

NODE STATE MULTICASTING IN WIRELESS AD HOC NETWORKS

Approved for Public Release; Distribution Unlimited
Case #05-0767

John A. Stine
The MITRE Corporation
McLean, VA 22102-7508

ABSTRACT

Multicasting is an essential service in wireless ad hoc networks. Applying multicasting concepts developed for wireline networks is inappropriate since they create stateful solutions (i.e. nodes learn to react to the receipt of a multicast packet.) which are short lived when topology varies. They require frequent exchanges among nodes with an attendant overhead that increases dramatically with the size of the network, its volatility, and the number of multicast groups. Nevertheless, most proposals for multicasting in ad hoc networks follow the same approach. Our proposal creates a stateless solution that uses the network state information that is already disseminated as part of the Node State Routing (NSR) protocol. Node State Multicasting (NSM) uses this information and various packet formats to enable a rich set of multicasting capabilities. Multicast routing is implemented by explicitly listing end destinations or regions in packet headers. Intermediate nodes assume responsibility for the delivery of packets to the end destinations or regions listed in the header. Routing decisions are based on the NSR routing tables. This approach is very generic and can support both traditional wireline multicast scenarios and additional scenarios typical of wireless applications.*

INTRODUCTION

Multicasting is a necessary service in military applications of ad hoc networking. Command and control, situation awareness, warnings and alerts are all source to many destination communications. Multicasting, however, has been one of the more difficult services to implement. We contend that this difficulty is the result of applying concepts conceived for wireline networks. Their shortcoming is that they attempt to create forwarding state in routers that requires a stable topology to work. The volatility of ad hoc networks make such state tenuous and attempts to keep the state current can be overwhelming in overhead. Further, stateful solutions do not support tracking group membership and so do not support call admission, service billing, or end-to-end transmission control. We offer as an alternative Node State Multicasting (NSM) which is made possible by Node State Routing (NSR) [1]. NSM, on its

own, has no overhead but uses information already collected and disseminated by NSR. The mechanics of multicasting is handled in a multi-destination packet format. The source of the multicast determines the final destinations and uses a single packet for all destinations that share the same next hop. Each router on the path does the same for the list of destinations in the arriving packet. Not only does this simplify the mechanics of multicasting it enables interesting and useful multicasting functions that are not possible with wireline concepts.

This paper begins with a broad review of multicasting algorithms and protocols in wireline and ad hoc networks. It then describes multicast scenarios emphasizing useful multicast capabilities that are not possible with stateful solutions. We begin our presentation of our alternative with an overview of NSR emphasizing the relevant features that affect multicasting. Finally, we describe the algorithms and mechanisms of NSM, describe how NSM supports the additional multicast capabilities needed, and identify the problems that must be resolved for its implementation.

MULTICASTING

Multicasting protocols have three components: a mechanism for nodes to become members of a multicast group, a mechanism for sources to reach a group, and a mechanism to build a tree or mesh across which to distribute packets. We review these components by briefly describing multicast routing algorithms, multicasting in wireline networks, and current work on multicasting in ad hoc networks.

MULTICAST ROUTING ALGORITHMS

Multicast algorithms determine a multicast tree or mesh to connect the source(s) to the group members with the intent that multicast packets traverse the edges of the tree just once. The origin of the tree may be shared or source specific with the terms core based tree (CBT) and source based tree (SBT) distinguishing the two respectively. When a CBT is used, sources route packets to the core from which they are distributed. When SBTs are used, each source for a group will have its own tree.

Given the network topology, group membership, network state, and constraints on performance and use of resources, the multicast routing algorithm attempts to maximize a performance objective. An overview in [2] provides a tax-

* Patent pending

onomy with two dimensions, type of constraints and optimization objective. There are two types of constraints, link and tree, that may be applied individually, combined or not at all. Link constraints are restrictions on the use of links. Tree constraints are bounds on the performance metric for packet delivery. There are two types of optimization, link and tree. The former optimizes the performance of all links in the tree and the latter optimizes the total cost of using the tree. Tree optimization problems are NP complete although there are numerous proposed heuristics.

The better known tree formation problems are the shortest path tree (SPT) which minimizes the origin to receiver cost for all group members, the minimum spanning tree which minimizes the tree cost when all nodes in the tree are group members, and the Steiner tree which is the minimum cost tree problem that is NP complete. Most multicast protocols provide SPT solutions.

