

Performance evaluation of the MoM mobile multicast protocol

Carey L. Williamson, Tim G. Harrison*, Wayne L. Mackrell and Richard B. Bunt

Department of Computer Science, University of Saskatchewan, 57 Campus Drive, Saskatoon, SK, Canada S7N 5A9

This paper presents a performance study of a mobile multicast protocol called MoM, which is designed to support IP multicast for mobile hosts in an IP internetwork. The protocol uses the basic unicast routing capability of IETF Mobile IP, and leverages existing IP multicast to provide multicast services for mobile hosts as well. A key feature of the MoM protocol is the use of designated multicast service providers (DMSPs) to improve the scalability of mobile multicast. Discrete-event simulation is used in the performance evaluation of the protocol. The performance study focuses on the scalability, routing efficiency, fairness, and overhead of the MoM protocol, as well as on DMSP selection policies and the deliverability of multicast messages. The simulation results suggest distinct performance advantages for the MoM protocol over other approaches for mobile multicast, such as bi-directional tunnelling, particularly as the number of mobile group members increases. Furthermore, even simple policies for choosing a DMSP from possible candidates provide reasonable tradeoffs between handoff rates, routing efficiency, deliverability of messages, and protocol overhead.

1. Introduction

Providing multicast support for mobile hosts in an IP internetwork is a challenging problem, for several reasons. First, the IETF Mobile IP [17] protocol (as currently defined for IPv4) concentrates on unicast delivery to mobile hosts; additional mechanisms must be added to support multicast delivery within or on top of IETF Mobile IP. Second, the addition of mobility to the host group model [9] implies that multicast routing algorithms must now deal not only with dynamic group membership, but also with dynamic member location (i.e., the routes used to reach specific group members are themselves transient in nature). Third, many of the algorithms used in multicast routing protocols, such as DVMRP [18], MOSPF [16] and PIM [10], implicitly assume static hosts when setting up a multicast delivery tree. Reconstructing the delivery tree every time a multicast source moves is not viable, because of the overhead involved, yet leaving the tree unchanged can cause inefficient or incorrect multicast message delivery.

Several approaches to mobile multicast have been proposed in the literature [1–4,6,8]. A good discussion of the issues involved in mobile multicast support is presented in [19]. In this paper, we restrict our attention to mobile multicast approaches based on IETF Mobile IP [17], which proposes two approaches to support mobile multicast that we call *remote subscription* and *bi-directional tunnelling*.

In remote subscription, the onus is on the mobile host to resubscribe to its desired multicast groups while at a foreign network. The remote subscription approach is simple, and works well if the mobile host spends a relatively long time at each foreign network, compared to the join and graft latencies [3]. It has the further advantage of offering good (i.e., shortest path) routes for delivery of multicast datagrams to mobile hosts. However, the approach implicitly assumes that mobile hosts are only *recipients* of multicast

messages, or that they have a co-located address on the foreign network. If the mobile host sends a multicast datagram with its home address as the source, the incoming interface check [10] of most multicast routing algorithms may discard datagrams intended for the multicast group. For remote subscription to work, rebuilding the multicast tree may be necessary as a result of source movement, or some form of tunnelling used to handle mobile multicast sources. Finally, the approach assumes the existence of a multicast router at the visited network, an assumption that should, but may not always hold. Without such a multicast router, multicast message delivery can be achieved only by using some form of tunnelling.

With bi-directional tunnelling, mobile hosts send and receive all multicast datagrams by way of their home network, using unicast Mobile IP tunnels from their Home Agents. This approach handles source mobility as well as recipient mobility, and in fact hides host mobility from all other members of the group. The drawbacks, however, are two-fold. First, the routing path for multicast delivery can be far from optimal (in the worst case, the source and the recipient can be on the same network, while all multicast messages between the two hosts must traverse the entire internetwork twice). Second, the approach offers limited scalability. Home agents with multiple mobile host group members away from home must replicate and deliver tunnelled multicast datagrams to each of them, regardless of at which foreign networks they reside.

In this paper, we evaluate a new approach to mobile multicast first introduced in [6]. We call our protocol MoM, for Mobile Multicast. The basic idea in MoM is to use the home agent functionality of IETF Mobile IP for delivery of multicast datagrams to mobile hosts, achieving scalability through the use of a designated multicast service provider (DMSP) optimization per multicast group for each foreign network. The high-level design of the MoM protocol was presented in an earlier paper [8]; the present

* Now at Teradyne, Inc., Agoura Hills, California, USA.

paper focuses primarily on the performance characteristics of the MoM protocol, extending the preliminary results reported in [12].

We evaluate the MoM protocol through simulation, focusing on its scalability with respect to multicast group size, number of multicast groups, and number of mobile hosts. In particular, we demonstrate the performance characteristics of remote subscription, bi-directional tunnelling, and the MoM approach across a range of network and workload parameters, and argue the suitability of our approach for the ever-growing world of mobile networking. We then proceed to address several practical matters, such as DMSP selection policies, handoff rates, fairness, routing efficiency, deliverability of multicast messages, and protocol overhead.

The remainder of this paper is organized as follows. Section 2 reviews the design and operation of the MoM mobile multicast protocol. Section 3 presents our simulation setup for a performance study of mobile multicast, and section 4 presents the results from the study. Finally, section 5 presents our conclusions.

2. A mobile multicast protocol: MoM

Our protocol addresses the problem of multicast to mobile hosts in an IP internetwork. Adapting a static host multicast mechanism such as that provided by IGMP and DVMRP to support mobile hosts as well would seem non-trivial. However, a deployed, proven multicast mechanism that already exists for static hosts on the Internet is of great utility if it can be leveraged to provide multicast services for mobile hosts as well.

Our protocol relies heavily on the foundation provided by IETF Mobile IP [17]. That is, we rely on the Home Agent (HA) and Foreign Agent (FA) to effect delivery of messages from a Correspondent Host (CH) to a Mobile Host (MH) away from its home network. The IETF Mobile IP terminology is used throughout the discussion of our protocol. Further background on Mobile IP and its related protocols is provided in [7,11,13–15].

2.1. Assumptions and design goals

Our protocol assumes that the service to be provided is the *unreliable, best effort, connectionless* delivery of multicast datagrams, and that multicast support must conform to the host group model [9]. We assume that a mobile host that wishes to receive multicast datagrams is capable of receiving them on its home network using existing (static) multicast techniques. In the proposed architecture, it is assumed that a multicast router is co-resident with the HA. Furthermore, we assume that foreign agents¹ are being used at the foreign networks. Home agents and foreign agents are assumed to be static (not mobile) hosts.

¹ Extension of our protocol to handle co-located addresses remains for future work.

We make no *a priori* assumptions about the size of multicast groups, the geographic distribution of group members, the number of mobile hosts in the network, the location of the mobile hosts, or the frequency of mobile host movement. In a practical scenario, however, it is unlikely that a given mobile host will move more frequently than on a minute-to-minute basis.

Our design goals include:

- *Scalability.* The approach should work well even when the number of mobile hosts in the internetwork is large (which it soon will be). Clearly, the approach should work for both small and large multicast groups.
- *Robustness.* The disruption of multicast service due to movement of a host from one network to another must be minimal.
- *Simplicity.* We would like the scheme to be as simple as possible, in the sense that it be able to interoperate with existing Internet protocols and mechanisms, with as few changes as possible.

2.2. Protocol overview

The MoM protocol relies on HAs to forward multicast traffic to MHs through the Mobile IP tunnel via the FA. In this approach, the FA need not join groups on behalf of mobile hosts that are visiting its network, and mobile hosts that are members of a multicast group G are not subject to join and graft delays every time they move [3]. While these join and graft latencies may not be much more than the handoff latencies needed by our protocol, our protocol confines the setup overhead to the mobile support entities, which expect frequently moving mobile hosts, and reduces the impact on multicast routers, which do not.

