

Agent Based Seamless IP Multicast Receiver Handover

Jiang Wu and Gerald Q. Maguire Jr.

Department of Teleinformatics, Royal Institute of Technology (KTH), Sweden
{jiang, maguire}@it.kth.se

Key words: The fast development and deployment of the Internet and mobile

Abstract: The fast development and deployment of the Internet and mobile communication has boosted personal computing and communication. The ever-increasing number of personal devices accessing the Internet demands IP mobility support ubiquitously. However, the current TCP/IP protocol suite is not able to fully support IP mobility, especially when real-time applications are concerned. This paper discusses the IP mobility support in general and IP multicast receiver mobility support in particular. The current Internet is capable of providing basic mobility support for IP multicast when mobile hosts act as IP multicast receivers. However, the worst case handover latency introduced by IP layer handover is usually unacceptable to the real-time applications running in the mobile receivers. In order to overcome this, we propose a Mobility Support Agent (MSA) architecture and a set of protocols to help achieve seamless IP multicast mobility. Our testbed runs IGMPv2 and PIM-SM. Measurements have shown that by using the MSA architecture, the handover latency introduced by the current multicast membership management protocols and multicast routing protocols is nearly negligible.

1. INTRODUCTION

Today's mobile communication systems are primarily designed to provide cost effective wide area coverage for a rather limited number of users with moderate bandwidth demands (voice + low rate data). The users of tomorrow will expect much more than today's technology and infrastructure can offer. Research in [1] shows that pico-cells (very small radio cells) will be the dominant wireless access networks with high

bandwidth capacity and low error rate. Both data applications and real-time applications will run over those pico-cells.

The new pico-cell based infrastructure will have IP running directly over the wireless networks. Figure 1 shows such an example network with link layer wireless networks connected via IP routers. With the help from IP, mobile communication can be maintained despite moves among heterogeneous wireless networks. Mobility support will not be limited to handover between the same link layer wireless network, rather, IP layer handover (vertical handover [2]) can hide changing of accessing network from applications as a mobile host roams.

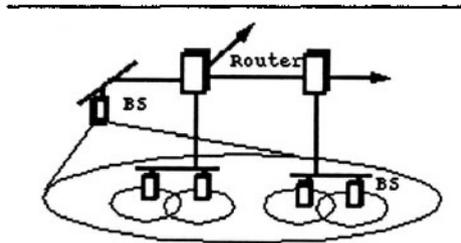


Figure 1. A pico-cell based wireless network infrastructure. Handover may occur in link layer or IP layer.

The key challenge for IP layer mobility support is from the TCP/IP protocol suite itself, due to the lack of handover support mechanisms. IP technology historically was developed and optimized for the fixed networks, though since its inception satellite networks were part of the network. Besides location dependent routing problems for IP unicast routing, Both IP unicast and IP multicast introduce significant amount of traffic loss during a handover (we call it "handover latency" in the following text), which makes it almost impossible to support real-time applications.

The unicast IP routing schemes use only the network part of the IP address to make routing decision for each packet, hence IP routing achieves very good scalability. However, those routing schemes implicitly assume that Internet hosts must attach to the network that has the same network prefix (the home network). Otherwise, it is impossible for the mobile hosts' home router to deliver the packets successfully to it.

The location dependent nature of unicast IP routing constrains the mobility of a host. Mobile IP [3] solves this problem by introducing an extension to the original unicast routing mechanism. When a mobile host visits a foreign network, the packets destined to that mobile host are intercepted by the Home Agent in the mobile host's home network. A tunnel is built between the Home Agent and a corresponding entity (either a

Foreign Agent in the foreign network or the mobile host itself) to deliver the packets to the mobile host. Mobile IP makes IP unicast handover possible. However, by using Mobile IP alone, it still introduces handover latency in the order of several seconds [4].

Unlike the unicast routing mechanism, IP multicast routing is location independent. An IP multicast address is a logical address representing a set of participating hosts in a multicast session. While in IP unicast, the address is used to indicate the point of attachment of a host to the Internet. A host simply propagates its interest in participating in a multicast session to the nearest upstream multicast router (not necessarily the one in its home network), in turn the intermediate multicast routers establish a new branch of the multicast tree which delivers the traffic to the host.

Handover latency for IP multicast varies from zero up to minutes. This paper examines the problem of inefficient IP multicast handover in the current Internet, and proposes an agent based architecture to help eliminate potential handover latency. A set of protocols were designed for this new architecture. Within our own testbed, simple implementation of the protocols and measurement of handover performance were done.

