Mobile IP-based Multicast as a Service for Mobile Hosts

Vineet Chikarmane Rick Bunt Carey Williamson

Department of Computer Science
University of Saskatchewan
Saskatoon, Canada S7N 5A9

Abstract

There is an increasing need to support user mobility in today's computing environments. The incorporation of multicast routing support in a mobile computing environment presents several challenges. Of special interest are mobile multicast schemes suitable for TCP/IP networks in general, and the Internet in particular.

This paper outlines a design of network layer modules in a 4.3 BSD UNIX^{TM} system to provide mobile routing support. This basic routing capability is then 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 effectively provide multicast services for mobile hosts as well.

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. This convergence of a need and products to satisfy that need has created an explosion of interest in problems related to the mobile computing paradigm [FZ94].

Efficient and effective *multicast* support could be very valuable in providing services to mobile users. Multicast is a mechanism for *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. A number of mobile computing services could be provided effectively over a multicast facility, such as the dissemination of information like 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 good use of a general purpose multicast facility. An approach to multicasting for mobile hosts was proposed by Acharya and Badrinath [AB93]. Their scheme manages location information on a *per-group* basis using a *host view membership protocol, HVMP,* executed within the static network. In this paper, we propose an alternative approach to mobile multicast based on unicast routing support for mobile hosts and existing Internet multicast.

This paper is organised as follows: Section 2 examines the problem of multicast to mobile hosts. Section 3 examines a new approach that we are exploring to provide multicast service over a raw unicast routing facility for mobile hosts called Mobile IP [Per95]. Finally, Section 4 concludes with an evaluation of our approach, a report on the status of our work, and an indication of the issues we intend to examine in the future.

2 Multicast Routing

Wide area multicast has been widely implemented on the Internet. In this section, we discuss multicast in the Internet today, and the problem of host mobility in the context of Internet multicast.

2.1 Internet Multicast

Deering [DC90] designed a form of WAN multicast that has been incorporated into the TCP/IP suite. Conventions have been defined for recognising 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 [Dee89], is used by hosts to inform multicast routers about the multicast addresses the hosts wish to receive. A group

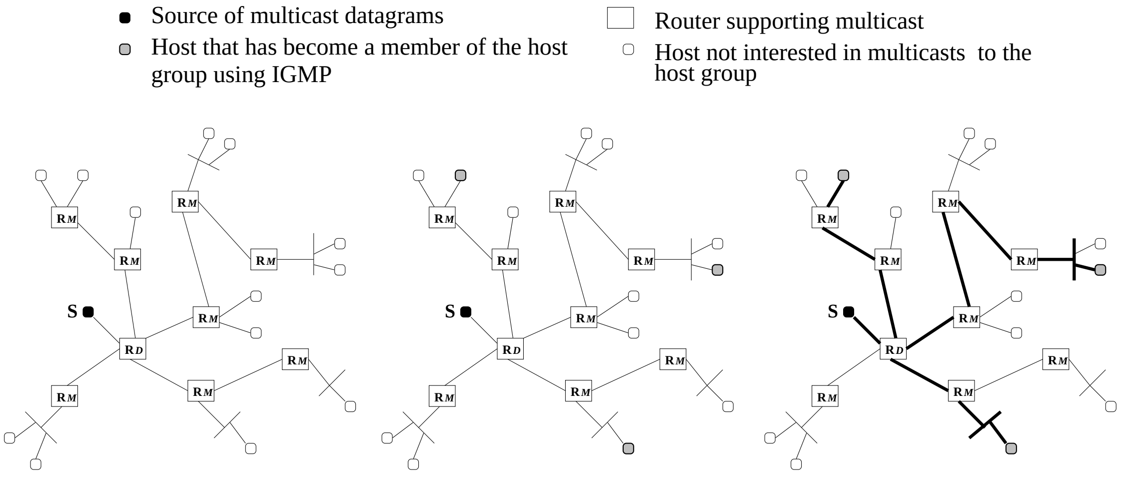

(a) Spanning tree for S **(b)** Listeners for multicast address M **(c)** Pruned spanning tree for datagram to destination IP address = M

Figure 1: Multicast in a TCP/IP internet

of hosts that might wish to communicate with each other through a given multicast address is called a *host group*. 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. DVMRP [WPD88], a distance-vector routing protocol, has been implemented as an extension to RIP [Mal94], and MOSPF [Moy94a], a link-state routing protocol, has been implemented as an extension to OSPF [Moy94b].

