An Exponential Improvement on the MST Heuristic for Minimum Energy Broadcasting in Ad Hoc Wireless Networks*

Ioannis Caragiannis[1], Michele Flammini[2], and Luca Moscardelli[2]

[1] Research Academic Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece
[2] Dipartimento di Informatica, Università di L'Aquila
Via Vetoio, Coppito 67100, L'Aquila, Italy

Abstract. In this paper we present a new approximation algorithm for the *Minimum Energy Broadcast Routing* (MEBR) problem in ad hoc wireless networks that has exponentially better approximation factor than the well-known Minimum Spanning Tree (MST) heuristic. Namely, for any instance where a minimum spanning tree of the set of stations is guaranteed to cost at most ρ times the cost of an optimal solution for MEBR, we prove that our algorithm achieves an approximation ratio bounded by $2 \ln \rho - 2 \ln 2 + 2$. This result is particularly relevant for its consequences on Euclidean instances where we significantly improve previous results.

1 Introduction

Over the last years the usage of *wireless networks* has seen a huge increase mostly because of the recent drop in equipment prices and due to the features provided by the new technologies. In particular, considerable attention has been devoted to the so-called *ad hoc* wireless networks, due to their potential applications in emergency disaster relief, battlefield, etc. [16, 25]. Ad hoc networks do not require any fixed infrastructure. The network is simply a collection of homogeneous radio stations equipped with omnidirectional antennas for sending and receiving signals. Communication occurs by assigning to each station a transmitting power. In the most common power attenuation model [22, 24], the signal power P_s of a station s decreases as a function of the distance in such a way that at any station t at distance $dist(s,t)$ it is received with a power $\frac{P_s}{dist(s,t)^\alpha}$, where $\alpha \geq 1$ is a constant called the *distance-power gradient.* While for practical purposes the constant α is usually assumed to be between 2 and 5, from a theoretical point of view general values of α are also of interest. The signal is correctly received by t if $\frac{P_s}{dist(s,t)^\alpha} \geq \beta$, where $\beta \in \mathbb{R}^+$ is the *transmission*

* This work was partially supported by the EU COST Action 293 "Graphs and Algorithms in Communication Networks" (GRAAL) and by the EU IST FET Integrated Project 015964 AEOLUS.

quality threshold. Therefore, if s correctly transmits within a given maximum distance $r(s)$, called the range of s, the transmission power P_s of s is at least $\beta \cdot r(s)^\alpha$. Usually, the transmission quality threshold β is normalized to 1. We remark that, due to the nonlinear power attenuation, multi-hop transmission of messages through intermediate stations may result in energy saving.

A naturally arising issue in ad hoc wireless networks is that of supporting communication patterns that are typical in traditional networks, such as broadcasting (one-to-all), multicasting (one-to-many), and gossiping (all-to-all), with a minimum total energy consumption. This problem is generally called *Minimum Energy Routing* (MER) and defines different optimization subproblems according to the connectivity requirements (see for instance [4, 5, 7, 8, 11, 12, 15, 17, 19, 20] for related results). In this paper, we are interested in the broadcast communication from a given source node.

Formally, given a set of stations S, let $G(S)$ be the complete weighted graph whose nodes are the stations of S and in which the weight $w(x, y)$ of each edge $\{x, y\}$ is the power consumption needed for a correct communication between x and y. A power assignment for S is a function $p : S \rightarrow \mathbb{R}^+$ assigning a transmission power $p(x)$ to every station in S. A power assignment p for S yields a directed communication graph $G^p = (S, A)$ such that, for each $(x, y) \in S^2$, the directed edge (x, y) belongs to A if and only if $p(x) \geq w(x, y)$, that is if x can correctly transmit to y. In this case y is also said to fall within the transmission range of x. The total cost of a power assignment is

$$cost(p) = \sum_{x \in S} p(x).$$

The MEBR problem for a given source $s \in S$ consists in finding a power assignment p of minimum cost $m^*(S, s)$ such that G^p contains a directed spanning tree rooted at s (and directed towards the leaves).

In general, the problem is unlikely to have polynomial-time approximation algorithms with approximation ratio $o(\ln n)$ [9], where n is the number of stations. Logarithmic (in the number of stations) approximation algorithms have been presented in [4, 6, 7].

An important case of practical interest is when stations lie in a $d-$dimensional Euclidean space. Then, given a constant $\alpha \geq 1$, the power consumption needed for a correct communication between x and y is $dist(x, y)^\alpha$, where $dist(x, y)$ is the Euclidean distance between the locations of x and y. The problem has been proved to be NP-hard for $\alpha > 1$ and $d > 1$, while it is solvable in polynomial time for $\alpha = 1$ or $d = 1$ [6, 9, 3].

