

New Results for Energy-Efficient Broadcasting in Wireless Networks^{*}

Ioannis Caragiannis¹, Christos Kaklamanis¹, and Panagiotis Kanellopoulos¹

Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece
{caragian, kakl, kanellop}@ceid.upatras.gr

Abstract. Motivated by the problem of supporting energy-efficient broadcasting in ad hoc wireless networks, we study the *Minimum Energy Consumption Broadcast Subgraph* (MECBS) problem. We present the first logarithmic approximation algorithm for the problem which uses an interesting reduction to Node-Weighted Connected Dominating Set. We also show that an important special instance of the problem can be solved in polynomial time, solving an open problem of Clementi et al. [2].

1 Introduction

Wireless networks have received significant attention during the recent years. Especially, *ad hoc wireless networks* emerged due to their potential applications in battlefield, emergency disaster relief, etc. [8]. Unlike traditional wired networks or cellular wireless networks, no wired backbone infrastructure is installed for ad hoc wireless networks.

A node (or station) in these networks is equipped with an omnidirectional antenna which is responsible for sending and receiving signals. Communication in these networks is established by assigning to each station a transmitting power. In the most common power attenuation model [8], the signal power falls as $1/r^\alpha$, where r is the distance from the transmitter and α is a constant which depends on the wireless environment (typical values of α are between 1 and 6). So, a transmitter can send a signal to a receiver if

$$\frac{P_s}{d(s,t)^\alpha} \geq \gamma$$

where P_s is the power of the transmitting signal, $d(s,t)$ is the Euclidean distance between the transmitter and the receiver, and γ is the receiver's power threshold for signal detection which is usually normalized to 1.

So, communication from a node s to another node t may be established either directly if the two nodes are close enough and s uses adequate transmitting

^{*} This work was partially supported by the European Union under IST FET Project ALCOM-FT, IST FET Project CRESCCO, and RTN Project ARACNE.

power, or by using intermediate nodes. Observe that due to the nonlinear power attenuation, relaying the signal between nodes may result in energy conservation.

A crucial issue in ad hoc wireless networks is to support communication patterns that are typical in traditional networks. These include broadcasting, multicasting, and gossiping (all-to-all communication). The important engineering question to be solved is to guarantee a desired communication pattern minimizing the total energy consumption. In this work, motivated by the problem of supporting energy-efficient broadcasting in ad hoc wireless networks, we study the *Minimum Energy Consumption Broadcast Subgraph* (MECBS) problem.

Consider a complete directed graph $G = (V, E)$ with a symmetric cost function $c : E \rightarrow R$ associated with its edges ($c(u, v) = c(v, u)$) and a special node $r \in V$. Given a weight assignment $w : V \rightarrow R$ to the nodes of G , the *transmission graph* G_w is the directed graph defined as follows. It has the same set of nodes as G and a directed edge (u, v) belongs to G_w if the weight assigned to node u is at least the cost of the edge (u, v) , i.e., $w(u) \geq c(u, v)$. The Minimum Energy Consumption Broadcast Subgraph (MECBS) problem is to assign weights to the nodes of V so that for any node $u \in V - \{r\}$, the transmission graph G_w has a directed path from r to u . The objective is to minimize the sum of weights.

Note that by changing the connectivity requirements for the transmission graph, several interesting combinatorial problems arise. Problems of this kind are examined in [1,4,6]. In these works, the objective is to establish a strongly connected transmission graph (with or without restrictions to its diameter).

MECBS has been proved to be inapproximable within $(1 - \epsilon) \ln n$ unless $\mathcal{NP} \subseteq DTIME(n^{O(\log \log n)})$ (where n denotes the number of nodes in the graph) using approximation preserving reductions to SET COVER [2] and CONNECTED DOMINATING SET [7]. However, to our knowledge, no logarithmic approximation algorithms for the problem were known prior to this work.

An important special case of MECBS (which reflects the properties of the engineering problem in practice) is when the nodes of the graph are points in the d -dimensional Euclidean space and the cost function is defined as a α -th power of the distance function. Following the notation of [2], we denote this special case of the problem with $\text{MECBS}[N_d^\alpha]$. Clementi et al. [2] have proved that $\text{MECBS}[N_d^\alpha]$ is \mathcal{NP} -hard for $d \geq 2$ and $\alpha > 1$. Note that, for $\alpha \leq 1$ and $d \geq 1$, $\text{MECBS}[N_d^\alpha]$ has a trivial optimal solution (the transmission graph is a star rooted at r). The complexity of $\text{MECBS}[N_1^\alpha]$ for $\alpha > 1$ was left open.

