Finding Almost Tight Witness Trees

Dylan Hyatt-Denesik
Eindhoven University of Technology, The Netherlands

Afrouz Jabal Ameli
Eindhoven University of Technology, The Netherlands

Laura Sanità
Bocconi University, Milano, Italy

Abstract

This paper addresses a graph optimization problem, called the Witness Tree problem, which seeks a spanning tree of a graph minimizing a certain non-linear objective function. This problem is of interest because it plays a crucial role in the analysis of the best approximation algorithms for two fundamental network design problems: Steiner Tree and Node-Tree Augmentation. We will show how a wiser choice of witness trees leads to an improved approximation for Node-Tree Augmentation, and for Steiner Tree in special classes of graphs.

2012 ACM Subject Classification Theory of computation → Routing and network design problems

Keywords and phrases Algorithms, Network Design, Approximation

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.79

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2211.12431` [15]

Funding The second and the third authors are grateful for the support received from the NWO-VIDI grant VI.Vidi.193.087.

Acknowledgements The authors would like to thank Haris Angelidakis for his valuable discussion on this project. Furthermore, the authors would like to acknowledge the 2021 Hausdorff trimester program "Discrete Optimization", during which this work was started.

1 Introduction

Network connectivity problems play a central role in combinatorial optimization. As a general goal, one would like to design a cheap network able to satisfy some connectivity requirements among its nodes. Two of the most fundamental problems in this area are *Steiner Tree* and *Connectivity Augmentation*.

Given a network $G = (V, E)$ with edge costs, and a subset of terminals $R \subseteq V$, Steiner Tree asks to compute a minimum-cost tree T of G connecting the terminals in R. In Connectivity Augmentation, we are instead given a k-edge-connected graph $G = (V, E)$ and an additional set of edges $L \subseteq V \times V$ (called *links*). The goal is to add a minimum-cardinality subset of links to G to make it $(k + 1)$-edge-connected. It is well-known that the problem for odd k reduces to $k = 1$ (called Tree Augmentation), and for even k reduces to $k = 2$ (called Cactus Augmentation) (see [9]). All these problems are NP-hard, but admit a constant factor approximation. In the past 10 years, there have been several exciting breakthrough results in the approximation community on these fundamental problems (see [5, 13, 4, 16, 17, 6, 19, 14, 1, 7, 8, 11, 2, 18, 20]).

50th International Colloquium on Automata, Languages, and Programming (ICALP 2023).
Editors: Kousha Etessami, Uriel Feige, and Gabriele Puppis; Article No. 79; pp. 79:1–79:16
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

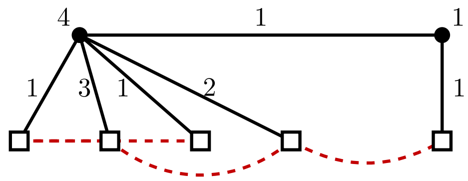

Figure 1 In black, the tree $T = (R \cup S, E)$. The dashed edges represent a witness tree W. The labels on edges of E and vertices of S indicate $\bar{w}(e)$ and $w(v)$, respectively. We have $\nu_T(W) = (H_4 + H_1)/2 = 1.541\bar{6}$. Assuming unit cost on the edges of E, we have $\bar{\nu}_T(W) = (4H_1 + H_2 + H_3)/6 = 1.\bar{2}$.

Several of these works highlight a deep relation between Steiner Tree and Connectivity Augmentation: the approximation techniques used for Steiner Tree have been proven to be useful for Connectivity Augmentation and vice versa. This fruitful exchange of tools and ideas has often lead to novel results and analyses. This paper continues bringing new ingredients in this active and evolving line of work.

Specifically, we focus on a graph optimization problem which plays a crucial role in the analysis of some approximation results mentioned before. This problem, both in its edge- and node-variant, is centered around the concept of *witness trees*. We now define this formally (see Figure 1 for an example).

Edge Witness Tree (EWT) problem. Given is a tree $T = (V, E)$ with edge costs $c : E \to \mathbb{R}_{\geq 0}$. We denote by R the set of leaves of T. The goal is to find a tree $W = (R, E_W)$, where $E_W \subseteq R \times R$, which minimizes the non-linear objective function $\bar{\nu}_T(W) = \frac{1}{c(E)} \sum_{e \in E} c(e) H_{\bar{w}(e)}$, where $c(E) = \sum_{e \in E} c(e)$, the function $\bar{w} : E \to \mathbb{Z}_{\geq 0}$ is defined as

$$\bar{w}(e) \coloneqq |\{pq \in E_W : e \text{ is an internal edge of the } p\text{-}q \text{ path in } T\}|$$

and H_ℓ denotes the ℓ^{th} harmonic number ($H_\ell = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{\ell}$).

Node Witness Tree (NWT) problem. Given is a tree $T = (V, E)$. We denote by R the set of leaves of T, and $S = V \setminus R$. The goal is to find a tree $W = (R, E_W)$, where $E_W \subseteq R \times R$, which minimizes the non-linear objective function $\nu_T(W) = \frac{1}{|S|} \sum_{v \in S} H_{w(v)}$, where $w : S \to \mathbb{Z}_{\geq 0}$ is defined as

$$w(v) \coloneqq |\{pq \in E_W : v \text{ is an internal node of the } p\text{-}q \text{ path in } T\}|$$

and again H_ℓ denotes the ℓ^{th} harmonic number.

We refer to a feasible solution W to either of the above problems as a *witness tree*. We call $\bar{w}$ (resp. w) the *vector imposed on* E (resp. S) by W. We now explain how these problems relate to Steiner Tree and Connectivity Augmentation.

EWT and relation to Steiner Tree

Currently, the best approximation factor for Steiner Tree is $(\ln(4) + \varepsilon)$, which can be achieved by three different algorithms [13] [5] [20]. These algorithms yield the same approximation because in all three of them, the analysis at some point relies on constructing witness trees.

More in detail, suppose we are given a Steiner Tree instance $(G = (V, E), R, c)$ where $c : E \to \mathbb{R}_{\geq 0}$ gives the edge costs. We can define the following:

$$\gamma_{(G,R,c)} \coloneqq \min_{\substack{T^* = (R \cup S^*, E^*):\ T^* \text{ is} \\ \text{optimal Steiner tree of } (G,R,c)}} \ \min_{\substack{W:\ W \text{ is a} \\ \text{witness tree} \\ \text{of } T^*}} \bar{\nu}_{T^*}(W)$$

We also define the following constant γ:

$$\gamma \coloneqq \sup\{\gamma_{(G,R,c)} : (G, R, c) \text{ is an instance of Steiner Tree}\}.$$

Byrka et al. [5] were the first to essentially prove the following.

▶ **Theorem 1.** *For any $\varepsilon > 0$, there is a $(\gamma + \varepsilon)$-approximation algorithm for Steiner Tree.*

Furthermore, the authors in [5] showed that $\gamma \leq \ln(4)$, and hence they obtained the previously mentioned $(\ln(4) + \varepsilon)$-approximation for Steiner Tree.

NWT and relation to Connectivity Augmentation

Basavaraju et al [3] introduced an approximation-preserving reduction from Cactus Augmentation (which is the hardest case of Connectivity Augmentation)[1] to special instances of Node-Steiner Tree, named *CA-Node-Steiner-Tree* instances in [2]: the goal here is to connect a given set R of terminals of a graph G via a tree that minimizes the number of non-terminal nodes (Steiner nodes) in it. The special instances have the crucial property that each Steiner node is adjacent to at most 2 terminals.

Byrka et al. [4] built upon this reduction to prove a 1.91-approximation for CA-Node-Steiner-Tree instances. This way, they were the first to obtain a better-than-2 approximation factor for Cactus Augmentation (and hence, for Connectivity Augmentation). Interestingly, Nutov [16] realized that a similar reduction also captures a fundamental node-connectivity augmentation problem: the *Node-Tree Augmentation* (defined exactly like Tree Augmentation, but replacing edge-connectivity with node-connectivity). This way, he could improve over an easy 2-approximation for Node-Tree Augmentation that was also standing for 40 years [12]. Angelidakis et al. [2] subsequently explicitly formalized the problem at the heart of the approximation analysis: namely, the NWT problem.

