Journal of Interconnection Networks
© World Scientific Publishing Company

NETWORK MOBILITY
FROM THE INTERNETCAR PERSPECTIVE

Thierry ERNST, Koshiro MITSUYA and Keisuke UEHARA
{ernst|mitsuya|kei}@sfc.wide.ad.jp
WIDE Project Keio University
Faculty of Environmental Information
5322 Endo, Fujisawa
Kanagawa 252-8520 Japan

Received (April 15, 2003)
Revised (June 15, 2003)

A number of devices, including sensors, mobile phones, and various computers will be deployed in next generation vehicles, and interconnected through an embedded in-vehicle network. These vehicles will be connected to the Internet as both a step toward ubiquitous computing and as a means to meet Intelligent Transportation Systems (ITS) needs. At first, a communication system is required to connect vehicles to the Internet. Such communication system is investigated in our InternetCAR project. For flexibility and ease of use, we advocate IPv6. This paper particularly focus on network mobility since network mobility support is mandated to maintain ongoing sessions as the in-vehicle embedded network changes its point of attachment to the Internet topology. We outline our testbed specifically designed for the purpose of demonstrating our proposed communication system and we detail our implementation based on Prefix Scope Binding Updates, the initial network mobility support solution proposed at the IETF before the NEMO working group was set up.

Keywords: Network Mobility, NEMO, IPv6, Mobile IPv6, ITS, in-vehicle network

1. Introduction

Recent advances in computer miniaturization and wireless technology promise increasingly powerful, light, small and functional wireless devices. This hardware miniaturization together with the improvement of wireless communication technologies drive the need for even more mobile communications. As more and more people are traveling with a laptop, a PDA, a mobile phone, a digital camera, or any other high-tech device, there is a desire to connect permanently to the Internet from anywhere, at anytime, without any disruption of service, particularly for those who spend a significant amount of their time for commuting. Connecting vehicles to the Internet is a means to achieve this trend toward ubiquitous computing.

Besides this, vehicles are more and more embedded with a wide range of devices, like sensors or navigation facilities. Some, not to say all, will need Internet access, as driven by Telematics and ITS (Intelligent Transportation Systems) applications. So far, conventional technologies have been providing information to drivers and passengers via signboards, radio system, mobile phones and more sophisticated technologies. With the advent of the Internet and IPv6, these systems can be integrated into a general digital communication system. Once connected to the Internet, not only Internet data can be accessed from the vehicle, but vehicles can also be monitored and various environmental information generated by the embedded sensors can be advertised to the Internet.

For a number of reasons (network management, security, performance, signaling overhead, value added services, ...), it is desirable to interconnect in-vehicles devices through a local in-vehicle network, therefore exhibiting the need to displace an entire network. Such a network would migrate in the Internet topology, and is referred to as a *mobile network*. A typical example of a mobile network is a vehicle connected to the Internet via multiple wireless medium, as investigated in the InternetCAR project (Internet Connected Automobiles Research) [1,16]. Other cases of mobile networks include networks attached to people (Personal Area Network, or PAN for short)[a], networks of sensors deployed in vehicles (aircrafts, boats, buses, trains) that need permanent Internet connectivity, and access networks deployed in public transportation (taxis, trains, aircrafts, personal cars) to provide Internet access to devices carried by their passengers (laptop, camera, mobile phone, and PAN).

In this paper, we first outline the recent IETF (Internet Engineering Task Force) effort conducted on network mobility (NEMO). Then, we outline the InternetCAR project and we detail the communication system requirements for connecting vehicles to the Internet. The following section describes our InternetCAR testbed, the protocols used to demonstrate our communication system, particularly network mobility, and our implementation.

2. Network mobility at the IETF

Mobility in the Internet arises when a portion of the network changes its point of attachment to the overall topology. This could be as easy as shutting down a computer, unplugging it, and plugging it in another network, switching on, and then running a protocol such as DHCP to acquire an IP address at the new point of attachment. However, ongoing data transfers would be lost during the migration. Indeed, the Internet is hardly tuned to allow mobility in the midst of data transfers because protocols are not conceived for devices that change their point of attachment in the topology. There is typically a change of the physical IP address each time a mobile node (MN) changes its point of attachment and thus its reachability to the Internet topology. This results in losing packets in transit and breaking transport

---

[a] PAN: a small network composed of Internet appliances like PDAs, mobile phones, digital cameras, etc.

protocols connections if mobility is not handled by specific services. The protocol stack must therefore be upgraded with the ability to cross networks in the midst of data transfers, without breaking the communication session and with minimum transmission delays and signaling overhead. This is commonly referred to as *mobility support.*