In particular, DVMRP has been deployed in hundreds of regions connected by the MBONE [Eri94]. We restrict our discussion of IP multicast in this paper to DVMRP since it is widely implemented. DVMRP uses the *truncated reverse-path broadcast* routing algorithm [DC90]. DVMRP differs from unicast distance vector routing protocols such as RIP in an important way: while RIP thinks in terms of possible *paths* to a particular *destination*, DVMRP keeps track of the *return paths* to the *source* of multicast datagrams. Routers on the multicast network that receive multicast datagrams are said to be *downstream* from the routers that deliver the datagrams to them. Directly connected multicast routers use link-level multicast mechanisms such as Ethernet multicast to deliver multicast datagrams to downstream routers in the tree. When multicast routers are separated by arbitrary subnets that do not support network-level multicast, they communicate using *tunnels*. Tunnels are implemented using IP-in-IP encapsulation or source routing [WPD88]. Once a multicast datagram reaches a leaf network that has host group members, link-level multicast is used to perform delivery to the listening hosts.

The IP routing algorithm is modified so that each multicast router, R_M, informs all other multicast routers about which multicast addresses are being listened to by R_M's neighbours. Each router also calculates a spanning tree per source as shown in Figure 1(a). Any multicast datagram sent by a source is picked up by a designated multicast router, R_D, for the network on which the datagrams were transmitted. This router assumes membership of all multicast routers of the spanning tree that are downstream. A datagram from the source is propagated downstream to all the multicast routers including those that have joined the multicast group using IGMP, as shown in Figure 1(b). 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. 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). Clearly, the cost of tree set-up is significant in terms of overall bandwidth consumed by datagrams sent by the sending host.

2.2 The Implications of Mobility

The following terms are used in this paper to identify the principal network entities that are of interest. A *mobile node, MN*, is a node that may move through an IP internet. A *correspondent host, CH,* is any arbitrary host communicating with the MN. The network that corresponds to the network component in the MN's IP address is called the *home network* of the MN. Any other network that a MN may visit is referred to as a *foreign network*. When the MN moves 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 the MN's home network. Since the MN is no longer present on its home network, unsuccessful delivery results.

When a MN wishes to *receive* multicast datagrams, the routing problem inherent in unicast does not exist. Because these datagrams have a class D IP multicast address specified as their destination address, no network is encoded, and multicast routing of these datagrams can take place even 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 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.

Unless a MN is at its home network, multicast datagrams it wishes to *send* could possibly be dropped. This is because if downstream multicast routers receive the multicast datagrams on a different interface from that used to send datagrams to the MN, then as mentioned in Section 2.1, these datagrams are dropped. Further, 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.

3 Multicast as a Mobile Service

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

3.1 Overview of the Proposed Scheme

We consider the problem of multicast to dynamic groups in an 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. We propose an intuitive scheme that cleanly separates the location management and multicast functions. We rely on a unicast routing protocol to provide location management, and provide multicast as a service accessed from the home network of the mobile host. The goal of our scheme is to transparently provide multicast capability to mobile hosts in a manner similar in function to that available on the Internet today. The following assumptions have been made:

- The disruption of multicast service due to movement of a host from one network to another must be minimised.
- 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 the robustness of the Internet and is largely responsible for its success.
- Dynamic group management is a necessary feature of multicast.
- Both static and mobile hosts can be members of multicast groups.
- A mobile host that wishes to receive multicast datagrams is capable of receiving them on its home network without any additional support, using existing multicast techniques.