More in detail, given a CA-Node-Steiner-Tree instance $(G = (V, E), R)$, we can define the following:

$$\psi_{(G,R)} \coloneqq \min_{\substack{T^*=(R\cup S^*,E^*):\, T^* \text{ is} \\ \text{optimal Steiner tree of } (G,R)}} \; \min_{\substack{W:\, W \text{ is a} \\ \text{witness tree} \\ \text{of } T^*}} \nu_{T^*}(W),$$

We also define the constant ψ:

$$\psi \coloneqq \sup\{\psi_{(G,R)} : (G, R) \text{ is an instance of CA-Node-Steiner-Tree}\}.$$

Angelidakis et al. [2] proved the following.

▶ **Theorem 2.** *For any $\varepsilon > 0$, there is a $(\psi+\varepsilon)$-approximation algorithm for CA-Node-Steiner Tree.*

Furthermore, the authors of [2] proved that $\psi < 1.892$, and hence obtained a 1.892-approximation algorithm for Cactus Augmentation and Node-Tree Augmentation. This is currently the best approximation factor known for Node-Tree Augmentation (for Cactus Augmentation there is a better algorithm [6]).

[1] Tree Augmentation can be easily reduced to Cactus Augmentation by introducing a parallel copy of each initial edge.

Our results and techniques

Our main result is an improved upper bound on ψ. In particular, we are able to show $\psi < 1.8596$. Combining this with Theorem 2, we obtain a 1.8596-approximation algorithm for CA-Node-Steiner-Tree. Hence, due to the above mentioned reduction, we improve the state-of-the-art approximation for Node-Tree Augmentation.

▶ **Theorem 3.** *There is a 1.8596-approximation algorithm for CA-Node-Steiner-Tree (and hence, for Node-Tree Augmentation).*

Our result is based on a better construction of witness trees for the NWT problem. At a very high level, the witness tree constructions used previously in the literature use a *marking-and-contraction* approach, that can be summarized as follows. First, root the given tree T at some internal Steiner node. Then, every Steiner node v chooses (*marks*) an edge which connects to one of its children: this identifies a path from v to a terminal. Contracting the edges along this path yields a witness tree W. The way this marking choice is made varies: it is random in [5], it is biased depending on the nature of the children in [4], it is deterministic and taking into account the structure of T in [2]. However, all such constructions share the fact that decisions can be thought of as being taken "in one shot", at the same time for all Steiner nodes. Instead, here we consider a bottom-up approach for the construction of our witness tree, where a node takes a marking decision only after the decisions of its children have been made. A sequential approach of this kind allows a node to have a more precise estimate on the impact of its own decision to the overall non-linear objective function cost, but it becomes more challenging to analyze. Overcoming this challenge is the main technical contribution of this work, and the insight behind our improved upper-bound on ψ.

We complement this result with an almost-tight lower-bound on ψ, which improves over a previous lower bound given in [2].

▶ **Theorem 4.** *For any $\varepsilon > 0$, there exists a CA-Node-Steiner-Tree instance $(G_\varepsilon, R_\varepsilon)$ such that $\psi_{(G_\varepsilon, R_\varepsilon)} > 1.841\bar{6} - \varepsilon$.*

The above theorem implies that, in order to significantly improve the approximation for Node-Tree Augmentation, very different techniques need to be used. To show our lower-bound we prove a structural property on optimal witness trees, called *laminarity*, which in fact holds for optimal solutions of both the NWT problem and the EWT problem.

As an additional result, we also improve the approximation bound for Steiner Tree in the special case of *Steiner-claw free* instances. A Steiner-Claw Free instance is a Steiner-Tree instance where the subgraph $G[V \setminus R]$ induced by the Steiner nodes is claw-free (i.e., every node has degree at most 2). These instances were introduced in [10] in the context of studying the integrality gap of a famous LP relaxation for Steiner Tree, called the *bidirected cut relaxation*, that is long-conjectured to have integrality gap strictly smaller than 2.

▶ **Theorem 5.** *There is a ($\frac{991}{732} + \varepsilon < 1.354$)-approximation for Steiner Tree on Steiner-claw free instances.*

We prove the theorem by showing that, for any Steiner-Claw Free instance (G, R, c), $\gamma_{(G,R,c)} \leq \frac{991}{732}$. The observation we use here is that an optimal Steiner Tree solution T in this case is the union of components that are caterpillar graphs[2]: this knowledge can be

[2] A caterpillar graph is defined as a tree in which every leaf is of distance 1 from a central path.

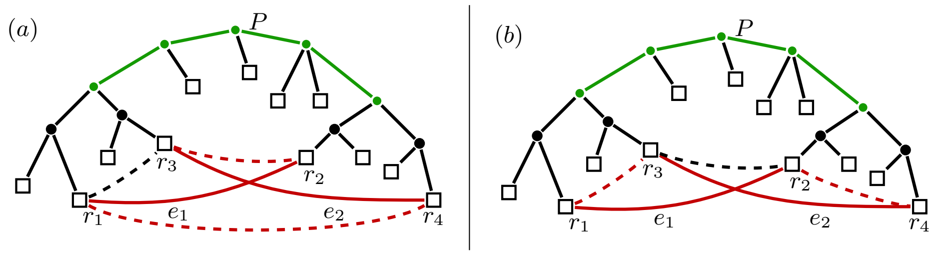

Figure 2 In both figures we have a tree, T, shown with black edges and green edges, with leaves, R, denoted by squares. Crossing edges e_1 and e_2 are shown with solid red edges. The green edges denote the path P. Figure (a): In this case, r_1 and r_3 are in the same component of $W\backslash\{e_1, e_2\}$, represented by the dashed black edge. We can replace e_1 with r_2r_3 or replace e_2 with r_1r_4 (red dashed edges). Figure (b): In this case, r_3 and r_2 are in the same component, denoted by the black dashed edge. We can replace e_1 and e_2 with r_1r_3 and r_2r_4 (red dashed edges).

exploited to design ad-hoc witness trees. Interestingly, we can also show that this bound is tight: once again, the proof of this lower-bound result relies on showing laminarity for optimal witness trees.

▶ **Theorem 6.** *For any $\varepsilon > 0$, there exists Steiner-Claw Free instance $(G_\varepsilon, R_\varepsilon, c_\varepsilon)$ such that $\gamma_{(G_\varepsilon, R_\varepsilon, c_\varepsilon)} > \frac{991}{732} - \varepsilon$.*

As a corollary of our results, we also get an improved bound on the integrality gap of the bidirected cut relaxation for Steiner-Claw Free instances (this follows directly from combining our upper bound with the results in [10]). Though these instances are quite specialized, they serve the purpose of passing the message: exploiting the structure of optimal solutions helps in choosing better witnesses, hopefully arriving at tight (upper and lower) bounds on γ and ψ.

2 Laminarity

In this section, we prove some key structural properties of witness trees. We assume to be given a Node (Edge) Witness Tree instance $T = (V, E)$ with leaves R (and edge costs $c : E \to \mathbb{R}_{\geq 0}$), where R denotes the leaves of T, we will show that we can characterize witness trees minimizing $\nu_T(W)$ ($\bar{\nu}_T(W)$) using the following notion of *laminarity*. Given a witness tree $W = (R, E_W)$, we say edges $f_1f_2, f_3f_4 \in E_W$ *cross* if the f_1-f_2 and f_3-f_4 paths in T share an internal node but not an endpoint. We say that W is *laminar* if it has no crossing edges. For nodes $u, v \in V$, we denote by T_{uv} the path in T between the nodes u and v. Similarly, for $e \in E_W$, we denote by T_e the path in T between the endpoints of e.

The following Theorem shows that there is always a witness tree minimizing $\nu_T(W)$ that is laminar.