Traditional work in this topic is to provide continuous and uninterrupted Internet access to *mobile hosts.* Host mobility support is handled by Mobile IPv6 [15,13] and specified by the IETF Mobile IP working group. On the other hand, the question of networks that frequently change their point of attachment to the Internet has only gained the deserved attention from the Internet community for about two years. In this case, a *mobile router* (MR) changes its point of attachment, but there is a number of nodes behind the MR (MNNs). In [11] we demonstrated that Mobile IPv6 is unable to keep connections open between arbitrary correspondent nodes (CNs) and MNNs. Network mobility indeed exhibits specific characteristics and new mobility scenarios which prevent to apply Mobile IPv6 as is. This led to set up the IETF NEMO (NEtwork MObility) working group in October 2002, in which we deeply contribute. NEMO's primary objective is to preserve session continuity between CNs and all MNNs behind MR while the MR changes its point of attachment. The working group is currently defining the problem scope of network mobility, the requirements that must or should be fulfilled by the solutions and the constraints that may limit the deployment of potentially good solutions [9]. A number of propositions, including PSBU [11] on which is built our implementation presented in the last section of this paper, have been submitted to the IETF before the working group was actually created. As such, these proposals do not tend to address all the issues nor do they meet all the later defined requirements. The working group has therefore started to work on a common proposal [7] that should meet the requirements for all the desirable scenarios.

The NEMO terminology used in this paper is defined in [10] and summarized in Fig.1. A MR has at least two interfaces, the *egress interface* is attached to the visited link, and the *ingress interface* is attached to an *internal link* in the mobile network. A mobile network is *multi-homed* when the Internet connectivity is provided simultaneously by more than one MR or when the MR has more than one *egress interface* (likely on distinct physical access medium). MNNs are either fixed nodes (LFNs) or mobile nodes. LFNs are unable to change their point of attachment while keeping their connections open, whereas mobile nodes have this ability. If a mobile node is indeed a MR with a number of nodes behind it, the aggregated network is said to be *nested.*

All documents related to network mobility can be found on the NEMO web page [4].

3. The InternetCAR Project

The InternetCAR project [1,16] was launched in July 1996 to investigate how ve-

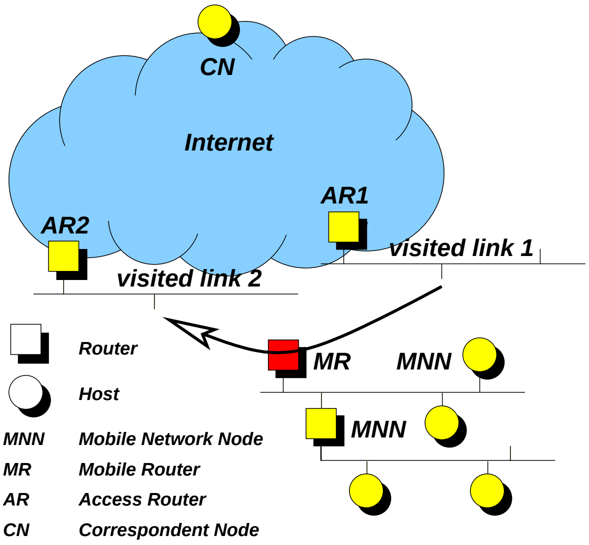

Fig. 1. NEMO Terminology

hicles could be connected to the Internet in a transparent manner, permanently, regardless of their location, mobility frequency and available communication medium. The initial project was divided into three phases and a few field trials were run. Phase 1 was started in July 1996 and ended in September 1997. We designed a special car which included a 19 inch rack mount computer, a wireless LAN, a cellular phone and an extra dynamo to supply electric power. Mobile IPv4 was run on the computer and the two media could be switched according to the wireless LAN's radio availability. Phase 2 was started in October 1997 and ended up in March 1998. 10 normal vehicles were used for this experiment, each loaded with a laptop computer, a data collection box, a GPS receiver and a cellular phone. In April 1998, phase 3 was started based on results of previous phases. The hardware was re-designed, focusing toward mass production. In 2001, a probe information system was tested with about 270 vehicles (including taxis, trucks, buses and so on) connected to the Internet by means of a mobile phone. Each vehicle had some sensors connected to a in-vehicle computer. In 2002, a larger scale trial was held involving 1640 vehicles. Vehicles were connected to the Internet via a cellular phone (PDC-P) or Dedicated Short Range Communication (DSRC). 140 vehicles had two of them and the rest had cellular phone only. A number of Internet applications were demonstrated, particularly the probe information system. This system aims at providing traffic condition information by collecting various data of a large number of vehicles. The traffic condition information was based on velocity of vehicles and weather prediction was based on the status of wipers. Other typical applications demonstrated in the project include screens delivering location-based information to the passengers, an accounting system in car parks, a ramp approach guidance in a gas station, a taxi control system and more conventional applications like multimedia and navigation. Many more could be contemplated.

