Energy-Efficient Multicast with Directional Antennae and Localized Tree Reconfiguration

Carmen M. Yago
Department of Engineering of Information and Communications
Facultad de Informatica, Campus de Espinardo
Murcia, Spain
carmen@dif.um.es

Pedro M. Ruiz
Department of Engineering of Information and Communications
Facultad de Informatica, Campus de Espinardo
Murcia, Spain
pedrom@dif.um.es

ABSTRACT

Ad hoc nodes are usually battery-powered. Thus, energy-efficiency is of paramount importance in wireless ad hoc and sensor networks. The use of directional antennae helps to reduce energy consumption by aiming the beam only to the intended target using less energy. In addition, these kind of networks have limited resources in terms of bandwidth. Thus, the use of multicast can help very much at maximizing the throughput that can be obtained. Directional MIP (D-MIP) is a well-known solution to the problem of energy-efficient multicast routing with directional antennae. However, it is a centralized scheme whose applicability to these scenarios is somehow limited. Thus, we present a new distributed energy-efficient multicast protocol with directional antennae, which incorporates a localized tree reconfiguration procedure. Our simulations show that our proposed scheme (MLDP) is able to outperform the centralized scheme under a variety of scenarios and energy models.

1. INTRODUCTION AND RELATED WORK

Ad hoc and sensor networks are usually formed by battery powered devices. So finite power supplies has become one of the main challenges in ad hoc network deployment. As the same time, multicast traffic is crucial to this kind of networks to efficiently facilitate group communication and service discovery. Therefore it becomes necessary to make an efficient use of the scarce energy available, to provide an efficient multicast service. During the last years, and regarding multicast routing, the proposals to make it energy efficient have basically focused either on maximizing network lifetime, or minimizing energy consumption of a multicast tree. The former refers to finding at each stage the multicast tree that maximizes network lifetime. The latter computes the minimum energy multicast tree, which cannot be the best alternative to maximize lifetime.

Both problems usually pursue conflicting goals. Maximizing lifetime does not lead to minimum routes, and minimizing energy consumption can provoke the exhaustion of intermediate nodes. Thus, they are usually addressed separately. There has been a number of proposals addressing these NP-complete problems. For example L-REMiT [14] and MLT-MD [7] offer solutions to maximize network lifetime, whereas SPF [2] or S-REMiT[15] reduce the total energy consumption. There are also proposals like MIP [18], RB-MIP[17] and D-MIP [17] which can be used to tackle any of them depending on a cost function.

In this paper we focus our effort on minimizing the total energy consumed. This problem is also known as the MEMT problem (Minimum Energy Multicast Tree Problem)[2],[6]. Existing approximations to the MEMT problem are either to build the multicast tree via centralized protocols like MIP [18] and SPF [2], or to refine a previously created multicast tree by means of distributed procedures such as the REMiT [14], [15] techniques.

Wieselthier *et al.* proposed MIP (Multicast Incremental Protocol) which produces a multicast tree by pruning the broadcast tree resultant after running BIP [16]. SPF is an incremental protocol which begins with an empty tree which is gradually extended to connect the source with all the receivers. To obtain low energy consumption, SPF starts from a tree consisting only of the source and builds iteratively the final multicast tree. At every step, the tree is extended with the lowest energy consuming path connecting a receiver out of the tree with the tree. Both MIP and SPF are centralized schemes, whose application to ad hoc networks is very difficult. There are other distributed schemes such as REMiT (Refining Energy efficiency of Multicast Trees), which is a suite of distributed techniques to refine an initial solution. The refinement operation is also iterative. In S-REMiT [15] at every iteration, a node decides whether to replace its current parent node in the tree with another neighbor, if the new parent offers a lower energy consumption.

The aforementioned protocols assume omnidirectional antennae. However, the use of directional antennae can help very much at reducing power consumption, and increasing spatial reuse. [10]. There is a number of energy efficient protocols which make use of directional antennae to broadcast traffic. Among them, we find both centralized and localized approaches. Representative examples of the first category are RB-BIP [19] and DBIP [19] which are directional versions of BIP. In addition to them, there is an algorithm proposed by Spyropoulos and Raghavendra [12] that computes a flow matrix to approximate a good solution. Regarding localized algorithms, Cartigny, Simplot-Ryl, and Stojmenovic have published solutions using underlying graphs which are built in a localized way. DLBOP [3] is based on LMST [9], OMDLBOP [4] also uses LMST but exploits the adaptability of the antenna and lastly, ADLBOP [4] an adaptive protocol in which each node decides by itself to broadcast using DLBOP or OMDLBOP depending of its local network density.

However,there are only a few number of papers studying the use of directional antennae to improve multicast delivery. An interesting contribution is MODA, proposed in [13], although its use of directional antennae is not intended to get an energy-efficient protocol but to reduce data overhead. This is a backbone protocol which builds a multicast tree, and it is based on the idea of covering two hops in the tree when transmitting, so that the number of forwarding nodes is kept low.

Regarding energy efficiency, MLT-MD proposed by Hou *et al.* maximizes network lifetime [8] This protocol tries to maximize network lifetime for a sequence of multicast requests, each of them originating a new multicast tree. To do this, for every request it establishes a first multicast solution. Afterwards, it improves it by identifying the node with less lifetime, and revises the routing topology.

Wieselthier *et al.* proposed directional variations of MIP, RB-MIP and D-MIP [17]. RB-MIP simply builds MIP and then reduces the beam properly since D-MIP builds the broadcast tree in a directional way adding at each step the node which causes less energy consumption but in a directional way (that is incrementing the range, beam or choosing another orientation). Another work related with directional antennae and the MEMT problem is [6] by Guo and Yang. In that work, the MEMT problem is transformed in a set of linear constraints in terms of MILP (Mixed Linear Programming), which can be used for assessing the performance of MEMT algorithms. That paper applies an heuristic approximation algorithm based on the linear optimization process to enhance RB-MIP and D-MIP.

However these proposals are either centralized (like D-MIP or RB-MIP) or do not address the problem of minimizing the total energy consumption like MODA or MLT-MD. Therefore, in this context, we consider focus our research in the provision of an algorithm using directional antennae and being fully distributed. In fact, it incorporates even a localized refinement phase, to furtheris both minimize the total energy consumption.

We propose MDLP (Multicast Distributed with Localized reconfiguration Protocol), a new energy-efficient multicast protocol which, making use of directional antennae is able to build an energy-efficient multicast tree basically in a localized way. Using distributed communications only to create the the initial tree. Also, unlike many previous works, we account for the amount of energy required to build the tree, which must also be kept low.

