

One Tree Suffices: A Simultaneous $O(1)$ -Approximation for Single-Sink Buy-at-Bulk

Ashish Goel*

Department of Management Science and Engineering
Stanford University
Stanford, USA
ashishg@stanford.edu

Ian Post†

Department of Computer Science
Stanford University
Stanford, USA
itp@stanford.edu

Abstract—We study the single-sink buy-at-bulk problem with an unknown cost function. We wish to route flow from a set of demand nodes to a root node, where the cost of routing x total flow along an edge is proportional to $f(x)$ for some concave, non-decreasing function f satisfying $f(0)=0$. We present a simple, fast, combinatorial algorithm that takes a set of demands and constructs a single tree T such that for all f the cost $f(T)$ is a 47.45-approximation of the optimal cost for that f . This is within a factor of 2.33 of the best approximation ratio currently achievable when the tree can be optimized for a specific function. Trees achieving simultaneous $O(1)$ -approximations for all concave functions were previously not known to exist regardless of computation time.

Keywords-approximation algorithms; network design; buy-at-bulk;

I. INTRODUCTION

Many natural network design settings exhibit some form of economies of scale that reduce the costs when many flows are aggregated together. We may benefit from cheaper bandwidth when laying high capacity network links [1], reduced infrastructure costs or bulk discounts for shipping large amounts of goods together [2], or summarization and compression of correlated information flows [3]. These scenarios are known in the literature as buy-at-bulk problems. In a general buy-at-bulk problem we are given a graph and a set of demands for flow between nodes. The cost per unit length for routing a total of x flow along an edge is $f(x)$ for some function f . To model the economies of scale, we assume f is concave and monotone non-decreasing.

We will focus on single-sink (or single-source) case, where all demands must be routed to a given root. When f is known, the problem becomes the well-studied single-sink buy-at-bulk (SSBaB) problem. SSBaB is NP -hard—it generalizes the Steiner tree problem—but constant-factor approximations are known for any given f (e.g. [4], [5]). The special case where f has the form $f(x) = \min\{x, M\}$ for some M (edges can be “rented” for linear cost or “bought” for a fixed cost) is known as the single-sink rent-or-buy

(SSRoB) problem and has also received significant attention (e.g. [6], [7]).

Buy-at-bulk algorithms produce trees that are heavily tailored to the particular function at hand, but in some scenarios f may be unknown or known to change over time. One setting where this arises is in aggregation of data in sensor networks. The degree of redundancy among different sensor measurements may be unknown, or the same network may be used for aggregating different types of information with different amounts of redundancy. In other situations rapid technological advancement may cause bandwidth costs to change drastically over time. Further, in the interest of simplifying the design process and building robust networks, it may be useful to decouple the problem of designing the network topology from that of determining the exact characteristics of the information or goods flowing through that network. In these settings it is desirable to find a single tree that is simultaneously good for all cost-functions, and from a theoretical perspective, the existence of such trees would reveal surprising structure in the problem.

There are two natural objective functions which capture the idea of simultaneous approximation for multiple cost functions. Let \mathcal{F} be the set of all concave, monotone non-decreasing cost functions satisfying $f(0) = 0$, $f(T)$ be the cost of a routing tree T under function f , and T_f^* be the optimal routing graph for f . Note that due to the concavity of f we may assume that T_f^* is a tree. Let \mathcal{R} be a randomized algorithm that returns a feasible routing tree T . First, we could try to minimize the quantity

$$\sup_{f \in \mathcal{F}} \frac{\mathbf{E}_{\mathcal{R}}[f(T)]}{f(T_f^*)} \quad (1)$$

which we call the *oblivious* approximation ratio. If the oblivious ratio is small, then \mathcal{R} returns a distribution that works well in expectation for any f . However, there may be no sample from this distribution that works for everything: for any tree T there may be functions for which $f(T)$ is expensive.

To circumvent this problem we can work with the much stronger *simultaneous* approximation ratio. For a deterministic or randomized algorithm \mathcal{A} that returns a tree $T_{\mathcal{A}}$ the

*Research funded by an NSF IIS grant, by funds from Google, Microsoft, and Cisco, a 3-COM faculty fellowship, and by the ARL Network Science CTA.

†Research funded by an NSF IIS grant.

simultaneous ratio of \mathcal{A} is defined as:

$$\mathbf{E}_{\mathcal{A}} \left[\sup_{f \in \mathcal{F}} \frac{f(T_{\mathcal{A}})}{f(T_f^*)} \right] \quad (2)$$

A bound on the simultaneous ratio subsumes one on the oblivious ratio and proves there exists a single tree that is simultaneously good for all f .

We emphasize that the distinction between the simultaneous and oblivious objectives is not a technicality in the objective but rather a fundamental difference and that the gap between these ratios can be large. Consider the problem of embedding arbitrary metrics into tree metrics, another problem that requires bounding the cost under many different functions (i.e. distortion of each edge). It is well-known that distributions over trees can achieve $O(\log n)$ expected distortion for all edges [8] but that even for simple graphs like the n -cycle no single tree can do better than $\Omega(n)$ distortion [9]. Therefore, the ratio between maximum expected distortion and the expected maximum distortion is $\Omega(n/\log n)$ in this case.

Goel and Estrin [10] introduced the problem of simultaneous SSBaB and gave an algorithm with an $O(\log D)$ bound on the simultaneous ratio (2), where D is the total amount of demand. Goel and Post [11] recently improved the oblivious ratio (1) to $O(1)$ for a large constant. Trees for which the simultaneous ratio was $O(1)$ were not known to exist regardless of computation time.