The InternetCAR project is now investigating the communication system requirements to efficiently connect vehicles to the Internet [12]. These requirements are detailed below and summarized in Tab.1[b].

Table 1. InternetCAR requirements

User requirements	System requirement	Type
interconnected devices	in-vehicle network	F
continuous Internet access	network mobility support	F
	multiple access media	F
	multi-homing	F
large number of vehicles /devices	IPv6	F, D
use standard IP devices and low-cost network appliances	mobility transparency for fixed nodes (backward compatibility)	D
connecting mobile phone, PDA	support mobile nodes	F
connecting PAN	support nested mobility	F
application quality	Fast handoff	P
	low overhead	P
	vertical handoff	P
	horizontal handoff	P
low costs	low overhead	P
	multiple access media	F
	multi-homing	F
	optimal routing	P
maintenance and ease of configuration	zero-conf	F,D

A typical vehicle has about 50 computers and 120 sensors. They need to be interconnected through an in-vehicle local network (LAN) in order to allow them to communicate with one another. Internet Protocol (IP) is the protocol that effectively allows any two Internet nodes to do so. There presently exists two versions of the protocol, the getting old IPv4, and the new generation IPv6 [6] designed to address its shortcomings. A number of reasons drive us to favor IPv6 over IPv4. First, IPv6 has built-in features that allow to support more effectively the new services requested by new demanding applications. This includes support for mobility, multicast, traffic reservation, security, etc. A second reason is that IPv6 offers a generous number of addresses compared to IPv4. As the number of automobiles is expected to reach 700 millions, each one embedding tens of computers, it is necessary to have a sufficiently large number of addresses. Not to say, it is also necessary to ensure that the system under development will be compatible with future devices. Basically, any kind of device, like a mobile phone or a watch will ultimately be connected to the Internet,

---

[b] Column on the left shows the requirements as expressed by users; column in the middle shows how it translates into system requirements for network mobility support. Column on the right classifies them into three categories: 'F' for "functional", 'P' for "performance", and 'D' for "deployment".

and this will be IPv6, not IPv4, as advocated by cellular operators, also concerned by the available address space and the built-in features of IPv6 since they are willing to allocate an IP address to each mobile phone. Lastly, deploying IPv6 devices in vehicles should make application development costs cheaper since these applications and devices wouldn't only be developed for telematics.

For safety reasons, it is necessary to isolate traffic types from one another. Devices are thus interconnected by three types of networks, as illustrated on Fig.2: audio and navigation systems are attached to the multimedia network; devices controlled by passengers (head lights, power windows, etc) are attached to the body network; whereas the control network takes care of the operation of the vehicle (engine and breaks, sensors that monitor the air pressure in the tyre or amount of fuel left in the tank, etc). The first network that need to be shifted to an IP-based network is obviously the multimedia network since many existing devices could straight be brought up into the vehicle.

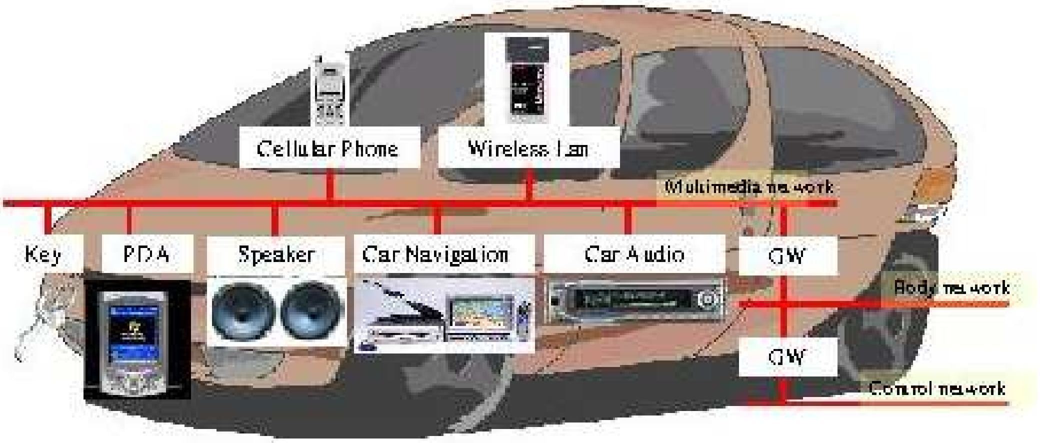

Fig. 2. The InternetCAR Prototype

Assuming all devices to connect directly to the Internet and to manage individually their own mobility with respect to the Internet is unrealistic for a few reasons. First, this would prevent most devices that were not conceived with mobility support, like most, from exchanging data from and to the Internet. Mobility thus needs to be managed transparently to those devices that do not have mobility support facilities (*mobility transparency*). Second, even if they would be embedded with such facilities, this would probably not be very efficient. It is better if mobility of the network is dealt with as a whole and single unit, i.e. the sole MR(s). For doing so, a specific *network mobility support* mechanism is needed.

