

IP Flow Mobility in PMIPv6 Based Networks: Solution Design and Experimental Evaluation



Telemaco Melia · Carlos J. Bernardos ·
Antonio de la Oliva · Fabio Giust · Maria Calderon

Abstract The ability of offloading selected IP data traffic from 3G to WLAN access networks is considered a key feature in the upcoming 3GPP specifications, being the main goal to alleviate data congestion in cellular networks while delivering a positive user experience. Lately, the 3GPP has adopted solutions that enable mobility of IP-based wireless devices relocating mobility functions from the terminal to the network. To this end, the IETF has standardized Proxy Mobile IPv6 (PMIPv6), a protocol capable to hide often complex mobility procedures from the mobile devices. This paper, in line with the mentioned offload requirement, further extends PMIPv6 to support dynamic IP flow mobility management across access wireless networks according to operator policies. Considering energy consumption

The research of Carlos J. Bernardos, Antonio de la Oliva, Telemaco Melia and Fabio Giust leading to these results has received funding from the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement n. 258053 (MEDIEVAL project). Carlos J. Bernardos and Maria Calderon have also received funding from the Ministry of Science and Innovation of Spain, under the QUARTET project (TIN2009-13992-C02-01).

T. Melia
Bell Labs France, Paris, France
e-mail: Telemaco.Melia@alcatel-lucent.com

C. J. Bernardos · A. de la Oliva · M. Calderon (✉)
Department of Telematics Engineering, University Carlos III of Madrid, Madrid, Spain
e-mail: maria@it.uc3m.es

C. J. Bernardos
e-mail: cjb@it.uc3m.es

A. de la Oliva
e-mail: aoliva@it.uc3m.es

F. Giust
Department of Telematics Engineering, University Carlos III of Madrid, Madrid, Spain
e-mail: fabio.giust@imdea.org

F. Giust
Institute IMDEA Networks, University Carlos III of Madrid, Madrid, Spain

as a critical aspect for hand-held devices and smart-phones, we assess the feasibility of the proposed solution and provide an experimental analysis showing the cost (in terms of energy consumption) of simultaneous packet transmission/reception using multiple network interfaces. The end-to-end system design has been implemented and validated by means of an experimental network setup.

Keywords Proxy Mobile IPv6 · Flow mobility · Experimental evaluation

1 Introduction

The exponential growth in mobile data applications and the resultant increase of traffic volume in 3G data networks has placed mobile operators in the challenging position—particularly when licensed spectrum is limited—of supporting large amounts of traffic chunks. With much of this increased IP data traffic directly attributable to the availability of affordable smart-phones featuring both 3G and WLAN access, mobile operators are now looking at WLAN networks as a low cost alternative to offload data from their 3G infrastructure. Offloading alleviates data congestion in cellular networks while delivering a positive user experience.

A first approach to the problem could be to perform an inter-technology handoff whenever WLAN connectivity becomes available, with all the traffic routed through the WLAN access. However having the capability to move selected IP traffic (i.e., HTTP, video, etc.) while supporting simultaneous 3G and WLAN access seems a more appealing solution. In this environment, mobile operators can develop policies for IP flow mobility, and control which traffic is routed over the WLAN and which one is kept on the 3G. For example, it seems reasonable that some IP flows (e.g., related to VoIP) are sent over 3G to benefit from its QoS capabilities, while IP flows related to “best-effort” Internet traffic can be moved to the WLAN access. Inter-working between 3G and WLAN access networks is not a new topic by itself, however the availability of smart-phones to the mass market and the proliferation of new applications renewed the interest by mobile operators in the subject.

Lately, we have been assisting to the development of new solutions that enable IP mobility of wireless devices within a local domain by means of special purpose functions installed in network components. We refer to these solution as network-based mobility management, as opposed to host-based mobility management (e.g., Dual Stack Mobile IP [1]).

Network-based Localized Mobility Management (NetLMM) [2] allows conventional IP devices to roam across wireless access networks without the support of mobility clients. This is an appealing feature from the service provider’s viewpoint, since it enables mobility support without strong dependence on software and complex mobility related configuration in the user terminals. To this end, the IETF has standardized Proxy Mobile IPv6 (PMIPv6) [3,4]. However, current specifications only provide mobility management at the granularity of interfaces, meaning that the network is only able to move all the communications associated with a particular interface of a mobile node, but they do not consider more granular management strategies.

This paper focuses on the design and implementation of flow mobility extensions for PMIPv6. It describes the functional components required in the network to support smart traffic steering while minimizing the impact on the mobile devices and augmenting user Quality of Experience (QoE). In our proposal, the network (in particular the mobility anchor) is the decision control entity. It performs flow mobility based on network operator policies, which may dynamically react upon the network load. We consider two different types of mobile devices: (i) terminals with a single interface visible from the IP stack, where the link-layer hides the use of multiple physical interfaces as in [5,6], and (ii) terminals with

multiple IP interfaces visible to the upper layers where the IP stack behaves according to the *weak host* model [7,8]. Our customized PMIPv6 protocol stack has been extended to support both types of terminals and an experimental evaluation has been carried out. The experimental results demonstrate the viability of performing flow mobility in network-based mobility management scenarios.