The way to do it is the following: we start from LMST [9], a localized graph which approximates the MST; this graph has a double purpose: firstly it is the starting point for building the multicast tree, and secondly, the control messages sent to connect the source with the receivers are sent directionally following this graph. Then, the usage of both directional antennae and the LMST graph makes it possible to build an initial multicast tree in an energy-efficient way. This initial tree is then refined by a localized process in which each node considers excluding itself from the tree, if doing so improves the overall energy efficiency. Finally, every sending node decides the way of sending the message to its children. It can either send a copy of the message separately to each child or to send a single message to all of them.

The remainder of the paper is organized as follows: in the next section, we introduce some preliminary concepts which set the basis for our explanation. Then, Section 3 describes the proposed protocol itself, showing how data is delivered through the tree, how such multicast tree is built and an approximation of the energy expended in creating it. Section 4 evaluates the performance of our proposal, comparing it with D-MIP, the state-of-the-art protocol. Finally, conclusions and future work are provided in the last section.

2. PRELIMINARIES

2.1 Network Model and Multicast Tree

An ad hoc network can be modeled as a graph $G = (V, E)$. The vertex set V represents nodes and the edges represent the omnidirectional links. Two nodes are connected by an edge if their relative distance is lower than the maximal range of the antennae. That is:

$$E = \{(v_x, v_y) \mid v_x \in V,\ v_y \in V,\ dist(v_x, v_y) \leq R_{max}\}.$$

In this graph G, we consider $s \in V$ the multicast source and a set of nodes $M = \{m_1, ..., m_k\}$ the multicast receivers $M \subset V$, then, we call multicast nodes those belonging to $M_{set} = \{s\} \cup M$.

Taking these definitions into account and assuming static networks with no node exhaustion, we formalize MEMT problem as follows: given a cost function $K : E \to \mathbb{R}^+$ which represents energy consumption, the goal is to find a multicast tree MT connecting the source s with every element of M with minimum energy consumption.

Now we present some concepts related to the multicast tree that will be used in the rest of the paper. To build this multicast tree MT we start from an underlying locally built graph. The used graph is LMST and we denote it as $G_{lmst} = (V, E_{lmst})$ where $E_{lmst} \subseteq E$. We call LMST neighbors to those nodes $\{a, b \in V | e_{a,b} \in E_{lmst}\}$

From this graph we build and initial tree $IMT(V_{imt}, E_{imt})$ where $M_{set} \subseteq V_{imt},\ V_{imt} \subseteq V\ and\ E_{imt} \subseteq E_{lmst} \subseteq E$. This initial tree is refined to get the definitive multicast tree $MT(V_{mt}, E_{mt})$ where $V_{mt} \subseteq V_{imt}\ and\ E_{mt} \nsubseteq$

E_{imt} *but* $E_{mt} \subseteq E$. Then, we call relays to those nodes which have to forward the message.

2.2 Directional Antennae and Energy Model

This work has been carried out considering an idealized adaptive directional antenna. This kind of antennae can aim a beam in any orientation, varying range (from 0 to R_{max}) and varying width(from 0 to 2π). Then, a node can receive the message if it lies on the area defined by the range, width and orientation.

This capacity of adaptation allows the sender to transmit to only one receiver using a narrow beam or, to transmit to a set of receivers using a wider beam. In the former case we are talking about a one-to-one communication model and in the latter about a one-to-many communication model. In our protocol nodes conveniently switch from one to another with the aim of saving energy.

Then, taking into account adaptive antennae features, eq. 1 shows the energy spent in each sending. This energy model is the same used in [4] which is derived form the one developed by Rodoplu *et al.* in [11].

$$e(\theta, r) = \begin{cases} \frac{\theta}{2\pi}(r^{\alpha} + C_1) + C_2 & \text{if } r \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

In (1), θ represents the width beam in radians, r represents the transmission range, α is the power attenuation coefficient (whose value ranges from 2 to 4), C_1 is a constant representing the cost of aiming the beam and C_2 is a constant overhead for each sending, representing the minimum needed energy for signal processing and MAC control mechanism.

But sometimes, it is also necessary to compute the additional required amount of energy needed by a device to receive a message. Although for some technologies as IEEE 802.15, energy in reception do not increase from that spent in idle mode, in the general case, receiving a message requires a relevant increment from the energy required to listen to the channel. So, it is necessary to extend the model with a new constant C_r representing the energy expended by a node in receiving a message, and also with the number of nodes receiving such message (n). Thus (1) becomes:

$$e(\theta, r) = \begin{cases} \frac{\theta}{2\pi}(r^{\alpha} + C_1) + C_2 + n * C_r & \text{if } r \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

As we see from 2), depending on the value of α and the rest of the constants C_1, C_2,and C_r, the configuration of the tree varies. C_1 and C_2 are both associated to sending a beam but they contribute differently to the total cost: C_1 is factored by the beam width, so the wider the beam is, the more representative C_1 becomes, which means the contribution of C_1 to the consumption is greater. C_2 has always the same impact independently on the width and range of the beam and the only way to minimize its impact is to reduce the number of messages sent. C_r is associated to reception as it represents the energy expended in receiving a message, so the only way of minimizing its impact is to reduce the number of messages received by the nodes.

If the three constants C_1, C_2, and C_r are null, the energy consumption is proportional to r^{α}. Therefore, the protocol will try to send many short ranged messages. This is because energy increases exponentially with range, but neither sending more messages nor receiving them involve any additional overload.

When the receiving constant C_r is null but the sending constants C_1 and C_2 are not, the autoconfiguration process tries to reach a trade off between the beam range, the number of messages sent (which increases the impact of C_2) and the width of the beam (which increases the impact of C_1).

And when the three constants C_1, C_2, and C_r are not null, there is a cost in both aiming and receiving the beam. Then, the tree will be configured to send few beams and many times, one-to- many model will be cheaper. It is worth mentioning that there is a representative amount of energy expended in receiving the message. Thus, the autoconfiguration of the tree pursues reducing the energy in reception as much as possible: firstly by trying to build the tree with the minimum number of nodes, and secondly by trying to avoid that messages reach not intended nodes (duplicate messages, or nodes out of the tree).

Our proposal takes into account all of these features, and the configuration of the multicast tree is directly dependant upon both the particular features of the propagation media and the way the devices deal with the reception. Section 4 shows the results obtained by simulation, which prove that the configuration of the multicast tree depends on the energy model under consideration.

