Mobile Multicast (MoM) Protocol: Multicast Support for Mobile Hosts

Tim G. Harrison
Carey L. Williamson
Wayne L. Mackrell
Richard B. Bunt
Department of Computer Science
University of Saskatchewan
Saskatoon, Canada

Abstract

This paper[1] describes a new protocol to support IP multicast for mobile hosts in an IP internetwork. It uses the basic unicast routing capability of IETF Mobile IP as a foundation, and leverages existing IP multicast to provide multicast services for mobile hosts as well. We believe that the resulting scheme is simple, scalable, transparent, and to the extent possible, independent of the underlying multicast routing facility.

Discrete-event simulation was used as the vehicle for a "proof of concept" debugging of the protocol, as well as to determine its performance characteristics. A key feature of the new protocol is the use of designated multicast service providers (DMSPs) to address the scalability issues of mobile multicast. Our simulation results suggest distinct performance advantages of our protocol using DMSPs over two other approaches proposed for the mobile multicast problem, namely remote subscription and bi-directional tunnelling, particularly as the number of mobile group members increases.

1 Introduction

Providing multicast support for mobile hosts in an IP internetwork is a challenging problem, for several reasons. First, the IETF Mobile IP [18] protocol (as currently defined for IPv4) concentrates on *unicast* delivery to mobile hosts; additional mechanisms must be added to those proposed to efficiently 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 [21], 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 always a viable option, because of the overhead involved, yet leaving the multicast delivery tree unchanged can cause inefficient, incorrect, or even failure of multicast message delivery.

An approach for mobile multicast was proposed in [1] (and revised in [2]), but this proposal is not based on IETF Mobile IP and is limited in terms of scalability [8]. The current version of Mobile IP [18] proposes two approaches to support mobile multicast, which 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. Finally, the approach assumes the existence of a multicast router at the visited network, an assumption that may not always hold in an IP internetwork. 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 to effect delivery of multicast datagrams to mobile hosts, while maintaining the scalability of the approach through the use of a designated multicast service provider (DMSP) optimization per multi-

[1]To appear in *Proceedings of the Third Annual International Conference on Mobile Computing and Networking (MobiCom'97)*, September 26-30, 1997, Budapest, Hungary.

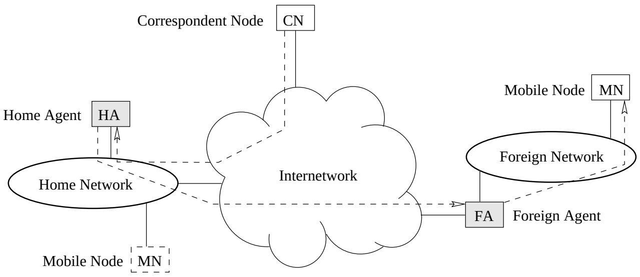

Figure 1: Mobile IP Entities and their Inter-relationships

cast group for each foreign network, and the use of dynamic multicast tunnels to foreign networks.

We evaluate our approach through simulation, focusing primarily on its scalability with respect to multicast group size, number of multicast groups, and number of mobile hosts. In particular, we demonstrate the performance trade-offs between remote subscription, bi-directional tunnelling, and the MoM approach across a wide range of network and workload parameters, and argue the suitability of our approach for the ever-growing world of mobile networking.

The remainder of this paper is organized as follows. Section 2 provides a brief background on the IETF Mobile IP protocol, and how its functionality can be leveraged to support mobile multicast. Section 3 discusses the design of the MoM mobile multicast protocol, identifying several issues raised by mobile multicast. These are the issues that MoM is intended to address. Section 4 presents a more detailed description of the MoM protocol. Section 5 presents our simulation setup for a performance study of mobile multicast, and Section 6 presents the results from the study. Finally, Section 7 presents the conclusions from our work.

2 Background

2.1 IETF Mobile IP

The IETF Mobile IP Working Group has defined a protocol called Mobile IP [18] to support (unicast) IP routing for mobile hosts in an IP internetwork. This protocol has roots in a few previously developed protocols for network mobility [13] [19] [20] [17] [14] [15]. The ideas in these protocols were discussed in the Internet forum and, with additions and refinements, resulted in the IETF Mobile IP protocol.

