



Algorithms for Energy-Efficient Multicasting in Static Ad Hoc Wireless Networks*

JEFFREY E. WIESELTHIER and GAM D. NGUYEN

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

ANTHONY EPHEMIDES

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

Abstract. In this paper we address the problem of multicasting in ad hoc wireless networks from the viewpoint of energy efficiency. We discuss the impact of the wireless medium on the multicasting problem and the fundamental trade-offs that arise. We propose and evaluate several algorithms for defining multicast trees for session (or connection-oriented) traffic when transceiver resources are limited. The algorithms select the relay nodes and the corresponding transmission power levels, and achieve different degrees of scalability and performance. We demonstrate that the incorporation of energy considerations into multicast algorithms can, indeed, result in improved energy efficiency.

Keywords: multicasting, energy efficiency, ad hoc network

1. Introduction

The problem of energy-efficient communications is a many-faceted one. Whereas traditional approaches to this problem have emphasized the development of improved batteries, low-power electronics, efficient coding and modulation, signal processing techniques, and antenna design, it has recently been recognized that networking techniques can also have a strong impact on the energy efficiency of such systems. A variety of networking-based approaches to energy-efficiency are possible. For example, protocols can be designed that minimize the occurrence of destructive collisions or the transmission of unnecessary (e.g., redundant) information. Also, asymmetrical protocols can be designed to reduce the energy expenditure of disadvantaged users. The issue of minimum power topology for a stationary ad hoc network was addressed in [1]. Additionally, energy-efficient routing schemes [2] can be developed, in some cases in conjunction with adaptive coding/modulation schemes that incorporate knowledge of link characteristics into network-level decisions [3].

Our approach to energy-efficient communication departs from the traditional layered structure in that we jointly address the issues of transmitted power levels (and, hence, network connectivity, a Physical layer function) and multicast tree formation (a routing function, associated with the Network layer). We argue that such joint decisions on connectivity and routing can result in improvement in energy efficiency, as compared to a rigid layered structure that makes these decisions independently.¹ Here we consider only the energy used for transmission, neglecting for the present the

energy associated with reception and signal processing; the joint study of all forms of energy expenditure and the associated trade-offs are not considered here.

It is clear that turning “on” and “off” the transmitter and/or the receiver and choosing the transmission power prescribes a schedule for the “amortization” of the stored battery energy.² Another level of complexity is added to the control and scheduling of the network by the fact that, while in the “off” state, a node’s transceiver cannot participate in a distributed control algorithm. In fact, this complication has potentially serious consequences on the delay performance and quality of service requirements of the network.

We have chosen the problem of multicasting (one-to-many communication) as the focus of our energy-efficient networking studies. 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 additional 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. 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) versus reaching fewer nodes in that single hop by using lower power (but at a lower interference cost). By contrast, the unicast (one-to-one) communication problem (although challenging in its own right) is characterized by

² Recent studies have shown that the total energy capacity of a battery is not fixed, but rather depends on the way in which the battery energy is used. For example, more energy can be obtained from a battery by means of pulsed, rather than continuous, operation [6,7].

* This work was supported by the Office of Naval Research.

¹ Other studies that have also made this observation include [4,5].

a minimum-energy solution in which multihop relaying at low power is generally favored over higher-power transmissions because of the nonlinear attenuation properties of radio signals.³ Such generalizations cannot be made for the multicasting problem, however. In some of our examples it is better to transmit at low power, whereas in other cases high power is better.

Few studies have addressed the crucial problem of multicasting in wireless networks. For example, the problem of multicast scheduling in cellular mobile networks was studied in [8], and a forwarding multicast protocol for noncellular networks was studied in [9]. Virtually all multicasting studies have been limited to the case of stationary networks that are not wireless (e.g., [10–12]). Ad hoc networks lack a fixed cellular infrastructure, and thus cannot effectively use multicast algorithms that are based on the availability of fixed topologies.

Our focus is on the source-initiated multicasting of “session” or connection-oriented traffic. To assess the complex trade-offs one at a time, we make two simplifying assumptions, namely that there is ample bandwidth (i.e., unlimited number of frequencies, time slots, or orthogonal CDMA codes, so that contention for the channel is not an issue) and no mobility.

To implement any network-layer function in a wireless environment, we need an underlying medium-access-control (MAC) protocol that ensures the collision-free coordination of transmissions by neighboring nodes. This is an important component of ad hoc network operation that is usually addressed separately. Although ideally it should be studied in conjunction with the higher-layer issues, the complexity of a joint consideration is far too great and, in addition, obscures the principal components of the energy trade-offs in wireless multicasting. Since the main thrust of the paper is to contribute toward the understanding of these trade-offs, it is prudent to separate them from the underlying access control considerations. The availability of a suitable MAC protocol is a reasonable assumption, especially under our assumption of ample frequency resources.

Although we assume that the nodes are stationary, 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. We neglect the overhead associated with the control messages exchanged during tree setup. Their energy expenditure is negligible for the case of session traffic because typical session duration is much longer than the time required for tree setup.

After a discussion of the basic issues of multicasting in wireless networks, we consider the problem of broadcasting (i.e., transmission to all nodes in the network), in which the goal is the determination of the minimum-energy broadcast

tree. We then return to the multicasting problem, in which we model the network’s resources by means of “node capacity” (namely, by assuming finite numbers of transceivers at each node). Our performance results demonstrate that the incorporation of energy considerations into the multicast algorithms can, indeed, result in energy saving.

2. Architectural issues in all-wireless networks

The ad hoc wireless networks studied here are quite different from the cellular systems that have been developed in the commercial domain. Cellular systems have fixed base stations, which communicate among themselves using dedicated non-wireless lines; thus, the only multicast problems that are new in those systems involve tracking the mobile users. However, in ad hoc wireless networks it is possible to establish a wireless link between any pair of nodes, provided that each has a transceiver available for this purpose and that the signal-to-noise ratio at the receiving node is sufficiently high. Thus, unlike the case of wired networks, the set of network links and their capacities are not determined a priori, but depend on factors such as distance between nodes, transmitted power, error-control schemes, other-user interference, and background noise. Thus, even when the physical locations of the nodes are fixed, many of the factors that affect network topology are influenced by the actions of the network nodes (either those directly participating in a link, or those contributing to the interference that affects a link). Furthermore, in ad hoc networks no distinction can be made between uplink and downlink traffic, thus greatly complicating the interference environment.

In this paper, we focus on wireless networks in which the node locations are fixed, and the channel conditions unchanging. The wireless channel is distinguished by its *broadcast* nature; when omnidirectional antennas are used, every transmission by a node can be received by all nodes that lie within its communication range. Consequently, if the multicast group membership includes multiple nodes in the immediate communication vicinity of the transmitting node, a single transmission suffices for reaching all these receivers. Hence, there is an incentive to perform a multicast by using high transmitter power. Of course, doing so results in interference with more nodes than if reduced power were used. Thus, there is a trade-off between the long “reach” of a single transmission and the interference (and/or delay) it creates.