3. THE NEW PROTOCOL, MDLP

Considering a source s which has to deliver multicast traffic to a set $M = \{m_1, m_2, ..., m_k\}$ of receivers, MDLP is a new distributed protocol, based on adaptive antennae, which builds a multicast tree connecting the source s with the set of receivers M. The goal of this protocol is to minimize the energy consumed by the multicast tree in a localized manner, after an initial distributed configuration. To do this, the proposed algorithm is based on the use of adaptive antennae and an underlying graph approximating the minimum spanning tree (MST). The use of adaptive antennae allows both one-to-one and one-to-many communication models. This is important because by reaching only those nodes belonging to the multicast tree, we spend less energy by using a narrower beam, and also we prevent nodes not taking part in the multicast tree from receiving not desired messages. Additionally, the use of a locally built underlying graph approximating the MST allows to deliver control messages in an energy-efficient way, and serves as starting point to build the multicast tree.

MDLP builds the multicast tree following three steps:

1. Building the underlying graph in a localized manner. The graph we have chosen is LMST [9] which requires

to know the location of one hop neighborhood or two hops if G_{0-} is used. To obtain the LMST graph, every node computes the minimum spanning tree of its one-hop neighborhood, then E_{lmst} is formed by the edges which belong to the MST of its both endpoints. This graph has a degree bounded by six with 2.04 [9] in average which is close to 1.98 [9], the degree of MST.

2. After building the graph, it is pruned to get an initial multicast tree. In this phase we obtain a multicast tree $IMT(V_{imt}, E_{imt})$ where $M_{set} \subseteq V_{imt}$, $V_{imt} \subseteq V$ and $E_{imt} \subseteq E_{lmst} \subseteq E$.

3. The last phase refines MT through a procedure performed by the set of relay nodes which are not multicast nodes $\{rl_i \in V_{mt} \mid rl_i \nsubseteq M_{set}\}$. In this procedure, we have called auto-excluding, every node computes whether the energy consumption can be decreased if it leaves the multicast tree. That is, their childs are assigned to its parent.

Finally, multicast data is delivered from source to receivers following the tree but, every forwarding node with more than two children has to evaluate if sending under a one-to-one communication model (that is a beam per child) or a one-to-many (only one beam covering all children)consumes a lower energy. Finally the node will aim the beam in the less expensive way.

The way the two last steps are carried out are explained in the following subsection.

3.1 Building the Multicast Tree

The building of MT is made using two processes: the first process which we call "Pruning" removes from the graph those branches which do not have multicast nodes. The second process, called "Aggregating" has the mission of establishing the initial multicast tree IMT and we call it "Aggregating" because it tries to make receivers share their paths towards the source in order to save energy.

The "Pruning" procedure carries out an initial refining of the graph whose purpose is to reduce the number of control messages delivered subsequently. It is initiated by those nodes $n_i \notin M_{set}$ which only have one neighbor. They send a message MDLP_PRUNED_MSG to their LMST neighbors so that they know the node is not going to take part on the multicast tree. This is a recursive process; if a node $n \mid n \notin M_{set}$ receives MDLP_PRUNED_MSG from all its LMST neighbors except from one, it will send the MDLP_PRUNED_MSG to the remaining LMST neighbors.

The second procedure"Aggregating", whose pseudocode is shown in Fig. 1(b), is initiated by the source to establish the multicast tree (that is, the tree which connects the source with the receivers) by creating low cost paths to a multicast node, either the source or a receiver, to which the message must arrive.

The source sends a message we call MDLP_QUERY_MSG to every LMST neighbor using a one-to-one communication model. The LMST neighbors propagate the message to its LMST neighbors (except the one which sent the message and the nodes pruned in the former process) and then MDLP_QUERY_MSG is propagated through the network.

The purpose of MDLP_QUERY_MSG is to establish low cost paths to a multicast node. Then, the source annotates in each packet sent the energy required to transmit it. When a node receives the MDLP_QUERY_MSG that it must propagate, it adds to the energy field of the received message the energy required send the message to the intended next hop. If the node has received more than one MDLP_QUERY_MSG it selects the one with the lowest value in the energy field, and annotate the node from which it received the best value as its parent towards the source. For example, node A receives MDLP_QUERY_MSG with a value of energy e and has to propagate it to its LMST neighbors B and C. Then A will send MDLP_QUERY_MSG to B with an energy $e+e_B$ and to C with an energy $e+e_C$, where e_B is the energy necessary to send the message from A to B and e_C is the energy necessary to send the message from A to C. These MDLP_QUERY_MSG messages serve as the basis to create a return path towards the source, which will be established by MDLP_JOIN_MSG messages as described below.

This scheme establishes low cost paths to the source, but the purpose is to minimize the consumed energy by establishing low cost multicast tree. To do that, it is benefitial to configure paths which are shared by multiple receivers. To achieve that, we use a simple mechanism proposed by Ruiz and Skarmeta in [5] to reduce data overhead. It simply works by making those nodes $n_i \in M_{set}$ propagate the MDLP_QUERY_MSG with the energy field set to 0. Thus, that particular path from that receiver to the source becomes attractive to other receivers. Continuing with the previous example, if A were a multicast node, the MDLP_QUERY_MSG sent to B, only will carry e_B because this is the only additional energy necessary to reach B after it has reach A.

As we said, the purpose of resetting the energy count when the message MDLP_QUERY_MSG arrives a multicast node is to establish the low cost path to a multicast node where the message has to arrive by making more interesting the path to a receiver. In that way, paths from source to receivers share part of their route resulting in lower energy consumption even if the paths are shorter separately. For example in Fig. 2 we have two receivers R_1 and R_2, and one source S. Sending multicast traffic from the source to the receivers using their minimum energy path consumes $1+2+3=6$ to reach R_1 and $2+3=5$ to reach R_2, totaling 11 energy units. On the other hand, if R_2 sends the message to R_1 (aggregating both paths) the total cost is $2+3+1=6$ energy units.

To answer the request of the source and establishing the first multicast tree, when MDLP_QUERY_MSG message reaches a multicast receiver $n_i \in M$ it starts a timer to address the possibility of receiving other MDLP_QUERY_MSG with less energy consumption. When the timer expires, the node sends MDLP_JOIN_MSG to the previous hop from which received the MDLP_QUERY_MSG with lower energy consumed. The purpose of this message is to join the multicast

Procedure Pruning

```
if (multicast_node) { return; }
else {
  if (PRUNED_MSG_received)
     non_pruned_neighbours--;
}
if (non_pruned_neighbours ==1){
  send(PRUNED_MSG)
  pruned=1;
}
```

(a)

Procedure Aggregating

```
If (node ϵ MU{s}) { e0=0; }
else { e0=getEnergy(PATH_MSG) }
for i in not_pruned_lmst_neighbour_set
{
     e=e0+compute_energy(i);
     send (PATH_MSG(e));
}
```

(b)

Figure 1: Pseudocode for building the multicast tree.

