

Balls and Funnels: Energy Efficient Group-to-Group Anycasts

Jennifer Iglesias¹, Rajmohan Rajaraman², R. Ravi¹, and Ravi Sundaram²(✉)

¹ Carnegie Mellon University, Pittsburgh, PA, USA
{jiglesia,ravi}@andrew.cmu.edu

² Northeastern University, Boston, MA, USA
{rraj,koods}@ccs.neu.edu

Abstract. We introduce group-to-group anycast (g2g-anycast), a network design problem of substantial practical importance and considerable generality. Given a collection of groups and requirements for directed connectivity from source groups to destination groups, the solution network must contain, for each requirement, an omni-directional down-link broadcast, centered at any node of the source group, called the ball; the ball must contain some node from the destination group in the requirement and all such destination nodes in the ball must aggregate into a tree directed towards the source, called the funnel-tree. The solution network is a collection of balls along with the funnel-trees they contain. g2g-anycast models DBS (Digital Broadcast Satellite), Cable TV systems and drone swarms. It generalizes several well known network design problems including minimum energy unicast, multicast, broadcast, Steiner-tree, Steiner-forest and Group-Steiner tree. Our main achievement is an $O(\log^4 n)$ approximation, counterbalanced by an $\log^{(2-\epsilon)} n$ hardness of approximation, for general weights. Given the applicability to wireless communication, we present a scalable and easily implemented $O(\log n)$ approximation algorithm, Cover-and-Grow for fixed-dimensional Euclidean space with path-loss exponent at least 2.

Keywords: Network design · Wireless · Approximation

1 Introduction

1.1 Motivation

Consider a DBS (Digital Broadcast Satellite) system such as Dish or DIRECTV in the USA (see Fig. 1). The down-link is an omni-directional broadcast from constellations of satellites to groups of apartments or neighborhoods serviced by one or more dish installations. The up-link is sometimes a wired network but in remote areas it is usually structured as a tree consisting of point-to-point wireless links directed towards the network provider's head-end (root). The high availability requirement of such services are typically satisfied by having multiple head-ends and anycasting to them. The same architecture is found in CATV

© Springer International Publishing Switzerland 2016

T.N. Dinh and M.T. Thai (Eds.): COCOON 2016, LNCS 9797, pp. 235–246, 2016.

DOI: 10.1007/978-3-319-42634-1_19

(originally Community Antenna TV), or cable TV distribution systems as well as sensor networks where an omni-directional broadcast from a beacon is used to activate and control the sensors; the sensors then funnel their information back using relays. Moreover, this architecture is also beginning to emerge in drone networks, for broadcasting the Internet, by companies such as Google [10] and Facebook's Connectivity Labs [8]. The Internet is to be broadcast from drones flying fixed patterns in the sky to a collection of homes on the ground. The Internet up-link from the homes is then aggregated using wireless links organized as a tree to be sent back to the drones. Anycasting is an integral part of high-availability services such as Content Delivery Networks (CDNs) where reliable connectivity is achieved by reaching some node in the group. What is the common architecture underlying all these applications and what is the constraining resource that is driving their form?

The various distribution systems can be abstractly seen to consist of a down-link *ball* and an up-link *funnel-tree* (see Fig. 1). The ball is an omni-directional broadcast from the publisher or content-producer to a large collection of subscribers or content-consumers. At the same time, the consumers have information that they need to dynamically send back to the publisher in order to convey their preferences and requirements. The funnel-tree achieves this up-link efficiently in terms of both time and energy. Aggregation of information and use of relays uses less energy as compared to omni-directional broadcasts by each node back to the publisher and also avoids

the scheduling needed to avoid interference. In this work, we focus primarily on total energy consumption. The application scenarios mentioned in the opening paragraph are all energy sensitive. Sensor networks [11] and drone fleets [12] are particularly vulnerable to energy depletion. For the purpose of energy conservation, generally each wireless node can dynamically adjust its transmitting power based on the distance of the receiving nodes and background noise. In the most common power-attenuation model [14], the signal power falls as $\frac{1}{r^\kappa}$ where r is the distance from the transmitter to the receiver and κ is the path-loss exponent - a constant between 2 and 4 dependent on the wireless environment. A key implication of non-linear power attenuation is that relaying through an intermediate node can sometimes be more energy efficient than transmitting directly - a counter-intuitive violation of the triangle inequality - e.g., in a triangle ABC with obtuse angle ABC , where $d_{AB}^2 + d_{BC}^2 < d_{AC}^2$.