Several attempts in the literature were made to find good approximation algorithms for Euclidean cases. One fundamental algorithm to provide an approximate solution of the MEBR problem is the MST heuristic [24]. It is based on the idea of tuning ranges so as to include a minimum spanning tree of the cost graph $G(S)$. More precisely, denote by $T(S)$ a minimum spanning tree of $G(S)$. The MST heuristic considers $T(S)$ rooted at the source station s, directs the edges of $T(S)$ towards the leaves, and sets the power $p(x)$ of every internal station x of $T(S)$ with $k > 0$ children $x_1, \ldots, x_k$ in such a way that $p(x) = max_{i=1,\ldots,k} w(x, x_i)$. In other words, p is the power assignment of minimum cost inducing the directed tree derived from $T(S)$ and is such that $cost(p) \leq c(T(S))$, where $c(T(S))$ denotes the total cost of the edges in $T(S)$). Therefore, the approximation ratio of the heuristic is bounded by the ratio between the cost of a minimum spanning tree of $G(S)$ and the optimal cost $m^*(S, s)$.

The performance of the MST heuristic has been investigated by several authors [1, 6, 9, 13, 18, 21, 23]. The analysis in all the above papers focuses on the case $\alpha \geq d$ (MST has unbounded approximation ratio when $\alpha < d$) and is based on elegant geometric arguments. The best known approximation ratios are 6 for $d = 2$ [1], 18.8 for $d = 3$ [21] and $3^d - 1$ for every $d > 3$ [13]. Moreover, in [9, 23], a lower bound on the approximation ratio of the MST heuristic has been proven, upper-bounding it by the d-dimensional kissing number n_d, i.e., the maximum number of d-dimensional unit spheres that touch a unit sphere but are mutually non-overlapping (but possibly touch each other). This number is 6 for $d = 2$ (and, hence, the upper bound of [1] is tight), 12 for $d = 3$ and, in general, $n_d = 2^{cd(1+o(1))}$ with $0.2075 \leq c \leq 0.401$ for large d [10]. Despite the considerable research effort in the area during the past years, no algorithm has been theoretically shown so far to outperform the MST heuristic in the Euclidean case, and the improvement of the corresponding ratios is a long standing open question.

Several other heuristics have been shown to perform better than MST in practice, at least for 2-dimensional instances (e.g., see [2, 3, 8, 14, 24]). The most famous among them is probably algorithm BIP (broadcast incremental power [24]). Starting from the source, it builds a tree in steps as follows: at each step, it includes the edge to an uncovered station that requires the minimum increase of power. BIP has been shown to be at least as good as MST while the best known lower bound on its approximation ratio is 4.33 [23]. All other heuristics that seem to work well in practice are either very complicated to analyze or have high lower bounds in terms of their approximation ratio.

In this paper we present a new approximation algorithm for the MEBR problem. For any instance of the problem where the minimum spanning tree of the cost graph $G(S)$ is guaranteed to cost at most ρ times the cost of an optimal solution for MEBR, our algorithm achieves an approximation ratio bounded by $2 \ln \rho - 2 \ln 2 + 2$ if $\rho > 2$ and bounded by ρ if $\rho \leq 2$, which exponentially improves upon the MST heuristic. Surprisingly, our algorithm and analysis does not make use of any geometric arguments and still our results significantly improve the previously best known approximation factor for Euclidean instances of the problem. The corresponding approximation ratio is reduced (when $\alpha \geq d$) from 6 [1] to 4.2 for $d = 2$, from 18.8 [21] to 6.49 for $d = 3$ and in general from $3^d - 1$ [13] to $2.20d + 0.61$ for $d > 3$. In the 2-dimensional case, the achieved approximation is even less than the lower bound on the approximation ratio of the BIP heuristic. In arbitrary (i.e., non-Euclidean) cost graphs, it is not difficult to see that the cost of the minimum spanning tree is at most $n-1$ times the cost of an optimal solution for MEBR; hence, our algorithm also slightly improves the logarithmic approximations of [4, 6, 7]. We also prove that our analysis is tight by showing that there are instances in which the ratio among the cost of the solution returned by the algorithm and the cost of the optimal solution is arbitrarily close to $2\ln\rho - 2\ln 2 + 2$.

