A Polylogarithmic Approximation for Computing Non-Metric Terminal Steiner Trees

Iftah Gamzu* Danny Segev†

Abstract

The main contribution of this short note is to provide improved bounds on the approximability of constructing terminal Steiner trees in arbitrary undirected graphs. Technically speaking, our results are obtained by relating this computational task to that of computing group Steiner trees. As a secondary objective, we make a concentrated effort to distinguish between the factor by which constructed trees exceed the optimal backbone cost and between the deviation from the optimal terminal linking cost.

*Blavatnik School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Email: iftgam@tau.ac.il. Supported by the Israel Science Foundation, by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848, by a European Research Council (ERC) Starting Grant, and by the Wolfson Family Charitable Trust.

†Department of Statistics, University of Haifa, Haifa 31905, Israel. Email: segevd@stat.haifa.ac.il.

1 Introduction

An instance of the *terminal Steiner tree* problem consists of an undirected graph $G = (V, E)$ on n vertices, with non-negative edge costs specified by $c : E \to \mathbb{R}_+$, and a subset of vertices $\mathcal{T} = \{t_1, \ldots, t_k\}$, which we refer to as terminals. The objective is to identify a minimum cost tree $\mathcal{H} \subseteq G$ spanning $\mathcal{T}$; however, some Steiner trees will not do the trick, since we also have to take into account the following structural requirement: terminals cannot serve as intermediate vertices in $\mathcal{H}$, or in other words, each terminal is required to be a leaf. This structural constraint is motivated by real-life applications in computational biology, VLSI design, and networking. We refer the reader to directly related papers [9, 10, 3] and to the references therein for a comprehensive review of these applications.

The metric case. There has been a growing line of work on *metric* terminal Steiner trees, in which we are given a complete graph whose edge costs satisfy the triangle inequality. Lin and Xue [9] seem to have been the first to consider this variant, for which they obtained a $(2+\rho)$-approximation; here, ρ denotes the best approximation ratio attainable in polynomial time for the standard Steiner tree problem. Later on, several authors [2, 5, 3] independently established an improved performance guarantee of 2ρ. Currently, the best known approximation ratio is $2\rho - \rho/(3\rho-2)$, which was achieved by Martinez, de Pina and Soares [11]. This factor evaluates to roughly 2.14 by plugging in $\rho = \ln 4 + \epsilon < 1.39$ [1].

The non-metric case. We study the approximability of computing terminal Steiner trees in *arbitrary undirected graphs*. That is, the underlying graph is no longer required to be complete, and its edge costs do not necessarily satisfy the triangle inequality. In spite of appearance, this setting cannot be sensibly reduced to the metric case. To better understand the latter statement, note that one cannot straightforwardly utilize the shortest-path metric induced by (G, c) since replacing single edges by their corresponding shortest paths may quite possibly violate the degree constraints of terminal vertices. Degree bounds are not intrinsic requirements in the standard Steiner tree problem, for which such edge-to-path replacements are possible (see, for instance, [12, Sec. 3]).

To our knowledge, the only non-trivial hardness result for the non-metric case is due to Drake and Hougardy [3], who proved that this problem cannot be approximated within a factor of $(1-\epsilon)\ln n$, for any fixed $\epsilon > 0$, unless NP $\subseteq$ DTIME$(n^{O(\log\log n)})$. On the positive side, Fukunaga [6, Sec. 6.5] obtained a performance guarantee of 2Δ, where Δ is the maximum degree of any terminal. However, simple examples demonstrate that this factor may be $\Omega(n)$ in the worst case. Consequently, approximating the terminal Steiner tree problem when one makes no structural assumptions about edge costs is still an open question.

Main results. The primary contribution of this short note is to provide improved bounds on the approximability of constructing terminal Steiner trees by relating this computational task to that of computing group Steiner trees. As a secondary objective, we make a concentrated effort to distinguish between the factor by which constructed trees exceed the optimal backbone cost and between the deviation from the optimal terminal linking cost; both of these measures will be defined later on. Our findings can be briefly summarized as follows:

- We present a randomized algorithm that constructs, with constant probability, a terminal Steiner tree whose cost is within a factor of $O(\log n \log k \log \Delta)$ of optimal[1]. The specifics of our approach, along with a refined statement of the latter bound (in terms of backbone and terminal linking costs), appears in Section 2.

