

Classification and Evaluation of Multicast-Based Mobility Support in All-IP Cellular Networks

Andreas Festag, Holger Karl, Adam Wolisz

TKN, Technical University of Berlin
festag|karl|wolisz@ee.tu-berlin.de
Sekt. FT5-2, Einsteinufer 25, 10587 Berlin, Germany

Abstract. To solve the IP mobility problem the use of multicast has been proposed in a number of different approaches, applying multicast in different characteristic ways. In this paper we provide a framework to classify such approaches by analyzing requirements, options for using multicast protocols, and mobility functionalities augmenting the multicast. Within this framework we identify promising combinations of mechanisms and derive four classes of multicast protocols. These classes include both the standard any-source IP multicast model as well as non-standard multicast models. In particular, the use of source-specific multicast is a new approach to support mobility and turns out to be especially beneficial. The paper describes a corresponding network architecture and a flexible software environment that allows to easily implement these and other classes of mobility-supporting multicast protocols. Based on this software environment, an implemented prototype and measurements quantify handover-specific metrics, namely handover and paging latency and packet loss and duplication rates, showing that a multicast-based mobility solution has comparable performance to standard Mobile IP/hierarchical Mobile IP. Moreover, it is shown that a non-standard multicast model has reduced implementation and deployment complexity and security risks than a standard multicast model.¹

1 Introduction

The growing importance of mobile communication deserves support for host mobility in IP-based cellular networks. A common architectural approach is to assign a mobile node, at every wireless access point with which it associates, a new IP address that is topologically correct. Such an address assignment allows to use standard IP routing, the disadvantage is that existing sessions are interrupted as a session is associated with the previous, now invalid IP address. This dichotomy—an IP address represents both node identification and location—is the fundamental problem of mobile IP-based networks and needs to be overcome by mobility concepts.

¹ This work was partly supported by a research contract with SIEMENS.

The mobility concept currently considered to be the most likely candidate for short-term deployment is Mobile IP (1), along with a number of improvements and variations. Mobile IP uses additional agents to separate node identification from current location and ensures that arbitrary nodes can communicate with a mobile node in an uninterrupted way even while the node moves around. Despite this achievement, Mobile IP has been criticized for its performance problems and for not matching all possible requirements for a mobility concept. Hence, other mobility concepts have been proposed to separate a node's identity and location. One approach that inherently uses such a separation is multicasting and has already been considered in several related works: here the identity of a node is the identifier of a multicast group to which a message is destined, the location is a set of conventional addresses. This set of conventional addresses could represent the address(es) under which the mobile node is currently reachable, and as this set can be modified, multicast provides a natural mapping from a fixed identity (equating a mobile node with a multicast group) to a changing set of locations.

The utilization of multicast for mobility support poses a number of challenges. i) Today's IP multicast faces some problems which have prevented its commercial deployment until today (necessity for global address allocation, lack of receiver and transmission authorization, complexity). ii) New challenges arise from the use of multicast for mobility support (dynamic join and leave operations, scalability with respect to number of multicast groups, etc.). iii) Multicast does not offer all functionalities which are required or useful for mobility-related performance.

It is widely accepted that multicast has some attractive features for mobility support, particularly in highly mobile environments with very small cells. Nevertheless, a systematic effort for investigating multicast-based mobility support is still missing. Hence, a number of interesting proposals exist which mostly have different motivations and requirements and are based on different assumptions about the networking architecture. A variety of solutions have been proposed (2) – (10) each providing certain mobility-related functionalities under particular assumptions. The next two sections will derive both the requirements under which to judge such solutions as well as a common framework to classify them and to derive new classes of solutions. In particular, a source-specific multicast approach is a consequence of the application of our framework and has not yet been described in the literature. The before mentioned challenges can be overcome in several different ways, and hence, considerable freedom remains in choosing multicast options to solve the mobility problem. The first main contribution of this paper is an attempt to structure these protocol options and functionalities in an abstract framework. This framework serves as a template from which new concrete solutions can be derived. On the basis of the requirements, we will make our second contribution, a set of preferable instantiations of this template. One of these instantiations is a new multicast-based mobility concept: It relies on modifying the standard *any source* IP service model (ASM) by using the *source-specific* service model (SSM) for mobility support in access networks. Its main benefits lie in reduced implementation and deployment complexity and im-

proved security while achieving a performance that is comparable with that of more complicated multicast models. In order to substantiate this claim we have developed a software environment that turns our abstract framework into a practical platform. While this platform is fixed in some architectural assumptions, it augments a number of different multicast concepts by mobility functionalities in a flexible and extendable manner. Based on this software environment we have evaluated – as the third contribution of this paper – the performance of different multicast schemes demonstrating that the performance is comparable to the classical mobility solution Mobile IP and its hierarchical variant.