Another undesirable impact of the use of high transmitter power is that it results in increased energy usage. Since the propagation loss varies nonlinearly with distance (at somewhere between the second and fourth power), in unicast applications it is best (from the perspective of transmission energy consumption) to transmit at the lowest possible power level, even though doing so requires multiple hops to reach the destination. However, in multicast applications it is not prudent to draw such conclusions because the use of higher power may permit simultaneous connectivity to a

³ However, even in this case the overhead traffic and the increased interference at the local level may reduce or even reverse the savings achieved by the use of multiple short hops.

sufficiently large number of nodes, so that the total energy required to reach all members of the multicast group may be actually reduced. Furthermore, even for unicast applications, the use of lower power (and, hence, multiple hops) necessitates the complex coordination of more signals and, therefore, may actually result in higher total energy expenditure.

3. Multicasting in wireless networks

To date, virtually all of the research and development work on multicasting has centered on tethered, point-to-point (typically high speed) networks and on methods of bandwidth-efficient maintenance of multicast group addresses and routing trees. There are two basic approaches to multicast tree construction. The first is the use of Source-Based Trees (SBT), which are rooted at the sender and which are designed to minimize the number of transmissions needed to reach all of the members of the multicast group. The second is the use of Core-Based Trees (CBT) [13], under which the same tree is used for all communication within the multicast group. The Sparse Mode of the Protocol Independent Multicasting (PIM) protocol [14] can be used with either SBTs or CBTs, whereas the PIM Dense Mode is based on the use of SBTs.

As pointed out earlier, the characteristics of the wireless medium and the ad hoc network architecture may render multicasting techniques developed for nonwireless applications inappropriate, or at least unable to provide acceptable performance in some scenarios. There are numerous and complex issues that must be addressed in wireless multicasting (e.g., see [4,15]). In this paper we focus on a single aspect of the multicasting problem, namely the incorporation of energy considerations into the construction of multicast trees and the choice of transmission power levels.

3.1. A model for wireless multicast

We consider source-initiated, circuit-switched, multicast sessions. The network consists of N nodes, which are randomly distributed over a specified region. 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 needed as relays to provide connectivity to all members of the multicast group. Also, the use of relays (even when not absolutely necessary to provide connectivity) may result in lower 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*.

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

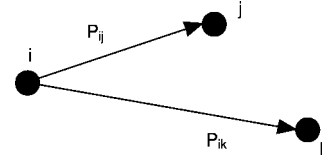


Figure 1. The “wireless multicast advantage”: $P_{i,(j,k)} = \max\{P_{ij}, P_{ik}\}$.

A constant bit rate (CBR) traffic model is assumed; thus, one transceiver is allocated to support each active multicast session at every node participating in the multicast tree throughout the duration of the session. Each node has T transceivers, and can therefore participate in up to T multicast sessions simultaneously. Since, as noted earlier, we assume in this paper that ample bandwidth is available, the only hard constraints we consider are the number of transceivers and the maximum permitted transmitter power p_{\max} .

We assume that the received signal power varies as $r^{-\alpha}$, where r is the range and α is a parameter that typically takes on a value between 2 and 4, depending on the characteristics of the communication medium. Based on this model the transmitted power required to support a link between two nodes separated by range r is proportional to r^α . Without loss of generality, we set the normalizing constant equal to 1, resulting in:

$$\begin{aligned} p_{ij} &= \text{power needed to support link} \\ &\quad \text{between nodes } i \text{ and } j \\ &= r^\alpha, \end{aligned}$$

where r is the distance between nodes i and j . If the maximum permitted transmitter power p_{\max} is sufficiently large, the nodes will be able to transmit at sufficiently high power so that the network is fully connected.

We assume an isotropic medium and omnidirectional antennas; thus, all nodes within communication range of a transmitting node can receive its transmission. It is important to note how the broadcast property of wireless communication can be exploited in multicast applications. Consider the example shown in figure 1, in which a subset of the multicast tree involves node i , which is transmitting 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 , based on our assumption of omnidirectional antennas. This situation is fundamentally different from wired applications, in which the cost of node i 's transmission to nodes j and k would be the sum of the costs to the individual nodes.⁴ The ability to exploit this property of wireless communication, which we refer to as the “wireless multicast advantage”, makes mul-

⁴ 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.

roadcasting an excellent setting in which to study the potential benefits of energy-efficient protocols.

4. Construction of minimum-energy broadcast trees

We first address the problem of constructing the minimum-energy broadcast tree for each newly arriving multicast request. Doing so involves a choice of transmitter-power levels and relay nodes. In addition to our assumption throughout this paper that ample bandwidth is available, we assume in this section that each node has a sufficient number of transceivers to accommodate all service requests. An insufficient quantity of either of these resources can result in the construction of trees that do not reach all destinations, use more than the minimum energy (because only suboptimal trees can be constructed), or both.

We start with simple examples with two, and then three, destinations, and discuss how our results can be extended to larger examples by means of a recursive technique. Our examples in this section are based on the *broadcasting* problem, in which all nodes in the network (other than, of course, the source) are destinations. In section 5 we return to the problem of multicasting in which only a subset of the network nodes need to be reached.

In wired networks, the broadcasting problem can be formulated as the well-known Minimum 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, discussed in section 3, 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; the generalization of the MST problem to wireless networks is a possible approach, although we do not pursue it further here. In this paper we introduce one heuristic for the formation of low-energy broadcast trees, which takes into account the wireless multicast advantage. We use low-energy broadcast trees (including versions based on both link-based and node-based versions) as the basis for some of our heuristics for the construction of suboptimal multicast trees in wireless networks.

4.1. Minimum-energy broadcasting: Two destinations

We consider a source node S (located at the origin) and two destination nodes D_1 (located along the x -axis, without loss of generality) and D_2 , as shown in figure 2. The topology is specified by the coordinates of D_1 and D_2 , which determine the angle θ . The distance between S and D_1 is r_1 , the distance between S and D_2 is r_2 , and the distance between D_1 and D_2 is r_{12} . It is assumed, without loss of generality, that $r_2 > r_1$. We define:

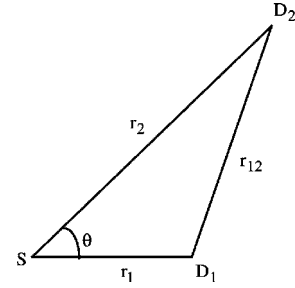


Figure 2. Broadcasting to two destinations.

- $P_{S1} = r_1^\alpha =$ power needed to support a link between S and D_1 ,
- $P_{S2} = r_2^\alpha =$ power needed to support a link between S and D_2 ,
- $P_{12} = r_{12}^\alpha =$ power needed to support a link between D_1 and D_2 .

In this simple example, there are two alternative strategies:

- (a) S transmits using P_{S2} : both D_1 and D_2 are reached,
- (b) S transmits using P_{S1} : only D_1 is reached. D_1 then transmits to D_2 with power P_{12} , resulting in a total power of $P_{S1} + P_{12}$.