[1] This improves on the upper bound of Fukunaga [6, Sec. 6.5] for large enough values of Δ, or more precisely, whenever $\Delta/\log\Delta = \omega(\log n \log k)$.

- We establish an $\Omega(\log^{2-\epsilon} k)$ hardness of approximation, for any fixed $\epsilon > 0$, unless $\text{NP} \subseteq \text{ZTIME}(n^{\text{polylog}(n)})$. Further details are given in Section 3.

Group Steiner trees: The bare necessities. Prior to delving into technicalities, some background on the *group Steiner tree* problem is necessary. In this setting, we are given an edge-weighted undirected graph $G = (V, E)$ and a family of vertex groups $g_1, \ldots, g_k \subseteq V$. The goal is to identify a minimum cost subgraph spanning at least one representative from each group. For our purposes, it would be sufficient to mention that Garg, Konjevod and Ravi [7] were the first to achieve polylogarithmic approximability results, through randomized LP-rounding on tree instances, which attains an approximation ratio of $O(\log k \log N)$, where $N = \max_i |g_i|$. Combining this approach with the probabilistic embedding method of Fakcharoenphol, Rao and Talwar [4] implies a performance guarantee of $O(\log n \log k \log N)$ for arbitrary graphs. From an inapproximability point of view, Halperin and Krauthgamer [8] demonstrated that the group Steiner tree problem cannot be approximated within a factor of $O(\log^{2-\epsilon} k)$, for any fixed $\epsilon > 0$, unless $\text{NP} \subseteq \text{ZTIME}(n^{\text{polylog}(n)})$.

2 A Polylogarithmic Approximation

To provide an accurate description of our main algorithmic result, we note that the optimal terminal Steiner tree consists of two separate components: Edges incident on terminals, whose combined cost is denoted by OPT_L (henceforth, the optimal *terminal linking cost*), and the tree obtained by removing these edges, whose cost is denoted by OPT_B (the optimal *backbone cost*). The primary finding of this section can be formally stated as follows.

Theorem 2.1. *There is a randomized algorithm that constructs, with constant probability, a terminal Steiner tree whose backbone cost is $O(\log n \log k \log \Delta) \cdot \text{OPT}_B$. At the same time, the terminal linking cost is only $O(\log k) \cdot \text{OPT}_L$.*

2.1 Preliminaries

We assume without loss of generality that the collection of terminals $\mathcal{T}$ forms an independent set in G. This assumption can be easily justified by observing that an edge joining two terminals cannot be included in any feasible solution, unless $|\mathcal{T}| = 2$. However, this special case can be solved to optimality by applying a shortest path algorithm. We now introduce some notation and terminology:

- Let $\mathcal{H}^* \subseteq G$ denote the optimal terminal Steiner tree, of cost $\text{OPT} = \text{OPT}_B + \text{OPT}_L$.
- We assume that an arbitrary root vertex $r \in V(\mathcal{H}^*) \setminus \mathcal{T}$ in the backbone of $\mathcal{H}^*$, and a constant factor estimate $\mathcal{E}_L \in [\text{OPT}_L, 2 \cdot \text{OPT}_L]$ of the optimal terminal linking cost, are known in advance. This assumption can be enforced by considering polynomially-many potential candidates for r and $\mathcal{E}_L$.
- For every $t_i \in \mathcal{T}$, let $\mathcal{G}_i$ be the *group* of this terminal, consisting of vertices reachable from t_i via a single edge of cost at most $2\mathcal{E}_L/k$, that is, $\mathcal{G}_i = \{v \in V : c(t_i, v) \leq 2\mathcal{E}_L/k\}$.

2.2 Overview

In the remainder of this section, we describe a randomized procedure that constructs, with constant probability, a tree $\mathcal{H} \subseteq G$ satisfying the following properties:

1. $\mathcal{H}$ spans the root r, as well as $\Omega(k)$ terminals.
2. All terminals spanned by $\mathcal{H}$ are necessarily leaves.

3. $\mathcal{H}$ has a backbone cost of $O(\log n \log \Delta) \cdot \text{OPT}_B$ and a terminal linking cost of $O(1) \cdot \text{OPT}_L$.

A procedure of this nature allows one to instantly validate our main result.

