UC Santa Cruz
UC Santa Cruz Previously Published Works

Title

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

Permalink

https://escholarship.org/uc/item/4rn4h5g9

Author

Garcia-Luna-Aceves, J.J.

Publication Date

2011-03-01

Peer reviewed

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

IEEE/ACM TRANSACTIONS ON NETWORKING

1

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

J. J. Garcia-Luna-Aceves, *Fellow, IEEE, ACM*, and Rolando Menchaca-Mendez

Abstract—A framework for integrated multicast and unicast routing in mobile ad hoc networks (MANETs) is introduced. It is based on interest-defined mesh enclaves that are connected components of a MANET spanning the sources and receivers of unicast or multicast flows. The Protocol for Routing in Interest-defined Mesh Enclaves (PRIME) is presented to implement the proposed framework for integrated routing in MANETs. PRIME establishes meshes that are activated and deactivated by the presence or absence of interest in individual destination nodes and groups and confines most of the signaling overhead within regions of interest (enclaves) in such meshes. The routes established in PRIME are shown to be free of permanent loops. Experimental results based on extensive 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). The experiments also show that signaling in PRIME is far more scalable than the one used by traditional multicast and unicast routing protocols such as AODV, OLSR, or ODMRP.

Index Terms—Interest-driven routing, mobile ad hoc network (MANET), multicast routing, unicast routing.

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, and hence supporting point-to-point and many-to-many communication in a MANET requires running a unicast and a multicast routing protocol in

J. J. Garcia-Luna-Aceves is with the Department of Computer Engineering, University of California, Santa Cruz, Santa Cruz, CA 95064 USA, and also with the Palo Alto Research Center, Palo Alto, CA 94304 USA (e-mail: jj@soe.ucsc.edu).

R. Menchaca-Mendez is with the Computer Research Center, National Polytechnic Institute, Mexico City 07738, Mexico (e-mail: rmen@cic.ipn.mx).

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.

The main contribution of this paper is the introduction of 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) [12], which implements our integrated routing framework. In PRIME, the routes 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 the signaling needed 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 shows that the routes established and maintained with PRIME are free of long-term loops.

Section V describes the results of simulation experiments used to compare the performance of PRIME to that of relevant multicast and unicast routing protocols for MANETs. Our comparison addresses the performance of routing protocols in support of unicast traffic, multicast traffic, and the combination of unicast and multicast traffic. Although PRIME supports routing that takes into account link and node characteristics, the focus of our study is on minimum-hop routing due to space limitations and the fact that minimum-hop routing is still the baseline for unicast and multicast routing in the MANET routing protocols proposed in the IETF. We compare PRIME to ODMRP [11] and PUMA [18] to determine the effectiveness of PRIME as a multicast routing protocol and consider different numbers of sources, groups, and the use of group and random waypoint mobility models. For the case of unicast traffic, we compare PRIME against OLSR [7] and AODV [15] and vary the number of concurrent unicast flows and the node density. We compare PRIME against the use of AODV and ODMRP and the use of OLSR

and ODMRP in medium-size scenarios (100 nodes) as well as in a large scenario (1000 nodes) to consider the case of combined unicast and multicast traffic. The results show that PRIME is a very efficient multicast and unicast 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. Lastly, we present simulation results illustrating that the performance of PRIME improves when link quality and queue lengths are taken into account in the selection of routes within enclaves. The limited results we present provide some justification to the intuition that PRIME benefits much more from the use of link and node parameters in route selection than traditional on-demand or proactive routing protocols because it contains route signaling while selecting less congested and more robust paths. However, a detailed study of quality-of-service routing and the use of link and node parameters for route selection in the context of interest-driven routing enclaves compared to traditional routing is beyond the scope of this paper. This type of study is the subject of additional papers (e.g., [4] and [13]) and future work.

II. Related Work

There have been a large number of routing protocols proposed and implemented to date for MANETs. However, due to space limitations, we can only address a very small sample of them and consider only the case of minimum-hop routing, which is still the baseline for the routing protocols proposed in the IETF. Our summary is intended simply to highlight the facts that: 1) existing routing protocols for MANETs support either unicast routing or multicast routing; and 2) 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 [14]) 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 [15] and 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 [5] and NEST [16]). In these hybrid schemes, however, proactive signaling is applied within areas or zones of the network [5] independently of the interest for destinations in such zones, or for specific destinations [16] and on-demand signaling propagates throughout the network.

Multicast routing protocols can be 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) [17] maintains a shared tree for each multicast group consisting of receivers and relays. Sources acquire routes to multicast groups on-demand in a way similar to the *ad hoc* on-demand distance vector protocol (AODV) [15]. 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-based tree is maintained by periodic keep-alive packets from the source. Like ADMR, MZR [1] maintains source-based trees, but performs zonal routing, and hence the dissemination of control packets is less expensive.

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) [2], and the Protocol for Unified Multicasting through Announcements (PUMA) [18].

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, the latter 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. However, the salient feature of ODMRP and its extensions is that multiple nodes produce some flooding for each group.

CAMP [2] 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 toward a core of the desired group. PUMA [18] also uses a receiver-initiated approach in which receivers join a multicast group using the address of a core that is broadcast to the network proactively. 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 the interest nodes may have in the group.

III. PRIME

A. Overview

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) 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 an MR establish and maintain a routing mesh spanning the active sources and the destination (a single node if the destination is unicast, and a dynamic set of nodes if the destination is multicast). In the case of a multicast flow, the receivers of the multicast group run a distributed election using *mesh announcements* (MAs) 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 after a finite time, 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. 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 the flows of a destination is a connected component of the network spanning all the receivers that compose the destination, the interested active sources, and the relay nodes needed to connect them. The frequency with which MAs for a given destination are sent within its enclave is much higher than the frequency with which MAs are sent outside it, and depending on the flow type (e.g., bidirectional unicast or multicast), MAs are not propagated outside enclaves.

In order to integrate the signaling for unicast and multicast routing, destinations are treated as connected *destination meshes* containing one or more nodes. In the case of a unicast data flow, a destination D is a singleton that contains a node with identifier D. By contrast, in the case of a multicast data flow, destination D contains the members of a multicast group as well as a set of nodes needed to keep D connected. We refer to this set of nodes used to maintain D connected as the multicast mesh of D (MM_D). Destinations (multicast groups or unicast destinations) and relays needed between them remain active for as long as there are interested 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 four consecutive MA-periods. The reception of data packets is used to distinguish between a network partition and the deactivation of multicast meshes.

It is important to note that although PRIME also uses meshes and a similar core-election mechanism to that introduced in PUMA for receiver-initiated multicast routing, the two protocols are very different because of the use of enclaves in PRIME (see Section III-F). In particular, PRIME integrates unicast and multicast routing into the same signaling, which is unique to
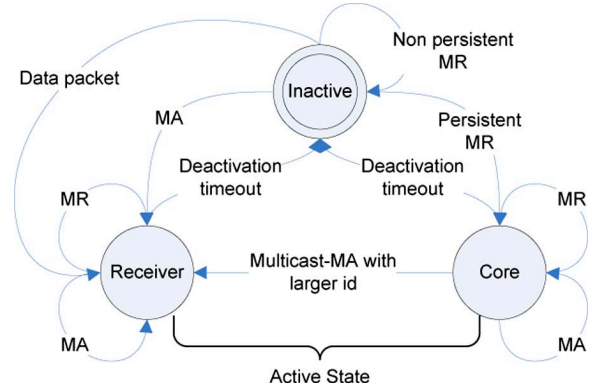

Fig. 1. Finite state automaton of the activation and deactivation process.

PRIME, establishes meshes and enclaves for destinations on demand, uses MAs primarily within enclaves, and disseminates MAs with exponentially decreasing rates outside enclaves for the case of multicasting and uses adaptive strategies (see Section III-I) and context awareness (see Section III-J) to make its signaling more efficient.

B. Mesh Activation and Deactivation

PRIME maintains routing information only for those destinations for which there is interest. Accordingly, it activates and deactivates the routing structures (meshes) used to support unicasting and multicasting. Meshes are activated using mesh-activation requests (MRs), 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 mesh-activation request (MR_D^S) generated by source S for destination D and transmitted by node B is a seven-tuple of the form

$$(\mathrm{type},\ \mathrm{horizon},\ \mathrm{persistent},\ \mathrm{id}^S, d_S^B, \mathrm{id}^D, sn^S)$$

where type states the type of message, $\mathrm{horizon}$ is an application-defined threshold used to limit the dissemination of the MR up to a preset horizon, $\mathrm{persistent} \in \{\mathrm{true}, \mathrm{false}\}$ is a flag that indicates the persistence of the interest, id^S is the sender's identifier, d_S^B is B's distance to the sender S, id^D is the destination's identifier (unicast destination or multicast group), and sn^S is an identifier for the message. The first MR sent by a source has a data packet piggybacked with it.