▶ **Theorem 7.** *Given an instance of the Node Witness Tree problem $T = (V, E)$, let $\mathcal{W}$ be the family of all witness trees for T. Then there exists a laminar witness tree W such that $\nu_T(W) = \min_{W' \in \mathcal{W}} \nu_T(W')$.*

Proof. We first show that there is a witness tree W minimizing $\nu_T(W)$ such that the induced subgraph of W on any maximal set of terminals that share a neighbour in $V\backslash R$ is a star. We assume for the sake of contradiction that there is a maximal set of terminals $S \subseteq R$ sharing a neighbour $v \in V\backslash R$, such that the induced subgraph of W on S is a set of connected components $W_1, \ldots, W_i$ for $i > 1$. Without loss of generality, suppose the shortest path between two components is from W_1 to W_2, and let e denote the edge of this path incident to W_2. We define $W' := W \cup \{f\} \backslash \{e\}$, where f is an arbitrary edge between W_1 and W_2. Since $\{v\} = T_f \backslash R \subsetneq T_e \backslash R$, we have $\nu_T(W') < \nu_T(W)$, contradicting the minimality of W. Therefore, the induced subgraph on S is connected. We can rearrange the edges of this subgraph to be a star as this will not affect $\nu_T(W)$, so we assume this holds on W for any such S.

For a maximal set of terminals $S \subseteq R$ that share a neighbour, by a slight abuse of notation, we denote by S the induced star subgraph of W on S, and denote its center by $s \in S$. We will assume without loss of generality that edges of W incident to S have endpoint s. To see this, as S is a connected subgraph of W, any pair of edges incident to S cannot share an endpoint outside of S, otherwise we have found a cycle in W. Furthermore, for any edge of W incident to S where s is not an endpoint, we can change the endpoint in S of that edge to be s and maintain the connectivity of W since S is connected. Edges changed in this way will have the same interior nodes between their endpoints, so this does not increase $\nu_T(W)$.

We assume for the sake of contradiction that the witness tree W minimizing $\nu_T(W)$ is not a laminar witness tree. As W is not laminar, there exist distinct leaves $r_1, r_2, r_3, r_4 \in R$ such that $e_1 = r_1r_2, e_2 = r_3r_4 \in E_W$ are crossing. We denote the path $T_{e_1} \cap T_{e_2}$ by P. We denote by P_i the (potentially empty) set of internal nodes of the shortest path from P to r_i in T.

Since e_1 and e_2 are crossing edges, one of $T_{r_1r_3}$ or $T_{r_1r_4}$ contains exactly one node of P. The same is true for r_2. Without loss of generality, let us assume that the paths $T_{r_1r_3}$ and $T_{r_2r_4}$ contain exactly one node of P. We consider by cases which component of $W \backslash \{e_1, e_2\}$ contains two nodes among r_1, r_2, r_3 and r_4. See Figure 2 for an example.

- Case: r_1 and r_3 (or similarly, r_2 and r_4) are in the same component of $W \backslash \{e_1, e_2\}$. If $P_1 = P_3 = \emptyset$, then r_1 and r_3 share a neighbour and thus, as shown above, e_1 and e_2 are assumed to share an endpoint, and are thus not crossing.
  Consider $W' := W \cup \{r_2r_3\} \backslash \{e_1\}$ and $W'' := W \cup \{r_1r_4\} \backslash \{e_2\}$. If $\nu_T(W) - \nu_T(W') > 0$, this contradicts the minimality of $\nu_T(W)$. Therefore, we can see
$$0 \le |V \backslash R|(\nu_T(W') - \nu_T(W)) = \sum_{u \in P_3} \frac{1}{w(u)+1} - \sum_{u \in P_1} \frac{1}{w(u)}$$
$$< \sum_{u \in P_3} \frac{1}{w(u)} - \sum_{u \in P_1} \frac{1}{w(u)+1} = |V \backslash R|(\nu_T(W) - \nu_T(W''))$$
  Clearly, we have $\nu_T(W'') < \nu_T(W)$, contradicting minimality of $\nu_T(W)$.
- Case: r_2 and r_3 (or similarly, r_1 and r_4) are in the same component of $W \backslash \{e_1, e_2\}$. Without loss of generality we can assume that $|V(P)| > 1$, because if $|V(P)| = 1$ then we can reduce to the previous case by relabelling the nodes r_1, r_2, r_3 and r_4. In this case, consider $W' := W \cup \{r_1r_3, r_2r_4\} \setminus \{e_1, e_2\}$. Therefore, we can see
$$|V \backslash R|\,(\nu_T(W') - \nu_T(W)) \le -\sum_{u \in P} \frac{1}{w(u)} < 0$$

Thus, we have $\nu_T(W') < \nu_T(W)$, contradicting the minimality of $\nu_T(W)$. ◀

The following theorem, similar to Theorem 7, shows that there are laminar witness trees that are optimal for the EWT problem. The proof is deferred to the full version of the paper.

▶ **Theorem 8.** *Given an instance of the Edge Witness Tree problem $T = (V, E)$ with edge costs c, let $\mathcal{W}$ be the family of all witness trees for T. Then there exists a laminar witness tree W such that $\bar{\nu}_T(W) = \min_{W' \in \mathcal{W}} \bar{\nu}_T(W')$.*

We now show that laminar witness trees are precisely the set of trees that one could obtain with a marking-and-contraction approach. The proof of this Theorem can be found in the full version of the paper.

▶ **Theorem 9.** *Given a tree $T = (V, E)$ with leaves R, a witness tree $W = (R, E_W)$ for T can be found by marking-and-contraction if and only if W is laminar.*

Incidentally, this has the following side implication. The authors of [13] gave a dynamic program (that is also a bottom-up approach) to compute the best possible witness tree obtainable with a marking-and-contraction scheme. Our structural results imply that their dynamic program computes an optimal solution for the EWT problem (though for the purpose of the approximation analysis, being able to compute the best witness tree is not that relevant: being able to bound ψ and γ is what matters).

3 Improved approximation for CA-Node-Steiner Tree

The goal of this section is to prove Theorem 3. We will achieve this by showing $\psi < 1.8596$, and by using Theorem 2. From now on, we assume we are given a tree $T = (R \cup S^*, E^*)$, where each Steiner node is adjacent to at most two terminals.

3.1 Preprocessing

We first apply some preprocessing operations as in [2], that allow us to simplify our witness tree construction. The first one is to remove the terminals from T, and then decompose T into smaller components which will be held separately. We start by defining a *final* Steiner node as a Steiner node that is adjacent to at least one terminal. We let $F \subseteq S^*$ denote the set of final Steiner nodes. Since we remove the terminals from T, we will construct a spanning tree W on F with edges in $F \times F$. With a slight abuse of notation, we refer to W as a witness tree: this is because [2, Section 4.1] showed that one can easily map W to a witness tree for our initial tree T (with terminals put back), and the following can be considered the vector imposed on S^* by W:

$$w(v) := |\{pq \in E_W : v \text{ belongs to the } p\text{-}q \text{ path in } T[S^*]\}| + \mathbb{1}[v \in F] \tag{1}$$

where $\mathbb{1}[v \in F]$ denotes the indicator of the event "$v \in F$", and $T[S^*]$ is the subtree of T induced by the Steiner nodes. See Figure 3.

So, from now on, we consider $T = T[S^*]$. The next step is to root T at an arbitrary final node $r \in F$. Following [2] we can decompose T into a collection of rooted components $T_1, \dots T_\tau$, where a component is a subtree whose leaves are final nodes and non-leaves are non-final nodes. The decomposition will have the following properties: each T_i is rooted at a final node r_i that has degree one in T_i, $r_1 := r$ is the root of T_1, $\cup_{j<i} T_j$ is connected, and $T = \cup_{i=1}^{\tau} T_i$. We will compute a witness tree W_i for each component T_i, and then show that we can join these witness trees $\{W_i\}_{i \geq 1}$ together to get a witness tree W for T.

3.2 Computing a witness tree W_i for a component T_i

Here we deal with a component T_i rooted at r_i, and describe how to construct a witness tree W_i. If T_i is a single edge $e = r_i v$, we simply let $W_i = (\{r_i, v\}, \{r_i v\})$.

