Providing Multicast for Mobile Hosts in a TCP/IP Internetwork

Vineet Chikarmane* Carey L. Williamson Richard B. Bunt
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}@cs.usask.ca

April 19, 1996

Abstract

As today's users become increasingly nomadic, support for mobility, or location-independent computing, is rapidly emerging as a fundamental requirement of network design. In addition to the basic support required for mobile operation, such as location management and data management, mobile users expect support for the same services that they receive when they are at their home location. Multicast is one such feature that network users will likely require, even when they are mobile.

In this paper, we consider the problem of multicast to dynamic groups in a TCP/IP internetwork with mobile hosts. The fact that the network must deal not only with dynamic group membership but also with the dynamic location characteristics of mobile nodes makes multicast in a mobile environment a challenging problem.

The paper describes an architecture to support IP multicast for mobile hosts in a TCP/IP internetwork. The basic unicast routing capability of IETF Mobile IP is used to serve as the foundation for the design of a multicast service facility for mobile hosts. The scheme presented leverages existing IP multicast for static hosts to provide multicast services for mobile hosts as well. We believe that the scheme is 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.

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

1 Introduction

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 and wireless networks a reality [Imielinski and Badrinath 94]. This convergence of a need for mobility and products to satisfy that need has created an explosion of interest in problems related to the mobile computing paradigm [Forman and Zahorjan 94].

Concurrent with this growth in the popularity of mobile computing, there 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 initially emerged in local area network (LAN) distributed system environments as a convenient means to provide *query* and *notification* services (e.g., locating a particular server from a well-known *server group*, or updating state information at multiple (replicated) sites), but has since expanded to the wide area network (WAN) environment through the use of *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 (swb) and desktop video/audio conferencing (vic, vat) on the Internet multicast backbone (MBONE) [Eriksson 94].

While few portable computers will have either the processing capacity or the communications bandwidth to run all of these multicast applications, there are nonetheless several mobile computing services that could be effectively provided over an underlying multicast facility. Among these are the dissemination of information such as weather reports, travel information, and stock market reports. In addition, applications like mail, multi-person communication, service location and distributed systems functions like cache consistency could also make use of a general purpose multicast facility.

In other words, multicast is a valuable mechanism that network users will likely require, even when they are mobile. Providing an efficient and effective multicast facility for mobile users is thus an important step.

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.[1] 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.[2] Finally, the

[1] An IETF Working Group has defined a protocol, Mobile IP [Perkins 95], to provide (unicast) IP routing for mobile hosts. Mobile IP is described in more detail in Section 3.2. To date there are very few public-domain implementations of Mobile IP (see [Chikarmane *et al.* 96] for one example), and very little experience with Mobile IP in the larger Internet. The existing Mobile IP protocol specification currently has no support for IP multicast, although the Mobile IP working group continues to address this issue.

[2] For example, if the source of a multicast should change location, then the *incoming interface check* conducted in

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.

This paper describes an architecture to provide multicast support for mobile hosts in a TCP/IP internetwork. The approach is attractive because it ties together Mobile IP and IP multicast, while requiring minimal changes to either protocol. While this approach seems obvious at first glance, combining multicast *and* mobility introduces several subtle issues and complications, which are discussed in more detail in the paper itself.

The resulting scheme is simple, flexible, robust, and scalable with the number of mobile nodes in an internetwork. Furthermore, the scheme can interoperate with DVMRP, MOSPF, CBT, or PIM in the current Internet.

The remainder of this paper is organized as follows. Sections 2 and 3 provide some background on IP multicast, Mobile IP, and the challenges introduced for multicast routing by host mobility. Section 4 outlines our proposed mechanism to support multicast for mobile hosts. Section 5 discusses the strengths and weaknesses of the approach, identifying its limitations and some outstanding issues. Finally, Section 6 concludes the paper.

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* 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 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 three of these algorithms here.

most multicast routing algorithms could result in discarding datagrams, rather than propagating them downstream to the multicast recipients.