Fig. 1. Pictogram of Digital Broadcast Satellite System with 2 satellite groups and 4 ground station groups on left with associated demand graph on the right. The broadcast balls are denoted by dotted black lines, and the funnel trees by solid yellow lines. (Color figure online)

1.2 Problem Formulation and Terminology

In this paper, we consider a general formulation that encompasses a wide variety of scenarios: given a collection of groups (of nodes) along with a directed demand graph over these groups the goal is to design a collection of balls and associated funnel-trees of lowest cost so that every demand requirement is met - meaning that if there is an arc from a source group to a destination group then the solution must have a ball centered at a node of the source group that includes a funnel-tree containing a node of the destination group.

Formally, we define the *group-to-group anycast* problem, or *g2g-anycast*, as follows: as input we are given n nodes along with a collection of source groups S_1, S_2, \dots, S_p and a collection of destination groups T_1, T_2, \dots, T_q which are subsets of these nodes; a demand graph on these groups consisting of directed arcs from source groups S_i to destination groups T_j . A nonnegative cost c_{uv} is specified between every pair of nodes; when a node u incurs a cost C in doing an omni-directional broadcast it reaches all nodes v such that $c_{uv} \leq C$. A metric d_{uv} is also specified between every pair of nodes and when a node u connects to node v in the funnel-tree using a point-to-point link it incurs a cost d_{uv} . A solution consists of a broadcast ball around every source node s (we give a radius which the source can broadcast to), and a funnel tree rooted at s . A demand S_i, T_j is satisfied if there is a broadcast ball from some $s \in S_i$ which contains some $t \in T_j$ and the funnel tree of s also includes t . The cost of the solution is the sum of the ball-radii around the source nodes (under the broadcast costs c) and the sum of the costs of the funnel trees (under the funnel metric d) that connect all terminal-nodes used to cover the demands to the source nodes within whose balls they lie. We do not allow funnel trees to share edges (even if they are going to the same source group), and will pay for each copy of an edge used.

- First, the bipartite demand graph is no less general than an arbitrary demand graph since a given group can be both a source group and destination group.
- Second, since funnel trees sharing the same edge pay separately, solutions to the problem decompose across the sources and it is sufficient to solve the case where we have exactly one source group $S = \{s_1, s_2, \dots, s_k\}$ and destination groups T_1, T_2, \dots, T_q (i.e. the demand graph is a star consisting of all arcs $(S, T_j), 1 \leq j \leq q$). This observation also enables parallelized implementations.
- Lastly, there is no loss of generality in assuming a metric d_{uv} for funnel-tree costs; even if the costs were arbitrary their metric completion is sufficient for determining the optimal funnel-tree.

We refer collectively to the (ball) costs c_{uv} and (funnel-tree) metric distances d_{uv} as *weights*. In this paper we consider two cases - one, the general case where the weights can be arbitrary and two, the special case where the nodes are embedded in a Euclidean space and all weights are induced from the embedding.

1.3 Our Contributions

Our main results on the minimum energy g2g-anycast problem are as follows (Fig. 2):

	g2g, any metric	g2s, any metric	g2g, ℓ_2^2 norm
Upper	$O(\log^4 n)$	$2 \ln n$	$O(\log n)$
Lower	$\Omega(\log^{2-\epsilon} n)$	$\Omega(\log n)$	$(1 - o(1)) \ln n$

Fig. 2. A summary of upper and lower bounds achieved in the different problems. The lower bound holds for every fixed $\epsilon > 0$

