Local Versus Global Power Adaptive Broadcasting in Ad Hoc Networks

Kyriakos Karenos, Asheq Khan, Srikanth Krishnamurthy, Michalis Faloutsos and Xiaohu Chen
Department of Computer Science and Engineering
University of California, Riverside
{kkarenos, kasheq, krish, michalis, xchen}@cs.ucr.edu

Abstract— **In many ad hoc deployments, making broadcasting, a critical function component, power efficient is extremely important. Current power-adaptive approaches proposed for broadcasting, can be grouped into a) centralized or omniscient schemes, and b) decentralized or localized schemes. The latter assume only local neighborhood information. Due to the absence of global information, localized algorithms may not produce optimal or near-optimal solutions. On the other hand, global knowledge is not typically available to nodes and may be extremely expensive to disseminate. In this paper, we examine the importance of the lack of global information on the performance of localized approaches. Towards this, we perform extensive simulations and compare the performance of localized power adaptive broadcasting with the performance of the well-known Broadcast Incremental Power or BIP, an omniscient algorithm. We analyze the behaviors of the two protocols and identify the reasons for the differences in behavior. We observe that while global state does provide better performance in terms of energy efficiency, the localized scheme outperforms the global scheme in terms of the latency incurred in the broadcast. Based on the observed behavioral traits of the two protocols, we suggest changes by which the energy consumption with the localized scheme is reduced by as much as 20%, while incurring almost no penalty in terms of latency.**

I. Introduction

Broadcasting in ad hoc networks is required for the dissemination of control information (like route queries in on-demand routing protocols [4]) or for the distribution of data to all the nodes in the network [2], [5], [12]. Typical ad hoc network deployments that may require *broadcast* of information include electronic classrooms, tactical networks and disaster recovery missions. Broadcasting is expensive in terms of the power consumed. Minimizing the total energy consumed during a broadcast is critical especially in the aforementioned deployment scenarios and could extend the lifetime of the network [1], [3], [6], [8], [10], [11].

Prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

Several methods have been proposed to make broadcasting energy efficient. They can be broadly classified into global or local schemes. Global approaches require nodes in the network to be aware of the entire network topology in terms of the graph that depicts the connectivity between nodes and the power budgets required for each link in the graph. As an example, in [6], the Broadcasting Incremental Power (**BIP**) algorithm utilizes global knowledge to achieve energy efficient broadcasting. On the other hand, local schemes attempt to perform local power optimizations. These schemes require that each node possesses partial topology information that is constrained to within its local neighborhood. As an example, in [13], the *Power Adaptive Broadcasting with LOcal information* (**PABLO**) algorithm uses only the node's local two-hop neighborhood information. With an omniscient algorithm, near-optimal energy-efficient broadcasting is potentially possible. In contrast, localized algorithms that rely on localized optimizations and *greedy* decisions. However, the centralized approaches do not scale in terms of communication overhead, and they are not deployable in highly mobile or large networks. Therefore, the localized approaches are more relevant in practice; however the centralized approaches can provide performance benchmarks that the local protocols can strive to reach.

The goal of this paper is to quantify the trade-offs between the use of a global versus a local approach: we compare a centralized approach (BIP) and a local approach (PABLO). While BIP has been popularly studied as an omniscient adaptive broadcasting method, PABLO is one of the very few currently proposed approaches and is representative of the general class. With our studies, we asses the various behavioral aspects of both the local and global schemes and suggest methods to further improve the performance of the local scheme. To the best of our knowledge, this is the first comparative study that attempts to provide an understanding of power adaptive broadcasting in ad hoc networks. First, we simulate both algorithms with similar constraints and assumptions to allow for a fair comparison. Our studies suggest that BIP, as expected, consumes significantly lower energy for a broadcast compared to that consumed with the localized power adaptive algorithm. However we observe that these savings of BIP are achieved at the expense of a much higher broadcast termination time. Second, we interpret the results and analyze the reason for the observed behavior. The analysis leads us to modify PABLO to improve its energy efficiency. The modification is based on the utilization of *neighborhood pruning* techniques to reduce redundancy in rebroadcasting. The key observation that leads to the proposed modifications is that a node should not be responsible for reaching a neighbor if the neighbor has a high chance of being reached by a rebroadcast from other nodes. This reduces the extent of *redundant considerations* of nodes (as we discuss later) during local power optimizations. The proposed mechanism can reduce the energy consumption of the initial protocol by up to 20%, while maintaining low latency and similar node coverage as with the previous algorithm.

