

A New Multicasting-based Architecture for Internet Host Mobility

Jayanth Mysore Vaduvur Bharghavan
Coordinated Sciences Laboratory
University of Illinois at Urbana-Champaign
{babu, bharghav}@timely.crhc.uiuc.edu

Abstract

IP multicasting provides a mechanism for location independent addressing and packet delivery to a group of hosts that belong to a multicast group. It also provides efficient mechanisms for hosts to join and leave multicast groups. The problem of supporting macro and micro level host mobility in the Internet involves similar issues of location independent addressing, address translation, packet forwarding and location management of mobile hosts.

In order to exploit the commonality of goals in the two problems, we have investigated a new approach for supporting host mobility using IP multicasting as the sole mechanism for addressing and routing packets to mobile hosts. Mobile hosts are identified by unique multicast IP addresses, and use the IP multicasting infrastructure to route packets destined for mobile hosts. We have implemented a mobile computing environment which uses this routing approach, and our performance analysis shows that it is very effective in supporting seamless mobility for both handoffs and interface changes within and across networks.

We discuss a number of issues which need to be resolved in order to make our approach viable and practical for supporting host mobility in the Internet.

1 Introduction

In the Internet addressing format, the unicast IP address of a host identifies the host uniquely and also identifies the network to which it is attached. Thus, a unicast IP address plays the dual role of both end-point identification and location identification. While the former role is used by TCP for connection management (a TCP connection is uniquely identified by the 4-tuple $\langle \text{src addr, src port, dst addr, dst port} \rangle$), the latter role is used by IP for routing packets to the host. Thus, when a mobile host changes its point of attachment from one network to another, its unicast address must be modified in order to reflect this change. The evolution of TCP/IP networks

did not consider the problem of mobility until recently. The TCP connection identifier is set up during connection establishment, and is not expected to change during the lifetime of the connection. Special mechanisms are therefore needed to support mobility while still allowing ongoing TCP connections to transmit and receive packets transparently. While simple changes in the TCP protocol which allow the connection identifier to change dynamically during a connection lifetime can solve the mobility problem, this will involve changing the protocol stack on every host and notifying all connection end-points when a mobile host moves. Thus, Mobile IP and its variants seek to provide a network-level solution which enables static hosts to remain oblivious to the dynamic address changes in the mobile hosts. In Mobile IP, a mobile host has two addresses: an invariant *home address*, which is used for end-point identification, and a dynamic *care-of address*, which is used for location identification. While TCP uses the home address for connection identifiers, IP uses the care-of address for routing. Mobile IP maps the home address to the care-of address, and hence provides location independent addressing and routing mechanisms for mobile hosts.

Interestingly, the problem of *multicasting* faces the identical challenge of providing location independent addressing and routing support, though in a different context. In IP multicasting, hosts spread across different networks can join a multicast group, and can receive packets that are sent to the group. A message sent to a multicast group does not embed the location of the destination hosts. Hence, location independent addressing and routing mechanisms are necessary to support multicasting.

The similarity in the fundamental nature of the two problems has resulted in the development of almost identical architectures for supporting the two services. The architectures are compared in Table 1. The commonality of the issues that need to be addressed in mobility support and multicasting caused us to explore the possibility of using the evolving Internet multicasting infrastructure in order to support host mobility. In particular, our goal is *to use multicasting as the sole mechanism to provide addressing and routing services for mobile hosts*. We assign a unique multicast address to each mobile host - consequently, all packets from a correspondent host to a mobile host are treated as multicast packets, while packets from a mobile host to a static host are treated as standard unicast packets. *The use of multicast addresses to identify mobile hosts enables us to support mobility using the multicasting infrastructure with minimal or no changes to the*

backbone networks. This is a major departure from contemporary approaches for supporting mobility [3], which typically use multicasting only as a support mechanism for reducing packet loss during handoff.

In this paper, we first show that the fundamental issues faced by mobile host support are a sub-set of the issues faced by any network level multicasting approach. We then present an architecture that uses multicasting to support mobility. We describe our implementation test bed and present performance results which show that our approach performs well for both handoffs and interface changes within and across subnetworks. This leads us to conclude that our approach could provide a viable alternative to current state-of-the-art solutions. Finally, we discuss several key issues that need to be resolved before our approach can be deployed in a widespread manner in the Internet. Our experiences indicate that the multicasting technology is still not mature enough to support ubiquitous host mobility in the Internet; however, the fundamental approach of using multicasting to support host mobility not only solves the problem of supporting mobility within and across networks efficiently, but also has several desirable properties in terms of resource reservation in future Integrated Services Networks, reducing handoff dropping, and advance registration/routing.

It is important to note that we are *not* claiming that the current IP multicasting infrastructure is sufficient to adequately support mobility in the Internet. Rather, the goal of this paper is to show that any viable multicasting infrastructure for the next generation Internet will need to address all the issues that are germane to the problem of supporting host mobility in the Internet (besides other issues that are unique to multicasting). Hence, we suggest using a multicasting infrastructure to solve the problem of mobility support essentially for free.

The rest of the paper is organized as follows. Section 2 presents the equivalence between the multicasting and mobility problems. Section 3 describes the multicasting architecture to support mobility. Section 4 describes our test bed implementation and presents the performance results. Section 5 discusses key issues which need to be resolved before an IP multicasting infrastructure can be used in practice to support Internet host mobility. Section 6 compares our approach with related work, and Section 7 concludes the paper.

2 Equivalence of Multicasting and Mobility Support

2.1 The IP-Multicasting Approach

The purpose of multicasting is to enable transmitters to send packets to multiple hosts in an efficient manner. Multicasting thus provides for a common communication channel with multiple receivers. In order to provide this service, the following issues need to be addressed.

1. Multicast IP addresses should identify location independent communication channels rather than location specific host addresses.

2. Packets with multicast destination addresses should be forwarded to all hosts in the multicast group irrespective of their location.
3. Mechanisms for locating, updating, and managing multicast groups, addresses and membership need to be developed.
4. Hosts should have the ability to dynamically join and leave multicast groups.

The IP multicasting approach addresses the above issues by assigning a multicast address to a group of receivers. Senders simply use the multicast address to send packets to the group. Specialized multicast routers are created in every subnet that supports multicast recipients. A group membership protocol called IGMP is used for routers to learn the existence of members of multicast groups in their directly attached subnets. Multicast routers establish a virtual network among themselves, on top of the underlying internetwork. ‘Links’ in the virtual network between any two multicast routers are tunnels between the two routers which may span multiple unicast links. Multicast routers exchange information between themselves in order to maintain and propagate membership information of different multicast groups. When a host addresses a packet to the multicast group, the multicast router on its local subnet picks up the packet, and assumes the burden of delivering the packet to all members of the multicast group. Several candidate multicast routing protocols, such as DVMRP [23], MOSPF [15], PIM [8], CBT [2], etc. are available to perform the routing among the multicast routers. Eventually, every multicast router which has at least one member receives the packet, and locally multicasts it in the subnet. While a detailed discussion of the candidate multicast routing algorithms is beyond the scope of this paper, the common distinguishing feature of all these algorithms is that they use the multicast router virtual network in order to tunnel packets to the destinations. End hosts receive multicast packets as a result of a directory-based packet forwarding mechanism rather than the use of location-specific unicast destination addresses. Figure 1.a illustrates the routing mechanism.