The remainder of this paper is organized as follows: The requirements are listed in Section 2 and the abstract framework is introduced in Section 3. Section 4 contains the performance evaluation and Section 5 concludes the paper.

2 Requirements

Identifying the requirements of mobility support is essential for selecting appropriate functionalities for mobility support and for choosing among multicast protocol options. Most of the requirements are evident for mobility support in general (11), therefore, we emphasize requirements that have consequences for the design of a multicast-based mobility solution.

- *Low-end, light-weight end systems* will have only limited memory and processing power, ruling out solutions that have substantial requirements.
- In *heterogeneous access networks*, mobiles will switch between wireless networks with different characteristics regarding bandwidth, etc. (vertical handover). Hence, a multicast scheme cannot assume that network parameters or even organizations remains constant before and after a handover.
- Handovers with *short latency and small packet loss* require a fast execution of multicast join operations, i.e. an unsolicited re-join to the multicast group after a handover.
- In order to support a very high number of mobile hosts, a multicast that is *scalable* with the number of groups with typically a few members is required.
- A full-scale communication network requires *reliable services*, in particular, TCP or a reliable multicast transport protocol should be supported.
- Besides general security concerns for multicast (especially, denial of service and eavesdropping), *location privacy* is important in mobile networks yet constantly violated by today’s mobility schemes as temporary IP addresses are tracked and communicated to correspondent nodes.
- *Signaling load* by location updates and handover should be low. Most IP multicast routing protocols provide soft state maintenance where the routing state needs to be refreshed and expires otherwise. Hard state maintenance reduces the signaling overhead, but is less robust for stale states which are likely to occur in error-prone wireless and mobile environments.
- *Redundant data transmissions* caused by delays in maintaining the multicast tree (branches not being removed immediately) should be kept low.

3 Framework Description

Supporting the requirements outlined in Section 2 with multicast requires answers to two questions: The first question is the choice among a number of different multicast protocol options, the second question pertains to different functionalities that are required to support mobility and that can be efficiently implemented using multicast protocols.

3.1 Protocol Options for Multicast-Based Mobility

Multicast protocols can be used in a number of different ways in order to support mobility. The main alternative options are the following:

Multicast endpoint. Selecting the mobile host as multicast endpoint enables multicast protocols to work across the wireless link. This requires multicast protocols which are optimized for small groups and for economical usage of the scarce wireless resources. Alternatively, the access point might be selected as the multicast endpoint and act as a multicast proxy. This option facilitates the usage of optimized signaling protocols on the wireless link.

Micro vs. macro mobility. Using multicast for macro-mobility allows a uniform solution but requires global scalability of the multicast scheme. If, on the other hand, multicast is only used for micro-mobility, an additional solution for handover between access networks is necessary but the scalability requirements are reduced. However, coupling two different mobility solutions will (in general) necessitate some form of address translating (between unicast and multicast).

Multicast tree directionality. A multicast scheme can provide either unidirectional or bidirectional trees. A unidirectional tree is setup to transport downlink packets from a correspondent node/gateway as the root of the tree towards the mobile node while uplink packets use unicast.

X+Multicast. Multicast based on location-independent addressing and routing can be applied as a sole mechanism for mobility support, but not all multicast schemes enjoy this property (e.g. SGM/XCast (12; 13)). Nevertheless, these schemes might be utilized for augmenting specific functionalities in other mobility approaches, e.g. Mobile IP.

Multicast type. Main options are network layer (including IP multicast and unicast-based solutions) and link-layer multicast (especially, ATM multicast). The multicast type determines the service model, such as the receiver-oriented model of IP, the sender-oriented model of ATM, or the call-model in (14).

Dynamic tree. The multicast tree can be static or dynamic. In the first case, the access points belonging to a pre-established multicast tree cover a geographical area. In the second case, the tree follows the current location (i.e. footprint) of the mobile.

Multicast adaptation. Existing multicast protocols can be used “as is,” without modifications. However, the protocols might be adapted to better meet the requirements of mobility support.

