

Multicast support for mobile hosts using Mobile IP: Design issues and proposed architecture

Vineet Chikarmane*, Carey L. Williamson, Richard B. Bunt and Wayne L. Mackrell

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

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.

1. Introduction

This paper considers the problem of multicast to groups with dynamic membership 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 IETF Mobile IP protocol [25] 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 much interest in problems related to the mobile computing paradigm [14,17].

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 [8,9]. Examples of multicast applications currently used on the Internet include resource discovery, as well as desktop audio/video conferencing and shared whiteboard applications on the Internet multicast backbone (MBONE) [12].

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 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 [27], MOSPF [23], CBT [3], and PIM [10]) implicitly assume *stationary* hosts when configuring the multicast delivery tree. The delivery trees established for static multicast cannot be changed easily or efficiently in all cases due to the propagation of these trees to many routers, and the potentially large cost associated with making changes to the trees' structure. In addition, 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,

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.

* Now at IBM Canada, Toronto, Ontario, Canada.

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 [25] defines a method for routing packets to mobile hosts. It also defines two multicast support options, which we will call *remote subscription* and *bi-directional tunneled multicast*. We describe a third option, similar to bi-directional tunneled 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 unicast² IP routing for mobile hosts in a TCP/IP internetwork. In Mobile IP, a mobile host is assigned an IP address on its home network, called the mobile host's *home address*. Packets from a correspondent host to the mobile host are always addressed to the home address of the mobile host. If the correspondent host is aware that the mobile host is mobile, then the correspondent host can encapsulate its packets and forward them directly to the mobile host's new location without traversing the home network [19].

In general, when the mobile host connects to a foreign network, it identifies and registers with a *foreign agent*, or registers directly with its *home agent*. When registering, the mobile host acquires a care-of address defining its current location. The combination of the mobile host's home address and the care-of address is known as a *binding*. The mobile host can acquire its care-of address either from a foreign agent or through autoconfiguration methods (such as DHCP [11]) designed for assigning temporary IP ad-

resses. If the care-of address is assigned by a foreign agent, it will be the address of the foreign agent itself. If the address is assigned via autoconfiguration, it is said to be *co-located* with the mobile host, and will be some available address on the foreign network. When the care-of address for the mobile host is not co-located, the foreign agent is responsible for decapsulating and forwarding packets to the mobile host on the foreign network using link-layer protocols. The foreign agent must also serve as the mobile host's default router. If the address is co-located, the mobile host itself must decapsulate datagrams. An example scenario is shown in figure 1.

Foreign agents responsible for managing mobile host registrations maintain a list of all currently registered mobile hosts that are visiting the network served by the foreign agent. A lifetime is negotiated during the registration process for each visiting mobile host and is stored in the visitor list entry for that mobile host. A mobile host must reregister before this period has expired to ensure uninterrupted service from the foreign agent. A timestamp originating from the mobile host and updated at each reregistration attempt is associated with every binding. This timestamp can be used to determine timeouts for registrations. If a mobile host has not reregistered within the timeout period, the foreign agent assumes that the mobile host has moved to another network.

A home agent on the mobile host's home network serves as a binding cache and forwarding agent for the mobile host. Extensions for optimized routing allow these forwarding capabilities to also exist at the correspondent hosts themselves. Lifetimes are associated with all bindings, and when a binding expires, the home agent or correspondent host assumes the mobile host has returned to its home network. During the lifetime of a binding, the home agent is responsible for intercepting packets addressed to the mobile host and forwarding them to the care-of address. These packets are forwarded using IP-in-IP tunneling, which encapsulates each IP packet to be forwarded inside another IP packet using a forwarding header [26]. Facilities are required at the care-of address for decapsulating the packets when they arrive and forwarding the packets to the mobile host.

2.2. Internet IP multicast

WAN multicast [9] has been incorporated into the TCP/IP suite using the host group model [8]. 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. Hosts use the *Internet Group*

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

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.

