Interoperation of Mobile IPv6 and Protocol Independent Multicast Dense Mode

Christian Bettstetter
Technische Universität München
Inst. of Communication Networks
Munich, Germany
Christian.Bettstetter@ei.tum.de

Anton Riedl
Technische Universität München
Inst. of Communication Networks
Munich, Germany
Anton.Riedl@ei.tum.de

Gerhard Geßler
IABG
Ottobrunn, Germany
Gessler@iabg.de

Abstract

This paper discusses different approaches of providing multicast traffic for mobile hosts. Mobile IPv6 is used for mobility support. The network employs Protocol Independent Multicast Dense Mode (PIM–DM) for multicast routing and Multicast Listener Discovery (MLD) to collect multicast group membership information. We identify and analyze interoperation problems concerning membership control for mobile hosts and efficient multicast packet transfer from/to mobile hosts. We discuss four multicast delivery mechanisms for mobile senders and receivers, and compare them using criteria such as join delay, routing optimality, protocol overhead, network bandwidth, and system load. In particular, we suggest timer optimizations for MLD to support highly mobile receivers.

1 Introduction

Demand for multimedia group communication, audio and video streaming, and collaborative engineering in the Internet is rapidly increasing. These applications are one–to–many or many–to–many, where one or multiple sources are sending to multiple receivers. IP multicast efficiently supports this type of transmission. Instead of sending an individual copy of a datagram to each receiver, a multicast source sends out a single copy addressed to a multicast group of receivers, which explicitly want to receive the information. The network routers are running a multicast routing protocol to distribute the datagrams to each receiver. *Protocol Independent Multicast Dense Mode (PIM–DM)* is such a multicast routing protocol and is presently being standardized by the Internet Engineering Task Force (IETF) [7]. It builds a source–rooted distribution tree for multicast traffic. Moreover, it does not rely on any specific protocol to discover the network topology or to collect information about local group memberships.

The increasing demand for mobility in the Internet has created the need for a routing protocol that allows a host to roam in the network. *Mobile IP* is a solution that enables an IP–host to leave its home link while transparently maintaining all of its present connections and remaining reachable to the rest of the Internet. *Mobile IPv4* has been standardized by the IETF in [13]. *Mobile IPv6* is currently an Internet draft [12].

This paper examines the problem of providing IP multicast to mobile hosts using Mobile IPv6 in a network running PIM–DM as multicast routing protocol. Possible approaches for interoperation are presented, and the assets and drawbacks are discussed. In Section 2, we review the fundamental functionality of Mobile IPv6. In Section 3, we show how PIM–DM works and how local group members are discovered with MLD (the Multicast Listener Discovery protocol). Section 4 is the main part of this paper. It discusses the problems and different approaches of providing multicast for mobile hosts. We examine different scenarios where we consider both mobile multicast senders and receivers. The approaches are evaluated qualitatively using the criteria join delay, routing optimality, system load, protocol overhead, additional network bandwidth, and the need to modify existing drafts. Finally, Section 5 concludes the paper and gives an outlook on further work.

Earlier work on the topic "IP multicast for mobile hosts" considers IPv4 and respective mobility support extensions. In [2], Acharya and Badrinath bring up the problem of providing multicast in a network with mobile hosts. In [3], they propose an IP multicast extension, which helps to deliver multicast messages to and from mobile hosts. Their approach is based on the "Columbia Mobile*IP" protocol [11] and mainly focuses on Distance Vector Multicast Routing (DVMRP). More recent work, such as by Chikarmane *et al.* [5] and by Xylomenos *et al.* [15], considers Mobile IPv4 for mobility support. In [1], Acharya and Badrinath give a more general guideline for structuring protocols for delivering multicast messages in a mobile environment. Our paper differs from these publications in two issues. First, we focus on the multicast protocol PIM–DM and investigate PIM–DM specific interoperation issues. Second, we use Mobile IPv6 for mobility support. Providing multicast to/from mobile IPv6 hosts is in various aspects different than with mobile IPv4 hosts.[1] For example, [5] describes the tunnel convergence problem, which occurs when multiple Mobile IP tunnels end at the same foreign agent. This problem is not relevant in Mobile IPv6, since it does not employ foreign agents.

2 Mobile IPv6

Today's version of the Internet Protocol does not support any mobility of hosts. This is because an IP address identifies the link (subnet) on which the host resides. If a host moves to a different link without changing its IP address, there is no information in its IP address about the new point of attachment. Existing routing protocols are therefore not able to deliver datagrams to the mobile host correctly, but always route them to its home link.

The purpose of Mobile IPv6 [12] is to enable a mobile host to change its point of attachment to the Internet while still maintaining transport–layer connectivity. The basic functionality is as follows:

A host is always identified by its *home address*, regardless of its current point of attachment to the Internet. This address is the usual IPv6 address, which has been assigned to the host on its home link. When the host is away from its home link and attaches to a foreign link, it obtains a so–called *care–of address*.[2] This is an additional temporary IPv6 address, which has the same network prefix as the visited foreign link and therefore provides information about the current location (the current network) of the mobile host. The mobile host registers its current care–of address with its *home agent* on its home link, using a BINDING UPDATE message. To do so, it sends the home agent an IPv6 packet containing a BINDING UPDATE IPv6 destination option.[3] The home agent stores the information about the current care–of address of the mobile host in its *binding cache* and acts as a proxy for the mobile host. To keep up connectivity, the mobile host must update its binding entry periodically.

