

Resource-Limited Energy-Efficient Wireless Multicast of Session Traffic

Jeffrey E. Wieselthier
Code 5521
Information Technology Division
Naval Research Laboratory
Washington, DC 20375
wieselthier@itd.nrl.navy.mil

Gam D. Nguyen
Code 5521
Information Technology Division
Naval Research Laboratory
Washington, DC 20375
nguyen@itd.nrl.navy.mil

Anthony Ephremides
Electrical and Computer Eng. Dept.
and Institute for Systems Research
University of Maryland
College Park, MD 20742
tony@eng.umd.edu

Abstract

In this paper, we address the impact of resource limitations on the operation and performance of the broadcasting and multicasting schemes developed for infrastructureless wireless networks in our earlier studies. These schemes, which provide energy-efficient operation for source-initiated session traffic, were previously studied without fully accounting for such limitations. We discuss the “node-based” nature of the all-wireless medium, and demonstrate that improved performance can be obtained when such properties are exploited by networking algorithms. Our broadcast and multicast algorithms involve the joint choice of transmitter power and tree construction, and thus depart from the conventional approach that makes design choices at each layer separately. We indicate how the impact of limited frequency resources can be addressed. Alternative schemes are developed for frequency assignment, and their performance is compared under different levels of traffic load, while also incorporating the impact of limited transceiver resources.

1. Introduction

In our earlier studies ([1], [2], [3], [4]) we developed energy-efficient algorithms for the construction of broadcast and multicast trees for all-wireless (i.e., infrastructureless, or ad hoc) multihop networks, and evaluated their performance under the assumption that ample transceiver resources and/or bandwidth are available. In this paper, we extend our previous results by incorporating the limitations imposed by the joint constraints of a finite number of transceivers at each network node and a finite number of available frequencies.

A novel feature of our approach is that instead of viewing energy efficiency from the perspective of low-power equipment or highly efficient batteries, we address it as a network design problem. We argue that it may be necessary to abandon the traditional layered network architecture in favor of new approaches that permit the vertical coupling of protocol layer functionality, thereby permitting improved energy efficiency; e.g., the routing

algorithm (multicast tree construction) should be coordinated with the choice of transmitter power levels because the latter determine the connectivities that are available to the former. Furthermore, new paradigms must be developed to reflect the “node-centric” nature of wireless communications, which provides a vastly different communications environment from the “link-centric” nature of wired networks.

Multicasting in wireless networks is fundamentally different from multicasting in “wired” or “tethered” networks. In addition to node mobility (and, hence, variable connectivity in the network), there are trade-offs between the “reach” of wireless transmission (namely the simultaneous reception by many nodes of a transmitted message) and the resulting interference by that transmission. Furthermore, there are trade-offs between reach and energy expenditure. We assume that the power level of a transmission can be chosen within a given range of values. Therefore, there is a trade-off between reaching more nodes in a single hop by using higher power (but at a higher interference cost and energy cost) versus reaching fewer nodes in that single hop by using lower power (but at a lower interference cost and energy cost). The wireless medium presents an environment that is vastly different from that of wired networks. Therefore, novel approaches are needed to exploit the properties of the wireless channel, while satisfying the additional constraints that it imposes.

Few studies have addressed multicasting in wireless networks. For example, the problem of multicast scheduling in cellular mobile networks was studied in [5], a forwarding multicast protocol for noncellular networks was studied in [6], and the performance of several multicasting protocols for ad hoc wireless networks is compared in [7]. Virtually all multicasting studies have been limited to the case of stationary networks that are not wireless (e.g., [8], [9], [10]).

In [1], we discussed the fundamental issues associated with energy-efficient multicasting, and proposed and evaluated several multicasting schemes. In [2] and [3] we developed the Broadcast Incremental Power (BIP) and Multicast Incremental Power (MIP) algorithms for energy-efficient tree formation, and demonstrated that they perform better than previously studied schemes. In [4] we studied the impact of limited bandwidth on the performance of the MIP scheme. In the present paper we

This work was supported by the Office of Naval Research.

extend our study to the case in which there are limitations on both bandwidth and transceiver resources.

After a brief discussion of the all-wireless medium, we describe several algorithms we have developed for wireless broadcasting and multicasting of “session” (or connection-oriented) traffic, and indicate how these algorithms exploit the properties of the wireless channel. We discuss the incorporation of limited bandwidth into our algorithms, which were originally developed and evaluated under the assumption that ample frequency resources are available. We evaluate the trade-offs between algorithm complexity (and hence scalability) and performance. Our performance results demonstrate that the incorporation of energy considerations into the multicast algorithms can, indeed, result in energy saving.

To assess the complex trade-offs one at a time, we assume in this paper that there is no mobility. Nevertheless, the impact of mobility can be incorporated into our models because transmitter power can be adjusted to accommodate the new locations of the nodes, as necessary. In other words, the capability to adjust transmission power provides considerable “elasticity” to the topological connectivity, and hence may reduce the need for hand-offs and tracking. But this issue is not addressed in this paper.