Although most in-vehicle devices are fixed, some of them may be brought in and off the vehicle, without disrupting on-going sessions. Vehicle users (driver, passengers) would probably like to bring in their PDA or even a PAN, and get Internet access on these appliances through the in-vehicle network. There is therefore a desire to achieve two levels of mobility (*nested mobility*). Also, connecting such devices must be easy since usual people are not familiar with Internet technologies. No specific configuration must be necessary (*auto-configuration*).

To ensure continuous connectivity to the Internet, the vehicle is best connected via several access technologies, preferably simultaneously and to distinct access networks (*multi-homing*). The system must thus be able to deal with both *horizontal handovers* (between different point of attachments using the same communication medium) and *vertical handovers* (between different communication medium). As a wide coverage cannot be ensured by a single Internet Service Provider (ISP), handovers may need to be performed between distinct administrative domains and thus topologically distant parts of the Internet (*global mobility*). This either may occur when vehicles cross country boundaries, or when different ISP offer distinct access technologies. The latter is effective to keep communication costs low from a user's point of view. One can use 802.11b instead of a cellular phone near parking lots or when there are traffic jams. At this point, media switching time must be kept short for sensitive applications (*fast handoffs*). Finally, *overhead* must be minimized to save bandwidth resources and thus the operation cost.

4. InternetCAR Experimentation

We have designed the testbed illustrated in Fig.3 to demonstrate how vehicles could be connected to the Internet. A network is deployed in one of the prototype vehicles described in the previous section. The vehicle gets Internet access via a MR. The MR (PentiumIII 500MHz, 192M) is running NetBSD (version 1.5.2 with KAME snap-20010921) and has 5 distinct links: 2 are Personal Digital Cellular (PDC) (PDC-P 9.6 Kbps, PacketOne 28.8 Kbps), 2 are Personal Handy-Phone System (PHS) (AirH" 32 Kbps, P-in 64 Kbps) and 1 is IEEE 802.11b. Since PDC and PHS currently only provide IPv4 connectivity, we use IPv6 over IPv4 tunneling and Dynamic Tunnel Configuration Protocol. A Dedicated Short Range Communication (DSRC) access is provided by another router, running Linux under IPv4, because there is no DSRC driver available in NetBSD for IPv6. 10 LFNs (PentiumII 500MHz,192M; 5 are running FreeBSD, 4 Win2K, and 1 Linux) are located besides MR, connected via FastEthernet. A CN, the HA and a tunnel server are located on the SFC campus at Keio University. The CN (PentiumIII 500MHz, 256M) is running a SNMP client under FreeBSD 4.4 or NetBSD 1.5.2.

The vehicle is moving in order to get in and out of the coverage area of the wireless access router. We only perform *vertical handoff* between 802.11b and PHS and only one interface is used at once. The current interface is switched automatically when the vehicle is getting out of the coverage area. Each I/F has a different priority (802.11b is 1, PHS is 2, Cellular is 3). A daemon monitors the status of the interfaces