3.2 Providing Basic Mobility Support

Several proposals including [IDM91, TCT92] have been made for network layer protocols to effect unicast routing in the face of host mobility. An IETF

Working Group has been chartered to coordinate efforts towards the specification of a protocol called *Mobile IP* [Per95] to handle host mobility in the Internet. This subsection gives a brief overview of the protocol interactions relevant to our work that take place between the functional entities defined by the IETF Mobile IP Working Group. The 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.

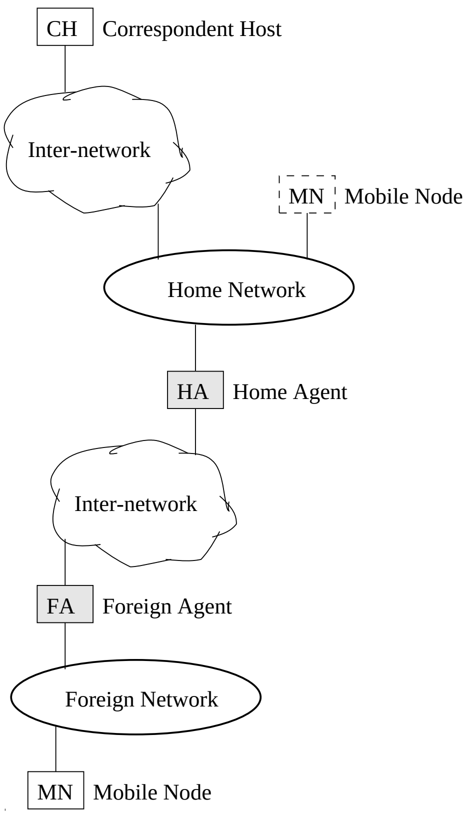

Figure 2: Mobile IP entities and their interrelationships

If the MN moves to a foreign network, a registration procedure is used by the MN to register with the HA through the FA. The HA arranges to intercept packets addressed to the MN and tunnel the packet, by encapsulating and sending it to the FA that currently has the MN registered. When a FA receives a tunneled packet addressed to a MN on its visitor list, it decapsulates the packet and delivers it locally to the MN. We have implemented the Mobile IP datagram routing functionality in 4.3 BSD UNIX. The home and foreign agents are basically routing entities that must be capable of handling bursts of packets with minimal packet loss. A user-level implementation of these agents would imply: a context switch overhead on every packet processed, and expensive copying to/from buffers for every packet during encapsulation and decapsulation. A kernel implementation of the routing mechanism would eliminate the context switch overhead and permit the use of UNIX *mbufs* for header manipulation. Since management packets and packets carrying location information are infrequent, it was decided to implement their processing in a user-level mobility daemon. Hence, a split-level approach was selected and the routing mechanism was implemented as shown in Figure 3.