2.2 The Mobility Support Approach

In order to support host mobility, the IP address needs to be location independent and invariant. However, for IP to route packets, the host address needs to reflect the network address of the current point of attachment. In order to reconcile these two conflicting requirements, most contemporary approaches provide a two-level addressing scheme. While the different approaches for providing mobility support are outlined in Section 6, the generic paradigm is to use packet forwarding and address translation *agents* in order to tunnel packets to the current location of the mobile host. Each subnet that supports mobility provides an *agent* (or at least a mechanism to provide a location specific temporary address). For a mobile host, the agent in its home network (identified by its invariant home address) is its home agent and the agent

Functionality	Mobile Host Support	Multicasting
Registration	A mobile host that enters a new network must register with the local agent.	A host that wishes to receive multicast datagrams must register with the local multicast router.
Connectivity	Connectivity to the rest of the internet is provided by the foreign agent.	A multicast router provides connectivity to the rest of the virtual multicast network to a host.
Data forwarding	The foreign agents forward datagrams to the mobile host.	The multicast router multicasts datagrams to the members in its subnet.
Address translation	Home agents translate the home address to the mobile host's care-of address.	Multicast messages need no address translation
Routing	The home agent tunnels messages destined to the mobile host to the current foreign agent.	Messages sent to a destination multicast address are forwarded to multicast routers with members using a multicast routing protocol.

Table 1: Comparison of Multicasting and Host Mobility Architectures

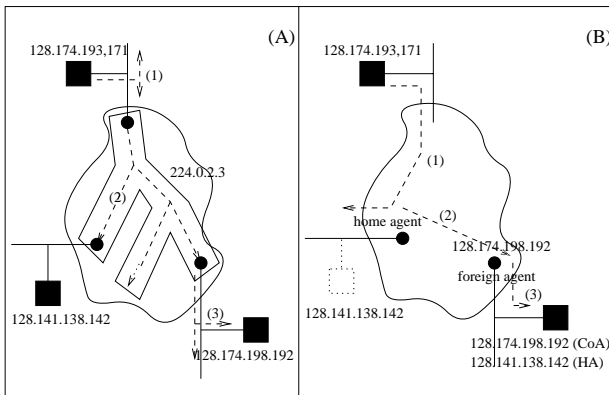


Figure 1: Figure (a) shows the multicast model. In Step (1), a host transmits a multicast packet. In Step (2), the multicast router tunnels the multicast packet along the multicast distribution tree. In Step (3), each destination multicast router decapsulates the tunneled packet and multicasts it locally in its subnet. Figure (b) shows the mobility model. In Step (1), the correspondent host transmits the packet to the home address of the mobile host. In Step (2), the home agent tunnels the packet to the foreign agent. In Step (3), the foreign agent decapsulates the tunneled packet and forwards it to the mobile host.

in its current network (identified by its care-of address) is its foreign agent.

When a mobile host moves to a different network, it registers the care-of address (foreign agent) with the home agent. Hosts on the Internet are aware only of the home address of the mobile host. When the home agent sees a packet addressed to a mobile host that is registered with a foreign agent, it tunnels the packet to the foreign agent, which then forwards the packet to the mobile host. Uplink transmissions from the mobile host to the correspondent host may be directly forwarded to it using standard IP mechanisms. In the above architecture, the agents form a virtual network among themselves, es-

entially providing a similar type of location independent service to mobile hosts, like the multicast routers provide to the members of a multicast group.

The basic motivation of our work is the observation that the mobility problem is subsumed in the multicasting problem. Both problems require the following key issues to be addressed: *providing the abstraction of location independent addressing, packet forwarding, and location management*. Several contemporary approaches for providing mobility support have been proposed. Bhagwat et al [3] provide a survey of some of these approaches, and compare them based on their solutions to the above issues. Among these approaches, Mobile IP is emerging as the standard for mobility support in IPv4. In Mobile IP, address translation is performed by the home agent (mapping home address to the care-of address); packet forwarding is done by the home agent and foreign agent; location management is done by the home agent.

It has been noted in both [17] and [19], that while Mobile IP (and IPv6) is designed to address the macro-level mobility problem such as supporting host mobility over wide area networks, it does not address micro-level mobility issues such as packet loss and delay due to hand-offs. We address the issues at the two levels in an integrated approach, by using the emerging multicasting infra-structure to provide support for host mobility in the Internet. In our approach, we *assign unique Class-D addresses to each mobile host. Messages sent to a mobile host will therefore be treated like multicast messages, and routed through a network of multicast routers to the host.*

In the rest of this paper, we refer to our approach as MSM-IP, which stands for Mobility Support using Multicasting in IP.

2.3 MSM-IP versus Mobile IP

While the problems of supporting multicasting and mobility are very similar, there are several differences in the state-of-the-art approaches for IP multicasting and Mobile IP. We first contrast IP multicasting with Mobile IP, and then compare it to IPv6.

The MSM-IP approach for supporting mobility differs from Mobile IP in five important ways: addressing, packet

forwarding, location management, service disruption, and advance reservation/routing.

Addressing: The IP multicasting approach to supporting mobility issues a single invariant address to a mobile host. This eliminates the need for explicit address translation as in Mobile IP. On the other hand, a mobile host does not have a home network. In particular, the behavior of the mobile host and the office desktop machine of a mobile user will be different (for example, the mobile host will not receive broadcast packets on the user's office network, it belongs to a different administrative domain, etc). An obvious limitation of this approach is the size of the Class D address space in IPv4.

Packet Forwarding: In Mobile IP, since only the home agent is aware of the care-of address of the mobile host, all packets from a correspondent host to a mobile host suffer from triangular routing. In the MSM-IP approach, the multicast router on the local subnet of the correspondent host will pick up the packet and forward it directly to the multicast router of the mobile host's subnet. Thus, packets will not suffer triangular routing. On the other hand, Mobile IP does not require a mobility-aware agent on a subnet that does not support mobile hosts, while IP multicasting requires a multicast router on every subnet that supports either senders or receivers of multicast packets.

