



Energy-Efficient Multicasting of Session Traffic in Bandwidth- and Transceiver-Limited Wireless Networks

JEFFREY E. WIESELTHIER and GAM D. NGUYEN

Information Technology Division, Naval Research Laboratory, Washington, DC 20375, USA

ANTHONY EPHREMIDES

Electrical and Computer Engineering Department and Institute for Systems Research, University of Maryland, College Park, MD 20742, USA

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. The performance results include the comparison of our algorithms to alternative “link-based” algorithms for broadcasting and multicasting.

Keywords: wireless multicast, energy efficient, ad hoc network, network algorithm

1. Introduction

In our earlier studies [1–4] 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, namely, we show that network protocol choices have an impact on energy efficiency. Moreover, our approach exploits the “node-centric” nature of wireless communications, which provides a vastly different communications environment from the “link-centric” nature of wired networks. Specifically, we utilize the “wireless multicast advantage”, which permits a node’s single transmission to reach several neighbors simultaneously. Finally, we abandon the traditional layered network architecture in favor of a new approach that permits the vertical coupling of protocol layer functionality, thereby permitting improved energy efficiency. Specifically, we coordinate the routing algorithm (i.e., multicast tree construction) with the choice of transmitter power levels. Doing so can provide better performance because the power levels determine the connectivities that are available for establishing routing paths.

Wireless networks are characterized not only by node mobility (and, hence, variable connectivity in the network), but also by 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, because we assume that the power level of a transmission can be chosen within a given range of values as part of the multicast tree construction process.

Some prior work on multicasting in wireless networks includes the study of multicast scheduling in cellular mobile networks [5], the development of a forwarding multicast protocol for noncellular networks [6], and the performance analysis of several multicasting protocols for ad hoc wireless networks [7]. Almost all multicasting studies until now have focused on fixed, non-wireless networks (e.g., [8–10]).

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 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 that we have developed for wireless broadcasting and multicasting of “session” (or connection-

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2002		2. REPORT TYPE		3. DATES COVERED 00-00-2002 to 00-00-2002	
4. TITLE AND SUBTITLE Energy-Efficeint Multicasting of Session Traffic in Bandwidth- and Transceiver-Limited Wireless Networks				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory,Information Technology Division,4555 Overlook Avenue, SW,Washington,DC,20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 14	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

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, up to a point, the agility provided by the adjustment of transmission power permits the maintenance of the existing connectivities despite the change in node location. 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. However, this issue is beyond the scope of this paper.

We use a performance “yardstick” measure 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 last possibility in the present paper, the reader is referred to [11] for details. Performance is evaluated by means of simulation.

To avoid unnecessary complications, 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. Furthermore, we focus totally on energy expenditure associated with RF transmission, and ignore completely the energy consumption for processing at the transceiver. Although it will be necessary to include eventually all forms of energy use in our model, we initially emphasize RF transmission only, for simplicity and to understand the trade-offs involved. It is fairly straightforward to incorporate processing energy in our model, and we are currently investigating this possibility [12]. However, modeling the dependence of the energy consumption on the processor capabilities and the signal processing algorithms used would introduce unnecessary and obscuring complications.

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, de-

pending on the characteristics of the communication medium. We do not consider elaborate fading models for the wireless channel. Instead, we assume additive background Gaussian noise of constant amplitude throughout the network, which implies that, in the absence of other-user interference, 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:

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

where r_{ij} is the distance between nodes i and j . If the maximum permitted transmitter power $p_{\max}$ is sufficiently large, the network is fully connected. The use of a nonzero value of $p_{\min}$ (the minimum transmission power) provides a way to account for the fact that the $r^{-\alpha}$ 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). Of course, the connectivity map drawn based on this model is valid only if no other-user interference is present.

Since the total number, F , of frequencies available to the network is finite, we must be able to reuse them in parts of the network that are sufficiently separated. Thus, we use a simplified “binary” interference model, in which another node’s transmission will cause destructive interference at the receiving node (at a given frequency) if and only if the received signal is sufficiently strong to provide a reliable communication link. It would be straightforward to generalize this model, e.g., by considering the case in which several interfering signals add to a total interference level that causes the signal to other-user-interference ratio at the frequency of interest to fall below the required threshold value.