In implementation, multicasting protocols react to changes in group membership and topology. The relative significance of the two types of changes depends on the rate they occur and the number of nodes and routers involved. Changes in group membership are the greater concern in wireline networks while changes in topology are the greater concern in wireless ad hoc networks.

MULTICASTING IN WIRELINE NETWORKS

In anticipation of very large multicast groups and a mostly static topology, wireline multicast protocols provide maximum flexibility for stations and routers to join and leave multicast groups. They build state at routers (i.e. routers learn how they should forward multicast packets). This approach allows the acts of joining and leaving to be restricted to the regions of the network they occur and the parts of the multicast tree they affect.

The Internet Group Management Protocol (IGMP) [3] enables multicasting on IPv4 LANs. It allows routers to identify *whether* stations on a connected LAN are subscribing to a multicast. This is sufficient for routers to know whether to broadcast multicasts onto the LAN and whether they need to join or leave a multicast tree. Routers do not learn the specific members and so at the destination edge knowledge of group membership is lost.

Meanwhile, multicast protocols among routers attempt to build trees. The most basic technique is for sources to create SBTs by flooding the initial packet to a group and then pruning the tree to remove links to routers that lead to no destinations. The Distance Vector Multicast Routing Protocol (DVMRP) [4] and Protocol Independent Multicast (PIM) Dense Mode [5] protocols that use this technique use reverse path forwarding (a router only forwards a packet that arrives on a link that is on the shortest path to the source) which limits the number of transmissions and

creates an SPT. New routers can use join messages to graft themselves to an existing tree. Nevertheless, periodic flooding and pruning is necessary to maintain SPTs.

The multicast routing extension to OSPF (MOSPF) [6] tracks group membership of routers. Routers calculate the forwarding state for a specific source and multicast group with the first packet that arrives from the source. It requires all routers use the same view of network topology to determine their forwarding responsibilities to arrive at a non-looping source-based SPT. This common view becomes less likely as network size and volatility increases.

The PIM Sparse Mode (SM) [5] protocol uses CBTs. If the volume of traffic from a source exceeds a threshold the PIM-SM protocol reverts to an SBT for that source. The CBTs and SBTs created by PIM-SM are SPTs. Tree construction originates from destination routers and the trees are built using existing routing tables and reverse path forwarding techniques. The default is to build a CBT and downstream routers initiate the transition to an SBT by triggering its construction.

As described, these multicast protocols were not designed to respond to changing topologies. Changing topology requires the reconstruction of multicast trees and forwarding state at routers. Volatile networks risk not keeping up with changes and generate substantial overhead trying.

MULTICASTING IN AD HOC NETWORKS

Proposed protocols for ad hoc environments seek solutions that are more responsive to topology changes yet attempt to balance overhead with the reliability of packet delivery. There are many proposed protocols but most are very similar to their wireline cousins. They too attempt to build forwarding state in routers. Efficiencies are achieved by relaxing the typical multicast optimization goals and allowing non-optimal trees and redundant packet delivery.

As there is a very large number of proposed multicast protocols we will review this work from a taxonomical view. We use the approach presented in [7] which divides multicasting protocols first into two broad groups that divide by whether multicasting is application independent or dependent. The independent group is the typical multicasting approach where destinations indicate their membership by joining groups. The application dependent versions multicast packets to destinations based on their context: their location, activity or need for the information in the multicast. We provide more details on the need for these dependent approaches in the next section. We are aware of no multicasting protocols that attempt to implement both types of schemes together.

The taxonomy further divides the independent group by three characteristics, the multicast topology used, the ini-

tialization approach, and the topology maintenance approach. The topology can be the typical tree based topology or a mesh based topology. The latter is used to provide a more robust solution since meshes provide multiple paths from sources to destinations. Initialization refers to where the path creation originates, from the sources or the destinations. Maintenance can be soft or hard state with the distinction that hard state is the attempt to fix a broken link while soft state approaches continuously attempt to refresh the topology. The relative merits of different combinations of these mechanisms are not definitive. Also, these protocols can be divided along the same lines as routing protocols with there being reactive, proactive and hybrid variants. Multicast protocols are every bit as complex as routing protocols with similar issues on the generation of overhead. Just as there is no routing protocol that stands out for ad hoc networking, there is no multicasting protocol. In fact, with the drive to make ad hoc networks compatible with the larger networking infrastructure, wire-line multicast protocols are appearing as preferred approaches in military ad hoc networking proposals.