Location Management: In Mobile IP, correspondent hosts use the home address of a mobile host. Thus, it is sufficient for a mobile host to notify its home agent of its current care-of address. In MSM-IP, there is no notion of a home agent. When any multicast router wishes to forward a packet to a mobile host, it needs to locate the host using a distributed directory algorithm. Development of a scalable location management protocol (or session discovery protocol for standard multicast transmissions) is still ongoing work in IP multicasting research. We describe our approach to location management in Section 3.

Service Disruption: As mentioned in the IETF drafts, Mobile IP attempts to solve the problem of 'macro-level' mobility rather than 'micro-level' mobility [19]. Thus, critical issues such as minimizing packet loss for on-going TCP connections during mobility are not stressed. In the MSM-IP approach for supporting mobility, we are concerned with both macro-level and micro-level mobility. Our claim is that MSM-IP (in effective sparse mode operation [8]) could perform as well as Mobile IP for macro level mobility, while significantly reducing the disruption in service during handoffs. The main reason for this claim is the following: in Mobile IP, response to a mobile host's change in location occurs only after the home agent (correspondent host in IPv6) is made aware of the change. However, in MSM-IP, all joins and prunes are receiver-initiated and terminate at the earliest branch point. Thus, response to a mobile host's change in location happens as close to the mobile host as possible,

which significantly reduces the latency of response and disruption of service for long haul connections.

Advance Reservation/Routing: Several contemporary approaches have used multicasting in a more restricted form than we propose, for the purpose of reducing or eliminating packet loss during handoffs. Before a mobile host performs a hand-off, it may request the multicasting router in the new network to join the multicast group given by its address, in advance. Upon handoff, the mobile host will continue to receive packets uninterrupted. The overhead for such advance joins is expected to be tolerable, particularly in the common case of handoff between adjacent cells, since the old and the new base stations will be either in the same subnet or neighboring subnets. Similar to an advance join, it is also possible for the multicast router to initiate advance resource reservation, as described in [14].

MSM-IP Vs IPv6: IPv6 has the luxury of being able to change the IP behavior at correspondent hosts as well as mobile hosts. A current approach for supporting mobility in IPv6 is to make correspondent hosts aware of the care-of address of a mobile host as it moves. Thus, triangular routing is avoided. However, for long haul connections, service will still be disrupted upon a handoff until the correspondent host is notified about the new care-of address of the mobile host. The comparison with IPv6 is not the focus of this work, since our goal is to be as minimalist as possible in the changes we make at the backbone, and still provide mobility support and reduced packet loss/service disruption upon handoff. In the next section, we describe our IP multicasting architecture for the support of mobility.

3 Architecture

The starting premise of this work is that host mobility can be supported without making any special changes to an Internet multicasting architecture, though the state-of-the-art of IP multicasting is still inadequate for this purpose. With this in mind, we present an architecture for supporting mobility using state-of-the-art IP multicasting in this section, describe our experience with such an architecture in Section 4, and raise several issues that need to be addressed before the proposed architecture can be used in practice in Section 5.

The IP multicasting architecture consists of three types of hosts, *transmitters*, *multicast routers*, and *receivers*. A transmitter need not be a member of a multicast group to transmit to the group. Every subnet that supports either the transmission or reception of multicast packets must have a multicast router. The multicast router is typically assumed to be colocated with the base station, though it may be any host on the same subnet as the base station (for example, we can have multiple Wavelan access points attached to the same subnet served by a single multicast router). The multicast routers form a virtual network among themselves, in order to exchange group membership information, perform location/session discovery, and

forward packets.

When a receiver wants to join a multicast group, it initiates a join by sending an IGMP registration message to the multicast router in its subnet. Each multicast router maintains a distributed directory of group membership information. When a transmitter wants to send a packet to a multicast group, it sets the IP destination address to the group address and sends the packet as a local multicast. When it transmits the packet, the multicast router in its subnet picks up the packet. It is now the responsibility of this multicast router to propagate the packet to the recipients in the multicast group. Multicast routing within the network of multicast routers is performed in several ways. For example, using a distance vector protocol such as DVMRP [23], the goal is to construct a multicast distribution tree for each group based on Truncated Reverse Path Broadcasting [6]. While DVMRP works well in the *dense mode*, wherein several members of the multicast group are geographically close together, supporting mobility typically will involve multicasting in the *sparse mode* with few receivers (in this case, one, in the absence of advance registration). Several protocols such as PIM-SM [8] and CBT [2] address the issues of sparse mode operation. In this section, we present a generic architecture for IP multicasting, that can be used to support host mobility effectively. The routing protocol used in this architecture is beyond the scope of this paper, though we expect a protocol such as PIM-SM to work well so long as the Rendezvous points are chosen carefully (see Section 3.4 on location management).

We now describe the seven key components of our IP multicasting based host mobility support architecture: (a) addressing, (b) tunneling architecture, (c) join and prune mechanisms, (d) location management, (e) handoff and advance registration, (f) resource reservation, and (g) supporting and interacting protocols such as ARP, IGMP, ICMP, RARP, TCP and UDP.

3.1 Addressing

Static hosts are assigned standard unicast IP addresses. Mobile hosts are assigned unique multicast IP addresses. Thus, packets from a mobile host to a static host will be routed by standard unicast mechanisms while packets from a static host to a mobile host will be routed by the multicasting infrastructure. There are two key issues in this approach: (a) some routers will not forward packets that originate from hosts with addresses that do not correspond to the local network- such routers will drop packets from the mobile hosts, and (b) mobile hosts cannot address each other directly if the multicast routers use reverse path forwarding (RPF) (since RPF assumes that the source address is a unicast IP address in order for multicast routers to determine the shortest path to the source). Section 5 discusses these issues and provides initial solutions.

3.2 Tunneling Architecture

Each subnet that supports multicast transmission and reception has a multicast router. Each multicast router is

aware of the other routers in its neighborhood. Multicast routers exchange routing information on-demand, and build multicast packet distribution trees for each multicast group. The mobile host will initiate a group membership registration with the multicast router in its subnet, and potentially some other neighboring multicast routers (advance registration). Thus, each multicast group corresponding to a mobile host's address has a few clustered recipients. Transmitters to this group may be anywhere in the network. Using sparse mode multicast routing algorithms [8, 2], it is possible to create a multicast distribution tree for the multicast group. Each node in the tree has information about which nodes are downstream, i.e. when it receives a packet for a multicast group, to which multicast routers it should forward these packets. The design of algorithms for building efficient multicast distribution trees is still ongoing research. It should be noted that any algorithm that works well for sparse mode multicasting will work equally well for supporting mobility.

3.3 Join and Prune Mechanisms

In order to provide support for hosts to dynamically join and leave multicast groups, the IP multicasting architecture uses a soft-state based approach for membership registration. Hosts join a multicast group by registering with the multicast router through the IGMP Registration message. Periodically, hosts also send a registration refresh message to the router (when a host sees an IGMP registration message for the multicast group it belongs to, it resets its timer; thus, only one host on a subnet needs to send the refresh during a timeout period). When a multicast router (which is not a leaf of the multicast distribution tree) receives prune messages from all its children, and does not have any hosts in the multicast group in its own subnet (either it never had members in its subnet or their registration timed out), it prunes itself from the multicast distribution tree by sending a prune message to its parent(s).