Since the HA may be serving MHs at several FAs that wish to receive datagrams addressed to the multicast address for group G , it forwards a copy into each corresponding Mobile IP tunnel. Note that the HA need not forward a separate copy for each mobile host that it serves, but only one copy for each foreign network at which its mobile host group members reside. Link-level multicast is used by the FA at each such foreign network to complete the delivery. Of course, care must be taken not to create routing loops when tunneling multicast datagrams “upstream” to an agent that is also a multicast router on the multicast delivery tree. One solution is for agents that are also multicast routers to not reforward packets emerging from a Mobile IP tunnel.

Our approach avoids the unnecessary duplication of multicast packets on the foreign network in the event that the HA has multiple MHs present there. If bi-directional tunnelling is used, all multicast packets are forwarded individually to each MH by its HA. While duplicate datagrams do not constitute a violation of the IP multicast service assumptions, they would constitute an additional load on possibly low-bandwidth links.

This forwarding scheme is complicated by a phenomenon we call the *tunnel convergence problem* [6,8], wherein multiple Mobile IP tunnels (from different HAs) can terminate at a particular FA. Thus one copy of every multicast packet would be forwarded to the FA by each HA that is serving interested MHs. Since the FA would locally deliver every multicast datagram forwarded to it, the problem of duplicate multicast packet delivery to the MHs must again be handled.

To solve the tunnel convergence problem, the FA selects one HA as the *designated multicast service provider, DMSP*, for a given multicast group. HAs that are not the DMSP for a given multicast group can suppress delivery down the Mobile IP tunnel using negative caching, as described for PIM [10].

One drawback of the DMSP approach is that when a mobile host moves from a foreign network to another network (home or foreign), there is a possibility for a temporary disruption of multicast delivery for other mobile hosts on the (previous) foreign network. This temporary multicast service outage stems from the fact that in Mobile IP there is no explicit deregistration with the Foreign Agent when a host moves. The MH's Home Agent learns of the movement when the MH reregisters at the new network, but the FA at the old foreign network learns about the movement only through a timeout. In the case that the moving host's HA was the DMSP for a group at the (previous) foreign network, a DMSP handoff may be required to a different HA, to forward datagrams for the remaining multicast group members (if any) at the foreign network. Until this handoff completes, multicast delivery for group members at the foreign network may be disrupted.² An experimental study of the severity of this problem is presented later in the paper (see section 4.5).

2.3. Protocol data structures

Figure 1 illustrates the data structures needed for the description of the MoM protocol. Further details are given in [7,8]. Each Home Agent must maintain an away list to keep track of which of its own mobile hosts are away, where they are (i.e., which FA), and when their bindings expire. Similarly, each Foreign Agent maintains a visitor list to keep track of which mobile hosts are currently at its LAN, where the hosts came from (i.e., which HA), and when these bindings expire.

The MoM protocol also requires group membership information for the away and visiting mobile hosts, which is also shown in figure 1. This group information could reside at the multicast router for the network, or at the Home and Foreign Agents. In our protocol, we assume the latter. That is, each Home Agent keeps track of three things for each multicast group that it knows about: a list of away mobile hosts that are members of the group, a

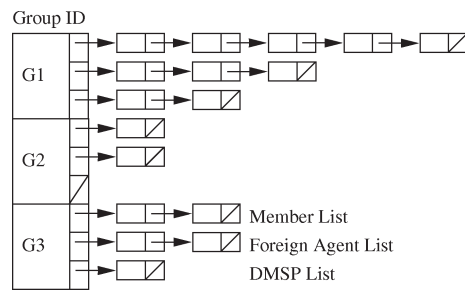
² A protocol extension that allows multiple DMSPs was proposed in [5] as a solution to this problem. Another solution would be to require a DMSP to continue forwarding until its last MH's lifetime has expired.

Home Agent Tables

Away Table

Host	FA	Timestamp

Group Information




(a)

Foreign Agent Tables

Visitor Table

Host	HA	Timestamp

Group Information



(b)

Figure 1. Data structures for the MoM protocol: (a) Home Agent; (b) Foreign Agent.

list of the FAs at which the away group members reside, and a list of the FAs for which the HA has DMSP responsibilities. Similarly, each Foreign Agent keeps track of three things on a per group basis: a list of visiting mobile hosts that are members of the multicast group, a list of the HAs to which these visiting group members belong, and a list of HAs that are currently serving as DMSPs for this group.

In simple terms, the main issue addressed by the MoM protocol is the lengths of the three lists that agents maintain for each multicast group. Consider a single HA, for example. Clearly, the number of multicast group members who are away may grow large, but the list of foreign networks at which they currently reside is likely to be shorter, and the list of these foreign networks for which the Home Agent has DMSP responsibilities shorter yet. The DMSP optimization can thus be effective in reducing the multicast message forwarding load for HAs, reducing the multicast message traffic on foreign networks, and improving the scalability of mobile multicast.

2.4. Protocol details

The events that must be handled in the MoM protocol are host arrivals, host departures, timeouts, DMSP handoffs, and multicast message arrivals. The main steps involved in the protocol are shown algorithmically in figures 2–5, and described briefly here.

When a mobile host leaves its home network to attach to a foreign network, there is no notification of this event. Thus there are no protocol steps required.

When a mobile host arrives at a foreign network, the steps listed in figure 2 are executed. These steps update the visitor list at the FA and the away list at the HA, and update group membership information at each agent so that

When mobile host MH arrives at foreign network:

1. Register with Foreign Agent (FA)
 - 1.1 Create Visitor Table entry for MH.
 - 1.2 Insert host name, HA info, and set timer.
 - 1.3 Notify FA of MH's current group memberships.

For each multicast group G that MH is in:

 - 1.3.1 Make entry in GroupInfoTable, if needed.
 - 1.3.2 Add MH to group membership list for G.
 - 1.3.3 If this is the first MH from that HA at this FA, then add the MH's HA to the HA list for group G, else increment the host count for the MH's HA.
 - 1.3.4 Select a DMSP for this group from HA list.
 - 1.3.5 If the chosen DMSP differs from the old DMSP then perform DMSP handoff.
2. Register with Home Agent (HA)
 - 2.1 Create or update Away Table entry for MH.
 - 2.2 Record oldFA, if any.
 - 2.3 Insert host name, FA info, and set timer.
 - 2.4 Notify HA of MH's group memberships.

For each multicast group G that MH is in:

 - 2.4.1 Make entry in GroupInfoTable, if needed.
 - 2.4.2 Add MH to G's membership list, if needed.
 - 2.4.3 If this is the first MH from this HA at that FA, then add the MH's FA to the FA list for group G, else increment the host count for the MH's FA.
 - 2.4.4 If the MH's new FA differs from oldFA then decrement host count for oldFA, discarding oldFA from list if count is zero.
 - 2.4.5 Record/update DMSP status (YES/NO) of HA for group G at FA (and oldFA, if needed).

Figure 2. Protocol steps when a Mobile Host arrives at a Foreign Network.

When a MH returns to its home network:

1. Notify the Home Agent (HA)
 - 1.1 Delete Away Table entry for MH, noting oldFA.
 - 1.2 For each multicast group G that MH is in:
 - 1.2.1 Delete MH from the membership list for G.
 - 1.2.2 Decrement the host count for MH's oldFA, discarding oldFA from FA list if count is zero, and deleting oldFA from DMSP list, if needed.

Figure 3. Protocol steps when a Mobile Host returns to its Home Network.