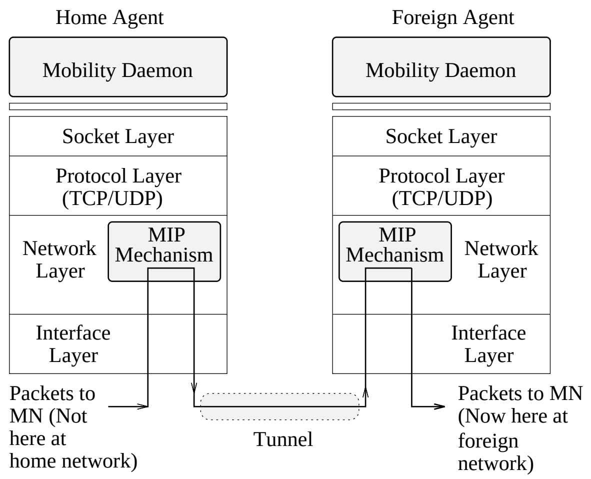

Figure 3: Architecture of the Mobile Routing Software

For the purpose of prototyping, we have chosen to implement our agents using *gratuitous proxy ARP* as described in [Per95] to enable the HA to grab packets addressed to the MNs it is serving. This results in all directly connected nodes on the MN's home network mapping the MN's IP address to the HA's link-level address in their ARP caches. Thus datagrams addressed to the MN are delivered to the HA which can then perform further processing. Location information is managed in user-space, but is also stored in kernel tables to enable datagram tunneling decisions to be made. The system call interface has been augmented with two new input-output control operations (`ioctl`s) to permit the daemon to update the kernel's tables. ICMP messages returned from the tunnel interior to the HA are used to update *soft state* [Per95] information on the tunnel. This state is then used to generate ICMP error messages for the CH.

This routing subsystem provides us with a basic foundation over which more sophisticated functionality can be incorporated. Given that the HA keeps track of the location of the MN, it is an ideal platform for providing multicast services to the MN. Hence, in our architecture, we assume that a multicast router is co-resident with the HA, which we denote as *mHA*.

3.3 Handling Multicast Source Mobility

One approach to solving the problem associated with a mobile *source* of multicast datagrams is to make the multicast network react to changes in the location of the source of multicast datagrams. This would potentially involve changing the algorithm currently used for IP multicast, possibly using some mechanism to invalidate old routes upon host movement. Such an approach seems prohibitively expensive in terms of both the cost of repeatedly modifying incorrect routes and the cost of repeated tree set-up. We would like to achieve reuse of the multicast tree once it has been set up. To ensure interoperability with DVMRP, our approach guarantees that the multicast tree is always rooted at the home network of the MN. When the MN wishes to send multicast datagrams,

- if it is on its home network, it uses link-level multicast to send the datagram. 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. The mHA then propagates the multicast downstream via all virtual interfaces, including the physical network interface corresponding to the MNs home network. The recipients of the multicast are unaware that the source is mobile.

Higher level protocols that implement retransmission mechanisms such as negative acknowledgements (NAKs) unicast these datagrams to the MN. Mobile IP software on the mHA provides for the delivery of these datagrams to the MN. The result is transparent operation in the face of multicast source mobility.

3.4 Handling Multicast Destination Mobility

Future public wireless networks such as those used for large scale information dissemination services would likely be served by one source or a few centralised sources of information, such as stock market reports originating from a few stock exchanges, or weather reports from a regional weather bureau. This information would be received by a very large number of MNs via a large number of access points. Thus in order to support mobile *destinations* of multicast datagrams the network must provide:

- a few access points to a general multicast infrastructure (the mHAs), and
- several access points to permit the dissemination of messages to the MNs (the FAs).

The underlying philosophy of our approach is that HAs are static and also know the location of the MNs they serve. Thus it is logically feasible for a multicast HA, i.e. a mHA, to forward multicast traffic to the MN through the MIP tunnel via the FA. This occurs transparently to software on the MN. By not maintaining group information at all, and rather passing this burden onto a tried and tested protocol, the resulting scheme is simple. Each FA needs to worry only about the MNs it is currently serving rather than all the agents serving all the members of all the groups represented on its local network. This information is mostly already known to the FA through the registration procedures used in Mobile IP.