Wieselthier et al. [8] study the behaviour of several greedy algorithms for $\text{MECBS}[N_2^\alpha]$. They present experimental results for three simple algorithms: MST (Minimum Spanning Tree), BIP (Broadcast Incremental Power), and SPT (Shortest Path Tree). By exploring geometric properties of Euclidean spanning trees, Clementi et al. [2] and Wan et al. [7] prove upper bounds on the approximation ratio of algorithm MST on instances of $\text{MECBS}[N_2^\alpha]$ for $\alpha \geq 2$. In [2], an upper bound of $(2\sqrt{10})^\alpha$ is proved, while an improved bound of 12 is proved in [7]. Wan et al. [7] also prove lower bounds on the performance of algorithms MST, BIP, SPT, and BAIP (Broadcast Average Incremental Power), a greedy algorithm based on the well-known greedy algorithm for SET COVER. They show that even when instances of $\text{MECBS}[N_2^\alpha]$ are considered, SPT and BAIP have approx-

ximation ratios $\Omega(n)$ and $\Omega(n/\log n)$, respectively. However, it is easy to verify that all the algorithms studied in [2,7,8] have approximation ratios $\Omega(n)$ when applied to general instances of MECBS. Furthermore, none of these algorithms can guarantee optimal solutions to MECBS $[N_1^\alpha]$ for $\alpha > 1$.

In this work, we present (in Section 2) the first logarithmic approximation algorithm for the problem which uses an interesting reduction to the *Node-Weighted Connected Dominating Set* (NWCDS) problem. The Node-Weighted Connected Dominating Set problem is the following: Given a graph $G = (V, E)$ with weights on the nodes, find the smallest weighted subset S of nodes that induce a connected subgraph and each node in $V - S$ is adjacent to at least one node in S . We also show (in Section 3) that MECBS $[N_1^\alpha]$ for $\alpha > 1$ can be solved in polynomial time, solving an open problem of Clementi et al. [2]. Very recently, Andrea Clementi informed us that a similar result was obtained independently in [3].

2 A Logarithmic Approximation Algorithm

In this section, we present a logarithmic approximation algorithm for MECBS. The algorithm uses a new reduction of instances of MECBS to instances of NWCDS.

First, we show how to transform any instance of MECBS to an instance of NWCDS. Let I_{MECBS} be an instance of MECBS which consists of a complete directed graph $G = (V, E)$ with $|V| = n$, a symmetric cost function $c : E \rightarrow R$, on the edges of G , and a special node $r \in V$. We will construct an instance I_{NWCDS} of the NWCDS which consists of a node-weighted undirected graph H .

We construct the graph H as follows. For each node v of G , the graph H has a set Z_v of n nodes $Z_{v,1}, Z_{v,2}, \dots, Z_{v,n}$ which we call a *supernode*. The supernodes of H corresponding to different nodes of G are disjoint. Let $c_{v,i}$ be the i -th largest cost among the costs of the edges directed out of v in G . The weight associated with the node $Z_{v,i}$, $i = 1, \dots, n - 1$ of the supernode Z_v is set to $c_{v,i}$. The node $Z_{v,n}$ of the supernode Z_v has infinite weight.

The set of edges in H is defined as follows. Nodes $Z_{v,1}, \dots, Z_{v,n-1}$ of the supernode Z_v form the complete graph K_{n-1} while node $Z_{v,n}$ is isolated from the other nodes of its supernode. Node $Z_{v,i}$, for $i = 1, \dots, n - 1$, is connected to all the nodes of those supernodes Z_u ($v \neq u$) corresponding to nodes u of G which are connected to v by edges of cost no more than the weight of $Z_{v,i}$ (i.e., $c(v, u) \leq c_{v,i}$). The reduction is depicted in Figure 1.

Now, we will show that an approximate solution to NWCDS for instance I_{NWCDS} can be used to obtain an approximate solution to the MECBS for instance I_{MECBS} with similar approximation guarantee. We first show that the cost $OPT(I_{\text{MECBS}})$ of an optimal solution of I_{MECBS} is close to the cost $OPT(I_{\text{NWCDS}})$ of the optimal solution of I_{NWCDS} .

Lemma 1. $OPT(I_{\text{NWCDS}}) \leq 2 \cdot OPT(I_{\text{MECBS}})$.

Proof. Consider a solution to MECBS for instance I_{MECBS} of minimum cost. Let w be the weight vector corresponding to this solution and G_w the transmission




Fig. 1. The reduction to NWCDS. (a) The graph G of an instance of MECBS. (b) The graph H of the corresponding instance of NWCDS. Each large cycle indicates a supernode. Only the edges incident to the node of weight 2 of the upper left supernode are shown. These edges are those which correspond to edges in (a) of cost at most 2, directed out of the left upper node. (c) The graph H of the corresponding instance of NWCDS. Thick edges from a node to a supernode represent edges connecting a node to all the nodes of a supernode.

graph defined by w . Since the solution is optimal, the weight of each node v equals the cost of some edge directed out of v .

