Multicast Support for Mobile Hosts Using Mobile IP: Design Issues and Proposed Architecture

Vineet Chikarmane* Carey L. Williamson Richard B. Bunt Wayne Mackrell
Department of Computer Science
University of Saskatchewan
57 Campus Drive
Saskatoon, SK, CANADA
S7N 5A9
Phone: (306) 966-8656
FAX: (306) 966-4884
Email: {carey,bunt,wlm125}@cs.usask.ca

May 16, 1997

Abstract

In this paper, we consider the problem of providing multicast to mobile hosts using Mobile IP for network routing support. Providing multicast in an internetwork with mobile hosts is made difficult because many multicast protocols are inefficient when faced with frequent membership or location changes. This basic difficulty can be handled in a number of ways, but three main problems emerge with most solutions. The *tunnel convergence problem*, the *duplication problem*, and the *scoping problem* are identified in this paper and a set of solutions are proposed.

The paper describes an architecture to support IP multicast for mobile hosts using Mobile IP. The basic unicast routing capability of Mobile IP is used to serve as the foundation for the design of a multicast service facility for mobile hosts. We believe that our scheme is transparent to higher layers, simple, flexible, robust, scalable, and, to the extent possible, independent of the underlying multicast routing facility. For example, our scheme could interoperate with DVMRP, MOSPF, CBT, or PIM in the current Internet. Where differences exist between the current version of IP (IPv4) and the next generation protocol (IPv6), these differences and any further optimizations are discussed.

*Presently at Glenayre Technologies Inc. in Burnaby, BC.

1 Introduction

This paper considers the problem of multicast to dynamic groups in a TCP/IP internetwork with mobile hosts. Although not common today, multicast operation with mobile hosts may become widespread in the future. Our approach is to adapt the Mobile IP protocol so that it can handle multicast forwarding with adequate scalability.

Today's user community is demanding a level of mobility not previously anticipated by designers of distributed systems and computer networks. Hardware technological advances have made inexpensive and powerful portable computers a reality. This convergence of a desire for mobility and products to satisfy that desire has created an explosion of interest in problems related to the mobile computing paradigm [Imielinski and Badrinath 94, Forman and Zahorjan 94].

Concurrent with this growth in the popularity of mobile computing has been a significant growth, in the Internet community at least, in the use of *multicast* network applications. Multicast is a mechanism for efficient *one-to-many* communication in which the source transmits a single datagram, and the network performs the task of delivering that datagram to the multiple destinations. By eliminating multiple transmissions, multicast results in significant savings in source host processing and network bandwidth.

Multicast operation on the Internet is now supported for fixed hosts through IP multicast [Deering and Cheriton 90, Deering 91]. Examples of multicast applications currently used on the Internet include resource discovery, as well as applications such as shared whiteboard and desktop video/audio conferencing on the Internet multicast backbone (MBONE) [Eriksson 94].

Multicast should be valuable for mobile hosts as well. Multicast services such as the dissemination of textual information like weather reports, travel information, and stock market reports may become widely used by mobile users. Applications such as e-mail, multi-person communication, service location and distributed systems functions like cache consistency could also make use of a general purpose multicast facility. As well, wired mobile hosts without the wireless bandwidth constraints will still require multicast support for the more expensive applications of today such as the MBONE.

Unfortunately, the provision of multicast services to mobile hosts proves to be a very challenging problem, for several reasons. First, even *unicast* routing for mobile hosts is a difficult problem, since the routing of datagrams intended for (or coming from) a mobile host changes whenever the mobile host changes location. Second, all existing multicast routing proposals (including DVMRP [Waitzman *et al.* 88], MOSPF [Moy 94a], CBT [Ballardie *et al.* 93], and PIM [Deering *et al.* 94]) implicitly assume *stationary* hosts when configuring the multicast delivery tree. Since the delivery trees established for static multicast cannot be changed easily or efficiently, the movement of a host (either a sender or a receiver) after the tree is constructed can create problems.[1] Finally, the mobile computing environment itself adds additional complexity to the problem. For example, in most wireless implementations of mobile computing, network bandwidth is scarce, error rates are higher, movement can be frequent, and the changing point of network attachment for a mobile user may mean that a multicast router is not always directly accessible.

The IETF Mobile IP specification [Perkins 96a] defines a method for routing packets to mobile

[1] For example, if the source of a multicast should change location, the *incoming interface check* conducted in most multicast routing algorithms could result in discarding datagrams, rather than propagating them downstream to the multicast recipients.

hosts. It also defines two multicast support options, which we will call *remote subscription* and *bi-directional tunnelled multicast*. We describe a third option, similar to bi-directional tunnelled multicast but with additional mechanisms for addressing specific problems associated with the scalability of mobile multicast. Our scheme makes use of the foundation provided by a Mobile IP implementation and IP multicast. The goal of the scheme is to provide multicast capability to mobile hosts transparently in a manner similar in function to that available on the Internet today, yet not suffering from problems specific to routing multicast packets to mobile hosts. This approach has advantages and disadvantages when compared to the IETF proposals.

The remainder of this paper is organized as follows. Section 2 provides some background on Mobile IP, IP multicast, the challenges introduced for multicast routing by host mobility, and the provisions for handling multicast in the current mobile IP specification. Section 3 outlines the approach of our protocol for supporting multicast for mobile hosts. Section 4 gives a more detailed description of the protocol, including some performance results and a discussion of advantages and limitations. Section 5 discusses similar work that has been done by other research groups. Section 6 summarizes our paper and presents our conclusions.

2 Background: Mobile IP and IP Multicast

2.1 IETF Mobile IP

The IETF Mobile IP Working Group has defined a protocol to support unicast[2] IP routing for mobile hosts in a TCP/IP internetwork. 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. If the CH is aware that the MH is mobile, then the CH can encapsulate its packets and forward them directly to the MH's new location without traversing the home network [Johnson 96].

In general, when the MH connects to a foreign network, it identifies and registers with a *foreign agent* (FA), or registers directly with its *home agent* (HA). When registering, the MH acquires a care-of address defining its current location. The combination of the MH's home address and the care-of address is known as a *binding*. The MH can acquire its care-of address either from a foreign agent or through autoconfiguration methods (such as DHCP [Droms 93]) designed for assigning temporary IP addresses. If the care-of address is assigned by a foreign agent, it will be the address of the FA itself. If the address is assigned via autoconfiguration, it is said to be *co-resident* with the MH, and will be some unused address on the foreign network. When the care-of address for the MH is not co-resident, the FA is responsible for decapsulating and forwarding packets to the MH on the foreign network using link-layer protocols. The FA must also serve as the MH's default router. If the address is co-resident, the MH itself must decapsulate datagrams. An example scenario is illustrated in Figure 1.

FAs responsible for managing MH registrations maintain a list of all currently registered MHs that are visiting the network served by the FA. A lifetime is negotiated during the registration process for each visiting MH and is stored in the visitor list entry for that MH. A MH must reregister before this period has expired to ensure uninterrupted service from the FA. A timestamp originating

[2]This proposal also includes two methods for handling multicast packets, discussed in Section 2.4.