We would like to choose the alternative that results in the smaller value of total power consumption. For the case of propagation that follows a $1/r^2$ law, it is very simple to derive the following result from simple geometrical considerations:

- use strategy (a) if $r_1 > r_2 \cos \theta$,
- use strategy (b) otherwise.

For the general case of propagation behaving as $1/r^\alpha$, algebraic manipulation results in the following:

- use strategy (a) if

$$x^\alpha - 1 < (1 + x^2 - 2x \cos \theta)^{\alpha/2}, \quad (1)$$

where $x = r_2/r_1$,

- use strategy (b) otherwise.

This result is shown graphically in figure 3. For example, in the region above the curve (for each particular value of α) it is best to use strategy (a). It is of special interest to note that for $\alpha \geq 3$ (which is characteristic of many realistic environments) the boundary separating these regions is quite steep; therefore a simple heuristic that uses strategy (a) whenever $\theta \geq 90^\circ$ and strategy (b) otherwise should be expected to provide nearly optimal performance. Thus, the incentive to use the shortest available links increases as α increases.⁵

⁵ We acknowledge that, in practical applications, the locations of the nodes generally will not be known precisely. Also, the propagation characteristics are often difficult to characterize. Nevertheless, heuristics such as the one described here, which can depend on estimates of these quantities, are expected to provide insight into the development of good (although suboptimal) broadcast trees.

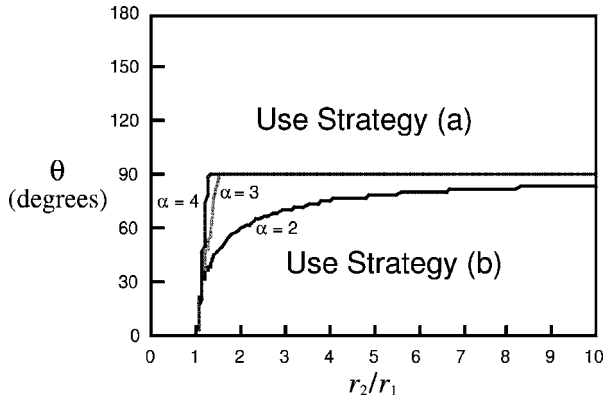


Figure 3. Transmission strategies for minimum-energy broadcasting to two destinations ($r_2 \geq r_1$).

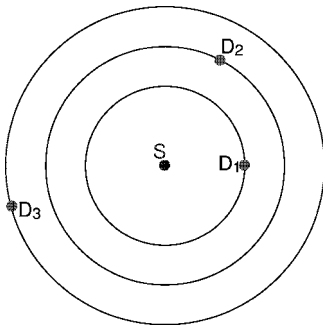


Figure 4. Broadcasting to three destinations.

4.2. Minimum-energy broadcasting: Three destinations

The minimum-energy broadcasting problem becomes more interesting as the number of destinations increases. In such cases, it is harder to make generalizations about the desirability of using the shortest available links because the use of a higher power transmission can often result in the ability to reach several nodes with a single transmission, thereby resulting in lower overall power in the complete tree. Figure 4 shows the case of broadcasting to three destinations.

We enumerate the alternative strategies:

- (a) S transmits using P_{S3} : all three destinations are reached.
- (b) S transmits using P_{S2} : destinations D_1 and D_2 are reached by this transmission. One of these nodes must then transmit to D_3 . The two alternatives are:
 - D_1 transmits to D_3 : total power = $P_{S2} + P_{13}$,
 - D_2 transmits to D_3 : total power = $P_{S2} + P_{23}$.
- (c) S transmits using P_{S1} : only D_1 is reached by this transmission. D_1 must then form a tree to nodes D_2 and D_3 . The three alternatives are:
 - D_1 transmits with sufficient power to reach D_2 and D_3 : total power = $P_{S1} + \max\{P_{12}, P_{13}\}$,
 - D_1 transmits to D_2 , which transmits to D_3 : total power = $P_{S1} + P_{12} + P_{23}$,
 - D_1 transmits to D_3 , which transmits to D_2 : total power = $P_{S1} + P_{13} + P_{32}$.

As in the case of two destinations, the strategy that minimizes total power is chosen.

4.3. A recursive formulation

The number of alternative strategies increases rapidly as the number of destinations increases. However, the effects of complexity can be mitigated somewhat by means of a recursive formulation. For example, let us consider alternative (c) in section 4.2. If the source transmits using power P_{S1} , it effectively delegates to D_1 the responsibility of reaching D_2 and D_3 . This is simply the problem of broadcasting to two destinations, which is precisely the problem solved in section 4.1. One can thus remap the origin to the location of D_1 , and use the already obtained solution to the problem of broadcasting to two destinations in the evaluation of strategies for the three-destination example. In general, the solution for N_D destinations can be expressed in terms of the solutions for various subsets of the solutions for a smaller number of destinations. Unfortunately, the complexity of this formulation is high, making it impractical except for small networks. One way to roughly estimate complexity is to evaluate the number of times that the solution for the two-destination problem is called during the course of the algorithm. For the case of four destinations, it is called three times, which is certainly easy to handle. However, the number of calls to this subproblem increases rapidly as N_D increases; e.g., for $N_D = 10$, more than 51,000 calls are needed, and for $N_D = 13$ more than 14 million calls are needed. Nevertheless, this approach may serve as the basis for a suboptimal heuristic that provides less than an exhaustive search of all possible trees.

5. A multicasting problem

We now address the problem of determining an appropriate multicast tree for each arriving multicast call 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 and the commitment of the needed transceiver resources throughout the duration of the multicast session. If there is no tree that can reach any of the desired destinations (because the needed resources are blocked), the call is rejected. If there are trees that can reach only a portion of the destination set, they are considered. In some cases, where one or more of the intended destinations is costly to reach, the “best” multicast tree may include only a subset of the reachable destinations.

In a wired network, the determination of the minimum-cost multicast tree is equivalent to the Steiner tree problem, which is NP-complete. By contrast, the MST problem (in wired networks) is polynomial in complexity. It would be

of great interest to formulate and develop heuristics for the Steiner tree problem in a node-based context.⁶

5.1. Admission-control policies

Recall that the establishment of a multicast session requires the allocation of a transceiver at every node participating in the session (source, relays, and destinations) throughout the duration of the session. A destination can be *reached* if there exists a path from the source to it, and provided that a transceiver is available (i.e., not already supporting another session) at each node along the path. There are two basic aspects to the admission-control problem, i.e., whether or not to establish a multicast session for a particular multicast request, and (assuming a session is, in fact, established) which of the desired destinations to include in the multicast tree. Most of the results presented in this paper are based on the use of the “admit-all” admission-control policy, under which all multicast requests are accepted as long as one or more of the intended destinations can be reached; furthermore, under such schemes paths are established to all reachable destinations (i.e., the potential energy savings from dropping a subset of the destination nodes is not an option). In some cases, however, we do consider admission-control policies in which “expensive” destinations are not included in the tree.