We construct a connected dominating set S for instance I_{NWCDS} as follows. For each node v of G_w which has at least one outgoing edge, we add to S the node $Z_{v,i(v)}$ of the supernode Z_v which has weight $w(Z_{v,i(v)}) = w(v) = \max_{(v,u) \in G_w} \{c(v,u)\}$.

Consider the induced subgraph G'_w of G_w which contains the nodes of G_w having at least one outgoing edge in G_w . By construction, each one of the nodes in S belongs to a different supernode of H mapping to a node in G'_w . Also, a directed edge between two nodes u and v in G'_w exists if and only if $w(u) \geq c(u,v)$. This means that the weight of the node $Z_{u,i(u)}$ is at least $w(u)$, and, thus, it is connected with node $Z_{v,i(v)}$ in H . This holds for any directed edge of G'_w and implies that the subgraph of H induced by S is isomorphic to G'_w (if we consider directed edges as undirected ones), and, thus, it is connected.

Also, S dominates all nodes of H except (possibly) for $Z_{r,n}$. The existence of a directed edge (u,v) in G_w implies that the node $Z_{u,i(u)}$ dominates all nodes of Z_v in H . Furthermore, all nodes in Z_r but $Z_{r,n}$ are dominated by $Z_{r,i(r)}$.

If $Z_{r,n}$ is adjacent to a node in S , then S is clearly a connected dominating set in H . The cost of S is

$$\sum_{Z_{v,i(v)} \in S} w(Z_{v,i(v)}) = \sum_{v \in G_w} w(v) = \text{OPT}(I_{\text{MECBS}}).$$

Now, if $Z_{r,n}$ is not adjacent to any node of S , we consider a node u such that the directed edge (r, u) exists in G_w . Let $Z_{u,j}$ be the node of Z_u with weight $c(u, r)$. Note that $Z_{u,j}$ is adjacent to both $Z_{u,i(u)}$ and $Z_{r,n}$. Thus, $S \cup \{Z_{u,j}\}$ is a connected dominating set in H . The cost of an optimal solution is at most the cost of $S \cup \{Z_{u,j}\}$ which is $OPT(I_{MECBS}) + w(Z_{u,j})$. Recall that $w(Z_{u,j}) = c(u, r) = c(r, u)$ and that $w(r) \geq c(r, u)$ since the edge (r, u) belongs to the transmission graph G_w . We obtain that

$$\begin{aligned} OPT(I_{NWCDS}) &\leq OPT(I_{MECBS}) + w(Z_{u,j}) = OPT(I_{MECBS}) + c(u, r) \\ &= OPT(I_{MECBS}) + c(r, u) \leq OPT(I_{MECBS}) + w(r) \\ &\leq 2 \cdot OPT(I_{MECBS}) \quad \square \end{aligned}$$

We now show that a solution to NWCDS for instance I_{NWCDS} can give a solution to MECBS for instance I_{MECBS} of similar cost.

Lemma 2. *Given a solution to NWCDS for instance I_{NWCDS} of cost $COST(I_{NWCDS})$, we can construct in polynomial time a solution to MECBS for instance I_{MECBS} of cost $COST(I_{MECBS})$, such that*

$$COST(I_{MECBS}) \leq 2 \cdot COST(I_{NWCDS})$$

Proof. Consider a connected dominating set S of H . We may assume that S contains at most one node from each supernode. If this is not the case and there are two nodes $Z_{u,i}$ and $Z_{u,j}$ ($i \neq j$) in S with weights $w(Z_{u,i}) \geq w(Z_{u,j})$, we remove $Z_{u,j}$ from S . By repeatedly executing this procedure, we end up with a connected dominating set having at most one node per supernode. Furthermore, in any solution of finite cost, no node with infinite weight is contained in S . For every supernode Z_u which has a node $Z_{u,i(u)}$ contained in S , we associate a weight $w(Z_{u,i(u)})$ to node u of G . Nodes of G whose corresponding supernodes of H have no node contained in S are assigned zero weight.

Now, consider the transmission graph G_w defined by w . We will first show that the graph G_w is loosely connected, i.e., its undirected counterpart is connected. Consider an edge $(Z_{u,i(u)}, Z_{v,i(v)})$ in H which connects two nodes of S . By the construction of graph H from G and the definition of the transmission graph G_w , we obtain that either (u, v) or (v, u) exists in G_w . This holds for any two nodes of S which are connected with an edge in H and, since the nodes of S are connected, this implies that the subgraph of G_w induced by the nodes corresponding to supernodes of H having a node in S is loosely connected. Also, observe that the nodes of any supernode Z_v containing no node in S are dominated by some node $Z_{u,i(u)}$ in S . By the construction of H from G and the definition of the transmission graph G_w , we obtain that for any node v of G_w with zero weight, there is a directed edge coming from a node with non-zero weight. We conclude that G_w is loosely connected.