Fig. 1 shows the finite state automation that summarizes the transitions between active and inactive states, which are dictated by the reception of MRs, MAs, and, for the case of multicasting, the election of a core.

Upon reception of an MR, a node determines whether it is an intended destination of the MR. If it is not, it scans for a hit in a data cache storing pairs of the form (*sender id*, *packet id*). If the pair is already in the cache, or if the horizon value of the MR 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. Otherwise, the destination

must start advertising its presence by establishing a mesh using MAs.

Nodes also store their distances to the source of the MR, which are computed from the distance values to the sources of the MRs stated in the MRs themselves. In the case of unicast destinations, this information is used to route mesh announcements back to the source and in general to include new sources into the destination enclaves (see Sections III-C and III-F).

The mesh activation and deactivation process also takes into account the election of cores for the case of multicasting. 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 one or more next-hops toward the core (see Section III-H).

A node that is a receiver of a multicast group is considered to be an "inactive receiver" until it receives an MR, an MA, or a data packet for the multicast group. The reception of an MA for a destination that is a multicast group by an inactive 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 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 multicast group, it assumes that it has missed an MR and that the group may already be active. Hence, it delays its participation in the core election process and sends an 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 short period of time (e.g., 1 s) 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 stated, meshes are maintained by means of *mesh announcements*. An (MA_D^{*B}) transmitted by node B for destination D is a six-tuple of the form

$$(\mathrm{id}^{*B}, \mathrm{core}_D^{*B}, \mathrm{sn}_D^{*B}, d_D^{*B}, \mathrm{mm}_D^{*B}, \mathrm{next}_D^{*B})$$

where id^{*B} is the identifier of B, and core_D^{*B} is either the identifier of the core of the multicast group D known by B or the identifier of the unicast destination. sn_D^{*B} is the largest sequence number known by B for destination D, d_D^{*B} is the distance of B to the core of D or to the destination itself in the case of a unicast destination. mm_D^{*B} is a multicast-specific flag that indicates if B is a mesh member, a multicast group member, both, or a regular node. In the case of a unicast flow, the 8 b 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. next_D^{*B} is the identifier of the preferred next-hop of node B toward the core of D. The core

of a unicast destination is always the destination itself (it is the only element of the singleton).

For a given destination D, nodes maintain a neighborhood list L_D that stores an ordered set composed of the MAs that the node has recently received from each of its neighbors regarding that destination. To differentiate from an MA that has just been received, an MA received from neighbor B that is already stored in L_D is denoted by MA_D^B (with the $*$ dropped).

Each MA stored in L_D is augmented with a timestamp (ts) obtained from the local clock. Announcements are ordered using a strict total order relation $\prec$, which is defined as follows:

$$
\begin{aligned}
\mathrm{MA}_D^B \prec \mathrm{MA}_D^A \Leftrightarrow & (\mathrm{sn}_D^B < \mathrm{sn}_D^A) \\
& \vee (\mathrm{sn}_D^B = \mathrm{sn}_D^A \wedge d_D^B > d_D^A) \\
& \vee (\mathrm{sn}_D^B = \mathrm{sn}_D^A \wedge d_D^B = d_D^A \wedge \mathrm{id}^B < \mathrm{id}^A). \quad (1)
\end{aligned}
$$

The proof that $\prec$ is antireflexive, asymmetric, and transitive is omitted for the sake of brevity.

In addition to L_D, a node x keeps track of the core of the destination (core_D^x), which is dynamic in the case of a multicast destination, the largest sequence number known for the destination (sn_D^x), its current distance to the core of the destination (d_D^x), its *feasible distance* to the core of the destination (fd_D^x), its preferred next-hop toward the core (next_D^x), and its mesh membership status flag mm_D^x that in the case of a unicast destination has no meaning. Nodes also keep track of the longest known distance to any source (ds_D^x). The initial value of the routing state is as follows: $L_D \leftarrow \emptyset$, $\mathrm{core}_D^x \leftarrow \mathrm{next}_D^x \leftarrow \mathrm{nil}$, $\mathrm{sn}_D^x \leftarrow 0, d_D^x \leftarrow \mathrm{fd}_D^x \leftarrow \infty, \quad \mathrm{ds}_D^x \leftarrow \infty, \quad \text{and} \quad \mathrm{mm}_D^x \leftarrow \mathrm{REG}$ (regular node).

D. Processing Mesh Announcements

Upon reception of MA_D^{*B} from neighbor B for destination D, node x updates its routing information using the following formulas.

Node x accepts the MA if it contains a sequence number equal to or larger than the current largest sequence number stored at node x, or if it is the first time that an MA is received from B (2). The MA is dropped otherwise (2)

$$
L_D \leftarrow \begin{cases}
L_D \cup \{\mathrm{MA}_D^{*B}\}, & \text{if } \mathrm{MA}_D^B \notin L_D \\
L_D - \{\mathrm{MA}_D^B\} \cup \{\mathrm{MA}_D^{*B}\}, & \text{if } \mathrm{sn}_D^x \leq \mathrm{sn}_D^{*B} \quad (2) \\
L_D, & \text{if } \mathrm{sn}_D^x > \mathrm{sn}_D^{*B}.
\end{cases}
$$

The feasible distance to the core of x (fd_D^x) is a nonincreasing function over time that can only be reset by a change of core or by a new sequence number (3). Feasible distances are used to select a *feasible set* of next-hops toward the core of the destination. As we show in Section IV, the way in which nodes select their next-hops suffices to guarantee instantaneous loop freedom

$$
\mathrm{fd}_D^x \leftarrow \begin{cases}
d_D^{*B}, & \text{if } \mathrm{sn}_D^{*B} > \mathrm{sn}_D^x \\
\min\{\mathrm{fd}_D^x, d_D^{*B}\}, & \text{if } \mathrm{sn}_D^{*B} = \mathrm{sn}_D^x \quad (3) \\
\mathrm{fd}_D^x, & \text{otherwise}.
\end{cases}
$$

The sequence number stored at node x for the core of destination D (sn_D^x) is a strictly increasing function over time that can only be reset by a change of core (4). The $\max$ function of (4) is the traditional $\max$ function defined over the natural numbers. PRIME uses sequence numbers consisting of a core-specific

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

GARCIA-LUNA-ACEVES AND MENCHACA-MENDEZ: PRIME 5

timestamp taken from the real-time clock of the node and an unsigned monotonically increasing counter. When the counter reaches its maximum value, the node uses a new timestamp in its sequence number and resets the counter to zero

$$\mathrm{sn}_D^x \leftarrow \max\{\mathrm{sn}_D^x, \mathrm{sn}_D^{*B}\}. \tag{4}$$

The distance to the core of destination D of node x (d_D^x) is computed using (5) and the relation $\prec$ defined in (1). By definition, the core of the group has a 0 distance to itself, and its feasible distance is always 0. Every link cost (lc) is assumed to be a positive real value

$$d_D^x \leftarrow \begin{cases} d_D^i + \mathrm{lc}_i^x : \max\limits_{i \in L_D : \mathrm{sn}_D^i = \mathrm{sn}_D^x} \{i\}, & \text{if such } i \text{ exists} \\ \infty, & \text{otherwise.} \end{cases} \tag{5}$$

The next-hop to the core of D (next_D^x) is also computed using the relation $\prec$ defined in (1), the current values of the feasible distance, and the sequence number

$$\mathrm{next}_D^x \leftarrow \begin{cases} \mathrm{id}^i : \max\limits_{i \in F_D^x} \{i\}, & \text{if such } i \text{ exists} \\ \mathrm{nil}, & \text{otherwise} \end{cases} \tag{6}$$

where $F_D^x = \{i : i \in L_D \wedge \mathrm{fd}_D^x = d_D^i \wedge \mathrm{sn}_D^i = \mathrm{sn}_D^x\}$ is the set of x's feasible neighbors for destination D.

Lastly, if a node x receives an MA advertising a core with a larger identifier ($\mathrm{core}_D^{*B} > \mathrm{core}_D^x$), then L_D is set to $\{\mathrm{MA}_D^{*B}\}$, core_D^x is set to core_D^{*B}, and the other parameters are set as follows: fd_D^x to d_D^{*B}, d_D^x to $d_D^{*B} + \mathrm{lc}_B^x$, sn_D^x to sn_D^{*B}, and next_D^x to id^B. Otherwise, if $\mathrm{core}_D^{*B} < \mathrm{core}_D^x$, then the MA is simply discarded.

The mesh membership flag $\mathrm{mm}_D^x \in \{\mathrm{RM}, \mathrm{MM}, \mathrm{RCV}, \mathrm{REG}, \mathrm{NIL}\}$ indicates whether D is a unicast destination (NIL) or a multicast destination, in which case it also states whether x is a regular node (REG), a group receiver (RCV), a mesh member (MM), or both group receiver and mesh member (RM).