The rest of the paper is organized as follows. In Section 2, we describe the new approximation algorithm and, in Section 3, we prove its correctness. In Section 4, we show that our analysis is tight, and, finally, in Section 5, we give some conclusive remarks.

2 The approximation algorithm

In this section, we describe our approximation algorithm. We begin with some necessary definitions.

Given a power assignment p and a given station $x \in S$, let $E(p, x) = \{\{x, y\} | w(x, y) \le p(x)\}$ be the set of the undirected edges induced by p at x, and $E(p) = \bigcup_{x \in S} E(p, x)$ the set of all the undirected edges induced by p.

In the following, for every subset of undirected edges $F \subseteq E$ of a weighted graph $G = (V, E)$, we will denote as $c(F)$ the overall cost of the edges in F, that is the total sum of their weights. For the sake of simplicity, we will identify trees with their corresponding sets of edges. Given an undirected tree T and two nodes u and v of T, $u, v \in V$, let $P(T, u, v) \subseteq T$ be the subset of the edges in the unique path from u to v in T.

A *swap set* for a spanning tree T of an undirected graph $G = (V, E)$ and a set of edges F with endpoints in V is any subset F' of edges that must be removed from the multigraph[3] $T \cup F$ in order to eliminate cycles and so that $T \cup F \setminus F'$ is a spanning tree of G.

We are now ready to describe the algorithm. Before going through the details, let us describe the basic underlying idea. Starting from a spanning tree $T(S)$ of $G(S)$, if the cost of T is significantly higher than the one of an optimal solution for performing broadcasting from a given source $s \in S$, then there must exist a cost efficient *contraction* of T. Namely, it must be possible to set the transmission power $p(x)$ of at least one station x in such a way that $p(x)$ is much lower than the cost of the swap set $A(p, x)$ for $T(S)$ and $E(p, x)$. The algorithm then repeatedly chooses at each step $p(x)$ in such a way that, starting from the current spanning tree, $c(A(p, x))/p(x)$ is maximized. The final tree will be such that, considering the correct orientation of the edges according to the final assignment p, some edges will be in the reverse direction, i.e., from the leaves towards the source s. However, the transmission powers can then be properly set with low additional cost in order to obtain the right orientation from s towards the other stations.

At a given intermediate step of the algorithm in which p and T are the current power assignment and maintained tree, respectively, consider a contraction at a given station x consisting of setting the transmission power of x to $p'(x)$, and let p' be the resulting power assignment. Then, a maximum cost swap set $A(p', x)$

[3] To the purposes of the algorithm, we need to maintain in $T \cup F$ at most two copies of the same edge with different weights.

to be accounted to the contraction can be trivially determined by letting $A(p', x)$ contain the edges that are removed when determining a minimum spanning tree in the multigraph $T \cup E(p', x)$ with the cost of all the edges in $E(p', x)$ set equal to 0. Let the ratio $\frac{c(A(p',x))}{p'(x)}$ be the cost-efficiency of the contraction.

The algorithm then performs the following steps:

- Set the transmission power $p(x)$ of every station in $x \in S$ equal to 0.
- Let $T = T(S)$ be a minimum spanning tree of $G(S)$.
- While there exists at least one contraction of cost-efficiency strictly greater than 2:
  - Perform a contraction of maximum cost-efficiency, and let $p'(x)$ be the corresponding increased power at a given station x, and p' be the resulting power assignment.
  - Set the weight of all the edges in $E(p', x)$ equal to 0.
  - Let $T' = T \cup E(p', x) \setminus A(p', x)$.
  - Set $T = T'$ and $p = p'$.
- Orient all the edges of T from the source s toward all the other stations.
- Return the transmission power assignment p that induces such a set of oriented edges.

Notice that if $\rho \leq 2$ the algorithm performs no contraction step and it returns the initial minimum spanning tree $T(S)$; thus, it guarantees the same approximation ratio ρ of the MST heuristic. Therefore, in the following of the paper we will assume that $\rho > 2$.

3 Correctness of the algorithm

Clearly, the algorithm has a running time polynomial in the size of the input instance since, at each step, the power of some station increases while the power of the remaining nodes does not decrease. We thus now focus on the proof of the achieved approximation ratio.

We first give two lemmas which are very useful in order to show the existence of good contractions. Due to lack of space, the corresponding proofs have been omitted from this extended abstract. They will appear in the final version of the paper.