We execute a spanning tree algorithm on G_w to compute a loosely connected spanning tree T_w of G_w . Next, we execute the following procedure that transforms T_w to an arborescence T'_w , i.e., a tree T'_w having a directed path from r

to any other node. Starting from r , we compute a Breadth-First-Search (BFS) numbering of the nodes of T_w . We visit the nodes of T_w according to this BFS numbering and, at the step associated with a node u , we transform ingoing edges coming to u from nodes with numbers greater than u 's to outgoing ones.

Now, by assigning weights $w'(u) = \max_{(u,v) \in T'_w} c(u, v)$ to the nodes of G , we obtain a transmission graph which contains T'_w as a subgraph and, thus, is a feasible solution to MECBS.

Let l_u be the node such that the edge (u, l_u) belongs to T'_w and has the maximum cost among all edges (u, v) directed out of u in T'_w . The transmission graph G_w contains either (u, l_u) or (l_u, u) . If it contains (u, l_u) , then $w(u) \geq c(u, l_u)$, otherwise $w(l_u) \geq c(l_u, u) = c(u, l_u)$. In any case,

$$w'(u) = \max_{(u,v) \in T'_w} c(u, l_u) \leq \max\{w(u), w(l_u)\} \leq w(u) + w(l_u).$$

The cost of the solution w' to MECBS is

$$\begin{aligned} COST(I_{MECBS}) &= \sum_{u \in V} w'(u) \leq \sum_{u \in V} (w(u) + w(l_u)) = \sum_{u \in V} w(u) + \sum_{u \in V} w(l_u) \\ &\leq 2 \cdot \sum_{u \in V} w(u) = 2 \cdot COST(I_{NWCDS}) \quad \square \end{aligned}$$

Given an instance I_{MECBS} of MECBS, we first transform it to the corresponding instance I_{NWCDS} as described above. Then, we run an algorithm for the Node-Weighted Connected Dominating Set problem and use the technique described in the proof of Lemma 2 to construct a solution to the original problem. Using Lemmas 1 and 2, we can show that, given a ρ -approximate solution to NWCDS for instance I_{NWCDS} of cost $COST(I_{NWCDS})$, we can obtain a solution to MECBS for instance I_{MECBS} with cost $COST(I_{MECBS})$ which is within 4ρ of optimal. Indeed, by Lemma 1, we have $OPT(I_{MECBS}) \geq OPT(I_{NWCDS})/2$, while Lemma 2 yields $COST(I_{MECBS}) \leq 2 \cdot COST(I_{NWCDS}) \leq 2\rho \cdot OPT(I_{NWCDS})$. Thus, the approximation ratio we obtain is

$$\frac{COST(I_{MECBS})}{OPT(I_{MECBS})} \leq 4\rho.$$

In [5], Guha and Khuller present a $1.35 \ln n$ -approximation algorithm for the Node-Weighted Connected Dominating Set problem, where n is the number of nodes in the graph. Given an instance I_{MECBS} of MECBS with n nodes, the corresponding instance I_{NWCDS} has n^2 nodes. Thus, the cost of the solution of I_{NWCDS} is within $4 \cdot 1.35 \ln(n^2) = 10.8 \ln n$ of the optimal solution. The next theorem summarizes the discussion of this section.

Theorem 1. *There exists a $10.8 \ln n$ -approximation algorithm for MECBS.*

3 A Polynomial Time Algorithm for MECBS[\mathbb{N}_1^α]

Assume that we have n points $x_{u_1}, \dots, x_r, \dots, x_{u_n}$ located on a line. An instance I of MECBS[\mathbb{N}_1^α], for some $\alpha > 1$, consists of a complete directed graph G in

which each node u corresponds to a point x_u and the cost of an edge (u, v) in G is defined as

$$c(u, v) = [d(x_u, x_v)]^\alpha,$$

where $d(x_u, x_v)$ is the Euclidean distance of the points x_u and x_v .

We will show that we can compute an optimal solution for the instance I of MECBS[\mathbb{N}_1^α] in polynomial time. We start with a few definitions.

We partition the nodes of G into two sets:

- the *left set* \mathcal{L} which contains those nodes of G corresponding to a point at the left of node x_r .
- the *right set* \mathcal{R} which contains those nodes of G corresponding to a point at the right of node x_r .