When a MH times out at a foreign network:

1. Delete MH's entry from visitor list, noting HA.
2. For each multicast group G that MH is in:
 - 2.1 Delete MH from the membership list for G.
 - 2.2 Decrement the host count for MH's HA, discarding the HA from HA list if count is zero, and deleting the HA from the DMSP list, if needed.
 - 2.3 Select a DMSP from HA list for this group.
 - 2.4 If chosen DMSP differs from the old DMSP then perform DMSP handoff.

Figure 4. Protocol steps when the Registration Timer for a Mobile Host at a Foreign Network expires.

When a unicast packet for MH arrives at MH's HA:

1. Look up FA information for MH in Away Table.
2. Encapsulate packet and tunnel it to the FA.

When a multicast packet for group G arrives at HA:

1. Forward multicast packet to local members.
2. Look up membership information for the away members of that group.
3. Encapsulate packet and forward to each FA for which the HA is the DMSP for group G. This could be done using a separate Mobile IP unicast tunnel to each such FA, or as a multicast tunnel to the set of FAs for which the HA is the DMSP for group G.

When a tunnelled packet arrives at FA from HA:

1. Decapsulate the packet.
2. If the packet is a unicast packet for a mobile host then forward to that host.
3. If the packet is a multicast packet for group G, then check for local members, and forward using link-level multicast if local members are found.

Figure 5. Protocol steps for message delivery to Mobile Hosts.

the DMSP status determined by the FA is known by the HA. The DMSP decision is made independently by each FA for each multicast group G .

When a mobile host leaves a foreign network, no notification is required. State updates are handled by timer expiration.

When a mobile host arrives at its home network, the steps in figure 3 are executed. These steps update the data structures to reflect the host's new "at home" status. Similar updates of the data structures at the (former) foreign network are handled by timer expiration.

Upon expiration of a timer at a foreign network, the steps in figure 4 are executed. These steps update the FA's information regarding the mobile host, and may result in a DMSP handoff for some multicast groups.

Delivery of multicast messages to mobile hosts is handled by forwarding from the designated DMSP HA to the FAs at which mobile group members reside. These steps are shown in figure 5.

3. Mobile multicast simulation model

We have evaluated our approach to mobile multicast using a discrete-event simulation tool constructed for this purpose. This simulator has served two distinct purposes. First, it has aided the design and debugging of the MoM protocol in its "proof of concept" stage. Second, it has helped to evaluate the performance of the MoM protocol, relative to the remote subscription and bi-directional tunnelling approaches, as network and workload parameters are varied.

3.1. Network and workload model

The network model for our simulation study is quite straightforward. We assume that there are N local area networks, each with H hosts. For simplicity, we assume that the LANs are pinpoint locations on an x - y coordinate system, with the x and y coordinates chosen uniformly at random for each LAN. This set of LAN locations is fixed for the duration of each simulation. We do not explicitly model the network topology between the LANs, but we do make use of the shortest-path Euclidean distances between the LANs. Each LAN has an associated Home Agent and Foreign Agent. A fraction p of the H hosts on each LAN are considered mobile hosts. The experiments in this paper all use $p = 1.0$. Although in a real LAN fewer than 100% of the hosts (excluding agents) would be mobile, we assume 100% for our simulations since no extra insight is gained by simulating stationary hosts.

Figure 6 illustrates the assumed model of host mobility. Very simply, hosts can be in one of two states (ignoring travel times for the moment): at the home network or at a foreign network (see figure 6(a)). Mobile hosts begin the simulation at their home network, and are allowed to roam about in the network at random. Foreign networks to visit

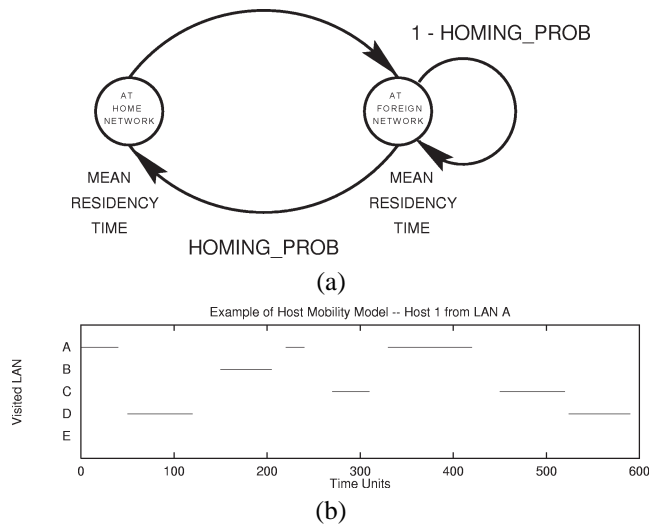


Figure 6. Host mobility model used for simulations: (a) State transition diagram; (b) Time plot illustrating Mobile Host behaviour.

are chosen equiprobably at random, while the homing probability after each visit to a foreign network is $h = 0.5$. The number of foreign networks visited on each trip away from home is thus geometrically distributed with mean $1/h = 2$. The residency time for each visit to a network (home or foreign) is drawn from an exponential distribution with a mean of 60 time units³, and the travel time for going between networks (regardless of distance) is exponentially distributed with a mean of 6 time units.⁴ Hosts thus spend 9.1% of their time in transit, and 90.9% of their time connected to a LAN (60.6% at foreign networks, and 30.3% at home). Each mobile host operates independently. The mobility behaviour of a typical mobile host is illustrated graphically in figure 6(b). The average "cycle time" for a mobile host in our simulation model is 198 time units (i.e., time spent at the home network, plus the time spent visiting two foreign networks, plus the travel times between networks).

Multicast group communication is also simulated, for M multicast groups. For each multicast group, group members are chosen equiprobably at random, with the size of the multicast group specified as a workload parameter. For each multicast group, there is a single point source for all multicast datagrams, located in the internetwork, and physically disjoint from all simulated LANs. Multicast messages are generated in a Poisson fashion, using a specified message arrival rate λ . We do not explicitly model the construction of the multicast delivery tree from the source to the recipients; we merely keep track of at which LAN the mobile hosts reside as the simulation progresses, and use this information as needed to simulate the delivery of each generated multicast message to the group recipients. The experiments in this paper all assume static membership of multicast groups, although the group members themselves may move about the internetwork. Extension of the

³ Since time is not an issue in our simulations, we use generic time units.

⁴ These parameter values have not been empirically validated, but rather were chosen to "exercise" the protocol.

Table 1
Network and workload parameters.

Parameter	Description	Value(s)
N	Number of LANs	5–20
H	Hosts per LAN	10–20
p	Fraction of hosts that are mobile	1.0
RES	Mean residency time per visit at each LAN visited (in time units)	60
TR	Mean travel time between LANs	6
TO	Registration timeout value	10
h	Homing probability	0.50
M	Number of multicast groups	1–8
g	Multicast group size	1–50
s	Sources per multicast group	1
λ	Multicast message generation rate (msgs/time unit)	0.5

simulator to support dynamic group membership is future work.

Table 1 summarizes the main network and workload parameters used in our simulation experiments. Many of these values were chosen so that the protocol could be adequately evaluated with simulation times that were manageable for the number of runs needed. In choosing the parameters the intent was to generate enough activity to adequately test the functionality and features of the protocol. The realistic modelling of specific mobile behaviour was not attempted, since a simple evaluation of protocol operation was desired.

3.2. Simulation methodology

The simulation experiments were conducted using a multi-factor experimental design. We used the simulator to assess the performance impacts of DMSP selection policy, multicast group size, network size, and number of mobile hosts.

The warmup period used for the simulations was 6000 time units (approximately 30 cycle times for each mobile host), following which we collect simulation statistics relating to mobile multicast until the end of the simulation run (20 000 additional time units, or approximately 100 cycle times for each mobile host).