Proof of Theorem 2.1. Suppose we repeatedly construct $O(\log k)$ such trees (say, $\mathcal{H}_1, \mathcal{H}_2, \ldots$), focusing in each new iteration only on terminals that have not been spanned up until now. That is, once a terminal t_i is spanned by some tree, we immediately remove it from the current graph. By amplifying the success probability of each iteration via $O(\log \log k)$ independent repetitions, we can ensure that, with constant probability, each and every terminal is indeed spanned by the union of all constructed trees, $\widetilde{\mathcal{H}} = \bigcup_j \mathcal{H}_j$. Needless to say, property 1 implies that $\widetilde{\mathcal{H}}$ must be connected, as each tree $\mathcal{H}_j$ spans r. In addition, property 2, in conjunction with the observation that every terminal may appear in at most one tree, guarantee that this subgraph can be converted into a terminal Steiner tree, since every possible cycle in $\widetilde{\mathcal{H}}$ consists solely of non-terminal vertices. Finally, property 3 ensures that the backbone cost of $\widetilde{\mathcal{H}}$ is $O(\log n \log k \log \Delta) \cdot \text{OPT}_B$, whereas its terminal linking cost is only $O(\log k) \cdot \text{OPT}_L$. ■

2.3 The randomized procedure

Tree embedding. Let $\widetilde{G}$ be the subgraph of G created by removing the set of terminals $\mathcal{T}$. We begin by probabilistically embedding the shortest-path metric induced by $(\widetilde{G}, c)$ into a random tree $T' = (V', E')$ using the method of Fakcharoenphol et al. [4], incurring an expected stretch of $O(\log n)$. More specifically, the distance in T' between any pair of original vertices is at least as large as the corresponding distance in $\widetilde{G}$, and the expected ratio between these distances is $O(\log n)$, where the expectation is with respect to the distribution over tree metrics.

Fractionally connecting $\boldsymbol{\Omega(k)}$ groups. We now focus our attention on identifying a low-cost subtree of T' that connects r to a fixed proportion of the groups $\mathcal{G}_1, \ldots, \mathcal{G}_k$. For this purpose, consider the following linear program, which is an adaptation of the cut-based LP-relaxation of group Steiner tree:

$$
\begin{array}{llll}
(\text{LP}_1) & \text{minimize} & \displaystyle\sum_{e \in E'} c(e) x_e & \\
& \text{subject to} & (1) \quad \displaystyle\sum_{e \in \delta(U)} x_e \geq y_i & \forall\, U \subseteq V',\ \forall i \in [k]: \\
& & & r \in U,\ \mathcal{G}_i \cap U = \emptyset \\
& & (2) \quad \displaystyle\sum_{i \in [k]} y_i \geq k/2 & \\
& & (3) \quad x_e,\ y_i \in [0,1] & \forall\, e \in E',\ \forall\, i \in [k]
\end{array}
$$

In an integral solution, the variable x_e indicates whether the edge e is picked, and y_i indicates whether r is connected to some representative of $\mathcal{G}_i$. Constraint (1) guarantees that, if r is connected to $\mathcal{G}_i$, at least one edge is picked from every cut $(U, V' \setminus U)$ separating r and $\mathcal{G}_i$; here, $\delta(U)$ denotes the set of edges crossing $(U, V' \setminus U)$. Constraint (2) ensures that at least $k/2$ of the groups are connected. We remark that although (LP_1) has exponentially many constraints, it admits a polynomial-time separation oracle based on a minimum cut procedure (see, for instance, [7]).

Lemma 2.2. *With constant probability,* $\text{OPT}(\text{LP}_1) = O(\log n) \cdot \text{OPT}_B$.