5.2. Performance metrics

We define:

- n_i : the number of intended destinations by i th multicast arrival,
- m_i : the number of destinations reached by i th multicast session,
- d_i : duration of i th multicast session (assumed exponentially distributed with mean = 1),
- p_i : sum of the transmitter powers used by all nodes in i th multicast session,
- E_i : total energy used by i th multicast session = $p_i d_i$,
- v_i : *multicast value* of i th multicast session.

Since the quantity of information delivered is proportional to the duration of a session and to the number of destinations reached, we define the multicast value of session i to be

$$v_i = m_i d_i. \quad (2)$$

A variety of performance measures can be defined for the multicasting problem, including the following.

⁶ Although [16] addressed the multicasting problem with a goal toward reaching efficient and near-minimum-cost algorithms for wireless networks, their approach was link-based, rather than node-based, and hence, does not take into consideration the wireless multicast advantage.

5.2.1. Average (per call) multicast value per unit energy

The average (per call) multicast value per unit energy V_E , observed over an interval with X multicast requests, is⁷

$$V_E = \frac{1}{X} \sum_{i=1}^X \frac{v_i}{E_i} = \frac{1}{X} \sum_{i=1}^X \frac{m_i d_i}{p_i d_i} = \frac{1}{X} \sum_{i=1}^X \frac{m_i}{p_i}. \quad (3)$$

We observed in [17] that use of this metric alone tends to favor the hoarding of energy because this metric can often be maximized by transmitting to only those destinations that can be reached with very little energy consumption. Thus, only a small fraction of the desired destinations would typically be reached in multihop networks when this metric is used as the basis for an admission-control policy.

5.2.2. Multicast efficiency

Also of interest is the *multicast efficiency* of the i th multicast session, which can be defined as the fraction of desired destinations of the i th multicast service request that are actually reached. Then, the overall multicast efficiency can be defined as

$$e = \frac{1}{X} \sum_{i=1}^X \left(\frac{m_i}{n_i} \right). \quad (4)$$

This metric is maximized when all possible destinations are reached, without regard to the energy required to do so.

5.2.3. The “Yardstick” metric

To take into consideration both of the criteria discussed above, namely reaching many destinations per unit energy and reaching a large fraction of the number of desired destinations, we define a *local yardstick* measure of multicast performance to be

$$y_i = \left(\frac{m_i}{p_i} \right) \left(\frac{m_i}{n_i} \right). \quad (5)$$

Our *global yardstick* Y is the average value of this quantity over the observation interval:

$$Y = \frac{1}{X} \sum_{i=1}^X y_i = \frac{1}{X} \sum_{i=1}^X \left(\frac{m_i}{p_i} \right) \left(\frac{m_i}{n_i} \right). \quad (6)$$

This metric is the primary performance metric we use to evaluate multicasting algorithms in this paper. Although maximization of y_i ($i = 1, \dots, X$) does not guarantee maximization of Y , the greedy approach of maximizing y_i is expected to perform well. In addition, we also consider the fraction of blocked calls as another metric.

5.2.4. Blocking probability

We define k_X to be the number of multicast sessions that are completely blocked during an interval with X multicast requests. A session will be completely blocked if resources are not available to reach any destinations, or alternatively if the

⁷ Totally unsuccessful multicast arrivals, in which no destinations are reached, do not contribute to either throughput or energy expenditure.

admission-control policy decides that it is not cost effective to form paths to any destinations. The *blocking probability* is defined as

$$P_B = \frac{k_X}{X}. \quad (7)$$

5.3. “Local” cost metrics

The problem of finding the multicast tree that maximizes the local yardstick for each new multicast request is highly complex, and not feasible, except for small examples. Therefore, we have found it necessary to take the heuristic 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 and/or link- or node-based) cost metrics.

5.3.1. Link-based costs

Consider link ij , which is established between nodes i and j . We define D_{ij} to be cost associated with link ij . In this paper we define the cost of a link as the power level needed to support it, provided that at least one transceiver is available at both nodes.⁸ If either node has no available transceiver (i.e., all are already committed to currently active sessions), the cost of the link is infinite. If the power required to support the link between nodes i and j is P_{ij} ,

$$D_{ij} = \begin{cases} P_{ij}, & \text{if there is at least one transceiver} \\ & \text{available at nodes } i \text{ and } j, \\ \infty, & \text{otherwise.} \end{cases} \quad (8)$$

The total cost of the tree can be defined (especially in wired networks) as the sum of the costs of all links in the tree.

The total cost to implement the multicast tree can be less than the sum of the link costs, however. It is sufficient for a node to transmit only once to reach all of its neighbors. The wireless multicast advantage applies here, since the total power required to reach several neighbors is the maximum power to reach any of them individually. However, the tree is selected without the ability to exploit the wireless multicast advantage in the choice of transmitting nodes.

5.3.2. Node-based costs

Since (under our assumptions of omnidirectional antennas, no interference, and perfect coordination of transmissions and receptions) a node’s transmission can be received by all of its neighbors, it would be best to design a tree that exploits the wireless multicast advantage. Tree formation would consist of a 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 the one that reaches all reachable nodes with minimum total

⁸ In [18] we discuss link cost functions that incorporate congestion. In this paper, however, we focus on the use of power (rather than congestion) as the cost metric to facilitate the comparison of several distinct algorithms without having to address subtle issues relating to the definition of alternative cost functions.

power. We noted in section 4 that 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.

6. Alternative algorithms

In this section we discuss several of the multicasting algorithms we have studied; full descriptions are available in [19]. In this paper we define the notion of the cost associated with the support of a multicast tree to be the power required to reach all destination nodes; thus, it is the sum of the powers at all transmitting nodes. This is a metric that is used as the basis of some of our algorithms. However, performance is always judged by the “yardstick” metric, the multicast efficiency, and the blocking probability.

Each transmission by a node is characterized by its transmitter power level, as well as a designation of which (possibly several) of the nodes receiving this transmission are to forward it toward which of the ultimate destination nodes. In all cases, we use greedy algorithms, which attempt to optimize performance on a “local” (call-by-call) basis.

When the number of transceivers at each node (T) is finite, it may not be possible to establish minimum-energy trees (even on a local basis) because of the lack of resources (transceivers) at one or more nodes. In this case, the greedy algorithms discussed here are applied to the subset of nodes that have non-zero *residual capacity*.⁹

6.1. A unicast-based multicast algorithm

A straightforward approach is the use of multicast trees that consist of the superposition of the best unicast paths to each individual destination (see, e.g., [20]). It is assumed that an underlying unicast algorithm (such as the Bellman–Ford or Dijkstra algorithm) provides “minimum-distance” paths from each node to every destination. However, the minimization of unicast costs does not necessarily minimize the cost of the multicast tree, as illustrated in figure 5, which shows a source and two destinations. Figure 5(a) shows the best unicast paths that reach the two destinations, and figure 5(b) shows the best multicast tree. The use of the best unicast paths fails to discover the path that reaches a neighbor of both destinations over a common path, thereby resulting in lower overall cost. Also, the use of the best unicast paths fails to incorporate the “multicast advantage,” which was discussed in section 3. Therefore, the trees obtained based on unicast information are not expected to provide optimal multicast performance. Nevertheless, they do perform reasonably well, and with considerably reduced complexity as compared to the calculation of truly optimal multicast trees.

