

Minimum-Energy Broadcast in All-Wireless Networks: NP-Completeness and Distribution Issues *

Mario Čagalj
LCA-EPFL
CH-1015 Lausanne
Switzerland
mario.cagalj@epfl.ch

Jean-Pierre Hubaux
LCA-EPFL
CH-1015 Lausanne
Switzerland
jean-
pierre.hubaux@epfl.ch

Christian Enz
CSEM
CH-2007 Neuchâtel
Switzerland
christian.enz@csem.ch

ABSTRACT

In all-wireless networks a crucial problem is to minimize energy consumption, as in most cases the nodes are battery-operated. We focus on the problem of power-optimal broadcast, for which it is well known that the broadcast nature of the radio transmission can be exploited to optimize energy consumption. Several authors have conjectured that the problem of power-optimal broadcast is NP-complete. We provide here a formal proof, both for the general case and for the geometric one; in the former case, the network topology is represented by a generic graph with arbitrary weights, whereas in the latter a Euclidean distance is considered. We then describe a new heuristic, Embedded Wireless Multicast Advantage. We show that it compares well with other proposals and we explain how it can be distributed.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design; C.2.2 [Computer-Communication Networks]: Network Protocols; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms

Algorithms, Design, Performance

Keywords

Wireless ad hoc networks, minimum-energy networks, energy efficiency, broadcast algorithms, distributed algorithms

*The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322. (<http://www.terminodes.org>)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MOBICOM'02, September 23-28, 2002, Atlanta, Georgia, USA
Copyright 2002 ACM 1-58113-486-X/02/0009 ...\$5.00.

1. INTRODUCTION

In recent years, all-wireless networks have attracted significant attention due to their potential applications in civil and military domains [17, 18, 11, 4]. An all-wireless network consists of numerous devices¹ that are equipped with processing, memory and wireless communication capabilities, and are linked via short-range ad hoc radio connections. This kind of network has no pre-installed infrastructure, but all communication is supported by multi-hop transmissions, where intermediate nodes relay packets between communicating parties. Each node in such a network has a limited energy resource (battery), and each node operates unattended. Consequently, energy efficiency is an important design consideration for these networks [23, 26].

The broadcast communication is an important mechanism to communicate information in all-wireless networks. This is because the network described above can be regarded as a *distributed system* (distributed hardware + distributed control + distributed data), where broadcast is an important communication primitive. In addition, many routing protocols for wireless ad-hoc networks need a broadcast mechanism to update their states and maintain the routes between nodes [21].

In this paper, we focus on source-initiated broadcasting of data in static all-wireless networks. Data are distributed from a source node to each node in a network. Our main objective is to construct a *minimum-energy broadcast tree*² rooted at the source node. Nodes belonging to a broadcast tree can be divided into two categories: *relay* nodes and *leaf* nodes. The relay nodes are those that relay data by transmitting it to other nodes (relaying or leaf), while leaf nodes only receive data. Each node can transmit at different power levels and thus reach a different number of neighboring nodes. Given the source node r , we want to find a set consisting of pairs of relaying nodes and their respective transmission levels so that all nodes in the network receive a message sent by r , and the total energy expenditure for this task is minimized. We call this broadcasting problem the *minimum-energy broadcast* problem.

We base our work on the so called *node-based* multicast model [25]. In this model there is a trade-off between reach-

¹Throughout the paper we refer to these devices as nodes.

²In this paper, the words energy and power are used interchangeably.

ing more nodes in a single hop using higher power and reaching fewer nodes using lower power. This trade-off is possible due to the *broadcast nature* of the wireless channel.

This paper is organized as follows. In Section 2, we overview related work concerning minimum-energy broadcast. In Section 3, we discuss the system model used. In Section 4, we prove that the minimum-energy broadcast problem is NP-complete and show that it cannot be approximated better than $O(\log d)$ for a general graph, where d is the maximum node degree in a network; we also give a proof of the NP-completeness of the geometric version of the minimum-energy broadcast problem. Then we describe an approximation algorithm and its distributed implementation in Section 5. Performance evaluation results are presented in Section 6. Finally we conclude in Section 7.

2. RELATED WORK

The problem of minimizing the energy consumption of all-wireless networks has received significant attention over the last few years [1, 20, 9, 13, 27, 14, 22]. We are inspired by exciting results related to the problem of minimum-energy broadcasting in all-wireless networks [23, 24, 5, 12], and in particular by the work of Wieselthier et al. [25, 26]. In this work they have introduced the node-based multicast model for wireless networks upon which they have built several broadcast and multicast heuristics. One of the most notable contributions of their work is the Broadcast Incremental Power (BIP) algorithm. The main objective of BIP is to construct a minimum-energy broadcast tree rooted at the source node. It constructs the tree by first determining the node that the source can reach with minimum expenditure of power. After the first node has been added to the tree, BIP continues by determining which uncovered node can be added to the tree at *minimum additional* cost. Thus at some iteration of BIP, the nodes that have already included some node in the tree can additionally increase their transmission power to reach some other yet uncovered node. The BIP algorithm can be regarded as Prim’s algorithm [3] for the formation of minimum spanning trees, but with the difference that weights, with BIP, are dynamically updated at each step. Also notice that BIP is a centralized algorithm.

In [24] Wan et al. have given the first analytical results for minimum energy broadcast. By exploring geometric structures of an Euclidean minimum spanning tree (MST), they have proved that the approximation ratio of MST is between 6 and 12, and the approximation ratio of BIP is between $\frac{13}{6}$ and 12. Wan et al. have also found that for some instances BIP fails to use the broadcast nature of the wireless channel. This happens because BIP adds just one node at each iteration, the one that can be added at minimum additional cost. Thus BIP, although centralized, doesn’t use all available information about the network. For this reason it may end up in a broadcast tree that coincides with the shortest path tree of a network graph, where the broadcast nature of the media is completely ignored. A possible approach to cope with this is to allow an algorithm to add to a tree more than one node at each iteration, and not necessarily at minimum additional cost. However, in this case there must be another criterion for the selection of nodes in a broadcast tree. Another difficulty with BIP is that it is not obvious how to distribute it, and according to the authors of BIP and the authors in [24] the development of distributed algorithms is the major challenge considering the minimum energy broad-

cast. However, Wan et al. [24] and Wieselthier et al. [25] do not really address this challenge. In Sections 5.1 and 5.2 we will describe a possible approach to the above problems.

Li et al., in another closely related work [12], also have recognized weaknesses of BIP and proposed another centralized heuristic to attack the broadcasting problem. However, they haven’t considered the issue of developing a distributed algorithm for a minimum energy broadcast. Li et al. [12] have also given a sketch proof of the NP-hardness of a general version of the minimum energy broadcast.

A proof of NP-hardness of the minimum energy broadcast problem in metric space has been given by Egecioğlu et al. [5]. However, in their interpretation of the minimum energy broadcast problem, they restrict a node to select the transmission radius only from a set of *integers*, which captures very few instances of the problem in metric space.

Very recently, it was brought to our attention that other researchers have also studied this same problem of minimum-energy broadcasting in all-wireless networks [2, 6, 15].

In the following section, we describe a system model for all-wireless networks that will be used throughout the paper.

3. SYSTEM MODEL

We first give a wireless communication model and then, based on it we develop a graph model, which will be used to assess the complexity of the minimum-energy broadcast problem and to develop an approximation algorithm.

