An Improved Approximation Algorithm for the Minimum Energy Consumption Broadcast Subgraph

Vittorio Bilò and Giovanna Melideo

Dipartimento di Informatica, Università di L'Aquila,
Via Vetoio loc. Coppito, I-67100 L'Aquila, Italy
{bilo,melideo}@di.univaq.it

Abstract. In an ad-hoc wireless network each station has the capacity of modifying the area of coverage with its transmission power. Controlling the emitted transmission power allows to significantly reduce the energy consumption and so to increase the lifetime of the network. In this paper we focus on the Minimum Energy Consumption Broadcast Subgraph (MECBS) problem [1, 2, 6], whose objective is that of assigning a transmission power to each station in such a way that a message from a source station can be forwarded to all the other stations in the network with a minimum overall energy consumption. The MECBS problem has been proved to be inapproximable within $(1 - \epsilon) \ln n$ unless $NP \subseteq DTIME(n^{O(\log \log n)})$ [2, 6], where n is the number of stations. In this work we propose a $2H_{n-1}$-approximation greedy algorithm which, despite its simplicity, improves upon the only previously known ratio of $10.8 \ln n$ [1] and considerably approaches the best-known lower bound on the approximation ratio.

1 Introduction

Ad hoc wireless networks have received significant attention during the recent years. In particular, they emerged due to their potential applications in emergency disaster relief, battlefield, etc [5, 7]. Unlike traditional wired networks or cellular networks, they do not require the installation of any wired backbone infrastructure. The network is a collection of transmitter/receiver stations each equipped with an omnidirectional antenna which is responsible for sending and receiving radio signals. A communication is established by assigning to each station a transmitting power.

A fundamental problem in ad hoc wireless networks is to support *broadcasting*, that is to allow a source station to transmit a message to all stations in the network. A communication from a station s to another t occurs either through a single-hop transmission if the transmitting power of s is adequate, or through relaying by intermediate stations, otherwise. One of the main advantages of ad-hoc networks is the ability of the stations to vary the power used in a transmission in order to reduce the power consumption and so to increase the lifetime of the network.

M. Danelutto, D. Laforenza, M. Vanneschi (Eds.): Euro-Par 2004, LNCS 3149, pp. 949–956, 2004.

In this paper we focus on the "energy-efficient" broadcasting, where the objective is to designate the transmission powers at which each station i has to transmit in such a way that a communication from a source station s to all the other stations can be established and the overall energy consumption is minimized. This problem is referred to as *Minimum Energy Consumption Broadcast Subgraph* (MECBS, for short) [1, 2].

An ad hoc wireless network is usually modelled by (i) a complete graph $G(S)$ whose n vertices $S = \{1, \dots, n\}$ represent radio stations and (ii) a symmetric *cost function* $c : S \times S \mapsto I\!R^+$ which associates each pair of stations i and j with its *transmission cost*, that is the power necessary for exchanging messages between i and j. Clearly, $c(i, i) = 0$ for every station $i \in S$. A power assignment $\omega : S \mapsto I\!R^+$ to the stations induces a directed weighted graph $G_\omega = (S, E)$, called *the transmission graph*, such that an edge $\langle i, j\rangle$ belongs to E if and only if the transmission power of i is at least equal to the transmission cost from i to j, i.e., $\omega(i) \geq c(i, j)$. The *cost* of a power assignment ω is the overall power consumption yielded by ω, i.e., $cost(\omega) = \sum_{i=1}^{n} \omega(i)$.

The *Minimum Energy Consumption Broadcast Subgraph* (MECBS) problem described above is then defined as follows:

– *Input.* A set S of n sender/receiver stations and a source station $s \in S$.
– *Output.* A power assignment ω such that the overall energy consumption $cost(\omega)$ is minimized and the induced transmission graph G_ω contains a directed spanning tree rooted at s.