Last but not least, 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, nodes j and 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 nodes j and k .

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 nodes j and 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. 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, and we suspect and conjecture that the problem is NP-complete. Related studies of complexity of tree construction and energy-efficient connectivity establishment, which do not exactly apply to our model, can be found in [13–15].

In this paper we compare the performance of the new, node-based multicasting scheme MIP, which we proposed in [2,3], 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 (transmitted power) required to do so (subject to the restriction that the transmitted power does not exceed $p_{\max}$ at any node). We have recently investigated admission-control policies that, when used in conjunction with our tree-formation algorithms, can improve performance based on the criteria discussed below [11].

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.

3.2.1. 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:

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

3.2.2. 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 have recently extended our model by considering the implementation of our algorithms under the assumption that each node has a finite quantity of energy; in that case, we use node-based cost metrics that reflect the “residual energy” that is available at each node at any given time [11,12].

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 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. 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 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,” defined below. 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 incremental cost is evaluated:

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

where P_{ij} is the link-based cost of a transmission between nodes i and 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.

Unlike BIP, these algorithms do not exploit the wireless multicast advantage in the tree-construction process, i.e., link costs are evaluated independently under BLU and BLiMST. Nevertheless, since the wireless multicast advantage is a fundamental property of the omnidirectional wireless medium, its beneficial effects do have an impact on these algorithms as well. Once a tree is constructed, its power is evaluated as the sum of the transmitter powers at each transmitting node. Regardless of the algorithm used to construct the tree, a node’s transmitter power is the maximum power required to reach any of its downstream neighbors individually.

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 [16, p. 111].

The complexity of BLiMST, when implemented by means of Prim’s algorithm, is $O(N^3)$ when a straightforward implementation is used [16, 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, a straightforward implementation of it also has complexity $O(N^3)$. However, unlike BLiMST, 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 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.

SW1 can be used with any broadcast or multicast algorithm, whereas SW2 is appropriate for algorithms that add nodes to the tree one at a time, such as BIP and BLiMST. We have observed that, when used with BIP, SW1 typically provides better performance than SW2. We believe that this is because SW1 begins the sweep only after a complete 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.

In our implementation of SW1, we examine each transmitting node (there are $O(N)$ of them), in sequence, to see whether its transmitting power can be reduced without losing members from the tree (some nodes may be reassigned to new father nodes). This operation involves a search of all of its neighboring nodes that are not upstream of it in the tree (again, there are $O(N)$ of them). Therefore, the complexity of the sweep operation, which is performed once (i.e., after the tree is formed) is $O(N^2)$. Thus, the complexity of BLU and BLiMST remain $O(N^2)$ after the SW1 sweep, and that of BIP remains $O(N^3)$.

The implementation of SW2 requires its integration into BIP. Each of the $N - 1$ steps of BIP (without the sweep) has complexity $O(N^2)$. Once these calculations are performed at any step, the SW2 sweep requires additional computation, which has complexity $O(N)$; thus, the complexity at each step remains $O(N^2)$. Therefore, the complexity of BIP with SW2 is still $O(N^3)$.

In [2] we demonstrated that the sweep can provide further improvement. 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 [1–3]). 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 some of the desired destinations, then those 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. 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. Specifically, the lowest-numbered non-interfering frequency is used. Thus, either of these schemes can result in unreached destinations, even though they might be reachable through a better frequency assignment. Such nonoptimal performance is a common characteristic of heuristic procedures. In this paper, we use the same greedy approach to frequency assignment for all three algorithms that we study.

Let us consider tree construction using BIP with FA2 when the number of frequencies F is finite. The cost of a transmission is set to infinity if no frequency is available. Also, when evaluating the incremental cost of equation (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 [17]. 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

The Multicast Incremental Power (MIP) algorithm 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 MLiMST [2]).