Receiver-initiated joins and prunes, and the use soft-state to time out registration work very well to support host mobility. When a host enters a new domain, it initiates a join for its (multicast) address. The multicast router in its subnet then joins the multicast distribution tree. In the common case of handoff, we anticipate that either the multicast router would already have joined the multicast group through advance registration, or at least is a neighbor of the previous multicast router serving the mobile host. Thus, most receiver-initiated joins terminate at the lowest levels of the distribution tree hierarchy. Once a mobile host enters a new cell and registers with its new multicast router, it ceases to send refreshes to the old multicast router. Eventually, this router will time out its registration and initiate a prune. As in the case of joins, we expect most prunes to terminate at the lowest levels of the distribution tree hierarchy. Thus, handoff using the IP multicast architecture is very efficient. As shown in Section 4, and also in related work [20], using multicasting for handoff can reduce or eliminate transient packet loss.

3.4 Location Management

A critical issue that needs to be addressed if IP multicasting is to be deployed in a widespread manner is multicast session discovery. The first problem relates to discovering what sessions are ongoing in the Internet at any time such that hosts can join these sessions - this is a resource discovery problem. The second problem relates to discovering the nearest multicast router to join in order to receive multicast packets addressed to a group - this is a location update problem. In the current IP multicasting architecture, these problems are not adequately handled. In particular, the resource discovery issue is handled by means of distributed session directories while the location management problem is handled by flooding the network to find the nearest multicast router in the worst case [23]. In order to support mobility in the IP multicasting archi-

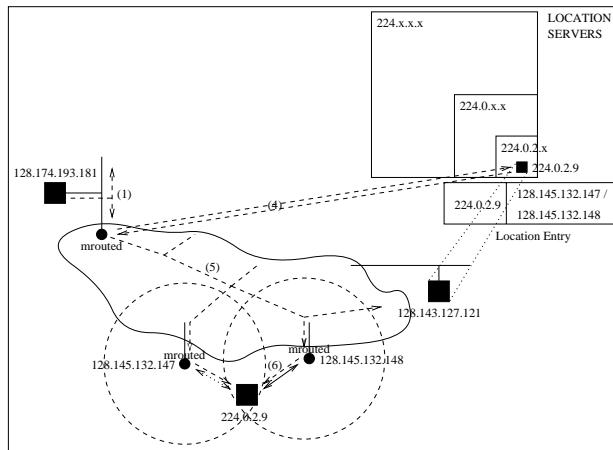


Figure 2: The location entry shows the contents at a location server corresponding to a mobile host. In Step (1), the correspondent host transmits a packet to the mobile host. In Step (2), the multicast router picks up the packet. In Step (3), it discovers the location server corresponding to the mobile host. In Step (4), it contacts the location server and retrieves the address of a multicast router which can forward packets to the mobile host. In Step (5), the transmission multicast router joins the multicast distribution tree and transmits the packet. In Step (6), each receiving multicast router decapsulates the packet and forwards it to the mobile host.

ture, the location update protocol needs to be sophisticated. We have designed a location directory protocol which is similar to the Mobile IP approach for location discovery. We divide the space of multicast addresses into *hierarchical virtual domains*, where each domain is served by a location server. The location server that serves the multicast address of a mobile host contains the location of a multicast router which ‘knows’ how to reach the mobile host (described in detail below). Route discovery to a mobile host in our environment is thus a 5 step process, as shown in Figure 2: (i) the correspondent host (CH) sends a multicast packet addressed to the mobile host

(MH), (ii) the multicast router (M1) in the transmitter’s subnet picks up the packet, (iii) M1 discovers the location server (LS) of the mobile host based on its multicast address (through a location discovery protocol that maps multicast addresses to location servers), (iv) M1 contacts LS and discovers the current multicast router entry (M2) corresponding to MH, (v) M1 communicates with M2 in order to join the multicast distribution tree.

Note, that while the architecture specifies the overall scheme for location discovery and updates, it does not specify the specific algorithms for either location server discovery in step (iii), the multicast distribution tree in step (v), or which multicast router is associated with a mobile host in the location server.

In a PIM-SM approach, for example, the rendezvous point would be identified as the multicast router corresponding to the mobile host. In our approach, we identify the current (or a recent) multicast router of the mobile host. When a mobile host moves, it periodically refreshes the multicast router in its location server. Note, that even after a mobile host moves out of its domain, the old multicast router will retain membership registration for a timeout period. So long as the mobile host sends location updates messages to its location server at a frequency greater than twice each timeout period, new transmitters which are not aware about the location of a mobile host can reach the host by looking at the multicast router information from its location server.

We believe that our location management architecture is a generic approach which can be integrated with several emerging approaches for IP multicast routing.

3.5 Handoff and Advance Registration

As mentioned above, the IP multicasting support for receiver initiated join/prune and soft-state based timeout provides effective mechanisms for initiating handoffs. One of the key advantages of multicasting is the ability to use advance registration with the new cell (or multicast router in the new domain) in order to minimize packet loss during handoff. We have used this mechanism in concert with our resource reservation approach in order to provide advance resource reservation in mobile computing environments [14].

3.6 Resource Reservation

One of the major advantages of using multicasting as the sole mechanism for mobility support is our ability to use RSVP for resource reservation and advance reservation. As noted in [24], one of the goals of RSVP is to support resource reservations in the presence of dynamically joining and leaving receivers. While recent approaches have tried to tune RSVP in order to run effectively in mobile computing environments (given that they take the Mobile IP approach for supporting mobility and still seek to perform advance reservations using RSVP), we can run RSVP without any modifications in our architecture for supporting mobility.

3.7 Interacting Protocols

Conceptually, the IP multicasting architecture can support host mobility without any changes to the existing infrastructure by simply assigning a unique multicast IP address to a mobile host. However, a number of interacting protocols make the practical implementation of such an architecture difficult without making any changes to any component either at the mobile host or the multicast routers. The following are the key protocols which interact with the IP routing protocol: ARP, RARP, ICMP, IGMP, TCP, and UDP. Among these, ARP replies sent to multicast IP addresses are not processed by the ARP module on the portable. ICMP messages cannot be routed to hosts with multicast addresses. TCP needs to change at the mobile host in order to accept packets with a multicast destination address. Unless TCP checks for a multicast source address (which some kernels do [21]), TCP at the static host does not need to change. UDP does not need to change in order to operate using multicast addresses, though a higher layer protocol needs to address the issue of lost packets and duplicates in UDP packets upon handoff or interface change. Details of the problems and the initial solutions are provided in Section 6.

