An Interest-Driven Approach to Integrated Unicast and Multicast Routing in MANETs

Rolando Menchaca-Mendez[§]
[§]Computer Engineering Department
University of California, Santa Cruz
Santa Cruz, CA 95064, USA
Email: menchaca@soe.ucsc.edu

J.J. Garcia-Luna-Aceves [§,*]
[*]Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304, USA
Email: jj@soe.ucsc.edu

***Abstract*—This paper introduces an integrated framework for multicast and unicast routing in mobile ad hoc networks (MANET) based on interest-defined mesh enclaves. Such meshes are connected components of a MANET that span the sources and receivers of unicast and multicast flows. We present the Protocol for Routing in Interest-defined Mesh Enclaves (PRIME), which establishes meshes that are activated and deactivated by the presence or absence of interest in destinations and groups, and which confines most of the signaling overhead within regions of interest (enclaves) in such meshes. Experimental results based on simulations show that PRIME attains similar or better data delivery and end-to-end delays than traditional unicast and multicast routing schemes for MANETs (AODV, OLSR, ODMRP), and that PRIME incurs only a fraction of the signaling overhead of traditional routing schemes.**

I. INTRODUCTION

The price, performance, and form factors of processors, radios and storage elements are such that mobile ad hoc networks (MANETs) can finally support distributed applications on the move. These applications (e.g., disaster relief) require point-to-point and many-to-many communication, and very few destinations or groups are such that a large percentage of the nodes in the network have interest in them. As Section II outlines, these application requirements are in stark contrast with the way in which today's MANET routing protocols operate. First, they have been tailored to support either unicast routing or multicast routing. Hence, supporting both point-to-point and many-to-many communication in a MANET involves running a unicast and a multicast routing protocol in parallel, which is very inefficient from the standpoint of bandwidth utilization. Second, the proactive and on-demand routing protocols for unicasting and multicasting proposed to date are such that the network is flooded frequently with link-state updates, distance updates, route requests, or multicast updates. This is the case even when the protocols maintain routing information on-demand (e.g., AODV and ODMRP).

The main contribution of this paper is to introduce a new framework for routing in MANETs. In this new approach to routing, the same control signaling is used to support unicast and multicast routing, and the distinction between on-demand and proactive signaling for routing is eliminated and interest-driven signaling is used instead. A node (router) maintains routing information proactively for those unicast or multicast destinations for which it has interest (user traffic) or for which other nodes have interest and the node can serve as relay. To attain this, interest-defined mesh enclaves are established and maintained, and such meshes are connected components of a MANET over which control signaling and data packets for unicast or multicast flows are disseminated.

Section III presents the Protocol for Routing in Interest-defined Mesh Enclaves (PRIME), which implements our integrated routing framework. In PRIME, the routing structure needed to forward packets for multicast and unicast flows are established using the same mechanisms. PRIME establishes enclaves for flows of interest on-demand, and signaling to update routing information within enclaves is sent proactively. Those regions of the network with interest in the destinations of flows receive timely updates, while the rest of the network receives updates about the flows with far less frequency, or not at all.

Section IV describes the results of simulation experiments used to study the performance of PRIME and compare it with that of relevant multicast and unicast routing protocols for MANETs. Our comparison addresses the performance of the protocols purely for multicast routing, and their performance in supporting unicast and multicast routing. We compare PRIME with ODMRP [11] and PUMA [17] to determine PRIME's effectiveness as just a multicast routing protocol, and consider different numbers of sources, groups, node densities and the use of group and random waypoint mobility models. We also compare PRIME against the use of AODV and ODMRP, and the use of OLSR and ODMRP. The results show that PRIME is a very efficient multicast routing protocol and provides substantial performance improvements over the traditional approach to supporting unicast and multicast routing. PRIME attains similar or better delivery ratios and significantly lower delays and communication overhead than the traditional approaches.

II. RELATED WORK

There have been a large number of routing protocols proposed and implemented to date for MANETs, and we can only address a very small sample of them in this paper due to space limitations. Our summary is intended simply to highlight the facts that (a) existing routing protocols for MANETs

1-4244-2507-5/08/$20.00 ©2008 IEEE

support either unicast routing or multicast routing, and (b) the dissemination of signaling traffic in MANETs is not closely linked to the interest that nodes have on destinations, and is structured as either strictly on-demand, strictly proactive, or the use of both types of signaling by dividing the network into zones.

Unicast routing protocols for MANETs are typically classified into proactive and on-demand. Proactive routing protocols maintain routing information for all destinations independently of the interest in them, i.e., regardless of the unicast flows in the network. There have been proposals based on distance (e.g., WRP [12]) or link-state information (e.g., OLSR [7]), and many approaches to reduce the amount of overhead incurred in disseminating routing information proactively. On-demand routing protocols (e.g., AODV [13], DSR [9]) maintain routes for only those destinations for which there is interest, which makes them attractive when not all the destinations are very popular.

There have also been proposals based on a combination of proactive and on-demand routing (e.g., ZRP [4], NEST [14]). In these hybrid schemes, however, proactive signaling is applied within areas or zones of the network [4] independently of the interest for destinations in such zones, or for specific destinations [14], and on-demand signaling propagates throughout the network.

Multicast routing protocols can also be classified on proactive and on-demand. However, they are typically classified based on the type of routing structure they construct and maintain; namely, tree-based and mesh-based protocols. A tree-based multicast routing protocol constructs and maintains either a shared multicast routing tree or multiple multicast trees (one per each sender) to deliver packets from sources to receivers. Several tree-based multicast routing protocols have been reported. The multicast ad hoc on-demand distance vector protocol (MAODV) [15] maintains a shared tree for each multicast group consisting of receivers and relays. Sources acquire routes to the group on-demand in a way similar to the ad hoc on-demand distance vector protocol (AODV) [13]. The adaptive demand-driven multicast routing protocol (ADMR) [8] maintains a source-based multicast tree for each sender of a multicast group. In ADMR, a new receiver performs a network-wide flood of a multicast solicitation packet when it needs to join the multicast group. Each source replies to the solicitation and the receiver sends a receiver join packet to each source that answered the solicitation. Each source-based tree is maintained by periodic keep-alive packets from the source. Like ADMR, MZR [2] maintains source-based trees, but performs zonal routing; and hence the dissemination of control packets is less expensive.

In the context of MANETs, establishing and maintaining a tree or a set of trees in the presence of frequent topology changes incurs substantial exchange of control messages. A simpler and more robust approach to maintaining trees is to maintain a routing mesh consisting of a connected sub-graph of the network containing all receivers of a particular group and the relays needed to maintain connectivity. Three basic approaches of mesh-based multicast routing are characterized by the On-Demand Multicast Routing Protocol (ODMRP) [11], the Core Assisted Mesh Protocol (CAMP) [3], and the Protocol for Unified Multicasting through Announcements (PUMA) [17].