6.2. An approach based on unicast paths

Recall that under BLU a minimum-cost path is established between the source and every destination. Under the Multicast Least-Unicast-cost (MLU) algorithm, the same approach is used, except that a minimum-cost path is established only to the desired destinations. 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 [18]. 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 interaction among links (e.g., the beneficial effects of the wireless multicast advantage and the harmful effects of other-user interference) while guaranteeing shortest paths.

Note, in this regard, that the assignment of a “local” (link or node) cost (e.g., transmitter power) to the components of a path enables the use of standard shortest-path algorithms for session traffic in wireless networks. By contrast, in wired networks such algorithms are normally not appropriate for session traffic because the typical performance metric for session traffic is blocking, which is a path metric, rather than a local metric (energy consumption is not relevant in wired networks). Expected delay (or queue size), a local metric that

is typically used for data traffic, is inappropriate in session-oriented applications (wired or wireless) because sufficient resources are reserved throughout the duration of the session; thus, there is normally no appropriate local metric for session traffic in wired networks. However, in wireless networks where energy consumption is of importance, there is a well-motivated local metric, namely transmitter power. Our use of shortest-path algorithms, in itself, represents an important contribution to optimal routing of session calls in wireless environments.

7. Performance results

In this section, we first compare the energy efficiency of multicast trees produced by our node-based algorithm MIP with those produced by the link-based algorithms MLU and MLiMST. We then address the impact of a limited number of transceivers and frequencies by evaluating performance over an extended period of time for a range of multicast-session arrival rates.

7.1. Energy-efficiency of multicast trees

In this subsection we summarize performance results from [2] and [3], which compare the effectiveness of MIP, MLU, and MLiMST in constructing low-power trees. We have simulated the performance of these three algorithms for 100 randomly generated networks of 100 nodes that are randomly located in a region with dimensions 5×5 (arbitrary units of distance). Here, we assume that the necessary transceivers and frequencies are available at all nodes. Thus, we are able to focus directly on the properties of the tree-construction algorithms, and therefore compare their effectiveness under comparable conditions, without addressing the impact of limited transceiver and frequency resources. We discuss results for propagation constant values of $\alpha = 2$ and 4 , which results in required transmitter power values of r^2 and r^4 , respectively, to support a link between two nodes that are separated by distance r . To enrich the search space of all possible trees, we do not set limits on the minimum or maximum permitted value of transmitted power, i.e., $p_{\min} = 0$ and $p_{\max} = \infty$.

Performance results in [2] and [3] indicate that multicasting schemes based on pruning broadcast trees (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 MLU, which is based on the superposition of unicast paths, works well when the fraction of nodes that are destinations is small (e.g., 10% or less). For example, for a propagation constant of $\alpha = 2$ and a group size of 5 nodes, the average total transmission power of trees generated by MIP was 14.2% greater than that of trees generated by MLU. For a group size of 10 nodes, the difference decreases to 1%. For a group size of 25 nodes, the situation reverses and the average power of trees produced by MLU is 9% greater than that of MIP. For the broadcast case (i.e., a group size of all 100 nodes), use of MLU (which becomes

BLU) produces trees with average power 14.4% greater than MIP (which becomes BIP). The superior performance of MIP and MLI_{MST} (as compared with MLU) for large multicast groups is easy to explain. When multicast groups are large, the structure obtained by first establishing a broadcast tree is highly beneficial. However, when multicast groups are small, many energy-inefficient paths are established; such behavior would be expected even for truly optimal broadcast trees, and is a consequence of the suboptimal nature of the pruning operation.

Comparing the relative performance of MIP and MLI_{MST}, both of which are based on the pruning of broadcast trees, we have seen that MIP provides better performance than MLI_{MST} over the complete range of network examples that we have studied, based on the criterion of mean tree power. Specifically, the average power of trees produced by MLI_{MST} was between 6.75 and 7% greater than that produced by MIP, with little variation in the percentage difference as the group size varied between 5 and 100. We attribute this improved performance to the fact that MIP exploits the node-based wireless multicast advantage property as it constructs the trees, whereas MLI_{MST} ignores this property as the trees are based solely on link costs. However, as noted in section 4.2, once the tree is constructed, the wireless multicast advantage does benefit the overall power expenditure.