As usual, packets addressed to the home address of the mobile host are routed to its home link. When the host is away from home, the home agent intercepts these packets and transparently forwards (tunnels) them to the current location of the mobile host. The tunnel is built using an IPv6 routing header or using IPv6 encapsulation [6]. This means that the home agent places the original (incoming) IPv6 packet in the payload part of a new IPv6 datagram and sends it to the mobile host, using the current care–of address of the mobile host as destination address. The mobile host decapsulates the received packet and processes it.

If a mobile host sends packets to other hosts, it will send them on a direct path, using its care–of address as IPv6 source address. It includes a HOME ADDRESS destination option to inform the correspondent host of its home address.

[1] See [12] for the differences between Mobile IPv4 and Mobile IPv6.

[2] The mobile node acquires its care–of address, for example, through stateless [14] or statefull (e.g. DHCPv6 [4]) autoconfiguration.

[3] The Mobile IPv6 draft [12] defines four new destination options for IPv6 packets: BINDING UPDATE, BINDING ACKNOWLEDGEMENT, BINDING REQUEST, and HOME ADDRESS.

3 Multicast routing with PIM–DM and MLD

3.1 Protocol Independent Multicast Dense Mode

The specification of PIM–DM [7] defines a multicast routing algorithm which is efficient for multicast groups that are densely distributed across a network. It uses a broadcast–and–prune mechanism to build a source–rooted distribution tree.

When a multicast source starts sending, PIM–DM initially assumes that all downstream hosts want to receive multicast datagrams. Thus, the datagrams are flooded to all links of the network. As the first datagram arrives at a PIM–DM router, the router checks if the datagram has arrived at the correct interface. This is the interface it usually uses to reach the multicast sender by unicast; it is denoted as *incoming interface*. The router now creates a (Source, Group) entry in a data base and stores (among other things) the incoming interface. The datagram is then forwarded over all other interfaces with attached PIM–DM routers or group members (*outgoing interfaces*). Afterward, an expiration timer for the (S, G) entry[4] is set to the so–called *data–timeout* value. This is the time after which an (S, G) state for a silent source will be deleted. Its default value is 210s. The timer will be restarted when a datagram is forwarded by the PIM router.

If a PIM–DM router receives a multicast datagram on an outgoing interface of the corresponding (S, G) entry, it will assume that there is a loop in the distribution tree and will start an assert process by sending ASSERT messages over this interface. All neighboring PIM–DM routers on this link that forward datagrams of this (S, G) pair onto the link, will also start an assert process and a single forwarder will be elected. Downstream PIM–DM routers listen to the ASSERT messages and store the elected forwarder for later PIM–DM protocol actions.

A PIM–DM router does not need to forward multicast datagrams for an (S, G) pair if it does not have any attached group members or other PIM–DM routers which need to forward the datagrams. In this case, it sends a PRUNE message over the incoming interface of this (S, G) pair. The next

[4] We abbreviate "(Source, Group) entry" as "(S, G) entry."

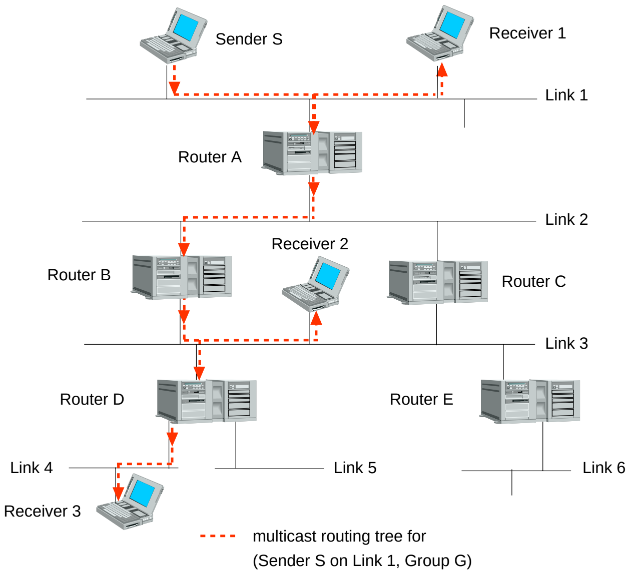

Figure 1. Multicast distribution tree

upstream PIM–DM router receives this message and prunes the link after a certain time. This time is called *Prune Delay Time* $T_{PruneDel}$. It gives other PIM–DM routers on this link the chance to send a JOIN message in case they still need multicast datagrams of this (S, G) pair.

The resulting multicast distribution tree for this (S, G) pair is now loop–free and connects the Sender S to all members (receivers) of the Multicast Group G. Figure 1 gives an example.