Lemma 1. *Given two rooted spanning trees T_1 and T_2 over the same set of nodes V, there exists a one-to-one mapping $f : T_1 \to T_2$, called the swap mapping, such that, if $v_1, \ldots, v_k$ are all the children of a same parent node u in T_1, then the set $\{f(\{u, v_1\}), \ldots, f(\{u, v_k\})\}$ of the edges assigned to $\{u, v_1\}, \ldots, \{u, v_k\}$ by f is a swap set for T_2 and $\{\{u, v_1\}, \ldots, \{u, v_k\}\}$.*

Lemma 2. *Given any tree T, and k edges $\{u, v_1\}, \ldots, \{u, v_k\}$ not necessarily belonging to T, if $\{\{w_1, y_1\}, \ldots, \{w_k, y_k\}\}$ is a swap set for T and $\{\{u, v_1\}, \ldots, \{u, v_k\}\}$, then $\{\{w_1, y_1\}, \ldots, \{w_k, y_k\}\}$ is the subset of a swap set for T and $\{\{u, v_1\}, \ldots, \{u, v_k\}, \{u, z_1\}, \ldots, \{u, z_l\}\}$, for every set of l newly added edges $\{\{u, z_1\}, \ldots, \{u, z_l\}\}$.*

We are now ready to prove the fundamental property that our algorithm exploits.

Lemma 3. *Let T be any spanning tree for $G(S)$ with an arbitrary weighting of the edges, and let $\gamma = c(T)/m^*(S,s)$ be the ratio among the cost of T and the one of an optimal transmission power assignment p^*. Then there exists a contraction of T of cost-efficiency γ.*

Proof. Consider a spanning tree T^* of G^{p^*}, and let f be the swap mapping for T^* and T derived from Lemma 1 considering T^* rooted at s.

Then, by Lemma 1, f assigns to all the descending edges $D(T^*,x)$ in T^* of every station $x \in S$ a subset of edges $SS(x) \subseteq T$ forming a swap set for T and $D(T^*,x)$. All such subsets $SS(x)$ form a partition of T, and since $\frac{c(T)}{m^*(S,s)} = \frac{\sum_{x\in S} c(SS(x))}{\sum_{x\in S} p^*(x)} = \gamma$, there must exist at least one station x such that $\frac{c(SS(x))}{p^*(x)} \geq \gamma$.

Since $D(T^*,x) \subseteq E(p^*,x)$, by Lemma 2, $SS(x) \subseteq A(p^*,x)$, where $A(p^*,x)$ is a swap set for T and $E(p^*,x)$. Therefore, there exists a contraction of T of cost-efficiency $c(A(p^*,x))/p^*(x) \geq c(SS(x))/p^*(x) = \gamma$. □

By exploiting Lemma 3, we can prove the following upper bound on the approximation ratio of our algorithm.

Theorem 1. *Given an instance of MEBR consisting of a cost graph $G(S)$ and a source station $s \in S$, the algorithm has approximation ratio $2\ln\rho - 2\ln 2 + 2$, where $\rho > 2$ is the ratio of the cost of the minimum spanning tree over $G(S)$ over the cost of an optimal solution for the MEBR instance.*

Proof. Let $T_0 = T(S)$ be the minimum spanning tree for $G(S)$ computed at the beginning of the algorithm, $T_1, \ldots, T_k$ be the sequence of the trees constructed by the algorithm after the contraction steps, that for the sake of clarity we assume numbered from 0 to $k-1$, and $\gamma_i = c(T_i)/m^*(S,s)$ be the ratio among the cost of T_i and the one of an optimal transmission power assignment p^*.

By applying Lemma 3, since the algorithm always considers contractions of maximum cost-efficiency, at each step $i = 0, 1, \ldots, k-1$ it performs a contraction having cost-efficiency at least γ_i.

If $\gamma_0 \leq 2$, the algorithm performs no contraction step and it returns the initial minimum spanning tree T_0; thus, the achieved approximation ratio is $\gamma_0 \leq 2\ln\rho - 2\ln 2 + 2$ as $\rho > 2$. In the remaining part of the proof we will assume $\gamma_0 > 2$.

Let x_i be the node involved in the contraction of step i and p_i be the resulting transmission power assigned to x_i. Let $t_i = c(T_i) - c(T_{i+1})$ be the cost of the edges belonging to the original spanning tree T_0 removed by the algorithm at step i, i.e. included in the maximum cost swap set $A(p_i, x_i)$ (only such edges have non-zero weights and thus contribute to the cost of $A(p_i, x_i)$).

