

Power Adaptive Broadcasting with Local Information in Ad hoc networks

Xiaohu Chen, Michalis Faloutsos, Srikanth V. Krishnamurthy*
Dept of Computer Science and Engineering
University of California, Riverside
{*xchen, michalis, krish*}@cs.ucr.edu

Abstract

Network wide broadcasting is an energy intensive function. In this paper we propose a new method that performs transmission power adaptations based on information available locally, to reduce the overall energy consumed per broadcast. In most of the prior work on energy efficient broadcasting it is assumed that the originator of the broadcast has global network information (both topology information as well as the geographical distance between nodes). This can be prohibitive in terms of the consumed overhead. In our protocol, each node attempts to tune its transmit power based on local information (of up to two hops from the transmitting node). We perform extensive simulations to evaluate our protocol. Our simulations take into account the possible loss of packets due to collision effects and the additional re-broadcasts that are necessary due to lower power transmissions. We show that our protocol achieves almost the same coverage as other non power-adaptive broadcast schemes but with a reduction of approximately 40 % in terms of the consumed power as compared to a scheme that does not adapt its power.

1. INTRODUCTION

In this paper, we propose and evaluate a new protocol for performing energy-efficient *broadcasting* in ad hoc networks, assuming that the nodes in the network can tune their transmission power levels. Broadcasting is often a necessary function in ad hoc networks, but it is also an expensive process in terms of power consumption. Reducing the overall energy consumption is extremely important in increasing the longevity of ad hoc networks [8] [12] [7] [15] [16] [13]. Therefore, it becomes important to ensure that the protocols that are used at various layers of the stack are *power-aware* and lightweight in terms of the energy consumed. Network wide broadcasting is an operation that is often invoked for various applications. A broadcast may be required for the dissemination of control

messages (such as route queries in on-demand routing protocols [2]) or for the transmission of actual data to all the nodes in the network [18] [17] [9]. Note that this is different from a *local broadcast* wherein a node simply wishes to reach its immediate neighbors via a single transmission.

There have been several approaches that have been proposed to reduce the overhead incurred in broadcasting and for reducing the overall number of *rebroadcasts* of the initial broadcast message from the originator [18] [17] [9]. In all of these schemes an attempt is made to intelligently select an appropriate set of nodes to perform rebroadcasts. This in turn, reduces the overall number of transmissions and can therefore reduce the overall energy consumed in the broadcast. Note that in these protocols, the nodes always use the default maximum power level for transmissions.

If the nodes in the ad hoc network were able to adaptively *tune* their transmission power levels, one can envision that further savings in terms of the consumed energy during a broadcast may be achieved. Several protocols have been proposed to minimize the total energy consumption by allowing nodes to transmit using heterogeneous power levels during a broadcast in an ad hoc network. However, almost all of these protocols require the detailed topology map of network, i.e., require *global state* [8] [10]. Clearly, this approach works if the network size is very small but the overhead is prohibitive even for medium size networks (of the order of 100 nodes). Thus, it is desirable that one be able to tune the transmission power level of a node based on certain local criteria. One such approach is presented in [19]; the authors attempt to tune the transmission power levels so as to reach only a sub-set of neighbors. The authors show an overall reduction in energy consumption. However, they do not consider realistic medium access control and the effects of collisions. Thus, the effectiveness of their algorithm in a realistic setting wherein the energy efficiency and coverage are also effected by channel access effects are not known.

We propose a protocol that allows each node in the broadcasting process to perform *local optimizations* in order to decide its own transmission power while ensuring that it reaches all its neighbors. We show that this local optimization leads to a *global* reduction in the overall energy consumed by the broadcast. In our scheme, each node must be aware of only its two

* This work was supported from grants from Telcordia Technologies and ARL no: 100833196 and from DARPA FTN Grant no: F30602-01-2-0535.

hop neighborhood. A node then uses this information to choose an appropriate power level for performing its transmission. The key idea in our scheme is to have nodes reduce their power range such that they reach only a sub-set of their neighbors. These neighbors are then used as relays to broadcast packets to the more distant neighbors. We perform a local optimization to determine this sub-set of neighbors. We show that our protocol can significantly improve the performance of the broadcast in terms of energy savings as compared to a scheme that uses fixed transmission power levels for the broadcast. To quantify the improvements, our protocol reduces the energy consumption by as much as 40% while achieving the same coverage (in terms of the number of nodes reached) as that of a scheme in which every node uses the default maximum power level for transmissions. We also study the sensitivity of our protocol to various parameters such as the density of nodes in the network and the default maximum power level.

We organize our paper as follows. In Section 2, we describe the problem in more detail and we provide and justify some of the underlying assumptions that we make. More details on relevant prior work are reported in Section 3. In Section 4 we present a detailed description of our algorithm. The performance results from our simulations are presented and discussed in Section 5. We present our conclusions in Section 6.

2. BACKGROUND AND MODEL

Broadcasting typically involves the initiation of a broadcast by a source node and the objective of the broadcast is to send the message to all other nodes in the network. Typically, each node is assumed to transmit using the same default transmission power level. In a network in which nodes can tune their transmission power levels, in order to optimize (minimize) the total power consumption, appropriate nodes will have to be selected to perform the rebroadcasts and will have to be assigned the appropriate transmission power levels so as to determine their range. This problem is referred to in literature as the minimum energy broadcast tree problem. Kirousis *et al*, [12] examined this problem and showed that this problem is NP-hard for a network in three-dimensional space. A recent publication [10] shows that this minimum-energy broadcasting tree problem is NP-complete. A recent paper [11] provides a formal proof that shows that power optimal broadcasting is an NP-complete problem. Consequently, approximation algorithms for performing energy efficient broadcasting are needed for practical purposes.