In this paper we give the first constant guarantee on the simultaneous ratio, resolving the major open question of Goel and Estrin and Goel and Post [10], [11]. Several aspects of our algorithm and analysis bear mentioning:

- We achieve a simultaneous approximation ratio of 47.45. This is within a factor of 2.33 of the current best approximation for normal SSBaB of 20.42 [5] and substantially smaller than the $O(1)$ oblivious bound [11], which we estimate to be around 15 million.
- The algorithm is entirely combinatorial, in contrast to the result of Goel and Post [11], which uses separation oracles to prove an $O(1)$ -oblivious distribution exists but reveals little of its structure.
- Our analysis is short and simple, no more complex than the analysis of a normal SSBaB algorithm.
- The runtime is only $O((t(n, m) + m + n \log n) \log D)$ for a graph with n nodes, m edges, and D demand, where $t(n, m)$ is the runtime of an SSRoB approximation.

The algorithm is quite simple. We first find approximate trees for a set of rent or buy basis functions, prune this set to obtain a subset L of trees whose total rent costs are increasing geometrically while total buy costs are dropping geometrically, and then prove it suffices to approximate every tree in L . The set of bought nodes for each tree in L defines a series of tree layers, which we stitch together

using light approximate shortest-path trees (LASTs) [12] to approximate both the minimum spanning tree (MST) and shortest-path tree. Finally, we consider any layer in the tree. Using the geometrically changing costs and the properties of the LAST construction, we conclude that everything within the layer is an approximate MST, and everything outside approximates the shortest-path tree cost.

A. Related Work

The SSBaB problem was first posed by Salman et al. [2], and the first general approximation algorithm was given by Awerbuch and Azar [13], who used metric tree embeddings [14] to achieve an $O(\log^2 n)$ ratio, later improved to $O(\log n)$ using better embeddings by Bartal [15] and Fakcharoenphol et al. [8]. Guha et al. [4] gave the first constant approximation, and follow-up work by Talwar [16], Gupta et al. [17], Jothi and Raghavachari [18], Grandoni and Italiano [19], and Grandoni and Rothvoß [5] has since reduced the constant to 20.42. Most recent algorithms for SSBaB (and several related problems) are based on the sample and augment framework of Gupta et al. [17]. Many algorithms using this framework have been derandomized by van Zuylen [20].

The special case of SSRoB has also been extensively studied, often as a special case of the connected facility location problem. The first constant factor approximation was given by Ravi and Salman [21] as a special case of the traveling purchaser problem. Karger and Minkoff [6] gave an alternate algorithm and introduced connected facility location. Gupta et al. improved the approximation to 9.01 [22], Swamy and Kumar to 4.55 [23], and Gupta et al. to 3.55 [17]. Gupta et al. [24] derandomized the 3.55-approximation to achieve a 4.2-approximation. Eisenbrand et al. [7] developed a randomized 2.92-approximation, which recently improved to 2.8 using the 1.39-approximation for Steiner tree of Byrka et al. [25]. Both Williamson and van Zuylen [26] and Eisenbrand et al. [7] independently derandomized the 2.92-approximation to achieve a deterministic 3.28-approximation.

The problem of simultaneous approximation for multiple cost functions has been studied using both the oblivious and simultaneous objectives. Goel and Estrin [10] were the first to explicitly pose the question of simultaneous approximations and gave an algorithm with an $O(\log D)$ simultaneous guarantee. Prior to that Khuller et al. [12] gave an algorithm to simultaneously approximate the two extreme cost functions $f(x) = 1$ and $f(x) = x$ —a result which plays an important role in this paper—and metric tree embeddings had been applied to SSBaB [13], [8] to achieve an $O(\log n)$ bound for the oblivious objective. Enachescu et al. [27] gave an $O(1)$ simultaneous guarantee for the special case of grid graphs with some spatial correlation. Goel and Post [11] proved that an oblivious guarantee of $O(1)$ is achievable for all graphs. Gupta et al. [28] and Englert and Räcke [29] have studied several generalizations of the problem where

both the demands and function are unknown, and multiple sinks are allowed. In these settings the guarantee is generally $O(\log n)$ or $O(\text{polylog } n)$.

II. NOTATION AND PRELIMINARIES

Formally, we are given a graph $G = (V, E)$ with edge lengths l_e for $e \in E$, a root node r , and a set of demand nodes $\mathcal{D} \subseteq V$ with integer demands d_v . The total demand is $D = \sum_v d_v$. We want to route d_v flow from each v to r as cheaply as possible, where the cost of routing x_e flow along edge e is $l_e f(x_e)$ for some unknown, concave, monotone increasing function f satisfying $f(0) = 0$. Not knowing f , our objective is to find a feasible tree T minimizing $\sup_f f(T)/f(T_f^*)$, where T_f^* is the optimal graph for f .

We first show that we can restrict our analysis to a smaller class of basis functions. Let $\epsilon > 0$ be a small constant which will trade off the runtime and the approximation ratio, and $K = \lceil \log_{1+\epsilon} D \rceil$. For $0 \leq i \leq K$, define $M_i = (1 + \epsilon)^i$, $A_i(x) = \min\{x, M_i\}$, and T_i^* as the optimal tree for A_i . By the monotonicity and concavity of f , whenever $M_i \leq x \leq M_{i+1}$ we have $f(M_i) \leq f(x) \leq f(M_{i+1}) \leq (1 + \epsilon)f(M_i)$, so with a loss of only a factor of $1 + \epsilon$ we can interpolate between $f(M_i)$ and $f(M_{i+1})$ and assume f is piecewise linear with breakpoints only at powers of $1 + \epsilon$. A nondecreasing concave function that is linear between powers of $1 + \epsilon$ can be written as a nonnegative linear combination of $\{A_i\}_{0 \leq i \leq K}$ by setting coefficients equal to the changes in slope: if the slope drops from δ_i to δ_{i+1} at $(1 + \epsilon)^i$ it induces the term $(\delta_i - \delta_{i+1})A_i(x)$. Now for a linear combination $\sum_i a_i A_i(x)$ and a tree T

$$\begin{aligned} \frac{\sum_i a_i A_i(T)}{\sum_i a_i A_i(T_f^*)} &\leq \frac{\sum_i a_i A_i(T)}{\sum_i a_i A_i(T_i^*)} \\ &= \frac{\sum_i a_i A_i(T_i^*) \frac{A_i(T)}{A_i(T_i^*)}}{\sum_i a_i A_i(T_i^*)} \leq \max_i \frac{A_i(T)}{A_i(T_i^*)} \end{aligned}$$

so it suffices to upper bound $\max_i A_i(T)/A_i(T_i^*)$.