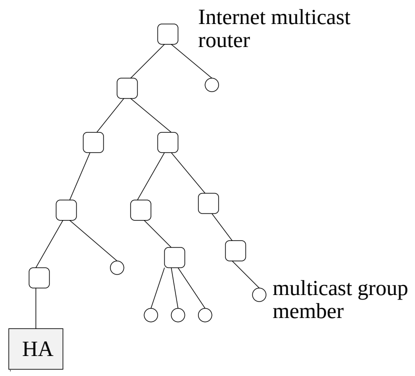

Figure 4: IP Multicast Tree

As shown in Figure 4, the HA receives multicast datagrams via the IP multicast distribution tree set up using DVMRP. Since the HA may be serving MNs at several FAs that wish to receive datagrams addressed to multicast address X, it forwards a copy into each corresponding MIP tunnel as shown in Figure 5. This scheme is complicated by the *tunnel convergence problem*, shown in Figure 6, resulting from the fact that multiple MIP tunnels (from different HAs) can terminate at a given 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 our assumptions, they constitute an additional load on a possibly low-bandwidth network. To solve this, a simple selection based on maximum number of MNs served is performed by the FA to appoint one HA as the *designated multicast service provider, DMSP* for a given multicast group. The DMSP forwards only one multicast datagram into the tunnel even if it serves multiple 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.

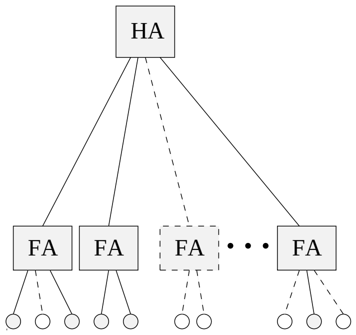

Figure 5: Multicast Propagation Tree

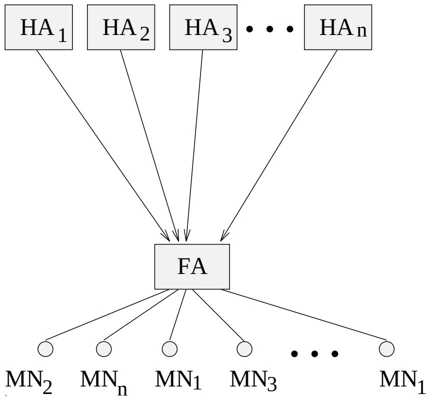

Figure 6: Tunnel Convergence Problem

3.5 Service Mechanism

All multicast transmissions by the MN are tunneled directly to the mHA, instead of using hardware multicast. Any recipients on the foreign network receive the datagrams through the multicast tree rather than directly via link-level multicast by the MN. The steps involved in providing multicast service for mobile recipients are:

Service Request: The 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 X, 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. If the MN is the first mobile on the FA's network to request service for address X, the FA also informs the mHA of the mHA's status as the *designated multicast service provider, DMSP*, for address X. The mHA does not prune itself from the multicast distribution tree, and receives all multicast datagrams addressed to X.

Service Execution: The multicast service is provided when the DMSP forwards all datagrams received corresponding to multicast address X into the MIP 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 specifies that the FA eventually times out and discards the MN's registered status when it fails to receive a re-registration 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 X, 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 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 [HKM+88] 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 X, a procedure called a *callback break*. This triggers the selection of a new DMSP for address X by the FA if any MNs remain that are still members of the group corresponding to address X. This callback mechanism ensures that the break in service for MNs is minimised 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.

4 Evaluation and Current Status

Our scheme has significant differences from that presented in [AB93]. Their scheme assumes static multicast groups, i.e. that the membership of any group does not change during the group's lifetime. Global state is managed using a centralised coordinator. An implementation of the scheme would use direct physical links or tunnels between a root multicast router and each of the leaf multicast routers with group members. The scheme implies that even if the multicast tunnels share common physical links, multiple datagrams will still be transmitted on those links. The scheme assumes that all the system components (hosts, routers, and communication links) are reliable and that routers have sufficient memory to buffer multicast messages as needed in order to provide *exactly once* delivery semantics. The multicast problem is approached as being distinct from, and independent of, any particular solution for mobile unicast routing. They focus solely on host mobility and consider multicast messages that are addressed only to mobile hosts. If a host is not a member of a group, it can multicast only if it is local to a multicast router that belongs to the host view.

