

Design and Experimental Evaluation of a Route Optimisation Solution for NEMO

María Calderón, Carlos J. Bernardos, Marcelo Bagnulo, Ignacio Soto, Antonio de la Oliva

Abstract—An important requirement for IP networks to achieve the aim of ubiquitous connectivity is Network Mobility (NEMO). With Network Mobility support we can provide Internet access from mobile platforms, such as public transportation vehicles, to normal nodes that do not need to implement any special mobility protocol. The Network Mobility Basic Support protocol has been proposed in the IETF as a first solution to this problem, but this solution has severe performance limitations. This paper presents MIRON: Mobile IPv6 Route Optimisation for NEMO, an approach to the problem of Network Mobility support that overcomes the limitations of the basic solution by combining two different modes of operation: a Proxy-MR and an address delegation with built-in routing mechanisms. The paper describes the design and rationale of the solution, with an experimental validation and performance evaluation based on an implementation.

Index Terms—Network Mobility, Route Optimisation, Mobile Router, Mobile IPv6, PANA.

I. INTRODUCTION

THE Internet is evolving towards a more ubiquitous network, accessible anytime, anywhere. Forthcoming 4G [1] networks are expected to make possible the access through different and heterogeneous technologies, enabling true user mobility.

Triggered by these needs and the fact that deployed Internet protocols did not support mobility of any kind, the technical community designed several solutions that addressed the problem of mobility. Protocols such as Dynamic Host Configuration Protocol (DHCP) [2], [3] enabled the *portability* of terminals, but this was not enough to achieve real and transparent mobility, as it required ongoing transport sessions to be restarted after a change of the point of attachment. Terminal *mobility* support was first enabled by Mobile IP protocols [4], [5] and nowadays there are lots of proposals that extend and improve this support [6], [7] or address it in a different way [8], [9], [10].

As the Internet access becomes more and more ubiquitous, demands for mobility are not restricted to single terminals anymore. Supporting the roaming of networks that move as a whole is required in order to enable the transparent provision of Internet access in mobile platforms, such as trains, planes,

buses, etc [11]. The Network Mobility (NEMO) Basic Support protocol [12] enables complete networks to roam among different access networks, without disrupting network nodes' ongoing sessions and without requiring any specific mobility capability in the hosts. Nevertheless, it has some important limitations in terms of performance, due to the increased path length and the packet overhead that this solution introduces. Such limitations trigger the need for what has been called Route Optimisation (RO) for NEMO. These solutions try to overcome some of the limitations of the basic solution currently defined for Network Mobility (NEMO) [12].

This paper presents a Route Optimisation solution called Mobile IPv6 Route Optimisation for NEMO (MIRON). An initial version of MIRON was presented in ASWN 2004 [13]. This paper represents many refinements and extensions to our original work from ASWN 2004, extending the scope of the original solution and providing significant contributions over [13], not only in the design and scope of the protocol, but also performing an important practical evaluation, based on an implementation. MIRON enables direct communication between nodes belonging to a mobile network and other nodes of the Internet. MIRON is composed of two main modes:

- For those nodes of the mobile network that do not have any mobility capability, the Mobile Router (MR) performs all the Route Optimisation and mobility tasks on their behalf (what some authors [14] have called *Proxy MR*).
- For those nodes and (mobile) routers with standard Mobile IPv6 support, an address delegation mechanism, based on PANA (Protocol for Carrying Authentication for Network Access) [15] and DHCP [3], provides these nodes with topologically meaningful addresses (i.e. addresses that are directly reachable without requiring special rendezvous points, such as Home Agents, to be deployed to re-route any packet towards the actual location of the node). This enables these nodes to manage their own mobility and to perform the Route Optimisation by themselves.

These two different key modes of operation of MIRON combined give as a result a complete Route Optimisation solution for mobile networks, enabling traffic from any kind of node (with and without mobility support) and network configuration (including nesting) to be optimised. This is achieved without requiring changes on the operation of any node except the Mobile Routers.

The paper is structured as follows. A brief summary of the basic concepts of network mobility and an introduction to

Authors are with the Department of Telematics Engineering of Universidad Carlos III de Madrid

The work described in this paper is based on results of IST FP6 Integrated Project DAIDALOS. DAIDALOS receives research funding from the European Community's Sixth Framework Programme. Apart from this, the European Commission has no responsibility for the content of this paper. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

PANA are described in Section II, as well as the motivation for Route Optimisation. Section III provides a detailed description of MIRON. Section IV presents a performance evaluation of MIRON and compares it with the NEMO Basic Support protocol, by means of real-life experiments as well as analytical studies. In Section V, we explore different alternatives for Route Optimisation in NEMO and compare them with MIRON. Finally, Section VI concludes the paper.

II. BACKGROUND AND MOTIVATION

This section summarises some of the concepts, terminology and related protocols that are used along the paper, as well as the motivation for the necessity of Route Optimisation solutions for NEMO. A brief description of network mobility and the NEMO Basic Support protocol [12] is provided first. A more detailed, but still summarised, description of PANA is provided afterwards, as this is a novel protocol that is used as part of our proposal.

A. Network Mobility

Users demand Internet access not only from fixed locations (e.g., at home, at work, in hotels, cafeterias, universities, etc.) but also in public transportation systems (e.g., planes, trains and buses). In order to satisfy such demands, the technical community worked on the design of the required protocols to provide Network Mobility support. In particular, a working group called NEMO was created within the IETF¹ to extend the basic end-host mobility support protocol, Mobile IP [4], [5], to provide network mobility support [12].

In the IETF NEMO solution, a mobile network (known also as Network that Moves - NEMO²) is defined 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) [16]. 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. Naturally, these addresses only have topological meaning when the NEMO is at home. When the NEMO is away from home, packets addressed to the Mobile Network Nodes (MNNs) will still be 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 from the visited network, called the Care-of Address (CoA), where the routing architecture can deliver packets without additional mechanisms.

The goal of the network mobility support mechanisms [17] is to preserve established communications between the MNNs and external Correspondent Nodes (CNs) despite movement. Packets of such communications will be addressed to the MNNs addresses, which belong to the MNP, so additional mechanisms to forward packets between the Home Network

and the NEMO are needed. The basic solution for network mobility support [12] essentially creates a bi-directional tunnel between a special node located in the Home Network of the NEMO (the Home Agent), and the Care-of Address of the MR.

This solution is quite similar to the solution proposed for host mobility support, Mobile IPv6 (MIPv6) [5], without including the Route Optimisation (RO) support. Actually, the protocol extends the existing Binding Update (BU) message to inform the Home Agent (HA) of the IP address of the NEMO side of the tunnel (i.e. the CoA of the MR), through which the HA has to forward the packets addressed to the MNP. A new BU option is also defined to convey information about the MNPs.

B. Network authentication and access control: PANA

Nowadays, most of the public wireless access networks deploy some authentication and access control mechanisms in order to avoid unauthorised clients gaining access to the network. Nevertheless, these mechanisms are either limited to specific access media technologies (e.g., 802.1X for IEEE 802 links) or based on proprietary solutions (e.g., web access-based authentication methods) [18]. This fact, together with the expectation that future mobile devices will have several access technologies to gain network connectivity, triggered the creation of a new working group within the IETF, called PANA (Protocol for Carrying Authentication for Network Access), aimed at the definition and specification of a standard network-layer solution for authenticating clients for network access.

Our Route Optimisation protocol, MIRON, assumes that (i) an authentication protocol is used in the access network and that (ii) this protocol is PANA [19]. Although it may seem that the assumption that an authentication protocol is run in the network limits the theoretical scope of application of our solution, we argue that it does not limit its practical usability, since it is not realistic to assume public access networks to be open and not to require any kind of authentication. The second assumption upon which our solution is based is that it is PANA and no other the protocol used to carry out the authentication and access control in the network. We believe that assuming the existence of PANA in the access network does not limit the generality and applicability of our solution.

The PANA protocol [19], [15] is designed to facilitate authentication and authorisation of clients in access networks. Basically, it is a link-layer agnostic network access authentication protocol - encapsulating Extensible Authentication Protocol (EAP) [20] authentication methods - that runs between an entity (called PANA Client, PaC) in a node that wants to gain access to the network and an agent (called PANA Authentication Agent, PAA) in a server on the network side [15]. PANA is responsible for enabling the authentication process between these two entities, but it is just a part of the overall process of Authentication, Authorisation and Accounting (AAA) and access control. The complete picture, with AAA and access control functions, comprises four entities (Fig. 1):

¹<http://www.ietf.org/>

²NEMO can mean Network MObility or Network that MOves according to the context

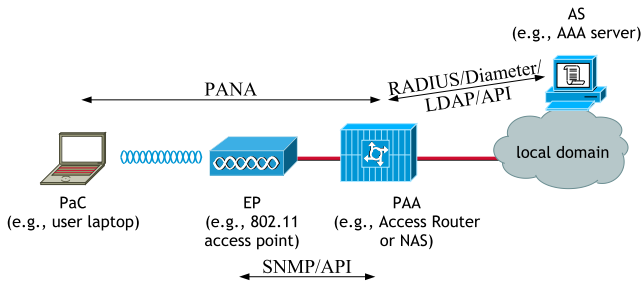


Fig. 1. PANA functional model overview. Some entities can be also collocated on a same physical node

- **PANA Client (PaC)**. Entity residing in the node that requests network access and implements the client part of the PANA protocol.
- **PANA Authentication Agent (PAA)**. Entity implementing the server part of the PANA protocol that interacts with the PaCs for authenticating and authorising them to access the network. The PAA consults the Authentication Server (AS) in order to verify the credentials and rights of a PaC and also updates the access control state, such as filters, in the Enforcement Points (EPs) in the network. The PAA usually resides in the Network Access Server (NAS) node, but can be hosted in any node that is in the same subnet (within 1-hop distance) as the PaC.
- **Authentication Server (AS)**. Server-side entity in charge of verifying the credentials of a PaC requesting access (sent by the PAA on behalf of the PaCs).
- **Enforcement Point (EP)**. Entity implementing the access control function by allowing access to authorised clients and preventing access from others.

