Proceedings of European Wireless 2000 together with ECRR 2000, Sept. 12-14, 2000, Dresden, Germany

MOMBASA: Mobility Support – A Multicast-based Approach

A. Festag, A. Wolisz, Telecommunication Networks Group (TKN), Technical University Berlin, Germany

Abstract

In general IP multicast supports location-independent addressing and routing for a group of hosts. This ability is similar to the requirement for mobility support in IP-based networks though in a different context. MOMBASA stands for “MObility support – a Multicast BASed Approach”. It is intended to utilise multicast in order to support network level mobility. The basic mechanism of a multicast-based handover can be composed of a multicast *join*- and *leave*-operation of the new and old multicast router, respectively. Additionally to the basic mechanism, MOMBASA investigates extended multicast functionalities which improves the handover in terms of the handover latency, the management of the multicast group and adaptation of the multicast data traffic to the wireless transmission capabilities. In particular the extended multicast functionalities are useful for vertical handover. Unfortunately today’s IP multicast suffers from scalability problems and is not generally available. Nevertheless it is commonly expected that the future Internet will provide an optimised multicast. MOMBASA works out the requirements of handover on future multicast protocols and investigates alternate *IP-style* multicast protocols. As a first approach a multicast protocol is considered which realises a dynamic multipoint-to-multipoint communication in switched networks. An experimental handover testbed is set up which proofs the basic approach and provides extended multicast functionalities. The testbed is described and performance results are given. Other IP-style multicast protocol will be investigated in further steps.

1 Introduction[1]

Handover describes a mechanism in wireless cellular networks that transfers the association of a mobile end system from one base station – which is presently active – to a new base station. In general handover is applied when a user moves through the coverage of a cellular network and crosses cell boundaries. The handover between wireless cells of the same type (in terms of coverage, data rate, and mobility) is often referred to as horizontal handover, whereas the handover between wireless cells of different type are characterised as vertical handover. The vertical handover poses new requirements on handover design.

Nevertheless the fundamental mobility problem in IP-based networks still remains: IP protocols were designed for stationary end systems. The IP address of an end system identifies: a host uniquely and also identifies the IP subnet to which the host is attached. Therefore the meaning of the IP address is twice: end point identification and location identification. When a host changes its point of attachment the IP address must be modified in order to route packets to the mobile's new subnetwork. In order to solve the mobility problems two main approaches can be identified: *Address Translation and Indirect Routing* (e.g. IETF Mobile IP) or *Location-Independent Addressing and Routing*.

Interestingly, there is already an approach which supports addressing and routing independently of the host’s location - multicast. Its ability is similar to the requirement of mobility support, though in a different context. This paper introduces MOMBASA. MOMBASA stands for ‘MObility support – a Multicast-BASed Approach’. MOMBASA intends to utilise multicast in order to support network-level mobility. In relation to the classical approach of *Address Translation and Indirect Routing*, MOMBASA has three main advantages:

- Rerouting for handover is done in a network node where the path to the old and from the new base station diverge (and not in a software agent in the mobile’s home network according to the Mobile IP approach).
- No “handover-specific” signalling and infrastructure is required, instead multicast is reused for mobility purposes.
- MOMBASA minimises the handover latency. In the utmost case packets are distributed in advance to potential new base stations which buffer the packets. In this case the usage of multicast provides the efficient distribution of data to multiple base stations.

Even for vertical handover the usage of predictive mechanisms, such as registration in advance and predictive data distribution, decreases the handover

---

[1] This work was supported by a grant from the BMBF (German Ministry for Science and Technology) within the project IBMS).

latency and packet losses. The realization of these mechanisms with multicast is very natural.

The today's IP multicast protocols fit well for 'broadcast-like' applications although there are a number of open issues to solve. But handover has requirements similar to 'narrow-cast' applications. Such 'narrow-cast' multicast protocols are matter of research and it is expected that the future Internet will provide an optimised Multicast. MOMBASA works out the requirements of handover on future multicast protocols and investigates alternate multicast protocols for their application to host mobility.