In our model of a wireless network, nodes are stationary. We assume the availability of a large number of bandwidth resources, i.e. communication channels. This is so because, in this paper, we are focused only on minimum energy broadcast communication and do not consider issues like contention for the channel, lack of bandwidth resources etc. We also assume that nodes in a network are equipped with omnidirectional antennas. Thus by a single transmission of a transmitting node, due to the broadcast nature of wireless channels, all nodes that fall in the transmission range of the transmitting node can receive its transmission. This property of wireless media is called *Wireless Multicast Advantage*, which we refer to as WMA [25].

In this model each node can choose to transmit at different power levels, which do not exceed some maximum value p_0 . Let P denote the set of power levels at which a node can transmit³. When a node i transmits at some power level $p \in P$, we assign it a weight, which we call a *node power*, that is equal to the power at which node i has transmitted, that is p . The connectivity of the network depends on the transmission power. Node i is said to be *connected* to node j if node j falls in the transmission range of node i . This link is then assigned a *link cost* c_{ij} , which is equal to the minimum power that is necessary to sustain link (i, j) .

Next we give a graph model for wireless networks that captures important properties of wireless media (including the wireless multicast advantage). An all-wireless network can be modeled by a directed graph $G = (V, E)$, where V represents the finite set of nodes and E the set of communication links between the nodes. Each edge (arc) $(i, j) \in E$

³We assume the cardinality of P to be finite; this does not reduce the generality of our approach, as this cardinality can be arbitrarily large.

has link cost $c_{ij} \in \mathbb{R}_+$ assigned to it⁴, while each node $i \in V$ is assigned a *variable node power* p_i^v . The variable node power takes a value from the set P defined above. Initially, the variable node power assigned to a node is equal to zero, and is set to value $p \in P$ after the node has transmitted at p . Let V_i denote the set of *neighbors* of node i . Node j is said to be a neighbor of node i if node j falls in the maximum transmission range of node i , which is determined by p_0 . All nodes $j \in V_i$ that satisfy $c_{ij} \leq p_i^v$ are said to be *covered* by node i . Thus, if node i transmits at power p_0 , all the nodes from V_i will be covered.

Now that we have the model, we study in detail the intrinsic complexity of the minimum-energy broadcast problem in the following section.

4. COMPLEXITY ISSUES

The problem of finding a minimum energy broadcast tree in wireless networks appears to be hard to solve [25]. For example, a simple analysis can show that given an instance of the minimum-energy broadcast problem, the number of possible broadcast trees is exponential in the number of nodes $|V|$ (when each node can reach all other nodes). This is easy to see by assigning each node a binary variable, which indicates whether the node transmits or not, and then by calculating the number of possible combinations of transmitters. An even more difficult problem is obtained when nodes are allowed to transmit at $|P|$ different power levels. Hence, acquiring insights into the complexity of the minimum-energy broadcast problem is of great importance. In what follows we give an in-depth analysis of the complexity of the minimum-energy broadcast problem.

Let us first briefly remind a few concepts from complexity theory [8]. The problems polynomially solvable by *deterministic* algorithms belong to the P class. On the other hand, all the problems solvable by *nondeterministic* algorithms belong to the NP class. It can easily be shown that $P \subseteq NP$. Also, there is widespread belief that $P \neq NP$. The theory of complexity is designed to be applied only to *decision problems*, i.e., problems which have either *yes* or *no* as an answer. Notice that each optimization problem can be easily stated as the corresponding decision problem. Informally, a decision problem Π is said to be NP-complete if $\Pi \in NP$ and for all other problems $\Pi' \in NP$, there exists a polynomial transformation from Π' to Π (we write $\Pi' \propto \Pi$) [8]. There are two important properties of the NP-complete class. If any NP-complete problem could be solved in polynomial time, then all problems in NP could also be solved. If any problem in NP is intractable⁵, then so are all NP-complete problems. Presently, there is a large collection of problems considered to be intractable.

In this section, we consider the problem of minimum-energy broadcast in two different graph models, namely a general graph and a graph in Euclidean metric space. In general graphs, links are arbitrarily distributed, and have arbitrarily weights chosen from the set P . This graph model is well suited for modeling wireless networks in indoor environments. On the other hand, for graphs in Euclidean metric space, the existence and the weight of the link between two nodes depends exclusively on the distance between the

⁴We designate with \mathbb{R}_+ strictly positive reals.

⁵We refer to a problem as intractable if no polynomial time algorithm can possibly solve it.

nodes and their transmission levels. This graph model fits well for outdoor scenarios.

4.1 General graph version

In the following we show that a general graph version of the minimum-energy broadcast problem is intractable, that is, it belongs to the NP-complete class. Because of its similarity to the well known *Set Cover* problem [10], which aims at finding the minimum cost cover for a given set of nodes, we call it the *Minimum Broadcast Cover* and refer to it as MBC. A decision problem related to the minimum broadcast cover problem can be described as follows:

MINIMUM BROADCAST COVER (MBC)

INSTANCE: A directed graph $G = (V, E)$, a set P consisting of all power levels at which a node can transmit, edge costs $c_{ij} : E(G) \rightarrow \mathbb{R}_+$, a source node $r \in V$, an assignment operation $p_i^v : V(G) \rightarrow P$ and some constant $B \in \mathbb{R}_+$.

QUESTION: Is there a node power assignment vector $A = [p_1^v \ p_2^v \ \dots \ p_{|V|}^v]$ such that it induces the directed graph $G' = (V, E')$, where $E' = \{(i, j) \in E : c_{ij} \leq p_i^v\}$, in which there is a path from r to any node of V (all nodes are covered), and such that $\sum_{i \in V} p_i^v \leq B$?

Notice that the above question is equivalent to asking if there is a broadcast tree rooted at r with total cost B or less, and such that all nodes in V are included in the tree (covered).

We prove NP-completeness of MBC for a general graph by showing that a special case of it is NP-complete. In order to obtain a special case of MBC, we specify the following restrictions to be placed on the instances of MBC. Each node is assigned just one power level $p \in P$ at which it can transmit. Consequently, the power level assigned to each node is either 0 (the node doesn't transmit) or p . We call this special case SINGLE POWER MBC. We prove NP-completeness of the SINGLE POWER MBC problem by reduction from the SET COVER (SC) problem, which is well known to be NP-complete [8].

SET COVER (SC)

INSTANCE: A set I of m elements to be covered and a collection of sets $S_j \in I$, $j \in J = \{1, \dots, n\}$. Weights w_j for each $j \in J$, and a constant $B \in \mathbb{R}^+$.

QUESTION: Is there a subcollection of sets C that form a cover, i.e., $\cup_{j \in C} S_j = I$ and such that $\sum_{j \in C} w_j \leq B$?

First we describe the construction of a graph G that represents any instance of the set cover problem. The graph G has a vertex set $I \cup \{v_1, v_2, \dots, v_n\}$, that is, G consists of elements of I and set vertices v_j representing sets $S_j \in I$, $j \in J = \{1, \dots, n\}$. There is an edge between an element $e \in I$ and a set node v_i if the set S_i contains the element. Each set node v_i is assigned the weight w_i of the set S_i the node represents. All other nodes and all edges are not weighted, that is, they have zero weight. Thus, $G = (V, E)$ is a bipartite graph, as is illustrated in Figure 1a.