PANA is a UDP-based protocol [19] (with its own re-transmission mechanism to provide reliability while delivering messages), consisting of a series of requests and responses. Each message can carry zero or more Attribute Value Pairs (AVPs) as payload. The main payload of PANA is EAP, which is responsible for performing authentication (PANA just helps the PaC and PAA establish an EAP session). Messages are sent between PaC and PAA as part of a PANA session, that consists of five different phases:

- 1) **Discovery and handshake phase**. This phase starts the PANA session. The PaC discovers the PAA(s) by either explicitly soliciting advertisements to the PAA(s) or receiving unsolicited advertisements. The PaC's answer, sent in response to an advertisement, starts a new session.
- 2) **Authentication and authorisation phase**. EAP execution between the PAA and PaC, by carrying an EAP method inside the EAP payload.
- 3) **Access phase**. If the authentication and authorisation phase is successful, the host gains access to the network and can send and receive IP data traffic through the EP(s).
- 4) **Re-authentication phase**. This phase is usually initiated by the PAA before the session lifetime expires (carrying EAP to perform authentication), although this phase may

be triggered by either the PaC or PAA regardless of the session lifetime.

- 5) **Termination phase**. The PaC or the PAA may choose to discontinue the access service at any time, by sending an explicit disconnect message.

C. Route Optimisation

The NEMO Basic Support protocol [12] has the following limitations, due to the fact that packets of MNNs' communications traverse - in both directions - the MR's HA, through the MR-HA bidirectional tunnel:

- It forces suboptimal routing (known as angular or triangular routing), i.e. packets are always forwarded through the HA following a suboptimal path and therefore adding a delay in the packet delivery.
- It introduces non-negligible packet overhead, reducing the Path MTU (PMTU). Specifically, an additional IPv6 header (40 bytes) is added to every packet because of the MR-HA bidirectional tunnel.
- The HA becomes a bottleneck of the communication as well as a potential single point of failure. Even if a direct path is available between a MNN and a CN, if the HA (or the path between the CN and the HA or between the HA and the MR) is not available, the communication is disrupted.
- These problems are exacerbated when considering nested mobility (i.e. a mobile network gains connectivity through other mobile networks), since in this case the packets are forwarded through all the HAs of all the upper level mobile networks involved (known as multi angular or pinball routing). This is because each sub NEMO obtains a CoA that belongs to the Mobile Network Prefix of its parent NEMO. Such a CoA is not topologically meaningful in the current location, since the parent NEMO is also away from home, and packets addressed to the CoA are tunnelled - thus increasing packet overhead - to the HA of the parent NEMO.

Because of all the limitations identified above, it is highly desirable to provide Route Optimisation support [14], [21], [22] for NEMO that enables direct packet exchange between a CN and a MNN without passing through any HA and without inserting extra IPv6 headers. In Mobile IPv6 [5], the Route Optimisation is achieved by allowing the Mobile Node (MN) to send Binding Update messages also to the CNs. In this way the CN is also aware of the CoA address where the MN's Home Address (HoA) is currently reachable. The Return Routability (RR) procedure is defined to protect a CN to change the IPv6 destination address (using the MN's CoA) of packets addressed to the MN's HoA [23].

III. MIRON: MOBILE IPV6 ROUTE OPTIMISATION FOR NEMO

In this section we present a novel solution that provides Route Optimisation for NEMO, enabling direct path communication between any kind of Mobile Network Node and a Correspondent Node in the Internet. An overview of the protocol is first provided, before describing in detail how the proposed solution works.

A. Protocol Overview

MIRON aims at improving the overall performance of communications involving nodes within a NEMO, by both avoiding packets passing through the MR's HA and reducing the packet overhead due to the additional IPv6 headers introduced by the NEMO Basic Support protocol. MIRON does not introduce any change on the operation of the Correspondent Nodes and the Mobile Network Nodes, but only of the Mobile Routers.

Fig. 2 shows a possible Route Optimisation target scenario for MIRON. It considers a mobile network deployed in a train, consisting of different types of MNNs:

- *Fixed nodes* in the train - without mobility support, called Local Fixed Nodes (LFNs) [16] -, such as internal servers or passengers' laptops that use an IPv6 address from the mobile network.
- *Mobile devices* - called Visiting Mobile Nodes (VMNs) - such as passengers' laptops that keep their IPv6 address.
- Nested mobile networks, such as Personal Area Networks (PANs), e.g., the passenger's laptop acts as a MR of his devices and is connected to the train's MR.

All of these devices access the Internet through the train's MR. This scenario includes almost every possible mobile network communication, involving LFNs, VMNs and nested NEMOs. Fig. 2 also shows the different components every entity is composed of. Both the components and the way they work together to construct a complete Route Optimisation solution will be described in detail later in this section. First, a brief protocol overview of MIRON is provided.

MIRON addresses two different Route Optimisation aspects:

- *Angular routing.* Angular routing is caused by the MR-HA bidirectional tunnel introduced by the NEMO Basic Support protocol, since packets of a communication involving a MNN has to be forwarded through the HA of the NEMO. MIRON addresses this problem in two different ways, depending on whether the MNN that is communicating with a CN has mobility support or not. If the MNN has no Mobile IPv6 capabilities (i.e. an LFN), the approach followed by MIRON consists in delegating the Route Optimisation functionality to the MR, that performs all the RO signalling on behalf of the LFNs and packet handling. Therefore, the MR is a kind-of "Proxy MR" [14] for the LFNs of the NEMO. On the other hand, if the MNN is a Mobile IPv6 MN (i.e. a VMN) that is visiting the mobile network, MIRON takes advantage of the already available mobility support that the MN has. In this case, by using PANA and DHCPv6, the MR provides a topologically meaningful IPv6 address (that is, an address belonging to the network that the MR is visiting) to every VMN attached to the NEMO and updates it every time the NEMO moves. This, in addition to a routing mechanism that enables these addresses to be routed inside the NEMO, allows the VMN to make use of its own Mobile IPv6 Route Optimisation functionality, therefore avoiding traversing the HA and reducing the packet overhead.

- *Multi-angular routing.* Multi-angular routing is caused in nested NEMOs by the chain of nested MR-HA bidirectional tunnels that packets should traverse. MIRON addresses this problem by using PANA and DHCPv6 to provide topologically meaningful IPv6 addresses to every MR in the nested NEMO hierarchy. In this way, every MR has an IPv6 address belonging to the network that the root-MR (that is, the MR of the NEMO at the top of the hierarchy) is visiting. This, in addition to a routing mechanism that enables these addresses to be reachable, makes it possible to avoid traversing any HA.

The set of mechanisms of MIRON enables direct path communication between a MNN (LFN or VMN) and a CN, avoiding the suboptimal MR-HA path. The recursive tunnelling due to nesting is also eliminated, therefore optimising the traffic in every possible configuration of a mobile network. MIRON only introduces changes in the MR (see Fig. 2) and MNNs and CNs remain unchanged, thus facilitating the deployment of the solution. The next sections provide a detailed protocol walk-through of MIRON.

B. Angular route optimisation

If no Route Optimisation mechanism is used, all the traffic sent/directed to a MNN goes through the bidirectional tunnel set up between the MR and its HA. MIRON enables direct communication - without traversing the MR's HA - by following one of the next approaches, depending on the type of MNN:

- *Local Fixed Node (LFN).* LFNs do not have mobility support, so any mechanism that attempts to optimise their traffic should be implemented without requiring support from the LFN itself. The MIRON mechanism for LFNs is basically a proxy-MR approach, in which the MR performs the Mobile IPv6 Route Optimisation [5] on behalf of the LFN.
- *Visiting Mobile Node (VMN).* VMNs are Mobile Nodes that are visiting the mobile network, managing their own mobility. By default, the Care-of Address obtained and used by a VMN attached to a NEMO belongs to the Mobile Network Prefix of that NEMO, so although these mobile nodes may be performing Route Optimisation with the CNs they are communicating to, there still exists a tunnel - between the NEMO's MR and the MR's HA - introduced by the NEMO Basic Support protocol. In this case, our proxy-MR approach is not feasible, so a different mechanism is used. MIRON takes advantage of the mobility support that VMNs already have. Basically, we propose a mechanism, using PANA and DHCPv6, that enables the VMNs to configure topologically valid IPv6 addresses (i.e. those addresses that belong to the address space of the foreign network the NEMO is visiting) as CoAs, and letting the VMNs manage their mobility and Route Optimisation tasks.