1. We present a polynomial-time $O(\log^4 n)$ approximation algorithm for the g2g-anycast problem on n nodes with general weights. We complement this with an $\Omega(\log^{2-\epsilon} n)$ hardness of approximation, for any $\epsilon > 0$ (Sect. 2).
2. One scenario with practical application is where every destination group is a singleton set while source groups continue to have more than one node; we refer to this special case of g2g-anycast as *g2s anycast*. We present a *tight* logarithmic approximation result for g2s-anycast (Sect. 3).
3. For the realistic scenario where the nodes are embedded in a 2-D Euclidean plane with path-loss exponent $\kappa \geq 2$, we design an efficient $O(\log n)$ -approximation algorithm Cover-and-Grow, and also establish a matching logarithmic hardness of approximation result (Sect. 4).
4. Lastly, we compare Cover-and-Grow with 4 alternative heuristics on random 2-D Euclidean instances; we discover that Cover-and-Grow does well in a wide variety of practical situations in terms of both running time and quality, besides possessing provable guarantees. This makes Cover-and-Grow a go-to solution for designing near-optimal data dissemination networks in the wireless infrastructure space (Sect. 5).

1.4 Related Work

A variety of power attenuation models for wireless networks have been studied in the literature [14]. Though admittedly coarse, the model based on the path loss exponent (varying from 2, in free space to 4, in lossy environments) is the standard way of characterizing attenuation [13]. The problems of energy efficient multicast and broadcast in this model have been extensively studied [9, 16–18]. Two points worth mentioning in this context are: one, we consider the funnel-tree as consisting of point-to-point directional transmissions rather than an omnidirectional broadcast since the nonlinear cost of energy makes it more economical to relay through an intermediate node, and two, we consider only energy spent in transmission but not in reception.

Network design problems are notoriously NP-hard. Over time sophisticated approximation techniques have been developed, ranging from linear programming and randomized rounding to metric embeddings [19]. The g2g-anycast problem with general weights is a substantial generalization including problems such as minimum spanning trees, multicast trees, broadcast trees, Steiner trees and Steiner forests. Even the set cover problem can be seen as a special case where the destination groups are singletons. The g2g-anycast also generalizes the much harder group Steiner tree problem [4, 5].

2 Approximating g2g-anycast

In this section, we present an $O(\log^4 n)$ -approximation for the g2g-anycast problem with general weights by a reduction to the generalized set-connectivity problem. We then give a reduction from the group Steiner tree problem that demonstrates that there is no polynomial-time $\log^{2-\epsilon} n$ -approximation algorithm for g2g-anycast unless $P = NP$.

2.1 Approximation Algorithm for g2g-anycast with General Weights

The generalized set-connectivity problem [2] takes as input an edge-weighted undirected graph $G = (V, E)$, and collection of demands $\{(S_1, T_1), \dots, (S_k, T_k)\}$, each pair are disjoint vertex sets. The goal is to find a minimum-weight subgraph that contains a path from any node in S_i to any node in T_i for every $i \in \{1, \dots, k\}$. Without loss of generality, the edge weights can be assumed to form a metric. Chekuri et al. [2] present an $O(\log^2 n \log^2 k)$ -approximation for this problem using minimum density junction trees.

We show a reduction from the g2g-anycast problem with general weights to the generalized set-connectivity problem. Recall that without loss of generality, we may assume that in the g2g problem, we are given a single source group S , a collection of destination groups T_1, \dots, T_q , nonnegative (broadcast) costs c_{uv} , and (funnel-tree) metric costs d_{uv} .

2.2 The Reduction

The main idea of the reduction is to overload the broadcast cost of the ball radius around each node in the source group S into a larger single metric in which we use the generalized set-connectivity algorithm. In particular, for every source node $s_i \in S$, we sort the nodes in $T_1 \cup \dots \cup T_q$ in increasing order of broadcast cost from s_i to get the sorted order, say t_1^i, \dots, t_r^i where t_j^i is at distance c_{ij} from s_i , and we have $c_{i1} \leq c_{i2} \dots \leq c_{ir}$, where $|T_1 \cup \dots \cup T_q| = r$. We now build r different graphs $G(i, 1), \dots, G(i, r)$ where $G(i, j)$ is a copy of the metric completion of G under the funnel tree costs d induced on the node set $\{s_i, t_1^i, \dots, t_r^i\}$, with the copies denoted as $\{s_i(j), t_1^i(j), \dots, t_r^i(j)\}$. (Note that the terminal node t_a^i appears in copies a through r .) Finally, we take the r copies of the node s_i denoted $s_i(1), s_i(2), \dots, s_i(r)$ and connect them to a new node $s_i(0)$ where the cost of the edge from $s_i(j)$ to $s_i(0)$ is c_{ij} . Thus these r different copies $G(i, 1), \dots, G(i, r)$ all connected to the new node $s_i(0)$




