Minimum-Energy Broadcast with Few Senders*

Stefan Funke, Sören Laue, and Rouven Naujoks

Max-Planck-Institut für Informatik,
Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany
{funke,soeren,naujoks}@mpi-inf.mpg.de

Abstract. Broadcasting a message from a given source node to all other nodes is a fundamental task during the operation of a wireless network. In many application scenarios the network nodes have only a limited energy supply, hence minimizing the energy consumption of any communication task prolongs the lifetime of the network. During a broadcast operation using intermediate nodes to relay messages within the network might decrease the overall energy consumption since the cost of transmitting a message grows super-linearly with the distance. On the other hand using too many intermediate nodes during a broadcast operation increases both latency as well as the chances that some transmission could not properly received (e.g. due to interference).

In this paper we consider a constrained broadcast operation, where a source node wants to send a message to all other nodes in the network but at most k nodes are allowed to participate actively, i.e. transmit the message. Restricting the number of transmitting nodes helps in reducing interference, latency and increasing reliability of the broadcast operation, of course at the cost of a slightly higher energy consumption. For the case of network nodes embedded in the Euclidean plane we provide a $(1+\epsilon)$-approximation algorithm which runs in time linear in n and polynomial in $1/\epsilon$ but with an exponential dependence on k. As an alternative we therefore also provide an $O(1)$-approximation whose running time is linear in n and polynomial in k. The existence of a $(1+\epsilon)$-approximation algorithm is in stark contrast to the unconstrained broadcast problem where even in the Euclidean plane no algorithm with approximation factor better than 6 is known so far.

1 Introduction

In contrast to wired or cellular networks, ad hoc wireless networks a priori are unstructured in a sense that they lack a predetermined interconnectivity. An ad hoc wireless network is built of a set of radio stations P, each of which consists of a receiver as well as a transmission unit. A radio station v can send a message by setting its *transmission range* $r(v)$ and then by starting the transmission process. All other radio stations at distance at most $r(v)$ from v will be able to receive

* This work was supported by the Max Planck Center for Visual Computing and Communication (MPC-VCC) funded by the German Federal Ministry of Education and Research (FKZ 01IMC01).

J. Aspnes et al. (Eds.): DCOSS 2007, LNCS 4549, pp. 404–416, 2007.

v's message (we are ignoring interference for now). For transmitting a message across a transmission range $r(v)$, the power consumption of v's transmission unit is proportional to $r(v)^\alpha$, where α is the *transmission power gradient*. In the idealistic setting of empty space, $\alpha = 2$, but it may vary from 2 to more than 6 depending on the environment conditions of the location of the network. Given some transmission range assignment $r : P \to \mathbb{R}_{\geq 0}$ for all nodes in the network, we can derive the so-called *communication graph* $G^{(r)} := G(P, E)$. $G(P, E)$ is a directed graph with vertex set P which has a directed edge (p, q) iff $r(p) \geq |pq|$, where $|pq|$ denotes the Euclidean distance between p and q. The cost of the transmission range assignment r is

$$\text{cost}(r) := \sum_{v \in P} r(v)^\alpha$$

Numerous optimization problems can now be considered by looking for the minimum cost transmission range assignment r such that the respective communication graph G satisfies some property Π, see [5] for an overview. One classic property Π is defined as follows: *Given a specific source node s we want the communication graph G to contain a directed spanning tree rooted at s.* This problem is called the *energy-minimal broadcast problem (EMBC)*, since the respective transmission range assignment allows the source node s to distribute a message over the whole network at minimum total energy cost.

It is known, that if the points are located in the Euclidean plane, the minimum spanning tree (MST) of the point set induces a transmission range assignment which has cost at most a factor of 6 above the optimum solution [1]. On the other hand, there are point sets where the MST-based solution is a factor 6 worse than the optimal solution, so the bound of 6 is tight [15].