1) *Detection of the type of node:* In order to apply the appropriate Route Optimisation mechanism, the MR should first be able to determine which kind of node (LFN or VMN) every node that is communicating is. The MR performs such

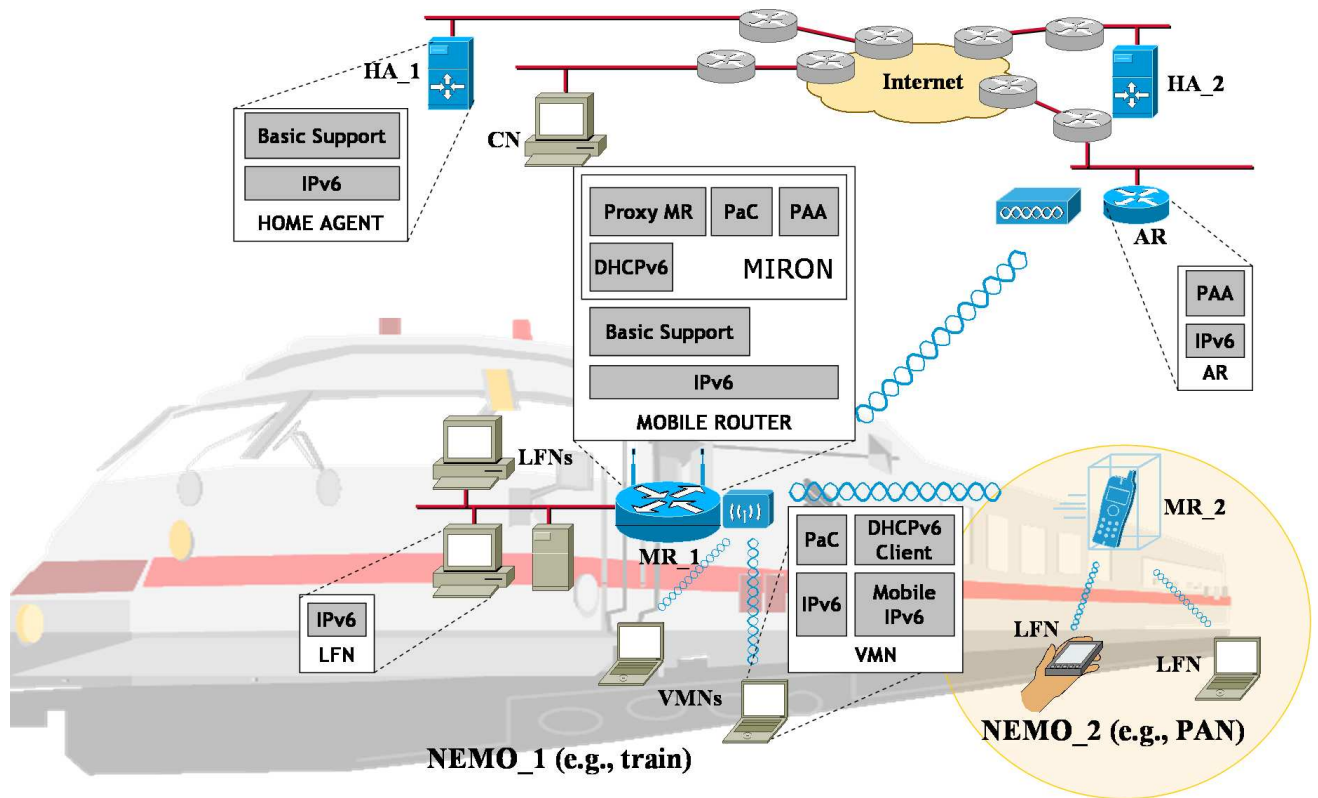


Fig. 2. Overview of the MIRON architecture in a practical scenario

a task by looking for Binding Update messages received at its ingress interfaces, since an MN right after gaining connectivity to a foreign network and configuring a new CoA (from the MNP), has to send a Binding Update to its HA to inform it about its new location (i.e. MN's CoA).

2) *Route Optimisation mechanism for LFNs*: Local Fixed Nodes are nodes without any mobility support running, therefore a mechanism that optimises their traffic cannot rely on any mobility function implemented by them. MIRON puts this LFN mobility support into the MR, that performs all the required mobility and Route Optimisation tasks on behalf of the LFNs attached to it.

The mechanism basically consists in enabling a MR to behave as a proxy for the LFN, performing the Mobile IPv6 Route Optimisation signalling and packet handling [5] on behalf of the LFN. In order to do that, the MR first tracks the different communications that LFNs have established and decides which of those will be optimised, since optimising a traffic flow involves a cost - in terms of signalling and computation resources at the MR - that may not be worth for some kinds of flows (e.g., DNS queries). This decision (that is, whether to perform Route Optimisation for each flow or not) is out of the scope of this paper, although we are working on different algorithms and heuristics to evaluate their performance based on real tests.

For those LFN-CN pairs whose traffic is to be optimised, the MR starts to send the RO signalling described for standard Mobile IPv6 in [5]:

- The BU is sent by the MR.

- The BU contains the LFN address as the MN's Home Address (HoA) and the MR's CoA as the MN's CoA (since the MR's CoA is the only topologically meaningful address available).

The Route Optimisation mechanism defined by Mobile IPv6 [5] requires an additional procedure to be performed before sending the BU message, in order to mitigate possible attacks [23]. Basically, this mechanism, called Return Routability (RR), verifies that the node that is reachable at the HoA is able to respond to packets sent to a given address (different to the HoA of the node). This mechanism can be deceived only if the routing infrastructure is compromised or if there is an attacker between the verifier and the addresses (that is, HoA and CoA) to be verified. With these exceptions, the test is used to ensure that the MN's Home Address (HoA) and MN's Care-of Address (CoA) are collocated.

In our solution we adopt the above procedure. For this purpose, the MR has to perform the Mobile IPv6 Return Routability procedure [5] on behalf of the LFN. Such procedure involves sending the Home Test Init (HoTI) and Care-of Test Init (CoTI) messages to the CN and processing the replies (Home Test message HoT - and Care-of Test message - CoT). These messages are sent as specified in [5], using the LFN's address as the source address in the HoTI message - which is sent encapsulated through the MR's HA -, and the MR's CoA as the source address in the CoTI message. With the information contained in the HoT and CoT messages, sent by the CN in response to the HoTI and CoTI messages respectively, the MR is able to build a BU message to be sent

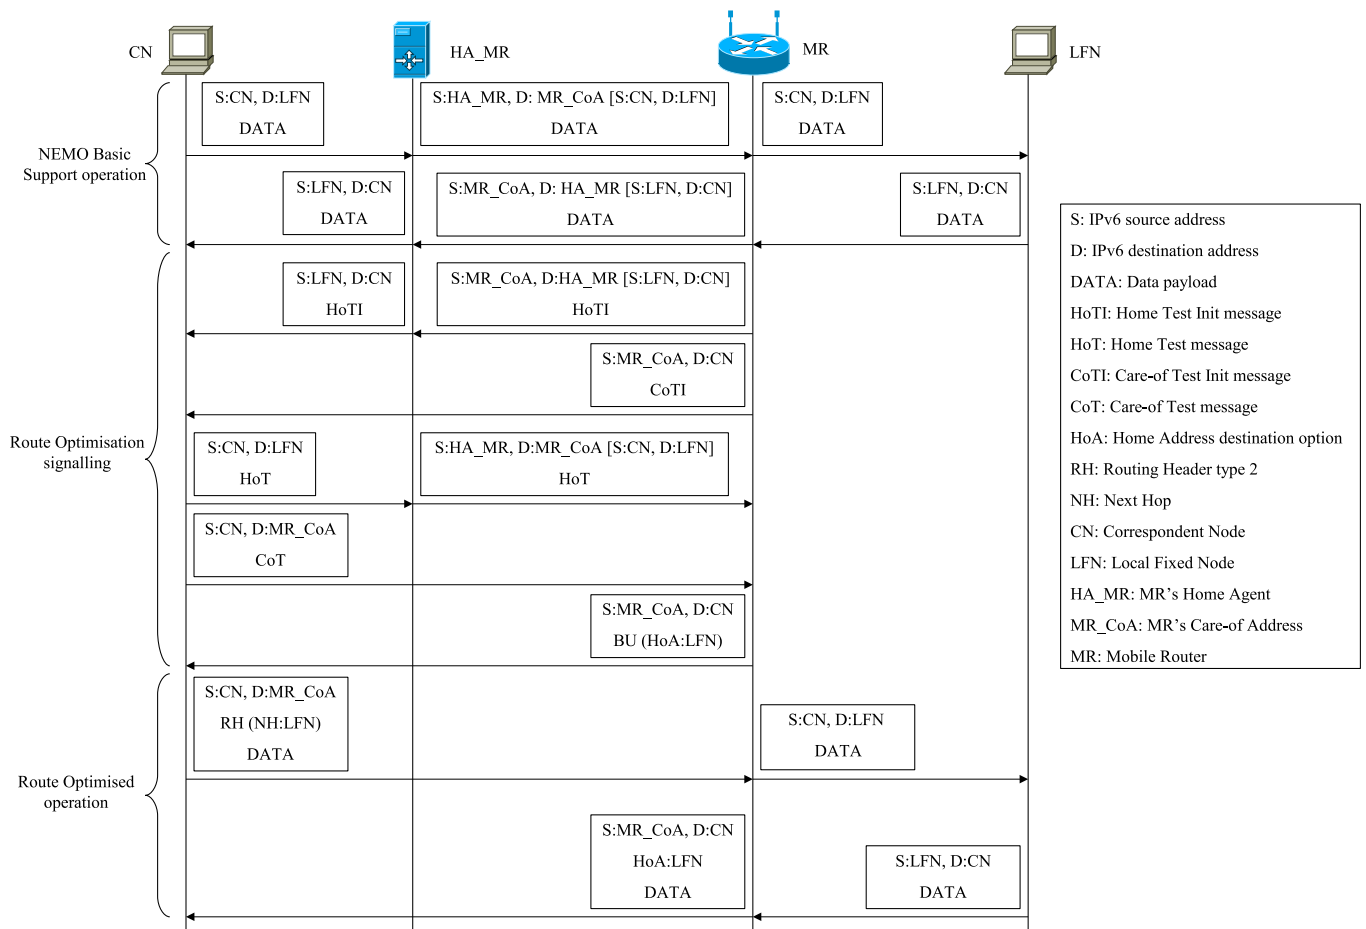


Fig. 3. Route Optimisation mechanism for LFNs: Proxy-MR operation

to the CN on behalf of the LFN. This message is sent using the MR's CoA as the packet source address and carries a Home Address destination option set to the LFN's address.