In Mobile IP, a mobile host (MH) is assigned an IP address on its home network, called the MH's *home address*. Packets from a correspondent host (CH) to the MH are always addressed to the home address of the MH, where they can be received and rerouted by the *home agent* (HA) if the MH is away.

When the MH connects to a foreign network, it identifies and registers with a *foreign agent, FA*, which also serves as its default router. An example scenario is illustrated in Figure 1. During the registration process, the FA provides the MH with a *care-of address*, defining the current location of the MH. The combination of the MH's home address and the care-of address is known as a *binding*. A timestamp originating from the MH and updated at each reregistration attempt is associated with every binding.

Each FA maintains a list of all currently registered MHs that are visiting the network it serves. A lifetime is associated with each visitor list entry and is negotiated during the registration process. A MH must reregister before this period has expired to ensure uninterrupted service from the FA. The home agent attached to the MH's home network maintains a *server cache* that identifies all MHs that it is configured to serve. The HA maintains a current binding for all the MHs in its server cache, and hence when a MH registers with a new FA, it must also register with its HA.

If the MH is at a foreign network, the HA must arrange to intercept packets addressed to the MH and *tunnel* the packet, by *encapsulating* and sending it to the FA that currently has the MH registered. When a FA receives a tunnelled packet addressed to a MH on its visitor list, the FA *decapsulates* and delivers the packet locally to the MH. This re-routing is illustrated by the dashed line in Figure 1.

The IETF Mobile IP also describes MHs operating independent of FAs using co-located addresses. In this mode of operation, the FAs are not required for routing, and the MH is itself a decapsulating and encapsulating agent. Addresses are obtained on the foreign network using DHCP [11] or another suitable method. Our protocol does not currently handle this mode of operation because of our use of foreign agents for various tasks. It is possible that MHs themselves could serve as proxy agents, but this option has not yet been fully explored.

2.2 The Challenges of Mobile Multicast

Host movement complicates multicast routing in a mobile environment, since the network must not only manage multicast group membership and establish the necessary routes, but must also contend with the fact that the established routes are themselves transient in nature.

For a MH that wishes to *receive* multicast datagrams, the routing problem is slightly different from the unicast routing problem addressed by Mobile IP for mobile hosts. Multicast datagrams have a class D IP address specified as their destination address. These addresses do not encode any specific network, and multicast routing of these datagrams can take place if the MH is not at its home network. However, the foreign network that the MH visits may not have a multicast router, but even if it does, multicast routers do not track movement, and therefore multicast packet delivery is disrupted by a location change.

Problems also arise for a MH that wishes to *send* multicast datagrams while away from its home network. These datagrams could possibly be dropped because downstream multicast routers receiving datagrams that are not topologically correct in the multicast delivery tree must drop them. Furthermore, downstream routers continue to track reverse paths to the MH assuming that the MH is on its home network. Thus the multicast routes that are established are always with reference to the MH's home network and are incorrect when the MH is at a foreign network.

A new approach to mobile multicast is described in [8]. Its details are summarized in the following two sections.

3 Our Approach to Mobile Multicast: 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 potentially of great utility if it can be leveraged to provide multicast services for mobile hosts as well.

3.1 Assumptions and Design Goals

The following assumptions have been made in the design of our approach:

- The service to be provided is the *unreliable, best effort, connectionless* delivery of multicast datagrams, i.e., datagrams may be lost, duplicated, delayed or delivered out of order. Higher level protocols are responsible for handling such conditions.
- Multicast support must conform to the host group model [9].
- 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.
- Foreign agents are being used at the foreign networks.
- The home agents and foreign agents are static (not mobile) hosts.

We make no *a priori* assumptions about the size of multicast groups, the geographic distribution of the multicast 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.* Finally, 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.

3.2 Design Overview

Since HAs are static and know the location of the MHs that they serve, it is logically feasible for a multicast HA to forward multicast traffic to the MH through the Mobile IP tunnel via the FA. As shown in Figure 2(a), the HA receives multicast datagrams via the IP multicast distribution tree. In this diagram, the (single) source of the multicasts is assumed to be at the root (top) of the tree, and the HA is one of many group members.