We now define some notation and subroutines that will be important for our algorithm. The problem of finding a good aggregation tree for the function $A_i(x) = \min\{x, M_i\}$ is an instance of the SSRoB problem, and we can find a λ -approximate tree T_i , where λ is the best approximation ratio known, currently equal to 2.8 using the algorithm of Eisenbrand et al. [7] and Byrka et al. [25]. We will assume the algorithm is deterministic. If not (as in the case of the 2.8-approximation) we repeat it a polynomial number of times and pick the best tree, so we are close to a λ -approximation with very high probability. In this case, our simultaneous approximation algorithm will have some tiny probability of failure.

The cost $A_i(T_i)$ can be broken into two pieces, the rent cost and the buy cost, based on whether A_i is maxed out at M_i :

Definition 2.1: For an aggregation tree T_i for cost function A_i with x_e flow on edge e , the *rent cost* R_i and *normalized buy cost* B_i are defined as

$$\begin{aligned} R_i &= \sum_{e \in T_i, x_e < M_i} l_e A_i(x_e) \\ B_i &= \sum_{e \in T_i, x_e \geq M_i} l_e \frac{A_i(x_e)}{M_i} = \sum_{e \in T_i, x_e \geq M_i} l_e \end{aligned}$$

Note that edge costs composing R_i are weighted by the amount of flow they carry, but edges in B_i are not; they use unweighted edge costs. The total cost of T_i is given by $A_i(T_i) = R_i + M_i B_i$. The rent and buy costs also partition the nodes of T_i into two sets:

Definition 2.2: The *core* C_i of tree T_i consists of r and all nodes spanned by bought edges and the *periphery* contains all vertices outside C_i .

If we condition on the nodes in C_i then the rent-or-buy problem becomes easy: demands outside the core pay linear cost until they reach C_i , so they should take the shortest path, whereas within C_i we pay a fixed cost per edge length, so the best strategy is to follow the min spanning tree. The cost R_i is therefore at least the sum of shortest path distances to C_i , while B_i is at least the weight of the MST of C_i .

In addition to the SSRoB approximation, we will also employ the light, approximate shortest-path tree algorithm of Khuller et al. [12]:

Definition 2.3 ([12]): For $\alpha \geq 1$ and $\beta \geq 1$, an (α, β) -*light, approximate shortest-path* or (α, β) -*LAST* is a spanning tree T of G with root r such that

- For each vertex v , the distance from v to r in T is at most α times the shortest path distance from v to r in G .
- The edge weight of T is at most β times the weight of an MST of G .

Khuller et al. show how to construct an (α, β) -LAST for any $\alpha > 1$ and $\beta \geq \frac{\alpha+1}{\alpha-1}$. Roughly, the algorithm performs a depth-first traversal of the MST of G starting from r , checking the stretch of the shortest path to each node. If the path to some v has blown up by at least an α factor, then it updates the tree to take the shortest path from v to r , adjusting other tree edges and distances accordingly. See the paper [12] for a full description and analysis.

Finally, we define four parameters α , β , γ , and δ used by our algorithm whose values we will optimize at the end.

- $\alpha > 1$ is the approximation ratio for shortest paths used in our LAST.
- $\beta \geq \frac{\alpha+1}{\alpha-1}$ is the corresponding approximation to the MST in the LAST.
- $\gamma > 1$ is the factor by which normalized buy costs B_i increase from layer to layer in our tree.
- $\delta > 1$ is the factor by which rent costs R_i drop from layer to layer.

We now turn to a more thorough explanation of tree layers.

III. TREE LAYERS

In the normal SSBaB problem, the cost function is defined as $f(x) = \min_j \{\sigma_j + \delta_j x\}$, the cheapest of a collection of different “pipes” or “cables” given to the algorithm, each with an affine cost function $\sigma_j + \delta_j x$. It is common (e.g. [4], [17]) to first prune these pipes to a smaller set with geometrically decreasing δ_j 's and geometrically increasing σ_j 's and then build a solution in layers where each layer routes with a different pipe.

We perform an analogous procedure. We would like to build our simultaneous tree T in a series of nested layers defined by the cores C_i of each tree T_i , so that the core of T under A_i is similar to C_i , but we have no guarantees on the relationships between different cores: C_i and C_k may be entirely disjoint except for r . However, we will show that as long as normalized buy costs B_i and B_k are within a constant factor of each other, the same core can be used for both trees. Consequently, we are able to define nested layers by choosing one C_i for each order of magnitude of B_i .

After finding λ -approximate trees T_i for each A_i , we loop through the costs B_i and R_i , discarding i whenever B_i does not drop by γ or R_i does not grow by δ . We are left with a subset L of the T_i where the B_i 's are dropping by a factor of γ and the R_i 's are growing by a factor of δ . The cores C_i for each $i \in L$ will define the layers of our tree. Algorithm 1 describes the procedure in more detail.

Intuitively, as it becomes more expensive to buy edges the optimum will buy fewer edges and rent more. In the case of approximations, the progression becomes muddled because for some i the approximation guarantee may be tight while for $i+1$ we may get lucky and find the optimum, resulting in both rent and normalized buy costs dropping. We first show that the monotonicity in buy and rent costs still holds as long as each T_i is better for A_i than both T_{i-1} and T_{i+1} .

Lemma 3.1: After line 8 of Algorithm 1, for every i we have $B_i \geq B_{i+1}$ and $R_i \leq R_{i+1}$.

