

LAST but not least: online spanners for buy-at-bulk

Citation for published version (APA):

Gupta, A., Ravi, R., Talwar, K., & Umboh, S. W. (2017). LAST but not least: online spanners for buy-at-bulk. In P. N. Klein (Ed.), *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms* (pp. 589-599). Society for Industrial and Applied Mathematics (SIAM). <https://doi.org/10.1137/1.9781611974782.38>

DOI:

[10.1137/1.9781611974782.38](https://doi.org/10.1137/1.9781611974782.38)

Document status and date:

Published: 01/01/2017

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

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

[Link to publication](#)

General rights

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

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

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

www.tue.nl/taverne

Take down policy

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

openaccess@tue.nl

providing details and we will investigate your claim.

LAST but not Least: Online Spanners for Buy-at-Bulk

Anupam Gupta* R. Ravi† Kunal Talwar‡ Seeun William Umboh§

November 2, 2016

Abstract

The online (uniform) buy-at-bulk network design problem asks us to design a network, where the edge-costs exhibit economy-of-scale. Previous approaches to this problem used tree-embeddings, giving us randomized algorithms. Moreover, the optimal results with a logarithmic competitive ratio requires the metric on which the network is being built to be known up-front; the competitive ratios then depend on the size of this metric (which could be much larger than the number of terminals that arrive).

We consider the buy-at-bulk problem in the least restrictive model where the metric is not known in advance, but revealed in parts along with the demand points seeking connectivity arriving online. For the single sink buy-at-bulk problem, we give a deterministic online algorithm with competitive ratio that is logarithmic in k , the number of terminals that have arrived, matching the lower bound known even for the online Steiner tree problem. In the *oblivious* case when the buy-at-bulk function used to compute the edge-costs of the network is not known in advance (but is the same across all edges), we give a deterministic algorithm with competitive ratio polylogarithmic in k , the number of terminals.

At the heart of our algorithms are optimal constructions for online Light Approximate Shortest-path Trees (LASTs) and spanners, and their variants. We give constructions that have optimal trade-offs in terms of cost and stretch. We also define and give constructions for a new notion of LASTs where the set of roots (in addition to the points) expands over time. We expect these techniques will find applications in other online network-design problems.

*Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Research partly supported by NSF awards CCF-1319811, CCF-1540541 and CCF-1617790.

†Tepper School of Business, Carnegie Mellon University, USA, ravi@cmu.edu; This material is based upon research supported in part by the U. S. Office of Naval Research under award number N00014-12-1-1001, and the U. S. National Science Foundation under award number CCF-1527032.

‡Google Research

§Department of Mathematics and Computer Science, Eindhoven University of Technology, Netherlands. Email: seeun.umbogh@gmail.com. This work was supported in part by ERC consolidator grant 617951 and NSF grant CCF-1320854. Part of this work was done while a student at the University of Wisconsin - Madison, and while visiting the Simons Institute for the Theory of Computing.

Contents

1	Introduction	2
2	Preliminaries	6
3	LASTs	8
4	Multi-Sink LASTs	9
5	Non-Oblivious Buy at Bulk	13
6	Multi-Commodity Spanners	18
7	Randomized Oblivious Buy at Bulk	22
8	Deterministic Oblivious Buy at Bulk	27
9	Conclusions	30

1 Introduction

The model of (uniform) buy-at-bulk network design captures economies-of-scale in routing problems. Given an undirected graph $G = (V, E)$ with edge lengths $d : E \rightarrow \mathbb{R}_{\geq 0}$ —we can assume the lengths form a metric—the cost of sending x_e flow over any edge e is $d(e) \cdot f(x_e)$ where f is some concave cost function. The total cost is the sum over all edges of the per-edge cost. Given some traffic matrix (a.k.a. demand), the goal is now to find a routing for the demand to minimize the total cost. This model is well studied both in the operations research and approximation algorithms communities, both in the offline and online settings. In the offline setting, an early result was an $O(\log k)$ -approximation due to Awerbuch and Azar, one of the first uses of tree embeddings in approximation algorithms [AA97]—here k is the number of demands. For the single-sink case, the first $O(1)$ -approximation was given by [GMM09]. In fact, one can get a constant-factor even for the “oblivious” single-sink case where the demands are given, but the actual concave function f is known only after the network is built [GP12].

The problem is just as interesting in the *online context*: in the online single-sink problem, new demand locations (called *terminals*) are added over time, and these must be connected to the central root node as they arrive. This captures an increasing demand for telecommunication services as new customers arrive, and must be connected via access networks to a central core of nodes already provisioned for high bandwidth. The Awerbuch-Azar approach of embedding G into a tree metric T with $O(\log n)$ expected stretch (say using [FRT04]), and then routing on this tree, gives an $O(\log n)$ -competitive randomized algorithm even in the online case. But *this requires that the metric is known in advance*, and the dependence is on n , the number of nodes in the metric, and not on the number of terminals k ! This may be undesirable in situations when $n \gg k$; for example, when the terminals come from a Euclidean space \mathbb{R}^d for some large d . Moreover, we only get a *randomized* algorithm (competitive against oblivious adversaries).¹

In this paper, we study the Buy-at-Bulk problem in the online setting, in the least restrictive model where the metric is not known in advance, so the distance from some point to the previous points is revealed only when the point arrives. This forces us to focus on the problem structure, since we cannot rely on powerful general techniques like tree embeddings. Moreover, we aim for deterministic algorithms for the problem. Our first main result is an asymptotically optimal deterministic online algorithm for single-sink buy-at-bulk.

Theorem 1.1 (Deterministic Buy-at-Bulk). *There exists a deterministic $O(\log k)$ -competitive algorithm for online single-sink buy-at-bulk, where k is the number of terminals.*

Note that the guarantee is best possible, since it matches the lower bound [IW91] for the special case of a single cable type encoding the online Steiner tree problem.

En route, we consider a generalization of the *Light Approximate Shortest-path Trees* (LASTs). Given a set of “sources” and a sink, a LAST is a tree of weight close to the minimum spanning tree (MST) on the sources and the sink, such that the tree distance from any source to the sink is close to the shortest-path distance. Khuller, Raghavachari, and Young [KRY95] defined and studied LASTs in the offline setting and showed that one can get constant stretch with cost constant times the MST. Ever since their work, LASTs have proved very versatile in network design applications. We

¹The tree embeddings of Bartal [Bar96] can indeed be done online with $O(\log k \log \Delta)$ expected stretch, where Δ is the ratio of maximum to minimum distances in the metric. Essentially, this is because the probabilistic partitions used to construct the embedding can be computed online. This gives an $O(\log^2 k)$ -competitive randomized algorithm, alas sub-optimal by a logarithmic factor, and still randomized.

first give (in §3) a simple construction of LASTs in the online setting where terminals arrive online. We get constant stretch, and cost at most $O(\log k)$ times the MST, which is the best possible in the online case.

For our algorithms, we extend the notion of LASTs to the setting of MLASTs (*Multi-sink LASTs*) where both sources *and* sinks arrive over time. We have to maintain a set of edges so that every source preserves its distance to the closest sink arriving before it, at minimum total cost. We provide a tight deterministic online algorithm also for MLASTs, which we think is of independent interest. This construction appears in §4. Then we use MLASTs to prove Theorem 1.1 in §5.

Oblivious Buy-at-Bulk. We then change our focus to the *oblivious* problem. Here we are given neither the terminals nor the buy-at-bulk function f in advance. When the terminals arrive, we have to choose paths for them to the root, so that for every concave cost function f , the chosen routes are competitive to the optimal solution for f . Our first result for this problem is the following:

Theorem 1.2 (Oblivious Buy-at-Bulk, Randomized). *There exists a randomized online algorithm for the buy-at-bulk problem that produces a routing P such that for all concave functions f ,*

$$\mathbb{E}[\text{cost}_f(P)] \leq O(\log^2 k) \text{OPT}_f.$$

This randomized algorithm has the same approximation guarantee as one obtained using Bartal’s tree-embedding technique. The benefit of this result, however, is in the ideas behind it. We give constructions of low-stretch spanners in the online setting. Like LASTs, spanners have been very widely studied in the offline case; here we show how to maintain light low-stretch spanners in the online setting (in §6). Then we use the spanners to prove the above theorem in §7. Moreover, building on these ideas, we give a deterministic algorithm in §8.

Theorem 1.3 (Oblivious Buy-at-Bulk, Deterministic). *There exists a deterministic online algorithm for the buy-at-bulk problem that produces a routing P such that for all concave functions f ,*

$$\text{cost}_f(P) \leq O(\log^{2.5} k) \text{OPT}_f.$$

A question that remains open is whether there is an $O(\log k)$ -competitive algorithm for this problem. The only other deterministic oblivious algorithm we know for the buy-at-bulk problem is a derandomization of the oblivious network design algorithm from [GHR06], which requires the metric to be given in advance.

LASTs and Spanners. A central contribution of our work is to demonstrate the utility of online spanners in building networks exploiting economies-of-scale. We record the following two theorems on maintaining LASTs and spanners in an online setting, since they are of broader interest. These results are near-optimal, as we discuss in the respective sections.

Theorem 1.4 (Online LAST). *There exists a deterministic online algorithm for maintaining a tree with cost $O(\log k)$ times the MST, and a stretch of 7 for distances from terminals to the sink.*

Theorem 1.5 (Online Spanner). *There exists a deterministic online algorithm that given k pairs of terminals, maintains a forest with cost $O(\log k)$ times the Steiner forest on the pairs, and a stretch of $O(\log k)$ for distances between all given pairs of terminals. Moreover, the total number of edges in the forest is $O(k)$, i.e. linear in the number of terminals.*

Our results on online LAST and multicommodity spanners give us an optimal $O(\log k)$ -competitive deterministic algorithms for the “single cable version” of the buy-at-bulk problems, where the function is a single-piece affine concave function [MP98].

1.1 Our Techniques

We now outline some of the ideas behind our algorithms and the role of spanner constructions in them. All our algorithms share the same high-level framework: (1) each terminal v is assigned an integer *type* $\tau(v)$; (2) it is then routed through a sequence of terminals with increasing types; (3) the routes are chosen using a spanner-type construction. In particular, each algorithm is specified by a *type-selection rule*, the sequence of terminals to route through, and whether to use a spanner or MLAST.

The analysis for our deterministic algorithms also has a common thread: while these algorithms are not based on tree-embeddings, the analysis uses tree-embeddings crucially. Both analyses are based on the analysis framework of [Umb15] and follow the same template: we fix an HST embedding T of the metric, and charge the cost of the algorithm to the cost of the optimal solution on this tree (losing some factor α). Since HSTs can approximate general metrics to within a factor of $O(\log k)$, this gives us an $O(\alpha \log k)$ -competitive algorithm.

Functions vs. Cables. As is common with buy-at-bulk algorithms, we represent the function $f(x)$ as the minimum of affine functions $\min_i \{\sigma_i + \beta_i x\}$. And when we route a path, we even specify which of the linear functions we use on each edge of this path. Each i is called a “cable type”, since one can think of putting down cable- i with an up-front cost of σ_i , and then paying β_i for every unit of flow over it.

Non-oblivious algorithm. In the non-oblivious setting where we know the function f , we will want to route each terminal v through a path $P(v)$ with non-decreasing cable types. So $\tau(v)$ should simply be the cable of lowest type that we will install on $P(v)$ —how should we choose this value? It makes sense to choose $\tau(v)$ based on the number of other terminals close to v . Intuitively, the more terminals that are nearby, the larger the flow that can be aggregated at v , making it natural to select a larger type for v .

Once we have chosen the type, the route selection is straightforward: first we route v ’s demand to the nearest terminal of higher type using cable type $\tau(v)$. Then we iterate: while v ’s demand is at a terminal w (which is not the root r), we route it to a terminal w' of type higher than $\tau(w)$ that is nearest to w , using cable type $\tau(w)$.

Finally, how do we select the path when routing v ’s demand from w to w' ? This is where our Multi-sink LAST (MLAST) construction comes in handy. We want that for each cable type i , the set of edges H_i on which we install cable type i has a small total cost while ensuring that each terminal of type i has a short path to its nearest terminal of higher type. We can achieve these properties by having H_i be an MLAST with the sources being the terminals of type exactly i , and the sinks being terminals of type higher than i .

Randomized Oblivious Algorithm. While designing oblivious algorithms seems like a big challenge because we have to be simultaneously competitive for all concave cost functions, Goel and Estrin [GE05] showed that functions of the form $g_i(x) = \min\{x, 2^i\}$ form a “basis” and hence we just have to be good against all these so-called “rent-or-buy” functions. This is precisely our goal.

Note that the optimum solution for the cost function g_0 is the optimal Steiner tree, and that for the cost function g_M , for $M \gg k$, is the shortest path tree rooted at the sink. Thus being competitive against g_0 and g_M already requires us to build a LAST. Thus it is not surprising that our online spanner algorithm is a crucial ingredient in our algorithm.

There are two key ideas. The first is that to approximate a given rent-or-buy function, it suffices to figure out which terminals should be connected to the root via “buy” edges—i.e., via edges of cost 2^i regardless of the load on them—these terminals we call the “*buy*” terminals. The rest of the terminals are simply connected to the buy terminals via shortest paths. One way to choose a good set of buy terminals is by random sampling: if we wanted to be competitive against function g_i , we could choose each terminal to be a “buy” terminal with probability 2^{-i} . (See, e.g., [AAB04].) Since in the oblivious case we don’t know which function g_i we are facing, we have to hedge our bets. Hence we choose each terminal v to have $\tau(v) = i$ with probability 2^{-i} . Thus terminal v is a good buy terminal for all g_i with $i \leq \tau(v)$.

Next, we need to ensure that the path $P(v)$ we choose for v simultaneously approximates the shortest path from v to the set of terminals with type at least i for all i . This is where our online spanner construction comes in handy. For each type i , we will build a spanner F_i on terminals of type at least i .

Deterministic Oblivious Algorithm. To obtain our deterministic oblivious algorithm, we make two further modifications. First, we remove the need for randomness by using the deterministic type selection rule of the non-oblivious algorithm. This modification already yields an $O(\log^3 k)$ -competitive algorithm. A technical difficulty that arises is that our type-classes are no longer nested. Thus it is not obvious how to route from a node of type i to one of a higher type, as these nodes may not belong to a common spanner. Adding nodes to multiple spanners can remedy this, but leads to a higher buy cost; being stingy in adding nodes to spanners can lead to higher rent cost. By carefully balancing these two effects and using a more sophisticated routing scheme, we are able to improve the competitive ratio to $O(\log^{2.5} k)$.