A generalised definition of broadcasting: Typically, the definition of broadcasting implies total node coverage, but this may not be as interesting in practice. First, in large networks, network-wide broadcasting may be prohibitively expensive. Second, even if we want to, we might not be able to guarantee that all nodes will receive the information due to mobility or due to the presence of obstructions or network partitions. This sug-

gests that a more relaxed requirement of broadcasting is *being able to send data to a part of the network*.

We prefer this more general definition of broadcasting, which includes network-wide broadcasting as a special case. In practice, it might only be required to disseminate a broadcast message to only a certain sub-set of nodes in the network. An example would be a route discovery procedure that uses the expanding ring search process (ERS). ERS is used in the Ad Hoc On Demand Distance Vector (AODV) protocol; with ERS, a node progressively increases the range of its broadcast query; each broadcast query may not reach the entire network [3]. Another example of a broadcast limited in scope might be one that is invoked when a node wants to contact a server and it only needs to search whether the nearest one is available within a particular distance in terms of hops from itself. With this refined definition wherein *broadcasting* would refer to the dissemination of information within a particular region, we let the application or the user select the appropriate extent of the broadcast. The requirement is then, to disseminate the content up to the desired scope while minimizing the total power.

The maximum power range. We assume that each node has a default maximum power level P_{max} , and that this value is the same for all nodes. In the absence of any power management, a node will use this default power level to communicate with its neighbors. When we use power adaptive broadcasting, each node is able to select the power-range for a particular local broadcast. The selected power level could take on any value larger than zero but less than or equal in magnitude to P_{max} . In a power adaptive broadcast scenario, every node in the network may use a different power level during the course of the broadcast.

The channel model: The channel model used in this paper will reflect the power law model used often in literature [1],[3],[9]: $P_{rec} = P_{tx}/r^n$; here, P_{rec} is the received power, P_{tx} is the transmission power, r is the transmission range and n is a positive real value greater than 2 [1]. Typically, n takes on a value between 2 and 4 is often used to model the power attenuation with distance. Note that this simple model is sufficient for the purposes of this paper. In communications, there exist more detailed models that reflect antenna gains and shadowing effects [1],[4]. Our requirement is that a node simply needs to know only the power levels to reach its neighbors and *not* the actual distances. A node empirically determines the power range that it needs to reach a neighbor by observing the received signal strength indicator (RSSI) of messages from its neighbors. A message would typically carry information to indicate the power that the sender used for the transmission. Thus, the recipient node can simply estimate the attenuation experienced and since the wireless channel is reciprocal [4], it can estimate the power that it would need to use in order to reach the original transmitting node.

With our simplified model, for a node to receive a packet successfully, a minimum power P_{min} is required when the signal is received. Thus the maximum range r_{max} within which

a packet can be received can be determined using: $P_{min} = P_{max}/r_{max}^n$. As a result, the energy needed to reach a node at a distance r is proportional to r_{max}^n .

Assumptions: We make the following assumptions about the underlying network. The nodes share a single common channel and use the carrier sense multiple access protocol (CSMA with no collision detection) for the transmission of broadcast messages. These assumptions are generally made in most of the prior papers on broadcasting in ad hoc networks [18] [9] [17].

We assume that each node maintains information with regards to its neighborhood in a table which we call the *neighborhood table*. The neighborhood table is updated periodically and in order to do so, nodes exchange periodic HELLO messages with their neighbors. The table can also be updated whenever a node either receives a packet from one of its neighbors (not necessarily a HELLO packet) or overhears a neighbor's transmission. Other previous work on broadcasting in ad hoc networks assumes the exchange of such periodic messages [14] [9] [18].

Broadcast efficiency metrics: We define the following performance metrics in order to evaluate our protocol and to compare its performance with that of a non power adaptive broadcast scheme:

Total power consumption: This metric represents the total power consumed during a broadcast. The amount of energy consumed is directly reflected by the transmission power of a node. The lower the transmission power, the lower will be the energy consumed. We assume for simplicity that (since each broadcast packet is assumed to be of the same length in terms of bits¹ the energy consumed is representable by the transmission power used by the node. Thus, the total consumed power is represented by a simple sum of the transmission powers that are used by the nodes to perform rebroadcasts. For energy efficiency, it is desirable to reduce the total power consumed during the broadcast.

Duration of the broadcast: The time interval between the instant when the the first local broadcast packet from the originator up to the conclusion of the broadcast is referred to as the *duration of the broadcast*. One would want the duration of the broadcast to be as low as possible.

An interesting trade off appears here between the duration of a broadcast and reducing the possibility of collisions during the broadcast. We want to avoid the synchronization of rebroadcasts among neighbors to reduce the possibility of collisions and eliminating redundant rebroadcasts. Thus, we require that when a node hears a packet it will wait a random amount of time, before performing the rebroadcast. We shall

refer to this time as the *rebroadcast back-off* time. If the range in which a node can choose the random back-off time is large, we minimize the possibility that neighbor nodes will rebroadcast at the same time. The drawback is that this increases the duration of the broadcast. Fixing a small range for limiting the possible back-off times has the opposite effect: short duration, but a higher probability of collisions. In our work, we examine strategies to set this back off time in an intelligent way.

Contention level: Our goal is to compare our power adaptive broadcast scheme with the non power adaptive broadcast scheme in terms of the contention induced in the channel. As we discuss later, we suggest methods that can help suppress redundant broadcasts and use these methods with both our power adaptive broadcast scheme and the non power adaptive broadcast scheme. Our objective is to compare the traffic induced due to the two schemes (in terms of the number of broadcast messages). Since we use CSMA, a node prior to transmitting a packet, will first sense the channel. If channel is busy, the transmission is deferred in accordance to a backoff policy for a random time. However, there could be potential collisions due to hidden terminal problems [5]. It is important to ensure that appropriate coverage is achieved in spite of these collisions. We consider this in our performance evaluations.