3.3. Experimental factors and levels

The two main factors considered were multicast group size and DMSP selection policy. The multicast group size was varied from 1 to 50 (only one multicast group at a time was used in the simulation, since each group is handled independently). Several different DMSP selection policies were studied: age-based, proximity-based, count-based, and random. These selection policies are described as follows:

- **Random:** The DMSP is selected at random from the entries in the HA list.
- **Age-based algorithms:**

- * **Oldest-HA:** The HA entry that has been in the HA list the longest time is chosen as the DMSP.
- * **Newest-HA:** The HA entry that has been in the HA list the shortest time is chosen as the DMSP.
- * **Oldest-MH:** The HA of the MH that has been visiting the FA the longest time is chosen as the DMSP.
- * **Newest-MH:** The HA of the MH that has been visiting the FA the shortest time is chosen as the DMSP.

- **Count-based:** The HA that presently has the most visitors at this FA is chosen as the DMSP. Ties are broken by choosing the oldest HA with the largest count.

- **Proximity-based algorithms:**

- * **Closest-to-FA:** The HA that is closest to the FA is chosen as the DMSP.
- * **Closest-to-Source:** The HA that is closest to the multicast source is chosen as the DMSP.

Among these policies, the Oldest-HA policy is the most intuitively obvious, but the range of other policies was selected to evaluate performance tradeoffs in the MoM protocol (see section 4.4).

4. Simulation results

This section presents the results from our mobile multicast simulation study.

4.1. Scalability with group size

The first simulation experiment compares the performance of the MoM protocol to bi-directional tunnelling. These results are illustrated in figure 7 for a five-LAN network scenario using the Oldest-HA DMSP policy.

Figure 7 illustrates how various aspects of the mobile routing environment scale as the multicast group size is increased. The values plotted are averages calculated on a

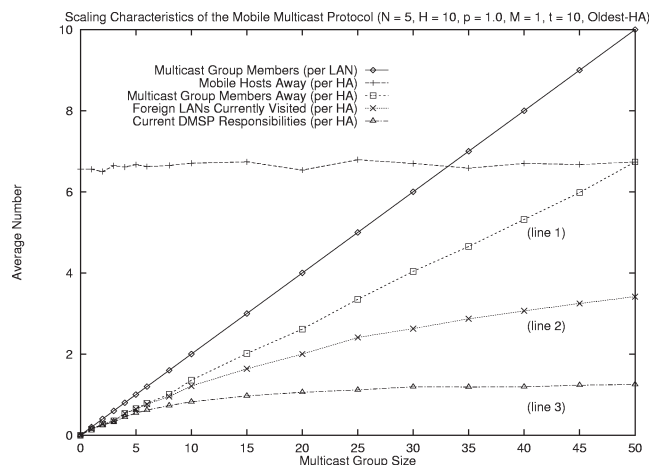


Figure 7. Scaling characteristics of the mobile multicast protocol with multicast group size (Oldest-HA DMSP selection policy).

per-HA basis from totals obtained for all HAs in the simulated network. The steepest line in figure 7 represents the expected number of multicast group members whose home network is the HA's LAN: this is simply a linear function of the multicast group size. The roughly horizontal line shows the average number of the HA's hosts that are away at any time; this is clearly independent of the multicast group size. The remaining three lines, from top to bottom, are: (1) the average number of the HA's mobile hosts that are multicast group members who are away from the home network; (2) the average number of foreign networks (i.e., FAs) at which these mobile multicast group members reside; and (3) the average number of these foreign networks for which the HA is the DMSP.

Since each MH has a set of multicast groups to which it is subscribed, the overhead in terms of packet density on the internetwork is greatly affected by the way in which these multicast packets are forwarded to the mobile hosts served by an HA. The most efficient method is remote subscription with a shared link multicast approach. Our simulation does not show this method. However, the cost of bi-directional tunnelling, direct forwarding by each HA, and DMSP forwarding can be inferred directly.

Bi-directional tunnelling requires that each HA forwards all multicast packets from groups to which its MHs are subscribed, to each MH individually. The number of packets transmitted in this scenario corresponds to the average number of MH group members away from home (line 1).

The number of packets delivered using direct forwarding by each HA (line 2) scales with the number of FAs visited. This is because each HA will forward multicast packets to each FA where its MHs are visiting regardless of whether or not the multicast packets are duplicated. DMSP forwarding (line 3) improves upon this by restricting the number of forwarding HAs for each foreign network to some small constant number.

Our simulations show that MoM has a clear advantage over bi-directional tunnelling in terms of the network traffic generated by multicast delivery. Bi-directional tunnelling performance (line 1) increases in a roughly linear fashion with the multicast group size. Both direct forwarding and DMSP forwarding increase sublinearly. At the largest multicast group size considered in the five-LAN scenario (50), direct forwarding halves the number of messages sent, while DMSP forwarding shows a further 50% reduction.

The following analysis also illustrates the advantage of MoM over bi-directional tunnelling:

- Let N be the number of networks in the system.
- Let G be the number of multicast groups that are forwarded to.
- Let c be the average number of MHs at each foreign network.
- Let k be the number of DMSPs forwarding packets.

Bi-directional tunnelling:
 $\#Messages = O(cN^2G)$.

In the worst case each home network will forward one packet to each foreign network for each multicast group for each host.

MoM:

$$\#Messages = O(kNG).$$

At most some constant k number of packets will be forwarded to each network for each group. Normally, each FA with visiting multicast group members chooses one DMSP, so $k \leq N$. Multiple DMSPs can be used, if desired, for robustness.

Our simulation results show that this performance advantage is evident even for the moderate number of mobile hosts and LANs used in our simulations.

4.2. Scalability with numbers of LANs and hosts

The MoM protocol scales well as the numbers of LANs and mobile hosts in the internetwork are increased. These results are illustrated in figure 8. Figure 8(a) shows the results for $N = 20$ LANs, each with $H = 10$ mobile hosts,

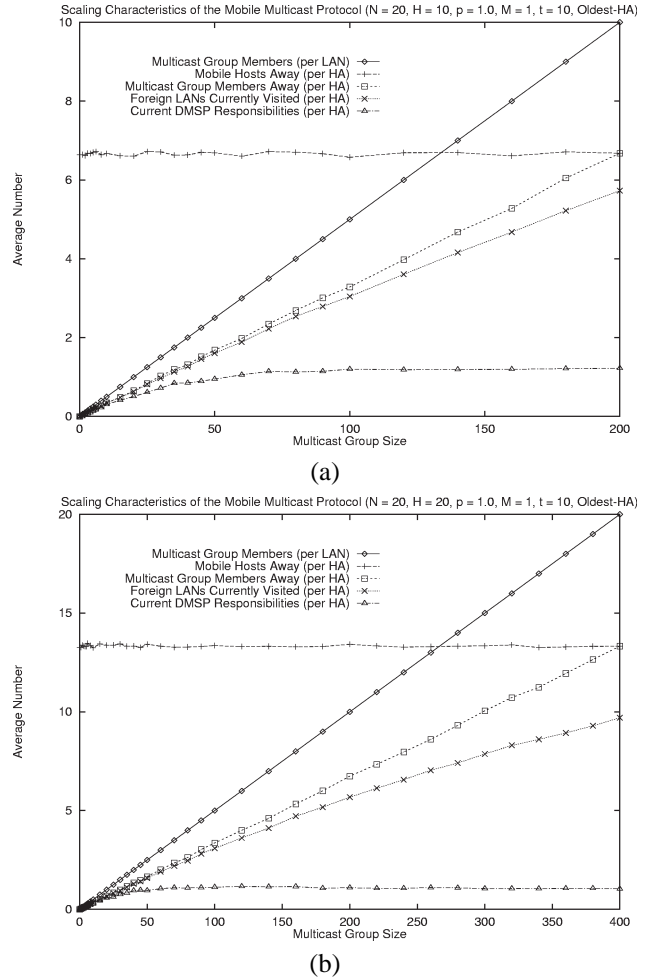


Figure 8. Scaling characteristics of the mobile multicast protocol with number of LANs and hosts (Oldest-HA DMSP selection policy): (a) Number of LANs $N = 20$, hosts per LAN $H = 10$; (b) Number of LANs $N = 20$, hosts per LAN $H = 20$.