1.2 Other Related Work

Offline approximation algorithms for the (uniform) buy-at-bulk network design problem were initiated in [SCRS97]—here *uniform* denotes the fact that all edges have the same cost function $f(\cdot)$, up to scaling by the length of the edge. Early results for approximation algorithms for buy-at-bulk network design [MP98, AA97] already observed the relationship to spanners, and tree embeddings. Using the notion of probabilistic embedding of arbitrary metrics into tree metrics [AKPW95, Bar96, Bar98, Bar04, FRT04], logarithmic factor approximations are readily derived for the buy-at-bulk problem in the offline setting. A hardness result of $O(\log^{1/4-\epsilon} n)$ shows we cannot do much better in the worst case [And04].

For the offline single-sink case, new techniques were developed to get $O(1)$ -approximations [GMM09], as well as to prove $O(1)$ -integrality gaps for natural LP formulation [GKK⁺01, Tal02]; other algorithms have been given by [GKR03, GI06, JR09]. Apart from its inherent interest, the single-sink buy-at-bulk problem can also be used to solve other network design problems [GRS11]. The oblivious single-sink version was first studied in [GE05], and $O(1)$ -approximations for this very general version was derived in [GP12].

In the setting of online algorithms, the Steiner tree and generalized Steiner forest problems have tight $O(\log k)$ -competitive algorithms [BC97, IW91], where k is the number of terminals. These algorithms work in the model where new terminals only reveal their distances to previous terminals when they arrive, and the metric is not known *a priori*. It is well-known that the tree-embedding result of Bartal [Bar96] can be implemented online to give an $O(\log k \cdot \min(\log \Delta, \log k))$ -competitive algorithm for the online single-sink oblivious buy-at-bulk problem, where Δ is the ratio of maximum to minimum distances in the metric. For online rent-or-buy, Awerbuch, Azar, and Bartal [AAB04]

gave an $O(\log^2 k)$ -deterministic and an $O(\log k)$ -randomized algorithm; recently, [Umb15] gave an $O(\log k)$ -deterministic algorithm.

If the metric is known a priori, the results depend on n , the size of the metric, and not on k , the number of terminals.² E.g., tree-embedding results of [FRT04] give a randomized $O(\log n)$ -competitive algorithm, or a derandomization of oblivious network design from [GHR06] gives an $O(\log^2 n)$ -competitive algorithm.

A generalization which we do not study here is the *non-uniform* buy-at-bulk problem, where we can specify a different concave function on each edge. A poly-logarithmic approximation for this problem was recently given by Ene et al. [ECKP15]; see the references therein for the rich body of prior work.

2 Preliminaries

Formally, in the online buy-at-bulk problem, we have a complete graph $G = (V, E)$, and edge lengths $d(e)$ satisfying the triangle inequality. In other words, we can treat (V, d) as a metric. We have M cable types. The i -th cable type has *fixed cost* σ_i and *incremental cost* β_i , with $\sigma_i > \sigma_{i-1}$ and $\beta_i < \beta_{i-1}$. Routing x units of demand through cable type i on some edge e costs $(\sigma_i + \beta_i x)d(e)$.

In the single-sink version, we are given a root vertex $r \in V$. Initially, no cables are installed on any edges. When a terminal v arrives, we install some cables on some edges and choose a path $P(v)$ on which to route v 's unit demand. (This routing has to be unsplittable, i.e., along a single path.) We are allowed to install multiple cables on the same edge; if $e \in P(v)$ has multiple cables installed on it, v 's demand is routed on the one with highest type, i.e., the one with least incremental cost. The choice of $P(v)$ and cable installations are irrevocable. Given a routing solution, if $\text{load}_i(e)$ is the total amount of demand routed through cable type i on edge e , the total cost is $\sum_{e \in E} [\sum_{i: \text{load}_i(e) > 0} \sigma_i + \beta_i \text{load}_i(e)] d(e)$. We call $\sum_{e \in E} \sum_{i: \text{load}_i(e) > 0} \sigma_i d(e)$ the *fixed cost* of the solution and $\sum_{e \in E} \sum_{i: \text{load}_i(e) > 0} \beta_i \text{load}_i(e) d(e)$ the *incremental cost* of the solution. Let OPT denote the cost of the optimal solution. We assume the cable costs satisfy the conditions of the following Lemma 2.1.

Lemma 2.1. [GMM09] *We can prune the set of cables such that for all the retained cable types i , (a) $\sigma_i \geq 3\sigma_{i-1}$, and (b) $\beta_i \leq (1/3^2)\beta_{i-1}$, so that the cost of any solution using only the pruned cable types increases only by an $O(1)$ factor.*

2.1 HST embeddings

Let (X, d) be a metric over a set of points X with distances d at least 1.

Definition 2.2 (HST embeddings [Bar96]). *A hierarchically separated tree (HST) embedding T of metric (X, d) is a rooted tree with height $\lceil \log_2(\max_{u, v \in X} d(u, v)) \rceil$ and edge lengths that are powers of 2 satisfying:*

- a. *The leaves of T are exactly X .*
- b. *The length of the edge from any node to each of its children is the same.*
- c. *The edge lengths decrease by a factor of 2 as one moves along any root-to-leaf path.*

For $e \in T$, we use $T(e)$ to denote the length of e and say that e is a level- j edge if $T(e) = 2^j$. Furthermore, we write $T(u, v)$ to denote the distance between u and v in T . We also use $L(e)$

²Note that k is at most n , and it can be much smaller.

denote the leaves that are “downstream” of e , i.e. they are the leaves that are separated from the root of T when e is removed from T .

The crucial property of HST embeddings that we will exploit in our analyses is the following proposition, which follows directly from Properties (2) and (3) of Definition 2.2.

Proposition 2.3. *Let T be a HST embedding T of (X, d) . For any level- j edge $e \in T$, we have $d(u, v) < 2^j$ for any $u, v \in L(e)$.*

Corollary 2.4 ([AA97, FRT04]). *For any online buy-at-bulk instance with k terminals X and distances d , there exists a HST embedding T of (X, d) such that $\text{OPT}(T) \leq O(\log k) \text{OPT}$, where $\text{OPT}(T)$ is the cost of the optimal solution for online buy-at-bulk with terminals X in the tree T .*

2.2 Decomposition into rent-or-buy instances

Since buy-at-bulk functions can be complicated, it will be useful to deal with simpler *rent-or-buy* functions where to route a load of x on an edge, we can either *buy* the cable for unlimited use at its buy cost, or pay the rental cost times the amount x . Given a buy-at-bulk instance as above, for each i , define the rent-or-buy instance with the rent-or-buy function $f_i(x) = \min\{\sigma_i, \beta_{i-1}x\}$. Let OPT_i to be the cost of the optimum solution with respect to this function f_i . Note that under this function when the load aggregates up to $\frac{\sigma_i}{\beta_{i-1}}$, it becomes advantageous to switch from renting to buying.

The following lemma will prove very useful, since we can charge different parts of our cost to different $\text{OPT}_i(T)$ s for some HST T , and then sum them up to argue that the total cost is $O(1) \text{OPT}(T)$, and hence $O(\log k) \text{OPT}$ using Corollary 2.4.

Lemma 2.5 (OPT Decomposition on Trees). *For every tree T , we have $\sum_i \text{OPT}_i(T) \leq O(1) \text{OPT}(T)$.*

Proof. For an edge $e \in T$, let $\phi_i^*(e)$ and $\phi^*(e)$ be the costs incurred by $\text{OPT}_i(T)$, and $\text{OPT}(T)$ respectively, on e . We will show that for every edge $e \in T$, we have $\sum_i \phi_i^*(e) \leq O(1)\phi^*(e)$. Let $X(e)$ be the set of terminal-pairs whose tree paths (in the single-sink case, terminals whose paths to r) in T include e , and $T(e)$ denote the length of the tree edge e . So, $\phi_i^*(e) = T(e) \cdot \min\{\sigma_i, \beta_{i-1}|X(e)|\}$ and $\phi^*(e) = T(e) \cdot \min_i\{\sigma_i + \beta_i|X(e)|\}$. For any fixed j , we have

$$\begin{aligned} \sum_i \min\{\sigma_i, \beta_{i-1}|X(e)|\} &\leq \sum_{i \leq j} \sigma_i + \sum_{i > j} \beta_{i-1}|X(e)| \\ &\leq O(1)(\sigma_j + \beta_j|X(e)|), \end{aligned}$$

where the last inequality follows from the fact that the fixed costs σ_i increase geometrically with i and the incremental costs β_i decrease geometrically with i , as assumed in Lemma 2.1.

Thus, $\sum_i \phi_i^*(e) \leq O(1)\phi^*(e)$ for every edge $e \in T$ and so $\sum_i \text{OPT}_i(T) \leq O(1) \text{OPT}(T)$. \square

The next lemma proves that the rent-or-buy functions form a “basis” for buy-at-bulk functions. For $i \in \mathbb{Z}_{\geq 0}$, define the rent-or-buy function $g_i(x) = \min\{x, 2^i\}$. (See, e.g., [GP12, Section 2].)

Lemma 2.6 (Basis of Rent-or-Buy Functions). *Fix some routing of demands. If for every i , the cost of this routing under g_i is within a factor of ρ of the optimal routing for g_i , then for any monotone, concave function f with $f(0) = 0$, the cost of the routing under f is within $O(\rho)$ of the optimal cost of routing under f .*

3 LASTs

In the *light approximate shortest-path tree* (LAST) problem, we are given a graph $G = (V, E)$ with edge costs $d(u, v)$ and a root r , and we seek a low-cost spanning subtree T that preserves distances to the root. Let MST denote the cost of the minimum spanning tree. We say that T is an (α, β) -LAST if $c(T) \leq \alpha \text{MST}$ and $d_T(v, r) \leq \beta d(v, r)$ for all $v \in V$. Without loss of generality, G is complete and d is a metric on V . In the online setting, terminals arrive one by one; when a terminal arrives, its distances to previous terminals are revealed and the algorithm picks an edge that connects the new terminal to a previous one. We prove the following theorem.

Theorem 1.4 (Online LAST). *There exists a deterministic online algorithm for maintaining a tree with cost $O(\log k)$ times the MST, and a stretch of 7 for distances from terminals to the sink.*

High-level overview. Our algorithm maintains a greedy online Steiner tree T plus a set of edges A which we call the “direct” edges. The subgraph H produced by the graph consists of a subset of edges from $T \cup A$. When terminal v arrives, the algorithm considers the shortest (v, r) -path P_v contained in $T \cup A$; if $d(P_v) > 7d(v, r)$, then it adds (v, r) to H and A , otherwise it adds P_v to H . The formal algorithm is presented as Algorithm 1.

Algorithm 1 Online LAST

- 1: $T \leftarrow \emptyset; H \leftarrow \emptyset; A \leftarrow \emptyset$
 - 2: **while** terminal v arrives **do**
 - 3: Let u be closest previous terminal
 - 4: Add (u, v) to T
 - 5: Let P_v be the shortest (v, r) -path in $T \cup A$
 - 6: **if** $d(P_v) > 7d(v, r)$ **then**
 - 7: Add (v, r) to H and A
 - 8: **else**
 - 9: Add P_v to H
-

3.1 Analysis

Since $H \subseteq T \cup A$ and $c(T) \leq O(\log n) \text{MST}$, it suffices to bound $c(A)$.

Lemma 3.1. $c(A) \leq 2c(T)$.

Proof. Define $Z = \{v : (v, r) \in A\}$. Note that $c(A) = \sum_{e \in A} d(e) = \sum_{v \in Z} d(v, r)$. For each $v \in Z$, define L_v to be the segment of the (v, r) -path in T that is contained in a ball centered at v of radius $d(v, r)/2$. The lemma follows if we can show that these segments are disjoint. The following claim implies that they are disjoint.

Claim 3.2. *For every $v, v' \in Z$, we have $d_T(v, v') > d(v, r)/2 + d(v', r)/2$.*

Proof. Suppose, towards a contradiction, that for some $v, v' \in Z$, we have $d_T(v, v') \leq d(v, r)/2 + d(v', r)/2$ and v arrived after v' . Consider the point in time when the v arrived. Since $(v', r) \in A$, a possible (v, r) -path in $T \cup A$, is to take the (v, v') -path in T followed by (v', r) . So

$$d(P_v) \leq d_T(v, v') + d(v', r) \leq \frac{d(v, r)}{2} + \frac{d(v', r)}{2} + d(v', r).$$

Furthermore, $v \in Z$ implies $d(P_v) > 7d(v, r)$. Therefore, $d(v, r) < d(v', r)/4$.

On the other hand,

$$\begin{aligned}
d(v', r) &\leq d(v', v) + d(v, r) \\
&\leq d_T(v', v) + d(v, r) \\
&\leq \frac{d(v, r)}{2} + \frac{d(v', r)}{2} + d(v, r) \\
&< \frac{d(v', r)}{8} + \frac{d(v', r)}{2} + \frac{d(v', r)}{4} \\
&< d(v', r).
\end{aligned}$$

This gives the desired contradiction. □

Therefore, the segments L_v are disjoint, so $c(A) = 2 \sum_{v \in Z} |L_v| \leq 2c(T)$. □

4 Multi-Sink LASTs

Recall that LASTs were Light Approximate-Shortest-path Trees, i.e., trees where we maintain the shortest-path distance of sources to a sink, using a light tree. (This was traditionally done offline, but in §3 we show how to maintain this online—though the reader is not required to know the offline construction for this section.) In this section, we are interested in the *multi-sink LAST* (MLAST) problem. The input is a sequence of terminals in the underlying metric G , each of which is either a *source* or a *sink* (we assume the first terminal is always a sink), and our algorithm has to maintain a subgraph H such that (a) the distance of any source to its closest sink in H should be comparable to the distance of that source to its closest sink in G , and moreover (b) the cost $d(H)$ should not be “too large”.

Note that when a new sink s arrives, it may be close to many sources, and hence we may need to add “shortcut” edges to reduce their distance to s . Moreover, what does it mean for the cost of H to not be “too large”, since the distance from a source to its closest sink can fall dramatically over time. Our notion is the following. Note that when a *source* v arrives, it has to pay at least the distance to its closest terminal (source or sink) at that time, just to maintain connectivity. Loosely, we want our cost $d(H)$ to be not much more than the sum of these distances. (N.b.: this is the intuition, formally we will pay $O(2^{\text{class}(v)})$ which will be defined soon.)