⁹ The residual capacity at node j is the number of transceivers at node j that are not currently supporting traffic, and hence, are available to support new sessions.

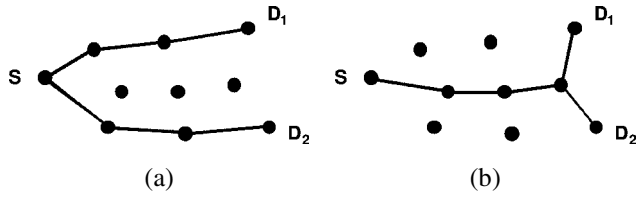


Figure 5. Unicast-based vs multicast-based trees: (a) best unicast paths, (b) best multicast tree.

Summarizing the above, we have:

Algorithm 1 (Least unicast cost). A minimum-cost path to each reachable destination is established. The multicast tree consists of the superposition of the unicast paths. Paths to all reachable destinations are established, regardless of the cost required to do so. This algorithm is scalable.

6.2. Algorithms based on pruning MSTs

One approach we have taken in the development of heuristics for multicasting is the pruning of broadcast spanning trees. 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.

We noted earlier that, for the case of wired networks, the determination of minimum-cost broadcast (spanning) trees is considerably easier than the determination of minimum-cost multicast trees. Nevertheless, the determination of minimum-cost broadcast trees for wireless networks remains a difficult problem for which no scalable solutions appear to be available at this time. In small network examples we have determined minimum-energy spanning trees by using the recursive technique of section 4.3; in moderate to large networks it is necessary to use heuristics. In this subsection we discuss the main features of three algorithms that are based on the technique of pruning. Further details are provided in [19].

Algorithm 2 (Pruned link-based MST). This algorithm is based on the use of the standard MST formulation in which a link cost is associated with each pair of nodes (i.e., the power to sustain the link); thus, the “wireless multicast advantage” is ignored in the construction of the MST. Since the MST problem is of polynomial complexity, this algorithm is scalable. To obtain the multicast tree, the MST is pruned by eliminating all transmissions that are not needed to reach the members of the multicast group. Once the MST is constructed in this manner, the evaluation of its cost (i.e., the total power needed to sustain the broadcast tree) does take into consideration the wireless multicast advantage.

Algorithm 3 (Pruned node-based MST). This algorithm requires the determination of the minimum-energy spanning tree that is rooted at the Source node. Unlike algorithm 2, the wireless multicast advantage is taken into consideration in the determination of the power needed to sustain the tree.

The recursive algorithm of section 4.3 can be used to determine the MST. Thus, this method is not scalable. Once the MST has been determined in this manner, it is pruned as in algorithm 2.

Algorithm 4 (Pruned node-based spanning tree). A heuristic is used to determine a suboptimal spanning tree.¹⁰ Once the spanning tree has been determined in this manner, it is pruned as in algorithm 2.

Construction of the spanning tree begins at the Source node. Its transmission power is chosen to maximize the following “ n/p ” metric:

$$\frac{n}{p} = \frac{\text{Number of “new” destinations reached}}{\text{Total power required to reach them}}. \quad (9)$$

At the next stage, each of the nodes that has been “covered” (i.e., the Source node plus all nodes within its communication range based on the calculation in the first stage) evaluates the n/p metric for all possible sets of neighbors (however, in computing this metric, only “new” nodes, i.e., nodes not previously covered, are included in the number of destinations). Note that it is possible to increase the transmission power that was assigned to a node in an earlier stage. This procedure is repeated until all nodes are covered. Full details are provided in [19].

6.3. Additional algorithms with high complexity

The following algorithms require an exhaustive search, and are thus not scalable. Nevertheless, they provide a useful benchmark that permits us to evaluate the performance of the other algorithms for small network examples.

Algorithm 5 (Least multicast cost). As in algorithm 1, paths to all reachable destinations are established, regardless of the cost required to do so. An exhaustive search of all multicast trees that reach all reachable destinations is performed. The tree with the lowest cost is chosen.

Algorithm 6 (Maximum local yardstick). The local yardstick function y_i is computed for each arriving multicast request i . Multicast trees are formed to all subsets of intended destinations. The tree that results in the maximum value of y_i is chosen. This tree does not necessarily include all reachable destinations.

7. Performance results

We have simulated the performance of the six algorithms for the 8-node network shown in figure 6. The connectivities shown are based on a maximum permitted transmitter power

¹⁰ In small examples, this heuristic typically provides a lower-energy broadcast tree than that produced by the link-based method of algorithm 2, but does not provide the true minimum that can be obtained by means of the recursive scheme.

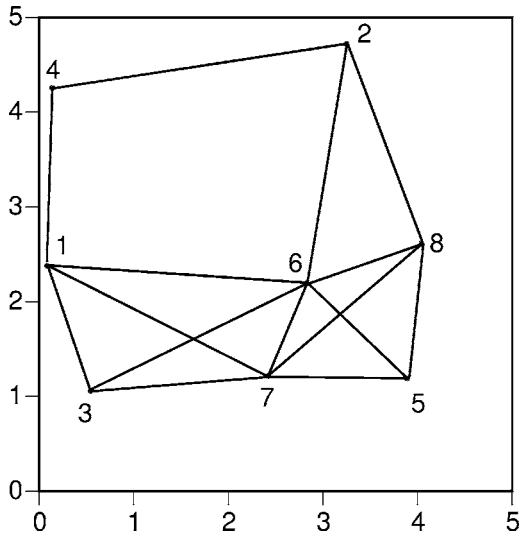


Figure 6. An example eight-node network ($p_{\max} = 10$ when $\alpha = 2$; or $p_{\max} = 100$ when $\alpha = 4$).

value of $p_{\max} = 10$ when $\alpha = 2$, which result in a maximum communication range of 3.16 (where the overall dimensions of the region are 5×5). A node's transmitter power (r^2) depends on the distance (r) to the farthest neighbor to which it is transmitting. The same connectivities apply to the case of $p_{\max} = 100$ and $\alpha = 4$, in which case the power used is r^4 .

In our simulations, multicast requests arrive randomly; interarrival times are exponentially distributed with rate λ . Service durations are exponentially distributed with mean 1 (i.e., $\mu = 1$). The multicast group, which must have at least two nodes (the source plus at least one destination), is chosen randomly at each arrival instant as follows. In a network with N nodes, one of the $2^N - N - 1$ subsets with at least two members is chosen with probability $(2^N - N - 1)^{-1}$. Then, a uniform distribution is used to choose one of the members of the multicast group to be the source. Each simulation run consists of $X = 1,000$ multicast session requests. We have run our simulations for numerous values of arrival rate λ .