Now we assume that T_i is not a single edge. We will construct a witness tree with a bottom-up procedure. At a high level, each node $u \in T_i \backslash r_i$ looks at the subtree Q_u of T_i rooted at u, and constructs a portion of the witness tree: namely, a subtree $\overline{W}^u$ spanning

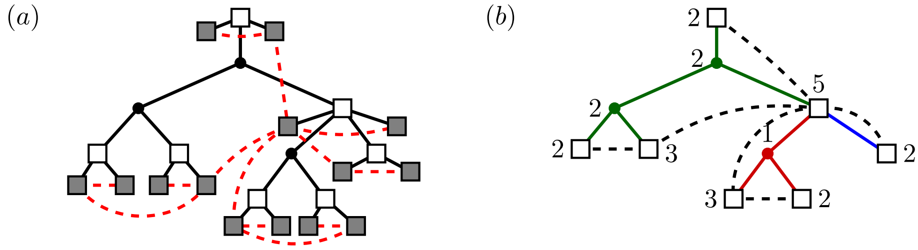

Figure 3 Figure (a): A tree T is shown by black edges. The terminals are shown by grey squares. The final Steiner nodes are shown by white squares, non-final Steiner nodes are shown by black dots. Figure (b): The tree T after the terminals have been removed. The color edges indicate the three components. A witness tree W is shown by the black dashed lines. The numbers indicate the values of w imposed on T computed according to (1). Red dashed lines in Figure (a) show how W can be mapped back.

the leaves of Q_u (note that, in case the degree of u is 1 in Q_u, we do not consider u to be a leaf of Q_u but just its root). Assume u has children $u_1, \ldots, u_k$. Because of the bottom-up procedure, each child u_j has already constructed a subtree $\overline{W}^{u_j}$. That is, u has to decide how to join these subtrees to get $\overline{W}^u$.

To describe how this is done formally, we first need to introduce some more notation. For every node $u \in T_i \backslash F$, we select one of its children as the "marked child" of u (according to some rule that we will define later). In this way, for every $u \in T_i$ there is a unique path along these marked children to a leaf. We denote this path by $P(u)$, and we let $\ell(u)$ denote the leaf descendent of this path. For final nodes $u \in F$, we define $\ell(u) := u$ and $P(u) := u$. For a subtree Q_u of T_i rooted at u and a witness tree $\overline{W}^u$ over the leaves of Q_u, let $\overline{w}^u$ be the vector imposed on the nodes of Q_u by $\overline{W}^u$ according to (1). Next, we define the following quantity (which, roughly speaking, represents the cost-increase incurred after increasing $\overline{w}^u(v)$ for each $v \in P(u) \backslash \ell(u)$ for the $(j+1)^{th}$ time):

$$C_j^u := \sum_{v \in P(u) \backslash \ell(u)} \left(H_{\overline{w}^u(v)+j+1} - H_{\overline{w}^u(v)+j} \right) = \sum_{v \in P(u) \backslash \ell(u)} \frac{1}{\overline{w}^u(v) + j + 1}$$

Algorithm 1 Computing the tree $\overline{W}^u$.

u has Steiner node children $u_1, u_2, \ldots, u_k$, and $\overline{W}^{u_j}$ have been defined
if $u_1, \ldots, u_k$ *are all non-final,* **then**
  The *marked* child is u_m, minimizing $C_1^{u_m}$
else
  Assume $\{u_1, \ldots, u_{k_1}\}$, $1 \leq k_1 \leq k$, are final node children of u
  if $k_1 = k$*, or, for all* $j \in \{k_1+1, \ldots, k\}$, $C_1^{u_j} \geq \phi - \delta - H_2$ **then**
    The *marked* child of u is u_m for $1 \leq m \leq k_1$ such that $C_1^{u_m}$ is minimized.
  if *There is a* $j \in \{k_1+1, \ldots, k\}$ *such that* $C_1^{u_j} < \phi - \delta - H_2$ **then**
    The *marked* child of u is u_m for $k_1 < m \leq k$ such that $C_1^{u_m}$ is minimized.
$\overline{W}^u \leftarrow \left(\bigcup_{j=1}^k V[Q_{u_j}], \bigcup_{j=1}^k \overline{W}^{u_j} \bigcup_{j \neq m} \{\ell(u_m)\ell(u_j)\} \right)$
Return $\overline{W}^u$

We can now describe the construction of the witness tree more formally. We begin by considering the leaves of T_i; for a final node (leaf) u, we define a witness tree on the (single) leaf of Q_u as $\overline{W}^u = (\{u\}, \emptyset)$. For a non-final node u, with children $u_1, \ldots, u_k$ and corresponding witness trees $\overline{W}^{u_1}, \dots, \overline{W}^{u_k}$, we select a marked child u_m for u as outlined in Algorithm 1, setting $\phi = 1.86 - \frac{1}{2100}$ and $\delta = \frac{97}{420}$. With this choice, we compute $\overline{W}^u$ by joining the subtrees $\overline{W}^{u_1}, \dots, \overline{W}^{u_k}$ via the edges $\ell(u_m)\ell(u_j)$ for $j \neq m$. Finally, let v be the unique child of r_i. We let W_i be equal to the tree $\overline{W}^v$ plus the extra edge $\ell(v)r_i$, to account for the fact that r_i is also a final node.

3.3 Bounding the cost of W_i

It will be convenient to introduce the following definitions. For a component T_i and a node $u \in T_i \setminus r_i$, we let W^u be the tree $\overline{W}^u$ plus one extra edge e^u, defined as follows. Let $a(u)$ be the first ancestor node of u with $\ell(a(u)) \neq \ell(u)$ (recall $\ell(r_i) = r_i$). We then let the edge $e^u := \ell(u)\ell(a(u))$. We denote by w^u the vector imposed on the nodes of Q_u by $W^u := \overline{W}^u + e^u$. Note that, with this definition, $W_i = W^v$ for v being the unique child of r_i.

We now state two useful lemmas. The first one relates the functions w^u and w^{u_j} for a child u_j of u. The statements (a)-(c) below can be proved similarly to Lemma 4 of [2]. We defer its proof to the full version of the paper.

▶ **Lemma 10.** *Let $u \in T_i \setminus r_i$ have children $u_1, \dots, u_k$, and u_1 be its marked child. Then:*

a $w^u(u) = k$.

b *For every $j \in \{2, \dots, k\}$ and every node $v \in Q_{u_j}$, $w^u(v) = w^{u_j}(v)$.*

c *For every $v \in Q_{u_1} \setminus P(u_1)$, $w^u(v) = w^{u_1}(v)$.*

d $\sum_{v \in P(u_1) \setminus \ell(u_1)} H_{w^u(v)} = \sum_{v \in P(u_1) \setminus \ell(u_1)} H_{w^{u_1}(v)} + \sum_{j=1}^{k-1} C_j^{u_1}$.

Next lemma relates the "increase" of cost C_j^u to the *degree* of some nodes in T_i.

▶ **Lemma 11.** *Let $u \in T_i \setminus r_i$ have children $u_1, \dots, u_k$, and u_1 be its marked child. Then, $C_1^u = C_k^{u_1} + \frac{1}{k+1}$. Furthermore, if u_1 is non-final and has degree d in T_i, then:*
1) $\sum_{j=1}^{k} (C_j^{u_1} - C_1^{u_j}) \leq \sum_{j=1}^{k-1} \left(\frac{1}{d+j} - \frac{1}{d}\right)$; 2) $H_{w^u(\ell(u_1))} - H_{w^{u_1}(\ell(u_1))} \leq \sum_{j=1}^{k-1} \frac{1}{d+j}$

Proof.

1. First observe that since $C_1^{u_1} = \min_{j \in [k]} C_1^{u_j}$, we have $C_j^{u_1} - C_1^{u_j} \leq C_j^{u_1} - C_1^{u_1}$. Consider $j \geq 1$, $C_j^{u_1} - C_1^{u_1}$ is equal to

$$= \sum_{v \in P(u_1) \setminus \ell(u)} \left(H_{w^{u_1}(v)+j} - H_{w^{u_1}(v)+j-1} - H_{w^{u_1}(v)+1} + H_{w^{u_1}(v)}\right)$$
$$= \sum_{v \in P(u_1) \setminus \ell(u)} \left(\frac{1}{w^{u_1}(v)+j} - \frac{1}{w^{u_1}(v)+1}\right) \leq \frac{1}{w^{u_1}(u_1)+j} - \frac{1}{w^{u_1}(u_1)+1}$$

Where the inequality follows since every term in the sum is negative. We know that $w^{u_1}(u_1) = d-1$ by Lemma 10.(a), therefore, $C_j^{u_1} - C_1^{u_1} \leq \frac{1}{d+j-1} - \frac{1}{d}$, and the claim is proven by summing over $j = 1, \dots, k$.