Proof: First we show that for each i , $A_i(T_i) \leq \min\{A_i(T_{i+1}), A_i(T_{i-1})\}$. After the loop on lines 3–4, we have $A_i(T_i) \leq A_i(T_{i-1})$, and after the second loop on lines 5–6 we have $A_i(T_i) \leq A_i(T_{i+1})$, so we only need to show that the second loop does not break the first condition. If the second loop updates T_i then $A_i(T_i)$ will only shrink, and if it changes T_{i-1} it does this by setting $T_{i-1} \leftarrow T_i$ which preserves $A_i(T_i) \leq A_i(T_{i-1})$.

Now consider $A_i(T_k)$ for any k . By definition $A_i(x) \leq x$ and $A_i(x) \leq M_i$, so to upper bound $A_i(T_k)$ we may assume edges within C_k pay M_i per unit length, which sums to $M_i B_k$, and edges outside C_k pay linear cost, or R_k total, implying $A_i(T_k) \leq R_k + M_i B_k$. Therefore

$$\begin{aligned} R_i + M_i B_i &= A_i(T_i) \leq A_i(T_{i+1}) \leq R_{i+1} + M_i B_{i+1} \\ \implies M_i(B_i - B_{i+1}) &\leq R_{i+1} - R_i \end{aligned}$$

Input: Graph G and demands \mathcal{D}

Output: Set L and cores C_i for each $i \in L$

```

1: for  $i \leftarrow 0$  to  $K$  do
2:    $T_i \leftarrow \lambda$ -approximate tree for  $A_i(x)$ 
3: for  $i \leftarrow 1$  to  $K$  do
4:   if  $A_i(T_{i-1}) < A_i(T_i)$  then  $T_i \leftarrow T_{i-1}$ 
5: for  $i \leftarrow K - 1$  down to  $0$  do
6:   if  $A_i(T_{i+1}) < A_i(T_i)$  then  $T_i \leftarrow T_{i+1}$ 
7: for  $i \leftarrow 0$  to  $K$  do
8:   calculate  $C_i, B_i, R_i$ 

9:  $L_B \leftarrow \emptyset$ 
10:  $B \leftarrow \infty$ 
11: for  $i \leftarrow 0$  to  $K$  do
12:   if  $B_i < \frac{1}{\gamma} B$  then
13:      $L_B \leftarrow L_B \cup \{i\}$ 
14:      $B \leftarrow B_i$ 
15:  $L \leftarrow \emptyset$ 
16:  $R \leftarrow \infty$ 
17: foreach  $i \in L_B$  in decreasing order do
18:   if  $R_i < \frac{1}{\delta} R$  then
19:      $L \leftarrow L \cup \{i\}$ 
20:      $R \leftarrow R_i$ 
21: return  $L$  and  $C_i$  for each  $i \in L$ 

```

Algorithm 1. Finding tree layers

Similarly,

$$\begin{aligned} R_{i+1} + M_{i+1} B_{i+1} &= A_{i+1}(T_{i+1}) \\ &\leq A_{i+1}(T_i) \leq R_i + M_{i+1} B_i \\ \implies R_{i+1} - R_i &\leq M_{i+1}(B_i - B_{i+1}) \end{aligned}$$

Combining the inequalities,

$$M_i(B_i - B_{i+1}) \leq R_{i+1} - R_i \leq M_{i+1}(B_i - B_{i+1})$$

If $B_i - B_{i+1} < 0$ the inequality is false because $M_{i+1} > M_i$, so we conclude $B_i \geq B_{i+1}$. And using the first inequality, $0 \leq M_i(B_i - B_{i+1}) \leq R_{i+1} - R_i$, so $R_i \leq R_{i+1}$. ■

We need to show that we can restrict our attention to T_i for $i \in L$. Suppose $i < k$ but $B_i \leq \gamma B_k$. Using Lemma 3.1, observe that $A_k(T_i) \leq R_i + M_k B_i \leq R_k + \delta M_k B_k \leq \delta A_k(T_k)$. Note this is independent of the size of the intersection of the cores C_i and C_k and any differences in routing. The following lemma generalizes this simple but key observation and proves that approximating each $i \in L$ is sufficient.

Lemma 3.2: Suppose there exists a tree T and constants c_B and c_R such that for all $i \in L$ there exists a partition of the edges of T into two sets T_{B_i} and T_{R_i} satisfying $A_0(T_{B_i}) \leq c_B B_i$ and $A_K(T_{R_i}) \leq c_R R_i$, then for all $k \in \{0, \dots, K\}$, $A_k(T) \leq \max\{c_B \gamma, c_R \delta\} \lambda A_k(T_k^*)$.

Proof: Let $k \in \{0, \dots, K\}$. Let $j = \max\{j \in L_B \mid j \leq k\}$. Either $j = k$ or k was discarded due to j on lines 11–14 because $B_k \geq \frac{1}{\gamma}B_j$, and either way $B_j \leq \gamma B_k$. Now let $i = \min\{i \in L \mid i \geq j\}$. Again $i = j$ or j was pruned due to i on lines 17–20, and $R_i \leq \delta R_j$. Applying Lemma 3.1 with $i \geq j$ and $j \leq k$, we have $B_i \leq B_j \leq \gamma B_k$ and $R_i \leq \delta R_j \leq \delta R_k$.

This is sufficient to bound the cost of $A_k(T)$:

$$\begin{aligned} A_k(T) &= A_k(T_{R_i}) + A_k(T_{B_i}) \leq A_K(T_{R_i}) + M_k A_0(T_{B_i}) \\ &\leq c_R R_i + c_B M_k B_i \\ &\leq c_R \delta R_k + c_B \gamma M_k B_k \\ &\leq \max\{c_R \delta, c_B \gamma\} A_k(T_k) \\ &\leq \max\{c_R \delta, c_B \gamma\} \lambda A_k(T_k^*) \end{aligned}$$

The equality follows because T_{B_i} and T_{R_i} partition the edges of T . The first inequality is because $A_K(x)$ and $M_k A_0(x)$ both upper bound $A_k(x)$, the second is by assumption, the third is from the derivation above, the fourth uses $A_k(T_k) = R_k + M_k B_k$, and the last is because T_k is a λ -approximation. ■

We will primarily assume that our SSRoB algorithm is a generic approximation, but Lemma 3.2 can easily be improved to take advantage of an SSRoB algorithm with a stronger guarantee that separately bounds R_i and $M_i B_i$ in terms of the optimal costs R_i^* and $M_i B_i^*$.