The paper is organised as follows: In the next section related work is presented. Then problems of today's IP multicast are discussed and an outlook on recent approaches for future IP multicast is given. Finally, the usage of CMAP (one of the alternate approaches) for handover is described and the testbed setup and performance measurements are presented.

2 Related Work

The authors of [2] were the first who utilised multicast for handover. They proposed a *Multicast-based Re-establishment scheme* which reroutes connections in a crossover point near the base stations. The scheme may use radio hints which identifies the potential new base station in advance. The authors can be regarded as the inventors of a multicast-based handover, although their work is more generic since IP multicast was in an early stage at the time of writing.

The work in [3] explored IP multicast as it is available today for host mobility, without any special changes on the multicast. Their investigations covered also supporting and interacting protocols, such as ARP, ICMP, IGMP, TCP and UDP. They concluded that multicast is able to support host mobility. Nevertheless they stated that several design issues and implementation constraints prevent the wide deployment of today's multicast for handover. They proposed several solutions for the implementation constraints, but argued that some of the issues are also unresolved in the context of IP multicast currently. The approach in [3] is very similar to MOMBASA, since it intends to utilise multicast as a sole mechanism for handover. The difference is that MOMBASA investigates alternate IP multicast approaches.

In [4] IETF Mobile IP is extended by multicast. The Mobile IP Foreign Agents carry multicast addresses and packets are distributed from the Home Agent to several Foreign Agents by multicast. The authors argue that the multicast extension improves the handover latency and packet loss for a Mobile IP handover, in particular for a vertical handover. Contrary to [4], MOMBASA does not base on Mobile IP.

In [8] a multicast routing protocol called Distributed Core Multicast (DCM) with application to host mobility is proposed. DCM is designed for multicast with a high number of multicast groups and a low number of receivers. DCM avoids multicast group state information in backbone routers, it avoids triangular routing across expensive backbone links and scales well with the number of multicast groups. The authors argue that their protocol performs better than the existing sparse-mode multicast routing protocols. The approach of DCM and MOMBASA are very similar. Nevertheless the focus of DCM is on the design of the multicast routing protocol, whereas MOMBASA stresses the mobility aspects. Moreover in contrast to MOMBASA, DCM retains the classical IP multicast service model, whereas in MOMBASA an extended service model is assumed.

3 Discussion of Today's and Future IP Multicast

In general IP multicast assumes 3 types of hosts: receivers, sender and multicast routers. A multicast router is located at least in every IP subnet which supports multicast transmission or reception. Multicast routers form a virtual network to exchange membership information of hosts and to forward multicast packets. The core concept of the today's multicast is that of a host group: A receiver registers with the multicast router in its IP subnet in order to join the multicast group. In the today's multicast a sender is not required to be a member of the multicast group. When a transmitter intends to send a packet to a group it sets the IP destination address of that multicast packet to the IP multicast address. The multicast router intercepts the packet and forwards it towards the distribution tree.

In today's Internet, multicast is realised by employing three types of protocols: multicast group management protocols (e.g. IGMP), routing protocols within a autonomous system (AS) (e.g. DVMRP, MOSPF, CBT, PIM) and routing protocols across ASs (e.g. BGMP). The classical IP multicast service model is common to all of the routing protocols. Although this model is very flexible the application of multicast suffers from scalability problems because

- it has no indication of the group size
- it has no restriction of the allowed senders
- it requires a global address allocation mechanism.

Therefore it is not generally available in the today's Internet. Moreover MOMBASA has specific requirements on multicast which can be summarised as “narrowcast” –type of multicast:

- Small multicast group size
- Very high number of multicast groups
- Frequent *join*- and *leave*-operations

These issues were also identified as requirements for typical “narrowcast” applications, such as IP telephony with conferencing, video-conferencing, multiparty networked games, etc. and have been tackled by several recent IP multicast approaches. Therefore, MOMBASA investigates alternate multicast approaches which extend the IP multicast service model, such as

- Explicitly Requested Single Source Multicast (EXPRESS) [5]
- Small Group Multicast (SMG) [6]
- Simple Multicast [7]
- Connection Management Access Protocol (CMAP) [9]

After proposing the generic handover scheme, in the second part of this paper the CMAP approach is discussed in more detail. The other approaches will be studied in further steps.