In our 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]. These join and graft latencies may not be much more than any handoff latencies needed by our protocol in some cases, but our protocol restricts the setup overhead to the mobile support entities and reduces the impact on internetwork routers not concerned with mobility support. Frequently moving mobile hosts are expected by the mobility support agents, but not by multicast routers, and the presence of many mobile hosts may overwhelm multicast routers not prepared for the rapidly changing delivery trees.

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, as shown in Figure 2(b). 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.

Unfortunately, our forwarding may result in a routing loop if the MH is residing on a network that is above its HA in the multicast routing tree. In this case, the packet from the HA gets re-forwarded down the tree to the HA which re-forwards it back up again. This can be fixed by enforcing a time-to-live of 1 hop on any multicast packet delivered from an FA (or using some other method for preventing the packet from being forwarded beyond that network).

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 scheme is complicated by a phenomenon called the *tunnel convergence problem* [6, 8] resulting from the fact that multiple Mobile IP tunnels (from different HAs) can terminate at a particular FA. This problem is illustrated in Figure 2(c), where multiple HAs (at different parts of the internetwork) all happen to have mobile hosts that are members of the same multicast group at the same foreign network, managed by the foreign agent 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 dealt with. In the worst case, the number of duplicate copies delivered increases with the number of mobile hosts present, a serious concern.

To solve the tunnel convergence problem, a selection is performed by the FA to appoint one HA as the *designated multicast service provider, DMSP*, for a given multicast group. The DMSP forwards only one multicast datagram into the tunnel even if it serves multiple MHs at the foreign network. 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]. Thus our scheme provides *at-most-once* delivery of multicast datagrams, which is identical to the semantics of IP multicast.

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 is the DMSP for a group at the (previous) foreign network, a DMSP handoff will be required to a different HA, to forward datagrams for the remaining

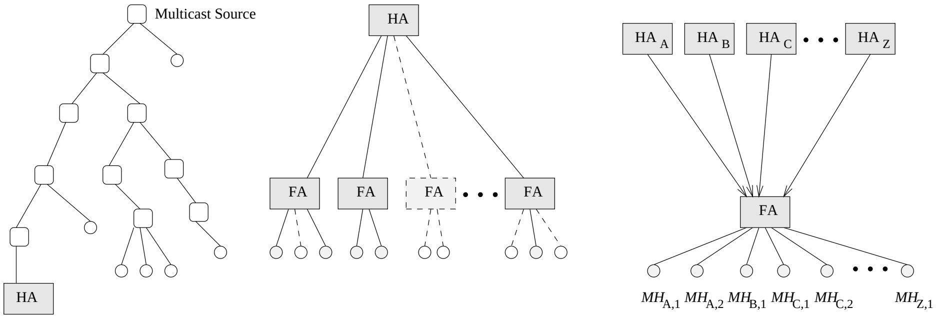

(a) Multicast tree (b) Propagation Tree (c) Tunnel Convergence problem

Figure 2: Multicast Packet Paths under the Proposed Scheme

multicast group members at the foreign network. Until this handoff completes, multicast delivery for group members at the foreign network may be disrupted. [2]

4 MoM Protocol Details

This section provides a more detailed description of the MoM protocol, in terms of the events handled and the data structures required.

4.1 Data Structures

Figure 3 illustrates the data structures needed to simulate the MoM protocol. Further details are given in [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 these mobile 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 3. 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 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 the list of HAs that are currently serving as DMSPs for this group.

In simple terms, the main issue addressed in this paper is the relative length of the three lists that agents maintain for each multicast group. Consider a single Home Agent, 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 much shorter, and the list of these foreign networks for which the Home Agent has DMSP responsibilities may be shorter yet. The DMSP optimization can thus be effective in reducing the multicast message forwarding load for Home Agents, reducing the multicast message traffic on foreign networks, and improving the scalability of mobile multicast.

[2] A protocol extension which 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 it's last MH's lifetime has expired.