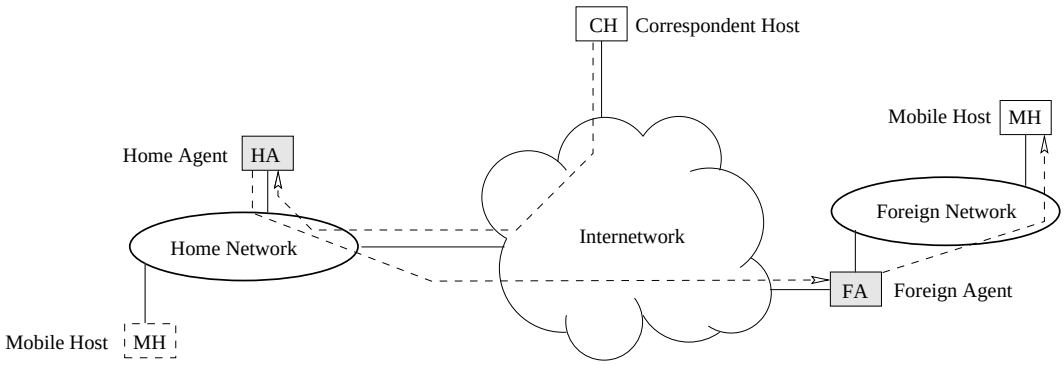

Figure 1: Mobile IP entities and their inter-relationships

from the MH and updated at each reregistration attempt is associated with every binding. This timestamp is used to determine timeouts for registrations. If a MH has not reregistered within the timeout period, the FA assumes that the MH has moved to another network.

A HA on the MH's home network serves as a binding cache and forwarding agent for the MH. Optimized routing extensions allow these forwarding capabilities to also exist at the CHs themselves. Lifetimes are associated with all bindings, and when bindings expire, the HA or CH assumes the MH has returned to its home network. During the lifetime of a binding, the HA is responsible for intercepting packets addressed to the MH and forwarding them to the care-of address. These packets are forwarded using IP-in-IP tunnelling, which encapsulates IP packets inside other IP packets using a forwarding header [Perkins 96b]. Facilities are required at the care-of address for decapsulating the packets when they arrive and forwarding the packets to the MH.

A good discussion of Mobile IP and other routing options available to both the mobile host and the correspondent host can be found in [Cheshire and Baker 96].

2.2 Internet IP Multicast

Deering [Deering and Cheriton 90, Deering 91] designed a form of WAN multicast that has been incorporated into the TCP/IP suite. Conventions have been defined for recognizing a particular class of IP addresses, the Class D IP addresses, as being *multicast addresses* [3] and for mapping these addresses to MAC layer multicast addresses. The *Internet Group Management Protocol*, IGMP, is used by hosts to inform multicast routers about the multicast addresses that the hosts wish to receive. A multicast routing protocol, in the form of an extension to an existing routing protocol, is executed on routers to construct datagram delivery paths and to accomplish datagram forwarding to the multicast recipients.

Internet IP multicast is designed to support the *host group* model [Deering 91]. Important properties of the host group service model are: group membership can be dynamic (i.e., hosts can join or leave the group at any time); membership is receiver-initiated; the sender need not know

[3] In the proposed IPv6 address space, a large class of addresses has also been set aside for multicast. Embedded within the IPv6 addresses are 4 bits which are used as flags to determine the *scope* of the multicast. Multicast scopes include organization, site, link, and node. These scopes are commonly emulated in IPv4 by restricting the *Time To Live* field for the multicast packet so that it is unlikely to propagate beyond its desired scope.

the location or identity of any group member (i.e., only the group address needs to be known); and the sender itself need not be a member of the group.

While the host group model defines the interface for users of the internetwork multicast service, the host group model does not define how that underlying multicast service is implemented. Several algorithms have been proposed for internetwork multicast, including DVMRP [Waitzman *et al.* 88], MOSPF [Moy 94a], Core Based Trees (CBT) [Ballardie *et al.* 93], and Protocol Independent Multicast (PIM) [Deering *et al.* 94]. We briefly review these algorithms here.

MOSPF: Multicast open shortest path first (MOSPF) is a multicast extension of the OSPF V2 link-state unicast routing protocol [Moy 94a]. MOSPF routers monitor system topology using IGMP messages, and propagate changes by flooding the internetwork with updates to other MOSPF routers. When multicast datagrams are encountered by a router it will calculate a shortest-path tree rooted at the source network of the multicast datagram, and reaching all member networks in the multicast group. As a result of the MOSPF design, the path taken for any multicast datagram will be the least cost path available in the internetwork. Also, owing to the MOSPF design, all routers in the internetwork will produce the same shortest-path tree, and therefore route datagrams consistently.

DVMRP: DVMRP [Waitzman *et al.* 88] is a distance-vector multicast routing protocol that has been implemented as an extension to RIP [Malkin 94]. DVMRP has been widely implemented, and is deployed in hundreds of regions connected by the MBONE [Eriksson 94]. DVMRP uses the *truncated reverse-path broadcast* [Deering and Cheriton 90] routing algorithm. In DVMRP, a datagram from a multicast source is initially propagated downstream (using either link-level multicast or a multicast *tunnel*) by a designated multicast router R_D to *all* other multicast routers, regardless of whether they have multicast group members or not. Multicast routers without downstream or local members send explicit *prune* messages upstream to remove themselves from the distribution tree. The net effect is a source-specific shortest path tree with the members forming the leaves of the tree. Once the multicast tree is set up, multicast routers keep track of the reverse path to the multicast source. If an arriving datagram does not come through the interface that the router uses to send datagrams to the source of the multicast, then the arriving datagram is dropped.

Core Based Trees (CBT): One of the drawbacks of DVMRP is that multicast delivery trees are constructed *for each source* in each multicast group. That is, the state information required at multicast routers scales with S x N, where S is the number of multicast sources, and N is the number of multicast groups. Ballardie *et al.* [Ballardie *et al.* 93] propose Core Based Trees (CBT) as a solution to this problem. The CBT approach constructs a single multicast delivery tree, centered around a *core* node, for each multicast group. This tree is then *shared* by all sources for the multicast group. As a result, state information at multicast routers scales only with N, the number of multicast groups, and not with the number of sources S. Thus, for wide area network multicast, the CBT approach has much greater scalability for wide area network multicast than DVMRP. Other advantages of the approach include simplicity, flexibility, independence from underlying unicast routing protocols, and robustness (if multiple cores are used).

Protocol Independent Multicast (PIM): More recently, Deering *et al.* [Deering *et al.* 94] defined a flexible architecture for multicast that they call Protocol Independent Multicast (PIM). As the name implies, the multicast architecture is independent of the protocol employed for unicast routing. Their paper shows that PIM can scale to wide-area networks, and is particularly attractive for *sparse* multicast groups. Essentially, PIM can use either the shared tree approach of CBT or the shortest-path tree approach of DVMRP, with the appropriate choice made on a per group or per host basis [Wei and Estrin 94, Wei and Estrin 95]. The PIM architecture relies upon choosing a suitable Rendezvous Point (RP), similar to a core in CBT, when constructing the multicast delivery tree for a group. The RP provides a place for multicast sources to "meet" multicast recipients.

2.3 The Challenges of Mobile Multicast

In a mobile environment, the network not only must manage multicast group membership and establish the necessary routes, but also must contend with the fact that the established routes are themselves transient in nature. The fact that the network must deal not only with dynamic group membership, but also with the dynamic location characteristics of mobile hosts, makes multicast in a mobile environment a challenging problem.