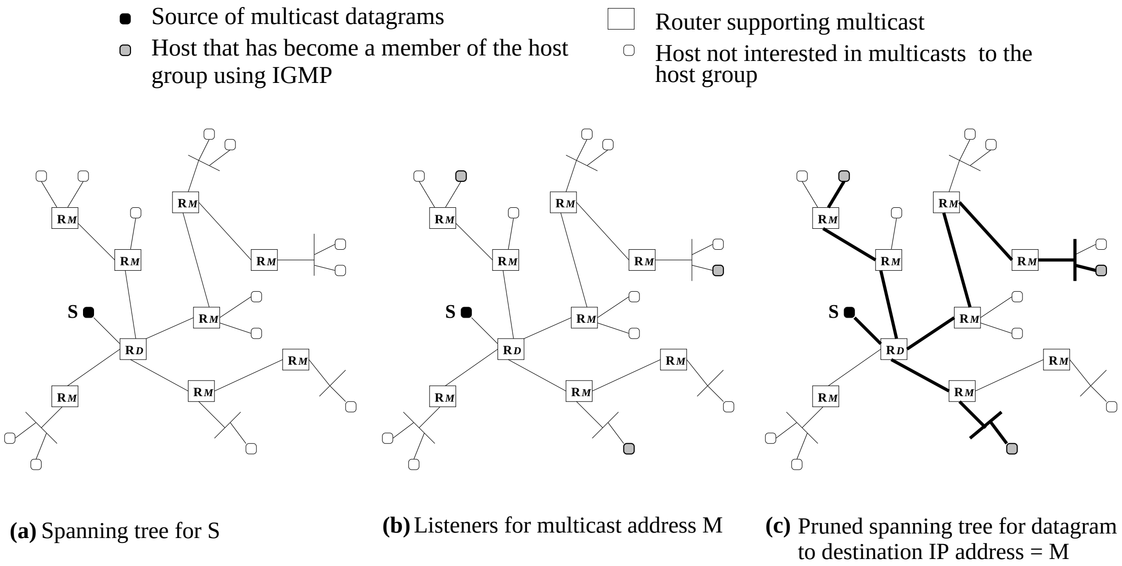

Figure 1: Multicast in a TCP/IP internetwork

2.1 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 (see Figure 1(a)), regardless of whether they have multicast group members (see Figure 1(b)) 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, as shown in Figure 1(c). 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.

Clearly, the cost of tree set-up is significant in terms of overall bandwidth consumed by datagrams sent by the sending host. It is thus not feasible to reconstruct the multicast delivery tree every time a mobile host (either the multicast source or a multicast recipient) relocates.

2.2 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. Furthermore, routers that are not on the multicast delivery tree must still incur overhead (processing plus state information) for each multicast group in the network.

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. Furthermore, routers not on the multicast delivery tree incur no overhead in tree construction and maintenance for that multicast group.

The CBT approach thus 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).

CBT still incurs significant overhead at the time of multicast tree construction, and thus implicitly assumes static hosts. For a multicast source that is *not* a member of the group, source movement causes no problem, since datagrams addressed to the group are unicast to the core for multicast delivery along the tree. For a multicast source that *is* a member of the group, however, source movement may require reconfiguration of the multicast tree. This makes CBT's adaptation to the mobile environment unlikely.

2.3 Protocol Independent Multicast (PIM)

More recently, Deering *et al.* [Deering *et al.* 94] define 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.

Even in PIM, however, there is an implicit assumption of stationary hosts, since the overhead of constructing the multicast delivery tree precludes rebuilding it each time a mobile host moves. Once again, this is not a feasible approach to mobile multicast.

3 The Challenges of Mobility

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

We illustrate several of these issues with a brief discussion of unicast and multicast routing when mobile hosts are present in an internetwork. Throughout the paper, we use the following IETF terminology to identify the principal network entities that are of interest: a *mobile node, MN*, is a node that may move through an IP internetwork; a *correspondent host, CH*, is any arbitrary host communicating with the MN; the network identifier in the network component in the MN's (permanent) IP address is called the *home network* of the MN; any other network that the MN may visit is referred to as a *foreign network*.

3.1 Routing Problems

Host movement creates problems for both unicast and multicast routing. In unicast routing, for example, when the MN is moved from one network to another, its IP address no longer correctly encodes the network to which the MN is attached. When a *remote* correspondent host, CH, attempts to send a datagram to the MN, the datagram is routed through the internetwork until it finally reaches a router that attaches directly to the MN's home network. If that router still has an ARP cache entry for the MN, then the datagram is encapsulated in a link-level frame addressed to the MAC address of the MN and transmitted. Since the MN is no longer there to receive the frame, the datagram is effectively lost. If no ARP cache entry exists at the router, then the router transmits an ARP request on the home network, asking for the MN. Since the MN is no longer there, no reply is received, and the router discards the datagram. The same result happens even if a *local* correspondent host attempts to send a datagram to the MN [Chikarmane 95].

