

A Comparative Study of Network Mobility Paradigms

Pedro Vale Pinheiro^{1,2}, Shivam Jain³, Fernando Boavida²

¹ Centro de Informática, Universidade de Coimbra, Portugal

² Faculdade de Ciências e Tecnologia da Universidade de Coimbra, Portugal

³ Indian Institute of Technology Roorkee, India

vapi@uc.pt, shivam.jain.iitr@gmail.com, boavida@uc.pt

Abstract. Effective solutions for the mobility of whole networks are essential for ubiquitous Internet access. Different network mobility paradigms have different consequences in terms of signalling load, performance and scalability. In this paper, a comparison of three network mobility paradigms is presented, namely the legacy-compatible, the infrastructure-centric and the end-node-centric network mobility paradigms. The study is unprecedented not only because the three paradigms encompass most of the existing network mobility solutions, but also because it involves scenarios with up to 22,800 routers, 11,250 networks, more than 27,000 nodes (fixed and mobile) and up to 16 levels of network mobility nesting. The obtained results point to the very good potential of the end-node-centric network mobility paradigm.

Keywords: network mobility; route optimisation.

1 Introduction

Moving entire networks raises problems that are not trivial, due to the potentially large number of flows involved in such movement. The problem can be considerably more complex in nested scenarios, where mobile networks exist inside other mobile networks.

Several solutions for network mobility have been proposed in the recent past. Despite their differences in detail, they fall into one of three paradigm classes: i) legacy class (LC), with no impact on end-nodes in order to maximize compatibility with legacy networks and nodes; ii) infrastructure-centric class (IC), in which mobile routers take up the management of most of the mobility-related tasks on behalf of the end-nodes, and iii) end-node-centric class (EC), in which end-nodes are aware of their mobility condition and perform most of the mobility functions, making mobility as transparent as possible to the networks.

The NEMO Basic Support Protocol, specified in RFC 3963 [1], belongs to the first paradigm class. NEMO has been developed in order to provide a basic solution for network mobility, having as main requirement complete mobility transparency, that is, requiring no need for any modifications to the nodes.

The well-known limitations of NEMO – of which lack of route optimization is the main one – have led to a whole class of more elaborate network mobility proposals,

with the common characteristic that they rely on additional functionality on the part of the networking infrastructure (e.g., mobile routers, routers, agents, etc.) which, for instance, take up the burden of performing route optimization on behalf of the mobile nodes they give access to. These proposals – which we collectively designate as infrastructure-centric paradigm – aim at minimal impact on the end-nodes, as opposed to the network.

The third network mobility paradigm is based on mobility awareness on the part of end nodes. In this case, nodes have enough ‘intelligence’ to know if they are mobile or not and to react accordingly.

Naturally, each of the above-mentioned paradigms has different consequences on the performance, scalability and signalling load of each particular mobility solution. These translate into different behaviour in terms of round trip time, route optimisation time and signalling overhead. In order to clearly identify and assess these consequences, a large-scale comparative study has been performed and the results are presented in this paper.

A clear option to compare paradigms instead of comparing particular network mobility solutions has been made for several reasons. First, it is not practical to compare the large number of specific solutions that exist. Second, the differences between solutions belonging to the same paradigm class are minor and can only have very limited impact on the results. Last, but not least, comparing paradigms provides a more independent and general view on the consequences of architectural choices, as opposed to detailed and often irrelevant views of very particular proposals and/or implementations.

Another important aspect of the presented study is the fact that it has an unprecedented scale. Typically, proposals are accompanied by some form of assessment involving tens of nodes and few networks. In contrast, in the present case we are interested in analyzing the consequences of the mobility paradigms at a very large scale, in generalised mobility environments, such as the one foreseen for the future Internet, in which it is anticipated that a very large percentage of nodes and networks will be mobile.

Before presenting the simulation study, a brief description of the various network mobility paradigms and a comparative analysis are made in section 2. In section 3 the simulation scenarios are detailed, followed by a presentation and discussion of extensive simulation results in section 4. The objective of this discussion is to clearly identify the advantages and drawbacks of each paradigm under study. Section 5 summarises the key findings and identifies guidelines for further work.

2 Network mobility paradigms

Network mobility can be achieved by implementations obeying to different architectural approaches or paradigms. The chosen paradigm will impact the effectiveness and efficiency of the actual solution and, thus, it is important to carefully analyse the existing paradigms before generalised, large-scale deployment takes place.