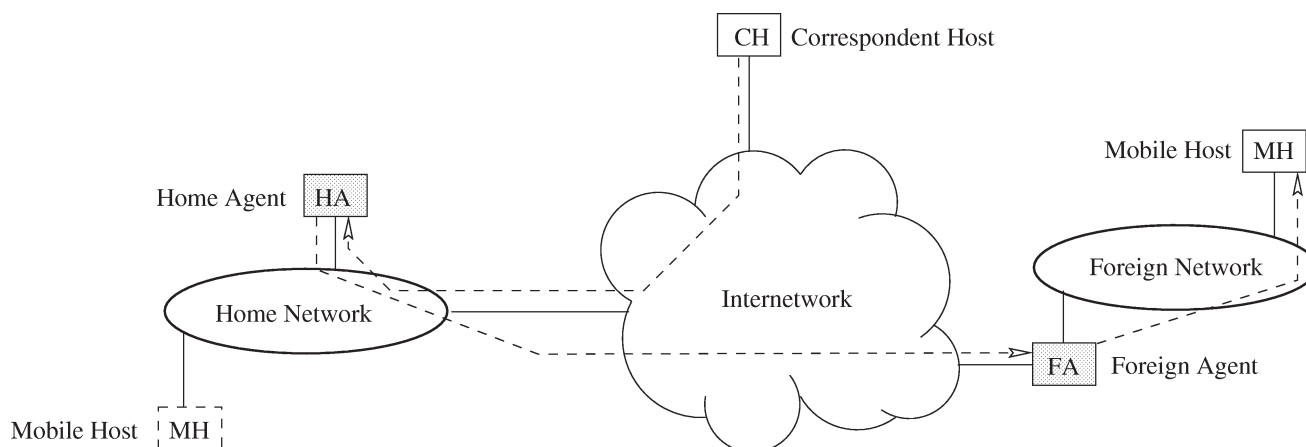


Figure 1. Mobile IP entities and their inter-relationships.

Management Protocol, IGMP, to inform multicast routers about the multicast groups that the hosts wish to join. 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.

The host group model used for incorporating multicast into TCP/IP has the following important properties:

- 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 the location or identity of any group member (i.e., only the group address needs to be known);
- 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 [27], MOSPF [23], Core Based Trees (CBT) [3], and Protocol Independent Multicast (PIM) [10]. 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 [23]. 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 [27] is a distance-vector multicast routing protocol that has been implemented as an extension

to RIP [21,22]. DVMRP has been widely implemented, and is deployed in hundreds of regions connected by the MBONE [12]. DVMRP uses the *truncated reverse-path broadcast* [9] 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 or not they have multicast group members. 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 $S \times N$, where S is the number of multicast sources, and N is the number of multicast groups. Ballardie et al. [3] 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 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. [10] 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 [28,29]. 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 locations of mobile hosts, makes multicast in a mobile environment a challenging problem.

For a mobile host 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 mobile host’s location. Some difficulties may be experienced, however, if the conventional IP multicast mechanism is used. The foreign network that the mobile host visits may not have a multicast router, a mobile host may experience unacceptable packet losses when resubscribing, packets will still be delivered to the previous foreign network after the mobile host leaves, and multicast routers may become overwhelmed by reconfiguration requests from many rapidly moving hosts.

For a mobile host that wishes to *send* multicast datagrams while away from its home network, multicast datagrams sent by it could possibly be dropped. This is because downstream multicast routers that receive the multicast datagrams on a different interface from that used to send datagrams to the mobile host must intentionally drop these datagrams. Furthermore, downstream routers continue to track reverse paths to the mobile host assuming that the mobile host is on its home network. Thus the multicast routes that are established are always with reference to the mobile host’s home network and are incorrect when the mobile host is at a foreign network. These problems appear when the source is a member of the multicast group and 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.

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 [18]. The details can be found in [2].

2.4. Current IETF Mobile IP multicast solutions

The current IETF Mobile IP specification proposes *remote subscription* and *bi-directional tunneled multicast* for handling multicast for mobile hosts. Both methods provide basic multicast support for mobile hosts, but do not address new problems that arise when multicast services are extended to mobile hosts (see section 3.4). This section reviews the basic operation of these two 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 mobile host is required to resubscribe to the multicast group on each foreign network, and must use a co-located care-of address. 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 as it moves. If the mobile host is highly mobile, however, packets will be lost owing to the set-up time associated with multicast subscription, and therefore in this case 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 does provide the most efficient delivery of multicast datagrams, but this service may 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 tunneled multicast: Bi-directional tunneled multicast is another option specified in the IETF Mobile IP standard. Bi-directional tunneling for unicast datagrams is discussed in [6,20]. This method is designed to solve the problem of topologically incorrect source addresses in datagrams by requiring traffic from the mobile host to be routed back to the home network through a foreign agent to home agent tunnel.

With bi-directional tunneled 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 mobile host 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 mobile hosts on the same foreign network belong to the same mul-

ticast group then duplicate copies of the multicast packets will arrive at 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 mobile host, the home agent must first encapsulate the multicast datagram in a unicast packet destined for the mobile host. This is done to ensure that the foreign agent will know to whom the packet is intended since it will likely not recognize the multicast address, or recognize the mobile host as belonging to that multicast group. Once the packet is unicast encapsulated, then it must be encapsulated again and addressed to the care-of address. Packets traveling the reverse route from the mobile host to the home agent are multicast packets encapsulated with a unicast header with the mobile host's home address as the source address. This multiple encapsulation increases the packet size substantially and can cause 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 home agent, and we refer to this as a multicast home agent.
- The home agents and foreign agents are static (not mobile) hosts.
- There is exactly one foreign agent per network visited. This is assumed in order to simplify some of the ensuing discussions. The issue of multiple foreign agents at the foreign network is dealt with in [5].