3. RELATED WORK

Energy efficient broadcast with global state: Most of the previous work in power-efficient broadcasting assumes that the initiator of the broadcast has global *state* information. Ramathan and Rosales-Hain [6] consider the problem of adjusting the transmit powers of nodes in a multihop wireless network to maintain a strongly connected network. They present several algorithms to choose the appropriate power levels for the various nodes such that the maximum total power used is reduced. Once these power levels are found, the same levels may be used for broadcasting. In [8], Wieselthier etc. proposed three greedy heuristic algorithms to perform energy efficient multihop broadcasting. The Broadcast Incremental Power (BIP) algorithm is the key contribution described in that paper and is used to construct an energy efficient broadcast tree. Wan *et al* [7] present the first analytical studies for the minimum energy broadcasting problem based on the algorithms proposed in [8]. In that paper, the authors have presented a quantitative characterization of the performance of BIP. From a more theoretical perspective, Liang proposed a heuristic based algorithm for constructing an approximate minimum energy broadcast tree in [10].

As mentioned earlier, the above approaches require a detailed knowledge of entire network topology, i.e., require global state, including the geographical distances between nodes in the network. First, this requires that each node be equipped with a GPS device. Furthermore, this information will have to be somehow made known to all the nodes in the network. This is clearly not scalable since it requires nodes to exchange

¹ Note that the energy consumed is in Joules and is represented as a product of the transmission power (in Watts) and the time duration of transmission. Since the broadcast packet is assumed to be of fixed size, the time taken for a single transmission is the same for every packet. The amount of energy consumed is just the transmission power times a constant factor.

link state, especially in the presence of mobility. The energy consumption of these messages should be considered in the scheme.

Intelligent selection of rebroadcasting nodes: There are other methods wherein nodes simply use the default maximum power level to perform rebroadcasts; however, instead of all the nodes performing the rebroadcast, a sub-set of the nodes in the network are chosen in some intelligent way to perform the rebroadcast [18] [9] [17]. These nodes could be chosen either randomly, in accordance to their distance from the node from which they receive the broadcast packet or by exchanging information in a local neighborhood to determine their reachability. These methods can provide some energy savings since they reduce the number of rebroadcasts. The policies proposed by some of these schemes are orthogonal to implementing power adaptability and could be combined with our power adaptive method to increase energy savings².

Localized schemes: To the best of our knowledge, the only work on allowing a node to choose a transmission power level for performing a rebroadcast based on local information (information with regards to the node's two hop neighborhood) appears in [19]. The authors assume that a node constructs what is known as the restricted neighborhood graph (RNG). A node then attempts to choose its transmission power level only to reach nodes within its restricted neighborhood graph. This approach *relies on using the distances between different nodes*. If we do not assume GPS, the estimation of the distance will need to rely on calculating the distance from the perceived power, which would be difficult to compute if an accurate channel model is not known a priori. Furthermore, it may be disadvantageous to have geographically closer neighbors relay packets to their more distant counterparts. A node that is geographically closer may be severely shadowed and in such a case, it might even be a good option to have a distant node relay the packet to the closer one. In our approach, we only need to estimate the *power needed to reach a node* and not an estimate of the distance. Thus, as mentioned earlier, the proposed protocol can be used even when transmissions experience shadowing. Furthermore, the work in [19] assumes a reliable MAC layer whereas we study the performance of our scheme in a more realistic setting.

4. POWER ADAPTIVE BROADCASTING

In this section we describe our protocol in detail. We allow each node to choose its transmission power level based on information with regards to its neighborhood within two hops of itself. As mentioned earlier, the goal of our protocol is to achieve a reduction in the consumed energy during the broadcast. In short, with our protocol, each node examines if it is bet-

ter, in terms of the energy consumed, to have some of its neighbors relay the broadcast packets to the other neighbors. If it is better to have the packets relayed for a certain sub-set of its neighbors the node would reduce its transmission power level such that the nodes in the sub-set are outside its range. The nodes within the *new* range will act as relays for these *external neighbors*. Note that in the extreme case, a node might reduce its transmission power to reach only its nearest neighbor. Furthermore, note that in the presence of shadowing, a node that is geographically further away may potentially act as a relay for a node that is closer but is heavily shadowed. However, in the absence of shadowing effects (the only attenuation is due to path loss), the neighbors that are geographically close to the rebroadcasting node will act as relays to its more distant neighbors. For example, in Figure 1 the node s reduces its range to reach only nodes r and l . It excludes node k since it now relays on node r to forward the packet to node k .

Exchanging Two Hop Neighborhood Information: We require that nodes exchange information with their neighbors by means of periodic HELLO messages. The information contained in the HELLO message transmitted by a node includes a list of the *one-hop* neighbors of the node and the transmission power levels required by the node in order to reach them. As mentioned earlier, this power level may be estimated based on the received power level from each neighbor. Note that the HELLO messages are always transmitted using the default maximum power level P_{max} and the one-hop neighbors are those neighbors that are within the corresponding maximum transmission range.

With this exchange, each node is aware of its two-hop neighborhood. As an example in Figure 2, node 5's HELLO message contains a list of its neighbors, viz., nodes 6 and 7 and the transmission power levels that node 5 requires in order to reach these nodes. Upon the receipt of this message, node 1 is now aware of nodes 6 and 7 and the total power that it takes in order to reach these nodes via node 5. The information thus obtained, is tabulated by each node in a *neighborhood table*.

Locally Optimizing the transmission power levels: Let us assume, without loss of generality, that an arbitrary tagged node (say node s) has n neighbors, $\{1, 2, 3 \dots n\}$. Let P_j indicate the transmission power that the node s has to use in order to reach neighbor j , where $1 \leq j \leq n$. We define the *furthest* node from the tagged node to be node k ($1 \leq k \leq n$) such that $P_k = \max_j P_j$. Note that if our simplified channel model is used, this node k is indeed the geographically furthest node from the node s . Clearly, node s does not need to use a transmission power larger than P_k in order to reach all the nodes within its range.

