

Improved Bounds for the Online Steiner Tree Problem in Graphs of Bounded Edge-Asymmetry

Spyros Angelopoulos

David R. Cheriton School of Computer Science, University of Waterloo
sangelop@cs.uwaterloo.ca

Technical Report CS-2006-36

David R. Cheriton School of Computer Science

Abstract

In this paper we consider the Online Steiner Tree problem in weighted directed graphs of bounded edge-asymmetry α . The edge-asymmetry of a directed graph is defined as the maximum ratio of the cost (weight) of antiparallel edges in the graph. The problem has applications in multicast routing over a network with non-symmetric links. We improve the previously known upper and lower bounds on the competitive ratio of any deterministic algorithm due to Faloutsos *et al.* [10]. In particular, we show that a better analysis of a simple greedy algorithm yields a competitive ratio of $O(\min\{k, \frac{\alpha \log k}{\log \log \alpha}\})$, where k denotes the number of terminals requested. On the negative side, we show a lower bound of $\Omega(\min\{k^{1-\epsilon}, \frac{\alpha \log k}{\log \log k}\})$ on the competitive ratio of every deterministic algorithm for the problem, for any arbitrarily small constant ϵ .

1 Introduction

The *Steiner Tree* problem in undirected graphs is defined as follows. Given an undirected graph $G = (V, E)$ with a cost function $c : E \rightarrow \mathbb{R}^+$ on the edges, and a subset of vertices $K \subseteq V$ with $|K| = k$, (also called *terminals*), the goal is to find a minimum-cost tree which spans all vertices in K . The cost of the tree is defined as the sum of the costs of its edges. When the input graph is *directed*, the input to the problem must specify, in addition to G and K a specific vertex $r \in V$ called the *root*. The problem then translates to finding a minimum cost *arborescence* rooted at r which spans all vertices in K .

In the *online* version of the problem, terminals in K are requested in an online, sequential fashion. Every time a request (terminal) $t \in V$ arrives, the algorithm must ensure that there is a path from an earlier requested terminal to t , for the undirected version, or a directed path from r to t in the directed version of the problem, respectively. We assume that the graph G is known to the algorithm. In the standard framework of competitive analysis (see, e.g., [6]), the goal is to design online algorithms of small competitive ratio. In the context of Steiner tree problems the competitive ratio is defined as the supremum of the ratio of the cost of the tree (or arborescence) produced by the algorithm over the optimal off-line cost (namely the cost of an offline algorithm which has complete knowledge of the request set K).

Both the offline and online versions of the problem have been studied extensively in the literature (c.f. section 1.1 for some representative results concerning the online version) and are often encountered in the context of several combinatorial optimization problems. In addition to interest from the point of view of theoretical analysis, the problem has important applications in the design of multicast protocols in computer networks, which involves distribution of the same information stream to the members of the multicast group over an existing network. Indeed, multicasting can be modeled as the problem of selecting communication links (edges of the underlying graph) so as to minimize the cost for supporting multicast routing through a tree, which is essentially identical to the Steiner tree problem formulation. For the interplay between the Steiner Tree problem and multicast applications, the interested reader is referred to the work of Faloutsos [9].

Most of the theoretical work on the Steiner tree problem and its generalizations is focused on undirected graphs (see eg [14] [1] [11] [5] [16] for some representative results concerning the offline version of the problem). On the other hand, research considering directed underlying graphs has been relatively limited (see e.g. [7] [15], once again for the offline case). However, a directed graph is a more appropriate and realistic representation of a real-life network. As argued in [13], [8] studies on network traffic on backbones have revealed marked asymmetry in link utilization. For instance, one should expect that a typical subscriber to a home internet-cable service will incur more traffic on the incoming link “download”), than the outgoing link (“upload”). Moreover [13] if the link is wireless, its quality/bandwidth is inherently asymmetric, due to differences in noise levels, power of transmission and mobility levels of its endpoints.

Motivated by such observations, Ramanathan [13] proposed several metrics which are meant to capture deviation from the symmetry observed in undirected graphs. Perhaps the most intuitive metric is the so-called *maximum edge-asymmetry* α of a directed graph $G = (V, E)$ (or simply *asymmetry*, for the rest of this paper) which is defined as the maximum ratio of the costs of antiparallel links in the graph. To define this measure formally, let A denote the set of pairs of vertices in V , such that if the pair u, v is in A , then either $(v, u) \in E$ or $(u, v) \in E$ (i.e, there is an edge from u to v or an edge from v to u or both). Then the edge asymmetry is defined as

$$\alpha = \max_{\{v,u\} \in A} \frac{c(v, u)}{c(u, v)}$$

Note that undirected graphs can be seen as graphs of asymmetry $\alpha = 1$, while directed graphs in which there is at least one pair of vertices v, u such that $(v, u) \in E$, but $(u, v) \notin E$ are graphs with unbounded asymmetry (meaning that $\alpha = \infty$). Between these extreme cases, graphs with relatively small asymmetry model networks which are relatively homogenous in terms of the quality/characteristics of antiparallel links.

Ramanathan presented a 2α -approximation algorithm for the offline Steiner problem in graphs of asymmetry α . In addition, Ramanathan showed that the same approximation ratio is guaranteed when a different metric for capturing the graph asymmetry is applied, namely the ratio, over all adjacent pairs of vertices in the graph, of the sum of the larger edge-pair costs over the sum of the smaller edge-pair costs. In this paper we focus exclusively on the maximum-edge asymmetry only, since it represents a more clean-cut and easy to estimate measure of the graph edge asymmetry.

In subsequent work, Faloutsos Pankaj and Sevcik [10] studied the online Steiner tree problem in graphs of asymmetry α . They showed that a simple greedy algorithm (which we denote by GREEDY) has a competitive ratio of $O(\min\{k, \alpha \log k\})$. In particular, when a new terminal t is requested, GREEDY will find the directed path of minimum cost from some vertex in the current arborescence to t , and buy such edges. Once an edge is bought it is assigned zero cost in subsequent iterations, to reflect that the edge has irrevocably become part of the solution. Intuitively, the upper bound is not too difficult to derive: first, note that the cost of GREEDY is at most k times the optimal cost. Second, since GREEDY is $O(\log k)$ -competitive for undirected graphs (see Section 1.1) when we move to directed graphs, the competitive factor is multiplied by at most α (to account for the fact that the connection paths may choose the “wrong”, expensive direction). On the negative side, Faloutsos *et al.* showed a lower bound on the competitive ratio of $\Omega(\min\{k, \frac{\alpha \log k}{\log \alpha}\})$ for every deterministic algorithm¹. The construction for the lower bound is interesting, since not only indicates that the problem is not trivial, but also provides some intuition about a better analysis of the greedy algorithm.

In this paper we narrow the gap between the upper and lower bounds for the problem. In particular, we first provide a more elaborate analysis of GREEDY and prove the following:

Theorem 1 GREEDY is $O(\min\{k, \frac{\alpha \log k}{\log \log \alpha}\})$ -competitive.

On the negative side, we show the following lower bound on the competitive ratio of deterministic algorithms:

Theorem 2 *The competitive ratio of every deterministic online algorithm is $\Omega(\min\{k^{1-\epsilon}, \alpha \frac{\log k}{\log \log k}\})$, for every constant $0 < \epsilon < 1$.*

Theorem 2 in conjunction with the lower bound of [10] yields

Corollary 3 *The competitive ratio of every deterministic online algorithm is $\Omega(\max\{\min\{k, \frac{\alpha \log k}{\log \alpha}\}, \min\{k^{1-\epsilon}, \alpha \frac{\log k}{\log \log k}\}\})$ for every constant $0 < \epsilon < 1$.*

Our results improve the known bounds in a variety of situations, depending on the parameters α and k . In particular, consider graphs which are highly asymmetric, in the sense that there is a constant $c > 1$ such that $k = \alpha^c$, for some constant c . In this case Theorem 1 and Theorem 2 yield a tight bound of $\Theta(\frac{\alpha \log \alpha}{\log \log \alpha})$ whereas the analysis of [10] only shows the trivial bounds of $\Omega(\alpha)$ and $O(\alpha \log \alpha)$.

An outline of the intuition behind the proof of Theorem 1 is given in Section 2. Section 3 is dedicated to the formal proof of Theorem 1. Section 5 outlines the proof of Theorem 2.

¹Note that when $\alpha \in \Omega(k)$ the lower bound on the competitive ratio due to [10] is $\Omega(k)$, which is obviously tight (using the trivial upper bound for the greedy algorithm). Thus the problem is interesting only when $\alpha \in o(k)$.

1.1 Some related results

The Steiner tree problem has been extensively studied in several settings and variations. We overview only some of the results which are of particular relevance to this work. For the online Steiner tree problem in *undirected* graphs, Imase and Waxman [12] showed a tight bound of $\Theta(\log k)$ on the competitive ratio of online Steiner Tree. In directed graphs of unbounded asymmetry, it is very easy to show that the competitive ratio of every algorithm, deterministic or randomized, is as large as the trivial bound, namely $\Theta(k)$. When the terminals are given as a sequence of points in the Euclidean plane, Alon and Azar [2] showed a lower bound of $\Omega(\log k / \log \log k)$ on the competitive ratio. For the so-called on-line *generalized* Steiner problem, in which pairs of terminals are requested sequentially and the algorithm must guarantee connectivity for every such requested pair, Berman and Coulston [4] proved a tight upper bound of $O(\log k)$, a result which improved the upper bound of $O(\log^2 k)$ due to Awerbuch *et al.* [3]. Both results apply to undirected graphs.

1.2 Preliminaries

Given a directed graph of bounded asymmetry and an edge $e = (v, u) \in E$, it is always the case that its *antiparallel* edge $\bar{e} = (u, v)$ is always in E as well. Let $T = (r', V', E')$ be an arborescence rooted at r' , we denote by \hat{T} the graph (V', E'') , with $E'' = E' \cup \{\bar{e} : e \in E'\}$. In words, \hat{T} is the subgraph of G induced by the vertices of T , and induces all edges in T as well as all their antiparallel edges. For arborescence T and vertices v, u in T , we denote by $p_T(u, v)$ (resp. $p_{\hat{T}}(u, v)$) the simple directed path from u to v using exclusively edges in T (resp. \hat{T}). Note that such paths are uniquely defined (provided that $p_T(u, v)$ exists).