Besides performing the Route Optimisation signalling on behalf of the LFN, the MR has to process the packets sent by and directed to the LFN. Packets sent by the CN follow a direct path to the MR, not traversing the HA, as a result of the Route Optimisation. These packets carry the MR's CoA as destination address, and also carry a Type 2 Routing Header with the LFN's address as next hop. The MR processes and removes the Routing Header of the packet, checking if the next hop address belongs to one of its LFNs and, if so, delivering the packet to the LFN. Current Mobile IPv6 specification [5] defines that IPv6 nodes which process a Type 2 Routing Header must verify that the address contained within is the node's own Home Address. This is done in order to prevent packets from being forwarded outside the node, but we change the processing of this kind of routing header - at the MRs - to verify that the address contained within is an address of one of the LFNs the MR is acting as Proxy-MR. In the opposite direction, the MR receives the packets sent by the LFN and performs the following actions on every packet:

- Set the MR's CoA as IPv6 source address.
- Insert an IPv6 Home Address destination option, carrying the address of the LFN.

Fig. 3 shows the signalling and data flows of the proposed Route Optimisation mechanism for LFNs, including at the top of the figure the NEMO Basic Support protocol data flow for comparison purposes.

From the security point of view, allowing the MR to perform some operations on behalf of the LFNs attached to it does not introduce any threat, because LFNs trust their MR for the routing of all their traffic. From the architectural point of view, the solution is also natural, as the Route Optimisation support defined by Mobile IPv6 [5] conceptually could be implemented in multiple boxes. MIRON just applies this mechanism, by dividing the functionalities among two different physical boxes, but actually the conceptual basis of the solution is the same as the one defined in [5].

3) *Route Optimisation mechanism for VMNs:* Visiting Mobile Nodes are nodes that support mobility (that is, nodes running Mobile IPv6 [5]) and are visiting a mobile network. Therefore the VMN is attached to an Access Router that is the NEMO's MR, and the address that the VMN obtains and configures as CoA belongs to the Mobile Network Prefix. In this case, our proxy-MR mechanism used for LFNs cannot be used, as the VMN itself may generate Route Optimisation signalling with its CNs. Besides, the MR cannot modify the RR and RO signalling sent by the VMN in order to make the MR's CoA the CoA that the CN uses to send the packets to the

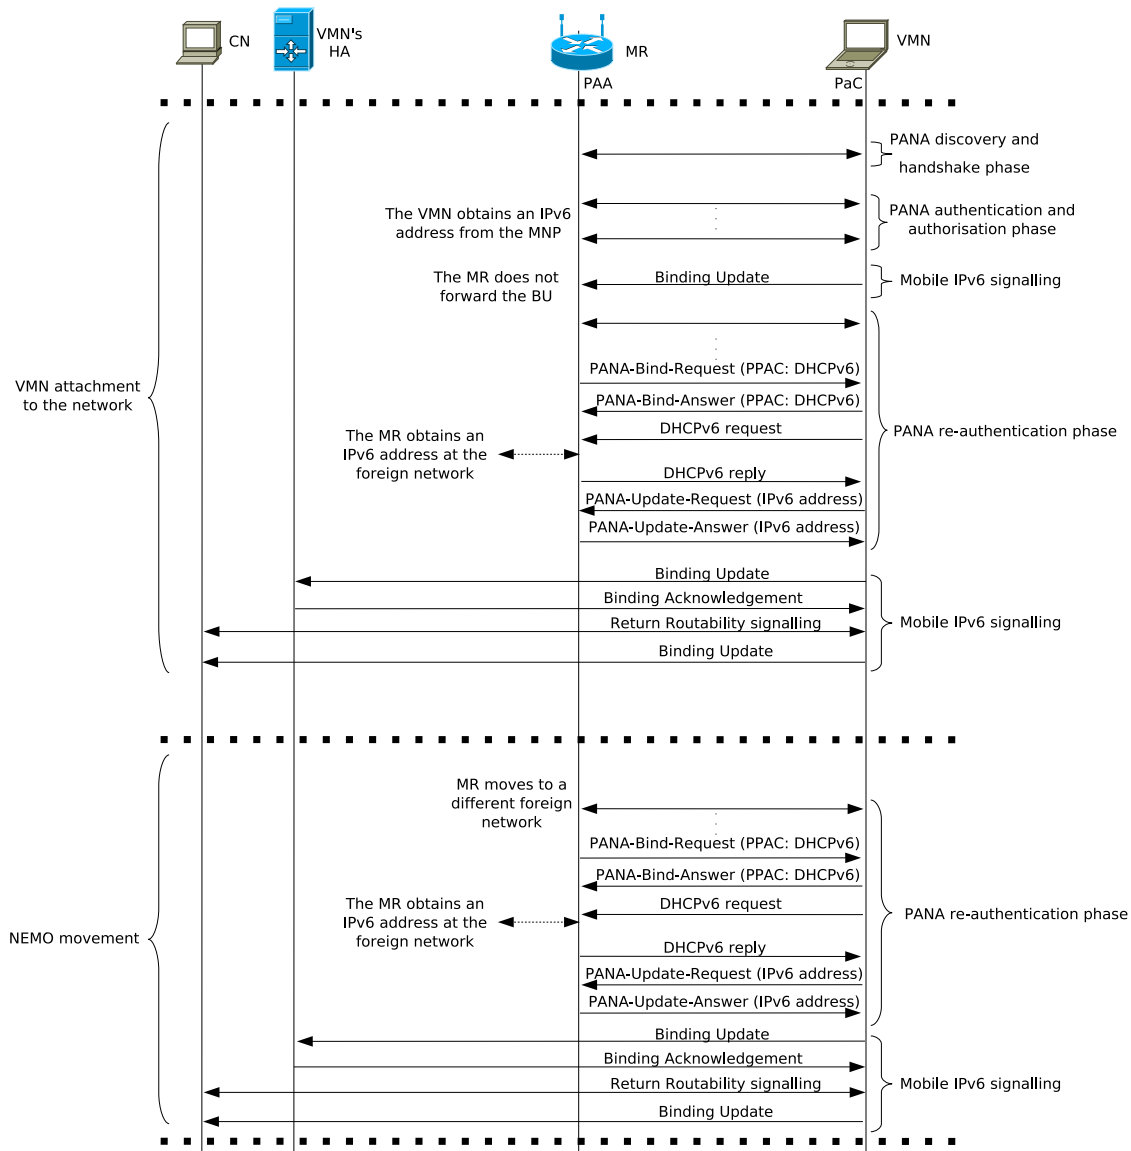


Fig. 4. Route Optimisation mechanism for VMNs

VMN, because part of the RR signalling is protected by IPsec (the HoTI message is sent through the VMN's HA protected by IPsec ESP).

The Route Optimisation approach that MIRON defines for this type of node is based on taking advantage of the mobility support that these nodes already have, providing the means to the VMN to perform the RO. In order to allow the VMN to manage its own mobility and enable it to perform Route Optimisation with the CNs (in a way that avoids the MR-HA bidirectional tunnel), we propose the following:

- Provide a topologically meaningful IPv6 address to the VMN. These addresses are those that belong to the network that the root-MR is visiting.
- Enable this address to be routable inside the NEMO, as it has only topological meaning in the visited network. The MR has to perform proxy neighbour discovery for this address in the egress interface that is attached to the network to which the address belongs. Besides, the MR

has to insert a host route for this address to be able to route packets destined to it.

- Perform source address routing in the MR in order to send directly (that is, avoiding the bidirectional MR-HA tunnel that still exists and is used for non-optimised traffic) packets sent by the VMN.
- Provide a new address to the VMN when the NEMO moves, replacing the previous one.

A mechanism that fulfils the previous goals should be able to allow VMNs, and only the VMNs - the mechanism must not affect other type of nodes -, to obtain a new IPv6 address to be used as the CoA, whenever the MR wants to, and in a secure way that does not introduce any new security threat.

Most of current access networks (such as hotspots deployed in airports and cafeterias) require users to authenticate to the network before gaining Internet access. As the number of hotspots continues growing in the coming years, authentication mechanisms will be more and more important in order to avoid

non-authorised users using and wasting the network resources. Using a standard protocol to perform such authorisation and authentication tasks would help in the deployment of ubiquitous access “anytime anywhere” networks. MIRON assumes that PANA [18], [15], [19] will play that role (thus the MR runs also PANA), and it makes use of its features (in the execution of the PANA protocol between a PaC entity residing in the VMN and a PAA entity residing in the MR) to provide Route Optimisation to VMNs visiting a NEMO.

The mechanism to provide an IPv6 address to the VMN using PANA works as follows (see Fig. 4): when a VMN connects to a NEMO, it initiates the PANA session (PANA discovery and handshake phases). Immediately after that, the actual authentication and authorisation phase (by executing EAP between the PAA and PaC) takes place. The VMN is then authorised to access the network and it has an IPv6 address. This address is obtained by using the address autoconfiguration mechanism available at the NEMO, first we assume that we are using stateless address configuration for addresses of the Mobile Network Prefix, but later we will see that we can also use stateful address configuration within the NEMO. The VMN then sends a Binding Update message to its Home Agent, informing about its current location. This BU is received at the MR, that then becomes aware that a new VMN is now attached to the NEMO. The MR does not forward this BU message to the next hop towards the HA and starts a PANA re-authentication phase.