We use a performance “yardstick” that reflects the desire to reach a large fraction of the desired destinations while maintaining energy efficiency. A destination may not be reached for any of the following reasons, which are discussed in greater detail in the paper: (1) lack of connectivity (i.e., excessive distance between nodes), (2) lack of equipment (i.e., all of the transceivers at one or more nodes in the multicast tree are already occupied with other traffic), or (3) lack of bandwidth (i.e., a node’s transmission would interfere with, or suffer interference from, the transmission of another node). Additionally, an admission-control process may be used to reject costly destinations, although we don’t address this possibility in the present paper. Performance is evaluated by means of simulation.

We do not address the protocol issues associated with determining connectivity and reserving resources, but instead focus on the fundamental issues associated with the determination of energy-efficient broadcast and multicast trees, assuming the existence of the underlying protocol that supplies the necessary topological connectivity information.

2. Wireless Communications Model

We consider source-initiated, circuit-switched, multicast sessions. The maintenance of a session requires the dedication of a transceiver at each participating node (source node, relay nodes, and destination nodes) throughout the duration of the session. The network consists of N nodes, which are randomly distributed over a specified region. Each node has several (say T) transceivers, and can thus support up to T multicast sessions simultaneously. We assume that there are F frequencies available to the network nodes. Frequencies can be reused, provided that doing so does not create interference, as discussed below. Thus, congestion (and

hence call rejection) may arise when either an insufficient number of transceivers or an insufficient number of frequencies are available.

Alternatively, it would be possible to consider a system that uses code-division multiple access (CDMA), rather than frequency-division multiple access (FDMA). Doing so would eliminate the difficult problem of assigning non-interfering frequencies because (at least in principle) quasi-orthogonal codes can be used. However, direct-sequence CDMA systems suffer from the near-far problem, and from an inability to support simultaneous transmission and reception in the same frequency band. Although frequency-hopped systems are less affected by the near-far problem, they are subject to spectral splatter, which can be especially troublesome when a node simultaneously transmits and receives at neighboring frequencies. By considering FDMA systems, we are able to assess the impact of limited bandwidth resources, and thereby to form the basis for future studies of specific systems, including those that use CDMA. It is also of interest to study systems that use time-division multiple access (TDMA), rather than multiple transceivers, to support multiple sessions simultaneously. In TDMA-based systems, the need to assign specific time slots creates a much more difficult problem than that of simply assigning any (of perhaps several available) transceiver to a new session. The study of TDMA-based systems is a topic for future research.

Any node is permitted to initiate multicast sessions. Multicast requests and session durations are generated randomly at the network nodes. Each multicast group consists of the source node plus at least one destination node. Additional nodes may be used as relays either to provide connectivity to all members of the multicast group or to reduce overall energy consumption. The set of nodes that support a multicast session (the source node, all destination nodes, and all relay nodes) is referred to as a *multicast tree*. Notice the difference between this definition and the conventional one that is based on links (or edges); here the links are incidental and their existence depends on the transmission power of each node. Thus it is the nodes (rather than the links) that are the fundamental units in constructing the tree.

The connectivity of the network depends on the transmission power. We assume that each node can choose its power level p , such that $p_{min} \leq p \leq p_{max}$. The nodes in any particular multicast tree do not necessarily have to use the same power levels; moreover, a node may use different power levels for the various multicast trees in which it participates.

We assume that the received signal power is equal to $pr^{-\alpha}$, where p is the transmission power, r is the distance and α is a parameter that typically takes on a value between 2 and 4, depending on the characteristics of the communication medium. We use a simplified interference model in which we assume that the interference level is independent of network traffic and the same at all nodes. Based on this model the transmitted power required to support a link between two nodes separated by distance r is proportional to r^α , since the received power must exceed some threshold (which depends on factors such as signal parameters, detector

structure, and noise levels). Without loss of generality, we set the threshold constant equal to 1, resulting in:

$$p_{ij} = \text{power needed for link between Node } i \text{ and Node } j \\ = r_{ij}^\alpha,$$

where r_{ij} is the distance between Node i and Node j . If the maximum permitted transmitter power p_{max} is sufficiently large, the network is fully connected. We also use a nonzero value of p_{min} (the minimum transmission power) as a way to account for the fact that the r^α dependence applies only in the far-field region (i.e., for nodes that are arbitrarily close, the minimum necessary transmission power to ensure connectivity is not arbitrarily small).

We assume the use of omnidirectional antennas; thus all nodes within communication range of a transmitting node can receive its transmission. In such cases, we can exploit the “wireless multicast advantage,” described in [2] and [3]. For example, consider a situation in which Node i transmits directly to its neighbors, Node j and Node k ; the power required to reach Node j is P_{ij} and the power required to reach Node k is P_{ik} . A single transmission at power $P_{i,(j,k)} = \max\{P_{ij}, P_{ik}\}$ is sufficient to reach both Node j and Node k (rather than the sum of these powers, as in wired applications).

2.1 Node-Based Communication Models

As a result of the wireless multicast advantage, an appropriate view of the omnidirectional wireless communication medium is as a *node-based* environment that is characterized by the following properties:

- A node’s transmission is capable of reaching another node if the latter is within communication range, which in turn means that the received signal-to-interference-plus-noise ratio exceeds a given threshold and that the receiving nodes have allocated (scheduled) receiver resources for this purpose.
- The total power required to reach a set of other nodes is simply the *maximum* required to reach any of them individually.