One problem that is particularly prominent for the MST–based solution is the fact that in the resulting transmission range assignment a very large fraction of the network nodes are transmitting (i.e. have non-zero transmission range). In the MST-based range assignment, at least $n/6$ nodes are actually senders during the broadcast operation (since the maximum degree of the minimum spanning tree of a set of points in the Euclidean plane is bounded by 6). This raises several critical issues: (a) The more network nodes are transmitting in the process of one broadcast operation, the more likely it is that some nodes in the network experience interference due to several nearby nodes transmitting at the same time (unless special precautions are taken that interference does not occur). (b) Every retransmission of a message implies a certain delay which is necessary to setup the transmission unit etc; that is, the more senders are involved in the broadcast operation, the higher the latency. This effect is even amplified by the previous problem if due to interference messages have to be resent. (c) Network nodes are not 100% reliable; if for example the probability for a network node to operate properly is 99.9%, the probability for a network broadcast to fail, i.e. not all nodes receiving the message, is $1 - 0.999^{(n/6)}$, which for a network of $n = 3000$ nodes is around 40%! This suggests to look for broadcast operations in the network that use only very few sending nodes. Of course, this comes at the cost of an increased total power consumption, but the behaviour with respect to the critical issues (a) to (c) can be drastically improved.

In this paper we suggest the following restricted broadcast operation: *Given a specific source node s we want to find a transmission range assignment r of minimum total cost such that the respective communication graph G contains a directed spanning tree rooted at s* **and at most k nodes have a non-zero transmission range assigned**. We call this problem the *k-set energy-minimal broadcast (k-SEMBC)* problem. Allowing only a small number k of sending nodes during the broadcast operation has several advantages: (a) The k transmissions can be easily scheduled in k different time slots, hence avoiding any interference at all. (b) The latency is obviously bounded by $O(k)$. (c) In the above scenario the probability of a broadcast operation to fail is $1-0.999^k$, which e.g. for $k=10$ is 1%.

Of course, the best behaviour in terms of interference, latency and reliability can be achieved by having the source node s transmit its message directly to all nodes in the network. Assume w.l.o.g. that by scaling the maximum distance of another node in the network to s is 1, this operation would cost one unit of power. Consider the scenario in Figure 1; here we have the n radio stations equally distributed on a segment of length 1. As mentioned above, the direct broadcast from s incurs a power cost of $1^\alpha = 1$. In case of an unbounded number of allowed sending nodes, essentially every node just forwards the message to the next node to the right. The total power consumption of this broadcast (which involves $n-1$ sending nodes) is $(\frac{1}{n-1})^\alpha \cdot (n-1) = (\frac{1}{n-1})^{\alpha-1}$. That is the power savings compared to the direct transmission is a factor of about $n^{\alpha-1}$, using $n-1$ sending nodes, though, which implies the above issues w.r.t. interference, latency, and reliability. If, on the other hand, we allow only k sending nodes, we could in the ideal case select $k-1$ stations at about equal distance from each other on the segment and incur a cost of $(\frac{1}{k})^\alpha \cdot k = (\frac{1}{k})^{\alpha-1}$. This simple example illustrates the (maximum) potential energy savings due to the use of intermediate stations when compared to a direct transmission from the source node s. In practice, though, the experienced energy savings are far from such high factors, and also the loss in energy efficiency when allowing only k stations to send compared to the unrestricted case is far less pronounced. In general, the advantage of using intermediate stations for a broadcast operation is greater if the nodes are distributed along 1-dimensional patterns and curves; for a very dense uniform distribution of the nodes e.g. within square, the direct transmission from the source is almost optimal (for $\alpha = 2$, for $\alpha > 2$ the gain of using intermediate stations grows). Unfortunately, the MST-based algorithm will always create a transmission range assignment where at least $n/6$ nodes are sending, even if there exist equally good or even better assignments with few sending nodes.

1.1 Our Contribution

In this paper we consider the k-set minimum energy broadcast problem from an analytical point of view. We show that somewhat surprisingly for any network of n radio stations there exists a subset S of the stations whose size is *independent*

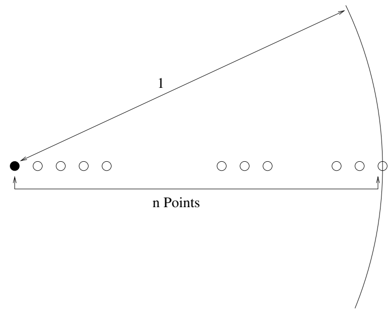

Fig. 1. Energy savings by using intermediate radio stations

of n and which preserves all the important characteristics of S with respect to a energy efficient k-set broadcast. We call S a *synopsis* or *core-set* of the network topology w.r.t. the k-SEMBC. We will show that using a synopsis of size $|S| = O((k/\epsilon)^2)$, any solution of the k-SEMBC for S translates to a solution for k-SEMBC for the original set P at a cost at most a $(1+\epsilon)$ factor away and vice versa. Since the size of this synopsis is independent of the network size, we can run even a brute-force algorithm to compute the optimum k-set broadcast. The running time of this algorithm is linear in n but still exponential in k. So we also present an $O(1)$-approximation algorithm whose running time is linear in n but polynomial in k. We want to emphasize that the focus of this paper is to examine the fundamental structure of the k-SEMBC problem rather than provide practical algorithms for direct use in a wireless network. Though we believe that with some engineering effort variants of our algorithms can be made practical, we have not conducted simulations yet to show practicability of our approaches.