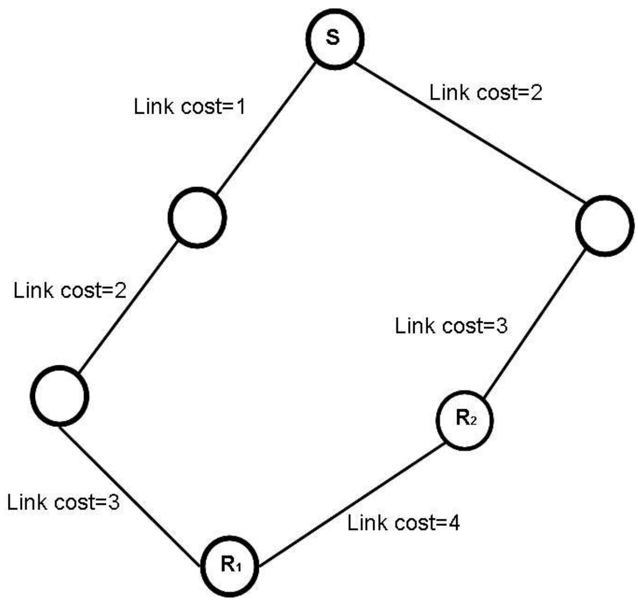

Figure 2: Example showing the aggregation of paths.

group and to create the multicast path. Then, when a node x receives a MDLP_JOIN_MSG from a node y, it forwards the MDLP_JOIN_MSG in the same way joining the multicast tree and it considers itself as a relay, which has to deliver multicast data to y. If x receives another MDLP_JOIN_MSG from z, it does not forward MDLP_JOIN_MSG to its parent but only keeps z as its child.

Then, nodes propagate MDLP_JOIN_MSG creating an initial tree which connects the source with the receivers.

3.2 Refining the Multicast Tree: Auto-excluding process

We have built a graph (LMST graph) which approximates the minimum spanning tree in a localized way. Over this graph, we have constructed a multicast tree which connects the source with the receivers by establishing low cost paths in a distributed way.

Now, we are going to refine this multicast tree by a localized procedure we have called "Auto-excluding". This procedure is performed by any relay $rl_i | rl_i \notin M_{set}$. Such a node checks if its father f can reach its children $c_1, .., c_k$. If this is the case, it computes the energy necessary to forward a multicast packet directly from its father to its children. That is, it compares $e_{frc} = e_{fr} + \sum_{c=1}^{k} e_{rc}$ (being e_{frc} the energy required to forward the packet from the father of the node to its children, e_{fr} the energy required to send the packet from its father to it and e_{rc} the energy required to send a packet from the node to a child) and the energy necessary to forward the message directly from its father to its children, that is $e_{fcc} = \sum_{c=1}^{k} e_{fc}$ (being e_{fcc} the energy required to forward directly the message from the father of the current node to the children of the current node and e_{fc} the energy required to send a packet from the father f to each of these children). Then, if $e_{fcc} < e_{frc}$ then the node communicates to the father and their children its intention to leave the tree. These steps are illustrated in Figure 3: node B computes e_{AB} (the energy necessary to send the packet from A to B), e_{BC} (from B to C) and e_{AD} (from A to D). Let us assume that $e_{AC} < e_{AB} + e_{BC}$, then the node notifies B and C its intention to leave the tree.

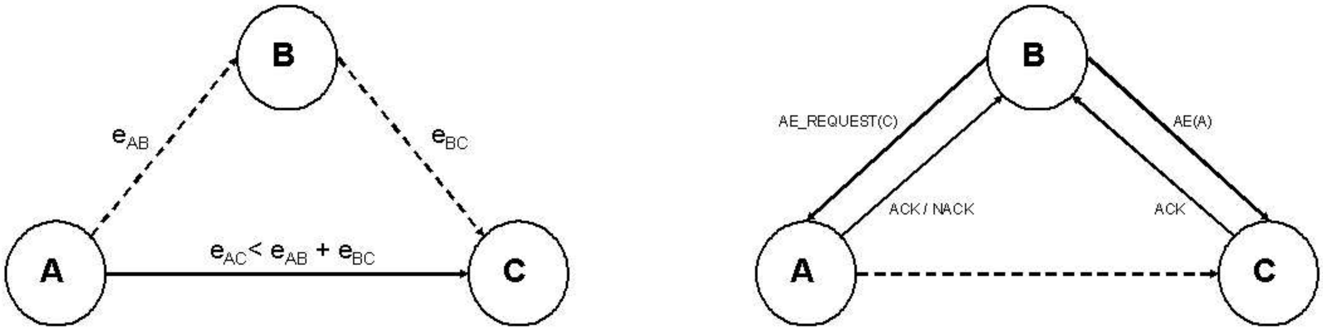

(*a*) (*b*)

Figure 3: Example which shows the autoExcluding of node B.

To avoid cycles if the father of a node is performing the Auto-excluding process, the node has to reset its own process and try later. Every relay (which is not a receiver) in the tree checks whether its presence is necessary. If it is not, it sends the EXCLUDING_REQ to its father, inserting in the message the identifier of its children. When the father receives the message, if it is not in the process it sends the EXCLUDING_ACK to its child and considers its grandchildren as new children. When the node receives EXCLUDING_ACK, it considers it has abandoned the tree, and communicates to its children the identifier of their new father (who has been its grandparent before).

The Auto-exclusion process is a localized process which directly depends on the additional energy in reception C_r and the values of the constants. If C_1,C_2, and C_r are null there are almost no auto-excluded nodes because it is not usually worth increasing the range. However, when C_1,C_2, or C_r are not null, the resulting tree ends up being almost completelly formed by multicast nodes if there is enough density of receivers in the network.