We have chosen to use PABLO for our comparisons since, to the best of our knowledge, it is the only power adaptive scheme proposed thus far. Note that an independent work [5] follows an approach similar PABLO: it develops a mechanism to tune the transmission power levels using local information. However, PABLO provides a more extensive and detailed deployment framework [13] than the aforementioned work.

The paper is organized as follows. In section II, an overview of prior work is provided. Specifically, we review BIP and PABLO. We present the simulation model in section III and the performance of the two schemes is compared in section IV. In section V, we introduce our pruning extension for the localized power adaptive algorithm and evaluate its performance. Our conclusions are presented in section VI.

II. Background

Broadcasting in ad hoc networks could be energy intensive if nodes transmit with a default, maximum power level. In order to make broadcasting energy-efficient, first, nodes must be able to intelligently tune their power levels. Furthermore, unlike in flooding where at each local broadcast instance, a node's neighbors rebroadcast a received packet, sets of nodes must be intelligently chosen at each such instance to perform the rebroadcasts. In other words, redundant broadcasts are required to be quelled by precluding the right set of nodes from performing rebroadcasts. This problem is referred to as the *minimum energy broadcast tree problem*. This has proven to be an *NP-complete* problem [9]. Approximate algorithms have been used to perform energy efficient broadcasting. Two such algorithms where a node can tune its transmission power level are the Broadcast Incremental Power (BIP) algorithm and the Power Adaptive Broadcasting (PABLO) algorithm.

BIP creates a broadcast structure by utilizing global information that reflects the state of the network. Each node is assumed to be aware of the distance between each pair of nodes in the network and the transmission power needed for the pair to communicate. There are two stages in BIP. In the first stage, a minimum power tree is built with a broadcast initiating source; the source is the root of the tree. The initiating source finds the neighbor node that it can transmit to with the minimum power and includes that node in the minimum power tree. Then, each node in the tree calculates the additional power it requires to reach a node that is not in the tree. The node that can be reached with the least additional power is then chosen for inclusion in the tree. The process is repeated until all the nodes are included in the minimum power tree. In the second stage, the initiating node begins the broadcast. The broadcast packets are then propagated down the constructed tree.

The performance of BIP can be improved with what is referred to as the *sweep* operation. The sweep operation takes advantage of omni-directional broadcast transmissions of nodes to prune the minimum power tree by eliminating unnecessary transmissions. Details of the sweep mechanism can be found in [6] and are omitted due to space constraints.

In **PABLO**, nodes do not have global knowledge about the network. Instead, each node has state information with regard to its two-hop neighborhood. In other words, each node knows the transmission power necessary in order to communicate with its neighboring nodes and the transmission powers needed by these neighbors in order to reach their other neighbors. The nodes exchange periodic HELLO messages to gather the aforementioned information with regard to their neighbors. The algorithm allows each node to use its default maximum power or tune down its power level according to an optimization function.

During the optimization stage, a node determines if it is power-efficient to have one or more neighbors relay the broadcast packet to its furthest neighbor. If the node finds such a relay, it removes the furthest neighbor from its neighborhood set. The process is repeated with the new furthest neighbor until no further optimization is possible. Every node in the network performs a similar optimization. By eliminating the furthest neighbor at each optimization step, a node reduces its transmission power level at every step. As a result, every node transmits with a locally optimal power level that is typically lower than that of its default maximum power level. Thus, the total energy needed for a broadcast is reduced.

III. Simulation

A. Model and Assumptions

In order to allow for a fair comparison between BIP and PABLO, we use the simulation framework described in this subsection. We assume the use of omnidirectional antennae with statically fixed node locations (no mobility). It is worth mentioning that in previous work in [13], it is argued that since the duration of a broadcast is typically much smaller than the time within which the topology of the network may be expected to change this is a reasonable choice. However, the frequency of HELLO messages in the power adaptive broadcasting and the overhead cost of maintaining the tree in BIP are affected by mobility. In order to ensure fairness, we ignore these costs. We assume that each node has a default maximum transmission power level, P_{max} and also requires a minimum power threshold P_{min} to receive a packet.