1.2 Related Work

The EMBC problem is known to be $\mathcal{NP}$-hard ([6,5]), For arbitrary, non-metric distance functions the problem can also not be approximated better than a log-factor unless $\mathcal{P} = \mathcal{NP}$ [14]. For the Euclidean setting in the plane, [6] and [15] have shown a lower bound of 6 for the approximation ratio of the MST-based solution. In a sequence of papers the upper bound for this solution was lowered from in several steps (e.g. [6,15,8]) to finally match its lower bound of 6 (Ambühl, [1]). While all these papers focused on analytical worst-case bounds for the algorithm performance, simulation studies e.g. in [7] show that the actual performance in "real-world" networks is much better. There has also been work on more restricted broadcast operations in the spirit of k-SEMBC. In [2] the authors consider a *bounded-hop* broadcast operation where the resulting communication graph has to contain a spanning tree rooted at the source node s of depth at most h. They show how to compute an optimal h-hop broadcast range assignment for $h = 2$ in time $O(n^7)$. For $h > 2$ they show how to obtain a $(1+\epsilon)$-approximation in time $O(n^{O(\mu)})$ where $\mu = (h^2/\epsilon)^{2^h}$, that is, their running time is triply exponential in the number of hops h and this shows up in the exponent of n. In [9] Funke and Laue show how to obtain a $(1+\epsilon)$ approximation for the h-hop broadcast problem in time doubly exponential in h. Their approach is also based on a synopsis of the network, but in contrast to this paper they require a synopsis S that has size exponential in h. We note that bounded-hop broadcasts address the issue of latency since a message will arrive at any network node after at most h intermediate stations, still the reliability and interference problems remain as potentially very many network nodes might actively participate in the broadcast. General surveys of algorithmic range assignment problems can be found in [5,16,12]. Closely related in particular to the $O(1)$-approximation algorithm that we will present is the work by Bilò et al. [4]. They consider the problem of covering a set of n points in the plane using at most k disks such that the sum of the areas of the disks is minimized. They provide a $(1+\epsilon)$-approximation to this problem in time $O(n^{\alpha^2/\epsilon^2})$. They do not address the problem of enforcing connectivity which is part of the k-SEMBC problem.

1.3 Outline

Section 2 recaps a known complexity result for the the unconstrained broadcast problem EMBC and sketches a simple folklore-brute-force algorithm to solve the k-SEMBC problem. Section 3 contains the core contributions of our paper; we show how to extract a small synopsis of the network topology (Section 3.1) and how to use that to obtain a $(1+\epsilon)$-approximation algorithm. In Section 3.2 we show how a faster algorithm obtains an $O(1)$-approximation. Finally, in Section 4 we point out directions for future research.

2 Preliminaries

The unconstrained broadcast problem EMBC is known to be $\mathcal{NP}$-hard and for non-metric distance functions even not well approximable ([6,14]. Since the unconstrained broadcast problem is a special case of the k-set broadcast problem with $k = n$ these hardness results carry over to the k-set broadcast problem, if k is not treated as a constant. If k is regarded a constant, the problem can be solved in polynomial time as we will see in the following.

2.1 A Naive, Brute-Force Algorithm

The k-set broadcast problem can be solved in a brute force manner. Essentially, one can try out all $\binom{n}{k-1}$ different subsets for the $k-1$ active senders apart from the source s. For each of those (and the source node s), one then assigns all possible $n-1$ ranges. In total we have then $O(n^{k-1}(n-1)^k) = O(n^{2k})$ potential power assignments. For each of those we can check in $O(n^2)$ time whether it is a valid k-set broadcast.