This section provides an overview of existing network mobility paradigms. It also places existing proposals for network mobility in the presented paradigm landscape, thus providing a clear high-level view of relevant work. As the purpose of the paper is to assess and compare paradigms, we will not go into the detail of each particular proposal, as this detail can easily be found in the literature.

2.1 Legacy paradigm

The idea behind the Legacy-compatible (LC) network mobility paradigm is to readily allow network mobility without the need to change mobile network nodes (MNN) and correspondent nodes (CN). The NEMO Basic Support Protocol [1] is the reference solution for this paradigm.

Whenever a packet destined to the mobile network prefix (MNP) arrives at the home network, a home agent (HA) encapsulates the packet and sends it to the care-of-address (CoA) of the mobile router (MR) over a tunnel (the MRHA tunnel). The MR will then decapsulate the packet and deliver it to the MNN. Conversely, packets from an MNN to a CN are encapsulated, sent from the MR to the HA over the MRHA tunnel, decapsulated at the HA and routed to the CN.

Although extremely simple and fully compatible with legacy devices, the NEMO Basic Support solution suffers from several important problems, such as triangular routing, potential bottleneck in the home network and amplified sub-optimality in nested mobile networks, as thoroughly discussed in [2], [3] and [4]. Virtually all problems affecting the LC paradigm have to do with lack of route optimization.

2.2 Infrastructure-centric paradigm

In the case of the infrastructure-centric (IC) network mobility paradigm, most of the tasks inherent to network mobility – with emphasis on route optimisation – are carried out by infrastructure elements, such as mobile routers, home agents and correspondent entities. The main idea behind this paradigm is that end nodes are as unaffected as possible by mobility, at the cost of higher infrastructure complexity. Nevertheless, in contrast with the LC paradigm, changes to correspondent nodes are allowed, in order to support route optimisation.

Route optimisation mechanisms differ, depending on the type of mobile network node. Typically, in the case of local fixed nodes (LFN) and local mobile nodes (LMN), all traffic to/from these nodes must be optimized by some network element, such as a mobile router, a transit router, or a special purpose agent, which, in this context, we will designate MR for simplicity. This means that MRs must keep track of all LFN-CN optimisations and must perform route optimisation, using the return routability mechanism, whenever the mobile network moves. For visiting mobile nodes (VMN), the MR must somehow provide a temporary address to the node and enable that address to be routable inside the moving network. Also, typically, nested networks are treated by MRs as VMNs.

Several network mobility proposals fall into the infrastructure-centric paradigm class. Examples are Optimised Route Cache (ORC) [5] [6], Path Control Header

(PCH) [7] [8], Global HA to HA [9] and Mobile IPv6 Router Optimisation for Network Mobility (MIRON) [10] [11] [12].

2.3 End-node-centric paradigm

In the case of the end-node-centric (EC) network mobility paradigm, nodes are aware of their mobility condition. This means that end nodes play an active part in mobility management tasks, such as route optimisation, relieving mobile routers from them. For this to be possible, when a mobile router acquires a new care-of-address (CoA) it informs its inside network, so that all the mobile network nodes may then use the MR's CoA as their own CoA. From this point onward, mobile nodes decide when to optimise routes and perform optimisation by themselves. On the other hand, mobile routers largely limit their role to plain routing.

Up to the present, only one network mobility proposal uses the EC paradigm. Optimised Mobility for Enhanced Networking (OMEN) was proposed by the authors and first presented in [13]. In addition to detailed description of OMEN, reference [13] also provides preliminary validation and assessment of the proposal, in a limited way and for small-scale scenarios.

Although it is based on an innovative paradigm, OMEN exclusively relies on existing protocols. Instead of creating a new protocol for informing MNNs of their MR's CoA, OMEN uses Neighbor Discovery (RFC4861) [14]. Thus, either as response to a router solicitation message or by its own initiative, an MR can send router advertisement messages that will be used by MNNs to learn their CoA. The CoA will be carried in a new option. As the definition of new options is already accounted for in RFC4861, there is no need to change the protocol. On the other hand, route optimisation uses standard MIPv6 mechanisms.

2.4 Preliminary analysis

The end-node-centric paradigm eliminates several problems inherent to the legacy and infrastructure-centric paradigms.

One of the issues with the IC paradigm is to decide if and when to perform route optimisation, as it is recognized that not all traffic needs it. In fact, in some cases route optimisation may not be advisable, as the process takes longer than the traffic flow itself, as is the case of DNS traffic.

