IPv6-based OverDRiVE Moving Networks[1]: Mobility Management and Testbed Implementation

Alexandru Petrescu[+], Hong-Yon Lach[+], Christophe Janneteau[+], Michael Wolf[++]
Tim Leinmüller[++], Christoph Barz[+++], Markus Pilz[+++], Matthias Frank[+++], Ralf Tönjes[++++]

[+]Motorola Labs, Paris, France, {petrescu,lach,jannetea}@crm.mot.com
[++]DaimlerChrysler, Telematics Research, {michael.m.wolf,tim.leinmueller}@dcx.com
[+++]University of Bonn, Germany, {barz, pilz, matthew}@cs.uni-bonn.de
[++++]Ericsson Eurolab Deutschland GmbH, Ralf.Toenjes@ericsson.com

ABSTRACT

Mobile IPv6 has no native support for mobility of entire networks, such as vehicular networks or personal area networks. This paper suggests protocol extensions for IPv6 moving networks. These extensions rely on a bi-directional tunnel between the moving network's Mobile Router and its Home Agent, providing uninterrupted connectivity for moving networks while changing access networks. Several supported configurations of moving networks are presented. The paper describes a proof-of-concept implementation using the Motorola LIVSIX IPv6 stack. Drawbacks such as crossover tunnels and excessive nested tunnelling are discussed.

I. INTRODUCTION

Continuous Internet connectivity of moving vehicles opens up new opportunities for applications, enhancing the driving experience while streamlining the vehicle maintenance processes. In addition to Web browsing and multimedia streaming for passengers, new services are possible, such as Internet-based region-specific driver information, remote supervision and control, telematics services as well as automatic software updates to vehicle equipment (PCs, head screens, engine computers and sensors).

Today IPv6 provides mobility support for individual network nodes. However, support for moving networks, e.g. cars, buses, trains, ships and also IP based Personal Area Networks (PANs), is missing. In the context of the IST project OverDRiVE we developed, implemented and demonstrated IPv6 protocol enhancements to provide seamless IP connectivity to whole IPv6 networks while changing access systems (cf. [1]).

The paper is structured as follows: the following section analyses the problem of mobility for IPv6 networks. Based on the requirements for moving networks described in section 3, section 4 proposes a solution for uninterrupted connectivity of moving networks while changing access networks and discusses open issues, such as crossover tunnels and asymmetric paths. Section 5 presents the LIVSIX implementation and finally section 6 concludes the paper.

II. PROBLEM STATEMENT

A moving network is a network segment that can move and change its point of attachment to the Internet. It is a leaf network and can only be accessed through specific gateways, called Mobile Routers (MRs). Mobile IPv6 solely supports the mobility of a Mobile Router and not the mobility of the whole leaf network – comprised of nodes that move along. These nodes are no longer reachable when the MR (and thus the moving network) moves away from the home network as described below.

A. Mobile IP Supports Mobility of Hosts

The main problem addressed by Mobile IPv6 is that, in the current Internet design, every IPv6 address corresponds to a fixed "location" in the routing fabric, the route information towards this location being maintained by all the intermediary routers. A Mobile Host (MH) or Mobile Router (MR) is permanently assigned a *Home Address* (HoA) that is valid in the home domain, i.e. the official administrative domain to which this MH or MR belongs to. If a host moves from its home location to another it gets a new routable address – the *Care-of-Address* (CoA). In order to maintain the ongoing applications (e.g. a video stream session is not interrupted by the address change) Mobile IPv6 introduces a Home Agent (HA) that intercepts all packets addressed to the home address and forwards them to the CoA. When the MH is not at home, the HA acts on its behalf using proxy Neighbour Discovery, intercepting all packets directed towards the Home Address and subsequently encapsulating them towards the MH's current CoA. The HA maintains a binding cache with value pairs (*Home Address, Care-of Address)*. In a symmetric manner, the MH in a foreign domain first encapsulates all its outgoing packets towards the HA that decapsulates and resends them towards the original destination. These two encapsulation mechanisms are currently referred to as *bi-directional tunnelling* and together with binding cache management. They constitute

[1]This work has been performed in the framework of the IST project IST-2001-35125 OverDRiVE (Spectrum Efficient Uni- and Multicast Over Dynamic Radio Networks in Vehicular Environments), which is partly funded by the European Union. The OverDRiVE consortium consists of Ericsson (co-ordinator), DaimlerChrysler, France Telecom, Motorola and Radiotelevisione Italiana as well as Rheinisch-Westfälische Technische Hochschule RWTH Aachen, Universität Bonn and the University of Surrey. The authors acknowledge the contributions of their colleagues in the OverDRiVE consortium.