If a host joins a multicast group on a link that is not part of the multicast distribution tree, the local PIM–DM router sends a GRAFT message over the incoming interface of the (S, G) entry to connect the new member to the tree. The GRAFT message is processed by the next upstream PIM–DM router and reinstates forwarding state for the interface over which the GRAFT message was received.

3.2 Multicast Listener Discovery (MLD)

The idea of multicast is that routers forward multicast datagrams only onto links on which at least one member of that multicast group resides. The *Multicast Listener Discovery (MLD)* protocol [8] is used by IPv6 multicast routers to learn the existence of multicast group members on their connected links. It does so by periodically sending MULTICAST LISTENER QUERIES (for short, QUERIES) and having hosts sending a MULTICAST LISTENER REPORT (for short, REPORT) about their multicast group memberships. The collected multicast group information is then provided to the multicast routing protocol. MLD is derived from version 2 of the *Internet Group Management Protocol (IGMPv2)* [9] in IPv4. Its functionality is as follows:

On each link, one multicast router is elected to act as a querier. It periodically sends out a QUERY onto this link. The default value for the *Query Interval* T_{Query} is 125 s. When a host receives a QUERY, it initializes a *response delay timer* for each group in which it is a member on the interface where it received the QUERY. Each timer is set to a different random value, selected from the range $[0s \ldots T_{RespDel}]$, where the default value of the *Maximum Response Delay* $T_{RespDel}$ is 10 s. When a timer expires, the host responds to the QUERY by sending a REPORT to the associated multicast group address with a hop limit of 1. It will be received by the multicast routers on the local link, which add the reported group to the list of multicast group memberships on the interface on which they received the REPORT. A *group membership timer* is initialized with the *Multicast Listener Interval*, which is $T_{MLI} = 2 \cdot T_{Query} + T_{RespDel}$. Using the default values for T_{Query} and $T_{RespDel}$ results in a default value $T_{MLI} = 260$ s.

Repeated REPORTS refresh the group membership timer. If no REPORTS for a group are received within this interval, the router will assume that there are no longer any members of this multicast group on this link. Whenever a router adds or deletes a multicast group membership for a link, it notifies the multicast routing protocol to add or delete this link from the multicast distribution tree.

When a host initially joins a multicast group, it should immediately send out an unsolicited REPORT for that group, in case there are no other group members on this link. When a host leaves a multicast group, it should send a DONE message to the link–scope all–routers multicast address.

4 Interoperation of PIM–DM and Mobile IPv6

4.1 Problem statement and motivation

We now consider an IPv6 network that employs PIM–DM as multicast routing protocol and offers mobility support using Mobile IPv6. The combination of these two protocols leads to several problems, for both mobile multicast senders and receivers.

For example, consider a sender of multicast traffic which has recently moved to a new link. Then, multicast datagrams in general do not reach all members (static or mobile) of the multicast group anymore. A new source–rooted distribution tree will not be built until the mobile multicast sender obtains a new care–of address at the visited foreign link and uses it as source address for its multicast datagrams. PIM–DM will then interpret the movement of the multicast sender as a new multicast sender on the foreign link and will thus create a completely new multicast distribution tree, including the initial flooding.

Another example: A host which has joined a multicast group moves to a new link (mobile receiver). If there do not exist any other members of the same multicast group at this link, the host will not receive multicast traffic until the multicast distribution tree is extended to the new link.

4.2 Four approaches to support multicasting for mobile hosts

The IETF draft of Mobile IPv6 mentions the basic ideas how mobile hosts can receive and send multicast datagrams. To receive datagrams sent to a multicast group, a host must join this group. It may send the necessary multicast group membership information: (A) either to its local multicast router on the visited foreign link, or (B) to its home agent on the home link. In the first case, incoming multicast datagrams will be delivered locally to the mobile host. In the second case, they will be delivered to the home agent, which forwards them via a tunnel to the mobile host. Analogously, a mobile sender may send multicast datagrams: (A) directly on the foreign link, or (B) via a tunnel to its home agent.

In this section, we investigate these cases in more detail. To simplify the scenarios, we consider either mobile senders or mobile receivers. The general case that a mobile host is both sender and receiver for a specific multicast group can be derived by combining these scenarios.

Figure 1 illustrates the network we use for our investigations. The five routers act as PIM–DM routers and home agents. Router A is home agent on Link 1, Router B on Link 2, Router C on Link 3, Router D on Link 4 and 5, and Router E on Link 6. In the beginning, each mobile host is attached to its home link. Sender S is sending multicast datagrams to a multicast group to which the other hosts (Receiver 1, 2, and 3) are subscribed. The initial distribution tree for the multicast datagrams is shown.

4.2.1 Mobile receiver

Approach A: Local group membership on foreign link In order to receive multicast traffic, a mobile host joins a multicast group via its local multicast router on the visited link. It uses its care–of address as source address. Afterward, the distribution tree will be extended to this link.