The transformation from SC to SINGLE POWER MBC first consists in adding a source (root) node r to G and making it adjacent to all the set nodes v_j . Then, a zero weight is assigned to the root node r while all other weights are kept the same. The resulting graph, which we denote with $G_b = (V_b, E_b)$, is illustrated in Figure 1b. It is easy to see that the transformation can be done in polynomial time. Notice that without any loss of generality we can use undirected graphs for our purposes. This is because we can easily transform an undirected graph to a directed one by simply exchanging

each undirected edge with two edges directed in opposite directions.

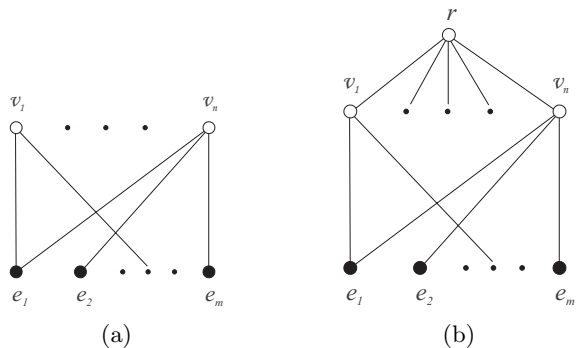


Figure 1: The reduction of (a) SET COVER to (b) SINGLE POWER MINIMUM BROADCAST COVER

Next we prove the following theorem.

THEOREM 1. *SINGLE POWER MBC is NP-complete.*

PROOF. The proof consists first in showing that SINGLE POWER MBC belongs to the NP class, and then in showing that the above polynomial transformation (Figure 1) reduces SC to SINGLE POWER MBC.

It is easy to see that SINGLE POWER MBC belongs to the NP class since a nondeterministic algorithm need only guess a set of transmitting nodes ($p_i^v > 0$) and check in polynomial time whether there is a path from the source node r to any node in a final solution, and whether the cost of the final solution is $\leq B$.

We continue the proof by showing that given the minimum broadcast cover C_b of G_b with cost $cost(C_b)$, the set $C_b - \{r\}$ always corresponds to the minimum set cover C of G of the same cost ($cost(C) = cost(C_b)$), and vice versa.

Let C denote the minimum set cover of G . Let $cost(C) = \sum_{j \in C} w_j$ denote the cost of this cover. It is easy to see that all nodes of G_b can also be covered with total cost $cost(C)$. This can be achieved by having the source node r cover all the set nodes v_j , $j \in J = \{1, \dots, n\}$ at zero cost, and then by selecting among the covered nodes those corresponding to the nodes of G that satisfy $v_j \in C$ as new transmitting nodes, which we refer to as $C_b - \{r\}$. Therefore the minimum broadcast cover of G_b is C_b with total cost $cost(C_b) = cost(C)$.

Conversely, suppose that we have the minimum broadcast cover C_b of G_b with total cost $cost(C_b)$. Then the minimum set cover C of G consists of nodes corresponding to those nodes of G_b that satisfy $v_j \in C_b - \{r\}$. We prove this by contradiction. Let C' denote the minimum set cover of G such that $C' \neq C$. In this case, by the same reasoning as before, G_b can be covered by some C'_b that satisfies $cost(C'_b) \leq cost(C_b)$. However this contradicts the preceding assumption that C_b is the minimum broadcast cover of G_b and concludes the proof. \square

Since the SINGLE POWER MBC problem is a special case of the MBC problem, and MBC belongs to the NP class, which can be shown along the similar lines as for THE SINGLE POWER MBC problem, we have the following corollary.

COROLLARY 1. *MINIMUM BROADCAST COVER (MBC) is NP-complete.*

We saw in the proof of Theorem 1 that the minimum cover sets of SC and SINGLE POWER MBC differ in only one item, namely, the source node r . However, the weight assigned to r is zero and thus the costs of the minimum cover sets of SC and SINGLE POWER MBC are the same. Hence, the transformation from SC to SINGLE POWER MBC preserves approximation ratios that can be achieved either for SC or MBC (generalization of SINGLE POWER MBC). It is known that no polynomial-time approximation algorithm for SC achieves an approximation ratio smaller than $O(\log d)$ if $P \neq NP$, where d is the size of largest set S_j [10]. Thus, for a general graph and arbitrary weights, we cannot expect to obtain an approximation algorithm for MBC that achieves the approximation ratio better than $O(\log d')$, where d' represents the maximum node degree in a graph.

Fortunately, this is not necessarily true for all instances of the minimum-energy broadcast problem. By exploring the geometric structure of the minimum-energy broadcast problem, Wan et al. in [24] were able to show that the Euclidean minimum spanning tree approximates the minimum-energy broadcast problem within a factor of 12. However, whether the geometric instances of the minimum-energy broadcast problem can be solved in polynomial time was left as an open question. We provide answer in this section.

4.2 Geometric version

In this section, we prove that the minimum-energy broadcast problem in two-dimensional Euclidean metric space is intractable. In metric space, the distance between points (nodes) obey triangle inequality, that is, $d_{ij} \leq d_{ik} + d_{kj}$, where d_{ij} is the Euclidean distance between nodes i and j . We have seen that given the graph version of the minimum-energy broadcast problem we could have arbitrary costs of links between nodes. This is because we haven't had to worry about the distances between nodes, and yet all links have been dictated by a given graph. However, in metric space links and their respective costs are dictated by the distances between nodes and their transmission energies. The cost c_{ij} between two nodes i and j is given as

$$c_{ij} = kd_{ij}^\alpha$$

where $k \in \mathbb{R}^+$ is constant depending on the environment, d_{ij} is the distance between the node i and j , and α is a propagation loss exponent that takes values between 2 and 5 [19].

We refer to this instance of the minimum-energy broadcast problem as to the *Geometric Minimum Broadcast Cover* problem and denote it with GMBC. A decision problem related to the GMBC problem can be formulated as follows: **GEOMETRIC MINIMUM BROADCAST COVER (GMBC)**

INSTANCE: A set of nodes V in the plane, a set P consisting of all power levels at which a node can transmit, a constant $k \in \mathbb{R}_+$, costs of edges $c_{ij} = kd_{ij}^\alpha$ where d_{ij} is Euclidean distance between i and j , a real constant $\alpha \in [2..5]$, a source node $r \in V$, an assignment operation $p_i^v : V(G) \rightarrow P$ and some constant $B \in \mathbb{R}_+$.

QUESTION: Is there a node power assignment vector $A = [p_1^v, p_2^v, \dots, p_{|V|}^v]$ such that it induces the directed graph $G = (V, E)$, with an edge (arc) directed from node i to node j if and only if $c_{ij} \leq p_i^v$, in which there is a path from r

to any node of V (all nodes are covered), and such that $\sum_{i \in V} p_i^v \leq B$?

We prove NP-completeness of GMBC by reduction from the *planar* 3-SAT problem, which is known to be NP-complete [16].

PLANAR 3-SAT (P3SAT)

INSTANCE: A set of variables $V = \{v_1, v_2, \dots, v_n\}$ and a set of clauses $C = \{c_1, c_2, \dots, c_m\}$ (Boolean formulae) over V such that each $c \in C$ has $|c| \leq 3$. Furthermore, the bipartite graph $G = (V \cup C, E)$ is planar, where $E = \{(v_i, c_j) | v_i \in c_j \text{ or } \bar{v}_i \in c_j\} \cup \{(v_i, v_{i+1}) | 1 \leq i < n\}$ ⁶.

QUESTION: Is there an assignment for the variables so that all clauses are satisfied?

THEOREM 2. *GEOMETRIC MINIMUM BROADCAST COVER (GMBC) is NP-complete.*

PROOF. The GMBC problem belongs to the NP-class for the same reason as the SINGLE POWER MBC (see the proof of Theorem 1).