For all six algorithms, figure 7 shows the global yardstick metric Y as a function of arrival rate λ for $T = 4$ transceivers at each node. The ordering of the algorithms in the legend of the figure is based on their relative performance (at low to moderate traffic loads); e.g., the best performance is provided by algorithm 6 (the top curve), the second best by algorithm 5 (second from the top) and the worst by algorithm 2 (the bottom curve). It is not surprising that the best performance is obtained by algorithm 6, which was designed specifically to maximize the local yardstick. The next best performance is obtained by algorithm 5, which (like algorithm 6) is based on an exhaustive search of all possible multicast trees. Thus, the two most highly complex algorithms provide the best performance. Although these two algorithms are too complex for practical applications, they are being studied because they can provide a benchmark of the performance that is achievable through appropriate choice of transmitter power levels and multicast trees.

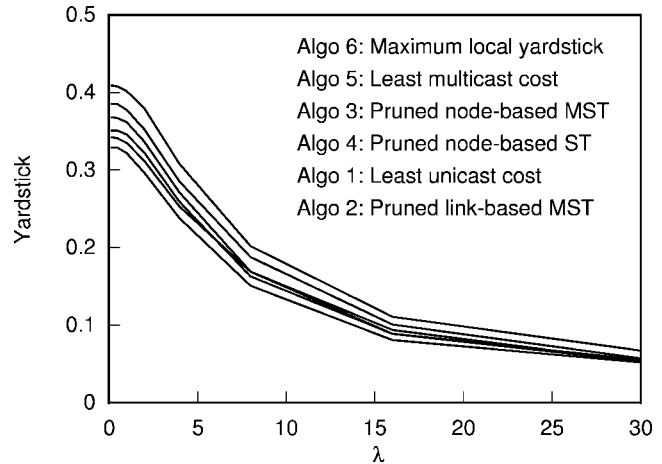


Figure 7. Global yardstick Y versus λ for eight-node network ($T = 4$, $\alpha = 2$).

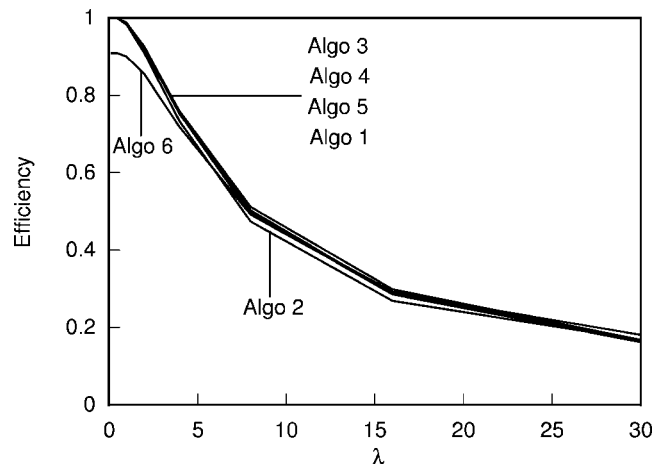


Figure 8. Efficiency e versus λ for eight-node network ($T = 4$, $\alpha = 2$).

Three of the four other algorithms are scalable. It is interesting to compare the performance of algorithm 6 with that of algorithm 1 (the first algorithm we studied and one of the simplest to implement). The fact that algorithm 6 provides approximately 19% better yardstick performance than algorithm 1 suggests that improvement can, in fact, be obtained through the exploitation of wireless networking properties, i.e., the choice of transmitter powers and relay nodes. On the other hand, the fact that simple algorithms can provide relatively good performance and the relatively small differences in performance among algorithms 1–4 indicates a high degree of robustness in that a variety of well-motivated algorithms can provide similar, and possibly acceptable performance.

Figure 8 shows the multicast efficiency e as a function of λ . At low traffic levels, there is little difference in performance among algorithms 1–5; however, algorithm 6 provides considerably lower values of e (it is the only algorithm for which e is not very close to 1 for very low values of λ). The low multicast efficiency provided by algorithm 6 results from the fact that it, unlike the other algorithms, does not provide paths to costly destinations. Algorithm 2 provides

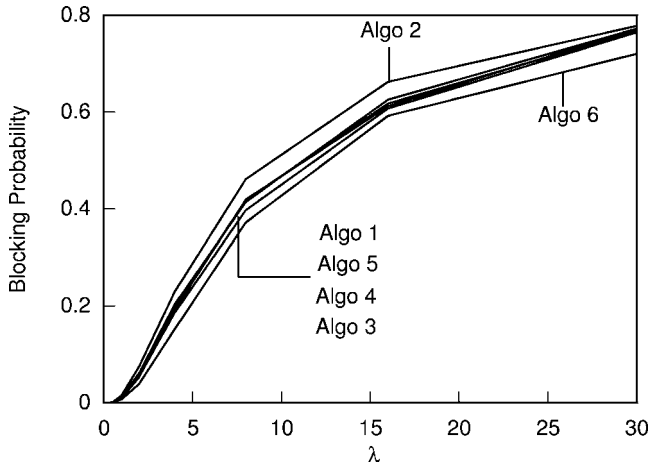


Figure 9. Blocking probability P_B versus λ for eight-node network ($T = 4$, $\alpha = 2$).

the lowest values of e for moderate through high values of λ , even though it provides high values of e at low values of λ . Small differences in performance should not be considered to be significant. These results, which are representative of many examples we have studied, illustrate the trade-off between the conflicting goals of providing high values of Y and high values of e .

Figure 9 shows the blocking probability P_B (the probability that no destinations are reached) as a function of λ . (The ordering of the algorithms in the legend is from highest P_B to lowest.) Algorithm 6 provides the best performance (lowest blocking probability) based on this metric. This behavior can be explained by the fact that since algorithm 6 does not provide paths to costly destinations, fewer resources are committed to support ongoing calls than with the other algorithms (for a given traffic level λ). Thus, more resources tend to be available, and fewer calls are totally blocked. Algorithm 2 provides the highest values of P_B . Curves for the other algorithms are clustered relatively close to each other.

Table 1 shows our performance results for algorithm 6 for the eight-node network with $\alpha = 2$. The three entries in each cell are global yardstick Y , efficiency e , and blocking probability P_B (respectively, from top to bottom). The multicast session arrival rate varies from $\lambda = 0.125$ to 64, and the number of transceivers at each node varies from $T = 1$ to 16. Thus, congestion is greatest near the lower left corner (where there are few transceivers and high arrival rate), and least near the upper right corner (where there are many transceivers and low arrival rate). For example, in the lower left cell of table 1 ($\lambda = 64$ and $T = 1$), only a fraction $e = 0.018$ of the intended destinations are reached, and $P_B = 95.6\%$ of all multicast requests are blocked completely. At low congestion levels, e approaches 1 and P_B approaches 0.

Figures 10–12 show Y , e , and P_B for the same network, but with $\alpha = 4$ and $p_{\max} = 100$. Qualitatively, the plots are similar to those for $\alpha = 2$, in the sense that algorithm 6 provides the best performance on the basis of metrics Y and P_B and the worst performance based on e . However, the difference in performance between algorithm 6 and the others

Table 1
Performance results for maximum local yardstick algorithm (algorithm 6), $\alpha = 2$.