In ODMRP [11], group membership and multicast routes are established and updated by the sources on-demand. Each multicast source broadcasts Join Query (JQ) packets periodically, and these are disseminated to the entire network to establish and refresh group membership information. When a JQ packet reaches a multicast receiver, it creates and broadcasts a Join Reply (JR) to its neighbors stating a list of one or more forwarding nodes. A node receiving a JR listing it as part of forwarding groups forwards the JR stating its own list of forwarding nodes. Several extensions to ODMRP have been proposed to reduce the signaling overhead it incurs. DCMP [1] designates certain senders as cores and reduces the number of senders performing flooding. NSMP [10] aims to restrict the flood of control packets to a subset of the entire network. MMARP [16] builds its multicast mesh as the union of a set of trees that approximate Steiner trees rooted at each source. The salient feature of ODMRP and its extensions is that multiple nodes produce some flooding for each multicast group.

CAMP [3] avoids the need for network-wide disseminations from each source to maintain multicast meshes by using one or more cores per multicast group. Only cores flood the network with signaling information about multicast groups, and a receiver-initiated approach is used for receivers to join a multicast group by sending unicast join requests towards a core of the desired group. In our view, the main limitations of CAMP are that it requires a unicast routing protocol to maintain routing information about the cores, and that information about all multicast groups is maintained proactively throughout the network.

PUMA [17] also uses a receiver-initiated approach in which receivers join a multicast group using the address of a core. PUMA eliminates the need in CAMP for an independent unicast routing protocol by implementing a distributed algorithm to elect one of the receivers of a group as the core of the group, and to inform each router in the network of at least one next-hop to the elected core of each group. The limitation of PUMA is that all nodes must receive periodic signaling packets regarding each multicast group, regardless of whether nodes have interest in the group.

Directed diffusion [6], which has been proposed for sensor networks, is closest to our interest-driven approach to signaling for unicasting. In directed diffusion, sinks disseminate interest in information objects, sources with the information forward it towards the sinks, and some reinforcement mechanisms are used between sources and sinks. The directed diffusion approach is in some ways similar to the way in which ODMRP operates for multicast routing. The framework we advocate in PRIME also focuses on using interest to tailor the signaling incurred for routing, limits the reach of the signaling to regions where interest exists, and limits the amount of signaling packets sent within such regions.

III. PRIME

A. Overview

The Protocol for Routing in Interest-defined Mesh Enclaves (PRIME) establishes and maintains a routing mesh for each active multicast group, i.e., for each group with active sources and receivers and for each unicast destination with at least one active source. The first source that becomes active for a given unicast or multicast destination sends its first data packet piggybacked in a Mesh Request (MR) packet that is flooded up to a horizon threshold. If the interest expressed by the source spans more than the single data packet, the intended receiver(s) of a MR will establish and maintain a routing mesh spanning the active sources and the destination (a single node in the case of unicast and a set of nodes in the case of multicast). In the case of a multicast flow, the receivers of the multicast group run a distributed election using Mesh Announcement (MA) packets to elect a core for the group, which is the only receiver that continues to generate MAs for the group. No such election is needed for a unicast destination. An elected core or unicast destination continues sending MAs with monotonically increasing sequence numbers for as long as there is at least one active source interested in it. When no active sources are detected for a flow, the destination or core of the flow stops generating MAs, which causes the routing information corresponding to the mesh of the flow to be deleted. To save bandwidth, MAs for different unicast and multicast flows are grouped opportunistically in signaling packets. Furthermore, to confine control traffic to those portions of the network that need the information, an enclave (or region of interest) is defined for an established mesh. The enclave of a flow is a connected component of the network spanning all the receivers and sources of the flow and the relay nodes needed to connect them. The frequency with which MAs for a given flow are sent within an enclave is much higher than the frequency with which MAs are sent for a flow outside it, and depending on the flow type (e.g., bidirectional unicast or multicast) MAs are not propagated outside enclaves.

B. Mesh Activation and Deactivation

PRIME maintains routing information only for those destinations for which there is interest. Accordingly, it must activate and deactivate the routing structures (meshes) used to support unicasting and multicasting. Meshes are activated using mesh-activation requests (MR), which make receivers (unicast destinations or receivers of multicast groups) change their states from inactive to active and to start the mesh creation and maintenance process.

A MR states the type of message, an application-defined horizon threshold that is used to define the scope of the dissemination of the MR, the persistence of the interest, the sender, the intended unicast destination or multicast group, and an identifier for the message. Given that a data packet is piggybacked on an MR, only the first three fields are needed beyond what a normal data packet already specifies.

Upon reception of an MR, a node determines if it is an intended destination of the MR. If it is not, it scans for a hit in a data cache storing tuples of the form *(sender id, packet id)*. If the *(sender id, packet id)*-pair is already in the cache, or if the horizon value is reached, the MR is not forwarded. If the node is a destination of the MR, it considers itself either an active unicast destination or the core of a multicast group. If the MR states no persistence in the interest, the destination only needs to process the data packet included in the MR. If the MR states persistent interest, then the destination must start advertising its presence by establishing a mesh for the destination using mesh announcements (MA). If the destination is a multicast group, the receivers of the group participate in a distributed election that informs all the nodes located in the same connected component of the network about the active multicast group, the identity of the core of the group and of at least one next hop towards the core. The details of the core election process are described in Section III-G.

Destinations (multicast groups or unicast destinations) and relays needed between them and interested sources remain active for as long as there are active sources in the connected component of the network. Cores and unicast destinations send MAs with newer sequence numbers every mesh-announcement period (MA-period), unless they stop receiving data packets for two consecutive MA-periods. The soft state regarding the routing structures is timed out and deleted if neither MAs nor data packets are received in three consecutive MA-periods. The reception of data packets is used to distinguish between a network partition and the deactivation of multicast meshes. The details about the procedure to handle partitions is also described in Section III-G.

A node that is a receiver of a multicast group is considered to be an "inactive receiver" until it receives a MR, a MA, or a data packet for the multicast group. The reception of a MA whose destination is a multicast group by a receiver of the same group prevents the receiver from participating in the core election for the group, because the core of the group has been selected or at least agreement has started to emerge. In this case, the multicast group receiver simply accepts the core advertised in the MA and changes its state to "active receiver" of the group. If an inactive receiver receives a data packet for its multicasts group, it assumes that it has missed a MR and that the group may already be active. Hence, it delays its participation in the core election process and sends a MA for the multicast group without a proposed core, which serves as a request to its neighbors for their latest state regarding the multicast group. The receiver waits for a sensible period of time (e.g., 1 sec.) to collect MAs from its neighbors. If it receives fresher MAs for the group, it adopts the core advertised in those MAs. Otherwise, it considers itself to be the core of the group and participates in the distributed election using MAs. Lastly, if an active receiver receives an MA, it stays in the active state.

