

NEMO-enabled Localised Mobility Support for Internet Access in Automotive Scenarios

Ignacio Soto*, Carlos J. Bernardos*, Maria Calderon*, Albert Banchs*, Arturo Azcorra*[†]

* University Carlos III of Madrid

Email: {isoto, cjbc, maria, banchs, azcorra}@it.uc3m.es

[†]IMDEA Networks

Email: arturo.azcorra@imdea.org

Abstract—This paper overviews major existing approaches and proposes a novel architecture to support mobile networks in network-based localised mobility domains. Our architecture enables conventional terminals without mobility support to obtain connectivity either from fixed locations or mobile platforms (e.g., vehicles) and move between them, while keeping their ongoing sessions. This functionality allows to offer broadband Internet access in automotive scenarios such as public transportation systems, where users spend time both in vehicles and stations. The key advantage of our proposal, as compared with current alternatives, is that the described mobile functionality is provided to conventional IP devices that lack any mobility functionality, but we have also performed an experimental evaluation of our proposal that shows that our architecture improves the quality perceived by the end-users.

I. INTRODUCTION

Nowadays, users increasingly demand Internet access everywhere. The current trend in hand-held terminals is towards devices that move away from the traditional phone service model and incorporate a large number of different data applications. Equipping terminals with multiple technologies – for example 3G and WLAN – is a widely used solution to provide ubiquitous Internet access. Internet access in automotive scenarios is a particularly relevant case, especially because people in modern cities spend a lot of time in vehicles. Although 3G is a possible option, it suffers from a number of drawbacks, such as capacity constraints from the point of view of the operator, as well as cost issues from the end-user perspective.

In the above context, there is a need for an alternative solution to 3G that provides efficient broadband Internet access in automotive scenarios. Public transportation systems, such as undergrounds, suburban trains and city buses, represent one relevant scenario because of the large number of users and the time spent by them both in vehicles and stations. In fact, communications in these environments are receiving a lot of attention by a number of research and standardisation activities¹. Other relevant scenarios with similar requirements

are those in which users move around big areas (e.g., airports, exhibition sites or fairgrounds). In these areas, attachment points to the Internet may be available both at fixed locations (such as coffee shops or airport terminals) but also in mobile platforms such as vehicles (e.g., buses to change pavilion in a fair or a train to change terminal in an airport). Users demand to keep their ongoing communications while changing their point of attachment to the network as they move around (e.g., when a user leaves a coffee shop and gets on a bus).

NETwork MObility (NEMO) solutions are currently being developed by the IETF and the research community to offer Internet access from vehicles. Special devices (called Mobile Routers) located in the vehicles take care of the communication with the fixed infrastructure and provide access to passengers' devices using a convenient short-range radio technology. However, in the scenarios mentioned above, users spend only part of their time in the vehicles, since they also move from vehicles to fixed platforms (e.g., the stations in the public transportation scenario or the terminals in the airport scenario). Therefore, an integrated solution for these scenarios, that considers Internet access not only from vehicles but also from associated fixed platforms, is a better approach.

Traditional IP mobility mechanisms (see section II) [1], [2] have been based on functionality residing both in the moving terminals and in the network. Lately, there is a new trend towards solutions that enable mobility of IP devices within a local domain with only the support from the network. This approach, called Network-based Localised Mobility Management (NetLMM) [3], allows conventional IP devices to benefit from this mobility support. This is very interesting from the point of view of operators, because it allows them to provide mobility support without depending on software and complex mobility related configuration in the terminals. The IETF has standardised Proxy Mobile IPv6 (PMIPv6) [4], a protocol to provide this functionality. But this solution has the limitation of not fully supporting mobile networks.

In this paper we propose a novel architecture (see section III), called NEMO-enabled PMIPv6 (N-PMIPv6), which fully integrates mobile networks in PMIPv6 localised mobility domains. With our approach, users can obtain connectivity either from fixed locations or mobile platforms (e.g., vehicles) and can move between them, while keeping their ongoing sessions. N-PMIPv6 architecture exhibits two remarkable characteristics. First, N-PMIPv6 is totally network-based – therefore no

¹Some examples are the work in the IETF MEXT WG (<http://www.ietf.org/html.charters/mext-charter.html>), the extension by the ETSI Technical Committee Railways Telecommunications (http://portal.etsi.org/rt/summary_06.asp) of the original GSM-R standard to benefit from the evolution of the GSM technology, or the PATH initiative (<http://www.path.berkeley.edu/PATH/Research/currenttransit.html>) that among other goals conducts research in technologies for innovating and enhancing public transportation solutions.

mobility support is needed in the terminals –, and second, the handover performance is improved (see section IV), both in terms of latency and signalling overhead.

II. OVERVIEW OF MAJOR EXISTING APPROACHES

This section provides an overview of existing mechanisms developed by the IETF that are relevant to provide Internet access in vehicular environments. Operators have shown a great interest in network-based localised mobility solutions. Additionally, NEMO approaches are a key element to provide connectivity from vehicles. Combining both brings the advantages of network-based localised mobility solutions to vehicular scenarios. This section reviews the work of the IETF in this area and highlights the limitations of current solutions.