Host movement also affects multicast routing. For a MN that wishes to *receive* multicast datagrams, the routing problem is slightly different from the unicast routing problem for mobile hosts. This is because these datagrams have a class D IP multicast address specified as their destination address. These addresses do not encode any specific network, and multicast routing of these datagrams can take place if the MN is not at its home network. However, the following difficulties may be experienced if the conventional IP multicast mechanism were used: the foreign network that the MN visits may not have a multicast router; and, even if the MN is able to inform a multicast router about the multicast addresses that it wishes to receive, when the MN moves *again*, the datagrams will not be delivered at the new network since the multicast routers have no knowledge of this movement.

For a MN that wishes to *send* multicast datagrams while away from its home network, multicast datagrams sent by it could possibly be dropped.[3] This is because downstream multicast routers that receive the multicast datagrams on a different interface from that used to send datagrams to the MN must intentionally drop these datagrams (Section 2.1). Furthermore, downstream routers continue to track reverse paths to the MN assuming that the MN is on its home network. Thus the multicast routes that are established are always with reference to the MN's home network and are incorrect when the MN is at a foreign network.

Several proposals have been made for network layer mobile host protocols to allow transparent interoperation across a TCP/IP internetwork. The proposals differ in the manner in which location information is represented and propagated, and the details of the mechanisms used to forward packets to mobile hosts. Myles and Skellern [Myles and Skellern 93] evaluated several proposals [Ioannidis *et al.* 91, Perkins 95] based on compatibility with existing infrastructure, route optimality, migration procedures and other considerations. Subsequent work [Myles and Perkins 93, Johnson 94, Johnson *et al.* 94] sought to address the problems of the earlier proposals in areas such as robustness, scalability and security. An IETF Working Group has been chartered to coordinate

[3]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.

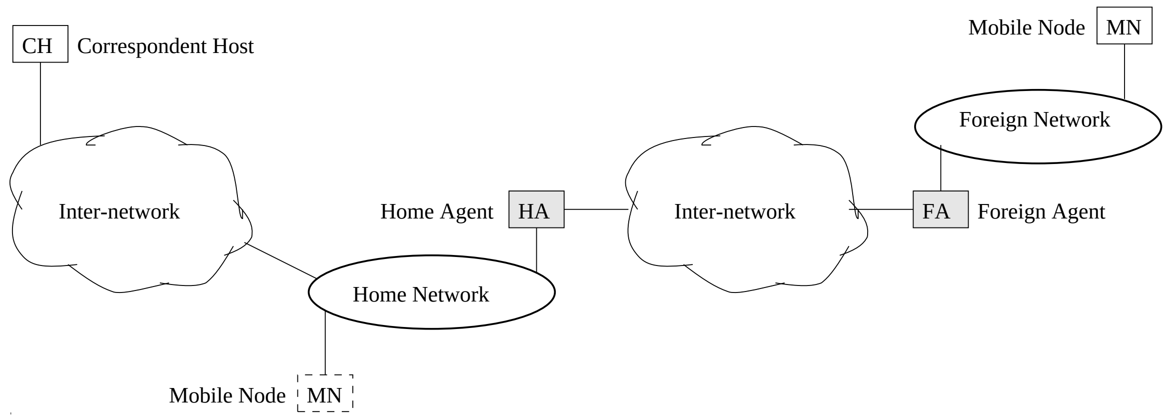

Figure 2: Mobile IP entities and their inter-relationships

efforts towards the specification of a protocol, Mobile IP [Perkins 95], to handle host mobility in the Internet. In the next subsection, we consider the relevant aspects of the (draft) Mobile IP protocol.

3.2 IETF Mobile IP

The IETF Mobile IP Working Group has defined a protocol called Mobile IP to support (unicast) IP routing for mobile hosts in a TCP/IP internetwork. In Mobile IP, a mobile node MN is assigned an IP address on its home network, called the MN's *home address*. Packets from the CH to the MN are always addressed to the home address of the MN.