4.1 The Algorithm

The first terminal to arrive is a sink we denote as s^* . We use R and S to denote the sets of sources and sinks that have arrived. For a terminal u , let $R(u)$ and $S(u)$ be the sources and sinks that arrived strictly before u . Our algorithm MLAST-ALG maintains a subgraph H that consists of two parts: a forest F which is a “backbone” connecting each source to some sink cheaply, and an edge set A which “augments” F to ensure that each source is not too far from the sinks in $H = A \cup F$.

The forest F is constructed using nets which we define now. Let $\Delta > 0$ be some distance scale. Then, a Δ -net Z is a subset of terminals (called *net points*) such that every terminal v has $d(v, Z) < \Delta$ and for every pair of net points $u, v \in Z$, we have $d(u, v) \geq \Delta$. Initially, both F and A are empty. The algorithm maintains a 2^j -net Z_j on the entire set of terminals, for each distance scale $j \in \mathbb{Z}$. When a terminal v arrives, for every distance scale j , it is added to Z_j if $d(u, Z_j) \geq 2^j$. The *class* of v is defined to be $\text{class}(v) := \max\{j : v \in Z_j\}$, i.e., $\text{class}(v)$ is the largest distance scale such that v belongs to the net of that scale. Hence the class of the first terminal s^* is ∞ .

Let u be the nearest terminal in $\bigcup_{j:j>\text{class}(v)} Z_j$, i.e., u is the nearest net-point lying in any net at a higher scale. If the current vertex v is a source, then the edge (u, v) is added to F ; otherwise, F is left untouched. Next, the algorithm goes through every source u (including v if it is also a source), checks if $d_{A \cup F}(u, S) \leq 3d(u, S)$, and otherwise adds the edge (u, u') to A where $u' \in S$ is the nearest sink to u . (See Algorithm 2.)

Algorithm 2 MLAST-ALG(R, S)

```

1:  $Z_j \leftarrow \emptyset$  for all  $j$ ,  $A \leftarrow \emptyset$ ,  $F \leftarrow \emptyset$ ,  $\Delta = 0$ ;
2: while terminal  $v$  arrives do
3:   If  $d(v, S(v)) > \Delta$ , set diameter  $\Delta \leftarrow d(v, S(v))$ 
4:   for all  $j \in \mathbb{Z}$  do
5:     if  $d(v, Z_j) \geq 2^j$  then
6:       Add  $v$  to  $Z_j$ 
7:    $\text{class}(v) \leftarrow \max\{j : v \in Z_j\}$ 
8:   if  $v \in R$  then
9:     Add edge  $(u, v)$  to  $F$  where  $u$  is a terminal of higher class nearest to  $v$ 
10:  for  $x \in R$  do
11:    if  $d_{A \cup F}(x, S) > 3d(x, S)$  then
12:      Add  $(x, x')$  to  $A$  where  $x'$  is the terminal in  $S$  nearest to  $x$ 

```

High-Level Idea of the Analysis. We first state the main properties of our MLAST algorithm that we will use in the remainder of the paper.

Lemma 4.1. *Given an online sequence of sources R and sinks S , the algorithm MLAST-ALG maintains a subgraph H and assigns classes to terminals such that*

- a. $d(H) \leq O(1) \sum_{v \in R} 2^{\text{class}(v)}$,
- b. $d_H(u, S) \leq 3d(u, S)$ for every $u \in R$,
- c. if $\text{class}(u) = \text{class}(v) = j$, then $d(u, v) \geq 2^j$,
- d. $d(u, R(u) \cup S(u)) \leq 2^{\text{class}(u)} \leq d(u, s^*)$ for every $u \in R$, where s^* is the first terminal that arrived, and $R(u)$ and $S(u)$ are the sources and sinks that arrived before u .

Property (a) is the formal bound on the cost of the MLAST. One should think of $2^{\text{class}(v)}$ as being the radius of some “dual” ball around source v , that we will later use to give lower bounds on our buy-at-bulk instances. Property (b) guarantees distance preservation. Properties (c) and (d) ensure that terminals of the same class are well-separated, and that $\text{class}(u)$ is not too large. Out of these, property (a) is the most non-trivial one to prove. The cost of F is easy to bound, since each source v adds in a single edge of length at most $O(2^{\text{class}(v)})$. For the cost of edges in A , the argument at a high level is that the sources adding two different edges of the same length must be far from each other in H (compared to the length of the edges added) else the later one could use a path to the earlier one, and then the added edge, to fulfill the distance guarantee.

4.2 Analysis of Mlast-Alg

We devote the remainder of this section to proving Lemma 4.1. First, note that by construction, MLAST-ALG ensures a short path in H from each source $x \in R$ to the sinks S : $d_H(x, S(x)) \leq 3d(x, S(x))$. Secondly, two terminals u and v having the same class j must both belong to the 2^j -net Z_j so their distance $d(u, v) \geq 2^j$. Thus, we have proved Properties (b) and (c).

Next, we prove Property (d).

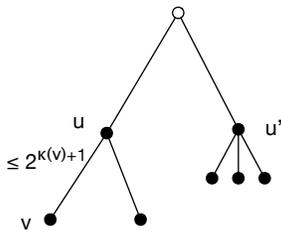


Figure 1: An example tree of F . The hollow vertex at the top is the leader and a sink, and the other are terminals. When v arrived, the algorithm added the edge (u, v) , making v a child of u , and the length of the edge is at most $2^{\text{class}(v)+1}$. Note that it is possible for two terminals u and u' to have the same class, and yet u 's children have much larger class than u' 's.

Proof of Lemma 4.1 (d). Let $\text{class}(u) = j$. Since every terminal that arrived before u is at distance $d(u, R(u) \cup S(u))$ from u , we have that u belongs to the net Z_{j^-} where $j^- = \lfloor \log d(u, R(u) \cup S(u)) \rfloor$. Thus, $2^j \geq d(u, R(u) \cup S(u))/2$. On the other hand, s^* belongs to every net since it was the first terminal to arrive. Thus, u cannot belong to any net $Z_{j'}$ where $j' > \lfloor \log d(u, s^*) \rfloor$. \square

To prove Property (a), we start with some useful properties of the forest F that will allow us to relate $d(H)$ to the classes of sources. Root each tree T in F using the terminal of highest class in that tree, which we call the *leader* of T .

Claim 4.2. *Each tree T of the forest F satisfies the following properties:*

1. *Its leader is a sink and all other terminals are sources.*
2. *On every leaf-to-leader path of T , the classes of terminals on that path strictly increase.*
3. *If $v \in T$ is a source, then when it arrived, the edge (u, v) was added to F , where u is a terminal in T with $\text{class}(u) > \text{class}(v)$ and $d(u, v) \leq 2^{\text{class}(v)+1}$.*
4. *If u is the parent of v and v is an ancestor of t in T , then the length of the (t, u) -path P in T is $d_T(t, u) \leq 2^{\text{class}(v)+2} \leq 2^{\text{class}(u)+1}$.*

See Figure 1 for an illustration of properties (1), (2), and (3); and Figure 2 for property (4).

Proof. Properties (1) and (2) follow from the fact that each source has a single edge in T to a terminal of higher class, and sinks are not connected to other terminals when they arrive.

We now prove property (3). Let (u, v) be the edge added when v arrived. The algorithm chose u to be the nearest terminal of some higher class. By definition, all terminals in $\bigcup_{j: j > \text{class}(v)} Z_j$ have class at least $\text{class}(v) + 1$. Moreover, the closest net-point to v in $Z_{\text{class}(v)+1}$ is within distance $2^{\text{class}(v)+1}$, so u can be no further. Hence, we have $d(v, u) \leq d(v, Z_{\text{class}(v)+1}) \leq 2^{\text{class}(v)+1}$.

Finally, we prove property (4). Suppose u is the parent of v and v is an ancestor of t in the tree T , and let P be the path in T from t to u . Property (2) implies that the classes of terminals in P are strictly increasing as we move from t to u . Property (3) now implies that the length of the t - u path P is at most $\sum_{x \in P: x \neq u} 2^{\text{class}(x)+1} < 2^{\text{class}(v)+2} \leq 2^{\text{class}(u)+1}$. This concludes the proof of the claim. \square

Since each edge in F is added by a distinct source, the total cost of F (using Claim 4.2(3)) is:

$$d(F) \leq 2 \sum_{v \in R} 2^{\text{class}(v)}. \quad (1)$$

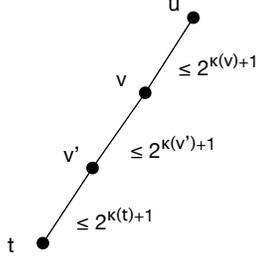


Figure 2: An illustration of the proof of property (4). Since $\text{class}(t) < \text{class}(v') < \text{class}(v) < \text{class}(u)$ by properties (2) and (3), the length of the path from t to u is at most $2^{\text{class}(v)+2} \leq 2^{\text{class}(u)+1}$.

It remains to bound $d(A)$ in terms of the classes of sources. The following claim allows us to find sources to charge $d(A)$ to. Define $A_j = \{(u, v) \in A : d(u, v) \in [2^j, 2^{j+1})\}$. Observe that every edge $(u, v) \in A$ has one endpoint being a source terminal and another being a sink terminal.

Claim 4.3. *Suppose $(u, v) \in A_j$ with u being the source endpoint and T is the tree in F containing u . Then there exists an ancestor $\varphi(u) = z$ of u in T such that z is a source, has class $\text{class}(z) \geq j - 2$, and the length of the (u, z) -path in T is $d_T(u, z) \leq 2^{j-1}$.*

Proof. Let a be the sink that is the leader of T . By definition of the algorithm, v is the nearest sink to u at the moment that (u, v) was added to A , so $d(u, v) \leq d(u, a) \leq d_T(u, a) \leq 2^{\text{class}(a)+1}$, where the last inequality follows from Claim 4.2(4). Since (u, v) belongs to A_j , we also have $d(u, v) \geq 2^j$. Putting the previous two inequalities together, $2^j \leq 2^{\text{class}(a)+1}$ and so the leader a has $\text{class}(a) \geq j - 1$.

If u has $\text{class}(u) \geq j - 2$ we are done, so assume $\text{class}(u) \leq j - 3$. Let t be the highest ancestor of u such that $\text{class}(t) \leq j - 3$ (clearly u is a candidate). Observe that $t \neq a$. Let z be t 's parent; we claim z satisfies the desired conditions. It has $\text{class} \geq j - 2$ (by the maximality of t), and $d(u, z) \leq 2^{\text{class}(t)+2} \leq 2^{j-1}$ by Claim 4.2(4). But we'd already observed that $d(u, a) \geq 2^j$, so $z \neq a$; hence z is a source by Claim 4.2(1). \square

Proof of Lemma 4.1 (a). To upper bound $d(A)$, we charge the length of each edge $(u, v) \in A$ to a source vertex, and then show that the charge received by each source z is at most $O(1)2^{\text{class}(z)}$. For each edge $(u, v) \in A$, the charging scheme is defined as follows: if $(u, v) \in A_j$ (i.e. its length $d(u, v) \in [2^j, 2^{j+1})$) and u is the source endpoint and contained in the tree T of F , charge $2^{j+1} \geq d(u, v)$ to the ancestor $\varphi(u)$ given by Claim 4.3. Hence $d(A)$ is at most the total charge received by all the sources.

Next, we show that the charge received by every source z is at most $O(1)2^{\text{class}(z)}$. Terminal z is only charged by edges belonging to A_j for $j \leq \text{class}(z) + 2$; each such edge charges z charges 2^{j+1} . We claim that for each j , at most one edge in A_j charges z : this means the total charge received by z is at most $\sum_{j \leq \text{class}(z)+2} 2^{j+1} \leq O(1)2^{\text{class}(z)}$, so $d(A) \leq \sum_{z \in R} O(1)2^{\text{class}(z)}$, which proves the lemma.

To prove the claim, suppose that there exist two edges $(u, v), (u', v') \in A_j$ that charge z . Let u and u' be sources, and v and v' be sinks. Since both edges charge z , $\varphi(u) = \varphi(u') = z$. Since both edges belong to A_j , their lengths must satisfy

$$2^j \leq d(u, v) \leq 2d(u', v'). \quad (2)$$

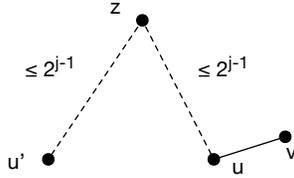


Figure 3: An illustration of the situation when two edges (u, v) and (u', v') in A_j charge the same terminal z , and u and u' are the source endpoints. The dashed edges represent paths from u' to z and u to z in T . There is a path in H from u' to S following the paths in T represented by the dashed edges and the edge (u, v) .

By definition of the charging scheme,

$$d_T(u', z) \leq 2^{j-1} \quad \text{and} \quad d_T(u, z) \leq 2^{j-1}. \quad (3)$$

Suppose (u', v') was added to A after (u, v) . At the time when we consider adding (u', v') , the sink $v \in S(u')$ and so $d_H(u', S) \leq d_H(u', v)$. Moreover, there is a path in H from u' to v : the path in T from u' to u , followed by the edge (u, v) , which implies

$$d_H(u', S) \leq d_T(u', z) + d_T(z, u) + d(u, v).$$

(See Figure 3 for an illustration.) Plugging in the inequalities (2) and (3), we get that $d_H(u', S) \leq 2^j + 2d(u', v') \leq 3d(u', v')$, where the last inequality follows from the fact that $2^j \leq d(u', v')$. However, this contradicts the condition for adding (u', v') to A . Hence only one edge in A_j can charge any z ; this proves the claim and hence the lemma. \square

5 Non-Oblivious Buy at Bulk

In this section we prove our main result for online buy-at-bulk for the case when the function f is known.

Theorem 1.1 (Deterministic Buy-at-Bulk). *There exists a deterministic $O(\log k)$ -competitive algorithm for online single-sink buy-at-bulk, where k is the number of terminals.*

For a high-level intuition behind the algorithm, see the discussion in §1.1. Here we explain precisely how to assign types, update layers and route demands.

Type assignment. Consider a terminal v . At the time of its arrival, for each type i , let $d_i(v)$ be the distance to the nearest type- i terminal, and define the ball $B_i(v) := \{u : d(u, v) \leq d_i(v)/2^3\}$. Let $n_i(v) := |B_i(v)|$ be the number of terminals in this ball. Terminal v is assigned type

$$\tau(v) := \max \left\{ i : n_i(v) \geq \frac{\sigma_i}{\beta_{i-1}} \right\}.$$

To make sense of this, observe that the threshold for being of a certain type is the same as the threshold for the load at which it is advantageous to buy rather than rent for that type. This means vertices assigned a certain type can collect enough “witnesses” whose paths to the root in the optimal solution can be used to pay for edges constructed by v in the online solution.