Node s , next, examines if node k is reachable by any of its other neighbors i such that $i \neq k$. If it is possible to do so for a sub-set of the nodes within this set, it computes the power required in order to reach node k via each of the nodes in the sub-set. Let P_{ik} represent the power required by node i in order to reach node k . Clearly, if P_{sik} is the power required by

² We note here that we already do this to some extent since, as we describe later, we adjust the rebroadcast back-off time to favor certain nodes intelligently.

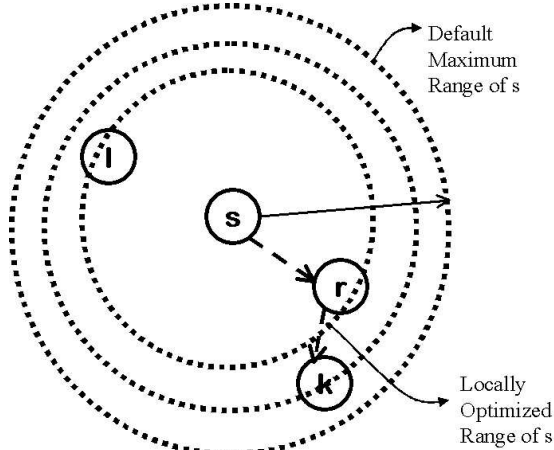


Fig. 1 Power Adaptive Broadcasting




Fig. 2 The two hop neighborhood of a node.

node s to reach node k via node i , then,

$$P_{sik} = P_i + P_{ik}. \quad (1)$$

If node s is thus able to find a neighbor r such that it takes lower power in order to send a packet to node k via r than it would take to reach node k directly, it would exclude node k from its set of neighbors. It assumes that node r 's broadcast will reach node k (Figure 1). If it is unable to find any such node r then, it would choose to transmit directly to node k . In this case, the transmission power is chosen to be P_k . In other words, if $r \in (1, n)$ is a node such that $P_r + P_{rk}$ is minimum among all $i \in (1, n)$ and $i \neq k$, node s does the following:

$$\begin{cases} \text{if } P_k < P_r + P_{rk}, & \text{choose the single hop transmission;} \\ \text{if } P_k > P_r + P_{rk}, & \text{choose the two-hop path.} \end{cases}$$

Let us assume that node s does indeed find a relay r for node k , and now, let node l be the *new furthest node* from node s (Figure 1). Now, node s can reduce its transmission power level to P_l . Now, it attempts to repeat the optimization process by finding neighbors that can relay packets to node l with a lower power budget. If this is not feasible, then node s chooses P_l for its broadcast. If, on the other hand, this is feasible, node s then, chooses the appropriate relay for node l as before and further reduces its power level so as to reach only its *new furthest neighbor* that is determined after the exclusion of node l from its neighborhood set.

When a node receives a broadcast packet, it will perform a similar local computation to determine its power level. It is possible that an intermediate rebroadcasting node will consider a neighbor that has already received the broadcast packet (via some other node) while performing its optimization. This may lead to a choice of sub-optimal power level, since, clearly it is unnecessary for the node to ensure that the particular neighbor receives the broadcast packet. While this problem cannot be

completely eliminated without global state, we consider strategies later, that can help alleviate this effect.

Note that a node performs a greedy search for the locally *optimal* energy-efficient power level. In order to perform the search, a node would use the information stored in its neighborhood table described earlier. We re-iterate that we do not attempt to construct a globally optimal minimum energy broadcast tree which requires the exchange of global state. We require the node to possess information only with regards to its two-hop neighborhood. We require that a node learn the two-hop neighborhood instead of simply the one-hop neighborhood since it is essential for the node to know how its neighbors are interconnected as opposed to simply its own connectivity with its neighbors. As an example, in Figure 1, it is essential for node s to know that node r and node k are neighbors and the transmission power that node r requires in order to reach node k .

Our power adaptive scheme will consume at most the energy consumed in a non power adaptive scheme in which the nodes simply use the default maximum power levels if we assume that each node possesses perfect information with regards to its two-hop neighborhood. We state this as a theorem:

Theorem 1: Power adaptive broadcasting consumes at most the energy consumed when nodes transmit with the default maximum power.

Proof: By construction. As described earlier, a node does not choose a power level higher than the default maximum power level. In the worst case, all the nodes will transmit at the default power level. Thus, the overall power consumption, in the worst case, is equal to that of the non power adaptive scheme.

In the ideal case, wherein the MAC layer is assumed to be perfectly reliable we can state the following theorem:

Theorem 2: The coverage achieved by the power adaptive broadcasting scheme is the same as that of the scheme wherein nodes transmit with the default maximum power.

Proof: This is again by construction. Each node ensures that




Fig. 3 Eliminating redundant transmissions

for every neighbor that is excluded from its range there is an appropriate relay node that is identified. Thus, the so called *external neighbors* receive the broadcast packet when the appropriate relays perform the rebroadcast. Thus, with the power adaptive scheme every node within a rebroadcasting node's maximum range receives the broadcast packet. Hence, we achieve the same coverage achieved as that of the scheme wherein nodes use the default maximum power.

Eliminating redundant transmissions: As mentioned earlier, a node has information with regards to its two hop neighborhood. Thus, based on the transmission power level chosen by the predecessor node from which it receives its broadcast, it can determine the nodes within its neighborhood that also will have potentially received the same broadcast from its predecessor. Furthermore, it chooses a random delay (sets a timer) before it performs its own rebroadcast. If in the interim, it overhears one of its neighbors performing a rebroadcast, it further eliminates neighbors that are reached by means of that rebroadcast (the two-hop neighborhood information is again used for making this determination). If, when its random timer expires, the node does not have any neighbors that have not received the broadcast message either from the node's predecessor or from any of its other neighbors, it simply aborts its rebroadcast.