4 Generic Handover Scheme

In MOMBASA a unique IP multicast address is assigned to the mobile. The mobile registers with a multicast router and a multicast distribution tree is constructed. A packet from a correspondent host is sent to the mobile’s multicast IP address. The local multicast router intercepts the packet and directs it towards the multicast distribution tree. In the reverse direction the mobile uses the correspondent host’s unicast IP address.

When a handover occurs the new base station might be connected to a different multicast router. In that case the mobile registers the same IP multicast address with the new multicast router. The new multicast router in turn joins the distribution tree at a point which is closed to the new base station. In the best case the new and the old base stations belong to the same multicast router. As an additional operation the mobile has to leave the multicast group via the old multicast router.

As stated above, in MOMBASA a handover can be composed of two basic operations: a multicast *join*- and a *leave*-operation.

In MOMBASA typically all three types of handover are supported (Fig.1). The first type is hard handover where the mobile performs the *leave*-operation to the old multicast router and then performs the *join*-operation. The second type is soft handover. In this case the mobile joins the multicast group and after a short duration of time the mobile leaves the multicast group. Soft handover has the advantage that the data are multicast to the mobile via the old and the new base station. The service interruption and the packet loss are decreased. Soft handover requires that the mobile is able to determine the new base station in advance and the wireless cells must overlap. For vertical handover this is usually granted. Assuming that the mobile is equipped with multiple interfaces, the mobile receives packets on its active interfaces.

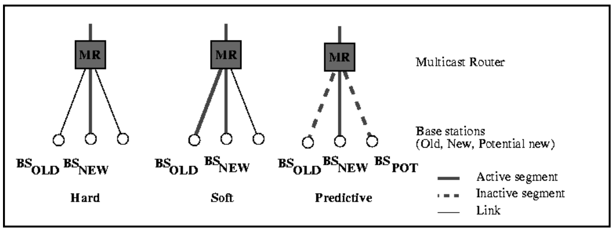

Figure 1: General handover schemes

The third scheme is the predictive handover. In this scheme base stations which are potential candidates for handover of mobile, form a set. When a mobile registers with a base station the base stations belonging to the set are added in advance to the multicast group. These base stations buffer the packets. When the mobile registers with the new base station the buffered packets are forwarded to the mobile. Otherwise the buffered packets will be dropped. The predictive handover reduces the service interruption to an absolute minimum. Nevertheless the approach causes a remarkable overhead in terms of buffers size and of transmission capacity, in particular of the wireless link. In order to decrease the overhead two solutions are considered: On the one hand the overhead can be reduced by making the base station set smaller, e.g. by utilising environmental information (floors, doors, rails, etc.). On the other hand it might be useful to add the potential new base station to the multicast group without multicasting data to them.

Additionally to the basic multicast *join*-and *leave*-operations some extended multicast functionalities have been identified which arise from the utilisation of multicast for handover:

- **Topological knowledge**: The prediction of a potential new base station may require topological knowledge. On the one hand for predictive horizontal handover this knowledge can be used to optimise the set of base station. On the other hand classical methods for handover initiation fail for vertical handover from a larger to a smaller wireless cell. Even for this case topological knowledge offers the possibility of network-initiated handover.
- **Third party signalling**: For soft and predictive handover at least one potential base stations is added to the multicast group in advance. Third party signalling enables the actual base station to

perform the signalling operations on behalf of the new base stations. This avoids additional signalling between the base stations and improves the management of the multicast groups.

- **Subcasting**: In order to reduce the overhead caused by predictive handover it is desirable that the packets might be sent to a subset of the multicast group.