3.2.1 *Evaluating the Energy Consumption of Control Messages*

In previous subsections we have seen that the multicast tree is fully configured with low cost. Here we present an estimation of this cost considering the LMST formed has a degree of 2.04 [9] in average. To build the initial tree the sum of MDLP_PRUNED_MSG (messages sent by nodes pruned) and MDLP_QUERY_MSG (messages sent to announce low cost paths) results in one message per node because pruned nodes sent one message to its remaining LMST neighbor, and MDLP_QUERY_MSG is sent by every node to all LMST neighbors except the one from which the message was received and those nodes previously pruned. So, given that the average degree of LMST is 2.04, it results approximately in one message in average. Regarding the MDLP_JOIN_MSG, it depends on network configuration but in the worst case every node except the source sends one message. Then, assuming a network with n nodes and considering we are using a one to one communication model with a beam width β, the energy consumed in this process is $2n(\frac{\beta}{2\pi} * (r^{\alpha} + C_1) + C_2 + C_r)$ As we are using a one to one communication model with narrow beams β could be approximately $\pi/9$ resulting them $\frac{n}{9} * (r^{\alpha} + C_1) + 2n(C_2 + C_r)$

Table 1: Values of the constants

α	C_1	C_2	C_r
2	0	0	0
4	0	0	0
2	8000	2000	10000
4	8e7	2e7	1e8
2	8000	2000	0
4	8e7	2e7	0

Regarding the Auto-excluding process it depends on the tree previously built and the number of receivers, but we can estimate that in the worst case a node which wants to be auto-excluded has to send a message to all its neighbors.

4. EXPERIMENTAL RESULTS

In this section, we present the performance results of MDLP. To perform these experiments, the well known ns-2 [1] network simulator has been used. To perform the tests, we have enhanced the ns-2 network simulator with the code modeling directional antennae, our protocol and the state-of-the-art protocol DMIP [17] for comparison.

We have carried out exhaustive experiments to evaluate the proposed protocol over a wide variety of scenarios. Thus, we have performed the simulations varying the parameters of the energy model, the density of the network and the density of multicast receivers.

We have used six different energy models where the parameter values have been chosen according to [4] and are summarized in Table 1.

As seen in Table 1, in the two first cases $C_1 = C_2 = C_r = 0$ and only the propagation constant varies: $\alpha = 2$ and $\alpha = 4$. In the next two, C_1 and C_2 constants are different from null having the following parameters: $\alpha = 2,\ C_1 = 8000,\ C_2 =$

2000, *and* $C_r = 10000$ and $\alpha = 4$, $C_1 = 8e7$, $C_2 = 2e7$, *and* $C_r = 1e8$. These two last cases have also been simulated with $C_r = 0$ to address no additional cost in reception.

We varied the density of nodes from 50 $nodes/Km^2$ to 1000 $nodes/Km^2$, and for every tested density we generated ten different scenarios by randomly placing static nodes in a 1000 m x 1000 m area. In these scenarios we have varied the density of multicast nodes from 5 % to 50 % of the nodes of the ad hoc network. Every node in the network uses an IEEE 802.11 MAC layer and is equipped with a directional antenna which can aim a beam with a variable range, width and orientation. The antenna has a maximal range of 250 m and a minimal beam width of $\pi/9$ $radians$.

We firstly show how the tree is differently configured depending on the energy model, to do so we show the average number of messages sent to deliver a packet from the source to the receivers and the average transmission range used. Secondly we evaluate the performance of the protocol by comparing it to DMIP, to to this, we evaluate the energy consumed to deliver a data message through the multicast tree.

4.1 Configuration of the Multicast Tree vs. Energy Models

After applying the three phases of the protocol, the configuration of the final multicast tree is influenced by the energy model: when constants are null, there is no additional cost for aiming the beam and the energy consumed is only dependant upon the transmission range r^α. So, the energy consumption is reduced by transmitting with a low range. Therefore, the auto-Excluding process results in low number of excluded nodes and the multicast tree is almost a subgraph of the LMST graph.

When C_1, C_1 and C_r are not null, three new costs are introduced: there is C_2, a cost for aiming the beam regardless of its characteristics, $\frac{\phi * C_1}{2\pi}$ (being ϕ the beam width), a cost which depends on the width of the beam, and also $n * C_r$ (being n the number of nodes which receive the message), a cost associated to the reception which depends on the number of nodes listening to the beam. Then, in that case, it is interesting to reduce the number of beams (we use beams with minimal width) and to reduce the number of extra nodes which receive the message. To do that, our protocol uses two mechanisms: aggregating paths and the Auto-excluding process. This results in a multicast tree which does not follow the primarily LMST graph, in which the majority of the forwarding nodes belong to M and are connected by long edges representing beams of high range. Thus, the multicast tree obtained reduces both the number of messages sent and the number of nodes receiving the message.

Fig. 4 shows the influence of the energy models in the number of messages sent and their range. Fig. 4a shows the number of messages sent in a network with 500 nodes in a 1000 m x 1000 m area and Fig. 4b shows their average transmission range. When $C_1 = C_2 = C_r = 0$ the number of messages is high since the average transmission range approximates the average LMST distance among the nodes [4]. In addition to this, when $\alpha = 4$ the number of messages sent is greater than when $\alpha = 2$ because, due to the fact that the energy consumption $e \propto r^\alpha$, when α increases, the protocol tries to keep energy consumption low by reducing the range.

In the cases where constants are null, the number of auto-excluding nodes is small and therefore, the range remains small and almost constant approximating the average LMST distance [4], due to the fact that LMST is the underlying graph. This is translated in a bigger tree and thus an increment in the number of messages sent Fig. 4a.

When all the constants are not null, we can see in Fig. 4a that the number of messages forwarded are substantially lower. In fact, Fig. 4a shows a linear graph which is only a bit greater than the number of receivers because the protocol tries to not add to the tree other nodes in addition to the source and the receivers, to avoid unnecessary reception costs. In this case the transmission range decreases with the number of receivers. This is because for reducing the impact of C_1, C_2 and basically C_r, the number of nodes which belong to the multicast tree and do not belong to M_{set} is really low, as said before. Then, the tree basically contains the nodes in $M_{set} = M \cup \{s\}$, and when the cardinality of M, $|M|$, increases, the average distance between them decreases.

4.2 Energy Efficiency of the Protocol

We evaluate the efficiency of the protocol by comparing it to the state-of-the-art centralized protocol DMIP [17].

We have evaluated both protocols by varying the key parameters involved in their performance: the energy model, the network density and the density of multicast nodes. We evaluate the energy consumed by the multicast tree to send the multicast message.

Fig. 5 shows the energy consumption as a function of the network density for the four tested models and for 20% of receivers. When constants are null, the energy consumption is proportional to $number\_messages\_sent * r^\alpha$, then, the protocol tries to use short ranges, approximating the average range to average LMST distance between nodes, because the underlying graph is LMST. When density increases, the average range decreases at the same time as the number of messages gets higher .

Although DMIP is a centralized approach, MDLP consumes less energy in almost all energy models, except when C_1 and C_2 are not null, $C_r = 0$ and the network has a high density of both nodes and receivers. The reason is that in that particular case, it is highly benefitial to start from the initial BIP tree, because reaching a huge number of receeivers approximates a broadcast. Thus, D-BIP performs particularly well with those parameters.