Layers. Let R_i be the current set of terminals of type exactly i , and $X_i := \cup_{i' \geq i} R_{i'}$ be the current set of terminals of type at least i . Let $X_i(v)$ (resp., $R_i(v)$) be the set of terminals of type at least (resp., exactly) i that arrived strictly before v . Layer H_i is maintained by running the online algorithm MLAST-ALG with R_i as sources and X_{i+1} as sinks. This ensures that every $u \in R_i$ has $d_{H_i}(u, X_{i+1}(v)) \leq 3d(u, X_{i+1}(v))$.

Routing. We route v 's demand to the root on path $P(v)$ that is constructed as follows. The path $P(v)$ is constructed iteratively and consists of M different segments $P_i(v)$, one per type. Initially, v 's demand is located at v and the segments are all empty. While the current terminal w containing v 's demand is not the root, choose w' to be the terminal in $H_{\tau(w)}$ with $\tau(w') > \tau(w)$ nearest to w (nearest in terms of distances in $H_{\tau(w)}$) and route v 's demand to w' along the shortest path between w and w' in $H_{\tau(w)}$.

This completes the description of the algorithm. See Algorithm 3 for a summary.

Algorithm 3 Online Algorithm for Single-Sink Buy-at-Bulk

```

1:  $\tau(r) \leftarrow \infty$ 
2:  $H_i \leftarrow \emptyset$  for all  $i$ 
3: while terminal  $v$  arrives do
4:   // Determine type of  $v$ 
5:   Let  $d_i(v) = d(v, X_i)$ ,  $B_i(v) = \{u : d(u, v) \leq d_i(v)/2^3\}$ , and  $n_i(v) = |B_i(v)|$  for  $1 \leq i \leq M$ 
6:    $\tau(v) \leftarrow \max\{i : n_i(v) \geq \frac{\sigma_i}{\beta_{i-1}}\}$ 
7:   // Update layers  $H_i$ 
8:   for  $i \leq \tau(v)$  do
9:     Update  $H_i$  using MLAST-ALG with  $R_i(v)$  as sources and  $X_{i+1}(v)$  as sinks
10:    Install cable of type  $i$  on new edges of  $H_i$ 
11:    // Find routing path  $P(v)$  to route  $v$ 's demand to the root
12:     $P_i(v) \leftarrow \emptyset$  for all  $i$ ;  $P(v) \leftarrow \emptyset$ 
13:    while  $v$ 's demand is not yet at  $r$  do
14:      Let  $w$  be the terminal containing  $v$ 's demand and  $i = \tau(w)$ 
15:      Let  $P_i(v) \subseteq H_i$  be the shortest path from  $w$  to  $X_{i+1}(v)$  in  $H_i$ 
16:      Route  $v$ 's demand from  $w$  to  $X_{i+1}(v)$  along the path  $P_i(v)$ 

```

5.1 Analysis

Let us now consider the cost of the algorithm's solution due to each cable type. The fixed cost due to cables of type i is $\sigma_i d(H_i)$ since they are installed on layer H_i . For each terminal v of type $\tau(v) \leq i$, $P_i(v)$ is the segment of its routing path consisting of type i cables. Hence the total cost of the algorithm's solution is

$$\text{ALG} := \sum_i \left(\sigma_i d(H_i) + \sum_{v: \tau(v) \leq i} \beta_i d(P_i(v)) \right). \quad (4)$$

Proof outline. We want to show that $\text{ALG} \leq \text{OPT}(T)$ for any HST embedding T , since Corollary 2.4 would then imply that $\text{ALG} \leq O(\log k) \text{OPT}$. The key is to decompose OPT on trees as follows: in §2.2 we defined the rent-or-buy functions $f_i(x) := \min\{\sigma_i, \beta_{i-1}\}$, and defined OPT_i to be the cost of the optimal solution with respect to f_i . Our proof (specifically Lemmas 5.3 and 5.2)

will use the structure of the HST to give a charging argument showing that

$$\sigma_i d(H_i) + \sum_{v:\tau(v)<i} \beta_{i-1} d_i(v) \leq O(1) \text{OPT}_i(T). \quad (5)$$

Roughly speaking, we argue that $d(H_i)$ can be charged to the optimal cost of the witnesses of the terminals of type exactly i , and moreover, the type selection rule forces terminals of type less than i to be spread out.

Thus, if we could prove that $d(P_i(v)) \leq O(1)d_{i+1}(v)$, then we would be done, because we could decompose the expression for ALG in (4) as follows:

$$\text{ALG} \leq O(1) \cdot \sum_i \left(\sigma_i d(H_i) + \sum_{v:\tau(v)<i} \beta_{i-1} d_i(v) \right). \quad (6)$$

Inequality (5) and Lemma 2.5 would imply that $\text{ALG} \leq O(1) \text{OPT}(T)$ for any HST embedding T , as desired.

However, $d(P_i(v))$ can be much larger than $d_{i+1}(v)$. This is because of the ‘‘selfishness’’ of the routing scheme: when a terminal w receives v ’s demand, it simply routes it to the terminal w' of higher type nearest to w , without any regard as to how far w' is from v . Fortunately, the fact that the incremental costs β_i are geometrically decreasing allows us to show that the total incremental cost incurred over all segments $P_i(v)$ are bounded. In particular, we show that $\sum_{v:\tau(v)\leq i} \beta_i d(P_i(v)) \leq O(1) \sum_i \sum_{v:\tau(v)\leq i} \beta_i d_{i+1}(v)$. The rest of the proof will follow this general outline.

Bounding incremental cost. We begin by proving the above bound on the incremental cost.

Lemma 5.1. *The incremental cost $\sum_i \sum_{v:\tau(v)\leq i} \beta_i d(P_i(v)) \leq O(1) \sum_i \sum_{v:\tau(v)\leq i} \beta_i d_{i+1}(v)$.*

Proof. For terminal v , suppose its type- i segment $P_i(v)$ is non-empty: let w_i be the starting terminal of the segment. Since $P_i(v)$ starts with a terminal of type i and ends with a terminal of type at least $i+1$, so $\tau(w_i) = i$. Now since $P_i(v)$ is a shortest path in H_i from w_i to $X_{i+1}(v)$, Lemma 4.1(b) gives us $d(P_i(v)) \leq 3 \cdot d(w_i, X_{i+1}(v))$. Since $d_{i+1}(v)$ is the distance from v to $X_{i+1}(v)$, $d(w_i, X_{i+1}(v)) \leq d(w_i, v) + d_{i+1}(v)$. Moreover, the segments $P_{i'}(v)$ for $i' < i$ form a path from v to w_i , so $d(w_i, v) \leq \sum_{i'<i} d(P_{i'}(v))$. Combining these inequalities, we get $d(P_i(v)) \leq 3(d_{i+1}(v) + \sum_{i'<i} d(P_{i'}(v)))$. Unrolling the recursion, this gives $d(P_i(v)) \leq \sum_{i'\leq i} 3^{i-i'+1} d_{i'+1}(v)$. Hence, we get the following bound on the incremental cost of v :

$$\sum_{i\geq\tau(v)} \beta_i \cdot d(P_i(v)) \leq \sum_{i\geq\tau(v)} \beta_i \cdot \sum_{i'\leq i} 3^{i-i'+1} d_{i'+1}(v) = \sum_{i\geq\tau(v)} d_{i+1}(v) \cdot \sum_{i'\geq i} 3^{i'-i+1} \beta_{i'},$$

where the equality follows from rearranging the sums. But the β_i s decrease geometrically by a factor of 3^2 , so the last sum is dominated by the first term, and is $O(\beta_i)$. Hence the proof. \square

Lemma 5.1 implies Inequality (6). Define $\text{ALG}_i = \sigma_i d(H_i) + \sum_{v:\tau(v)<i} \beta_{i-1} d_i(v)$. The rest of this section will show that on any HST embedding T , $\text{ALG}_i \leq O(1) \text{OPT}_i(T)$ for every i . Lemma 2.5 then implies that $\text{ALG} \leq O(1) \text{OPT}(T)$ for any HST embedding T and so $\text{ALG} \leq O(\log k) \text{OPT}$ by Corollary 2.4.

Charging to HST embeddings. For the following, fix an HST embedding T . Recall that for edge $e \in T$, $L(e)$ denotes the leaves below e , $T(e)$ the length of e , and $T(u, v)$ the distance between u and v in T . Also, observe that an edge e such that $r \notin L(e)$, the terminals in $L(e)$ either have to buy the edge at cost $\sigma_i T(e)$, or rent it at cost $\beta_{i-1} |L(e)| T(e)$. We record this lower bound for later.

$$\text{OPT}_i(T) \geq \sum_{e \in T: r \notin L(e)} T(e) \min\{\sigma_i, \beta_{i-1} |L(e)|\}. \quad (7)$$

To upper-bound ALG_i , we bound both $\sigma_i d(H_i)$ and $\sum_{v: \tau(v) < i} \beta_{i-1} d_i(v)$ separately by $\text{OPT}_i(T)$. In each case, we proceed by developing an appropriate charging scheme that charges to the edges of T and then arguing that the total charge received by each edge $e \in T$ is at most a constant times its contribution to the lower bound of $\text{OPT}_i(T)$ in Inequality 7.

First, we bound $\sum_{v: \tau(v) < i} \beta_{i-1} d_i(v)$. The charging scheme is as follows: for each terminal v , charge β_{i-1} to an edge in T whose length is proportional to $d_i(v)$ and which contains v as a leaf. Then we argue that no edge is overcharged; in particular, for every edge $e \in T$, the total number of terminals that can charge e is at most σ_i / β_{i-1} . Finally, we use the fact that terminals charging to e are close together (by the bounded diameter property of HSTs), so if there were more than σ_i / β_{i-1} terminals charging e , then the one that arrived last would have been assigned a type of at least i .

Lemma 5.2. $\sum_{v: \tau(v) < i} \beta_{i-1} d_i(v) \leq O(1) \text{OPT}_i(T)$.

Proof. Consider the following charging scheme. For each terminal v with type $\tau(v) < i$, if $d_i(v) \in [2^j, 2^{j+1})$, charge $2^{j+1} \beta_{i-1}$ to the length 2^{j-4} edge $e \in T$ whose leaves $L(e)$ contain v but not r . Such an edge must exist since otherwise $d_i(v) \leq d(v, r) < 2^j$. The total charge received by the edges of T is at least $\sum_{v: \tau(v) < i} \beta_{i-1} d_i(v)$, and only edges e with $r \notin L(e)$ were charged.

Consider an edge $e \in T$ of length 2^{j-4} . Let $C(e) \subseteq L(e)$ be the set of terminals charging e . We claim that $|C(e)| \leq \sigma_i / \beta_{i-1}$. Note that the total charge received by e is

$$2^{j+1} \beta_{i-1} |C(e)| = 2^5 T(e) \beta_{i-1} |C(e)| \leq 2^5 T(e) \min\{\sigma_i, \beta_{i-1} |L(e)|\}.$$

By (7), this proves the lemma.

Now to prove the claim. Suppose for a contradiction, that $|C(e)| > \sigma_i / \beta_{i-1}$. Let v be the last-arriving terminal of $C(e)$. Since e is a length 2^{j-4} edge and v charged e , we have $d_i(v) \geq 2^j$ and $\tau(v) < i$. Moreover, since $C(e) \subseteq L(e)$, we have that $\text{diam}(C(e)) < 2^{j-3} \leq d_i(v) / 2^3$. Thus, every terminal $u \in C(e)$ is within distance $d_i(v) / 2^3$ from v so $C(e) \subseteq B_i(v)$. This implies that $|B_i(v)| \geq |C(e)| \geq \sigma_i / \beta_{i-1}$ so v should have been assigned a type that is at least i , contradicting the fact that $\tau(v) < i$. Therefore, $|C(e)| \leq \sigma_i / \beta_{i-1}$, as desired. \square

Bounding the Fixed Cost of H_i . There are three steps to the proof. The first step is to use Lemma 4.1 to charge the fixed cost to the terminals of type i . The second is to argue that since each terminal v of type i has at least σ_i / β_{i-1} witnesses nearby, v can use the cost incurred by $\text{OPT}_i(T)$ in routing its witnesses to pay off the charge it accumulated. Finally, we use the fact that terminals of type i that accumulate similar charges must be far apart to argue that no witness is overcharged.

Lemma 5.3. $\sigma_i d(H_i) \leq O(1) \text{OPT}_i(T)$.

Proof. The layer H_i is an MLAST whose set of sources is R_i (the terminals of type exactly i) and sinks $S_i = X_{i+1}$, the terminals of higher type. Let $\text{class}_i(v)$ be the class assigned by this

MLAST algorithm to terminal v . By Lemma 4.1, we have $\sigma_i d(H_i) \leq O(1) \sum_{v \in R_i} \sigma_i 2^{\text{class}_i(v)}$. Define $R_i(j) = \{v \in R_i : \text{class}_i(v) = j\}$. Let E_j be the set of length 2^{j-4} edges of T . To prove the lemma, we will show that for each j ,

$$\sigma_i |R_i(j)| \leq \sum_{e \in E_j : r \notin L(e)} \min\{\sigma_i, \beta_{i-1} |L(e)|\}. \quad (8)$$

This would then imply that

$$\sigma_i d(H_i) \leq O(1) \sum_{v \in R_i} \sigma_i 2^{\text{class}_i(v)} \leq O(1) \sum_j 2^j \sigma_i |R_i(j)| \leq O(1) \sum_j 2^{j-4} \sum_{e \in E_j : r \notin L(e)} \min\{\sigma_i, \beta_{i-1} |L(e)|\}.$$

But $T(e) = 2^{j-4}$ for edges $e \in E_j$, so (7) bounds the cost by $O(1) \text{OPT}_i(T)$ to complete the proof.

Now to prove (8). We will show that for each $v \in R_i(j)$ (i.e., having type i , and class j in the MLAST H_i corresponding to type- i terminals), every terminal u in its ball $B_i(v)$ has to be routed on some level- j edge $e \in E_j$ with $r \notin L(e)$, and that the level- j edges used by $B_i(v)$ is disjoint from the edges used by $B_i(v')$ for any other $v' \in R_i(j)$. More formally, for each terminal v , define $e_j(v)$ to be the unique edge of E_j such that $v \in L(e_j(v))$ and define $E_j(v) := \{e_j(u) : u \in B_i(v)\}$. We need the following claims.