If one of the sets \mathcal{L} and \mathcal{R} is empty (this means that all the nodes correspond to points located at the right or the left of x_r), then there is a trivial optimal solution. For example, if $r = u_1$ and $\mathcal{R} = \{u_2, \dots, u_n\}$, it is easy to verify that $w(u_i) = c(u_i, u_{i+1})$, for $i = 1, \dots, n - 1$, and $w(u_n) = 0$ is an optimal solution. So, in the following we assume that both \mathcal{L} and \mathcal{R} are non-empty.

Consider an arborescence T of G rooted at r . A node u of T different than r is called *root-crossing* if it belongs to \mathcal{L} and has a child in \mathcal{R} or if it belongs in \mathcal{R} and has a child belonging in \mathcal{L} . The node r is called root-crossing if it contains children in both \mathcal{L} and \mathcal{R} . Note that, if both \mathcal{L} and \mathcal{R} are non-empty, then any arborescence of G rooted at r contains at least one root-crossing node.

A *chain* from u to v is a directed path from u to v containing only edges between consecutive nodes v' which correspond to point $x_{v'}$ located between the points x_u and x_v on the line. We also say that a node u is a trivial chain from u to u .

An arborescence of G rooted at r is called *single root-crossing* if

- it has exactly one root-crossing node u ,
- it contains a chain from r to u , a chain from the node u_L , which is the node among u and its children corresponding to the leftmost point, to u_1 , and a chain from the node u_R , which is the node among u and its children corresponding to the rightmost point, to u_n , and
- u is the parent of all nodes corresponding to points between x_{u_L} and x_{u_R} except for nodes in the chain from r to u .

Examples of arborescences are shown in Figure 2.

Let T be an arborescence of G rooted at r . By setting $w(u) = \max_{(u,v) \in T} c(u, v)$, the transmission graph G_w contains T as a subgraph and, thus, it is a feasible solution to the problem with cost $\sum_{u \in V} w(u)$. We will show that given an arborescence T which is not single root-crossing, there exists a feasible solution w' of cost $\sum_{u \in V} w'(u) \leq \sum_{u \in V} w(u)$ such that the transmission graph $G_{w'}$ contains a single root-crossing arborescence T'' as a subgraph. We prove it by constructing from T , a single root-crossing arborescence T'' with

$$\sum_{u \in V} \max_{(u,v) \in T''} c(u, v) \leq \sum_{u \in V} \max_{(u,v) \in T} c(u, v)$$




Fig. 2. Points on the line and three arborescences rooted at the node represented by the rectangle. The first arborescence is not single root-crossing because it has two root-crossing nodes. The second and third arborescences have one root-crossing node but, according to the definition, only the third one is single root-crossing.

and setting $w'(u) = \max_{(u,v) \in T''} c(u, v)$. This is proved in Lemmas 3 and 4. In this way, we will obtain that there exists an optimal solution \hat{w} to instance I which defines a transmission graph $G_{\hat{w}}$ containing a single root-crossing arborescence as a subgraph.

Lemma 3. *Let T be an arborescence of G rooted at r containing $k \geq 2$ root-crossing nodes in T . There exists an arborescence T' of G rooted at r with $k - 1$ root-crossing nodes such that*

$$\sum_{u \in T'} \max_{(u,v) \in T'} c(u, v) \leq \sum_{u \in T} \max_{(u,v) \in T} c(u, v).$$

Proof. Consider an arborescence T of G rooted at r with $k \geq 2$ root-crossing nodes. Let u_1 and u_2 be two root-crossing nodes such that u_1 is at the same or higher level than u_2 in T . Again, we denote by x_v the point on the line which corresponds to the node v . We call *left* (resp. *right*) children of a node v the children of v in T which belong to set \mathcal{L} (resp. \mathcal{R}). Also, we denote by v_L and v_R the node between v and its children which correspond to the leftmost and rightmost point on the line, respectively.

We distinguish between the following four cases

Case 1. If $u_1 = r$, then assume without loss of generality that $u_2 \in \mathcal{R}$. We claim that either all left children of u_2 in T are within distance $[\max_{(u_1,v) \in T} c(u_1, v)]^{1/\alpha}$ from u_1 or all right children of u_1 in T are within distance $[\max_{(u_2,v) \in T} c(u_2, v)]^{1/\alpha}$ from u_2 . Assume otherwise, i.e., $d(x_{u_1}, x_{u_1^R}) > [\max_{(u_2,v) \in T} c(u_2, v)]^{1/\alpha}$ and $d(x_{u_2}, x_{u_2^L}) > [\max_{(u_1,v) \in T} c(u_1, v)]^{1/\alpha}$. Then,