That is, overall we have the following corollary:

Corollary 1. *For n points we can compute the optimal k-set broadcast in time $O(n^{2k+2})$.*

For most practical applications, we expect k to be a small constant, but unfortunately not small enough that this naive algorithm can be applied to networks of not too small size (e.g. several thousand nodes). In the following Section we lower our expectations and aim for *approximate* solutions to the k-set broadcast problem. This allows for more efficient algorithms as we will see.

3 Algorithms

3.1 Small Synopsis of the Network Topology

We say a range assignment r is *valid* if the induced communication graph $G^{(r)}$ contains a directed spanning tree reaching all nodes $p \in P$ and rooted at s, and at most k nodes have non-zero transmission range assigned; otherwise we call r *invalid*.

Definition 1. *Let P be a set of n points, $s \in P$ a designated source node. Consider another set S of points (not necessarily a subset of P). If for any valid range assignment $r : P \to \mathbb{R}_{\geq 0}$ there exists a valid range assignment $r' : S \to \mathbb{R}_{\geq 0}$ such that $\text{cost}(r') \leq (1+\epsilon) \cdot \text{cost}(r)$ and for any valid range assignment $r' : S \to \mathbb{R}_{\geq 0}$ there exists a valid range assignment $r : P \to \mathbb{R}_{\geq 0}$ such that $\text{cost}(r) \leq (1+\epsilon) \cdot \text{cost}(r')$ then S is called* $(1+\epsilon)$-synopsis *for (P, s).*

A $(1+\epsilon)$-synopsis for a problem instance (P, s) can hence be viewed as a problem sketch of the original problem. If we can show that a $(1+\epsilon)$-synopsis of small size (independent of n) exists, solving the k-SEMBC problem on this problem sketch immediately leads to an $(1+\epsilon)^2$-approximate solution to the original problem. The former can be even done using a brute force algorithm. Of course, the transformation from range assignment r' for the synopsis S to range assignment r for the input point set P has to be practical in order to derive a solution for the original problem. We will see that this can be done in linear time.

The definition of a synopsis can be seen as a generalization of core-sets defined in previous papers. For example, the term core-set has been defined for k-median [11] or minimum enclosing disk [13]. However, in the case of the k-SEMBC problem we have to consider two more issues. The first is feasibility. While any solution to the k-median problem is feasible not every solution is feasible for the k-SEMBC problem. The second issue is monotonicity. For the problem of the smallest enclosing disk the optimal solution does not increase if we remove points from the input. We do not have this property here. An optimal solution can increase or decrease if we remove input points. Hence, the above definition of a $(1+\epsilon)$-synopsis can be seen as a generalization of core-sets to any optimization problem.

We will now show that we can find a small synopsis to the original problem. We assume that the maximum distance from the source node s to another node is 1.

Lemma 1. *For any k-SEMBC instance there exists a* $(1+\epsilon)^\alpha$*-synopsis of size* $O(\frac{k^2}{\epsilon^2})$.

Proof. We place a grid of grid width $\Delta = \frac{1}{\sqrt{2}}\frac{\epsilon}{k}$ on the plane. Notice, that the grid has to cover an area of radius 1 around the source only because the furthest distance from node s to any other node is 1. Hence its size is $O(\frac{k^2}{\epsilon^2})$. Now we assign each point in P to its closest grid point. Let S be the set of grid points that had at least one point from P snapped to it.

It remains to show that S is indeed a $(1+\epsilon)^\alpha$-synopsis. We can transform any given valid range assignment r for P into a valid range assignment r' for S. We define the range assignment r' for S as

$$r'(p') = \max_{p \text{ was snapped to } p'} r(p) + \sqrt{2}\Delta.$$

Since each point p is at most $\frac{1}{\sqrt{2}}\Delta$ away from its closest grid point p' we certainly have a valid range assignment for S. It is easy to see that the cost of r' for S is not much larger than the cost of r for P. We have:

$$\begin{aligned}\sum_{p'\in S}(r'(p'))^\alpha &= \sum_{p'\in S}(\max_{p \text{ was snapped to } p'} r(p) + \sqrt{2}\Delta)^\alpha \\ &\leq \sum_{p'\in S}(\max_{p \text{ was snapped to } p'} r(p) + \frac{\epsilon}{k})^\alpha \\ &\leq \sum_{p\in P}(r(p) + \frac{\epsilon}{k})^\alpha.\end{aligned}$$

The relative error satisfies

$$\frac{\text{cost}(r')}{\text{cost}(r)} \leq \frac{\sum_{p\in P}(r(p)+\frac{\epsilon}{k})^\alpha}{\sum_{p\in P}(r(p))^\alpha}.$$