When the MN connects to a foreign network, it identifies and registers with a *foreign agent, FA*, which also serves as its default router. An example scenario is illustrated in Figure 2. During the registration process, the FA provides the MN with a *care-of address*, defining the current location of the MN. The combination of the MN's home address and the care-of address is known as a *binding*. A timestamp originating from the MN and updated at each reregistration attempt is associated with every binding.

Each FA maintains a list of all currently registered MNs that are visiting the network served by the FA. A lifetime is associated with each visitor list entry and is negotiated during the registration process. A MN must reregister before this period has expired to ensure uninterrupted service from the FA. The FA establishes a permanent IP to link-level address mapping in its ARP cache by recording the MAC address of the MN from the registration request packet. A *home agent*, *HA*, attached to the MN's home network maintains a *server cache* which identifies all MNs that it is configured to serve. The HA maintains a current binding for all the MNs in its server cache and hence when a MN registers with a new FA, it must also register with its HA.

Bindings and lifetimes are associated with all MNs registered with a HA. In case a MN does not reregister upon expiry of the lifetime, the home agent assumes that the MN is connected to its home network. If the MN is at a foreign network, the HA must arrange to intercept packets addressed to the MN and *tunnel*[4] the packet, by *encapsulating* and sending it to the FA that currently has

[4] Note that unlike IP multicast tunnels, Mobile IP tunnels are dynamically set up and torn down, and are not read from a configuration file.

the MN registered. When a FA receives a tunneled packet addressed to a MN on its visitor list, the FA *decapsulates* and delivers the packet locally to the MN.

4 Our Approach to Mobile Multicast

This paper considers the problem of multicast to dynamic groups in a TCP/IP internetwork with mobile hosts. Adapting a static host multicast mechanism such as that provided by IGMP and DVMRP in the TCP/IP suite to support mobile nodes as well would seem non-trivial. However, a deployed, proven multicast mechanism that already exists for static hosts on the Internet is potentially of great utility if it can be leveraged to provide multicast services for mobile hosts as well.

An intuitive scheme for providing multicast for mobile nodes is proposed, which 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.

4.1 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 they are members of the group or not.

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

Note that we make no *a priori* assumptions about the size of multicast groups, the geographic distribution of the multicast group members, the number of mobile nodes in the network, the location of the mobile nodes, or the frequency of mobile node movement.[6]

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

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

The design goals for our multicast routing mechanism include:

- *Scalability.* The approach should work well even when the number of mobile nodes 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.

- *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.

In the remaining subsections, we discuss the operation of our proposed approach for providing multicast service for mobile hosts, and the issues that arise in the support of mobile multicast.

4.2 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.

To ensure interoperability with DVMRP, the proposed approach always roots the multicast tree at the home network of the MN. When the MN wishes to send multicast datagrams, the MN operates as follows:

- if it is on its home network, it 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, it 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 MN's home network.

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

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

[8] Note that care must be taken not to create routing loops when tunneling 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.

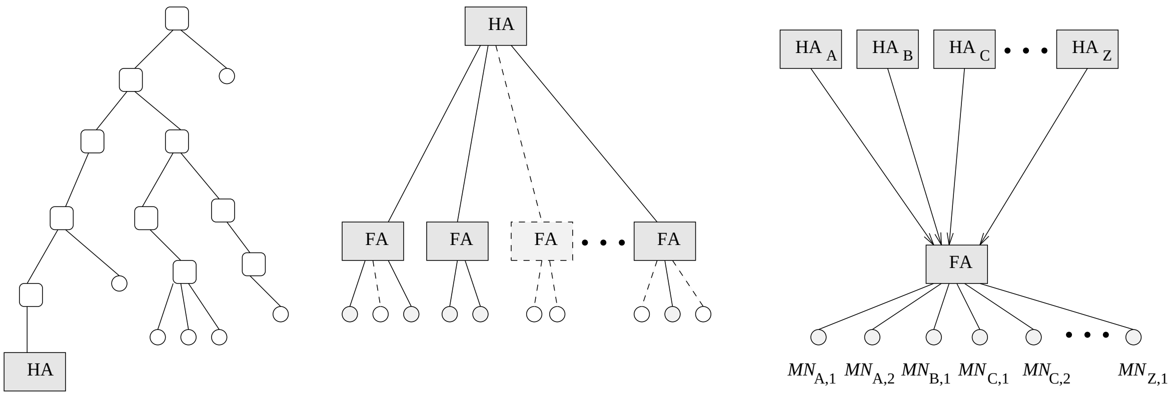

(a) Multicast tree (b) Propagation tree (c) Tunnel convergence problem

Figure 3: Multicast packet paths under proposed scheme

4.3 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 MNs.