Proof. One can easily argue that the backbone of $\mathcal{H}^*$ spans representatives from at least $k/2$ of the groups $\mathcal{G}_1, \ldots, \mathcal{G}_k$. Otherwise, the individual linking cost of at least $k/2$ terminals is strictly greater than $2\mathcal{E}_L/k$, implying that the terminal linking cost of $\mathcal{H}^*$ is greater than $\mathcal{E}_L \geq \text{OPT}_L$. Consequently, when each backbone edge (u, v) is translated to the unique u-v path in T', we obtain a subtree that connects r to at least $k/2$ groups. Moreover, since the expected stretch is $O(\log n)$, a straightforward application of Markov's inequality implies that the subtree cost is $O(\log n) \cdot \mathrm{OPT}_B$ with constant probability. The lemma follows by observing that $\mathrm{OPT}(\mathrm{LP}_1)$ provides a lower bound on the cost of any feasible integral solution. ■

Creating a group Steiner tree instance. Let (x^*, y^*) be an optimal fractional solution to (LP_1), and let $I^* = \{i : y_i^* \geq 1/4\}$ be the index set of groups that are fractionally connected to the extent of at least 1/4. Note that constraint (2) implies $|I^*| \geq k/3$, since

$$\frac{k}{2} \leq \sum_{i \in [k]} y_i^* = \sum_{i \in I^*} y_i^* + \sum_{i \notin I^*} y_i^* \leq |I^*| + \frac{(k - |I^*|)}{4} .$$

We proceed by setting up a group Steiner tree instance in which the groups are $\{\mathcal{G}_i : i \in I^*\}$; the objective is to connect r to at least one representative of each and every group. Consider the cut-based relaxation of this instance, formally defined as follows.

$$\begin{array}{llll} (\mathrm{LP}_2) & \text{minimize} & \displaystyle\sum_{e \in E'} c(e) x_e & \\ & \text{subject to} & \displaystyle\sum_{e \in \delta(U)} x_e \geq 1 & \begin{array}{l}\forall\, U \subseteq V', \forall i \in I^* : \\ r \in U,\, \mathcal{G}_i \cap U = \emptyset\end{array} \\ & & x_e \in [0,1] & \forall\, e \in E' \end{array}$$

Notice that the main constraint in (LP_2) is nearly identical to the one in (LP_1), with an additional restriction stating that $y_i = 1$ if $i \in I^*$, and $y_i = 0$ otherwise. With this observation in mind, it is easy to verify that the vector $\hat{x}$, defined by $\hat{x}_e = \min\{4x_e^*, 1\}$, constitutes a feasible solution to (LP_2), since $y_i^* \geq 1/4$ for every $i \in I^*$. As a result, we immediately have $\mathrm{OPT}(\mathrm{LP}_2) \leq 4 \cdot \mathrm{OPT}(\mathrm{LP}_1)$.

Putting it all together. At this point in time, we remark that, given an optimal solution to (LP_2), the randomized rounding procedure of Garg et al. [7] generates an r-rooted subtree of T' satisfying the following properties:

1. For every $i \in I^*$, at least one representative of $\mathcal{G}_i$ is spanned with probability $\Omega(1/\log|\mathcal{G}_i|)$.

2. The expected cost of this subtree is $\mathrm{OPT}(\mathrm{LP}_2) = O(1) \cdot \mathrm{OPT}(\mathrm{LP}_1)$.

Now let $\mathcal{H} \subseteq T'$ be a subtree created by unifying $O(\log \Delta)$ such random subtrees. Since $\Delta \geq \max_{i \in I^*} |\mathcal{G}_i|$ and $|I^*| \geq k/3$, it follows that the expected number of spanned groups is $\Omega(k)$. In addition, the expected cost of $\mathcal{H}$ is $O(\log \Delta) \cdot \mathrm{OPT}(\mathrm{LP}_1)$. Therefore, based on Lemma 2.2 and simple probabilistic arguments, we can safely claim that, with constant probability, $\Omega(k)$ groups will be spanned at a combined cost of $O(\log n \log \Delta) \cdot \mathrm{OPT}_B$. We proceed by translating $\mathcal{H}$ to an analogous tree in the original graph G, creating a backbone of no greater cost. Finally, we identify a representative of each group $\mathcal{G}_i$ spanned by this backbone, and connect it to the corresponding terminal t_i. As each terminal can be linked to any representative of its group by an edge of cost at most $2\mathcal{E}_L/k$, we conclude that the terminal linking cost is $O(k) \cdot 2\mathcal{E}_L/k = O(1) \cdot \mathrm{OPT}_L$.

3 Polylogarithmic Hardness

In what follows, we establish a polylogarithmic hardness of approximation for the terminal Steiner tree problem on arbitrary undirected graphs. This result is obtained via a simple reduction from the group Steiner tree problem.