Corollary 3.3: Let T , c_B , and c_R be as in Lemma 3.2, and suppose

- $R_i \leq \mu_R R_i^* + \mu_B M_i B_i^*$
- $M_i B_i \leq \nu_R R_i^* + \nu_B M_i B_i^*$

then for all k ,

$$A_k(T) \leq \max\{c_R \delta \mu_R + c_B \gamma \nu_R, c_R \delta \mu_B + c_B \gamma \nu_B\} A_k(T_k^*).$$

Proof: We change the inequalities in the proof above as follows:

$$\begin{aligned} A_k(T) &\leq c_R \delta R_k + c_B \gamma M_k B_k \\ &\leq c_R \delta (\mu_R R_k^* + \mu_B M_k B_k^*) + c_B \gamma (\nu_R R_k^* + \nu_B M_k B_k^*) \\ &\leq \max\{c_R \delta \mu_R + c_B \gamma \nu_R, c_R \delta \mu_B + c_B \gamma \nu_B\} A_k(T_k^*) \end{aligned}$$

IV. CONSTRUCTING THE TREE

The construction of the tree itself is quite simple. We have a set of indices L and core sets C_i for $i \in L$. Starting with the largest $i \in L$, i.e. smallest B_i , and working downward, we connect each C_i to T , the tree so far, with a LAST. Algorithm 2 describes the procedure more formally. The notation G/T represents contracting T to a single node in G , and $G[C_i]$ is the induced subgraph on C_i , so $(G/T)[C_i]$ denotes first contracting T and then restricting to nodes in C_i .

Lemma 4.1: The graph T constructed by Algorithm 2 is a tree and spans all demand nodes.

Input: Graph G . Set L and accompanying C_i for each $i \in L$

Output: Aggregation tree T

- 1: $T \leftarrow \{r\}$
- 2: **foreach** $i \in L$ in decreasing order **do**
- 3: $T' \leftarrow (\alpha, \beta)$ -LAST of $(G/T)[C_i]$ with root T
- 4: $T \leftarrow T \cup T'$
- 5: **return** T

Algorithm 2. Constructing the tree

Proof: Observe that $0 \in L$, and $R_0 = 0$, so C_0 covers all demands. Therefore after the last iteration T spans \mathcal{D} . Each iteration only adds edges spanning new vertices, so no cycles are created. ■

The tree may contain paths connecting Steiner nodes that carry no flow. Such edges can be safely pruned or just ignored because they contribute nothing to the cost.

All that remains is to define the partitions T_{B_i} and T_{R_i} and prove the bounds needed in Lemma 3.2. The set T_{B_i} contains all edges present in T after connecting C_i , and T_{R_i} contains the rest. Both cost bounds will follow easily from the geometrically changing costs: the cost $A_0(T_{B_i})$ is dominated by the cost of the C_i layer, an approximate MST, and $A_K(T_{R_i})$ is dominated by the rent costs of the next layer, an approximate shortest-path tree. First, we bound the normalized buy cost of T_{B_i} :

Lemma 4.2: Let $i \in L$, and T_{B_i} be the tree T after the round when C_i is added. Then the edge cost $A_0(T_{B_i})$ is at most $\frac{\beta\gamma}{\gamma-1} B_i$.

Proof: The proof is by decreasing induction on i , i.e. in the order in which the layers are built. Let c be a constant to be chosen at the end. The base case is the largest $i \in L$, which is the smallest i such that $B_i = 0$. In this case, $C_i = \{r\}$, $T_{B_i} = \{r\}$, and the edge cost is 0.

Now let $i \in L$, $k = \min\{k \in L \mid k > i\}$ be the previous (inner) layer, and suppose the edge cost of T_{B_k} is at most $c B_k$. By the construction of L , we know $B_k < \frac{1}{\gamma} B_i$, implying T_{B_k} costs at most $\frac{c}{\gamma} B_i$. The cost of an MST of C_i in G is at most B_i , and T_{B_k} may already span part of C_i , so connecting the rest with an MST¹ of $(G/T_{B_k})[C_i]$ costs at most B_i . Using an (α, β) -LAST scales the cost by at most β .

The total cost of edges laid so far is at most $\frac{c}{\gamma} B_i + \beta B_i$, so the proof is complete as long as $c B_i \geq \frac{c}{\gamma} B_i + \beta B_i$. Set $c = \frac{\beta\gamma}{\gamma-1}$:

$$c \geq \beta + \frac{c}{\gamma} \iff c \left(1 - \frac{1}{\gamma}\right) \geq \beta \iff c \geq \frac{\beta\gamma}{\gamma-1}$$

Now we bound the rent costs of T_{R_i} by a similar proof. ■

¹We could allow Steiner nodes and use a Steiner tree approximation, but this would not improve the worst-case bound.

Lemma 4.3: Let $i \in L$ and $T_{R_i} = T/T_{B_i}$, i.e. all edges outside of T_{B_i} . Then the rent cost $A_K(T_{R_i})$ is at most $\frac{\alpha\delta}{\delta-\alpha-1}R_i$.

Proof: We prove by increasing induction on $i \in L$ (the reverse of Lemma 4.2) that $A_K(T_{R_i}) \leq cR_i$ for some c to be determined. Since $0 \in L$, and T_{B_0} covers everything, the base case T_{R_0} costs 0 too.