A node x is a mesh member if and only if

$$\exists y \in L_D : \mathrm{mm}_D^y \neq \mathrm{REG} \wedge \mathrm{mm}_D^y \neq \mathrm{NIL} \wedge d_D^y$$
$$> d_D^x \wedge \mathrm{next}_D^y \leq \mathrm{id}^x \wedge \mathrm{ts}_D^y + \mathrm{MA\_period} \geq \mathrm{ct} \tag{7}$$

where ts_D^y is the timestamp assigned to y when it was stored in L_D, ct is the current value of the clock of x, and $\mathrm{MA\_period}$ is the value of the MA-period.

E. Transmission of Mesh Announcements

Nodes transmit MAs to inform other nodes about updates in their routing state for a destination. These updates can be originated by such internal events as a change in the group membership status (a node joining or leaving a multicast group) that modify the value of mm_D^x and the generation of a new sequence number in the case of the core, or by such an external event as the reception of an MR generated by a source that has just become active or the reception of a MA_D^{*B} from a neighbor B. Therefore, whenever the core of the destination generates a new MA with a larger sequence number, the latter is disseminated along the network advertising the new sequence number (4) and establishing next-hop pointers toward the core [core_D, (6)]. The

mesh composed of these next-hop pointers from a given source to the core_D is the *routing mesh* of that source.

In the case of a multicast destination, an MA transmitted by a multicast group member R forces the next-hop of R $(n = \mathrm{next}_D^R)$ to update its mesh membership status according to (7). If this changes the value of mm_D^n, then n transmits a new MA to advertise its new state. This way, nodes that lay in a path $p = R, n, n_1, \ldots, n_k, \mathrm{core}$ with $\mathrm{next}_D^R = n, \mathrm{next}_D^n = n_1, \ldots, \mathrm{next}_D^{n_k} = \mathrm{core}_D$ are forced to become multicast-mesh members, creating a connected component that contains all the receivers of a multicast group and that we call the *multicast destination*. The set of nodes $M^D = \{y : \mathrm{mm}_D^y = \mathrm{MR} \vee \mathrm{mm}_D^y = \mathrm{MM}\}$ form the *multicast mesh* of the multicast destination D.

To augment the reliability of routing, nodes can force more than one next-hop (if available) to become mesh members or to route packets to the destination. For instance, to force two next-hops to join a multicast mesh, (6) has to be modified so that it returns the second to the maximum $i \in L_D : \mathrm{fd}_D^x = d_D^i \wedge \mathrm{sn}_D^i = \mathrm{sn}_D^x$.

To reduce the number of control packets transmitted 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 (e.g., 50 ms) 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 that occurs during this period of time is also advertised in the control bundle. Accordingly, when the control bundle is transmitted, it may include many MAs regarding multicast groups and unicast destinations with recent updates in their associated routing states.

F. Enclaves Versus Meshes

An enclave or region of interest of a unicast or multicast destination D is a connected component of the network that contains those nodes relevant to the dissemination of information for the flows with destination D, namely nodes that compose the destination, senders, and relay nodes located in the paths connecting the sources to the destination. Because all the nodes in the enclave of a destination have interest in the destination, they participate proactively in the signaling needed to maintain routing information for the destination. 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 unnecessary overhead for them.

For the case of unicast flows for a given destination, the nodes with interest in such flows are the unicast destination, the active sources with traffic for the destination, and nodes laying in shortest paths from sources to 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.

A multicast destination is a dynamic set of nodes, even though it is represented by a core. Furthermore, nodes may send to a multicast group without being part of the group. Accordingly, to support a receiver-initiated method for multicast receivers to

join multicast groups and to allow non-group members to 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 multicast 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, destination (including receivers and the multicast mesh), 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)

---

1 **if** AddressType(MA.D) = multicast **then**
2   **if** rc $\vee$ sd $\vee$ mm $\vee$ pn **then**
3   **else**
4   **if** $r \bmod R = 0$ **then**
5        $r + +$;
6   **else**
7        $r + +$;
8   return false;
9 **else**
10   **if** pn $\vee$ sd **then**
11   **else**
12   return false;
13 return true;

Regions of interest are established with the support of the data packets that flow from sources to destinations. Algorithm 1 is used to decide if a node has to relay an MA for destination D. For a unicast destination, the *Enclave* algorithm returns true if the node is either a sender (sd) or a path node (pn), i.e., if according to (8), it has relayed a packet transmitted from any source to D during the previous MA-period seconds, and false otherwise. For a multicast destination, *Enclave* returns true if the node is a receiver (rc), a sender (sd), a mesh member (mm), or if the node is a path node (pn). As in the previous case, a node is a path node if, according to (8), it has relayed a packet transmitted from any source to D during the last MA-period. 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 *1-extended enclave* as the union of the enclave of a flow with those nodes that are located 1 hop away from the enclave. The objective of the 1-extended enclave is to provide some degree of redundancy to cope with node mobility. Nodes located inside of the 1-extended enclave forward fresh MAs with the same frequency as the nodes located inside of the enclave.

Fig. 2 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 lay in shortest paths from the sender s to the core and have recently been used as relays. Nodes like w and x are part of the 1-extended enclave
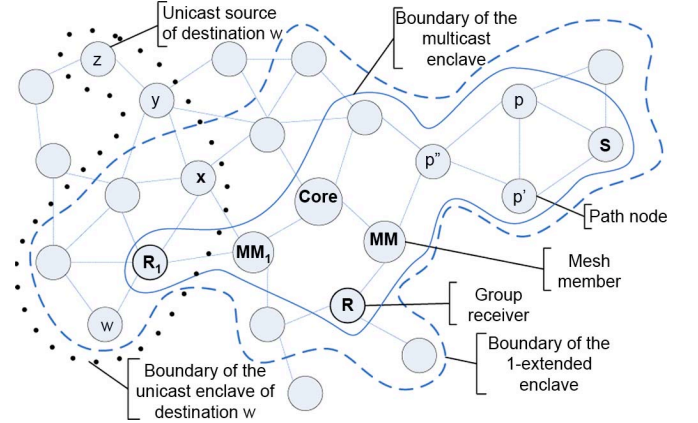

Fig. 2. Examples of a multicast enclave, its associated 1-extended enclave, and the enclave of a unicast flow.

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 (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. 2 also shows a unicast enclave for source z and destination w.

A multicast source S that becomes active and that is located outside of the enclave of its intended destination D can take advantage of the state established by the dissemination of MAs to acquire a route toward D. However, if the source is located far away from the enclave of D, the state maintained at S regarding D may be stale. In the latter case, S has to transmit a mesh activation request (MR_D^S). When MR_D^S reaches a node x that is part of the enclave, x has to reply with an MA. The state established by the MR_D^S (see Section III-B) is then used to route the MA generated by x toward the source S. This way, the new source acquires routes toward D and is included into the enclave. The state used to route MAs toward the source is short-lived and is deleted within an MA-period. Similarly, nodes located either inside the enclave of D or not to far from it can join to D simply by transmitting an MA stating the new value of their mesh membership flag mm_D^x. On the other hand, a node r that is located far way from the enclave of D, wishing to join to D and whose state regarding D is stale, has to transmit MA_D^r with $\mathrm{core}_D^r \neq r$ and d_D^r equal to the latest known distance to the core of D. This MA is disseminated at most $d_D^r + 1$ hops away from r, establishing a state similar to the one established by the dissemination of an MR. Upon reception of MA_D^r, mesh members reply with MAs that are routed toward r so that r can join to the multicast group D and be included in the enclave.

G. Packet Forwarding and Local Repairs

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

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

GARCIA-LUNA-ACEVES AND MENCHACA-MENDEZ: PRIME
7

Upon reception of a data packet, nodes check for a hit in their data-packet cache, which stores the address of the sender 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 following predicate is used by node x to decide if it has to forward a data packet to destination D received from neighbor y:

$$\mathrm{mm}_D^x = \mathrm{RM} \vee \mathrm{mm}_D^x = \mathrm{MM} \vee \exists y \in L_D : d_D^y \\ > d_D^x \wedge \mathrm{next}_D^y \leq \mathrm{id}^x. \quad (8)$$

In the case of a unicast destination, the first two terms of (8) are always false. Hence, nodes only forward data packets if they are part of a source–destination shortest path. For a multicast destination, (8) states that node x forwards a data packet received from node y if x is part of the multicast mesh or if x was selected by the previous relay (y) as one of its next-hops to the core. This way, data packets travel along directed routing meshes until they reach either the first mesh member or the destination itself and then, in the case of a multicast destination, they are flooded along the multicast mesh.