Since HAs are static and know the location of the MNs that they serve, it is logically feasible for a multicast HA to forward multicast traffic to the MN through the Mobile IP tunnel via the FA. As shown in Figure 3(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 forwarding occurs transparently to the MN and the FA. The FA need not join groups on behalf of mobile nodes that are visiting its network, and mobile nodes that are members of a group G are not subject to *join* and *graft* delays every time they move [Acharya *et al.* 96]. The information needed for datagram forwarding is known through the registration procedures used in Mobile IP. Furthermore, the extra delay incurred due to routing multicast packets via the mHA (probably measured in milliseconds) is much less than the delay incurred to rebuild a multicast tree (probably several seconds, depending on the size and geographic distribution of the multicast group).

Since the HA may be serving MNs at several FAs that wish to receive datagrams addressed to multicast address G, it forwards a copy into each corresponding Mobile IP tunnel, as shown in Figure 3(b). Note that the HA need not forward a separate copy for each mobile node that it serves, but only one copy for each foreign network at which its mobile node group members reside (the solid lines). Link-level multicast is used by the FA at each such foreign network to complete the delivery.

The approach as described here avoids the unnecessary duplication of multicast packets on the foreign network in the event that the HA has multiple MNs present there. 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.

4.4 Other Problems and Issues

4.4.1 The Tunnel Convergence Problem

The foregoing 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 3(c), where multiple HAs (at different parts of the internetwork) all happen to have mobile nodes (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 MNs. Since the FA would locally deliver every multicast datagram forwarded to it, the problem of duplicate multicast packet delivery to the MNs 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 nodes 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 MNs 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]

4.4.2 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 FA selects a HA as a DMSP if a MN is the first *mobile* node to request subscription to group G at the foreign network. It is possible for a local *static* node at the foreign network to already be a member of group G when a mobile node (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 datagrams from the DMSP could create duplicate multicast packets on the current network.

The 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]).

---

[9] For example, suppose that MNs 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 nodes present from each HA. 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].

4.4.3 The Scoping Problem

The DMSP optimization suffers from a scoping problem for *small scope groups*: not all group members are at the same "distance" anymore, and multicast packets that are transmitted on Home Networks other than the one the DMSP is on may not make it to intended mobile nodes at a foreign network.

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. However, the precise nature of the support that needs to be provided for mobile hosts is still very much in the process of being defined. For example, it was only recently, in April 1995, that the IETF Mobile IP Working Group decided that LAN broadcasts must be forwarded to roaming MNs. Since broadcasts are to be tunneled, it seems logical that network support also be provided for small scope multicast groups, such as that used in preference to the IP broadcast address.

One ramification of this scoping problem is that it completely confuses the notions of "local" for *well-known groups*. For example, IP multicast 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. 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 colour)
- 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 semantics, these semantics can only be achieved by either 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 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.

Our 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.

[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.

4.4.4 Disruptions of Multicast Service

When a mobile node moves, there is a possibility that it moves to a foreign network that does not have an associated multicast router. In such a case, 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 node 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 MN's Home Agent learns of the movement immediately when the MN 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 at the foreign network. Until this handoff completes, multicast delivery for group members at the foreign network may be disrupted.

This problem can be alleviated if mobile nodes, by courtesy, notify FAs of their impending departure, in order to expedite DMSP handoffs. We recommend adding such a mechanism to Mobile IP, even if it is not explicitly required for the correct operation of Mobile IP.

4.5 Summary

We now summarize the basic operation of our multicast scheme, for both sending and receiving multicast datagrams.

For sending, all multicast transmissions by the MN are tunneled 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 MN.