$$d(x_{u_1}, x_{u_1^R}) > \left[\max_{(u_2,v) \in T} c(u_2, v) \right]^{1/\alpha} \geq d(x_{u_2}, x_{u_2^L}) > \left[\max_{(u_1,v) \in T} c(u_1, v) \right]^{1/\alpha},$$

a contradiction since $(u_1, u_1^R) \in T$. So, we can construct an arborescence T' either by removing all edges (u_2, v) from u_2 to its left children (in this way, u_2

is not root-crossing) and adding edges (u_1, v) from u_1 to the left children of u_2 or by removing all edges (u_1, v) from u_1 to its right children (in this way, u_1 is not root-crossing) and adding edges (u_2, v) from u_2 to the right children of u_1 . The arborescence T' has at most $k - 1$ root-crossing nodes and, clearly,

$$\sum_{u \in T'} \max_{(u,v) \in T'} c(u, v) \leq \sum_{u \in T} \max_{(u,v) \in T} c(u, v).$$

Case 2. If $u_1 \neq r$ and u_2 belongs to the same set with u_1 (wlog, we assume that $u_1, u_2 \in \mathcal{L}$) then either all right children of u_2 in T are within distance $[\max_{(u_1,v) \in T} c(u_1, v)]^{1/\alpha}$ from u_1 or all right children of u_1 in T are within distance $[\max_{(u_2,v) \in T} c(u_2, v)]^{1/\alpha}$ from u_2 . Assume otherwise, i.e., $d(x_{u_1}, x_{u_1^R}) > [\max_{(u_2,v) \in T} c(u_2, v)]^{1/\alpha}$ and $d(x_{u_2}, x_{u_2^R}) > [\max_{(u_1,v) \in T} c(u_1, v)]^{1/\alpha}$. Then,

$$d(x_{u_1}, x_{u_1^R}) > \left[\max_{(u_2,v) \in T} c(u_2, v) \right]^{1/\alpha} \geq d(x_{u_2}, x_{u_2^R}) > \left[\max_{(u_1,v) \in T} c(u_1, v) \right]^{1/\alpha},$$

a contradiction since $(u_1, u_1^R) \in T$. So, we can construct an arborescence T' by removing all edges (u_1, v) from u_1 to its right children (in this way, u_1 is not root-crossing) and adding edges (u_2, v) from u_2 to the right children of u_1 or by removing all edges (u_2, v) from u_2 to its right children (in this way, u_2 is not root-crossing) and adding edges (u_1, v) from u_1 to the right children of u_2 . Again, the arborescence T' has at most $k - 1$ root-crossing nodes and, clearly,

$$\sum_{u \in T'} \max_{(u,v) \in T'} c(u, v) \leq \sum_{u \in T} \max_{(u,v) \in T} c(u, v).$$

Case 3. If $u_1 \neq r$, and u_1, u_2 belong to different sets and u_2 is not a child of u_1 , then assume without loss of generality that $u_1 \in \mathcal{L}$ and $u_2 \in \mathcal{R}$. Again, using the same argument with the previous two cases, we can show that either all left children of u_2 in T are within distance $[\max_{(u_1,v) \in T} c(u_1, v)]^{1/\alpha}$ from u_1 or all right children of u_1 in T are within distance $[\max_{(u_2,v) \in T} c(u_2, v)]^{1/\alpha}$ from u_2 . So, we can construct an arborescence T' either by removing all edges (u_2, v) from u_2 to its left children (in this way, u_2 is not root-crossing) and adding edges (u_1, v) from u_1 to the left children of u_2 or by removing all edges (u_1, v) from u_1 to its right children (in this way, u_1 is not root-crossing) and adding edges (u_2, v) from u_2 to the right children of u_1 . Again, we obtain an arborescence T' with at most $k - 1$ root-crossing nodes, such that

$$\sum_{u \in T'} \max_{(u,v) \in T'} c(u, v) \leq \sum_{u \in T} \max_{(u,v) \in T} c(u, v).$$

Case 4. Now, we consider the case where $u_1 \neq r$, and u_1, u_2 belong to different sets and u_2 is a child of u_1 in T . Without loss of generality, we assume that $u_1 \in \mathcal{L}$ and $u_2 \in \mathcal{R}$. If all left children of u_2 are within distance $[\max_{(u_1,v) \in T} c(u_1, v)]^{1/\alpha}$

from u_1 , then we construct T' by removing the edges (u_2, v) from u_2 to every left child v of u_2 and adding edges from u_1 to v . In T' , u_2 is not root-crossing and we obtain an arborescence with $k - 1$ root-crossing nodes. Again, we have

$$\sum_{u \in T'} \max_{(u,v) \in T'} c(u, v) \leq \sum_{u \in T} \max_{(u,v) \in T} c(u, v).$$