We make no 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.⁴

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 in the Internet). 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 small or nonexistent.
- *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. If the mobile source is expected to move frequently, maximal reuse of the multicast tree is desirable to reduce overhead on the multicast routers.

To ensure interoperability with DVMRP, our approach always roots the multicast tree at the home network of the mobile host. The sending of datagrams from the mobile host 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 mobile host wishes to send multicast datagrams, the mobile host operates as follows:

- If it is on its home network, the mobile host uses link-level multicast to send the datagram. The home agent propagates the multicast downstream normally.

⁴ 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.

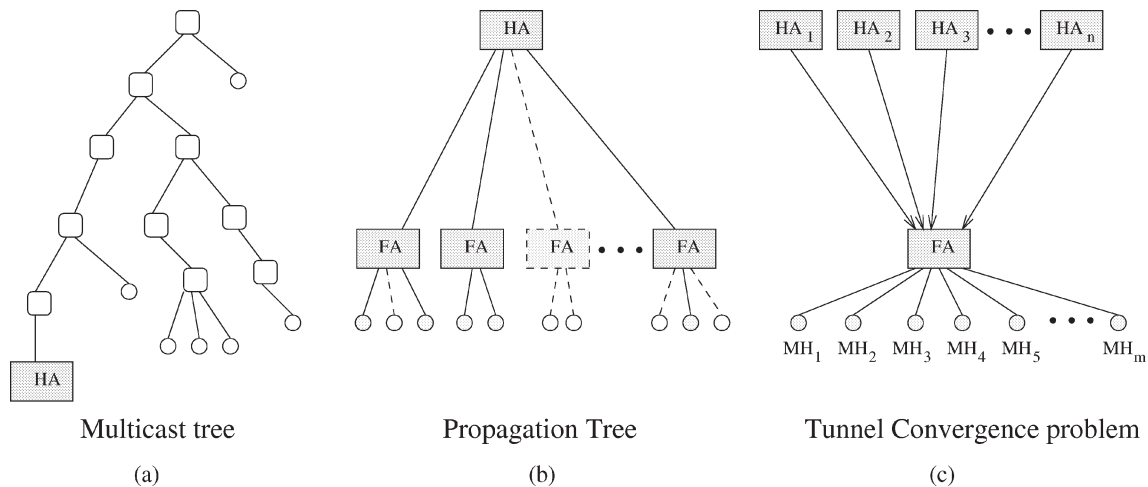


Figure 2. Multicast packet paths under the proposed scheme.

- If it is on a foreign network, the mobile host uses a tunnel to deliver the datagram to its home agent.⁵ The multicast home agent then propagates the multicast datagram downstream via all interfaces, including the physical network interface corresponding to the mobile host's home network, as well as virtual interfaces for tunnels to other MHs. Care must be taken, however, not to create routing loops when tunneling multicast datagrams “up-stream” to an agent that is also a multicast router. For this reason, the time to live values on forwarded datagrams should restrict the delivery of those datagrams to the foreign network only.

In both cases, the source address in the multicast packets is the mobile host's home address.

3.3. Handling multicast destination mobility

Supporting mobile *destinations* of multicast datagrams requires that an internetwork have: (1) home access points (the multicast home agents) to a general multicast infrastructure, and (2) foreign access points (the foreign agents) to permit the dissemination of messages to the mobile hosts.

Since home agents are static and know the location of the mobile hosts that they serve, it is logically feasible for a multicast home agent to forward multicast traffic to the mobile host through the Mobile IP tunnel via the foreign agent. As shown in figure 2(a), the home agent 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 home agent is one of many group members.

The resulting scheme is simple, since the foreign agent 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.

⁵ One drawback of this approach is that it may preclude using link-level multicast to bootstrap services that are truly local to the mobile host. This issue is still under investigation.

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 multicast home agent (typically measured in milliseconds) could be much less than the *join* and *graft* delays [2] incurred when rebuilding a multicast tree (often several seconds, depending on the size and geographic distribution of the multicast group).

Since the home agent may be serving mobile hosts at several foreign agents 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 home agent 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 foreign agent at each such foreign network to complete the delivery. This is where our approach differs from the IETF bi-directional tunneling approach in which multicast packets are delivered as unicast packets to each mobile host. 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.