2. To prove the second inequality, first observe that $w^u(\ell(u_1)) = w^{u_1}(\ell(u_1)) + k - 1$. This follows by recalling that W^u is equal to $\overline{W}^{u_1}, \dots, \overline{W}^{u_k}$ plus the edges $\ell(u_1)\ell(u_j)$ for $j \neq 1$, and e^u. Thus, $H_{w^u(\ell(u_1))} - H_{w^{u_1}(\ell(u_1))} = H_{w^{u_1}(\ell(u_1))+k-1} - H_{w^{u_1}(\ell(u_1))} = \sum_{i=1}^{k-1} \frac{1}{w^{u_1}(\ell(u_1))+i}$. Recall u_1 is not a final node, so $w^{u_1}(\ell(u_1)) > d$. Therefore,

$$\sum_{i=1}^{k-1} \frac{1}{w^{u_1}(\ell(u_1))+i} \leq \sum_{i=1}^{k-1} \frac{1}{d+i}.$$ ◀

3.4 Key Lemma

To simplify our analysis, we define $h_{W^u}(Q_u) \coloneqq \sum_{\ell \in Q_u} H_{w^u(\ell)}$, and we let $|Q_u|$ be the number of nodes in Q_u. The next lemma is the key ingredient to prove Theorem 3.

▶ **Lemma 12.** *Let $\delta = \frac{97}{420}$ and $\phi = 1.86 - \frac{1}{2100}$. Let $u \in T_i \setminus r_i$ and k be the number of its children. Let $\beta(k)$ be equal to 0 for $k = 0, \dots, 8$ and $\frac{1}{3} - \delta$ for $k \geq 9$. Then*

$$h_{W^u}(Q_u) + C_1^u + \delta + \beta(k) \leq \phi \cdot |Q_u|$$

Proof. The proof of Lemma 12 will be by induction on $|Q_u|$. The base case is when $|Q_u| = 1$, and hence u is a leaf of T_i. Therefore, W^u is just the edge e^u, and so by definition of w^u we have $w^u(u) = 2$. We get $h_{W^u}(Q_u) = 1.5$, $C_1^u = 0$, $\beta(k) = 0$ and the claim is clear.

For the induction step: suppose that u has children $u_1, \dots, u_k$. We will distinguish 2 cases: (i) u has no children that are final nodes; (ii) u has some child that is a final node (which is then again broken into subcases). We report here only the proof of case (i), and defer the proof of the other case to the full version of the paper as the reasoning follows similar arguments.

Case (i): No children of u are final

According to Algorithm 1, we mark the child u_m of u that minimizes $C_1^{u_j}$. Without loss of generality, let $u_m = u_1$. Furthermore, let $\ell \coloneqq \ell(u_1)$. We note the following.

$$h_{W^u}(Q_u) = \sum_{j=1}^{k} h_{W^u}(Q_{u_j}) + H_{w^u(u)}$$

By applying Lemma 10.(a) we have $H_{w^u(u)} = H_k$. By Lemma 10.(b) we see $h_{W^u}(Q_{u_j}) = h_{W^{u_j}}(Q_{u_j})$ for $j \geq 2$. Using Lemma 10.(c) and (d) we get $h_{W^u}(Q_{u_1}) = h_{W^{u_1}}(Q_{u_1}) + \sum_{j=1}^{k-1} C_j^{u_1} + H_{w^u(\ell)} - H_{w^{u_1}(\ell)}$. Therefore:

$$h_{W^u}(Q_u) = \sum_{j=1}^{k} h_{W^{u_j}}(Q_{u_j}) + \sum_{j=1}^{k-1} C_j^{u_1} + H_k + H_{w^u(\ell)} - H_{w^{u_1}(\ell)}$$

We apply our inductive hypothesis on $Q_{u_1}, \dots, Q_{u_k}$, and use $\beta(j) \geq 0$ for all j:

$$\begin{aligned} h_{W^u}(Q_u) \leq & \sum_{j=1}^{k} \left(\phi|Q_{u_j}| - \delta - C_1^{u_j}\right) + \sum_{j=1}^{k-1} C_j^{u_1} + H_k + H_{w^u(\ell)} - H_{w^{u_1}(\ell)} \\ =& \phi(|Q_u| - 1) - k\delta - C_k^{u_1} + \sum_{j=1}^{k} \left(C_j^{u_1} - C_1^{u_j}\right) + H_k + H_{w^u(\ell)} - H_{w^{u_1}(\ell)} \end{aligned}$$

Using Lemma 11, we get

$$\begin{aligned} \leq & \phi(|Q_u| - 1) - k\delta - C_1^u + \sum_{j=1}^{k-1} \left(\frac{1}{d+j} - \frac{1}{d}\right) + H_{k+1} + \sum_{j=1}^{k-1} \frac{1}{d+j} \\ \leq & \phi|Q_u| - \delta - C_1^u - \beta(k) \end{aligned}$$

where the last inequality follows since one checks that for any $k \geq 1$ and $d \geq 2$ we have $-\phi - (k-1)\delta + \sum_{j=1}^{k-1} \left(\frac{1}{d+j} - \frac{1}{d}\right) + H_{k+1} + \sum_{j=1}^{k-1} \frac{1}{d+j} \leq -\beta(k)$. We show this inequality the full version of the paper. ◀

3.5 Merging and bounding the cost of W

Once the $\{W_i\}_{i\geq 1}$ are computed for each component T_i, we let the final witness tree be simply the union $W = \cup_i W_i$. Our goal now is to prove the following.

▶ **Lemma 13.** $\nu_T(W) \leq \phi = 1.86 - \frac{1}{2100}$.

Proof. Recall that we decomposed T into components $\{T_i\}_{i=1}^{\tau}$, such that $\cup_{j\leq i} T_j$ is connected for all $i \in [\tau]$. For a given i, define $T' = \cup_{j<i} T_j$, $W' = \cup_{j<i} W_i$, and let w' be the vector imposed on the nodes of T' by W' (for $i = 1$, set $T' = \emptyset$, $W' = \emptyset$, and $w' = 0$). Finally, define $W'' = W_i \cup W'$ and let w'' be the vector imposed on the nodes of $T'' := T' \cup T_i$. By induction on i, we will show that $\nu_{T''}(W'') \leq \phi$. The statement will then follow by taking $i = \tau$. Recall that, for any i, r_i is adjacent to a single node v in T_i, and $W_i = W^v$.

First consider $i = 1$. Hence, $W'' = W_1 = W^v$ and $w''(r_1) = 2$. By applying Lemma 12 to the subtree Q_v we get

$$\sum_{u\in T''} H_{w''(u)} = h_{W^v}(Q_v) + H_{w''(r_i)} \leq \phi(|Q_v|) + H_2 \leq \phi(|Q_v| + 1) \Rightarrow \nu_{T''}(W'') \leq \phi$$

Now consider $i > 1$. In this case, $w''(r_i) = w'(r_i) + 1 \geq 3$. Therefore:

$$\sum_{u\in T''} H_{w''(u)} = \sum_{u\in T_i\setminus r_i} H_{w^v(u)} + \sum_{u\in T'} H_{w'(u)} - H_{w'(r_i)} + H_{w'(r_i)+1}$$
$$= \sum_{u\in T_i\setminus r_i} H_{w^v(u)} + \sum_{u\in T'} H_{w'(u)} + \frac{1}{w'(r_i)+1} \leq \sum_{u\in T_i\setminus r_i} H_{w^v(u)} + \sum_{u\in T'} H_{w'(u)} + \frac{1}{3}$$

If v is a final node, then $\sum_{u\in T_i\setminus r_i} H_{w^v(u)} = H_{w^v(v)} = H_2$ and by induction

$$\sum_{u\in T''} H_{w''(u)} \leq H_3 + \sum_{u\in T'} H_{w'(u)} \leq \phi|T''| \Rightarrow \nu_{T''}(W'') \leq \phi$$

If v is not a final node, then by induction on T' and by applying Lemma 12 to the subtree Q_v, assuming that v has k children, we can see

$$\sum_{u\in T''} H_{w''(u)} \leq \phi|T''| - C_1^v - \delta - \beta(k) + \frac{1}{3} \leq \phi|T''| - \frac{1}{k+1} - \delta - \beta(k) + \frac{1}{3}$$

If $1 \leq k \leq 8$, then $\beta(k) = 0$, but we have $\frac{1}{3} < 431/1260 = \frac{1}{9} + \delta \leq \frac{1}{k+1} + \delta$. If $k \geq 9$, $\beta(k) = \frac{1}{3} - \delta$ and $\frac{1}{3} - \delta - \beta(k) = 0$. In both cases, $\nu_{T''}(W'') \leq \phi$. ◀

Note that we did not make any assumption on T, other than being a CA-Node-Steiner-Tree. Hence, Lemma 13 yields the following corollary.