The MECBS problem has been proved to be inapproximable within $(1-\epsilon)\ln n$ unless $NP \subseteq DTIME(n^{O(\log\log n)})$ [2, 6] where n is the number of stations. The only known logarithmic approximation algorithm for the problem is due to Caragiannis et al. who presented a $10.8 \ln n$-approximation algorithm which uses a reduction to the Node-Weighted Connected Dominating Set problem [1].

This paper takes an important step toward the reduction of the gap between these two bounds. Indeed, we propose a $2H_{n-1}$-approximation greedy algorithm for the problem that is returning assignments of cost at most twice the cost of the best possible approximated solution. Besides the reduction of the existing gap, our algorithm confirms the strict relationship occurring between the MECBS problem and the Set Cover problem. In fact, as also suggested by the negative results in [2], our result reinforces the intuition that MECBS can be seen as a set covering problem with additional connectivity requirements.

The paper is organized as follows. In section 2 we first introduce the necessary notations and definitions, then we describe our algorithm. Section 3 is dedicated to the correctness proof and to the performance analysis of the proposed algorithm, and, finally, Section 4 presents future directions.

2 An Improved Logarithmic Approximation Algorithm

In this section we present our logarithmic approximation algorithm which, given (i) the complete weighted graph $G(S)$ with vertices $S = \{1, \dots, n\}$ representing stations and with a cost function $c : S \times S \mapsto I\!R^+$ which associates each pair of stations i and j with its transmission cost, and (ii) a vertex s, returns a power assignment $\omega : S \mapsto I\!R^+$ which induces a transmission graph G_ω containing a directed spanning tree rooted at s.

Before presenting the algorithm, we introduce a few preliminary notations.

Notations. For any station i, let $\mathcal{C}_i = \bigcup_{j \in S}\{c(i,j)\}$ denote the set of all the possible level powers allowing i to transmit to any other station j in the network through a single-hop transmission. It is important to notice that even if, by definition, for any station i, the transmission power $\omega(i)$ can range in the interval $[0, \infty]$, the set of possible power assignments is actually the discrete set $\mathcal{C}_i$ of at most $n - 1$ values.

For any transmission power $\varpi \in \mathcal{C}_i$ assigned to a station i, $S_\varpi(i) = \{j \in S \mid j \neq i \text{ and } \varpi \geq c(i,j)\}$ is the set of all the stations j distinct from i into the area of coverage of i, and $\mathcal{P}_\varpi(i) = \{X \in \mathcal{P} \mid X \cap S_\varpi(i) \neq \emptyset\}$ is the family of sets belonging to a given partition $\mathcal{P}$ and containing stations covered by i.

Finally, given a power assignment ω associating a power transmission $\omega(i)$ to each station i, $\mathit{\Pi}(s, G_\omega)$ is the set of vertices i connected to a vertex s through a directed path from s to i, while $G_\omega[X]$ is the subgraph of G_ω induced by a set $X \subseteq S$.

2.1 The Greedy-Assignment Algorithm

We propose the "Greedy-Assignment" algorithm, which uses a greedy approach for solving the MECBS problem. Indeed, in order to compute a power assignment ω, it initially assigns a transmission power zero to every station in S; then it iteratively uses a specific greedy rule to select a station and increases its transmission power until ω induces on S a weakly connected transmission graph G_ω. Finally, this assignment ω is re-adjusted in order to grow the directed spanning tree rooted at s, one station at a time, without destroying any connection previously established.

The algorithm works as follows. The loop 3 iteratively manages a partition $\mathcal{P}$ of stations initially composed of n singleton sets and a power assignment ω identically zero at the beginning. At each stage, $\mathcal{P}$ represents all the weakly connected components of G_ω. This invariant is guaranteed by the "Merge" procedure which, given the pair $\langle i, \varpi \rangle$ chosen at line 3.(a), merges all the sets in $\mathcal{P}_\varpi(i)$ together with $S_\varpi(i) \cup \{i\}$ in such a way that, after line 3.(b), their union induces on G_ω a new weakly connected subgraph. At the end of loop 3, all the sets contained in the partition $\mathcal{P}$ are merged together, denoting that the transmission graph G_ω is weakly connected.