A power assignment inducing all the edges of T_k oriented from s towards all the other stations can be obtained by assigning to all the nodes a power assignment equal to the maximum weight of its outgoing edges in T_k with non-zero weight. Moreover, given the power assignment p determined by the algorithm right at the end of the contraction steps, orienting all the corresponding edges of T_k with zero weight in the right direction requires at most doubling the cost of p. Indeed, consider a node x which is connected through edges of zero weight to $\ell > 0$ children $x_1, ..., x_\ell$ in T_k (according to the orientation of T_k from the source node s). Then, in order to make the power assignment induce these edges with direction from x to the children $x_1, ..., x_\ell$, it suffices to increase the power of node x by at most $\max_{i=1,...,\ell} p(x_i) \leq \sum_{i=1}^{\ell} p(x_i)$. Hence, the final power assignment has overall cost upper-bounded by $2\sum_{i=0}^{k-1} p_i + c(T_k)$.

Since by the definition of the algorithm the last contraction has cost-efficiency $\gamma_{k-1} > 2$, and afterwards no contraction of cost-efficiency greater than 2 exists, by Lemma 3 the cost $c(T_k)$ of the final tree is at most $2m^*(S,s)$. Denoting by m the cost of the final power assignment returned by the algorithm we obtain

$$\begin{aligned} m &\leq 2\sum_{i=0}^{k-1} \frac{t_i}{\gamma_i} + c(T_k) \\ &\leq 2\sum_{i=0}^{k-2} \frac{t_i}{\gamma_i} + 2\frac{t_{k-1} - \delta}{\gamma_{k-1}} + \delta + c(T_k) \\ &= 2\sum_{i=0}^{k-2} \frac{t_i}{\gamma_i} + 2\frac{t_{k-1} - \delta}{\gamma_{k-1}} + 2m^*(S,s), \end{aligned}$$

where $\delta = 2m^*(S,s) - c(T_k)$.

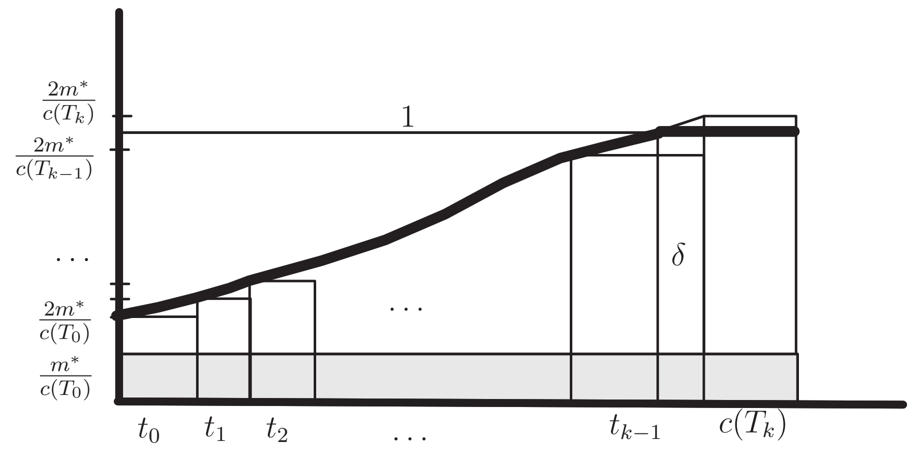

Fig. 1. The cost of the optimal solution (gray area); the cost of the MST heuristic solution (area of the rectangle of height 1); the cost of the solution returned by the algorithm (area below the bold line).

By recalling that $\gamma_i = c(T_i)/m^*(S,s)$, we finally have (see Figure 1)

$$
\begin{aligned}
m &\leq 2\sum_{i=0}^{k-2}\frac{t_i}{\gamma_i} + 2\frac{t_{k-1}-\delta}{\gamma_{k-1}} + 2m^*(S,s) \\
&= 2m^*(S,s)\left(\sum_{i=0}^{k-2}\frac{t_i}{c(T_i)} + \frac{t_{k-1}-\delta}{c(T_{k-1})} + 1\right) \\
&\leq 2m^*(S,s)\left(\sum_{i=0}^{k-2}\int_0^{t_i}\frac{\mathrm{d}t}{c(T_i)-t} + \int_0^{t_{k-1}-\delta}\frac{\mathrm{d}t}{c(T_{k-1})-t} + 1\right) \\
&= 2m^*(S,s)\left(\sum_{i=0}^{k-2}\int_{c(T_0)-c(T_i)}^{c(T_0)-c(T_{i+1})}\frac{\mathrm{d}t}{c(T_0)-t} + \int_{c(T_0)-c(T_{k-1})}^{c(T_0)\left(1-\frac{2}{\gamma_0}\right)}\frac{\mathrm{d}t}{c(T_0)-t} + 1\right) \\
&= 2m^*(S,s)\left(\int_0^{c(T_0)\left(1-\frac{2}{\gamma_0}\right)}\frac{\mathrm{d}t}{c(T_0)-t} + 1\right) \\
&= m^*(S,s)\left(2\ln\gamma_0 - 2\ln 2 + 2\right).
\end{aligned}
$$