Claim 5.4. *For every $v \in R_i(j)$ and $u \in B_i(v)$, we have $d(u, v) \leq 2^{j-3}$.*

Proof. By definition of $B_i(v)$, we have the following bound: $d(u, v) \leq d_i(v)/2^3 = d(v, X_i(v))/2^3$. Now observe that $X_i(v) = R_i(v) \cup S_i(v)$, where $R_i(v)$ and $S_i(v)$ are the sources and sinks that arrived before v in the MLAST for type i . By Lemma 4.1(d), $d(v, X_i(v)) \leq 2^{\text{class}_i(v)} = 2^j$. Combining this with the above bound on $d(u, v)$ gives us $d(u, v) \leq 2^{j-3}$, as desired. \square

Claim 5.5. *For every $v \in R_i(j)$, we have $\bigcup_{e \in E_j(v) : r \notin L(e)} L(e) \supseteq B_i(v)$. Moreover, $E_j(v) \cap E_j(v') = \emptyset$ for distinct $v, v' \in R_i(j)$.*

Proof. To prove the first part of the claim, observe that by definition of $E_j(v)$, we have $\bigcup_{e \in E_j(v)} L(e) \supseteq B_i(v)$. Thus, it suffices to show that for every $v \in R_i(j)$ and each terminal $u \in B_i(v)$, the edge $e_j(v)$ does not contain r as a leaf. Suppose, towards a contradiction, that there exists $u \in B_i(v)$ such that $r \in L(e_j(u))$. Since $e_j(u)$ has length 2^{j-4} , we have $\text{diam}(L(e_j(u))) \leq 2^{j-3}$; moreover, $r, u \in L(e_j(u))$ and so $d(u, r) \leq 2^{j-3}$. Now, Claim 5.4 implies that $d(u, v) \leq 2^{j-3}$. Therefore, $d(v, r) \leq d(u, v) + d(u, r) < 2^j$. However, this contradicts Lemma 4.1(d), which implies that $d(v, r) \geq 2^j$. Thus, $r \notin L(e_j(u))$.

To prove the second part of the claim, suppose towards a contradiction, that there exist $v, v' \in R_i(j)$ and $u \in B_i(v)$ and $u' \in B_i(v')$ such that $e_j(u) = e_j(u')$. By triangle inequality, $d(v, v') \leq d(u, v) + d(u, u') + d(u', v')$. Claim 5.4 implies that $d(u, v), d(u', v') \leq 2^{j-3}$. Since u, u' are leaves of the same edge of length 2^{j-4} , we get $d(u, u') \leq 2^{j-3}$. Therefore, $d(v, v') < 2^j$. On the other hand, Lemma 4.1(c) says that $d(v, v') \geq 2^j$. This gives us our desired contradiction. \square

With these claims in hand, we have

$$\begin{aligned} \sum_{e \in E_j : r \notin L(e)} \min\{\sigma_i, \beta_{i-1} |L(e)|\} &\geq \sum_{v \in R_i(j)} \sum_{e \in E_j(v) : r \notin L(e)} \min\{\sigma_i, \beta_{i-1} |L(e)|\} \\ &\geq \sum_{v \in R_i(j)} \min\{\sigma_i, \beta_{i-1} \sum_{e \in E_j(v) : r \notin L(e)} |L(e)|\} \end{aligned}$$

$$\begin{aligned}
&\geq \sum_{v \in R_i(j)} \min\{\sigma_i, \beta_{i-1} |B_i(v)|\} \\
&\geq \sum_{v \in R_i(j)} \sigma_i = \sigma_i |R_i(j)|,
\end{aligned}$$

where the first inequality follows from the second part of Claim 5.5, the third inequality from the first part, and the last inequality from the fact that $|B_i(v)| \geq \sigma_i/\beta_{i-1}$. \square

Having proved these two lemmas, we have that for every HST embedding T and every i , we have $\text{ALG}_i \leq O(1) \text{OPT}_i(T)$. Summing over all i and using Lemma 2.5, we have $\text{ALG} \leq O(1) \text{OPT}(T)$ for any HST embedding T , and so $\text{ALG} \leq O(\log k) \text{OPT}$. This proves Theorem 1.1.

6 Multi-Commodity Spanners

In the *online α -MCS* problem, we are given a graph $G = (V, E)$ with edge costs $d(u, v)$ satisfying the triangle inequality. Terminal pairs (s, t) arrive one-by-one. The goal is to maintain a minimum-cost subgraph H such that $d_H(s, t) \leq \alpha d(s, t)$ for all terminal pairs (s, t) . We say that H is β -competitive if its cost $d(H) = \sum_{(u,v) \in H} d(u, v)$ is at most β of the cheapest Steiner forest connecting the terminal pairs. We recall our main theorem about multi-commodity spanners.

Theorem 1.5 (Online Spanner). *There exists a deterministic online algorithm that given k pairs of terminals, maintains a forest with cost $O(\log k)$ times the Steiner forest on the pairs, and a stretch of $O(\log k)$ for distances between all given pairs of terminals. Moreover, the total number of edges in the forest is $O(k)$, i.e. linear in the number of terminals.*

Note that this result is tight in several ways. Firstly, the solution is also a feasible Steiner forest, and we know that any online Steiner forest algorithm has cost $\Omega(\log k)$ times the cheapest Steiner forest, hence we cannot improve the cost even at the expense of greater stretch. Moreover, there exists a k -vertex unweighted graph with girth $\Omega(\frac{\log k}{\log \log k})$ and size $\Omega(k \log k)$ [LPS88]. Thus, we cannot improve the stretch by much either. Finally, $\Omega(k)$ edges are needed to connect k terminals.

High-level overview. Our algorithm (Algorithm 4) is based on the Berman-Coulston online Steiner forest algorithm. There are two key ingredients to the algorithm. First, it classify terminals according to the distance to their mates: $\text{class}(s_i) = \text{class}(t_i) = \lfloor \log_2 d(s_i, t_i) \rfloor$. Let Y_j be the set of terminals with class at least j . Second, the algorithm maintains a clustering of Y_j as follows. It designates a subset $Z_j \subseteq Y_j$ that are $2^j/16$ -separated from each other as *centers*, and assigns each $v \in Y_j$ to a *cluster* $C(z)$ where z is the nearest center of Z_j to u .

The rule for adding edges is as follows: if for any $u, v \in Y_j$ with $d(u, v) \in [2^j, 2^{j+1})$ has $d_H(u, v) > 4 \log k \cdot d(u, v)$ (i.e., the pair has large stretch in the current graph H), the algorithm adds the edge (u, v) , and connects u and v to their respective centers; we call the former an *augmentation edge* and the latter *bridge edges* and account for them separately.

Let A_j be the set of augmentation edges (u, v) with $d(u, v) \in [2^j, 2^{j+1})$. First, we show that it suffices to bound the cost and number of augmentation edges (Lemma 6.1). Since the diameter of each cluster is at most $2^j/8$, each augmentation edge must connect two different clusters. Thus we can think of a meta-graph where the nodes correspond to clusters and meta-edges correspond to A_j . The bridge edges allow us to argue that if there is a short path in the meta-graph between nodes corresponding to clusters $C(z)$ and $C(z')$, then there is a short path in H between every $u \in C(z)$ and $v \in C(z')$ (Lemma 6.2). We then use this to show that there cannot be a cycle of

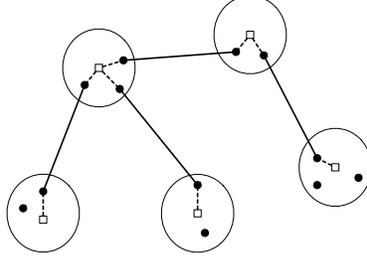


Figure 4: An example of the meta-graph. The circles represent clusters that are nodes of the meta-graph. The dashed edges and solid edges represent bridge and augmentation edges, respectively. Note that not all terminals in a cluster has a bridge edge to the center.

length less than $\log k$ in the meta-graph and so the total number of meta-edges is at most $O(1)$ times the number of nodes, i.e. $|A_j| \leq O(|Z_j|)$. Finally, we use the separation of Z_j to show that $2^j|Z_j| = O(\text{OPT})$ via a ball-packing argument (Lemma 6.5), and complete the analysis of the algorithm’s cost. To bound the total number of edges $\sum_j |A_j|$, we use a charging scheme that charges $|A_j|$ to a carefully-chosen set of terminals and show that across all distance scales j , the total charge received by each terminal is $O(1)$.

Algorithm 4 Online $(4 \log k)$ -MCS

- 1: $A_j \leftarrow \emptyset, B_j \leftarrow \emptyset, Z_j \leftarrow \emptyset$ for all j ; $H \leftarrow \emptyset$;
 - 2: **while** (s_i, t_i) arrives **do**
 - 3: Define $\text{class}(s_i) = \text{class}(t_i) = \lfloor \log d(s_i, t_i) \rfloor$
 - 4: **for** $j = 0$ up to $j = \lfloor \log d(s_i, t_i) \rfloor$ **do**
 - 5: **for** $u \in \{s_i, t_i\}$ **do**
 - 6: **if** $d(u, z) \geq 2^j/16$ for every center $z \in Z_j$ **then**
 - 7: Add u to Z_j
 - 8: Define a new cluster $C(u) = \{u\}$
 - 9: **else**
 - 10: Choose $z \in Z_j$ closest to u and add u to $C(z)$
 - 11: **for** u, v such that $\min\{\text{class}(u), \text{class}(v)\} \geq j$ and $d(u, v) \in [2^j, 2^{j+1})$ **do**
 - 12: **if** $d_H(u, v) > 4 \log k \cdot d(u, v)$ **then**
 - 13: Add augmentation edge (u, v) to A_j and H
 - 14: Let $z, z' \in Z_j$ such that $u \in C(z), v \in C(z')$
 - 15: Add bridge edges (u, z) and (v, z') to B_j and H
-

6.1 Analysis

To analyze Algorithm 4, we use a “shadow” algorithm that builds a family of $O(\log k)$ meta-graphs F_j , one for each length scale 2^j . Whenever the algorithm adds a new cluster—i.e., it adds a new center z to Z_j —the shadow algorithm creates a new node in F_j corresponding to $C(z)$. We use $C(z)$ to refer to both the cluster as well as the corresponding node in F_j . Whenever Algorithm 4 adds an edge (u, v) to A_j , the shadow algorithm adds the meta-edge $(C(z), C(z'))$ where $u \in C(z)$ and $v \in C(z')$. This is well-defined since u, v have class at least j and every terminal of class at least j belongs to some cluster $C(z)$ for $z \in Z_j$. Furthermore, since the diameter of a cluster is less than $2^j/8$, and each edge of A_j has length at least 2^j , there are no self-loops. See Figure 4 for an illustration.

Note that it suffices to consider the augmentation edges to bound the cost of H as well as the number of edges in H . Indeed, every bridge edge in B_j has a corresponding augmentation edge in A_j , and each augmentation edge corresponds to at most two bridge edges (of no greater cost). This proves the following lemma.

Lemma 6.1. $d(H) \leq \sum_j 2^{j+1} \cdot 3|A_j|$ and $|E(H)| \leq \sum_j 3|A_j|$.

The following lemma relates distances in the meta-graph F_j to distances in H and G , and is key to bounding $d(H)$ and $|E(H)|$.

Lemma 6.2. *If there is a path of ℓ meta-edges in F_j between nodes $C(z)$ and $C(z')$, then*

1. $d_H(z, z') < (3\ell + 1) \cdot 2^j$;
2. $d(u, v) < 2^j/8 + (3\ell + 1) \cdot 2^j$ for every $u \in C(z)$ and $v \in C(z')$.

Proof. Let the nodes along the path in F_j be $C(z_1), C(z_2), \dots, C(z_{k+1})$, where $z = z_1$ and $z_{k+1} = z'$. Let the corresponding augmentation edges be $(u_1, v_1), \dots, (u_k, v_k)$ where $u_1 \in C(z_1), v_k \in C(z_{k+1})$ and $v_i, u_{i+1} \in C(z_i)$ for $1 \leq i \leq k$. Note that each augmentation edge (u_i, v_i) has two corresponding bridge edges (u_i, z_i) and (v_i, z_{i+1}) .

We will now show that there is a short path in H from u to v via these augmentation and bridge edges. Consider the path P in H from z_1 to $z_{\ell+1}$ consisting of these augmentation edges and their corresponding bridge edges. There are ℓ augmentation edges and they have length at most 2^{j+1} each. There are 2ℓ bridge edges and they have length at most $2^j/16$. Therefore, $d(P) \leq 3\ell \cdot 2^{j+1}$. This proves the first part of the lemma.

We now turn to proving the second part of the lemma. Observe that $d(u, v) \leq d(u, z) + d(z, z') + d(z', v)$. Now $d(z, z') \leq d_H(z, z') \leq (3\ell + 1) \cdot 2^j$, from the first part of the lemma. We also have $d(u, z), d(z', v) \leq 2^j/16$ by definition of the clusters. Thus, $d(u, v) \leq 2^j/8 + (3\ell + 1) \cdot 2^j$, as desired. \square

With this lemma in hand, we can lower-bound the girth of F_j .

Lemma 6.3 (Girth Bound). *The girth of the meta-graph F_j is at least $\log k$.*

Proof. Suppose, towards a contradiction, that the meta-graph has a cycle of length less than $\log k$. Suppose that $(C(z), C(z'))$ is the meta-edge of the cycle that was added last, and $(u, v) \in A_j$ is the corresponding augmentation edge with $u \in C(z)$ and $v \in C(z')$. The edge (u, v) was added during the j -th iteration of the outer for loop. Consider the state of Algorithm 4 and the meta-graph right before (u, v) was added to H and $(C(z), C(z'))$ was added to F_j . At this time, there is a path between nodes $C(z)$ and $C(z')$ in F_j with less than $\log k$ edges, so by Lemma 6.2, we have $d_H(u, v) < 4 \log k \cdot 2^j \leq 4 \log k \cdot d(u, v)$. However, this contradicts the condition for adding (u, v) to A_j and so F_j has girth at least $\log k$. \square

Combining the previous lemma with the following result from extremal graph theory will allow us to bound $|A_j|$ later.

Lemma 6.4. [Bol04] *Every graph on N vertices and girth k has at most $O(N^{1+2/(k-1)})$ edges.*

Let OPT be the cost of the optimal Steiner forest for the n terminal pairs.