Line 3.(a) is the greedy decision-making step: a station i and a transmission power $\varpi \in \mathcal{C}_i$ are chosen such that $|\mathcal{P}_\varpi(i)|$ components of G_ω are connected with i with a minimum average power consumption.

Definition 1. *Given a partition $\mathcal{P}$ of S, a station $i \in S$ and a transmission power $\varpi \in \mathcal{C}_i$, we define the* cost-effectiveness $\varepsilon_\varpi(i, \mathcal{P})$ of ϖ at i, with respect to $\mathcal{P}$, *as the average transmission power at which i covers sets in $\mathcal{P}_\varpi(i)$, i.e.,*

$$\varepsilon_\varpi(i,\mathcal{P}) = \frac{\varpi}{|\mathcal{P}_\varpi(i)|}.$$

If $\mathcal{P}_\varpi(i) = \emptyset$ *we set* $\varepsilon_\varpi(i,\mathcal{P}) = \infty$.

According to Def. 1, at line 3.(a) a pair $\langle i, \varpi\rangle$ is chosen such that, after setting $\omega(i) = \varpi$ (see line 3.(b)), the cost-effectiveness $\varepsilon_\varpi(i,\mathcal{P})$ is minimum. Then all the components in $\mathcal{P}_\varpi(i)$ are merged together with $S_\varpi(i) \cup \{i\}$ at line 3.(c).

Finally, The loop 5 re-adjusts the transmission powers of stations i not connected to s through a directed path, without destroying any connection previously established, until the transmission graph contains a directed spanning tree rooted at s.

Greedy-Assignment $(G(S), s)$

1. $\mathcal{P} \leftarrow \{\{1\},\{2\},...,\{n\}\}$
2. $\omega(i) \leftarrow 0$ for every station $i \in S$.
3. **while** $|\mathcal{P}| > 1$ **do**
   (a) Choose the pair $\langle i, \varpi\rangle$ with the lowest cost-effectiveness $\varepsilon_\varpi(i,\mathcal{P})$, for $i \in S$ and $\varpi \in \mathcal{C}_i$
   (b) $\omega(i) \leftarrow \varpi$
   (c) $\mathcal{P} \leftarrow$ Merge $(G(S), \mathcal{P}, i, \varpi)$
4. $\nu = \omega$
5. **while** $\Pi(s, G_\omega) \neq S$ **do**
   (a) choose a vertex $i \in \Pi(s, G_\omega)$ and a vertex $j \notin \Pi(s, G_\omega)$ such that there exists a directed edge $\langle j, i\rangle$ in G_ν
   (b) $\omega(i) \leftarrow \nu(j)$
6. **return** ω

Merge $(G(S), \mathcal{P}, i, \varpi)$

– $S^+_\varpi(i) \leftarrow (\bigcup_{X\in\mathcal{P}_\varpi(i)} X) \cup S_\varpi(i) \cup \{i\}$
– $\mathcal{P}^-_\varpi(i) \leftarrow \mathcal{P}_\varpi(i) \cup \{S_\varpi(i) \cup \{i\}\}$
– **return** $\mathcal{P} \setminus \mathcal{P}^-_\varpi(i) \cup \{S^+_\varpi(i)\}$

The correctness proof of the algorithm is based on the two following invariant properties.

2.2 Invariant Properties

Let us denote by $\mathcal{P}^{(k)}$ the partition of stations at the end of the k-th iteration of loop 3 and by $\omega_1^{(k)}$ and $\omega_2^{(k)}$ the power assignments returned at then end of the k-th iteration of loop 3 and loop 5, respectively.

The first important property states that, although different transmission powers $\varpi_1, \ldots, \varpi_m$ may be assigned to a same station i, such a sequence is increasing. This guarantees that re-adjustments (see lines 3.(b) and 5.(b)) never destroy connections previously established.