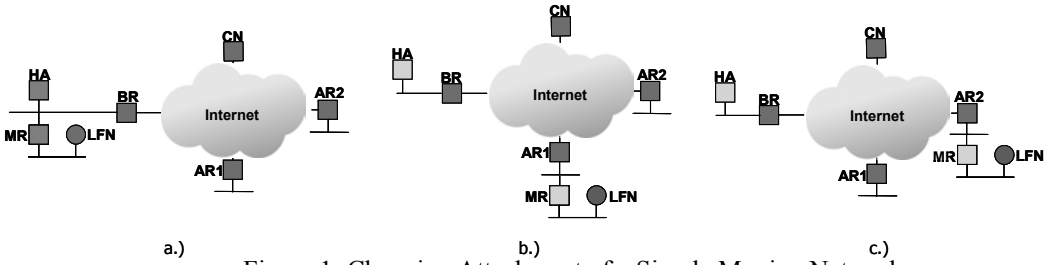

Figure 1: Changing Attachment of a Simple Moving Network

the essence of the Mobile IPv6 protocol. Some extensions worth mentioning include security mechanisms for MH-HA communication, route optimization with return routability tests and hierarchical mobility management.

B. Missing Support for Moving Networks

Mobility of entire networks (see Figure 1) constituted of fixed nodes is not supported by the current Mobile IPv6 protocol. When a Correspondent Node (CN) sends a packet to a Local Fixed Node (LFN) in the moving network, this packet is routed to the Border Router (BR) of the LFN's home network. The BR checks its routing table for the next hop towards the LFN. The routing table lists the MR's home address as next hop. The BR sends NDP messages to discover the MR's MAC address. If the MR is at home it returns its address and the LFN is reached via the MR.
If the MR is not at home the HA sends gratuitous neighbour advertisement messages on behalf of the MR. Hence the packets are sent to the HA. The HA checks its binding cache but has no entry for the LFN. (The LFN is mobility unaware and only the MR has registered a CoA). The HA forwards the packet to the default router, e.g. the BR. This router checks its binding cache and finds again the MR as next hop which are intercepted by the HA. The packet enters a loop between the HA and the BR playing *ping-pong* until the TTL expires. Experimental trials reported in [2] expose this inappropriate routing table management of Mobile IPv6.

III. REQUIREMENTS

The requirements for a network mobility protocol can be stated shortly as:

- Reachability of nodes regardless of the current point of attachment of MR (nodes can be contacted at their well known, fixed address).
- Session continuity for nodes while MR changes its point of attachment to the Internet (when MR changes its Care-of Address, applications on a network node should not be forced to re-establish sessions).
- If the size of the moving network is relatively large, it is important for the Mobile Router to be capable of using a dynamic routing protocol, even if Mobile IPv6 involves switching between various IP addresses (CoA) on the egress interface.

In order to address the first two requirements, it is necessary to adapt the Mobile IPv6 protocol to support moving networks (instead of MHs).

The third requirement imposes a different type of problem that we present only briefly. In short, no known implementation of a dynamic routing protocol (e.g. Zebra for RIPng and OSPF) is capable to dynamically launch new instances when *tunnel interfaces* change. In addition, if a Mobile Router – as we'll see in section 4 – would dynamically sets-up and tears down the bi-directional tunnel towards the Home Agent whenever a new CoA is acquired, the resulting flipping effect of tunnels is subject to induce instabilities in the inner workings of the respective dynamic routing protocol. Furthermore, the management of the *egress* interface of a MR must not allow for participation in the routing protocol exchanges with routers of the visited domain, but at the same time it must participate in a such exchanges with routers at home (MR should not inject new routes on the visited domain. Otherwise, it risks inducing uncontrollable growth of routing tables in the core Internet routers and hence induces important instabilities to Internet.)
While we do not intend to present a detailed solution to the MR dynamic routing problem here, a simple method can be sketched as follows. A simple Mobile Router with two physical interfaces (ingress and egress) must be able to run a dynamic routing protocol in two states: (1) when the MR is at home run normally on both physical interfaces, (2) when MR is in a foreign network, stop running on the egress interface and run on the tunnel interface (and continue on the ingress). In addition, the MR must be able to switch between the two states dynamically when changing attachment between the home and the visited network, and when changing attachment between various visited networks.
Finally, the centralized mobility management handled by the MR is attractive from an efficiency point of view, e.g. hosts within the moving network are relieved from the need to send individual Binding Updates (BUs) – only the MR sends a unique BU when the entire network changes attachment point.