We continue the proof by showing that P3SAT polynomially reduces to GMBC. Our proof of the NP-completeness of GMBC follows Lichtenstein’s proof of the NP-completeness of the GEOMETRIC CONNECTED DOMINATING SET [16]. We encode a Boolean formula C of P3SAT by a network representing an instance of GMBC such that given the source node, all nodes in the network can be covered at minimum cost *if and only if* C is satisfiable. We first describe the

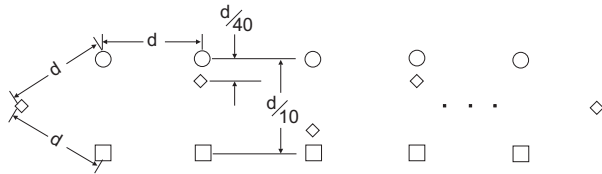


Figure 2: The structure representing a variable of P3SAT

structures we will be using in the rest of the proof. Let d denote the distance that corresponds to the maximum transmission range $p_0 = \max\{p : p \in P\}$, that is d is the farthest distance that can be reached by any node. We encode the variables of C by the structure shown in Figure 2. Let us call a group consisting of one round node, one square node and one rhombus node laying on the same line a *variable triplet*. Assume now that one variable triplet is covered. Then there are just two ways to cover the structure representing the variable at minimum cost, specifically, either all the round nodes or all the square nodes transmit. Notice that the minimum cost equals $\frac{1}{3}p_0$ times the total number of nodes in the variable. If all the round or square nodes transmit, this corresponds to the variable being set to *true* or *false*, respectively. The rhombus nodes force at least one of the two nodes adjacent to it to transmit. This structure can be arbitrary long. The distances $(\frac{d}{10}, \frac{d}{40}, \dots)$ are selected such that they ensure required properties of all the structures we will be using in the proof. Notice that these distances are not unique in this regard.

⁶We removed the edge (v_n, v_1) without any change in difficulty with the problem. See [16].

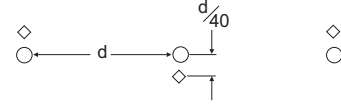


Figure 3: The line that connects variables

The variables are linked together by the connector shown in Figure 3. The connector passing through a variable is shown in Figure 4. The connector will follow the path taken by the arcs $\{(v_i, v_{i+1}) | 1 \leq i < n\}$ from the P3SAT. The arrangement of the rhombus nodes here ensures that the connector end nodes transmit at k_0 and thus cover one round and one square node belonging to a variable triplet. The rhombus node belonging to the variable triplet is moved outside the variable, since otherwise it would be covered by the connector end node, and would not force at least one of the two nearby nodes from the variable triplet to transmit.

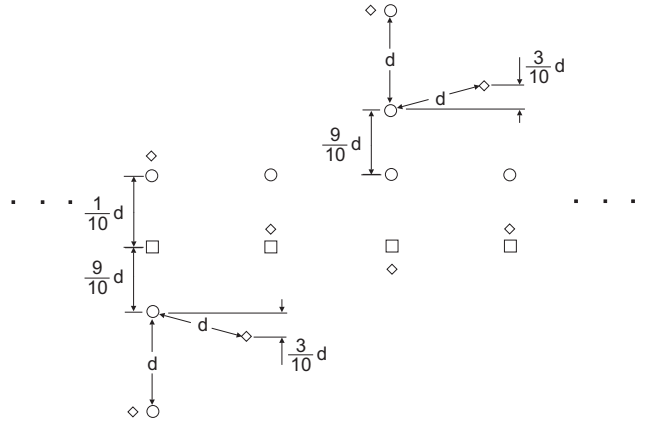


Figure 4: The “connecting” line passing through a variable

Clauses are represented by the kind of structure shown in Figure 5. Notice here that the rhombus nodes only force the round nodes from the clause to transmit. For a clause $c = (v_1 \vee v_2 \vee \bar{v}_3)$, one black node is at distance d of a round node in the structure representing the variable v_1 , a second one is at distance d of a round node representing v_2 , and a third one is at distance d of a square node representing the variable v_3 . It is important to emphasize that the black nodes are placed so that they are always in the transmission range of only a single round (square) node.

Finally, as the source node in an instance of the GMBC we choose one round node from the connector lines.

Let us introduce the following notation: N_{var} is $\frac{1}{3}$ the number of nodes in all the structures representing variables; N_{conn} the number of nodes that are forced to transmit in all the connectors; N_{cls} the number of nodes that are forced to transmit in all the clauses. Let $e_{min} = (N_{var} + N_{conn} + N_{cls} + m)p_0$. Recall that m is the total number of clauses and p_0 the maximum transmission power. Now that we have described the structure for reducing the P3SAT to GMBC, we prove that GMBC has a minimum broadcast cover of the cost e_{min} *if and only if* C is satisfiable.

Let us assume that we have an assignment of the vari-

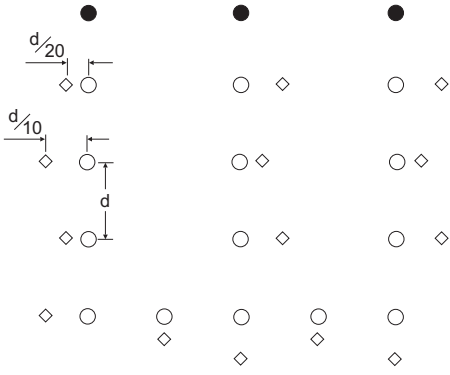


Figure 5: The structure that encodes a clause

ables from V that satisfies the Boolean formula C . Then the corresponding instance of the GMBC can be covered with e_{min} . This can be achieved by selecting the following nodes in the set of the transmitting node. We select the round (square) nodes in variables according to whether the variable is true or false in the given assignment. Then we select one black node in each clause that lies at the distance d of a round (square) node already chosen. Finally, we select all the round nodes that are forced to transmit by the corresponding rhombus nodes in each clause and in each connector.

Conversely, let us assume that we have an instance of the GMBC with a minimum cover of cost e_{min} . We will show that in this case all the structure representing variables look right, that is, no variables switch from true to false or vice versa (i.e. some round nodes and some square nodes of the same variable transmit in the same instance of the GMBC). As an immediate consequence, all the clauses have to be satisfied, otherwise they could not be covered. Let us assume that in the given instance of GMBC a variable switches from true to false. However, this would incur a larger cost to cover the variable than in the case when either all the round nodes would transmit or all the square nodes (the total cost of the cover would be $\geq e_{min}$). Consequently either all the round nodes of the variable transmit or all the square nodes do, that is, the variable looks right.

The building blocks used in our construction are of polynomial size and it requires polynomial time to put all nodes at consistent coordinates (i.e. $d, \frac{d}{10}, \frac{d}{20}, \dots$). Hence, our transformation can be done in polynomial time. This concludes the proof. \square

We have seen that the problem of minimum energy broadcast is intractable, even in two-dimensional Euclidean metric space. For this reason, in the following section, we devise a heuristic algorithm that enables us to find good solutions to the problem at reasonable computation costs.

5. PROPOSED ALGORITHMS

In this section, we will first provide the description of a centralized heuristic algorithm. We will then show that it can easily be distributed.