A. Network-based localised mobility

Unlike host-based localised mobility [1], where Mobile Terminals (MTs) signal a location change to the network to update routing states, Network-based Localised Mobility Management (NetLMM) [3] approaches provide mobility support to moving hosts without their involvement. This is achieved by relocating relevant functionality for mobility management from the MT to the network. In a Localised Mobility Domain (LMD), the network learns through standard terminal operation, such as router and neighbour discovery or by means of link-layer support, about an MT's movement and coordinates routing state updates without any mobility specific support from the terminal. While moving inside the LMD, the MT keeps its IP address, and the network is in charge of updating its location in an efficient manner. Proxy Mobile IPv6 (PMIPv6) [4] is the NetLMM protocol proposed by the IETF. This protocol is based on Mobile IPv6 (MIPv6) [2] – it extends MIPv6 signalling messages and reuses the Home Agent (HA) concept.

The core functional entities in the PMIPv6 infrastructure are (see Fig. 1):

- **Mobile Access Gateway (MAG).** This entity performs the mobility related signalling on behalf of an MT that it is attached to its access link. The MAG is usually the access router for the MT, i.e. the first hop router in the Localised Mobility Management infrastructure. It is responsible for tracking the MT's movements in the access link. There are multiple MAGs in an LMD.
- **Local Mobility Anchor (LMA).** This is an entity within the backbone network that maintains a collection of routes for individual MTs within the LMD. The routes point to MAGs managing the links in which the MTs are currently located. Packets for an MT are routed to and from the MT through tunnels between the LMA and the corresponding MAG.

Once an MT enters an LMD and attaches to an access link, the MAG in that access link, after identifying the MT, performs mobility signalling on behalf of the MT. The MAG sends to the LMA a Proxy Binding Update (PBU) associating its own address with the MT's identity (e.g., its MAC address or an ID related with its authentication in the network). Upon receiving this request, the LMA assigns a prefix to the MT. Then, the

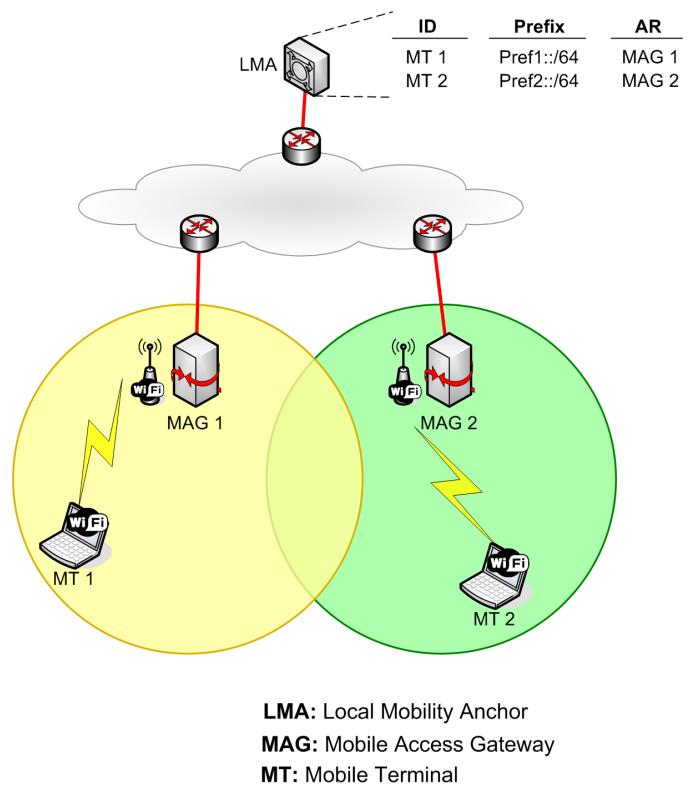


Fig. 1. Proxy Mobile IPv6 domain

LMA sends to the MAG a Proxy Binding Acknowledgement (PBA) including the prefix assigned to the MT. It also creates a Binding Cache entry and establishes a bi-directional tunnel to the MAG. Whenever the MT moves, the new MAG updates the MT's location in the LMA and advertises the same prefix to the MT (through unicast Router Advertisement messages) thereby making the IP mobility transparent to the MT. The MT can keep the address configured when it first entered the LMD, even after changing its point of attachment within the network.

B. Network Mobility support

To address the requirement of transparent Internet access from vehicles, the IETF standardised the NEMO Basic Support (NEMO B.S.) protocol [5]. This protocol defines a Mobile Network (or Network that MOves, NEMO) as a network whose attachment point to the Internet varies with time. The router within the NEMO that connects to the Internet is called the Mobile Router (MR). It is assumed that the NEMO has a Home Network where it resides when it is not moving. Since the NEMO is part of the Home Network, the Mobile Network has configured addresses belonging to one or more address blocks assigned to the Home Network: the Mobile Network Prefixes (MNPs). These addresses remain assigned to the NEMO when it is away from home, although they only have topological meaning when the NEMO is at home. So, when the NEMO is away from home, packets addressed to the Mobile Network Nodes (MNNs) will be still routed to the Home Network. Additionally, when the NEMO is away from home, i.e. it is in a visited network, the MR acquires an address