By contrast, in wired models, as long as there is a wire or cable link connecting two nodes, the reception is ensured over that link, and the cost of Node i ’s transmission to Node j and Node k would be the *sum* of the costs to the individual nodes.¹ Thus, wired networks can be viewed correctly as *link-based*.

The node-based nature of wireless networks necessitates the development of new networking techniques, because the models developed for wired networks do not adequately capture the characteristics of the wireless medium. For example, in wired networks, the broadcasting problem can be formulated as the well-known minimum-cost spanning tree (MST) problem.

¹ In wired networks, energy is not a concern; the cost of a link would typically be related to bandwidth and congestion (and hence delay) considerations. The case of wireless applications with highly directive antennas is similar to the case of wired networks in the sense that multiple beams may be needed to reach multiple destinations; thus the total cost of a node’s transmissions to its neighbors would be equal to the sum of the cost of the individual beams needed to reach each individual destination.

This formulation is based on the existence of a cost associated with each link in the network; the total cost of the broadcast tree is the sum of the link costs. The situation in wireless networks is different, however, because of the “wireless multicast advantage” property, which permits all nodes within communication range to receive a transmission without additional expenditure of transmitter power. Therefore, the standard MST problem, which reflects the *link-based* nature of wired networks, does not capture the *node-based* nature of wireless networks. We do not know of any scalable solutions to the node-based version of this problem. Related studies of complexity of tree construction and energy-efficient connectivity establishment, which do not exactly apply to our model, can be found in [11], [12], and [13].

In this paper we compare the performance of a new node-based multicasting scheme with that of two other schemes, which are adapted from those used for conventional link-based wired networks. We demonstrate that the use of node-based schemes can, in fact, provide improved performance.

3. A Multicasting Problem

We now address the problem of determining an appropriate multicast tree for each arriving multicast session request, so that a reward function (which incorporates both throughput and energy efficiency) is maximized. The establishment of a multicast tree requires the specification of the transmitted power levels, the frequencies used by each node, and the commitment of the needed transceiver resources throughout the duration of the multicast session.

3.1 Admission-Control Policies

We say that a destination can be *reached* if the following conditions are satisfied:

- there exists a path from the source to it (i.e., the transmitted power required to support the path does not exceed p_{max} at any node);
- a transceiver is available (i.e., not already supporting another session) at each node along the path;
- a suitable frequency assignment can be found to support the path (i.e., a non-interfering frequency is available to support the link between each node pair in the network along the path; these frequency assignments must not interfere with, or suffer interference from, currently ongoing sessions).

The results presented in this paper are based on the use of the “admit-all” admission-control policy, in which all multicast requests are accepted as long as one or more of the intended destinations can be reached, and in which paths are established to all reachable destinations, regardless of the cost required to do so (subject to the restriction that the transmitted power does not exceed p_{max} at any node). We are currently investigating admission-control policies that, when used in conjunction with our tree-formation algorithms, can improve performance based on the criteria discussed below.

3.2 Performance Metrics

Our performance measure must incorporate the characteristics of the multicast problem, as well as the need to conserve energy. In view of the fact that partial multicast sessions may take place, the performance metric should provide a reward that reflects the number of destinations that are actually reached. We define

- n_i = # of intended destinations by i th multicast session.
- m_i = # of destinations reached by i th multicast session.
- p_i = sum of the transmitter powers used by all nodes in i th multicast session.

The following performance metrics are studied in this paper.

Multicast efficiency

We define the *multicast efficiency* of the i th multicast session to be the fraction of desired destinations of that session request that are actually reached. Then, the overall multicast efficiency over an observation interval of X multicast requests can be defined as:

$$e = \frac{1}{X} \sum_{i=1}^X \left(\frac{m_i}{n_i} \right). \quad (1)$$

The "Yardstick" metric

To take into consideration the often-conflicting objectives of reaching as many destinations as possible and of maximizing the number of destinations reached per unit energy, we define a *local yardstick* measure:

$$y_i = \left(\frac{m_i}{p_i} \right) \left(\frac{m_i}{n_i} \right). \quad (2)$$

Our *global yardstick* Y is the average value of y_i :

$$Y = \frac{1}{X} \sum_{i=1}^X y_i = \frac{1}{X} \sum_{i=1}^X \left(\frac{m_i}{p_i} \right) \left(\frac{m_i}{n_i} \right). \quad (3)$$

In this paper we do not place a hard limit on the energy resources at the individual nodes, but instead evaluate performance based on the metrics discussed above, which we have found to be useful in the development of energy-aware protocols. We are currently at the early stages of implementing our algorithms under the assumption that each node has a finite quantity of energy; we use node-based cost metrics that reflect the "residual energy" that is still available at each node.

3.3 "Local" Cost Metrics

The problem of finding the multicast tree that maximizes the local yardstick y_i for each new multicast request is highly complex, and not feasible, except for small examples. Moreover, maximizing y_i for each i does not guarantee the maximization of the global yardstick Y . Therefore, we have found it necessary to take the approach of minimizing a cost function that is related to the ultimate objective, but only indirectly, and which is based on the use of local (i.e., per multicast request) cost metrics. Our BIP algorithm (see Section 4) uses node-

based metrics rather than on the more-conventional link-based metrics.