The following assumptions are made about the underlying network to provide a realistic simulation setting: the nodes share a single common channel and use the carrier sense multiple access (CSMA) protocol with no provision for collision detection (as in most prior literature) for the transmission of broadcast messages. With CSMA, we set a random rebroadcast back-off time (**RBT**) [13] which is computed as follows:

$$\mathrm{RBT} = \left(\frac{\mathrm{MaxReach} - \mathrm{unreached}}{\mathrm{MaxReach}}\right) \times \mathrm{rand} \times \mathrm{MaxDel};$$

MaxDel is the maximum permissible back-off delay, *rand* is a random number between 0 and 1, and the number of neighbors of a node that have not been reached as yet by the broadcast is denoted by *unreached.* In order to ensure a uniform RBT, as in [13], *MaxReach* is set to a large constant that approximates an upper limit on the maximum number of neighbors a node can have in accordance to the chosen network size.

B. Metrics and Setup

We compare the performance of BIP and PABLO in terms of the following performance metrics:

Total power consumption: This represents the total power needed to complete a broadcast. For simplicity, we assume that the transmission power of a node reflects the energy spent by the node as in [13]. Therefore, the total energy needed for a broadcast is the sum of the transmission powers used by the nodes to perform rebroadcasts. Hence, the broadcast is more energy-efficient when nodes can rebroadcast with lower transmission powers.

Duration of the broadcast: This metric represents the time interval between the initiation of the broadcast from the originating node and the termination of the broadcast. We consider the actual time spent for rebroadcast as the latency and ignore the time spent to build the minimum power tree in BIP. With PABLO, as described in an earlier section, there is no need for the construction of a global structure (or tree).

Total number of rebroadcasts: We measure the total number of rebroadcasts invoked by each algorithm to complete the broadcast. This is the sum of the number of rebroadcasts performed by the intermediate nodes and the single local broadcast performed by the initiating node.

Node coverage:In addition to performing an energy-efficient broadcast, it is also important that the broadcast fulfills its goal of providing coverage, i.e., reaching all the nodes in the network. This metric, then, represents the number of nodes in the network that receive the broadcast.

We have compared BIP and PABLO based on the performance metrics defined above by varying the *network size* or *node density*. The node density is defined as the total number of nodes in the simulation area. The network is deployed in a 1000*m* $\times$ 1000*m* area and the node density is varied by placing 30 to 90 nodes in the area. Upon receiving a broadcast, prior to performing its own rebroadcast, each node waits for a random time that is computed according to the random back-off time (RBT) formula (in section III-A). The maximum possible RBT before a rebroadcast, is set to 1000*msec*. The maximum power range of each node is fixed at 250*m*. The power attenuation coefficient due to distance is 4 i.e., the power attenuation at a distance d from the transmitter is proportional to d^4. For each chosen value of node density, we run 100 simulations and compute the average values. We also compare the performance results of BIP and PABLO with that of a scheme in which nodes deploy identical, fixed power levels. The scheme is a variant of SBA scheme in [2] and has also been considered in [13].

IV. Results

Total energy consumed: Figure 1 shows the total energy consumed by each algorithm as the node density is varied. As expected with global knowledge, BIP outperforms PABLO. We also notice that with BIP the total energy consumed decreases as the node density increases. Since BIP is based on the minimum spanning tree algorithm, (a minimum power tree is built), with an increase in the number of nodes in the network, the distance between nodes decreases; thus, the power budget on each link is reduced. PABLO outperforms the fixed power level based scheme since it is more efficient to route packets over smaller hops than use a long single-hop transmission. We also observe that the “sweep” operation does not improve BIP’s performance significantly. Due to reduced transmission range with increased node density, fewer nodes are likely to be covered by a single broadcast; thus, sweeping is less likely to eliminate any of the rebroadcasts. Clearly, the presence of global knowledge helps BIP. The power consumed on each link is significantly lower as compared to PABLO. The local properties of PABLO cause many nodes to be covered by more than one re-broadcast and hence, the power consumed is much higher than in BIP.