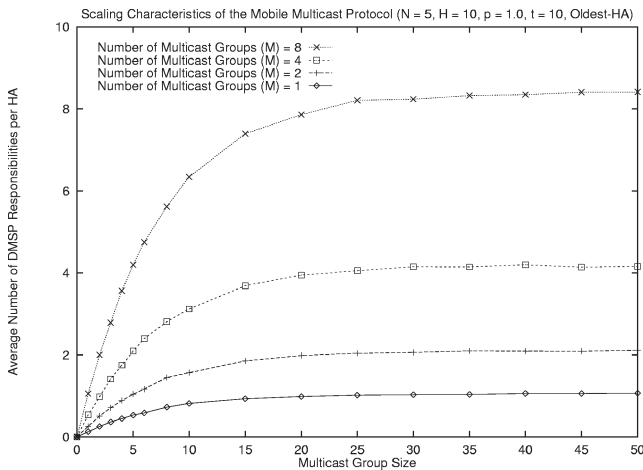


Figure 9. Scaling characteristics of the mobile multicast protocol with number of multicast groups (Oldest-HA DMSP policy).

and figure 8(b) shows the results for $N = 20$ LANs, each with $H = 20$ mobile hosts.

Figure 8(a) shows that with more LANs available to visit, the average number of foreign LANs visited at any time is close to the number of mobile host group members away (i.e., each away group member is at a different foreign network). However, the number of DMSP responsibilities per HA still remains fairly constant, and is only marginally higher than that in figure 7.

When the number of mobile hosts per LAN is increased (see figure 8(b)), the mobility model produces more away group members, with a resulting increase in number of foreign LANs visited. Nevertheless, the number of DMSP responsibilities per HA still remains low, scaling much less than linearly with the number of hosts and LANs in the network.

4.3. Scalability with number of multicast groups

Since each multicast group in the MoM protocol is handled independently, the overhead of the protocol scales linearly with the number of multicast groups. This behaviour is illustrated in figure 9, which shows the average number of DMSP forwarding responsibilities per home agent, as a function of multicast group size, for four different numbers of multicast groups in the network ($M = 1, 2, 4, 8$). Each multicast group is assumed to have the same number of members. As shown in the figure, the expected number of DMSP responsibilities increases additively with the number of multicast groups present in the network. For the largest multicast group sizes considered, there is an average of one DMSP for each multicast group in the network.

4.4. DMSP selection policies

The next set of simulation experiments focuses on the importance of the DMSP selection policy. If only one HA is to act as the forwarder for a multicast group G at a

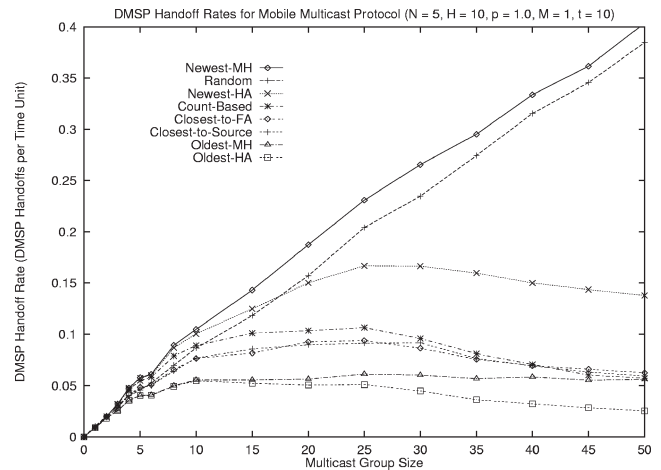


Figure 10. DMSP handoff rates, as a function of multicast group size and DMSP selection policy.

given foreign network, then how can this forwarder best be chosen? Three main issues must be considered: DMSP handoffs, route-optimality, and fairness.

4.4.1. DMSP handoffs

Figure 10 shows the number of DMSP handoff events that take place during the simulation, for each of the DMSP selection algorithms listed in section 3.3. The results are plotted as a function of multicast group size.

The DMSP selection policies fall into two categories: those for which the number of DMSP handoffs scales linearly with group size, and those for which DMSP handoffs scale sublinearly with group size. The linear-growth policies are Random and Newest-MH.⁵ Newest-MH is in fact worse than Random, since a FA that receives a MH arrival from a newly seen HA will always hand off, while Random in this situation will hand off only part of the time. Oldest-MH is much better than Newest-MH since it postpones handoff decisions as long as possible. Among the policies that exhibit sublinear growth, Oldest-HA clearly outperforms Newest-HA, for the same reason, with Count-Based falling in between these two. The two proximity-based policies provide low and stable handoff rates across a wide range of the multicast group sizes considered. The Oldest-HA policy has the lowest handoff rate of all.

Overall, the proximity-based policies perform extremely well. The explanation for this lies in the static nature of the network topology, and the law of large numbers. That is, once the number of visiting group members is large enough, there is likely to be a large list of HAs from which to choose a DMSP. For proximity-based policies, this means there is likely to be at least one visitor from a well-positioned HA, and, once this HA is chosen as the DMSP, it will likely remain as the DMSP for some time.

⁵ The Newest-MH policy can also be interpreted as a lower-bound estimate on the Mobile IP registration rate in the simulated network and the rate of multicast tree reconfiguration for remote subscription.

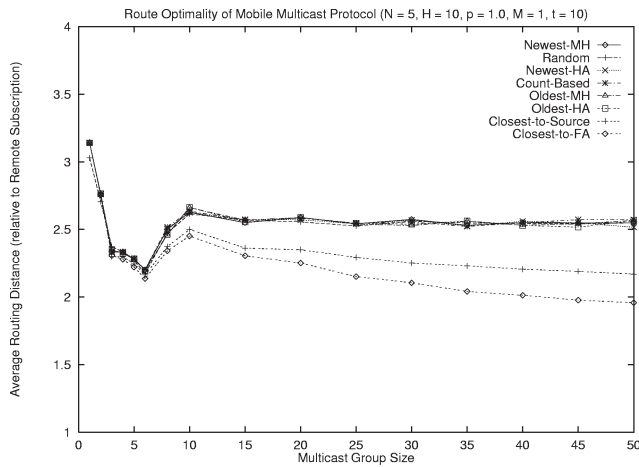


Figure 11. Routing efficiency as a function of multicast group size, for different DMSP selection policies.

Random policies make good choices on small lists (e.g., 1 or 2 HAs to choose from), and arbitrarily bad choices on large lists. Furthermore, the DMSP choice can change with each mobile host arrival and departure (since SelectDMSP is called). Count-based and proximity-based policies are much less sensitive to the dynamics of host arrival and departure, particularly in large groups. The separation between “random” and “good” DMSP selection policies takes place around group sizes of 15, in our simulations.

4.4.2. Route optimality

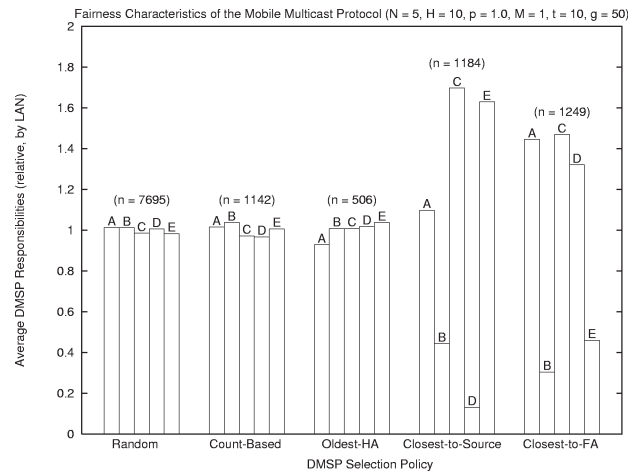
A second issue related to DMSP selection is the optimality of the routes used for multicast packets. That is, how long is the routing path to get a multicast packet to a group recipient via the DMSP, compared to routing the packet directly (remote subscription) or via the Home Agent (bi-directional tunnelling)?