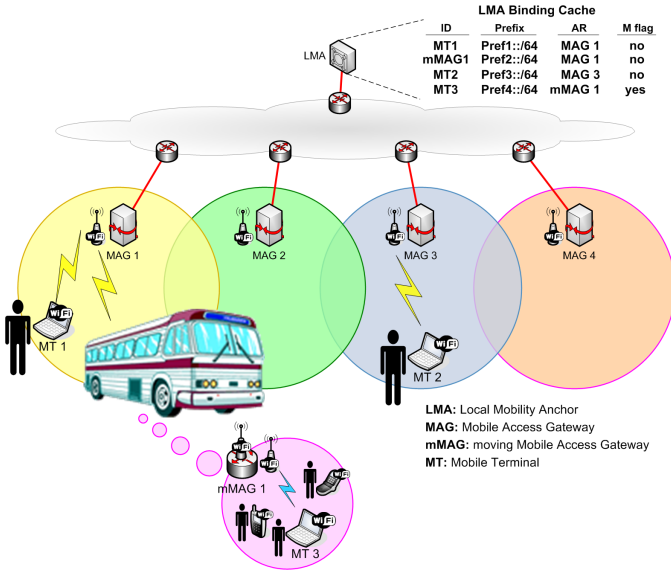


Fig. 2. Architecture overview of an N-PMIPv6 domain

from the visited network, called the Care-of Address (CoA), where the routing infrastructure can deliver packets without additional mechanisms.

The basic solution for network mobility support is quite similar to the solution proposed for host mobility (Mobile IPv6 [2]) and essentially creates a bi-directional tunnel between a special node located in the Home Network of the NEMO (the Home Agent, HA), and the CoA of the MR. Route Optimisation support is currently being researched, with special attention to the requirements of the vehicular scenario (see e.g., [6]).

C. Current solution for combining NEMO and PMIPv6

Both Network Mobility and Network-based Localised Mobility Management solutions provide some interesting features that could be combined in an integrated architecture. Nowadays, it is possible to partially benefit from the following advantages by using NEMO B.S. and PMIPv6:

- Transparent network mobility support: MRs take care of the management of the mobility of a network composed by a set of devices moving together.
- Transparent localised mobility support without node involvement: MRs and MTs can roam within a PMIPv6 domain without changing their IP addresses.

Although current mechanisms (i.e. NEMO B.S. and PMIPv6) can be combined to bring the advantages described above, this is not a full integration, since an MT cannot roam between an MR and a MAG of the fixed infrastructure without changing its IP address. This is because the addresses used within the mobile network belong to the Mobile Network Prefix and not to the prefixes used by PMIPv6. This means that in order to support – in a transparent way – MTs roaming between MRs and MAGs without any restriction, MTs are required to run MIPv6 to manage mobility (that is, the change of IP address) by themselves. If MTs have to use MIPv6, the mobility support provided within the PMIPv6 domain is no

longer fully network-based, since some mobility operations are performed by MTs.

III. N-PMIPv6 ARCHITECTURE

In this section we propose a novel architecture that overcomes the shortcomings identified in the previous section for the current solution for NEMO support in PMIPv6. Our architecture, called NEMO-enabled PMIPv6 (N-PMIPv6), enables a seamless and efficient integration of Mobile Networks within a Network-based Localised Mobility Management solution based on Proxy Mobile IPv6 (PMIPv6) without adding any extra mobility support on terminals (i.e. mobility is totally managed by the network) and improving handover performance. First, an overview of the architecture is provided and subsequently its operation is presented in further detail.

A. Overview

The key idea of N-PMIPv6 consists in extending the PMIPv6 domain to include also mobile networks. Both the fixed infrastructure (i.e. MAGs) and the mobile networks (i.e. MRs) belong to the same network operator. With N-PMIPv6, an MT attached to a mobile network is also part of the PMIPv6 domain. Hereinafter, we refer to an N-PMIPv6 enabled LMD as N-PMIPv6 domain. This enables conventional IP nodes to roam between fixed MAGs and also between fixed MAGs and MRs, without changing the IPv6 addresses they are using. As a result, the handover-related signalling load is reduced and the handover performance (i.e. the associated latency) is improved, when compared to traditional global IP mobility solutions (e.g., MIPv6).

While the NEMO B.S. protocol requires MRs to manage their own mobility, this is not required in N-PMIPv6, in the same way as N-PMIPv6 does not require mobility related functionality in MTs. This is a consequence of the fact that mobility of MRs and MTs in N-PMIPv6 is managed by the network (i.e. it is network-based). With N-PMIPv6, MTs do not require any additional functionality, and MRs require functionality to extend the PMIPv6 domain to mobile networks so that an MT that attaches to a mobile network does not need to change its IPv6 address. Since MRs in N-PMIPv6 perform similar functions to MAGs in PMIPv6, while being mobile, hereinafter we refer to them with the name moving MAGs (mMAGs).