A unique multicasting scheme that is similar in some ways to what we are proposing is Differential Destination Multicasting (DDM) [8]. Rather than building state at routers, DDM allows the sources to specify the multicast destinations in a variable length destination header. The protocol uses the unicast routing tables to determine how to forward the packets. If a multicast is persistent, the protocol can revert to a soft state forwarding approach where the destination list in the header is dropped. Implementation of DDM requires destinations to subscribe directly with sources. This feature is both a strength (enables admission control) and a weakness (disables traditional multicasting).

MULTICASTING SCENARIOS

IP multicasting has two main features: sources send packets to a multicast group by using the multicast address and routers forward these packets using the forwarding state created for that multicast address. Sources cannot control group membership. If there is a need for a source to reach a set of destinations not already part of a group, then the source must either unicast the messages to each destination or advertise a multicast address and wait for the destinations to subscribe. The first option is not multicasting and the second is slow with no safeguards against illicit membership. It is this ability of sources to control access to their multicasts and to choose destinations that characterizes many military multicasting scenarios.

In tactical networks operational requirements will determine how information is multicast. At minimum, capacity constraints will require dissemination to be selective. Further, operational scenarios are likely to require multicast groups that cannot be anticipated. Dissemination may be

based on need to know or the context of the destinations (e.g. their location or role). Below we list five multicast scenarios that are useful to military networks but impractical to implement using the standard IP approach.

CONTROLLED ACCESS

A controlled access scenario occurs when the multicast data must be limited to those nodes that have a need to know. In this scheme, destinations would petition the source to subscribe to the multicast data and the source would admit the destination based on some criteria it evaluates.

GEOGRAPHIC WARNING

Geographic warnings are multicasts that have a geographic limit to their relevance and so are limited to all nodes or a subset of nodes in the specified region.

GEOGRAPHIC DISTRIBUTION

Geographic distribution is distinguished from geographic warning in that it defines a distribution approach rather than operational objective. Sources send the packet to a node in the midst of the region from which a tree is created. The packet contains a definition of the geographic region and this node determines the members within the region and creates a tree to those members.

AD HOC COLLABORATION

Ad hoc communities of interest are organized to coordinate activities. The objective of the service to support ad hoc collaboration is to allow a user to solicit group membership. The initial multicast is sent to user specified destinations. Any response to this message would also be disseminated to the specified group members. This type of multicast is likely to be used in command and control situations where synchronized activities are required of actors in disparate organizations.

THROTTLED DISSEMINATION

In cases of congestion or when there is a need to preserve capacity for certain QoS purposes it is desirable for there to be control of the volume of traffic offered to the network. Publish and subscribe information systems are likely to be a culprit of congestion. Throttling can reduce congestion. There are two implementations: reducing the rate information is disseminated or limiting dissemination to a subset of the subscribers. The latter approach is enabled by the user being able to specify the destinations.

NODE STATE ROUTING (NSR)

Node State Routing (NSR) is an alternative to the standard link driven approaches to routing. The distinction is that rather than discovering and explicitly disseminating connectivity in terms of links, node states are disseminated

and connectivity is inferred from their pairwise use. Articulating network state information in node states allows NSR to support other functions such as quality of service [1] and, as we describe here, multicasting. NSR is implemented beneath IP and is very much a part of the link layer. It is intended for a homogeneous wireless network. Fig. 1 illustrates that additional routing functionality above IP is needed for heterogeneous networks. IP routing exchanges information with NSR routing and does not offer load to the wireless network.

OVERVIEW

In lieu of links, there are two different routing constructs used in NSR, a node and a wormhole. The node construct is modeled as a point in space and is assumed to have connectivity with other nodes through the use of wireless connections. In many cases nodes may be connected using a dedicated link such as a cable. To use these links within the node state routing protocol we define a second routing construct called a wormhole. We define our wormhole construct as a directed path between two points in the network. The basic algorithm used to select which routing constructs to use in a route considers the cost of sending a packet to a construct and the cost of using the construct. These costs are derived from the states of the nodes and the wormholes.