Theorem 3.1. *The terminal Steiner tree problem cannot be approximated within a factor of $O(\log^{2-\epsilon} k)$, for any fixed $\epsilon > 0$, unless $\mathrm{NP} \subseteq \mathrm{ZTIME}(n^{\mathrm{polylog}(n)})$.*

Proof. Given a group Steiner tree instance, as described in Section 1, we define an instance of the terminal Steiner tree problem as follows. The underlying graph is augmented with k new vertices, $t_1, \ldots, t_k$, which constitute the set of terminals. In addition, zero-cost edges join each terminal t_i to every vertex in its corresponding group g_i. Figure 1 provides a concrete example of this construction.

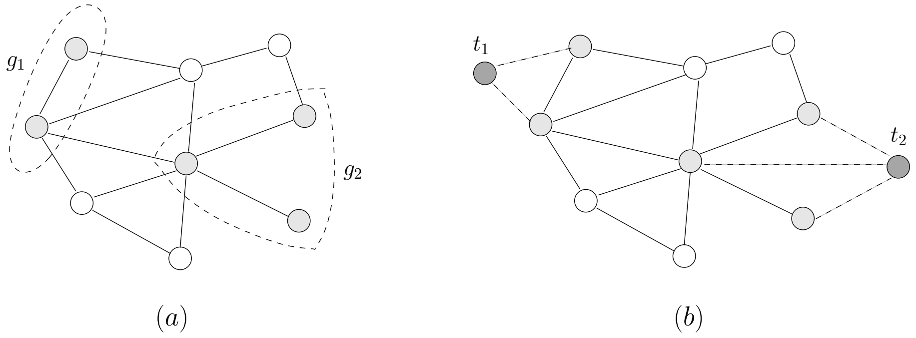

Figure 1: (a) An instance of the group Steiner tree problem with two groups; (b) The resulting terminal Steiner tree instance (edge costs are not indicated).

We begin by pointing out that a group Steiner tree $\mathcal{H}$ in the original instance can be converted to a terminal Steiner tree of identical cost in the newly-created instance. For this purpose, simply add an edge connecting each terminal t_i to some representative of g_i spanned by $\mathcal{H}$. Conversely, it is not difficult to verify that, given a terminal Steiner tree $\mathcal{H}$, we can perform a similar cost-preserving transformation in the opposite direction. Clearly, $\mathcal{H}$ must span at least one representative from g_i, as there is no other way to access t_i; furthermore, noting that every terminal is a leaf of $\mathcal{H}$, connectivity is still maintained when each terminal is discarded along with its incident edge. ■

References

[1] J. Byrka, F. Grandoni, T. Rothvoß, and L. Sanità. An improved LP-based approximation for Steiner tree. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 583–592, 2010.

[2] Y. H. Chen, C. L. Lu, and C. Y. Tang. On the full and bottleneck full Steiner tree problems. In *Proceedings of the 9th Annual International Computing and Combinatorics Conference*, pages 122–129, 2003.

[3] D. E. Drake and S. Hougardy. On approximation algorithms for the terminal Steiner tree problem. *Information Processing Letters*, 89(1):15–18, 2004.

[4] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 69(3):485–497, 2004.

[5] B. Fuchs. A note on the terminal Steiner tree problem. *Information Processing Letters*, 87(4):219–220, 2003.

[6] T. Fukunaga. *Approximation algorithms to the network design problems.* PhD thesis, Graduate School of Informatics, Kyoto University, 2007.

[7] N. Garg, G. Konjevod, and R. Ravi. A polylogarithmic approximation algorithm for the group Steiner tree problem. *Journal of Algorithms*, 37(1):66–84, 2000.

[8] E. Halperin and R. Krauthgamer. Polylogarithmic inapproximability. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 585–594, 2003.

[9] G.-H. Lin and G. Xue. On the terminal Steiner tree problem. *Information Processing Letters*, 84(2):103–107, 2002.

[10] C. L. Lu, C. Y. Tang, and R. C.-T. Lee. The full Steiner tree problem. *Theoretical Computer Science*, 306(1-3):55–67, 2003.

[11] F. V. Martinez, J. C. de Pina, and J. Soares. Algorithms for terminal Steiner trees. *Theoretical Computer Science*, 389(1-2):133–142, 2007.

[12] V. V. Vazirani. *Approximation Algorithms*. Springer, 2001.