λ	$T = 1$	2	4	8	16
0.125	0.365	0.405	0.409	0.409	0.409
	0.818	0.903	0.909	0.909	0.909
	0.086	0.007	0.000	0.000	0.000
0.25	0.332	0.400	0.409	0.409	0.409
	0.748	0.892	0.909	0.909	0.909
	0.155	0.017	0.000	0.000	0.000
0.5	0.278	0.380	0.408	0.409	0.409
	0.638	0.853	0.909	0.909	0.909
	0.262	0.052	0.001	0.000	0.000
1	0.216	0.338	0.402	0.409	0.409
	0.503	0.770	0.900	0.909	0.909
	0.403	0.130	0.007	0.000	0.000
2	0.152	0.267	0.380	0.409	0.409
	0.354	0.613	0.856	0.909	0.909
	0.549	0.280	0.039	0.000	0.000
4	0.090	0.176	0.308	0.403	0.409
	0.222	0.429	0.718	0.901	0.909
	0.676	0.447	0.152	0.006	0.000
8	0.052	0.098	0.202	0.341	0.408
	0.123	0.250	0.492	0.788	0.909
	0.786	0.643	0.372	0.095	0.000
16	0.028	0.054	0.111	0.226	0.380
	0.069	0.141	0.289	0.564	0.865
	0.866	0.768	0.592	0.297	0.031
32	0.018	0.027	0.061	0.116	0.244
	0.039	0.068	0.166	0.307	0.592
	0.924	0.878	0.738	0.545	0.254
64	0.009	0.015	0.028	0.056	0.124
	0.018	0.036	0.072	0.149	0.336
	0.956	0.924	0.863	0.757	0.523

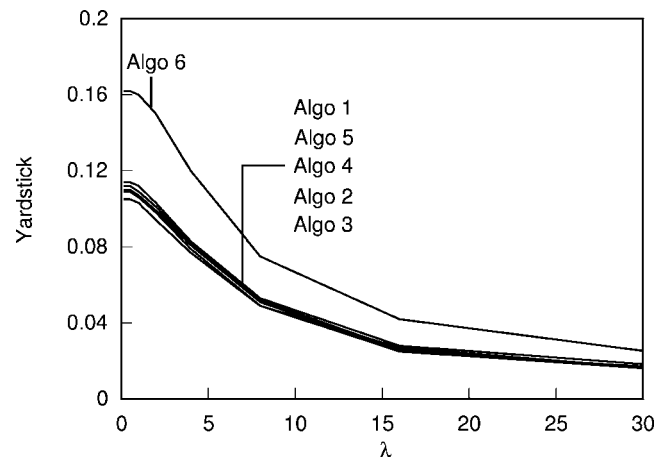


Figure 10. Global yardstick Y versus λ for eight-node network ($T = 4$, $\alpha = 4$).

is much greater for $\alpha = 4$. As α increases, the incentive to use the shortest possible links increases. Also, as α increases, the cost of including distant destinations in the tree increases rapidly. Thus, there is an incentive to exclude such costly destinations from the multicast tree, which is an option only for algorithm 6, but not for the other algorithms we

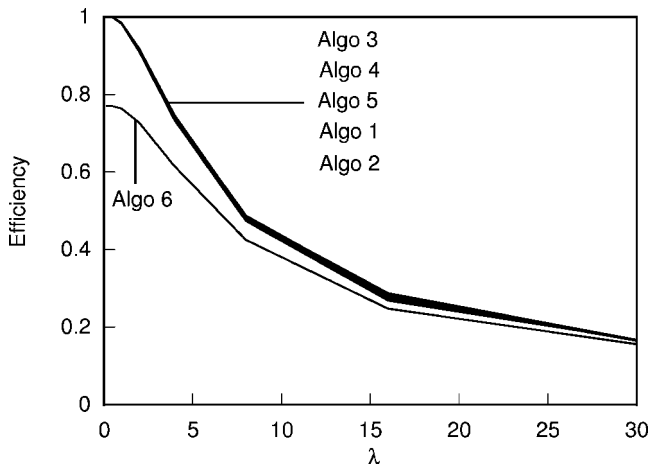


Figure 11. Efficiency e versus λ for eight-node network ($T = 4, \alpha = 4$).

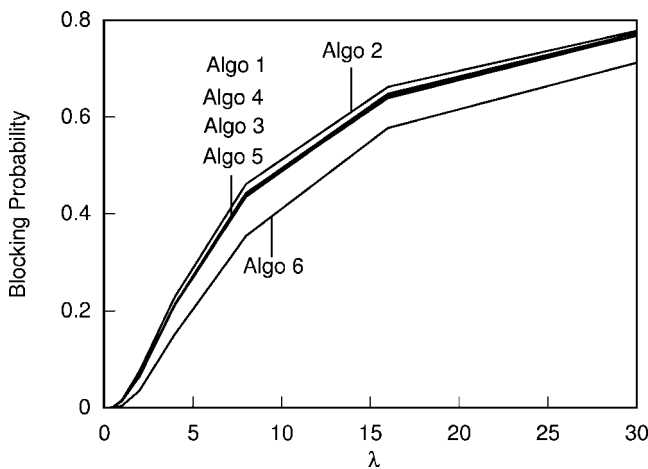


Figure 12. Blocking probability P_B versus λ for eight-node network ($T = 4, \alpha = 4$).

have evaluated thus far. Based on these observations, our future studies will develop algorithms that do not necessarily use the admit-all admission-control policy.

7.1. A larger network example: 100-node network

We have also applied three of our scalable algorithms (1, 2, and 4) to a network consisting of $N = 100$ nodes. Like the eight-node network of figure 6, all nodes are located randomly in a square region of dimensions 5×5 , but here we let $p_{\max} = \infty$. Figures 13 and 14 show the global yardstick Y for $\alpha = 2$ and 4, respectively. Our first observation is that the yardstick values are considerably higher for the 100-node network than for the eight-node network. The higher yardstick values result from the generally much smaller communication ranges that can be implemented by means of extensive relaying in the considerably denser 100-node network (since 100 nodes are now located in the same region as eight nodes in our earlier examples). The impact of these smaller ranges is especially apparent for $\alpha = 4$.

For $\alpha = 2$, the two algorithms that are based on the pruning of spanning trees (algorithms 2 and 4) provide better per-

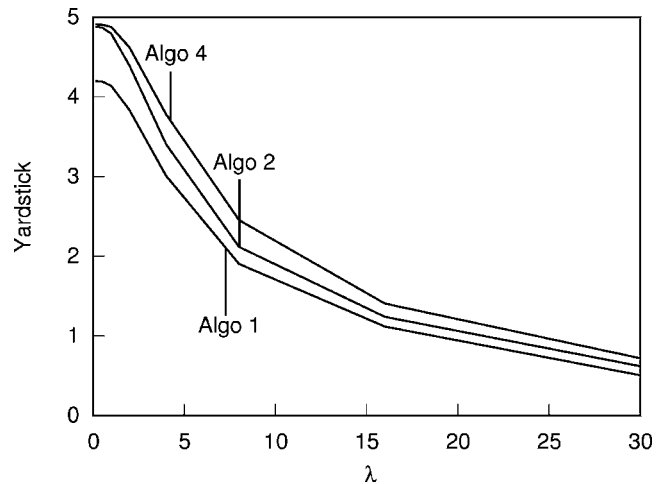


Figure 13. Global yardstick Y versus λ for 100-node network ($T = 4, \alpha = 2$).