For the inductive case, let $i \in L$, $k = \max\{k \in L | k < i\}$ be the next (outer) layer, and T_{R_k} have rent cost at most cR_k . As before, note $R_k < \frac{1}{\delta}R_i$, so $A_K(T_{R_k}) \leq \frac{c}{\delta}R_i$. Tree T_{B_i} spans C_i and possibly more, so if all demands outside T_{B_i} took the shortest path from their sources to T_{B_i} the shortest-path cost would be at most R_i . However, the edges of T_{R_k} have moved some demands around, and by the time they reach the current layer they may be farther from T_{B_i} than their original sources were. But by the triangle inequality the cost of sending all demands from their current locations to T_{B_i} via shortest paths is at most $\frac{c}{\delta}R_i + R_i$, the cost of sending all flow in T_{R_k} back to its source and from there to T_{B_i} using shortest paths. The LAST algorithm guarantees α -approximate shortest paths, multiplying the cost by α .

Consequently the total rent cost for T_{R_i} is bounded by $\frac{c}{\delta}R_i + \alpha(\frac{c}{\delta}R_i + R_i)$, which needs to be at most cR_i . We can set $c = \frac{\alpha\delta}{\delta-\alpha-1}$:

$$c \geq \frac{c}{\delta} + \frac{c\alpha}{\delta} + \alpha \iff c \left(1 - \frac{1}{\delta} - \frac{\alpha}{\delta}\right) = c \frac{\delta - \alpha - 1}{\delta} \geq \alpha$$

$$\iff c \geq \frac{\alpha\delta}{\delta - \alpha - 1}$$

■

We note that Lemma 4.3 explains how we circumvent a major obstacle to an $O(1)$ -simultaneous approximation—the $\Omega(\log n)$ distortion lower bound for embedding arbitrary metrics into tree metrics [14]. If we needed to maintain distances between many pairs of nodes the task would be hopeless, but Lemma 4.3 shows that it suffices to preserve the distance of each node to the next layer, so the graph of distances to be maintained forms a tree.

We can now complete the proof of our main theorem and choose the optimal parameters.

Theorem 4.4: The tree T achieves a simultaneous approximation ratio of $(1 + \epsilon)\lambda(8 + 4\sqrt{5})$ using a λ -approximation to SSRoB. In particular,

- There is a randomized polynomial time algorithm that finds a 47.45 simultaneous approximation with high probability.
- There is a deterministic polynomial time algorithm that finds a 55.58 simultaneous approximation.
- There exists a tree that is a 16.95 simultaneous approximation.

Proof: Applying Lemma 3.2 with $c_B = \frac{\beta\gamma}{\gamma-1}$ (Lemma 4.2) and $c_R = \frac{\alpha\delta}{\delta-\alpha-1}$ (Lemma 4.3), the final approximation

ratio for an arbitrary cost function f is

$$(1 + \epsilon)\lambda \max \left\{ \frac{\beta\gamma^2}{\gamma-1}, \frac{\alpha\delta^2}{\delta-\alpha-1} \right\} \quad (3)$$

where the extra $1 + \epsilon$ comes from the approximation of f by a combination of A_i 's. Now it is a simple matter of applying calculus to find the optimal values for α , β , γ , and δ .

The easiest parameters to fix are γ and δ . For γ :

$$\frac{d}{d\gamma} \left[\frac{\beta\gamma^2}{\gamma-1} \right] = \beta \frac{(2\gamma)(\gamma-1) - \gamma^2}{(\gamma-1)^2} = 0$$

$$\implies \gamma(\gamma-2) = 0 \implies \gamma = 2$$

For δ :

$$\frac{d}{d\delta} \left[\frac{\alpha\delta^2}{\delta-\alpha-1} \right] = \alpha \frac{2\delta(\delta-\alpha-1) - \delta^2}{(\delta-\alpha-1)^2} = 0$$

$$\implies \delta^2 - 2\alpha\delta - 2\delta = \delta(\delta - 2\alpha - 2) = 0$$

$$\implies \delta = 2\alpha + 2$$

Plugging $\gamma = 2$ and $\delta = 2\alpha + 2$ into (3), $\beta \frac{\gamma^2}{\gamma-1} = 4\beta$, and

$$\frac{\alpha\delta^2}{\delta-\alpha-1} = \frac{\alpha(2\alpha+2)^2}{(2\alpha+2)-\alpha-1} = 4\alpha(\alpha+1)$$

so (3) is now $\max\{4\beta, 4\alpha(\alpha+1)\}$. The constraints on α and β require $\beta \geq \frac{\alpha+1}{\alpha-1}$ [12], so one term blows up if the other shrinks. To minimize the maximum set the two expressions to be equal:

$$\beta = \frac{\alpha+1}{\alpha-1} = \alpha(\alpha+1)$$

$$\implies \alpha(\alpha-1) = 1 \implies \alpha^2 - \alpha - 1 = 0$$

$$\implies \alpha = \frac{1 \pm \sqrt{5}}{2}$$

Using $\alpha = \frac{1+\sqrt{5}}{2}$, we get

$$\beta = \frac{\alpha+1}{\alpha-1} = \frac{3+\sqrt{5}}{-1+\sqrt{5}} = \frac{(3+\sqrt{5})(1+\sqrt{5})}{4}$$

$$= \frac{8+4\sqrt{5}}{4} = 2 + \sqrt{5}$$

and $\delta = 2\alpha + 2 = 3 + \sqrt{5}$.

Finally, $\frac{\beta\gamma^2}{\gamma-1} = 4(2 + \sqrt{5})$ so the simultaneous approximation ratio is

$$(1 + \epsilon)\lambda(8 + 4\sqrt{5}).$$

Now,

- Using the best randomized approximation $\lambda = 2.8$, and the ratio is 47.45 with high probability.
- Using the best deterministic approximation $\lambda = 3.28$, and the ratio is 55.58.
- If the algorithm is allowed to run in exponential time $\lambda = 1$, and the ratio is 16.95.