Duration of the broadcast: The latency (i.e. the time interval between the initiation of the broadcast from the originating node and the termination of the broadcast) incurred by the two schemes is depicted in Fig. 2. It is seen that BIP requires a much longer time to complete a broadcast compared to the localized power adaptive algorithm. This is because, in BIP, nodes receive the broadcast packet almost in sequential order over a long path of rebroadcasts (an example of the broadcast sequence is shown in Fig. 3(a)) whereas, with PABLO, a node induces *broadcast flashes* to reach multiple nodes simultaneously (Fig. 3(b)). In these figures we have chosen a topology and depict an instance of the broadcast as seen with BIP and PABLO respectively. Clearly the trade-off of using a localized scheme is then a reduction in the total time taken at the expense of consuming higher energy (as compared to the global scheme). This is due to the fact that a node may be considered multiple (redundant) times by different neighbors when these neighbors perform local optimizations for rebroadcasting. All but one of these neighbors might potentially use higher powers than needed for performing rebroadcasts.

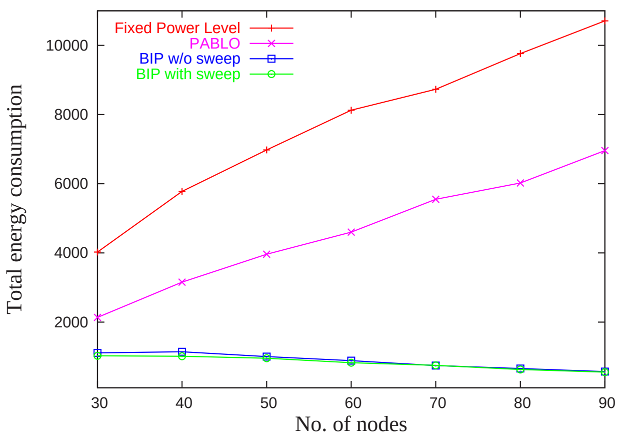

Fig. 1. Total Energy Consumed Vs Network Size

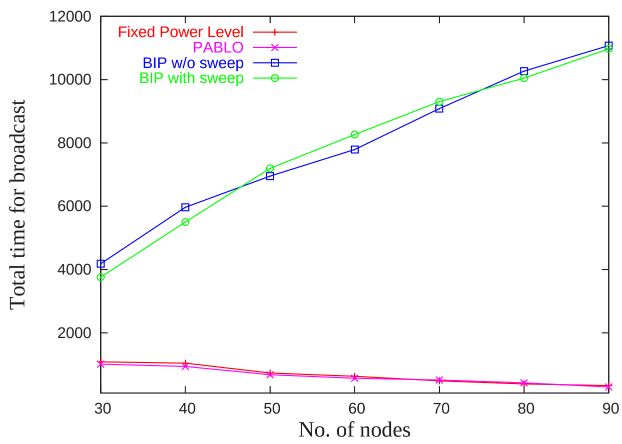

Fig. 2. Duration of Broadcast Vs Network Size

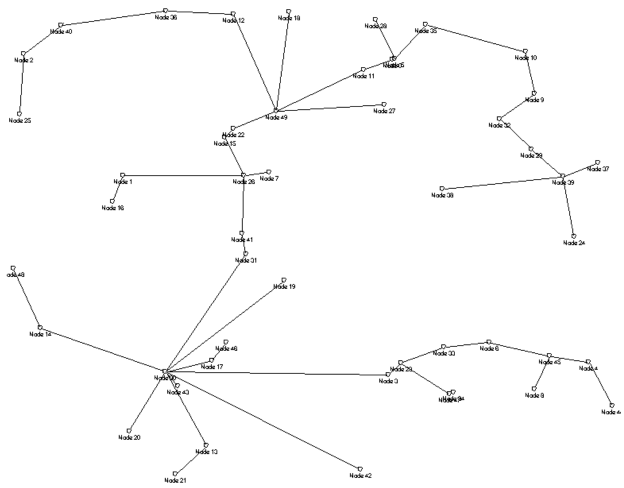

(a)BIP tree (no sweep)