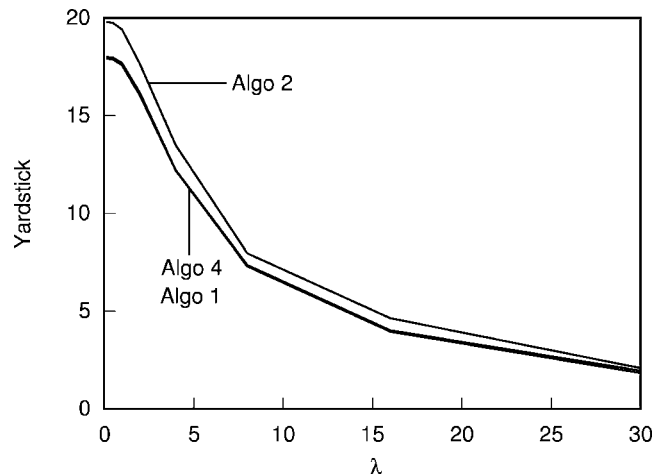


Figure 14. Global yardstick Y versus λ for 100-node network ($T = 4, \alpha = 4$).

formance than the unicast-based algorithm, especially at low levels of offered load. Algorithm 4 provides the best yardstick performance over the entire throughput range, while algorithm 1 provides the worst.

For $\alpha = 4$, the best performance is provided by algorithm 2 (the pruned link-based spanning tree), while there is no visible difference in the curves for algorithms 1 and 4. We suspect that the relatively poor performance of algorithm 4 may be a result of the implementation of equation (9), which can result in long communication ranges when their use results in reaching many destinations (many of which may not be helpful to the construction of the eventual multicast tree). The “penalty” associated with using long communication ranges is especially severe for large values of α .

8. Conclusions

In this paper, we have addressed some of the issues associated with energy-efficient multicasting in ad hoc wireless networks, and we have presented preliminary algorithms for

the solution of this problem. Our studies to date indicate that improved performance can be obtained when exploiting the properties of the wireless medium, e.g., the wireless multicast advantage. The fact that improved performance can be obtained by jointly considering physical layer issues and network layer issues suggests that novel approaches to wireless networking, which incorporate the vertical integration of protocol layer functions, may provide advantages over traditional network architectures.

However, further study of our algorithms is needed, including the study of additional (and larger) networks. Additionally, further research is needed to develop scalable algorithms that can achieve nearly optimal performance. Although some of the algorithms we have studied are non-scalable, they provide a useful benchmark for the evaluation of the performance of scalable algorithms. On the other hand, we have demonstrated that reasonably good performance can be obtained by using simple, scalable heuristics. Thus, we have demonstrated not only that simple algorithms can perform well, but that there is a potential for improvement through the development of improved, scalable algorithms.

References

- [1] V. Rodoplu and T.H. Meng, Minimum energy mobile wireless networks, *IEEE Journal on Selected Areas in Communications* 17 (August 1999) 1333–1344.
- [2] T.J. Shepard, A channel access scheme for large dense packet radio networks, in: *Proceedings of ACM SIGCOMM'96* (1996) pp. 219–230.
- [3] M.B. Pursley, H.B. Russell and J.S. Wycarski, Energy-efficient routing in frequency-hop networks with adaptive transmission, in: *Proceedings of IEEE MILCOM'99*, Atlantic City, NJ (November 1999) paper No. 46.1.
- [4] NSF, Research priorities in wireless and mobile communications and networking, Report of a Workshop (March 1997).
- [5] C.D. Rais and S.K. Tripathi, Studying vertical dependence to improve NFS performance in wireless networks, *Cluster Computing* 1 (December 1998) 225–235.
- [6] M. Zorzi and R.R. Rao, Error control and energy consumption in communications for nomadic computing, *IEEE Transactions on Computers* 46 (March 1997) 279–289.
- [7] E.J. Podlaha and H.Y. Cheh, Modeling of cylindrical alkaline cells VI: Variable discharge conditions, *Journal of the Electrochemistry Society* 141 (January 1994) 28–35.
- [8] M. Nagy and S. Singh, Multicast scheduling algorithms in mobile networks, *Cluster Computing* 1 (December 1998) 177–185.
- [9] C.-C. Chiang, M. Gerla and L. Zhang, Forwarding group multicast protocol (FGMP) for multihop, mobile wireless networks, *Cluster Computing* 1 (December 1998) 187–196.
- [10] C. Diot, W. Dabbous and J. Crowcroft, Multipoint communication: A survey of protocols, functions, and mechanisms, *IEEE Journal on Selected Areas in Communications* 15 (1997) 277–290.
- [11] M.H. Ammar, G.C. Polyzos and S.K. Tripathi, *IEEE Journal on Selected Areas in Communications* 15, Special issue on Network Support for Multipoint Communication (April 1997).
- [12] L. Zhang, S. Deering, D. Estrin, S. Shenker and D. Zappala, RSVP: A new resource ReSerVation protocol, *IEEE Network Magazine* (September 1993) 8–18.
- [13] A.J. Ballardie, P.F. Francis and J. Crowcroft, Core Based Trees, in: *Proceedings of ACM SIGCOMM'93*, San Francisco, CA (1993).
- [14] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C.-G. Liu and L. Wei, Protocol independent multicast (PIM): Motivation and architecture, Internet draft (1995).
- [15] J.P. Macker, Reliable multicast transport and integrated erasure-based forward error correction, in: *Proceedings of MILCOM'97*, Monterey, CA (November 1997) paper No. 26-2.
- [16] K. Makki, N. Pissinou and O. Frieder, Efficient solutions to multicast routing in communication networks, *Mobile Networks and Applications* 1 (1996) 221–232.
- [17] J.E. Wieselthier, G.D. Nguyen and A. Ephremides, Multicasting in energy-limited ad-hoc wireless networks, in: *Proceedings of IEEE MILCOM'98*, Bedford, MA (October 1998) pp. 723–729.
- [18] A. Ephremides, J.E. Wieselthier, G.D. Nguyen and R. Bhattacharya, Energy-efficient wireless multicasting – a beginning, in: *Proceedings of the Third Annual ATIRP Conference* (February 1999) pp. 37–45.
- [19] J.E. Wieselthier, G.D. Nguyen and A. Ephremides, Energy-efficient networking – a first look, NRL report, Naval Research Laboratory (2000) (to appear).
- [20] D. Bertsekas and R. Gallager, *Data Networks* (Prentice-Hall, Englewood Cliffs, NJ, 1992).



Jeffrey E. Wieselthier 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 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. 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 a Professor in the Electrical Engineering Department and the Institute of 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 won the IEEE Donald E. Fink Prize Paper Award (1992) and he was the first recipient of the Sigmobile Award of the ACM (Association of Computer Machinery) 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