IV. NETWORK MOBILITY WITH MRHA TUNNEL

A. The MRHA Tunnel Concept

The approach taken in OverDRiVE for network mobility has been to use vanilla Mobile IPv6 wherever possible in order to keep the impact on existing implementations small. Therefore, only changes to the MR and its HA have been suggested. The approach employs a bi-directional tunnel between the MR and its HA. The basic idea is to include a 'R' –flag for the MR in the HA's binding cache,

indicating that a complete moving network (and not only a host) moved. The HA forwards packets through the MR-HA tunnel for all nodes within the network of the MR. The HA acts on behalf of the link-local address of MR's moving interface (when the MR is in a foreign network). As in the mobile hosts case, the MR is using BUs and BAcks with the HA to maintain the MR-HA bidirectional tunnel. For a detailed description the reader is referred to [1].

The modifications to Mobile IPv6 HA and MH specifications allow for co-existence with existing implementations of Mobile IPv6 for hosts. These modifications involve an appropriate binding cache and routing table management on the HA and MR. Other proposals to support network mobility with an MR/HA-like approach (cf. [6] for IPv4) involve important modifications in routing table management *and* in protocol signalling, like using mobile network home prefixes instead of the full /128 address of the MR. Additionally, these alternative approaches have important security risks due to incapability of authenticating an entire prefix (instead of an address). To entirely avoid dealing with the prefix-ownership problem[2], our MR/HA approach considers only the movement of the MR's Home Address – a full /128. This is feasible as a result of the following analysis, again similar to the Mobile IPv6 for hosts case: traffic coming from outside the home link takes a certain network path, involving multiple routers. Of those, only the Mobile Router is affected by mobility; even if its network layer address is part of Border Router BR's routing table (and BR does not run Mobile IPv6), BR actually uses the link-layer address corresponding to IPv6 address, by discovering it via the Neighbour Discovery (ND) protocol. With proxy ND the HA advertises its own link-layer address to resolve MR's IPv6 address. Thus, a HA intercepts the traffic destined for the MR at the home link. In this way, it is sufficient for a HA to have a Binding Cache entry of the form MR_CoA-MR_HoA (as compared to MR_CoA-prefix). It is assumed that HA also maintains a routing table entry towards the mobile network's prefix through MR's Home Address (or a host-routing entry towards the Local Fixed Node LFN) [3]. Reversely, traffic coming from the mobile network (when the MR is in a foreign network) towards the Internet, is first forwarded by the MR through the reverse MR-HA tunnel to the HA. Then HA decapsulates and forwards to the original destination. Note that the absence of protocol modifications to Mobile IPv6 messages between MH and HA brings in the advantage of obtaining a MR implementation by having minimal modifications to an existing MH implementations (only the table management is modified, but not the format of messages). A unique MR implementation works for both: MHs and MRs. This is an essential advantage for OverDRiVE scenarios involving a PDA user (the PDA is the MH) getting in a car where a mobile network is deployed (the car's mobile network connects to the Internet with an MR)[4].

B. Signalling for Tunnel Setup

Figure 1a illustrates a generic initial setting for mobile networks. The following entities are depicted: a large Internet cloud, a Correspondent Node CN, two Access Routers AR1 and AR2 as well as the home network composed of a Border Router BR, a Home Agent HA and the mobile network comprising the Mobile Router and a Local Fixed Node LFN. For this scenario we assume that an application is running continuously between the CN and the LFN. Here, changes in addresses induced by the mobility of the MR do not affect the communication. In our particular trials, CN is continuously streaming towards the Home Address of LFN and the LFN displays that stream on a video screen, as the packet exchanges in left diagram of Figure 2 suggests.