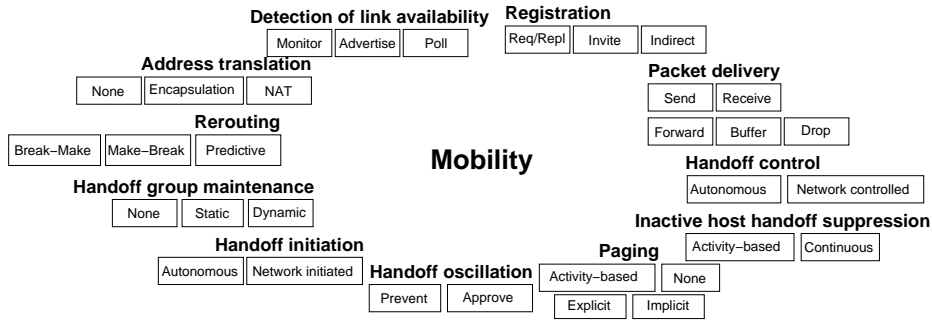


Fig. 1. Functional categories for mobility support.

3.2 Mobility Support Functionalities

The functionalities associated with mobility support can be classified into several categories, each providing a basis for constructing a variant for a multicast-based mobility protocol. The most relevant functionalities are discussed here, additional ones are shown in Fig. 1.

Detection of link availability. Access points may advertise their availability on their local links. A multicast management protocol can directly provide this functionality (e.g. IGMP membership query/report scheme). Optionally, a mobile node may also solicit advertisements from access points.

Registration. On top of existing link-layer connectivity, a mobile node registers with an access point to update its current location information, enabling tracking. Registration can be based on a request/reply scheme initiated by the mobile host or on an invitation by the access point. Alternatively, a mobile node can also be indirectly registered by another access point. Again, the multicast group management protocol can be used for implicit registration.

Rerouting. A rerouting operation changes the network path of packets for a mobile node in a certain access point. A rerouting operation is based on adding and pruning branches of an existing multicast tree. The appropriate multicast operations can be executed in a *break-make* and *make-break* order: new branches are added before old ones are deleted or vice-versa. Additionally, access points can be added to a multicast tree in advance, implementing predictive handover.

Handover initiation and control. Both the network or the mobile host could initiate and control a handover from old to new access point. Controlling a handover means orchestrating the sequence of multicast group manipulation functions. A related issue is the prevention of handover oscillation.

Paging. Inactive mobile nodes reduce their frequency of handover registration and location updates, saving wireless resources. Paging locates such mobile nodes and multicast can be used to efficiently distribute paging requests to a paging area identified by a multicast group. Paging can be done explicitly by sending paging requests to the access points in the paging area or implicitly when data packets are distributed to access points of the paging area.

3.3 Candidate Selections

This framework allows to put together a mobility concept as a combination of various multicast options and supporting functionalities. The application of traditional IP multicast is an evident solution. We follow the approach in (8) in its basic protocol option choices (e.g., placing the multicast endpoint in the access point) and extend it by some additional functionalities (especially support for inactive nodes) resulting in case study MB-ASM (Tab. 1)². We consider this extended version as a base case against which to compare other protocols in a unified experimental environment.

Inspecting this approach reveals that it does not exploit all capabilities of its underlying multicast service model. In particular, all actually necessary multicast functionalities can also be provided by a simplified multicast service model, e.g., source-specific IP multicast resulting in case study MB-SSM. In fact, for fixed networks a trend to a source-specific service model (SSM) can be identified (e.g. by driven by the *EXPRESS* (15) proposal and by availability in commercial products like Cisco IOS-based routers). The benefits of MB-SSM in comparison to MB-ASM are: a) Less protocol complexity and easy deployment, b) inhibits denial of service attacks from unwanted sources and c) averts the problem of address allocation. In this way, SSM alleviates some of the main problems associated with ASM-style multicast as a prerequisite for its utilization for host mobility. Moreover, there are other reasons which endorse the support of SSM: It is ideally suited for tree-like topologies of access networks with a gateway as the root. Since SSM sets up source-based forwarding trees, there is no need for a shared infrastructure with core routers. Finally, the problem of security is aggravated in mobile networks and SSM fairly solves the source access control problem by itself. Additionally, it provides the same actual protocol actions as would result from the use of case study MB-ASM. Hence, the mere performance of both these case studies is expected to be practically identical.