Link-based metrics assign a cost to the maintenance of each link, e.g., the power needed to maintain the link. The total cost of a multicast tree is then the sum of the costs of the links that form the tree. Such metrics do not reflect the wireless multicast advantage property, discussed in Section 2. By contrast, node-based costs (e.g., the power needed by a node to reach all of its neighbors in the tree, i.e., the maximum power needed to reach any individual neighbor) do reflect the wireless multicast advantage property. The total cost of a multicast tree is then the sum of the costs of the transmitting nodes that form the tree.

Since (under our assumptions of omnidirectional antennas and no interference) a node's transmission can be received by all of its neighbors, it is best to design a tree that exploits the wireless multicast advantage. Tree formation consists of the choice of transmitting nodes and their transmitting powers. The total cost of the tree is then the sum of the powers of all transmitting nodes. Here we consider only the energy used for transmission, neglecting for the present the energy associated with reception and signal processing. A minimum-cost tree is then one that reaches all reachable nodes with minimum total power. We know of no scalable algorithms for the minimum-cost broadcast tree problem, and certainly not for the presumably more difficult problem of minimum-cost multicasting.

4. Minimum-Energy Broadcast Trees

Before addressing the problem of multicasting, we discuss an algorithm for the more-fundamental (but simpler) problem of wireless broadcasting, in which the goal is to form a tree from the source to all other nodes. We then demonstrate how this broadcasting algorithm can be adapted to multicasting.

We consider the problem of constructing the minimum-energy, source-based broadcast tree for each newly arriving broadcast session request. Doing so involves the choice of transmitter-power levels, relay nodes, and transmission frequencies. The total energy of the broadcast tree is simply the sum of the energy expended at each of the transmitting nodes in the tree; leaf nodes (which do not transmit) do not contribute to this quantity. Since we are considering session traffic, all transmitting nodes transmit for the entire duration of each session. Therefore, the total transmission energy is proportional to the total power needed to maintain the tree. Hence, we evaluate performance in terms of the total power required to maintain the tree.

We assume that each node has T transceivers, and can thus participate in at most T multicast sessions simultaneously. If a node is already supporting T sessions, the cost of adding the node to the tree is set to ∞ .² It is more difficult to incorporate the effect of a limited number of frequencies, because doing so requires

² It is also possible to associate a higher cost with nodes that have low "residual capacity" (i.e., few available transceivers); however, we do not do so in this paper.

that one keep track of all frequencies in use at potentially interfering nodes (see Section 5). As a result of either an insufficient number of transceivers at one or more nodes, or the unavailability of a non-interfering frequency at one or more nodes, some trees may not reach all destinations and/or may use more than the minimum energy (because only suboptimal trees can be constructed).

4.1 A node-based algorithm: Broadcast Incremental Power (BIP)

In [2] and [3] we introduced the ‘‘Broadcast Incremental Power’’ (BIP) heuristic, a node-based algorithm that takes into account the wireless multicast advantage in the formation of low-energy broadcast trees. BIP is similar in principle to Prim’s algorithm for the formation of MSTs, in the sense that new nodes are added to the tree one at a time (on a minimum-cost basis) until all nodes are included in the tree. In fact, the implementation of this algorithm is based on the standard Prim algorithm, with one fundamental difference. Whereas the inputs to Prim’s algorithm are the link costs P_{ij} (which remain unchanged throughout the execution of the algorithm), BIP must dynamically update the costs at each step (i.e., whenever a new node is added to the tree) to reflect the fact that the cost of adding new nodes to a transmitting node’s list of neighbors is the incremental cost. Consider an example in which Node i is already in the tree (it may be either a transmitting node or a leaf node), and Node j is not yet in the tree. For all such Nodes i (i.e., all nodes already in the tree), and Nodes j (i.e., nodes not yet in the tree), the following is evaluated:

$$P'_{ij} = P_{ij} - P(i), \quad (4)$$

where P_{ij} is the link-based cost of a transmission between Node i and Node j (i.e., it is r_{ij}^α), and $P(i)$ is the power level at which Node i is already transmitting (prior to the addition of Node j ; if Node i is currently a leaf node, $P(i) = 0$). The quantity P'_{ij} represents the incremental cost associated with adding Node j to the set of nodes to which Node i already transmits. The pair $\{i, j\}$ that results in the minimum value of P'_{ij} is selected, i.e., Node i transmits at a power level sufficient to reach Node j . Thus, one new node is added to the tree at every step of the algorithm.

Unlike Prim’s algorithm, which guarantees the formation of minimum-cost spanning trees for link-based costs (as in wired networks), BIP does not necessarily provide minimum-cost trees for wireless networks. However, neither do any other scalable algorithms that we are aware of. The performance results of Section 7 demonstrate nonetheless that this algorithm does, in fact, provide satisfactory performance.

4.2 Link-Based Algorithms for Broadcasting

Two of the algorithms we have studied [2] are based on well-known techniques, namely the use of shortest unicast paths and the use of spanning trees, both of which use link-based costs. We summarize these schemes as follows.

Broadcast Least-Unicast-cost (BLU) Algorithm

A minimum-cost path from the source node to every other node is established. The broadcast tree consists of the superposition of these unicast paths.

Broadcast Link-based MST (BLiMST) Algorithm

A minimum-cost (minimum-power) spanning tree is formed using standard (link-based) MST techniques.