Let us give an example using Figure 1. Receiver 3 moves from Link 4 to Link 6, which is a pruned link of the tree. In order to receive again multicast traffic, Receiver 3 must join the multicast group via its current local multicast router (Router E) using its local care–of address. It may either wait for the next QUERY or send unsolicited REPORTS for the multicast group it wants to subscribe. Only when Router E receives a REPORT from Receiver 3, it will graft to the multicast distribution tree and start forwarding multicast traffic onto Link 6. Then, all multicast traffic will be delivered to

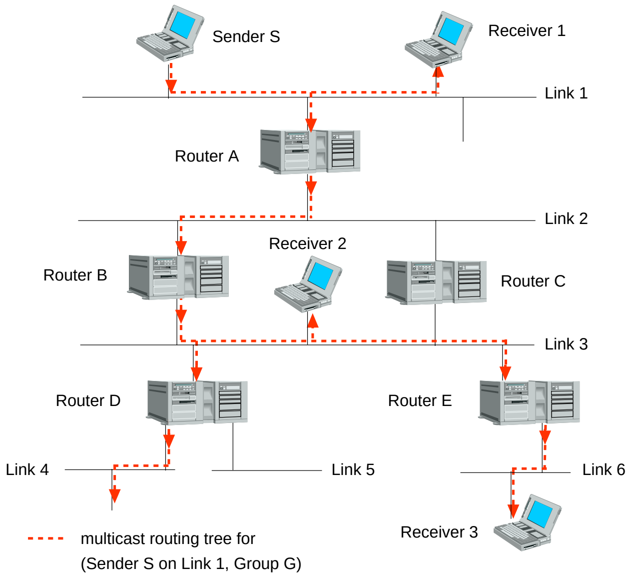

Figure 2. Mobile receiver: Group membership on foreign link

Receiver 3. The time between the attachment of the mobile receiver to the link and its ability to again receive multicast traffic is called *join delay*.

MLD detects the absence of group members (i.e. the group members have left the group or left the link) by not receiving REPORTS for this group within the Multicast Listener Interval T_{MLI} (by default 260 seconds). Therefore, Router D still "believes" that there is a group member on Link 4 and continues to forward multicast datagrams onto it. This unnecessarily consumes bandwidth. Only after expiration of the group membership timer (max. 260 seconds after Receiver 3 has left Link 4), MLD in Router D detects that all receivers on Link 4 have left and notifies PIM–DM. Router D now stops forwarding multicast datagrams onto Link 4. This delay is called *leave delay*.

Figure 2 illustrates the resulting multicast distribution tree after the link change. The MLD group timer in Router D has not expired yet.

Approach B: Group membership on home link (via tunnel from home agent to mobile host) Alternatively, a mobile host may join a multicast group on its home link via its home agent. To do so, it informs the home agent about the multicast groups it wants to subscribe. The home agent becomes group member of the requested groups on behalf of the mobile host and stores the group membership information. Datagrams addressed to this group will be routed to the home agent, which then forwards them through the tunnel to the mobile host.

Let us again consider the initial scenario in Figure 1,

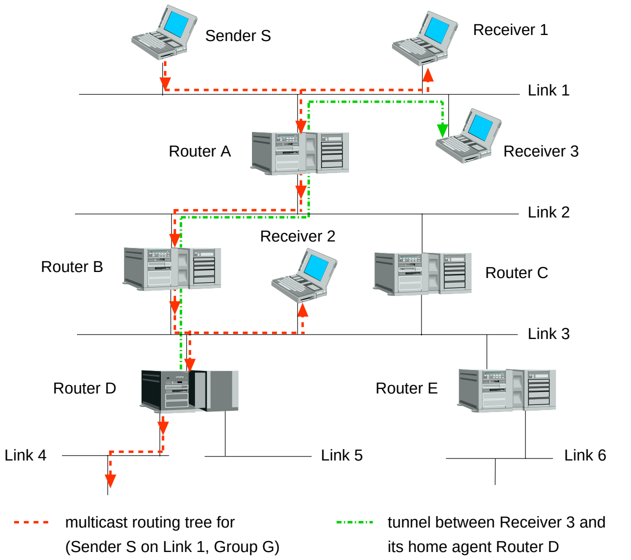

Figure 3. Mobile receiver: Group membership on home link (via tunnel from home agent to mobile host)

where all hosts reside at their home links and Sender S is sending multicast datagrams. Receiver 3 moves from Link 4 to Link 1, a link that is already part of the distribution tree. After connecting to the foreign link, it generates a care–of address and sends a BINDING UPDATE together with its group membership information to its home agent Router D. With this, a tunnel has been established. The multicast distribution tree and the tunnel are shown in Figure 3.

4.2.2 Mobile sender

Approach A: Local sending on foreign link We consider again the scenario shown in Figure 1. Now, Sender S disconnects from its home link and moves to Link 6. As soon as it recognizes that it has changed the link, it generates a care–of address and sends a BINDING UPDATE to its home agent Router A. It uses the care–of address as IPv6 source address. The first multicast datagram is interpret by the PIM–DM routers as a new multicast sender appearing on Link 6. Thus, the PIM–DM routers initially flood the datagrams to the entire network and build a new source–rooted multicast distribution tree. The old tree (Sender S on Link 1) is still stored in the PIM–DM routers. Only after expiration of the (S, G) timer, an (S, G) entry will be deleted if no further multicast datagrams are distributed over the tree. As mentioned in Section 3.1, the default value of the (S, G) timer is 210 s.