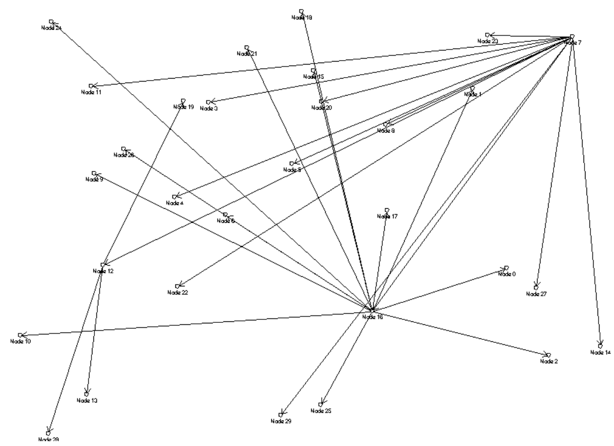

(b)Power Adaptive Broadcasting (PABLO)

Fig. 3. Visualization of Broadcasting Instances

Total number of rebroadcasts: The total number of rebroadcasts induced by each algorithm to complete a broadcast are shown in Fig. 4. As previously explained, with *broadcast flashes* (Fig. 3(b)), the localized power adaptive algorithm can reach multiple nodes in parallel and therefore, requires fewer rebroadcasts than BIP. On the other hand, BIP builds a path to reach, potentially, just a single additional node at each rebroadcast step (Fig. 3 (a)). Thus, although it does incur a higher number of rebroadcasts, the power consumed for each rebroadcast in BIP is consider-

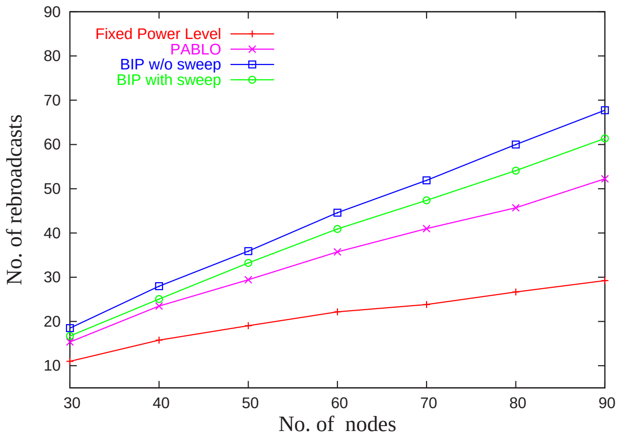

Fig. 4. Total Number of Rebroadcasts Vs Network Size

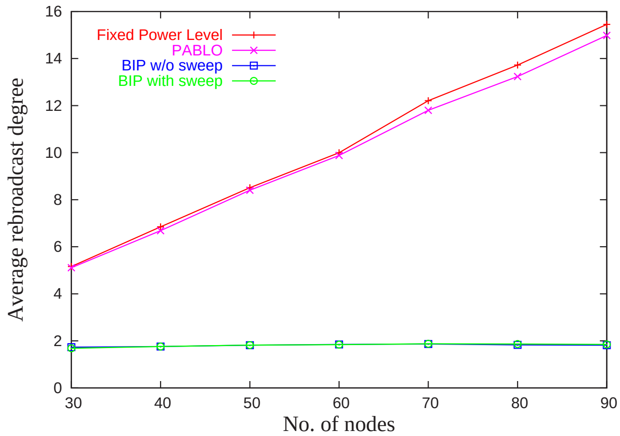

Fig. 6. Average Rebroadcast Degree Vs Network Size

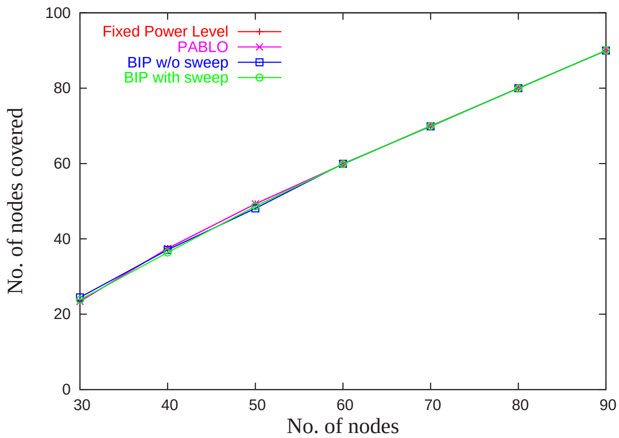

Fig. 5. Node Coverage

ably lower than PABLO.