mMAGs extend the PMIPv6 domain by providing IPv6 prefixes belonging to this domain to attached MTs and by forwarding their packets through the LMA. The basic operation of an mMAG is as follows. When an mMAG attaches to a fixed MAG, this fixed MAG informs its LMA about this event, by sending a PBU message that contains the mMAG's identity. The LMA then delegates an IPv6 prefix to the mMAG and creates a binding cache entry, associating the mMAG's identity with the delegated prefix and the fixed MAG to which the mMAG is attached. If the mMAG moves to another fixed MAG, the LMA updates the binding with the information of the new MAG. Note that this is basically the PMIPv6 behaviour when a conventional MT connects to a PMIPv6

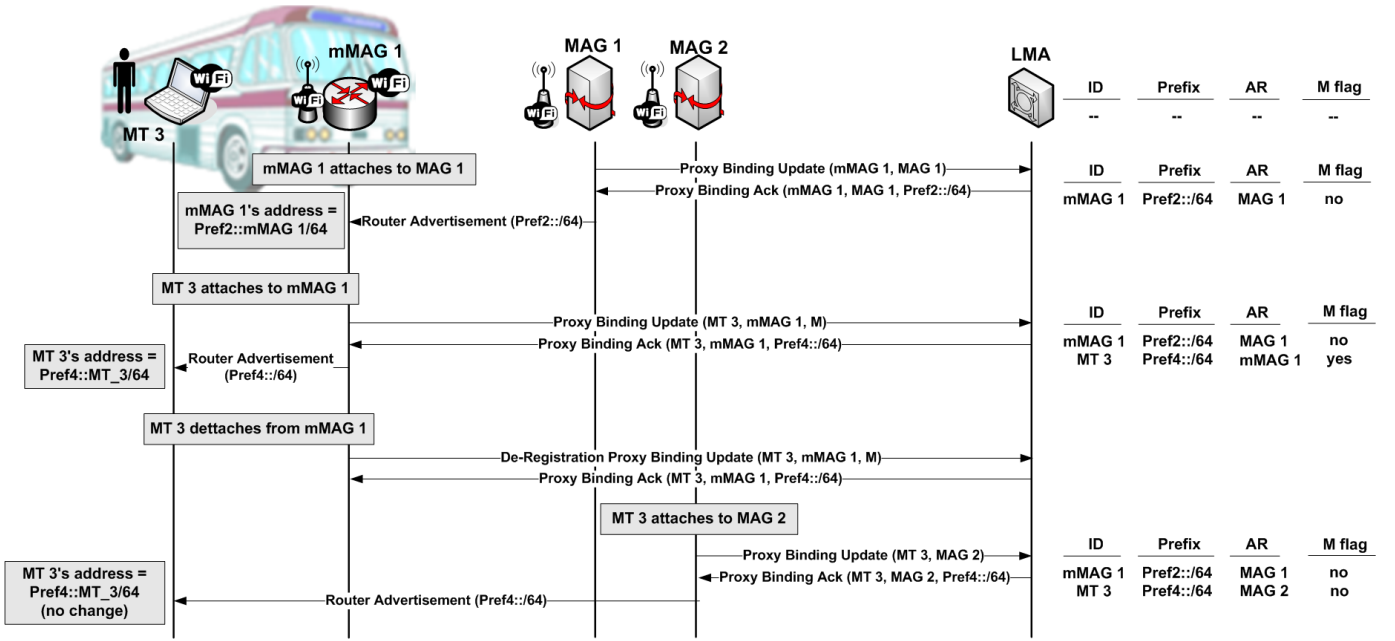


Fig. 3. Detailed operation signalling

MAG, i.e. our architecture manages the mobility of an mMAG in the same way as PMIPv6 manages the mobility of an MT.

From the point of view of an MT that attaches to an mMAG, this mMAG behaves as a fixed MAG of the N-PMIPv6 domain. In particular, when an MT attaches to an mMAG, the mMAG informs the LMA and, following PMIPv6 procedures, gets an IPv6 prefix for the MT. The LMA then adds a new binding cache entry, associating the MT's ID with the delegated prefix and the MAG IPv6 address to which it is attached (i.e. the mMAG address). The LMA cannot accept requests for these kinds of operations from any node, only from authorised MAGs. This implies that mMAGs must have a security association with the LMAs to be able to operate in the N-PMIPv6 domain. The way this association is created is out of the scope of this article, but notice that it is not different from the security association needed with any fixed MAG. This basically means that for practical purposes we assume scenarios in which the mMAGs, the fixed MAGs and the LMA belong to the same administrative domain, as it would be the case in the automotive scenarios described in the introduction.

In order to deliver IPv6 packets addressed to an MT attached to a connected mMAG, a change in the normal operation of a PMIPv6 LMA is introduced. Specifically we extend LMA functionality to support recursive look-ups in its binding cache as follows. In a first look-up, the LMA obtains the mMAG to which the MT is attached. After that, the LMA performs a second look-up searching for this mMAG in its binding cache, and finds the associated fixed MAG. With this information, the LMA is able to encapsulate the received packets towards the mMAG, through the appropriate fixed MAG. The mMAG is then able to forward data packets to the MT. Two nested tunnels are used to encapsulate data packets between the LMA and the mMAG: one between the LMA and the mMAG, and another one between the LMA and the fixed MAG. A new

field, called mMAG (M) flag, is added to the binding cache used by the LMA, in order to support recursive look-ups. The entries in the binding cache created/updated by PBUs received from mMAGs have the M flag set to 'yes'. On the other hand, entries created/updated by PBUs received from fixed MAGs have the M flag set to 'no'. The use of this flag avoids the LMA performing unnecessary recursive look-ups in its binding cache.