Fig. 3. A connected component $G(i)$ in the reduction of the g2g-anycast problem with general weights to the generalized set connectivity problem.

together form one connected component $G(i)$. We now repeat this process for every source node s_i for $i \in \{1, \dots, k\}$ to get k different graphs $G(1), \dots, G(k)$ (Fig. 3).

We are now ready to define the generalized set connectivity demands. We define a new super source set $SS = \{s_1(0), s_2(0), \dots, s_k(0)\}$. For each of the destination groups T_x , we define the terminal set TT_x to be the union of the copies of all corresponding terminal nodes in any of the copies $G(i)$. More precisely $TT_x = \{\cup_i t_a^i(j) | a \leq j \leq r, t_a^i \in T_x\}$. The final demand pairs for the set connectivity problem are $\{(SS, TT_1), \dots, (SS, TT_q)\}$.

Lemma 1. *Given an optimal solution to the g2g-anycast problem, there is a solution to the resulting set connectivity problem described above of the same cost.*

Proof. Suppose the solution of the g2g problem involved picking broadcast ball radii c_1, \dots, c_k from source nodes s_1, \dots, s_k respectively. We also have funnel trees H_1, \dots, H_k that connect terminals $T(H_1), \dots, T(H_k)$ to s_1, \dots, s_k respectively. Note that all terminals in $T(H_x)$ are within the thresholds that receive the broadcast from s_x , i.e. for every such terminal $t \in H_x$, the broadcast cost of the edge between s_x and t is at most the radius threshold c_x at which s_x is broadcasting.

Consider the tree H_x with terminals $T(H_x)$ connected to the root s_x , so that c_x is the largest weight of any of the edges from s_x to any terminal in $T(H_x)$. (If all of them were even closer, we can reduce the broadcast cost c_x of broadcasting from s_x and reduce the cost of the g2g solution.) Let the terminal in the funnel tree with this broadcast cost be $t(x)$ and in the sorted order of weights from s_x let the rank of $t(x)$ be p . We now consider the graph copy $G(x, p)$ and take a copy of the funnel tree H_x in this copy. To this we add an edge from the root $s_x(p)$ to the node $s_x(0)$ of cost c_{xp} . The total cost of this tree thus contains the funnel tree cost of H_x (denoted by $d(H_x)$) as well as the broadcast cost of c_{xp} from s_x . Taking the union of such funnel trees over all the copies gives the lemma.

Lemma 2. *Given an optimal solution to the set connectivity problem described above, there is a solution to the g2g-anycast problem from which it was derived of the same total weight.*

Proof. In the other direction, consider each copy $G(x)$ in turn and consider the set of edges in the tree containing the source node $s_x(0)$ in the solution to the generalized set-connectivity instance. First notice that it contains at most one of the edges to a copy $s_x(q)$ for some q . Indeed if we have edges to two different copies $s_x(p)$ and $s_x(q)$ from $s_x(0)$ for $p < q$, then since $G(x, p) \subset G(x, q)$, we can consider the tree edges in $G(x, p)$ and buy them in $G(x, q)$ where they also occur to cover the same set of terminals at smaller cost. In this way, we can save the broadcast cost of the copy of the edge from $s_x(0)$ to $s_x(p)$ contradicting the optimality of the solution. Now that we have only one of the edges, say to $s_x(q)$ from $s_x(0)$, we can consider all the edges of the tree in the copy $G(x, q)$ and include these edges in a funnel tree H_x' . The distance of the edge from $s_x(0)$ to

$s_x(q)$ pays for the broadcasting cost from s_x in the original instance and the cost of the rest of the tree is the same as the funnel tree cost of H'_q (Note that our observation above implies that edges in the metric completion in the tree can be converted to paths in the graph and hence connect all the nodes in the tree).