For a MH that wishes to *receive* multicast datagrams, the routing problem is slightly different from the unicast Mobile IP routing problem. This is because multicast datagrams are sent to *group* addresses that do not encode any specific network. In IPv4 and IPv6, multicast addresses are defined independent of location and use an address class separate from the normal unicast addresses. Multicast routing of datagrams to these addresses can therefore take place regardless of the MH's location. Some difficulties may be experienced, however, if the conventional IP multicast mechanism is used. The foreign network that the MH visits may not have a multicast router, a MH may experience unacceptable packet losses when resubscribing, packets will still be delivered to the previous foreign network after the MH leaves, and multicast routers may become overwhelmed by reconfiguration requests from large groups of rapidly moving hosts.

For a MH that wishes to *send* multicast datagrams while away from its home network, multicast datagrams sent by it could possibly be dropped.[4] This is because downstream multicast routers that receive the multicast datagrams on a different interface from that used to send datagrams to the MH must intentionally drop these datagrams. 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.

Research groups have begun to look at this problem. Some advocate using combinations of the existing IETF mechanisms and others advocate modifying the multicast protocols themselves. One such proposal for multicast routing to mobile hosts has been developed at Columbia University and is based upon Columbia Mobile*IP. The details can be found in [Acharya *et al.* 96].

[4] The discussion in this paragraph assumes that the source is a member of the multicast group, and that DVMRP is used as the underlying multicast protocol. The problems are less pronounced for shared-tree approaches (e.g., CBT and PIM), but can still arise depending on the exact scenario of host movement considered.

2.4 Current IETF Mobile IP Multicast Solutions

In the current version of the IETF Mobile IP specification there are two methods described for handling multicast for mobile hosts. We refer to these methods as *remote subscription* and *bi-directional tunnelled multicast*. Both of these methods address the basic multicast requirement for mobile hosts, but do not address new problems that arise when multicast services are extended to include mobile hosts (see Section 3.4). This section provides a brief introduction to the fundamental operation and capabilities of these proposed methods.

Remote Subscription: Subscription on the foreign network is the simplest option for obtaining multicast service since it has no special encapsulation requirements and operates using only existing protocols. With this option, the MH is required to resubscribe to the multicast group on each foreign network. This option will be the preferred option for some hosts depending on mobility characteristics and quality of service (QOS) requirements and the availability of multicast routers on the foreign network. If the MH is highly mobile, then packets will be lost owing to the set-up time associated with multicast subscription, and therefore this method is not preferred. If QOS requirements are crucial, or the host is likely to be stationary for an extended period of time (measured in hours or more), then this option is preferred, especially if the mobile host can unsubscribe before leaving.

Remote subscription, combined with multiple bindings, could theoretically provide respectable QOS, even for hosts that are constantly moving, but this service would come at a high price for the networks involved and the multicast routers that must manage the multicast tree. For hosts that want guaranteed two-way communication with the multicast group and are unable to acquire a co-located address, or hosts that are highly mobile, a different method is needed that will not overload the multicast routers.

Bi-directional Tunnelled Multicast: Bi-directional tunnelled multicast is another option specified in the IETF Mobile IP proposal. A bi-directional tunnelling method for unicast datagrams has recently been proposed in [Montenegro 96]. This method is designed to solve the problem of topologically incorrect source addresses in datagrams, and Montenegro presents a good discussion of the issues involved in reverse tunnelling. With bi-directional tunnelled multicast, the mobile routing agent on the home network (the home agent) must also be a multicast router. Using this option, subscriptions are done through the home agent. When the MH is away from home, a bi-directional tunnel to the home agent is set up. This allows both sending and receiving of multicast datagrams, with the same delivery guarantees given to fixed hosts. One disadvantage is that if multiple MHs on the same foreign network belong to the same multicast group then the multicast packets will be duplicated on that foreign network (see Section 3.4). This problem negates some of the advantages of using multicast in the first place.

When forwarding multicast datagrams to the MH, the HA must first encapsulate the multicast datagram in a unicast packet destined for the MH. This is done to ensure that the FA will know to whom the packet is intended since it will likely not recognize the multicast address, or recognize the MH as belonging to that multicast group. Once the packet is unicast encapsulated, then it must be encapsulated using IP-in-IP as usual. Packets travelling the reverse route from the MH to the HA are multicast packets encapsulated with a unicast header with the MH's home address

as the source address. This multiple encapsulation can dramatically increase the header overhead and induce fragmentation.

3 Our Approach to Mobile Multicast

3.1 Assumptions and Design Goals

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

- 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. This feature, a basic design tenet of the network layer in the Internet, has contributed to its robustness and is largely responsible for its success.

- Multicast support must conform to the host group model. That is, dynamic group membership is a necessary feature of multicast. Both static and mobile hosts can be members of multicast groups. Mobile hosts can join and leave multicast groups at any time, even when they are away from their home network. Both static and mobile hosts can send to multicast groups, whether or not they are members of the group.

- A mobile host that wishes to receive multicast datagrams is capable of receiving them on its home network using existing (static) multicast routing techniques. In the proposed architecture, it is assumed that a multicast router is co-resident with the HA and is denoted as *mHA*.

- The Home Agents and Foreign Agents are static (not mobile) hosts.

- There is exactly one HA and one FA per network.[5]

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.[6]

The design goals for our multicast routing mechanism 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.

[5] It turns out that this assumption is not necessary, but it does simplify the ensuing discussions.

[6] In practice, however, it is unlikely that a given mobile host will move more frequently than on a second-to-second or minute-to-minute basis.

- *Routing algorithm independence.* To the extent possible, the approach should be independent of how the underlying multicast service is provided in the internetwork. The multicast routing should, in turn, be as independent as possible from the underlying unicast routing algorithm.

- *Simplicity.* Finally, we would like the scheme to be as simple as possible, in the sense of it being able to interoperate with existing Internet protocols and mechanisms, with as few changes as possible.

We believe that our proposed approach meets all of these goals, with the exception of routing algorithm independence. That is, while our proposed approach is independent of how the multicast routing is implemented, there are distinct advantages if PIM is used. Furthermore, we rely on IETF Mobile IP (or something very similar to IETF Mobile IP) to provide unicast routing for mobile hosts.

3.2 Handling Multicast Source Mobility

One approach to solving the problem associated with a mobile *source* of multicast datagrams is to reconfigure the multicast delivery tree upon host movement, possibly using some mechanism to invalidate old routes. Such an approach seems prohibitively expensive in terms of the repeated tree set-up cost. Clearly, the reuse of the multicast tree once set up is desirable if frequent movement is to be experienced.

To ensure interoperability with DVMRP, our approach always roots the multicast tree at the home network of the MH. The sending of datagrams from the MH is implemented using a bi-directional tunnel. This method is very similar to that prescribed in the IETF proposal (see Section 2.4). When the MH wishes to send multicast datagrams, the MH operates as follows:

- if it is on its home network, the MH uses link-level multicast to send the datagram. The multicast home agent (the mHA) propagates the multicast downstream normally.

- if it is on a foreign network, the MH uses a tunnel to deliver the datagram to its mHA[7]. The mHA then propagates the multicast datagram downstream[8] via all virtual interfaces, including the physical network interface corresponding to the MH's home network.

In both cases, the source address in multicast packets is the MH's home address. Thus the recipients of the multicast are unaware that the source is mobile.

3.3 Handling Multicast Destination Mobility

Supporting mobile *destinations* of multicast datagrams requires that the network have: (1) a few access points (the mHAs) to a general multicast infrastructure, and (2) several access points (the FAs) to permit the dissemination of messages to the MHs.

[7] One drawback of this approach is that it may preclude using link-level multicast to bootstrap services that are truly local to the MH. This issue is still under investigation.