C. Mesh Establishment and Maintenance

As we have mentioned, once a destination becomes active by receiving MRs stating persistent interest from at least one source, it starts advertising its existence periodically using mesh announcements (MAs). A MA specifies seven items. The message type specifies a MA. A destination address is used to state the unicast node or multicast group. The address of the core is used to state the unicast destination itself, the core of a multicast group, the fact that the MA is a partition confirmation request, or a neighbor request; in both of the latter two cases the remaining fields have no meaning. A sequence number used to eliminate outdated MAs. A distance to the destination is included to state the distance to a core or a unicast destination. A preferred next hop to the destination that is used to prevent relays from leaving meshes prematurely. A membership code is used to indicate whether the node is a multicast mesh member, a receiver, both, or a regular node. In the case of a unicast flow, the eight bits that store the membership code are used to indicate the node's longest known distance to an active source for the unicast destination. These distances are used to route MAs back to sources that are not already included in the flow's enclave.

The MA sent by a destination (unicast destination or multicast core) states its latest sequence number and a 0 distance to itself. MAs propagate throughout the network at some rate to establish or refresh the routing structure that constitutes the mesh for the destination. Upon receiving a MA with a larger sequence number, nodes wait for a short period of time before generating their own MA that contains their current routing state regarding the destination.

The information received in valid MAs is stored in neighborhood lists. Each node selects one to three neighbors (if available) as next hops to the destination of a flow. Nodes select among their neighbors those with the larger sequence number and shortest distance to the destination. If two neighbors have the same values for these two metrics, the one with the larger identifier is selected. A node with address $addr$ considers one of its neighbors a *mesh child* if it has a shorter distance to the destination than that neighbor, and the neighbor stated that its next hop has address $addr_p$ such that $addr \geq addr_p$. For the case of multicast, nodes consider themselves *multicast mesh members* if they have at least one mesh child that is a receiver or a multicast mesh member (as stated in the membership codes of MAs).

This way, the routing mesh of a multicast group is composed of the union of its multicast mesh and a set of directed meshes that are composed of shortest paths from sources located outside of the multicast mesh to the core of the group. The routing mesh of a unicast destination is composed of the union of directed meshes that are composed of shortest paths from active sources to the unicast destination.

D. Opportunistic Transmission of MAs

To reduce the number of control packets sent and save bandwidth, nodes group MAs for different destinations opportunistically into control bundles. When a routing event is detected (namely, a change in the membership status, in the distance to the destination, or a change of next hop) nodes wait for a short period of time before transmitting the MA that informs other nodes about the change in the node's state. Any other change in the node's routing state, regarding the same or other multicast group or regarding a unicast destination that may happen during this period of time is also advertised during the transmission of the control bundle. This way, at the time the control bundle is transmitted, it includes as many MAs as groups and unicast destinations with recent updates in their associated routing states. Fig. 1 illustrates this idea. The figure also introduces a second class of routing event that is denominated as *urgent*. Urgent events, such as a change of core for multicast groups, are transmitted using a shorter delay than the one used for regular events, but the same policy of grouping MAs is applied to them. For the example of Fig. 1, all the events received between the reception of the event for multicast group 3 and the transmission of the control bundle are advertised with the transmission of the control bundle at time t'. Because no other event is received between t' and t'', no MA is transmitted at t''.

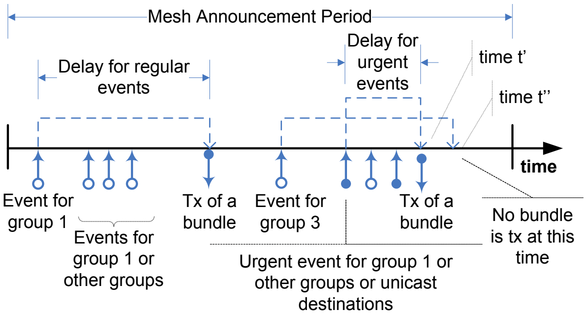

Fig. 1. Opportunistic grouping of MAs in control bundles.

E. Enclaves vs. Meshes

An enclave or region of interest of a unicast or multicast flow is a connected component of the network that contains those nodes relevant to the dissemination of information for the flow, namely receivers, senders, and relay nodes located in the paths connecting the sources to the receivers. Because all the nodes in the enclave of a flow have interest in the flow, they participate proactively in the signaling needed to maintain routing information for the flow. By the same token, nodes located outside of the enclave defined for a destination (unicast or multicast) do not participate in the process of routing data packets for that destination; hence, transmitting and receiving MAs regarding that destination is simply overhead to them.

For the case of unicast flows for a given destination, the nodes with interest in such flows are the unicast destination and the active sources with traffic for the destination. Because a unicast destination is a static singleton, nodes outside the enclave of a unicast destination simply stop the propagation of such MAs. In contrast, a multicast destination–even when represented by a core–is a dynamic set of nodes; furthermore, nodes may send to a multicast group without being part of the group. Accordingly, and to support a receiver-initiated method for multicast receivers to join groups and to let non-group members send data to multicast groups, the mesh of a multicast destination is not confined to its enclave. Instead, all nodes in the network receive information about the existence of the core for a group that has been activated by MRs, i.e., all nodes receive MAs about active multicast destinations. However, an enclave is defined for an active multicast destination that includes the sources, receivers (including the core) and relays between them. MAs are sent within a multicast enclave with much higher frequency than outside the enclave. This frequency decreases exponentially with respect to the distance in hops from a node to the boundary of the enclave.

Algorithm 1: ENCLAVE(MA)

```
if $AddressType(MA.destination) = multicast$ then
    if $rc \vee sd \vee mm \vee np$ then
    else
        if $r \bmod R = 0$ then
            $r++$;
        else
            $r++$;
            return $false$;
else
    if $np \vee sd$ then
    else
        return $false$;
return $true$;
```

Fig. 2. Pseudocode of the Eclave algorithm.

The algorithm used to decide if a node belongs to a enclave for a given destination (unicast or multicast) is presented in Fig. 2. For a unicast destination, the *Enclave* algorithm returns $true$ if the node is either a sender (sd) or a node path (np), i.e., it lies on a shortest path between a source and the destination; and $false$ otherwise. For a multicast destination, *Enclave* returns $true$ if the node is a receiver (rc), a sender (sd), or a mesh member (mm), or if the node is a node path (np), i.e., it lies on a shortest path from a sender to the core of the group. Otherwise, *Enclave* checks for the value of $r \bmod R$ and returns $true$ if it is equal to 0 and $false$ otherwise. The value of $r \bmod R$ is used to reduce the frequency with which a node located outside of the enclave transmits MAs. The value of r is initially set to 0.