If the initial minimum spanning tree $T_0 = T(S)$ guarantees a ρ-approximation of an optimal solution, the theorem follows by observing that $\gamma_0 \leq \rho$. □

4 A matching lower bound

In this section we present a matching lower bound on the approximation ratio of our algorithm, i.e., we show that our analysis is tight.

Theorem 2. *For any $\epsilon > 0$, there exists an instance of MEBR consisting of a cost graph and a source station for which the solution returned by the algorithm has cost at least $2\ln\rho - 2\ln 2 + 2 - \epsilon$ times the optimal cost, where $\rho > 2$ is the ratio between the cost of the minimum spanning tree over the cost graph and the cost of an optimal solution for MEBR.*

Proof. In order to describe the considered instance, it is useful to first describe the *building block* Q_x with $x > 2$, depicted in Figure 2.

The set of nodes and edges of Q_x are $V_x = \{v_x\} \cup \{u_{x,i} | i = 1, 2, \ldots, \lceil x \rceil\}$, and $E_x = \{\{v_x, u_{x,i}\} | i = 1, 2, \ldots, \lceil x \rceil\}$, respectively. The weight of the edge $\{v_x, u_{x,\lceil x \rceil}\}$ is equal to $1 + x - \lceil x \rceil$, and the weights of all the other edges are equal to 1. Notice that in Q_x there exists a contraction centered at node v_x having cost-efficiency equal to x.

We are now ready to describe the whole instance, whose minimum spanning tree is depicted in Figure 3. Let k be an integer parameter; the node set of the instance is obtained by sequencing $k(\rho - 2)$ building blocks $(Q_{2+\frac{1}{k}}, Q_{2+\frac{2}{k}}, \ldots, Q_3, \ldots, Q_\rho)$ in such a way that for two consecutive blocks Q_x and Q_y, $x > y$, nodes $u_{x,2}$ and $u_{y,1}$ coincide. Moreover, in the instance there are other 3 nodes: the source s, node v_1, and node v_2 which coincides with $u_{2+\frac{1}{k},2}$.

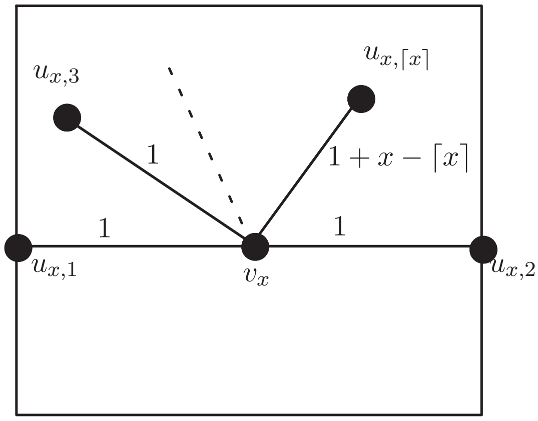

Fig. 2. The building block Q_x.

It remains to define the weights of the edges between the nodes. The weights of the edges connecting s to all the other nodes are equal to 1; moreover, $w(v_1, v_2) = 1$. The weights of the edges contained in the building blocks are properly scaled so that the sum of all the edges of each building block is equal to $\frac{1}{k}$. In particular, the weights of all the edges belonging to building block Q_x are divided by kx. For all the other pairs of nodes, we assume that the mutual power communication cost is very high.

Assume that the initial minimum spanning tree considered by the algorithm is the one depicted in Figure 3, whose cost is ρ.

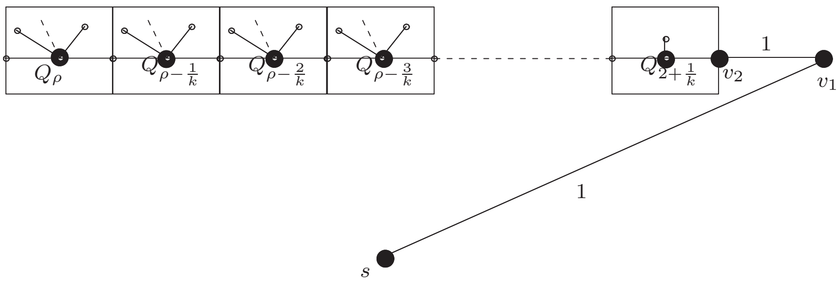

Fig. 3. A minimum spanning tree of the lower bound instance.