One could argue that the simultaneous use of two or more wireless interfaces can be a blocking factor to the wide adoption of seamless IP flow mobility management, due to the additional battery consumption. To show its feasibility we have analyzed the energy consumption of a simultaneous use of multiple network interfaces, focusing on WLAN and 3G access. The tests, conducted on an experimental platform, successfully demonstrate the feasibility of the approach.

The rest of the article is organized as follows. In Sect. 2, we provide an overview of the PMIPv6 protocol, highlighting the motivation to enable IP flow mobility in this scenario, and evaluating—from an energy point of view—the cost incurred by enabling IP flow mobility. Section 3 presents the details of our proposed flow mobility solution for PMIPv6. Next, Sect. 4 reports on the results of our experimental evaluation. Section 5 compares our solution with existing work. Finally, we conclude the article in Sect. 6.

2 Background and Motivation

2.1 Network-Based Localized Mobility Management: Proxy Mobile IPv6

Unlike client-based mobility, such as Mobile IPv6 [9], where Mobile Nodes (MNs) signal a location change to the network to update routing state and in this way maintain reachability, NetLMM [2] approaches provide mobility support to moving hosts (e.g., IP hosts changing its attachment to the network) without their involvement. This is achieved by relocating relevant functionality for mobility management from the MN to the network. In a Localized Mobility Domain (LMD), the network learns through standard terminal operation, such as router and neighbor discovery [10] or by means of link-layer support [11], about an MN's movement and coordinates routing state updates without any mobility specific support from the terminal. While moving inside the LMD, the MN keeps its IP address, and the network is in charge of updating its location in an efficient manner [12,13]. Proxy Mobile IPv6 [3] is the NetLMM protocol proposed by the IETF. This protocol is based on Mobile IPv6 (MIPv6) [9], extending the MIPv6 signalling messages and reusing the Home Agent (HA) concept.

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

- **Mobile Access Gateway (MAG).** This entity performs the mobility related signalling on behalf of an MN that it is attached to its access link. The MAG is usually the access router for the MN, i.e., the first hop router in the Localized Mobility Management infrastructure. It is responsible for tracking the MN's movements on the access link. There are multiple MAGs in an LMD.
- **Local Mobility Anchor (LMA).** This is an entity within the backbone network that maintains a collection of routes for individual MNs within the LMD (i.e., it is the entity that manages the MN's binding state). The routes point to MAGs managing the links in which the MNs are currently located. Packets for an MN are routed to and from the MN through tunnels between the LMA and the corresponding MAG. The LMA is also responsible for assigning IPv6 prefixes to MNs (e.g., it is the topological anchor point for the prefixes assigned to the MN). There may be more than one LMA in an LMD.




Fig. 1 Proxy Mobile IPv6 domain

Once an MN enters an LMD and attaches to an access link, the MAG in that access link, upon identifying the MN, performs mobility signalling on behalf of the MN. The MAG sends to the LMA a Proxy Binding Update (PBU) associating its own address with the MN's identity (e.g., its MAC address or an ID related with its authentication in the network). Upon receiving this request, the LMA assigns a prefix—called Home Network Prefix (HNP)—to the MN (i.e., allocate a prefix for the attached interface). Then, the LMA sends to the MAG a Proxy Binding Acknowledgement (PBA) including the prefix assigned to the MN. Then, the MN is able to configure one or more addresses from the assigned prefix. The LMA also creates a Binding Cache Entry (BCE) and establishes a bi-directional tunnel to the MAG (the end-point of this tunnel on the MAG side is called Proxy Care-of Address—Proxy CoA). Whenever the MN moves, the new MAG updates the MN's location in the LMA, advertises the same prefix to the MN (through unicast Router Advertisement messages) and shows the same layer-2 and layer-3 identifiers to the MN, thereby making the IP mobility transparent to the MN. Thus, the MN can keep the address configured when it first entered the LMD, even after changing its point of attachment to the network.

In the context of PMIPv6 specification, the term mobility session refers to the creation or existence of state associated with the mobile node's mobility binding on the LMA and on the serving MAG. If the mobile node connects to the PMIPv6 domain through multiple interfaces, simultaneously, each of the attached interfaces will be assigned a unique set of home network prefixes, and all the prefixes assigned to a given interface of a mobile node will be managed under one mobility session.

2.2 IP Flow Mobility

We are witnessing that the number of wireless mobile subscribers accessing data services does not stop increasing. This is motivated by a variety of different reasons: 3G access is widely available (coverage reaches almost 100% of dense populated areas in developed countries) and affordable by users (most mobile handsets are 3G capable, USB modems are quite cheap and operators offer flat rates to their customers). Besides, the number and popularity of applications designed for smart-phones that make use of Internet connectivity is getting larger every day, contributing to an increase of market penetration of such devices (e.g., iPhone, Android, Blackberry and Windows Mobile phones), which results in growing demands for 3G connectivity everywhere. Due to the increasing demand for 3G connectivity, operators are challenged to enhance their network deployments.