The change in TCP and IGMP at the mobile host are trivial and no new functionality has to be added¹. While the above set of problems seems like a large list of incompatible protocols, note that the key problems stem from a single fact - they require a location-specific host IP address for network management or network layer processing. In other words, while the use of a multicast address is very effective in order to provide location independence and support mobility, several IP level functionalities in fact require the location specific nature of unicast IP addressing to be explicitly exposed (e.g. ARP and ICMP). There are two ways to solve the above problems - (a) provide additional functionality in multicast routers, or (b) use DHCP in order to obtain a temporary unicast address at the mobile host for the specific purpose of handling network level management. We adopt the latter approach because it is much simpler, and provides a natural solution to the routing problems - use location independent addresses to support mobility, and use location dependent addresses to support network management. Thus, with minimal changes in the supporting protocols and no change to the basic IP multicasting architecture, we can support host mobility effectively, and provide several desirable features such as advance registration, resource reservation and no transient packet loss during handoff.

4 Implementation and performance results

The MSM-IP architecture has been implemented and operational in our laboratory testbed environment for about six months now. The testbed backbone consists of three switched Ethernet subnets, which are served by gateways *durga*, *atri* and *radha* respectively. *durga* is connected to *atri* and *radha*, which are not connected directly to each other. We have two Wavelan wireless subnets, whose

gateways are attached to the backbone subnets of *atri* and *radha* respectively. All the correspondent hosts used for testing are P6-200 Gateway Pentium Pros, and all the mobile hosts used for testing are P6-120 TI laptops. All the hosts run Linux 2.0.30. *durga*, *atri* and *radha* run the `mouted` 3.81 multicast router. In all our experiments, the correspondent hosts are located in the subnet of *durga*, while the mobile hosts move between the wireless subnets of *atri* and *radha*. All mobile hosts have both Ethernet and Wavelan interfaces. Note that a handoff or interface switch between subnets is accompanied with a change in path from *durga* (i.e. at least two hops). Additionally, we have instrumented *durga* to optionally add a 100ms delay to all packets passing through it in order to simulate network latency.

Each mobile host runs a `registration` module which can register with the current `mouted` and the `mouted` serving the next cell via the backbone network (for advance registration), and send periodic refresh messages to keep alive the registrations. The applications we run in the testbed include WWW Browsers, ftp, Mpeg video over the network, editors, and file system access as well as special workloads generated for the purpose of quantifying the performance of our approach. Since the standard TCP does not work when one of the end-points is a Class D address, we made simple changes to the kernel to disable these checks for the purposes of doing performance measurements. We also instrumented TCP to record round trip time measurements. Our TCP-based applications thus use this modified TCP protocol, while while our UDP based applications use the standard UDP protocol.

We have measured the performance of MSM-IP with variations in offered workload, artificial delay over the network, applications, types of handoff, and parameters of measurement. In this paper, we present a subset of these measurements. In terms of offered load, we used UDP flows with offered throughputs of 10Kbps, 100Kbps, 512 Kbps and 1Mbps, and TCP flows at peak traffic. In terms of packet sizes, we used 512 bytes and 1 KB packets for the user generated workloads, and also used the application programs (WWW browser, ftp, Mpeg) unchanged. In terms of offered traffic patterns, we used constant interpacket delay, packet bursts, and exponential interpacket delay distributions. In terms of handoffs, we performed two types of handoffs: (a) *hot* switches, in which the mobile host registered with the new `mouted` before handing off, and (b) *cold* switches, in which the mobile host handed off before registering with the new `mouted`. In each case, we experimented with handing off between cells, subnets, and switching between interfaces. The introduction of the 100ms delay at *durga* was used in our comparison of MSM-IP with Mobile IP, as described later in this section. The parameters we measured included packet loss upon cold switches and packet duplicates upon hot switches for UDP traffic at the receiving application, and packet traces using `tcpdump` for TCP traffic. In all cases, the measurements were taken for both uplink and downlink.

We report three sets of results in this paper. The first set of results measures packet loss and duplicates for

¹Note that the change in TCP still assumes that there is only one receiver with the multicast address.

UDP traffic with switching between two interfaces connected to the same network. The second set of results measures packet loss and duplicates for UDP traffic with handoff between subnets and switching between two interfaces connected to different subnets (subnet of *radha* to subnet of *atri*). Note that this involves path changes beyond *durga*. The third set of results measures TCP round-trip time variation with switching between two interfaces in different networks.

For the first set of results, we found that for both hot and cold switches within the same network, the effect of both handoff and interface switching is not noticeable even at 1Mbps. At end-to-end throughputs of up to 100Kbps, typically no packets are either lost or duplicated in cold/hot switches. The maximum packet loss observed was 3 packets on the downlink for a cold switch involving bursty traffic at 512 Kbps. Average loss for downlink traffic is 1 packet for a cold switch at 1Mbps, while average loss for the uplink is 1.6 packets for a cold switch at 1 Mbps. For hot switches, we did not observe any duplicates, since the time between advance registration and handoff was negligible. The results of our performance analysis for handoff between cells connected to the same subnet is presented in Table 2.

Mode	Flow	Size	T'put	Loss	Dup
Hot	Down	512	10-512	0	0
Hot	Down	512	1024	1	0
Hot	Up	512	10-1024	0	0
Cold	Up	512	10-100	0	0
Cold	Up	512	512	0.8	0
Cold	Up	512	1024	1.6	0
Cold	Down	512	10-100	0	0
Cold	Down	512	512	3	0
Cold	Down	512	1024	1	0
Hot	Down	1024	10-1024	0	0
Hot	Up	1024	10-1024	0	0
Cold	Up	1024	10-512	0	0
Cold	Up	1024	1024	2	0
Cold	Down	1024	10-100	0	0
Cold	Down	1024	512	0.4	0
Cold	Down	1024	1024	0	0

Table 2: Performance of MSM-IP For Handoff Within The Same Network. The traffic model for all experiments reported in this table was Constant. ‘Mode’, ‘Flow’, ‘Size’ and ‘T’put’ indicate the switching mode, direction of the flow, packet size (in bytes) and offered UDP throughput (in Kbps) respectively. ‘Loss’ and ‘Dup’ indicate the packet loss and duplicates that were measured by the receiving application.

Causing either interface changes or handoffs across two different subnets involves additional overhead since it involves the new multicast router to join the multicast distribution tree (in the case of handoff within the same subnet, the new base station will also receive the multicast packets automatically). We present the measurements when the subnet switching is associated with

interface change. The MSM-IP performance is good even for interface switching at relatively good throughputs for wireless networks. For the second set of results, the maximum packet loss we observed was 1, and we recorded a maximum of two duplicates. The performance was studied at throughputs of 100 Kbps and 400 Kbps. The two networks shared a common gateway (*durga*) two hops away. The results of our performance measurements for switching across different networks is presented in Table 3. Note, that the results presented represent not only a handoff across neighboring networks, but also an interface change.