Another issue is that the load put on mobile routers may be excessive, as all the route optimisation work for all MNNs must be performed by their MR. For small mobile networks this may be acceptable, but in large ones the burden may be too high for mobile routers. Of course, the processing power of router devices has been increasing over the time, but the lighter the mobile router the smaller and the more energy-efficient it can be, and these are important factors when mobility is concerned.

The advantage of the EC paradigm mobility-awareness is that the nodes themselves can take decisions on whether to perform route optimization or not, depending on their needs, and will not be subject to a global, general-purpose, undifferentiated 'route optimisation'. Whenever a node is informed that it is under a

mobile network, it can be given all the needed information for it to perform route optimisation, using the existing mechanisms of MIPv6. This, of course, does not preclude legacy nodes from using plain, sub-optimal NEMO Basic Support mechanisms or any other solution that they implement and understand.

The advantage of passing the responsibility of route optimisation to the mobile nodes is that it allows MRs to be extremely light, taking care of network access tasks and MRHA tunnel handling, thus leading to lower complexity and lower processing needs. The price to pay for this approach is very low. LMNs need to be updated, in order for them to take advantage of mobility. VMNs and MRs will also have to perform very few EC-specific tasks. On the other hand, there is no need to change any of the protocols in order to adapt them to the EC paradigm, as it either uses existing capabilities of existing protocols, or it resorts to local processing in MRs and MNs.

Of course, the potential advantages of the end-node-centric network mobility paradigm may only be apparent and significant in generalised mobility scenarios, such as the one anticipated for the future Internet, where a vast amount of networks and nodes will be mobile. Thus, it is essential to study and compare the behaviour of the three network mobility paradigms in large-scale scenarios, which, naturally, can only be done by simulation due to the necessarily high number of involved networks and nodes. The simulations set up and the results of such study are presented in the following sections.

3 Simulations setup

The comparison of the LC, IC and EC paradigms in a large-scale scenario was only possible because three critical requirements were met: large processing power, a realistic and flexible network mobility emulation tool capable of running in a parallel processing environment and, last but not least, a set of scenarios suitable for drawing conclusions on the behaviour of the various paradigms under study in terms of round trip time, route optimisation time and signalling overhead. In the next sub-sections we will briefly present the mobility emulator and describe the simulation scenarios. The simulation results are presented and discussed in section 4.

3.1 The mobSim network mobility emulator

A network mobility emulator tool – named mobSim [15] – was especially built for this study due to the following reasons: 1) the large-scale of the study required the use of a cluster and the tool should be built and optimised for parallel execution; 2) the simulation detail should be as high as possible, ranging from protocol fields implementation to high-level mobility mechanisms and allowing for a configuration granularity ranging from individual hosts to routers, networks and global scenarios; 3) total flexibility to define the simulation scenarios, from the global dynamic behaviour to topologies, with any mix of mobile and fixed networks and nodes, any number of networks and any number of levels of nested mobility; 4) ability to perform any

simulation using the exact same parameters, conditions and settings for the three paradigms under study.

Table 1 presents the standard functionality implemented by mobSim. This functionality is used to construct generic implementations of the paradigms under study. It should be noted that mobSim actually implements the specified mechanisms down to packet field level, and that all messages are actually generated and sent between emulated devices, thus maximising the fidelity of the tool. Configurable delays include the time for data link layer handoff, the time needed for acquiring a new IP address (through DHCP), the binding update procedure time, the binding acknowledgement procedure time, line speed, HoTi time, CoTi time, HoT time, CoT time, encapsulation and decapsulation times.

Table 1. mobSim standard functionality.

IPv6 basic support (RFC 2460)	MIPv6 (RFC 3775)
Hop limit	Binding update
Next Next header implementations	Binding acknowledgement (binding accept, reject)
ICMPv6 (RFC4443)	Return routability procedure
ICMP echo, reply, unreachable, time exceeded	Home test init (HoTi)
IPv6 encapsulation (RFC2473)	Care-of test init (CoTi)
Mobility Header (RFC3775)	Home test (HoT)
Type 2 routing header	Care-of test (CoT)
Home address option (specific for T2RH)	Nonce utilisation (RFC3775 section 5.2.2)
	Binding refresh
Neighbor Discovery (RFC 4861)	NEMO basic support (RFC 3963)
Router advertisement	Bidirectional tunnel (MRHA tunnel)
Router solicitation	Binding update
	Binding acknowledgement
	Home agent implementation
	Mobile router implementation

The LC paradigm implementation closely follows the NEMO basic support specification (RFC 3963). In the case of the IC paradigm implementation, mobility functions – including route optimisation – are performed by mobile routers, on behalf of mobile nodes. The actual approach is a mix of the most important features of the proposals in [5] and [12]. In the case of the EC paradigm implementation, mobility functions are carried out by end-nodes, closely following the specification presented in [13].