For receiving, the steps involved in providing multicast for mobile recipients are:

Service Request: A MN 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 G, corresponding to the group being joined, and multicast routers listen promiscuously to all LAN multicast traffic. The IGMP software on the MN is modified to perform a check to determine whether the MN is on its home network. If the MN is not on its home network, the membership report is addressed explicitly to the FA instead of being multicast on the foreign network.

Service Processing: If the FA supports multicast service, it propagates the request to the mHA. Also, if the MN is the first mobile node 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 MN. However, the mHA forwards multicast datagrams only into those Mobile IP tunnels leading to foreign networks for which it is the 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.

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 MN moves from one network to another, no notice is given to the previous FA, whereas the mHA receives explicit information regarding the MN's new location when the MN attempts to reregister through the new FA. Rather, Mobile IP [Perkins 95] specifies that the FA eventually times out and discards the MN's registered status when it fails to receive a reregistration request in a specified period. This implies that if the DMSP has only one MN at a foreign network interested in receiving multicast datagrams sent to address G, and that MN moves to another location, the tunnel to the old FA is discarded when a registration request by the MN is received through another FA. This leaves the other group members (if any) at the previous foreign network without a DMSP.

Service Handoff: It is imperative that the FA be 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 MN it serves at the foreign network that is also a group member moves elsewhere. At this point of 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 selection of a new DMSP for address G by the FA if any MNs remain that are still members of the group corresponding to address G. This callback mechanism ensures that the break in service for MNs is minimized due to the use of 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).

5 Discussion and Related Work

This section discusses the strengths and weaknesses of our proposed approach, identifies several outstanding issues, and discusses related research on protocol support for mobile multicast.

5.1 Advantages

The main advantage of our approach is its ability to "glue together" Mobile IP (which does not yet explicitly consider multicast support) and Internet IP multicast (which does not explicitly consider mobile host support) to provide multicast support for mobile hosts. 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. The approach extends the applicability of the host group model, providing support for dynamic multicast groups with mobile hosts as multicast sources and/or recipients.

To the extent possible, our approach is independent of the underlying multicast routing protocol (e.g., DVMRP, MOSPF, CBT, PIM).[13] 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 nodes in the internetwork by making several optimizations (with regard to multicast datagram delivery) at home and foreign agents.

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

5.2 Limitations

The limitations of our approach are:

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

We briefly discuss each of these issues here.

5.2.1 Non-Optimal Routing

Non-optimality may be defined as the degree to which the route taken by a packet deviates from the shortest possible path (or at least from the path IP would have chosen if the sender knew the location of the MN). 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 phenomenon 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. The added complexity of trying to optimize routing may only be justified for applications with stringent Quality of Service (QOS) requirements.

5.2.2 Possible Use of Multiple Unicasts

The multiple unicast problem arises if the HA is serving more than one mobile group member for a given group, and these MNs are at different foreign networks. The HA must therefore make multiple copies of the multicast packet[14], and then, depending on whether it is serving as the DMSP for a particular foreign network, tunnel the packet to the corresponding FAs.

Clearly, making multiple copies and transmissions per incoming packet has a high associated cost, one that reduces the number of hosts that a HA can serve, and, as a consequence, the scalability

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

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

of the system. In the case of multicast and broadcast packets, a large number of copies of whole packets and multiple transmissions of these copies may be needed, and the resulting impact on performance would be substantial.[15]

One possible solution is to use multicast itself (i.e., multicast Mobile IP tunnels) to propagate the multicast packet to the MNs using only a single transmission by the HA.[16] The basic idea would be to establish a dynamic multicast group, with a multicast tunnel leading from the HA to all FA's that currently have one or more of the HA's mobile nodes. Since such a group is likely to be sparse and dynamic, PIM would provide an excellent means of configuring this group. The HA could tell the FA the address of the group to join as part of the registration process for the mobile node.

5.3 Current Status and Future Work

We have implemented the Mobile IP (unicast) datagram routing functionality in 4.3 BSD UNIX, and are presently working on integrating support for our multicast scheme. Descriptions of our Mobile IP implementation and its performance are available in [Chikarmane *et al.* 96].

Our current efforts focus on a more formal specification (including state diagrams) of a protocol to support Mobile Multicast (MoM). At the moment, integration of our mobile multicast mechanism with PIM seems very attractive. For example, the rendezvous point (RP) in PIM would make a natural choice for the location of the Foreign Agent (FA) in our approach.