Mode	Flow	T’put	Model	Loss	Dup
Cold	Down	100	Constant	0.25	0
Cold	Up	100	Constant	1	0
Cold	Down	400	Constant	1	0
Cold	Up	400	Constant	1	0
Hot	Down	100	Constant	0	0.75
Hot	Up	100	Constant	1	0
Hot	Down	400	Constant	0	2.5
Hot	Up	400	Constant	1	0
Cold	Down	100	Poisson	0.75	0
Cold	Up	100	Poisson	1	0
Cold	Down	400	Poisson	1	0
Cold	Up	400	Poisson	1	0
Hot	Down	100	Poisson	0	0.75
Hot	Up	100	Poisson	0.75	0
Hot	Down	400	Poisson	0.33	1.75
Hot	Up	400	Poisson	1	0

Table 3: Performance of MSM-IP with Handoff Across Two Different Subnetworks. ‘Mode’, ‘Flow’, ‘T’put’ and ‘Model’ indicate the switching mode, direction of the flow, offered UDP throughput (in Kbps) and traffic model respectively. ‘Loss’ and ‘Dup’ indicate the packet loss and duplicates that were measured by the receiving application.

In order to measure the effect of handoffs and interface change on TCP, we traced the rtt of packets while transmitting at peak rate. An interesting phenomenon that we have observed when performing interface changes and/or handoffs is that for connections with a long delay, TCP over Mobile IP can experience loss of throughput for several seconds. Essentially, packets that were directed to the old interface are lost, and eventually, when the new interface sends duplicate ACKs, fast retransmit and recovery kicks in for a sequence of several lost packets even if the RTO does not expire. The work in [11] addresses this problem and proposes a solution which enables TCP to recognize handoff and interface change, and treat the consequent packet loss differently from congestion, which partially alleviates this problem. In this work, we use TCP unchanged (except for disabling checks for Class D source/destination addresses) over MSM-IP.

In our third set of results, we show the effect of a cold interface switch between two Ethernet interfaces between

two subnets. Figure 3 plots the `rtt` value for the sequence of TCP segments. The interface was changed at packet number 350, with the effect that the round trip time went up by about 2ms for about 20 packets before stabilizing to the old value.

When running the WWW browsing, ftp and Mpeg applications, we found no perceptible impact (qualitatively) upon the applications during or immediately after the handoff or interface change. The caveat is that we have only tested MSM-IP in our laboratory environment with handoffs and interface changes between subnets that share gateways two hops away. Further work needs to be done to explore how the performance scales to larger networks.




Figure 3: The variation of RTT with switching. The interface change occurred at segment 350, and caused the perceived round trip time of the next 20 segments to go up by approximately 2ms.

5 Key Issues and Problems

The MSM-IP architecture has several advantages and disadvantages as compared to the Mobile IP approach. Among the major advantages are the elegant integration of architectures to support multicasting and mobility, seamless mobility during both handoffs and interface changes, advance registration/reservation, and RSVP-based resource reservation for future Integrated Services networks. Among the major disadvantages are the inability of the state-of-the-art of IP multicasting to provide scalable location management and security/authentication. Over the course of our testbed experimentation, it became clear that the current MSM-IP architecture is still not sufficiently mature enough for widespread deployment. Unfortunately, most of the problems that we encountered in the MSM-IP approach are also unsolved in the context of IP multicasting currently. While it is true that the basic MSM-IP architecture should be able to support seamless host mobility without any changes in the backbone, several implementation constraints and design issues remain to be solved before an Internet-wide deployment of MSM-IP will be possible. This section

identifies the key issues and proposes initial solutions to some of them.

In terms of implementation, the following 6 issues need to be solved: (a) ARP reply processing, (b) TCP support, (c) IGMP membership registration, (d) ICMP message delivery, (e) Multicast TTL initialization in static hosts, and (f) location dependent addressing using DHCP. In terms of the design of the MSM-IP architecture, the following 2 issues need to be addressed: (a) security and authentication, and (b) scalability. We discuss each of these issues below.

1. **ARP reply processing:** ARP requests from the mobile host will have a multicast source address. The ARP response packet thus has a multicast destination IP address (but unicast link layer interface address). In the standard Linux kernel we used, the ARP response packet is thrown away, causing the host to hang at times. By modifying the kernel to recognize a multicast packet that matches the IP address of the host, we overcame this problem. An alternative solution is to acquire a dynamic location dependent address for use with ARP to overcome this problem (item 6).
2. **TCP support:** Currently, there is no support for TCP over IP-multicasting. In general, supporting a TCP-like protocol over multicasting can be exceedingly complex; however, in our environment, the multicast address corresponding to a mobile host is guaranteed to have only one destination host. In the unmodified TCP version of Linux 2.0.30, we found that when a mobile host initiates a TCP connection request, the SYN packet is acknowledged with a SYN-ACK by correspondent hosts, but the mobile host ignores the SYN-ACK packet since the destination address is a multicast address (`tcp_input()` drops the packet according to RFC 1122 4.2.3.10). The change that is required to the TCP code at the mobile host in order to make it work with multicast mobile host addresses is minimal. However, in several TCP implementations, the correspondent host drops a SYN request if the sender is a multicast address. Supporting TCP connections with such hosts will involve a change to the TCP code in both mobile and correspondent hosts. The emerging interest in reliable multicast applications will push for a variation in TCP that supports multiple communicating hosts, e.g. RMTP [13]. As in the case of other multicasting functionality identified in this paper, this is a superset of the requirements to support mobility within the MSM-IP framework.
3. **Initializing the multicast TTL:** In the current IP multicasting implementation, the default TTL for a multicast packet is set to 1. Thus, when a correspondent host transmits packets to a mobile host not in its own network, it needs to explicitly set the TTL (otherwise the local multicast router will ignore the packet). Application-transparent support is needed to set the value of the TTL to an appropriate value (as in unicast packets), so that end-user

applications can remain unaltered. This problem has more to do with the way DVMRP works than an inherent problem with the MSM-IP approach.