When constants are null MDLP offers better results than DMIP due to the use of a one-to-one communication model, since energy consumption only depends on the width and range of the beam, so sending the packet to two nodes using a wide beam usually is more expensive than sending two narrow beams. In addition, the consideration of the C_r

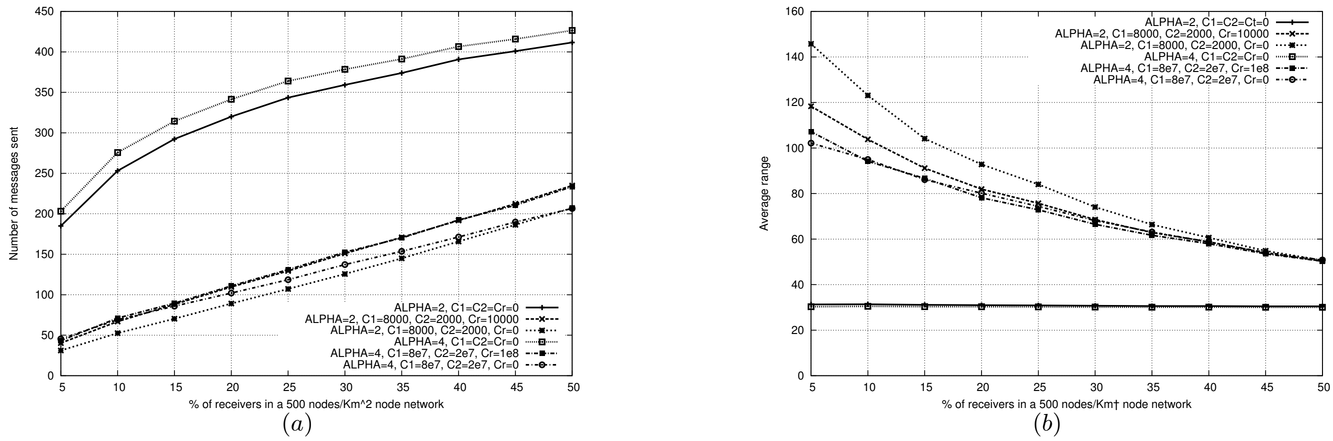

(a) (b)

Figure 4: Average number of received messages and average transmission range with several energy models in a network with 500 nodes in a 1000 m x 1000 m area

constant by MDLP to perform all the energy-related computations also allow it to outperform D-MIP.

When C_1, C_2 or C_r are not null, although transmission range is important, the energy consumption is mostly influenced by those constants. Here, the lower energy consumption of our protocol is basically due to the treatment given to C_r. That is, we try to minimize the number of nodes which receive messages. To do that, firstly, we aggregate paths favoring those paths which can be shared by many multicast receivers. Secondly, when carrying out the Auto-excluding process, the energy wasted by those nodes which receive the packet but do not belong to the multicast tree is also taken into account. However, as DMIP only adds one node at each step, it does not contemplate the effect of C_r, because it has always the same value. As a result, the number of received messages is almost twice than in our protocol.

Fig. 6 shows the energy consumption as a function of the percentage of receiver nodes when network density is kept fixed. When the number of receivers grows, the energy consumption gets higher as well. This increment is greater when all C_1, C_2 and C_r are not null. However, an important part of this expenditure is due to the energy used by multicast receivers to receive the message. Moreover, when the number of receivers exceeds 20%, the energy consumed by multicast receivers is greater than the energy consumed to deliver the message to them. This can be already seen in Fig. 6 which also shows the energy spent by multicast receivers in receiving the message (the curve "$|M| * C_r$"). This figure shows how the energy needed to deliver the multicast packet (the difference between the MDLP curve and the "$|M| * C_r$" curve) remains almost constant when the number of receivers increases. This is because the increment in the number of messages sent is compensated with the decrease of the beam range whereas the number of non-receiver nodes which listen to the message experiments little variation.

5. CONCLUSION AND FUTURE WORKS

One of the major limitations in ad hoc networks lies in the fact that the mobile nodes are usually battery powered. Thus, energy consumption is a key factor when designing protocols for this kind of networks. The usage of adaptive directional antennae helps in reducing energy consumption, although their deployment requires changes to be done on every layer of the network. However, few works exists on the usage of directional antennae with multicast traffic, whose correct treatment is crucial when deploying this kind of networks. In this paper we have presented MDLP, a multicast protocol using directional antennae, which manages to be energy efficient in most of the scenarios considered, with different energy models, using only distributed and localized procedures. Simulations results show that MDLP is more efficient than the state-of-the-art centralized protocol DMIP under a variety of scenarios, energy models and network configurations. As a future work, we are working towards the efficient support of scenarios involving mobility, as well as with the consideration of more realistic physical models, with probabilistic reception of messages.

6. REFERENCES

[1] Ns-2 network simulator.

[2] S. Athanassopoulos, I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Experimental comparison of algorithms for energy-efficient multicasting in ad hoc networks. In I. Nikolaidis, M. Barbeau, and E. Kranakis, editors, *ADHOC-NOW*, volume 3158 of *Lecture Notes in Computer Science*, pages 183–196. Springer, 2004.

[3] J. Cartigny, F. Ingelrest, D. Simplot-Ryl, and I. Stojmenovic. Localized lmst and rng based minimum-energy broadcast protocols in ad hoc networks. *Ad Hoc Networks*, 3(1):1–16, 2005.

[4] J. Cartigny, D. Simplot-Ryl, and I. Stojmenovic. An adaptive localized scheme for energy-efficient broadcasting in ad hoc networks with directional antennas. In I. G. Niemegeers and S. M. H. de Groot, editors, *9th IFIP Int. Conf. on Personal Wireless Communications PWC, September 21-23, 2004, Delft - The Netherlands; Lecture Notes in Computer Science*, volume 3260 of *Lecture Notes in Computer Science*, pages 399–413. Springer, 2004.