▶ **Corollary 14.** $\psi \leq 1.86 - \frac{1}{2100} < 1.8596$.

Combining Corollary 14 with Theorem 2 yields a proof of Theorem 3.

4 Improved Lower Bound on ψ

The goal of this section is to prove Theorem 4. For the sake of brevity, we will omit several details. (see the full version of the paper for a completed proof).

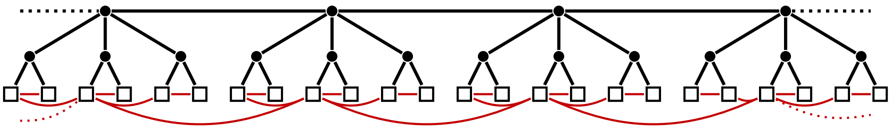

Figure 4 Lower bound instance shown in black. The white squares are terminals and black circles are Steiner nodes. Red edges form the laminar witness tree W^*.

Sketch of Proof of Theorem 4

Consider a CA-Node-Steiner-Tree instance (G, R), where G consists of a path of Steiner nodes $s_1, \dots, s_q$ such that, for all $i \in [q]$, s_i is adjacent to Steiner nodes t_{i1}, t_{i2}, t_{i3}, and each t_{ij} is adjacent to two terminals r^1_{ij} and r^2_{ij}. See Figure 4. We will refer to B_i as the subgraph induced by $s_i, t_{ij}, r^1_{ij}, r^2_{ij}$ $(j = 1, 2, 3)$. Since G is a tree connecting the terminals, clearly the optimal Steiner tree for this instance is $T = G$.

Let W^* be a witness tree that minimizes $\nu_T(W^*)$. Recall that we can assume W^* to be laminar by Theorem 7. We arrive at an explicit characterization of W^* in three steps. First, we observe that, without loss of generality, we can assume that every pair of terminals r^1_{ij} and r^2_{ij} are adjacent in W^* and that r^2_{ij} is a leaf of W^*. Second, using the latter of these observations and laminarity, we show that for all i, the subgraph of W induced by $r^1_{i1}, r^1_{i2}, r^1_{i3}$ can only be either (a) a star, or (b) three singletons, adjacent to a unique terminal $f \notin B_i$. We say that B_i is a *center* in W^* if (a) holds. Finally, we get rid of case (b), and essentially arrive at the next lemma, whose proof can be found in the full version of the paper.

▶ **Lemma 15.** *Let $\mathcal{W}$ be the family of all laminar witness trees over T, and let W^* be a laminar witness tree such that for every $i \in [q]$, B_i is a center in W^*. Then $\nu_T(W^*) = \min_{W \in \mathcal{W}} \nu_T(W)$.*

Once we impose the condition that all B_i are centers, one notes that the tree W^* essentially must look like the one shown in Figure 4. So it only remains to compute $\nu_T(W^*)$. For every B_i, we can compute $\sum_{v \in B_i} H_{w^*(v)}$, where w^* is the vector imposed on the set S of Steiner nodes by W^*. For $i \in \{2, \dots, q-1\}$, one notes that $\frac{1}{4}\sum_{v \in B_i} H_{w^*(v)} = \frac{1}{4}(2H_2 + H_4 + H_5) = 221/120 = 1.841\bar{6}$. Similarly, for $i = 1$ and q we have $\frac{1}{4}\sum_{v \in B_1} H_{w^*(v)} = \frac{1}{4}\sum_{v \in B_q} H_{w^*(v)} = \frac{1}{4}(2H_2 + H_3 + H_4) = \frac{83}{48} = 1.7291\bar{6}$. Therefore, we can see that $\nu_T(W^*) = \sum_{v \in S} \frac{H_{w^*(v)}}{|S|} = \frac{1.841\bar{6}q - 2(1.841\bar{6} - 1.7291\bar{6})}{q}$. Thus, for $q > \frac{1}{\varepsilon}$ we have $\nu_T(W^*) > 1.841\bar{6} - \frac{1}{q}$.

5 Tight bound for Steiner-Claw Free Instances

We here prove Theorem 5. Our goal is to show that for any Steiner-Claw Free instance (G, R, c), $\gamma_{(G,R,c)} \le \frac{991}{732}$, improving over the known $\ln(4)$ bound that holds in general. From now on, we assume that we are given an optimal solution $T = (R \cup S^*, E^*)$ to (G, R, c).

Simplifying Assumptions

As standard, note that T can be decomposed into components $T_1, \dots, T_\tau$, where each component is a maximal subtree of T whose leaves are terminals and internal nodes are Steiner nodes. Since components do not share edges of T, it is not difficult to see that one can compute a witness tree W_i for each component T_i separately, and then take the union of the $\{W_i\}_{i \ge 1}$ to get a witness tree W whose objective function $\bar{\nu}_T(W)$ will be bounded

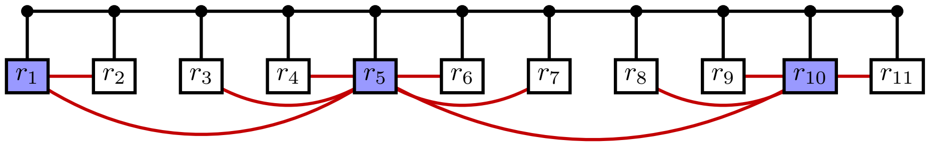

Figure 5 Edges of T are shown in black. Red edges show W. Here, $q = 11$, $t_\alpha = 5$ and $\sigma = 5$. Initially r_5 and r_{10} are picked as the centers of stars in W. Since $\sigma > \lceil \frac{t_\alpha}{2} \rceil$, r_1 is also the center of a star. Since $\sigma + t_\alpha \lfloor \frac{q-\sigma}{t_\alpha} \rfloor > q - \lceil \frac{t_\alpha}{2} \rceil$, r_q is not the center of a star.

by the maximum among $\bar{\nu}_{T_i}(W_i)$. Hence, from now on we assume that T is made by one single component. Since T is a solution to a Steiner-claw free instance, each Steiner node is adjacent to at most 2 Steiner nodes. In particular, the Steiner nodes induce a path in T, which we enumerate as $s_1, \ldots, s_q$. We will assume without loss of generality that each s_j is adjacent to exactly one terminal $r_j \in R$: this can be achieved by replacing a Steiner node incident to p terminals, with a path of length p made of 0-cost edges, if $p > 1$, and with an edge of appropriate cost connecting its 2 Steiner neighbors, if $p = 0$. We will also assume that $q > 4$. For $q \leq 4$, it is not hard to compute that $\gamma_{(G,R,c)} \leq \frac{991}{732}$. (For sake of completeness we explain this in the full version of the paper)

Witness tree computation and analysis

We denote by $L \subseteq E^*$ the edges of T incident to a terminal, and by $O = E^* \setminus L$ the edges of the path $s_1, \ldots, s_q$. Let $\alpha := c(O)/c(L)$. For a fixed value of $\alpha \geq 0$, we will fix a constant t_α as follows: If $\alpha \in [0, 32/90]$, then $t_\alpha = 5$, if $\alpha \in (32/90, 1)$, then $t_\alpha = 3$, and if $\alpha \geq 1$, then $t_\alpha = 1$. Given α (and thus t_α), we construct W using the randomized process outlined in Algorithm 2. At a high level, starting from a random offset, Algorithm 2 adds sequential stars of t_α terminals to W, connecting the centers of these stars together in this sequence. See Figure 5 for an example.