Since every terminal superset TT_j is connected to some source node of SS , all the demands of the g2g problem must be satisfied in the collection of funnel trees H'_x constructed in this way giving a solution to the g2g problem of the same cost.

The above two lemmas with the result of [2] gives us the following result.

Theorem 1. *The general weights version of the g2g-anycast problem with k destination groups admits a polynomial-time approximation algorithm with performance ratio $O(\log^2(k) \log^2 n)$ in an n -node graph.*

2.3 Hardness of Approximating g2g-anycast

We observe that the g2g-anycast problem with general weights can capture the group Steiner tree problem which is known to be $\log^{2-\epsilon} n$ -hard to approximate unless NP is contained in quasi-polynomial time [6].

In the group Steiner tree problem, we are given an undirected graph with metric edge costs, a root s and a set of subsets of nodes called groups, say T_1, \dots, T_g , and the goal is to find a minimum cost tree that connects the root with at least one node from each group. We can easily define this as a g2g-anycast problem with a singleton source group $S = \{s\}$ with the single root node. The terminal sets for the g2g-anycast problem are the groups T_1, \dots, T_g , with the demand graph $(S, T_1), \dots, (S, T_g)$. We can set the broadcast costs of any node in the graph from s to be zero; we use the given metric costs in the group Steiner problem as the funnel tree costs to capture the cost of the group Steiner tree. Any solution to the resulting g2g-anycast problem is a single tree connecting s to at least one node in each of the groups as required and its total weight is just its funnel tree cost that reflects precisely the cost of this feasible group Steiner tree solution. The hardness follows from this approximation-preserving reduction.

3 Approximating g2s-anycast

In this section, we consider g2s-anycast, a special case of the g2g-anycast, in which each destination group is a singleton set (i.e., has exactly one terminal). Let S denote the source-set and t_1, \dots, t_q denote the terminals.

The desired solution is a collection of broadcast balls and funnel trees T_v , each rooted at a source node v , so that for every demand (S, t_j) , there exists at least one node v in S such that $t_j \in T_v$.

We now present a $\Theta(\log n)$ -approximation algorithm for g2s-anycast problem. Our algorithm iteratively computes an approximation to a minimum density assignment, which assigns a subset of as yet unassigned terminals to a source node, and then combines these assignments to form the final solution.

Minimum Density Assignment. We seek a source s and a tree T_s rooted at s that connects s to a subset of terminals, such that the ratio $(c(T_s) + d(T_s))/|T_s|$ is minimized among all choices of s and T_s (here $c(T_s)$ denotes the minimum broadcast cost for s to reach the terminals in T_s , while $d(T_s)$ denotes the funnel-tree cost, i.e. the sum of the metric distances d_{uv} over all edges $uv \in T_s$). We present a constant-approximation to the problem, using a constant factor approximation algorithm for the rooted k -MST problem, which is defined as follows: given a graph G with weights on edges and a root node, determine a tree of minimum weight that spans at least k vertices. The best known approximation factor for the k -MST problem [15] is 2 [3]. We now present our algorithm for minimum density assignment.

- For each source $s \in S$, integer $k \in [1, n]$, and integer r drawn from the set $\{c_{st_j} | 1 \leq j \leq q\}$:
 - Let G' denote the graph with vertex set $\{s\} \cup \{t_j | c_{st_j} \leq r\}$, and edge weights given by d .
 - Compute a 2-approximation $T'(s, r, k)$ to the k -MST problem over the graph G' with s being the root.
- Among all trees computed in the above iterations, return a tree that minimizes $\min_{s,r,k} (d(T'(s, r, k)) + r)/k$.