Lemma 6.5 (Cost Bound). $\sum_j 2^{j+1}|A_j| \leq O(1) \sum_j 2^{j+1}|Z_j| \leq O(\log k) \text{OPT}$.

Proof. By Lemmas 6.3 and Lemma 6.4, we have $|A_j| \leq O(1)|Z_j|$. Now we bound OPT in terms of $|Z_j|$. Place a ball of radius $2^j/32$ around each $z \in Z_j$. Since for any $z, z' \in Z_j$, $d(z, z') \geq 2^j/16$, these balls are disjoint. For any $z \in Z_j$, the distance from z and its mate is at least 2^j so the optimal solution has a path of length at least $2^j/32$ that is contained in the ball around z . This implies that $\text{OPT} \geq (2^j/32) \cdot |Z_j|$. Putting all of this together, we have $\sum_j 2^{j+1}|A_j| \leq \sum_j 2^{j+1} \cdot O(1)|Z_j| \leq O(\log k) \text{OPT}$. \square

We now turn to bounding the number of edges $|E(H)| = \sum_j |A_j|$. The main idea is to find for each scale j , a subset of terminals X'_j such that $|A_j| \leq O(1)|X'_j|$, and each terminal belongs in X'_j for a constant number of scales j . This allows us to charge $|A_j|$ to the terminals in X'_j for each j , and each terminal only receives a charge of $O(1)$ across all j . Define $\mathcal{C}_j = \{C(z) : z \in Z_j\}$, the set of clusters corresponding to centers Z_j . First, we show that there cannot be any large subset of clusters in \mathcal{C}_j that are far apart from each other. In particular, for any such subset $\mathcal{C}_j^+ \subseteq \mathcal{C}_j$, $|A_j|$ is at most twice the number of remaining clusters $\mathcal{C}_j \setminus \mathcal{C}_j^+$.

Lemma 6.6. *Let \mathcal{C}_j^+ be any subset of \mathcal{C}_j such that for any clusters $C(z), C(z') \in \mathcal{C}_j$ and terminals $u \in C(z)$ and $v \in C(z')$, we have $d(u, v) \geq 8 \cdot 2^j$. Let $\mathcal{C}_j^- = \mathcal{C}_j \setminus \mathcal{C}_j^+$ be the remaining clusters. Then, we have $|A_j| \leq 2|\mathcal{C}_j^-|$.*

Proof. Without loss of generality, we assume that the meta-graph F_j is connected (otherwise we can apply the same argument on each component of F_j). We begin by showing that $|\mathcal{C}_j^+| \leq |\mathcal{C}_j^-|$. The second part of Lemma 6.2 implies that in F_j , if there exists a path consisting of at most 2 meta-edges between meta-nodes $C(z)$ and $C(z')$, then $d(u, v) < 8 \cdot 2^j$ for all $u \in C(z)$ and $v \in C(z')$. Thus, any path between two meta-nodes $C(z), C(z') \in \mathcal{C}_j^+$ has to contain at least 3 meta-edges. This implies that every meta-node in \mathcal{C}_j^+ can only have meta-nodes in \mathcal{C}_j^- as neighbors, and no two meta-nodes of \mathcal{C}_j^+ can share the same neighbor in \mathcal{C}_j^- . Since F_j is connected, we have that each meta-node in \mathcal{C}_j^+ must have at least one neighbor in \mathcal{C}_j^- . Putting all these facts together gives us $|\mathcal{C}_j^+| \leq |\mathcal{C}_j^-|$.

By Lemmas 6.3 and Lemma 6.4, we have $|A_j| \leq O(1)|Z_j|$. Since $|Z_j| = |\mathcal{C}_j|$ and $|\mathcal{C}_j^+| \leq |\mathcal{C}_j^-|$, we have $|A_j| \leq 2|\mathcal{C}_j^-|$. \square

Lemma 6.7 (Sparsity Bound). $\sum_j |A_j| \leq O(k)$.

Proof. We can easily construct a family of nets N_j , one for each j , such that N_j is a $2^j/32$ -net and is nesting, i.e. $N_j \subseteq N_{j-1}$ for all j .

First, we claim that each cluster $C(z) \in \mathcal{C}_j$ contains at least one net point $v \in N_j$. Suppose, towards a contradiction, that there exists a cluster $C(z) \in \mathcal{C}_j$ such that $C(z) \cap N_j = \emptyset$. Consider a net point $v \in N_j$. Since each terminal is added to the cluster of the closest center in Z_j and v is not assigned to z , it must be the case that $d(v, z) \geq 2^j/32$; otherwise, z would have been the closest center to v because $d(z, z') \geq 2^j/16$ for every $z' \in Z_j$. Therefore, every net point $v \in N_j$ has $d(v, z) \geq 2^j/32$, but this contradicts the fact that N_j is a $2^j/32$ -net. Thus, each cluster in \mathcal{C}_j contains at least one net point $v \in N_j$.

Next, for each terminal v , we define its *rank* to be the largest j such that $v \in N_j$. Let $N_j^+ = N_{j+9}$ and $X'_j = N_j \setminus N_j^+$. Note that because of the nesting property of the nets, the terminals in X'_j are exactly those with rank between j and $j+9$. Let \mathcal{C}_j^+ be the subset of \mathcal{C}_j that intersects with N_j^+ and $\mathcal{C}_j^- = \mathcal{C}_j \setminus \mathcal{C}_j^+$. Now, for every $C(z), C(z') \in \mathcal{C}_j^+$ and $u \in C(z)$ and $v \in C(z')$, we have

$d(u, v) \geq d(z, z') - d(u, z) - d(v, z') \geq 2^{j+9}/32 - 2 \cdot 2^j/16 \geq 2^{j+3}$. Therefore, by Lemma 6.6, we have $|A_j| \leq 2|\mathcal{C}_j^-|$.

Now, every cluster $C \in \mathcal{C}_j^-$ contains a terminal of X'_j ; this is because every cluster contains at least one terminal of N_j and clusters in \mathcal{C}_j^- do not contain any terminal of N_j^+ . Since the clusters \mathcal{C}_j^- are disjoint, we have $|\mathcal{C}_j^-| \leq |X'_j|$. We now have $\sum_j |A_j| \leq 2 \sum_j |X'_j|$. Observe that every terminal is contained in X'_j for at most a constant number of scales j (each terminal $v \in X'_j$ has $\text{rank}(v) \in [j, j+9]$). Since there are $2k$ terminals, we have $\sum_j |X'_j| \leq O(k)$ and so, $\sum_j |A_j| \leq O(k)$. \square

Lemmas 6.1, 6.5 and 6.7 give us Theorem 1.5.

Handling lack of knowledge of k . Our Algorithm 4 requires the knowledge of the value of k in determining when augmentation edges must be added, which is not available a priori in the online setting. However, it is relatively straightforward to get over this by noting that the subgraph H maintained by the algorithm can be decomposed into one set of edges for every distance scale j . At a specific distance scale, the number of centers in Z_j gives us a good lower bound on k_j which can be used in place of k in determining the stretch condition: in other words, we use $4 \log k_j$ in the stretch condition when processing the loop for j , and update it whenever the number of clusters in Z_j doubles.

7 Randomized Oblivious Buy at Bulk

In this section, we describe our randomized algorithm for the oblivious case, where we don't know the buy-at-bulk function in advance, and hence have to give a solution that is simultaneously good for all buy-at-bulk functions. Let us recall our main result for this section, Theorem 1.2:

Theorem 1.2 (Oblivious Buy-at-Bulk, Randomized). *There exists a randomized online algorithm for the buy-at-bulk problem that produces a routing P such that for all concave functions f ,*

$$\mathbb{E}[\text{cost}_f(P)] \leq O(\log^2 k) \text{OPT}_f.$$

Recall that this guarantee can also be achieved by performing Bartal's tree embeddings, but the ideas we develop here will be crucial to getting the deterministic algorithm in §8.

For this section and the next, we use g_i to denote the “rent-or-buy” concave function $g_i(x) = \min\{2^i, x\}$ and write OPT_i to be the cost of the optimal solution for g_i . Given a routing solution P , we denote by $\text{cost}_i(P)$ the cost of P under g_i . By Lemma 2.6 it suffices to show that simultaneously for all g_i ,

$$\text{cost}_i(P) \leq O(\log^2 k) \text{OPT}_i.$$

This is what we show in the rest of this section and the next.

7.1 High-Level Description

We begin by reviewing the main ideas for approximating a fixed rent-or-buy function $g_i = \min\{x, 2^i\}$ online. Essentially, it boils down to choosing a good set of “buy” terminals X_i (which includes the root) online. The terminal set X_i is connected by buying the edges of an online Steiner tree F_i on them. The remaining terminals are called “rent” terminals and are routed as follows: For each rent terminal u , let $X_i(v)$ be the set of bought terminals at the time of its arrival; the terminal u is then routed to the nearest buy terminal $v \in X_i(v)$ (the edge (u, v) is rented), and then via v 's path to the root in F_i . Observe that the cost of buying edges is $2^i d(F_i)$ and the cost of renting

edges is $\sum_{v \notin X_i} d(v, X_i(v))$. A key ingredient³ (see, e.g., [AAB04]) is the fact that if we choose each terminal v to be a buy terminal with probability 2^{-i} , then

$$2^i d(\text{Steiner}(X_i)) + \sum_{v \notin X_i} d(v, X_i(v)) \leq O(\log k) \text{OPT}_i. \quad (9)$$

But this approach requires knowing the particular function g_i , whereas we need to be oblivious, i.e., to be good for all g_i simultaneously. This suggests the following approach to designing an online oblivious algorithm. Firstly, we assign a *type* $\tau(v)$ to each terminal v at random: $\tau(v) = i$ with probability 2^{-i} . As before, we assign the root r to have type ∞ . Thus, if we define X_i to be the set of terminals with type at least i , then it satisfies Inequality (9).

Secondly, for each type i , we maintain an online spanner F_i on X_i using the algorithm from §6. We think of the edges that belong to a spanner of type at least i as the edges that our solution buys if it is given the i -th rent-or-buy function g_i . Then, each terminal v is routed using a path $P(v)$ that goes through spanners of progressively higher type until we reach the root. Let $P_i(v)$ be the prefix of $P(v)$ right before it enters a spanner of type at least i . This is the part of $P(v)$ that will be rented for g_i . At a high level, the lightness and low distortion properties of the spanners will allow us to prove that the routing P satisfies the following bounded cost properties for all i :

- Bounded buy cost: $2^i d(F_i) \leq O(\log k) \text{Steiner}(X_i)$;
- Bounded rent cost: $\sum_{v: \tau(v) < i} d(P_i(v)) \leq O(\log k) d(v, X_i(v))$, i.e. the distance from v to $X_i(v)$ on the path $P(v)$ is not much larger than the actual distance.

Some care needs to be taken in choosing which spanners to route through. Fix a terminal v , and for each $i \geq \tau(v)$, let w_i be the nearest terminal of type at least i and call it the i -th *waypoint* for v . A natural idea is to route from v to $w_{\tau(v)+1}$ using the spanner $F_{\tau(v)}$, then from $w_{\tau(v)+1}$ to $w_{\tau(v)+2}$ using $F_{\tau(v)+1}$, and so on, until we reach the root. However, the problem is that even if we were to route between consecutive waypoints w_i, w_{i+1} using the direct edge (w_i, w_{i+1}) , the total length of the path before we can reach a given $w_{i'}$ might be too large. To overcome this issue, we bypass waypoints that are of roughly the same distance from v and route through a subset of waypoints whose distances from v are geometrically increasing. See Algorithm 5 for a formal description of our algorithm.

7.2 Analysis

We now analyze $\text{cost}_i(P)$. For each terminal v , define $X_i(v)$ to be the subset of terminals of type at least i when v arrived, excluding v itself.

Lemma 7.1. $\text{cost}_i(P) \leq 2^i \sum_{i' \geq i} d(F_{i'}) + \sum_{v: \tau(v) < i} d(P_i(v))$.

Proof. Recall that $\text{cost}_i(P) = \sum_{e \in E} d(e) \min\{2^i, \text{load}_P(e)\}$, where $\text{load}_P(e)$ is the total number of terminals v whose paths $P(v)$ include e . We will divide the edges used by the routing into two sets \mathcal{R}_i and \mathcal{B}_i , and think of \mathcal{R}_i as the set of rented edges (i.e. we pay $d(e) \text{load}_P(e)$ on edges $e \in \mathcal{R}_i$), and \mathcal{B}_i as the set of bought edges (i.e. we pay $2^i d(e)$ per edge in \mathcal{B}_i). Define $\mathcal{R}_i = \bigcup_v P_i(v)$ and $\mathcal{B}_i = \bigcup_v P(v) \setminus P_i(v)$. We have

$$\text{cost}_i(P) = \sum_{e \in E} d(e) \min\{2^i, \text{load}_P(e)\} \leq 2^i d(\mathcal{B}_i) + \sum_{e \in \mathcal{R}_i} d(e) \text{load}_P(e).$$

³We will actually need a stronger version, Lemma 7.2

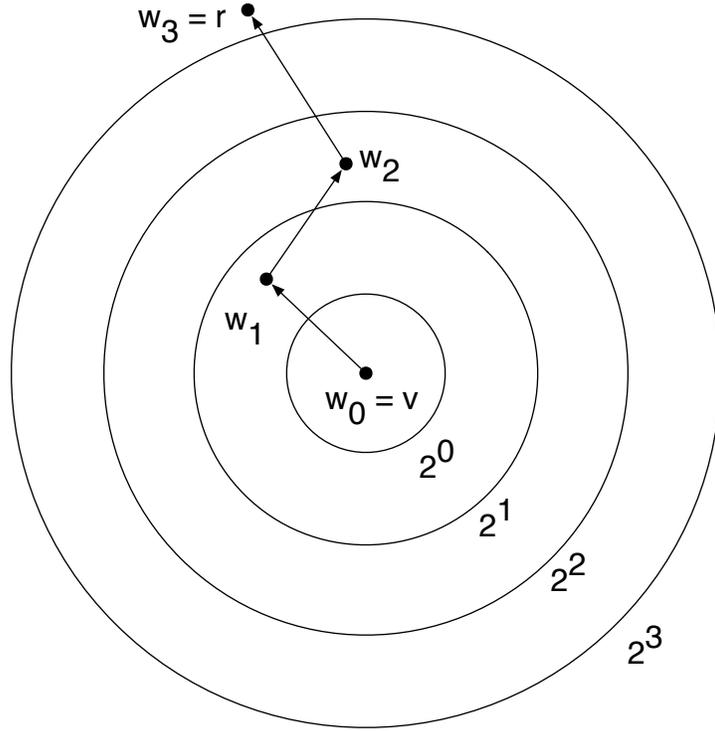


Figure 5: The arrow from w_i to w_{i+1} represents the shortest path between $w_i(v)$ and $w_{i+1}(v)$ in $F_{\tau(w_i)}$. Note that it is possible for a ring (in this figure, the fourth ring) to not have a waypoint.