B. Detailed operation

This section describes in more detail the operation of N-PMIPv6 architecture, using the network scenario that appears in Fig. 2 and the signalling sequence depicted in Fig. 3.

When an mMAG – mMAG 1 – attaches to a fixed MAG – MAG 1 –, this event is detected by MAG 1 and reported to its serving LMA, by means of a PBU message. If no existing entry for mMAG 1 is found in the LMA binding cache, the LMA assigns an IPv6 prefix to the mMAG 1 ($Pref2::/64$), and creates a new entry in the cache. This entry includes the information of the assigned IPv6 prefix and the IPv6 address of the fixed MAG to which mMAG 1 is attached (i.e. MAG 1). The LMA then replies with a PBA message, that includes the IPv6 prefix assigned to mMAG 1 ($Pref2::/64$). With this information, MAG 1 sends a unicast Router Advertisement (RA) message to mMAG 1, so it can form an IPv6 address and start sending/receiving traffic. While the mMAG moves within the same domain – roaming between different fixed MAGs – its IPv6 address does not change.

When an MT – MT 3 – attaches to mMAG 1, mMAG 1 sends a PBU message towards the LMA, which assigns an IPv6 prefix to MT 3 ($Pref4::/64$) and creates a new entry for this MT in its binding cache, setting the M flag of this

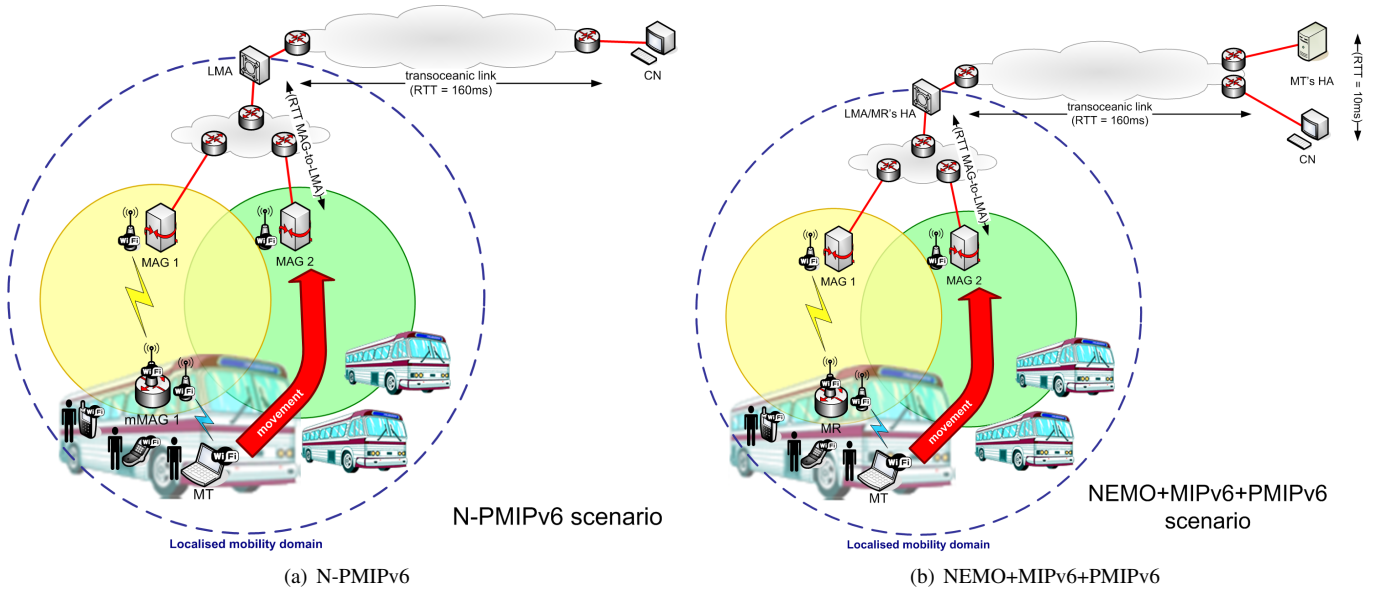


Fig. 4. Analysed scenarios

entry to 'yes'². The LMA then provides mMAG 1 with the assigned prefix. Finally, mMAG 1 informs MT 3 about the IPv6 prefix it has to use, by sending a unicast RA to the MT.

To hide the network topology and avoid changing the particular prefix assigned to an mMAG or an MT while they roam within the same domain, IP bi-directional tunnelling is used. Following our example, if the LMA receives a packet from a CN addressed to MT 3, it performs a recursive look-up at its binding cache. As a result of this look-up, the packet is sent through a nested tunnel, the inner header with source address set to the LMA and destination address the mMAG 1, and the outer header with source address the LMA and destination address the MAG 1. The outer header brings the packet to MAG 1, that then removes that header. Next, the inner header brings the packet to the mMAG 1. Finally, mMAG 1 removes the inner header and delivers the packet to MT 3.