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 home agents) can terminate at a particular foreign agent. This problem is illustrated in figure 2(c), where multiple home agents (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 same foreign agent. Thus one copy of every multicast packet would be forwarded to the foreign agent by each home agent that is serving interested mobile hosts. Since the foreign agent would locally deliver every multicast datagram forwarded to it, the problem of duplicate multicast packet delivery to the mobile hosts must be dealt with. While duplicate datagrams do not constitute a violation of the service assumptions, they again constitute unnecessary network load.⁶ 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, the foreign agent selects one home agent as the *designated multicast service provider, DMSP*, for a given multicast group.⁷ The DMSP forwards only one multicast datagram into the tunnel even if it serves multiple mobile hosts at the foreign network. Thus the scheme provides *at-most-once* delivery of multicast datagrams to all subscribing mobile hosts. These semantics may not hold if multiple foreign agents are located on the same network. Dealing with multiple foreign agents is discussed in [5].

This method solves the tunnel convergence problem, but it creates a handoff problem when the last mobile host 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 mobile hosts. If a constant number of DMSPs (a small number is recommended due to packet duplication) is 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, a foreign agent selects a home agent as a DMSP if a mobile host 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 foreign agent has no knowledge of this (unless it is also a multicast router for the network), forwarding of multicast datagrams from the DMSP could create duplicate multicast packets on the current network.

One solution to this problem is to require that the foreign agent be co-resident with the multicast router for the network. A natural choice for the foreign agent would be the rendezvous point RP in PIM (e.g., first hop PIM-speaking router [10]). If this is not possible, or if multiple FAs exist at the foreign network, then the foreign agent should at least monitor transmissions on the foreign network.

If there is only a single foreign agent, forwarding from the DMSP should stop when duplicate packets from another source are detected. If multiple foreign agents are present, duplication is inevitable unless the foreign agents explicitly exchange state information.

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 networks; thus 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 mobile hosts that are not from the same network as the DMSPs to receive local scope multicasts from their own networks. Second, when a packet addressed to one of these groups is received by the foreign agent, 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 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).

⁶ For example, suppose that mobile hosts 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 foreign agent (i.e., one for each home agent), then the foreign network would quickly be saturated.

⁷ 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 home agent [15]. Home agents 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]. The same approach can be used by the home agent to suppress link-level multicasts on the local network if no group members are currently at home.




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

It is not clear which semantics the network should implement. 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. 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 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 a roaming mobile host may request LAN broadcasts on its home network to be forwarded to it. The forwarding is done by encapsulating the broadcast packet within a packet from the original source to the mobile host. This new packet is in turn encapsulated in a packet from the home agent to the care-of address, and is forwarded normally. Small scope multicast is similar to broadcast, and thus 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 identifiers (i.e., network or subnet masks) when constructing group identifiers 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 identifiers into global identifiers, which hurts scalability.

A third solution (illustrated in figure 4) is to have the foreign agent multicast router assign groups of mobile hosts 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 home agent 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 foreign agent. All hosts served by the same home agent 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 home agent addresses onto local link layer multicast groups on the foreign network. This still leaves 2^{28} addresses for other use, but allows 28 bits of a home agent’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 home agents on subnets as small as 16 hosts. The previous statements assume IPv4 addressing. In IPv6 addressing, however, 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 mobile hosts could pre-compute their scoped multicast address and use it at all foreign networks.

When the home agent receives datagrams destined for a local multicast group, it forwards them to the foreign agent with a destination group, it forwards them to the foreign agent 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 middle header with the home agent’s address of the source and the home network broadcast address as the destination (128.233.128.79, 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 foreign agent 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 home agent’s home IP address in hexadecimal format (80.d9.80.4f). When the mobile hosts 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 mobile hosts.

Disruptions of multicast service: When a mobile host 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 foreign agent when a host moves. The mobile host’s home agent learns of the movement im-