4.2 Event Processing

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 MoM protocol are illustrated algorithmically in Figures 4 through 7, and are described below.

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 4 are executed. These steps update the visitor list at the Foreign Agent and the away list at the Home Agent, and update group membership information at each agent so that the DMSP status determined by the Foreign Agent is known by the Home Agent. Note that 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 5 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 protocol steps in Figure 6 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 7.

5 Mobile Multicast Simulation Model

We have evaluated our approach to the mobile multicast problem 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

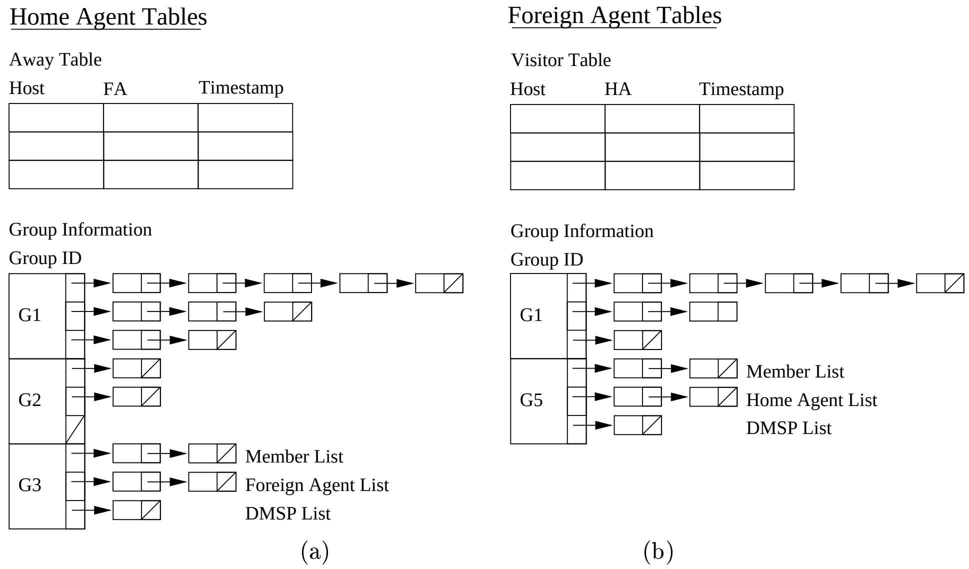

Figure 3: Data Structures for the MoM Protocol: (a) At the Home Agent; (b) At the Foreign Agent

When a mobile host (MH) arrives at a foreign network:

1. Register with Foreign Agent (FA)
   1.1 Create Visitor Table entry for MH.
   1.2 Insert host name, HA info, and expiry time.
   1.3 Notify FA of MH's current group memberships.
   For each multicast group G to which MH belongs:
      1.3.1 Make new entry in GroupInfoTable, if required.
      1.3.2 Add MH to the 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 from the HA list for this group.
      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 expiry time.
   2.4 Notify HA of MH's group memberships.
   For each multicast group G to which MH belongs:
      2.4.1 Make new entry in GroupInfoTable, if required.
      2.4.2 Add MH to G's group membership list, if required.
      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 is different from the oldFA then decrement the host count for the MH's oldFA, discarding oldFA from list if count is zero.
      2.4.5 Record/update DMSP status (YES/NO) of this HA for group G at FA (and for oldFA, if required).

Figure 4: Protocol Steps Required Upon Arrival of a Mobile Host 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, recording oldFA.
   1.2 For each multicast group G to which MH belongs:
      1.2.1 Delete MH from the group membership list for G.
      1.2.2 Decrement the host count for the MH's oldFA, discarding oldFA from the FA list if count is zero, and deleting oldFA from DMSP list, if required.

Figure 5: Protocol Steps Required Upon Return of a Mobile Host to its Home Network

When a MH times out at a foreign network:

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

Figure 6: Protocol Steps Required Upon Timer Expiration for a Mobile Host at a Foreign Network

When a unicast packet for a MH arrives at the MH's HA:
1. Look up FA information for MH in the Away Table.
2. Encapsulate the packet and tunnel it to the FA.