2. IP MULTICAST

IP multicast aims to provide a routing architecture to support group communication in a dynamic and efficient way. IP multicast consists of two parts: group membership management and multicast routing. The group management protocols work mainly in the leaf networks where the local membership information is collected, which is later used by the routing protocols to build up the multicast routing tree. The group management protocols should be able to quickly reflect the dynamic changing nature of membership in the leaf networks, while at the same time should consume as little as possible the network resources.

Different leaf networks usually have different group management mechanisms. One of the most popular protocols among them is the Internet Group Management Protocol version 2 (IGMP) [5], it is widely used in the LANs. Although other group management protocols are interesting, this paper focuses on IGMP. IGMP is used by the Querier Router to poll the presence of the multicast members in its directly attached LAN. The host may report its interest in a multicast group after receiving an IGMP Query from the Querier Router. The Querier Router consequently constructs a membership table indicating the presence of the members.

It is only important for the multicast routers to know whether a multicast group is active or not in the LAN (not the identities of the explicit hosts).

Thus the multicast group membership table in the router is not complicated. Report suppression is used to avoid report flooding when multiple participants report to the same group after an IGMP Query is received. Before sending an IGMP Report, a host backs off a random time interval. The host with the smallest timer will send out the report to that group, suppressing the other reports which are waiting for their timer expiration.

Multicast routing protocols can be generally divided into two categories: dense mode multicast protocols and sparse mode multicast protocols. As shown in Figure 2, the dense mode protocols and sparse mode protocols can be further divided into different groups. The resulting protocols: DVMRP [6], MOSPF [7], PIM-DM [11], PIM-SM [8] and CBT [9] are the ones currently getting most attention in the Internet. The inter-domain multicast protocol BGMP [10] is also under development nowadays.

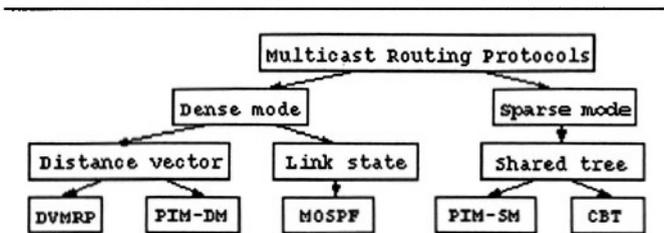


Figure 2. The multicast routing protocols

The dense mode and sparse mode routing protocols are different in many ways. Dense mode protocols are usually used in heavily populated member environments, while the sparse mode protocols are used in the situation when the group members are scattered over the Internet. Dense mode protocols build up the multicast tree usually by a data driven approach (MOSPF is an exception), while the sparse mode protocols requires an explicit join from the hosts.

When the first member joins a multicast group in the LAN, the Designated Routers will be informed by receiving an unsolicited IGMP Membership Report from that host. Immediately the Designated Routers will try to establish the multicast tree branch for the group. This is the mechanism used by IGMP and multicast routing protocols to reduce the joining latency for the first member.

The role of a multicast sender is quite different from the role of a multicast receiver. A host is not required to join in the multicast group when it acts only as a multicast sender. The host can directly send to the multicast group if there are multicast routers available in its directly attached LAN. Also the multicast routing protocols treat the sender differently from the receiver. This paper does not discuss sender mobility.

2.1 Handover for IP multicast receivers

Handover performance for IP multicast is determined by IGMP and multicast routing protocols. In this paper we use PIM-SM to illustrate the behavior of a routing protocol. Although PIM-SM is quite different from the other multicast routing protocols, they have similar behavior in handover.

Because IP multicast uses only the group address for membership reporting and routing, it doesn't require the hosts' identification in the network layer. The minimum requirement for a host to join a multicast group is that the host is attached to the same LAN as the corresponding Querier Router. A handover does not require new IP address as Mobile IP does. In principle, the current Internet is capable of handling IP multicast handover without any modification of the running protocols. However, the problem of handover for IP multicast comes from the potential heavy traffic loss. The handover latency in such a multicast handover is sometimes unacceptably long, i.e., 30 seconds.

After having established a steady multicast tree, IGMP and PIM-SM sends periodic protocol messages to maintain state in the host and multicast router. In IGMP, the Querier Router sends IGMP Query messages every IGMP Query Interval to probe the membership information in the LAN. In PIM-SM, routers send join/prune message upstream every Join/Prune-Period to refresh the routing states in the router. High granularity of the queries and join/prunes can make the protocols react quickly to the changing of the membership situation. However, the queries and join/prunes are sent only at a moderate frequency because otherwise it will lead to high protocol overhead. In order to respond quickly to a member's joining/leaving, both IGMP and PIM-SM introduce event triggered protocol messages.