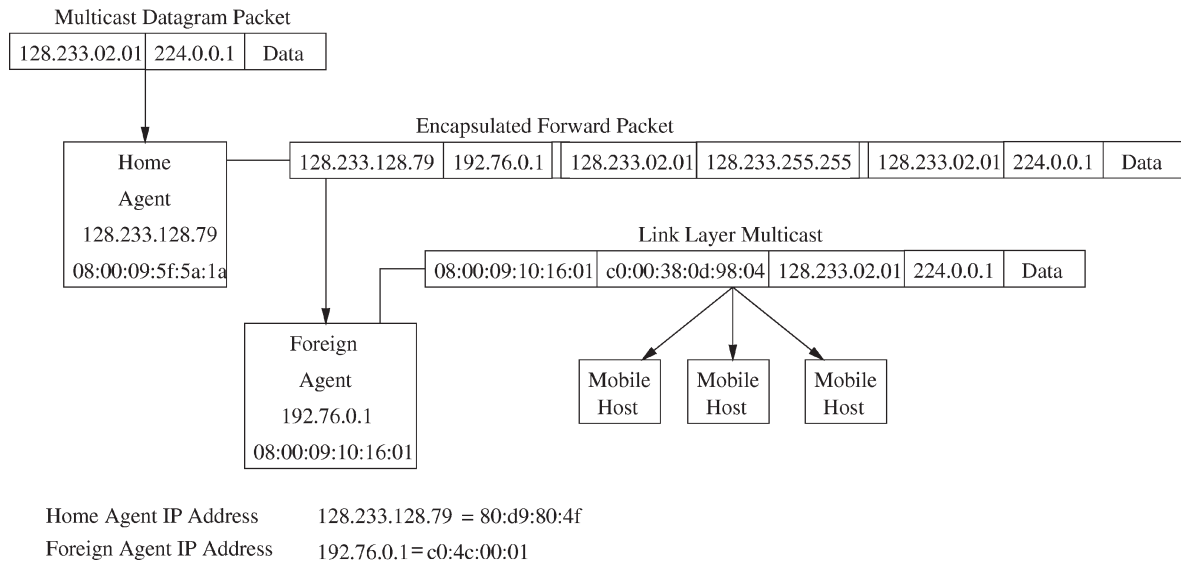


Figure 4. An example of using local link-layer multicast to solve the scoping problem.

mediately when the mobile host reregisters at the new network, but the foreign agent at the old foreign network learns about the movement only through a timeout. In the case that the moving host’s home agent is the DMSP for a group at the (previous) foreign network, a DMSP handoff will be required to a different home agent, 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. To avoid service disruption, the DMSP must continue forwarding to the foreign agent until the handoff is completed. 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.

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 mobile host are tunneled directly to the multicast home agent, 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 mobile host. For receiving, the steps involved in providing multicast for mobile recipients are described below.

Service request: A mobile host wishing to receive multicast datagrams informs the foreign agent 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 mobile host is modified to perform a check to determine whether the mobile host is on its home network. If the mobile host is not on its home network, the membership report is addressed explicitly to the foreign agent 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 (or possibly by defining a new registration extension). This forwarding address for the scoped datagrams would be the local broadcast address on the mobile host’s home network.

Service processing: If the foreign agent supports multicast service, it propagates any subscription requests from visiting MHs to their respective home agent. Also, if the mobile host is the first mobile host on the foreign agent’s network to request service for address G , the foreign agent also informs the home agent of that home agent’s status as the DMSP for address G . The home agent 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 home agent stays “registered” in group G and is always ready to forward datagrams on behalf of the mobile host. However, the home agent forwards multicast datagrams only into those Mobile IP tunnels leading to foreign networks for which it is a DMSP.

Negative caching [10] is used to suppress delivery on other interfaces (physical or virtual). This requires per-tunnel state information for each group in which the multicast home agent has mobile members. Any small-scope multicasts intercepted by the home agent are forwarded to the foreign agent with a destination address that is the home agent’s broadcast address. The foreign agent 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 mobile hosts serviced by the associated home agent 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 mobile host moves from one network to another, a notice should be given to the previous foreign agent. The multicast home agent receives explicit information regarding the mobile host's new location when the mobile host attempts to reregister through the new foreign agent. This information can be forwarded to the previous foreign agent, which can either start forwarding packets to the new foreign agent until the visitor cache entry times out, or simply update its visitor cache to reflect that the mobile host is no longer registered.

Mobile IP specifies that the foreign agent eventually times out and discards the mobile host's registered status when it fails to receive a reregistration request in a specified period [25]. This should happen only if the mobile host or home agent are unable to inform the foreign agent of a location change before the time out. This situation implies that if the DMSP has only one mobile host at a foreign network interested in receiving multicast datagrams sent to address G , and that mobile host moves to another location, the tunnel to the old foreign agent is discarded when a registration request by the mobile host is received through another foreign agent. This leaves the other group members (if any) at the previous foreign network without an active DMSP for a short duration. To avoid this, we require the DMSP to continue forwarding packets to the foreign agent until the handoff has completed. This requirement is mandatory for all DMSPs, but since the home agent for one mobile host should not be trusted by mobile hosts from other networks, redundant DMSPs may be needed to ensure smooth handoffs.

Redundant DMSPs provide the best mechanism for protection against packet loss when handoffs occur. When multiple home agents offer DMSP services, the foreign agent should pick at least two DMSPs (if minimal packet loss is desired). When the DMSP forwarding has been set up, the foreign agent 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 secondary DMSP becomes the primary DMSP, and the cache of received packets are forwarded onto the network. If there are other home agents to choose from, then a new secondary DMSP is chosen. Otherwise, the old primary becomes the secondary. These caches must be small (probably fewer than 10 packets) to conserve resources at the foreign agent, and to minimize the delay introduced through storing the packets in the cache.