NSR requires two capabilities: location awareness and the ability to measure signal strength. With this information, each node creates a pathloss map. Location and the pathloss maps of all nodes and wormhole endpoints provide sufficient information to determine connectivity between the constructs and then the overall topology.

NSR consists of three processes, propagation map discovery, node state dissemination, and a route calculation. On a periodic basis, each node in the network transmits node state update packets. These transmissions are used to discover propagation conditions and to disseminate the node states. Either on a periodic basis or as required, nodes use these states to determine topology. We now describe these processes in greater detail.

NODE STATES

The node states used in NSR may describe any type of state information for a node. As a minimum, it provides the node's location, the propagation conditions about the node, and a mapping between IP and MAC addresses. Table I provides the minimum states required to implement multicasting. The use of other node states for the purpose of QoS or energy conservation [1] can also be used to extend these services to multicasting.

Propagation maps are data structures in which pathloss conditions are recorded so that link budgets can be calcu-

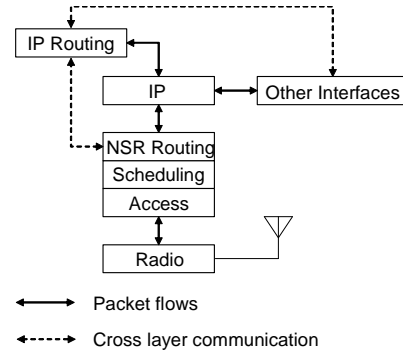


Fig. 1. NSR's multilayer routing functionality

Table 1. Proposed node states to support multicasting

STATE	DESCRIPTION
Address	MAC address of the node or the wormhole. In the case of the wormhole, the address is associated with the node at the front end.
1-meter Path loss	Pathloss of the first meter of propagation used with the log distance pathloss model.
Propagation map	Propagation conditions can vary based on the location of nodes and the direction of propagation. To accommodate this concern nodes measure and estimate a pathloss exponent for the pathloss model. We require each node that broadcasts a packet to announce the power level it is using. We assume that each destination node that hears a broadcast can determine the power level of the received signal and can then estimate a pathloss exponent using the attenuation of the signal and the separation distance from the source. When propagation characteristics vary to different destinations, these states can be broken up into different sectors that account for these differences.
Cost	A cost that is assigned to using a node or a wormhole that is considered when assigning a metric to a link.
IP Addresses	IP addresses that are used by the node. It includes multicast addresses.
Location	The location defines where the node or where the wormhole's endpoints physically exist in the network. Node state routing requires location awareness.
Time Stamp	This is the time that the reported state was measured. We assume time is absolute and synchronized throughout the network.

lated. There are many ways this information can be determined and placed into a data structure. Our approach is to empirically observe and record the pathloss from neighboring transmitters and then to record these into a data structure that differentiates different pathloss values by direction from the receiver. Pathloss in a particular direction is articulated using two values, a one meter pathloss value, PL_{1m} , and a pathloss exponent, n , which when used with the log distance pathloss model, $PL(\text{dB}) = PL(1m) + 10n \log(d)$, estimates a pathloss between a transmitter and receiver given the distance that separates them. We use a single 1-meter pathloss per node and differentiate different pathloss conditions by using different exponents. To differentiate pathloss by direction, we use a variable data structure that uses a series of words to specify path loss exponents on a directional basis. We use 8 bit words which allows us to specify 256 different pathloss exponents, in our case $n = 1.9$ to 7.0 in increments that provide equidistant changes in propagation range and to divide a sphere into 256 longitudes (θ) and, by choice, 180 latitudes (ϕ), providing 46,080 sectors. Not all sectors need to be explicitly specified. The propagation map would have the form $(0, 0, n_{00}, \theta_{01}, n_{01}, \dots, \theta_{0x}, n_{0x}, 255, \phi_1, \theta_{10}, n_{10}, \theta_{11}, n_{11}, \dots, 255, 180)$. Since $\phi = 0$, $\theta = 0$, $\theta = 255$




Fig. 2. Example propagation map

and $\phi = 180$ occur predictably we drop most from the structure. We still use $\theta = 255$ and 0 as delimiters in our abbreviated data structure. The value 255 delimits a horizontal sweep and the value 0 delimits the end of the map. Fig. 2 illustrates an example propagation map and the transmission ranges it predicts.