One could further eliminate certain redundant retransmissions by using the two-hop neighborhood information. In order to explain this we consider an example shown in Figure 3. In this figure, node 2 and node 3 are not within the transmission range of each other. Node 1 and 4 are their common neighbors. Node 1 is the originator of a broadcast message and directs this message to nodes 2 and 3. Both node 2 and node 3 will perform a local optimization and when doing so, they both take their common neighbor, node 4, into account. However, it could be potentially possible either (a) to completely eliminate either node 2 or node 3's transmission or (b) to reduce the transmission range of one of these nodes, if only one of these nodes were made to consider node 4 while performing their local optimizations.

This is achievable by using the two-hop neighborhood infor-

mation that is already possessed by nodes 2 and 3. With this, both node 2 and node 3 know that nodes 1 and 4 are their common neighbors. Thus, each of the nodes knows that the other is capable of broadcasting the message such that node 4 receives the message. Furthermore, we note that each of these nodes knows the power level that it requires, as well as the power level that the other node requires, in order to reach node 4. Thus, each node determines if it requires a *lower power level* in order to reach node 4 than the other. If it is the case, then it takes node 4 into account while performing its optimization. If not, it determines that the other node will do so and thus excludes node 4 from its neighborhood. Thus, for example, if node 4 is *closer* in terms of the power consumed to node 3, node 3 includes node 4 in its neighborhood during its local optimization while node 2 will exclude it.

However, due to possible loops, a node may be considered in multiple neighborhood sets and thus, may receive multiple copies of the broadcast. However, in order to eliminate such *overlaps* it is necessary to facilitate the exchange of more *state*.

5. SIMULATIONS AND DISCUSSION

We developed a simulation tool to simulate our algorithm in the ad hoc network environment. The channel model and the medium access control are implemented as described in Section 2.

We make sure that nodes are made aware of their 2-hop neighborhood. One way to do this is to have each node broadcast a HELLO message periodically to notify its neighbors of its presence and its current neighborhood information. If the nodes move fairly rapidly, then in order to obtain and maintain up to date information, HELLO messages are needed quite frequently. If on the contrary, the nodes are relatively static and the topology changes infrequently, then the HELLO messages could be less frequent. In order for ensuring that the local optimizations are done correctly, it is important that the information be relatively fresh. Typical routing protocols require these HELLO messages and the frequencies are chosen appropriately by the protocol. For the simulations presented here, we use a static topology with a periodicity of 1 second for the HELLO messages. In some preliminary experiments with mobility, we used the period of 1 second which gave satisfactory results. However, we note that this is a system parameter that has to be appropriately set depending upon the mobility. Note that the duration of a broadcast session for moderately sized ad hoc networks (as we consider here) is relatively independent of the speed of motion. This is because the duration of a broadcast session (for example in a network containing 500 nodes) is typically much smaller than the time it takes for a topology change to occur even when there is high mobility (nodes move at 20 m/s) [9].

In our simulation models, the network is deployed in a 1000m x 1000m area. Nodes are uniformly placed in this area of interest. We record the number of nodes reached during a broad-

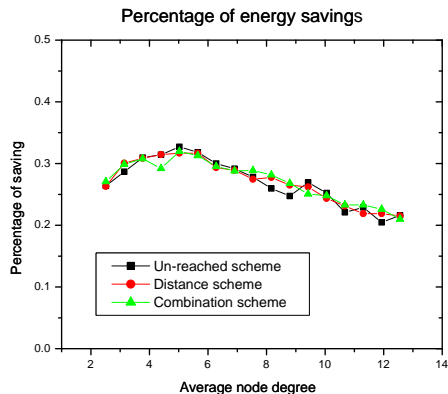


Fig. 4 Comparison of percentage of energy savings for different methods for computing the RBT.

cast session to reflect the *coverage* achieved by the broadcast and the corresponding number of rebroadcast packets and their transmission power levels to reflect the consumed energy for each configuration. We perform 200 simulation iterations in all and compute average results. We also use node density and the maximum default power as parameters and perform simulations to examine the effects of these parameters. We compare our results with a scheme that uses fixed power-range. The primary objective for doing such comparisons is that we want to examine if the gains that we achieve due to power adaptivity justify the use of our protocol. An additional advantage is that it is an established framework for comparison. In the non power adaptive scheme used for comparisons, we retain the features that enable the node to suppress redundant transmissions as described in the previous section. The fixed power scheme also requires nodes to possess two-hop neighborhood information.

We vary the maximum power range from 100m to 250m. For small values of the power range, the broadcast is unable to reach all the nodes in the network due to the presence of partitions. However, we find that our protocol and the fixed broadcast have the same coverage in terms of the number of nodes reached as we might expect (Theorem 2).

5.1. Choosing the rebroadcast back-off time

As mentioned earlier, a node upon the receipt of a packet would perform a rebroadcast; however, it does so after a random period of time. Upon the expiry of a timer after this randomly chosen rebroadcast back-off, the network layer sends the broadcast packet to the MAC layer for transmission. We examine three ways to choose the random distribution to specify the rebroadcast back-off. Our expectation is that this would in some way affect the efficiency of the broadcast. We define $MaxDel$ to be the maximum permissible back-off delay.

In our first scheme, we give priority (small back-off) to nodes

that have a larger number of neighbors that have not received the broadcast at the current time. Recall that if there are no unreached neighbors there is no need to rebroadcast at all. Furthermore, the maximum rebroadcast back-off time is limited to $MaxDel$ as mentioned above (We choose $MaxDel$ to be between 10 ms and 100 ms in our simulations. However note that this is a system parameter and can be set as appropriate.). Thus, for a node that has *unreached* number of neighbors that have not received the broadcast, the rebroadcast back-off time (RBT) is computed to be:

$$RBT = \frac{(MaxReach - unreached)}{MaxReach \times rand() \times MaxDel} \quad (2)$$

where $MaxReach$ is a parameter that we discuss below, and $rand()$ is a random number between 0 and 1. We can set the value of $MaxReach$ to obtain a behavior of our choice. A possible choice is to set this value to be the total number of neighbors of the node. As a consequence, if the number of unreached neighbors is high, one would expect to see a shorter back-off time. Another option is to set $MaxReach$ to a large constant that would be an estimate of the maximum number of neighbors that a node could have. An advantage of this approach, is that it provides a more uniform reference point for selecting a back-off time. In our simulations, we tried the second option and set the value to 15; note that this is typically much higher than the degree of a node in our simulation framework. Note that in practice, RBT should always have a positive value. We can easily guarantee this by taking the maximum of the calculated RBT and a preset positive value.