Driven by this continuous growth on the users' demand for connectivity and the high costs of 3G deployment (mainly caused because the radio spectrum is limited), the use of disparate heterogeneous access technologies—what is commonly referred to as 4G [14]—is considered as a mechanism to expand network capacity. This extension is not only achieved in terms of effective coverage (i.e., one particular access technology might not be offered in certain locations, while others could be deployed as an alternative way of accessing the network) but also in terms of simultaneously available bandwidth (i.e., the effective data rate that could be achieved by using two or more access technologies at the same time). User devices equipped with multiple radios (also known as multi-mode terminals) would be potentially capable of improving the connectivity experience they provide by simultaneously using more than one single access technology. Mobile operators see today an opportunity of reducing the average cost per offered Megabyte (and therefore an increase of their revenue) by introducing an intelligent resource management mechanism that allows to offload traffic from the 3G network into other access candidate networks (mainly WLAN due to its high penetration and rate) when available. This optimizes the operator's network use, while keeping the users' QoE.

Fully exploiting heterogeneity in the network access—e.g., enabling 3G offload—has proved to be difficult. Most of today's solutions enable the use of different technologies (e.g., 3G and WLAN) by adopting one of the following approaches (or a combination of them): (i) manual user-based switching, or (ii) application-based switching. In the former case, users decide to switch on a network interface based on their preferences (e.g., cost, required bandwidth for the applications being used, WLAN availability, etc.), while in the latter, applications decide to turn on and off interfaces based on predefined preferences and network availability. Both approaches involve a change on the IP address seen by the applications, and therefore rely on them surviving that change (or re-establishing the sessions). Operators are not satisfied with any of these approaches, as they leave the mobility control on the final users and/or the application developers. Additionally, the QoE obtained by users in this case may not be good enough, as it depends on the application behavior or requires the sessions to be restarted.

The 3GPP and IETF are currently working towards the definition and specification of much richer solutions which aim at enabling true flow mobility. Flow mobility refers to the movement of selected flows from one access technology to another, minimizing the impact on the users' QoE. Solutions for both Dual Stack Mobile IP (DSMIP) [1] and PMIPv6 are being explored, but here we focus on flow mobility extensions for PMIPv6, as this protocol does not require to install and configure a mobility stack on the user's terminal, and allows for a better mobility control by the network.

2.3 Flow Mobility for PMIPv6

A first step required in order to support flow mobility is the capacity to use several physical network interfaces. Proxy Mobile IPv6 allows an MN to connect to the same PMIPv6 domain through different interfaces, though in a very limited way. There are three possible scenarios [15]:

- *Unique set of prefixes per interface.* This is the default mode of operation in PMIPv6. Each attached interface is assigned a different set of prefixes, and the LMA maintains a mobility session (i.e., a binding cache entry) per MN's interface. PMIPv6 only allows to transfer all the prefixes assigned to a given interface to another one attaching to the same PMIPv6 domain, and does not fully specify how a MAG can figure out if a new mobile node wants to get a new set of prefixes assigned (i.e., having simultaneous access via multiple interfaces) or if the mobile node is performing a handover (i.e., the MN wants to transfer the prefixes bound to a previous interface to the new one).
- *Same prefix but different global addresses per interface.* In this case the same prefix is assigned to multiple interfaces, though a different address is configured on each interface. This mode is not completely supported by PMIPv6. It either requires two different mobility sessions (as in the previous scenario) or only one but two separate host route entries. In any case, this scenario creates a multi-link subnet as the same prefix is advertised over different point-to-point links. This kind of scenario presents some issues as documented in [16].
- *Shared address across multiple interfaces.* In this scenario, the MN is assigned the same IP address across multiple interfaces. This enables applications on the terminal to see and use only one address, and therefore the MN could be able to benefit from transparent mobility of flows between interfaces. This scenario is not supported by current PMIPv6, it requires one mobility session per terminal and some kind of flow filters/routes at the LMA to be able to forward packets via the appropriate MAG. Besides, ensuring that multiple IP interfaces of the same device configure the same IP address is not easy to achieve (e.g., IPv6 specs assume that unique IPv6 addresses are configured per interface, as guaranteed by running Duplicate Address Detection, DAD) nor to operate (not all Operating Systems support assigning the same IP address to multiple interfaces, and the multi-link subnet issue also appears here). One approach to mitigate this is to make use of link layer implementations that can hide the actually used physical interfaces from the IP stack [17]. For instance, the *logical interface* solution at the IP layer may enable packet transmission and reception over different physical media [5,6].

Proxy Mobile IPv6 as defined in [3] cannot provide flow mobility in any of the previously described scenarios. We next identify and describe what functionality is missing from PMIPv6 to support flow mobility, by making use of an example. Figure 2 shows a potential use case of interest involving a multi-mode terminal attached to a PMIPv6 domain. The MN is attached to MAG1 through its WLAN interface (*if1*), and to MAG2 through its 3G interface (*if2*). With current PMIPv6 specification (*plain* PMIPv6, see Fig. 2a), each interface is assigned a different prefix by the LMA (to allow simultaneous access), and two different mobility sessions (i.e., two separate binding cache entries) are maintained at the LMA. PBU/PBA signalling is used to keep alive the bindings at the LMA or to completely transfer the whole set of assigned prefixes from one interface to another. In order to support flow mobility, the state at the LMA needs to be extended (*extended* PMIPv6, see Fig. 2b), so the LMA is able to group mobility bindings referring to the same MN. Additionally, flow state should be introduced at the LMA, so it can forward packets differently (i.e., through different