[8] Care must be taken not to create routing loops when tunnelling multicast datagrams "upstream" to an agent that is also a multicast router. A simple solution is for agents that are also multicast routers to refrain from reforwarding packets emerging from a Mobile IP tunnel.

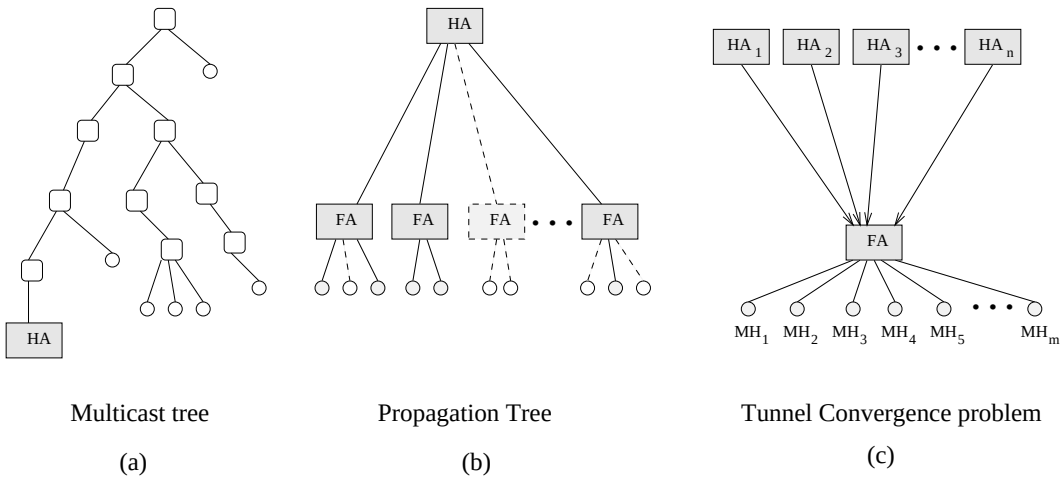

Figure 2: Multicast packet paths under the proposed scheme

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 (e.g., set up using DVMRP). 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.

The resulting scheme is simple, since the FA need not join groups on behalf of mobile hosts that are visiting its network, and mobile hosts that are members of a group are not required to resubscribe every time they move. The information needed for datagram forwarding is known through the registration procedures used in Mobile IP. Furthermore, the extra delay incurred from routing multicast packets via the mHA (probably measured in milliseconds) could be much less than the *join* and *graft* delays [Acharya *et al.* 96] incurred when rebuilding a multicast tree (possibly several seconds, depending on the size and geographic distribution of the multicast group).

Since the HA may be serving MHs at several FAs that wish to receive datagrams addressed to a given multicast address, 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 (the solid lines). Link-level multicast is used by the FA at each such foreign network to complete the delivery. This is where our approach differs from the IETF bi-directional tunnelling approach in which multicast packets are delivered as unicast packets to each MH. 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 is a potential problem with the bi-directional tunnelled multicast proposal from IETF.

3.4 Other Problems and Issues

This section discusses other problems and issues that arise when providing multicast services to mobile hosts. The problems discussed here are not addressed in the IETF proposals and will be

present in such implementations. Our proposal has been designed specifically to address these problems.

The Tunnel Convergence Problem: Our scheme is complicated by the *tunnel convergence problem* 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 (and 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 be dealt with. While duplicate datagrams do not constitute a violation of the service assumptions, they again constitute unnecessary network load.[9] 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.[10] The DMSP forwards only one multicast datagram into the tunnel even if it serves multiple MHs at the foreign network. Thus the scheme provides *at-most-once* delivery of multicast datagrams which is identical to the semantics of IP multicast.[11]

This method solves the tunnel convergence problem, but it creates a handoff problem when the last MH supported by the DMSP leaves the foreign network, and it reduces the robustness of the system by relying on one home agent for delivery. These problems can be solved by having redundant DMSPs. For example, if two DMSPs are chosen, the foreign agent can drop packets from one DMSP and deliver from the other. If one of these DMSPs stops forwarding packets, the foreign agent can rely on the other for packets, thereby greatly reducing any loss for the MHs. If a constant number of DMSPs (no more than three are recommended) are chosen, the number of duplicate packets is only O(1), which should be manageable by the foreign network.

The Duplication Problem: There is another subtle way in which duplicate multicast packets can be delivered to each group member on a network. As our approach is currently described, an FA selects an HA as a DMSP if a MH is the first *mobile* host to request subscription to group G at the foreign network. It is possible for a local *static* host at the foreign network to already be a member of group G when a mobile host (also a member of group G) arrives. Since the FA may have no knowledge of this (unless it is also a multicast router for the network), forwarding of multicast

[9] For example, suppose that MHs from several different home networks converge at the same foreign network, and are receiving an MBONE video feed that consumes 128 Kbps of bandwidth. If the video feed is multicast once, then a commercial wireless LAN with today's technology can handle the traffic. If several multicasts of the same stream were sent by the FA (i.e., one for each HA), then the foreign network would quickly be saturated.

[10] The choice of the DMSP can be based on geographic proximity, Quality of Service (QOS) requirements, or the number of mobile hosts present from each HA [Harrison *et al.* 97]. 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 [Deering *et al.* 94]. The same approach can be used by the HA to suppress link-level multicasts on the local network if no group members are currently at home.

[11] The at-most-once delivery semantics may not hold true if multiple FAs are located on the same network. Dealing with multiple FAs is discussed in [Chikarmane 95].

datagrams from the DMSP could create duplicate multicast packets on the current network.

One solution to this problem is to require that the FA be co-resident with the multicast router for the network. A natural choice for the FA would be the rendezvous point RP in PIM (e.g., first hop PIM-speaking router [Deering *et al.* 94]). If this is not possible, then the FA should at least monitor transmissions on the foreign network and if duplicate packets to the same multicast group are detected, forwarding from the DMSP should be suspended.

The Scoping Problem: A *small scope group* may be defined as a group such that multicasts addressed to that group are constrained (using the IP Time-To-Live (TTL) field) to remain within a relatively small region of an internetwork. In IPv6, group scope is controlled by a flag in the multicast address itself.

The existence of small scope groups introduces a scoping problem. This problem has two dimensions. First, small scope groups are not unique across multiple networks, and therefore a naive foreign agent may designate DMSPs to handle a small scope multicast group, ignorant of the fact that this may eliminate the ability of MHs that are not from the same network as the DMSPs to receive local scope multicasts from their own networks. Second, when a packet is received by the foreign agent addressed to one of these groups, it will be broadcast onto the foreign network and picked up by all hosts subscribing to that group. Since these groups are not unique across networks, hosts on the foreign network may receive packets that were never meant for them.

One ramification of this scoping problem is that it completely confuses the notion of "local" for *well-known groups*. For example, IP multicast (in IPv4) currently uses group id 224.0.0.1 to mean "all hosts on THIS local area network". Now consider an internetwork with mobile hosts, where two "blue" hosts leave their home (blue) local area network to attach at a foreign "green" local area network, one "red" host leaves its home network to attach to the green network, and two green hosts are mobile, one attaching to the blue network, and one to the red (see Figure 3). If an IP multicast to group 224.0.0.1 originates from a router on the blue LAN (e.g., an IGMP query), then what are the proper semantics for the multicast? Several possibilities exist:

- deliver to all hosts currently on the blue LAN (regardless of host color)
- deliver to all blue hosts (only) currently on the blue LAN (only)
- deliver to all blue hosts (regardless of their location in the internetwork)

It is not clear which semantics the network should implement.[12] While the latter semantics (all blue hosts, regardless of location) are likely the most commonly desired, these semantics can only be achieved either by doing multiple unicasts for each mobile blue host (which is inefficient if many hosts are mobile, particularly if they are at the same location) or by multicasting on all LANs where blue hosts reside (which means that non-blue hosts on that LAN also receive the multicast, since they are also members of well known group 224.0.0.1 for their own network). The same argument applies for multicasts to group 224.0.0.1 originating on the red and green LANs. Thus, in this

[12] The correct semantics to choose are likely to be application-dependent. For example, an impending file server shutdown notification, a load balancing mechanism for scheduling distributed applications, and a query to find a particular user normally resident on the LAN may all choose different semantics.