During the PANA re-authentication phase, the PAA located in the MR tells the PaC located in the MN that it should obtain and configure a new IPv6 address (Post-PANA address, POPA) and how to obtain it, by including available configuration methods in a Post-PANA-Address-Configuration (PPAC) AVP included in a PANA message (PANA-Bind-Request). DHCPv6 is the only available configuration mechanism listed in the message, and upon the reception of that, the VMN requests an address using DHCPv6. There is a DHCPv6 component located at the MR that receives the DHCPv6 requests from the VMN and then obtains (using one of the available auto configuration mechanisms at the foreign network) an IPv6 address. The DHCPv6 component generates a DHCPv6 reply - including this address - that is delivered to the requesting VMN. This DHCPv6 component implements the client part of DHCPv6 and also some reduced functionalities of the server part (e.g., the generation of DHCPv6 responses), but it is not a DHCPv6 server (for example, the DHCPv6 component does not have a pool of available addresses, each time an address is needed, it obtains it from the foreign network), although the implementation of this DHCPv6 component can be very easily performed from the code of a normal DHCPv6 client and server implementation. Once the MR has sent the DHCPv6 reply - including the (/128) delegated address - to the VMN, the PaC in the VMN conveys this newly configured IPv6 address to the PAA in the MR by sending the PANA-Update-Request message.

The use of stateful address configuration (DHCPv6) within the NEMO (to configure addresses from the MNP) is also possible, but it requires the DHCPv6 component at the MR to implement the complete server functionality and to check,

before providing an address, if the requesting node is an identified VMN or not, to know whether this address should belong to the MNP or to the visited network address space. Nodes are identified as VMNs at the MR according to the procedure described above.

In order to enable the VMNs’ addresses reachability inside the NEMO, the MR has to add a host route for each VMN’s address and perform proxy neighbour discovery on its egress interface (the interface that is connected to the link where the address has topological meaning), allowing the MR to forward packets to their final destinations. Both the delegated IPv6 addresses and the host routes have a lifetime that prevents this state to remain in the network after a sub-NEMO or a node leaves a parent-NEMO (the value depends on the lifetime of the address obtained by the root-MR).

The VMN, triggered by the change of address, starts the Mobile IPv6 location update process, sending first a Binding Update (BU) message to its HA. The VMN may then update the location information in the CNs it is communicating with (if the VMN is running Route Optimisation with its CNs). This process consists of the VMN performing the Return Routability process [5] and sending a BU to every CN whose traffic is to be route optimised.

When the NEMO moves to a different foreign network, the MR requests new IPv6 addresses and provides them to the VMNs attached to the NEMO by starting a new PANA re-authentication phase. The MR requests VMNs to configure a new IPv6 address using DHCPv6.

C. Multi angular route optimisation

The routing inefficiencies due to the MR-HA bidirectional tunnel are exacerbated when NEMOs are attached to other NEMOs, forming a nested NEMO. Packets belonging to a communication between a MNP of a nested NEMO and a CN have an additional IPv6 header per nesting level and traverse the HAs of every MR of the nested NEMO.

The problem of enabling RO for nested NEMOs (i.e. MRs visiting a NEMO) is very similar to that of VMNs (i.e. MNs visiting a NEMO). Both VMNs and MRs are nodes that are mobile-capable and can manage their own mobility. Routing inefficiencies arise from the fact of not using topologically meaningful addresses (i.e. addresses belonging to the NEMO MNP) as CoAs. Section III-B.3 describes an address delegation mechanism with a built-in routing system that is able to provide IPv6 addresses - belonging to the foreign network that the MR is visiting - to a VMN in a secure way, by using PANA facilities.

MIRON extends that solution, used for providing angular RO for VMNs, to enable multi angular RO in nested NEMOs. Basically, the solution consists in providing topologically meaningful addresses - that is, those that belong to the foreign network that the root-MR is visiting - to every MR in the nested NEMO. The same PANA-with-DHCPv6-based mechanism is used to provide an IPv6 address to a MR that attaches to a NEMO (and to change it when one of the parent NEMOs moves). MRs have both a PAA and a PaC component and also a DHCPv6 component, so when a MR connects to

a mobile network, they are able to get and configure a new IPv6 address.

The parent-MR indicates - using PANA - to the newly arrived sub-MR (or VMN) that it should configure an IPv6 address using DHCPv6. The parent-MR plays the role of a DHCPv6 server in the NEMO and receives the request of the sub-MR, requesting after that an address. In this way, DHCPv6 requests are received and generated recursively until they reach the root-MR. The needed address is then requested using one of the available address auto configuration mechanisms at the visited network. This IPv6 address is then sent back to the original requester, also in a recursive way, where each MR in the path receives the reply and generates a new reply. In order to avoid the delay that would be introduced by DHCPv6 in this recursive address request and delegation process, MRs implement the Rapid Commit option [3], so the MR terminates the waiting process as soon as a reply message generated by the parent-MR is received. Once the process is completed, the new MR has an IPv6 address that may be used as the MR's CoA in the NEMO Basic Support registration process, as well as in the MIRON angular RO mechanism by performing the Proxy-MR optimisation on behalf of the LFNs of its NEMO.

Providing topologically meaningful addresses is not the only required step to avoid the suboptimal multi angular routing in nested networks. Another requirement that needs to be met is that these addresses are globally reachable. To enable that, every MR in the nested NEMO keeps track of the address of the node requesting an IPv6 address using DHCPv6, so when the delegated address is received, it can insert a host route entry in its routing table that allows it to route packets destined to that address afterwards. This information is also used to perform source address based routing for the packets generated inside the NEMO, as every MR should know for each packet if it has to be sent directly to the router it is connected to (this way, avoiding the tunnel), or it has to be sent towards the HA, through the bidirectional tunnel (for traffic that is not being optimised).

This address delegation mechanism with built-in routing avoids the multi encapsulation and multi angular routing in nested networks. Besides, it enables angular MIRON route optimisations to work when applied to a NEMO located in any level of a nested NEMO.

IV. EVALUATION OF MIRON

This section provides both an experimental and analytical evaluation of MIRON. The main aim of this evaluation is to study the performance of MIRON and compare it with the NEMO Basic Support protocol.

A. Experimental evaluation

1) *MIRON implementation*: In order to be able to conduct real experiments that allowed us to evaluate the performance of the NEMO Basic Support protocol and the improvements provided by MIRON, we first implemented the NEMO Basic Support protocol [24]. A first prototype of MIRON was also implemented, providing all the Route Optimisation mechanisms. Packets belonging to a communication flow optimised

by MIRON must not traverse the bidirectional tunnel, so for outgoing traffic a host route towards the CN of the flow should be inserted at the MR to avoid the default route through the tunnel interface. Besides, there may be simultaneously communications in a NEMO - from different MNNs - with the same destination CN that are not all being optimised, thus source address based routing is necessary. IPv6 based source routing was not natively supported in Linux at the beginning of the MIRON implementation process, so this part was (temporally) tested without performing source address based routing, as we did not make any test involving optimised and non optimised traffic with a CN at the same time.

The required additional protocols and procedures (such as the Return Routability and DHCPv6) were completely implemented, with the exception of PANA, that is being implemented and integrated now. The fact of not having implemented the PANA signalling does not have an impact on the results obtained in the tests, as we have focused in the TCP throughput and PANA signalling is generated during handovers (and also periodically to renew the lifetime). In this paper, we have not been concerned about the performance of our solution during handovers as we address the problem of Route Optimisation, just in the same way that the Route Optimisation solution for Mobile IPv6 does. Improvements in the handover latency (like the ones designed for Mobile IPv6 [6], [7]) requires further study and will be addressed in future works.

The NEMO Basic Support protocol [24] and MIRON were mostly implemented in user space, because in this way the development was easier and quicker than doing that in the kernel. The main software characteristics are: a Linux machine with kernel linux-2.6.x (tested with linux-2.6.8.1) with support for IPv6-in-IPv6 tunnels (used for the HA-MR bidirectional tunnel) and Netlink sockets, and the *pcap* library (used for the capture and processing of the mobility related signalling).

2) *Studied scenarios*: Two different scenarios (Fig. 5) were deployed to allow us to experimentally test the performance of MIRON and compare it with the NEMO Basic Support solution. The first one (Fig. 5(a)) was used to evaluate the performance in a non-nested case, whereas the second (Fig. 5(b)) is an extension of the former to include nesting.

We describe next the second scenario, as it is an extension of the first one. This scenario (Fig. 5(b)) consists of thirteen Mandrake 10.0 Linux machines (all with linux-2.6.8.1 kernels, except 3 routers that run linux-2.4.22). Five of them act as *fixed* (i.e. non-mobile) routers (R1 to R5), two as Home Agents (HA1 and HA2), two as Mobile Routers (MR1 and MR2), one as Correspondent Node (CN), one as Local Fixed Node (LFN) and two as *fixed* nodes (Fixed nodes 1 and 2). This is part of the IST Daidalos³ project testbed at the Universidad Carlos III de Madrid.