Node coverage: In addition to minimizing the energy required for a broadcast, it is equally important to ensure that each node receives the packet. We have verified that both algorithms achieve an almost identical coverage as shown in Fig. 5.

Average Rebroadcast Degree: We also look at an interesting parameter: the average rebroadcast degree observed at participating nodes. We calculate average rebroadcast degree as the average number of neighbors that a transmitting node can reach (even if the transmission is not intended to reach some of these neighbors). One of the most important facts is that, as seen in Fig. 6, BIP maintains an almost stable average rebroadcast near the size of 2. Since a greater coverage was observed per rebroadcast in PABLO, this increased rebroadcast degree suggests that there exist redundant transmissions.

V. Improving Performance

A. Conceptual Description

Based on above observations, we explore the possibility of improving PABLO. Our primary goal is to reduce the power consumed by PABLO while ensuring that (i) reducing the power level will not result in reduction in node coverage and (ii) no additional *topological* information will be required. In order to achieve this objective, we propose a design modification based on the observation that the high average rebroadcast degree with PABLO suggests redundant coverage by causing nodes to use higher powers than needed. It is expected that reducing this rebroadcast degree will implicitly shift the behavior of PABLO towards the behavior of a globalized scheme.

The modification that we propose to PABLO is that, when a node computes its locally optimum power level to perform a rebroadcast, it could potentially *speculate* whether its neighbors might have been (or could be) considered in the local optimizations of another node. This other node under discussion is possibly not a common neighbor (else an agreement could be reached as to who should consider the common neighbor while performing optimizations as discussed in [13]). However, note that energy savings can be achieved only when excluding a neighbor that will allow for the transmission power to be reduced. This is possible if the *furthest* neighbor (after optimization) of the rebroadcasting node can be excluded from the transmission range.

A node needs to decide, after its preliminary optimizations, whether it should further reduce the power level to limit the transmission so as to exclude one or more of its furthermost neighbors. This is achieved by utilizing *neighborhood pruning* techniques. Pruning has been used previously in attacking the broadcasting problem mainly to reduce redundancy and wasteful transmissions [12], [7].