Assume now that not all left children of u_2 are within distance at most $[\max_{(u_1,v) \in T} c(u_1, v)]^{1/\alpha}$ from u_1 , i.e., $d(x_{u_1}, x_{u_2^L}) > [\max_{(u_1,v) \in T} c(u_1, v)]^{1/\alpha}$. Then, u_2^L is at the left of u_1 and

$$\left[\max_{(u_2,v) \in T} c(u_2, v) \right]^{1/\alpha} \geq d(x_{u_1}, x_{u_2^L}) + d(x_{u_1}, x_{u_2}) > \left[\max_{(u_1,v) \in T} c(u_1, v) \right]^{1/\alpha}.$$

Thus, since either u_1^R is at right of u_2 or $u_1^R = u_2$,

$$\begin{aligned} d(x_{u_2}, x_{u_1^R}) &= d(x_{u_1}, x_{u_1^R}) - d(x_{u_1}, x_{u_2}) \\ &\leq \max_{(u_1,v) \in T} c(u_1, v)^{1/\alpha} < \max_{(u_2,v) \in T} c(u_2, v)^{1/\alpha}. \end{aligned}$$

Now, we construct the arborescence T' as follows. First, we remove from T the edge between the parent of u_1 and u_1 and all edges (u_1, v) from u_1 to its children (in this way, u_1 is not root-crossing). We add to T' an edge from u_2 to u_1 and edges (u_2, v) from u_2 to the children of u_1 in T .

If r has right children in T , then we complete the construction of T' by adding edge (r, u_2) . Note that, in this way, we add a right child to r , so r is root-crossing in T' only if it was root-crossing in T . We obtain an arborescence T' with at most $k - 1$ root-crossing nodes such that

$$\begin{aligned} \sum_{u \in T'} \max_{(u,v) \in T'} c(u, v) &\leq \max_{(r,v) \in T'} c(r, v) + \max_{(u_1,v) \in T'} c(u_1, v) + \max_{(u_2,v) \in T'} c(u_2, v) \\ &\quad + \sum_{u \in T' - \{r, u_1, u_2\}} \max_{(u,v) \in T'} c(u, v) \\ &\leq \left(\max_{(r,v) \in T} c(r, v) + d(x_r, x_{u_2})^\alpha \right) + 0 + \max_{(u_2,v) \in T} c(u_2, v) \\ &\quad + \sum_{u \in T - \{r, u_1, u_2\}} \max_{(u,v) \in T} c(u, v) \\ &\leq \max_{(r,v) \in T} c(r, v) + d(x_{u_1}, x_{u_2})^\alpha + \max_{(u_2,v) \in T} c(u_2, v) \\ &\quad + \sum_{u \in T - \{r, u_1, u_2\}} \max_{(u,v) \in T} c(u, v) \\ &\leq \max_{(r,v) \in T} c(r, v) + \max_{(u_1,v) \in T} c(u_1, v) + \max_{(u_2,v) \in T} c(u_2, v) \\ &\quad + \sum_{u \in T - \{r, u_1, u_2\}} \max_{(u,v) \in T} c(u, v) \\ &= \sum_{u \in T} \max_{(u,v) \in T} c(u, v) \end{aligned}$$

If r has only left children in T , then we first remove all edges from r to its children. Then, we add edges (u_2, v) to connect to u_2 all children of r in T that are within distance $[\max_{(u_2, v) \in T} c(u_2, v)]^{1/\alpha}$ from u_2 while the rest of the children of r in T (if any) are connected to u_2^L . We complete the construction of T' by adding edge (r, u_2) to connect u_2 to r . Note that, now, r has only one (right) child, and, thus, it is not root-crossing. Furthermore, node u_1 has no children in T' and, thus, it is not root-crossing, while node u_2^L is a root-crossing node in T' only if it was a root-crossing node in T . We conclude that T' has at most $k - 1$ root-crossing nodes. We denote by r_l the child of r in T which is not within distance $[\max_{(u_2, v) \in T} c(u_2, v)]^{1/\alpha}$ from u_2 and corresponds to the leftmost point. If no such point exists, then the terms in the following expression referring to r_l can be removed. We have that