As α increases, the penalty for using longer links increases; thus low-power trees may consist of a larger number of shorter links. Performance results for $\alpha = 4$, which are qualitatively similar to those for $\alpha = 2$, are also presented in [2] and [3]. Again, MLU produces trees with the lowest average power for a group size of 5. However, for a group size as small as 10, the average power of trees produced by MIP is 6.8% less than that of trees produced by MLU, and this difference increases to 16% as the group size increases to 100. Again, MIP performs better than MLI_{MST}; however, the improvement in performance provided by MIP as compared with MLI_{MST} is only about 2% over the entire range of group sizes.

Based on the results summarized here (using the performance criterion of minimizing tree power), for moderate to large multicast groups, MIP provides the best performance, with MLI_{MST} and MLU following in that order. It is still difficult, however, to predict the relative performance of these algorithms in realistic environments that are characterized by limited resources and by many session arrival requests over a period of time. We address such issues in subsection 7.2.

7.2. Resource-limited multicast network performance

We now discuss the operation of a network over a period of time under our three algorithms, MIP, MLU, and MLI_{MST}. We consider a network of 50 nodes that are randomly located in a region with dimensions 5×5 (arbitrary units of distance). We present results for a propagation constant value of $\alpha = 2$, and we again set $p_{\min} = 0$ and $p_{\max} = \infty$. The impact of a finite number of transceivers and frequencies is now demonstrated.

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 the sweep rules and frequency assignment rules that were described earlier in sections 4 and 5; namely, we consider:

Scheme A: combination of SW1 and FA1,

and

Scheme B: combination of SW2 and FA2.³

7.2.1. A comparison of multicasting algorithms

Figure 1 shows the value of the yardstick Y as a function of the number of frequencies ($F = 1, 2, 4, 8, 16$, and 32) for the three algorithms we have studied, all operating using scheme A. The solid lines represent MIP, the dotted lines represent MLU, and the dashed-dotted lines represent MLI_{MST}. Figure 1(a) provides results for $T = \infty$, and figure 1(b) provides results for $T = 4$. Each curve corresponds to a constant value of offered load λ (from top to bottom the values are $\lambda = 0.25, 1, 4$, and 16). For all three algorithms, 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. As λ increases (for fixed values of T and F), congestion increases, resulting in reduced values of Y , since insufficient transceiver and/or frequency resources are available to reach all of the desired destinations. Figure 1(b) shows that the impact of the limited number of transceivers is significant for $\lambda > 1$.

The results for all three algorithms, for $T = 4$ transceivers at each node, show that at very low levels of offered load, the small number of transceivers has virtually no impact on performance. However, for $\lambda > 1$, 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 somewhat 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.

For $F \geq 2$ frequencies, MIP provides the best yardstick performance for all values of λ and F . It appears that, for




Figure 1. Yardstick Y vs. F for MIP, MLU, and MLiMST under scheme A (MIP: solid, MLU: dotted, MLiMST: dashed-dotted, top to bottom: $\lambda = 0.25, 1, 4$, and 16). (a) $T = \infty$ and (b) $T = 4$.

$T = \infty$, an asymptotic or “saturation” value of Y is reached as F increases. This value is about 10% greater for MIP than for MLU. The worst performance, for moderate to large values of F , is provided by MLiMST. The impact of a finite value of T is apparent in figure 1(b). Reducing T to 4 results in a decreased maximum value of Y for large values of F , and actually a decrease in Y as F increases at very high loads.

The superior performance of MIP, as compared to the other two algorithms, is not surprising in view of the performance results for energy-efficient tree construction discussed in section 7.1. Since the multicast group size is uniformly distributed between 2 and 50, the average group size is 26; we showed that MIP provides better performance than MLU when the group size is moderate to large. However, it is perhaps surprising that MLiMST provides the worst performance over a wide range of parameter values.