At the initial step (step 0), the algorithm can arbitrarily choose among two equivalent contractions, i.e., having the same cost-efficiency; the first choice is the contraction centered at the source and having transmission power equal to 1, and the second choice is the contraction centered at v_ρ and having transmission power equal to $\frac{1}{\rho}$. Both contractions have a cost-efficiency equal to ρ, and we assume that the algorithm chooses the contraction centered at v_ρ.

Using the same arguments, we can assume that the algorithm proceeds by performing other $k(\rho-2)-1$ steps of contractions (steps $1, 2, \ldots, k(\rho-2)-1$), choosing at step i the contraction centered at $v_{\rho-\frac{i}{k}}$, having transmission power equal to $\frac{1}{k\rho-i}$ and cost-efficiency equal to $\rho-\frac{i}{k}$, instead of the equivalent contraction centered at s.

At this point, no contraction having cost-efficiency at least 2 exists any longer. Notice that the sum of the costs of the transmission powers set in the contractions is $\sum_{i=2k+1}^{k\rho} \frac{1}{i} = H_{k\rho} - H_{2k}$, where $H_i = 1 + \frac{1}{2} + \ldots + \frac{1}{i}$ is the harmonic number. In order to orient the edges of the final tree from the source towards the node, we have to globally double the cost of the transmission powers set in the contraction steps. In particular, we have to assign to v_2 a transmission power equal to the one of $v_{2+\frac{1}{k}}$ and to $u_{x,2}$ a transmission power equal to the one of v_x, for $x = 2+\frac{2}{k}, 2+\frac{3}{k}, \ldots, \rho$. Thus, the final cost of the solution returned by the algorithm has cost $2H_{k\rho} - 2H_{2k} + 2$, while the optimal solution has cost 1 and is obtained by assigning to the source node a transmission cost equal to 1. As it can be easily checked, letting k go to infinity, the approximation ratio tends to $2\ln\rho - 2\ln 2 + 2$ from below. □

5 Conclusions

We have presented an approximation algorithm that exponentially outperforms the MST heuristic on any instance of the minimum energy broadcasting problem in ad hoc wireless networks. Our results are particularly relevant for their consequences on Euclidean instances where the achieved approximation ratio is linear in the number of dimension d instead of exponential. Therefore, the improvement becomes more and more significant as d increases. Some corresponding values are depicted in Figure 4.

Number of dimensions	1	2	3	4	5	6	7
MST heuristic approx. ratio	2 [9, 3]	6 [1]	18.8 [21]	80 [13]	242 [13]	728 [13]	2186 [13]
Our algorithm approx. ratio	2	4.2	6.49	9.38	11.6	13.8	16

Fig. 4. Comparison between the approximation factors of our algorithm and the MST heuristic in Euclidean instances for an increasing number of dimensions.

Several questions are left open. First of all, our analysis works on general instances, but further improvements might be possible for specific cases like the Euclidean ones. For such instances, it would be worth to determine exact results tightening the current gap between the lower and upper bounds on the approximation ratio. Another interesting issue is that of determining similar contraction strategies possibly leading to better approximate solutions. An important open question is also that of determining better approximation results of the MST heuristic on high-dimensional Euclidean instances. In particular, tightening the approximation ratio to the kissing number for any number of dimensions would also decrease the approximation of our algorithms, although the improvement would be restricted only to a constant multiplicative factor.