5.4 Related Work

To the best of our knowledge, only one other research group (at Rutgers University) is considering multicast support for mobile hosts.[17] We thus focus our discussion here on the Rutgers work [Acharya and Badrinath 93, Acharya and Badrinath 96, Acharya *et al.* 96].

An approach to multicasting for mobile hosts was first proposed by Acharya and Badrinath [Acharya and Badrinath 93]. Their 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. Thus the approach does not extend easily to IP multicast and the host group model.

More recently, however, Acharya *et al.* [Acharya *et al.* 96] have proposed extensions to IP multicast to support mobile hosts. Their 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. Their current implementation uses DVMRP, although the authors claim that the approach would work with other multicast routing algorithms.

[15] For example, measurements of our Mobile IP implementation [Chikarmane *et al.* 96] show that processing times for unicast packets almost *triple* (from 700 microseconds to 1900 microseconds) when IP fragmentation occurs, which causes two outgoing packet transmissions (one large, one small) for each incoming packet.

[16] The idea of using multicast tunnels to provide multicast delivery has been independently proposed by Acharya *et al.* [Acharya *et al.* 96]. However, their approach uses a *statically configured* multicast tunnel (MTUNNEL) that leads to *all* mobile support routers (MSRs) in the network.

[17] Several members of the IETF Mobile IP working group (e.g., Tony Li) are (of course) also working on support for mobile multicast. A separate RFC for mobile multicast will likely be forthcoming soon.

Our work differs from their 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 nodes in a campus area are assigned addresses from the *same* subnet, which is used *exclusively* for mobile nodes. The multicast problem then reduces to creating the illusion of link-level multicast delivery to this subnet, regardless of where the mobile nodes 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 first proposed in [Chikarmane *et al.* 95]). 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 nodes. 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 node 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 nodes, rather than multicast. Our approach is much more conscious of the bandwidth used for the delivery of multicast packets to mobile nodes.

6 Summary and Conclusions

This paper has described a new approach for 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 networks. 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. That is, the approach "glues together" Mobile IP (which does not (yet) explicitly consider multicast support) and IP multicast (which does not explicitly consider mobile host support) to provide multicast support for mobile hosts. 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.

The proposed approach has limitations, however, 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 is hoped 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) through research grants OGP0003707 and OGP0121969, by Telecommunications Research Laboratories (TR*Labs*), and by IBM Canada.

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 and Badrinath 96] A. Acharya and B. Badrinath, "A Framework for the Delivery of Multicast Messages in Networks with Mobile Hosts", to appear in *Wireless Networks*, 1996.

[Acharya *et al.* 96] A. Acharya, A. Bakre, and B. Badrinath, "IP Multicast Extensions for Mobile Internetworking", to appear at 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.

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

[Chikarmane *et al.* 95] 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.* 96] V. Chikarmane, C. Williamson, and R. Bunt, "Implementing Mobile IP in 4.3 BSD Unix: Design, Implementation, and Performance Evaluation", submitted to USENIX *Computing Systems*, April 1996.

[Deering 91] S. Deering, *Multicast Routing in a Datagram Internetwork*, Ph.D. Thesis, 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.

[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.

[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.

[Ioannidis *et al.* 91] J. Ioannidis, D. Duchamp, and G. Maguire, "IP-based Protocols for Mobile Internetworking", *Proceedings of the 1991 ACM SIGCOMM Conference*, Zürich, Switzerland, pp. 235-245, September 1991.

[Ioannidis 93] J. Ioannidis, "Protocols for Mobile Internetworking", Ph.D. Thesis, Columbia University, 1993.

[Johnson 94] D. Johnson, "Scalable and Robust Internetwork Routing for Mobile Hosts", *Proceedings of the 14th International Conference on Distributed Computing Systems*, June 1994.

[Johnson *et al.* 94] D. Johnson, A. Myles, and C. Perkins, "The Internet Mobile Host Protocol (IMHP)", *Internet Draft*, February 1994.

[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.

[Myles and Perkins 93] A. Myles and C. Perkins, *Mobile IP Extensions (MIP)*, Internet Draft, October 1993.

[Perkins 95] C. Perkins, *IP Mobility Support*, draft-ietf-mobileip-protocol-11.txt, Mobile IP Working Group, Internet Draft, July 1995.

[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, 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.