5 CMAP Approach

CMAP [9] stands for **C**onnection **M**anagement **A**ccess **P**rotocol. It is a signalling protocol for dynamic multipoint communication. The protocol provides network clients with dynamic multipoint, multi-connection communication channels, which are termed *calls*. In particular a client may

- Open, modify, close a call
- Add, modify, drop endpoints to/from a call
- Add, modify, drop connections to/from a call
- Trace calls, endpoints of a call

As an important feature, in CMAP a transmitter must be a member of the multicast group. CMAP can be considered as a user-network side protocol, whereas on the network side a multicast tree between the sources and the sinks is constructed: One of the network clients participating in the call is designated as the owner of the call and a single tree is established by incremental routing from the end points to the owner of the call.

Additionally to the “toolbox-like” multicast mechanisms CMAP offers a number of *extended* functionalities:

- Third party signalling
- Resource reservation
- Multiple parallel connections with different QoS parameters within a call
- Mechanisms to control the access, monitoring and modification of call attributes

The latter mechanisms allows to exclude base stations from receiving data although they are members of the multicast group and resources are provisioned.

Originally the CMAP approach was developed at the beginning of ATM standardization. Nevertheless we believe in its actuality due to several reasons:

First, CMAP has features which are recently applied in IP-based multicast. As an example, in CMAP a host which intends to send to a multicast group must be a member of the multicast group. This is different to the native IP multicast service model but is considered in several new IP multicast approaches, such as [3].

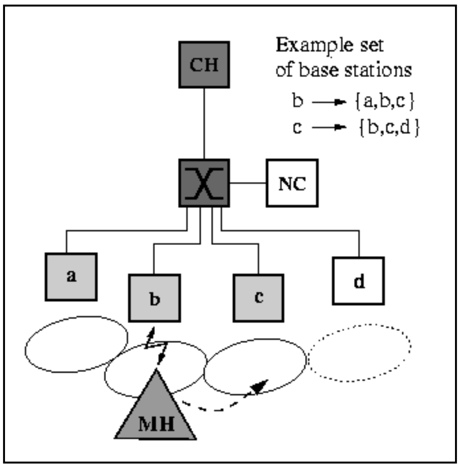

Figure 2: Selected Scenario

Second, the CMAP approach can be considered as a migration solution until scalable multicast is available: The cell pipes which constructs the multicast tree can be regarded as tunnels for IP traffic. The tunnelling of IP packets is well known from the today’s MBONE where multicast routers are interconnected by IP tunnels. In native IP multicast the multicast group is identified by the IP multicast address, whereas in the CMAP approach the call identifier plays a similar role as an identifier of the multicast tree.

Third, the CMAP approach is very flexible. In particular the extended functionalities as enumerated above qualifies it for experimental investigations.

6 Experimental Testbed

The testbed demonstrates the feasibility of the MOMBASA approach and provides extended multicast functionalities. In general the testbed is based on an open, programmable, non-proprietary networking environment. This section describes the hard- and software components and the selected scenario. Finally performance measures are presented. A more detailed description of the architecture and setup can be found in [10].

6.1 Hardware and Software Components

The testbed is consists of hardware and software components and tools for monitoring and measurement. The hardware elements are a multicast network node, a node controller, four base stations, a single mobile and a correspondent host.

The network node is an open, non-proprietary cell switch from the Washington University Gigabit Switch (WUGS) program [12] intended for experimental research. It is an 8 port cell switch supporting up to 2.4Mbps per port. The switch is equipped with two dual OC-3 line cards and six G-

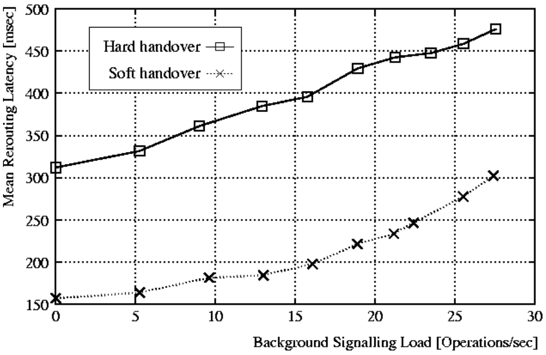

Figure 3: Rerouting latency (Mean) vs. Background Signalling Load

Link line cards (1.2Gbps). The cell switch supports multicast efficiently through cell recycling. This is a technique where cells arriving on an input port are sent to an output port and optionally recycled back to the input port. This recycling of cells can be repeated. The cell recycling adds a very small delay, but yields a scaling gain.