On the other hand, one might want the nodes that are at the geographical periphery of a previous broadcast to be the ones to rebroadcast first. Intuitively, if we denote as the **center**, the location of the predecessor node whose rebroadcasts a set of nodes hear, we want those nodes that are the furthest from the center to perform the rebroadcast. This can work well only if the received power level seen by a node is a good estimate of the distance of the node from the predecessor broadcasting node. If $MaxRange$ represents the maximum transmission range (200m in our simulations), and the distance of a given node from the originating node is estimated to be $distance$, the RBT is calculated to be:

$$RBT = \frac{(MaxRange - distance)}{MaxRange \times rand() \times MaxDel} \quad (3)$$

We consider a combination of the two schemes to harness the advantages of both. If we give both the approaches equal weights, the same weights (we assume a unit weight) the expression for computing the RBT now becomes:

$$RBT = \left\{ \frac{(MaxRange - distance)}{MaxRange} \right\} \quad (4)$$




Fig. 5 Energy consumption and percentage savings (power attenuation $n=2$)

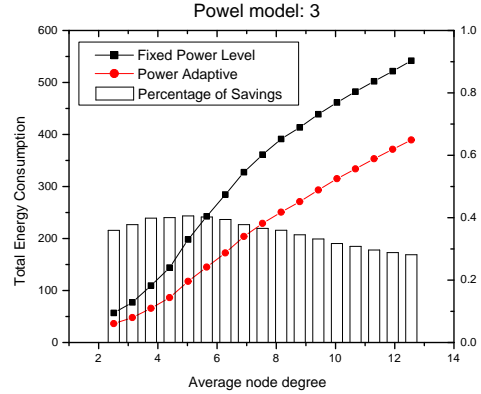


Fig. 6 Energy consumption and percentage savings ($n=3$)

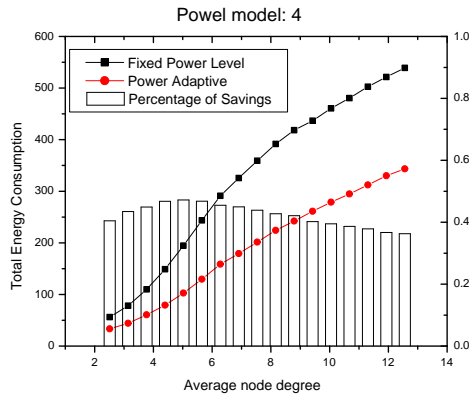


Fig. 7. Energy consumption and percentage savings ($n=4$)




Fig. 8. Energy consumption and percentage savings ($n=4$)

$$+ \left. \frac{(MaxReach - unreached)}{MaxReach} \right\} \\ \times rand() \times MaxDel.$$

We perform experiments to first determine the extent up to which the optimization of the back-off time is beneficial and compare the three methods. First, we find that with any of the three optimization schemes, we can reduce the total energy consumed by about 20-35% for a wide range of densities as shown in Figure 4. Second, we find the performance of the three schemes to be similar in terms of the percentage of energy savings. The main cause for this behavior is that we preclude a node from performing transmissions if it has no neighbors that are yet to receive the broadcast. Both the schemes described operate in a greedy way, i.e., they both attempt to provide a higher priority to nodes that most likely cover a high number of new

nodes. Let us consider a high degree node with a large number of neighbors. If the node is near the center of the previous rebroadcast, it probably has few *unreached* neighbors. Thus, it may be discouraged from performing the rebroadcast. Similarly, a node with a lower degree at the periphery of the previous broadcast may be encouraged to broadcast since it has a relatively high number of *unreached* neighbors. Thus, the distance and degree are not totally independent of each other. It turns out that, for a given scenario, most of the nodes that perform the rebroadcast are common irrespective of the scheme used for choosing the RBT.

5.2. Evaluating the benefits of power adaptivity

Varying the node density: We vary the number of nodes in

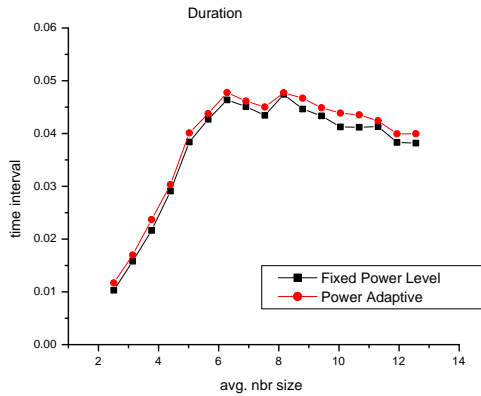


Fig. 9. Duration of the broadcast versus the average neighborhood size ($n=2$).

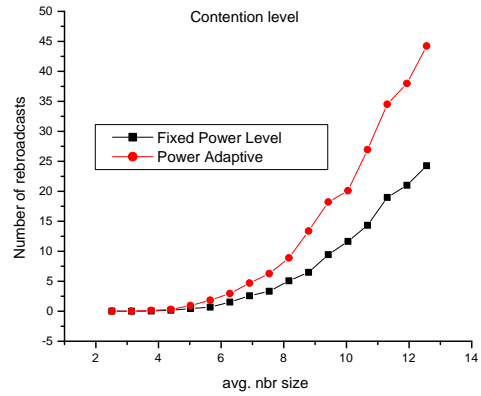


Fig. 10. Number of rebroadcasts versus average neighborhood size ($n=2$).




Fig. 11. Percentage of collisions experienced versus average neighborhood size ($n=2$, $MaxDel = 10$ ms).