In an attempt to explain the relatively poor performance of MLiMST, we look at the multicast efficiency e for the three




Figure 2. Multicast efficiency e vs. F for MIP, MLU, and MLiMST under scheme A (MIP: solid; MLU: dotted; MLiMST: dashed-dotted; top to bottom: $\lambda = 0.25, 1, 4$, and 16). (a) $T = \infty$ and (b) $T = 4$.

algorithms. Figures 2(a) and (b) show e for $T = \infty$ and 4, respectively. As noted above, congestion increases as λ increases; thus e decreases as λ increases. We now address the relative performance of the three protocols for a given value of λ . For a wide range of system parameters, there is little difference in the value of e provided by MIP and MLU. However, MLiMST provides the lowest values of e for all system parameters. Thus, it appears that MLiMST makes relatively inefficient use of available resources (transceivers and frequencies), which results in more destinations being blocked, and hence in decreased values of Y . In view of the superior performance provided by MIP, we focus the remainder of our discussion of performance on it.

7.2.2. Yardstick performance

Now that the superior performance of MIP has been verified, we present a more-detailed set of performance results for it. Figure 3 is similar to figure 1, except that it provides results

for a more extensive set of offered loads. 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, for $T = \infty$ the value of Y increases with increasing F , as expected. For $T = 4$, there is a significant decrease in the value of Y at high congestion levels ($\lambda \geq 32$); this behavior was explained in conjunction with figure 1.

We noted earlier that the yardstick exhibits saturation behavior when sufficient resources are available to handle the offered load. In the examples of figure 3, 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 . However, 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 4 shows the yardstick performance as a function of T ($= 1, 2, 4, 8, 16, 32$) for F fixed at $\infty, 16$ 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 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 .




Figure 3. Yardstick vs. F for scheme A: (a) $T = \infty$ and (b) $T = 4$.

Figure 4. Yardstick vs. T for scheme A: (a) $F = \infty$, (b) $F = 16$ and (c) $F = 8$.

Figure 5 shows Y vs. F when scheme B is used. Again, results are shown for $T = \infty$ and 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, shown in figure 4. Thus we have not included curves for this case.




Figure 5. Yardstick vs. F for scheme B: (a) $T = \infty$ and (b) $T = 4$.

7.2.3. Multicast efficiency

Figures 6(a) and (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 7 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.




Figure 6. Multicast efficiency for scheme A: (a) $T = \infty$ and (b) $T = 4$.




Figure 7. Multicast efficiency for scheme B: (a) $T = \infty$ and (b) $T = 4$.

The second, perhaps surprising, observation is that e decreases significantly as F increases from 1 to 2. This behavior is also easily explained. Since a second frequency is available, 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. However, because 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.

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.

7.2.4. 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 $\bar{P}_{\text{tree}}$. These are the two factors incorporated into the yardstick measure. In figures 8 and 9 we plot e vs. $\bar{P}_{\text{tree}}$ 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




Figure 8. Multicast efficiency vs. average tree power using scheme A. (a) $T = \infty$ and (b) $T = 4$.

$F = 1, 2, 4, 8, 16$, and 32; the minimum and maximum values of e that are obtained for each value of F can be inferred from figures 6 and 7). The quantity represented by the horizontal axis, $\bar{P}_{\text{tree}}$, 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 $\bar{P}_{\text{tree}}$ 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 $\bar{P}_{\text{tree}}$ are obtained as a function of them. Low values of λ correspond to the upper right portion of each curve; both e and $\bar{P}_{\text{tree}}$ decrease as λ increases.

Figure 8(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, $\bar{P}_{\text{tree}}$ also increases by the same factor. Figure 8(b) shows that when $T = 4$, similar behavior is observed for $F \leq 8$; however lower values of this ratio and a non-linear relationship are observed when more frequencies are available. This behavior is consistent with that of figure 1(b), where it was observed that the availability of few transceivers




Figure 9. Multicast efficiency vs. average tree power using scheme B. (a) $T = \infty$ and (b) $T = 4$.

but many frequencies can result in trees that are not energy efficient.