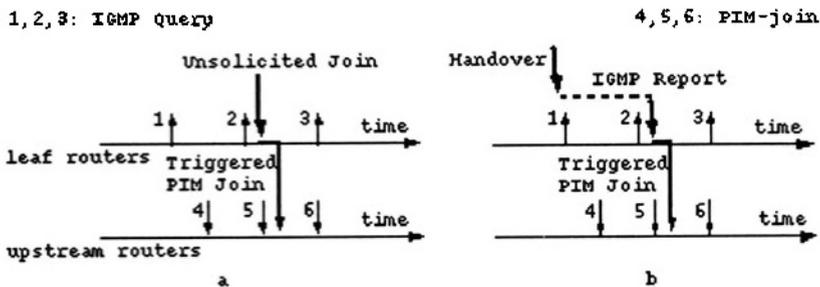


Figure 3. (a) An application first joins a multicast group. (b) An application joins after handover, in this case no member exists in the new IP network. A leaf router can be either Querier Router or Designated Router.

As can be seen in Figure 3a, When the application in the host wants to join a multicast group, an unsolicited IGMP Membership Report is sent. A routing entry for that multicast group, if not already existing, will be added in the PIM-SM Designated Router immediately. This new routing entry triggers the Designated Router to lookup the corresponding Rendezvous Point and send a join message to the upstream router. Thus when an application in the host first joins a multicast group, the joining delay is made very low by using the event triggered protocol messages.

After handover, a mobile host can continue to receive multicast traffic in the new IP network if there are other active members of the group in that network. However, it most probably will be the opposite case (the worst case) in a pico-cellular environment where the cell size is too small to accommodate many mobile hosts. Hence the probability of their already being a member of the multicast group in the cell is small. In order to receive the traffic quickly, a join message should be sent immediately regardless of the IGMP Queries. However, neither the multicast applications nor IGMP has the mechanism to detect the handover and trigger the unsolicited IGMP Membership Report.

Figure 3b shows that after a handover, the mobile host must wait for an IGMP Query in the new IP network. After receiving the IGMP Query, it sets up a random back off timer. An IGMP Membership Report will be sent when the timer expires.

The additional handover latency introduced by multicast routing protocols are not as severe as that due to IGMP, because the Designated Routers always use triggered join to the upstream router to establish the multicast tree branch. When a Designated Router discovers a new group is present in the LAN by IGMP, it will send immediately a join/prune message to the upstream router. It takes time to establish the tree branch. However, the tree establishment latency is low in most cases, i.e., when the distance between the Designated Router and the upstream router is short. In some cases, e.g., cross-ocean links, the propagation time to the upstream router is not negligible. But even in those cases it is less than the latency introduced by IGMP.

2.2 THE MSA ARCHITECTURE AND PRE-REGISTRATION

In this paper we propose a Mobility Support Agent (MSA) architecture to help achieve seamless IP mobility. The MSA architecture is designed to eliminate the potential traffic loss caused by handover. It is transparent to the applications running in the mobile hosts and the related protocols running in the network.

A set of protocols are designed for the MSA architecture. The MSA protocols run only between the MSAs and the mobile hosts. Although in this paper we only discuss the protocols related to IP multicast mobility, the MSA architecture can be used as general framework for IP mobility support. The architecture is shown in Figure 4a.

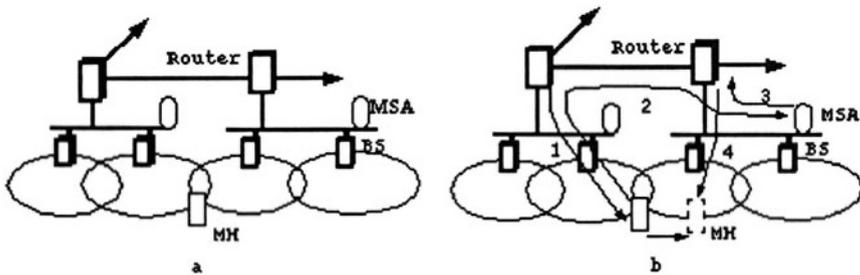


Figure 4. (a) The MSA architecture. (b) The procedure for the Pre-registration protocol. There is a MSA node located in each IP network which supports seamless IP mobility.