5.1 A heuristic based approach

Let us first provide an informal description of the algorithm we propose. We begin with a feasible solution (an

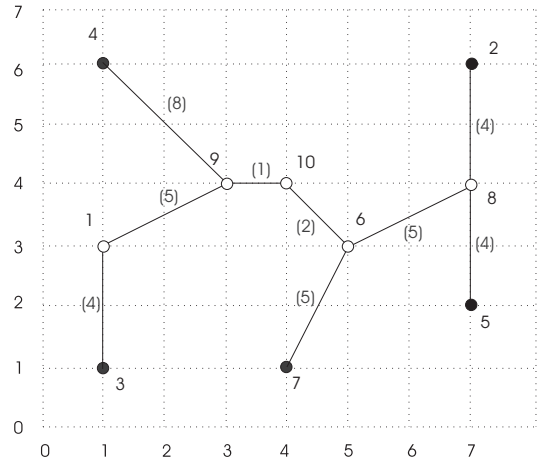


Figure 6: The network example and its MST ($e_{MST} = 23$)

initial feasible broadcast tree) for a given network. Then we improve that solution by exchanging some existing branches in the initial tree for new branches so that the total energy necessary to maintain the broadcast tree is lower. This is done so that the feasibility of the obtained solution remains intact. We call the difference in the total energies of the trees before and after the branch exchange a *gain*. In our heuristic, the notion of gain is used as the criterion for the selection of transmitting nodes in a broadcast tree.

We use the link-based minimum spanning tree (MST) as the initial feasible solution. The main reason we take MST is that it performs quite well even as a final solution to our problem, which can be seen from the simulation results in Section 6. Notice that although we use link-based MST, which doesn't exploit WMA, the evaluation of its cost takes into consideration the WMA [25]. We will now describe in detail our algorithm, which we call *Embedded Wireless Multicast Advantage* and refer to as EWMA. An example is provided in Figure 6. Let us first introduce some notations. Let C denote the set of covered nodes in a network, F the set of transmitting nodes of the final broadcast tree, and E the set of *excluded* nodes. Node i is said to be an *excluded node* if node i is the transmitting node in the initial solution but is not the transmitting node in the final solution (i.e. $i \notin F$). Notice that the contents of the above sets change throughout the execution of the EWMA, and that the sets do not hold any information about the MST. Initially, $C = \{r\}$, where r is the source node (node 10 in our example), and sets F and E are empty.

In this example, we assume a propagation constant $\alpha = 2$. After the MST has been built in the *initialization* phase, we know which nodes in the MST are transmitting nodes, and their respective transmission energies. In our example the transmitting nodes are 10, 9, 6, 1, 8, and their transmission energies are 2, 8, 5, 4, and 4, respectively. The total energy of MST is $e_{MST} = 23$. Notice here that we take into consideration the WMA in the evaluation of the cost of the MST. Notice also that $C = \{10\}$, and $F = E = \{\emptyset\}$. In the second phase, EWMA starts to build a broadcast tree from nodes in the set $C - F - E$ by determining their respective gains. The gain of a node v is defined as a decrease

in the total energy of a broadcast tree obtained by excluding some of the nodes from the set of transmitting nodes in MST, in exchange for the increase in node v 's transmission energy. Notice that this increase of node v 's transmission energy has to be sufficient for it to reach all the nodes that were previously covered by the nodes that were excluded. Consequently, the feasibility of a solution is preserved. At this stage of the algorithm the set $C - F - E$ contains just the source node 10. Thus for example, in order to exclude node 8, the source node 10 has to increase its transmission energy by (see Figure 6):

$$\Delta e_{10}^8 = \max_{i \in \{2,5\}} \{e_{10,i}\} - e_{10} = 13 - 2 = 11$$

The gain (g_{10}^8) obtained in this case is:

$$g_{10}^8 = e_6 + e_8 + e_9 - \Delta e_{10}^8 = 5 + 4 + 8 - 11 = 6$$

where $e_i, i = \{6, 8, 9\}$, is the energy at which node i transmits in MST. Notice that, in addition to node 8, the nodes 6 and 9 can also be excluded.

Likewise,

$$\begin{aligned} g_{10}^1 &= e_1 + e_6 + e_8 + e_9 - \Delta e_{10}^1 = 5 \\ g_{10}^6 &= e_6 - \Delta e_{10}^6 = -2 \\ g_{10}^9 &= e_6 + e_8 + e_9 - \Delta e_{10}^9 = 6 \end{aligned}$$

Having the gains for all nodes from $C - F - E$, our algorithm selects a node with the highest positive gain in the set F . Our algorithm then adds all the nodes that this node excludes to the set E . Thus the source node 10 is selected in the set F to transmit with energy that maximizes its gain, that is:

$$e'_{10} = e_{10} + \arg \max_{\Delta e_{10}^i} \{g_{10}^i\}, \quad g_{10}^i \geq 0$$

The source node 10 transmits with energy $e'_{10} = e_{10} + \Delta e_{10}^8 = 2 + 11 = 13$ at which it can cover nodes 6, 8, 9 and all their *child* nodes in MST. Node j is said to be a *child node* of node i if node j is included in a broadcast tree by node i . Hence, at this stage we have $C = \{1, 2, 4, 5, 6, 7, 8, 9, 10\}$, $E = \{6, 8, 9\}$ and $F = \{10\}$. If none of the nodes from $C - F - E$ has a positive gain, EWMA selects among them the node that includes its child nodes in MST at minimum cost (energy).

The above procedure is repeated until all nodes in the network are covered. In our example there is still one node to be covered, namely node 3. Again, EWMA scans the set $C - F - E = \{1, 2, 4, 5, 7\}$ and at last selects node 1 to be the next forwarding node. When node 1 transmits with energy $e_1 = 4$, all nodes are covered ($C = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$) and the algorithm terminates. At the final stage we have $E = \{6, 8, 9\}$ and $F = \{1, 10\}$. The resulting tree, shown in Figure 7, has a cost $e_{EWMA} = 17$, ($e_{MST} = 23$). Notice that our algorithm always results in a broadcast tree with the total energy $\leq e_{MST}$, which is, in the case of Euclidean MST, less than $12e_{opt}$ [24].

The EWMA algorithm does not perform an exhaustive search over all possible combinations of transmitting nodes. Next we show that the running time of the EWMA algorithm is polynomial in the total number of nodes n . Let d denote the size of largest neighborhood (i.e. the maximum node degree), d_i the number of nodes that are newly covered in iteration i , and m the total number of transmitting nodes

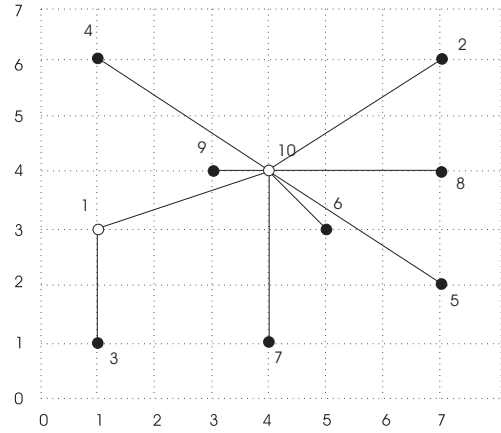


Figure 7: The broadcast-tree obtained by EWMA heuristic ($e_{EWMA} = 17$)

(i.e. $|F| = m$). Here we assume a straightforward implementation of our algorithm. Thus, in order for a node to check if it can exclude some neighbor, the node has to test all the neighbors of that neighbor, which takes $O(d)$ time. Now, in order to calculate the gain that can be attained by this exclusion, the node repeats the above procedure for all the remaining neighbors, which thus takes $O(d^2)$ time. Finally, the node repeats all the above steps for all its neighbors and decides to transmit with the energy that maximizes its gain, which takes $O(d^3)$ time. As this is repeated for, at most, all the covered nodes up to and including some iteration i , the running time of the EWMA algorithm is bounded by:

$$\begin{aligned} \sum_{i=1}^m O(d^3) \sum_{j=1}^i d_j &\leq O(d^3) \sum_{i=1}^m \sum_{j=1}^i d \\ &= O(d^3) d \sum_{i=1}^m i \\ &\leq O(d^4) m^2 \end{aligned}$$

Note that $O(d^4)m^2$ is not a tight bound on the running time of the EWMA algorithm.