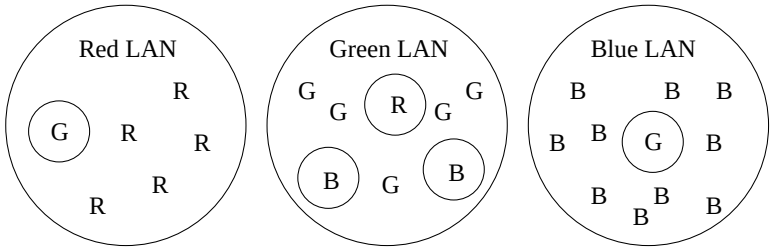

Figure 3: An illustration of the scoping problem: (circled letters represent mobile hosts on a foreign network)

example, the green hosts on the green LAN may receive multicasts for group 224.0.0.1 from red, blue, and green networks.

In April 1995, the IETF Mobile IP Working Group decided that LAN broadcasts must be forwarded to roaming MHs. The forwarding is done by encapsulating the broadcast packet within a packet from the original source to the MH. This new packet is in turn encapsulated in a packet from the HA to the care-of address and forwarded normally. Small scope multicast is similar to broadcast and it makes sense to handle it in the same way. This simple solution however reintroduces the problem of duplicating packets on the foreign network.

Another proposed fix to this scoping problem is to use network id's (i.e., network or subnet masks) when constructing group id's for well-known groups, similar to the way Ethernet multicast addresses are constructed from IP multicast addresses. With such an approach, "local" groups have globally unique identifiers. The confusion between local and global groups is then alleviated. This solution, however, has the drawback that it will make all group ID's into global identifiers which hurts scalability.

A third solution (illustrated in Figure 4) is to have the FA multicast router assign groups of MHs to local link-level multicast groups when they register and bind this link-level address to an address that can be created using the IP address of the HA itself. On an Ethernet, for instance, the addresses c0-00-00-00-80-00 through c0-00-40-00-00-00 are user-defined and should be configurable by the FA. All hosts served by the same HA should be assigned to the same link-level multicast group.

We propose using the addresses c0-00-30-00-00-00 through c0-00-3f-ff-ff-ff for mapping HA addresses onto local link layer multicast groups on the foreign network. This still leaves 2fff8000 addresses for other use, but allows 28 bits of an HA's IP address to map onto the multicast group address. If we use the first 28 bits, this would allow globally unique multicast groups for unique HAs on subnets as small as 16 hosts. The previous statements assume IPv4 addressing. If we are faced with IPv6 addressing, there is no way to guarantee uniqueness because of the large number of bits (128) in the IP address. A good hashing function, however, may allow IPv6 addresses to be mapped into 28 bits with reasonable uniqueness. As long as the hashing function was standard, the MHs could pre-compute their scoped multicast address and use it at all foreign networks.

When the HA receives datagrams destined for a local multicast group, it forwards them to the FA with a destination marked as its own subnet broadcast address. The resulting packet has an outer header with source and destination addresses (128.233.128.79, 172.16.0.1 in Figure 4), a

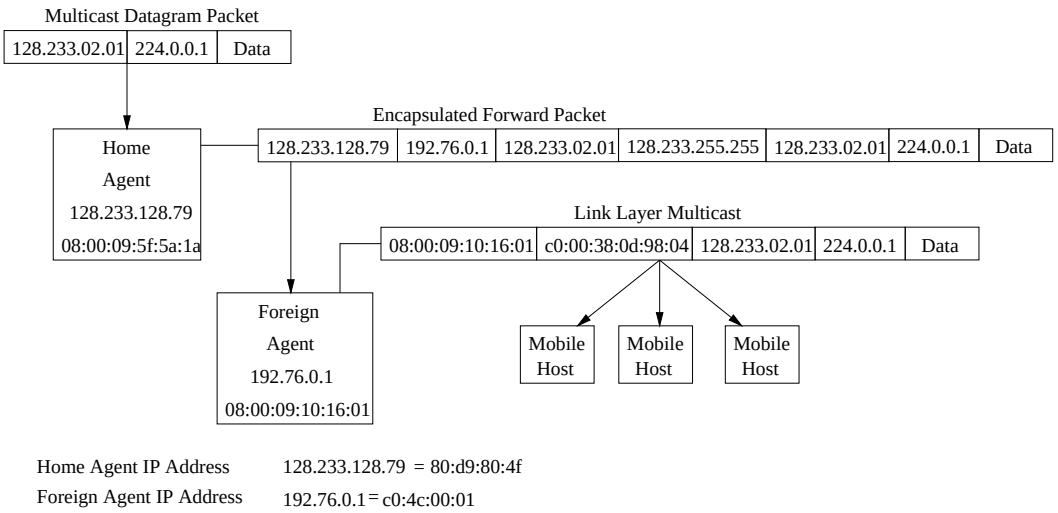

Figure 4: An example of using local link-layer multicast to solve the scoping problem

middle header with the original source and the home network broadcast address as the destination (128.233.02.01, 128.233.255.255 in Figure 4), and an inner header with the original source address and the original destination address (128.233.02.01, 224.0.0.1 in Figure 4). This address is mapped by the FA onto the local link-level multicast address for the mobile hosts (c0:00:38:0d:98:04 in Figure 4), and is broadcast onto the network. Note that in Figure 4 the final link-layer hardware address in hexadecimal format (c0:00:38:0d:98:04) contains the first 7 digits of the HA's home IP address in hexadecimal format (80.d9.80.4f). When the MHs receive the packets, they can check the destination multicast address in the inner IP header to determine if they want the packet or not. This solution does not necessitate multiple packet delivery on the foreign network, and it does allow local multicasts to be restricted to the correct MHs.

Disruptions of Multicast Service: When a MH moves, there is a possibility that it moves to a foreign network that does not have an associated multicast router. In such a case, if remote subscription is being used, multicast service (sending or receiving) may be disrupted until the host moves again to a network (home or foreign) with multicast capability.

Furthermore, 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. The reason for this temporary multicast service outage stems from the fact that in Mobile IP there is no explicit deregistration with the FA when a host moves. The MH's HA learns of the movement immediately 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 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. A discussion of this handoff is given in the next section.

Some techniques are available to minimize packet loss during handoff and relocation, and this issue is explored in Section 4.1. For hosts with a wired connection to the internetwork there is no easy method for minimizing relocation packet loss since many packets always get lost when physical network connections are torn down. However, packet loss at the foreign network, and to wireless hosts can be improved.