3.2 Simulation scenarios

Fig. 1 presents the base topology of the overall network used in all simulation scenarios. This base topology is made up of fixed and mobile networks, and can change in the course of a simulation, as mobile networks change their point of

attachment, in non-nested and/or nested configurations. For simplicity, networks are represented by their border router only. Moreover, end-nodes are not depicted.

The overall network consists of a tree-like structure with seven levels. This organisation was chosen due to its natural ability to model network mobility nesting in the mobile part of the network, as well as different path lengths in the fixed part of the network. Level 1 of the hierarchy comprises 50 routers. Every level 1 router is directly connected to the other 49 level 1 routers. Each of these routers is connected to 5 level 2 routers. In turn, each level 2 router is connected to 3 level 3 routers, and so on, according to the diagram shown in Fig. 1. With this configuration, there are 456 routers under each level 1 router. In addition, the number of end-nodes under each level 1 router was configured to be between 225 and nearly 54,000. In total, the simulated network had 22,800 routers, 11,250 networks, and more than 27,000 end nodes (fixed and mobile).

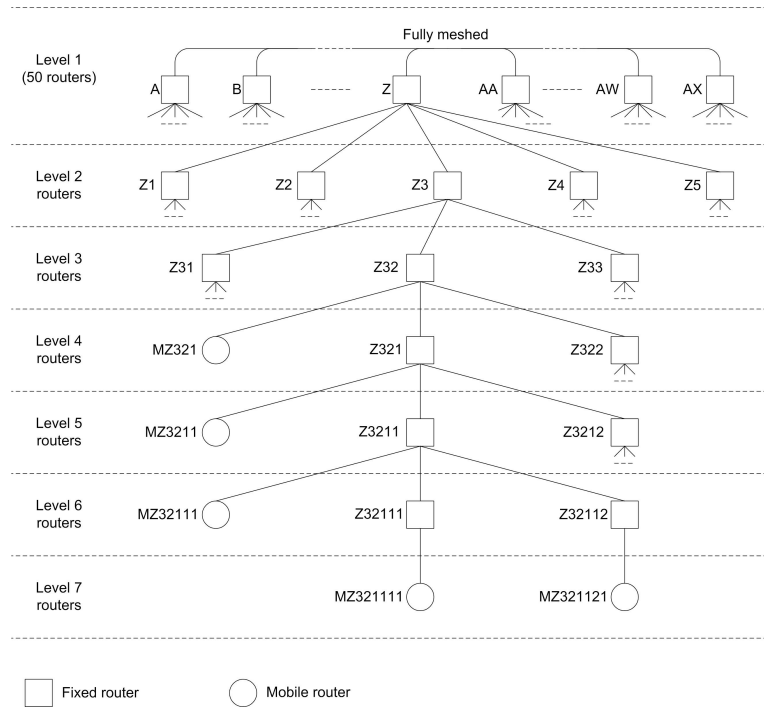


Fig. 1. Overall network topology.

Seventeen simulation scenarios were used. For scenario 1, no mobile nesting was used, i.e., all mobile networks were directly connected to fixed networks. For scenarios 2 to 17, one to sixteen levels of nesting (i.e., mobile networks inside mobile networks) were used, respectively. In all scenarios, up to 1,000 mobile nodes could be transmitting and/or receiving packets at the same time, although, in the course of the simulation runs, all mobile nodes were active at some point in time. Each mobile node transmitted a total of around 140 packets. The exact same conditions and settings were used for each of the paradigms under study.

In addition to the above, several delay parameters were used for all simulations, aiming to represent the time taken to perform the various actions associated with mobility procedures (e.g., table look up and update, security verifications, etc.). The chosen values are an approximation of actual values measured in a lab implementation, and were the following: DHCP delay – 300 ms; Return Routability delay – 200 ms; HoTi, CoTi, HoT or CoT messages processing delay – 100 ms; MAC-layer handoff delay – 500 ms; MR-HA tunnel setup delay – 10 ms; BU or BA messages processing delay – 10 ms.