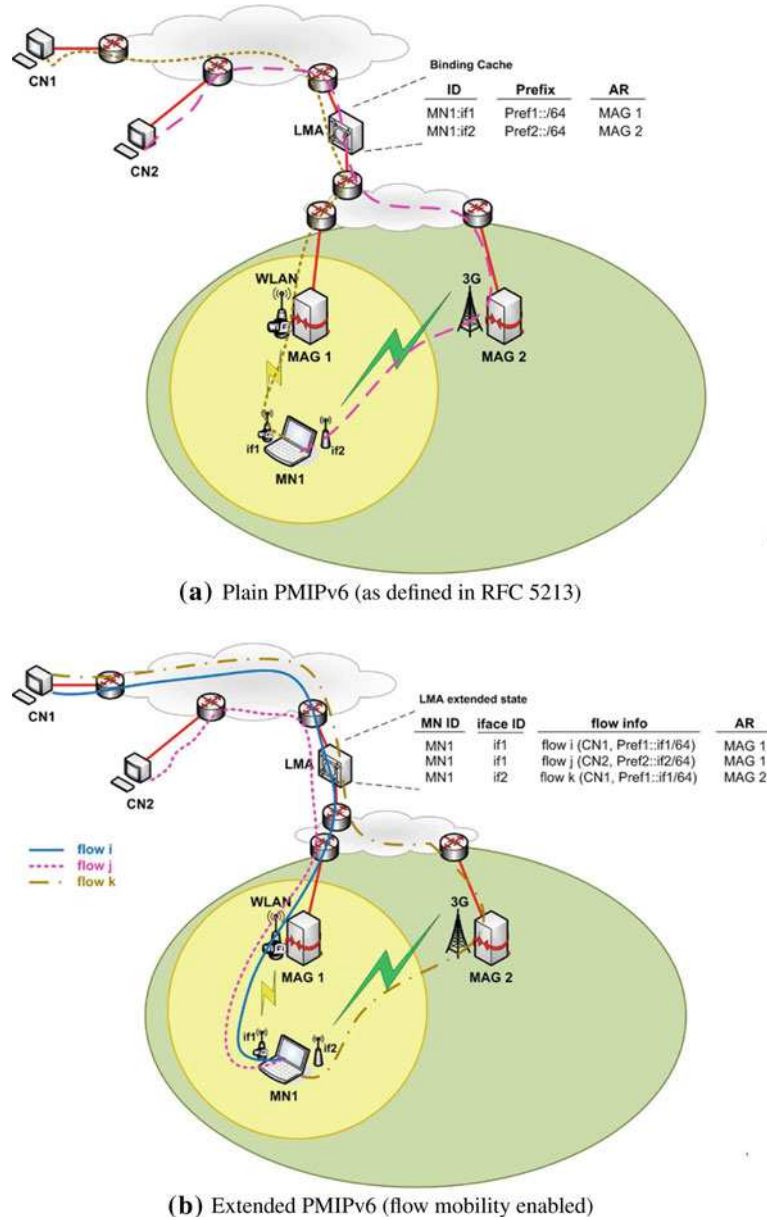


Fig. 2 Flow mobility in PMIPv6: what is missing?

MAGs) on a per-flow basis. The MAG behavior needs also to be modified, since the MAG should be aware of all the MNs' IP addresses that are reachable through the point-to-point link it has set up with the MN. In order to transfer this information, the PMIPv6 signalling between the MAG and the LMA has to be extended as well.

The mobile node behavior needs also to be considered. In the plain PMIPv6 scenario, the IPv6 addresses assigned to if1 (addr1) and if2 (addr2) are different

(`Pref1::if1/64` and `Pref2::if2/64`, respectively). Packets addressed to `addr1` will always arrive via `if1` (and the same for packets addressed to `addr2`, arriving via `if2`). In a flow mobility-enabled scenario, `addr1` and `addr2` may belong to different prefixes, belong to the same one, or even be the same IP address. Moreover, packets addressed to `addr1` may arrive at `if2` (and the other way around), and should be processed by the MN normally.

In Sect. 3 we describe in detail our PMIPv6 extensions to support flow mobility, from both the network (i.e., changes to the LMA and MAG operations) and the mobile node viewpoint.

2.4 Energy Cost of a Flow Mobility Solution

As it has been discussed, enabling flow mobility enhances overall satisfaction of both operators and users. However there are two main issues that should be analyzed to assess if a PMIPv6 and flow mobility enabled solution is feasible in a real deployment. First issue is—as in any communications system—the complexity of the solution, in terms of protocol overhead and ease of configuration and maintenance (we elaborate more on this in Sects. 3 and 4). Second issue is the energy cost associated with using multiple network interfaces simultaneously, which is the focus of this section.