Algorithm 5 Randomized Online Oblivious Algorithm for Buy-at-Bulk

- 1: Initialize $X_i \leftarrow \emptyset, F_i \leftarrow \emptyset$ for all i
 - 2: Set $\tau(r) \leftarrow \infty$
 - 3: **while** terminal v arrives **do**
 - 4: Choose a random type $\tau(v)$, with $\tau(v) = i$ w.p. 2^{-i}
 - 5: For $i \leq \tau(v)$, add v to X_i and update spanner F_i to include v
 - 6: For $i > \tau(v)$, define waypoint w_i to be terminal of X_i nearest to v
 - 7: Set initial waypoint to be v
 - 8: **for** $j = 0$ up to $\lfloor \log d(v, r) \rfloor$ **do**
 - 9: Define the ring $R_j = \{w_i : d(v, w_i) \in [2^j, 2^{j+1})\}$
 - 10: **if** $R_j \neq \emptyset$ **then**
 - 11: Choose waypoint w_i in R_j of highest type and route from current waypoint $w_{i'}$ to w_i along shortest path in $F_{i'}$
-

Since $P(v) \setminus P_i(v) \subseteq \bigcup_{i' \geq i} F_{i'}$, we have that $2^i d(\mathcal{B}_i) \leq 2^i \sum_{i' \geq i} d(F_{i'})$. Now observe that for each $e \in \mathcal{R}_i$, $\text{load}_P(e)$ is exactly the number of terminals v of type less than i with $P_i(v) \ni e$. This is because $P_i(v) \neq \emptyset$ only for terminals v of type less than i . Thus, we have $\sum_{e \in \mathcal{R}_i} d(e) \text{load}_P(e) = \sum_{v: \tau(v) < i} d(P_i(v))$. \square

We define $\text{Steiner}(S)$ to be the cost of the optimal Steiner tree for terminals S .

Lemma 7.2. *We have that*

1. $E[2^i \text{Steiner}(X_i)] \leq \text{OPT}_i$ and
2. $E[\sum_v d(v, X_i(v))] \leq O(\log k) \text{OPT}_i$.

Proof. We begin with the first part of the lemma. For terminal $v \in X$, let $P_i^*(v)$ be the path to the root in OPT_i . Define \mathcal{B}_i^* to be the set of edges over which at least 2^i terminals are routed, and for each v , define $\mathcal{R}_i^*(v) = P_i^*(v) \setminus \mathcal{B}_i^*$. The cost of the optimal solution is

$$\text{OPT}_i = 2^i d(\mathcal{B}_i^*) + \sum_{v \in X} d(\mathcal{R}_i^*(v)).$$

Define $H_i = \bigcup_{v \in X_i} P_i^*(v)$. Since H_i is a Steiner tree for terminals X_i ,

$$\begin{aligned} E[2^i \text{Steiner}(X_i)] &\leq 2^i d(\mathcal{B}_i^*) + E\left[2^i \sum_{v \in X_i} d(\mathcal{R}_i^*(v))\right] \\ &\leq 2^i d(\mathcal{B}_i^*) + 2^i \sum_v d(\mathcal{R}_i^*(v)) \cdot \Pr[v \in X_i] \\ &\leq 2^i d(\mathcal{B}_i^*) + \sum_v d(\mathcal{R}_i^*(v)) = \text{OPT}_i. \end{aligned}$$

This proves the first part of the lemma.

We now turn to the second part of the lemma. For this, we will show that $E\left[2^i \sum_{v \in X_i} d(v, X_i(v))\right] = E\left[\sum_v d(v, X_i(v))\right]$, then prove $\sum_{v \in X_i} d(v, X_i(v)) \leq O(\log k) \text{Steiner}(X_i)$, and apply the first part of the lemma. By linearity of expectation,

$$\begin{aligned} E\left[2^i \sum_{v \in X_i} d(v, X_i(v))\right] &= 2^i \sum_v E[d(v, X_i(v)) \cdot \mathbf{1}\{v \in X_i\}] \\ &= 2^i \sum_v E[d(v, X_i(v))] \cdot \Pr[v \in X_i] \\ &= E\left[\sum_v d(v, X_i(v))\right] \end{aligned}$$

where the second equality follows from the fact that whether $v \in X_i$ is independent of $d(v, X_i(v))$, and the last equality from the fact that $\Pr[v \in X_i] = 1/2^i$.

Next, we prove $\sum_{v \in X_i} d(v, X_i(v)) \leq O(\log k) \text{Steiner}(X_i)$. Consider the online Steiner tree instance defined on terminals X_i . The cost of the greedy algorithm, which connects each terminal to the nearest previous terminal, is exactly $\sum_{v \in X_i} d(v, X_i(v))$. The fact that the greedy algorithm is $O(\log k)$ -competitive for the online Steiner tree problem implies $\sum_{v \in X_i} d(v, X_i(v)) \leq O(\log k) \text{Steiner}(X_i)$.

Therefore,

$$E \left[\sum_v d(v, X_i(v)) \right] = E \left[2^i \sum_{v \in X_i} d(v, X_i(v)) \right] \leq O(\log k) E[2^i \text{Steiner}(X_i)] \leq O(\log k) \text{OPT}_i.$$

This completes the proof of the second part of the lemma. \square

In the following, denote by $P_i(v)$ the prefix of $P(v)$ right before it enters a spanner $F_{i'}$ with $i' \geq i$.

Lemma 7.3. *For every $i > \tau(v)$, we have $d(P_i(v)) \leq O(\log k)d(v, X_i(v))$.*

Proof. Note that w_i is the nearest terminal of $X_i(v)$ so $d(v, w_i) = d(v, X_i(v))$. Define W to be the subset of waypoints w_i such that w_i is the waypoint of highest type in some ring R_j , and $I = \{i : w_i \in W\}$. For every type $i' > \tau(v)$ that is not in I , there exists $i \in I$ such that $i > i'$ and $d(v, w_i) \leq 2d(v, w_{i'})$; moreover, since $i > i'$, we have $d(P_{i'}(v)) \leq d(P_i(v))$. Therefore, it suffices to prove $d(P_i(v)) \leq O(\log k)d(v, w_i)$ for types $i \in I$.

For brevity, we abuse notation and write w_ℓ to denote the waypoint in W with ℓ -th highest type; F_ℓ to denote the spanner of type $\tau(w_\ell)$. Let $P_\ell(v)$ be the prefix of $P(v)$ ending at w_ℓ . We now show that $d(P_\ell(v)) \leq O(\log k)d(v, w_\ell)$.

Since $P_\ell(v)$ consists of the union of the shortest path between $w_{\ell-1}$ and w_ℓ in $F_{\ell-1}$, for every $\ell' < \ell$, the length of $P_\ell(v)$ is $d(P_\ell(v)) = \sum_{\ell' \leq \ell} d_{F_{\ell'-1}}(w_{\ell'-1}, w_{\ell'})$. For every $\ell' < \ell$, since $F_{\ell'-1}$ is a spanner with stretch $O(\log k)$, we have

$$d_{F_{\ell'-1}}(w_{\ell'-1}, w_{\ell'}) \leq O(\log k)d(w_{\ell'-1}, w_{\ell'}) \leq O(\log k)(d(v, w_{\ell'-1}) + d(v, w_{\ell'})).$$

Thus,

$$d(P_\ell(v)) \leq O(\log k) \sum_{\ell' < \ell} (d(v, w_{\ell'-1}) + d(v, w_{\ell'})) \leq O(\log k)d(v, w_\ell),$$

where the last inequality follows from the fact that we chose W to contain at most one waypoint per ring and the rings have geometrically increasing distances. \square

Lemma 7.4. $E[\text{cost}_i(P)] = O(\log^2 k) \text{OPT}_i$.

Proof. By Lemma 7.1, the expected cost of the algorithm's solution is $E[\text{cost}_i(P)] \leq E[2^i \sum_{i' \geq i} d(F_{i'})] + E[\sum_{v: \tau(v) < i} d(P_i(v))]$. Let $i' > i$. Since $d(F_{i'})$ is a $(\log k)$ -competitive $(\log k)$ -spanner on X_i , we have $d(F_{i'}) \leq O(\log k) \text{Steiner}(X_{i'})$. Together with the fact that $\text{Steiner}(X_{i'}) \leq \text{Steiner}(X_i)$ (since $X_{i'}$ contains X_i), we have

$$E[2^i \sum_{i' \geq i} d(F_{i'})] \leq O(\log k) \sum_{i' \geq i} E[2^i \text{Steiner}(X_{i'})] \leq O(\log^2 k) E[2^i \text{Steiner}(X_i)] \leq O(\log^2 k) \text{OPT}_i,$$

where the last inequality follows from Lemma 7.2.

Next, we turn to bounding $E[\sum_{v: \tau(v) < i} d(P_i(v))]$. By Lemmas 7.3 and 7.2, we have

$$E \left[\sum_{v: \tau(v) < i} d(P_i(v)) \right] \leq O(\log k) E \left[\sum_{\tau(v) < i} d(v, X_i(v)) \right] \leq O(\log^2 k) \text{OPT}_i.$$

\square

8 Deterministic Oblivious Buy at Bulk

In this section, we present our deterministic algorithm for the oblivious setting, and prove Theorem 1.3.

Theorem 1.3 (Oblivious Buy-at-Bulk, Deterministic). *There exists a deterministic online algorithm for the buy-at-bulk problem that produces a routing P such that for all concave functions f ,*

$$\text{cost}_f(P) \leq O(\log^{2.5} k) \text{OPT}_f.$$

Our deterministic oblivious algorithm is based on the randomized version from §7 (Algorithm 5) with two modifications. First, we remove the need for randomness by using the deterministic type selection rule used by our non-oblivious algorithm from §5 (Algorithm 3). We can show that this modification already yields an $O(\log^3 k)$ -competitive algorithm. To get the competitive ratio down to $O(\log^{2.5} k)$, we need to use a different routing scheme. For clarity of presentation, we first present the simpler $O(\log^3 k)$ -competitive algorithm and later (in Section 8.2) show how to modify it to obtain the $O(\log^{2.5} k)$ -competitive algorithm of Theorem 1.3.

8.1 An $O(\log^3 k)$ -competitive Algorithm

We will assign a type $\tau(v)$ to each terminal v , with the root having type $\tau(r) = \infty$. As before, we use X_i to denote the set of terminals with type at least i . We will also maintain a $(\log k, \log k)$ -spanner F_i on X_i .

Algorithm 6 Simpler Online Deterministic Oblivious Algorithm for Buy-at-Bulk

- 1: Initialize $X_i \leftarrow \emptyset, F_i \leftarrow \emptyset$ for all i
 - 2: Set $\tau(r) \leftarrow \infty$
 - 3: **while** terminal v arrives **do**
 - 4: **for** each type i **do**
 - 5: Define $\text{class}_i(v)$ to be the highest j such that v is a level j center in spanner F_i if it were to be included in F_i
 - 6: Define the ball $B_i(v) = \{u : d(u, v) \leq 2^{\text{class}_i(v)} / 2^3\}$.
 - 7: Assign type $\tau(v) = \max \{i : |B_i(v)| \geq 2^i\}$.
 - 8: For $i \leq \tau(v)$, update spanner F_i to include v
 - 9: For $i > \tau(v)$, define waypoint w_i to be terminal of X_i nearest to v
 - 10: Set initial waypoint to be v
 - 11: **for** $j = 0$ up to $\lfloor \log d(v, r) \rfloor$ **do**
 - 12: Define the ring $R_j = \{w_i : d(v, w_i) \in [2^j, 2^{j+1})\}$
 - 13: **if** $R_j \neq \emptyset$ **then**
 - 14: Choose waypoint w_i in R_j of highest type and route from current waypoint $w_{i'}$ to w_i along shortest path in $F_{i'}$
-

8.1.1 Analysis

Fix a type i . We will now show that $\text{cost}_i(P) \leq O(\log^2 k) \text{OPT}_i$. As before, define $P_i(v)$ to be the segment of $P(v)$ right before it enters a spanner $F_{i'}$ with $i' \geq i$. (Note that $P_i(v) = \emptyset$ if $\tau(v) \geq i$.)

By Lemma 7.1, we have $\text{cost}_i(P) \leq 2^i \sum_{i' \geq i} d(F_{i'}) + \sum_{v: \tau(v) < i} d(P_i(v))$. Next, we bound $2^i \sum_{i' \geq i} d(F_{i'})$ and $\sum_{v: \tau(v) < i} d(P_i(v))$ separately.

Lemma 8.1. $\sum_{v:\tau(v)<i} d(P_i(v)) \leq O(\log^2 k) \text{OPT}_i$.

Proof. Using the same argument as in the proof of Lemma 7.3 gives us that for every terminal v , and type $i > \tau(v)$, we have $d(P_i(v)) \leq O(\log k) d(v, X_i(v))$. Thus,

$$\sum_{v:\tau(v)<i} d(P_i(v)) \leq O(\log k) \sum_{v:\tau(v)<i} d(v, X_i(v)).$$

Using a similar proof as that of Lemma 5.2, we also have

$$\sum_{v:\tau(v)<i} d(v, X_i(v)) \leq O(\log k) \text{OPT}_i.$$

Combining these two inequalities gives us the lemma. \square

Lemma 8.2. $2^i \sum_{i' \geq i} d(F_{i'}) \leq O(\log^2 k) \text{OPT}_i(T)$ for any HST embedding T .

Proof. Let $i' > i$. By Lemma 6.5 and the fact that $\text{class}_{i'}(v)$ is the highest level j such that v is a level- j center in $F_{i'}$, we have that $2^i d(F_{i'}) \leq O(1) 2^i \sum_{v \in X_{i'}} 2^{\text{class}_{i'}(v)}$. Now, we would like to say that the RHS of the previous inequality is at most $O(1) \text{OPT}_{i'}(T)$ for any HST embedding T . We can prove this using a similar proof as in Lemma 5.3 if for every v with $\tau(v) \geq i'$, we have $|B_{i'}(v)| \geq 2^{i'}$. Unfortunately, this is not necessarily the case. A terminal v might have $\tau(v) \geq i'$, i.e. it has at least $2^{\tau(v)} \geq 2^{i'}$ terminals in its $\tau(v)$ -th ball $B_{\tau(v)}(v)$, but its i' -th ball $B_{i'}(v)$ might have a radius much smaller than $B_{\tau(v)}(v)$ (i.e. $\text{class}_{i'}(v) < \text{class}_{\tau(v)}(v)$), and thus does not contain at least $2^{i'}$ terminals. In this case, roughly speaking, we will charge the cost of including v in $F_{i'}$ to the cost of including it in $F_{\tau(v)}$.