The key protocols for supporting IP multicast mobility are Agent Discovery Protocol and Pre-registration protocol. MSAs advertise their presence and services via the Agent Discovery protocol. The MSAs form their own multicast group to which they are both senders and receivers. Periodic protocol message exchanges update the MSAs' current information about each other. The advertised information may include the MSAs' IP addresses, the services they provide, etc. The advertisement between the MSAs and mobile hosts is exchanged locally via ICMP Router Advertisement Extensions.

The Pre-registration protocol is designed for supporting seamless IP multicast. A Pre-registration procedure is illustrated in Figure 4b.

1. A mobile host receives multicast traffic from the Internet in its currently attached network.
2. As soon as the mobile host decides to perform a handover, it sends a Pre-registration message using UDP to the neighboring network's MSA. This Pre-registration packet is routed according to the unicast routing protocols.
3. Enclosed in this packet are the multicast groups the mobile host is participating in. Once having received this packet, the MSA sends immediately an IGMP Membership Report to the Designated Router, which in turn triggers a join message to the upstream routers to establish the multicast tree.
4. When the mobile host arrives at the new network, multicast traffic is already available for it.

One advantage IP multicast has over IP unicast in handover is that the mobile host knows in advance the multicast addresses it will use in the new network are those same addresses as it used before the handover. The multicast traffic grafting is from the closest upstream router, while in Mobile IP, traffic has to be redirected by the fixed home agent.

The Pre-registration protocol can be extended to support seamless Mobile IP support. The additional functions that should be added in the MSA are: pre-negotiate the co-located care-of address and pre-register with the home agent for traffic redirecting, etc.

3. TESTBED AND MEASUREMENT

Figure 5 shows the testbed we use for testing the MSA architecture. The Pre-registration protocol is implemented in the testbed. Routers 1, 2 and 3 all support PIM-SM and all the nodes in the testbed support IGMPv2. In the testbed (as shown in 5a), A multicast sender "Msender" sends packets at a constant rate to a certain multicast group. The mobile host first joins the multicast group in network 2. At a certain moment the mobile multicast receiver "MH" handovers to network 3. Our measurement will compare the worst case handover performance with/without using the MSA in network 3.

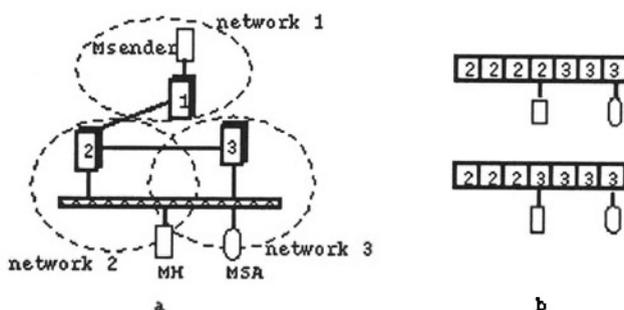


Figure 5. (a) The MSA testbed. (b) Handover emulation using Ethernet.

We use 10BaseT Ethernet in our testbed, each Ethernet segment is a separate IP network. The Ethernet segments emulate pico-cells. Handover is emulated by moving the mobile host from one Ethernet segment to another. The HP hub we use has a control module that can dynamically allocate the port to different Ethernet segments. As shown in Figure 5b, the port the mobile host is attaching to can be switched from network 2 to network 3 by using a SNMP control message.

3.1 Handover latency in handovers without pre-registration

Figure 6 illustrates the handover procedure for IP multicast, when there is no members in the new network. The handover latency (T_{ho}) consists of physical movement time when non-overlapped cells are used (T_{mv}), IGMP Query waiting time (T_{qw}), IGMP Report back off time (T_{rb}) and multicast tree establishing time (T_{te}).

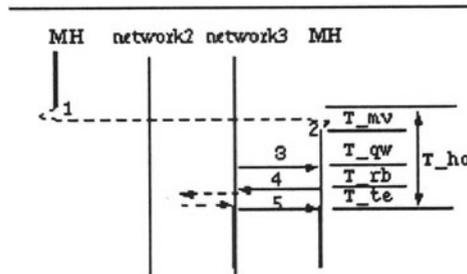


Figure 6. Mobile host handovers in normal IP multicast. The dash arrows between 4 and 5 indicate the tree establishment process. Thick line indicates multicast traffic.