Energy consumption is particularly critical for hand-held devices and smart-phones, which already suffer from reduced battery life compared to plain mobile phones. The use of 3G is known to drain battery life faster than 2G (actually, most mobile phones allow the user to disable the use of 3G). However, current smart-phones make an intensive use of 3G and stay almost “always-on” (this is particularly true for the case of Android phones). In 3GPP Rel-8 and next releases, the concept of *always-on*¹ is introduced and future terminals are expected to implement it. Enabling and turning on additional network interfaces leads to an increase of the energy consumption, and the question that needs to be answered is whether this increase is affordable by the user’s terminal.

To perform an experimental assessment of the energy cost derived from enabling IP flow mobility (i.e., use of multiple network interfaces at the same time) we perform real power consumption measurements on a multi-mode device, equipped with a WLAN IEEE 802.11a/b/g and a 3G UMTS (HSDPA capable) interface. In order to be able to control as much as possible the used devices, capture traffic sent/received at the network interfaces, as well as closely monitor the device, we decided to use a small residential router based on a Linux firmware (an Asus WL-500GP v1.0). We conjecture that the conclusions we learnt from these experiments are also valid for the case of smart-phone devices, as the key part is to use a device which energy consumption under regular operation is low enough to allow noticing the difference in energy cost when a network interface is activated and used.

The Asus WL-500GP v1.0 is equipped with a 266MHz processor, an IEEE 802.11b/g WLAN interface and an IEEE 802.3 Ethernet interface connected to a VLAN capable 5-port switch. This version of the router has a mini-PCI slot that allows to change the original wireless card. We remove the original Broadcom card and insert instead an Atheros based 802.11a/b/g (Alfa Networks AWPCI085S) one. This card is supported by the Madwifi² driver. In order to mitigate as much as possible the impact of collisions and interference in the power consumption measurements, we avoid the use of the 2.4 GHz band (IEEE 802.11b/g)—which

¹ In the context of 3GPP, “always-on” refers to the following: a default bearer is established after the terminal attaches to the network, meaning that a Packet Data Protocol (PDP) context is set up and an IPv6 address is configured. This best-effort QoS bearer is kept during all the MN’s network attachment lifetime.

² <http://www.madwifi.org/>

is very crowded in our lab, as reported in [18]—and configure the WLAN interface in 802.11a mode.

The firmware of this router can be replaced with an open source Linux-based firmware. We install the OpenWRT³ Kamikaze 8.09.2 distribution with a Linux-2.6 kernel in the routers. This firmware gives us more flexibility in the use and configuration of the routers than the original firmware, and allows for example the configuration and use of a 3G USB stick modem. For our tests, we use a Huawei E160 HSDPA USB stick.⁴

Power consumption is measured using a PCE-PA 6000 power analyzer.⁵ Measurement of power is done using a PCE-PA-ADP current adaptor where the power supply of the router is plugged in. Measurement data is transferred from the power analyzer to a computer via an RS-232 interface for its processing.

Using this setup we perform the measurements described next. We first calibrate the power analyzer by measuring the consumption when both the WLAN and 3G interfaces are switched off. All reported results are relative to this level. For the actual measurements, we are interested in the power consumption when the network interfaces are in the following states:

- OFF: the interface is switched off.
- IDLE: the interface is on but it does not send/receive any data traffic. For the case of WLAN, this means that the card is associated to an access point (so the card is receiving beacon frames) without sending/receiving any user data traffic. For the case of 3G, this means that the interface is up, a PDP context has been activated and a PPP interface has been set up, but no data is exchanged.
- ON: the interface is on and engaged in a data traffic exchange. In our tests, this means that a file is downloaded from a server using HTTP. By using TCP, the card is receiving at the maximum available rate, and traffic is sent in both directions (downlink: mostly data segments, uplink: mostly TCP acknowledgements).

We measure the power consumption for different possible states of the WLAN and 3G interfaces. Table 1 shows the obtained results (mean and 95% confidence interval obtained from five 300 sec experiments). We focus on the scenarios in which at least one of the interfaces is actively involved in sending/receiving traffic, as those are the cases in which it is important to evaluate the energy cost associated with having a second active interface. This second interface may be either receiving/sending traffic or just idle, ready to be used. Results show that the 3G interface consumes more energy than the WLAN one, and that the difference between the case of only using the 3G interface (which is currently the most common one) and the case of using simultaneously the 3G and the WLAN interfaces is only of 20%. Note that this additional cost is only incurred when both interfaces are actively engaged in a data transfer, and that by using them simultaneously the time required to send a given amount of data via WLAN would be shorter—since the throughput obtained via a WLAN network is typically higher than the one that can be obtained via a 3G network—and this would also contribute to a lower power consumption. The extra power consumption caused by activating the WLAN interface (IDLE state) is just of about 3%, which besides would only be needed when the mobile is sending/receiving traffic, as it is then when the network operator and the user may benefit from offloading traffic from the 3G infrastructure to a WLAN hotspot (if available).