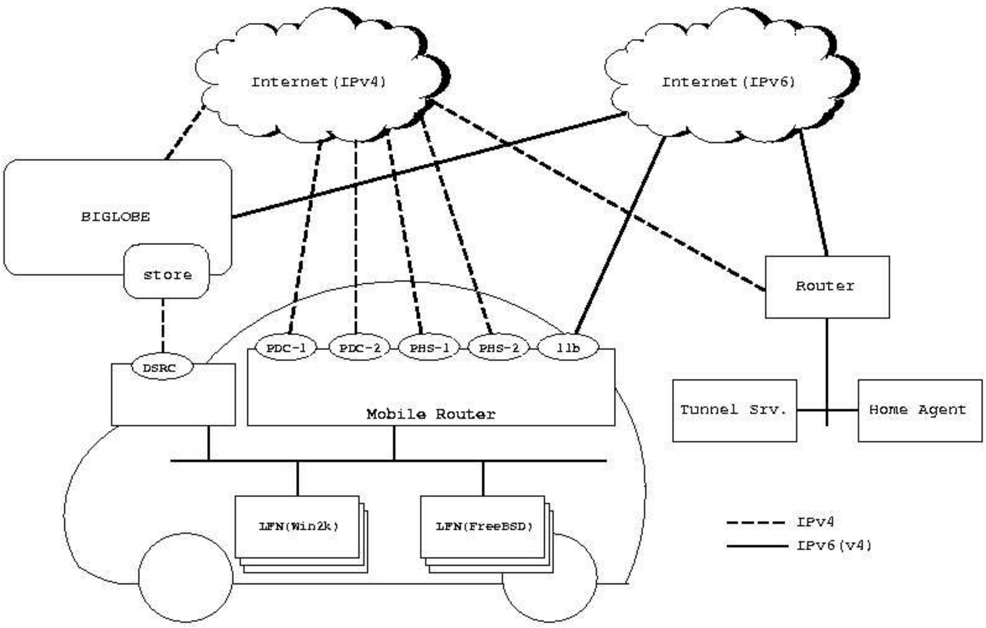

Fig. 3. InternetCAR Testbed

on the MR. The status is a link-layer information such as radio field intensity. PHS is usually used because the vehicle is not in the service area. Once the vehicle comes under the service area, the daemon decides to use 802.11b.

Although this testbed was used to demonstrate the operation of multiple protocols, including interface switching, this paper's focus is network mobility. Network mobility support relies on Prefix Scope Binding Update (PSBU) [11] which was the first and only available solution when we started our experimentation, and also because it fits most of the requirements listed in Tab.1, particularly mobility transparency at LFNs. PSBU re-uses most of the Mobile IPv6 features, so we assume the reader is familiar with Mobile IPv6 which is best explained in [15]. Its operation is nevertheless summarized below, and illustrated on Fig.4. PSBU is implemented on top of our own Mobile IPv6 stack (SFC-MIP) [5], based on Mobile IPv6 draft-13. Our implementation is called MADRID [2] and is run by MR, CNs and HA.

The results obtained so far show that the system is working fine. CN and LFNs are able to communicate efficiently in both directions. The system is evaluated in [14].

4.1. *Mobile IPv6*

Mobile IPv6 [13] associates a mobile node with two distinct addresses, a permanent home address MN_{ip}, obtained in the home network (home link), and a temporary care-of address MN_{coa}, obtained in the visited network (foreign link). The home address serves as a location invariant node identifier whereas the care-of address serves as a routing directive to the current point of attachment. The operation of the protocol is illustrated on Fig.4. The binding between MN_{ip} and MN_{coa} is registered with the home agent (HA) located on the home link by means of a Binding Update (BU) message (1). The binding is recorded in the HA's Binding Cache. As a result of this registration, HA adds a *host-specific route* for MN_{ip} (i.e. for a 128-bit IPv6 address) via the MN_{coa} through a tunnel. Packets originated from a CN are first sent to MN_{ip}, therefrom routed to the home link where they are intercepted by HA (2). MN_{ip} is then used as the key for searching the Binding Cache. MN_{coa} is return and the packet is encapsulated by HA to MN_{coa} (3). At this point, the packet is decapsulated by MN and a BU can be sent to CN (4,5) in order to avoid triangle routing via HA (6: route optimization).

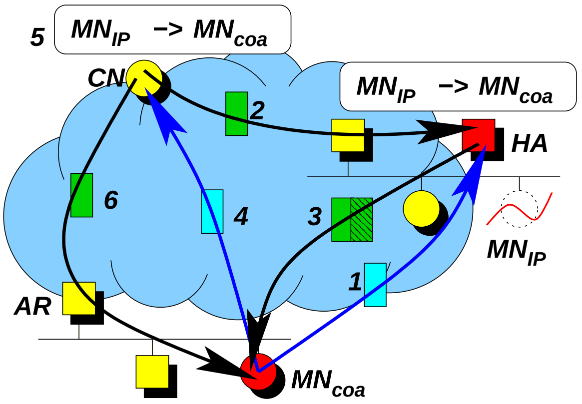

Fig. 4. Overview of Mobile IPv6

4.2. *Prefix Scope Binding Updates*

Prefix Scope Binding Update (PSBU) [11] was submitted to the Mobile IP Working Group in August 2000. It proposes straight forward extensions to Mobile IPv6 (draft version 13). A new BU format is defined to register an entire mobile network in one step as opposed to Mobile IPv6's BUs which only register a sole end-node. All MNNs are assumed to be named on a common prefix (NEMO Prefix). The MR is behaving similarly to a MN. Its egress interface has two addresses, a home address on the home link, and a care-of address on the foreign link. The ingress interface has an address taken from the NEMO Prefix. The binding associates MR's care-of address with the NEMO Prefix instead of the full 128-bit IPv6 home address as defined in the Mobile IPv6 specification[c]. The format of the BU Option is slightly modified