■

The 2.8-approximation of Eisenbrand et al. [7] actually provides a slightly stronger guarantee on R_i and B_i , and we can use Corollary 3.3 to get a tiny improvement in the approximation ratio at the cost of a more complex derivation:

Theorem 4.5: There is a randomized polynomial time algorithm that finds a 47.07 simultaneous approximation with high probability.

Proof: Lemma 2 and Theorem 5 in [7] prove that

$$\mathbf{E}[R_i] \leq 2R_i^* + \frac{.807}{x} M_i B_i^*$$

$$\mathbf{E}[M_i B_i] \leq \rho(x + \epsilon) R_i^* + \rho M_i B_i^*$$

where $\rho = 1.39$ is the Steiner tree approximation ratio and $x \in (0, 1]$ is a parameter. Applying Corollary 3.3 with

$$\mu_R = 2 \quad \mu_B = \frac{.807}{x} \quad \nu_R = \rho(x + \epsilon) \quad \nu_B = \rho$$

the simultaneous ratio is bounded by

$$(1 + \epsilon) \max \left\{ \frac{2\alpha\delta^2}{\delta - \alpha - 1} + \frac{\rho(x + \epsilon)\beta\gamma^2}{\gamma - 1}, \frac{.807\alpha\delta^2}{x(\delta - \alpha - 1)} + \frac{\rho\beta\gamma^2}{\gamma - 1} \right\}$$

Using

$$x = .5995 \quad \alpha = 1.5495 \quad \beta = \frac{\alpha + 1}{\alpha - 1} \quad \gamma = 2 \quad \delta = 2\alpha + 2$$

yields a simultaneous ratio of 47.07. We omit the tedious but straightforward derivation of these values. ■

We leave as an open question the problem of exploiting Corollary 3.3 to substantially improve the ratio.

A. Runtime

Let $t(n, m)$ be the running time of our SSRoB approximation on a graph with n vertices and m edges, which must be at least $\Omega(n)$ to write down the output. When ϵ is constant, running the SSRoB approximation for each i takes $O(t(n, m) \log D)$. Subsequent loops in Algorithm 1 take $O(n \log D) = O(t(n, m) \log D)$.

For each of the $O(\log D)$ iterations of Algorithm 2 we need to do a graph contraction and run the LAST algorithm, which requires computing the MST and shortest path trees. The computation of the shortest path tree takes $O(m + n \log n)$ and dominates the other steps. Combining the two algorithms, the total time is $O((t(n, m) + m + n \log n) \log D)$.

V. OPEN PROBLEMS

We have answered the open questions posed by Goel and Estrin and Goel and Post [10], [11], but there are several avenues for further work. Our simultaneous ratio of 47.45 already surpasses many algorithms for normal SSBaB and is only a factor of 2.33 away from the best. It would be nice to eliminate this gap or, alternately, prove that a gap exists between the approximation achievable for fixed f and the

best simultaneous ratio. We know of no lower bounds on what simultaneous ratio may be possible, so any progress in this direction would also be interesting. Generalizing the settings in which $O(1)$ simultaneous ratios are possible would be interesting, but may be unlikely given that one must contend with lower bounds for metric tree embedding [14] and multi-sink buy-at-bulk [30].

ACKNOWLEDGMENTS

We thank the anonymous reviewers for many helpful comments that improved the presentation and for suggesting Corollary 3.3 and Theorem 4.5.

REFERENCES

- [1] M. Andrews and L. Zhang, “The access network design problem,” in *Proceedings of the 39th IEEE Symposium on Foundations of Computer Science*, 1998, pp. 40–49. [Online]. Available: <http://dx.doi.org/10.1109/SFCS.1998.743427>
- [2] F. Salman, J. Cheriyan, R. Ravi, and S. Subramanian, “Buy-at-bulk network design: Approximating the single-sink edge installation problem,” in *Proceedings of the 8th annual ACM-SIAM Symposium on Discrete Algorithms*, 1997, pp. 619–628. [Online]. Available: <http://doi.acm.org/10.1145/314161.314397>
- [3] B. Krishnamachari, D. Estrin, and S. Wicker, “Modelling data-centric routing in wireless sensor networks,” in *IEEE INFOCOM*, 2002, pp. 1–11. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.11.7212>
- [4] S. Guha, A. Meyerson, and K. Munagala, “A constant factor approximation for the single sink edge installation problems,” in *Proceedings of the 33rd annual ACM Symposium on Theory of Computing*, 2001, pp. 383–388. [Online]. Available: <http://dx.doi.org/10.1145/380752.380827>
- [5] F. Grandoni and T. Rothvoß, “Network design via core detouring for problems without a core,” in *37th International Colloquium on Automata, Languages and Programming. To appear.*, 2010. [Online]. Available: <http://www.disp.uniroma2.it/users/grandoni/Pubblicazioni/GR10icalp.pdf>
- [6] D. Karger and M. Minkoff, “Building Steiner trees with incomplete global knowledge,” in *Proceedings 41st IEEE Symposium on Foundations of Computer Science*, 2000, pp. 613–623. [Online]. Available: <http://people.csail.mit.edu/karger/Papers/maybecast.ps>
- [7] F. Eisenbrand, F. Grandoni, T. Rothvoß, and G. Schäfer, “Connected facility location via random facility sampling and core detouring,” *Journal of Computer and System Sciences*, 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.jcss.2010.02.001>
- [8] J. Fakcharoenphol, S. Rao, and K. Talwar, “A tight bound on approximating arbitrary metrics by tree metrics,” in *Proceedings of the 35th annual ACM Symposium on Theory of Computing*, 2003, pp. 448–455. [Online]. Available: <http://dx.doi.org/10.1145/780542.780608>