Approach B: Sending on home link (via tunnel from mobile host to home agent) Alternatively, a mobile multicast sender tunnels outgoing datagrams to its home agent. This is done by using the home address as source address of the inner datagram (the actual multicast datagram) and the care–of address as source address of the outer datagram. The home agent then decapsulates the inner datagram and forwards it on the home link. From there, the datagram is distributed to all group members over the usual multicast distribution tree.

Let us give an example. Sender S disconnects from its home link and moves to Link 6. It generates a care–of address and sends a BINDING UPDATE to its home agent Router A. With this, a tunnel for multicast datagrams is established. The tunnel and the multicast distribution tree are displayed in Figure 4.

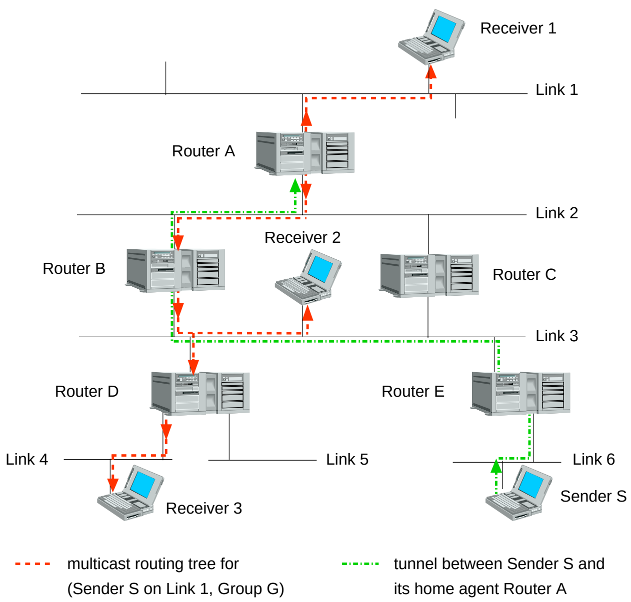

Figure 4. Mobile sender: Sending on home link (via tunnel to home agent)

4.2.3 The four combinations

Combining the different mechanisms for mobile senders and receivers yields four different approaches for a mobile host to participate in a multicast group (also see Table 1):

1. **Local group membership on foreign link**: The mobile host sends and receives multicast datagrams via the local multicast router on the visited foreign link.

2. **Bi–directional tunnel**: The mobile host sends and receives multicast datagrams with help of its home agent via a bi–directional tunnel between home agent and mobile host.

3. **Uni–directional tunnel from mobile host to home agent**: The mobile host sends multicast datagrams with help of its home agent via a uni–directional tunnel from the mobile host to its home agent and receives multicast datagrams via the local multicast router.

4. **Uni–directional tunnel from home agent to mobile host**: The mobile host sends multicast datagrams via the local multicast router and receives multicast datagrams via a uni–directional tunnel from its home agent.

Table 1. Four approaches to support multicast for mobile hosts

receive→ send↓	A: local (see Fig. 2)	B: via tunnel (see Fig. 3)
A: local	Local group membership	Uni–dir. tunnel HA → MH
B: via tunnel (see Fig. 4)	Uni–dir. tunnel HA ← MH	Bi–dir. tunnel HA ↔ MH

HA: home agent; MH: mobile host

4.3 Comparison

We now compare the four approaches, using the following criteria:

- **Join delay** of a mobile receiver attaching to a link.
- **Protocol overhead** and data transfer due to e.g. tunneling or building a new distribution tree.
- **Bandwidth consumption** for unnecessary multicast traffic and extra signaling traffic caused by mobility of senders or receivers.
- **Multicast routing**. Is it optimal or suboptimal?
- **System load**, i.e. additional processing and storage load on home agents, mobile hosts, and PIM–DM routers.

4.3.1 Local group membership on foreign link

The predominant delay factor in delivering multicast traffic to a moved group member (receiver) in this approach is the join delay caused by MLD. To receive traffic from a multicast group, a mobile host must inform its local multicast router about its demand for multicast groups by sending a REPORT. If the mobile host is configured to wait for the next QUERY, it may experience quite a long join delay. As shown in Section 3.2, the default value of the MLD Query Interval T_{Query} is 125 s. This is far too high, especially for real–time applications.

Let us now consider a mobile sender. The movement of a multicast sender results in two major overheads in the PIM–DM protocol. The first is an unwanted assert process. After moving to a new link, it takes the mobile sender a certain time to detect the link change and to generate a new care–of address. During this time, all outgoing multicast datagrams have an erroneous IPv6 source address. This might trigger a PIM–DM assert process on the foreign link. For example, if Sender S in Figure 1 moves to either Link 2, 3, or 4, the PIM–DM routers will receive multicast datagrams from Sender S on their outgoing interfaces and will therefore believe that there is a loop in the distribution tree. They will try to resolve the loop by sending ASSERT messages. The second major overhead is due to the fact that after each movement of the sender to a new link, a new source–rooted multicast distribution tree must be built by PIM–DM. Initially, multicast traffic is flooded onto every link of the network; afterward, the prune mechanism establishes a multicast distribution tree.