---

[c] Accustomed readers may have noted that MR's prefix on the home link (the home prefix) is not the same as the one on the internal link (the NEMO Prefix). In addition to the binding between

and comprises a new Destination Sub-Option (Mobile Network Prefix Sub-Option) (Fig.5 and 6). The MR is a new entity which performs most of the existing MN Operation with the additional ability to send PSBUs. The CN Operation and the HA Operation are slightly enhanced in order to process the PSBU. Upon reception of a valid PSBU, the record in the Binding Cache can be seen as a *network-specific route* for the NEMO Prefix.

Modified Binding Update Option format A new bit (P) is taken from the reserved set of bits in order to indicate that the registration worth for the prefix as indicated in the Mobile Network Prefix Sub-Option, and not for a single 128 bits IPv6 address. The format is described in Fig.5.

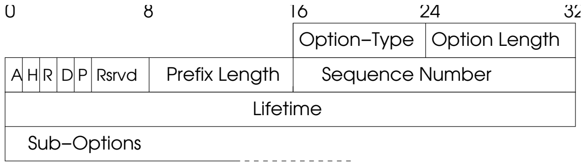

Fig. 5. PSBU Format: Modified Binding Update Option

Mobile Network Prefix Sub-Option: This new BU Option Sub-Option is defined to carry the NEMO Prefix. It contains the NEMO Prefix used as a netmask. This Sub-Option is inserted in the BU. The format is described in Fig. 6.

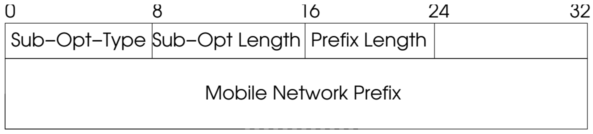

Fig. 6. PSBU Format: Mobile Network Prefix Sub-Option

Mobile Router Operation It corresponds to Mobile IPv6's MN Operation, enhanced with the ability to send PSBUs to the HA and arbitrary CNs corresponding with MNNs. Like Mobile IPv6's MNs, the MR obtains a care-of address on each subsequent point of attachment of its egress interface. Unlike Mobile IPv6's BU, the PSBU includes the Mobile Network Prefix Sub-Option. No changes are required to the standard MN Operation.

mobile router's home address and mobile router's care-of address, a binding between the NEMO Prefix and the mobile router's care-of address is needed.

Extended CN Operation and HA Operation Those are extended to process the Mobile Network Prefix Sub-Option. Before sending a datagram, the CN checks if the prefix of the destination address matches the NEMO Prefix recorded in the Binding Cache. Packets bearing a destination address that matches the NEMO Prefix are transmitted via the MR's care-of address using a Routing Extension Header. The HA Operation is enhanced in the same way. It only requires that the HA is able to encapsulate packets to the MR's care-of address for all destinations addresses that match the NEMO Prefix, and not only those that match solely MR_{ip}.

Evaluation PSBU ensures permanent connectivity to the Internet, session continuity between arbitrary CNs and MNNs, optimal routing between CNs and LFNs, and preserves MNNs from individual registrations (mobility support transparency). All the mobility management burden is entirely left to the MR, the node that effectively changes its point of attachment. The number of BUs is thus minimized by regarding MNNs collectively and scales to the size of the mobile network. This limits bandwidth consumption between MR and HA and avoids a CN to receive and register duplicate bindings. *Authentication* of PSBUs is performed as for standard BUs: the mobile router is the sender of the PSBU and is authenticated as such. Therefrom, recipients of PSBUs are not misled by the identity of the sender since the mobile router does not send PSBUs on behalf of its MNNs. However, this faces *authorization* concerns. Currently, PSBU doesn't consider multi-homed nor nested mobile networks. Other motivations and trade-offs that led to its design are explained in [8].

4.3. *MADRID Implementation*

MADRID is implemented into three distinct modules: the Header Processing unit, the Binding Manager unit, and CoA Manager unit. The Header Processing unit generates the Mobile IPv6 header. The Binding Manager unit provides an interface to manage the Binding Cache and the Binding Update List. The CoA Manager unit deals with movement detection and is only implemented on the MR. Fig.7 shows the modules on the MR entity. Two caches are actually used for convenience (specifications are still under development) and ease of implementation (less coding): one for the binding between MR's home address (MR_{ip}) and its care-of address (MR_{coa}), and one between a LFN's permanent address (LFN_{ip}) and MR_{coa}.