[5] P. M. R. A. F. Gomez-Skarmeta. Heuristic algorithms for minimum bandwidth consumption multicast trees

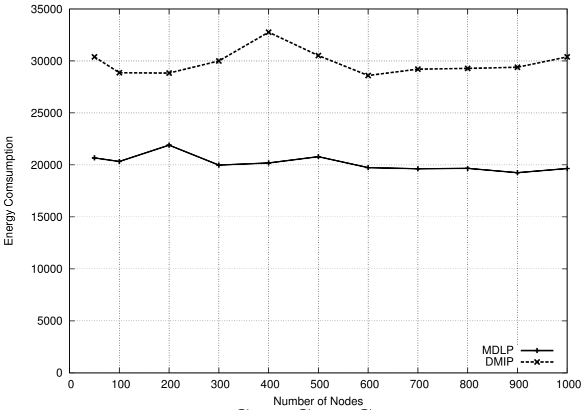

$\alpha = 2, C_r = C_1 = C_2 = 0$

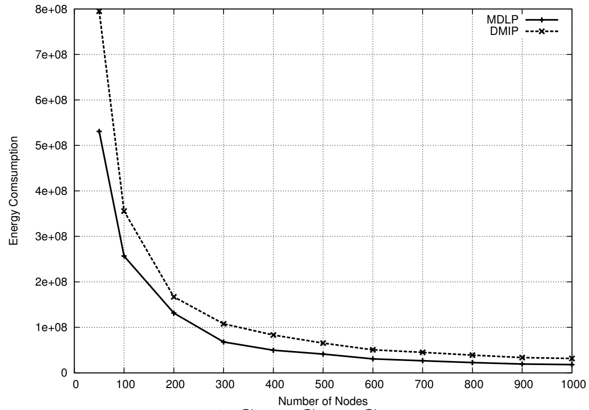

$\alpha = 4, C_r = C_1 = C_2 = 0$

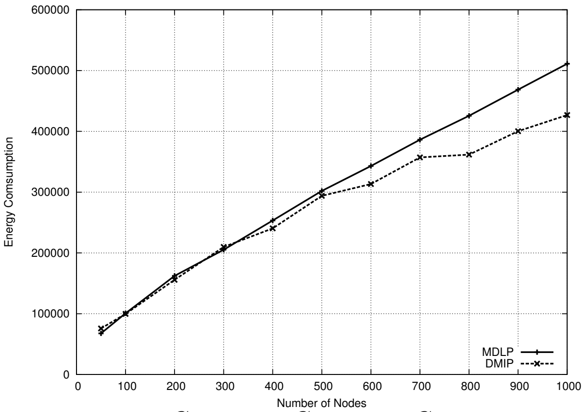

$\alpha = 2, C_1 = 8000, C_2 = 2000, C_r = 0$

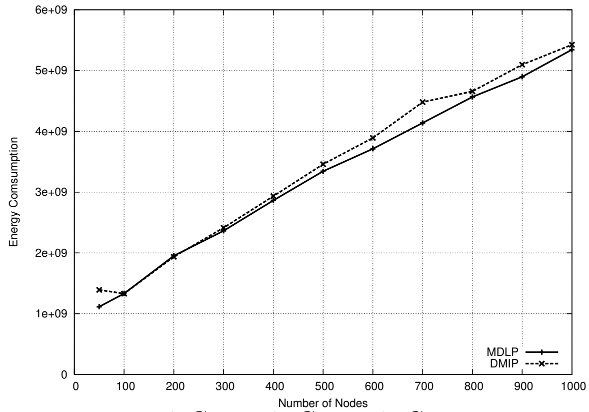

$\alpha = 4, C_1 = 8e7, C_2 = 2e7, C_r = 0$

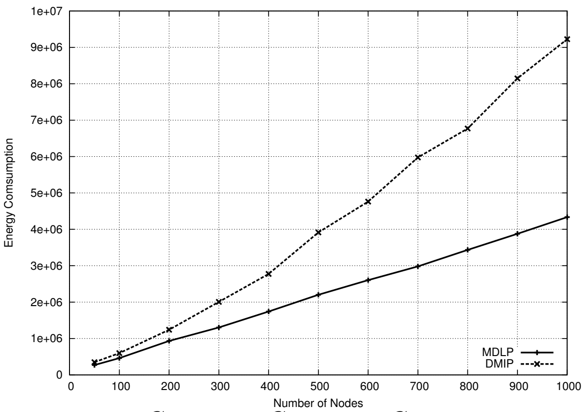

$\alpha = 2, C_1 = 8000, C_2 = 2000, C_r = 10000$

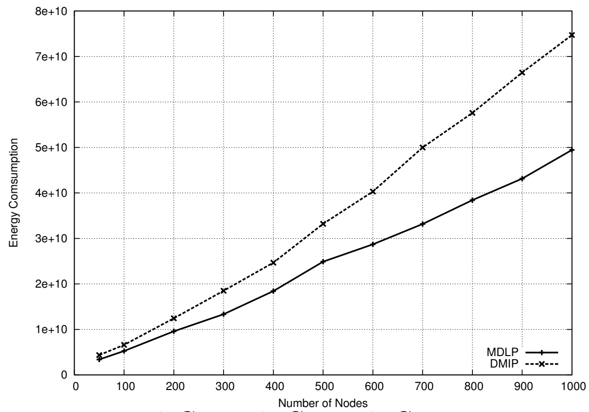

$\alpha = 4, C_1 = 8e7, C_2 = 2e7, C_r = 1e8$

Figure 5: Energy consumed as a function of network density with 20% of receivers.

