 < 2.7726 < H_9 > 2.8289$.

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

Appendix B. Details on the reduction to a Steiner tree.

DEFINITION B.1. Let $G = (V, E)$ be a connected graph and let L be a set of extra edges. Let $E_1 \subseteq E$ be an edge-cut of G . We say that L covers the cut E_1 if E_1 is not an edge-cut of $G' = (V, E \cup L)$.

LEMMA B.2. 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 by contradiction that $G_{ST}[A]$ is not connected. Then A can be partitioned into 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 into V_R and V_B , such that the endpoints of red links belong to V_R and the endpoints of blue links belong 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 w 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 w.l.o.g. 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 . We next show that the edge-cut $\{e_1, e_2\}$ 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 (which contradicts the choice of ℓ). \square

Proof of Lemma 2.1. \Leftarrow Let $A \subseteq L$ be such that $G_{ST}[T \cup A]$ is connected. 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 (node disjoint) connected components of G identified by this cut. Let also $t_L \in V_L$ and $t_R \in V_R$ be any two nodes of degree 2 in G . (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 , respectively. 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 Let $A \subseteq L$ be a feasible CacAP solution. We will show that $G_{ST}[T \cup A]$ is connected. We first observe that, w.l.o.g., we can replace each link ℓ with its projections $proj(\ell)$. Indeed, the feasibility of A is preserved. Furthermore, the number of connected components of $G_{ST}[T \cup A]$ does not change since the links in $proj(\ell)$ induce a path in G_{ST} . With this modification, all links in A have both their endpoints in the same cycle (since projections have this property by definition). 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 . Lemma B.2 shows that $G_{ST}[A_i]$ is connected. For every pair of cycles C_i and C_j that share a node 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. \square

Acknowledgments. This work is highly in debt to Saket Saurabh. During a visit of the second author to Bergen University a few years ago, Saket mentioned the possibility of using the reduction to a Steiner tree to approximate connectivity augmentation problems, possibly with an ad hoc analysis. The result in this paper follows precisely that high-level path; however, fixing the details in the analysis was highly nontrivial. The second author is also grateful to M. S. Ramanujan and L. Vegh for several helpful discussions on this topic.

REFERENCES

- [1] D. ADJIASHVILI, *Beating approximation factor two for weighted tree augmentation with bounded costs*, ACM Trans. Algorithms, 15 (2018), 19, <https://doi.org/10.1145/3182395>.
- [2] H. ANGELIDAKIS, D. HYATT-DENESIK, AND L. SANITÀ, *Node connectivity augmentation via iterative randomized rounding*, Math. Program., to appear.
- [3] M. BASAVARAJU, F. V. FOMIN, P. A. GOLOVACH, P. MISRA, M. S. RAMANUJAN, AND S. SAURABH, *Parameterized algorithms to preserve connectivity*, in Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Proceedings, Part I, Copenhagen, Denmark, Springer, Berlin, 2014, pp. 800–811, <https://doi.org/10.1007/978-3-662-43948-7>.
- [4] J. BYRKA, F. GRANDONI, AND A. JABAL AMELI, *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 2020, Chicago, IL, K. Makarychev, Y. Makarychev, M. Tulsiani, G. Kamath, and J. Chuzhoy, eds., ACM, New York, 2020, pp. 815–825.
- [5] J. BYRKA, F. GRANDONI, T. ROTHVOSS, AND L. SANITÀ, *Steiner tree approximation via iterative randomized rounding*, J. ACM, 60 (2013), 6.
- [6] F. CECCHETTO, V. TRAUB, AND R. ZENKLUSEN, *Bridging the gap between tree and connectivity augmentation: Unified and stronger approaches*, in Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2021, ACM, New York, 2021, pp. 370–383, <https://arxiv.org/abs/2012.00086>.
- [7] J. CHERIYAN AND Z. GAO, *Approximating (unweighted) tree augmentation via lift-and-project, part I: Stemless TAP*, Algorithmica, 80 (2018), pp. 530–559, <https://doi.org/10.1007/s00453-016-0270-4>.
- [8] J. CHERIYAN AND Z. GAO, *Approximating (unweighted) tree augmentation via lift-and-project, part II*, Algorithmica, 80 (2018), pp. 608–651, <https://doi.org/10.1007/s00453-017-0275-7>.
- [9] J. CHERIYAN, T. JORDÁN, AND R. RAVI, *On 2-coverings and 2-packings of laminar families*, in Proceedings of Algorithms - ESA '99, 7th Annual European Symposium, Prague, Czech Republic, Springer, Berlin, 1999, pp. 510–520, https://doi.org/10.1007/3-540-48481-7_44.