Property 1. For any station $i \in S$ and for any constant k, l, m, q such that $k < l$ and $m < q$ it holds: $\omega_1^{(k)}(i) \leq \omega_1^{(l)}(i) \leq \omega_2^{(m)}(i) \leq \omega_2^{(q)}(i)$.
Proof. Initially, $\omega(i) = \omega_1^{(0)}(i) = 0$ (line 2). Let $\omega(i)$ be set to ϖ at the k-th iteration of loop 3, i.e., $\omega_1^{(k)}(i) = \varpi$. The "MERGE" procedure (line 3.(c)) makes the union of all the sets in $\mathcal{P}_\varpi^{(k)}(i)$ and $S_\varpi(i) \cup \{i\}$. Hence, since for every $\varpi' \in \mathcal{C}_i$ such that $\varpi' \leq \varpi$ it holds $S_{\varpi'}(i) \subset S_\varpi(i)$, when $h > k$ it must be $\mathcal{P}_{\varpi'}^{(h)}(i) = \emptyset$ and, as a consequence, $\varepsilon_{\varpi'}(i, \mathcal{P}^{(h)}) = \infty$. This implies that it cannot be $\omega_1^{(k)}(i) > \omega_1^{(h)}(i)$ for every $h > k$. Let $\omega_2^{(0)}$ denote the assignment ω $(= \nu)$ at line 4. Now, consider the m-th iteration of loop 5, for any $m \geq 1$. At line 5.(b) the transmission power $\nu(j)$ is assigned to i only if the edge $\langle j, i\rangle \in G_\nu$ but $\langle i, j\rangle \notin G_{\omega_2^{(m)}}$. Then it must be $\nu(j) > \omega_2^{(m)}(i)$. □

Now we are able to prove that, for any k, at the end of the k-th iteration of loop 3, the partition $\mathcal{P}^{(k)}$ always contains sets of stations that induce on $G_{\omega_1^{(k)}}$ weakly connected subgraphs. In other words, $\mathcal{P}^{(k)}$ represents all the weakly connected components of $G_{\omega_1^{(k)}}$.

Property 2. For each set $X \in \mathcal{P}^{(k)}$, $G_{\omega^{(k)}}[X]$ is weakly connected.
Proof. This is obviously true for $k = 0$. For any $k \geq 1$, at the k-th iteration of loop 3, a pair $\langle i, \varpi\rangle$ is chosen in such a way to increase to ϖ the transmission power associated with i. By construction, each set in $\mathcal{P}_\varpi(i)$ contains at least one station belonging to $S_\varpi(i)$, so $S_\varpi^+(i)$ necessarily induces on the transmission graph $G_{\omega_1^{(k)}}$ a weakly connected subgraph. □

3 Correctness and Performance Analysis

In the following section we provide a correctness proof of our algorithm, by showing that the power assignment ω returned by it actually induces a transmission graph G_ω which contains a directed spanning tree rooted at a source station s. Due to space limitations, a few details are omitted. Finally, we show that it has an approximation ratio of $2H_{n-1}$.

Theorem 1 (Correctess). *The output ω returned by the* GREEDY-ASSIGNMENT *algorithm is feasible.*
Proof. The proof is made of two parts: first we prove that ω at the end of loop 3 (i.e., ν) induces on S a weakly connected transmission graph; then we show that the graph G_ω induced by the output ω admits a directed spanning tree rooted at s. The former is a trivial consequence of Prop. 2. By Prop. 1, no re-adjustment at line 5.(b) can destroy connections previously established, then the latter is clearly true if the algorithm terminates, as it is guaranteed by the while condition of loop 5. Thus, we only need to prove that loop 5 always terminates. At a generic iteration of loop 5, if $\Pi(s, G_\omega) \neq S$, since G_ν is weakly connected, there must exist two vertices i and j satisfying the conditions in line 5.(a). The claim follows by noting that adding the power assignment in 5.(b), by Prop. 1, we have that the set $\Pi(s, G_\omega)$ expands. □