Figure 9 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 $F = 1$ (see section 7.2). For $T = 4$, results are similar to those for $T = \infty$, but we observe the same nonlinear relationship between e and $\bar{P}_{\text{tree}}$ shown in figure 8(b).

8. Conclusions

The wireless networking environment raises many issues that are not encountered in wired networks, thus necessitating the development of novel techniques that exploit the properties of the wireless 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 the constraints of finite equipment and bandwidth resources are introduced.

Specifically, we first showed that our MIP algorithm provides better performance than link-based schemes that are adaptations of algorithms developed for wired applications. We then focused exclusively on MIP, and evaluated the dependence of yardstick and multicast efficiency on the number of transceivers and frequencies, as well as on offered load. Of particular interest is the behavior of multicast efficiency under scheme B when $F = 1$ and 2. We also discovered an interesting relationship between multicast efficiency and average tree power. Our performance results provide a basis for understanding energy-related behavior in wireless networks.

Acknowledgements

This work was supported by the Office of Naval Research.

Notes

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.
2. It is also possible to associate a higher cost with nodes that have low "residual capacity" (i.e., few available transceivers) or low "residual energy" (i.e., nearly depleted batteries); however, we do not do so in this paper.
3. When MIP and MLiMST are used with scheme A, the sweep operation takes place after pruning the broadcast trees produced by BIP and BLiMST, respectively. When scheme B is used with MIP, the sweep is performed during the execution of BIP, hence before pruning.

References

- [1] J.E. Wieselthier, G.D. Nguyen and A. Ephremides, Algorithms for energy-efficient multicasting in static ad hoc wireless networks, *Mobile Networks and Applications (MONET)* 6 (2001) 251–263.
- [2] J.E. Wieselthier, G.D. Nguyen and A. Ephremides, On the construction of energy-efficient broadcast and multicast trees in wireless networks, in: *Proceedings of IEEE INFOCOM 2000* (March 2000) pp. 585–594.
- [3] J.E. Wieselthier, G.D. Nguyen and A. Ephremides, Energy-efficient broadcast and multicast trees in wireless networks, *Mobile Networks and Applications (MONET)* (2002), to appear.
- [4] G.D. Nguyen, J.E. Wieselthier and A. Ephremides, The effect of bandwidth limitations on wireless broadcasting and multicasting, in: *Proceedings of 2000 Conference on Information Sciences and Systems*, Princeton University (March 2000) pp. TA7a-7–TA7a-12.

- [5] M. Nagy and S. Singh, Multicast scheduling algorithms in mobile networks, *Cluster Computing* 1(2) (December 1998) 177–185.
- [6] C.-C. Chiang, M. Gerla and L. Zhang, Forwarding group multicast protocol (FGMP) for multihop, mobile wireless networks, *Cluster Computing* 1(2) (December 1998) 187–196.
- [7] S.-J. Lee, W. Su, J. Hsu, M. Gerla and R. Bagrodia, A performance comparison study of ad hoc wireless multicast protocols, in: *Proceedings of IEEE INFOCOM 2000* (March 2000) pp. 565–574.
- [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) (1997) 277–290.
- [9] M.H. Ammar, G.C. Polyzos and S.K. Tripathi, eds., 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* (September 1993) 8–18.
- [11] J.E. Wieselthier, G.D. Nguyen and A. Ephremides, Energy efficiency, energy constraints and admission control in wireless networks: A case study in session multicasting, in: *Proceedings of the 38th Annual Allerton Conference on Communication, Control and Computing*, October 2000, pp. 103–112.
- [12] J.E. Wieselthier, G.D. Nguyen and A. Ephremides, Energy efficiency in energy-limited wireless networks for session-based multicasting, in: *Proceedings of the 2001 Spring Vehicular Technology Conference*, May 2001.
- [13] L. Viennot, Complexity results on election of multipoint relays in wireless networks, Report RR-3584, INRIA (December 1998).
- [14] W.-T. Chen and N.-F. Huang, The strongly connecting problem on multihop packet radio networks, *IEEE Transactions on Communications* 37(3) (1989) 293–295.
- [15] A.E.F. Clementi, P. Penna and R. Silvestri, Hardness results for the power range assignment problem in packet radio networks, in: *Proceedings 3rd International Workshop on Randomization and Approximation Techniques in Computer Science* (1999) pp. 197–208.
- [16] R.K. Ahuja, T.L. Magnanti and J.B. Orlin, *Network Flows* (Prentice Hall, Englewood Cliffs, 1993).
- [17] F.K. Hwang, D.S. Richards and P. Winter, *The Steiner Tree Problem* (North-Holland, Amsterdam, 1992).
- [18] D. Bertsekas and R. Gallager, *Data Networks* (Prentice Hall, Englewood Cliffs, 1992).