To easily support advanced functionalities such as *make-break* rerouting (in order to implement predictive handover), additional protocol mechanisms such as third-party registration, resource reservation in advance, sub-casting, and others are useful and may eventually improve performance. Such mechanisms provide a larger design space and increased possibilities for the design of multicast-based mobility concepts, but are not available in current IP-based multicast protocols. Yet these mechanisms do exist in some multicast protocols for connection-oriented backbones. In order to investigate such mechanisms, the case study MB-CMAP uses such a connection-oriented, link-layer type multicast protocol. This case study is not investigated in detail in this paper but described in (16).

One of the main objections to multicast is the scalability problem with the number of groups. With a multicast protocol that is to support mobility, a high number of multicast groups exists rather with only very few participants belong to a single group, namely either only the mobile node itself or the access points

² Abbreviations: MC = Multicast, MIP = Mobile IP, HMIP = Hierarchical Mobile IP, NWL = Network Layer, LL = Link layer, SBT = Source-based tree, NA = Not applicable.

with which the mobile node is currently associated (or immediately neighboring access points in case of predictive handover), depending on where the multicast tree is terminated. But the number of multicast groups is going to be very large. Therefore, a multicast protocol for mobility support should much rather scale with the number of groups, where scalability with the number of group participants is only a secondary concern. One example for multicast protocol optimized in this way is small-group multicast. However, this protocol does not separate location and identity and must hence be supplemented by a mobility mechanism. Choosing mobile IP or hierarchical mobile IP results in case the case study MIP-SGM.

	<i>MB-ASM</i>	<i>MB-SSM</i>	<i>MB-CMAP</i>	<i>MIP-SGM</i>
Micro vs. macro	Micro	Micro	Micro	Micro/Macro
Tree directionality	Uni	Uni	Bi	Uni
X+Multicast	Only MC	Only MC	Only MC	MIP/HMIP
Type	NWL	NWL	LL	NWL
	Shared	SBT	Shared	NA
Registration	Req/Reply & Indirect	Req/Reply & Indirect	Req/Reply & by surrogate reg.	Req/Reply
Address translation	Yes	Yes	Yes	Yes/No
Paging		Activity-based,	explicite	None

Table 1. Protocol options and functionalities (different values)

Protocol options	
Multicast endpoint	Access point
Dynamic tree	Yes
Adaptation	No
Functionality	
Detecting link availability	Advertise/solicitare
Packet delivery	Send, receive, forward, buffer, drop
Handover initiation	Autonomous
Handover control	Autonomous
Rerouting	Break-make, predictive

Table 2. Protocol options and functionalities (common to all case studies)

Considering the most important common protocol options in Tab. 2, we assume that the multicast terminates in the access point. Thereby, the mobile host does not need to have any knowledge about multicast. Additional advantages are that this approach better integrates with existing IP-based protocols such as TCP or ARP (3) and that it facilitates the deployment of performance-enabling proxies in the access point improving the protocol performance over wireless links.

4 Evaluation

In order to investigate the selected case studies, we have set up a testbed, designed a set of protocols, implemented a prototype and conducted experiments to quantitatively evaluate mobility-related performance metrics and to compare selected policies. The presented results comprise the case studies MB-ASM, MB-SSM and MB-CMAP including a comparison with basic Mobile IP and hierarchical Mobile IP.³ The protocols for the three case studies are based on a common protocol design. The full description and specification can be found in (17); a description of the prototype in (18). The software is implemented for Linux systems as daemons running in user space and is based on IPv4. It is worth noting that the prototype provides a generic interface to the multicast that facilitates the utilization of different multicast types. Moreover, the prototype provides hooks for policy handlers to control certain system behavior, such as time and destination of handover, buffering and flushing of packets, retrieving signal quality indicators for handover trigger and paging strategies.⁴

We investigate a network topology depicted in Fig. 2(a). The core components are mobile hosts, access points, multicast nodes and gateways running a mobile agent (MA), mobility-enabling proxy (MEP), multicast routing daemon (MRD) and a gateway proxy (GW-P), respectively. These components comprise an access network, which is a network under the control of a single authority. The gateway interconnects the access network with the global Internet.