To investigate the additional bandwidth consumption, we must consider mainly two issues. First, bandwidth is wasted due to a mobile sender that has attached to a new link, as just described. While pruning takes place, multicast traffic is unnecessarily transmitted onto links that will not be part of the multicast distribution tree. The wasted capacity depends mainly on the bit rate of the sender, the PIM–DM Prune Delay Time $T_{PruneDel}$ (default 3 s), the number of links to be pruned, and the mobility rate of the sender. The second bandwidth consumption that cannot be neglected is caused by the leave delay of mobile receivers. During the time it takes MLD to detect the absence of group members on a link and to notify PIM–DM, multicast datagrams may still be delivered to links where no group members are attached to (see e.g. Figure 2).

A great advantage of the local group membership approach is that PIM–DM multicast routing is performed in an optimal way, for both mobile receivers and mobile senders. There is no additional system load in home agents and mobile hosts, whereas PIM–DM routers need additional resources to store and maintain all multicast distribution trees, even for those trees that are not used anymore after a host has moved to a new link. To employ this approach, in principle, no modifications of the used protocols are required. However, mobile hosts should send unsolicited REPORTS after moving to a new link. This would decrease the join delay of mobile group members. Moreover, the values of timers in MLD should be decreased to reduce the join and leave delay. We discuss this issue in Section 4.4.

4.3.2 Bi–directional tunnel between home agent and mobile host

We now consider the second approach, a bi–directional tunnel between the mobile host and its home agent. This method is also mentioned in the Mobile IPv6 draft as one possibility for mobile hosts to send and receive multicast datagrams. However, a detailed specification concerning this issue is not given. In case of a mobile multicast sender, the tunnel approach works fine, and no modifications are needed in the draft. In case of a mobile multicast receiver, the situation is more complicated. To have the relevant data forwarded, the mobile host must inform its home agent about the multicast groups it is subscribed to (see Section 4.2.1 B), and the home agent must implement some extra functionalities. Two different solutions seem to be applicable.

The first solution is that home agents are also PIM–DM routers, which is also assumed e.g. in [5]. These routers then regard tunnels to mobile hosts as individual interfaces. This allows mobile hosts to subscribe to multicast groups (and update their memberships) by sending MLD REPORTS through the tunnel directly to their home agent / PIM–DM router. The home agents then tunnel host–specific as well as multicast relevant datagrams to their respective mobile hosts.

In the second (and more general) scenario, we assume that a home agent is not necessarily also a PIM–DM router. Here, the mobile host cannot send its REPORT directly to the home agent. We propose to employ an extended BINDING UPDATE message, which the mobile host can use to send multicast group membership information to its home agent. The home agent then becomes group member of the requested groups on behalf of the mobile host and stores the group membership information in its binding cache. Datagrams addressed to this group will be routed to the home agent, which then forwards them through the tunnel to the mobile host. As long as the home agent has a binding cache entry for the mobile host, it periodically sends REPORTS to its local PIM–DM router.

To allow the mobile host to send group membership information to its home agent, the Mobile IPv6 draft could be extended with a new sub–option for BINDING UPDATES. Currently, the draft defines two sub–options in the IPv6 destination options, namely the *Unique Identifier Sub–Option* and the *Alternate Care–of Address Sub–Option*. We propose a new sub–option called *Multicast Group List Sub–Option*. This option is valid only in a BINDING UPDATE sent to a home agent (*Home Registration* (*H*) is set). Its format is illustrated in Figure 5. The *Sub–Option Data* contains a list of all multicast groups requested by the mobile host. The *Sub–Option Len* fields must be set to $16N$, where N is the number of multicast group addresses included in the Sub–Option Data.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                |Sub-Option Type| Sub-Option Len|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                                                               |
+                   Multicast Group Address                     +
|                                                               |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
               zero or more of the following ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                                                               |
+                   Multicast Group Address                     +
|                                                               |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5. Multicast Group List Sub–Option

In both scenarios mobile hosts do not experience relevant join delays unless they are detached from the network for a certain amount of time. In case of PIM–DM enabled home agents, a long join delay will occur if a mobile host cannot send REPORTS, and thus the MLD group membership timer (default $T_{MLI} = 260$s) expires in its home agent. In the second case, missing extended BINDING UPDATES would let the home agent delete its binding cache entry (binding lifetime default MAX_BINDACK_TIMEOUT = 256 s [12]), and, thus, give up the representation of the host as member of its multicast group. In both cases, the join delay after reconnecting to the network is approximately the same as in the local group membership approach. Again, in both scenarios, the delay could be reduced to a much smaller value by employing unsolicited REPORTS.