Notice, that $\sum_{p\in P} r(p) \geq 1$ and r is positive for at most k points p. Hence, the above expression is maximized when $r(p) = \frac{1}{k}$ for all points p that are assigned a positive value. Thus

$$\frac{\text{cost}(r')}{\text{cost}(r)} \leq \frac{k\cdot(\frac{1}{k}+\frac{\epsilon}{k})^\alpha}{k\cdot(\frac{1}{k})^\alpha} = (1+\epsilon)^\alpha.$$

On the other hand we can transform any given valid range assignment r' for S into a valid range assignment r for P as follows. We select for each grid point $g \in S$ one representative g_P from P that was snapped to it. For the grid point to which s (the source) was snapped we select s as the representative. If we define the range assignment r for P as $r(g_P) = r'(g) + \sqrt{2}\Delta$ and $r(p) = 0$ if p does not belong to the chosen representatives, then r is a valid range assignment for P because every point is moved by the snapping by at most $\Delta/\sqrt{2}$. Using the same reasoning as above we can show that $\text{cost}(r) \leq (1+\epsilon)^\alpha \text{cost}(r')$. Hence, we have shown that S is indeed a $(1+\epsilon)^\alpha$-synopsis.

Once we have solved the k-SEMBC problem for the $(1+\epsilon)^\alpha$-synopsis S we can easily transform the obtained solution to a $(1+\epsilon)^{2\alpha}$-approximate solution to the original problem. Let us now concentrate on solving the k-SEMBC problem for the synopsis S. Since we were able to reduce the problem size to a constant independent of n, we can employ a brute-force strategy to compute an optimal solution for the reduced problem (S, s).

When looking for an optimal, energy-minimal solution for S, it is obvious that each node needs to consider only $|S|$ different ranges. Hence, naively there are at most $\binom{\frac{k^2}{\epsilon^2}}{k} \cdot \left(\frac{k^2}{\epsilon^2}\right)^k$ different range assignments to consider at all. We enumerate all these assignments and for each of them we check whether the range assignment is valid; this can be done in time $|S|$. Of all the valid range assignments we return the one of minimal cost.

Assuming the floor function a $(1+\epsilon)^\alpha$-synopsis S for an instance of the k-SEMBC problem for a set of n radio nodes in the plane can be constructed in linear time. Hence we obtain the following theorem:

Theorem 1. *A $(1+\epsilon)^{2\alpha}$-approximate solution to the k-SEMBC problem on n points in the plane can be computed in time $O(n + \left(\frac{k}{\epsilon}\right)^{4k})$.*

A simple observation allows us to improve the running time slightly. Since eventually we are only interested in an approximate solution to the problem, we are also happy with only approximating the optimum solution for the synopsis S. Such an approximation for S can be found more efficiently by not considering all possible at most $|S|$ ranges for each grid point. Instead we consider as admissible ranges only 0 and $\frac{\epsilon}{k} \cdot (1+\epsilon)^i$ for $i \geq 0$. That is, the number of different ranges a node can attain is at most $1 + \log_{1+\epsilon} \frac{k}{\epsilon} \leq \frac{2}{\epsilon} \cdot \log \frac{k}{\epsilon}$ for $\epsilon \leq 1$. This comes at a cost of a $(1+\epsilon)$ factor by which each individual assigned range might exceed the optimum. The running time of the algorithm improves, though, which leads to our main result in this section:

Theorem 2. *A $(1+\epsilon)^{3\alpha}$-approximate solution to the k-SEMBC problem on n points in the plane can be computed in time $O\left(n + \left(\frac{k^2 \log \frac{k}{\epsilon}}{\epsilon^3}\right)^k\right)$.*

Obviously, a $(1+\psi)$-approximate solution can be obtained by choosing $\epsilon = \theta(\psi/\alpha)$.

3.2 Faster $O(1)$-Approximations

We now show how to compute a constant approximation for the k-set broadcast problem. The idea is to first cluster the points into k clusters. Then we ensure connectivity of these point sets by increasing their cluster sizes. As clustering we define the *k-disk cover problem*:

Definition 2 (k-disk cover problem (k-DCP)). *Given a set P of n points in the Euclidean plane $\mathbb{R}^2$, find a subset $C \subseteq P$ of cardinality at most k and* *radii* $r_v \geq 0$ *associated with each element* $v \in C$ *such that* $\sum_{v \in C} r_v^\alpha$ *is minimized and all points in* P *are covered by the disks* $D_v^{r_v} := \{x \in \mathbb{R}^2 \mid |xv| \leq r_v\}$.