We illustrate these performance results in figure 11. As is to be expected, the proximity-based policies provide the best DMSP routing performance, generally requiring routes 2 to 2.2 times longer than those of remote subscription.⁶ Routing performance improves with group size for these two policies, again because of the law of large numbers (i.e., with more HAs available to choose from, better choices can be made). The Closest-to-FA policy provides the best performance, with the Closest-to-Source next best. This makes sense intuitively, since the Closest-to-Source policy has no notion of direction to the FA when the DMSP is chosen. The remaining policies (age-based, count-based, and random) all perform similarly, with routes approximately 2.5 times longer than the optimal route.

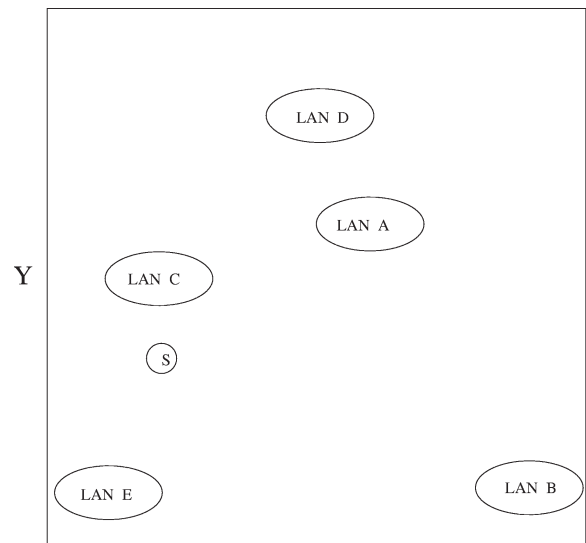
4.4.3. Fairness

A third issue to consider in terms of DMSP selection is fairness; i.e., how evenly is the DMSP forwarding task

⁶ The routing cost associated with DMSP forwarding is less severe on larger network scenarios (e.g., 1.6 for $N = 10$ LANs), but still significant, particularly since it represents a per-packet cost, rather than a per-host-movement cost like multicast tree reconfiguration.



(a)



(b)

Figure 12. Fairness characteristics of the mobile multicast protocol: (a) Load distribution of DMSP responsibilities among HA's, for different DMSP selection policies; (b) Network topology used for fairness study.

distributed amongst the HAs in the network. We illustrate these results in figure 12(a), focusing only on five DMSP selection policies: Random, Count-Based, Oldest-HA, Closest-to-Source, and Closest-to-FA. The graph illustrates the relative number of times each HA is chosen as the DMSP, as well as the total number (n) of DMSP selections made during the simulation. These results are for a single randomly generated network topology, shown in figure 12(b). The diagram shows the location of each LAN, as well as the location of the multicast source S .

Figure 12(a) shows that most of the policies provide fairly well balanced DMSP responsibilities, with the exception of the proximity-based policies, Closest-to-Source and Closest-to-FA. Closest-to-Source provides the most unbalanced loading since all foreign networks are likely to choose the same HAs as the DMSP, if opportunity permits.

With the Closest-to-FA policy, each FA is likely to choose one of its neighbours⁷ as the DMSP.

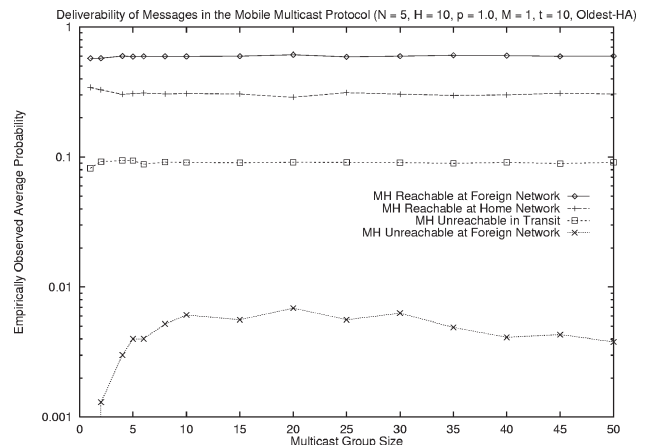
4.5. Deliverability of multicast messages

The next part of our simulation study addresses the effectiveness of multicast message delivery in the MoM protocol. That is, given that DMSP handoffs are occasionally required, how often is multicast message delivery disrupted for a mobile group member because of stale DMSP information? The answer to this question is shown in figure 13.

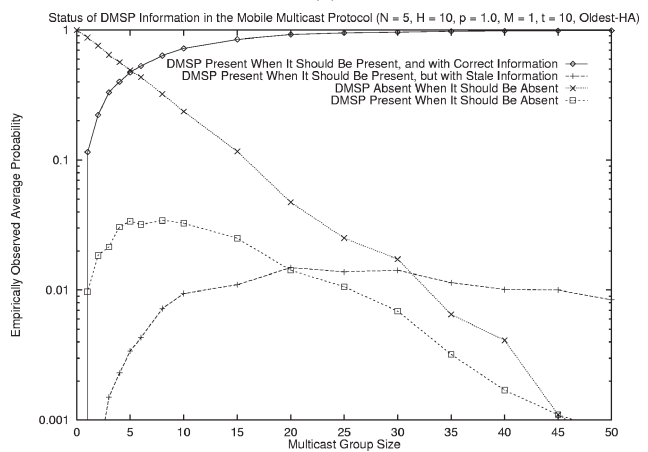
Figure 13(a) shows the empirically observed probabilities of successful and unsuccessful multicast message delivery, from simulation runs that generated approximately 10 000 multicast messages. There are four possible situations that can arise regarding delivery of a multicast message to a mobile host: (1) the host is at its home network; (2) the host is in transit; (3) the host is attached to a foreign network that has stale DMSP forwarding information, and (4) the host is attached to a foreign network that has up-to-date DMSP forwarding information. The information on the deliverability of messages was obtained using global (i.e., omniscient) information about host location at the times of message generation. Since the multicast message generation process is Poisson, and independent of the host movements in the simulation, it provides a statistically valid estimate of the deliverability of messages to each multicast group member. It also provides a means to assess the “correctness” of DMSP state information during the simulation (see figure 13(b)).

Figure 13(a) shows that approximately 30% of the time, group members are directly reachable at their home network, and 9% of the time group members are unreachable because they are in transit. These two values come directly from the host mobility model parameters, and are independent of multicast group size. The two other cases show a slight dependency on multicast group size. Approximately 60% of the time, a group member is successfully reachable at a foreign network via the DMSP, and approximately 1% of the time it is not. The unsuccessful deliveries arise because of stale DMSP information, wherein the foreign network and the home networks have inconsistent information about which HA is providing the DMSP service. This condition can arise because a departing mobile host provides no notification to the FA, but does inform its HA upon reattaching to a network (home or foreign). Until a timeout occurs, and a DMSP handoff completes, multicast message delivery can be disrupted for the mobile group members at the FA.

Figure 13(b) provides an indication of how these different deliverability conditions arise. In particular, the graph illustrates how often the DMSP information at a foreign network is up-to-date or not. Five possibilities exist:



(a)



(b)

Figure 13. Effectiveness of multicast message delivery in the MoM protocol: (a) Message deliverability results; (b) Status of DMSP state information.