When a multicast packet for group G arrives at an HA:
1. Forward the multicast packet to local members.
2. Look up membership information for away members for that group.
3. Encapsulate the 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 an FA from an 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 7: Protocol Steps Required for Message Delivery to Mobile Hosts

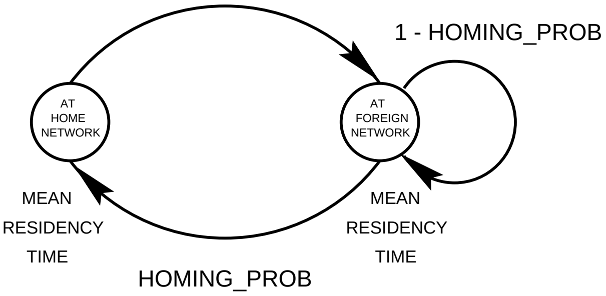

Figure 8: Host Mobility Model

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.

5.1 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, it is reasonable to assume 100% for our simulations since no extra insight is gained by simulating stationary hosts.

Figure 8 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. Mobile hosts begin the simulation at

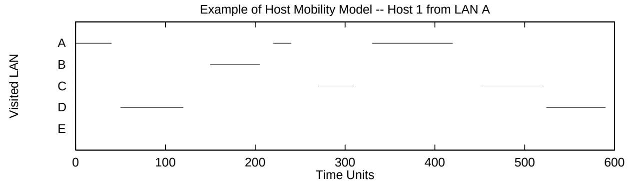

Figure 9: A Time Plot Illustrating Mobile Host Behaviour in the Simulation Model

their home network, and are allowed to roam about in the network at random. Foreign networks to visit 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 $\frac{1}{h} = 2$. The residency time for each visit to a network (home or foreign) is drawn from an exponential distribution with a mean of 60 time units [3], and the travel time for going between networks (regardless of distance) is exponentially distributed with a mean of 15 time units. Hosts thus spend 20% of their time in transit, and 80% of their time connected to a LAN (53.3% at foreign networks, and 26.7% at home) [4]. Each mobile host operates independently. The mobility behaviour of a typical mobile host is illustrated graphically in the timing diagram shown in Figure 9.

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, centrally 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. Note that 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 simulator to support dynamic group membership is future work. More complex simulations involving fixed nodes, existing multicast routers, and other complications is also 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.

5.2 Simulation Methodology

The simulation experiments were conducted using a multifactor experimental design. We used the simulator to assess

[3] Since time is not an issue in our simulations, we choose 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
N	Number of LANs	5
H	Hosts per LAN	10
p	Fraction of hosts that are mobile	1.0
RES	Percentage of time at a network	80%
TR	Percentage of time in transit	20%
TO	Registration timeout value (in time units)	60
h	Homing probability	0.50
M	Number of multicast groups	1
g	Multicast group size	1..50
s	Sources per multicast group	1
λ	Multicast message generation rate (msgs/time unit)	0.1

the relative performance impact of DMSP selection algorithm, multicast group size, network size, and number of mobile hosts.

The warmup period used for the simulations was 600 time units, following which we collect simulation statistics relating to mobile multicast until the end of the simulation run (6000 additional time units). Simulation runs were performed on 10 different randomly-generated network topologies for each set of workload parameters, with the results averaged for the purposes of plotting.

5.3 Experimental Factors and Levels

The two main factors considered were multicast group size and DMSP selection algorithm. The multicast group size was varied from 1 to 50 (only one multicast group at a time was used in the simulation). Several different DMSP selection algorithms were studied:

- **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-Newest-HA:** The DMSP is initially chosen as the oldest in the HA list. When this HA can no longer be the DMSP, the newest entry in the HA list is chosen to be 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.
  - **Oldest-Newest-MH:** The DMSP is initially chosen as the HA of the oldest MH at the FA. When this HA can no longer be the DMSP, 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.