$$\begin{aligned}
 \sum_{u \in T'} \max_{(u, v) \in T'} c(u, v) &= \max_{(r, v) \in T'} c(r, v) + \max_{(u_1, v) \in T'} c(u_1, v) + \max_{(u_2, v) \in T'} c(u_2, v) \\
 &\quad + \max_{(u_2^L, v) \in T'} c(u_2^L, v) + \sum_{u \in T' - \{r, u_1, u_2, u_2^L\}} \max_{(u, v) \in T'} c(u, v) \\
 &\leq d(x_r, x_{u_2})^\alpha + 0 + \max_{(u_2, v) \in T} c(u_2, v) + d(x_{u_2^L}, x_{r_l})^\alpha \\
 &\quad + \max_{(u_2^L, v) \in T} c(u_2^L, v) + \sum_{u \in T - \{r, u_1, u_2, u_2^L\}} \max_{(u, v) \in T} c(u, v) \\
 &\leq d(x_{u_1}, x_{u_2})^\alpha + \max_{(u_2, v) \in T} c(u_2, v) + d(x_r, x_{r_l})^\alpha \\
 &\quad + \max_{(u_2^L, v) \in T} c(u_2^L, v) + \sum_{u \in T - \{r, u_1, u_2, u_2^L\}} \max_{(u, v) \in T} c(u, v) \\
 &\leq \max_{(u_1, v) \in T} c(u_1, v) + \max_{(u_2, v) \in T} c(u_2, v) + \max_{(r, v) \in T} c(r, v) \\
 &\quad + \max_{(u_2^L, v) \in T} c(u_2^L, v) + \sum_{u \in T - \{r, u_1, u_2, u_2^L\}} \max_{(u, v) \in T} c(u, v) \\
 &\leq \sum_{u \in T} \max_{(u, v) \in T} c(u, v)
 \end{aligned}$$

This completes the proof of the lemma. □

By repeatedly executing the procedure used in the proof of Lemma 3, we obtain an arborescence T' with exactly one root-crossing node. We can also prove the following.

Lemma 4. *Let T' be an arborescence of G rooted at r containing exactly one single root-crossing node. There exists a single root-crossing arborescence T'' such that*

$$\sum_{u \in T''} \max_{(u, v) \in T''} c(u, v) \leq \sum_{u \in T'} \max_{(u, v) \in T'} c(u, v).$$

Note that, for any node v , there are $O(n^2)$ single root-crossing arborescences in which v is the root-crossing node. Thus, there are $O(n^3)$ single root-crossing

arborescences of G rooted at r . For each possible single root-crossing arborescence T_i , we set $w_i(u) = \max_{(u,v) \in T_i} c(u,v)$. Clearly, w_i is a feasible solution for instance I since the transmission graph G_{w_i} contains the arborescence T_i as a subgraph. The cost of the solution w_i is $\sum_{u \in T_i} w_i(u)$. We select that solution w_i which minimizes $\sum_{u \in T_i} w_i(u)$. By the discussion in this section, this is an optimal solution for I . Since every step in the procedure described above is performed in polynomial time, we have obtained the following theorem.

Theorem 2. *MECBS $[\mathbb{N}_1^\alpha]$ can be solved in polynomial time.*

References

1. G. Călinescu, I. Măndoiu, A. Zelikovsky. Symmetric Connectivity with Minimum Power Consumption in Radio Networks. In *Proc. of the 2nd IFIP International Conference on Theoretical Computer Science, 2002*, to appear.
2. A. E. F. Clementi, P. Crescenzi, P. Penna, G. Rossi, P. Vocca. On the Complexity of Computing Minimum Energy Consumption Broadcast Subgraphs. In *Proc. of the 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS '01)*, LNCS 2010, Springer, pp. 121–131, 2001.
3. A. E. F. Clementi, M. Di Ianni, R. Silvestri. The Minimum Broadcast Range Assignment Problem on Linear Multi-Hop Wireless Networks. *Theoretical Computer Science*, to appear.
4. A. E. F. Clementi, P. Penna, R. Silvestri. Hardness Results for the Power Range Assignment Problem in Packet Radio Networks. In *Proc. of Randomization, Approximation, and Combinatorial Optimization (RANDOM/APPROX '99)*, LNCS 1671, Springer, pp. 197–208, 1999.
5. S. Guha, S. Khuller. Improved Methods for Approximating Node Weighted Steiner Trees and Connected Dominating Sets. *Information and Computation*, Vol. 150 (1), pp. 57–74, 1999.
6. L. M. Kirousis, E. Kranakis, D. Krizanc, A. Pelc. Power Consumption in Packet Radio Networks. *Theoretical Computer Science*, Vol. 243(1-2), pp. 289–305, 2000.
7. P.-J. Wan, G. Călinescu, X.-Y. Li, O. Frieder. Minimum-Energy Broadcast Routing in Static Ad Hoc Wireless Networks. In *Proc. of IEEE INFOCOM 2001*, pp. 1162–1171.
8. J. E. Wieselthier, G. D. Nguyen, A. Ephremides. On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks. In *Proc. of IEEE INFOCOM 2000*, pp. 585–594.