We also define the *k-extended enclave* as the union of the enclave of a flow with those nodes that are located k hops away from the enclave. The objective of the k-extended enclave is to provide some degree of redundancy to cope with node mobility. Nodes located inside of the k-extended enclave forward fresh MAs with the same frequency as the nodes located inside of the enclave.

Fig. 3 presents an example of an enclave for a multicast group and its associated 1-extended enclave. Nodes labeled p, p' and p'' are part of the enclave, because they lie on shortest paths from the sender s to the core. Nodes such as w and x are part of the 1-extended enclave and *may* help to keep the enclave connected in case R_1 moves out of range of mesh member MM_1. Nodes like y, receive MAs every mesh-announcement period (MA-period), but they use *Enclave* to choose when to forward them. For instance, if R equals 2, y would send MAs at half of the frequency used inside of the enclave, while nodes located one hop away from the enclave, such as z, would send MAs at one quarter of the frequency used inside the enclave. Fig. 3 also shows a unicast enclave for source z and destination w.

The reception of MRs by nodes that are already active force them to store the address of the previous relay of the MR as well as their distance to the source of the MR. This state is used to route the next MA generated by the core or unicast destination towards the source that originated the MR. This way, the new source will acquire routes towards the destination or multicast core and will be included into the enclave. The state used to route MAs towards the source is short-lived and is deleted within a MA-period.

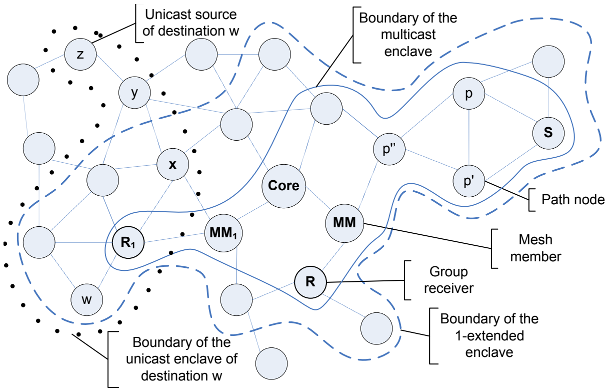

Fig. 3. Examples of a multicast enclave, its associated 1-extended enclave, and of the enclave of a unicast flow.

F. Packet Forwarding and Local Repairs

When a source has data to send, it first checks whether it has received a MA advertising the intended destination within the last three MA-periods. If not, it broadcasts a MR as described in Section III-B. Otherwise, the sender forwards the data packet according to its routing table.

Upon reception of a data packet, nodes first check for a hit in their data packet cache which stores the sender's address and sequence number of recently received data packets. If the *(sender's address, sequence number)* pair is already in the cache, the packet is silently dropped. Otherwise, the receiving node inserts the pair in its packet cache and determines whether it has to relay the data packet or not. The node also passes the packet to the upper layers if it is a receiver for the flow.

The two rules used to decide when to relay a *multicast* data packet are as follows: First, if the node is part of the multicast mesh (i.e., it has mesh children) it broadcasts the packet without further processing. Second, a node located outside of the multicast mesh relays a data packet it receives from a neighbor if it was selected by that neighbor as one of its next hops to the core. This way, and since nodes select up to three neighbors as next hops, data packets travel along directed meshes consisting of shortest paths from sources to the core of the group until they reach either the first mesh member or the core and then, the packets are flooded over the mesh of the multicast group.

Unicast data packets are also routed using directed meshes composed of shortest paths from sources to destinations. Nodes forward a unicast data packet they receive if they were selected as a next hop to the destination by the previous relay of the data packet.

Nodes located in a directed mesh employ the transmission of data packets by their next hops as implicit ACKs. If a node fails to receive three consecutive implicit ACKs from a neighbor, then it removes that node from the neighborhood list and takes one of three actions to locally repair the routing mesh:

Rapair 1: If the node is left with no paths to the core, then it broadcasts a neighbor request. Neighbor requests are replied by nodes with MAs that advertise their latest routing information regarding a given destination (unicast or multicast). This information can be further used to select a new next hop to the destination.

Repair 2: If the distance to the destination of the node increases, then it broadcasts a new MA that informs other nodes of its new state. This way, a new set of neighbors will be selected as this node's next hops and previous upstream nodes may select new nodes as their next hops to the destination.

Repair 3: If the distance to the destination of the node does not increase, then it checks its neighborhood list for other potential next hops (nodes with shorter distance) to the destination. If at least one of these potential nodes exists, then a MA is transmitted to inform the potential next hop that it has been selected as next hop. If no potential nodes are found, no further action is taken.

G. Core Election

Core elections are held only if the MR contains a multicast address. Upon reception of a MR, a group receiver first determines whether it has received a MA from the core of the multicast group within the last two MA-periods. If the node has, no further action in this regard is needed. Otherwise, the receiver considers itself the core of the group and starts transmitting MAs to its neighbors, stating itself as the core of the group. Nodes propagate MAs based on the best announcements they receive from their neighbors. A MA with a higher core ID is considered better than one with a lower core ID. Therefore, if a node receives an announcement advertising a core with a larger ID than the current core, then the new core is adopted and a new MA advertising the new core is transmitted. On the other hand, if a MA advertising a core with a smaller ID is received, then nodes check if they have recently broadcasted a MA with the current core, and if so, the MA is simply ignored. Otherwise, nodes send a MA that forces the neighbor with the smaller core to adopt the core with higher ID. Eventually, each connected component has only one core.

A core election is also held if the network is partitioned. The election is held in the connected component of the partition that does not have the old core. A node detects a partition if it does not receive a fresh MA from the core for three consecutive MA-periods and if it has received data packets within the last four MA-periods. Once a receiver detects a partition, it considers itself the core and participates in the core election.