The testbed setup is illustrated in Fig. 2(b). In addition to the prototype components MA, MEP and GW-P common to all case studies, for the case study MB-ASM the gateway and multicast node execute a multicast router daemon PIM-SMv2 and the access points use IGMPv2. For the case study MB-CMAP the multicast node is represented by a multicast-capable switch with a switch controller⁵ and the MEP and GW use CMAP as a multicast management protocol. In the testbed we replace the wireless link by standard Ethernet. The benefit is that the mobility-specific performance characteristics of various protocols can be studied in a controlled manner in isolation of the error-prone wireless link. Based on the protocol mechanisms, we do not expect any of the protocols to be overly sensitive to wireless errors, so this impact can be safely ignored.

In order to investigate the SSM-style multicast case study, we in fact used an ASM-style multicast implementation. This is indeed feasible and the measurement results are valid as the additional ASM functionality is not used and hence incurs no additional overhead; the differences between the usage of IGMPv2/PIM-SM and IGMPv3/PIM-SSM only marginally affect the performance. Moreover, we are able to corroborate the initial measurements for the handover latency published in (8) of the approach using also IGMPv2/PIM-SM with the same protocol options, but in an extended setup and with enhanced functionalities.

³ Results for MIP-SGM will be published separately due to space limitations.

⁴ The prototype is open software (www.tkn.ee.tu-berlin.de/research/mse.html)

⁵ These components belong to the WashU Gigabit Switch Kit, an open, non-proprietary networking equipment. <http://www.arl.wustl.edu/gigabitkits>.

We have also repeated the same set of experiments for basic and hierarchical Mobile IP (MIP and HMIP) using the Dynamics Mobile IP implementation (www.cs.hut.fi/Research/Dynamics) in a comparable testbed setup with MIP agents replacing the corresponding components of our prototype.

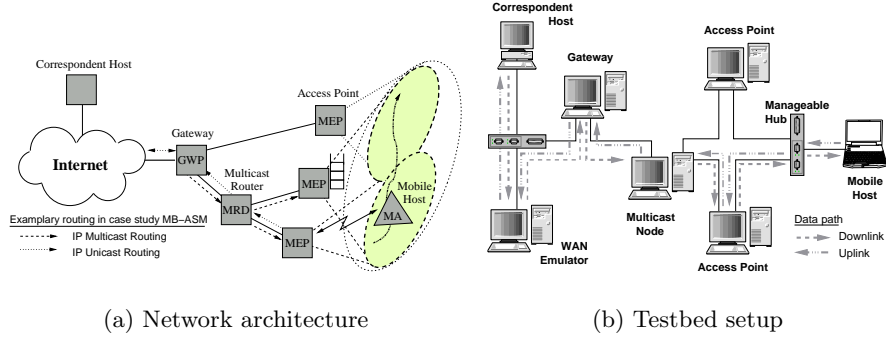


Fig. 2. Investigated scenario

The handover performance is evaluated in terms of handover latency, TCP throughput, and UDP packet loss and duplications.

Handover latency. The handover latency is measured in the mobile host at the IP layer. A continuous packet stream is sent downlink from the correspondent host to the mobile host. During the receive process, the mobile host executes periodic handovers between both access points. The duration between two subsequent handovers (cell dwell time of a mobile in a particular cell) is exponentially distributed with a mean duration of 10s plus an offset of 5s.

In Fig. 3(a), the mean handover latency is independent of the round-trip time between the correspondent host and the mobile host, except for Mobile IP. We examine three handover types: *Hard* means that the old AP is removed from the multicast tree before the new AP is added. *Soft* inverts this order. With *predictive handover* the neighboring AP are added in advance of a handover and buffer packets. In all cases, the handover latency is comprised of the handover detection time and the handover execution time. The handover detection time is determined by the advertisement lifetime of 300 ms sent with an interval of 100 ms and can be regarded as a constant offset of 250 ms. Both the predictive scheme of MB-ASM and all variants of MB-CMAP are superior to MIP/HMIP. In soft handover with MB-ASM the join operation contributes significantly to the handover latency resulting in a handover latency of about 600 ms.