- [10] J. CHERIYAN AND R. THURIMELLA, *Approximating minimum-size k -connected spanning subgraphs via matching*, SIAM J. Comput., 30 (2000), pp. 528–560, <https://doi.org/10.1137/S009753979833920X>.
- [11] N. COHEN AND Z. NUTOV, *A $(1+\ln 2)$ -approximation algorithm for minimum-cost 2-edge-connectivity augmentation of trees with constant radius*, Theoret. Comput. Sci., 489-490 (2013), pp. 67–74, <https://doi.org/10.1016/j.tcs.2013.04.004>.
- [12] E. A. DINITZ, A. V. KARZANOV, AND M. V. LOMONOSOV, *On the structure of a family of minimal weighted cuts in a graph*, Studies in Discrete Optimization., Nauka, Moscow, 1976, pp. 290–306.
- [13] G. EVEN, J. FELDMAN, G. KORTSARZ, AND Z. NUTOV, *A 1.8 approximation algorithm for augmenting edge-connectivity of a graph from 1 to 2*, ACM Trans. Algorithms, 5 (2009), 21.
- [14] S. FIORINI, M. GROSS, J. KÖNEMANN, AND L. SANITÀ, *Approximating weighted tree augmentation via Chvátal-Gomory cuts*, in Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, A. Czumaj, ed., New Orleans, LA, SIAM, Philadelphia, 2018, pp. 817–831, <https://doi.org/10.1137/1.9781611975031.53>.
- [15] G. N. FREDERICKSON AND J. JÁJÁ, *Approximation algorithms for several graph augmentation problems*, SIAM J. Comput., 10 (1981), pp. 270–283.
- [16] H. N. GABOW AND S. R. GALLAGHER, *Iterated rounding algorithms for the smallest k -edge connected spanning subgraph*, SIAM J. Comput., 41 (2012), pp. 61–103, <https://doi.org/10.1137/080732572>.
- [17] M. X. GOEMANS, A. V. GOLDBERG, S. A. PLOTKIN, D. SHMOYS, É. TARDOS, AND D. P. WILLIAMSON, *Improved approximation algorithms for network design problems*, in Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, Arlington, Virginia, SIAM, Philadelphia, 1994, pp. 223–232, <https://dl.acm.org/doi/10.5555/314464.314497>.
- [18] F. GRANDONI, A. JABAL AMELI, AND V. TRAUB, *Breaching the 2-approximation barrier for the forest augmentation problem*, in Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022, ACM, New York, 2022, pp. 1598–1611.
- [19] F. GRANDONI, C. KALAITZIS, AND R. ZENKLUSEN, *Improved approximation for tree augmentation: Saving by rewiring*, in Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, ACM, New York, 2018, pp. 632–645, <https://doi.org/10.1145/3188745.3188898>.
- [20] W. GÁLVEZ, F. GRANDONI, A. JABAL AMELI, AND K. SORNAT, *On the cycle augmentation problem: Hardness and approximation algorithms*, Theory Comput. Syst., 65 (2021), pp. 1–24, <https://doi.org/10.1007/s00224-020-10025-6>.
- [21] C. HUNKENSHRÖDER, S. S. VEMPALA, AND A. VETTA, *A $4/3$ -approximation algorithm for the minimum 2-edge connected subgraph problem*, ACM Trans. Algorithms, 15 (2019), 55, <https://doi.org/10.1145/3341599>.
- [22] J. IGLESIAS AND R. RAVI, *Coloring down: $3/2$ -approximation for special cases of the weighted tree augmentation problem*, Oper. Res. Lett., 50 (2022), pp. 693–698.
- [23] K. JAIN, *A factor 2 approximation algorithm for the generalized Steiner network problem*, Combinatorica, 21 (2001), pp. 39–60.
- [24] S. KHULLER AND R. THURIMELLA, *Approximation algorithms for graph augmentation*, J. Algorithms, 14 (1993), pp. 214–225.
- [25] P. KLEIN AND R. RAVI, *A nearly best-possible approximation algorithm for node-weighted steiner trees*, J. Algorithms, 19 (1995), pp. 104–115.
- [26] G. KORTSARZ AND Z. NUTOV, *L_p -relaxations for tree augmentation*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2016, 2016, Paris, France, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Wadern, Germany, 2016, 13, <https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2016.13>.
- [27] G. KORTSARZ AND Z. NUTOV, *A simplified 1.5-approximation algorithm for augmenting edge-connectivity of a graph from 1 to 2*, ACM Trans. Algorithms, 12 (2016), 23.
- [28] D. MARX AND L. A. VÉGH, *Fixed-parameter algorithms for minimum-cost edge-connectivity augmentation*, ACM Trans. Algorithms, 11 (2015), 27, <https://doi.org/10.1145/2700210>.
- [29] H. NAGAMACHI, *An approximation for finding a smallest 2-edge-connected subgraph containing a specified spanning tree*, Discrete Appl. Math., 126 (2003), pp. 83–113.
- [30] Z. NUTOV, *2-node-connectivity network design*, in Approximation and Online Algorithms - 18th International Workshop, WAOA 2020, Virtual Event, 2020, Revised Selected Papers, C. Kaklamani, and A. Levin, eds., Lecture Notes in Comput. Sci. 12806, Springer, Cham, Switzerland, 2021, pp. 220–235, https://doi.org/10.1007/978-3-030-80879-2_15.

- [31] Z. NUTOV, *Approximation algorithms for connectivity augmentation problems*, in Proceeding of International Computer Science Symposium in Russia, CSR 2021, Sochi, 2021, pp. 321–338, https://link.springer.com/chapter/10.1007/978-3-030-79416-3_19.
- [32] Z. NUTOV, *On the tree augmentation problem*, *Algorithmica*, 83 (2021), pp. 553–575, <https://doi.org/10.1007/s00453-020-00765-9>.
- [33] G. ROBINS AND A. ZELIKOVSKY, *Tighter bounds for graph Steiner tree approximation*, *SIAM J. Discrete Math.*, 19 (2005), pp. 122–134.
- [34] A. SEBŐ AND J. VYGEN, *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 (2014), pp. 597–629, <http://arxiv.org/abs/1201.1870>.
- [35] V. TRAUB AND R. ZENKLUSEN, *A better-than-2 approximation for weighted tree augmentation*, in IEEE 62nd Annual IEEE Symposium on Foundations of Computer Science, IEEE, Piscataway, NJ, 2021, pp. 1–12.
- [36] V. TRAUB AND R. ZENKLUSEN, *Local search for weighted tree augmentation and Steiner tree*, in Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, SIAM, Philadelphia, pp. 3253–3271, <https://doi.org/10.1137/1.9781611977073.128>.