³ <http://www.openwrt.org/>

⁴ <http://www.huawei.com/mobileweb/en/products/view.do?id=1960>

⁵ <http://www.industrial-needs.com/technical-data/power-analyser-PCE-PA-6000.htm>

Table 1 Power consumption results

3G ON		WLAN ON	
WLAN OFF	1.80 ± 0.10 W	3G OFF	1.03 ± 0.08 W
WLAN IDLE	1.86 ± 0.08 W	3G IDLE	1.21 ± 0.16 W
WLAN ON	2.16 ± 0.13 W	3G ON	2.16 ± 0.13 W

In order to better understand the impact of this additional power consumption, we next include a simple example of how battery life of a typical smart-phone would be affected. This is important for the users of this kind of device, who are typically concerned about the battery life of their devices. Let's make some assumptions, for the sake of the simplicity of this analysis, which aims at assessing if a typical mobile user could afford the additional power consumption introduced by the use of flow mobility extensions. Several studies, such as [19], point out that users of smart hand-held devices download an average of 20 MBytes per day via 3G. We have performed tests with two different models of 3G enabled hand-helds—an HTC Magic and an iPhone 3GS—in which we measure the energy consumed when a 20 Mbyte file is downloaded. For the case of the HTC device, an average of the 7.9% of the battery is consumed, while for the iPhone, it is just a 0.9% of the battery.⁶ In case a flow mobility solution was enabled in the terminals, the power consumption would have been increased at most by 20% during the download, meaning that the overall resulting reduction of the battery life would have been of 9.48 and 1.06%, respectively for the HTC and iPhone devices. If we just assume that a normal user is able to operate its terminal during a whole day without charging the battery (i.e., a full battery lasts for 24 h of use at least, including idle state), the extra power consumption incurred by the use of a flow mobility solution would just cause a reduction of the battery life of less than 23 min for the HTC device, and less than 3 min for the iPhone. This simple analysis does not aim at providing rigorous and precise figures, but just at roughly assessing if a flow mobility solution is affordable from the perspective of power consumption. Based on the obtained results, we can conclude that selectively switching on and using more than one network interface results in an affordable additional cost.

3 Solution Description

In this section we present the design of a solution enabling flow mobility for PMIPv6. An overview of the proposed mechanism is followed by the detailed description of the solution.

We first define the term *flow*. A flow is intended as a stream of packets that traverses the LMA to/from the MN, regardless of which entity started the communication or which transport protocol is being used. A flow is univocally identified by 6 parameters—also referred to as flow 6-tuple: (i) Source IP address; (ii) Destination IP address; (iii) IPv6 flow label field; (iv) IPv6 next header field (transport); (v) Source port; (vi) Destination port.

⁶ Note that capacity of the batteries of the HTC and iPhone are different, as well as the network operators (and therefore the performance provided by the 3G networks used in the experiments), so we cannot directly compare the values obtained for the battery consumption.

3.1 Protocol Overview

As outlined in Sect. 2.3, a solution enabling flow mobility for PMIPv6 requires, on the one hand, extensions on the mobility signalling between the LMA and the MAG and, on the other hand, modifications to the behavior and data structures maintained by the LMA and the MAG.

The LMA is the decision control entity in the proposed approach. It performs flow routing based on operator policies—which may be dynamic to allow performing flow balancing to adapt to the network load—and/or other external triggers. The LMA enforces in this way which interface is used by the MN to receive downlink data traffic. For the uplink traffic, there are potentially several different approaches that the MN may follow. For example, the decision can be taken by the MN itself, selecting which interface to use independently of the LMA, however this could lead to asymmetric routing in the uplink-downlink paths.⁷ So, we propose that the MN uses to send uplink traffic the same interface that is used to receive downlink packets belonging to the same flow. Following this approach, the MN *replicates* the decisions made by the LMA for the downlink traffic when sending uplink traffic, and consequently replicating any posterior changes that the LMA may perform during the flow lifetime.

Due to the fact that PMIPv6 does not require the MN to implement nor participate in any mobility protocol, considerations about how the terminal behaves are very relevant. In this paper we consider two different kinds of IPv6 mobile nodes:

1. *Terminals with a single interface visible from the IP stack.* Certain link-layer implementations can hide the use of multiple physical interfaces from the IP stack [17]. The *logical interface* [5,6] at the IP layer is the most complete approach, as it allows both sequential and simultaneous use of different physical media.

For this kind of terminal, our preferred solution is based on the LMA delegating the same prefix (or set of prefixes) to the MN, regardless of the physical interface that is getting attached to a MAG, since there is only one interface visible from the IP layer. In fact, this basically means that from the viewpoint of the network, the MN is sharing the same IP address(es) across multiple physical interfaces, although the addresses are not really configured on the physical interfaces but on the logical one. The LMA decides—on an IP flow basis—through which MAG data traffic is forwarded to the MN, and consequently through which physical interface the MN receives traffic.