If MT 3 performs an intra N-PMIPv6 domain handover from mMAG 1 to MAG 2 (see Fig. 3), MAG 2 informs the LMA, so it can update the binding entry accordingly (now MT 3 is attached to MAG 2, instead of mMAG 1, and the M flag is set to 'no'). The mMAG 1, upon detecting disconnection of MT 3, sends a De-Registration PBU (a PBU with the lifetime value of zero) to its LMA, following standard PMIPv6 operation. If the LMA does not receive any PBU about MT 3 after a pre-configured amount of time, the binding entry is deleted, in order to avoid stale state at the LMA binding cache.

C. Scalability of the solution

An additional advantage of our proposal as compared with PMIPv6 is that it increases the scalability. This is because mMAGs concentrate MTs, so when a vehicle moves, instead

of having a number of individual MTs changing their point of attachment to the network, with a control message per MT sent by the MAG to the LMA, we will have just one control message sent by the MAG to the LMA indicating the movement of the mMAG. The cost, from the point of view of the scalability, is having more entries (one per mMAG) in the binding cache of the LMA, but this is not a problem as it is always possible to distribute the LMA function among different nodes in the network.

IV. PERFORMANCE EVALUATION

In this section we evaluate the performance improvement achieved with N-PMIPv6 when compared against the existing approach for NEMO support in PMIPv6 domains (NEMO+MIPv6+PMIPv6) described in Section II-C. While the main benefit of N-PMIPv6 over NEMO+MIPv6+PMIPv6 is that N-PMIPv6 does not require mobility support on terminals, in this section we show that this benefit not only comes at no performance penalty but actually N-PMIPv6 provides better performance than NEMO+MIPv6+PMIPv6.

Fig. 4 shows the two scenarios we consider in this section. The left part shows an N-PMIPv6 domain consisting of two MAGs, one LMA, one mMAG and one MT. The right part shows a network deployment of the NEMO+MIPv6+PMIPv6 approach, consisting of two MAGs, one LMA, one MR and its HA (called MR's HA), and finally, one MT and its HA (MT's HA). In both scenarios, there is a Correspondent Node (CN) located on the Internet communicating with the MT.

From the point of view of performance, the key advantage of N-PMIPv6 over NEMO+MIPv6+PMIPv6 is that, upon executing an MT's handover to or from a mobile network, the corresponding signalling is only sent to the LMA, as opposed to NEMO+MIPv6+PMIPv6 which requires signalling down to the MT's HA. This results in a reduction of the signalling load in the backbone, as well as shorter handover latencies.

²In order to enable the LMA to know which is the value that the M flag of an entry should have, we extend the PBU message so it contains a new M flag (carrying this information). Only PBUs sent by mMAGs have this M flag set.

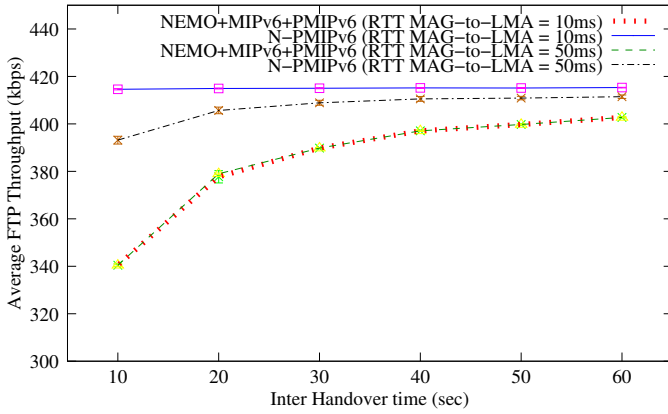


Fig. 5. FTP throughput obtained by N-PMIPv6 compared with NEMO+MIPv6+PMIPv6

In case of an mMAG/MR handover, since mobility is managed by PMIPv6 (i.e. the location of the mMAG/MR is updated at the LMA by the MAG to which the mMAG/MR is attached, and no further signalling is needed) in both N-PMIPv6 and NEMO+MIPv6+PMIPv6 solutions, the handover performance is the same.

In this section we concentrate on the performance analysis for the case of the mobile terminal handover, since this is the only case in which the performance of both approaches differs.

In the NEMO+MIPv6+PMIPv6 scenario, the mobile terminal and its HA are separated by a transoceanic link, in order to understand the impact of long RTTs on performance. The MT is communicating with a CN that is topologically close to the MT's HA. The N-PMIPv6 scenario is equivalent in terms of functionality and relevant network entities' location. The LMA of the N-PMIPv6 scenario is located in the same place that the MR's HA in the NEMO+MIPv6+PMIPv6 scenario in order to carry out a fair comparison. The location of MR's HA has an impact on the end-to-end delay of data traffic, since every packet sent by a node attached to the MR has to traverse the MR's HA (i.e. there is not standardised NEMO Route Optimisation solution yet).