H. Adaptive Strategies

PRIME adjusts the size and dynamics of its routing meshes depending on the perceived level of channel contention. Nodes employ information collected at the MAC layer to select the strategy that best fits the nodes' perceived channel conditions. We use channel contention as the metric to switch among operation modes in PRIME, because it has a significant impact on the performance of routing protocols that run on top of contention-based MAC protocols. To measure local contention, we use a simple and very intuitive metric based on the proportion of time in which the channel is perceived as busy. First, we define *instantaneous local contention* c as the ratio t_b/S_p, where t_b is the amount of time the channel was perceived busy during the last sampling period of S_p seconds. Then, we compute an exponential weighted moving average to avoid reacting too fast to sudden and short term changes in the instantaneous local contention and we get $\rho_n = (1 - \beta)\rho_{n-1} + \beta c$, where ρ_n is the current *level of local contention*, β is a constant used to assign weight to the level of local contention calculated at the previous sampling period (ρ_{n-1}) and, the current instantaneous local contention (c). The current value for β is 0.2. However, our simulation results showed that the performance of PRIME is not very sensitive to this parameter.

We use the following three strategies to take advantage of the information collected about the level of channel contention:

Adjust the size of the mesh: Nodes select the number of neighbors with shorter distance to the destination (if available) that are forced to join the routing meshes (multicast or directed meshes).

Adjust the mesh dynamics: Under light loads, nodes consider themselves multicast mesh members if they have had at least a mesh child during the last two MA-periods, whereas under high loads nodes consider themselves mesh members for as long as they have mesh children. The first approach leads to more stable meshes which are more resilient to mobility and attain higher delivery ratio under light loads. The second approach leads to more dynamic meshes that perform better under high loads.

Adjust timers: PRIME employs timers to check for implicit ACKs that are used to detect multicast mesh disconnections and link breakages on directed meshes. Setting adequate values for these timers is important because it allows timely actions to repair routing structures.

We defined three threshold values that are used to select among operation modes. Table I shows the actions taken depending on the value of the current level of local contention.

TABLE I
PRIME: OPERATION MODES

Level of local congestion	Operation Mode
$LOW > \rho_n$	Three parents
$MID > \rho_n$	Stable mesh
$LOW \leq \rho_n \leq HIGH$	Two parents
$HIGH < \rho_n$	One parent

The timers used for implicit ACKs are increased in a 30% when the level of local congestion passes the HIGH threshold.

IV. PERFORMANCE RESULTS

We present simulation results comparing PRIME against ODMRP and PUMA for the case of multicast traffic, as well as against AODV with ODMRP and OLSR with ODMRP for the case of combined unicast and multicast traffic. We use ODMRP, AODV, and OLSR in our experiments, because they are *de facto* baselines for performance comparisons of multicast and unicast routing protocols. PUMA was selected because it also uses core elections and meshes like PRIME, which allows us to highlight the performance benefits of the interest-based signaling used in PRIME. We use packet delivery ratio, generalized group delivery ratio, end-to-end delay, and total overhead as our performance metrics. The generalized group delivery ratio is an extension of the group reliability metric introduced in [18], in which a packet is considered as delivered, if and only if it is received by a given proportion of the receivers, for instance, if it is received by at least 80% percent of the receivers. This metric emphasizes the importance of group delivery by not considering packets that are received by a small subset of the group members.

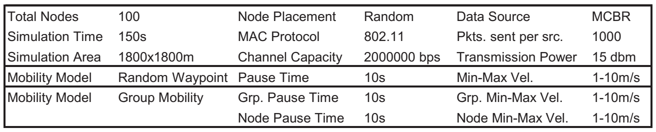

Total Nodes	100	Node Placement	Random	Data Source	MCBR
Simulation Time	150s	MAC Protocol	802.11	Pkts. sent per src.	1000
Simulation Area	1800x1800m	Channel Capacity	2000000 bps	Transmission Power	15 dbm
Mobility Model	Random Waypoint	Pause Time	10s	Min-Max Vel.	1-10m/s
Mobility Model	Group Mobility	Grp. Pause Time	10s	Grp. Min-Max Vel.	1-10m/s
		Node Pause Time	10s	Node Min-Max Vel.	1-10m/s

Fig. 4. Simulation environment.

The protocols are tested with IEEE 802.11 as the underlying MAC protocol, and all signaling packets are sent in broadcast mode for the multicast protocols. We use random waypoint and group mobility [5] as our mobility models. The first model allows us to test the protocols on general situations in which each node moves independently, and the latter models situations in which the members of a team tend to move in groups. We used the discrete event simulator Qualnet [19] version 3.9, that provides a realistic simulation of the physical layer, and well tuned versions of ODMRP, AODV and OLSR. For PUMA simulations we obtained the original code used in [17]. Each simulation was run for ten different seed values. To have meaningful comparisons, all the multicast protocols use the same period of three seconds to refresh their routing structures (join query periods for ODMRP and announcement periods for PUMA and PRIME). For ODMRP, the forwarding group timeout was set to three times the value of the join query (JQ) period, as advised by its designers. For all the experiments we set the value of k to 1. Hence, only nodes that are at most one hop away from an enclave belong to a k-extended enclave. The value of PRIME's horizon threshold was set to the same value as the TTL used in the ODMRP's JQs, which is the worst-case scenario for propagation of MRs in PRIME. Fig. 4 lists the details of the simulation environment.

A. Multicast Traffic with Increasing Number of Sources

We first focus on an experiment in which the number of concurrent active senders increases. Each sender transmits 10 packets of 256 bytes per second and the group is composed of 20 nodes. Sources are not group members. Figs. 5(a-d) present the results for the random waypoint mobility model. Fig. 5(a) shows the delivery ratio attained by the different protocols when the number of concurrent sources is increased. We observe that PRIME performs similar to or better than the other protocols. From Figs. 5(a, b and d) we observe that ODMRP performs particularly well for small numbers of sources but the other protocols scale better thanks to their reduced overhead. For up to 14 sources, PRIME attains higher delivery ratios than PUMA due mainly to the higher reliability of the directed meshes used to route packets from sources located outside of the multicast mesh to the core, in comparison with the single paths used by PUMA. In addition, the meshes established by PRIME are also locally repaired, which provides extra reliability. This situation is more evident when the group mobility model is used (Figs. 5(e-h)), where due to the physical proximity of group members, the paths from sources located outside of the multicast mesh tend to be longer than the ones observed with random waypoint, where the group members are spread all over the simulation area. Fig. 5(b) presents the generalized group delivery ratio attained by the protocols when the delivery threshold is set to 80%. In this case, we notice that, even when the delivery ratios attained by PUMA and PRIME are better that the one attained by ODMRP for 8 or more sources, the generalized group delivery ratios are not. This is due to the fact that under heavily loaded networks, maintaining large routing structures like the ones observed when random placement and random waypoint are used becomes a very hard task. We often observe weakly connected routing structures. In this scenario, the collaborative way in which sources build the ODMRP mesh (composed of the union of the meshes of the active sources) helps to cope with this problem. Fig. 5(c) shows that PUMA and PRIME achieve considerably less end-to-end delay than ODMRP. The increased delay shown by ODMRP is mainly due to the amount of packets that are injected into the network (Fig. 5(d)) which leads to high levels of congestion.