4. IGMP membership registration: Upon handoff or for advance registration, the mobile host needs to register its (multicast) address with the multicast router in its subnet so that the router will begin to receive the multicast packets addressed to the mobile host. Registration is performed through IGMP. There are two problems with the way IGMP works when the mobile host address is a multicast address: (a) at the mobile host, when the mobile host sends the IGMP registration message to the multicast group identified by its address, the address resolution module at the mobile maps the address to the link-level local interface address rather than the link-level multicast address; consequently, the multicast router does not look at the registration message, and (b) even if the above problem is fixed (so that the link-level address is set to the multicast link address), the multicast router is unable to determine on which subnet it must transmit the packets destined for the multicast group. Usually, by looking at the source address of the IGMP request, the router can determine from which subnet the request was made. However, when the mobile host sends a registration request, the source address is a subnet-independent multicast address. As with the ARP problem, acquiring a location-dependent dynamic unicast address solves this problem in a simple way.
5. ICMP message delivery: ICMP messages are issued by intermediate routers when a message or delivery error is detected. These messages are sent back to the original sender of the datagram. In MSM-IP messages from the mobile host have a multicast source address; consequently, intermediate routers will not be able to send ICMP packets back to the mobile host. Of course, tunneled packets by the multicast routers will still have the unicast address of the routers in their headers. Thus, the problem relates to ICMP over the local subnet of the mobile host. We are exploring simple support at the multicast router to overcome this problem as a part of ongoing work.
6. Location Independent Addressing and DHCP: While we propose MSM-IP as the architecture to support effect macro and micro-level host mobility in the Internet, several network management functions require a location dependent address for the mobile host. For these reasons, our support architecture uses DHCP [10] in order to provide dynamic location-specific addresses to mobile hosts. This approach works well for most of the network management related problems we have encountered thus far.
7. Security and Authentication: One of the major issues we need to resolve before the MSM-IP approach reaches maturity is the support for security

and authentication in such an environment. Inherently, multicasting is more open to security problems, since any secret or key will need to be shared among all the members of the multicast group. Besides, the shared key may need to be changed each time a host joins or leaves the group. The current design of IP multicasting allows any host to join a multicast group merely by registering with the multicast router in its subnet. While it is well possible to provide secure multicast networks, wherein the multicast routers are isolated from other multicast routers in the Internet, it is much harder to provide for secure multicast groups within the generic framework of IP multicasting.

Fortunately for us, the MSM-IP approach offers a much simpler context in which to provide security and authentication. In this discussion, our goal is to first provide mechanisms for authentication between communicating entities; data security is an issue that can be handled by encryption or digital signatures at a higher level. In the MSM-IP architecture, the location management hierarchy provides a convenient framework for supporting authentication. Essentially, the location server of a mobile host plays the role of the home agent in Mobile IP from the perspective of authentication. The mobile host and the location server share a secret key which can change over time. In effect, the same authentication framework that is applied to the mobile host, foreign agent and home agent interactions in Mobile IP can be applied to the mobile host, multicast router, and location server interactions in MSM-IP.

8. Scalability: One of the key features of Mobile IP is its scalability. There are no additional constraints in terms of addressing (except for dynamic address allocation where supported) imposed over standard unicast addressing. Home agents and foreign agents need be deployed only in networks that support mobile hosts. Location management is distributed and handled by the individual home networks of mobile hosts. Consequently, Mobile IP substantially retains the laissez-faire nature of the Internet. On the other hand, response to a handoff only occurs when the new care-of address propagates back to either the home agent or the correspondent host (in IPv6). Thus, in terms of support for seamless roaming with ongoing connections, the scale of the connections is likely to be severely restricted. In MSM-IP, management of a shared address pool is a long term issue that needs to be addressed. Multicast routers need to be deployed in every subnet that wishes to support communication with mobile hosts. Potentially, TCP in many flavors of operating systems needs to be changed to accommodate multicast end-point addresses. However, note that none of these problems are unique to mobility support in multicasting architectures. All of the above issues of scale need to be resolved even for standard multicasting services in the Internet. We believe that once this infrastructure is in place, support-

ing mobility in such an infrastructure is a trivial issue. The location management hierarchy emulates the approach in Mobile IP and is highly scalable. The multicasting approach for setting up and tearing down routes with receiver initiated joins/prunes and soft state is highly scalable, and can support seamless mobility effectively. Clearly, the benefits of the MSM-IP approach over Mobile IP and its variants are in superior performance, support for seamless mobility, and support for the requirements of future Integrated Services networks.

6 Related Work

In this section we compare the MSM-IP approach to other contemporary approaches to providing mobile host support. A common feature of existing approaches for providing network level support for host mobility is the use of a two tier addressing architecture, consisting of a home address and a care-of address. As described in [3], the key functions of any architecture for mobility support include address translation (mapping of the home address to the care-of address), packet forwarding (the tunneling of packets destined to the home address to the location of the care-of address), and location management (discovery and update of the mobile host's location - typically transparent to the correspondent host). Based on these functionalities, [3] compares the following approaches: Columbia [12], Sony [22], Loose Source Routing [16], Mobile IP [19], and IPv6 [17]. Other approaches to support host mobility include Mosquitonet [5] and Daedalus [20]. We briefly described the Mobile IP approach in Section 2 and compared it with MSM-IP.

The approach in Mosquitonet is very similar to the instance of Mobile IP in which a stripped down version of the foreign agent is co-located with the mobile host. Essentially, a mobile host may acquire a care-of address using a protocol such as DHCP, and then encapsulate or decapsulate packets itself, thereby performing the packet forwarding functionality of the foreign agent.

The approach in Daedalus is to multicast packets from the home agent to the cluster of foreign agents in the neighborhood of the mobile host. As the mobile host moves, this cluster changes. This approach is similar to the MSM-IP approach in its goal of using multicasting to reduce handoff latency. However, it still retains the two-level addressing hierarchy of Mobile IP, and thus incurs the same benefits and overheads (e.g. triangular routing) of Mobile IP. In this approach, multicasting is provided only as a convenience for the sole purpose of reducing packet loss upon handoff. In MSM-IP, multicasting is the *only* form of routing to a mobile host. Our whole point is that once an effective IP multicasting infrastructure is in place, it can be used as-is for supporting mobility. Thus, MSM-IP differs from Daedalus from both the theoretical perspective and the practical issues it needs to address.

The Columbia approach [12] was among the pioneering efforts to support mobility in the Internet. MSM-IP and the Columbia approach share several common goals and features; in fact, MSM-IP can be viewed as the multicasting analogue of the Columbia approach. In

the Columbia approach, mobile hosts belong to a virtual mobile network with a distinct network id. A collection of dedicated mobile support routers (MSRs) are used to provide packet forwarding and location management. MSRs communicate with each other by means of tunnels. Mobile host locations are updated and propagated by means of a distributed directory protocol. The key distinction between the Columbia approach and the MSM-IP approach is the use of multicasting in order to reduce handoff packet loss, the ability to advance register and perform resource reservation, and the use of the existing IP multicasting infrastructure to accomplish host mobility.

In the area of advance reservation, a number of recent efforts have proposed the use of multicasting for advance registration and reservation in the neighborhood of the mobile host [14]. The determination of the neighborhood ranges from simply reserving in every neighboring cell, to using sophisticated probabilistic techniques to determine 'shadow clusters', to profiling user movement in order to determine a subset of 'next-predicted-cells'. Again, as in the case of Daedalus, multicasting has been suggested as an enhancement to unicast routing rather than as a replacement of the basic routing scheme itself.