TCP. In this experiment we study the impact of handover on the throughput of TCP Reno. In order to observe the impact of multiple handovers on TCP where the handover events impact each other, we examine a single long-lived TCP connection with subsequently executed handover events. For an advertisement interval of 100 ms we have observed that the TCP throughput decreases only moderately, since the service interruption is relatively small. However, also

for TCP the throughput for predictive handover is larger than for soft handover. Unlike the UDP experiment, this is not an evident result, since the duplication of TCP segments trigger the mobile host to send duplicated acknowledgements, which in turn forces the correspondent host to retransmit TCP segments. However, in our scheme the mobile receives at most two copies of a TCP segment during the handover phase and the TCP retransmission mechanism is not activated. With an advertisement interval of 1 s (Fig. 3(b)) the throughput degrades as the handover frequency grows. Since the duration to detect the handover is relatively long, the TCP sender is forced into TCP's slow-start phase and with a high handover rate TCP has less time to recover from loss. Again, with predictive handover the TCP goodput can be improved even with very frequent handover.

UDP. In this experiment we examine the number of lost and duplicated packets due to handover (Fig. 4(a)). The correspondent host shown in Fig. 2(b) continuously generates 1024 byte UDP packets. We vary the inter-packet times to change the offered load; other parameters are set to the same value as in the handover latency experiment. We obtain the number of lost and duplicated packets by comparing the sent and received trace and marking lost and duplicated packets. For hard (MB-ASM) and soft handover (MB-ASM, MB-CMAP) packet loss increases constantly and reaches 25 lost packets per handover while it can be neglected for MB-ASM and MB-CMAP predictive handover. Hence, except MB-ASM (soft), all multicast schemes have less packet loss than MIP/HMIP. However, the reliability of predictive handover causes some costs in terms of duplications 4(b). For predictive handover the duplications grow with the offered load up to a certain point at 25 kbps with 75/90 (MB-ASM/MB-CMAP) duplicated packets per handover. Beyond a load of 25 kbps the number of duplicated packets decreases slowly. The reason for the shape of the curve is the limitation of the buffer size (set to 100 kByte per mobile) and the buffering policy that forwards only new (not older than 5 s) packets and drops outdated ones.

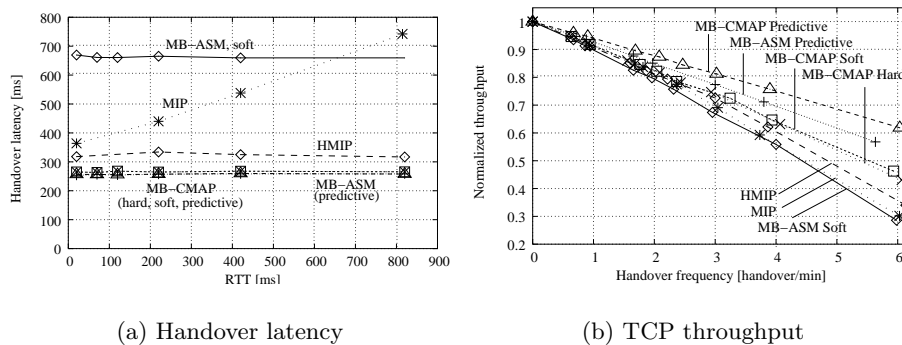


Fig. 3. Handover latency versus round trip time between CH and MH and TCP throughput versus handover frequency (≈ 230 handovers)

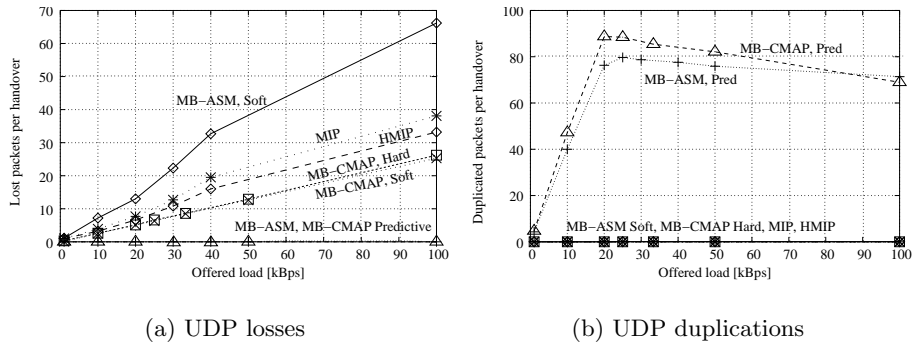


Fig. 4. UDP losses/duplications versus round trip time between CH and MH for horizontal handover (≈ 230 handovers)

5 Conclusions

The large amount of possible approaches to use multicast in order to solve the mobility problem in IP networks made it necessary to structure all these options. Based on a list of requirements for a mobility concept, we developed a list of possible options how to deploy multicast for mobility and a number of possible support functionalities for mobility that can be implemented using multicast. Requirements, options, and functionalities constitute a framework in which existing approaches can be classified and new approaches can be derived by sensible combinations of decisions for each individual aspect.