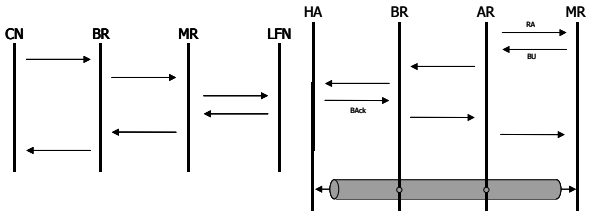

Figure 2: Initial Communication and Tunnel Setup for Simple Moving Networks

The streaming packets and the encapsulation/decapsulation actions are illustrated in Figure 3, where bended arrows at the left of vertical bars represent decapsulations, and at the right – encapsulations. In that figure, the detailed ND messaging between BR and HA is excluded; a complete description can be found in [5]. In the other direction, when LFN needs to send a packet to CN and the mobile network is not at home, it will first send the packet to its default route, the MR. MR encapsulates back to HA which decapsulates and forwards to original destination. The diagrams b and c of Figure 1 describe the movement of the mobile network away from home towards any of the foreign networks governed by AR1 or AR2. The movement from the home network to AR1 triggers a binding message exchange between HA and MR in order to setup the MR-HA bidirectional tunnel. This is illustrated in the right diagram of Figure 2. The MR receives a Router Advertisement (RA), configures a new CoA and default route. Subsequently, the MR sends a BU to the HA. Then, the HA sets up its endpoint of the MR-HA tunnel and replies with a Binding Acknowledgement (Back). At this moment the MR-HA tunnel endpoint at the MR is set up, too. Once that tunnel is up, the streaming between CN and LFN continues via the new CoA. The CN sends next application packet towards the BR and latter asks for the

[2] The classic address ownership problem [3] can be reformulated as a prefix-ownership problem, where MR needs to demonstrate it "owns" the prefix assigned to the mobile network to which it offers Internet connectivity.

[3] There are several ways in which HA can maintain these entries, mentioned in [2] and [5] (basically: manual configuration, dynamic routing protocols, ICMP Redirect).

[4] If OverDRiVE only involved MR's then this would not have been an advantage.

link-layer address corresponding to the IPv6 address of MR. The HA replies on behalf of the MR. Afterwards, the HA encapsulates and forwards this packet through the MRHA tunnel towards the current CoA. MR decapsulates the received packet and forwards it to the LFN. A symmetry can be noticed by drawing an imaginary horizontal axis through the centre of Figure 3: packets from CN to LFN take exactly the same path backwards, without state maintained neither on CN nor on LFN, as a consequence of bidirectional tunnelling.

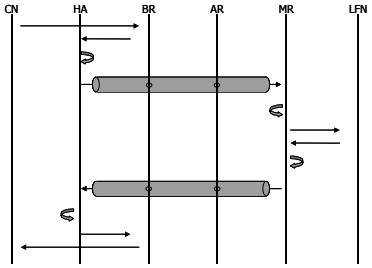

Figure 3: CN-LFN Exchanges, Simple Moving Networks

The simplicity of this approach is advantageous to supporting moving networks. The topology presented in Figure 1 can be easily augmented with a large number of moving networks, ARs and CNs. All involved packet exchanges are similar (if not the same) to the packet exchange in Figure 3.

B. Nested Moving Networks

As presented in section 2, another important OverDRiVE scenario is exhibited by a Mobile Host attaching to a moving network (a person with a mobile device entering a vehicle). A moving network setup for this scenario is depicted in Figure 4. The diagram a. is the initial configuration with MR and MH at home. The diagrams b., c. and d. illustrate the movements, in that order. Diagrams at the top of Figure 5 present the packet exchange between different entities of the initial configuration.