We estimate the MT's handover latency for both N-PMIPv6 – handovers from a mMAG to a MAG, or vice versa – and NEMO+MIPv6+PMIPv6 – handovers from MAG to MAG. We assume that in the NEMO+MIPv6+PMIPv6 case, the MT is performing MIPv6 Route Optimisation (RO) with the CN, so data packets do not traverse the MT's HA. The MT's handover latency can be estimated for this case following [7], according to which latency is approximately equal to one MT-HA RTT plus one MT-CN RTT, which is roughly two MT-HA RTTs (we take the RTT measurements of [8]), because of the Return Routability signalling required to perform RO with the CN. For the N-PMIPv6 case, the handover latency is approximately just one mMAG-LMA RTT (for the case of an MT's handover from a fixed MAG to an mMAG, or one MAG-LMA RTT, for the case of a handover from an mMAG to a fixed MAG), since updating the LMA with the new location of the MT is the only required signalling. We further consider a frequency of handovers ranging from one handover every 10 seconds (highly dynamic scenarios) to one handover every 60 seconds

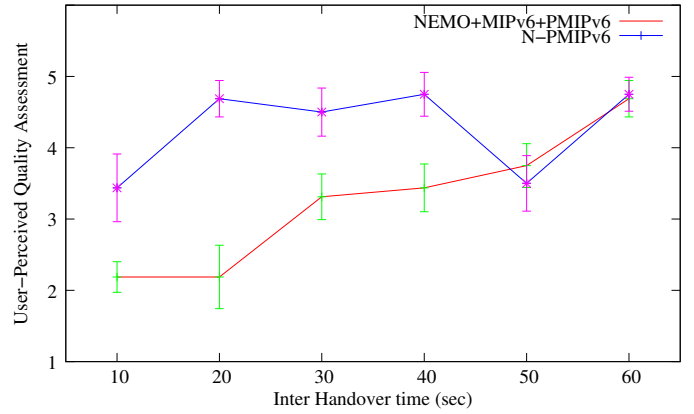


Fig. 6. User-Perceived Video Quality Assessment

(slowly changing scenarios).

We first analysed the performance of a TCP data transfer by measuring the average throughput experienced when transferring a 20 MB data file from the CN to the MT. Experiments were performed via simulation with the OPNET tool³. Two different values of RTT between the LMA and the MAGs ($RTT_{MAG-to-LMA}$) were used in the simulations: 10 ms (usual case) and 50 ms (extreme case). This allowed us to evaluate the impact of the size of the N-PMIPv6 domain on the overall performance. The results obtained from the experiments with our approach and with NEMO+MIPv6+PMIPv6 are illustrated in Fig. 5. It can be observed that N-PMIPv6 improves the average throughput. Indeed, with NEMO+MIPv6+PMIPv6 each handover causes a severe interruption due to the latency associated with the signalling, degrading thus TCP performance. With N-PMIPv6, interruptions are much shorter, since only local signalling is required, and as a result handovers do not degrade the throughput performance of TCP as much as in the case of NEMO+MIPv6+PMIPv6.

The second application whose performance we analysed is video streaming, in particular VLC⁴ which transmits video over RTP/UDP. The performance of this application was evaluated by means of real-life experiments with the following setup. Video was streamed from one PC to another crossing a third PC. The iptables⁵ software was configured in this third PC to introduce interruptions of a duration and frequency equal to the ones caused by handovers (for the usual case).

We conducted experiments with 16 real users that assessed the subjective video quality they perceived for each experiment. Following ITU recommendations for the subjective evaluation of video and audio quality [9], [10], we asked the users to rate the quality of each video with a scale ranging from 5 (excellent quality) to 1 (bad quality). Fig. 6 depicts the results obtained, in terms of average subjective quality and 95% confidence intervals.

The obtained results show that N-PMIPv6 clearly outperforms NEMO+MIPv6+PMIPv6, specially for highly dynamic environments (i.e. those in which an MT performs

³OPNET University Program, <http://www.opnet.com/services/university/>

⁴Videolan Client: <http://www.videolan.org/>

⁵<http://www.netfilter.org/>

handovers very often). It can be seen that there is one point in the figure (one handover every 50 seconds) where the subjective quality with N-PMIPv6 drops down to the level of NEMO+MIPv6+PMIPv6. The reason for this anomaly is that this particular experiment involved the unfortunate drop of some key packets that significantly degraded video quality despite the small number of lost packets. Nonetheless, results show that N-PMIPv6 performs significantly better due to the longer latency of NEMO+MIPv6+PMIPv6 handovers.

V. CONCLUSIONS

In this paper we provide an overview of major existing approaches to support mobile networks in Network-based Localised Mobility domains. Then, we propose N-PMIPv6, a novel architecture that extends these domains to include not only fixed points of attachment, but also mobile ones, achieving a better integration of mobile networks. N-PMIPv6, like PMIPv6, bases mobility support on network functionality, allowing thus conventional (i.e. not mobility-enabled) IP devices to change their point of attachment within a Localised Mobility Domain, without disrupting ongoing communications. As a result, N-PMIPv6 enables off-the-shelf IP devices to roam within the fixed infrastructure, attach to a mobile network and move with it, and also roam between fixed and mobile points of attachment, while keeping the same IP address.