All the mobility-aware nodes run the network mobility software, that is, the NEMO Basic Support protocol (at the HA and MR) and MIRON (at the MR only) developed by us. The CN runs MIPL⁴ 2.0 RC1, with the support of Route

³<http://www.ist-daidalos.org>

⁴Mobile IPv6 for Linux, available at <http://www.mobile-ipv6.org/>

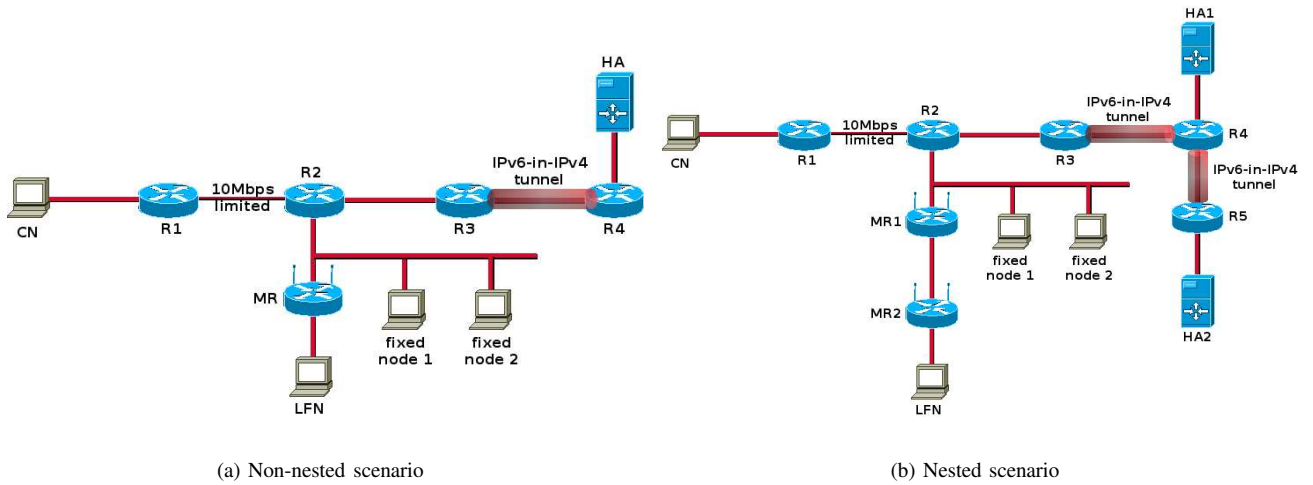


Fig. 5. Network mobility testbed employed during the experimental evaluation

Optimisation enabled.

We need the ability to modify the delay in the path followed by packets of a communication between a CN and a MNN (that is, the path between the CN and the MNN’s HA and/or the path between the MNN’s HA and the foreign network the MR is currently attached to). This allows us evaluate how the performance of a particular network mobility solution is affected by network characteristics, such as the particular location of mobile networks, home networks and correspondent nodes. For this purpose, we used the NIST Net emulator⁵. NIST Net allows a single Linux PC, set up as a router, to emulate a wide variety of network conditions (e.g., latency, jitter, packet loss, . . .).

We were interested in studying how the delay (and also the packet overhead) introduced by the MR-HA bidirectional tunnel affects the performance of applications. TCP performance is heavily dependent on the round trip time (RTT) between the communication peers. Taking this into consideration and the fact that 85% of the traffic in the Internet is generated by TCP connections [25], the TCP study case becomes very interesting to be performed and analysed. Therefore, we set up an scenario that allowed us to modify the delay in the CN-HA-MR path.

Other network characteristics, besides the delay, that do not have an special effect in the TCP performance and that are also present in non-mobile networks, were not modified.

NIST Net software runs only in IPv4 and with linux-2.4.x kernels. Therefore, in order to use it in our testbed, we had to set up an IPv6-in-IPv4 tunnel – between R3 and R4 and between R3 and R5 – for use in our IPv6 scenario. In the first, non-nested scenario setup (Fig. 5(a)), every packet in the CN-HA-MR path traverses the IPv6-in-IPv4 tunnel, which allows us to modify the network behaviour by changing the parameters of the NIST Net emulator running in R3 and R4. In the rest of the path followed by packets, native IPv6 is used, so the tunnel inclusion does not affect the overall test performance except for the small added delay due to

IPv6-in-IPv4 tunnelling and the reduction of the PMTU (the situation is not different from having a change of the transport link technology in the path and it is transparent to the IPv6 behaviour). Actually, the IPv4 tunnel clearly shows the current status of IPv6 networks in the Internet, with lots of IPv4 clouds connecting IPv6 native networks. In the second, nested scenario, a second IPv6-in-IPv4 tunnel was set up – between R4 and R5 – to allow us to modify the network delay between the two different home networks (i.e. between HA1 and HA2).

To avoid the influence of the wireless media characteristics and interferences from other neighbouring wireless networks, the performance tests were conducted using wired mobile routers, although experiments using wireless mobile networks were also performed to check the correctness of our solution in a more realistic scenario.

3) *Impact of network mobility on the TCP performance:*

The suboptimal routing introduced by the NEMO Basic Support protocol [12] makes packets not follow the direct CN-MR-MNN path, but the usually longer CN-HA-MR-MNN path. This adds a delay in the packet delivery that can significantly reduce the performance of certain applications. Furthermore, packets are encapsulated between the HA and the MR, thus reducing the PMTU. Both effects, increased delay and reduced PMTU, have an impact in the performance of applications.

TCP is the predominant type of traffic in the Internet nowadays [25], so analysing how the network mobility support impacts the TCP throughput is important in order to comprehend what price users would have to pay in terms of performance if they access the Internet through a mobile network. Both the NEMO Basic Support protocol and MIRON solution are tested for the purpose of comparing them and justifying the necessity of a Route Optimisation solution that mitigates the poor performance of the basic solution.

The test consists of measuring the average TCP throughput of an MNN (in the tests, an LFN) downloading a file located at a CN, while two other non-mobile network hosts (Fixed nodes 1 and 2), attached to the same network the NEMO is visiting, simultaneously download the same file, both in a non-

⁵<http://www-x.antd.nist.gov/nistnet/>

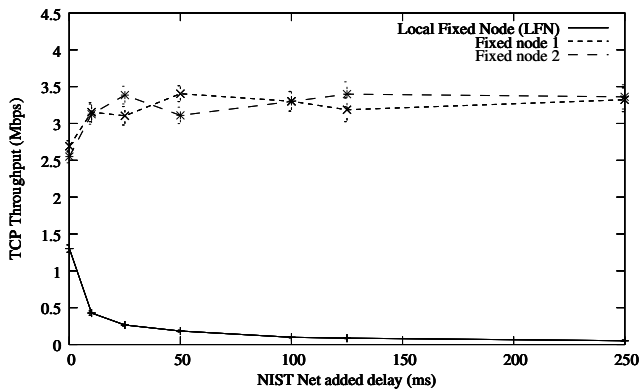


Fig. 6. Impact of NEMO Basic Support protocol on the TCP throughput

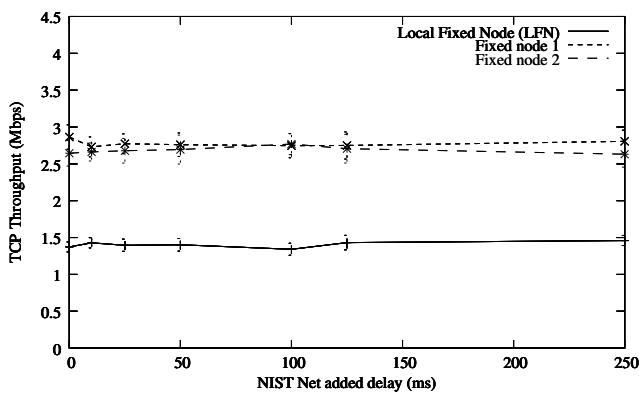


Fig. 7. Impact of MIRON on the TCP throughput

nested and in a nested scenario (see Figs. 5(a) and 5(b)). The available bandwidth between the CN and the network that the mobile network is visiting was limited to 10Mbps, by setting the R1-R2 link to 10Mbps Half-Duplex. The tool used for the download was *scp* (secure copy) and the size of the file was 50 MBytes.

Each average TCP throughput sample was calculated over a 20 seconds independent interval of download and at least 30 samples were obtained for each test (to guarantee the statistical validity of the measurements).

For the non-nested scenario, the unidirectional NIST Net added delay of the link R3-R4 – *delay1* – was varied between 0ms (i.e. home network, visited network and CN locate very close each other) and 250ms (this value represents a high, but still common RTT value of 500ms in the Internet nowadays). *Delay1* is part of the CN-HA and HA-MR delays, thus affects the overall delay in the CN-HA-MR-LFN path followed by packets of the CN-LFN communication. Results for the case of using the NEMO Basic Support protocol are shown in Fig. 6. Results for the case of using MIRON are shown in Fig. 7. Confidence limits (95%) are also shown in both figures.

If the NEMO Basic Support protocol is used, the effect of a higher value of *delay1* in the performance of TCP application is clear: the effective throughput decreases as the delay increases (Fig 6). The LFN would obtain a much higher effective throughput if it was connected directly to the foreign network instead of the NEMO. This difference in the

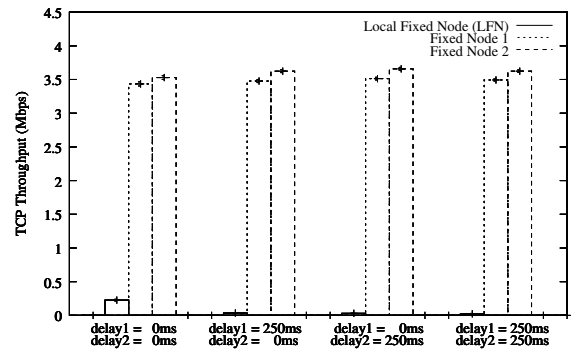


Fig. 8. Impact of NEMO Basic Support protocol on the TCP throughput in a 2-level nested Mobile Network

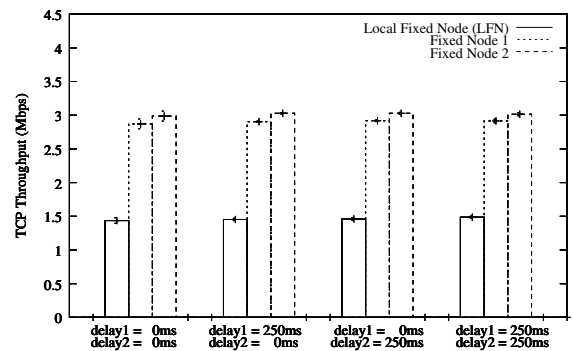


Fig. 9. Impact of MIRON on the TCP throughput in a 2-level nested Mobile Network

throughput increases with the delay in the CN-HA-MR-LFN path. Therefore, nodes of a mobile network located far from its home network and/or from the CN they are communicating with, would obtain extremely low TCP throughput when competing with other TCP flows, because of the suboptimal path introduced by the NEMO Basic Support protocol. Even for a value of *delay1* equal to 0ms the throughput obtained by the MNN is almost a half of the one obtained by the non-mobile nodes. Although *delay1* is 0ms, the RTT between CN and MNN is bigger than the RTT between CN and fixed nodes, because the path is not direct and there are more hops, and this difference, even though small, has an important effect on the TCP fairness. Moreover, there exists a difference in the PMTU because of the overhead that also has an influence in the TCP performance.