- (1) DMSP information is absent when it should be absent (a correct state);
- (2) DMSP information is absent when it should be present (an incorrect state, which never happens in our simulations);
- (3) DMSP information is present when it should be absent (an incorrect state);
- (4) DMSP information is present when it should be present, but with stale information (an incorrect state); and
- (5) DMSP information is present when it should be present, and with up-to-date information (a correct state).

As can be seen, the correct states ((1) and (5)) dominate, particularly at small multicast group sizes (when few group members are mobile, and thus there is little DMSP information in the network) and at large multicast group sizes (when many group members are mobile, and DMSP information is updated frequently throughout the network). At intermediate group sizes (e.g., 10–20 group members), there are non-negligible probabilities (e.g., 1–5%) of incorrect DMSP state information existing in the network. It

⁷ Note, of course, that an FA can only choose from its current HA list when selecting a DMSP.

is this (transient) incorrect state that disrupts the message deliverability for some mobile group members.

The results for multicast message deliverability are largely independent of the DMSP selection policy used. The only notable exception is the Newest-MH policy (not shown here), which updates DMSP state information for multicast group members upon each mobile host arrival. As a result, this policy rarely has out-of-date DMSP information, particularly for large multicast group sizes. However, the improvements in message deliverability with the Newest-MH policy are marginal, and its excessive hand-off rate makes it unsuitable for the DMSP selection policy.

4.6. Timeouts

A more effective way to improve upon message deliverability in the MoM protocol is to shorten the lifetime of DMSP state information in the network. One way to do this is by reducing the timeout value used for mobile host registration (and re-registration). For example, figure 14

shows the message deliverability results (figure 14(a)) and the status of DMSP information (figure 14(b)) observed in the simulated network for five different timeout values, ranging from 60 time units (the mean host residency time at a LAN) down to 1 time unit. These results are shown for a multicast group size of $g = 10$.

Clearly, reducing the timeout value reduces the existence of stale DMSP state information in the network, and improves message deliverability. In fact, with a timeout value of 1 time unit, the probability of having stale DMSP information (i.e., states other than state 1 or state 5) in the network is negligible, since FA's initiate DMSP handoffs (when necessary) within 1 time unit of a mobile host's departure. The improvements in multicast message deliverability do not come for free, however, since short timeout values imply a significant increase in the number of re-registration events for mobility support agents.

4.7. Overhead

Our final comments concern the messaging overheads for the MoM mobile multicast protocol. Clearly, appointing DMSP's for multicast groups implies explicit messages between FA's and HA's at the time of DMSP selection. However, with a good choice of DMSP selection policy, such as Oldest-HA, the number of explicit messages required for DMSP assignment scales much less than linearly with multicast group size (see the DMSP handoff rates illustrated in figure 10). Thus the DMSP messaging overhead is significantly lower than the registration messaging overhead of Mobile IP.

Furthermore, it is possible to piggyback DMSP assignment information on the registration messages from a MH to its HA since, in many cases, the DMSP assignment is made to the HA of the first MH group member to arrive at a LAN. Figure 15(a) illustrates this point. In particular, the figure shows the fraction of DMSP assignment messages that could be piggybacked on HA registrations (since they occur "synchronously" with mobile host arrival), as well as those that could not be piggybacked (since they occur "asynchronously" as notifications during a MH's residency at the foreign LAN). The total number of DMSP assignments involved is also shown in the graph. As can be seen in figure 15(a), most of the DMSP assignment messages for small multicast group sizes can be piggybacked, since an arriving mobile host is likely to be the first mobile group member at the foreign LAN. For large multicast group sizes, only about one-tenth of the messages can be piggybacked. Thus the DMSP messaging overhead does increase a bit with multicast group size (at least up to $g = 25$).

The amortization of this messaging overhead actually improves, however, with multicast group size. This fact is illustrated in figure 15(b), which shows the cumulative distributions of the DMSP lifetimes observed in the simulated network for different multicast group sizes. The DMSP life-

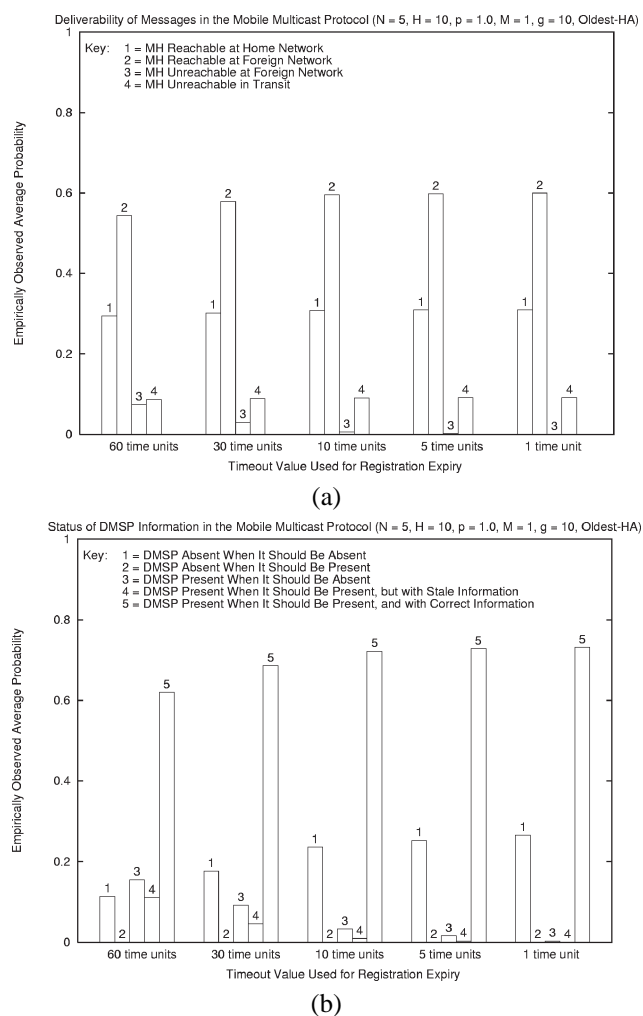


Figure 14. Effect of timeout value on multicast message delivery in the MoM protocol ($g = 10$): (a) Message deliverability results; (b) Status of DMSP information.

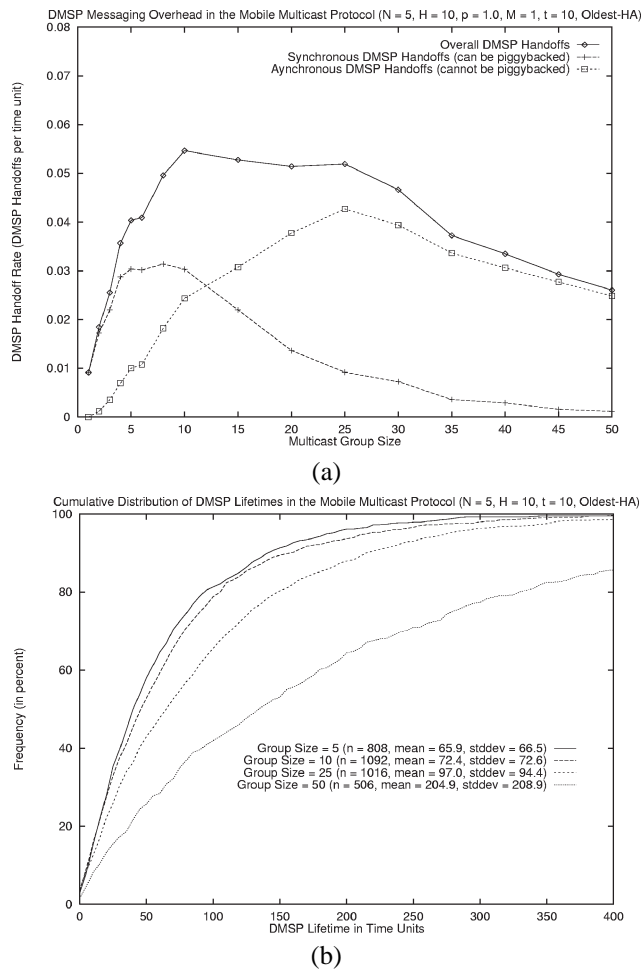


Figure 15. Assessment of DMSP messaging overhead in the MoM protocol: (a) DMSP assignment messages; (b) DMSP lifetime distributions.

time is defined as the consecutive time period during which an HA appointed as the DMSP remains as the DMSP for a given group for a given LAN. Figure 15(b) shows that the DMSP lifetimes are relatively short (about the same as the average mobile host residency time) for small multicast groups, but significantly longer for large multicast groups. In other words, for large multicast groups, the appointed DMSP is likely to remain the DMSP for a significant period of time (much longer than the average mobile host residency time). Long DMSP lifetimes reduce the number of DMSP handoffs, and also serve to amortize the cost of DMSP assignment messages.