Forwarders located in directed routing meshes employ the transmission of data packets by their next-hops as implicit acknowledgments (ACKs). If a node fails to receive three consecutive implicit ACKs from a neighbor, then it removes that node from its neighborhood list L_D and updates the value of next_D^x using (6) and the value of d_D^x using (5).

As discussed in Section III-E, a change in any of these two values forces node x to transmit a new MA to advertise its new routing state and, for instance, force a newly selected next-hop node to route data packets toward the destination. If the value of next_D^x equals nil after removing the neighbor from L_D, then the current feasible distance fd_D^x is included in the MA (instead of the new value of the distance d_D^x). An MA with $\mathrm{next}_D^{*x} = \mathrm{nil}$ and $\mathrm{id}^{*x} \neq \mathrm{core}_D^{*x}$ is interpreted as a *neighbor request* that informs other nodes that x no longer has a route toward the destination. Upon reception of a neighbor request from x, nodes eliminate x's entry from L_D and update their routing state accordingly. Additionally, a node z also transmits an MA as a response to a neighbor request if it complies with the predicate

$$\mathrm{next}_D^z \neq \mathrm{nil} \wedge (\mathrm{sn}_D^z > \mathrm{sn}_D^{*x} \vee \mathrm{sn}_D^z = \mathrm{sn}_D^{*x} \wedge d_D^z \leq d_D^{*x}). \quad (9)$$

This is because z is in fact a feasible next-hop for x and can be used to reach the destination through a loop-free path.

H. Core Election

Core elections are held only if the MR contains a multicast address. Upon reception of an MR, a group receiver first determines whether it has received an MA from the core of the multicast group within the last two MA-periods, and no further action in this regard is needed if that is the case. 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. An MA with a larger core id

is considered better than one with a smaller core id. Therefore, if a node receives an MA 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 an MA advertising a core with a smaller id is received, then nodes check if they have recently broadcast an MA with the current core, and if so, the MA is simply ignored. Otherwise, nodes send an MA that forces the neighbor with the smaller core to adopt the core with higher id. Eventually, each connected component has only one core.

I. 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 their perceived channel conditions. Nodes 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 strategy similar to the one proposed in [20], where nodes use a simple and intuitive metric based on the proportion of time in which the channel is perceived as busy. First, nodes compute their *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 s. Then, nodes compute an exponential weighted moving average to avoid reacting too fast to sudden and short-term changes in the instantaneous local contention and get $\rho_n = (1 - \beta)\rho_{n-1} + \beta c$, where ρ_n is the current *level of local contention*, and 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 show that the performance of PRIME is not very sensitive to this parameter.

There are of course practical limitations to accurately measuring the time in which the channel is perceived busy. However, even when using off-the-shelf hardware, it can be estimated from the packet lengths, the data rate, and the total number of transmissions of each node in the one-hop neighborhood [3], [10]. However, to measure the impact of an imperfect estimation of this metric, we performed a sensitivity analysis whose results are shown in Section V-F.

Nodes use the following three strategies to take advantage of the information collected about the level of channel contention.

1) *Adjust the size of the mesh.* Nodes select the number of feasible next-hop nodes (if available) that are forced to join the routing meshes (multicast or directed meshes).
2) *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.
3) *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. The

values of these timers depend on the time it takes to access the channel, which is a function of the current channel contention.

J. Context Awareness in Prime

The basic scheme used in PRIME to select paths within enclaves consists of simply preferring nodes with higher node ids [see (6)]. However, the way in which the routing paths are selected could be enhanced considerably by considering the quality of each individual link, the quality of a neighbor node as a potential forwarder, and hence the quality of the different paths. A detailed treatment of how to use link and node parameters for path selection within enclaves is beyond the scope of this paper; the details of a promising approaches are presented elsewhere [4], [13].

To establish an ordering over the enclaves, we define the *context* of a node x with respect to destination D as a time varying n-dimensional vector $C_D^x(t)$. One or more components of $C_D^x(t)$ can be time-varying functions that depend on the local conditions of node x. These conditions can be either *internal* (e.g., battery power, length of data queues) or *external* (e.g., link quality, local contention, relative mobility). Some components can also depend on the topology (e.g., number of feasible next-hops to D). In this paper, we define $C_D^x(t)$ as a three-dimensional vector $(\rho_n, f_s(\varphi_n), \phi)$, where ρ_n is the current *level of local contention*, φ_n is the average queue length and ϕ is inversely proportional to the cardinality of the set F_D^x of x's feasible neighbors for destination D as defined in Section III-D.

From the context of a node and as it is shown in (10), the cost of using node x to get to D (fc_D^x) is computed as a convex combination of $C_D^x(t)$ plus the minimum cost reported by a feasible next-hop to destination D

$$fc_D^x = \begin{cases} c_1\rho_n + c_2 f_s(\varphi_n) + c_3\phi + \text{cost}_D^i \\ \quad \text{with cost}_D^i : \max_{i \in F_D^x}\{i\}, & \text{if such } i \text{ exists} \\ \infty, & \text{otherwise} \end{cases}$$

(10)

where the $\max$ function is computed using the relation $\prec$ defined in (1) and f_s is a smoothing function defined as $f_s(\varphi_n) = 1 - e^{-0.02\varphi_n}$.

Section V-E presents the results of a series of simulation experiments used to asses the positive impact of using context information in PRIME.

IV. CORRECTNESS IN PRIME

We show that no directed loops can be formed in the nodal routing tables. In our proofs, we assume that the network is connected and that link costs (lc) are positive and larger than 0. Let the current routing state stored at node n_i regarding destination D be $s_D^{n_i} = (\text{id}^{n_i}, \text{core}_D^{n_i}, \text{sn}_D^{n_i}, d_D^{n_i}, \text{fd}_D^{n_i}, \text{mm}_D^{n_i}, \text{next}_D^{n_i})$, and let $\prec_s$ be a total order relation defined in the same way as the relation of (1) but over the routing state of nodes. Again, in the case of a unicast destination, $\text{core}_D^{n_i}$ stores the address of the destination itself.

Theorem 1: Any successor path $p = \{n_0, n_1, \ldots, n_k\}$ with $\text{next}_D^{n_0} = n_1, \text{next}_D^{n_1} = n_2, \ldots, \text{next}_D^{n_{k-1}} = n_k$ established using the data structures and procedures described in Section III is loop-free.

Proof: From (1)–(6), we have that $\text{next}_D^{n_i} = n_j \implies s_D^{n_i} \prec_s s_D^{n_j}$, and hence, for any path $p = \{n_0, n_1, \ldots, n_k\}$, we also have $s_D^{n_0} \prec_s s_D^{n_1} \prec_s \cdots \prec_s s_D^{n_{k-1}} \prec_s s_D^{n_k}$. Now, let us proceed by contradiction and assume that a loop is formed in p when a node n_i selects n_x as its next-hop. Then, we would have $s_D^{n_x} \prec_s s_D^{n_{i-j}} \prec_s \cdots \prec_s s_D^{n_i} \prec_s s_D^{n_x}$, which is a contradiction. ∎

The proof of Theorem 1 makes the implicit assumption that the distance from nodes to the destination cannot be increased for a given sequence number. Theorem 2 relaxes this assumption by allowing nodes to remove neighbors from their L_D and then update d_D^x according to (5).

Theorem 2: Considering the case of Theorem 1, let $p = \{n_0, n_1, \ldots, n_k\}$ be a successor path with $\text{next}_D^{n_0} = n_1, \text{next}_D^{n_1} = n_2, \ldots, \text{next}_D^{n_{k-1}} = n_k$, then any change of successor along that path originated by a local repair as described in Section III-G does not create loops.

Proof: Assume that a cycle is formed when node n_i selects n_x (after removing n_y from L_D or after receiving an update that disqualifies n_y as a feasible successor) as its next-hop. From (6), we can observe that in order to be selected by n_i as a next-hop, n_x has to be an element of the feasible set of next-hops of n_i for destination D, i.e., an element of $F_D^{n_i} = \{n : \text{fd}_D^{n_i} = d_D^n \wedge \text{sn}_D^n = \text{sn}_D^{n_i}\}$. Now, if $n_x \in F_D^{n_i}$, and since $\text{lc}_{n_x}^{n_i} > 0$, we have $d_D^{n_x} = \text{fd}_D^{n_i} < d_D^{n_i}$. Therefore, either $s_D^{n_i} \prec_s s_D^{n_x}$ or $\text{next}_D^{n_i} = n_x = \text{nil}$. In the first case, we reach to the same contradiction as in Theorem 1, i.e., that $s_D^{n_x} \prec_s \cdots \prec_s s_D^{n_i} \prec_s s_D^{n_x}$. In the second case, we would have a cycle with a nil element, which is also a contradiction. ∎

The proof of Theorem 2 shows that, to guarantee instantaneous loop freedom, nodes have to be very conservative when selecting a new next-hop. This could lead to situations in which nodes have a nil value in their next-hop field even when they actually have a valid route toward the destination. However, PRIME alleviates this problem by periodically refreshing the routing structure, which allows nodes to find new routes to the destination. It is important to emphasize that, thanks to the enclave-based way in which control packets are disseminated, the scalability of PRIME is not compromised by the periodical transmission of mesh announcements as it is confirmed by our experimental results presented in Section V.