Given a k-disk cover $D := (C,\ (r_v)_{v \in C})$ for P with center points C and radii r_v, we associate with D a range assignment r_D on P as follows:

$$\forall v \in P: \quad r_D(v) := \begin{cases} r_v & \text{if } v \in C \\ 0 & \text{otherwise} \end{cases}$$

By D_i we denote a disk in D. Note that the k-DCP with the additional constraint that the communication graph $G^{(r_D)}$ is connected is exactly the k-set broadcast problem and that an instance of one problem is an instance of the other. Thus, the cost of an optimal solution for an instance I of the k-DCP is a lower bound on I for the k-set broadcast problem. Unfortunately k-DCP is NP-hard (see [4]) but it admits a PTAS as shown in [4] by Bilò et al. A direct consequence of their results is:

Corollary 2. *There exists an algorithm for the* k*-DCP that computes* $(1+\epsilon)$*-approximate solutions in time* $n^{(\frac{\alpha}{\epsilon})^d}$ *for a constant* d.

By setting ϵ to 1 we obtain a 2-approximation algorithm for the k-DCP that runs in time $n^{c'_\alpha}$ for a constant c'_α. Note that their algorithm can easily be modified such that the source s is the center of one of the disks.

Our approximation algorithm works as follows: First we compute an approximate k-disk cover $D := (C,\ (r_v)_{v \in C})$ over P. Then we determine for the center points in C an approximate broadcast with range assignment r_B by using an MST based algorithm (see [1]) that has an approximation guarantee of 6. Now we construct a range assignment r_A for P in the following way:

$$\forall v \in P: \quad r_A(v) := \begin{cases} \max\{r_v,\ r_B(v)\} & \text{if } v \in C \\ 0 & \text{otherwise} \end{cases}$$

Note that $G^{(r_A)}$ is connected and therefor induces a valid k-set broadcast since only k stations are sending. We still have to show that we have computed an approximate solution:

Theorem 3. $\text{cost}(r_A) \leq 36 c_\alpha \cdot \text{cost}(r_{opt})$, *where* r_{opt} *is the range assignment of an optimal* k*-set broadcast and* c_α *is a constant depending only on* α.

Proof. The proof idea is as follows: Assuming knowledge about an optimal range assignment r_{opt} for the k-set broadcast we transform the range assignment r_D into r'_D such that

a) r'_D is a valid k-set broadcast
b) the sending nodes in r'_D are exactly the center points of D and
c) r'_D is a constant factor approximation of r_{opt}.

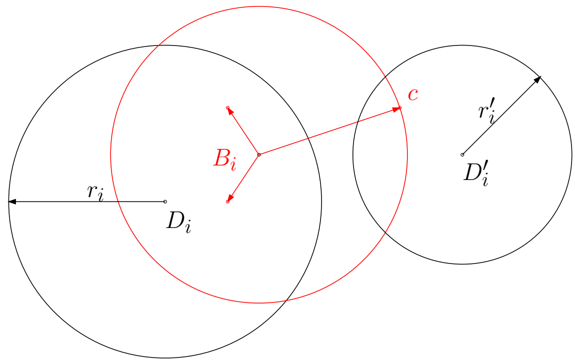

Fig. 2. Proof illustration for the constant factor approximation algorithm

If we know that such a broadcast r'_D exists, we can simply compute an optimal broadcast r_B over the center points of D. Then we know that $\text{cost}(r_B) \leq \text{cost}(r'_D)$ and r_B must also be a constant factor approximation of r_{opt}.

Consider now the communication tree T which is defined as a subtree of $G^{(r_{opt})}$ spanning P. The idea of the construction of r'_D is to replace the inner nodes of T (i.e. the sending stations of r_{opt}) by increasing the radii of the disks in C appropriately so that r'_D is valid.