Algorithm 2 Computing the witness tree W.

Initialize $W = (R, E_W = \emptyset)$
Sample uniformly at random σ from $\{1, \ldots, t_\alpha\}$.
$E_W \leftarrow \{r_\sigma r_{\sigma+k} | 1 \leq |k| \leq \lfloor \frac{t_\alpha}{2} \rfloor, 1 \leq \sigma + k \leq q\}$
Initialize j=1
while $j \leq \frac{q-\sigma}{t_\alpha}$ **do**
  $\ell := \sigma + t_\alpha j$
  $E_W \leftarrow E_W \cup \{r_\ell r_{\ell+k} | 1 \leq |k| \leq \lfloor \frac{t_\alpha}{2} \rfloor, 1 \leq \ell + k \leq q\}$
  $E_W \leftarrow E_W \cup \{r_{\sigma+t_\alpha(j-1)} r_{\sigma+t_\alpha j}\}$
  $j \leftarrow j + t_\alpha$
if $\sigma > \lceil \frac{t_\alpha}{2} \rceil$ **then**
  $E_W \leftarrow E_W \cup \{r_1 r_k | 2 \leq k \leq \sigma - \lceil \frac{t_\alpha}{2} \rceil\} \cup \{r_1 r_\sigma\}$
$j \leftarrow \lfloor \frac{q-\sigma}{t_\alpha} \rfloor$
if $\sigma + t_\alpha j \leq q - \lceil \frac{t_\alpha}{2} \rceil$ **then**
  $E_W \leftarrow E_W \cup \{r_k r_q | \sigma + t_\alpha j + \lceil \frac{t_\alpha}{2} \rceil \leq k \leq q - 1\} \cup \{r_{\sigma+t_\alpha j} r_q\}$
Return W

Under this random scheme, we define $\lambda_L(t_\alpha) := \max_{e \in L} \mathbb{E}[H_{\bar{w}(e)}]$, and $\lambda_O(t_\alpha) := \max_{e \in O} \mathbb{E}[H_{\bar{w}(e)}]$.

▶ **Lemma 16.** *For any $\alpha \geq 0$, $\lambda_L(t_\alpha) \leq \frac{1}{t_\alpha} H_{t_\alpha+1} + \frac{t_\alpha-1}{t_\alpha}$, and $\lambda_O(t_\alpha) \leq \frac{1}{t_\alpha} + \frac{2}{t_\alpha}\sum_{i=2}^{\lceil \frac{t_\alpha}{2} \rceil} H_i$.*

Proof. Let $W = (R, E_W)$ be a witness tree returned from running Algorithm 2 with α and $t := t_\alpha$, and let w be the vector imposed on E^* by W. If Algorithm 2 samples $\sigma \in \{1, \dots, t\}$, then we say that the terminals $r_{\sigma+tj}$ are *marked* by the algorithm. Moreover, if $\sigma > \lceil \frac{t_\alpha}{2} \rceil$ (resp. $\sigma + t_\alpha \lfloor \frac{q-\sigma}{t_\alpha} \rfloor \leq q - \lceil \frac{t_\alpha}{2} \rceil$) then r_1 (resp. r_q) is also considered marked.

1. Consider edge $e = s_j s_{j+1} \in O$, with $j \in \{\lceil \frac{t}{2} \rceil, \dots, q - \lceil \frac{t}{2} \rceil\}$. Let $m \in \{j - \lfloor \frac{t}{2} \rfloor, \dots, j + \lfloor \frac{t}{2} \rfloor\}$, such that $\sigma \mod t = m \mod t$. Observe that in this case r_m is marked. If $m = j - x$ for $x \in \{0, \dots, \lfloor \frac{t}{2} \rfloor\}$, then $w(s_j s_{j+1}) = \lceil \frac{t}{2} \rceil - x$. Similarly if $m = j + x$ for $x \in \{1, \dots, \lfloor \frac{t}{2} \rfloor\}$, then $w(s_j s_{j+1}) = \lceil \frac{t}{2} \rceil - x + 1$. Since $m \mod t = \sigma \mod t$ with probability $\frac{1}{t}$, we have $\mathbb{E}[H_{w(s_j s_{j+1})}] = \frac{1}{t} + \frac{2}{t}\sum_{k=2}^{\lceil \frac{t}{2} \rceil} H_k$.
   Now assume $j < \lceil \frac{t}{2} \rceil$ (the case $j > q - \lceil \frac{t}{2} \rceil$ can be handled similarly). Recalling that since t is odd it is not hard to determine the value of $w(s_j s_{j+1})$ by cases, depending on the value of σ.
   a. $1 \leq \sigma \leq j$: Then $w(s_j s_{j+1}) = \lceil \frac{t}{2} \rceil + \sigma - j$.
   b. $j + 1 \leq \sigma \leq \lceil \frac{t}{2} \rceil$: Then $w(s_j s_{j+1}) = j$.
   c. $\lceil \frac{t}{2} \rceil + 1 \leq \sigma \leq j + \lfloor \frac{t}{2} \rfloor$: Then $w(s_j s_{j+1}) = \lceil \frac{t}{2} \rceil - \sigma + j + 1$.
   d. $j + \lceil \frac{t}{2} \rceil \leq \sigma \leq t$: Then $w(s_j s_{j+1}) = \sigma - j - \lceil \frac{t}{2} \rceil + 1$.

$$
\begin{aligned}
&\mathbb{E}[H_{w(s_j s_{j+1})}] = \\
&= \frac{1}{t}\left(\sum_{\sigma=1}^{j} H_{\lceil \frac{t}{2} \rceil + \sigma - j} + \sum_{\sigma=j+1}^{\lceil \frac{t}{2} \rceil} H_j + \sum_{\sigma=\lceil \frac{t}{2} \rceil + 1}^{j + \lfloor \frac{t}{2} \rfloor} H_{\lceil \frac{t}{2} \rceil - \sigma + j + 1} + \sum_{\sigma=j+\lceil \frac{t}{2} \rceil}^{t} H_{\sigma - j - \lceil \frac{t}{2} \rceil + 1}\right) \\
&= \frac{1}{t}\left(\sum_{i=\lceil \frac{t}{2} \rceil - j + 1}^{\lceil \frac{t}{2} \rceil} H_i + \left(\left\lceil \frac{t}{2} \right\rceil - j\right) H_j + \sum_{i=2}^{j} H_i + \sum_{i=1}^{\lceil \frac{t}{2} \rceil - j} H_i\right) \\
&= \frac{1}{t}\left(\sum_{i=1}^{\lceil \frac{t}{2} \rceil} H_i + \left(\left\lceil \frac{t}{2} \right\rceil - j\right) H_j + \sum_{i=2}^{j} H_i\right) < \frac{1}{t}\left(1 + 2\sum_{i=2}^{\lceil \frac{t}{2} \rceil} H_i\right).
\end{aligned}
$$

2. Consider edge $e = s_j r_j \in L$. We first show the bound for $j \in \{1, \dots, q\}$. Algorithm 2 marks terminal r_i with probability $\frac{1}{t}$. If r_i is marked, then $w(e) \leq t$. If r_i is not marked, then $w(e) = 1$. Therefore, $\mathbb{E}[H_{w(e)}] \leq \frac{1}{t} H_{t+1} + \frac{t-1}{t}$
   Now consider edge $e = s_1 r_1$ (the case $e = s_q r_q$ can be handled similarly). We consider specific values of $\sigma \in \{1, \dots, t\}$ sampled by Algorithm 2. With probability $\frac{1}{t}$, we have $\sigma = 1$, so r_1 is marked initially and $w(e) = \lceil t/2 \rceil$. For $\sigma = 2, \dots, \lceil t/2 \rceil$, r_1 is unmarked and $w(e) = 1$. If $\sigma > \lceil t/2 \rceil$, then r_1 is marked by the algorithm and $w(e) = \sigma - \lceil t/2 \rceil$. Therefore, we can see

$$
\mathbb{E}[H_{w(r_1 s_1)}] = \frac{1}{t}\left(H_{\lceil t/2 \rceil} + \left\lfloor \frac{t}{2} \right\rfloor + \sum_{k=1}^{t - \lceil t/2 \rceil} H_k\right)
$$