Corollary 1: Data packets are routed from sources to destinations along loop-free paths.

Proof: From (8), we can observe that a node n_x, which is located outside of the multicast mesh $(\text{mm}_D^{n_x} \neq \text{RM} \wedge \text{mm}_D^{n_x} \neq \text{MM})$, will forward a data packet received from node n_y only if $\text{next}_D^{n_y} = n_x$. Hence, data packets travel hop by hop following successor paths until they reach either the unicast destination or the first mesh member. Now, since successor paths are loop-free (Theorems 1 and 2), we get the desired conclusion. ∎

V. PERFORMANCE COMPARISON

We present simulation results comparing PRIME against ODMRP and PUMA for the case of multicast traffic, against AODV and OLSR for the case of unicast traffic, and against AODV with ODMRP and OLSR with ODMRP for the case of

TABLE I
SIMULATION ENVIRONMENT

Total nodes	100	Node placement	Random
Simulation area	1800×1800^2	Simulation time	150s
MAC Protocol	802.11	Tx. power	15dbm
		Tx. Rate	2000000bps
Data source	MCBR and CBR	Pkts. per src.	1000
Mobility model	Random waypoint	Pause time	10s
		Min.-Max. Vel.	1-10m/s
Mobility model	Group mobility	Node pause time	10s
		Grp. pause time	10s
Grp. Min-Max Vel.	1-10m/s	Node Min-Max Vel.	1-10m/s

combined unicast and multicast traffic. We use ODMRP, AODV, and OLSR in our experiments because they are widely used 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 based on enclaves used in PRIME. We use packet delivery ratio, generalized group delivery ratio, end-to-end delay, control overhead, and total overhead as our performance metrics. The control overhead is the average number of control packets generated by the routing protocols, and the total overhead is the average number of packets that are actually transmitted by the physical layer. The generalized group delivery ratio is an extension of the group reliability metric introduced in [19], in which a packet is considered as delivered if and only if it is received by a given proportion of the multicast group members. This metric emphasizes the importance of group delivery by not considering packets that are received by a small subset of the group members.

The routing protocols are tested with IEEE 802.11 DCF 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 [6] 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 [21] ver. 3.9, which 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 [18]. Unless stated otherwise, each simulation was run for 10 different seed values. To have meaningful comparisons, all the multicast protocols use the same period of 3 s 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 JQ period, as advised by its designers. 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. Unless otherwise stated, Table I lists the details of the simulation environment.

A. Performance With Multicast Traffic

1) Increasing Number of Sources: We first focus on an experiment in which a single multicast group is present and the number of concurrent active senders increases. Each sender
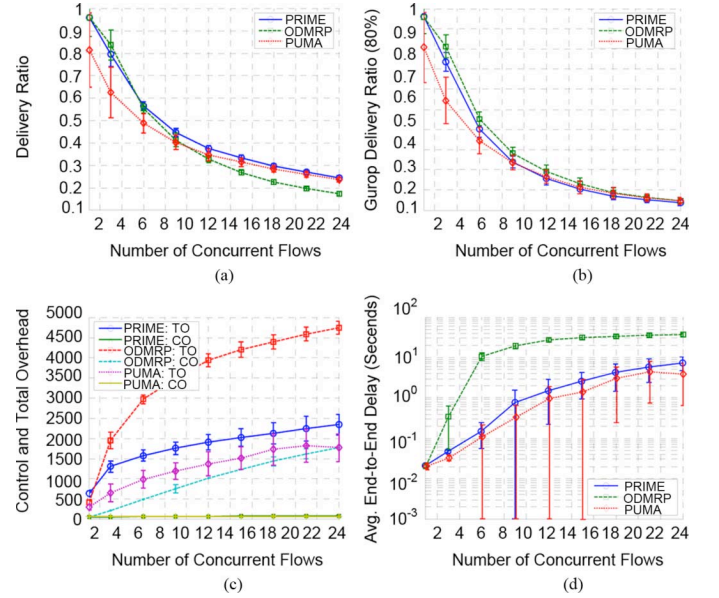

Fig. 3. Random waypoint mobility model: performance with increasing number of MCBR sources. (a) Delivery ratio. (b) Group delivery ratio. (c) Average number of packets transmitted per node and average number of control packets generated per node. (d) End-to-End delay.

transmits 10 packets of 256 B/s, and the group is composed of 20 nodes. In this scenario, sources are not group members.

Fig. 3(a)–(d) presents the results when nodes move according to the random waypoint mobility model. Fig. 3(a) shows the delivery ratio attained by the different protocols. We observe that PRIME performs similar to or better than the other protocols. From Fig. 3(a)–(c), 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 to the single paths used by PUMA. In addition, the meshes established by PRIME are repaired locally, which provides extra reliability. This situation is more evident when the group mobility model is used [Fig. 4(a)–(d)], 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. 3(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 eight or more sources, the generalized group delivery ratios are not. This is due to the fact that maintaining large routing structures like the ones observed when random placement and random waypoint are used becomes a very difficult task in heavily loaded networks. 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.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.
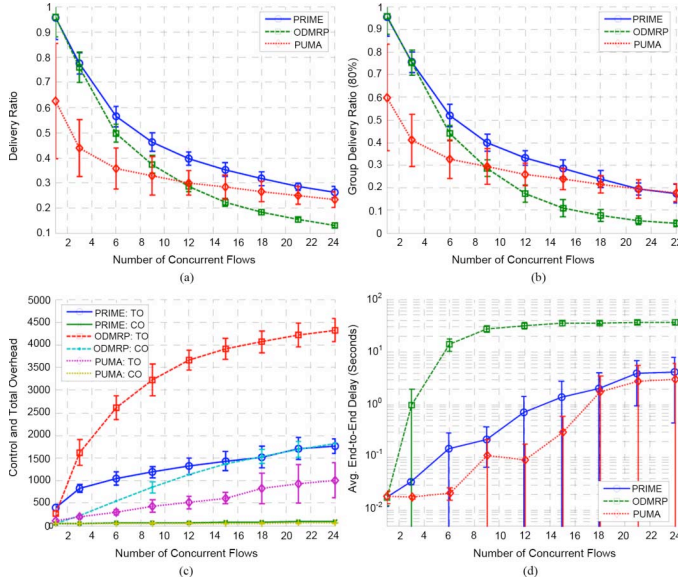
10

IEEE/ACM TRANSACTIONS ON NETWORKING

Fig. 4. Group mobility model: performance with increasing number of MCBR sources. (a) Delivery ratio. (b) Group delivery ratio. (c) Average number of packets transmitted per node and average number of control packets generated per node. (d) End-to-End delay.

Fig. 5. Three sources per group: performance with increasing number of active groups and group areas of 900×900 m^2. (a) Delivery ratio. (b) Group delivery ratio. (c) Average number of packets transmitted per node and average number of control packets generated per node. (d) End-to-end delay.

Fig. 3(c) shows that as the number of sources is increased, the control overhead (CO) induced by PUMA and PRIME remains constant. The latter result is expected because a single node (the core) disseminates control information per multicast group in both protocols, whereas all the sources disseminate control information in ODMRP. The same figure also shows that the average total number of packets transmitted by nodes (TO) running PUMA and PRIME is much smaller than that of ODMRP.

Fig. 3(d) shows that PUMA and PRIME achieve considerably less end-to-end delay than ODMRP. The increased delay shown by ODMRP is due to the large amount of high-priority control packets that are injected into the network, which tend to stall data packets [Fig. 3(c)]. The delay graphs are plotted using logarithmic scale.