We increase the nonzero values of r_D in the following way: with each of the inner nodes B_i of T we associate arbitrarily one disk D_i in which B_i is contained. Note that there must be at least one such disk for each B_i since the disks cover P. We now update r_D in a breath first search manner on T starting from source node s (see figure 2):

Given an inner node B_i of T if all children of B_i in T lie in the associated disk D_i then all of them can be reached from node C_i without increasing r_i. The interesting case is if there are children of B_i that are not contained in D_i but contained in a disk D'_i whose center is not covered so far. Assume that there is exactly one such child c. We then set the radius of D_i to $r_i + r'_i + r_{opt}(B_i)$. If there is more than one such child, let c be the one that maximizes r'_i so that each child of B_i and the centers of the disks in which the children of B_i are contained in can be reached by D_i. Note that it can happen that two different inner nodes B_i and B_j are associated with the same disk D_k, so that D_k is updated more than once in the process. In such a case we update D_k only if r_k is increasing. Now let us assume that for a disk D_i the last update involved disk D'_k then we call disk D'_k the target disk of D'_k.

By induction $G^{(r_D)}$ is connected after these updates. furthermore note that the sending stations are still exactly the center points of D. Let $D^* \subseteq D$ be the set of disks that are updated and let B_i be the node in T in the update step for disk $D_i \in D^*$. Summing over all disks, the total cost of the broadcast is therefor bounded by:

$$\underbrace{\sum_{D_i \in D \setminus D^*} r_i^\alpha}_{\leq \text{cost}(D) \leq 2\,\text{cost}(r_{opt})} + \underbrace{\sum_{D_i \in D^*} (r_i + r_i' + r_{opt}(B_i))^\alpha}_{(**)}$$

Before we bound the second term, note that a disk appears as a target disk only once in the process of updating the disk radii since once its center point is covered it is never considered as a target disk again. Thus each r_i' in the above sum can also appear only once. Thus,

$$\begin{aligned}(**) &\leq c_\alpha \Big[\sum_{D_i \in D^*} r_i^\alpha + \sum_{D_i \in D^*} {r_i'}^\alpha + \\ &\qquad \sum_{D_i \in D^*} r_{opt}(B_i)^\alpha \Big] \\ &\leq c_\alpha \Big[\underbrace{2 \sum_{D_i \in D} r_i^\alpha}_{\leq 4\,\text{cost}(r_{opt})} + \underbrace{\sum_{D_i \in D} r_{opt}(B_i)^\alpha}_{=\text{cost}(r_{opt})} \Big] \\ &\leq 5 \cdot c_\alpha \cdot \text{cost}(r_{opt})\end{aligned}$$

where the constant c_α can be bounded by $3^{\alpha-1}$. Thus there exists a broadcast over the center points C with total cost upper bounded by

$$\begin{aligned}&2\,\text{cost}(r_{opt}) + 5 \cdot c_\alpha \cdot \text{cost}(r_{opt}) \\ \leq\ &6 \cdot c_\alpha \cdot \text{cost}(r_{opt})\end{aligned}$$

Since we use a 6-approximate broadcast, the algorithm has an approximation ratio of $36c_\alpha$.

Theorem 4. *There exists a constant factor approximation algorithm for the k-set broadcast problem over n points in the Euclidean plane that runs in $O(n^{c'_\alpha})$.*

The theorem can be further improved by using the results of the previous section. By setting ϵ to 1 we obtain a synopsis of size k^2. Using theorem 4 we obtain directly a constant factor approximation algorithm whose running time is only linear in n and polynomial in k:

Theorem 5. *There exists a constant factor approximation algorithm for the k-set broadcast problem over n points in the Euclidean plane that runs in time linear in n and polynomial in k, i.e. in $O(n + k^{2 \cdot c'_\alpha})$.*

4 Future Work

4.1 Simple, Distributed Algorithms with Good Performance

In the introduction we have pointed out why we believe that a broadcast operation with a bounded number of senders can be of great benefit for the efficient operation of a wireless network. The main part of the paper considered the k-SEMBC problem from a purely analytical point of view, though. The most important research direction to follow in the near future is to design simple and *distributed* algorithms for the k-SEMBC problem. One possible idea could be to construct distributedly a low-weight k-disk-cover of the network and then connect the components in some way such that the overall cost does not increase by too much – very much in the spirit of our $O(1)$-approximation algorithm. We also plan to look at the k-SEMBC problem from a more empirical point of view. Several heuristic solutions could be thought of and examined using extensive simulations on different network deployments.