The cost of a directed path p is the total cost of all directed edges in p , and will be denoted by $c(p)$. We denote by $c(T)$ the cost of arborescence T , namely the sum of the cost of the directed edges in T . We emphasize that only edges in T and none of their antiparallel edges contribute to $c(T)$. We will always use T^* to denote the optimal arborescence on input (G, K) , with $|K| = k$, with $OPT = c(T^*)$. For any $K' \subseteq K$, we let $c_{GR}(K')$ denote the cost that GREEDY pays on the subset K' of the input (in other words, the contribution of terminals in K' to the total cost of GREEDY).

For convenience, we will be using the term “tree” to refer to a (rooted) arborescence.

2 Outline of the proof of Theorem 1 and intuition

First, note that when $\alpha \in \Omega(k)$ the lower bound on the competitive ratio due to [10] is $\Omega(k)$, which is obviously tight (using a trivial upper bound for the greedy algorithm). Thus the problem is interesting only when $\alpha \in o(k)$. We will be assuming that α is integral since we can round α to the closest integer without affecting, asymptotically, the bounds. Let l be such that $\alpha^l = k$, which means that $l = \frac{\log k}{\log \alpha}$.

There are two main components in the proof. The first addresses the following question. Suppose that a subset $K' \subseteq K$ of $O(\alpha)$ terminals belongs in a (rooted) subtree T' of T^* . Suppose also that GREEDY has been charged already the cost for serving a single terminal in K' , but no other terminals in K' have arrived yet. Can we bound $c_{GR}(K')$ as a function of $c(T')$? Of course we can give trivial upper bounds: for every terminal $t \in K'$ (excluding the first terminal in K') we have $c_{GR}(t) = O(\alpha)c(T')$, hence $c_{GR}(K') = O(\alpha^2)c(T')$; even better we can use the fact that GREEDY is log-competitive in undirected graphs, which implies that $c_{GR}(K') = O(\alpha \log |K'| c(T')) = O(\alpha \log \alpha \cdot c(T'))$. Note that it is not true that one can claim that $c_{GR}(K') = \alpha \cdot c(T')$; this would be true if the

root of T' had already become part of the current tree GREEDY builds (borrowing terminology from the literature on the Steiner problem in undirected graphs, GREEDY should include appropriate *Steiner vertices* in the tree it builds). There is no easy way to guarantee this; in fact Theorem 2 shows this does not hold. However, we can still improve the bound to $O(\alpha \frac{\log \alpha}{\log \log \alpha} c(T'))$, as shown in Lemma 4.

The second component of the proof provides a framework for a recursive application of the previous observation. In particular, suppose that T^* can be partitioned into (roughly) α edge-disjoint trees T_1, \dots, T_α of (roughly) the same number of terminals, namely α^{l-1} . Let v_i ($i \in [1, \alpha]$) denote the first terminal, among all terminals in T_i , to be requested; V_1 denote the set $\{v_i : i \in [1, \alpha]\}$ and K_1, \dots, K_α denote the set of all remaining terminals in T_1, \dots, T_α respectively. Using Lemma 4 we have that $c_{GR}(K) = c_{GR}(V_1) + \sum_{i=1}^\alpha c_{GR}(K_i) = O(\alpha \frac{\log \alpha}{\log \log \alpha} OPT) + \sum_{i=1}^\alpha c_{GR}(K_i)$. We then proceed recursively² at each subtree T_i . Note that trees T_i are edge-disjoint, hence at each level of the recursion the cost of GREEDY increases by $O(\alpha \frac{\log \alpha}{\log \log \alpha})OPT$. Since there are roughly l levels of recursion, we derive the required upper bound on the cost of GREEDY.

Naturally, the previous argument relies upon the ability to provide a decomposition of T^* into (roughly) α trees of (roughly) the same size. Moreover, the decomposition should be hierarchical, in the sense that we should be able to further decompose the resulting trees while upholding the above property. Lemma 14 proves the existence of such a “balanced”, hierarchical decomposition of T^* .

3 Proof of Theorem 1

The following is a key Lemma in the analysis of GREEDY.

Lemma 4 *Let T' be a subtree of T^* rooted at vertex r' and let $K' \subseteq K$, with $|K'| = O(\alpha)$ be a subset of K such that every terminal in K' is a vertex in T' . Let w denote the terminal which was requested the earliest among all terminals in K' . Then³ $c_{GR}(K') = c_{GR}(w) + O(\alpha \frac{\log \alpha}{\log \log \alpha})c(T')$.*

In order to show Lemma 4, we will first prove the lemma for the case in which T' and K' have a relatively simple structure (c.f. Lemma 6). The proof of Lemma 4 will then become substantially easier.

Definition 5 *Let T', K' and r' be as defined in the statement of Lemma 4. We call the triplet $\mathcal{C} = (T', K', r')$ a comb instance, or comb for simplicity if the following hold: T' consists of a directed path P from r' to a certain vertex v_1 , which visits vertices $v_{k'}, \dots, v_1$ in this order (but possibly other vertices too); there are also directed paths p_i from v_i to u_i . No other edges are in T' . Finally the set K' is defined as the set of vertices $\{u_1, \dots, u_{k'}\}$. We call P the backbone of \mathcal{C} , and the paths p_i the terminal paths of the comb. The vertex set of \mathcal{C} is the set of vertices in T' .*

Figure 1 illustrates the structure of a comb. Note that the definition allows the paths p_i to be empty, in which case $v_i \equiv u_i$; in addition, the directed paths from v_i to v_{i-1} (as well as the path from r' to $v_{k'}$) may also be empty, in which case $v_i \equiv v_{i-1}$. For the proof of Lemma 6 we will assume, *a fortiori*, that such degeneracies do not arise (see Appendix A for a justification).

² In the first level of the recursion, we could instead claim that $c_{GR}(V_1) = O(\alpha \cdot OPT)$, since we know that r is the root of T^* . However, this is true only for the first level, and for all subsequent levels we have to rely to Lemma 4.

³An alternative statement of the lemma is that $c_{GR}(K' \setminus \{w\}) = O(\alpha \frac{\log \alpha}{\log \log \alpha})c(T')$. In fact, in the proof of Theorem 1 we will use the latter statement. The same applies to the statement of Lemma 6.




Figure 1: The structure of a comb instance. Wavy lines indicate directed paths between vertices.

We will make use of some auxiliary definitions. For terminal u_i in comb \mathcal{C} we say that its *index* is i . We say that u_j *precedes* u_i in \mathcal{C} (denoted by $u_j \prec u_i$) iff $j < i$. We say that u_j is *between* u_i and $u_{i'}$ iff $u_i \prec u_j \prec u_{i'}$. Given terminals u_i and u_j such that $u_i \prec u_j$ we call the path $p_{\hat{T}}(v_i, v_j)$ the *segment* between u_i and u_j .

The following lemma is a version of Lemma 4 in which T' and K' are restricted to form a comb.

Lemma 6 *Given the comb $\mathcal{C} = (T', K', r')$, with $|K'| = k' = O(\alpha)$, let $w \in K'$ denote the terminal requested the earliest among all terminals in K' . Then $c_{GR}(K') = c_{GR}(w) + O(\alpha \frac{\log \alpha}{\log \log \alpha})c(T')$.*

Proof. Let π denote a permutation of $\{1, \dots, k'\}$ such that $u_{\pi_1}, \dots, u_{\pi_{k'}}$ is the sequence of the requests in K' in the order in which they are requested ($w = u_{\pi_1}$). Also let x be such that $x^x = |K'|$ which implies that $x = O(\frac{\log \alpha}{\log \log \alpha})$. For convenience we will assume that x is an integer.

In order to bound $c_{GR}(K' \setminus u_{\pi_1})$ we will determine an assignment for every terminal u_{π_i} with $2 \leq i \leq k'$ to a *specific* terminal $\bar{u}_{\pi_i} \in \{u_{\pi_1}, \dots, u_{\pi_{i-1}}\}$. We call terminal \bar{u}_{π_i} the *mate* of u_{π_i} . Let q_i denote the directed path in \hat{T} from \bar{u}_{π_i} to u_{π_i} , also called the *connection path* for u_{π_i} . We will show that

$$C \stackrel{\text{def}}{=} \sum_{i=2}^{k'} c(q_i) = O(\alpha \frac{\log \alpha}{\log \log \alpha})c(T'). \quad (1)$$

Since $c_{GR}(K' \setminus u_{\pi_1}) \leq \sum_{i=2}^{k'} c(q_i)$ the lemma will then follow.

In order to simplify the proof we will ignore the contribution to the cost C of all (directed) edges in the connection paths q_i 's which belong in the tree T , and will only consider the contribution of edges in \hat{T} but not in T . This is because the total contribution to C of edges in q_i which belong to T is at most $k'c(T) = O(\alpha c(T))$, which does not exceed asymptotically the bound we want to prove.

We first aim towards grouping together terminals as they are being requested; this will also facilitate our assignment of terminals to their mates. To this end we introduce the concept of a *run* and the concept of the *representative* of a run. The first terminal to be requested, namely u_{π_1} begins run 1. At the time u_{π_i} ($i \geq 2$) is requested, for every run j generated thus far, let $u_{h(j)}$ denote the terminal with the highest index in the comb among all terminals in run j requested so far. If $u_{h(1)} \prec u_{\pi_i}$ in the comb, then u_{π_i} becomes a member of run 1. Otherwise let π_l, π_r be the highest and lowest indices, respectively, with $l, r \leq i - 1$, such that u_{π_i} is between u_{π_l} and u_{π_r} in the comb (if such a π_l does not exist, then we set it to 0). If either i) $\pi_l = 0$; or ii) there is no run j with the property that $u_{\pi_l} \equiv u_{h(j)}$ **and** u_{π_r} is the representative of run j , then u_{π_i} starts a new run with the representative of the new run to be u_{π_r} . Otherwise, u_{π_i} becomes a member of the same run⁴ as u_{π_l} , namely run j , and its *predecessor* in the run is u_{π_l} .

⁴Note that the representative of a run is not a member of the run.

The above process produces a partition of K' into runs, and determines a representative of all runs other than run 1. We will be denoting by $rep(r)$ the representative of a run r ; we will also be using the notation $rep(u_{\pi_i})$ to denote the representative of the run to which u_{π_i} is assigned. An example of the partition of K' into runs is given in Appendix B.

We introduce some additional definitions. Let $\sigma(i, j) \in \{1, \dots, k'\}$ be such that $u_{\sigma(i, j)}$ is the terminal of the j -th lowest index which belongs in run i . The *size of a run* is the number of terminals in the run. For a certain terminal u_{π_i} , denote by $s(u_{\pi_i})$ the segment between u_{π_i} and $rep(u_{\pi_i})$; we call $s(u_{\pi_i})$ the *segment of u_{π_i}* . Denote by $d(r)$ the cost of the segment of the first terminal in run r , namely $d(r) = c(s(u_{\sigma(r, 1)}))$.