TOPOLOGY DETERMINATION

Given a set of node states, each node determines topology in three steps. First, connectivity between constructs is inferred using their propagation maps and locations. For each pair of constructs both propagation maps are applied and the worst case predicted pathloss is used as the measure of connectivity. If this predicted pathloss is below some designated threshold, a connection is inferred. Second, for all inferred links a metric is assigned. These metrics are formed from the node states and include the cost of transmitting the packet and using the destination construct. Finally, Dijkstra's algorithm is used with the weighted set of inferred links to find the shortest paths to all destinations. The power of this approach is that a whole assortment of filters and weighting techniques can be used to affect the routing tables that are calculated without having to change the state dissemination mechanism.

SCALABLE NODE STATE DISSEMINATION

Nodes distribute the node states using a diffusion mechanism. On a periodic basis a node will broadcast a node state packet (NSP) which will include its own state and other states in its list restricted in number by the maximum packet size. The states that are included in these updates are selected by two criteria, a threshold that indicates whether an update is needed and a prioritization criterion to enable selection amongst several states that meet the update threshold. In the diffusion process, the update threshold depends on the distance between the node that owns the state and the node doing the rebroadcast.

Scaling is forced using a minimum interval between NSP updates, i.e., a node may send one NSP per interval. However, NSP updates are accelerated when routing failures are observed. Loops do not occur in link state routing protocols if all nodes use the same states. In NSR, nodes may

have different node state information and loops may occur. The observation of a loop triggers accelerated updates. The goal of these updates is to synchronize the node state tables of all the nodes in the loop so it can be broken. After identifying a looping condition, a node in the loop broadcasts a relevant subset of its node state table that covers the region of interest, recalculates its routing tables and then forwards the packet that was looping. This process is repeated so long as the packet remains in the loop. Ultimately, all nodes in the loop will have a common picture of the network and the packet will progress.

Through diffusion and forced scaling, NSR aggressively employs fisheye scope [9] and the distance effect [10] to mitigate stale states. Fisheye scope refers to the effect distance has on the accuracy of a node's view of the network's topology. It is most accurate close to the node and so route accuracy improves as packets progress toward their destination. Then, because location is a part of the routing calculation, the distance effect mitigates routing errors. The distance effect refers to the effect that the further nodes are apart from each other, the less effect their relative movement has on the direction between the two nodes. The next hop in routing a packet between the two, even with the stale information for distant nodes, is likely to be correct. The use of loop detection and the subsequent accelerated node state distribution correct the situation when the network is too volatile.

NODE STATE MULTICASTING (NSM)

Node state multicasting (NSM) is a feature built directly on top of NSR and so requires no additional overhead. It supports both the standard multicasting approach where destinations subscribe to a multicast address and the special scenarios where the source specifies the destinations. Multicasting in NSM is accomplished with four features, the mapping of IP addresses to MAC addresses in the node state tables, special multicasting packet formats, algorithms for forming and routing packets, and cross layer communications.

ADDRESS MAPPING

Two of the states that NSR disseminates are the MAC address of the wireless modem and a list of all IP addresses of the hosts serviced by the modem including multicast addresses to which they subscribe. Thus, IP addresses and MAC addresses are mapped to each other. This scheme eliminates the need for the address resolution protocol and provides a means to join multicast groups. Nodes join a multicast group by adding its address to their states.

PACKET FORMATS

A standard NSR MAC packet will have three MAC addresses in the header, the source, the next hop and the final


destination address. To support NSM, we add an additional optional extension that is an address list. NSM populates this list with the final destinations for the packet and the final destination in the standard header is set to a MAC multicast address. The next hop node upon receiving this type of multicast packet has responsibility for forwarding the packet so that it can reach each of the destinations in the list. This node may transmit it once with the same destination list or multiple times each with different non-intersecting subsets of the original destination list.