Figs. 5(e-h) present results for the group mobility model in which the 20 nodes that belong to the multicast group move around inside of a square region of $900 \times 900m$. In the group mobility model, each group decides its group mobility direction and speed randomly. Each node then decides its internal mobility randomly and computes its actual mobility by

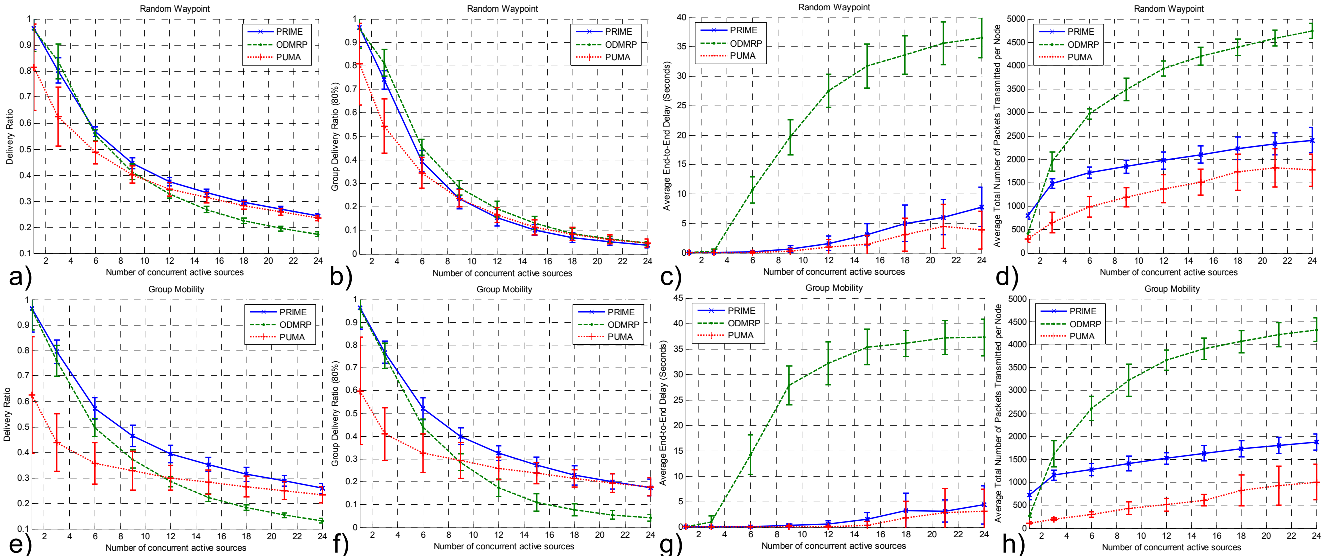

Fig. 5. Performance with increasing number of MCBR sources. (a-d) Random waypoint mobility model. (e-h) Group mobility model. (a and e) Delivery ratio. (b and f) Group delivery ratio. (c and g) End-to-End delay. (d and h) Average total number of packets transmitted per node.

summing the two mobility vectors [5]. The remaining 80 nodes move following the random waypoint mobility model. From Figs. 5(e-f) we can observe that PRIME clearly outperforms the other protocols in both delivery ratio and generalized delivery ratio. It is interesting to observe that, constrasting with the previous case in which random waypoint was used, the collaborative construction of the ODMRP mesh does not help to improve the generalized delivery ratio when group members move following the group mobility model. This result is intuitive. Given that receivers tend to be concentrated in a particular region of the simulation area, the establishment of a mesh by a source that is located at the opposite side of a second source has just a marginal benefit for the establishment of the mesh of the second source. Moreover, data packets generated by a given source are also routed towards the other sources in ODMRP; hence, the concentration of group members also increases the probability of routing packets towards places where no receiver is located. For the end-to-end delay (Fig. 5(g)) and total overhead (Fig. 5(h)), the three protocols show behaviors similar to the ones observed in the previous case.

B. Multicast Traffic with Increasing Number of Groups

The second set of experiments evaluates the performance of the routing protocols as the number of concurrent active multicast groups increases. These scenarios try to model situations where the interaction among team members is the predominant communication pattern; hence, sources are also group members. Group members follow the group mobility model, whereas the remaining nodes move according to the random waypoint model. Fig. 6 presents the results obtained when the group members are located inside of a square area of $600 \times 600m$. From Figs. 6(a-b) we can observe that PRIME attains similar or better delivery and generalized delivery ratios than the other protocols when each group has one active source. In these two figures we can also notice that the strategy of PUMA of reducing as much as possible the number of control packets is not well suited to lightly loaded networks, where the available bandwidth can be used to establish more robust routing structures. The delay attained by the different protocols is shown in Fig. 6(c) and the overhead is shown in Fig. 6(d). In this scenario, sources are also group members and where groups have only one active source, the three routing protocols should establish similar routing structures (source-based threes). This allows us to highlight the benefits of using enclaves, adaptive meshes and the concept of control bundles. Figs. 6(e-h) present the performance of the protocols when the number of active sources per group is increased to 3. As in the previous case, PRIME attains similar or higher delivery and generalized delivery ratio than the other protocols (Figs. 6(e-f), in particular for four or more groups. For the end-to-end delay, this scenario with increased traffic load is particularly disadvantageous for ODMRP. Fig. 6(g) shows that, as the number of groups increases, the delay attained by PUMA and PRIME is close to an order of magnitude smaller that the one attained by ODMRP. Again, this is due to the extra overhead incurred by ODMRP, as can be seen in Fig. 6(h).

Fig. 7 presents the results obtained when the group members are located inside of a square area of $900 \times 900m$. For one (Figs. 7(a-b)) and three (Figs. 7(e-f)) sources per group PRIME attains similar or better delivery and generalized delivery ratios than the remaining protocols. In addition, PRIME also attains the lowest delays, as it is shown in Figs. 7(c and g). Lastly, this scenario clearly shows how PRIME adjusts its overhead to the current conditions of the network by adapting how control