4.3 Complexity Considerations

The complexity of BLU, when implemented by means of the Dijkstra algorithm, is $O(N^2)$, where N is the number of nodes in the network [14] (p. 111).

The complexity of BLiMST, when implemented by means of Prim’s algorithm, is $O(N^3)$ when a straightforward implementation is used [14] (p. 524). However, a more-sophisticated implementation using a Fibonacci heap yields complexity $O(M + N \log N) = O(N^2)$, where $M = N(N - 1)/2$ is the number of links (in a fully connected network).

Since BIP is based on Prim’s algorithm, it also has complexity $O(N^3)$. Because of the need to update the costs P'_{ij} at each step of the algorithm, it is not yet clear whether the Fibonacci heap technique is applicable here.

4.4 The Sweep: Removing Unnecessary Transmissions

In [2] we note that the performance of our broadcast algorithms can be improved somewhat by using what we call the ‘‘sweep’’ operation, which detects redundant transmissions as well as transmissions that can be reduced in power. The sweep is used in the numerical results presented in this paper.

We have studied the two following sweep rules:

SW1: Construct the tree first; then sweep at each non-leaf node according to some ordered sequence.

SW2: Sweep at each step during the tree construction.

We have observed that SW1 typically provides better performance than SW2. We believe that this is because SW1 begins the sweep only after a complete multicast tree is formed. Therefore, any changes produced by the sweep can potentially affect major portions of the network. By contrast, under SW2 the sweep can affect only those nodes that have already been added to the tree, which is a small subset of the network in the early steps of the algorithm.

The complexity of the sweep, as currently implemented, appears to be bounded by approximately N^2 . In [2] we demonstrated that the sweep can improve performance modestly. For example, the percentage improvement achieved by the sweep is somewhat greater for BLU and BLiMST (typically 5 – 20%) than for BIP (typically 5 – 10%), but BIP typically provides better performance than the other schemes (both pre- and post-sweep).

5. Incorporation of Bandwidth Limitations

The discussion of BIP in the previous section assumes the availability of an infinite number of frequencies. However, in realistic situations the number of frequencies is finite, and poses a limitation to overall network throughput. Although, as noted in the previous section, it is straightforward to incorporate the impact of a finite number of transceivers into the implementation of BIP (i.e., by setting the cost of the node to ∞), the modeling of finite frequency resources is much more complicated.

Let us consider the case in which Node m wants to transmit to Node n . Any particular frequency f may be unusable for one of the following reasons:

- f is already in use (for either transmission or reception) at either Node m or Node n ;
- f is being used by one or more nodes that create interference at Node n , thereby preventing the reception of f ;
- the use of f by Node m would interfere with ongoing communications at other nodes.

In this section, we discuss the following basic greedy approaches for frequency assignment in our broadcast and multicast algorithms:

FA1: Assume the availability of an infinite number of frequencies when forming the tree (the approach used in [1], [2] and [3]). Then attempt to assign the available frequencies to the tree. The assignment process is complete when either frequencies have been assigned to all transmissions, or when no additional frequencies are available to support portions of the tree.

FA2: At each step of the tree-construction, the frequency is chosen along with the transmission power level.

Under FA1, the tree construction process ignores the possibility that frequencies may not be available to provide the required connectivity. Thus, if appropriate frequencies cannot be found along the paths to all desired destinations, then some destinations will not be reached. By contrast, under FA2 the tree is formed using only nodes that do, in fact, have frequencies available. (The cost of adding a node is set to ∞ if a non-interfering frequency is not available.) Again, there is no guarantee that all destinations will be reached. However, FA2 provides a richer search space than FA1.

Note that FA1 and FA2 actually represent classes of frequency assignment policies, rather than single well-defined schemes. We have used greedy versions, in which frequencies are assigned using an orderly procedure, without the possibility of backtracking to change assignments and without the use of exhaustive search (or other scheme) to determine whether a consistent frequency assignment is possible. Thus, either of these schemes can result in unreached destinations, even though they might be reachable through a better frequency assignment. But this is a common characteristic of all heuristic procedures. In this paper, we compare the performance of our algorithms with two

others, which also use the same greedy approach to frequency assignment.

Let us consider the operation of BIP for the case in which the number of frequencies F is finite. Under FA2 the cost of a transmission is set to infinity if no frequency is available. Also, when evaluating the incremental cost of Eq. (4), the multicast advantage applies only when the same frequency can be used by Node i to reach all of its intended neighbors. Typically, the use of FA2 permits the construction of trees that reach a larger number of the desired destinations.

6. Algorithms for Multicasting

It is well known that the determination of a minimum-cost multicast tree in wired networks is a difficult problem, which can be modeled as the NP-complete Steiner tree problem. This problem appears to be at least as hard in wireless networks as it is in wired networks. As we noted earlier, we know of no scalable algorithms for the minimum-energy broadcast problem. Thus, heuristics are needed.

We have considered two basic approaches for multicasting:

- Pruning the broadcast tree;
- Superposing the minimum-cost unicast paths to each individual destination.

Examples of these approaches are discussed below.