Another important aspect of the simulations was that, for the EC paradigm, different ratios of number of nodes not performing route optimisation and number of nodes performing it were used. For instance, a ratio of 1:2 means that for each node not performing route optimisation two other nodes perform it. On the other hand, by design, in the case of the IC paradigm, route optimisation is performed for all nodes, whether their traffic requires it or not. This, of course, has influence in the performance of the mobile routers and on the amount of necessary signalling.

Finally, it is important to note that the objective of all simulations was to obtain results that could clearly show the impact of each paradigm on networks and routers. This impact can only be measured if no other factors distort the results. This is why all scenarios assumed unloaded networks. Naturally, if a given paradigm has poorer performance than another in an unloaded network, its behaviour will be even worse in networks with background traffic for which, possibly, some links will have little available capacity and some network elements will be overloaded. Moreover, there was no intention to perform stress tests to the various paradigms under study, in order to identify their limits and, thus, the number of packets transmitted per node was relatively low. Stress tests will be performed in future work.

4 Simulation results

Several simulation sets were run, each with a particular feature in mind, namely the effect of each paradigm on the average round-trip time (RTT), average route optimisation time and signalling overhead. The following sub-sections present and discuss the results of each set in detail

4.1 Average round-trip time

This set comprised a considerable number of simulations, each for each of the paradigms under study and for each scenario (1 to 17). In addition, each simulation was run for the following ratios of nodes not performing route optimisation and nodes performing route optimisation: 1:1, 1:2, 1:10, 1:100 and 1:1000. Fig. 2 presents the results for the 1:10 ratio.

It is possible to see that, as expected, the average round-trip time increases with the level of nesting for the case of the LC paradigm – as opposed to the case of the IC and EC paradigms – as in the LC paradigm there is no route optimization: as the level of nesting increases so increases the number of tunnels inside tunnels and, consequently,

the number of networks that have to be traversed by the packets in order to reach the destination. Note that, for visibility purposes, the maximum value in Fig. 2 is cut at 2 seconds, although the LC RTT values go up to 23.4 seconds.




Fig. 2. Average RTT, for the three paradigms, all scenarios and optimization ratio of 1:10.

Fig. 3 presents the evolution of the average RTT, for three levels of nested mobile networks, as the number of nodes performing route optimization increases, according to the ratios 1:1, 1:2, 1:10, 1:100 and 1:1000.




Fig. 3. Average RTT for three levels of nested mobile networks and different route optimization ratios

As the number of nodes requiring optimization increases, the EC paradigm gets better than the IC paradigm in terms of RTT. The reason is that the impact on MRs gets lower and lower in the case of the EC paradigm as MRs simply become routing devices, as opposed to the case of the IC paradigm, for which MRs have to perform all routing optimization operations for all nodes. This extra load on mobile routers leads to poorer IC RTT in all cases, except in the 1:1 case, for which the EC RTT is

worse due to the fact that half of the packets do not follow an optimal path, as opposed to the IC case where all packets follow optimal paths.

4.2 Average route optimisation time

The second batch of simulations aimed at determining the average route optimisation time as a function of the number of nodes requiring it. Fig. 4 presents this evolution, for up to 51 nodes.




Fig. 4. Average time for route optimization, as a function of the number of nodes requiring it

It is apparent that, as more nodes requiring route optimisation are present, the average route optimisation time increases steadily in the case of the IC paradigm and remains more or less constant in the case of the EC paradigm. This is due to the fact that in the case of the IC paradigm the mobile routers' processing load is higher, as all the tasks concerning route optimisation are performed by these network elements, as opposed to the EC paradigm, for which route optimisation tasks are performed by the end-nodes themselves.

4.3 Signalling overhead

The final set of simulations addressed the issue of signalling overhead, resulting from the transmission of HoTi, CoTi, HoT, CoT, BU and BA messages, according to RFC 3775, during route optimisation processes. Signalling overhead increases with the number of moves of mobile routers and with the number of route optimisation operations.

Fig. 5 presents the signalling overhead for a scenario in which the point of attachment of a mobile router was changed 5 times, for different optimisation ratios. As can be seen, the simulation results clearly confirm that the signalling overhead of the EC paradigm is always lower than the one for the IC paradigm, due to the fact that IC performs route optimisation for all nodes, regardless their need for route

optimisation. As the percentage of nodes requiring route optimisation increases, the signalling overhead of the EC paradigm approaches that of the IC paradigm.




Fig. 5. Signalling overhead for 5 moves and for different optimization ratios

5 Conclusion

We have presented a comparison of three network mobility paradigms – the legacy-compatible paradigm, the infrastructure-centric paradigm, and the end-node-centric paradigm – that span the existing solutions for network mobility. The comparison targeted a very large scenario, having in mind generalised mobility environments anticipated for the future Internet.