4 Mobile Multicast Protocol: MoM

4.1 MoM Overview

We now summarize the basic operation of our multicast scheme, for both sending and receiving multicast datagrams. For sending, all multicast transmissions by the MH are tunnelled directly to the mHA, instead of using link-level multicast on the foreign network. Any recipients on the foreign network receive the datagrams through the multicast tree rather than directly via link-level multicast by the MH. For receiving, the steps involved in providing multicast for mobile recipients are:

Service Request: A MH wishing to receive multicast datagrams informs the FA that it wishes to join a specific multicast group. Ordinarily, to join a group, an Internet host sends an IGMP *membership report* in reply to the multicast router's *membership query*. This is addressed to the multicast address corresponding to the group being joined, say G, and multicast routers listen promiscuously to all LAN multicast traffic. The IGMP software on the MH is modified to perform a check to determine whether the MH is on its home network. If the MH is not on its home network, the membership report is addressed explicitly to the FA instead of being multicast on the foreign network.

If support for small scope multicast is desired, this can be indicated by setting the broadcast flag in the registration request. In order for our link-level multicast option to be used, the foreign agent needs to negotiate a suitable proxy address (usually the HA's broadcast address) for forwarding these packets on the foreign network. This allows link-level multicast groups to be set up at the FA for MHs serviced by the same HA.

Service Processing: If the FA supports multicast service, it propagates any subscription requests to the mHA. Also, if the MH is the first mobile host on the FA's network to request service for address G, the FA also informs the mHA of the mHA's status as the DMSP for address G. The mHA adds itself to the multicast distribution tree (if it is not already on the tree), and receives all multicast datagrams addressed to G. In other words, the mHA stays "registered" in group G and is always ready to forward datagrams on behalf of the MH. However, the mHA forwards multicast datagrams only into those Mobile IP tunnels leading to foreign networks for which it is a DMSP. Negative caching [Deering *et al.* 94] is used to suppress delivery on other interfaces (physical or virtual). This requires per-tunnel state information for each group in which the mHA has mobile members. Any small-scope multicasts intercepted by the HA are forwarded to the FA with a destination address that is the HA's broadcast address. The FA should have an entry in its visitors list for this broadcast address, and map the destination onto the appropriate link-level multicast address. These entries do not time out, but are deleted when all MHs serviced by the associated HA leave the foreign network.

Service Execution: The multicast service is provided when the DMSP forwards all datagrams received corresponding to multicast address G into the Mobile IP tunnel. When a MH moves from one network to another, a notice should be given to the previous FA. The mHA receives explicit information regarding the MH's new location when the MH attempts to reregister through the new FA, and this information can be forwarded to the FA, which can act as a binding cache agent, or simply update its visitor cache. Mobile IP specifies that the FA eventually times out and discards the MH's registered status when it fails to receive a reregistration request in a specified period [Perkins 96a]. This should happen only if the MH or HA are unable to inform the FA of a location change before the time out. This situation implies that if the DMSP has only one MH at a foreign network interested in receiving multicast datagrams sent to address G, and that MH moves to another location, the tunnel to the old FA is discarded when a registration request by the MH is received through another FA. This leaves the other group members (if any) at the previous foreign network without an active DMSP for a short duration.

Redundant DMSPs provide the best mechanism for protection against packet loss. When multiple HAs offer DMSP services, the FA should pick at least two DMSPs (if minimal packet loss is desired). When the DMSP forwarding has been set up, the FA picks a primary DMSP. When packets arrive from the primary DMSP they are forwarded onto the network. When packets arrive from a secondary DMSP, they are cached. If at any time, packets are not received from the primary DMSP within a certain timeout period, yet packets are received from secondary DMSPs, then the primary DMSP is changed to the secondary DMSP which is forwarding the packets, and the cache of received packets from that secondary DMSP are forwarded onto the network. Needless to say, these caches must be small (probably fewer than 10 packets).

Service Handoff: It is helpful if the FA is informed about the inability of a mHA to continue as the DMSP. A mechanism similar to that used for large-scale distributed file system consistency is used. For the first group member at a foreign network, the corresponding mHA is the DMSP. As part of the service setup procedure, the mHA issues a *callback promise* to the FA. It retains this status until it relinquishes it when the last MH it serves at the foreign network that is also a group member moves elsewhere. At this point in time, the DMSP informs the FA that it is no longer willing to provide the multicast service for address G, a procedure called a *callback break*. This triggers the activation of a new DMSP for address G by the FA if any MHs remain that are still members of the group corresponding to address G.

A HA providing DMSP service for network N will continue forwarding the multicast packets until one of the following conditions applies:

1. The FA no longer requires DMSP service from the HA and informs it so.
2. A registration message is received from a new network for the HA's last MH at network N.
3. The registration of the HA's last MH at N times out.
4. The HA is physically unable to continue as DMSP for some reason (e.g., it is turned off).

In all but the final case, the HA informs the FA that it is no longer willing to function as DMSP.

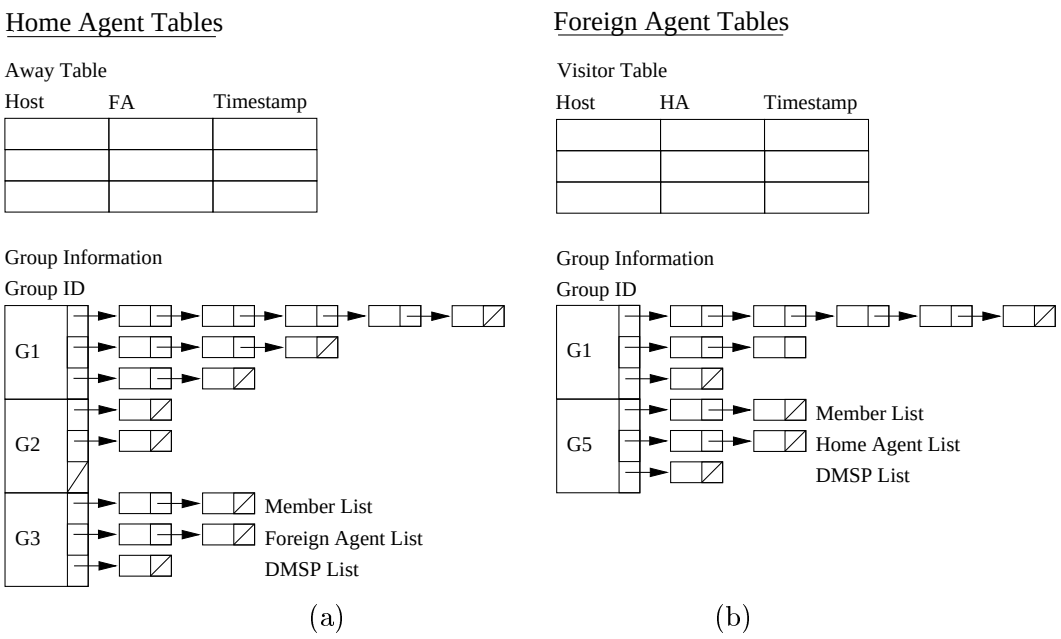

Figure 5: Data structures for the MoM protocol: (a) Data structures at the Home Agent; (b) Data structures at the Foreign Agent

This callback mechanism ensures that the break in service for MHs is minimized by using a handoff procedure. Since the next DMSP is receiving the multicast datagrams, it can simply enable the forwarding of these datagrams into the tunnel, with minimal added delay (i.e., no reconfiguration of the multicast delivery tree is required). Lost packets should be reduced to none if either the old DMSP is required to forward packets until the FA unsubscribes to the service or multiple DMSPs are defined so that the FA always has a redundant source for multicast packets. In the case of redundant DMSPs, the FA would be required to cache some packets from the redundant DMSPs to guarantee no losses during DMSP handoff.