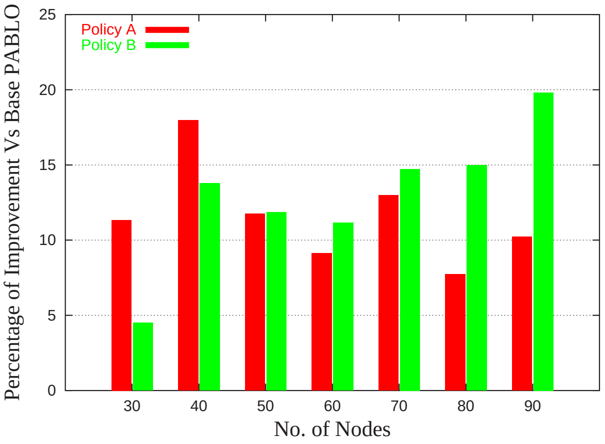

Fig. 7. Percentage of Improvement Vs Network Size

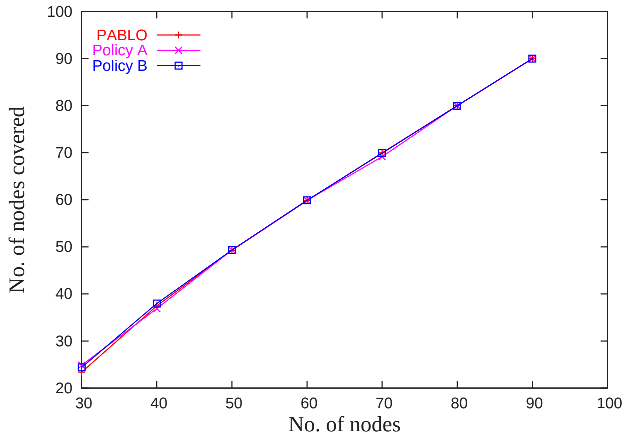

Fig. 8. Node Coverage Vs Network Size

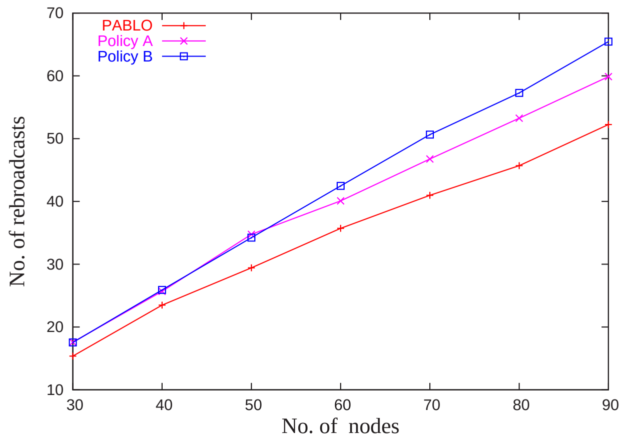

Fig. 9. Total Number of Rebroadcasts Vs Network Size

BIP's sweeping procedure can also be considered to be one type of pruning. In our work, however, the pruning process is used in a different context: pruning takes place *after* the local optimization computations are completed. Our goal herein is not to compare the performance of potential pruning techniques but rather to explore whether a speculative pruning would lead to beneficial results as expected due to the behavioral shift discussed above.

B. Neighborhood Pruning

With neighborhood pruning in PABLO, after performing local optimizations, a node would choose to exclude one or more furthest neighbors from its range. Our initial goal was to choose a probability with which the furthermost neighbor is dropped. We note that, this can potentially cause a lapse in coverage if the network is sparse. A node with simply two neighbors may exclude one of the neighbors from consideration which could sometimes prevent the broadcast from reaching a plurality of the nodes in the network that can only be reached via this excluded neighbor. In order to remedy this, we modify the scheme so as to prune neighbors based on the degree (after local optimizations) of a node. Thus, if after local optimizations, a node is of high degree, the node would almost certainly exclude its furthest neighbor. If, on the other hand, the degree is small, the node would not perform pruning. The actual policies that we consider for pruning are discussed in the following subsection.

C. Performance Revisited

We now examine the results from our simulation experiments, observed after modifying PABLO with neighborhood pruning. The metrics of evaluation remain the same as before. In order to perform pruning we consider the following policies: (i) *Policy A*: A node simply prunes its furthest neighbor after local optimizations if it has at least two neighbors (ii) *Policy B*: A node would prune 25% of its neighbors if it has a degree of 4 or larger, after performing its local optimizations. We compare the performance of PABLO in its native state and with modifications in accordance with Policy A and Policy B.

Total energy consumed: Figure 7 shows the percentages of reduction in energy consumed with the modified versions as per Policy A and Policy B with respect to the original PABLO version. We observe that energy savings could range from about 5% up to approximately 20%. At low densities, Policy A seems to be better in terms of energy savings. The small neighborhood size in sparse settings precludes efficient pruning via Policy B. However, pruning larger numbers at higher densities via Policy B provides better savings. Our simulation results verify that the node coverage is not affected by our modifications (Fig. 8). In other words, the pruning policies, almost always, only eliminate redundant broadcast receptions.

Duration of the broadcast: We observe in Fig. 10 that PABLO, modified as per our policies, is completed within a duration that is almost identical to that of its original counterpart. The important point to re-emphasize is that

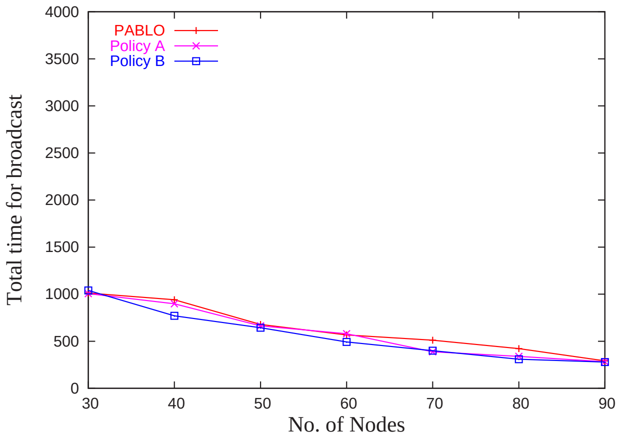

Fig. 10. Total Broadcast Duration Vs Network Size

this is an advantage over BIP which takes significantly longer time to complete a broadcast. This is especially important for time-critical situations (such as route error messages).