Given a sequence of pairs $\mathcal{F} = \langle i_1, \varpi_1 \rangle, \ldots, \langle i_r, \varpi_r \rangle$, let $cost(\mathcal{F}) = \sum_{j=1}^{r} \varpi_j$. Each transmission power ϖ_j can be spread among the sets in $\mathcal{P}_{\varpi_j}(i_j)$ by assigning each set a cost equal to the cost-effectiveness of ϖ_j at i_j, with respect to $\mathcal{P}^{(j)}$. That is, for every $X \in \mathcal{P}_{\varpi_j}(i_j)$, $cost(X) = \varepsilon_{\varpi_j}(i_j, \mathcal{P}^{(j)})$. If $\mathcal{T}_{\mathcal{F}}$ denote the family of sets $\mathcal{T}_{\mathcal{F}} = \bigcup_{j=1}^{r} \mathcal{P}_{\varpi_j}(i_j)$, one has: $cost(\mathcal{F}) = \sum_{j=1}^{r} \varpi_j = \sum_{X \in \mathcal{T}_{\mathcal{F}}} cost(X)$. Observe that if we consider the sequence $\mathcal{A}$ of the pairs chosen by our algorithm, clearly $|\mathcal{T}_{\mathcal{A}}| = n - 1$. Number these sets in the order in which they are merged by the algorithm breaking ties arbitrarily.

Let OPT be the cost of an optimal solution ω^*. At any iteration k_i when T_i is merged the following lemma holds.

Lemma 1. *For any $i \leq n-1$ there exists a pair $\langle j, \varpi \rangle$ having a cost-effectiveness $\varepsilon_{\varpi}(j, \mathcal{P}^{(k_i)}) \leq \frac{OPT}{n-i}$.*

Proof. Let $\mathcal{F}^{(i)}$ be the sequence of the pairs chosen by the optimal algorithm such that $\omega^*(j) > \omega_1^{(k_i)}(j)$. Clearly, $OPT \geq cost(\mathcal{F}^{(i)}) = \sum_{j | \langle j, \varpi \rangle \in \mathcal{F}^{(i)}} \omega^*(j)$. Now, suppose by contradiction that for every pair $\langle j, \varpi \rangle \in \mathcal{F}^{(i)}$, it holds $\varepsilon_{\varpi}(j, \mathcal{P}^{(k_i)}) > \frac{OPT}{n-i}$. Since by definition, the cost-effectiveness of a pair $\langle j, \varpi \rangle$ cannot decrease as i increases and $|\mathcal{T}_{\mathcal{F}^{(i)}}| = n-i$, we have that $OPT \geq cost(\mathcal{F}^{(i)}) > (n-i)\frac{OPT}{n-i} = OPT$. □

Theorem 2 (Performance Guarantee). *The* Greedy-Assignment *algorithm has an approximation ratio of* $2H_{n-1}$.

Proof. We first show that $cost(\nu) \leq H_{n-1} \cdot OPT$, then that $cost(\omega) \leq 2 \cdot cost(\nu)$.

Since the algorithm chooses the power assignment with the lowest cost-effectiveness, from the above lemma we have that $cost(T_i) \leq \frac{OPT}{n-i}$. Thus, the solution at line 4 is such that $cost(\nu) = \sum_{i=1}^{|\mathcal{T}|} cost(T_i) = \sum_{i=1}^{n-1} cost(T_i) \leq \sum_{i=1}^{n-1} \frac{OPT}{n-i} = OPT \sum_{i=1}^{n-1} \frac{1}{i} = H_{n-1} \cdot OPT$.