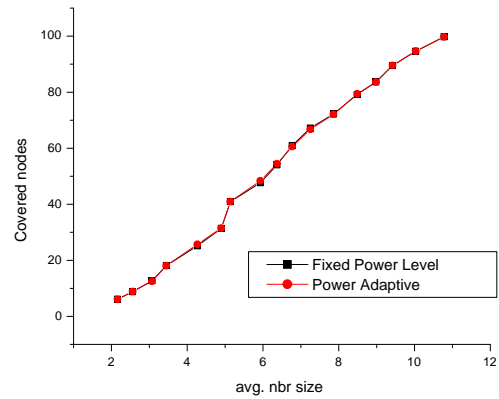


Fig. 12. Coverage achieved versus average neighborhood size ($n=2$, $MaxDel = 10$ ms).

the simulation area between 50 and 100. This translates to a variation in the average degree from 2.5 to 12.6. We also vary the path loss exponent in the channel model: the attenuation varies with distance r as $1/r^n$, and we vary the value of n from 2 to 4 [1]. In this series of simulations, we set the maximum default power range to 200m. In Figures 5, 6 and 7, we notice that our power adaptive scheme outperforms the non power adaptive scheme irrespective of the path loss exponent. Notice that as we increase the path loss exponent, the relative performance of the power adaptive scheme as compared with the non power adaptive scheme (fixed maximum power level) improves. This is because, it is more efficient in terms of energy to route packets over shorter hops as opposed to making a single-hop long transmission as the path loss exponent increases. Furthermore,

notice that at smaller densities, it is probably infeasible to perform local power optimizations since a node may not have other neighborhood relays through which packets can be forwarded to distant neighbors. Thus, at low densities there is little advantage to using the power adaptive scheme and this advantage grows as the density increases. However, beyond a certain density we see that the percentage of savings actually dips a bit. This is because even if the number of neighbors of a rebroadcasting node increases the advantage in terms of finding a relay to reach distant nodes does not increase beyond a certain limit. But now, due to the increased number of nodes attempting to perform rebroadcasts, it is likely that the RBT variations would creep in leading to sub-optimal transmissions. This in turn causes a slight decrease in the percentage of energy sav-

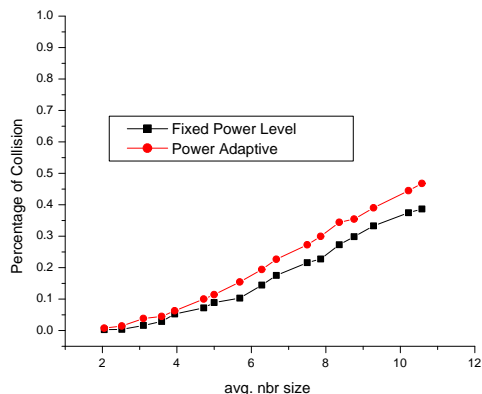


Fig. 13 Percentage of collisions experienced versus average neighborhood size ($n=2$, $MaxDel = 100$ ms).




Fig. 14 Coverage achieved versus average neighborhood size ($n=2$, $MaxDel = 100$ ms).

ings.

Note that in all the figures below, the *bar plots* refer to the percentage of savings and are quantified by the ordinate on the right. The ordinate on the left refers to the actual energy consumption that is depicted by the linear plots. Note that these results were similar for various values of $MaxDel$ since the subset of nodes that performed the rebroadcast did not change in most cases.

Varying the Default Power Range. In Figure 8, we plot the total energy consumed by the power adaptive scheme and that consumed by the non adaptive scheme versus the default maximum power range P_{max} . Clearly for small values of P_{max} there is not much local optimization that is possible and the two schemes almost perform identically. However, the savings are evident for larger values of P_{max} . Occasionally we see only marginal savings. However, notice that the power adaptive scheme never consumes more energy than the non power adaptive scheme. Note that these results were valid for all values of $MaxDel$ chosen.

Duration: The duration of a broadcast session is of importance. The number of rebroadcasts in the power adaptive scheme could be potentially higher since the range is reduced. Note that the fixed power scheme (also referred to as the non power adaptive scheme) is also able to eliminate redundant transmissions. In fact since a node transmits with P_{max} , a single transmission could potentially eliminate many of the potential redundant transmissions. Furthermore, due to the increased range, the non power adaptive scheme could potentially take a smaller time in completing the rebroadcast. We see from Figure 9 that this increase in time due to using the power adaptive scheme is almost negligible. In this reported experiment $MaxDel$ was chosen to be 10 milliseconds. However, the result were similar in behavior for other values of $MaxDel$ as well.

Contention and Coverage: Since we are using a small power

range for each local broadcast to reduce the total energy consumption, the could potentially increase the number of rebroadcasts in the network. The additional number of rebroadcasts incurred are shown in Figure 10. However, note that many of these rebroadcasts are with smaller transmission powers. Thus, even though there are a larger number of rebroadcasts we reduce the overall energy efficiency while ensuring the same coverage as that of the non power adaptive scheme. Furthermore, since the transmission range is considerably reduced, the interference is reduced. At any give time, a smaller part of the network is “busy”, and this can lead to a higher spatial reuse of the channel. In Figures 11 and 13, we plot the fraction of packets that collided i.e., the ratio of the receptions that experienced collisions to the total number of receptions versus the average node degree (represents density). Note that a node may have multiple receptions due to transmissions from more than one of its neighbors and hence, experiencing a collision does not mean that a node did not receive the broadcast packet. If it did receive the broadcast packet from any of its neighbors it is said to be *covered*. The performance reported is in fact dependent on the value of $MaxDel$ chosen earlier. When a small value of $MaxDel$ is chosen (10 milliseconds) we note that there are fewer collisions as compared to when a larger value of $MaxDel$ is used (100 milliseconds). An increased value of $MaxDel$ spreads out the transmission instants and thereby reduces the number of collisions to some extent. We note that, regardless of the value chosen for $MaxDel$, the percentage of collisions experienced due to power adaptive broadcasting is fairly close to that in the case of the non power adaptive broadcast scheme. In fact, considering the fact that there are a significantly larger number of broadcasts with power adaptive broadcasting at high densities this increase is very small. Due to the reduction in power levels there is a tendency for the number of collisions to decrease; however, due to the increase in the num-