The corresponding events during a handover are:

1. The mobile host leaves network 2.
2. The mobile host enters network 3.
3. The first IGMP Query is received from the Querier Router.
4. An IGMP Membership Report is sent after the back off timer expires.
5. The first multicast packet received in network 3.

In our test, the mobile host was in a stable situation receiving the multicast packets before the handover. Because handover is emulated by switching the Ethernet hub port by using SNMP, T_{mv} is very small (around several milliseconds) so that it can be neglected. Our focus will be on T_{qw} , T_{rb} and T_{te} .

In this measurement, the Mender sends 50-Byte packets once per second. Because the handover latency in a normal handover is in the order of seconds, this packet rate is high enough to capture the handover latency characteristics. "Tcpdump" is used in network 3 to capture the events during a handover.

Router 2 and router 3 send all IGMP queries every 60 seconds (this is done using Cisco 7000 series routers). The mobile host may handover at any moment during the IGMP Query Interval, so the waiting time for the next IGMP Query can be anywhere between 0 and 60 seconds. When the mobile

host receives the query, it sets a random timer between 0 and 10 seconds, for sending IGMP Membership Report.

Figure 7 shows the measurement results. We observe that T_{qw} and T_{rb} contribute most to the overall handover latency. The maximal T_{qw} in our test is 51 seconds, and the average is around 27.2 seconds. The average T_{rb} is around 5.2 seconds. Due to IGMP, a mobile host needs to wait for 32.4 seconds on average to continue to receive the multicast traffic after a handover.

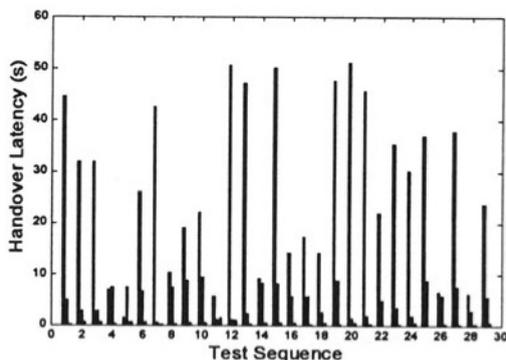


Figure 7. Measurement results of handover latency. The bars represent T_{qw} , T_{rb} and T_{te} from the left to right order for each group data, T_{mv} is negligible and not shown in this figure.

T_{te} is considerably smaller when compared to T_{qw} and T_{rb} , with an average of 0.5 second. In the result T_{te} is calculated by comparing the time difference between step 5 and step 4. The real T_{te} could be even smaller if the Msender sends out packets at higher rate than in these measurements. On the other hand, the small T_{te} results from the short distance between the upstream router (router 2) and the leaf network (network 3). It will be larger when the upstream router is far away or the network is highly loaded. However, we expect that T_{te} will not have the same significance in overall IP multicast handover latency.

We can see from the measurements that because of the large handover latency during an IP multicast handover, the current Internet simply can not support any of the applications that require real-time communication.

3.2 Handover latency in handovers with pre-registration

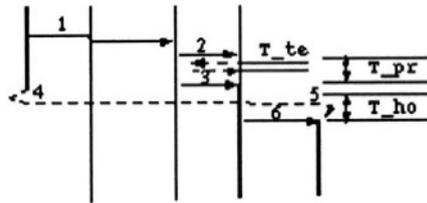


Figure 8. Mobile host handovers with pre-registration.

Pre-registration eliminates the need for waiting for T_{qw} and T_{rb} by asking the MSA to send an unsolicited IGMP Membership Report in advance of the handover. In consequence, T_{te} is also made to happen before the mobile host moves to the new network. The detailed handover procedure with pre-registration is illustrated in Figure 8:

1. The mobile host sends a pre-registration message to the MSA in network 3.
2. The MSA joins the multicast groups as a proxy for the mobile host.
3. The first multicast packet arrives in network 2.
4. The mobile host leaves network 2.
5. The mobile host enters network 3.
6. The mobile host starts to receive multicast traffic in network 3.

As can be seen, multicast traffic arrives at network 3 before the arrival of the mobile host. Thus the latency the mobile host suffers from mainly comes from the physical movement between the non-overlapped networks and other possible disturbances by the handover, e.g., network overload caused by additional multicast traffic. In some cases when the upstream router in the new network is closer to the Rendezvous Point, receiving duplicated packets is possible, but this is not a big problem for IP multicast receiving.