Service handoff: It is helpful if the foreign agent is informed about the inability of a multicast home agent 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 multicast home agent is the DMSP. As part of the service set-up procedure, the foreign agent registers a *callback* with the DMSP. When the last mobile host served by the DMSP at the foreign network moves elsewhere, the DMSP performs the callback to notify the foreign agent that it is no longer willing to provide the multicast service for address G . This triggers the activation of a new DMSP for address G by the foreign agent if any mobile hosts remain that are still members of group G .

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

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

In all but the final case, the home agent informs the foreign agent that it is no longer willing to function as the DMSP.

This callback mechanism ensures that the break in service for mobile hosts 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). There should not be any lost packets if either the old DMSP is required to forward packets until the foreign agent unsubscribes to the service or multiple DMSPs are defined so that the foreign agent always has a redundant source for multicast packets. In the case of redundant DMSPs, the foreign agent should be required to cache 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 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 care-of address), 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 home agent), and when these bindings expire.

The away list and visitor list are already required by the Mobile IP standard, but 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

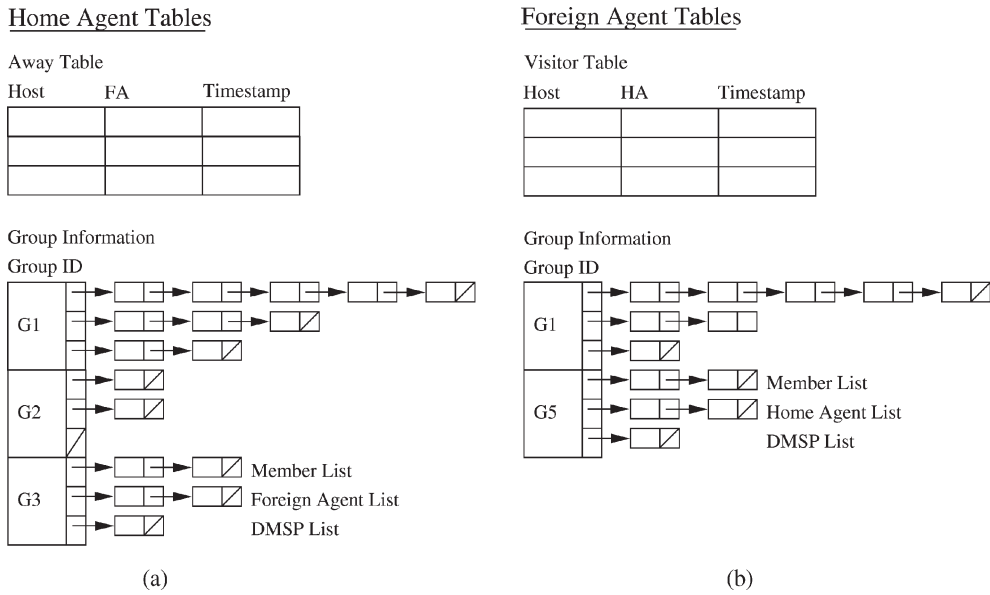


Figure 5. Data structures for the MoM protocol: (a) Data structures at the Home Agent; (b) Data structures at the Foreign Agent.

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 care-of addresses at which the away group members reside; and a list⁸ of the care-of addresses for which the home agent 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 home agent’s to which these visiting group members belong; and the list of home agent’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 home agent, for example. The number of multicast group members who are away may grow large. However, if the group members visit some of the same foreign networks, then the list of the foreign networks visited will be shorter than the list of the group members. Also, assuming that there are mobile hosts from other networks residing at the foreign networks, the list of networks for which the home agent has DMSP responsibilities should be even shorter. 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 multicast. The results of a simulation study reported in [15] and [16] suggest significant advantages over remote subscription and

⁸ This need not be implemented as a separate list. It can simply be a flag in the foreign agent structures used. We present it as a separate list to simplify presentation of the protocol.

bi-directional tunneling, particularly as the number of mobile group members increases. Because of space limitations we can provide only a sample of these results here.

In our simulations we made the following assumptions: there was only one fixed source for multicast datagrams, there was static membership in the multicast groups, 100% of the receiving nodes were mobile, local multicast routing was not available, and multicast datagram arrivals were Poisson. The simulation was run with 5 LANs, 10 mobile hosts per LAN, and one multicast group. In our simulations, we only studied delivery to mobile hosts, and issues with having mobile hosts as the source of datagrams were not studied.

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 our simulation experiments and the results from them are available in [16].