The following claim describes some properties related to a run.

Claim 7 (i) *Every terminal in a run other than run 1 precedes its representative in the comb. Furthermore, each terminal is the representative of at most one run.*

(ii) *Suppose r, r' ($r \neq r'$) denote two runs such that there exists a terminal u in run r' such that $u_{\sigma(r, j)} \prec u \prec u_{\sigma(r, j+1)}$, i.e., u is between two consecutive terminals of run r . Then the whole run r' is contained between $u_{\sigma(r, j)}$ and $u_{\sigma(r, j+1)}$.*

Proof. We only prove the second statement (the first statement is straightforward). First, note that at the time $u_{\sigma(r, j+1)}$ is requested, no member of run r' which is between $u_{\sigma(r, j)}$ and $u_{\sigma(r, j+1)}$ has yet been requested, otherwise $u_{\sigma(r, j+1)}$ would not have become a member of run r with $u_{\sigma(r, j)}$ as its predecessor. Hence, at the time u is requested, since it becomes member of the run r' , $rep(r')$ has to be either $u_{\sigma(r, j+1)}$, or a terminal u' which is between u and $u_{\sigma(r, j+1)}$ in the comb. In either case, $rep(r')$ will be between $u_{\sigma(r, j)}$ and $u_{\sigma(r, j+1)}$.

Suffices then to show that $u_{\sigma(r', 1)}$ is such that $u_{\sigma(r, j)} \prec u_{\sigma(r', 1)}$. By way of contradiction, suppose that this is not true, then let u'' denote the terminal of smallest index in run r' which is between $u_{\sigma(r, j)}$ and $u_{\sigma(r, j+1)}$ (such a terminal exists, by the hypothesis of the claim). When u'' is requested, $u_{\sigma(r, j+1)}$, and thus $u_{\sigma(r, j)}$ as well have been requested, as argued early. This would imply, however, that u'' cannot be a member of run r' , a contradiction. \square

Once the runs and the representatives have been determined, we can proceed with assigning mates to the terminals, using the following rules.

1. Terminal $u_{\sigma(1, j)}$ is always assigned $u_{\sigma(1, j-1)}$ as its mate, for $j \geq 2$.
2. The terminal of smallest index in run $i > 1$, i.e., $u_{\sigma(i, 1)}$, is always assigned the representative of the run $rep(i)$ as its mate.
3. If the size of run $i > 1$ is at most x , then each terminal in the run i is assigned the representative of the run as its mate.
4. Otherwise (i.e., if the size of run $i > 1$ is larger than x), then
 - (a) If $c(p_{\hat{T}}(v_{\sigma(i, j-1)}, v_{\sigma(i, j)})) \geq d(i)/x$ ($j \geq 2$) then $u_{\sigma(i, j)}$ is assigned the representative of the run as its mate;
 - (b) Otherwise, the mate of $u_{\sigma(i, j)}$ is set to be $u_{\sigma(i, j-1)}$.

Note that for every i, j $u_{\sigma(i, j)}$ is requested *after* $u_{\sigma(i, j-1)}$. Likewise, every terminal in any run (other than run 1) is requested *after* the representative of the run. This means that our assignment of terminals to mates is feasible, in the sense that the mate of a terminal is requested prior to the terminal itself. Recall that once terminals are assigned mates, the connection paths q_i 's are uniquely determined.

Recall that C denotes the total cost due to the assignment of terminals to mates (see the definition in Eq (1)) via the connection paths q_i 's. In order to bound C we observe that we can express C as the sum of two partial costs, which we call C_1 and C_2 . Here, C_1 denotes the cost due to edges \bar{e} such that e belongs in some terminal path p_i in T , and C_2 denotes the cost due to edges \bar{e} such that e belongs in the backbone P of the comb. Indeed, our particular assignment of terminals to mates is motivated by the main objective to balance the contributions of C_1 and C_2 to the overall cost C . At an intuitive level, if a run other than run 1 has small size (at most x), then its representative can “afford” to act as the mate of all terminals in the run (see assignment rule 3), since this does not contribute much to the C_1 cost, and it definitely does not affect the C_2 cost. However, if the size of a run is large we must be more careful as assignment rule 4 suggests. For instance, we can no longer afford to assign all terminals in the run to the representative of the run as their mate: this would implode the C_1 cost. Instead, as long as a terminal in a run is “not too far away” with respect to the backbone cost (distance) from its predecessor in the run, we let the predecessor be its mate: the C_2 cost will not increase by much, in the sense that the average contribution of such terminals to C_2 will be kept low (rule 4b). Otherwise, i.e., when the terminal is indeed far away from its predecessor in the run, then we will choose the representative as the mate (rule 4a); since the later case will not arise too often for any given run (namely at most x times) the C_1 cost will be kept low. In the remainder of the proof we elaborate on this intuitive explanation.

First we show how to bound C_1 . Denote by $C_{1,i}$ the contribution of the connection path q_i for terminal u_{π_i} to this cost which means that $C_1 = \sum_{i=2}^{k'} C_{1,i}$. Note that u_{π_i} contributes to C_1 only in two cases: i) If it is the mate of its successor in the run to which it belongs (rule 1, 4b). In this case, it contributes at most $\alpha c(p_i)$ to $C_{1,i}$; ii) If it is the mate of certain terminals in the run for which it is the representative (the remaining rules). Recall that from Claim 7 (i), u_{π_i} can be the representative of at most one run; let r denote this specific run. Then u_{π_i} can be the mate of $u_{\sigma(r,1)}$ (rule 2), as well as either at most x terminals in r (as follows when either rule 3 or 4a applies). The total contribution to $C_{1,i}$, in this case, is then bounded by $(x + 1)\alpha c(p_i)$. Summarizing,

$$C_{1,i} \leq (x + 2)\alpha \cdot c(p_i) = O(x\alpha \cdot c(p_i)). \quad (2)$$

Next define $C_{2,i}$ as the contribution of u_{π_i} to C_2 . We say that u_{π_i} *contributes* the directed edge e , with $\bar{e} \in P$ when the path q_i includes e . For the remainder of this proof, we will call such edges *expensive*. Also, let \bar{P} denote the directed path from v_1 to r in \hat{T} and q'_i denote the intersection of \bar{P} with q_i , namely the subpath of q_i which consists of expensive edges only, which means that $C_{2,i} = c(q'_i)$. Let X denote the subset of K' which consists of terminals with non-zero contribution to C_2 . We can think of the assignment of terminals in X to their mates as being done as the terminals in X are requested over time; more precisely, we can think of all edges in q'_i being “bought”, as the connection path between the terminal and its mate is *established*, at the precise moment $u_{\pi_i} \in X$ is requested. In this view, every time an expensive edge in \bar{P} is contributed due to such an assignment, we say that the *depth* of the edge increases by 1 (initially, i.e., before any terminals have been requested, all expensive edges have depth zero).

In addition, observe that u_{π_i} contributes to C_2 when it is assigned to its mate as a result of either rule 1 or rule 4b only. In other words, we are only considering cases in which \bar{u}_{π_i} is the predecessor of u_{π_i} in its run.

Claim 8 *For a terminal $u_{\pi_i} \in X$, all expensive edges in q'_i have the same depth, right after q_i is established.*

Proof. By way of contradiction. Consider the first terminal $u_{\pi_i} \in X$, in the order in which terminals are requested, which does not have the required property. This implies that there would exist a terminal $u' \neq \bar{u}_{\pi_i}$ such that u' is requested earlier than u_{π_i} , and for which $\bar{u}_{\pi_i} \prec u' \prec u_{\pi_i}$ in the comb. Indeed, if this was not the case, then for all pairs $(\bar{u}_{\pi_j}, u_{\pi_j})$ such that u_{π_j} is requested before u_{π_i} and $u_{\pi_j} \in X$, we would have that either i) $\bar{u}_{\pi_j} \prec \bar{u}_{\pi_i}$ and $u_{\pi_i} \prec u_{\pi_j}$; or ii) $\bar{u}_{\pi_j} \prec u_{\pi_j} \prec \bar{u}_{\pi_i}$ or $u_{\pi_i} \prec \bar{u}_{\pi_j} \prec u_{\pi_j}$. This would imply that either q'_i is a subpath of q'_j (case (i)), or q'_j and q'_i are edge-disjoint (case (ii)). But since all edges in q'_j have the same depth after q'_j is established (by our choice of u_{π_i}), the same would be true for q'_i , a contradiction. Hence there is a terminal u' between \bar{u}_{π_i} and u_{π_i} in the comb which is requested earlier than u_{π_i} , which is also a contradiction since u_{π_i} would not become the successor of \bar{u}_{π_i} in the latter's run. \square

Claim 8 asserts that it is meaningful to say that terminal u_{π_i} is of depth δ if right after it is assigned to its mate, and the connection path q_i is established, the depth of all expensive edges at the connection path becomes equal to δ . This implies that we can further partition X into sets $X_1, X_2 \dots$ such that X_i consists of all terminals of depth i . Note that for all i with $u_{\pi_i} \in X_j$, the paths q'_i are edge-disjoint.

The following lemma shows that the contribution of a terminal to C_2 decreases exponentially with its depth.

Lemma 9 For a terminal $u_{\pi_i} \in X_j$, $C_{2,i} \leq \frac{oc(P)}{x^{j-1}}$.

Proof. By induction on j . The claim is trivially true for $j = 1$ from the disjointness of all q'_i 's for terminals in X_1 . Suppose the claim holds for j , we will show it holds for $j + 1$. Consider the set of terminals X_{j+1} . Recall that every terminal in X_{j+1} will buy expensive edges of current depth exactly j prior to the assignment of the said terminal to its mate, then right after the assignment the depth of such edges increases by one. Let u_{π_i} be a terminal in X_{j+1} , q_i its connection path, and $r > 1$ the run to which it belongs. We want to show that the whole run r is contained between two terminals \bar{u}_{π_l} and u_{π_l} with the property that $u_{\pi_l} \in X_j$, and that in addition $\bar{u}_{\pi_l} \prec rep(r) \preceq u_{\pi_l}$ (here \preceq denotes either precedence or equivalence).

We begin with a simple observation. Consider the set Q of paths of the form q'_l such that $u_{\pi_l} \in X_j$. As noted earlier, any two paths in Q are edge disjoint. We claim that q'_i is a subpath of one of the paths in Q . Note first that every edge $e \in q'_i$ must belong in some path $q \in Q$, since the depth of all edges in q'_i become $j + 1$ once q'_i is established. Then we can use an argument along the lines of the proof of Claim 8: If q'_i was not a subpath of a path in Q , then there would exist a terminal $u' \in X_j$ other than \bar{u}_{π_i} which is requested earlier than u_{π_i} and such that $\bar{u}_{\pi_i} \prec u' \prec u_{\pi_i}$ in the comb, a contradiction, since that would mean that \bar{u}_{π_i} cannot be the predecessor of u_{π_i} in r .