At any iteration of loop 5 a transmission power $\nu(j)$ is assigned to i only if the conditions of line 5.(a) are enjoyed by the pair of stations i and j. Such an assignment establishes a directed path from s to j, then j can be chosen at step 5.(a) at most once. Consequently, $cost(\omega) \leq 2 \cdot cost(\nu)$. □

Notice that the approximation ratio $2 \cdot H_{n-1}$ is essentially tight, as shown by the following example.

Example. When the Greedy-Assignment algorithm runs on the input graph in Fig. 1.a) it computes the partial solution ν depicted in Fig. 1.c). Thus, in the case in which n is an odd number:

$$cost(\nu) = \sum_{j=0}^{\lceil \frac{n-1}{2} \rceil - 1} \frac{1}{\lceil \frac{n-1}{2} \rceil - j} = \sum_{j=1}^{\lceil \frac{n-1}{2} \rceil} j = H_{\lceil \frac{n-1}{2} \rceil} = H_{n-1} - 1.$$

Next, loop 5 exactly doubles all the transmission powers, hence $cost(\omega) = 2H_{n-1} - 2$. Since the cost of the optimal solution in Fig. 1.b) is $1 + \epsilon$, we have that the approximation factor of our algorithm is tight up to low order terms.

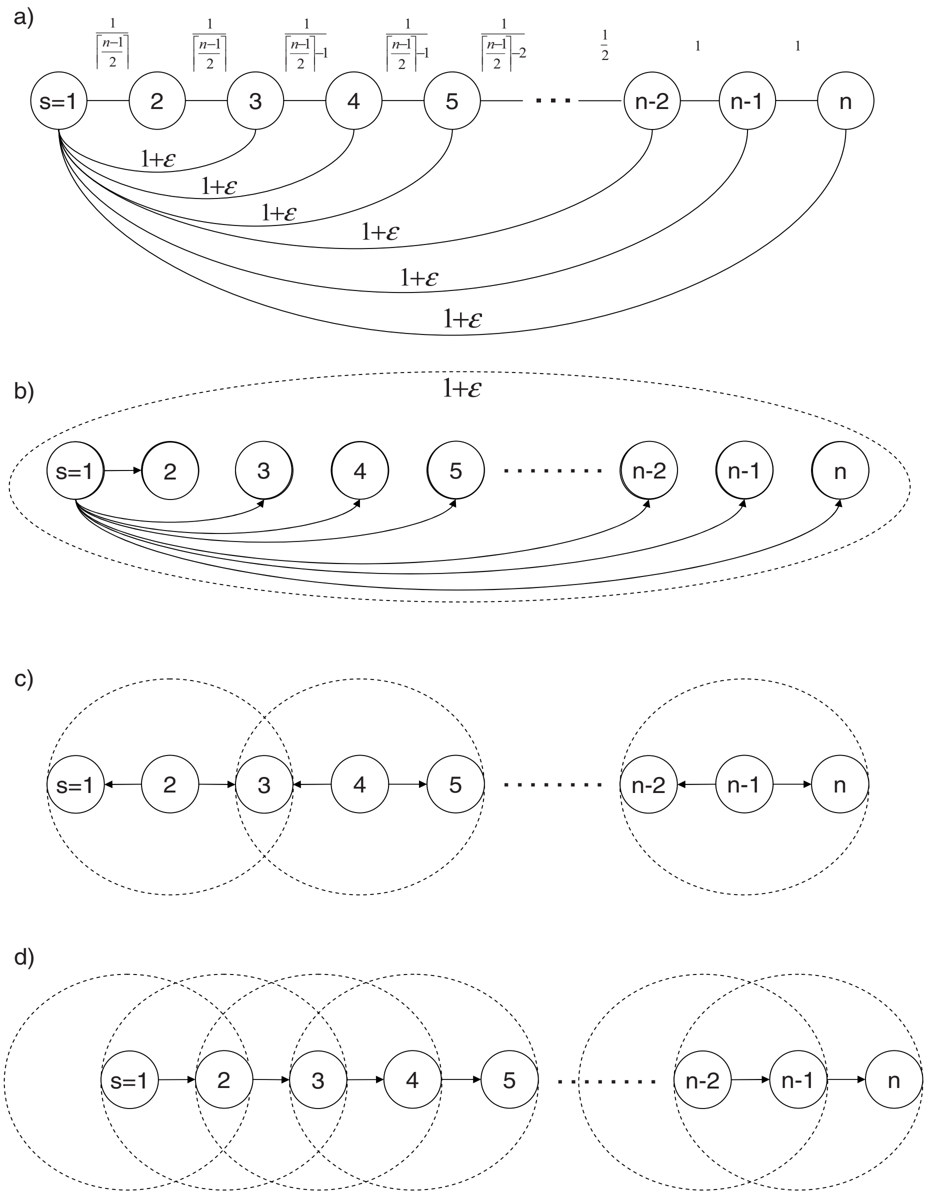

Fig. 1. a) The input graph. Missing edges are assumed to have an infinite cost. b) The optimal solution. c) The partial solution ν. d) The solution ω.

4 Future Directions

We presented a $2H_{n-1}$-approximation algorithm for the MECBS problem [1, 2, 6], which, despite its simplicity, improves the only known algorithm with a ratio of $10.8 \log n$ [1] and considerably approaches the best-known lower bound of $\log n$ on the approximation ratio [2, 6]. We leave the natural open problem of bridging this gap. Moreover, once broadcasting models are developed for node-based models, future studies could address the impact of mobility and limited resources (both bandwidth and equipment). A significant restriction of the MECBS problem, denoted MECBS[$\mathtt{N}_d^\alpha$], consists in considering the stations located in the d-dimensional Euclidean Space, and a cost function s.t. $c(i,j) = dist(i,j)^\alpha$, where α is the *distance-power gradient*. It has been proved that the MECBS[$\mathtt{N}_d^\alpha$] problem is NP-Hard for $\alpha > 1$ and $d > 1$, while it is in P if $\alpha = 1$ or $d = 1$ [1–3, 6]. The best known approximation algorithm, called MST, has been presented and compared with other heuristics (SPT, BIP) through simulations on random instances in the case $d = \alpha = 2$ [7]. Its performance has been investigated by several authors [2, 6, 4] and the evaluation of its approximation ratio progressively reduced till $3^d - 1$ for every $\alpha \geq d$ [4]. Starting from these results, a further left open question is that of comparing through simulations our algorithm with MST, SPT and BIP when restricting to the special case of MECBS[$\mathtt{N}_d^\alpha$].

Acknowledgements

The authors would like to thank Michele Flammini for several useful discussions.

References

1. I. Caragiannis, E. Kaklamanis, and P. Kanellopoulos. A logarithmic approximation algorithm for the minimum energy consumption broadcast subgraph problem. *Information Processing Letters*, 86:149–154, 2003.
2. A.E.F. Clementi, P. Crescenzi, P. Penna, G. Rossi, and P. Vocca. On the complexity of computing minimum energy consumption broadcast subgraph. In *Proceedings of the 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2010 of *Lecture Notes in Computer Science*, pages 121–131. Springer-Verlag, 2001.
3. A.E.F. Clementi, M. Di Ianni, and R. Silvestri. The minimum broadcast range assignment problem on linear multi-hop wireless networks. *Theoretical Computer Science*. to appear.
4. M. Flammini, R. Klasing, A. Navarra, and S. Perennes. Improved approximation results for the minimum energy broadcasting problem in wireless networks. Manuscript, 2003.
5. T.S. Rappaport. *Wireless communications: principles and practice*. Prentice-Hall, Englewood Cliffs, NY, 1996.
6. P. J. Wan, G. Calinescu, X. Li, and O. Frieder. Minimum energy broadcast routing in static ad hoc wireless networks. *Wireless Networks*, 8(6):607–617, 2002.
7. J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. On the construction of energy-efficient broadcast and multicast trees in wireless networks. In *Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 585–594. IEEE Computer Society, 2000.