5.2 Distributed implementation of EWMA

One of the major research challenges, with respect to the broadcasting problem, is the development of a distributed algorithm [25, 24]. In the following we describe our solution.

Let us first introduce the notation we will be using. Let node i transmit at power level $p \in P$. We denote the set of nodes that are covered by this transmission with V_i^p . Let node j be a neighbor of i , that is, $j \in V_i$. We denote with O_{ij}^p the set of nodes belonging to $V_i^p \cap V_j$ and call it the *overlapping* set. We assume that each node knows the cost of each edge adjacent to itself, and the identity of its neighbors. A node maintains this information in a *cost matrix*. Once node j receives a message from node i , it can learn which of the nodes from its neighbor set V_j have also received the message by calculating the overlapping set O_{ij}^p . The neighbors of node j that have not yet received the message are said to be *uncovered*, and we denote this set with U_j where $U_j = V_j - O_{ij}^p$. If node j is a forwarding node in the MST, then the set of yet uncovered children nodes of node j in

the MST is denoted with U_j^{mst} where $U_j^{mst} = V_j^{mst} - O_{ij}^p$. Here, V_j^{mst} is the set comprising all the children nodes of j in the MST. Finally, we denote with e_j^{mst} the energy with which node j transmits in the MST.

Now we describe our distributed algorithm. The algorithm is divided into two phases. In the first phase, all nodes run a distributed algorithm proposed by Gallager et al. [7] to construct a *minimum-weight spanning tree*. The total number of messages required for a graph of $|V|$ nodes and $|E|$ edges is at most $5|V| \log_2 |V| + 2|E|$, and the time until completion is $O(|V| \log |V|)$ [7]. Notice that Gallager et al. considered the link-based model, while we use the node-based multicast model, which captures the wireless multicast advantage property [25]. As a consequence, the total number of messages required in our model may be considerably lower. We require that at the end of the first phase, each node has information about the cost of its two-hop neighbors related to the built MST. This can be achieved by piggybacking information about the cost on regular messages.

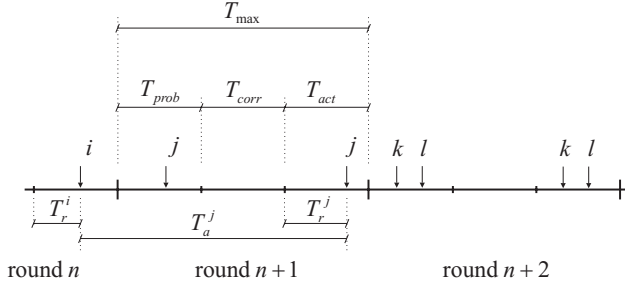


Figure 8: Synchronization of the second phase

In the second phase, the final broadcast tree is built up. The main difficulty in this distributed setting is the unavailability of information about which nodes have been covered up to a certain moment. In order to cope with this problem we apply two techniques. First, we organize this second phase in rounds. Second, we require that the identities of the nodes on the transmission chain from the source to some node and their respective transmission powers are propagated along that chain to the node in question (source routing like technique).

Each round of the second phase is T_{max} long. Rounds are additionally divided into three time periods, namely, a *probation* period (T_{prob}), a *correction* period (T_{corr}), and an *active* period (T_{act}), which are all known by network nodes (Figure 8). Let node i transmit at T_r^i time from the beginning of the active period of round n . Node j receives this message and starts the update procedure shown in Figure 9. Thus, node j calculates the overlapping set for the sender i and for other transmitters on this chain of transmitting nodes for which node j has neighbors in common (recall that this information is propagated along the chain). If node j is a forwarding node in the MST and it finds that the set of uncovered nodes U_j^{mst} is empty for the received message, it will not re-broadcast the message. Otherwise, (namely if U_j^{mst} is non-empty or j was a leaf node in the MST), it calculates the gains it can achieve by covering yet uncovered nodes (based on locally available information), and selects the maximum gain $g_{j \max}$. In the case $g_{j \max} > 0$, node j

$$U_j^{mst} \leftarrow V_j^{mst} \quad e_j \leftarrow e_j^{mst}$$

LABEL B: upon receive a msg from node i

$$O_{ij}^p \leftarrow V_i^p \cap V_j^{mst} \quad U_j^{mst} \leftarrow U_j^{mst} - O_{ij}^p$$

if ($U_j^{mst} = \emptyset$ AND $e_j > 0$) **then**

HALT

else

for all $l \in V_j$ **do**

 calculate-gains g_j^l

$g_{j \max} \leftarrow \max_l \{g_j^l\}$

if $g_{j \max} > 0$ **then**

$e_j \leftarrow e_j + \arg \max_{\Delta e_j^l} \{g_j^l\}$

$$T_r^j \leftarrow \frac{\Delta_1}{g_{j \max}}$$

else if $e_j > 0$ **then**

$$T_r^j \leftarrow \Delta_2 \cdot e_j$$

else

HALT

 wait $T_a^j \leftarrow T_{max} + T_r^j - T_r^i$

if during T_{act} and before expiration of T_a^j same msg received **then**

goto B:

else

 broadcast the msg at power e_j

Figure 9: Update session at node j - message reception from node i

can contribute to the decrease of the total cost of the broadcast tree and its transmission energy increases as follows: $e_j = e_j + \arg \max_{\Delta e_j^l} \{g_j^l\}$, otherwise ($g_{j \max} \leq 0$) its transmission energy remains unchanged. Notice here that the leaf nodes re-broadcast a message only if they can achieve a positive gain.

At this stage, node j waits for some time period T_a^j before possibly re-broadcasting the message. The waiting period is given as follows:

$$T_a^j = T_{max} + T_r^j - T_r^i$$

where $T_r^j = \frac{\Delta_1}{g_{j \max}}$ if $g_{j \max} > 0$, and $T_r^j = \Delta_2 \cdot e_j$ if $g_{j \max} \leq 0$ and $e_j > 0$. In the first case the waiting period T_a^j is reciprocal to the gain, in order to give advance to nodes with higher positive gains over nodes with lower positive gains. In the second case, the waiting period T_a^j is proportional to the transmission energy in order to give advance to nodes with lower transmission energies over nodes with higher transmission energies. Additionally, the nodes with positive gains are preferable to the nodes with low transmission energies (i.e. $\frac{\Delta_1}{g_{j \max}} \ll \Delta_2 \cdot e_j$). This fact is captured by setting appropriately the constants Δ_1 and Δ_2 .