References

1. C. Ambühl. An optimal bound for the MST algorithm to compute energy efficient broadcast trees in wireless networks. In *Proceeding of the 32nd International Colloquium on Automata, Languages and Programming (ICALP '05)*, LNCS 3580, Springer, pp. 1139–1150, 2005.
2. S. Athanassopoulos, I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Experimental comparison of algorithms for energy-efficient multicasting in ad hoc networks. In *Proc. of the 3rd International Conference on Ad-Hoc Networks & Wireless (ADHOC NOW '04)*, LNCS 3158, Springer, pp. 183-196, 2004.
3. M. Čagalj, J. Hubaux, and C. Enz. Minimum-energy broadcast in all-wireless networks: NP-completeness and distribution issues. In *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking (MobiCom '02)*, ACM Press, pp. 172–182, 2002.
4. G. Calinescu, S. Kapoor, A. Olshevsky, and A. Zelikovsky. Network lifetime and power assignment in ad-hoc wireless networks. In *Proceedings of the 11th Annual European Symposium on Algorithms (ESA '03)*, LNCS 2832, Springer, pp. 114–126, 2003.
5. G. Calinescu, I. Mandoiu, and A. Zelikovsky. Symmetric connectivity with minimum power consumption in radio networks. In *Proceedings of the 2nd IFIP International Conference on Theoretical Computer Science*, Kluwer, B.V., pp. 119–130, 2002.
6. I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. New results for energy-efficient broadcasting in wireless networks. In *Proceedings of the 13th Annual International Symposium on Algorithms and Computation (ISAAC '02)*, LNCS 2518, Springer, pp. 332–343, 2002.
7. I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Energy-efficient wireless network design. *Theory of Computing Systems*, 39(5), pp. 593–617, 2006.
8. J. Cartigny, D. Simplot, and I. Stojmenovic. Localized minimum-energy broadcasting in ad-hoc networks. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '03)*, Vol. 3, IEEE Computer Society Press, pp. 2210–2217, 2003.
9. A. Clementi, P. Crescenzi, P. Penna, G. Rossi, and P. Vocca. On the complexity of computing minimum energy consumption broadcast subgraphs. In *Proceedings of the 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS '01)*, LNCS 2010, Springer, pp. 121–131, 2001.
10. J. H. Conway, and N. J. A. Sloane. "The kissing number proble" and "Bounds on kissing numbers". Ch. 2.1 and Ch. 13 in: *Sphere Packings, Lattices, and Groups.* Springer-Verlag, New York, 3rd edition, 1998.
11. Q. Dong, S. Banerjee, M. Adler, and A. Misra. Minimum energy reliable paths using unreliable wireless links. In *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc '05)*, ACM Press, pp. 449–459, 2005.

12. S. Doshi, S. Bhandare, and T. X. Brown. An on-demand minimum energy routing protocol for a wireless ad hoc network. *SIGMOBILE Mobile Computing and Communication Review*, 6(3), pp. 50–66, 2002.
13. M. Flammini, R. Klasing, A. Navarra, and S. Perennes. Improved approximation results for the minimum energy broadcasting problem. In *Proceedings of ACM Joint Workshop on Foundations of Mobile Computing (DIALM-POMC)* ACM Press, pp. 85–91, 2004.
14. M. Flammini, A. Navarra, and S. Perennes. The 'real' approximation factor of the MST heuristic for the minimum energy broadcasting. *ACM Journal of Experimental Algorithmics*, 11, Article no. 2.10, pp. 1-13, 2006.
15. E. Gilbert, and H. Pollak. Steiner minimal trees. *SIAM Journal on Applied Mathematics*, 16, pp. 1–29, 1968.
16. A. Hac. *Wireless sensor network designs.* John Wiley & Sons, Ltd, 2004.
17. L. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc. Power consumption in packet radio networks. *Theoretical Computer Science*, 243(1-2), pp. 289–305, 2000.
18. R. Klasing, A. Navarra, A. Papadopoulos, and S. Perennes. Adaptive broadcast consumption (ABC), a new heuristic and new bounds for the minimum energy broadcast routing problem. In *Proceedings of the 3rd IFIP-TC6 International Networking Conference*, LNCS 3042, Springer, pp. 866–877, 2004.
19. X. Li, and P. Wan. Constructing minimum energy mobile wireless networks. *SIGMOBILE Mobile Computing and Communication Review*, 5(4), pp. 55–67, 2001.
20. W. Liang. Constructing minimum-energy broadcast trees in wireless ad hoc networks. In *Proceedings of the 3rd ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc '02)*, ACM Press, pp. 112–122, 2002.
21. A. Navarra. 3-d minimum energy broadcasting. In *Proceedings of the 13th Colloquium on Structural Information and Communication Complexity (SIROCCO '06)*, LNCS 4056, Springer, pp. 240–252, 2006.
22. T. Rappaport. *Wireless communications: principles and practice.* Prentice-Hall, Englewood Cliffs, NY, 1996.
23. P.-J. Wan, G. Calinescu, X. Li, and O. Frieder. Minimum energy broadcast routing in static ad hoc wireless networks. *Wireless Networks*, 8(6), pp. 607–617, 2002.
24. 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 '00)*, IEEE Computer Society Press, pp. 585–594, 2000.
25. F. Zhao, and L. Guibas. *Wireless sensor networks: an information processing approach.* Morgan Kaufmann, 2004.