The obtained simulation results clearly point to several conclusions. The first conclusion is that the legacy paradigm has several performance limitations and, consequently, will hardly scale. Clearly, this paradigm was developed for allowing network mobility despite the presence of legacy nodes.

The simulation results concerning the infrastructure-centric paradigm clearly show that the concentration of functionality in network elements affects the overall performance of these elements and, consequently, the quality characteristics of end-to-end communication.

On the other hand, the performed simulations have clearly shown that the end-node-centric paradigm has very good potential and should be explored in the future. By relieving mobile routers and other infrastructure elements from several mobility tasks, significant benefit can be achieved in terms of round trip time, route optimisation time and signalling overhead. All simulation results confirmed this.

The presented comparative study opens up a large field for further research. Naturally, stress tests to the paradigms can be performed. The study of their behaviour under different network loads will surely be carried out in the near future. Last but not least, the implementation and detailed study of an end-node-centric network mobility

solution – as a follow up of the preliminary solution proposed in [13] – is a natural step.

Acknowledgments. This work was partially financed by project PTDC/EIA-EIA/116173/2009 (CoFiMoM) of the Portuguese Foundation for Science and Technology (FCT) and the Centre for Informatics and Systems of the University of Coimbra (CISUC). The authors would also like to thank Pedro Vieira Alberto, Luis Pinto and Pedro Almeida, from the Laboratory of Advanced Computing (LCA) of the University of Coimbra, for providing access the Milipeia cluster, without which this work would not have been possible.

References

1. V. Devarapalli et. al, “Network Mobility (NEMO) Basic Support Protocol”, RFC3963, Internet Engineering Task Force, January 2005
2. Chan-Wah Ng et. al, “Network Mobility Route Optimization Problem Statement”, draft-ietf-nemo-ro-problem-statement-03, Internet Engineering Task Force, September 2006
3. Chan-Wah Ng et. al, “Network Mobility Route Optimization Solution Space Analysis”, draft-ietf-nemo-ro-space-analysis-03, Internet Engineering Task Force, September 2006
4. Carlos J. Bernardos et. al, “NEMO: Network Mobility in IPv6”, Upgrade Vol IV, issue no. 2, April 2005
5. Ryuji Wakikawa et. al, "Optimized Route Cache Protocol (ORC)", draft-wakikawa-nemo-orc-01, work in progress, Internet Engineering Task Force, November 2004
6. Ryuji Wakikawa et. al, "ORC: Optimized Route Cache Management Protocol for Network Mobility", 10th International Conference on Telecommunications, vol 2, pp 1194-1200, February 2003
7. Jongkeun Na et. al, "Route Optimization Scheme based on Path Control Header", draft-na-nemo-path-control-header-00, work in progress, Internet Engineering Task Force, April 2004
8. Jonkeun Na, Seoul National University, “Supporting Route Optimization in Network MObility (NEMO)”, September 2004
9. Pascal Thubert et. al, “Global HA to HA protocol”, draft-thubert-nemo-global-haha-01.txt, work in progress, Internet Engineering Task Force, October 2005
10. Carlos Bernardos et. al, "MIRON: MIPv6 Route Optimization for NEMO", 4th Workshop on Applications and Services in Wireless Network, August 2004
11. Carlos Bernardos et. al, "Mobile IPv6 Route Optimisation for Network Mobility (MIRON)", draft-bernardos-nemo-miron-00, work in progress, Internet Engineering Task Force, July 2005
12. Carlos Bernardos, “Route Optimisation for Mobile Networks in IPv6 Heterogeneous Environments”, PhD thesis, Universidad Carlos III de Madrid, Spain, September 2006
13. Pedro Vale Pinheiro, Fernando Boavida, “OMEN – A New Paradigm for Optimal Network Mobility”, in Proceedings of WWIC 2008 - 6th International Conference on Wired/Wireless Internet Communications, Tampere, Finland, May 28-30, 2008
14. T. Narten et. al, “Neighbor Discovery for IP version 6 (IPv6)”, RFC 4861, Internet Engineering Task Force, September 2007
15. Pedro Vale Pinheiro, Fernando Boavida, “mobSim – A Network Mobility Simulation Tool for Very Large-Scale Scenarios”, in Proceedings of NTMS 2011 - 4th IFIP International Conference on New Technologies, Mobility and Security, Paris, France, 7 - 10 February 2011