Lemma 3. *The above algorithm is a polynomial-time 2-approximation algorithm for the minimum density assignment problem.*

Proof. We first show that the algorithm is polynomial time. The number of different choices for the source equals the size of the source set, the number of choices for k is n , and the number of different values for r is the number of different broadcast costs, which is at most n . Thus the number of iterations in the for loop is at most n^3 . Consider an optimal solution T to the minimum density assignment problem, rooted at source s . It is a valid solution to the k -MST problem in the iteration given by $s, r = c(T), k = |T|$. For this particular iteration, the tree $T'(s, r, k)$ satisfies $(d(T'(s, r, k)) + r)/k \leq (2d(T) + r)/k \leq 2 \cdot (d(T) + r)/k$. Since our algorithm returns the tree that has the best density, we have a 2-approximation for the minimum density assignment.

Approximation Algorithm for g2s-Anycast. Our algorithm is a greedy iterative algorithm, in which we repeatedly compute an approximation to the minimum density assignment problem, and return an appropriate union of all of the trees computed.

- For each source s , set T_s to $\{s\}$.
- While all terminals are not assigned:
 - Compute a 2-approximation T to the minimum density assignment problem using any source s and the unassigned terminals.
 - If T is rooted at source s , then set T_s to be the minimum spanning tree of the union of the trees T and T_s .
- Return the collection $\{T_s\}$.

Theorem 2. *The greedy algorithm yields an approximation algorithm with performance ratio $2 \ln n$ to the g2s-anycast problem.*

Proof. Let OPT denote the cost of the optimal solution to the problem. Any solution is composed of at most m trees, one for each of the sources, with each singleton group being included as a node in one of these trees. Let T_s^* denote the tree rooted at source s in an optimal solution.

Consider any iteration i of our algorithm. Let n_i denote the number of unassigned terminals at the start of the iteration i . By an averaging argument, we know there exists a source s such that

$$\frac{d(T_s^*) + c(T_s^*)}{|T_s^*|} \leq \frac{OPT}{n_i},$$

By Lemma 3, it follows that in the i th iteration of the greedy algorithm, if T_i is the tree computed in the step, then

$$\frac{d(T_i) + c(T_i)}{|T_i|} \leq \frac{2 \cdot OPT}{n_i},$$

Adding over all steps, we obtain that the total cost is

$$\sum_i (d(T_i) + c(T_i)) \leq 2 \cdot OPT \cdot \sum_i \frac{|T_i|}{n_i} \leq 2 \cdot OPT \cdot H_n \leq 2OPT \ln n.$$

Hardness of Approximation. We complement the positive result with a matching inapproximability result which shows that the above problem is as hard as set cover.

Theorem 3. *Unless $NP = P$ there is no polynomial-time $\alpha \ln n$ approximation to the g2s-anycast problem, for a suitable constant $\alpha > 0$.*

We defer the proof of this theorem to Appendix A of the full version [7].

4 Euclidean g2g-anycast

In this section, we present a $\Theta(\log n)$ -approximation for the more realistic version of the g2g-anycast problem in the 2-D Euclidean plane. We achieve our results by a reduction to an appropriately defined set cover problem.

In detail, all the points in both the source group S and destination groups T_1, \dots, T_q lie in the 2-D Euclidean plane. The cost of an edge (u, v) is the Euclidean distance between u and v raised to the path loss exponent κ . For the rest of this section, we assume that $\kappa = 2$. (The corresponding results for $\kappa > 2$ follow with very simple modifications.) First we show that even this special case of the g2g-anycast problem does not permit an approximation algorithm with ratio $(1 - \epsilon) \ln n$ on an instance with n nodes unless NP is in quasi-polynomial time. Next, we present *Cover-and-Grow*, an $O(\log n)$ -approximation algorithm that applies a greedy heuristic to an appropriately defined instance of the set covering problem.

Hardness of 2-D g2g-anycast. Again we can prove a hardness via a reduction from set cover.

Theorem 4. *The 2-D Euclidean version of the g2g problem on n nodes does not permit a polynomial-time $(1 - o(1)) \ln n$ approximation algorithm unless $NP = P$.*

The proof of this can be found in Appendix B of the full version [7].

4.1 Cover-and-Grow

We now describe a matching $O(\log n)$ -approximation for the problem. For this we first need the following property of minimum spanning trees of points in the 2-D Euclidean plane within a unit square, when the costs of any edge in the tree are the squared Euclidean distances between the edge's endpoints.