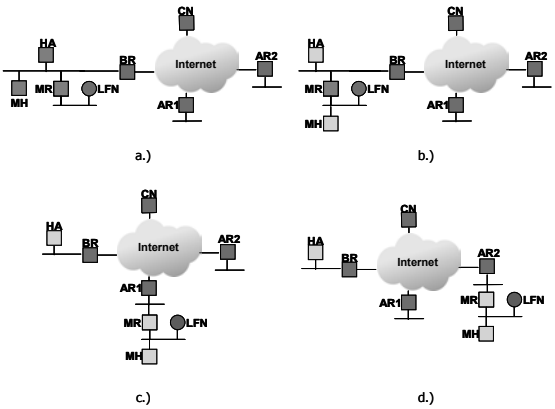

Figure 4: Mobile Host and Mobile Router, unique HA

The packet exchanges between CN and MH are corresponding to diagram c. in Figure 4, to the top-left diagram in Figure 5 and to Figure 7. Remark that HA performs a double encapsulation for each forwarded packet. Among the various paths, the link AR1–MR is the most influenced segment by the multiple encapsulations. Remark also that initial communication between CN and MH involves 3 IP entities; while after moving from home, it involves 6 entities (excluding each time the IP entities between CN and BR, and between BR and AR). The LFN–MH packet exchange corresponds to diagram c. in Figure 4, to the top-middle diagram in Figure 5 and Figure 6.

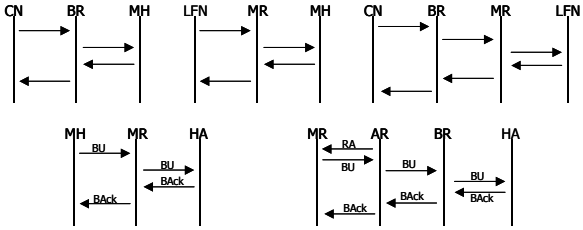

Figure 5: Initial Exchanges and Tunnel Setup, Mobile Network and Mobile Host

Remark that even if LFN and MH are physically close to each other, many other out-of-path entities are forwarding their packets. This configuration lacks the symmetry of the simple MR case in Figure 3, the asymmetry (4 BR-HA exchanges when LFN sends to MH vs. 2 BR-HA exchanges when MH sends to LFN, see Figure 6) is due to attaching both MH and MR to a unique HA. The CN-LFN message exchanges corresponding to diagram c in Figure 4 and to top-left diagram in Figure 5 is not depicted here due to space constraints, but it is similar to the CN-MH exchanges, except that the last de-capsulation (or the first encapsulation) is performed by MR instead of MH. Note finally that it is possible to replace the MH with another MR. This leads to a more generic nested moving network configuration, where the message exchange is similar.

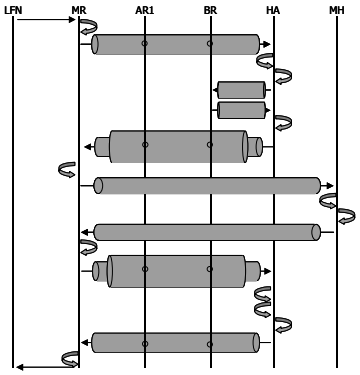

Figure 6: LFN-MH Exchanges, when MH and MR

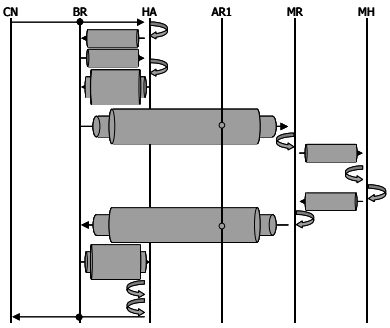

Figure 7: CN-MH Exchanges, when MR and MH

V. TESTBED IMPLEMENTATION

As described in section IV, network mobility support with the MR-HA tunnel requires an IPv6 stack implementation that supports the appropriate MR and HA behaviour. In OverDRiVE, we used the open source LIVSIX IPv6 stack ([10]). It offers all the basic TCP/IPv6 functionalities in a Linux environment. It supports Mobile IPv6 Home Agent and Mobile Host implementations. We developed the Mobile Router functionality on top of LIVSIX and performed several experiments that helped us to reveal some of the issues at the protocol level. The OverDRiVE LIVSIX testbed includes several core routers, access routers, home agents and mobile networks.

The test setup consisted of a DaimlerChrysler E-Class equipped with a prototypical Mobile Router running LIVSIX IPv6 mobility stack, realizing bi-directional Internet access via wireless LAN and UMTS and uni-directional DVB-T reception for specific web pages and video streaming. Applications like a car web server allows authorized users to monitor the status of the car and to perform certain actions remotely. For example, certified OEM software could be downloaded to the car. In addition adaptive multimedia streaming was shown to demonstrate the great benefits of seamless connectivity for IPv6 networks. A detailed description of the demonstrator can be found in [4] where the core network and the moving network are described as part of the final project demonstration in Torino, Italy, in December 2003.