We know, therefore, that there exists a terminal $u_{\pi_l} \in X_j$ for which q'_i is a subpath of q'_l . Since u_{π_l} is requested earlier than u_{π_i} , and $u_{\pi_i} \prec u_{\pi_l}$ in the comb, we have that $u_{\pi_i} \prec rep(r) \preceq u_{\pi_l}$, hence every terminal in r must precede u_{π_l} . In addition, if $r' \neq r$ denotes the run in which the (consecutive) members \bar{u}_{π_l} and u_{π_l} belong, we know that there exists a member of run r , namely u_{π_i} which is between \bar{u}_{π_l} and u_{π_l} . Claim 7, in particular statement (ii), will then guarantee that all elements in r are between \bar{u}_{π_l} and u_{π_l} .

We thus derive that $C_{2,i} \leq d(r)/x \leq C_{2,l}/x$. The first inequality follows from the assignment of terminals to mates, when applying rule 4b (Note that rule 1 cannot apply since $u_{\pi_i} \in X_{j+1}$ has to belong to a run $r > 1$). The second inequality follows from the previously shown property concerning the containment of run r , including $rep(r)$, between \bar{u}_{π_l} and u_{π_l} . By the induction hypothesis, $C_{2,l} \leq \frac{oc(P)}{x^{j-1}}$ and the lemma follows. \square

We now proceed to bound C_2 . Denote by $c_2(X_j)$ the contribution of X_j to C_2 . Recall that paths of the form q'_l , for all $u_{\pi_l} \in X_j$ are edge-disjoint, for fixed j . Hence $c_2(X_j) \leq \alpha \cdot c(P)$. Combining this fact with Lemma 9 we have

$$c_2(X_j) = \min\{\alpha c(P), \frac{\alpha c(P)}{x^{j-1}} |X_j|\}$$

therefore,

$$C_2 = \sum_j c_2(X_j) = \sum_j \min\{\alpha c(P), \frac{\alpha c(P)}{x^{j-1}} |X_j|\} \quad (3)$$

Note that (3) is maximized when $|X_j| = x^{j-1}$, for all $j \geq 2$, which implies that

$$C_2 \in O(x\alpha c(P)) \quad (4)$$

We are now ready to conclude the proof of Lemma 6, observe that

$$\begin{aligned} C &= C_1 + C_2 = \sum_{i=1}^{k'} C_{1,i} + C_2 \\ &= O(x\alpha \sum_{i=1}^{k'} c(p_i)) + O(x\alpha c(P)) \quad (\text{From Eq (2) and Eq (4)}) \\ &= O(x\alpha (\sum_{i=1}^{k'} c(p_i) + c(P))) = O(x\alpha c(T')) = O(\alpha \frac{\log \alpha}{\log \log \alpha} C(T')). \end{aligned}$$

□

As mentioned earlier in the section, Lemma 6 will facilitate the proof of Lemma 4. The idea behind the proof is to decompose the set K' into a collection of “near-disjoint” comb instances.

Proof of Lemma 4 We will show how to partition K' into two collections of sets: The first collection, denoted by I , will consist of terminals which asymptotically do not affect much the overall cost of GREEDY; we will call such terminals *ignored*. The second collection will consist of a partition of $K' \setminus I$ into near-disjoint comb-instances. In particular, by near-disjoint we mean that any edge in T' will be shared by at most two comb instances.

The decomposition is determined by visiting terminals in K' in the order they are requested. Let σ' denote the sequence of terminals in K' , and $\sigma'[i]$ the i -th requested terminal in K' . We initialize the set I as well as the collection of comb instances, denoted by \mathcal{C} , to empty sets. For any terminal $t \in \sigma'$ let q_t denote the path $p_{T'}(r', t)$, namely the path from r' to t in T' and \bar{q}^t denote the path $p_{\bar{T}}(t, r')$, namely the path from t to r' which follows all edges antiparallel to edges in q_t .

The first terminal in σ' , namely $\sigma'[1] = w$, induces a (trivial) comb instance of the form $\mathcal{C}_1 = (q_w, w, r')$, with backbone q_w , and a single, empty terminal path⁵. We say that we *assign* w to \mathcal{C}_1 , and we add \mathcal{C}_1 to \mathcal{C} .

Consider now terminal $t = \sigma'[i]$ with $i > 1$. Focus on the sequence of vertices, in the order they are visited, in the directed path \bar{q}_t . Let v denote the first vertex in this sequence which belongs to the vertex set of some $\mathcal{C}_l \in \mathcal{C}$, with the convention that in the case where t itself is in the vertex set of \mathcal{C}_l , we consider v to be vertex t (if v belongs to more than one combs, then we choose any of such combs to be \mathcal{C}_l .) Note that the sequence of vertices must include r' which is in \mathcal{C}_1 , so such a v always exists. We consider the following cases, and make appropriate decisions *in this order*:

⁵Recall that our definition of a comb instance allows combs in which the terminal paths are empty. This is one case where such an instance arises.

- *Case 1.* v is a terminal $\sigma'[j]$ with $j < i$. In this case we add t to the set of ignored terminals I .
- *Case 2.* v is a vertex in the backbone of \mathcal{C}_l . In this case we assign t to \mathcal{C}_l and we update \mathcal{C}_l by adding the corresponding terminal path $p_{T'}(v, t)$ (possibly empty) to \mathcal{C}_l .
- *Case 3.* If none of the above happen, then it must be the case that v belongs in one of the terminal paths for \mathcal{C}_l . In particular, there must exist some $j < i$, such that terminal $s = \sigma'[j]$ is a terminal already assigned to \mathcal{C}_l , and v is a vertex in the terminal path corresponding to s in \mathcal{C}_l which does not belong in the backbone of \mathcal{C}_l and is not s either. Let s' denote the vertex of the terminal path for s in \mathcal{C}_l which also belongs in the backbone of \mathcal{C}_l . In this particular case, we say that t initiates a new comb $\mathcal{C}_{\rho+1}$, where ρ is the highest current index of combs in the collection \mathcal{C} . More precisely, we create a new comb $\mathcal{C}_{\rho+1}$, rooted at s' , with $p_{T'}(s', s)$ as its backbone, and the path $p_{T'}(v, t)$ as the terminal path for t . We assign t to $\mathcal{C}_{\rho+1}$; we also say that terminal s pays for initiating comb \mathcal{C}_{l+1} (the context of this will become clear later in the proof).

Figure 2 illustrates the application of the decomposition.




Figure 2: An example of the comb decomposition. Here, terminals $\sigma'[1]$, $\sigma'[2]$, $\sigma'[11]$ and $\sigma'[12]$ are assigned to \mathcal{C}_1 . Terminals $\sigma'[3 \dots 10]$ and $\sigma'[13, \dots 19]$ which belong in the subtrees of $\sigma'[2]$ and $\sigma'[11]$, respectively, are all ignored terminals. Terminal $t = \sigma'[20]$ initiates \mathcal{C}_2 , with terminal $s = \sigma'[11]$ paying for initiating \mathcal{C}_2 .

Once the whole sequence σ' is processed as above, we are left with a set of ignored terminals I and a partition of terminals in $K' \setminus I$ into sets K'_i such that all terminals in K'_i are assigned to comb \mathcal{C}_i . Denote by t_i^1 the terminal which initiates comb \mathcal{C}_i , and s_i^1 the terminal which pays for initiating that comb. From the decomposition we have

$$c_{GR}(K') = c_{GR}(I) + \sum_i c_{GR}(t_i^1) + \sum_i c_{GR}(K'_i \setminus \{t_i^1\}). \quad (5)$$

Suffices then to bound $c_{GR}(K')$, as expressed by (5). First, observe that for every ignored terminal $i \in I$, there is a terminal which has been requested earlier in σ' and which is an ancestor

of i in T' . It follows that $c_{GR}(i) \leq c(T')$, therefore

$$c_{GR}(I) \leq |I'|c(T') = O(\alpha)c(T'). \quad (6)$$

Next, denote by T'_i the edge set for \mathcal{C}_i . The decomposition has the following property

Property 10 *A (non-empty) terminal path for any comb $\mathcal{C}_i \in \mathcal{C}$ is the backbone for at most one \mathcal{C}_j , with $j \neq i$. In addition, all (non-empty) terminal paths for all combs in \mathcal{C} are edge-disjoint.*

Property 10 implies that any edge in T' appears in at most two combs in \mathcal{C} . We can combine this fact with Lemma 6 to derive that

$$\begin{aligned} \sum_i c_{GR}(K'_i \setminus \{t_i^1\}) &= O\left(\alpha \frac{\log \alpha}{\log \log \alpha}\right) \sum_i c(T'_i) \\ &\leq O\left(\alpha \frac{\log \alpha}{\log \log \alpha}\right) 2c(T'), \end{aligned} \quad (7)$$

Finally, consider terminal $t = t_i^1$. Let v be the vertex on the path \bar{q}_t , as defined in the statement of case (3), which describes the initiation of comb \mathcal{C}_i . For convenience, let p_i^1 denote the terminal path for s_i^1 in the comb to which s_i^1 is assigned. Observe that a terminal s_i^1 will always initiate only one comb, namely \mathcal{C}_i . It follows that for every $i > 1$, we have

$$c_{GR}(t_i^1) \leq c(p_{\hat{T}'}(s_i^1, v)) + c(p_{T'}(v, t_i^1)) \leq \alpha c(p_i^1) + c(T').$$

Therefore

$$\begin{aligned} \sum_i c_{GR}(t_i^1) &= c_{GR}(t_1^1) + \sum_{i \geq 1} c_{GR}(t_i^1) \leq c_{GR}(w) + \sum_{i \geq 1} (\alpha c(p_i^1) + c(T')) \\ &= c_{GR}(w) + \alpha \sum_{i \geq 1} c(p_i^1) + \sum_{i \geq 1} c(T') \\ &\leq c_{GR}(w) + \alpha c(T') + O(\alpha)c(T') = c_{GR}(w) + O(\alpha)c(T'), \end{aligned} \quad (8)$$

where, in the last line of inequalities, we used the fact that a terminal s_i^1 will always initiate only one comb, in conjunction with the edge-disjointness of terminal paths (Property 10) and the fact that $O(\alpha)$ combs are initiated.