2. *Terminals with multiple IP interfaces.* In case the mobile terminal does not implement the logical interface concept (or an alternative link-layer approach that hides the use of multiple media to the IP layer), it is still possible to enable full flow mobility if the terminal follows the *weak host* model [7,8]. This model does not limit the traffic reception at a host to only those IP packets whose destination address matches the IP address assigned to the interface receiving the packets, but allows the host to receive and process packets whose IP destination address corresponds to that of any of the local interfaces of the host. We have performed some tests with different operating systems, and the results show that both Linux (tested with Linux-2.6.26) and Mac OS X (tested with Leopard version) implement the weak host model for both IPv4 and IPv6 traffic. We have not performed tests with Windows, but some results have been reported in [20]. Windows XP and Windows Server 2003 use the weak host model for all IPv4 interfaces and the

⁷ The main problem here would not be the asymmetry in the paths followed by packets—IP routing does not guarantee symmetric routing—but the different access network delays imposed by different technologies, which could have an impact on the performance, e.g., of TCP flows.

strong host model for all IPv6 interfaces, not being possible to modify this behavior. The Next Generation TCP/IP stack in Windows Vista and Windows Server 2008 supports the strong host model for both IPv4 and IPv6 by default on all interfaces but in this case, the stack can be configured to use the weak host model.

For this kind of terminal, our solution is based on the LMA delegating a unique prefix (or set of prefixes) per interface (as in plain PMIPv6). The LMA performs flow-based routing while the MN is able to process received packets at any of its interfaces, thanks to the use of the weak host model.

In the next sections, we elaborate more on the specific protocol extensions that are required to enable flow mobility in a PMIPv6 domain for the two kinds of terminals supported by our solution.

3.2 PMIPv6 Extensions

3.2.1 Single IP Interface Case: Logical Interface Model

When an MN uses a logical interface to connect to the same LMD via multiple physical interfaces, it appears to the rest of the network as a set of different endpoints with the same Layer-2 and Layer-3 addresses. In PMIPv6, once an MN has attached one of its interfaces and has been registered in the LMA, subsequent attachments via different interfaces to different MAGs might be identified as handover requests. Our approach: (i) extends the original PMIPv6 to allow the MAG specify—upon attachment of a mobile node—that the attaching physical interface belongs to a logical interface, and (ii) modifies the conceptual data structure at the LMA, so it stores information about all the MAGs that lead to the same host (that is, the Proxy CoAs and the tunnel-IDs). One extra instance of these parameters should be added for each physical interface (grouped under the same logical interface), so that the LMA is able to create tunnels and routes without deleting the existing ones.

The above description (to simplify the explanation of the protocol procedures) takes into account the assignment of a single HNP per logical IP interface. In case the LMA assigns a pool of HNPs to the logical IP interface (from the LMA perspective this is a standard IP interface) all the logic still holds. The LMA will need to store all the HNPs for the specific mobility session. From a MAG point of view there may be different protocol choices:

- *Unique HNP (or set of HNPs) per physical interface.* In this case the LMA, upon attachment of each physical interface, assigns a different HNP (or set of HNPs). That is, the MAGs providing network connectivity to the MN know only the on-link prefix(es). To enable flow mobility, the LMA—during the PBU/PBA protocol exchange—should inform the MAGs about all the HNPs associated to the MN. The PBA should carry the HNPs that should be reachable via the on-link HNP. This procedure is similar to the one described in the weak host section allowing the MN to receive packets to any HNP (irrespective of the on-link configuration) as long as they are properly assigned to the logical IP interface. The PBA message contains a specific option and upon parsing, the MAG installs the required routing state.
- *Multiple shared HNPs per physical interface.* In this case the LMA behaves according to the original PMIPv6 specification [3] and assigns a pool of HNPs to the logical physical interface. The same prefixes will be assigned when the MN attaches a second physical interface.

The experimental results presented in Sect. 4 describes the single HNP per logical IP interface. We argue that, from a session continuity point of view, this is the most interesting

scenario: the node configures a single global, always-on reachable IP address from that HNP. Moreover, in a 3GPP context the HNP is the IP prefix assigned by the mobility anchor to the MN upon network attachment allowing seamless mobility of IP flows across heterogeneous access.⁸

3.2.2 Multiple IP Interfaces Case: Weak Host Model

With regular PMIPv6, when an MN attaches to an LMD via more than one interface, it receives a different prefix for each one of them. Each interface is treated as if it was a completely different MN (i.e., separated mobility sessions). Our solution solves this issue by enabling the LMA to group together all the mobility state that it has referring to the same MN in a new conceptual structure called *flow-mob list*.

The MAG, upon detecting MN attachment, checks whether the MN is authorized for PMIPv6 service. If so, the MAG prepares the PBU with the acquired MN-ID⁹ in the MN-ID option and the MAC address in the Link Layer ID (LL-ID) option. When the PBU is received, the LMA registers a new BCE following the PMIPv6 standard procedure (because the HNP and the LL-ID are new), and in addition it checks whether the MN-ID is already present in the flow-mob list. The LMA then builds a PBA with the prefix assigned to the new interface (standard PMIPv6 behavior), including a new extra option—which has the same format of the HNP prefix option—that carries the prefix(es) assigned to the previously attached interface(s). This allows the MAG to install routes to all the prefixes assigned to the MN for each of its interfaces attached to the same LMD.

It should be noted that the above behavior is similar to the one described for the logical IP interface when multiple HNPs are delegated to the MN.