VI. CONCLUSIONS AND FUTURE WORK

This paper described a simple yet scalable approach to mobile networks: the MR-HA bidirectional tunnel with Mobile IPv6. At protocol level, several issues have been raised with respect to the described MR-HA approach. Most topics relate to Neighbour Discovery issues between a HA and a BR, route discovery, link-local addressing and security delegation between a LFN and a MR [5]. Overall, the presented MR-HA approach to network mobility has the following advantages:

- Little or no modification to current Mobile IPv6 protocol specification, allowing simple co-existence with Mobile IPv6 for mobile hosts. It inherits from Mobile IPv6 the session continuity and ubiquitous reachability at a permanent Home Address.
- Support of Mobile Hosts and networks that visit mobile networks (nested mobility).
- Inheritance of the strong security mechanisms used by Mobile IPv6, based on the IPsec architecture for securing the MR-HA tunnel.
- Applicability to a large number of mobile networks that roam around the edge of the Internet and no introduction of non-scalable routing table updates to the core network (prefix-augmented host-based routing inefficiencies are avoided).

In addition, other issues of the MR-HA approach such as excessive tunnelling, cross-over tunnelling, disconnected operation and under-optimal paths have been described in [7]. A proof-of-concept implementation with the LIVSIX IPv6 stack has been demonstrated and briefly described.

Future work will address the identified MRHA drawbacks. Using various types of header compression can ease excessive tunnelling (ROHC or Deering-Zill compression). The last four drawbacks will be addressed from a generic route optimization perspective, i.e. pure routing functionalities (other than Mobile IPv6) are combined with the mobile networks with MR-HA functionality. In addition, future work will consider multi-access aspects for Mobile Routers.

REFERENCES

[1] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", IETF Internet Draft, draft-ietf-nemo-basic-support-02, (Work in Progress), December 2003.

[2] [5]T. Ernst, A. Olivereau, L. Bellier, C. Castellucia, H.-Y. Lach, "Mobile Networks Support in Mobile IPv6 (Prefix Scope Binding Updates)", IETF Internet Draft, draft-ernst-mobileip-v6-network-03.txt, (Work in Progress), March 2002.

[3] H.-Y. Lach, C. Janneteau, A. Olivereau, A. Petrescu, T. Leinmueller, M. Wolf, M. Pilz, "Laboratory and Field Experiments with IPv6 Moving Networks in Vehicular Environments", IETF Internet Draft, draft-lach-nemo-experiments-overdrive-01.txt, Work in Progress, October 2003.

[4] T. Leinmueller (Ed.), "Description of Demonstrator for Mobile Multicast and the Vehicular Router", OverDRiVE Deliverable D14, IST OverDRiVE project, March 2004.

[5] [5]P. Nikander, "An Address Ownership Problem in IPv6", IETF Internet Draft, draft-nikander-ipng-address-ownership-00.txt, (Work in Progress), September 2001.

[6] C. Perkins, "Mobile IP", IEEE Communications Magazine, May 2002.

[7] A. Petrescu, ed., M. Catalina-Gallego, C. Janneteau, H.-Y. Lach, A. Olivereau, "Issues in Designing Mobile IPv6 Network Mobility with the MRHA Bidirectional Tunnel (MRHA)," IETF Internet Draft, draft-petrescu-nemo-mrha-03.txt, (Work in Progress), November 2003.

[8] M. Pilz, C. Barz (Eds.), "Functional Description and Validation of Mobile Router and Dynamic IVAN Management", OverDRiVE Deliverable D17, IST OverDRiVE project, March 2004.

[9] M.A. Ronai, R. Tönjes, M. Wolf, A. Petrescu, "Mobility Issues in OverDRiVE Mobile Networks", IST Mobile Summit 2003, Aveiro, Portugal, June 2003.

[10] Livsix, http://www.nal.motlabs.com/livsix

[5] Note: References [2] and [5] are expired Internet Drafts. However, they contain important concepts for our paper. The interested reader may obtain these IDs from the authors of this paper.