   We let $g(t)$ be equal to the equality above. It remains to show that $g(t) \leq \frac{1}{t} H_{t+1} + \frac{t-1}{t} := f(t)$ for $t \in \{1, 3, 5\}$.

$$
\begin{aligned}
g(1) &= H_1 = 1 < H_2 = f(1) \\
g(3) &= \frac{1}{3}(H_2 + 1 + H_1) = 1.1\bar{6} < 1.36\bar{1} = \frac{1}{3}(H_4 + 2) = f(3) \\
g(5) &= \frac{1}{5}(H_3 + 2 + H_1 + H_2) = 1.2\bar{6} < 1.29 = \frac{1}{5}(H_6 + 4) = f(5)
\end{aligned}
$$

Combining these two facts gives us the bound on $\lambda_{L_i}(t)$, for $t \in \{1, 3, 5\}$. ◀

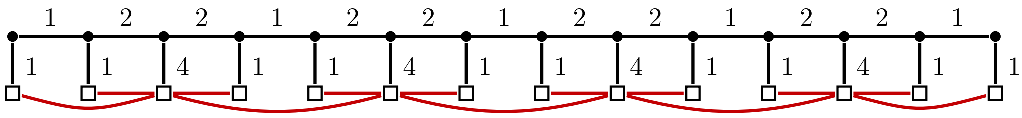

Figure 6 Lower bound instance shown in black with $c(e) = 1$ for all the edges in L and $c(e) = \alpha$ for all the edges in O, for $\alpha = \frac{32}{90}$. The white squares are terminals and black circles are Steiner nodes. Red edges form the laminar witness tree W^*, with the numbers next to each edge the value of w imposed on T.

The following Lemma is proven in the full version of the paper.

▶ **Lemma 17.** *For any $\alpha \geq 0$, the following bounds holds:*

$$\frac{1}{\alpha+1}\left(\frac{1}{t_\alpha}H_{t_\alpha+1} + \frac{t_\alpha - 1}{t_\alpha} + \alpha\left(\frac{1}{t_\alpha} + \frac{2}{t_\alpha}\sum_{i=2}^{\lceil \frac{t_\alpha}{2} \rceil} H_i\right)\right) \leq \frac{991}{732}$$

We are now ready to prove the following:

▶ **Lemma 18.** $\mathbb{E}[\bar{\nu}_T(W)] \leq \frac{991}{732}$.

Proof. One observes:

$$\sum_{e \in L \cup O} c(e)\mathbb{E}[H_{\bar{w}(e)}] \leq \sum_{e \in L} c(e)\lambda_L(t_\alpha) + \sum_{e \in O} c(e)\lambda_O(t_\alpha) = (\lambda_L(t_\alpha) + \alpha\lambda_O(t_\alpha))\sum_{e \in L} c(e)$$

Therefore $\mathbb{E}[\nu_T(W)]$ is bounded by:

$$\frac{\sum_{e \in L \cup O} c(e)\mathbb{E}[H_{\bar{w}(e)}]}{\sum_{e \in L \cup O} c(e)} \leq \frac{(\lambda_L(t_\alpha) + \alpha\lambda_O(t_\alpha))\sum_{e \in L} c(e)}{(\alpha+1)\sum_{e \in L} c(e)} = \frac{\lambda_L(t_\alpha) + \alpha\lambda_O(t_\alpha)}{\alpha+1} \leq \frac{991}{732}.$$

where the last inequality follows using Lemma 16 and 17. ◀

Now Theorem 5 follows by combining Lemma 18 with Theorem 1 in which γ is replaced by the supremum taken over all Steiner-claw free instances (rather than over all Steiner Tree instances).

Tightness of the bound

We conclude this section by spending a few words on Theorem 6. Our lower-bound instance is obtained by taking a tree T on q Steiner nodes, each adjacent to one terminal, with $c(e) = 1$ for all the edges in L and $c(e) = \alpha$ for all the edges in O, for $\alpha = \frac{32}{90}$. Similar to Section 3, a crucial ingredient for our analysis is in utilizing Theorem 8 stating that there is an optimal laminar witness tree. See Figure 6. We use this to show that there is an optimal witness tree for our tree T, whose objective value is at least $\frac{991}{732} - \varepsilon$. Details can be found in the full version of the paper.

References

1 David Adjiashvili. Beating approximation factor two for weighted tree augmentation with bounded costs. *ACM Trans. Algorithms*, 15(2):19:1–19:26, 2019.

2 Haris Angelidakis, Dylan Hyatt-Denesik, and Laura Sanità. Node connectivity augmentation via iterative randomized rounding. *Mathematical Programming*, pages 1–37, 2022.

3 Manu Basavaraju, Fedor V. Fomin, Petr A. Golovach, Pranabendu Misra, M. S. Ramanujan, and Saket Saurabh. Parameterized algorithms to preserve connectivity. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 800–811, 2014.

4 Jaroslaw Byrka, Fabrizio Grandoni, and Afrouz 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)*, pages 815–825, 2020.

5 Jaroslaw Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. Steiner tree approximation via iterative randomized rounding. *J. ACM*, 60(1):6:1–6:33, 2013.

6 Federica Cecchetto, Vera Traub, and Rico 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)*, pages 370–383. ACM, 2021.

7 Joseph Cheriyan and Zhihan Gao. Approximating (unweighted) tree augmentation via lift-and-project, part I: stemless TAP. *Algorithmica*, 80(2):530–559, 2018.

8 Joseph Cheriyan and Zhihan Gao. Approximating (unweighted) tree augmentation via lift-and-project, part II. *Algorithmica*, 80(2):608–651, 2018.

9 Efim A Dinitz, Alexander V Karzanov, and Michael V Lomonosov. On the structure of the system of minimum edge cuts in a graph. *Issledovaniya po Diskretnoi Optimizatsii*, pages 290–306, 1976.

10 Andreas Emil Feldmann, Jochen Könemann, Neil Olver, and Laura Sanità. On the equivalence of the bidirected and hypergraphic relaxations for steiner tree. *Mathematical programming*, 160(1):379–406, 2016.

11 Samuel Fiorini, Martin Groß, Jochen Könemann, and Laura Sanità. Approximating weighted tree augmentation via chvátal-gomory cuts. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 817–831. SIAM, 2018.

12 Greg N. Frederickson and Joseph JáJá. Approximation algorithms for several graph augmentation problems. *SIAM J. Comput.*, 10(2):270–283, 1981.

13 Michel X. Goemans, Neil Olver, Thomas Rothvoß, and Rico Zenklusen. Matroids and integrality gaps for hypergraphic steiner tree relaxations. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 1161–1176, New York, NY, USA, 2012. Association for Computing Machinery. doi:10.1145/2213977.2214081.

14 Fabrizio Grandoni, Christos Kalaitzis, and Rico Zenklusen. Improved approximation for tree augmentation: saving by rewiring. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 632–645. ACM, 2018.

15 Dylan Hyatt-Denesik, Afrouz Jabal Ameli, and Laura Sanità. Finding almost tight witness trees, 2023. arXiv:2211.12431.

16 Zeev Nutov. 2-node-connectivity network design. In *Proceedings of the 18th International Workshop on Approximation and Online Algorithms (WAOA)*, volume 12806 of *Lecture Notes in Computer Science*, pages 220–235. Springer, 2020.

17 Zeev Nutov. Approximation algorithms for connectivity augmentation problems. In *Proceedings of the 16th International Computer Science Symposium in Russia (CSR)*, volume 12730, pages 321–338. Springer, 2021.

18 Vera Traub and Rico Zenklusen. A $(1.5+\varepsilon)$-approximation algorithm for weighted connectivity augmentation, 2022. doi:10.48550/arXiv.2209.07860.

19 Vera Traub and Rico Zenklusen. A better-than-2 approximation for weighted tree augmentation. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1–12. IEEE, 2022.

20 Vera Traub and Rico Zenklusen. Local search for weighted tree augmentation and steiner tree. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3253–3272. SIAM, 2022.