In particular, we have derived four classes of mobility support protocols, focusing on micro mobility. As a main result, the newly suggested source-specific multicast is superior to multicast based on the standard service model in a micro-mobility context. The multicast protocols themselves were unchanged but augmented with support functionalities for multicast.

In order to quantify relevant performance metrics for multicast-based mobility protocols, we developed a software environment that allows to easily implement a variety of multicast protocols and to experiment with them. Experiments showed that our protocols provide efficient micro mobility and location management along with smooth, fast handovers.

Applying multicast for mobility raises some common objections. We believe that with a judicious choice of protocol options, most of them are unfounded: Scalability problems with the number of multicast groups can be solved by single-source multicast and the application of paging. The (advocated) use of source-specific multicast simplifies the design of sender authentication and authorization; the small number of group members enables the use of key management protocols for securing multicast even in mobile environments.

Bibliography

- [1] Perkins, C.: IPv4 Mobility Support. RFC 2002 (1996)
- [2] Keeton, K., Mah, B., Seshan, S., Katz, R., Ferrari, D.: Providing Connection-oriented Network Services to Mobile Hosts. In: Proc. of the USENIX Symposium on Mobile and Location-Independent Computing. (1993)
- [3] Mysore, J., Bharghavan, V.: A New Multicast-based Architecture for Internet Mobility. In: ACM MOBICOM 97. (1997)
- [4] Ghai, R., Singh, S.: An Architecture and Communication Protocol for Pico-cellular Networks. IEEE Personal Communications **1** (1994) 36–46
- [5] Acampora, A., Naghshineh, M.: An Architecture and Methodology for Mobile-Executed Handoff in Cellular ATM Networks. IEEE JSAC **12** (1994) 1365–1375
- [6] Seshan, S.: Low-latency Handoffs for Cellular Data Networks. PhD thesis, University of Berkeley at California (1996)
- [7] Stemm, M., Katz, R.: Vertical Handoffs in Wireless Overlay Networks. ACM MONET, Special Issue on Mobile Networking in the Internet (1998)
- [8] Wu, J., Maguire, G.: Agent Based Seamless IP Multicast Receiver Handover. In: Proc. of PWC 2000. (2000) 213–225
- [9] Helmy, A.: A Multicast-based Protocol for IP Mobility Support. In: Proc. of ACM NGC 2000. (2000)
- [10] Mihailovic, A., Shabeer, M., Aghvami, A.: Multicast for Mobility Protocol (MMP) for Emerging Internet Networks. In: Proc. of PIMRC 2000. (2000)
- [11] Festag, A., Karl, H., Schaefer, G.: Current Developments and Trends in Handover Design for ALL-IP Wireless Networks. Technical Report TKN-00-007, TKN, TU Berlin (2000)
- [12] Boivie, R.: A New Multicast Scheme for Small Groups. Research Report RC21512(97046)29Juni1999, IBM T.J.Watson Research Center (1999)
- [13] Boivie, R.e.a.: Explicit Multicast (Xcast) Basic Specification. Internet Draft (2000)
- [14] Gaddis, E., Bubenik, R., DeHart, J.: A Call Model for Multipoint Communication in Switched Networks. In: Proc. of IEEE ICC'92. (1992) 609–615
- [15] Holbrook, H., Cheriton, D.: IP Multicast Channels: EXPRESS Support for Large-scale Single-source Applications. In: Proc. of ACM SIGCOMM 1999. (1999) 65–78
- [16] Festag, A., Westerhoff, L., Assimakopoulos, A., Wolisz, A.: Rerouting for Handover in Mobile Networks with Connection-Oriented backbones – An Experimental Testbed. In: Proc. of ICATM'2000. (2000) 491–499
- [17] Festag, A., Westerhoff, L.: Protocol Specification of the MOMBASA Software Environment. Technical Report TKN-01-014, TKN, TU Berlin (2001)
- [18] Festag, A., Westerhoff, L., Wolisz, A.: The MOMBASA Software Environment – A Toolkit for Performance Evaluation of Multicast-Based Mobility Support. In: Proc. of Performance Tools 2002, London, GB (2002)