ber of transmissions there is a tendency for this number to increase. As a result of these conflicting effects we see that the fraction of colliding packets increases to a small extent. However, this might be acceptable especially since the power savings achieved can be significant (from 25 % to 40 %). Furthermore, note that the two schemes result in the same coverage (in terms of the nodes reached) irrespective of the value chosen for $MaxDel$ (Figures 12 and 14). The plots for the power adaptive broadcasting and the non power adaptive broadcasting almost lie on top of each other. This is in accordance to our expectations (Theorem 2).

6. CONCLUSIONS

In this paper, we propose a new scheme in which nodes perform local optimizations so as to reduce the overall energy consumption during a broadcast session. Unlike most previous algorithms that require global state to perform power adaptive broadcast, our algorithm achieves significant energy savings by exchanging local two-hop neighborhood information. The key idea is to have neighbors who can be reached with low transmission power, relay the broadcast packets to those that require higher transmission powers. If there are no fading effects and the signal is attenuated only due to path loss, this corresponds to relaying broadcast packets to the more distant neighbors via the closer neighbors. Thus, the overall range is reduced to include only those neighbors for whom it is better to use a direct single hop link as opposed to relaying via other neighbors. We also propose modifications that can suppress redundant transmissions and provide some co-ordination among the rebroadcasting nodes. We perform extensive simulations to compare our power adaptive scheme with a non power adaptive scheme and we show a drastic improvement in terms of the energy consumption (up to 40 % savings). We examine the sensitivity of the scheme to various parameters such as the node density, the pathloss exponent and the maximum default transmission power level and discuss the results. Our protocol results in a larger number of low power rebroadcasts and a slight increase in the fraction of colliding packets. However, this is not an issue since the coverage is maintained and in light of the significant power savings. In summary we conclude that our scheme is a viable option for broadcasting in ad hoc networks and can result in almost the same coverage while consuming much lower energy.

References

- [1] T.S. Rappaport, *Wireless Communications, Principles and Practices*, Prentice Hall, 1996.
- [2] E. Royer and C-K. Toh, "A Review of Current Routing Protocols for Ad hoc Wireless Networks", *IEEE Personal Communications Magazine*, April 1999.
- [3] C. E. Perkins and E. M. Royer, "Ad hoc On-Demand Distance Vector Routing," *Ad hoc Networking*, Addison-Wesley Publishers, Ed. C. Perkins, 2000.
- [4] J.P. Linnartz, *Narrow Band Land-Mobile Radio Networks*, Artech House, 1993.
- [5] V.Bharghavan, A.Demers, S.Shenker and L.Zhang, "MACAW - A media access protocol for Wireless LANs", in *Proc. of ACM SIGCOMM*, 1994.
- [6] R. Ramanathan and R. Rosales-Hain. "Topology Control of Multihop Wireless Networks Using Transmit Power Adjustment", *Proc. IEEE INFOCOM 2000*, Tel Aviv, Israel, March 2000, pp. 404-413.
- [7] P.-J. Wan, G. Calinescu, X.-Y. Li, and O. Frieder. "Minimum-energy broadcast routing in static ad hoc wireless networks", *Proc. IEEE INFOCOM*, 2001.
- [8] J.E. Wieselthier, G. D Nguyen and A. Ephremides. "On construction of energy-efficient broadcast and multicast trees in wireless networks", *Proc. of IEEE INFOCOM*, 2000.
- [9] X. Chen, M. Faloutsos, S.V. Krishnamurthy, "Distance Adaptive (DAD) Broadcasting for Ad Hoc Networks", *Proc. of IEEE MILCOM 2002*.
- [10] W. Liang, "Constructing Minimum-Energy Broadcast Trees in Wireless Ad Hoc Networks", *Proc. of ACM MOBIHOC*, 2002.
- [11] M. Cagalj, J. P. Hubaux and C. Enz, "Minimum-Energy Broadcast in All-Wireless Networks: NP-Completeness and Distribution Issues", *Proceedings of ACM MobiCom 2002*, Atlanta, September 2002.
- [12] L.M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc, "Power consumption in packet radio networks", *Proc. 14th Symp. on Theoretical Computer Science (STACS)*, 1997.
- [13] S. Singh, M. Woo and C. S. Raghavendra "Power-Aware Routing in Mobile Ad Hoc Networks", *Proceedings of ACM MOBI-COM*, 1998
- [14] H. Lim and C. Kim, "Multicast Tree Construction and Flooding in Wireless Ad Hoc Networks", *Proceedings of ACM MSWiM*, 2000.
- [15] A.Muqattash and M.Krunz, "Power Controlled Dual Channel (PCDC) Medium Access Protocol for Wireless Ad Hoc Networks", *Proc. of IEEE INFOCOM 2003*.
- [16] E.Jung and N.H.Vaidya, "A Power Control MAC Protocol for Ad Hoc Networks", *Proc. of ACM MOBICOM 2002*.
- [17] B.Williams and T.Camp, "Comparison of Broadcasting Techniques for Mobile Ad Hoc Networks", *Proc. of ACM MOBIHOC 2002*.
- [18] S.Y.Ni, Y.C.Tseng, Y.S.Chen and J.P.Sheu, "The Broadcast Storm Problem in a Mobile Ad Hoc Network", *Proc. of ACM MOBICOM*, 1999.
- [19] J.Cartigny, D.Simplot and I.Stojmenovic, "Localized Minimum-Energy Broadcasting in Ad-hoc Networks", *Proc. of IEEE INFOCOM*, 2003.