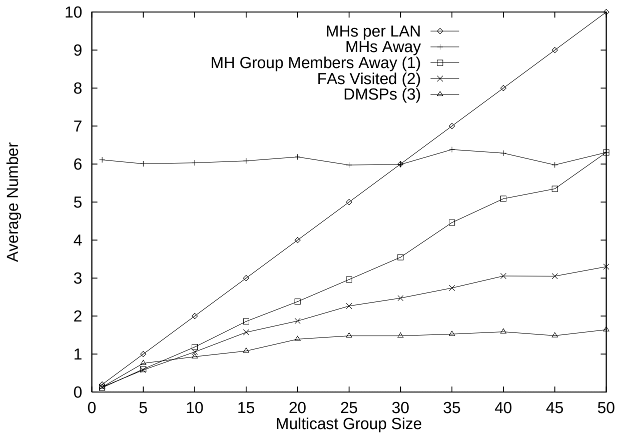

Figure 10: Scaling Characteristics of the Mobile Multicast Protocol with Multicast Group Size (for a single HA and the Oldest-HA DMSP selection policy)

- Proximity-based algorithms:
  - **Closest-to-FA:** The HA that is closest to the FA (i.e., shortest Euclidean distance) is chosen as the DMSP.
  - **Closest-to-Source:** The HA that is closest to the multicast source is chosen as the DMSP.

In total, 100 combinations of factors and levels were employed in this experiment.

6 Simulation Results

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

6.1 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 10, for one selected Home Agent from the five-LAN simulation, using the Oldest-HA DMSP selection policy. The results for other HAs and other DMSP algorithms are similar.

Figure 10 illustrates how various aspects of the mobile routing environment scale as the multicast group size is increased. The steepest line in Figure 10 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.

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 MHs 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 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 (50), direct forwarding halves the number of messages sent, while DSMP 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 redundant DMSPs forwarding packets.
>
> Bi-directional Tunnelling:
> $\#Messages = \mathrm{O}(c * N^2 * G)$
> In the worst case each home network will forward one packet to each foreign network for each multicast group for each host.
>
> MoM:
> $\#Messages = \mathrm{O}(k * N * G)$
> At most some constant k number of packets will be forwarded to each network for each group.

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

6.2 DMSP Selection Algorithms

The second simulation experiment focuses on the importance of the DMSP selection algorithm. If only one HA is to act as the forwarder for a multicast group G at a given foreign network, then how can this forwarder best be chosen? Two main issues must be considered, DMSP handoffs and route-optimality.

6.2.1 DMSP Handoffs

Figure 11 shows the number of DMSP handoff events that take place during the simulation, for each of the DMSP selection algorithms listed in Section 5.3. The results are plotted as a function of multicast group size.

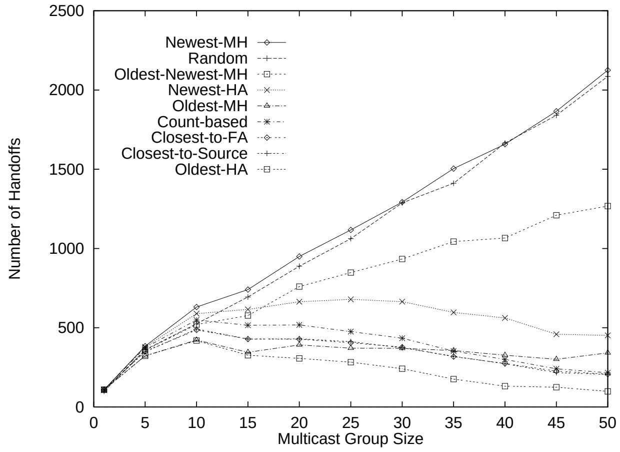

Figure 11: Number of DMSP Handoffs that Occur, as a Function of Multicast Group Size and DMSP Selection Algorithm

The resulting DMSP selection algorithms 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 all policies based directly on MH state information (e.g., MH arrival time). Among these policies, Oldest-MH performs best, since it always postpones handoff decisions as long as possible, and Newest-MH performs worst, since it forces many unnecessary handoffs. 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.

The sublinear-growth policies are those based on HA state information: age-based, count-based, and proximity-based. Among these, Oldest-HA clearly outperforms Newest-HA, for the reasons stated above, 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.