If MIRON is used, the performance improvement is substantial (Fig. 7). The TCP throughput remains constant despite the value of *delay1*. This result is as expected, because with MIRON data packets do not follow the CN-HA-MR-LFN suboptimal path, but the direct CN-MR-LFN path. Part of the difference in the TCP throughput of the fixed nodes and the LFN is due to the packet overhead (MIRON introduces a 24-byte per packet overhead, because of the Routing Header type 2 and the Home Address destination option). The performance of the MIRON prototype used during the tests (completely implemented in user space) may also have something to do with the obtained difference, although this difference could be reduced by improving the implementation (e.g., by implementing it in kernel space, or at least those tasks that have an

strong impact in the overall performance).

For the nested scenario (Fig. 5(b)), besides evaluating the effect of the varying $delay1$, that is, the delay of the path CN-HA-MR-LFN, a second adjustable delay – $delay2$ – was introduced between R4 and R5, allowing us to evaluate also the effect of the distance between the home networks of two different mobile networks that are nested. Fig. 8 shows the obtained throughput results for the NEMO Basic Support protocol and Fig. 9 for MIRON.

As in the non-nested test, the improvement achieved by MIRON is clear. The NEMO Basic Support protocol performs worse than in the non-nested scenario, even for the null added delay case. This is because the actual RTT is bigger for the LFN than for the fixed nodes due to the longer path that packets have to traverse (CN-HA2-HA1-MR1-MR2-LFN) and the reduced PMTU. On the other hand, the performance obtained with MIRON is the same as in the non-nested scenario, as packets follow the optimal direct path and the overhead remains the same, no matter what number of nesting levels the mobile network has. As in the non-nested scenario, the TCP throughput of the LFN is lower than the one achieved by the fixed nodes because of higher RTT (packets go through more intermediate hops – MR1 and MR2) and the impact of implementing MIRON completely in user space.

B. Analytical evaluation

We have analysed how the added delay due to the sub-optimal CN-HA-MR-MNN path introduced by the use of the NEMO Basic Support protocol affects the performance of TCP applications. In addition to the severe effect that the overall RTT has in the TCP performance, and the obvious effect that the delay itself has on real time applications⁶, there is another effect that impacts performance: the packet overhead (and the associated PMTU reduction).

A 40-byte IPv6 header is added to every packet in the MR-HA bidirectional path due to the NEMO Basic Support protocol. Moreover, an IPv6 additional header is added per nesting level. The effect of this overhead can be negligible for non real time applications, but it can be very important for real time ones, such as VoIP applications. In order to quantitatively evaluate this effect, we analyse next the effects of the NEMO Basic Support protocol and MIRON, comparing it with plain IPv4 and IPv6, in a VoIP communication using the widely utilised Skype⁷ application. Skype uses the iLBC (internet Low Bitrate Codec) [27] codec, which is a free speech codec suitable for robust voice communication over IP. The codec is designed for narrow band speech and results in a payload bit rate of 13.33 kbps with an encoding frame length of 30 ms and 15.20 kbps with an encoding length of 20 ms.

Table I shows the packet overhead and the bandwidth consumed by a VoIP communication using UDP/RTP and the iLBC codec, for plain IPv4, plain IPv6, the NEMO Basic Support protocol and MIRON. The overhead of MIRON is less than the one introduced by the NEMO Basic Support protocol

Protocol	Bitrate (kbps)	Packet Overhead (%)
IPv4	31.2	51.28
IPv6	39.2	61.22
NEMO (without nesting)	55.2	72.46
NEMO (2 nesting levels)	71.2	78.65
NEMO (3 nesting levels)	87.2	82.57
MIRON (without nesting)	48.8	68.85
MIRON (2 nesting levels)	48.8	68.85
MIRON (3 nesting levels)	48.8	68.85

TABLE I

iLBC BITRATES AND PACKET OVERHEAD (20MS ENCODING LENGTH)

and remains constant though the number of nesting levels. The reader should notice that a nested mobile network connected to the Internet through a 64 kbps connection would not be able to support this kind of VoIP traffic (VoIP applications are expected to be very important in forthcoming 4G networks). In [24] a more detailed analysis of the packet overhead in network mobility environments is presented.

V. COMPARISON WITH PREVIOUS WORK

This section presents two different approaches of Route Optimisation for NEMO and compares them with MIRON, in terms of performance, signalling load and complexity.

A possible approach to achieve Route Optimisation for NEMO [11] is to allow the Mobile Router directly to inform the CN about the location of the MNP using the NEMO Prefix Option. So far, this is simply a direct extension of the MIPv6 Route Optimisation procedure to the NEMO case. However, the security mechanism used for securing Route Optimisation in MIPv6 cannot be directly applied to this case. In MIPv6, Binding Update messages are secured through the Return Routability procedure, that verifies the collocation of the HoA and the CoA. In the case of a prefix, it is unfeasible to verify that all the addresses contained in the prefix (2^{64} addresses) are collocated with the CoA contained in the Binding Update message. In order to overcome this difficulty, a Return Routability Procedure for Network Prefix (RRNP) [28] is proposed, which consists of performing the MIPv6 Return Routability procedure with a randomly selected address of the NEMO prefix.

We will next compare this proposal (hereafter BU for Network Prefixes) with MIRON. In particular, we will consider the benefits and the costs associated with each one of them. With respect to the costs, the main difference concerns the deployment effort associated with the different proposals. MIRON, as we have already mentioned, uses the existent MIPv6 protocol unchanged. This means that the deployment of MIRON only implies modifications to the MRs. CNs do not need any upgrade since they do not require any MIRON-specific mechanism. On the other hand, BU for Network Prefixes requires not only upgrading the MRs but also upgrading all the potential correspondent nodes, i.e., all the nodes in the Internet. This is huge deployment cost, which may not be worth depending on the resulting benefits, which will be considered next.

The benefit resulting from the adoption of any of the proposals is the optimised path through which packets are routed

⁶There are analytical studies [26] that say that the maximum tolerable delay in a voice communication is about 50 ms

⁷<http://www.skype.com>

between the MR and the CN. However, the approach based on BU for Network prefixes requires less signalling than MIRON. We will next quantify the difference in order to evaluate if this overhead reduction can justify the deployment cost previously identified. Consider a moving network with N MNNs. Suppose that each MNN communicates simultaneously with M CNs in average. This means that with MIRON, $N \cdot M$ Binding Updates messages will be required to optimise these communications. On the other hand, if the approach based in BU for network prefixes is used, the number of BU required depends only of the number of different CNs that are communicating with at least one MNN. This is so, because the BU message refers to the whole MNP, implying that if two or more MNNs are communicating with the same CN, only one BU message is needed. The net benefit resulting from the adoption of BU for Network Prefixes with respect to MIRON is a reduction in the amount of BU messages required proportional to the number of MNNs that are simultaneously communicating with a common CN. It should be noted that this only applies for those CNs that do not belong to the Home Network, since those nodes residing in the Home Network already benefit from a direct routing with the mobile network thanks to NEMO Basic Support protocol. So, the benefits provided by an approach based on BU for Network Prefixes heavily depend on the expected number of MNNs that will communicate with a common CN outside the Home Network. The costs, on the other hand, are objective and account for the upgrading of all the nodes of the Internet to support the new option. Any approach requiring modifications in CNs is very unlikely to be deployed, even if it has reduced signalling cost. MIRON, on the other hand, is compatible with standard Mobile IPv6 CNs.

The NEMO Basic Support protocol when applied to the case of nested mobile networks is quite inefficient as was mentioned in section II-C. [29] proposes a solution to alleviate these inefficiencies. The proposal requires modifications in MRs and HAs, but not in LFNs, VMNs, or CNs.

The idea is the following: for packets going out of the nesting, the first MR in the path, in addition to tunnelling the packet to its HA with a header with source address its own CoA and destination address its Home Agent address, it also includes in the outer header of the packets a new type of Routing Header called Reverse Routing Header (RRH), where it introduces its own Home Address and empty slots where the rest of the MRs in the path can introduce their respective CoAs. This proposal requires the use of Tree Discovery [30] to allow the MRs to find out the level of hierarchy in the nesting (i.e. the number of slots required).

The rest of the MRs change the source address of the outer header and include their own CoA, but put the old source address (the CoA of the previous MR) in the Reverse Routing Header. When the packets leave the nesting, they are forwarded to the HA of the first MR in the path. The HA decapsulates the packets and sends them to their destination (it uses the Home Address included in the RRH to find out or create the right Binding Cache Entry), but also keeps associated to the respective Binding Cache Entry the information contained in the Reverse Routing Header. This information

allows the HA to include in the outer header of the packets destined to a node in the nesting, a Routing Header indicating how the packet must be routed inside the nesting (the CoAs of the MRs in the nesting in the order that must be traversed). The final result is that the packets in each direction only go through one tunnel and one Home Agent, although some processing is added in the HA and MRs plus the overhead of the information added to the packets.