Multicasting may be reliable or unreliable at the MAC. A reliable implementation requires each packet recipient to acknowledge receipt in the same frame. An unreliable implementation tries to take advantage of the broadcast nature of the wireless medium and a single packet is transmitted to multiple recipients making acknowledgment in the same frame impractical. In the reliable implementation, the next hop address is a standard peer address and in the unreliable packet it is the limited broadcast address. For an unreliable multicast the list of multicast addresses is delimited by some unique address with the first address after each delimiter being the specific next hops that must receive and handle the packet. The addresses between the next hop addresses and the next delimiter are the destinations to which it must forward the packet. Fig. 3 illustrates the differences in these approaches.


We offer specialized packet formats for geographical multicasting. The basic concept is for the packet's source to route a packet to a distant node that is given the responsibility to multicast the packet to all nodes in a region. Regions may be described as a center point and range or as a rectangle using three coordinates. The header must also specify the qualifier to be a destination, either a multicast address, a broadcast address, a node state, or an explicit list of destinations. Fig. 4 illustrates the different header information for the packets. The main header appears as an ordinary peer-to-peer packet header with a flag set for the appropriate region based routing. The additional header contains the region definition and the destination qualifier. The destination node in the main header has responsibility to identify the final destinations and to create the multicast packets of the type shown in Fig. 3 to distribute the packet.

FORMING AND ROUTING PACKETS


Table 2 illustrates the fields of the NSR routing tables. For each destination there is an entry for the next hop, the node this node (i.e. the node that owns the table) should forward the packet for delivery and an entry for the previous hop which is the node this node thinks will ultimately forward the packet to the destination. The previous node entry allows this node to reconstruct the full path in the reverse.



a. The multicast scenario. Node 1 is the source and nodes 2, 4, 5, 6, and 7 are the final destinations.



b. Packet headers for the reliable MAC multicast. Node 1 sends two packets to 2 and 3 and then those nodes forward the packet. The MC destinations of the packets are listed in the MC DST list



c. Packet headers for the unreliable MAC multicast. Node 1 sends just 1 packet and the MC DST List identifies all the next hops and the destinations to which those nodes must forward the packet. Nodes 2 and 3 forward the packet and all multicast destinations are covered.

Fig. 3. Multicast packet headers

SRC Addr	NH Addr	DST Addr	MC Qualifier
----------	---------	----------	--------------

SRC Addr – The node transmitting the packet
 NH Addr – The next hop destination of the packet
 DST Addr – The destination node that subsequently multicasts the packet
 MC Qualifier – Method use to identify the final destinations, 0 = none, 1 = DST list, 2 = MC or broadcast addr, 3 = State values

a. Packet header

DST List	Broadcast Addr Region description
Final destinations are those in the list	Final destinations are all nodes in the region
MC Addr Region description	State values Region description
Final destinations are all nodes in the region that subscribe to the specified multicast address	Final destinations are all nodes in the region that have state values that match the listed values

b. Qualifier criteria for multicast destinations

Fig. 4. Geographical multicasting packet headers

Table 2. Typical NSR routing table (Table excerpt of Node 2 in Fig. 3a.)

Destination	Next Hop	Previous Hop	Cost	Distance
1	1	2	.8	1
3	1	1	1.7	2
4	1	3	2.8	3
5	1	3	2.5	3
6	1	3	3.0	3
...				

Given a complete listing of destinations for a multicast packet the source node sorts the destinations into groups that share the same next hop. In the reliable approach a packet is created for each next hop while in the unreliable approach the groups of destinations are listed in blocks as illustrated in Fig. 3. In some implementations there may be a limit to the number of destinations that can fit in the destination list. In this case the list must be subdivided and multiple packets formed. The list is subdivided based on the expected follow-on hops that packets will be forwarded by the next hop node. These sublists are combined if together they have fewer destinations than the list limit. If a sublist remains too large the process is repeated using the third hop as the criteria for dividing, and so on. Sublists are never divided for the sake of filling a destination list. This type of division should only be necessary at the source of the multicast since lists never grow as they are forwarded. Intermediate forwarding nodes base all their actions on the destinations specified in the MC list.

CROSS LAYER COMMUNICATIONS

NSM is primed to support the many useful multicasting tasks that are described in the multicast scenarios. Any listing of destinations can be used to define the destinations of a multicast packet. At present it is not clear how applications will articulate and inform NSM of these destinations. Parts of the functionality of deciding the destinations to send a packet will reside in applications and in the IP routing protocol illustrated in Fig. 1.