Fig. 4(a)–(d) presents 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 900×900 m^2. 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 summing the two mobility vectors [6]. The remaining 80 nodes (including sources) move following the random waypoint mobility model. From Fig. 4(a) and (b), we can observe that PRIME clearly outperforms the other protocols in both delivery ratio and generalized delivery ratio. It is interesting to observe that, unlike 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 routed toward the other sources in ODMRP. Hence,
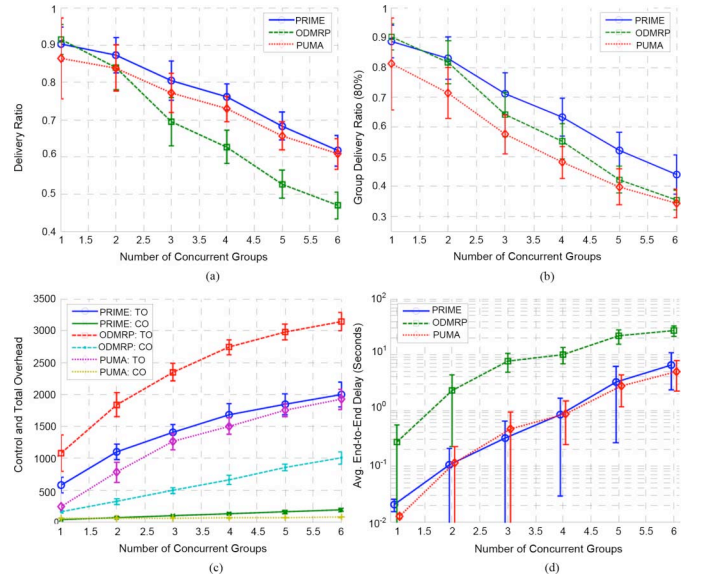
the concentration of group members increases the probability of routing packets toward places where no receiver is located. For the end-to-end delay [Fig. 4(d)] and overhead [Fig. 4(c)], the three protocols show behaviors similar to the ones observed in the previous case.

2) 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. 5(a)–(d) presents the results obtained when group members are located inside of a square area of 900×900 m^2. For each group, three members are randomly selected as sources. We can observe that PRIME consistently attains similar or better delivery and generalized delivery ratios than the remaining protocols. In addition, PRIME attains very low delays, as it is shown in Fig. 5(d). Lastly, this scenario shows how PRIME adjusts its TO to the current conditions of the network by adapting how control and data packets are sent according to the perceived congestion.

B. Performance With Unicast Traffic

1) Increasing Number of Flows: In this scenario we evaluate the performance of PRIME as a unicast routing protocol by including only unicast flows and by increasing the number of CBR sources from 5 to 35. Sources and destinations are selected randomly, and each CBR source generates a total of 1000 data packets of 256 B at a rate of 10 packets per second. In these experiments, all nodes move according to the random waypoint mobility model with parameters as described in Table I. We compare PRIME against AODV and OLSR.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

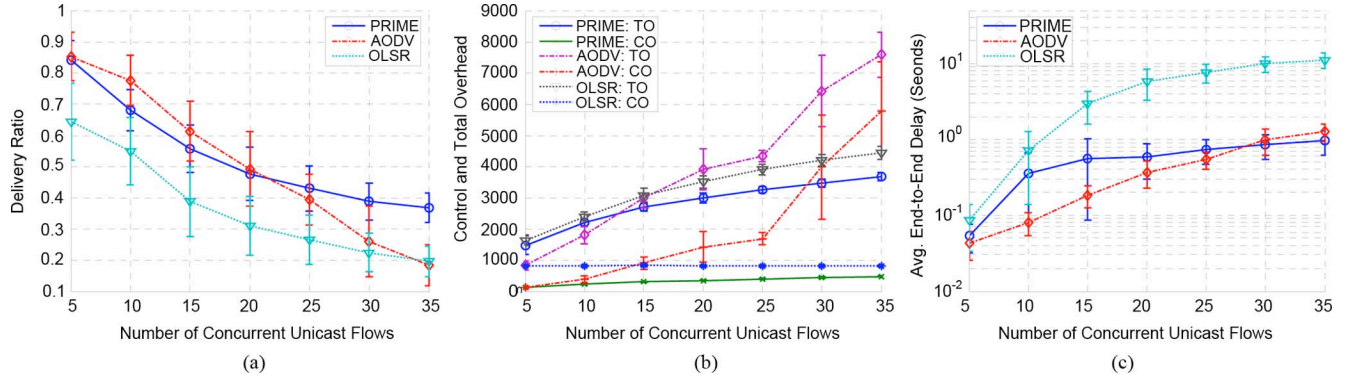GARCIA-LUNA-ACEVES AND MENCHACA-MENDEZ: PRIME

11

Fig. 6. Only unicast traffic: performance with increasing number of CBR flows. (a) Delivery ratio. (b) Average number of packets transmitted per node and average number of control packets generated per node. (c) End-to-end delay.
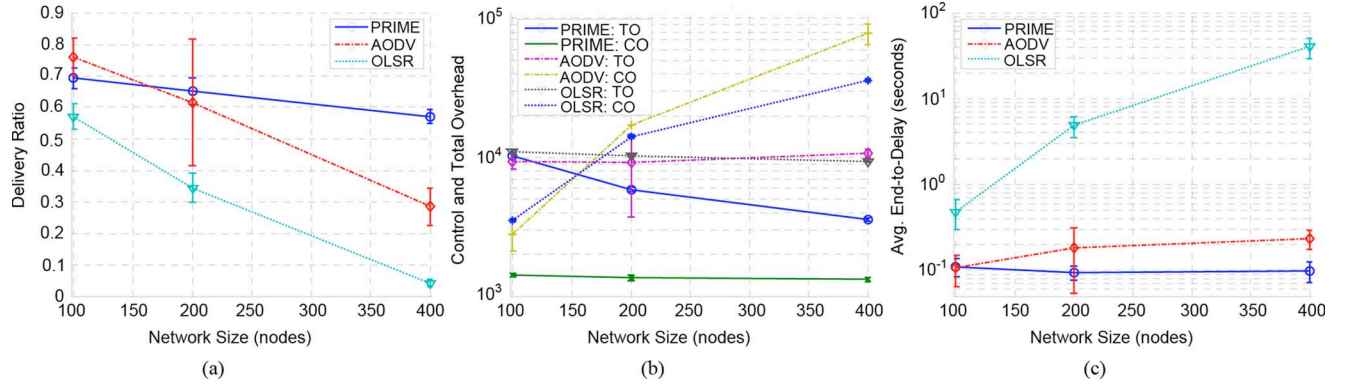

Fig. 7. Only unicast traffic: performance with increasing node density. (a) Delivery ratio. (b) Average number of packets transmitted per node and average number of control packets generated per node. (c) End-to-end delay.

The results are shown in Fig. 6(a)–(c). From Fig. 6(a), we observe that PRIME performs close to AODV for up to 20 flows. However, as the number of flows is further increased, PRIME scales better than AODV and is capable of delivering close to 20% more packets than AODV for 35 flows. We also observe that PRIME consistently outperforms OLSR. Fig. 6(b) shows the control and total overhead induced by the different protocols. We notice that PRIME induces very low control overhead, which contrasts with the one induced by AODV that experiences a steep increase for 25 or more flows. These results allow us to highlight the effectiveness of the enclaves to reduce control overhead. The reduced control overhead is also reflected on the total overhead induced by the protocols, where we observe that the most efficient protocol is PRIME followed by OLSR and AODV in that order. The end-to-end delay attained by the protocols is presented in Fig. 6(c). The figure shows that the best performing protocols are PRIME and AODV, with PRIME attaining delays that are always lower than one second, and with AODV attaining delays similar or better than the other protocols.

2) Increasing Network Size: In this scenario, we evaluate the performance of the protocols as the network size increases. We maintain the same simulation area as described in Table I and increase the number of nodes from 100 to 400. Instead of having flows that last the whole simulation time, in this experiment we use flows that were generated following an exponential distribution. CBR flows are established among randomly selected nodes

with a mean interarrival time of $1/\lambda = 10$ s and a mean flow duration time of $1/\mu = 200$ s. The latter corresponds to 33.3% of the total simulation time, which is equal to 600 s in this scenario. CBR sources generate data packets of 256 B at a rate of five packets per second. Under these parameters, the total number of packets sent in each run is 42 191. The remaining simulation parameters are as defined in Section V-B1.