More formally, for each type i' , define $A_{i'} = \{v \in X_{i'} : |B_{i'}(v)| \geq 2^{i'}\}$, i.e. this is the set of terminals who can “pay” to be included in spanner $F_{i'}$.

Claim 8.3. *Let v be a terminal with type $\tau(v) > i'$. If $v \notin A_{i'}$, then $\text{class}_{i'}(v) < \text{class}_{\tau(v)}(v)$.*

Proof. The balls $B_{\tau(v)}(v)$ and $B_{i'}(v)$ have radii $2^{\text{class}_{\tau(v)}(v)-3}$ and $2^{\text{class}_{i'}(v)-3}$, respectively. By definition of types, $v \in A_{\tau(v)}$ and so $|B_{\tau(v)}(v)| \geq 2^{\tau(v)} \geq 2^{i'}$. Since $v \notin A_{i'}$, the radius of $B_{i'}(v)$ must be smaller than the radius of $B_{\tau(v)}(v)$ and so $\text{class}_{i'}(v) < \text{class}_{\tau(v)}(v)$. \square

Define $N(v) = \{i' < \tau(v) : v \in X_{i'} \setminus A_{i'}\}$, i.e. the set of types i' for which v is included in spanner $F_{i'}$ but could not pay to be included in $F_{i'}$. Combining this with the above claim, we get

$$\begin{aligned} \sum_{i' \geq i} d(F_{i'}) &\leq O(1) \sum_{i' \geq i} \left(\sum_{v \in A_{i'}} 2^{\text{class}_{i'}(v)} + \sum_{v \in X_{i'} \setminus A_{i'}} 2^{\text{class}_{\tau(v)}(v)} \right) \\ &= O(1) \sum_{i' \geq i} \sum_{v \in A_{i'}} 2^{\text{class}_{i'}(v)} + \sum_{v \in X_i} |N(v)| \cdot 2^{\text{class}_{\tau(v)}(v)} \\ &\leq O(1) \sum_{i' \geq i} \sum_{v \in A_{i'}} 2^{\text{class}_{i'}(v)} + O(\log k) \sum_{v \in X_i} 2^{\text{class}_{\tau(v)}(v)} \\ &\leq O(\log k) \sum_{i' \geq i} \sum_{v \in A_{i'}} 2^{\text{class}_{i'}(v)} \end{aligned}$$

where the last inequality follows from the fact that $v \in A_{\text{class}_{\tau(v)}(v)}$ and $\tau(v) \geq i$.

Using a similar proof as in Lemma 5.3, we have $2^i \sum_{v \in A_{i'}} 2^{\text{class}_{i'}(v)} \leq O(1) \text{OPT}_i(T)$. Thus,

$$2^i \sum_{i' \geq i} d(F_{i'}) \leq O(\log^2 k) \text{OPT}_i(T),$$

as desired. \square

With these two lemmas and Lemma 2.5 in hand, we have that $\text{ALG} \leq O(\log^3 k) \text{OPT}_i$ and thus the algorithm has competitive ratio $O(\log^3 k)$.

8.2 Improving to $O(\log^{2.5} k)$

We now describe the modified routing scheme. Suppose we know the number of terminals k of the online instance in advance (we will show how to remove this assumption at the end). Let $L = \sqrt{\log k}$. The main idea is that instead of including each terminal v in spanner F_i for every $i \leq \tau(v)$, we only include v in F_i for $\tau(v) - L \leq i \leq \tau(v)$. This helps us bound the number of spanners that v is included but could not pay for by L . However, we now cannot reach all terminals of type higher than i using spanner F_i . To cope with this, we will group the L^2 different types into consecutive intervals of length L . The routing for a terminal v now proceeds in several phases, one per interval, starting from the interval containing $\tau(v)$. During phase t (corresponding to the interval $[tL, (t+1)L]$), we find a path Q_t from the waypoint w_{tL} (or v if the interval contains $\tau(v)$) to $w_{(t+1)L}$ by passing through intermediate waypoints that belong to rings geometrically increasing distances.

In the following, we write $u \in F_i(v)$ to mean that u is a terminal included in spanner F_i before v arrived.

Algorithm 7 Improved Online Deterministic Oblivious Algorithm for Buy-at-Bulk

- 1: Initialize $X_i \leftarrow \emptyset, F_i \leftarrow \emptyset$ for all i
 - 2: Set $\tau(r) \leftarrow \infty$
 - 3: **while** terminal v arrives **do**
 - 4: **for** each type i **do**
 - 5: Define $\text{class}_i(v)$ to be the highest j such that v is a level j center in spanner F_i if it were to be included in F_i
 - 6: Define the ball $B_i(v) = \{u : d(u, v) \leq 2^{\text{class}_i(v)}/2^3\}$.
 - 7: Assign type $\tau(v) = \max \{i : |B_i(v)| \geq 2^i\}$.
 - 8: For $\tau(v) - L \leq i \leq \tau(v)$, update spanner F_i to include v
 - 9: For $i > \tau(v)$, define waypoint w_i to be terminal of F_i nearest to v
 - 10: Set initial waypoint to be v
 - 11: **for** t such that $\tau(v) - 1 \leq tL \leq \log k$ **do**
 - 12: Define the ring $R_j(t) = \{w_i : tL \leq i \leq (t+1)L \wedge d(v, w_i) \in [2^j, 2^{j+1})\}$
 - 13: Initialize $Q_t \leftarrow \emptyset$
 - 14: **while** current waypoint is not $w_{(t+1)L}$ **do**
 - 15: **if** $R_j(t) \neq \emptyset$ **then**
 - 16: Choose waypoint w_i in $R_j(t)$ of highest type and route from current waypoint $w_{i'}$ to w_i along shortest path in $F_{i'}$
 - 17: Add shortest path to Q_t
-

Lemma 8.4. For each terminal v , we have

$$d(P_i(v)) \leq O(\log k) \left(\sum_{t: \tau(v) < (t+1)L \leq i} d(v, F_{(t+1)L}(v)) + d(v, F_i(v)) \right).$$

Proof. Let Q_t be the segment of $P(v)$ starting from $w_{tL}(v)$ to $w_{(t+1)L}(v)$ and $Q_t(i)$ be the segment of Q_t right before it enters a spanner $F_{i'}$ with $i' \geq i$. Using a similar proof as in Lemma 7.3, for every t and $i \in [tL, (t+1)L]$, we get that $d(Q_t(i)) \leq O(\log k)d(v, w_i(v))$ and $d(Q_t) \leq O(\log k)d(v, w_{(t+1)L}(v))$. For each type i , define t_i to be the integer such that $i \in [t_iL, (t_i+1)L]$. We have

$$d(P_i(v)) = \sum_{t: \tau(v) < (t+1)L < i} d(Q_t) + d(Q_{t_i}(i)) \leq O(\log k) \left(\sum_{t: \tau(v) < (t+1)L < i} d(v, w_{(t+1)L}(v)) + d(v, w_i(v)) \right).$$

□

Lemma 8.5. $\sum_{v: \tau(v) < i} d(P_i(v)) \leq O(\log^{2.5} k) \text{OPT}_i$.

Proof. Using a similar proof as Lemma 5.2, we get that $\sum_{v: \tau(v) < i'} d(v, F_{i'}(v)) \leq O(\log k) \text{OPT}_{i'}$. Combining this with Lemma 8.4 gives us

$$\begin{aligned} \sum_{v: \tau(v) < i} d(P_i(v)) &\leq O(\log k) \left(\sum_{t: (t+1)L < i} \sum_{v: \tau(v) < (t+1)L} d(v, F_{(t+1)L}(v)) + \sum_{v: \tau(v) < i} d(v, F_i(v)) \right) \\ &\leq O(\log^2 k) \left(\sum_{t: (t+1)L < i} \text{OPT}_{(t+1)L} + \text{OPT}_i \right) \\ &\leq O(\log^{2.5} k) \text{OPT}_i, \end{aligned}$$

where the last inequality follows from the fact that $\text{OPT}_{i'} \leq \text{OPT}_i$ if $i' < i$ and that the number of integers t such that $(t+1)L < i$ is at most $L = \sqrt{\log k}$. □

Lemma 8.6. $2^i \sum_{i' \geq i} d(F_{i'}) \leq O(\log^{1.5} k) \text{OPT}_i(T)$ for any HST embedding T .

Proof. The proof of this lemma is similar to that of Lemma 8.2 except that the total number of types i' such that v is included in spanner $F_{i'}$ is at most $L = \sqrt{\log k}$. □

Removing the need to know k in advance. First, instead of using intervals of length $\sqrt{\log k}$, we will use intervals between perfect squares. In particular, the t -th interval is $[t^2, (t+1)^2]$. Second, each terminal v is included in spanners F_i where $\tau(v) \geq i \geq \tau(v) - \sqrt{i}$.

9 Conclusions

Several open questions remain: can we do better than $O(\log^2 k)$ for the oblivious randomized case, and $O(\log^{2.5} k)$ for the oblivious deterministic case? Can we get online constructions of tree embeddings with stretch better than $O(\log k \log \Delta)$?

References

- [AA97] Baruch Awerbuch and Yossi Azar. Buy-at-bulk network design. In *38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997*, pages 542–547, 1997.
- [AAB04] Baruch Awerbuch, Yossi Azar, and Yair Bartal. On-line generalized steiner problem. *Theor. Comput. Sci.*, 324(2-3):313–324, 2004.
- [AKPW95] Noga Alon, Richard M. Karp, David Peleg, and Douglas West. A graph-theoretic game and its application to the k -server problem. *SIAM J. Comput.*, 24(1):78–100, 1995.
- [And04] Matthew Andrews. Hardness of buy-at-bulk network design. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 115–124, 2004.
- [Bar96] Yair Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 184–193, 1996.
- [Bar98] Yair Bartal. On approximating arbitrary metrics by tree metrics. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 161–168, 1998.
- [Bar04] Yair Bartal. Graph decomposition lemmas and their role in metric embedding methods. In *Algorithms - ESA 2004, 12th Annual European Symposium, Bergen, Norway, September 14-17, 2004, Proceedings*, pages 89–97, 2004.
- [BC97] Piotr Berman and Chris Coulston. On-line algorithms for steiner tree problems (extended abstract). In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 344–353, 1997.
- [Bol04] Béla Bollobas. *Extremal Graph Theory*. Dover Publications, Incorporated, 2004.
- [ECKP15] Alina Ene, Deeparnab Chakrabarty, Ravishankar Krishnaswamy, and Debmalya Panigrahi. On-line buy-at-bulk network design. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 545–562, 2015.
- [FRT04] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004.
- [GE05] Ashish Goel and Deborah Estrin. Simultaneous optimization for concave costs: Single sink aggregation or single source buy-at-bulk. *Algorithmica*, 43(1-2):5–15, 2005.
- [GHR06] Anupam Gupta, Mohammad Taghi Hajiaghayi, and Harald Räcke. Oblivious network design. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 970–979, 2006.
- [GI06] Fabrizio Grandoni and Giuseppe F. Italiano. Improved approximation for single-sink buy-at-bulk. In *Algorithms and Computation, 17th International Symposium, ISAAC 2006, Kolkata, India, December 18-20, 2006, Proceedings*, pages 111–120, 2006.
- [GKK⁺01] Naveen Garg, Rohit Khandekar, Goran Konjevod, R. Ravi, F. Sibel Salman, and Amitabh Sinha II. On the integrality gap of a natural formulation of the single-sink buy-at-bulk network design problem. In *Integer Programming and Combinatorial Optimization, 8th International IPCO Conference, Utrecht, The Netherlands, June 13-15, 2001, Proceedings*, pages 170–184, 2001.
- [GKR03] Anupam Gupta, Amit Kumar, and Tim Roughgarden. Simpler and better approximation algorithms for network design. In *35th STOC*, pages 365–372, 2003.
- [GMM09] Sudipto Guha, Adam Meyerson, and Kamesh Munagala. A constant factor approximation for the single sink edge installation problem. *SIAM J. Comput.*, 38(6):2426–2442, 2009.

- [GP12] Ashish Goel and Ian Post. One tree suffices: A simultaneous $o(1)$ -approximation for single-sink buy-at-bulk. *Theory of Computing*, 8(1):351–368, 2012.
- [GRS11] Fabrizio Grandoni, Thomas Rothvoss, and Laura Sanità. From uncertainty to nonlinearity: Solving virtual private network via single-sink buy-at-bulk. *Mathematics of Operations Research*, 36(2):185–204, 2011.
- [IW91] Makoto Imase and Bernard M. Waxman. Dynamic steiner tree problem. *SIAM J. Discrete Math.*, 4(3):369–384, 1991.
- [JR09] Raja Jothi and Balaji Raghavachari. Improved approximation algorithms for the single-sink buy-at-bulk network design problems. *J. Discrete Algorithms*, 7(2):249–255, 2009.
- [KRY95] Samir Khuller, Balaji Raghavachari, and Neal E. Young. Balancing minimum spanning trees and shortest-path trees. *Algorithmica*, 14(4):305–321, 1995.
- [LPS88] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
- [MP98] Yishay Mansour and David Peleg. An approximation algorithm for minimum-cost network design. In *Robust Communication Networks: Interconnection and Survivability, Proceedings of a DIMACS Workshop, New Brunswick, New Jersey, USA, November 18-20, 1998*, pages 97–106, 1998.
- [SCRS97] F. Sibel Salman, Joseph Cheriyan, R. Ravi, and S. Subramanian. Buy-at-bulk network design: Approximating the single-sink edge installation problem. In *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, 5-7 January 1997, New Orleans, Louisiana.*, pages 619–628, 1997.
- [Tal02] Kunal Talwar. The single-sink buy-at-bulk LP has constant integrality gap. In *Integer Programming and Combinatorial Optimization, 9th International IPCO Conference, Cambridge, MA, USA, May 27-29, 2002, Proceedings*, pages 475–486, 2002.
- [Umb15] Seun Umboh. Online network design algorithms via hierarchical decompositions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1373–1387, 2015.