EXPERIMENTS AND FUTURE WORK

We have implemented a crude form of NSM, only the reliable multicast, in OPNET and tested its performance in a specialized simulation environment that has been developed for military scenarios. [11] The scenarios in this environment heavily use multicasting. The traffic is threaded such that success leads to more exchanges. Our implementation using a 1 Mbps physical layer delivered three times more goodput than the default implementation using a 600 Mbps data rate. There were many differences in the characteristics of the wireless nodes and so this observation serves only to substantiate the suitability of our approach.

Further research is needed into how to integrate the specialized multicast approaches in wireless ad hoc networks with the applications and the heterogeneous networks with which they are expected to be used. Many interesting multicasting approaches are possible.

A criticism of this type of explicit multicasting is that it is impractical with large groups. We envision using NSM's geographical multicasting mechanisms to overcome these shortcomings. The network may be divided into geographic regions with the explicit multicasting being con-

tained within those regions. Both choosing regions and the nodes responsible to route within those regions are interesting research problems.

CONCLUSION

We reviewed the issues of multicasting in ad hoc networks. Multicasting is challenging not only because of the volatile topology of these types of networks but also their unique multicast requirements. We provided a quick review of wireline approaches and demonstrated that the core mechanism that they use is to build forwarding state at routers. Changing topologies may prevent convergence and networks may be overcome by the resulting flood of administrative traffic of these protocols trying. We further provided an overview of research on multicasting in ad hoc networks explaining that these approaches follow the footsteps of their wireline cousins and also build state at routers. Although better, they too may not converge and are large sources of overhead. We point out that the stateful solutions that these protocols create are not sufficient for military networks and are not designed to support heterogeneous networking. We describe NSR and how its state dissemination mechanism supports joining multicast groups and tracking group membership. We described NSM, how it solves the traditional multicasting problem, and how it can also support the multicasting scenarios that require source selection of destinations. Finally, we point out that exploiting the capabilities of NSM will require the development of protocols that applications can use to get access to these unique multicast services.

REFERENCES

- [1] J. Stine and G. de Veciana, "A paradigm for quality of service in wireless ad hoc networks using synchronous signaling and node states," *IEEE J. Selected Areas of Communications*, Sep. 2004, pp. 1301-1321.
- [2] B. Wang and J. Hou, "Multicast routing and its QoS extension: problems, algorithms, and protocols," *IEEE Network*, Jan-Feb 2000, pp. 22-36.
- [3] B. Cain, S. Deering, I. Kouvelas, B. Fenner, and A. Thyagarahan, IETF RFC 2236, Internet Group Management Protocol, Version 3, Oct 2002.
- [4] D. Waitzman, C. Partridge, and S. Deering, IETF RFC 1075, Distance vector multicast routing protocol, Nov. 1988.
- [5] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C. Liu, and L. Wei, "The PIM architecture for wide-area multicast routing," *IEEE/ACM Trans. on Networking*, Apr. 1996, pp. 153 - 162.
- [6] J. Moy, "Multicast routing extension for OSPF," *Comm. of the ACM*, Aug 1994, pp. 61-66.
- [7] C. Murthy and B. Manoj, *Ad Hoc Wireless Networks: Architectures and Protocols*, Prentice Hall, Upper Saddle River, NJ, 2004.
- [8] L. Ji and M. Corson, "Differential destination multicast - a MANET multicast routing protocol for small groups," IEEE INFOCOM 2001, pp 1192 - 1201.
- [9] G. Pei, M. Gerla, and T-W. Chen, "Fisheye Scope Routing: a routing scheme for wireless ad hoc networks," *Proc. of the Int. Communications Conf.*, June 2000, pp. 70 - 74.
- [10] S. Basagni, I. Chlamatac, V. Syrotiuk, and B. Woodward, "A distance effect routing protocol with group mobility," *Proc. of the 4th Annual IEEE/ACM Conf. on Mobile Computing and Networking*, 1998, pp. 76 - 84.
- [11] G. Comparetto, E. Lindy, M. Mirhakkak, and N. Schult, "Overview and application of a modeling and simulation environment to support protocol performance evaluations in mobile communications networks," *Proc. Int. Conf. on Communications in Computing*, 2004