Theorem 5 [1]. *The weight of a minimum spanning tree of a finite number of points in the 2-D Euclidean plane within a unit square, where the weight of any edge is the square of the Euclidean distance between its endpoints, is at most 3.42.*

We can apply this theorem to bound the cost of the funnel trees within any demand ball in the solution within a factor of at most 3.42 of the cost of the ball. Indeed, by scaling the diameter of the demand ball to correspond to unit distance, the above theorem shows that for any finite set of terminal nodes (i.e. nodes in the destination group) within the ball, a funnel tree which is an MST that connects these terminal nodes to the center of the ball has total cost at most 3.42. The cost of the demand ball is the square of the Euclidean distance of the ball radius which, in the scaled version, has cost $(\frac{1}{2})^2 = \frac{1}{4}$. This shows that the funnel tree has cost at most 13.68 times the cost of the funnel ball. This motivates an algorithm that uses balls of varying radii around each source node as a “set” that has cost equal to the square of the ball radius (the ball cost) and covers all the terminal nodes within this ball (which can be connected in a funnel tree of cost at most 13.68 times that of the demand ball).

Algorithm Cover-and-Grow

1. Initialize the solution to be empty.
2. While there is still an unsatisfied demand edge
 - For every source node s_i , for every possible radius at which there is a terminal node belonging to some destination group T for which the demand (S, T) is yet unsatisfied, compute the ratio of the square of the Euclidean radius of the ball to the number of *as yet unsatisfied* destination groups whose terminal nodes lie in the ball.
 - Pick the source node and ball radius whose ratio is minimum among all the available balls, and add it to the solution (both the demand ball around this node and a funnel tree from one node of each destination group whose demand is unsatisfied at this point). Update the set of unsatisfied demands accordingly.

Theorem 6. *Algorithm Cover-and-grow runs in polynomial time and gives an $O(\log n)$ -approximate solution for the 2-D g2g-anycast problem in an n -node graph.*

Proof. We will use a reduction from the given 2-D g2g-anycast problem to an appropriate set cover problem as described in the algorithm. The elements of the set cover problem are the terminal sets T_j such that the demand graph has the edge (S, T_j) . For every source node $s_i \in S$, and for every possible radius r at which there is a terminal node belonging to some destination group T for which there is a demand (S, T) , we consider a set $X(s_i, r)$ that contains all the destination groups T_j such that some node of T_j lies within this ball. The cost of this set is r^2 .

First, we argue that an optimal solution for the 2-D g2g-anycast problem of cost C^* gives a solution of cost at most C^* to this set cover problem. Next, we show how any feasible solution to the set cover problem of cost C gives a feasible solution to the 2-D g2g-anycast problem of cost at most $14.68C$. These two observations give us the result since the algorithm we describe is the standard greedy approximation algorithm for set cover.

To see the first observation, given an optimal solution for the 2-D g2g-anycast problem of cost C^* , we pick the sets corresponding to the demand balls in the solution for the set cover problem. Since these demand balls are a feasible solution to the anycast problem, they together contain at least one terminal from each of the destination groups T_j for which there is a demand edge (S, T_j) . These balls form a solution to the set cover problem and the demand ball costs of the anycast solution alone pay for the corresponding costs of the set cover problem. Hence this feasible set cover solution has cost at most C^* .

For the other direction, given any feasible solution to the set cover problem of cost C , note that this pays for the demand balls around the source nodes in this set cover solution. Now we can use the implication in the paragraph following Theorem 5 to construct a funnel tree for each of these demand balls that connects all the terminals within these balls to the source node at the center of the ball with cost at most 13.68 times the cost of the demand ball around the source node. Summing over all such balls in the solution gives the result.

5 Empirical Results

We conducted simulations comparing Cover-and-Grow with four different natural heuristics for points embedded in a unit square in the 2-D Euclidean plane. These simulations allow us gain perspective on the real-world utility of Cover-and-Grow vis a vis alternatives that do not possess provable guarantees but yet have the potential to be practical. The specifics of the simulation and the details of the results are discussed in Appendix C of the full version [7]. Cover-and-Grow performs comparably to the heuristics in performance; and the runtime of Cover-and-Grow was better than the heuristics except for the T-centric approach.