Since node j calculates the gains based on only locally available information, it can happen that in the calculation of the gains, it tries to exclude already excluded nodes. In order to prevent this, node j transmits a probe message during the probation period T_{prob} of round $n + 1$. Note that by knowing T_{act} and T_r^i , node j actually knows when round $n + 1$ starts. The probe message carries the addresses of all the nodes by exclusion of which node j attains $g_{j \max} > 0$, and it carries the starting time of the correction period. If some of these nodes have already been excluded, they will respond back to node j during the correction period. Node

j will accordingly update its gain and the waiting period T_a^j by taking into account the already elapsed time of the waiting period. The duration of the probation and correction periods should be such that any potential forwarding node is given the chance to test its prospect of actually being the forwarding node. It is important to stress that only those nodes that obtain $g_{j \max} > 0$ do the probation. Note that throughout the execution of the second phase, the waiting period T_a^j is counted down.

Finally, node j enters into the active period. Again, node j , based on the knowledge of T_{prob} and T_{corr} , knows when the active period of round $n + 1$ starts. If during that period and before expiration of the waiting period T_a^j node j receives a duplicate message, it repeats the update procedure in Figure 9, otherwise, upon expiration of T_a^j , it re-broadcasts the message with energy e_j , stores this value and marks itself as the forwarding node. In our example shown in Figure 8, node j decides to be the forwarding node and broadcasts a message at power e_j . By doing so, it initiates the update procedure at nodes k and l , which repeat the whole process. Recall that node j also sends information about all the transmitting nodes on the chain from the source node to node j .

Next we show under which conditions the waiting period T_a^j expires solely during the active period of round $n + 1$. From Figure 8 we can see that this happens if T_a^j conforms to the following conditions:

$$\begin{aligned} T_a^j &\geq T_{\max} - T_r^i \\ T_a^j &\leq T_{\max} + T_{act} - T_r^i \end{aligned}$$

From the first inequality and the definition of T_a^j we obtain that $T_r^j \geq 0$, which is always satisfied. Along the same lines, from the second inequality we obtain that $T_r^j \leq T_{act}$. Consequently, we define the active period as follows:

$$\begin{aligned} T_{act} &= \max_{j \in F} \{T_r^j\} \\ &= \max_{j \in F} \{\Delta_2 \cdot e_j\} \end{aligned}$$

where the second equality follows from the fact that $\frac{\Delta_1}{g_{j \max}} \ll \Delta_2 \cdot e_j$. Now, since we already have decided on Δ_1 and Δ_2 , we only have to find the cost of the most expensive edge in the MST. Note that this information can be obtained from the first phase of the algorithm. This in addition to the appropriate selection of the periods T_{prob} and T_{corr} , ensures a synchronous execution of the second phase of the distributed algorithm.

The duration of the second phase is bounded by $|F| \cdot T_{\max}$, where F is the set of the forwarding nodes at the end of the second phase. Thus, at the end of the second phase, the broadcast tree is built (i.e. we have a set of forwarding nodes F and their respective transmission energies for a given source node). Any subsequent broadcast message can be disseminated along the tree in an asynchronous way (i.e. forwarding nodes may re-broadcast a message immediately upon receiving it).

6. PERFORMANCE EVALUATION

We performed a simulation study to evaluate our centralized algorithm (EWMA) and its distributed version.

We compared the centralized version of our algorithm (EWMA) with BIP and MST algorithms. The simulations were performed using networks of four different sizes: 10, 30,

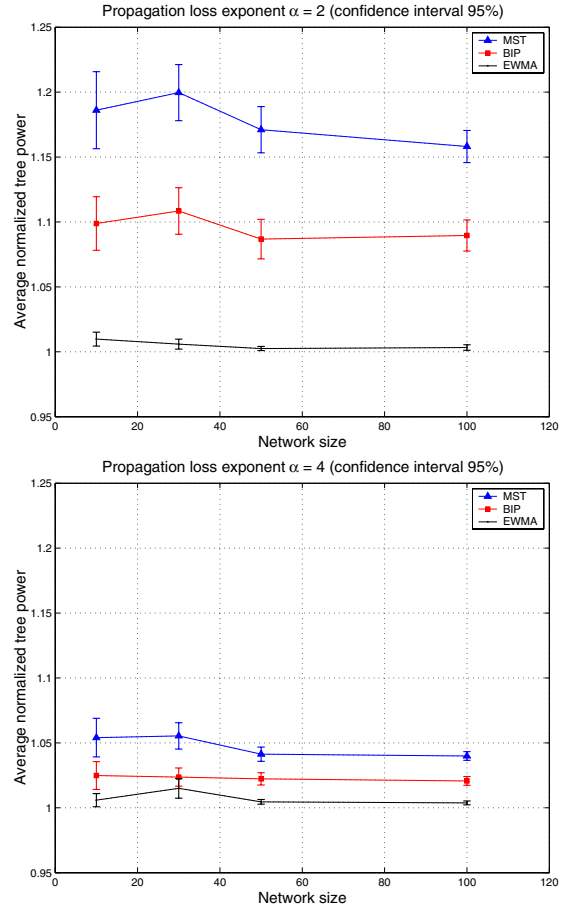


Figure 10: Normalized tree power for 100 network instances (confidence interval 95%) and propagation loss exponent $\alpha = 2$ (above) and $\alpha = 4$ (below)

50 and 100 nodes. The nodes in the networks are distributed according to a spatial Poisson distribution over the same deployment region. Thus, the higher the number of nodes, the higher the network density. The source node for each simulation is chosen randomly from the overall set of nodes. The maximum transmission range is chosen such that each node can reach all other nodes in the network. The transmission power used by a node in transmission (d^α) depends on the reached distance d , where the propagation loss exponent α is varied. Similarly to Wieselthier et al. in [25], we ran 100 simulations for each simulation setup consisting of a network of a specified size, a propagation loss exponent α , and an algorithm. The performance metric used is the total power of the broadcast tree. Here we use the idea of the *normalized tree power* [25]. Let $p_i(m)$ denote the total power of the broadcast tree for a network instance m , generated by algorithm i ($i = \{EWMA, BIP, MST\}$). Let p_0 be the power of the lowest-power broadcast tree among the set of algorithms performed and all network instances (100 in our case). Then the normalized tree power associated with algorithm i and network instance m is defined as follows: $p'_i(m) = \frac{p_i(m)}{p_0}$.

Let us first consider the performance of the algorithms

shown in Figure 10. In the figure we can see the average normalized tree power (shown on the vertical axis) achieved by the algorithms on networks of different sizes (the horizontal axis). To estimate the average power, we have used an interval estimate with the confidence interval of 95%. The figure shows that the solutions for the broadcast tree obtained by EWMA have, on the average, lower costs than the solutions of BIP and MST. (This is also true for $\alpha = 3$, which is not shown in the figure). However, we can notice that for the propagation loss exponent of $\alpha = 4$, the confidence intervals of the algorithms overlap for certain cases, which means that the solutions provided by the algorithms are not significantly different. Thus the figure also reveals that the difference in performance decreases as the propagation loss exponent increases. The main reason for such behaviour is that by increasing the propagation loss exponent, the cost of using longer links increases as well. Consequently, EWMA and BIP select their transmitting nodes to transmit at lower powers, which is typical for the transmitting nodes of MST. Hence, in a sense, EWMA and BIP's broadcast trees *converge* to the MST tree when α increases. This indicates that in scenarios where α takes higher values, MST performs quite well.

We also conducted a simulation study of the distributed algorithm presented in Section 5.2. The performance metric used here is the same as in the case of the centralized algorithm, and based again on the normalized tree power. However, here we do not consider the cost of building a broadcast tree, but only the cost of the final tree produced by the distributed algorithm. The performance of the distributed algorithm is compared to that of the centralized algorithms, and is shown in Figure 11. We can see that broadcast trees produced by distributed EWMA have, on the average, lower costs than those obtained by the centralized BIP and MST. Also, we can see that distributed EWMA performs slightly worse than its centralized counterpart. Note that the results for the centralized algorithms differ between Figure 10 and Figure 11. This is because here we run another set of simulations for all the algorithms, and for each network the source node is chosen at random.