In terms of complementary work in the area of resource reservation, RSVP [24] provides an effective and dynamic resource reservation scheme for multicast communication. Since one of the main goals of RSVP is to provide effective resource reservation mechanisms in the presence of dynamic host joins and prunes, RSVP works well as a resource reservation mechanism in a mobile computing environment. Also, since another main goal of RSVP is to be independent of the underlying multicast routing algorithm, it works well with the MSM-IP architecture. In effect, we can use RSVP without modifications as the resource reservation mechanism for mobile flows. Given that arbitration of scarce and shared wireless networking resources will be a major concern of future Integrated Services mobile computing environments, the seamless interoperation of RSVP with MSM-IP is extremely beneficial.

7 Conclusions

In this paper, we have proposed a new architecture for supporting host mobility in the Internet using IP multicasting as the sole mechanism for routing packets to the mobile hosts. Our approach, called MSM-IP, assigns a unique location independent (multicast) address to a mobile host. Consequently, we eliminate the problem of mapping location independent to location dependent addresses for mobile hosts, which is the key focus of almost every contemporary approach for supporting host mobility in the Internet. Since our mobile hosts are identified by multicast addresses, packets from the correspondent host to the mobile host will be tunneled through a sequence of multicast routers in the multicast distribution tree and reach the mobile host, rather than go through a home network as in Mobile IP.

As a result of using multicasting for supporting host mobility, performing advance registration and having pack-

ets delivered to the next cell in advance of a handoff is very natural. Likewise, advance resource reservation and the use of RSVP-like mechanisms for resource reservation are also very natural in our approach. IP multicasting seeks to support efficient joins and prunes of receivers in a multicast group. This helps us perform handoffs in an efficient and graceful manner.

We have implemented a testbed and extensively tested its performance for handoffs and interface changes within and across networks, with hot and cold switches, with TCP, UDP and a variety of communication-intensive applications. Our results indicate that MSM-IP performs very well for all the above tests, and is definitely a viable alternative to contemporary approaches for supporting mobility from the perspective of performance.

However, a number of factors currently prevent the deployment of the IP multicasting infrastructure as-is to support host mobility. Location management is an issue that needs to be resolved effectively - we provide a solution to this problem. Interacting and supporting protocols such as ARP, IGMP, ICMP and TCP (at least for the mobile host) either require changes or modifications to the IP multicasting infrastructure in order to work seamlessly with the routing protocol. We have identified a list of issues that need to be solved, and provided preliminary solutions to most of these issues. However, future work will involve a more detailed study of these issues and attempt to provide better and more scalable solutions to the problems that are currently in the way of allowing an MSM-IP architecture to be widely used for supporting host mobility in the Internet.

References

- [1] H. Balakrishnan, V.N. Padmanabhan, S. Seshan, and R.H. Katz. *A Comparison of Mechanisms for Improving TCP Performance over Wireless Links*, **Proceedings of ACM SIGCOMM**, 1996.
- [2] A. Ballardie, P. Francis, and J. Crowcroft. *Core Based Trees*, **Proceedings of ACM SIGCOMM**, August 1993.
- [3] P. Bhagwat, C. Perkins, and Satish K. Tripathi. *Network Layer Mobility : an Architecture and Survey*, **IEEE Personal Communication** v.3 no. 3, June 1996.
- [4] R. Caceres and V. Padmanabhan. *Fast and Scalable Handoffs for Wireless Internetworks*, **Proceedings of ACM Mobicom**, November 1996.
- [5] S. Cheshire and M. Baker. *Internet Mobility 4x4*, **Proceedings of ACM SIGCOMM**, 1996.
- [6] Y. Dalal and R. Metcalfe. *Reverse Path Forwarding of Broadcast Packets*, **Communications of the ACM**, v. 21, no. 12, December 1978.
- [7] S. Deering and J. Hinden. *Internet Protocol, Version 6 (IPv6) Specification*, **RFC 1883**, December 1995.
- [8] S. Deering, D. Estrin, et al. *Protocol Independent Multicast - Sparse Mode (PIM-SM): Motivation and Architecture*, **draft-ietf-idmr-pim-arch-04.ps**, October 1996.
- [9] S. Deering and D. Cheriton. *Multicast Routing in Datagram Internetworks and Extended LANs*, **ACM Transactions on Computer Systems**, May 1990.
- [10] R. Droms. *Dynamic Host Configuration Protocol*, **RFC 1541**, October 1993.
- [11] M. Stangel, V. Gupta, and V. Bharghavan. *TCP Connection Management for Lossy Mobile Handoff*, In preparation.
- [12] J. Ioannidis, D. Duchamp, and G.Q. Maguire Jr.. *IP-based Protocols for Mobile Internetworking*, **Proceedings of ACM SIGCOMM**, September 1991.
- [13] J.C. Lin and S. Paul. *RMTP: A Reliable Multicast Transport Protocol*, **newblockProceedings of IEEE Infocom**, April 1996.
- [14] S. Lu and V. Bharghavan. *Adaptive Resource Management Algorithms in Mobile Computing Environments*, **Proceedings of ACM SIGCOMM**, August 1996.
- [15] J. Moy. *Multicast Extension to OSPF*, **Internet Draft**, September 1992.
- [16] C. Perkins and P. Bhagwat. *A Mobile Networking System based on Internet Protocol*, **IEEE Personal Communications Magazine**, v.1, no. 1, February 1994.
- [17] C. Perkins and D. Johnson. *Mobility Support in IPv6*, **draft-ietf-mobileip-ipv6-00.txt**, January 1996.
- [18] C. Perkins and D. Johnson. *Mobility Support in IPv6*, **Proceedings of ACM Mobicom**, November 1996.
- [19] C. Perkins, Editor. *IP Mobility Support*, **Network Working Group, RFC 2002**, October 1996.
- [20] S. Seshan, H. Balakrishnan and R. H. Katz. *Handoffs in Cellular Wireless Networks: The Daedalus Implementation and Experience*, **ACM/Baltzer Journal on Wireless Networks**, 1995.
- [21] G. Wright and W. Stevens. *TCP/IP Illustrated, Volume 2*, **Addison Wesley**, 1995.
- [22] F. Teraoka, Y. Yokote, and M. Tokoro. *A Network Architecture Providing Host Migration Transparency*, **Proceedings of ACM SIGCOMM**, 1991.
- [23] D. Waitzman, S. Deering, and C. Partridge. *Distance Vector Multicast Routing Protocol*, **RFC 1075**, November 1988.
- [24] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala. *RSVP: A new Resource reSerVation Protocol*, **IEEE Network**, September 1993.