5. Summary and conclusions

This paper evaluates a new approach for providing multicast to mobile hosts in an IP internetwork. The proposed scheme called MoM extends the reach of IP Multicast by using Mobile IP tunnels to constitute the “last mile” for delivery of multicast datagrams to mobile hosts. MoM has several features that make it practical as a solution for mobile hosts on IP internetworks. In particular, the use of designated multicast service providers (DMSPs) provides the

ability to support large, mobile multicast groups, and provides minimal break in service as a result of host movement.

Simulation results demonstrate the performance characteristics of the protocol, including its advantages over other approaches to mobile multicast, such as bi-directional tunnelling. While the approach does require modification to the home and foreign agents in Mobile IP, there are significant benefits to be had, particularly as the number of mobile group members increases. Furthermore, even simple DMSP selection policies, such as Oldest-HA, provide reasonable performance in terms of handoff rates, routing efficiency, fairness, deliverability of messages, and protocol overhead.

Acknowledgements

Financial support for this research came from the Natural Sciences and Engineering Research Council of Canada (NSERC) through grants OGP0121969, OGP0003707 and IOR152668, and from Telecommunications Research Laboratories (TRLabs). We also acknowledge the substantial contributions of our former colleague Vineet Chikarmane, from whom came the original ideas leading to the MoM protocol.

References

- [1] A. Acharya and B. Badrinath, Delivering multicast messages in networks with mobile hosts, in: *Proceedings of the 13th International Conference on Distributed Computing Systems*, Pittsburgh, PA (May 1993) pp. 292–299.
- [2] A. Acharya and B. Badrinath, A framework for the delivery of multicast messages in networks with mobile hosts, *Wireless Networks* (1996).
- [3] A. Acharya, A. Bakre and B. Badrinath, IP multicast extensions for mobile internetworking, in: *Proc. INFOCOM '96*, San Francisco, CA (March 1996).
- [4] K. Brown and S. Singh, RelM: Reliable multicast for mobile networks, Technical Report, Department of Computer Science, University of South Carolina (September 1995).
- [5] V. Chikarmane, Network support for mobile hosts in a TCP/IP internetwork, M.Sc. thesis, Department of Computer Science, University of Saskatchewan (August 1995).
- [6] V. Chikarmane, R. Bunt and C. Williamson, Mobile IP-based multicast as a service for mobile hosts, in: *Proceedings of the Second International Workshop on Services in Distributed and Networked Environments*, Whistler, BC, Canada (June 1995) pp. 11–18.
- [7] V. Chikarmane, C. Williamson and R. Bunt, Implementing mobile IP routing in 4.3 BSD Unix: A case study, *Internetworking: Research and Experience* 6(4) (1995) 209–227.
- [8] V. Chikarmane, C. Williamson, R. Bunt and W. Mackrell, Multicast support for mobile hosts using mobile IP: Design issues and proposed architecture, *Mobile Networks and Applications* (1998, to appear).
- [9] S. Deering and D. Cheriton, Multicast routing in datagram internetworks and extended LANs, *ACM Transactions on Computer Systems* 8(2) (1990) 85–110.
- [10] S. Deering, D. Estrin, D. Farinacci and V. Jacobson, An architecture for wide-area multicast routing, in: *Proceedings of the 1994 ACM SIGCOMM Conference*, London, UK (August 1994) pp. 126–135.
- [11] R. Droms, Dynamic host configuration protocol, RFC 1541, Bucknell University (October 1993).

- [12] T. Harrison, C. Williamson, W. Mackrell and R. Bunt, Mobile Multicast (MoM) protocol: Multicast support for mobile hosts, in: *Proceedings of the Third Annual ACM/IEEE Conference on Mobile Computing and Networking (MobiCom '97)*, Budapest, Hungary (September 1997) pp. 151–160.
- [13] D. Johnson, Scalable support for transparent mobile host internet-working, *Wireless Networks* 1(3) (1995) 311–321.
- [14] D. Johnson, Scalable and robust internetwork routing for mobile hosts, in: *Proceedings of the 14th International Conference on Distributed Computing Systems* (June 1994).
- [15] D. Johnson, A. Myles and C. Perkins, The Internet Mobile Host Protocol (IMHP), Internet Draft (February 1994).
- [16] J. Moy, Multicast routing extensions for OSPF, *Communications of the ACM* 37(8) (1994) 61–66.
- [17] C. Perkins, ed., IP mobility support, RFC 2002, IBM (October 1996).
- [18] D. Waitzman, C. Partridge and S. Deering, eds., Distance vector multicast routing protocol, RFC 1075, BBN STC and Stanford University (November 1988).
- [19] G. Xylomenos and G. Polyzos, IP multicast for mobile hosts, *IEEE Communications* 35(1) (1997) 54–58.




Carey L. Williamson is a Professor in the Department of Computer Science at the University of Saskatchewan, in Saskatoon, Canada, and an Adjunct Professor at Telecommunications Research Laboratories (TRLabs) in Saskatoon. He received a B.Sc. (Honours) in computer science from the University of Saskatchewan in 1985, and a Ph.D. in computer science from Stanford University in 1991. Dr. Williamson's general research interests are in computer networks and computer systems performance evaluation. More specific interests include network traffic measurement, workload characterization, network simulation, protocol performance, mobile computing, high speed networking, and ATM. He is a member of ACM.
E-mail: carey@cs.usask.ca



Tim G. Harrison received a Bachelor's degree in mechanical engineering from the University of Saskatchewan in 1986, and has been a licensed Professional Engineer with the Professional Engineers of Ontario (Canada) since 1989. In 1991, Tim returned to the University of Saskatchewan to pursue studies in computer science, and received an M.Sc. in computer science in 1995. From 1995 to 1997, Tim worked as a Research Assistant with the Department of Computer Science at the University of Saskatchewan. He is now working at the Memory Test Division of Teradyne, Inc., Agoura Hills, California.
E-mail: timh@std.teradyne.com



Wayne L. Mackrell received his B.Sc. (Honours) in computer science from the University of Saskatchewan in 1995. He is currently pursuing his M.Sc. in computer science at the University of Saskatchewan. His research interests include mobile computing, operating systems, and object-oriented languages.
E-mail: wlm125@cs.usask.ca



Richard B. Bunt is a Professor and former Department Head in the Computer Science Department at the University of Saskatchewan, and an Adjunct Professor at Telecommunications Research Laboratories (TRLabs). He holds a B.Sc. degree from Queen's University and M.Sc. and Ph.D. degrees from the University of Toronto. His general research interests are in the performance evaluation of computer systems and networks. Recently he has been working on issues relating to systems and networking support for mobile computing. He is a member of ACM.
E-mail: bunt@cs.usask.ca