Determining the handover latency in the measurements, by sending a packet once per second and using "Tcpdump" in network 3 is not feasible. Hence, we conducted a second test, in this test the Msender sends 50-Byte packets every millisecond, producing a 400kbps stream traffic. The Msender puts a sequence number in each packet when sending. The mobile host receiving these packets will record the sequence number and calculate how many packets are missed during the handover, the packet loss corresponds to the T_{ho} in Figure 8. In this way we can measure the handover latency with resolution of 1 millisecond.

Measurement results are shown in Figure 9. From these results we observe that 34 percent of the handovers suffer no traffic loss. The handover

latency is up to 11 milliseconds, the average is 5.1 milliseconds. This result is significantly less than the result in the previous measurements (32.4 seconds) when no MSAs were used. We believe that by using the MSA architecture and Pre-registration, real-time applications can operate despite making a handover.

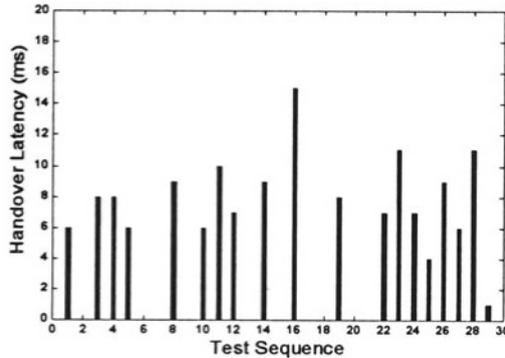


Figure 9. Measurement results of the handover latency with pre-registration.

In this scenario, physical movement is the major factor contributing to the overall handover latency. In our testbed, even though the “cells” are not overlapped, the physical movement time is almost negligible because of the mechanism we use. In other cases when the distance between the cells are large, the physical movement will cause noticeable disturbance to the applications. However, good network planning can solve this problem. Another possible solution is to add features to the Pre-registration protocol to ask the MSA to buffer the incoming multicast packets which can be later delivered to the mobile host after it has handedover to the new network.

4. CONCLUSIONS AND FUTURE WORK

Expecting a pico-cell based high capacity wireless infrastructure in the future, we find that handover will frequently happen at the IP layer. By introducing the MSA (Mobility Support Agent) architecture to the existing Internet routing architecture, seamless mobility support can be achieved. In this paper, we show that no additional processing or elements are needed for supporting mobility for multicast receivers. However, in order to make low latency handover, MSAs are used. Tests and measurements have shown that the worst case handover latency drops to a level that is sufficient to support even quite demanding real-time communication during a handover.

We are going to explore the handover procedure in more detail, better implementation and more measurement will be done in the near future. Real-time video will be used in our testbed to test the feasibility of mobility support for the applications.

In order to provide better performance and to be flexible in pre-registration, movement detection and prediction need to be studied more. How to reduce the departure latency in the previously visited network is also an interesting problem. Finally, the handover behavior as a multicast sender will be studied to make the multicast mobility research complete.

REFERENCES

- [1] M. Flament, F Gessler, F Lagergren, O Queseth, R Stridh, M Unbehaun, J Wu, J Zander, "An Approach to 4th Generation Wireless Infrastructures - Scenarios and Key Research Issues", VTC 99, May 1999.
- [2] Mark Stemm, Randy H.Katz, "Vertical handoffs in wireless overlay networks", ACM Mobile Networks and Applications", Volume. 3, 1998, pp. 335-350.
- [3] C. Perkins, "IP mobility support", RFC2002, October 1996.
- [4] Jon-Olov Vatn and Gerald Q. Maguire Jr. "The effect of using co-located care-of addresses on macro handover latency", Fourteenth Nordic Tele-traffic Seminar, August 1998.
- [5] W. Fenner, "Internet group management protocol, version 2", RFC2236, November 1997.
- [6] D. Waitzman, C. Partridge, S. Deering, "Distance vector multicast routing protocol", RFC1075, November 1988.
- [7] J. Moy, "Multicast Extensions to OSPF", RFC1584, March 1994.
- [8] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma and L. Wei, "Protocol independent multicast-sparse mode (PIM-SM): protocol specification", RFC2362, June 1998.
- [9] A. Ballardie, "Core based trees multicast routing, protocol specification", RFC2189, September 1997.
- [10] D. Thaler, U. Michigan, D. Estrin, D. Meyer, U. Oregon, "Border Gateway Multicast Protocol (BGMP): Protocol Specification", Internet Draft, March 2000.
- [11] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, A. Helmy, and L. Wei, "Protocol independent multicast version 2, dense mode specification", Internet Draft, May 1997.