Using (6), (7) and (8), Eq (5) gives $c_{GR}(K') = c_{GR}(w) + O\left(\alpha \frac{\log \alpha}{\log \log \alpha}\right)c(T')$. \square

As outlined in Section 2, we aim towards applying Lemma 4 in a recursive manner. To this end, we now proceed to decompose the optimal arborescence T^* into a “balanced” family of arborescences. We are interested in decompositions which are edge-disjoint. Vertex-disjointness is not critical, however, in the event a terminal vertex belongs to more than one tree in the decomposition, we insist that the terminal itself is *assigned* to exactly one tree in the decomposition (c.f. proof of Lemma 11). Effectively, this will allow us to treat the trees as if they were “terminal-disjoint”, even though the trees are not necessarily vertex-disjoint. This is a convention we make and its only purpose is to simplify the proofs.

Lemma 11 *Let T be a Steiner tree for a set of k terminals, and x a number such that $1 \leq x \leq k$. Then T can be decomposed in at most $\lceil \frac{k}{x} \rceil$ and at least $\frac{k}{2x}$ edge-disjoint arborescences, with each arborescence assigned between x and $2x$ terminals.*

Proof. There is a simple algorithm that provides the decomposition. Let T_v denote the subtree of T rooted at node $v \in T$. Let $k(T')$ denote the number of terminals in a tree T' . Starting with T , find a node v of maximum depth such that $k(T_v)$ is at least x . Let T_1, \dots, T_l denote those subtrees, in increasing order of terminals they contain. Denote by T'_j , $j \in [1, l]$ the subtree rooted at v which has $T_1 \dots T_j$ as its children, i.e., in T'_j there are directed edges from v to each of T_1, \dots, T_j . Find the smallest $j \leq l$ such that $k(T'_j)$ is at least x and let T' denote this particular T'_j . Clearly, $x \leq k(T') \leq 2x$. We input T' in the decomposition, and repeat the process with the tree $T \setminus T'$. We emphasize that once a terminal t is assigned to a tree in the decomposition, the vertex of terminal t is treated, in subsequent iterations of the algorithm as a non-terminal vertex. This will guarantee that the decomposition assigns each terminal to a unique tree in the decomposition, even though the decomposition itself is not vertex-disjoint. Since every time a tree is inserted in the decomposition we decrease the number of terminals in T by at least x and by at most $2x$, the decomposition consists of at most $\lceil \frac{k}{x} \rceil$ and at least $\frac{k}{2x}$ trees. \square

Definition 12 Let $D(T)$ denote the set of trees obtained by applying an edge-disjoint decomposition D to tree T with k terminals. We say that D is balanced for tree T if every tree in $D(T)$ is assigned between $\frac{k}{4\alpha}$ and $\frac{4k}{\alpha}$ terminals, and $\frac{\alpha}{4} \leq |D(T)| \leq 4\alpha$.

In what follows, let $k(T')$ denote the number of terminals assigned to a tree T' , as the result of a decomposition. Recall that K denotes the set of all terminal requests, with $|K| = k$. Let τ be such that $\alpha^\tau = k$, which means that $\tau = \frac{\log k}{\log \alpha}$.

Definition 13 A balanced hierarchical decomposition H of height L of a Steiner tree T consists of L families of sets, denoted by $\mathcal{F}_1, \dots, \mathcal{F}_L$, defined recursively as follows. Let D_1^T be a balanced decomposition for tree T . First, we define \mathcal{F}_1 to be the set $D_1^T(T)$, and every tree in \mathcal{F}_1 is said to be at level 1. We define \mathcal{F}_{i+1} in terms of \mathcal{F}_i as $\mathcal{F}_{i+1} = \{D_{i+1}^{T'}(T') \mid T' \in \mathcal{F}_i\}$, where $D_{i+1}^{T'}$ is a balanced tree decomposition for tree T' (with respect to the terminals assigned to T'). We also say that every tree in \mathcal{F}_{i+1} is at level $i+1$. Finally, H must be such that every tree at level L is assigned at most $c \cdot \alpha$ terminals, for some constant c .

Lemma 14 For every tree Steiner tree T with k terminals there exists a balanced hierarchical decomposition H of T such that for every tree T' at level i , we have $\frac{1}{4}\alpha^{\tau-i} \leq k(T') \leq 4\alpha^{\tau-i}$, and the height of the decomposition is $L = O(\tau)$.

Proof. We prove the first statement by induction on i , the second statement concerning the bound on L follows then from the fact that all trees at level $\lceil \tau - 1 \rceil$ in H have $O(\alpha)$ terminals. We need to define an appropriate balanced tree decomposition, for every tree and every level in the decomposition. For the first level, we invoke Lemma 11 with parameter $x = \alpha^{\tau-1}$. Given tree T' at level i , we define $D_{i+1}^{T'}$ as follows:

- If $k(T') > \alpha^{\tau-i}$, then invoke Lemma 11 with parameter $x = 2\alpha^{\tau-(i+1)}$.
- otherwise, invoke Lemma 11 with parameter $x = \frac{\alpha^{\tau-(i+1)}}{2}$.

It is easy to verify, by induction on the levels that for every tree T' at level i , $D_{i+1}^{T'}$ is a balanced decomposition for T' , and to show by induction that for every tree in $T'' \in D_{i+1}^{T'}(T')$ (i.e., a tree at level $i+1$) we have $\frac{1}{4}\alpha^{\tau-(i+1)} \leq k(T'') \leq 4\alpha^{\tau-(i+1)}$. \square

Let H^* denote the balanced decomposition of T^* which satisfies the conditions of Lemma 14, and let $\mathcal{F}_1, \dots, \mathcal{F}_L$ be its levels. We can now proceed with the main result of this section:

Proof of Theorem 1. Let σ denote the input sequence of terminals. Define the *depth* of a terminal $t = \sigma[i]$ in σ , as the smallest index j such that t is assigned to some tree $T \in \mathcal{F}_j$ **and** no terminal in $\sigma[1 \dots i-1]$ is assigned to T ; we also say that t is *associated* with tree T . If such a j does not exist, then the depth of the terminal t is defined to be $L+1$, and t is associated with the (unique) tree of level L to which it is assigned in H^* . Note that every terminal is associated with a unique tree. We can then partition K' into disjoint sets S_1, \dots, S_{L+1} , where S_i denotes the set of terminals of depth i . Let $c_{GR}(S_i)$ denote the cost induced by GREEDY on S_i .

Lemma 15 $c_{GR}(S_i) = O(\alpha \frac{\log \alpha}{\log \log \alpha} OPT)$.

Proof. Consider first S_1 . There are at most 4α trees in \mathcal{F}_1 , hence at most 4α terminals in S_1 . The first terminal ($\sigma[1]$) incurs a cost of at most OPT . We can now invoke Lemma 4 to get that⁶ $c_{GR}(S_1) = O(\alpha \frac{\log \alpha}{\log \log \alpha} OPT)$.

In general, consider set S_i , with $2 \leq i \leq L$. Recall that \mathcal{F}_i is derived by applying a balanced decomposition to each tree in \mathcal{F}_{i-1} , namely $\mathcal{F}_i = \bigcup_{T \in \mathcal{F}_{i-1}} D_i^T(T)$. For each $T \in \mathcal{F}_{i-1}$ there are $|D_i^T(T)|$ trees in \mathcal{F}_i , and each such tree is assigned at most one terminal in S_i . Denote by S_i^T the subset of S_i which consists of terminals assigned to trees in $D_i^T(T)$. Since D_i^T is a balanced decomposition for T by Definition 12 we know that $|D_i^T(T)| \leq 4\alpha$, hence $|S_i^T| \leq 4\alpha$. In addition, by the definition of the depth of a terminal, we know that there is one terminal in S_j with $j \leq i-1$, say terminal w , which is assigned to T and has been requested earlier than all terminals in S_i^T (if this was not the case, one terminal in S_i^T would have depth smaller than i , a contradiction). In addition, the contribution of w to the cost of GREEDY has already been taken into account when bounding $c_{GR}(S_j)$. This means that we can apply Lemma 4 for the tree T , the set of terminals S_i^T and w as a terminal requested earlier than all terminals in S_i^T , and get that

$$c_{GR}(S_i^T) = O(\alpha \frac{\log \alpha}{\log \log \alpha} c(T)),$$

Finally, since every terminal in S_i is associated with a tree of level i , note that $S_i = \bigcup_{T \in \mathcal{F}_{i-1}} S_i^T$. In addition, \mathcal{F}_{i-1} consists of edge-disjoint trees, therefore

$$c_{GR}(S_i) = \sum_{T \in \mathcal{F}_{i-1}} c_{GR}(S_i^T) = O(\alpha \frac{\log \alpha}{\log \log \alpha} \sum_{T \in \mathcal{F}_{i-1}} c(T)) = O(\alpha \frac{\log \alpha}{\log \log \alpha} c(T^*)).$$

For the case $i = L+1$ the bound on $c_{GR}(S_{L+1})$ follows along the same lines and the fact that from the definition of H^* , $O(\alpha)$ terminals are associated with each tree at level $L+1$. \square

To conclude the proof, we have that the total cost of GREEDY is bounded as follows:

$$\begin{aligned} c_{GR}(K) &= \sum_{i=1}^{L+1} c_{GR}(S_i) = (L+1) \cdot O(\alpha \frac{\log \alpha}{\log \log \alpha} OPT) && \text{(From Lemma 15)} \\ &= O(\tau) O(\alpha \frac{\log \alpha}{\log \log \alpha} OPT) && \text{(From Lemma 14)} \\ &= O(\frac{\log k}{\log \alpha}) O(\alpha \frac{\log \alpha}{\log \log \alpha} OPT) = O(\alpha \frac{\log k}{\log \log \alpha} OPT) \end{aligned}$$

hence $c_{GR}(K') = O(\min\{k, \alpha \frac{\log k}{\log \log \alpha}\}) OPT$. \square

⁶For S_1 a better bound is $c(S_1) \leq 4\alpha OPT$. See also footnote 2.

4 Outline of the proof of Theorem 2 and intuition

Consider the graph G illustrated in Figure 3: this graph will provide the motivation behind the definition of the actual adversarial input graph. Graph G is such that all “downwards” edges are cheap, whereas all “upwards” edges are expensive, and each has cost α times the cost of its antiparallel upwards edge. In particular, there is a path from the root r to a vertex t such that the cost of the directed path $r \rightarrow t$ is equal to 1 (hence the cost of the path $t \rightarrow r$ is α). Call P the path from s to t .