This overhead can be quantified in one IPv6 header plus one routing header plus one IPv6 address per level of nesting of the Mobile Network, i.e. $(40 + 8 + n \cdot 16)$ bytes = $(48 + n \cdot 16)$ bytes, where n is the number of levels in the nesting (at least 2). This overhead is required in all the packets that go to and from the mobile network. It could be eliminated from some packets in the way out of the mobile network only at some cost in functionality (ability to detect changes in the nesting) and security. Notice that the solution of MIRON for nested mobile networks only requires the 40 bytes of the tunnelling and even that is avoided when an end-to-end optimisation of the path between the mobile network and the CN is used.

The additional need for using Tree Discovery [30] implies changes in MRs and routers included in the nesting, because Router Advertisements must support the functionality of Tree Discovery. This implies also an overhead in signalling because Router Advertisements in the nesting must have a minimum of 32 bytes more than normal Router Advertisements. This must be compared with the signalling load to distribute topological valid addresses to MRs in MIRON.

VI. CONCLUSIONS

The NEMO Basic Support protocol [12] enables whole networks to move and change their point of attachment, transparently to the nodes of the network. This solution introduces some limitations and problems in terms of performance (increased delay in packet delivery and packet overhead, decrease in available PMTU, the HA becoming a bottleneck, etc). To overcome these limitations we have designed and implemented a Route Optimisation solution: MIRON, that enables direct path communication between a node of the mobile network - supporting any kind of node, with and without mobility capabilities - and a Correspondent Node. MIRON does not require any change on the Correspondent Nodes (i.e. any potential peer that a node of the mobile network may want to communicate with) nor in the Mobile Network Nodes, as it only introduces changes on the MRs, thus facilitating the deployment of the solution.

MIRON has two modes of operation: the MR performing all the Route Optimisation tasks on behalf of those nodes that are not mobility capable - thus working as *Proxy MR* [14] - and an additional mechanism, based on PANA and DHCP, enabling mobility-capable nodes (i.e. Mobile Nodes attached to a NEMO) and routers (i.e. nested Mobile Routers) that actually have mobility and Route Optimisation capabilities to manage their own Route Optimisation.

To validate the design of the solution and evaluate the actual performance of it, a prototype of MIRON was implemented in Linux. The NEMO Basic Support was also implemented

so we could compare the results obtained with MIRON with the basic solution for network mobility. Tests involving TCP applications showed that the increased RTT perceived by the nodes of a NEMO (due to the suboptimal path followed by packets) has a severe impact on the performance (in terms of effective throughput, when sharing some link with traffic from other active non-mobile TCP nodes). This effect is exacerbated when NEMOs are nested. On the other hand, the same tests conducted with MIRON showed a better performance, by obtaining much higher effective TCP throughputs than in the case of the NEMO Basic Support, also in the case of nested networks.

The effect that packet overhead may have was described by means of a quantitative analytical study of the overhead that several protocols add to packets belonging to a VoIP application, such as Skype. These results show that the packet overhead introduced by the NEMO Basic Support protocol is significant for this kind of application, specially when there is nesting.

In conclusion, this paper provides a Route Optimisation for NEMO solution, that provides significant performance improvements over the NEMO Basic Support protocol. The validity of the solution has been proven by making experiments and tests with an implementation for Linux.

ACKNOWLEDGEMENT

The authors would like to thank Pascal Thubert and Erik Nordmark for their helpful comments that contributed to improve MIRON design. We also would like to thank Albert Banchs, Manuel Urueña, Huw Oliver and Isáías Martínez for reviewing and helping to improve this paper. Also, the authors would like to thank Carlos Izquierdo for the development of the DHCPv6 implementation used in the MIRON prototype.

REFERENCES

- [1] D. Yi and K. H. Yeung, "Challenges in the Migration to 4G Mobile Systems," *IEEE Communications Magazine*, December 2003.
- [2] R. Droms, *Dynamic Host Configuration Protocol*, RFC 2131, March 1997, Internet Engineering Task Force.
- [3] R. Droms, J. Bound, B. Volz, T. Lemon, C. Perkins, and M. Carney, *Dynamic Host Configuration Protocol for IPv6*, RFC 3315, July 2003, Internet Engineering Task Force.
- [4] C. Perkins, *RFC 3344: IP Mobility Support for IPv4*, August 2002, Internet Engineering Task Force.
- [5] D. Johnson, C. Perkins, and J. Arkko, *RFC 3775: Mobility Support in IPv6*, June 2004, Internet Engineering Task Force.
- [6] R. Koodli, *Fast Handovers for Mobile IPv6*, draft-ietf-mipshop-fast-mipv6-03.txt (work-in-progress), October 2004, Internet Engineering Task Force.
- [7] H. Soliman, C. Catelluccia, K. E. Malki, and L. Bellier, *Hierarchical Mobile IPv6 mobility management (HMIPv6)*, draft-ietf-mipshop-hmipv6-04.txt (work-in-progress), December 2004, Internet Engineering Task Force.
- [8] R. Moskowitz and P. Nikander, *Host Identity Protocol Architecture*, draft-ietf-hip-arch-02 (work-in-progress), January 2004, Internet Engineering Task Force.
- [9] P. Nikander, J. Ylitalo, and J. Wall, "Integrating Security, Mobility and Multi-homing in a HIP Way," in *Proceedings of Network and Distributed Systems Security Symposium (NDSS'03)*, San Diego, USA, February 2003, pp. 87–99.
- [10] T. R. Henderson, "Host Mobility for IP Networks: A comparison," *IEEE Network*, November/December 2003.
- [11] H.-Y. Lach, C. Janneteau, and A. Petrescu, "Network Mobility in Beyond-3G," *IEEE Communications Magazine*, July 2003.
- [12] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, *Network Mobility (NEMO) Basic Support Protocol*, RFC 3963, January 2005, Internet Engineering Task Force.
- [13] C. J. Bernardos, M. Bagnulo, and M. Calderón, "MIRON: Mobile IPv6 Route Optimization for NEMO," in *4th Workshop on Applications and Services in Wireless Networks. ASWN 2004*. IEEE, August 2004.
- [14] C. Ng, P. Thubert, H. Ohnishi, and E. Paik, *Taxonomy of Route Optimization models in the NEMO Context*, draft-thubert-nemo-ro-taxonomy-03 (work-in-progress), October 2004, Internet Engineering Task Force.
- [15] P. Jayaraman, R. Lopez, Y. Ohba, M. Parthasarathy, and A. Yegin, *PANA Framework*, draft-ietf-pana-framework-03 (work-in-progress), December 2004, Internet Engineering Task Force.
- [16] T. Ernst and H.-Y. Lach, *Network Mobility Support Terminology*, draft-ietf-nemo-terminology-02 (work-in-progress), October 2004, Internet Engineering Task Force.
- [17] T. Ernst, *Network Mobility Support Goals and Requirements*, draft-ietf-nemo-requirements-03 (work-in-progress), October 2004, Internet Engineering Task Force.
- [18] A. Yegin, Y. Ohba, R. Penno, G. Tsirtsis, and C. Wang, *Protocol for Carrying Authentication for Network Access (PANA) Requirements*, draft-ietf-pana-requirements-09 (work-in-progress), August 2004, Internet Engineering Task Force.
- [19] D. Forsberg, Y. Ohba, B. Patil, H. Tschofenig, and A. Yegin, *Protocol for Carrying Authentication for Network Access (PANA)*, draft-ietf-pana-pana-07 (work-in-progress), December 2004, Internet Engineering Task Force.
- [20] L. Blunk and J. Vollbrecht, *PPP Extensible Authentication Protocol (EAP)*, RFC 2284, March 1998, Internet Engineering Task Force.
- [21] F. Zhao, S. F. Wu, and S. Jung, *NEMO Route Optimization Problem Statement, Requirements and Evaluation Considerations*, draft-zhao-nemo-ro-ps-00.txt (work-in-progress), October 2004, Internet Engineering Task Force.
- [22] E. Perera, V. Sivaraman, and A. Seneviratne, "Survey on Network Mobility Support," *ACM SIGMOBILE Mobile Computing and Communications Review*, 2004.
- [23] P. Nikander, J. Arkko, T. Aura, G. Montenegro, and E. Nordmark, *Mobile IP version 6 Route Optimization Security Design Background*, draft-ng-nemo-rrnp-00 (work-in-progress), December 2003, Internet Engineering Task Force.
- [24] A. de la Oliva, C. J. Bernardos, and M. Calderón, "Practical evaluation of a network mobility solution," in *EUNICE*, 2005.
- [25] S. McCreary and K. Claffy, "Trends in wide area IP traffic patterns - A view from Ames Internet Exchange," CAIDA, Tech. Rep., 2000.
- [26] M. K. M and F. Tobagi, "Analysis of the delay and jitter of voice traffic over the internet," in *INFOCOM*, April 2001.
- [27] S. Andersen, A. D. Telio, H. Astrom, R. Hagen, W. Kleijn, and J. Linden, *Internet Low Bitrate Codec*, RFC 3951, December 2004, Internet Engineering Task Force.
- [28] C. Ng and J. Hirano, *Extending Return Routability Procedure for Network Prefix (RRNP)*, draft-ng-nemo-rrnp-00 (work-in-progress), October 2004, Internet Engineering Task Force.
- [29] P. Thubert and M. Molteni, *IPv6 Reverse Routing Header and its application to Mobile Networks*, draft-thubert-nemo-reverse-routing-header-05 (work-in-progress), June 2004, Internet Engineering Task Force.
- [30] P. Thubert and N. Montavont, *Nested Nemo Tree Discovery*, draft-thubert-tree-discovery-01 (work-in-progress), October 2004, Internet Engineering Task Force.