4.2 MoM Data Structures

Figure 5 illustrates the data structures needed for the MoM protocol. Each HA 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 FA 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 5. 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 HA 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 FA's at which the away group members

reside; and a list[13] of the FA's for which the HA has DMSP responsibilities. Similarly, each FA 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 HA's to which these visiting group members belong; and the list of HA's that are currently serving as the DMSP for this group.

In terms of protocol overhead, the main issue is the relative length 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 much shorter, and the list of these foreign networks for which the HA has DMSP responsibilities is 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.

4.3 Performance

We have used several techniques to understand the performance implications of our approach to mobile muilticast. The results of a simulation study reported in [Harrison *et al.* 97] suggest significant advantages over remote subscription and bi-directional tunnelling, particularly as the number of mobile group members increases. Because of space limitations we can provide only a sample of these results here.

Scalability is an important requirement, and Figure 6 shows how a number of key factors scale with the size of the multicast group in our simulations. Although the number of mobile hosts to be serviced grows quickly, the number of DMSPs required to forward multicast messages grows quite slowly. Were every mobile host attached to a foreign network to have its multicast messages forwarded directly by its home agent, the traffic on the network would be much higher. Forwarding through DMSPs can significantly reduce network load, and thus improve performance.

More details on the design of our simulation experiments and further results from them are available in [Harrison *et al.* 97].

4.4 Discussion

In this section we address strengths and weaknesses of our approach and identify several outstanding issues. Table 1 provides a brief comparison with remote subscription and bi-directional tunnelling, over a range of issues, in a manner similar to [Xylomenos and Polyzos 97]. We make no claim to optimal routing, where the "tailored" routing of remote subscription can't be beaten, but our approach offers the advantage of being transparent to higher-level protocols and applications, as does bi-directional tunnelling. Our protocol does require modifications to both foreign agent and home agent code, but presents performance benefits from reduced packet traffic.

To the extent possible, our approach is independent of the underlying multicast routing protocol (e.g., DVMRP, MOSPF, CBT, PIM).[14] Our approach inherits the scalability, flexibility, efficiency, and robustness of the underlying multicast routing scheme, and attempts to provide scalability

[13] Clearly, this need not be implemented as a separate list. It can simply be a flag in the FA structures used. We present it as a separate list to simplify presentation of the protocol.

[14] Clearly, however, PIM offers distinct advantages for the implementation of our approach.

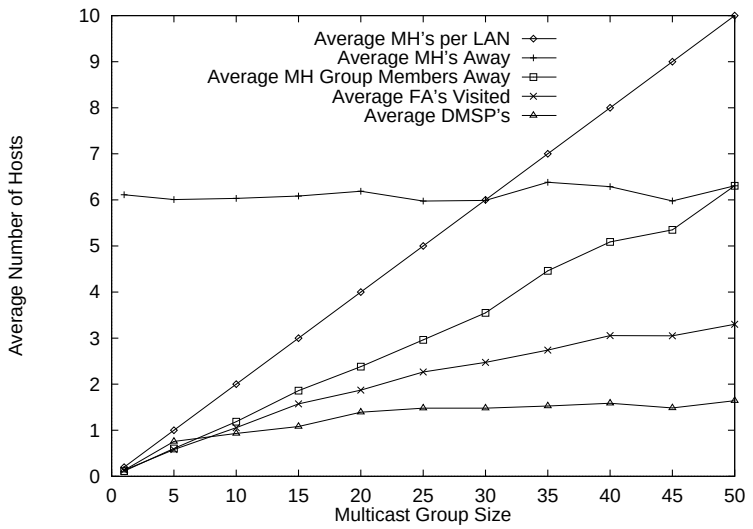

Figure 6: Scaling characteristics of the MoM protocol as a function of multicast group size

Table 1: A comparison of mobile multicast options

Category	Remote Subscription	Bi-Directional Tunnelling	MoM
Optimal routing	Yes	No	No
Transparency	No	Yes	Yes
Redundant packet delivery	No	Yes	Minimal
Delivery of scoped multicast	No	Yes	Yes
Multicast protocol independent	Yes	Yes	Yes
Join & graft delays	Yes	No	No
FA modification	No	No	Yes

with respect to the number of mobile hosts in the internetwork by making several optimizations (with regard to multicast datagram delivery) at home and foreign agents.

The main protocol-dependent part of our approach is the assumption of Mobile IP (or something *very* similar to Mobile IP) available on mobile hosts, HAs, and FAs. We see this requirement as quite crucial to the efficient support of multicast for mobile hosts.

In order to implement our link-level multicast solution to the scoping problem, an Ethernet LAN or a LAN supporting Ethernet addresses is required. We have not explored solutions for other types of networks (such as token ring).

Two limitations of our approach are:

- Packets that are sent and received by MHs must always traverse the home network, making routing non-optimal.
- Multiple unicasts are used by the HA to tunnel multicast packets to FAs of MHs that are group members.

Non-optimality is the extent to which the route taken by a packet deviates from the shortest possible path to the MH. For example, following host movement for a multicast source, it is entirely possible for a source and a recipient of a given multicast group to reside on the same physical network, while all communication between the hosts must traverse the entire internetwork, via the source's HA. This is unfortunate because of the desire for low latency delivery of datagrams in many multicast applications [Ballardie *et al.* 93, Deering *et al.* 94, Wei and Estrin 94, Wei and Estrin 95]. Non-optimality is not seen as a major drawback of our approach because neither IP nor Mobile IP guarantees optimal paths. In fact, the added complexity to optimize routing may be justified only for applications with stringent QOS requirements.

The multiple unicast problem arises if the HA is serving more than one mobile group member for a given group, and these MHs are at different foreign networks. The HA must make multiple copies of the multicast packet[15], and then, if it is serving as the DMSP for a particular foreign network, tunnel the packet to the corresponding FAs. The inefficiencies caused by multiple unicast forwarding can be solved only if a method using multicast is used to deliver the datagrams to the foreign agents involved. This solution results in multicast trees being reconfigured upon every move, however, which is part of the problem we were seeking to eliminate. A more thorough discussion of this problem can be found in [Chikarmane 95].

5 Related Work

Several other research groups have considered multicast support for mobile hosts. For example, Brown and Singh [Brown and Singh 95] specified a protocol called RelM for providing reliable 1 X N multicast service for mobile hosts, but it is focused mainly on hosts using wireless interfaces. Their protocol operates by introducing three new entities: a supervisor host (SH), a memory management module (MMM), and a local management module (LMM). These entities interact to reduce the demands on the mobile support stations (MSSs), and still provide reliable service to the MHs.

The DATAMAN research group at Rutgers University was one of the first to propose a multicast protocol for mobile hosts [Acharya and Badrinath 93]. Their initial scheme, designed to support *exactly-once* multicast delivery, assumes static multicast groups (i.e., the membership of a group does not change during the group's lifetime) and sender knowledge of the group membership, and thus does not extend easily to IP multicast and the host group model. More recent work [Acharya *et al.* 96] proposed extensions to IP multicast to support mobile hosts. This approach is based on Columbia University Mobile*IP, and uses mobile support routers (MSRs, which are similar to but not the same as the agents in IETF Mobile IP) to provide multicast datagram delivery to mobile group members. The implementation uses DVMRP, although the authors claim that the approach will work with other multicast routing algorithms.