In addition, there are $\Theta(k)$ pairs of vertices, of the form (v_i, u_i) , defined in a recursive manner as follows. Let x be such that $x^x = k$, hence $x = \Theta(\frac{\log k}{\log \log k})$, and assume wlog that x is integral. The first group of vertices, namely group K_1 , consists of x pairs such that the v_i 's are all evenly distributed over the path P (namely the cost of the path from v_i to v_{i+1} , over edges of P is the same for all i 's such that $v_i \in K_1$, and is also the same as the cost of the path $r \rightarrow v_1$ and $v_x \rightarrow t$). In addition, the cost of the edge (v_i, u_i) is equal to $\frac{1}{x^2}$ whereas its antiparallel edge has cost $\alpha \frac{1}{x^2}$.

Suppose now that groups K_1, \dots, K_j have been defined, we will show how to define K_{j+1} . For every pair of consecutive vertices (v_i, v_{i+1}) in P such that each of v_i, v_{i+1} belongs in some K_m with $m \leq j$, we insert x vertices of the form v_i^1, \dots, v_i^x , all distributed evenly over the path $v_i \rightarrow v_{i+1}$. In addition, we insert x vertices of the form $u_i^l, l \in [1, x]$ such that the cost of the edge (v_i^l, u_i^l) is $\frac{1}{x^{j+2}}$, while the antiparallel edge (u_i^l, v_i^l) has cost $\alpha \frac{1}{x^{j+2}}$. We do the same with the pairs (r, v_1) and (v_{last}, t) , where v_{last} is the bottom vertex among all groups K_m with $m \leq j$. We continue until x groups have been defined. Note that the total number of u -vertices is then $\Theta(k)$ ⁷

The adversary will request u -vertices as terminals in *rounds*. In particular, in round i , with $1 \leq i \leq x$, the adversary will request the u -vertices of group K_i , in a bottom-up manner (i.e., starting from the u -vertex which is the farthest away from r and ending with the one that is the closest).




Figure 3: The structure of the adversarial graph G , for the case $x = 2$. Only “cheap” edges are shown while “expensive” antiparallel edges are omitted

To illustrate the intuition behind our argument, we will make the following assumption: suppose that every time the algorithm establishes a new connection path for a certain request (ie a path from a previously requested terminal), *several new edges must be bought*, in the sense that there is little

⁷Note the close similarity between this construction and the concepts of the *comb* and the *run* which we used in the proof of the upper bound. In fact, their definition was much motivated by this adversarial construction, and vice versa.

overlapping between the new connection path and the ones which have already been established. Of course this is not the case in G itself, and enforcing this requirement is by no means a trivial task. In fact, much of the details in the formal proof of Theorem 2 is dedicated to this issue.

Consider then the l -th u -vertex (terminal) requested in round j . There are three options concerning the connection path for this u -vertex: either i) will originate in r ; or ii) originate in a “higher” vertex which was requested in an earlier round; or iii) originate in a “lower” vertex which was requested before u . In the second and third cases, a cost (roughly) of at least $\alpha \frac{1}{(x+1)^j}$ will be incurred. It is easy to show that round j consists of $\Theta((x+1)^j)$ vertices. Thus, if the majority of the requests in round j fall in the last two cases, a cost of $\Omega(\alpha)$ is incurred for the round. Otherwise, the majority of the requests in the round are for terminals u which incur cost roughly the cost of the directed path from r to u , which translates to a total cost of $\Omega(k_j)$, where k_j is the number of requests in K_j . This argument shows that each round contributes a cost of $\Omega(\min\{k_j, \alpha\})$. Since there are x rounds, the result will then follow by combining the contribution of each round to the overall cost, and the observation that the the optimal algorithm will buy the path P and all edges of the form (v, u) , for a total cost which can be shown that it is bounded by a constant.

Details on the formal proof can be found in Section 5.

5 Proof of Theorem 2

5.1 Construction of the adversarial graph

The construction of the adversarial input is based on a recursive definition of a suitable graph G . In order to define G we will need first to define certain auxiliary constructions which will facilitate the description of the input graph. First, let v, u be two vertices (in the vertex set of a certain graph) with the property that edge $e = (v, u)$ has cost $c(e) = c$, while the antiparallel edge $\bar{e} = (u, v)$ has cost αc , where α is the asymmetry of the graph in which v, u belong. We then say that we *insert a vertex w at height h in e* with $h < c$ if we introduce a new vertex w and replace e, \bar{e} with new edges of costs $c(v, w) = c - h$, $c(w, u) = h$, $c(w, v) = \alpha(c - h)$, $c(u, w) = \alpha h$ (for the sake of visualization, we should think of v as being located higher than u). Note that the insertion maintains the asymmetry of the graph.

Second, let $T_1 = \{v_1, \dots, v_l\}$ and $T_2 = \{u_1, \dots, u_l\}$ be two disjoint sets of l vertices each (again, we may think of vertices of T_1 as located higher than vertices of T_2). In addition, we require that T_1 and T_2 have the property that all edges of the form $e_i = (v_i, u_i)$ have the same cost, say c , while all antiparallel edges \bar{e}_i have cost αc . Let E denote the set $\{e_i | i \in [1, l]\}$. We call index $i \in [1, l]$ the *i -th column*. On the collection of edges E we define a construction denoted by (E, q, g) which we call *layered component* or simply *component* and which, informally, adds vertices in edges in E in *layers*. Here, q and g are parameters used in the construction. The construction is as follows (see also Figure 4). Layer 1 consists of l vertices inserted at height c/q , one for each edge in E . Call those vertices $w^{1,1}, \dots, w^{1,l}$. We group these vertices into $2^1 = 2$ groups, namely $S^{1,1} = \{w^{1,1}, \dots, w^{1,l/2}\}$ and $S^{1,2} = \{w^{1, \frac{l}{2}+1}, \dots, w^{1,l}\}$. We also add two new vertices, $u^{1,1}$ and $u^{1,2}$, such that for every vertex $w \in S^{1,i}$, with $i \in \{1, 2\}$, there is an edge of cost g from $w^{1,i}$ to $u^{1,i}$ and an edge of cost αg from $u^{1,i}$ to $w^{1,i}$. Recursively, suppose that layers $1, 2, \dots, j-1 \leq q-2$ have been defined, we show how to derive layer j . Let E' be the collection of edges of the form $(v_i, w^{j-1,i})$, after layers $1, \dots, j-1$ have been created. By construction, all of them have the same cost, which we denote by c' . Again, we insert l new vertices $w^{j,1}, \dots, w^{j,l}$ at height c/q for each edge in E' which we then partition (left-to-right) in 2^j groups $S^{j,1}, \dots, S^{j,2^j}$, all of the same size; namely, group $S^{j,i}$ consists

of vertices $\{w^{j,i(2^{l-j})+1}, \dots, w^{j,(i+1)(2^{l-j})}\}$. We also add 2^j new vertices $u^{j,1}, \dots, u^{j,2^j}$ such that for every vertex $w \in S^{j,i}$, there is an edge from w to $u^{j,i}$ of cost g , and an edge from $u^{j,i}$ to w of cost αg . We stop when $q - 1$ layers in total have been defined.




Figure 4: The structure of a component for the case $q = 3$, and $l = 8$.

We will call a set of the form $S^{i,j}$ an s -set. We will also say that a set of vertices S *crosses* or *intersects* a certain set of columns if and only if the set of columns in which the vertices of S lie intersect the set of columns in question. Two sets of vertices *cross each other* iff the intersection of the columns each one crosses is non-empty. Similar definitions apply for a set of edges: a directed edge which lies on column i is said to cross column i . Note also that there is a 1-1 correspondence between sets $S^{i,j}$ and vertices $u^{i,j}$, which means that several properties/definitions pertaining to S -sets carry over to their corresponding u -vertices: we will use this correspondence to make the discussion more accessible (at a slight abuse of notation).

Note that in the layered component, a set $S^{j,i}$ at layer j crosses a set of columns $J = J_1 \cup J_2$ (where J_1 and J_2 are disjoint sets of columns) such that one s -set at layer $j + 1$ crosses J_1 (and only J_1) and another s -set in layer $j + 1$ crosses J_2 (and only J_2). We call these two s -sets the *children* of $S^{j,i}$ (or we say that $S^{j,i}$ is their *parent*). We extend this definition to the u -vertices to which the s -sets described above correspond. Namely, the children of $u^{j,i}$ are the u -vertices corresponding to the children of set $S^{j,i}$. By convention, we will also say that the children of T_2 are $S^{1,1}$ and $S^{1,2}$.

We now proceed with the description of the adversarial input graph G . At a high level, the construction is based on a recursive series of insertions of appropriately defined layered components. In particular, we will show how to construct a family of graphs G_0, G_1, \dots, G_ρ : G will then be defined as G_ρ for a suitable choice of the value ρ . First, denote by G_0 the graph which consists of two sets of vertices $T_1 = \{v_1, \dots, v_l\}$ and $T_2 = \{u_1, \dots, u_l\}$, as well as vertex r . The root of the Steiner tree is set to be r , and the same will hold for every G_i . It will be convenient to think of l as very large compared to k , although suffices to set $l = 2^k$. Edges from r to vertices in T_1 have all zero cost (or, more precisely, infinitesimally small cost), and the same holds for their antiparallel edges. In addition, $c(v_i, u_i) = 1$ and $c(u_i, v_i) = \alpha$. Informally, the “downwards” direction is the cheap one, while the “upwards” direction is the expensive; this property will be maintained throughout the recursive construction. We let E denote the set of edges (v_i, u_i) . We will use the words “down” and “up” to distinguish between edges in the graphs, with the natural meaning of the words.

For the sake of clarity, we first present the construction of G_1 , the construction of G_i for larger values of i will follow next. Suppose that x' is such that $(x')^{(x')} = k$, which means that $x' = \Theta(\frac{\log k}{\log \log k})$ (recall that k denotes the total number of requests). We let x denote $\lfloor x' \rfloor$. G_1 is

derived by inserting only one component, namely $C_1 = (E, x, 1/x^2)$. The *height* H_1 of G_1 is defined as the total number of layers added by the construction, plus one (to account for T_1 and T_2 which we may think of as layers on their own; we define the height of T_2 to be zero). After the insertion of the component, we partition the set of all l columns into a collection of $\frac{l}{2^x}$ disjoint sets of columns $R = \{R_1, \dots, R_{\frac{l}{2^x}}\}$, with the property that column $y \in [1, l]$ is in class R_i if and only if it crosses $S_1^{x-1, i}$. In other words, there is a class in R for each of the highest s -sets in the component C_1 .

We will also use the notation $L_1^{j, m}$ to denote the m -th (from left to right) s -set in G_1 at height j . By convention we consider $L_1^{0, 1} \equiv T_2$, while $L_1^{x, 1} = T_1$, to reflect that T_1 and T_2 constitute the bottom and the top level, respectively.