We continue our discussion about advantages and drawbacks of the bi–directional tunnel approach: Every mobile host uses its home agent to receive and/or send multicast datagrams. Thus, home agents experience a high system load due to the number of multicast datagrams they must process. The system load of a single home agent increases with the number of mobile hosts it must support, the number of multicast groups its mobile hosts need to receive, and the amount of traffic in the groups. In addition to this, home agents must collect and maintain group membership information for their mobile hosts. Mobile hosts also have extra system load, as they must encapsulate and decapsulate every multicast datagram they send or receive. Moreover, they must supply their home agent with their group membership information. In addition, two relevant protocol overheads occur. First, extended BINDING UPDATES are needed to collect and record the group membership information in the home agents. Second, every datagram must be encapsulated to be tunneled from home agents to mobile hosts and vice versa.

As illustrated in Figures 3 and 4, multicast routing using a bi–directional tunnel is suboptimal. Multicast datagrams for mobile receivers are routed from the sender to the home link over the distribution tree with optimal multicast routing. Then they are encapsulated by the home agent and sent over the tunnel to the mobile host. Multicast datagrams from a mobile sender are first tunneled to the home agent, then decapsulated and forwarded onto the home link. From there, optimal multicast routing based on PIM–DM is performed. We observe that in both cases the tunnel consists of some links and routers that are also used for the multicast distribution tree. Thus, the datagrams are crossing some links and routers twice.

Perhaps the most important advantage of this approach compared with the first approach is that a mobile receiver does not experience any significant join delay when moving to a new link. Another advantage is that the multicast distribution tree does not need to be modified when a mobile host (sender or receiver) moves to a new link. However, the datagram delivery to mobile receivers through the tunnel reduces the benefit of multicasting. If several mobile members of the same multicast group are located on the same foreign link, they will all receive group traffic via their tunnel. This means, that the same multicast datagrams will be sent via unicast to each group member on the foreign link.

4.3.3 Uni–directional tunnel from mobile host to home agent

For this approach, we again suggest that mobile receivers should send unsolicited REPORTS after moving to a new link. Moreover, we recommend to optimize timers in MLD to limit the join and leave delay (see Section 4.4). Multicast routing from stationary senders to mobile receivers is optimal (see Figure 2); but multicast routing from mobile senders to a multicast group is suboptimal, as datagrams are first tunneled to the home link (see Figure 4). Thus, a mobile receiver does not cause any protocol overhead to join a multicast group, but a mobile sender causes a tunnel overhead. In this approach, it is neither required that home agents implement PIM–DM functionality nor extended BINDING UPDATES, as proposed for the bi–directional tunnel, are necessary.

As in the bi–directional tunnel approach, a great advantage of this method is that the multicast distribution tree does not need to be re–built when a mobile sender moves to a new link. However, bandwidth might get wasted when a mobile receiver that is the only group member on a link moves to a new link.

Concerning the join delay, the same problems arise as in the local group membership approach (Section 4.3.1). Home agents have additional system load as they have to decapsulate and forward multicast datagrams from their mobile hosts onto the home link. However, the overall system load for home agents is smaller than using a bi–directional tunnel. Mobile senders have also an extra system load, as they must encapsulate every multicast datagram they send to the group. PIM–DM routers experience no significant additional system load.

4.3.4 Uni–directional tunnel from home agent to mobile host

Each time a mobile sender moves to a new link, a new source–rooted distribution tree must be build (see Section 4.3.1). In addition, it triggers unwanted ASSERT and JOIN messages. If a mobile receiver moves away from its home link, it must supply its home agent with its multicast group membership information (see Section 4.3.2). Furthermore, all multicast traffic from the group must be tunneled to the mobile host, which causes a protocol overhead and suboptimal routing.

4.4 Timer optimization in MLD for better support of mobile hosts

Let us now discuss the problem of long join and leave delays of mobile receivers in more detail. MLD has a main disadvantage in detecting the presence and absence of mobile group members. It uses relatively long timeouts to detect changes in the local group memberships. Mobile hosts cannot use the DONE message when they leave a link. In order to learn the presence of group members, MLD uses periodic QUERIES. They are sent every Query Interval T_{Query}, which is by default 125 s. The absence of group members is detected by not receiving REPORTS for this group before the group membership timer (default Multicast Listener Interval T_{MLI} = 260 s) expires. If we reduce T_{Query}, REPORTS will be sent out more frequently. This will lead to a lower join and leave delay of mobile receivers.

Therefore we propose the following: For early detection of the presence and absence of mobile multicast group members, administrators should speed up the MLD group membership registration process by decreasing the Query Interval T_{Query}.[5] They must check how often mobile hosts with demand for multicast groups connect to their network and how much system resources (bandwidth, processing time, etc.) are needed by QUERIES and REPORTS. Mobile receivers will experience a lower join delay when they move to a new link where previously no group members have been located. The bandwidth cost for this tuning step is small, compared with the bandwidth saving due to a lower leave delay.

[5]It must not be smaller than the Maximum Response Delay $T_{RespDel}$, which is by default 10 s.

5 Conclusions and further work

This paper described some interoperation problems between Mobile IPv6 and PIM–DM. In particular, we dealt with the problem of long join and leave delays if default timer values in MLD are used. To decrease these delays, MLD timers should be set to lower values. To limit the join delay, mobile hosts should send unsolicited REPORTS (for the multicast groups they subscribe) when moving to a new link.