References

1. Aichholzer, O., Allen, S., Aloupis, G., Barba, L., Bose, P., de Varufel, J.L., Iacono, J., Langerman, S., Souvaine, D., Taslakian, P., Yagnatinsky, M.: Sum of squared edges for MST of a point set in a unit square. In: Japanese Conference on Discrete and Computational Geometry (JCDCG) (2013)
2. Chekuri, C., Even, G., Gupta, A., Segev, D.: Set connectivity problems in undirected graphs and the directed Steiner network problem. *ACM Trans. Algorithms* **7**(2), 18:1–18:17 (2011)
3. Garg, N.: Saving an epsilon: a 2-approximation for the k-MST problem in graphs. In: *ACM Theory of Computing*, pp. 396–402 (2005)
4. Garg, N., Konjevod, G., Ravi, R.: A polylogarithmic approximation algorithm for the group Steiner tree problem. *J. Algorithms* **37**(1), 66–84 (2000)
5. Halperin, E., Kortsarz, G., Krauthgamer, R., Srinivasan, A., Wang, N.: Integrality ratio for group Steiner trees and directed Steiner trees. *SIAM J. Comput.* **36**(5), 1494–1511 (2007)
6. Halperin, E., Krauthgamer, R.: Polylogarithmic inapproximability. In: *ACM Theory of Computing*, pp. 585–594 (2003)
7. Iglesias, J., Rajaraman, R., Ravi, R., Sundaram, R.: Balls and funnels: energy efficient group-to-group anycasts. In: Dinh, T.N., Thai, M.T. (eds.) *COCOON 2016*. LNCS, vol. 558, pp. 235–246. Springer, Heidelberg (2016). CoRR <http://arxiv.org/abs/1605.07196>
8. Lapowsky, I.: Facebook lays out its roadmap for creating internet-connected drones. *Wired* (2014). <http://www.wired.com/2014/09/facebook-drones-2/>
9. Li, D., Liu, Q., Hu, X., Jia, X.: Energy efficient multicast routing in ad hoc wireless networks. *Comput. Commun.* **30**(18), 3746–3756 (2007)
10. McNeal, G.: Google wants Internet broadcasting drones, plans to run tests in New Mexico. *Forbes* (2014). <http://www.forbes.com/sites/gregorymneal/2014/09/19/google-wants-internet-broadcasting-drones-plans-to-run-tests-in-new-mexico/>
11. Miluykovski, V., Segal, M., Shpungin, H.: Location, location, location: using central nodes for efficient data collection in WSNs. In: *WiOpt*, pp. 333–340, May 2013
12. Olsson, P.M.: Positioning algorithms for surveillance using unmanned aerial vehicles. Licentiate thesis, Linköpings universitet (2011)
13. Path Loss Wikipedia. http://en.wikipedia.org/wiki/Path_loss
14. Rappaport, T.: *Wireless Communications: Principles and Practice*, 2nd edn. Prentice Hall PTR, Upper Saddle River (2001)
15. Ravi, R., Sundaram, R., Marathe, M.V., Rosenkrantz, D.J., Ravi, S.S.: Spanning trees short or small. In: *ACM-SIAM Discrete Algorithms, SODA 1994*, pp. 546–555. SIAM (1994)
16. Wan, P.J., Călinescu, G., Li, X.Y., Frieder, O.: Minimum-energy broadcasting in static ad hoc wireless networks. *Wirel. Netw.* **8**(6), 607–617 (2002)
17. Wieselthier, J.E., Nguyen, G.D., Ephremides, A.: On the construction of energy-efficient broadcast and multicast trees in wireless networks. In: *INFOCOM*, pp. 585–594 (2000)
18. Wieselthier, J.E., Nguyen, G.D., Ephremides, A.: Algorithms for energy-efficient multicasting in static ad hoc wireless networks. *MONET* **6**(3), 251–263 (2001)
19. Williamson, D.P., Shmoys, D.B.: *The Design of Approximation Algorithms*, 1st edn. Cambridge University Press, New York (2011)