4.2 Extension to Metrics of Bounded Doubling Dimension

While actual network deployments often are in a planar setting, the experienced metric for several reasons is typically not exactly of the Euclidean type (due to obstacles, interference etc.), but often in some sense 'close' as it still retains some correlation between geographic distance and distance in the metric. One way to measure similarity between metrics is the so called *doubling dimension* [10]. It remains to examine whether our algorithms also work on metrics of bounded doubling dimension.

4.3 $(1+\epsilon)$-Approximation with Running Time Polynomial in k (and $1/\epsilon$) ?

One major drawback of the $(1+\epsilon)$-approximation algorithm presented is that it still requires time exponential in the number of senders k. It is unclear whether this exponential dependence could be removed. One idea might be to use the approach by Arora ([3]) based on a shifted dissection which was already used successfully to obtain fast approximation algorithms for geometric problems. This might induce an exponential dependence on $1/\epsilon$, though.

References

1. Ambühl, C.: An optimal bound for the mst algorithm to compute energy-efficient broadcast trees in wireless networks. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, Springer, Heidelberg (2005)
2. Ambühl, C., Clementi, A.E.F., Di Ianni, M., Lev-Tov, N., Monti, A., Peleg, D., Rossi, G., Silvestri, R.: Efficient algorithms for low-energy bounded-hop broadcast in ad-hoc wireless networks. In: Diekert, V., Habib, M. (eds.) STACS 2004. LNCS, vol. 2996, Springer, Heidelberg (2004)
3. Arora, S.: Approximation schemes for geometric np-hard problems: A survey. In: Hariharan, R., Mukund, M., Vinay, V. (eds.) FST TCS 2001: Foundations of Software Technology and Theoretical Computer Science. LNCS, vol. 2245, Springer, Heidelberg (2001)

4. Bilò, V., Caragiannis, I., Kaklamanis, C., Kanellopoulos, P.: Geometric clustering to minimize the sum of cluster sizes. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, Springer, Heidelberg (2005)
5. Clementi, A., Huiban, G., Penna, P., Rossi, G., Verhoeven, Y.: Some recent theoretical advances and open questions on energy consumption in ad-hoc wireless networks. In: Proc. 3rd Workshop on Approximation and Randomization Algorithms in Communication Networks (ARACNE), pp. 23–38 (2002)
6. Clementi, A.E.F., Crescenzi, P., Penna, P., Rossi, G., Vocca, P.: On the complexity of computing minimum energy consumption broadcast subgraphs. In: Ferreira, A., Reichel, H. (eds.) STACS 2001. LNCS, vol. 2010, pp. 121–131. Springer, Heidelberg (2001)
7. Clementi, A.E.F., Huiban, G., Rossi, G., Verhoeven, Y.C., Penna, P.: On the approximation ratio of the mst-based heuristic for the energy-efficient broadcast problem in static ad-hoc radio networks. In : IPDPS, p. 222 (2003)
8. Flammini, M., Navarra, A., Klasing, R., Pérennes, S.: Improved approximation results for the minimum energy broadcasting problem. In: DIALM-POMC, pp. 85–91 (2004)
9. Funke, S., Laue, S.: Bounded-hop energy-efficient broadcast in low-dimensional metrics via coresets. In: Thomas, W., Weil, P. (eds.) STACS 2007. LNCS, vol. 4393, pp. 272–283. Springer, Heidelberg (2007)
10. Gupta, A., Krauthgamer, R., Lee, J.R.: Bounded geometries, fractals, and low-distortion embeddings. In: FOCS, pp. 534–543 (2003)
11. Har-Peled, S., Mazumdar, S.: Coresets for k-means and k-median clustering and their applications. In: STOC, pp. 291–300 (2004)
12. Kirousis, L.M., Kranakis, E., Krizanc, D., Pelc, A.: Power consumption in packet radio networks. Theor. Comput. Sci. 243(1-2), 289–305 (2000)
13. Kumar, P., Mitchell, J.B., Yildirim, E.A.: Approximate minimum enclosing balls in high dimensions using core-sets. J. Exp. Algorithmics, vol. 8(1.1) (2003)
14. Guha, S., Khuller, S.: Improved methods for approximating node weighted steiner trees and connected dominating sets. Information and Computation 150, 57–74 (1999)
15. Wan, P.-J., Calinescu, G., Li, X., Frieder, O.: Minimum-energy broadcast routing in static ad hoc wireless networks. In: INFOCOM, pp. 1162–1171 (2001)
16. 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)