6.1 Approaches based on Pruning

In this paper we present results based on the Multicast Incremental Power (MIP) Algorithm, which is a straightforward modification of BIP. First, a broadcast tree is formed using BIP. To obtain the multicast tree, the broadcast tree is pruned by eliminating all transmissions that are not needed to reach the members of the multicast group. More specifically, nodes with no downstream destinations will not transmit, and some nodes will be able to reduce their transmitted power (i.e., if their more-distant downstream neighbors have been pruned from the tree). The same pruning technique can also be applied to broadcast trees produced by alternative algorithms, such as BLiMST (resulting in the algorithm MLI_{MST} [2]).

6.2 Approaches based on Unicast Paths

A minimum-cost path is established between the source and every desired destination. The multicast tree consists of the superposition of the appropriate unicast paths. The three algorithms most often used for finding shortest paths are the Dijkstra, Bellman-Ford, and Floyd-Warshall algorithms [15]. Each of these will find the shortest paths when link costs are independent of each other. However, we do not know of any algorithms that can incorporate the effects of other-user wireless interference while guaranteeing shortest paths.

Note, in this regard, that by placing a link (or node) cost on the components of a path we are able to use standard shortest tree algorithms for session traffic, while

in non-wireless networks such algorithms are applicable only to data traffic. This, in itself, represents an important contribution to optimal routing of session calls in wireless environments.

6.3 On the Effectiveness of Alternative Schemes

Performance results in [2] and [3] (for an unlimited number of frequencies) indicate that multicasting schemes based on pruning (MIP and MLiMST) tend to work well when the number of destinations is a relatively large fraction of the total number of nodes (e.g., 25% or greater), whereas schemes based on unicast paths work well when the fraction of nodes that are destinations is small (e.g., 10% or less). For the case of broadcasting, our BIP algorithm performed better than not only unicast-based approaches, but it also performed better than a version of a MST algorithm that was based on link costs (rather than the node costs used by BIP). Because of its typically superior performance, we present results only for MIP in this paper.

7. Performance Results

We have simulated the performance of MIP for a network of 50 nodes that are randomly located in a region with dimensions 5×5 (arbitrary units of distance). We have obtained results for a propagation constant value of $\alpha = 2$, which results in required transmitter power values of r^2 to support a link between two nodes that are separated by distance r . We do not set limits on the maximum permitted value of transmitted power, i.e., $p_{max} = \infty$.

In our simulations, multicast requests arrive with interarrival times that are exponentially distributed with rate λ/N at each node. Session durations are exponentially distributed with mean 1. Multicast groups are chosen randomly for each session request; the number of destinations is uniformly distributed between 1 and $N-1$. Each simulation run consists of $X = 1,000$ multicast sessions, some of which may be (totally or partially) blocked because of lack of resources (i.e., transceivers and/or frequencies).

In this paper we present plots of yardstick and multicast efficiency as a function of the number of frequencies (for a fixed number of transceivers at each node) and as a function of the number of transceivers at each node (for a fixed number of frequencies). Frequencies can be reused at different locations in the network, provided that doing so does not create interference. We present results for two combinations of sweep rules and frequency assignment rules:

Scheme A: SW1 and FA1

Scheme B: SW2 and FA2

7.1 Yardstick Performance

Figure 1(a) shows the value of the yardstick Y as a function of the number of frequencies ($F = 1, 2, 4, 8, 16$, and 32) for $T = \infty$ and for operation under Scheme A;

each curve corresponds to a constant value of offered load λ (which starts at 0.125 at the uppermost curve, and doubles with each of the following curves). The value of Y is highest when λ is low and F is high. In such cases, most (if not all) destinations can be reached, and the more energy-efficient paths are almost always available. It is not clear why there is a slight decrease in the value of Y when F increases from 1 and 2 for λ between 1 and 8; otherwise the value of Y increases with increasing F , as expected.

Figure 1(b) shows similar results for the case of $T = 4$ transceivers at each node. At very low levels of offered load, the small number of transceivers has virtually no impact on performance. However, for $\lambda \geq 1$ Y is significantly lower than that observed for $T = \infty$ because of the insufficient number of transceivers. Little or no improvement is seen as F is increased past approximately 16 because performance is limited by the insufficient number of transceivers. In fact, Y decreases significantly as F increases past 8 for very large loads. In this region of operation, the unavailability of transceivers at many nodes, coupled with the availability of a large number of frequencies, results in the construction of trees with longer links; therefore, the power needed to maintain the tree increases and yardstick performance decreases.




Fig. 1 — Yardstick vs F for Scheme A.

An interesting property of the yardstick is that it exhibits saturation behavior when sufficient resources are available to handle the offered load. In the examples of Fig. 1, when $T = \infty$ and the offered load λ is very low, $F = 8$ is sufficient to achieve the maximum possible value of Y . As λ increases, it is necessary to increase F to reach the saturation value. The existence of such a saturation value (which is sensitive to λ and to the system resources T and F) suggests that our yardstick measure is, in fact, a reasonable measure of system performance. When $T = 4$, the same saturation value is reached at very low values of λ . But, as λ increases, the saturation value decreases, and increasing F past 16 does not result in improved performance.

We next consider the effect of varying T while keeping F fixed. Figure 2 shows the yardstick performance as a function of T for F fixed at ∞ , 16 and 8, again for operation under Scheme A. Performance is virtually identical for very low values of λ because few frequencies are needed when the traffic load is low. However, for $\lambda \geq 1$, increasing the value of F results in significantly improved performance, especially for large values of λ .