Random policies make relatively 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.

6.2.2 Route Optimality

A second issue related to DMSP selection is the optimality of the routes multicast packets take to their destination. That

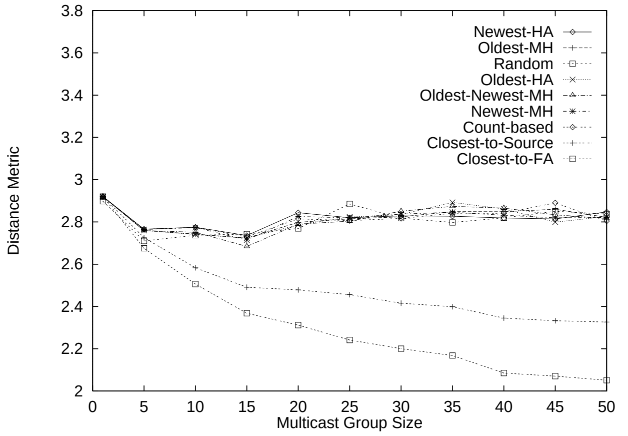

Figure 12: Illustration of Route Optimality as a Function of Multicast Group Size, for Different DMSP Selection Algorithms

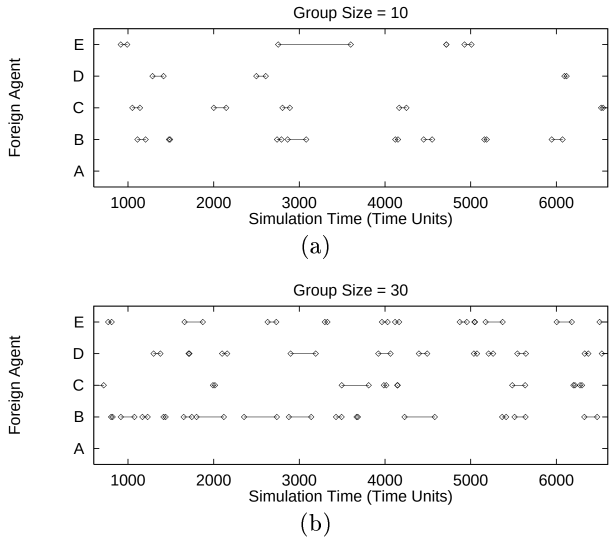

Figure 13: Dynamic Membership of DMSP Multicast Group G_{DMSP} Over Time using Oldest-HA for DMSP selection: (a) G_{DMSP} Membership for a Multicast Group of Size 10; (b) G_{DMSP} Membership for a Multicast Group of Size 30

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 12. As is to be expected, the proximity-based policies provide the best DMSP routing performance, generally requiring routes 2 to 2.5 times longer than those of remote subscription. 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 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 almost 3 times longer than the optimal route.

6.3 DMSP Groups

Our final experiment addresses the use of multicast tunnels to tunnel multicast packets to foreign networks. That is, rather than having the HA forward a tunnelled multicast packet (for group G) to each FA for which it is the DMSP for group G (using *unicast* Mobile IP tunnels), the HA can create a dynamic multicast group G_{DMSP} consisting of the FAs for which it is the DMSP for group G, and transmit a tunnelled multicast packet to G_{DMSP} using IP multicast. This approach has obvious advantages for the HA, in terms of the number of messages forwarded. However, the approach is worthwhile only if the average size[5] of group G_{DMSP} is larger than 1 (Figure 10 shows that we average 1.4 for groups of size 20 and larger, but this should increase with more LANs), and the membership of group G_{DMSP} is relatively static (compared to the dynamics of host arrivals and departures, and FA additions and deletions).

Figure 13 presents simulation results addressing this issue. This figure plots the membership of group G_{DMSP} as a function of time for one selected HA in our simulated internetwork. Figure 13(a) tracks the membership of G_{DMSP} for a multicast group of size 10, while Figure 13(b) tracks the membership for a group of size 30.

The DMSP multicast group approach appears worthwhile for the larger multicast group sizes because of the larger size and greater continuity of membership in the DMSP group. In this situation, there are performance advantages to the DMSP multicast group approach. This advantage increases with the number of LANs and mobile hosts in the internetwork.