A key scenario for our architecture is the provision of Internet access from urban public transportation systems, such as undergrounds, suburban trains and city buses. In these systems, providing connectivity from vehicles and stations is not the only requirement, since this connectivity also needs to be maintained while changing vehicles.

There exist protocols already defined by the IETF that could be combined to achieve a similar functionality to N-PMIPv6, although at the cost of introducing additional complexity at the user terminal. Furthermore, the experimental and simulation results given in this paper show that the performance of such a combination of protocols is substantially worse, from a user perspective, than with N-PMIPv6. Future plans include the implementation of N-PMIPv6 and the experimental evaluation of the state and processing overhead in the nodes of the architecture.

ACKNOWLEDGEMENT

We would like to thank the users that kindly participated in the video experiments. We also thank the anonymous reviewers of this paper for their valuable comments.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 214994. This work has also been partially supported by the Spanish Government through the POSEIDON project (TSI2006-12507-C03).

REFERENCES

- [1] H. Soliman, C. Catelluccia, K. E. Malki, and L. Bellier, "Hierarchical Mobile IPv6 Mobility Management (HMIPv6)," RFC 4140, August 2005.
- [2] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6," RFC 3775, June 2004.
- [3] J. Kempf, "Problem Statement for Network-Based Localized Mobility Management (NETLMM)," RFC 4830, April 2007.
- [4] S. Gundavelli, K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil, "Proxy Mobile IPv6," RFC 5213, August 2008.
- [5] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol," RFC 3963, January 2005.
- [6] C. J. Bernardos, I. Soto, M. Calderon, F. Boavida, and A. Azcorra, "VARON: Vehicular Ad-hoc Route Optimisation for NEMO," *Computer Communications*, vol. 30, no. 8, pp. 1765 – 1784, June 2007.
- [7] C. Vogt and M. Zitterbart, "Efficient and scalable, end-to-end mobility support for reactive and proactive handoffs in IPv6," *IEEE Communications Magazine*, vol. 44, no. 6, pp. 74–82, June 2006.
- [8] W. Matthews and L. Cottrell, "The PingER project: active Internet performance monitoring for the HENP community," *IEEE Communications Magazine*, vol. 38, no. 5, pp. 130–136, May 2000.
- [9] ITU-T P.800 Recommendation, "Methods for subjective determination of transmission quality," 1996.
- [10] ITU-R BT.500-11 Recommendation, "Methodology for the subjective assessment of the quality of television pictures," 2002.

PLACE
PHOTO
HERE

Ignacio Soto received a telecommunication engineering degree in 1993, and a Ph.D. in telecommunications in 2000, both from the University of Vigo, Spain. He was a research and teaching assistant in telematics engineering at the University of Valladolid from 1993 to 1999. In 1999 he joined University Carlos III of Madrid, where he has been an associate professor since 2001. He has published several papers in technical books, magazines and conferences, lately in the areas of mobility support in packet networks and heterogeneous wireless access

networks.

PLACE
PHOTO
HERE

Carlos J. Bernardos received a telecommunication engineering degree in 2003, and a Ph.D. in telematics in 2006, both from the University Carlos III of Madrid (UC3M), where he worked as a research and teaching assistant from 2003 to 2008. Since then he has worked as associate professor at UC3M. His Ph.D. thesis focused on Route Optimisation for Mobile Networks in IPv6 Heterogeneous Environments. His current work focuses on vehicular networks and IP-based mobile communication protocols. He has served as TPC chair of WEEDEV 2009.

PLACE
PHOTO
HERE

Maria Calderon is an Associate Professor in the Telematics Engineering Department of the University Carlos III of Madrid, Spain. She received the computer science engineering degree in 1991, and the PhD degree in computer science in 1996, both from the Technical University of Madrid (UPM), Spain. She has published over 25 papers in the fields of advanced communications, reliable multicast protocols, programmable networks, network mobility and IPv6 mobility, in outstanding magazines and conferences. Some of the recent European research projects in which she has participated are E-NEXT, LONG, GCAP, DAIDA-LOS and GEONET.

PLACE
PHOTO
HERE

Albert Banchs received his M.Sc. and Ph.D. degrees in telecommunications from the Technical University of Catalonia, Spain, in 1997 and 2002, respectively. He worked for ICSI in 1997, for Telefonica I+D in 1998, and for NEC Network Laboratories from 1998 to 2003. Since 2003 he is with University Carlos III of Madrid. A. Banchs is associated editor for IEEE Communications Letters and has been guest editor for IEEE Wireless Communications and Computer Networks. His research interests include performance evaluation and resource allocation of

wireless networks.

PLACE
PHOTO
HERE

Arturo Azcorra received his PhD in 1989 from UPM, Spain. In 1992 he graduated as MBA from Instituto de Empresa, one of the most prestigious Business School in the world. He is a full professor at the University Carlos III of Madrid since 1999. On 2006 he was appointed as Director of the international research institute IMDEA Networks. He has participated in a many European research projects since the 3rd Framework Program, and is currently the coordinator of the EU project CARMEN. He has published over 100 scientific papers in prestigious

international magazines and conferences.