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 tunneling, over a range of issues, in a manner similar to [30]. We make no claim to optimal routing; the location specific routing of remote subscription cannot be improved upon. Our approach offers the ad-




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 tunneling	MoM
Optimal routing	Yes	No	No
Transparency	No	Yes	Yes
No Redundant packet delivery	Yes	No	Minimal
Delivery of scoped multicast	No	Yes	Yes
Multicast protocol independent	Yes	Yes	Yes
Join & graft delays	Yes	No	No
Foreign agent modification	No	No	Yes

vantage of being transparent to higher-level protocols and applications, as does bi-directional tunneling. Our protocol requires modifications to both foreign agent and home agent code, but presents performance benefits from reduced packet traffic.

Our approach inherits the scalability, flexibility, efficiency, and robustness of the underlying multicast routing scheme, and attempts to provide scalability 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 that Mobile IP (or something *very* similar to Mobile IP) is available on mobile hosts, home agents, and foreign agents. We see this requirement as crucial to the efficient support of mobile multicast.

In order to implement our solution to the scoping problem a LAN supporting link-level multicast is required. Our

design has only considered Ethernet networks, but it should work on other networks that support multicast.

Two limitations of our approach are:

- Packets that are sent and received by mobile hosts must always traverse the home network, making routing non-optimal.
- Multiple unicasts are used by the home agent to tunnel multicast packets to foreign agents of mobile hosts 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 mobile host. 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 home agent (and possibly through a home agent

for the recipient as well). This is unfortunate because of the desire for low latency delivery of datagrams in many multicast applications [3,10,28,29]. 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 home agent is serving more than one mobile group member for a given group, and these mobile hosts are at different foreign networks. The home agent must make multiple copies of the multicast packet,⁹ and then, if it is serving as the DMSP for a particular foreign network, tunnel the packet to the corresponding foreign agents. 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 [5].

5. Related work

Several other research groups have considered multicast support for mobile hosts. For example, Brown and Singh [4] specified a protocol called RelM for providing reliable $1 \times N$ multicast service for mobile hosts, but it is focused mainly on hosts using wireless interfaces. Their protocol is based on the Columbia Mobile*IP protocol and operates by introducing three new entities: a supervisor host, a memory management module, and a local management module. These entities interact to reduce the demands on the mobile support stations, and still provide reliable service to the mobile hosts.

The DATAMAN research group at Rutgers University was one of the first to propose a multicast protocol for mobile hosts [1]. 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 [2] proposed extensions to IP multicast to support mobile hosts. This approach is also based on the Columbia University Mobile*IP protocol, and uses mobile support routers (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, 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., tunneling packets from a multicast source back to the home network, as proposed in [6]). 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* mobile support routers 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* mobile support routers, 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 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 simulation studies [15,16] 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

⁹ A similar approach is used by the multicast routers on the MBONE to tunnel multicast packets to multicast-capable islands in the Internet.

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 (TRLabs), and by IBM Canada. The authors thank the anonymous reviewers for their constructive comments on an earlier version of this paper, and Charles Perkins and David Johnson for their detailed feedback.