For $F = 16$ and $F = 8$, we observe that little improvement is seen when T is increased beyond more than approximately half the value of F . Again, we see the same type of saturation behavior that was observed for our examples with constant values of T and varying values of F .

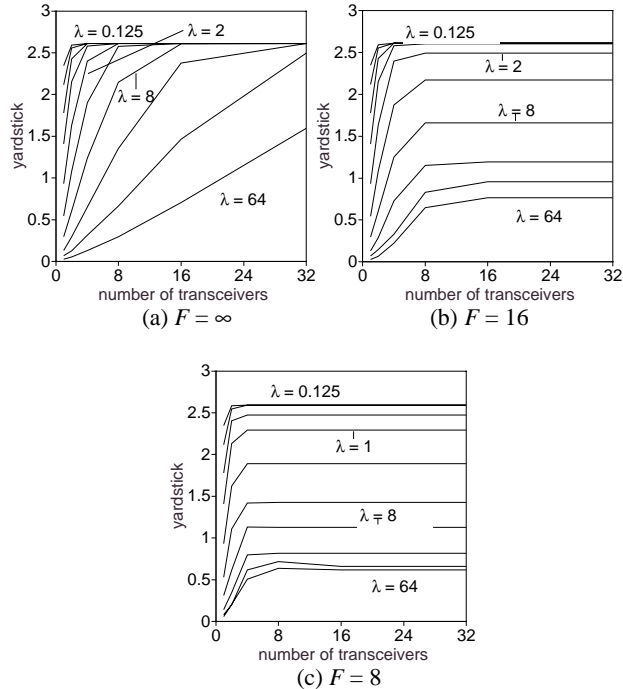


Fig. 2 — Yardstick vs T for Scheme A.

Figure 3 shows Y vs F when Scheme B is used. Again, results are shown for $T = \infty$ and $T = 4$. Results are qualitatively similar to those for Scheme A, except that there is no decrease in Y as F increases from 1 to 2. The use of Scheme B results in higher values of Y for low values of λ and small values of F than were observed for Scheme A. The improved performance in this region can be attributed to the fact that Scheme B verifies the availability of a frequency before adding a node to a tree. On the other hand, Scheme A provides somewhat higher saturation values of Y when F is large. We already commented in Section 4 that SW1 typically provides better performance than SW2 (in the sense of finding trees with lower total power, without regard to the availability of frequencies) because SW1 performs the sweep on the entire network. Additionally, when F is large it is best to find a complete low-cost tree before making frequency assignments (the approach of FA1) because the frequencies needed to implement the tree will always be available.

The yardstick performance as a function of T for fixed values of F , under Scheme B, is qualitatively similar to that for Scheme A. Thus we have not included curves for this case.




Fig. 3 — Yardstick vs F for Scheme B.

7.2 Multicast Efficiency

Figures 4(a) and 4(b) show multicast efficiency e vs F for Scheme A, for $T = \infty$ and $T = 4$, respectively. For $T = \infty$, the assured availability of transceivers permits e to increase as F is increased until the maximum possible value of 1 is reached; however, for $T = 4$, performance is limited by the insufficient number of transceivers.

Figure 5 shows similar results for the case of Scheme B. Some significant differences are apparent in the performance of these two schemes. Scheme A is “well behaved” in the sense that e increases monotonically with F . However, the behavior of Scheme B is considerably more interesting. First, we observe that a high value of e can be obtained when $F = 1$, i.e., when there is only a single frequency available. This is easily explained. Since there is only one frequency available, and since a node will not be added to the tree unless a frequency is available (because we are using FA2), the source node will gradually increase its power until as many of the desired destinations as possible are included in the network, resulting in a star configuration. The impact of a limited number of transceivers ($T = 4$) under Scheme B is similar to that under Scheme A.




Fig. 4 — Multicast Efficiency for Scheme A.

The second, perhaps surprising, observation is that e decreases significantly as F increases from 1 to 2. This behavior is also easily explained. Because of the availability of a second frequency, the source node will typically transmit with enough power to reach only those neighbors that are relatively close to it, and one or more of them will relay to their neighbors. Since there are only

two frequencies available, it is difficult to find a frequency assignment that will permit the construction of a tree to reach most of the destinations.




Fig. 5 — Multicast efficiency for Scheme B.

Finally, we note that, in the region of low congestion (low values of λ and high values of F), the multicast efficiency approaches 1.0 because all of the desired destinations are reached.

7.3 Multicast Efficiency per Unit Power

Although our yardstick incorporates aspects of both multicast efficiency and energy expenditure, it is also of interest to examine directly the relationship between e and average tree power, which we denote as \bar{P}_{tree} . These are the two factors of the yardstick measure. In Figs. 6 and 7 we plot e vs. \bar{P}_{tree} for Schemes A and B, respectively. Curves are shown for fixed values of F , as λ is varied between 0.125 and 64 (six curves are actually shown, i.e., for $F = 1, 2, 4, 8, 16,$ and 32). The quantity represented by the horizontal axis, \bar{P}_{tree} , is the average power needed to sustain the multicast tree for specified values of λ and F over the set of 1000 multicast requests. Hence, both e and \bar{P}_{tree} are dependent variables, which are obtained from the simulation results discussed above. Note that λ and F are the independent variables, and that e and \bar{P}_{tree} are obtained as a function of them. Low values of λ correspond to the upper right portion of each curve; both e and \bar{P}_{tree} decrease as λ increases.