7 Summary and Conclusions

This paper considers 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, it provides the ability to support large, mobile, dynamic multicast groups, and provides minimal break in service as a result of host movement.

We believe that such an approach offers important advantages. First, it requires minimal modifications to IP multicast and Mobile IP: primarily the addition of mechanisms to handle issues that arise only when multicast and host mobility are both present in an internetwork. Furthermore, it has scalability and performance advantages over other approaches to mobile multicast, such as bi-directional tunnelling. Compared to the remote subscription model, MoM offers greater generality and functionality, but must sacrifice on routing latency for multicast messages.

[5] We assume that the average size of the group G_{DMSP} will be significantly less than the number of foreign networks on the internetwork. Thus our approach is not the same as configuring a static multicast tunnel to all foreign agents in the internetwork [3].

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 (TR*Labs*). We also acknowledge the substantial contribution 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", *Proceedings of the 13th International Conference on Distributed Computing Systems*, Pittsburgh, PA, pp. 292-299, May 1993.

[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", Proc. INFOCOM'96, San Francisco, CA, March 1996.

[4] A. Ballardie, P. Francis, and J. Crowcroft, "Core Based Trees (CBT): An Architecture for Scalable Inter-Domain Multicast Routing", *Proceedings of the 1993 ACM SIGCOMM Conference*, San Francisco, CA, pp. 85-95, September 1993.

[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", *Proceedings of the Second International Workshop on Services in Distributed and Networked Environments*, Whistler, BC, Canada, pp. 11-18, June 1995.

[7] V. Chikarmane, C. Williamson, and R. Bunt, "Implementing Mobile IP Routing in 4.3 BSD Unix: A Case Study", *Internetworking: Research and Experience*, Vol. 6, No. 4, pp. 209-227, December 1995.

[8] V. Chikarmane, C. Williamson, R. Bunt and W. Mackrell, "Multicast Support for Mobile Hosts Using Mobile IP: Design Issues and Proposed Architecture", *ACM/Baltzer Mobile Networking and Applications*, 1997, (to appear).

[9] S. Deering and D. Cheriton, "Multicast Routing in Datagram Internetworks and Extended LANs", *ACM Transactions on Computer Systems*, Vol. 8, No. 2, pp. 85-110, May 1990.

[10] S. Deering, D. Estrin, D. Farinacci, and V. Jacobson, "An Architecture for Wide-Area Multicast Routing", *Proceedings of the 1994 ACM SIGCOMM Conference*, London, UK, pp. 126-135, August 1994.

[11] R. Droms, *Dynamic Host Configuration Protocol*, RFC 1541, Bucknell University, October, 1993.

[12] J. Ioannidis, D. Duchamp, and G. Maguire, "IP-based Protocols for Mobile Internetworking", *Proceedings of the 1991 ACM SIGCOMM Conference*, Zürich, Switzerland, pp. 235-245, September 1991.

[13] J. Ioannidis, "Protocols for Mobile Internetworking", Ph.D. Thesis, Columbia University, 1993.

[14] D. Johnson, "Scalable and Robust Internetwork Routing for Mobile Hosts", *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*, Vol. 37, No. 8, pp. 61-66, August 1994.

[17] A. Myles and C. Perkins, *Mobile IP Extensions (MIP)*, Internet Draft, October 1993.

[18] C. Perkins (editor), *IP Mobility Support*, RFC 2002, IBM, October, 1996.

[19] C. Perkins and Y. Rekhter, *Short-cut Routing for Mobile Hosts*, Internet Draft, IBM T.J. Watson Research Center, 1992.

[20] F. Teraoka, Keisuke Uehara, Hideki Sunahara, and J. Murai, "VIP: A Protocol Providing Host Mobility", Communications of the ACM, Vol. 37, No. 8, pp. 67-75, August 1994.

[21] D. Waitzman, C. Partridge, and S. Deering (editors), *Distance Vector Multicast Routing Protocol*, RFC 1075, BBN STC and Stanford University, November 1988.