Fig. 7(a)–(c) shows the results for these experiments. From Fig. 7(a), we can observe that PRIME scales much better than AODV and OLSR. For 400 nodes, PRIME delivers as much as twice the packets delivered by AODV, while OLSR delivers less than 10% of the packets. The reason behind this behavior can be observed in Fig. 7(b), which presents the total and control overhead induced by the protocols. The y-axis of this figure has logarithmic scale. From the figure, we notice that the control overhead induced by PRIME remains constant across the different values in the number of nodes, while that of AODV and OLSR experiences a steep increase. PRIME is able to keep the control overhead constant due to the opportunistic aggregation of control packets and to the restricted flooding enforced by the enclaves. Additionally, even when enclaves confine the dissemination of control packets, they still allow many nodes to acquire routes toward the destinations. On the other hand, the continuous arrival of new flows forces AODV to constantly flood the network with route request packets. As the node density is increased, flooding becomes more costly because it tends to easily make the network congested. Under high congestion,
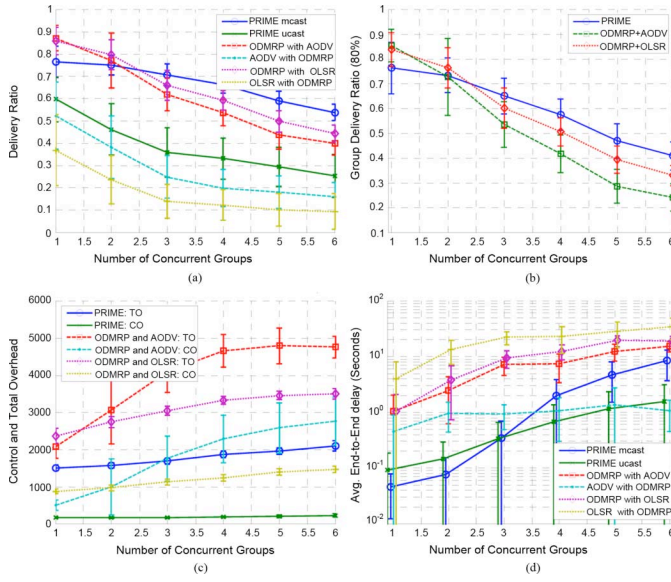
This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

12   IEEE/ACM TRANSACTIONS ON NETWORKING

Fig. 8. Performance with increasing number of active groups, three sources per group, group areas of 900×900 m^2, and five unicast flows. (a) Delivery ratio. (b) Group delivery ratio. (c) Average number of packets transmitted per node and average number of control packets generated per node. (d) End-to-end delay.

TABLE II
PERFORMANCE IN A LARGE COMBINED SCENARIO

	PRIME	ODMRP+OLSR	ODMRP+AODV
Delivery Ratio (m)	$\mathbf{0.90 \pm 0.01}$	0.35 ± 0.03	0.12 ± 0.01
Grp. Delivery Ratio	$\mathbf{0.83 \pm 0.01}$	0.28 ± 0.03	0.05 ± 0.01
Delivery Ratio (u)	$\mathbf{0.50 \pm 0.03}$	0.13 ± 0.01	0.08 ± 0.01
Ctr. Overhead	$\mathbf{1133 \pm 10}$	11067 ± 536	55023 ± 1741
Total Overhead	$\mathbf{4181 \pm 34}$	11909 ± 189	26017 ± 490
E2E Delay (mcast)	$\mathbf{0.051 \pm 0.01}$s	102.17 ± 8.27s	26.05 ± 3.52s
E2E Delay (ucast)	$\mathbf{0.07 \pm 0.01}$s	12.46 ± 4.88s	2.67 ± 1.68s

the same time, up to 15% 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 one group [Fig. 8(b)] and the lowest delays for both unicast and multicast traffic [Fig. 8(d)], while incurring far less CO and TO than the other protocols [Fig. 8(c)], almost five times less than OLSR+ODMRP, and nine times less than AODV+ODMRP.

D. Performance in a Large Network

We evaluate the performance of the protocols in a network of 1000 nodes that are randomly placed in a simulation area of 5000×5000 m^2. The scenario contains 62 concurrent unicast flows, 30 of them established among randomly selected *nearby* nodes, and the remaining 32 flows established among randomly selected nodes and four "popular" destination nodes. The endpoints of the nearby flows are placed into mobility groups of 1000×1000 m^2 so that sources and destinations can move around the entire simulation area while always keeping a maximum distance between them of 1414.2 m. For the flows toward popular destinations, we divided the simulation area into four regions (north–west, north–east, south–west, and south–east) of 2500×2500 m^2 and manually placed a popular node at the center of each region. For each popular destination, eight sources were randomly selected among nodes located in the same region. Popular nodes are static, and their corresponding sources never leave the region where they are placed originally. This scenario also includes 25 multicast groups, 15 of them composed of 15 group members, and the remaining 10 groups composed of 50 nodes. The members of the small-size groups (15 nodes) move inside of mobile regions of 900×900 m^2, and the members of the medium-size groups (50 nodes) move inside of mobile regions of 1500×1500 m^2. The medium-size groups have only one active source, whereas seven small-size groups have five active sources, and the remaining eight have three active sources. Each source (multicast and unicast) transmits a total of 1000 packets of 256 B at a rate of two packets per second. The parameters of the mobility models are as defined in Table I. Each simulation was run for 600 s.

The results shown at Table II are the average of 20 independent runs. From the table, we can notice that PRIME is the only protocol that provides a reasonably good performance in such a large scenario. PRIME delivers 54% more multicast packets than ODMRP and 36% more unicast packets than OLSR. At the same time, PRIME attains end-to-end delays that are three orders of magnitude smaller than those of ODMRP, AODV, and OLSR while inducing an order of magnitude less control overhead. We also observe that the parallel execution of ODMRP

many control packets are lost due to collisions. The latter is interpreted by AODV as broken links that need to be repaired, and hence nodes react to these packet lost by generating even more route request and route reply packets making the network even more congested. As can be observed in Fig. 7(b), the network reaches a saturation point where a large amount of the packets are dropped at the queues. An evidence of that is the fact that the number of control packets generated by the nodes is even larger than the number of packets that are actually transmitted. Something similar happens to OLSR that erroneously interprets control packets lost (due to collisions) as topology changes, triggering new topology control messages that are diffused though the network causing more congestion. The delay attained by the protocols is shown in Fig. 7(c). We observe that PRIME performs similar or better than the other protocols for all the network sizes. As for the previous metrics, the main reason behind the high delays shown by OLSR is the excessive overhead: The large number of high-priority control packets that are pushed into the queues tend to stall data packets.

C. Performance With Combined Multicast and Unicast Traffic

In these experiments we evaluate the performance of the routing protocols in a scenario with combined multicast and unicast traffic. We use the same settings as in the experiment with increasing number of multicast groups and three sources per group, but with the addition of five CBR flows between nodes that do not belong to a multicast group and that move following the random mobility model. As in the unicast scenario, unicast sources send a total of 1000 data packets of 256 B at a rate of 10 packets per second. The results are shown in Fig. 8(a)–(d). From Fig. 8(a), we observe that PRIME attains higher delivery ratios than the other protocols for both unicast and multicast traffic. PRIME delivers up to 18% more data packets than AODV and up to 28% more than OLSR and at
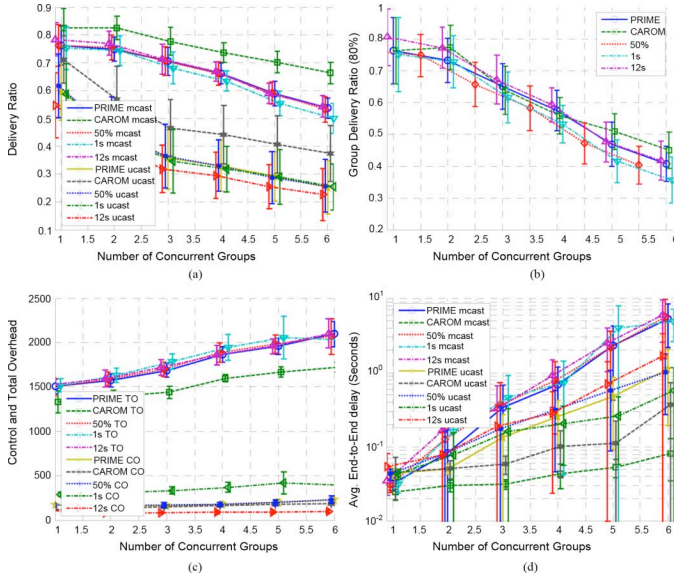
Fig. 9. Sensitivity analysis and context awareness. (a) Delivery ratio. (b) Group delivery ratio. (c) Average number of packets transmitted per node and average number of control packets generated per node. (d) End-to-end delay.

and AODV makes a large network like this unusable. ODMRP delivers only 12% of the multicast packets, while AODV delivers as few as 8% of the unicast packets. The main reason behind this poor performance is the combined control overhead generated by ODMRP and AODV.

As the results of our simulation experiments show, enclaves are a very efficient tool for reducing signaling overhead and hence improving the scalability of routing. At the same time, the establishment of an enclave allows PRIME to identify those nodes that need timely updates about the routes toward a given destination, and hence the quality of these routes is improved by using proactive signaling within enclaves. The use of enclaves significantly reduces the frequency with which nodes located far from the enclaves receive multicast updates. The frequency with which nodes receive a multicast update decreases exponentially with respect to their hop distance to the enclave. This contrasts with previous multicast protocols such as PUMA [18], where all the nodes in the network receive multicast updates with the same frequency per multicast group even if they are located far away from sources and receivers or if the multicast groups do not have interested sources.