Total number of rebroadcasts: In Fig. 9, we observe that our policies cause the number of rebroadcasts to increase. This is because, pruning the furthest node would cause lower power links and therefore, results in a higher number of uncovered nodes per rebroadcast. Nodes that had possibly quelled their transmissions (as per the protocol rules [13]) could now be potentially forced to transmit if they have uncovered neighbors. However, note that these additional rebroadcasts occur almost in parallel with previously existing broadcasts; thus, as seen in Fig. 10 the duration of the broadcast is largely unaffected.

To summarize, these results demonstrate that our pruning policies can maintain the desirable properties of PABLO (such as latency and coverage) while providing additional energy savings of up to 20%.

VI. Conclusions

Our objective in this paper is to examine and understand the trade-offs between global and local approaches to the broadcast problem. Towards this objective, we compared the performance of a global algorithm, BIP and a localized protocol, PABLO, that relies on two-hop neighborhood information. To the best of our knowledge, this is the first work that compares global and local power adaptive protocols. Our work leads to the following observations.

- The reduced power consumption is traded-off for increased broadcast duration in the global approach. In investigating this further, we find that the rebroadcasts follow lengthy paths, which minimize the power but delay the termination of the broadcast process. BIP nodes try to reach only a very few additional nodes at each rebroadcast step. Thus, the power consumed at each rebroadcast step is much smaller than with the localized scheme, PABLO, wherein, the reach per rebroadcast is much larger.
- We propose a neighborhood pruning mechanism to improve the energy conserved with PABLO based on reducing redundancy in the rebroadcast optimization while incorporating the behavioral traits of BIP. We show via simulations that total energy can be reduced by up to approximately 20% compared to the original PABLO scheme, while the node coverage and the latency profile remain essentially unchanged.

Although global approaches are of limited practical interest, their study can help us develop better local protocols. Through such studies, we obtain a deeper understanding of the fundamental trade-offs and limitations of the broadcast problem. In addition, we can comprehend and identify the weaknesses of the local approaches, and thus, improve them, as we do herein.

References

[1] A. Muqattash and M.Krunz, “Power Controlled Dual Channel (PCDC) Medium Access Protocol for Wireless Ad Hoc Networks”, *Proc. of IEEE INFOCOM, 2003.*

[2] B.Williams and T.Camp, “Comparison of Broadcasting Techniques for Mobile Ad Hoc Networks”, *Proc. of ACM MOBIHOC, 2002.*

[3] E. Jung and N.H. Vaidya, “A Power Control MAC Protocol for Ad Hoc Networks”, *Proc. of ACM MOBICOM, 2002.*

[4] E. Royer and C-K. Toh, “A Review of Current Routing Protocols for Ad hoc Wireless Networks”, *IEEE Personal Communications Magazine*, 1999.

[5] J.Cartigny, D.Simplot and I.Stojmenovic, “Localized Minimum-Energy Broadcasting in Ad-hoc Networks”, *Proc. of IEEE INFOCOM, 2003.*

[6] J. E.Wieselthier , G. D. Nguyen, A. Ephremides, “On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks”, *Proc. of IEEE INFOCOM, 2000.*

[7] J. Wu and F. Dai, “Broadcasting in Ad Hoc Networks Based on Self-Pruning,” *Proc. of the 22nd Annual Joint Conf. of IEEE Communication and Computer Society (INFOCOM), March 2003.*

[8] L.M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc, “Power consumption in packet radio networks”, *Proc. of STACS, 1997.*

[9] M. Cagalj, J. P. Hubaux and C. Enz, “Minimum-Energy Broadcast in All-Wireless Networks: NP-Completeness and Distribution Issues”, *Proc. of ACM MOBICOM, 2002.*

[10] 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.*

[11] S. Singh, M. Woo and C. S. Raghavendra “Power-Aware Routing in Mobile Ad Hoc Networks”, *Proc. of ACM MOBICOM, 1998.*

[12] 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.*

[13] X. Chen, M. Faloutsos, and S. Krishnamurthy, “Power Adaptive Broadcasting with Local Information in Ad Hoc Networks”, *Proc. of IEEE ICNP, 2003.*

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U. S. Government.