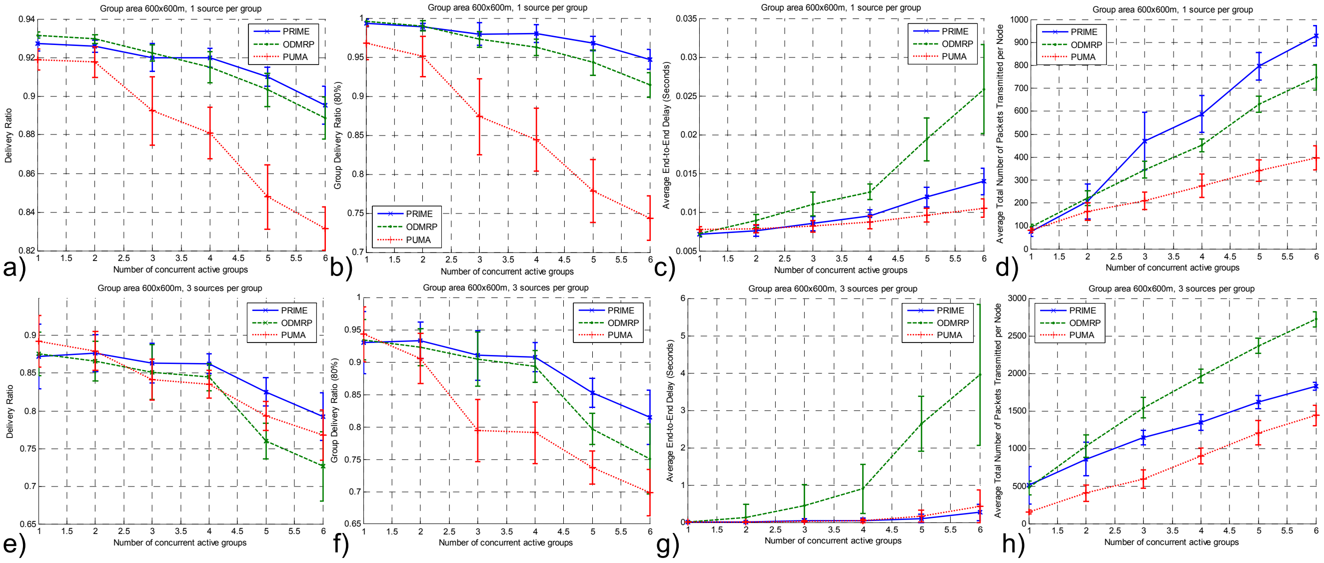

Fig. 6. Performance with increasing number of active groups and group areas of 600 × 600m. a-d) One source per group. e-h) Three sources per group. (a and e) Delivery ratio. (b and f) Group delivery ratio. (c and g) End-to-End delay. (d and h) Average total number of packets transmitted per node.

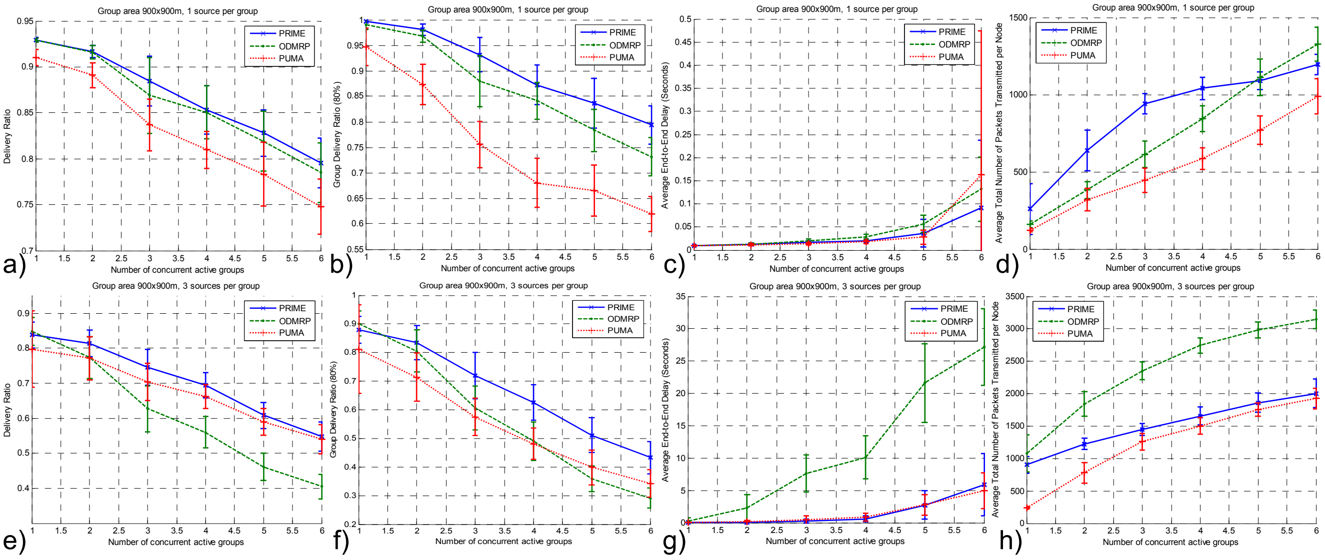

Fig. 7. Performance with increasing number of active groups and group areas of 900 × 900m. a-d) One source per group. e-h) Three sources per group. (a and e) Delivery ratio. (b and f) Group delivery ratio. (c and g) End-to-End delay. (d and h) Average total number of packets transmitted per node.

and data packets are sent according to perceived congestion (Figs. 7(d and h)).

C. Combined Multicast and Unicast Traffic

This set of experiments evaluates the performance of the routing protocols in a scenario with combined multicast and unicast traffic. We use the same settings as in the prior experiment but with the addition of 5 CBR flows between nodes that do not belong to a multicast group. Unicast sources send a total of 1000 data packets of 256 bytes at a rate of 10 packets per second. The results are shown in Figs. 8(a-d). From Fig. 8(a) we can observe that PRIME attains higher delivery ratios than the other protocols for both unicast and multicast traffic. PRIME delivers up to 10% more data packets than AODV and up to 20% more than OLSR, and at the same time, up to 10% more multicast data packets than ODMRP when it is used in conjunction with OLSR and up to 20% when it is coexisting with AODV. PRIME also attains higher generalized group delivery ratios than ODMRP for more than 1 group (Fig. 8(b)) and the lowest delays for both unicast and multicast traffic (Fig. 8(c)), while incurring far less control (CO) and total overhead (TO) than the other protocols (Fig.