- [9] Y. Rabinovich and R. Raz, “Lower bounds on the distortion of embedding finite metric spaces in graphs,” *Discrete and Computational Geometry*, vol. 19, no. 1, pp. 79–94, 1998. [Online]. Available: <http://dx.doi.org/10.1007/PL00009336>
- [10] A. Goel and D. Estrin, “Simultaneous optimization for concave costs: single sink aggregation or single source buy-at-bulk,” in *Proceedings of the 14th annual ACM-SIAM Symposium on Discrete Algorithms*, 2003, pp. 499–505. [Online]. Available: <http://doi.acm.org/10.1145/644108.644191>
- [11] A. Goel and I. Post, “An oblivious $O(1)$ -approximation for single source buy-at-bulk,” in *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science*, 2009. [Online]. Available: <http://arxiv.org/abs/0908.3740v1>
- [12] S. Khuller, B. Raghavachari, and N. Young, “Balancing minimum spanning trees and shortest-path trees,” *Algorithmica*, vol. 14, no. 4, pp. 305–321, 1995. [Online]. Available: <http://arxiv.org/abs/cs/0205045v1>
- [13] B. Awerbuch and Y. Azar, “Buy-at-bulk network design,” in *Proceedings of the 38th annual IEEE Symposium on Foundations of Computer Science*, 1997. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.46.6906>
- [14] Y. Bartal, “Probabilistic approximation of metric spaces and its algorithmic applications,” in *Proceedings of the 37th IEEE Symposium on Foundations of Computer Science*, 1996, pp. 184–193. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.31.8507>
- [15] —, “On approximating arbitrary metrics by tree metrics,” in *Proceedings of the 30th annual ACM Symposium on Theory of Computing*, 1998, pp. 161–168. [Online]. Available: <http://dx.doi.org/10.1145/276698.276725>
- [16] K. Talwar, “The single-sink buy-at-bulk lp has constant integrality gap,” in *Proceedings of the 9th International IPCO Conference on Integer Programming and Combinatorial Optimization*, 2002, pp. 475–486. [Online]. Available: http://dx.doi.org/10.1007/3-540-47867-1_33
- [17] A. Gupta, A. Kumar, M. Pal, and T. Roughgarden, “Approximation via cost sharing: Simpler and better approximation algorithms for network design,” *Journal of the ACM (JACM)*, vol. 54, no. 3, 2007. [Online]. Available: <http://dx.doi.org/10.1145/1236457.1236458>
- [18] R. Jothi and B. Raghavachari, “Improved approximation algorithms for the single-sink buy-at-bulk network design problems,” in *Proceedings of the 9th Scandinavian Workshop on Algorithm Theory*, 2004, pp. 336–348. [Online]. Available: <http://dx.doi.org/10.1016/j.jda.2008.12.003>
- [19] F. Grandoni and G. Italiano, “Improved approximation for single-sink buy-at-bulk,” *Lecture Notes in Computer Science*, vol. 4288, pp. 111–120, 2006. [Online]. Available: http://dx.doi.org/10.1007/11940128_13
- [20] A. van Zuylen, “Deterministic sampling algorithms for network design,” *Algorithmica*, pp. 1–42, 2009. [Online]. Available: <http://dx.doi.org/10.1007/s00453-009-9344-x>
- [21] R. Ravi and F. Salman, “Approximation algorithms for the traveling purchaser problem and its variants in network design,” *Lecture Notes in Computer Science*, vol. 1643, pp. 29–40, 1999. [Online]. Available: http://dx.doi.org/10.1007/3-540-48481-7_4
- [22] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener, “Provisioning a virtual private network: A network design problem for multicommodity flow,” in *Proceedings of the 33rd annual ACM Symposium on Theory of Computing*, 2001, pp. 389–398. [Online]. Available: <http://dx.doi.org/10.1145/380752.380830>
- [23] C. Swamy and A. Kumar, “Primal–dual algorithms for connected facility location problems,” *Algorithmica*, vol. 40, no. 4, pp. 245–269, 2004. [Online]. Available: <http://dx.doi.org/10.1007/s00453-004-1112-3>
- [24] A. Gupta, A. Srinivasan, and É. Tardos, “Cost-sharing mechanisms for network design,” *Algorithmica*, vol. 50, no. 1, pp. 98–119, 2008. [Online]. Available: <http://dx.doi.org/10.1007/s00453-007-9065-y>
- [25] J. Byrka, F. Grandoni, T. Rothvoß, and L. Sanità, “An improved LP-based approximation for Steiner tree,” in *Proceedings of the 42nd ACM Symposium on Theory of Computing*. ACM, 2010, pp. 583–592. [Online]. Available: <http://dx.doi.org/10.1145/1806689.1806769>
- [26] D. Williamson and A. van Zuylen, “A simpler and better derandomization of an approximation algorithm for single source rent-or-buy,” *Operations Research Letters*, vol. 35, no. 6, pp. 707–712, 2007. [Online]. Available: <http://dx.doi.org/10.1016/j.orl.2007.02.005>
- [27] M. Enachescu, A. Goel, R. Govindan, and R. Motwani, “Scale-free aggregation in sensor networks,” *Theoretical Computer Science*, vol. 344, no. 1, pp. 15–29, 2005. [Online]. Available: <http://dx.doi.org/10.1016/j.tcs.2005.06.023>
- [28] A. Gupta, M. Hajiaghayi, and H. Räcke, “Oblivious network design,” in *Proceedings of the 17th annual ACM-SIAM Symposium on Discrete Algorithms*, 2006, pp. 970–979. [Online]. Available: <http://dx.doi.org/10.1145/1109557.1109665>
- [29] M. Englert and H. Räcke, “Oblivious routing for the L_p -norm,” in *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science*, 2009, pp. 32–40. [Online]. Available: <http://www.dcs.warwick.ac.uk/~harry/pdf/obliviouslpnorm.pdf>
- [30] M. Andrews, “Hardness of buy-at-bulk network design,” in *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science*, 2004, pp. 115–124. [Online]. Available: <http://dx.doi.org/10.1109/FOCS.2004.32>