E. Impact of Context Awareness

We evaluate the impact of enhancing PRIME with context awareness as described in Section III-J. For this set of experiments, we used a scenario identical to the one of Section V-C. The results are shown in Fig. 9(a)–(d). From Fig. 9(a), we can observe that the enhanced version of PRIME (labeled as CAROM) delivers up to 10% more packets than regular PRIME for both unicast and multicast traffic. The reason is as follows. Thanks to the way in which the cost of the links are computed, CAROM tends to select less congested and more robust paths (those composed by nodes with many feasible next-hops that also have small levels of local congestion and contention), and

hence fewer packets are dropped due to link breakages or collisions. At the same time, because CAROM tends to avoid nodes with congested data queues and control queues, it consistently attains smaller delays than PRIME for unicast traffic and a reduction of up to two orders of magnitude for multicast traffic [Fig. 9(d)]—all these while inducing less total overhead.

F. Sensitivity Analysis

A series of experiments were used to evaluate the sensitivity of PRIME to the length of the MA-periods, the delay used to opportunistically aggregate control messages, and the accuracy of the estimation of the level of local contention (ρ; see Section III-I). For this set of experiments, we used the same settings as in Section V-E. To evaluate the sensitivity of PRIME to the length of the MA-periods, we varied this value from 1 to 12 s. The results are shown in Fig. 9(a)–(d), where the graphs labeled as "1 s" and "12 s" correspond to PRIME using MA-periods of 1 and 12 s, respectively. The intermediate values are not shown for the sake of legibility. As can be seen in the figures, using longer MA-periods in PRIME is not very attractive, given that it renders similar or lower delivery ratios and end-to-end delays [Fig. 9(c)].

For the case of the delay used to aggregate control messages, we varied this value from 10 to 100 ms and found that PRIME is not sensible to this parameter. These graphs are omitted due to space limitations. Lastly, we evaluate the robustness of PRIME to inaccurate estimations of ρ by adding a random perturbation (r) to the time in which nodes perceive the channel as busy (t_b). r is sampled from the continuous uniform distribution with parameters $(0, x \cdot t_b)$, where $x \in (0, 1)$ is the intensity of the perturbation. For these experiments, we increased the value of x from 0.1 to 0.5 and computed the instantaneous local contention c as $(t_b \pm r)/S_p$. The results are shown in Fig. 9(a)–(d), where the graphs labeled as "50%" correspond to a value of $x = 0.5$. From the figures, we can also conclude that PRIME is not sensible to this parameter.

VI. CONCLUSION

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 integrated routing is a much more attractive approach for MANETs than attempting to provide either on-demand or proactive unicast and multicast routing using separate protocols. 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 nodes that actually need the information. This property has a positive impact on the scalability of the protocol, particularly in medium to large networks in which the members of the same multicast group tend to be near one another. 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. Moreover, PRIME is the only protocol that provides adequate performance in a large network of 1000 nodes with combined multicast and unicast traffic.

REFERENCES

[1] V. Devarapalli and D. Sidhu, "Mzr: A multicast protocol for mobile ad hoc networks," in *Proc. IEEE ICC*, 2001, vol. 3, pp. 886–891.

[2] J. J. Garcia-Luna-Aceves and E. L. Madruga, "The core-assisted mesh protocol," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 8, pp. 1380–1394, Aug. 1999.

[3] M. Genetzakis and V. A. Siris, "A contention-aware routing metric for multi-rate multi-radio mesh networks," in *Proc. IEEE SECON*, 2008, pp. 242–250.

[4] R. Ghosh, R. Menchaca-Mendez, and J. J. Garcia-Luna-Aceves, "An interest-driven approach for unicast routing in MANETs with labeled paths and proactive path maintenance," in *Proc. IEEE WCNC*, 2011, pp. 1–6.

[5] Z. J. Haas, M. R. Pearlman, and P. Samar, "The Zone Routing Protocol (ZRP) for ad hoc networks," MANET Working Group, IETF, Internet draft, 1998.

[6] X. Hong, M. Gerla, G. Pei, and C.-C. Chiang, "A group mobility model for ad hoc wireless networks," in *Proc. ACM/IEEE MSWiM*, 1999, pp. 53–60.

[7] P. Jacquet, A. Laouiti, P. Minet, and L. Viennot, "Performance analysis of OLSR multipoint relay flooding in two ad hoc wireless network models," INRIA, Rocquencourt, France, Tech. rep. 4260, Sep. 2001.

[8] J. G. Jetcheva and D. B. Johnson, "Adaptive demand-driven multicast routing in multi-hop wireless ad hoc networks," in *Proc. ACM MobiHoc*, 2001, pp. 33–44.

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

[10] K. H. Kim and K. G. Shin, "On accurate measurement of link quality in multi-hop wireless mesh networks," in *Proc. ACM MobiCom*, 2006, pp. 38–49.

[11] S.-J. Lee, M. Gerla, and C.-C. Chiang, "On-demand multicast routing protocol," in *Proc. IEEE WCNC*, 1999, vol. 3, pp. 1298–1302.

[12] R. Menchaca-Mendez and J. J. Garcia-Luna-Aceves, "An interest-driven approach to integrated unicast and multicast routing in MANETs," in *Proc. ICNP*, 2008, pp. 248–257.

[13] R. Menchaca-Mendez and J. J. Garcia-Luna-Aceves, "Robust and scalable integrated routing in MANETs using context-aware ordered meshes," in *Proc. IEEE INFOCOM*, 2010, pp. 1–9.

[14] S. Murthy and J. J. Garcia-Luna-Aceves, "An efficient routing protocol for wireless networks," *Mobile Netw. Appl.*, vol. 1, no. 2, pp. 183–197, 1996.

[15] C. E. Perkins and E. M. Royer, "Ad-hoc on-demand distance vector routing," in *Proc. 2nd IEEE WMCSA*, Feb. 1999, pp. 90–100.

[16] S. Roy and J. J. Garcia-Luna-Aceves, "Node-centric hybrid routing for ad hoc networks," in *Proc. IEEE/ACM MASCOTS*, Oct. 12, 2002, pp. 63–71.

[17] E. M. Royer and C. E. Perkins, "Multicast operation of the ad-hoc on-demand distance vector routing protocol," in *Proc. ACM MobiCom*, 1999, pp. 207–218.

[18] R. Vaishampayan and J. J. Garcia-Luna-Aceves, "Efficient and robust multicast routing in mobile ad hoc networks," in *Proc. IEEE MASS*, Oct. 2004, pp. 304–313.

[19] K. Viswanath, K. Obraczka, and G. Tsudik, "Exploring mesh and tree-based multicast routing protocols for MANETs," *IEEE Trans. Mobile Comput.*, vol. 5, no. 1, pp. 28–42, Jan. 2006.

[20] Y. Yang and R. Kravets, "Contention-aware admission control for ad hoc networks," *IEEE Trans. Mobile Comput.*, vol. 4, no. 4, pp. 363–377, Jul.–Aug. 2005.

[21] Qualnet ver. 3.9, Scalable Network Technologies, Los Angeles, CA [Online]. Available: http://www.scalablenetworks.com

J. J. Garcia-Luna-Aceves (S'75–M'77–SM'02–F'06) received the B.S. degree from the Universidad Iberoamericana, Mexico City, Mexico, in 1977, and the M.S. and Ph.D. degrees from the University of Hawaii at Manoa, Honolulu, in 1980 and 1983, respectively, all in electrical engineering.

He holds the Jack Baskin Endowed Chair of Computer Engineering at the University of California, Santa Cruz (UCSC), is Chair of the Computer Engineering Department, and is a Principal Scientist with the Palo Alto Research Center (PARC), Palo Alto, CA. Prior to joining UCSC in 1993, he was a Center Director with SRI International (SRI), Menlo Park, CA. He has been a Visiting Professor with Sun Laboratories, Menlo Park, CA, and a Principal of Protocol Design with Nokia, Mountain View, CA.

Dr. Garcia-Luna-Aceves is a Fellow of the Association for Computing Machinery (ACM) and the American Association for the Advancement of Science (AAAS) and is listed in Marquis *Who's Who in America* and *Who's Who in The World*. He is the co-recipient of the IEEE Fred W. Ellersick 2008 MILCOM Award for best unclassified paper. He is also co-recipient of Best Paper Awards at the European Wireless Conference 2010, IEEE MASS 2008, SPECTS 2007, IFIP Networking 2007, and IEEE MASS 2005 conferences and of the Best Student Paper Award of the 1998 IEEE International Conference on Systems, Man, and Cybernetics. He received the SRI International Exceptional-Achievement Award in 1985 and 1989.

Rolando Menchaca-Mendez received the B.S. degree in electronic engineering from the Universidad Autonoma Metropolitana, Mexico City, Mexico, in 1997, the M.S. degree in computer science from the Mexican National Polytechnic Institute, Mexico City, Mexico, in 1999, and the Ph.D. degree in computer engineering from the University of California, Santa Cruz, in 2009.

He is a Professor and Head of the Communications and Networking Laboratory with the Computer Research Center, Mexican National Polytechnic Institute.