Initialization The mobile network is configured with a unique network prefix ($Prefix1$, 64 bits). MR's egress interface MR_{ip} is configured with the home prefix whereas the ingress interface is configured with $Prefix1$. The MR is operating PSBU.

MR registration with HA Each time the MR attaches to a new link, the egress interface is configured with a new care-of address MR_{coa}. It then sends a PSBU

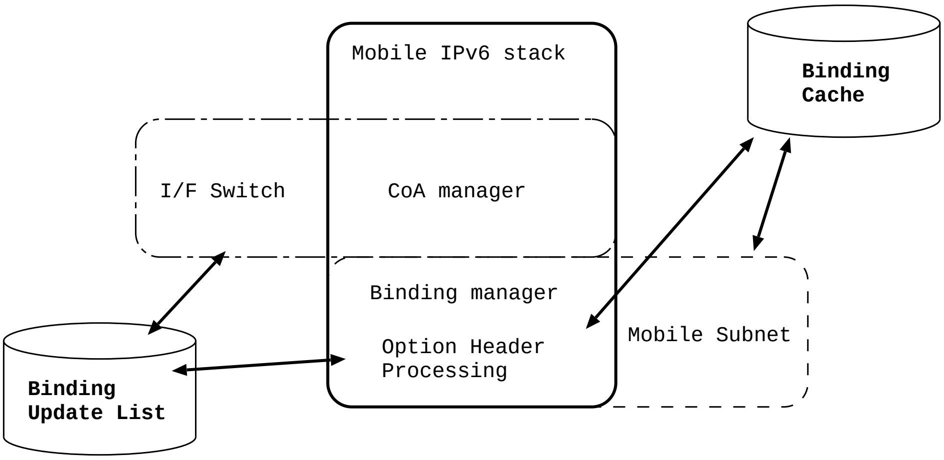

Fig. 7. MR Implementation

to its HA. The PSBU contains a BU Option with the Mobile Network Prefix Sub-Option (set to $Prefix1$), a Home Address Option (set to MR_{ip}), but no AH nor ESP header for authentication. The PSBU is sent at periodic time intervals to refresh the lifetime in the cache.

CNaddr1	Haddr1	Other Values
CNAddr2	Haddr2	Other Values

Fig. 8. MR: BU List

Prefix1	length	Haddress1

Fig. 9. MR: PSBU List

MR receiving packet on its egress interface Packets intended to LFN_{ip} are received on MR's egress interface. They are either encapsulated from HA to MR_{coa} or transmitted directly from CN to LFN_{ip} via MR_{coa}, using a Routing Extension Header. In the former case, packets are de-tunneled, a PSBU is sent to the CN and its address included in the BU List; in the latter, the Routing Extension Header is processed. In both cases, the packet is transmitted on the ingress interface toward the LFN.

MR receiving packet on its ingress interface Packets sent from LFNs are received on MR's ingress interface. The Binding Update List and the PSBU List are searched. If a match is found, the packet is encapsulated. As shown on Fig.8 and 9, if MR receives a packet on its ingress interface with $Prefix1$; the PSBU list is searched and $Haddress1$ is returned. Then, the MR searches CN's address in the Binding Update List which is corresponding to $Haddress1$.

CN receiving packet from LFN If the next protocol header is IPv6, the Binding Cache and the PSBU Cache are searched for an entry corresponding to the source address. If a match is found, the packet is decapsulated. If not, the packet is silently discarded. If the packet contains a PSBU with the P flag set and the Mobile Network Prefix Sub-Option, a Binding Cache entry and a PSBU Cache entry are created as illustrated on Fig.10 and 11. The binding between MR_{ip} and MR_{coa} is registered in the Binding Cache whereas the prefix is recorded in the PSBU Cache.

Haddress1	CoA1	Other Values
Haddress2	CoA2	Other Values

Fig. 10. CN: Binding Cache

Prefix1	length	Haddress1

Fig. 11. CN: PSBU Cache

CN sending a packet to a LFN The packet is constructed and the destination address is set to LFN_{ip}. The Binding Cache is then searched for the destination address LFN_{ip}. If no match is found, the PSBU Cache is searched for the longest match. If again no match is found, the packet is transmitted and gets routed up to MR's home link. On the home link, packets are intercepted by HA and tunneled to MR_{coa}. On the other hand, if a prefix matches LFN_{ip}, MR_{ip} is returned. The corresponding MR_{coa} is then searched in the Binding Cache. An IPv6 Routing Extension Header is then appended to the packet, LFN_{ip} moved into it, and MR_{coa} inserted in the IPv6 Header. The pending packet therefore gets directly routed to

MR_{coa} where the IPv6 Routing Extension Header is processed. As illustrated on Fig.10 and 11, the Binding Cache is searched for the destination $Prefix1 :: beef$. There is no match. $Prefix1$ is then searched in the PSBU Cache where a record is found. $HAddress1$ is returned, thus the Binding Cache is searched for the care-of address corresponding to $HAddress1$. $CoA1$ is returned.