Suppose now that graphs G_0, \dots, G_i have been defined. Let H_i denote the height of G_i , namely the total number of layers of s -sets inserted, plus one. We say that two s -sets of vertices (possibly including T_1 and T_2) are *consecutive* in G_i if and only if they cross each other, and in addition there is no other s -set of vertices between them which crosses either set. As expected, an s -set is “between” two s sets if it is located above one of the two sets and below the other.

We now describe how to define G_{i+1} . The construction is presented in pseudocode. Similar to the case of $i = 1$, we use the notation $L_i^{j, m}$ to denote the m -th (from left to right) s -set in G_i at height j . L_i^j is then defined simply as the union of all $L_i^{j, m}$.

```

for  $j = 0$  to  $H_i - 1$  do
  for every  $m$  such that  $L \equiv L_i^{j, m} \in L_i^j$  do
    Let  $C(L) \subseteq L_i^{j+1}$  denote the collection of  $s$ -sets in  $L_i^{j+1}$  with the additional property
    that each  $s$ -set in  $C(L)$  is consecutive to  $L$  in  $G_i$  ;
    for every  $L' \in C(L)$  do
      Partition the set of columns crossed by both  $L'$  and  $L$  into a collection of classes
       $P = \{P_1, \dots\}$  such that  $P_j$  consists of columns in one of the classes in  $R$  and only
      that class;
      for each class  $P_i$  do
        Let  $E'$  be the set of edges from  $L'$  to  $L$  which are on the same column with class
         $P_i$  ;
        Perform the layered insertion  $(E', x, 1/x^{i+2})$ . ;
      end
    end
  end
  Update the partition of the column space  $R$ : Redefine  $R$  such that all columns crossed by
  the same highest  $s$ -set inserted in this iteration(i.e., for the current value of  $j$ ) are placed
  in the same partition ;
end

```

An example of the construction is given in Figure 5
The following lemma is easy to show (proof omitted).

Lemma 16 *Let l_i denote the total number of levels in G_i (including T_1, T_2). Also let n_i denote the number of “new” levels added when obtaining G_i from G_{i-1} (i.e., $n_i = l_i - l_{i-1}$). Then $l_i, n_i \in \Theta((x+1)^i)$. In addition for $\rho = x$, there are $\Theta(k)$ levels in G_ρ .*

Graph G_ρ will be the adversarial input graph, i.e., $G \equiv G_\rho$. In the next section we present the adversarial strategy, namely the game between the adversary and the deterministic algorithm. We say that a u -vertex (resp. S -set) is a *vertex/set* of G_i if it is present in G_i but not in G_{i-1} . Note that such a vertex is present in all graphs $G_{j'}$ with $j' > i$. For simplicity we will assume wlog that




Figure 5: Example of the construction of the adversarial graph. Bold rectangles depict the layers of the only component of G_1 . The eight s -sets of the highest layer of this component induce a partition of the column space into 8 classes. Shaded rectangles correspond to components of the form $G_{2,0,id}$, with $id = 1 \dots 8$.

G_ρ has exactly k levels. (the proof is only slightly more complicated if this is not the case, e.g. we must set $l = 2^{\Theta(k)}$ instead of 2^k .)

Having defined the construction of G_i , we need a more convenient way of identifying all components present in graph G_i . In particular we will use the notation $C_{i,l,id}$ to refer to the component inserted during the construction of graph G_i between levels l and $l+1$ of G_{i-1} (which means that l ranges between 0 and $H_{i-1} - 1$) and has a certain component id , which is simply determined by numbering the components from left to right for each such fixed value of l . Within a specific component $C_{i,l,j}$, we use the notation $S_{i,l,j}^{m_1,m_2}$ to define the m_1 -th s -set from bottom to the top of the component (i.e., located at the m_1 -th layer of the component) which is also the m_2 -th s -set from left to right among all s -sets in that layer). The corresponding u vertices will be identified in the natural way, i.e. $u_{i,l,j}^{m_1,m_2}$ is the u -vertex corresponding to $S_{i,l,j}^{m_1,m_2}$.

5.2 The algorithm/adversary game

At a high level, the game proceeds in rounds. Only u vertices are requested. In round i , the algorithm will request vertices of G_i only in a bottom-up fashion, i.e., from lowest to highest height. The important property that the adversary will guarantee is that all requested u -vertices have corresponding s -sets which all lie on the same unique path of down-edges from r . This will ensure that the optimal cost is bounded by a constant. On the other side, we will show that the algorithm will have to pay a high cost per round (roughly $\Theta(\min(\{k_i, x\}))$, where k_i is the number of vertices requested at round i . This means that the algorithm will pay (roughly) either a constant cost per request, or a cost proportional to x per round.

We begin with certain preliminary definitions and conventions. Each time the adversary presents a new request, namely a certain u -vertex, the algorithm must guarantee there is a (directed) path from a previously requested vertex to u , possibly buying some new edges. Among all possible such

paths the adversary will fix/choose *one* such path, which we call a *connection path*. For this path, our analysis will *charge* only parts of it (i.e. specific edges) to the algorithm. If we ensure that each edge bought by the algorithm is charged at most once, then the total cost of all charged edges cannot exceed the actual cost paid by the online algorithm. Without loss of generality we will assume that the connection path is acyclic (otherwise we simply bypass all cycles).

Consider now a u -vertex requested in round i , and its connection path $p(u)$ chosen by the adversary. We observe that $p(u)$ can be chosen so that it is one of the following three types:

- Connection path *from r* : A connection path which originates from r and does not visit any previously requested vertex.
- Connection from above: A connection path which originates from a previously requested vertex which lies *higher* than u in G_i (and visits no other previously requested vertices).
- Connection from below: A connection path which originates from a previously requested vertex which lies *lower* than u in G_i (and visits no other previously requested vertices).

To illustrate the main ideas of the game, we describe the intuition behind the actions of the adversary during the first round of the game, namely when the adversary requests u -vertices in G_1 . Later we will describe the game in a more formal way. Since there is only one component in G_1 , we omit subscripts for simplicity, in this example.

The first round begins with request $u^{1,1}$ which corresponds to $S^{1,1}$. The algorithm will buy a connection path from r , say $p(u^{1,1})$. The adversary will then charge “down”-edges in this path (but no other edges). Observe that no matter what the connection path is, there is a child of $S^{1,1}$ in the second layer of the component, which does not cross any columns of the charged edges. The adversary will choose this child as its second request; call this child u , and its s -set S . At this point, there are two choices concerning the connection path $p(u)$:

- *Case 1.* $p(u)$ is of the form shown in Figure 6: namely, it consists of down-edges, from r down to (at least) the level of $u^{1,1}$. In this case we charge the algorithm with the cost of one down-edge per level, from r down to the level of $u^{1,1}$ in G_1 .
- *Case 2.* The connection path is as in Figure 6: namely the connection path consists of “up”-edges between the level of $u^{1,1}$ up to (at least) the level of u . We charge the algorithm with the cost of one up-edge per level from the level of $u^{1,1}$ up to the level of u .




Figure 6: Depictions of Case 1 (left) and Case 2 (right).

In either of the above cases, let $q(p)$ denote the part of the connection path $p(u)$ (i.e., collection of edges) which is charged. Hence depending on which case applies, different edges will be charged, and different costs will be incurred. The critical observation (along the lines of the earlier observation concerning $u^{1,1}$) is that exactly one of the children of u in C will correspond to an s -set which does not cross columns with the columns of $q(p)$; call this child u' . This provides a good strategy for the adversary to choose its next request; in fact u' will be the next request. In particular, once again, the connection path to u' will either result in a charge of “down”-edges from r down to the level of u' , or a charge of up-edges from the level of u up to the level of u' . We continue the game until we reach a vertex at the top layer (layer $x - 1$) of the component, at which point the first round concludes⁸ The algorithm yields a charge for (almost) all the layers in the component: if the majority of such requests were charged as in case (1), then the total charge is $\Omega(k_1)$, where k_1 is the number of requests in graph G_1 ; otherwise the total charge due to requests charged as in case (2) will clearly be $\Omega(\alpha)$. Therefore, the charge incurred during round 1 is $\Omega(\min\{k_1, \alpha\})$.

For round 2, the adversary will play a strategy defined along the same lines; however, instead of requesting vertices within a single component only, the adversary will request vertices in H_1 components in total, where H_1 is the height of G_1 . More precisely, round 2 begins with identifying the set of *critical columns*, namely those columns which cross the highest s -vertex requested in round 1 (which was also the last vertex requested in round 1); then the adversary identifies the unique component of the form $C_{2,0,id}$ for some id which crosses exclusively the critical set of columns. For this component, the algorithm plays a game very similar to the one corresponding to round 1 (which, recall, takes place within a single component). Once “done” with this component, the critical set of columns is updated to reflect the set of columns crossed by the last request (or more precisely, the s -set of the last request) so far; the algorithm then proceeds with the next unique lowest component in G_2 which crosses only edges of the critical set (a component of the form $C_{2,1,id'}$ for some id'). This repeats until the algorithm has completed requests within H_1 components, in which case round 3 begins. The game continues for ρ rounds.

In the rest of this section we give a more formal description of the game. Consider the input graph G and any two s -sets s_1 and s_2 such that s_1 is higher than s_2 and the two s -sets cross each other. We say that a directed path p *includes all down-edges between s_1 and s_2* (resp. all up-edges) if the path includes one down-edge (resp. up edge) per level, for all levels between s_1 and s_2 .

The following is a basic property concerning the connection paths which can be derived by observing the structure of the request graph (proof omitted).

Claim 17 *Let C be a certain component in G and let $S^{i,j}$, with $i < x - 1$ be an s -set in layer i of C with children S^{i+1,j_1}, S^{i+1,j_2} such that: the corresponding vertex $u^{i,j}$ has already been requested, and no other u -vertex in C located between the level of $S^{i,j}$ and the level of its children in G has been requested. Suppose that the next request, is a child of $u^{i,j}$, say u^{i+1,j_1} . Let $p = p(u^{i+1,j_1})$ be the connection path for this request. Then:*

- *If p is a path from r , then the connection path will always include all down-edges between r and S^{i+1,j_1} .*
- *If p is a path from below, then then the connection path will always include all upward edges between $S^{i,j}$ and S^{i+1,j_1} .*

⁸In order to ensure that no edge is charged more than once, we will not charge any cost incurred at the first and top $(x - 1)$ -th layers of each component. Since this only happens for 2 our of the $x - 1$ layers, this technical point will not affect the analysis.