The base stations are standard PCs equipped with common ATM Network Cards ENI-155 or custom PIC cards. The latter cards support a G-Link interface. They are open, non-proprietary and programmable as the cell switch. The wireless interface in the base station usually is replaced with standard Ethernet (10Mbps).[2] The mobile is a standard PC and is equipped with a Ethernet card.

The node controller, correspondent and mobile host, and base stations run Linux 2.2.10 and NetBSD 1.4.1, respectively. The base stations and the correspondent host run CMAP clients, the node controller run the CMAP session manager and underlying protocols, which belong to the WUGS signalling software environment. Between the mobile host and the base stations a lightweight signalling protocol performs registration at the base station. In the mobile a virtual driver solves operating-system problems related to handover (such as MAC address changes in a socket structure, etc.).

6.2 Selected Scenario

The following scenario has been investigated (Fig. 2): For the predictive scheme the base station BS1 has got assigned the set of base stations {BS1, BS2 and BS1a}. The set of base station of BS2 consists of {BS1, BS2, BS2a}. The mobile host has connectivity to the base stations BS1 and BS2, respectively and performs a *ping-pong* handover between base station BS1 and BS2. All base stations are connected to the same network node (switch with node controller NC)

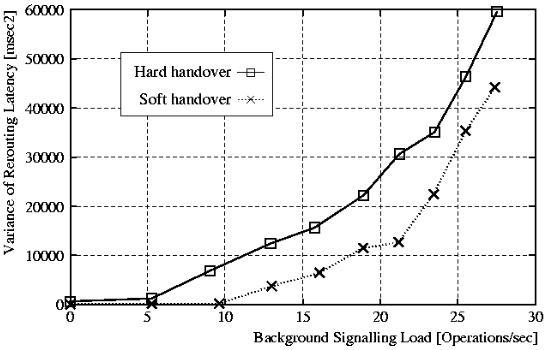

Figure 4: Rerouting latency (Variance) vs. Background Signalling Load

. The correspondent host (CH) opens and closes a multipoint call and remains fixed. When the mobile host (MH) registers with the one of the base stations the base station adds itself to the call. A handover is initiated automatically with a specific frequency: Then the mobile registers with the second base station. This new base station adds itself to the multipoint call and drops the old base station. The sequence of execution is according to the selected handover scheme.

6.3 Performance Results

In this section performance results are presented. We were interested in two parameters: First in the duration of service interruption caused by handover[3] and second in the signalling overhead. The results are realized by measurements and a number of factors have an impact on the measurements such as operation system, protocols and implementation issues.

In order to investigate the first parameter we have measured the rerouting latency in the base station. The rerouting latency is considered as the major part of the service interruption. It is defined as the duration from the reception of a handover request sent by the mobile until the reestablishment of the connectivity - in the case of a soft handover after the *join*-operation[4], in the case of a hard handover after the *leave*- and *join*-operation. Figure 3 shows the mean value of the rerouting latency dependent of a background signalling load which was generated by additional signalling clients executing a mixture of CMAP operations. At a low signalling load a reasonable mean rerouting latency can be observed. As commonly expected the mean latency increases as the load grows. But as depicted in figure 4 the latency

---

[2] The focus of the measurements is on backbone aspects. Therefore we are not interested in an error-prone wireless link.

[3] Predictive handover is not taken into account since the rerouting latency is theoretically zero.

[4] For soft handover the *leave*-operation does not contribute to the rerouting latency

variance increases. Thus the rerouting latency is difficult to predict due to its high spread.

To estimate the overhead caused by signalling we have measured the received and transmitted CMAP traffic in the node controller for a single rerouting operation which consists of several CMAP operations. Since CMAP uses TCP as a transport protocol we have measured at TCP level with *tcpdump*. In this setup there was no background signalling load. For hard and soft handover we have measured the same amount of signalling: 1.392bytes. For the predictive handover this measure depends strongly on the number of *join*-and *leave*-operations which are required to maintain the multicast tree. Consider an example where two *join*- and two *leave*-operations are executed. In this case, a signalling load of 2.768bytes is generated. This is about twice as much as for hard and soft handover.