5. Conclusion

This paper detailed the networking requirements for connecting vehicles to the Internet, as investigated in the InternetCAR project. The need to displace an entire IPv6 network and network mobility support were particularly exhibited. As such, we contributed in setting up the IETF NEMO working group which was recently created to deal with the specific issues raised by network mobility. To demonstrate network mobility in the transportation industry, we designed a testbed and we based our implementation on PSBU. Our implementation must be revised in order to comply with the NEMO Basic Support specification [7] recently released by the NEMO working group. This working group proposal inherits from PSBU and later proposals to support network mobility. It allows much more complex configurations of mobile networks, particularly nested mobile networks and multi-homed mobile networks. Our current work is targeted toward the validation of the network mobility support specification and multi-homing capabilities which are necessary to achieve the goals of the automobile industry and the ubiquitous Internet. For this purpose, we have recently set up the NAUTILUS [3] working group within the WIDE [17] organization and we are actively contributing to the NEMO working group.

Acknowledgments

Authors are grateful to InternetCAR team members from Keio University who also contributed in the experiments and discussions that led to this paper. We are also indebted to INRIA Rhône-Alpes (Grenoble, France, PLANETE team) and Motorola Labs (Paris, France, NAL team) for the first author's earlier work conducted there in network mobility before he moved away to join Keio University. At Keio University, Thierry Ernst was initially financed by a post-doctoral research fellowship granted by INRIA (French National Institute for Research in Computer Science and Control) while writing earlier versions of this paper.

References

1. InternetCAR web page at WIDE, May 2002. http://www.sfc.wide.ad.jp/InternetCAR.
2. MADRID: NEMO implementation at Keio University SFC, May 2002. http://www.sfc.wide.ad.jp/NEMO/.
3. NAUTILUS working group at WIDE: web page, April 2003 http://www.nautilus6.org.
4. NEMO web page and mailing list, June 2002. http://www.nal.motlabs.com/nemo.
5. SFC-MIP: Mobile IPv6 Implementation at Keio University SFC, May 2002.

http://www.sfc.wide.ad.jp/MIP6/.

6. S. Deering and R. Hinden. Internet Protocol Version 6 (IPv6) Specification. Request For Comments 2460, IETF, December 1998.
7. V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert. Nemo Basic Support Protocol Internet Draft draft-ietf-nemo-basic-support-00.txt, IETF, June 2003. Work in progress.
8. T. Ernst. *Network Mobility Support in IPv6.* PhD thesis, Université Joseph Fourier, October 2001. http://www.inria.fr/rrrt/tu-0714.html.
9. T. Ernst. Network Mobility Support Requirements. Internet Draft draft-ietf-nemo-requirements-01.txt, IETF, May 2003. Work in progress.
10. T. Ernst and H.-Y. Lach. Network Mobility Support Terminology. Internet Draft draft-ietf-nemo-terminology-00.txt, IETF, May 2003. Work in progress.
11. T. Ernst, A. Olivereau, L. Bellier, C. Castelluccia, and H.-Y. Lach. Mobile Networks Support in Mobile IPv6 (Prefix Scope Binding Updates). Internet Draft draft-ernst-mobileip-v6-network-03.txt, IETF, March 2002. Expired.
12. T. Ernst and K. Uehara. Connecting Automobiles to the Internet. In *ITST: 3rd International Workshop on ITS Telecommunications*, Seoul, South Korea, November 2002.
13. D. B. Johnson, C. Perkins and J. Arkko. Mobility Support in IPv6. Internet Draft draft-ietf-mobileip-ipv6-24.txt, IETF, June 2003. Work in progress.
14. K. Mitsuya, K. Uehara, and J. Murai. The In-vehicle Router System to support Network Mobility. In *International Conference on Information Networking (ICOIN)*, South Korea, Feb. 2003.
15. C. Perkins and D. B. Johnson. Mobility Support in IPv6. In *Proc. of the Second ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 27–37, Rye, New-York, USA, November 1996. IBM Corporation and Carnegie Mellon University.
16. K. Uehara, Y. Watanabe, H. Sunahara, O. Nakamura, and J. Murai. InternetCAR - Internet Connected Automobiles. In *INET*, July 1998. http://www.sfc.wide.ad.jp/InternetCAR/.
17. WIDE: Widely Integrated Distributed Environment, http://www.wide.ad.jp.