We investigated four possible approaches to support PIM–DM multicast for mobile IPv6 hosts. Each approach is a solution for some specific scenarios and demands, but no general solution can be presented.

The first approach, the *local group membership on foreign link*, is the simplest solution for providing multicast for mobile IPv6 hosts. It does not require special encapsulation and decapsulation mechanisms and can be employed using the current Mobile IPv6 and PIM–DM functionality without any modifications. Moreover, routing of multicast packets is optimal. It is a good solution if processing resources in home agents and mobile hosts are very low. It is not a good solution for highly mobile hosts, both receivers and senders. Mobile receivers must re–subscribe to the multicast group after each movement to a new link. Each time they change a link, they experience quite a long join delay, and, thus, datagrams will be lost. This could be improved by MLD timer optimization and unsolicited REPORTS. Each time a mobile sender changes its link, a new source routed distribution tree must be built.

A *bi–directional tunnel* is interesting for highly mobile hosts, since no significant join and leave delay occurs. However, for this approach, home agents must also have PIM–DM routing functionality or use the proposed Mobile IPv6 *Multicast Group List Sub–Option* in BINDING UPDATES. Furthermore, routing is suboptimal, and compared to the first approach, more processing and storage resources must be available in home agents and mobile hosts.

A *uni–directional tunnel from the mobile host to the home agent* is a combination of the two approaches mentioned above. It preserves network and system resources better than the bi–directional tunnel, routing to mobile receivers is optimal, and there is no additional bandwidth consumption due to mobile senders. Furthermore it is not required that home agents implement PIM–DM functionality. However, this method has some limitations similar to the local group membership, such as join and leave delays (which could be solved by MLD timer improvements and unsolicited REPORTS on the local link). Changes to Mobile IPv6 and PIM–DM are not required.

The last approach, the *uni–directional tunnel from the home agent to the mobile host,* seems to combine most disadvantages of the other approaches if a mobile host is both sender and receiver for a multicast group. Bi–directional tunneling offers more opportunities with the same implementation cost.

This paper focused on *identification and analysis* of the major advantages and drawbacks of the different approaches. We presented some ideas about modifications to the affected protocols that are needed to provide multicast to mobile IPv6 hosts in an efficient and effective way. More work is needed, however, to evaluate these approaches in more detail and to consider implementation issues (in particular with the proposed uni–directional tunnels). In addition, other features like security and load balancing (see our work in [10]) could be considered in combination with these mechanisms.

References

[1] A. Acharya and B. Badrinath. A framework for delivering multicast messages in networks with mobile hosts. *Wireless Networks, Special Issue on Routing in Mobile Communication Networks, to appear*.

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

[3] A. Acharya, A. Bakre, and B. Badrinath. IP multicast extensions for mobile internetworking. In *Proceedings IEEE INFOCOM'96*, San Francisco, Mar. 1996.

[4] J. Bound, M. Carney, and C. E. Perkins. Dynamic Host Configuration Protocol for IPv6 (DHCPv6). Internet Draft, draft-ietf-dhc-dhcpv6-15.txt, May 2000.

[5] V. Chikarmane, C. L. Williamson, R. B. Bunt, and W. L. Mackrell. Multicast support for mobile hosts using Mobile IP: Design issues and proposed architectures. *Mobile Networks and Applications*, 3(4), 1998.

[6] A. Conta and S. Deering. Generic packet tunneling in IPv6 specification. RFC 2473, Dec. 1998.

[7] S. Deering, D. Estrin, D. Farinacci, V. Jaconson, A. Helmy, D. Meyer, and L. Wei. Protocol Independent Multicast Version 2 Dense Mode Specification. Internet Draft, draft-ietf-pim-v2-dm-03.txt, June 1999.

[8] S. Deering, W. Fenner, and B. Haberman. Multicast Listener Discovery (MLD) for IPv6. RFC 2710, Oct. 1999.

[9] W. Fenner. Internet Group Management Protocol. RFC 2236, 1997.

[10] F. Heissenhuber, A. Riedl, and W. Fritsche. Home agent redundancy and load balancing in Mobile IPv6. In *Proceedings Broadband Communications*, Hong Kong, Nov. 1999.

[11] J. Ioannidis, D. Duchamp, and G. Q. Maguire. IP–based protocols for mobile internetworking. In *Proceedings ACM SIGCOMM'91*, pages 235–245, Sept. 1991.

[12] D. B. Johnson and C. E. Perkins. Mobility Support in IPv6. Internet Draft, draft-ietf-mobileip-ipv6-10.txt, Feb. 2000.

[13] C. E. Perkins. IP mobility support. RFC 2002, Oct. 1996.

[14] S. Thomson and T. Narten. IPv6 stateless address autoconfiguration. RFC 2462, Dec. 1998.

[15] G. Xylomenos and G. C. Polyzos. IP multicast for mobile hosts. *IEEE Communications Magazine*, Jan. 1997.