Our work differs from the DATAMAN work in three important ways. First, our approach uses IETF Mobile IP instead of Columbia University Mobile*IP, which changes several of the multicast issues. In Columbia University Mobile*IP, for example, all mobile hosts in a campus area are assigned addresses from the *same* subnet, which is used *exclusively* for mobile hosts. The multicast problem then reduces to creating the illusion of link-level multicast delivery to this subnet,

[15] A similar approach is used by the multicast routers on the MBONE to tunnel multicast packets to multicast-capable islands in the Internet.

regardless of where the mobile hosts reside. Second, their approach is primarily for handling host movement within a campus network (i.e., one home network), while ours is for internetwork-wide movement. While their approach can be extended to handle inter-campus movement, separate mechanisms are required to do so (e.g., tunnelling packets from a multicast source back to the home network, as proposed in [Chikarmane *et al.* 95a]). It is not clear how well their approach would scale to a large internetwork, since it may require static multicast tunnels with world-wide scope (i.e., to all agents, which may be anywhere in the internetwork). Finally, we believe that our approach is more scalable in the number of mobile hosts. For example, their approach requires that *all* MSRs in a campus area join the multicast delivery tree for group G even if there is only a *single* mobile host that is a member of G somewhere in the network. Furthermore, link-level multicasts for the group G are performed by *all* MSRs, even if they have no local members of G. In our opinion, this resembles broadcast to mobile hosts, rather than multicast. Our approach is much more conscious of the bandwidth used for the delivery of multicast packets to mobile hosts.

6 Summary and Conclusions

This paper has described a new approach to providing multicast to mobile hosts in a TCP/IP internetwork. The proposed multicast scheme is based on an expansion of the virtual network established for IP Multicast by using Mobile IP tunnels to constitute the "last mile" for delivery of multicast datagrams. The scheme has several features that make it practical as a solution for mobile hosts on TCP/IP internetworks. In particular, it provides the ability to support dynamic groups and provide minimal break in service as a result of host movement.

We believe that such an approach is possible and offers important advantages. The main advantage is that minimal changes are required to IP multicast and Mobile IP. Our approach requires minimal modifications to either protocol: primarily the addition of several mechanisms to handle issues that arise only when multicast and host mobility are both present in an internetwork.

Our preliminary simulation studies have shown promise in the method. More complete performance studies are needed, however, and experiments with a prototype implementation are planned.

In conclusion, we believe it is possible to leverage existing protocols to provide multicast services to mobile hosts in an efficient and effective manner. Although our approach has limitations, including non-optimal packet routing and the possibility of multiple unicasts being used by the Home Agents to propagate multicast datagrams to more than one Foreign Network, it scales well to permit widespread use over geographically dispersed internetworks. In closing, we hope that the issues identified in this paper (e.g., tunnel convergence, support for small scope groups) are addressed in future versions of the Mobile IP and IP multicast protocol specifications.

Acknowledgements

Financial support for this research was provided by the Natural Sciences and Engineering Research Council (NSERC) of Canada, through research grants OGP0003707 and OGP0121969, by Telecommunications Research Laboratories (TR*Labs*), and by IBM Canada. The authors thank the anonymous reviewers for their constructive comments on an earlier version of this paper.

References

[Acharya and Badrinath 93] 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.

[Acharya *et al.* 96] A. Acharya, A. Bakre, and B. Badrinath, "IP Multicast Extensions for Mobile Internetworking", *Proceedings of IEEE INFOCOM'96*, San Francisco, CA, March 1996.

[Ballardie *et al.* 93] 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.

[Brown and Singh 95] K. Brown, S. Singh, "RelM: Reliable Multicast for Mobile Networks", Technical Report, Department of Computer Science, University of South Carolina, September 1995.

[Cheshire and Baker 96] S. Cheshire, M. Baker, "Internet Mobility 4x4", *Proceedings of the 1996 ACM SIGCOMM Conference*, Palo Alto, CA, pp. 318-329, August 1996.

[Chikarmane 95] V. Chikarmane, *Network Support for Mobile Hosts in a TCP/IP Internetwork*, M.Sc. Thesis, Department of Computer Science, University of Saskatchewan, August 1995.

[Chikarmane *et al.* 95a] 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.

[Chikarmane *et al.* 95b] 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, Dec 1995.

[Deering 91] S. Deering, *Multicast Routing in a Datagram Internetwork*, Ph.D. Thesis, Department of Computer Science, Stanford University, 1991.

[Deering and Cheriton 90] 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.

[Deering *et al.* 94] 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.

[Droms 93] R. Droms, "Dynamic Host Configuration Protocol", *RFC1541*, Network Working Group, November 1993.

[Eriksson 94] H. Eriksson, "MBONE: The Multicast Backbone", *Communications of the ACM*, Vol. 37, No. 8, pp. 54-60, August 1994.

[Forman and Zahorjan 94] G. Forman and J. Zahorjan, "The Challenges of Mobile Computing", *IEEE Computer*, Vol. 27, No. 4, pp. 38-47, April 1994.

[Harrison *et al.* 97] T. Harrison, C. Williamson, R. Bunt and W. Mackrell, "Mobile Multicast (MoM) Protocol: Multicast Support for Mobile Hosts", Research Report, Department of Computer Science, University of Saskatchewan, May 1997.

[Imielinski and Badrinath 94] T. Imielinski and B. Badrinath, "Mobile Wireless Computing", *Communications of the ACM*, Vol. 37, No. 10, pp. 18-28, October 1994.

[Johnson 96] D. Johnson, and C. Perkins, "Route Optimization in Mobile IP", *Internet Draft*, Mobile IP Working Group, November 1996.

[Montenegro 96] G. Montenegro, "Bi-directional Tunneling for Mobile IP", *Internet Draft*, Mobile IP Working Group, Sun Microsystems Inc., September 1996.

[Malkin 94] G. Malkin (editor), *RIP Version 2 Carrying Additional Information*, RFC 1723, Xylogics, Inc., November 1994.

[Moy 94a] J. Moy, "Multicast Routing Extensions for OSPF", *Communications of the ACM*, Vol. 37, No. 8, pp. 61-66, August 1994.

[Myles and Skellern 93] A. Myles and D. Skellern, "Comparison of Mobile Host Protocols for IP", *Internetworking Research and Experience*, Vol. 4, No. 4, pp. 175-194, December 1993.

[Perkins 96a] C. Perkins, *IP Mobility Support*, RFC 2002, Mobile IP Working Group, October 1996.

[Perkins 96b] C. Perkins, *IP Encapsulation within IP*, RFC 2003, Mobile IP Working Group, October 1996.

[Waitzman *et al.* 88] D. Waitzman, C. Partridge, and S. Deering (editors), *Distance Vector Multicast Routing Protocol*, RFC 1075, BBN STC and Stanford University, November 1988.

[Wei and Estrin 94] L. Wei and D. Estrin, "The Trade-offs of Multicast Trees and Algorithms", *Proceedings of the 1994 International Conference on Computer Communications and Networks*, San Francisco, CA, September 1994.

[Wei and Estrin 95] L. Wei and D. Estrin, "Multicast Routing in Dense and Sparse Modes: Simulation Study of Tradeoffs and Dynamics", *Proceedings of the 1995 International Conference on Computer Communications and Networks*, Las Vegas, NV, pp. 150-157, September 1995.

[Xylomenos and Polyzos 97] G. Xylomenos and G. Polyzos, "IP Multicast for Mobile Hosts", *IEEE Communications*, pp. 54-58, January 1997.