Our 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. We believe our multicast scheme has several features that make it very practical as a solution for mobile hosts on TCP/IP networks. In particular, it provides the ability to support dynamic groups and provide a minimal break in service due to host movement. It also provides an element of robustness due to the availability of alternate multicast service providers in the event of a DMSP crash or a network partition.

We have implemented the necessary functionality for the support of Mobile IP in a 4.3 BSD UNIX-based kernel and are presently working on integrating support for our multicast scheme. We noticed very early that an obvious weakness of our scheme is the non-optimal routing that is employed for delivery of multicast datagrams. This is a consequence of the use of Mobile IP as the underlying foundation. We intend to attack this by conducting performance studies to evaluate

- mechanisms to rely on the FA itself to provide the service when possible
- DMSP selection strategies using metrics such as propagation delay through the multicast virtual network and Mobile IP tunnel.

Our multicast scheme relies primarily on the HA for providing the multicast service, with the FA serving to facilitate access to the service. The FA presents itself as the ideal platform for providing services that are *local* to the foreign network, such as print and local information services. By virtue of always being aware of the location of MN, the HA is ideally suited to provide services that are *global* in nature, such as mail, messaging, and the multicast service presented in this paper. In future we hope to examine further how the Mobile IP infrastructure can be leveraged to serve as a flexible platform for providing services in general for mobile hosts.

Acknowledgements

Financial support for this research was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC) through research grants OGP0003707 and OGP0121969, by Telecommunications Research Laboratories (TR*Labs*), and by IBM Canada. The authors would also like to thank Prithviraj Narayanan who provided valuable insight during the initial stages of the implementation work.

References

[AB93] A. Acharya and B. R. Badrinath. Delivering multicast messages in networks with mobile hosts. In *Proc. 13th International Conference on Distributed Computing Systems*, pages 292–299, Pittsburgh, PA, May 1993.

[DC90] S. Deering and D. Cheriton. Multicast routing in datagram internetworks and extended LANs. *ACM Transactions* *on Computer Systems*, 8(2):85–110, May 1990.

[Dee89] S. Deering, editor. *Host Extensions for IP Multicasting*, RFC 1112. Stanford University, August 1989.

[Eri94] H. Eriksson. MBONE: The multicast backbone. *Communications of the ACM*, 37(8):54–60, April 1994.

[FZ94] G. H. Forman and J. Zahorjan. The challenges of mobile computing. *IEEE Computer*, 27(4):38–47, April 1994.

[HKM+88] J. H. Howard, M. L. Kazar, S. G. Menees, D. A. Nichols, M. Satyanarayanan, R. N. Sidebotham, and M. J. West. Scale and performance in a distributed file system. *ACM Transactions on Computer Systems*, 6(1):51–81, October 1988.

[IDM91] J. Ioannidis, D. Duchamp, and G. Q. Maguire. IP-based protocols for mobile internetworking. In *Proc. of the ACM SIGCOMM '91 Conference: Communications Architectures and Protocols*, pages 235–245, Zurich, Switzerland, September 1991.

[Mal94] G. Malkin, editor. *RIP Version 2 Carrying Additional Information*, RFC 1723. Xylogics, Inc., November 1994.

[Moy94a] J. Moy. Multicast routing extensions for OSPF. *Communications of the ACM*, 37(8):61–66, August 1994.

[Moy94b] J. Moy, editor. *OSPF version 2*, RFC 1583. Proteon, Inc., March 1994.

[Per95] C. Perkins. IP mobility support. Internet Draft, Mobile IP Working Group, IETF, January 1995.

[TCT92] F. Teraoka, K. Claffy, and M. Tokoro. Design, implementation, and evaluation of virtual internet protocol. In *Proc. 12th International Conference on Distributed Computing Systems*, pages 170–177, Yokohama, Japan, June 1992.

[WPD88] D. Waitzman, C. Partridge, and S. Deering, editors. *Distance Vector Multicast Routing Protocol*, RFC 1075. BBN STC and Stanford University, November 1988.