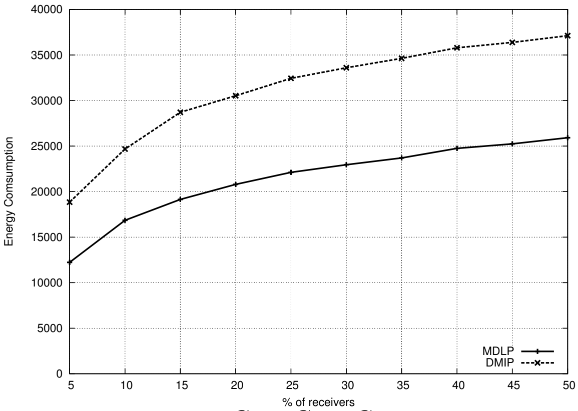

$\alpha = 2, C_r = C_1 = C_2 = 0$

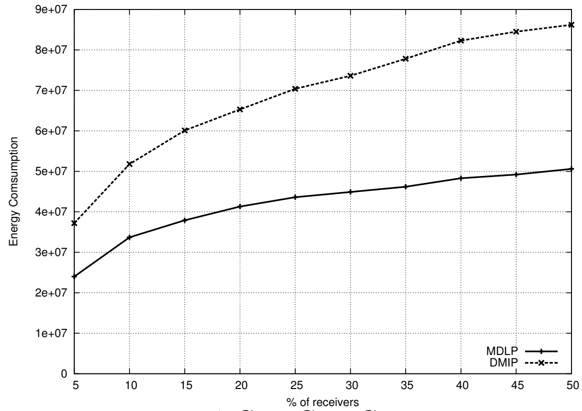

$\alpha = 4, C_r = C_1 = C_2 = 0$

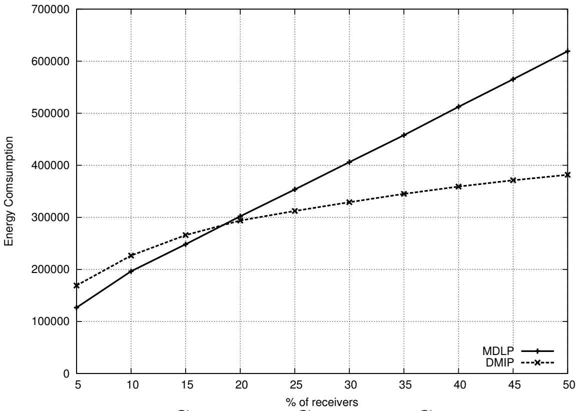

$\alpha = 2, C_1 = 8000, C_2 = 2000, C_r = 0$

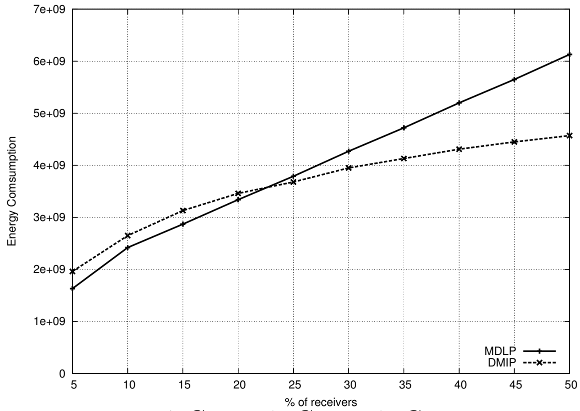

$\alpha = 4, C_1 = 8e7, C_2 = 2e7, C_r = 0$

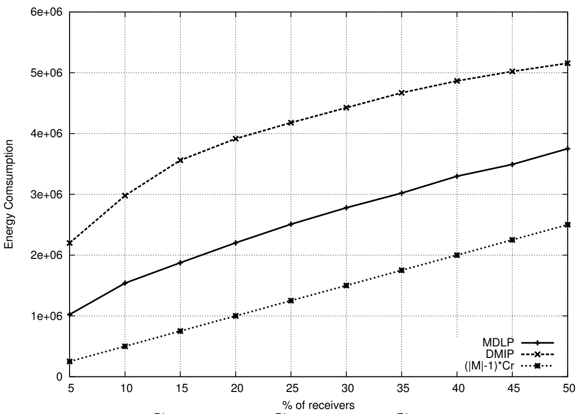

$\alpha = 2, C_1 = 8000, C_2 = 2000, C_r = 10000$

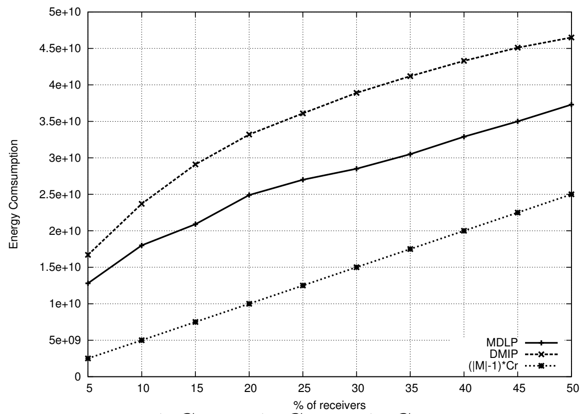

$\alpha = 4, C_1 = 8e7, C_2 = 2e7, C_r = 1e8$

Figure 6: Energy consumed as a function of the percentage of receivers nodes with a network density of $500\ nodes/Km^2$

in wireless mesh networks. *Lecture Notes in Computer Science*, 3738:258–270, 2005.

[6] S. Guo and O. W. W. Yang. Minimum energy multicast routing for wireless ad-hoc networks with adaptive antennas. In *ICNP*, pages 151–160. IEEE Computer Society, 2004.

[7] Y. T. Hou, Y. Shi, H. D. Sherali, and J. E. Wieselthier. Online lifetime-centric multicast routing for ad hoc networks with directional antennas. In *INFOCOM*, volume 1, pages 13–17, 2005.

[8] Y. T. Hou, Y. Shi, H. D. Sherali, and J. E. Wieselthier. Online lifetime-centric multicast routing for ad hoc networks with directional antennas. In *INFOCOM*, volume 1, pages 761–772, 2005.

[9] N. Li and J. C. Hou. Blmst: A scalable, power-efficient broadcast algorithm for wireless networks. In *QSHINE*, pages 44–51. IEEE Computer Society, 2004.

[10] R. Ramanathan, J. Redi, C. S. Member, D. Wiggins, , and S. Polit. Ad hoc networking with directional antennas: A complete system solution. *JSAC*, 23:496–506, MAR 2005.

[11] V. Rodoplu and T. H. Meng. Minimum energy mobile wireless networks. 17(8):1333–1344, Aug. 1999.

[12] A. Spyropoulos and C. S. Raghavendra. Energy efficient communications in ad hoc networks using directional antennas. In *INFOCOM*, 2002.

[13] R. Vaishampayan, J. Garcia-Luna-Aceves, and K. Obraczka. Efficient multicasting in multi-hop ad hoc networks using directional antennas. In *ICNP*, SEP 2005.

[14] B. Wang and S. Gupta. On maximizing lifetime of multicast trees in wireless ad hoc networks. In *International Conference On Parallel Processing (ICPP-03)*, pages 333–340, Kaohsiung, Taiwan, China, 2003.

[15] B. Wang and S. K. S. Gupta. Energy efficient broadcast with sectored antennas in wirelessad hoc networks. In *In Proc. of GLOBECOM - IEEE Global Telecommunications Conference*, volume 22, pages 3519–3524, Sardegna, Italy, DEC 2003.

[16] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. On the construction of energy-efficient broadcast and multicast trees in wireless networks. In *INFOCOM*, pages 585–594, 2000.

[17] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. Energy-aware wireless networking with directional antennas: The case of session-based broadcasting and multicasting. *IEEE Transactions on Mobile Computing*, (1-3):176–191, 2002.

[18] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. Energy-efficient broadcast and multicast trees in wireless networks. *Mob. Netw. Appl.*, 7(6):481–492, 2002.

[19] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides. Energy-limited wireless networking with directional antennas: The case of session-based multicasting. In *INFOCOM*, 2002.