Jeffrey E. Wieselthier was born in Brooklyn, NY, on March 14, 1949. He received the S.B. degree from the Massachusetts Institute of Technology, Cambridge, in 1969, the M.S. degree from the Johns Hopkins University, Baltimore, MD, in 1971, and the Ph.D. degree from the University of Maryland, College Park, in 1979, all in electrical engineering. He was employed at the Naval Surface Warfare Center, White Oak, Silver Spring, MD, from 1969 to 1979. Since 1979 he has been with the Information Technology Division of the Naval Research Laboratory, Washington, DC. He was Technical Program Co-Chair of the Third IEEE Symposium on Computers and Communications in Athens, Greece, in 1998 and Treasurer of the 1991 IEEE International Symposium on Information Theory in Budapest, Hungary. He won the IEEE Fred W. Ellersick Award for the best unclassified paper at MILCOM 2000. He has studied a variety of communication networking problems, including multiple access and routing in spread-spectrum networks and the use of neural

networks and other approaches for network performance evaluation, optimization and control. His current interests include wireless communication networks, with an emphasis on issues relating to energy-efficient operation.

Dr. Wieselthier is a member of Eta Kappa Nu and Sigma Xi, and is a Senior Member of IEEE.

E-mail: wieselthier@itd.nrl.navy.mil



Gam D. Nguyen received the Ph.D. in electrical engineering from the University of Maryland, College Park, MD in 1990. He has been at the Naval Research Laboratory, Washington, DC, since 1991. He won the IEEE Fred W. Ellersick Award for the best unclassified paper at MILCOM 2000. His research interests include communication systems, computer communication networks, and information processing.

E-mail: nguyen@itd.nrl.navy.mil



Anthony Ephremides received his B.S. degree from the National Technical University of Athens (1967), and M.S. (1969) and Ph.D. (1971) degrees from Princeton University, all in electrical engineering. He has been at the University of Maryland since 1971, and currently holds a joint appointment as Professor in the Electrical and Computer Engineering Department and the Institute for Systems Research (ISR). He is co-founder of the NASA Center for Commercial Development of Space on Hybrid and Satellite Communications Networks established in 1991 at Maryland as an off-shoot of the ISR.

He was a Visiting Professor in 1978 at the National Technical University in Athens, Greece, and in 1979 at the EECS Department of the University of California, Berkeley, and at INRIA, France. During 1985–1986 he was on leave at MIT and ETH in Zurich, Switzerland. He was the General Chairman of the 1986 IEEE Conference on Decision and Control in Athens, Greece and of the 1991 IEEE International Symposium on Information Theory in Budapest, Hungary. He was the Technical Program Co-Chair of the IEEE INFOCOM in New York City in 1999 and of the IEEE International Symposium on Information theory in Sorrento, Italy in 2000. He has also been the Director of the Fairchild Scholars and Doctoral Fellows Program, an academic and research partnership program in Satellite Communications between Fairchild Industries and the University of Maryland. He has received several awards, including the IEEE Donald E. Fink Prize Paper Award (1992) and the Sigmobility Award of the ACM for contributions to wireless communications (1997). He has been the President of the Information Theory Society of the IEEE (1987), and has served on the Board of the IEEE (1989 and 1990). He is a Fellow of IEEE.

Dr. Ephremides' interests are in the areas of communication theory, communication systems and networks, queueing systems, signal processing, and satellite communications.

E-mail: tony@eng.umd.edu