7 Conclusions

In the paper we have proposed MOMBASA. MOMBASA stands for MObility support – a Multicast-BASed Approach. MOMBASA intends to utilise multicast in order to support Internet host mobility. This approach eliminates the need for address translation for handover. The main advantages of MOMBASA are: a) rerouting is performed in a network node close to the base station. b) the multicast infrastructure is reused for mobility purposes and c) the handover latency is minimised. Unfortunately today's IP multicast is not generally available mainly due to scalability problems. Additionally handover has "narrowcast" requirements on multicast, such as small group size, a high number of groups and frequent *join*- and *leave*-operations. MOMBASA investigates alternate IP multicast approaches which extend the classical IP multicast service model. As a first step one of the approaches was investigated in an experimental testbed. The CMAP approach offers a number of extended multicast functionalities, such as third party signalling, resource reservation, support of multiple parallel connections with different QoS parameters and sophisticated mechanisms to control the access, monitoring and modifications of the multicast group. Finally results of performance measurements were presented which investigate rerouting latency and signalling overhead.

8 Literature

[1] Perkins, C.: IPv4 Mobility Support. IETF RFC 2002

[2] Keeton, M.; et al.: Providing Connection-oriented Network Services to Mobile Hosts. Proc. of the USENIX Symposium, Oct. 93

[3] Mysore, J; et al.: A New Multicasting-based Architecture for Internet Host Mobility. Proceedings of ACM Mobicom'97, Oct. 97

[4] Stemm, M.; et al.: Vertical Handoffs in Wireless Overlay Networks. ACM Mobile Networking (MONET), Special Issue on Mobile Networking in the Internet, Winter 1998

[5] Holbrook, H.; et al.: IP Multicast Channels: EXPRESS Support for Large-Scale Single-source Applications. Proceedings of ACM SIGCOMM'99, Sept. 99

[6] Boivie, R. : A New Multicast Scheme for Small Groups. IBM T.J. Watson Research Center, Technical Report RC21512(97046)29JUNI1999

[7] Perlman, R. et al.: Simple Multicast : A Design for Simple, Low Overhead Multicast. Internet Draft, Febr. 1999

[8] Blazevic, L.; et al.: Distributed Core Multicast (DCM): A Routing Protocol for IP with Application to Host Mobility. ACM SIGCOMM Computer Communication Review, Vol. 29, No. 5, Sept. 99

[9] Gaddis, M.; et al.: A Call Model For Multipoint Communication in Switched Networks. Proc. of ICC'92, June '92.

[10] Festag, A.; et al.: Rerouting for Handover in Mobile Networks with a Connection-Oriented Backbone – An Experimental Testbed. Proc. of ICATM'2000, June 2000

[11] Festag, A.; et al.: Interconnection of Wireless Cells - a Multicast-based Approach. Proceedings of ICUPC '97, Oct. 1997

[12] WUGS - Washington University Gigabit Switch Environment. Washington University, St. Louis, URL: http://www.arl.wustl.edu/gigabitkits/

Andreas Festag received his Diplom degree in Electrical Engineering from the Technical University Berlin, Germany in 1996. He is currently a research assistant of the Telecommunication Networks Group, Department of Electrical Engineering, TU Berlin. His research interests are in the area of mobile communication.

Adam Wolisz obtained his degrees Dipl.- Ing. in Control Engineering, Dr.-Ing. and Habilitation (both in Computer Engineering) respectively in 1972, 1976 and 1983, all at the Silesian Technical University in Gliwice, Poland. Since 1993 he is a Chaired Professor of Electrical Engineering and Computer Science (secondary assignment) at the University of Technology Berlin, where he is leading the Telecommunication Networks Group (TKN). In parallel he is also member of the Senior Board of GMD Fokus. His research interests are in architectures and protocols of communication networks as well as protocol engineering with impact on performance and Quality of Service aspects. Recently he is focusing mainly on mobile networking.