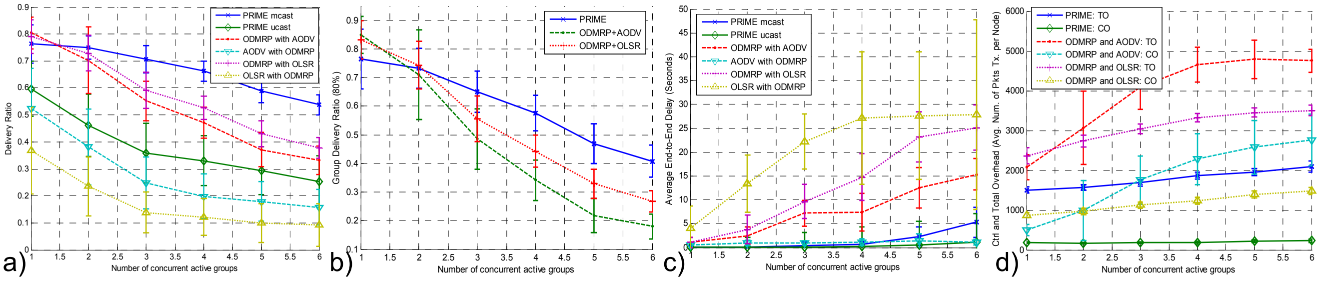

Fig. 8. Performance with increasing number of active groups, 3 sources per group, group areas of 900 × 900m and five unicast flows. (a) Delivery ratio. (b) Group delivery ratio. (c) End-to-End delay. (d) Average total number of total and control packets transmitted per node.

8(d)), almost five times less than OLSR+ODMRP and nine times less than AODV+ODMRP.

V. CONCLUSIONS

We have shown by example that it is possible and perhaps desirable to support the dissemination of information for end user applications using a single routing protocol, and that interest-driven routing should be adopted for MANETs, instead of attempting to provide either on-demand or proactive routing. We introduced the Protocol for Routing in Interest-defined Mesh Enclaves (PRIME). PRIME redefines how signaling is done for routing in MANETs by integrating unicast and multicast routing using interest-driven establishment of meshes and enclaves. PRIME establishes meshes (connected components of a MANET) that are activated and deactivated by the presence or absence of data traffic. Enclaves confine most of the dissemination of control packets to those that actually need the information. This property has a positive impact over the scalability of the protocol, particularly in medium to large networks in which the members of the same multicast group tend to be close by. The results of a series of simulation experiments illustrate that PRIME attains higher delivery ratios than ODMRP and PUMA for multicast traffic, and higher delivery ratios than AODV and OLSR for unicast traffic. At the same time, PRIME induces much less communication overhead and attains lower delays than the other routing protocols.

ACKNOWLEDGMENTS

Work partially sponsored by the U.S. Army Research Office (ARO) under grants W911NF-05-1-0246 and SC20070363, by the National Science Foundation under grant CNS-0435522, by DARPA through Air Force Research Laboratory (AFRL) Contract FA8750-07-C-0169, by the Baskin Chair of Computer Engineering, by the UC MEXUS-CONACyT program and by the Mexican National Polytechnic Institute.

REFERENCES

[1] S. K. Das, B. S. B. S. Manoj, and C. S. R. Murthy. A dynamic core based multicast routing protocol for ad hoc wireless networks. In *MobiHoc '02: Proc. of the 3rd ACM intl. symp. on Mob. ad hoc net. & comp.*, pages 24–35. ACM, 2002.

[2] V. Devarapalli and D. Sidhu. Mzr: A multicast protocol for mobile ad hoc networks. In *Proc. of the IEEE Intl. Conf. on Comm., 2001. ICC 2001.*, volume 3, pages 886–891 vol.3, 2001.

[3] J. J. Garcia-Luna-Aceves and E. L. Madruga. The core-assisted mesh protocol. *IEEE Journal on Selected Areas in Communications*, 17(8):1380–1394, Aug 1999.

[4] Z. J. Haas, M. R. Pearlman, and P. Samar. The zone routing protocol (zrp) for ad hoc networks.

[5] X. Hong, M. Gerla, G. Pei, and C.-C. Chiang. A group mobility model for ad hoc wireless networks. In *Proc. of ACM/IEEE MSWiM'99*, pages 53–60, 1999.

[6] C. Intanagonwiwat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva. Directed diffusion for wireless sensor networking. *IEEE/ACM Trans. Netw.*, 11(1):2–16, 2003.

[7] P. Jacquet, A. Laouiti, P. Minet, and L. Viennot. Performance analysis of OLSR multipoint relay flooding in two ad hoc wireless network models. In *The second IFIP-TC6 NETWORKING Conference*, may 2002.

[8] J. G. Jetcheva and D. B. Johnson. Adaptive demand-driven multicast routing in multi-hop wireless ad hoc networks. In *MobiHoc '01: Proc. of the 2nd ACM intl. symp. on Mob. ad hoc net. & comp.*, pages 33–44. ACM, 2001.

[9] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad hoc wireless networks. *Mobile Computing*, 353:153–181, 1996.

[10] S. Lee and C. Kim. Neighbor supporting ad hoc multicast routing protocol. In *MobiHoc '00: Proc. of the 1st ACM intl. symp. on Mob. ad hoc net. & comp.*, pages 37–44. ACM, 2000.

[11] S.-J. Lee, M. Gerla, and C.-C. Chiang. On-demand multicast routing protocol. In *Proc. of the IEEE Wireless Comm. and Net. Conf., 1999. WCNC.*, pages 1298–1302 vol.3, 1999.

[12] S. Murthy and J. J. Garcia-Luna-Aceves. An efficient routing protocol for wireless networks. *Mob. Netw. Appl.*, 1(2):183–197, 1996.

[13] C. E. Perkins and E. M. Royer. Ad-hoc on-demand distance vector routing. In *Proc. of the Second IEEE Workshop on Mob. Comp. Syst. and App., 1999. WMCSA '99.*, pages 90–100, Feb 1999.

[14] S. Roy and J. J. Garcia-Luna-Aceves. Node-centric hybrid routing for ad hoc networks. In *Proc. of 10th IEEE/ACM MASCOTS 2002, Workshop on Mobility and Wireless Access*, pages 63–71, October 12, 2002.

[15] E. M. Royer and C. E. Perkins. Multicast operation of the ad-hoc on-demand distance vector routing protocol. In *MobiCom '99: Proc. of the 5th annual ACM/IEEE intl. conf. on Mob. comp. and net.*, pages 207–218. ACM, 1999.

[16] P. M. Ruiz and A. F. Gomez-Skarmeta. Reducing data-overhead of mesh-based ad hoc multicast routing protocols by steiner tree meshes. In *Proc. of IEEE SECON 2004.*, pages 54–62, 4-7 Oct. 2004.

[17] R. Vaishampayan and J. J. Garcia-Luna-Aceves. Efficient and robust multicast routing in mobile ad hoc networks. In *Proc. of the IEEE Conf. on Mob. Ad-hoc and Sensor Syst., 2004*, pages 304–313, Oct. 2004.

[18] K. Viswanath, K. Obraczka, and G. Tsudik. Exploring mesh and tree-based multicast routing protocols for manets. *IEEE Transactions on Mobile Computing*, 5(1):28–42, 2006.

[19] Qualnet 3.9, scalable network technologies: http://www.scalablenetworks.com.