- If p is a path from above, say from a previously requested vertex u' , then the connection path will always include the edge (u', s') from u' to its corresponding s -set s' such that S^{i+1, j_1} crosses s' .

In addition, if $q(p)$ denote either i) the set of up-edges between $S^{i, j}$ and S^{i+1, j_1} ; or ii) the set of down-edges between r and S^{i+1, j_1} ; or iii) the edge (u', s') (depending on which case applies) then exactly one of the children of $S^{i+1, j}$ will not cross columns with the columns of $q(p)$.

The first part of the claim suggests that the connection paths have a certain structure. The second part of the claim suggests a way for the adversary to request the “appropriate” child of the last request as the next request to be presented. In particular, let us denote by \bar{u} the last requested vertex, and $p(\bar{u})$ the corresponding connection path for \bar{u} . Then, the second part of the claim shows that given $q(p(\bar{u}))$, as defined in the statement of the claim, there is a unique child of \bar{u} (within the component to which \bar{u} belongs) which does not cross the columns of $q(p(\bar{u}))$. We will use the notation $ch(q(p(\bar{u})), \bar{u})$ to denote this unique vertex.

To make the above more precise, the adversary will play the game as follows:

```

Initialize the critical set to the set of all columns  $\{1, \dots, l\}$  ;
for  $i = 1$  to  $\rho$  do
    REQUESTS( $G_i$ )
end

```

where REQUESTS(G_i) is the strategy for requesting vertices which belong in G_i , described in what follows:

```

for  $l = 0$  to  $H_i - 1$  do
    Find the unique component  $C_{i, l, y}$  (for some index  $y$ ) which crosses all columns of the critical set ;
    for  $j = 1$  to  $x - 1$  do
        if  $j = 1$  then
            request  $u = u_{i, l, y}^{1, 1}$  ;
            update  $\bar{u} \leftarrow u$  ;
        end
        if  $j = x - 1$  then
             $u = ch(q(p(\bar{u})), \bar{u})$  ;
            update the critical set to the set of columns crossed by the  $s$ -set to which  $u$  corresponds. ;
        end
        else
            request  $u = ch(q(p(\bar{u})), \bar{u})$  ;
            for the connection path  $p(u)$  chosen update  $\bar{u} \leftarrow u$  and  $p(\bar{u}) \leftarrow p(u)$  ;
        end
    end
end

```

From the structure of the graph, one can easily prove the following claim:

Claim 18 *Let M_i denote the number of columns with the property that they cross the s -sets of all vertices requested up to and including the i -th request. Then $M_i = \frac{l}{2^i}$.*

Proof sketch. First, observe that the set of columns crossed by the s -set which corresponds to the i -th request crosses all s -sets for vertices requested in iterations $1 \dots i - 1$. In addition, this set crosses, by construction, exactly $l/2^i$ columns. \square

Claim 18 implies that the adversary/algorithm game is feasible, in the sense that it can go on for k iterations, namely as many as the number of terminals requested by the adversary.

5.3 Analysis

Last, we need to upper-bound the cost of the optimal offline algorithm, and lower-bound the cost of any deterministic online algorithm. For the former, suffices to use Claim 18. An offline algorithm can buy a path with r as its origin, consisting exclusively of down-edges, and which visits all s -sets corresponding to all u -vertices requested by the adversary, at a cost of at most 1. In addition, it will buy all directed edges from the s -sets in question to all u -vertices requested. Recall that the adversary requests $\Theta((x+1)^i)$ vertices which belong to graph G_i (Lemma 16). Therefore, the total cost for such edges is bounded by

$$\sum_{i=1}^{\rho} \frac{1}{x^{i+1}} \Theta((x+1)^i) = \frac{1}{x} \sum_{i=1}^x \left(1 + \frac{1}{x}\right)^i \leq e,$$

where e denotes the base of the natural logarithm. Hence the algorithm's total cost is upper-bounded by $e + 1$.

In order to place a lower bound on the cost of any online algorithm, we will charge a subset of the edges of each connection path according to specific rules. The charging will be such that every edge is charged at most once, which guarantees that the total charged cost is a lower bound for the cost paid by the algorithm.

Let u be a vertex requested at round i (i.e., a vertex which belongs in graph G_i), and $p(u)$ its corresponding connection path. Recall that Claim 17 provides a classification of the connection paths. Our charging scheme is based on this classification. In particular, we will charge certain edges in $q(p)$, and only such edges (with the exception of the lowest and highest layers in a component as argued earlier). There are three cases:

- *Case 1:* $q(p)$ describes up-edges. In this case, the cost charged to the algorithm is $\alpha \cdot \frac{1}{(x+1)^i}$ (since $q(p)$ includes all up-edges between two consecutive levels at graph G_i).
- *Case 2:* $q(p)$ describes down-edges from s . In this case, the cost charged to the algorithm is the cost of all such down-edges, down to the level of u .
- *Case 3.* $q(p)$ describes an edge of the form (u', s') such that u' is a previously requested vertex. Since u' belongs to some G_j , with $j < i$, the charge in this case is at least $\frac{\alpha}{x^{j+1}} \geq \frac{\alpha}{x^i}$.

The second part of Lemma 17 guarantees that each edge in the graph is charged only once.

It remains to express the total charge incurred by the algorithm. Let k_i^1 and k_i^2 denote the number of requests which incur a charge of at least $\alpha \frac{1}{(x+1)^i}$ (Cases (1) and (3) above) and a charge which follows Case (2), respectively. Recall that k_i denote the number of requests in graph G_i . From Lemma 16 we have that $k_i^1 + k_i^2 = \Theta((x+1)^i)$. It follows that if $k_i^1 \geq k_i/2$ then the total charge for round i is $\Omega(\alpha)$, otherwise, it is not difficult to see that the total charge is $\Omega(k_i)$. This observation holds for all levels i ; let r_i be the indicator variable which is 0 if the charge of round i is $\Omega(\alpha)$ and 1 if the charge is $\Omega(k_i)$.

Let R_1 denote the set $\{i : r_i = 1\}$ and R_0 the set $\{i : r_i = 0\}$, for all $i \in [1, \rho]$. Suppose that $|R_0| \geq \rho\epsilon$, for some constant $\epsilon \geq 0$. In this case, the total charge is $\Omega(\alpha\rho) = \Omega(\alpha x)$. Otherwise, $|R_1| \geq \rho(1 - \epsilon)$. Therefore the total charge due to all rounds in R_1 is at least $\Omega(\sum_{i=1}^{\rho(1-\epsilon)} (x+1)^i) = \Omega(k^{1-\epsilon})$.

Summarizing, the total cost paid by the algorithm is at least the total charge, which in turn is bounded as shown above, and thus is

$$\Omega(\max\{k^{1-\epsilon}, \alpha x\}) = \Omega(\max\{k^{1-\epsilon}, \alpha \frac{\log k}{\log \log k}\}).$$

Acknowledgements. I am grateful to Michalis Faloutsos for communicating the problem to me, and for some very interesting preliminary discussions. Many thanks also to Alejandro López-Ortiz and Reza Dorrigiv for their helpful comments.

References

- [1] A. Agrawal, P. N. Klein, and R. Ravi. When trees collide: An approximation algorithm for the generalized steiner tree problem on networks. *SIAM Journal on Computing*, 24:440–456, 1995.
- [2] N. Alon and Y. Azar. On-line steiner trees in the euclidean plane. *Discrete and Computational Geometry*, 10:113–121, 1993.
- [3] Baruch Awerbuch, Yossi Azar, and Yair Bartal. On-line generalized steiner problem. In *Proceedings of the seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 68–74, 1996.
- [4] P. Berman and C. Coulston. Online algorithms for Steiner tree problems. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing*, pages 344–353, 1997.
- [5] M. Bern and P. Plassmann. The Steiner problem with edge lengths 1 and 2. *Information Processing Letters*, 32:171–176, 1989.
- [6] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998. Borodin.
- [7] M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li. Approximation algorithms for directed steiner problems. *Journal of Algorithms*, 1(33):73–91, 1999.
- [8] K. Claffy, G. Polyzos, and H.W. Braun. Traffic characteristics of the t1 nsfnet backbone. In *Proceedings of INFOCOM*, 1993.
- [9] M. Faloutsos. *The Greedy the Naive and the Optimal Multicast Routing—From Theory to Internet Protocols*. PhD thesis, University of Toronto, 1998.
- [10] M. Faloutsos, R. Pankaj, and K. C. Sevcik. The effect of asymmetry on the on-line multicast routing problem. *Int. J. Found. Comput. Sci.*, 13(6):889–910, 2002.
- [11] M.X. Goemans and D.P. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 6(24), 1995.
- [12] M. Imase and B. Waxman. The dynamic Steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384, 1991.
- [13] S. Ramanathan. Multicast tree generation in networks with asymmetric links. *IEEE/ACM Trans. Netw.*, 4(4):558–568, 1996.
- [14] G. Robins and A. Zelikovsky. Improved Steiner tree approximation in graphs. In *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pages 770–779, 2000.
- [15] A. Segev. The node-weighted steiner tree problem. *Networks*, (17):1–17, 1987.
- [16] M. Thimm. On the approximability of the Steiner tree problem. *Theoretical Computer Science*, 295(1):387–402, 2003.

Appendix

A Dealing with degeneracies in the proof of Lemma 6

There are many ways to argue about this. A formal way is as follows. Suppose first that $v_i \equiv u_i$, then we can create a new vertex u'_i , add a directed edge of infinitesimally small, but positive cost ϵ , treat u'_i as a terminal while the vertex v_i as a non-terminal vertex. Likewise, if $v_{i-1} \equiv v_i$, then we can add a directed edge of cost ϵ from v_i to v_{i-1} in the comb. We charge only the algorithm with the additional cost, but not the optimal solution. Since ϵ is negligible, the overhead does not affect asymptotically the algorithm's cost.

A different way to argue is that even if such degeneracies arise, they do not affect the proof of Lemma 6. The only reason they should be nevertheless considered is that Lemma 4 requires that Lemma 6 is correct even when such degeneracies arise.

B An example of the partition of the terminals in a comb into runs

Consider the comb instance of Figure 7, with K' consisting of 12 terminals (for simplicity, we represent terminal u_i , with $1 \leq i \leq 12$, by its index i). Suppose the terminals are given in the order

$$3,7,1,11,4,6,5,2,8,12,10,9$$

The decomposition is then as follows:

run	representative	terminals
1	–	3,7,11,12
2	3	1,2
3	7	4,6
4	6	5
5	11	8,10
6	10	9




Figure 7: An example of the comb decomposition.