References


- [1] A. Acharya and B. Badrinath, Delivering multicast messages in networks with mobile hosts, in: *Proc. 13th Int. Conf. Distributed Computing Systems*, Pittsburgh, PA (May 1993) pp. 292–299.
- [2] A. Acharya, A. Bakre and B. Badrinath, IP multicast extensions for mobile internetworking, in: *Proc. IEEE INFOCOM'96*, San Francisco, CA (March 1996).
- [3] A. Ballardie, P. Francis and J. Crowcroft, Core based trees (CBT): an architecture for scalable inter-domain multicast routing, in: *Proc. 1993 ACM SIGCOMM Conference*, San Francisco, CA (September 1993) pp. 85–95.
- [4] K. Brown and S. Singh, RelM: reliable multicast for mobile networks, Technical Report, Department of Computer Science, University of South Carolina (September 1995).
- [5] V. Chikarmane, Network support for mobile hosts in a TCP/IP internetwork, M.Sc. Thesis, Department of Computer Science, University of Saskatchewan (August 1995).
- [6] V. Chikarmane, R. Bunt and C. Williamson, Mobile IP-based multicast as a service for mobile hosts, in: *Proc. 2nd Int. Workshop on Services in Distributed and Networked Environments*, Whistler, BC, Canada (June 1995) pp. 11–18.
- [7] V. Chikarmane, C. Williamson and R. Bunt, Implementing Mobile IP routing in 4.3 BSD Unix: a case study, *Internetworking: Research and Experience* 6(4) (December 1995) 209–227.
- [8] S. Deering, Multicast routing in a datagram internetwork, Ph.D. Thesis, Department of Computer Science, Stanford University (1991).
- [9] S. Deering and D. Cheriton, Multicast routing in datagram internetworks and extended LANs, *ACM Trans. Computer Systems* 8(2) (May 1990) 85–110.
- [10] S. Deering, D. Estrin, D. Farinacci and V. Jacobson, An architecture for wide-area multicast routing, in: *Proc. 1994 ACM SIGCOMM Conference*, London, UK (August 1994) pp. 126–135.
- [11] R. Droms, Dynamic host configuration protocol, RFC 1541, Network Working Group (November 1993).
- [12] H. Eriksson, MBONE: The multicast backbone, *Commun. ACM* 37(8) (August 1994) 54–60.
- [13] W. Fenner, Internet group management protocol, version 2, RFC 2236, Network Working Group (November 1997).
- [14] G. Forman and J. Zahorjan, The challenges of mobile computing, *IEEE Computer* 27(4) (April 1994) 38–47.
- [15] T. Harrison, C. Williamson, W. Mackrell and R. Bunt, Mobile multicast (MoM) protocol: multicast support for mobile hosts, in: *Proc. 3rd Annual ACM/IEEE Conf. Mobile Computing and Networking (MobiCom'97)*, Budapest, Hungary (September 1997) pp. 151–160.
- [16] C. Williamson, T. Harrison, W. Mackrell and R. Bunt, Performance evaluation of the MoM mobile multicast protocol, *Mobile Networks and Applications* 3(2) (1998) 189–201.
- [17] T. Imielinski and B. Badrinath, Mobile wireless computing, *Commun. ACM* 37(10) (October 1994) 18–28.
- [18] J. Ioannidis, Protocols for mobile internetworking, Ph.D. Thesis, Columbia University (1993).
- [19] D. Johnson and C. Perkins, Route optimization in Mobile IP, Internet Draft, Mobile IP Working Group (November 1996) (work in progress).
- [20] G. Montenegro, Bi-directional tunneling for Mobile IP, Internet Draft, Mobile IP Working Group, Sun Microsystems Inc. (September 1996) (work in progress).
- [21] G. Malkin, ed., Routing information protocol, RFC 1058, Rutgers University (June 1988).
- [22] G. Malkin, ed., RIP version 2 carrying additional information, RFC 1723, Xylogics, Inc. (November 1994).
- [23] J. Moy, Multicast routing extensions for OSPF, *Commun. ACM* 37(8) (August 1994) 61–66.
- [24] A. Myles and D. Skellern, Comparison of mobile host protocols for IP, *Internetworking: Research and Experience* 4(4) (December 1993) 175–194.
- [25] C. Perkins, IP mobility support, RFC 2002, Mobile IP Working Group (October 1996).
- [26] C. Perkins, IP encapsulation within IP, RFC 2003, Mobile IP Working Group (October 1996).
- [27] D. Waitzman, C. Partridge and S. Deering, eds., Distance vector multicast routing protocol, RFC 1075, BBN STC and Stanford University (November 1988).
- [28] L. Wei and D. Estrin, The trade-offs of multicast trees and algorithms, in: *Proc. 1994 Int. Conf. Computer Communications and Networks*, San Francisco, CA (September 1994).
- [29] L. Wei and D. Estrin, Multicast routing in dense and sparse modes: simulation study of tradeoffs and dynamics, in: *Proc. 1995 Int. Conf. Computer Communications and Networks*, Las Vegas, NV (September 1995) pp. 150–157.
- [30] G. Xylomenos and G. Polyzos, IP multicast for mobile hosts, *IEEE Communications* (January 1997) 54–58.



Vineet Chikarmane is currently a Business Development Manager with IBM Software in Toronto, developing revenue-driving partnerships with California-based e-business software ventures. While involved with research and development, he spent 2 years in Vancouver as a Network Specialist with Glenayre Technologies, Inc., developing next-generation TCP-IP based 2-way wireless voice and data messaging network infrastructure. Vineet graduated with a Masters degree in computer science from the University of Saskatchewan, Saskatoon, in 1995. E-mail: vineet@vnet.ibm.com



Carey L. Williamson received a B.Sc. (Honours) in computer science from the University of Saskatchewan in 1985, and his Ph.D. in computer science from Stanford University in 1992. He is currently a Professor in the Department of Computer Science at the University of Saskatchewan, in Saskatoon, Canada. His research interests include congestion control, network traffic measurement and modeling, mobile computing, Internet Web servers, high speed networking, and ATM. E-mail: carey@cs.usask.ca



Richard Bunt has been a faculty member in the Department of Computer Science at the University of Saskatchewan since 1972 and holds the rank of Professor. He has a bachelors degree from Queen's University, in Kingston, Ontario, Canada, and masters and doctoral degrees from the University of Toronto. His research interests are in performance issues relating to distributed computing. E-mail: bunt@cs.usask.ca



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