Based on our simulation results, we conclude that EWMA utilizes the wireless multicast advantage property at least as well as BIP. The main problem with BIP is that it is not easy to distribute. We showed in Section 5.2 that EWMA can be easily distributed by using the mechanism of the waiting periods. For these reasons, we think EWMA to be preferable to BIP.

7. CONCLUSION

We have provided novel contributions on the two most relevant aspects of power-efficient broadcast in all-wireless networks. First, we studied the complexity of the problem. We discussed two configurations, represented each by a specific graph: a general graph and a graph in Euclidean space (geometric case). For both, we provided a proof that the problem is NP-complete.

Second, we elaborated an algorithm, called Embedded Wireless Multicast Advantage (EWMA). We showed that this algorithm outperforms one of the most prominent proposals provided in the literature, BIP. Moreover, we described a fully distributed version of EWMA, a feature that other authors have reckoned to be both necessary and challenging, and for which we could not find a solution in the

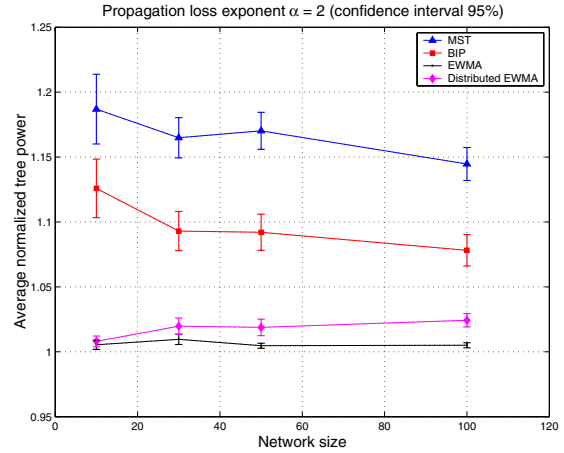


Figure 11: Distributed algorithm - normalized tree power for 100 network instances (confidence interval 95%) and propagation loss exponent $\alpha = 2$

literature.

In terms of future work, we intend to explore how other mechanisms can be used to further reduce power consumption. Moreover, we will explore how to extend our proposal to multicast. Finally, we intend to study how to cope with the mobility of the nodes.

8. ACKNOWLEDGMENTS

We thank to Naouel Ben Salem for helpful discussions and comments. We also thank to the anonymous referees for their comments.

9. REFERENCES

- [1] J. Chang and L. Tassiulas. "Energy Conserving Routing in Wireless Ad-hoc Networks". In *Proceedings of IEEE INFOCOM 2000*. ACM, 2000.
- [2] A. Clementi, P. Crescenzi, P. Penna, G. Rossi, and P. Vocca. On the complexity of computing minimum energy consumption broadcast subgraphs. In *Proceedings of 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS 2001)*, 2001.
- [3] T. Cormen, C. Leiserson, R. Rivest, and C. Stein. *Introduction to Algorithms. Second Edition*. The MIT Press, 2001.
- [4] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networks (MOBICOM 1999)*, Seattle, Washington, August 1999.
- [5] O. Egecioglu and T. Gonzalez. Minimum-energy broadcast in simple graphs with limited node power. In *Proceedings of IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS 2001)*, pages 334-338, Anaheim, CA, August 2001.
- [6] A. Faragó and V. Syrotiuk. Algorithmic problems in power controlled ad hoc networks. In *Proceedings of the 14th International Conference on Parallel and*

- Distributed Computing Systems (PDCS 2001)*, Dalas, Texas, August 2001.
- [7] R. Gallager, P. Humblet, and P. Spira. A distributed algorithm for minimum-weight spanning trees. *ACM Transactions on Programming Languages and Systems*, 5(1):66–77, January 1983.
- [8] M. Garey and D. Johnson. *Computers and intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
- [9] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocol for wireless microsensor networks. In *Proceedings of the 33rd Hawaii International Conference on System Sciences (HICSS 2000)*, Maui, Hawaii, January 2000.
- [10] D. S. Hochbaum. *Approximation Algorithms for NP-hard problems*. PWS Publishing Company, 1997.
- [11] J.-P. Hubaux, T. Gross, J.-Y. L. Boudec, and M. Vetterli. Towards self-organized mobile ad hoc networks: the terminodes project. *IEEE Communications Magazine*, 39(1), January 2001.
- [12] F. Li and I. Nikolaidis. On minimum-energy broadcasting in all-wireless networks. In *Proceedings of the 26th Annual IEEE Conference on Local Computer Networks (LCN 2001)*, Tampa, Florida, November 2001.
- [13] L. Li, V. Bahl, Y. Wang, and R. Wattenhofer. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *Proceedings of IEEE INFOCOM 2001*, April 2001.
- [14] Q. Li, J. Aslam, and D. Rus. Online power-aware routing in wireless ad-hoc networks. In *Proceedings of the Seventh Annual International Conference on Mobile Computing and Networks (MOBICOM 2001)*, Rome, Italy, July 2001.
- [15] 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 2002)*, Lausanne, Switzerland, June 2002.
- [16] D. Lichtenstein. Planar formulae and their uses. *SIAM Journal on Computing*, 11(2):329–343, May 1982.
- [17] G. Pottie and W. Kaiser. Wireless integrated network sensors. *Communication of the ACM*, 43(5):51–58, May 2000.
- [18] J. Rabaey, M. Ammer, J. da Silva Jr., D. Patel, and S. Roundy. Picoradio supports ad hoc ultra-low power wireless networking. *IEEE Computer Magazine*, 33(7):42–48, July 2000.
- [19] T. S. Rappaport. *Wireless communications: Principles and Practice*. Prentice Hall, 1996.
- [20] V. Rodoplu and T. H. Meng. Minimum energy mobile wireless networks. *IEEE Journal on Selected Areas in Communications*, 17(8), August 1999.
- [21] E. M. Royer and C.-K. Toh. A review of current routing protocols for ad hoc mobile wireless networks. *IEEE Personal Communications*, 6(2):46–55, April 1999.
- [22] C. Schurgers, V. Tsiatsis, S. Ganeriwal, and M. B. Srivastava. Optimizing sensor networks in the energy-latency-density design space. *IEEE Transactions on Mobile Computing*, 1(1):70–80, January-March 2002.
- [23] S. Singh, C. Raghavendra, and J. Stepanek. Power-aware broadcasting in mobile ad hoc networks. In *Proceedings of IEEE PIMRC'99*, Osaka, Japan, September 1999.
- [24] P.-J. Wan, G. Calinescu, and O. F. X.-Y. Li. Minimum-energy broadcast routing in static ad hoc wireless networks. In *IEEE INFOCOM 2001*, Anchorage, Alaska, April 2001.
- [25] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. On the construction of energy-efficient broadcast and multicast trees in wireless networks. In *IEEE INFOCOM 2000*, pages 586–594, Tel Aviv, Israel, 2000.
- [26] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. Algorithms for energy-efficient multicasting in static ad hoc wireless networks. *ACM Mobile Networks and Applications (MONET)*, 6(3):251–263, June 2001.
- [27] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the Seventh Annual International Conference on Mobile Computing and Networks (MOBICOM 2001)*, July 2001.