Figure 6(a) shows that, for Scheme A and $T = \infty$, efficiency per unit power is virtually constant over the entire range of values of λ and F . (We have observed similar behavior when the propagation constant is $\alpha = 4$.) This means that if the independent variables change so as to increase e by a certain factor, \bar{P}_{tree} also increases by the same factor. Figure 6(b) shows that when $T = 4$, similar behavior is observed for $F \leq 8$; however lower values of this ratio and a nonlinear relationship are observed when more frequencies are available. This behavior is consistent with that of Fig. 1(b), where it was observed that the availability of few transceivers but many frequencies can result in trees that are not energy efficient.

Figure 7 shows that a similar result applies for Scheme B with $T = \infty$, but only for $F \geq 2$. For Scheme B with $F = 1$, efficiency per unit power is virtually constant over the entire range of λ , but is roughly half that for larger values

of F . The lower values of e per unit power are a consequence of the star configuration (and hence high transmission power) that results from the use of FA2 when $F = 1$. For $T = 4$, results are similar to those for $T = \infty$, but we observe the same nonlinear relationship between e and \bar{P}_{tree} shown in Fig. 6(b).




Fig. 6 — Multicast efficiency vs average tree power using Scheme A.




Fig. 7 — Multicast efficiency vs average tree power using Scheme B.

8. Conclusions

The wireless networking environment raises many issues that are not encountered in conventional networks, thus necessitating the development of novel techniques that exploit the properties of the communication medium. In addition, energy conservation is of paramount importance in wireless networks. Incorporating energy savings into the performance measures, as we have done in this paper, permits the definition of meaningful problems for routing and multicast tree construction. We have demonstrated the improvement that can be obtained by using a node-based algorithm for multicasting, and we have extended our earlier work by incorporating the impact of limited bandwidth and transceiver resources. Although our algorithms are by no means the only possible approaches, they are among the first to address this problem. Our simulation results have demonstrated some of the trade-offs that arise when finite equipment and bandwidth resources are available.

References

- [1] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides, "Algorithms for Energy-Efficient Multicasting in Static Ad Hoc Wireless Networks," to appear in *Mobile Networks & Applications (MONET)*, 2001.
- [2] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides, "On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks," *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, ISRAEL, pp. 585-594, March 2000.
- [3] J. E. Wieselthier, G. D. Nguyen, and A. Ephremides, "Energy-Efficient Broadcast and Multicast Trees in Wireless Networks," to appear in *Mobile Networking and Applications (MONET)*, 2001.
- [4] G. D. Nguyen, J. E. Wieselthier, and A. Ephremides, "The Effect of Bandwidth Limitations on Wireless Broadcasting and Multicasting," *2000 Conference on Information Sciences and Systems*, Princeton University, pp. TA7a-7—TA7a-12, March 2000.
- [5] M. Nagy and S. Singh, "Multicast Scheduling Algorithms in Mobile Networks," *Cluster Computing*, **1-2**, pp. 177-185, December 1998.
- [6] C.-C. Chiang, M. Gerla, and L. Zhang, "Forwarding Group Multicast Protocol (FGMP) for Multihop, Mobile Wireless Networks," *Cluster Computing*, **1-2**, pp. 187-196, December 1998.
- [7] S.-J. Lee, W. Su, J. Hsu, M. Gerla, and R. Bagrodia, "A Performance Comparison Study of Ad Hoc Wireless Multicast Protocols," *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, ISRAEL, pp. 565-574, March 2000.
- [8] C. Diot, W. Dabbous, and J. Crowcroft, "Multipoint Communication: A Survey of Protocols, Functions, and Mechanisms," *IEEE Journal on Selected Areas in Communications*, **15-3**, pp. 277-290, 1997.
- [9] M. H. Ammar, G. C. Polyzos, and S. K. Tripathi, "Special Issue on Network Support for Multipoint Communication," *IEEE Journal on Selected Areas in Communications*, **15-3**, April 1997.
- [10] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: A New Resource ReSerVation Protocol," *IEEE Network Magazine*, pp. 8-18, September 1993.
- [11] L. Viennot, "Complexity Results on Election of Multipoint Relays in Wireless Networks," Report RR-3584, INRIA, December 1998.
- [12] W.-T. Chen and N.-F. Huang, "The Strongly Connecting Problem on Multihop Packet Radio Networks," *IEEE Transactions on Communications*, **37-3**, pp. 293-295, March 1989.
- [13] A. E. F. Clementi, P. Penna, and R. Silvestri, "Hardness Results for the Power Range Assignment Problem in Packet Radio Networks," *Proc. Third International Workshop on Randomization and Approximation Techniques in Computer Science*, pp. 197-208, 1999.
- [14] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows*, Englewood Cliffs: Prentice Hall, 1993.
- [15] D. Bertsekas and R. Gallager, *Data Networks*, Englewood Cliffs: Prentice Hall, 1992.