4 Validation and Experimental Evaluation

4.1 Testbed Description

In order to be able to conduct real experiments that allow us to evaluate the feasibility and performance of our proposed solution, we implemented the basic PMIPv6 protocol as well as our flow mobility extensions. Figure 3 depicts the functional boxes in our testbed and the associated software modules. The network setup features one LMA, three MAGs, a machine acting as network server connected to the LMA and two mobile nodes: one implementing the weak host model (*weak host MN*) and one implementing a particular realization of the logical interface concept: the bonding interface (*bonding MN*). These nodes are Ubuntu 9.04 Linux machines (with Linux-2.6.31). PMIPv6 mobility support is enabled on the LMA and the MAGs. Two real access points (APs) are deployed to provide WLAN access, attached to MAG2 and MAG3 via an Ethernet cable. These APs are Linksys WRT54GL v1.1 routers (configured to operate in AP mode), running OpenWRT Kamikaze 7.09 distribution. 3G access is also provided (MAG1), via the 3G Alcatel Lucent in-house network.

⁸ It should be noted that the 3GPP SA2 working group will be standardizing for Rel-10 mechanisms for seamless WLAN offload from the LTE wireless access. Such technologies are currently based on DSMIP, but studies show the strong interest of mobile operators in the deployment of network-based solutions.

⁹ We use the MAC address as MN-ID because this is what it is supported by our current implementation. Nevertheless, a different approach, such as the use of Network Access Identifiers (NAIs) could be followed instead, and in this case a conversion mechanism would not be necessary.

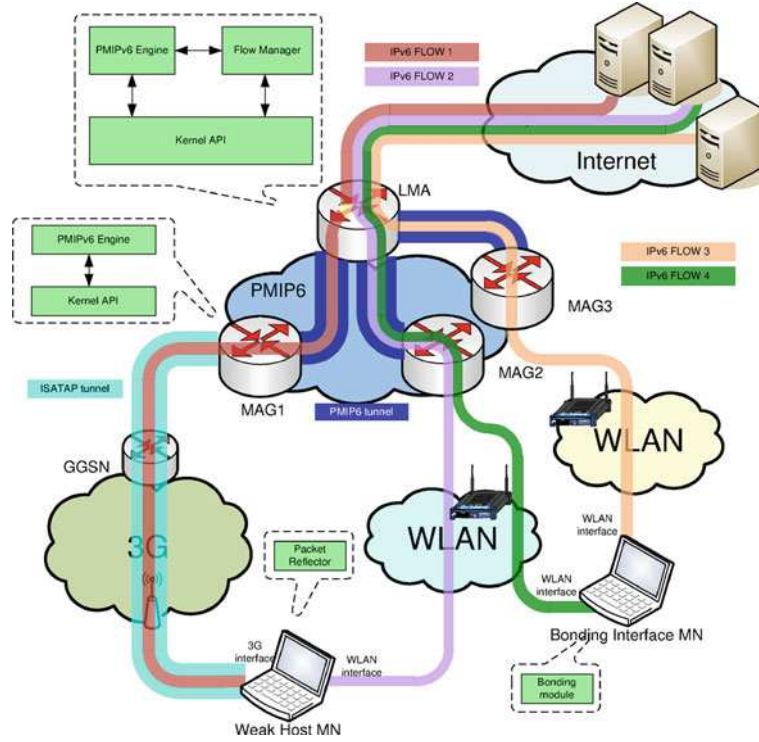


Fig. 3 Testbed setup

The weak host MN has one WLAN interface and one 3G interface (Novatel USB dongle). Since the 3G network only provides IPv4 connectivity, we setup an Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) [21] connection to convey IPv6 packets over the point-to-point IPv4 3G connection. That is, the in-house Gateway GPRS Support Node (GGSN) has been connected to MAG1 and upon ISATAP establishment, the Router Solicitation generated by the MN is conveyed to the MAG through the ISATAP tunnel. Upon Router Solicitation reception, the MAG triggers the PBU/PBA protocol exchange with the LMA. From a protocol behavior and flow management point of view the use of the ISATAP tunnel has no impact. When the weak host MN performs network attachment it receives two HNPs, one on each interface (e.g., 3G and WLAN) and the *packet reflector* module assures that uplink (UL) and downlink (DL) packets are sent through the same interface. This small module takes care of identifying IP flows, monitoring at which interface IP packets belonging to a particular flow arrives (downlink), and replicating that behavior in the uplink (i.e., using the same interface when sending packets belonging to this flow).

The bonding MN features the Linux bonding module modified to install specific transmitting policies. The bonding device is created “enslaving” two wireless network interfaces, each of them connected to the WLAN access point attached to MAG2 and MAG3. It should be noted that the access points feature special purpose software (code runs on top of the OpenWRT distribution) to perform network attachment/detachment detection of WLAN stations. That is, upon successful Layer-2 association, the AP sends to the MAG an *AttachmentTrigger* to bootstrap the PMIPv6 registration procedure. After the attachment of the two wireless

physical interfaces, the MN has an HNP configured on the bonding device and can receive packets on any of the two physical interfaces.

The MAGs implement the PMIPv6 engine to form PBUs, parse PBAs and install the required routing state for packet delivery. MAG2 and MAG3 as mentioned before, and in addition to Router Solicitation messages, are able to receive Layer-2 attachment triggers from the AP and start the PBU/PBA protocol exchange. There are no further required components to perform flow mobility.

The LMA plays a key role in the flow mobility procedure. It runs the PMIPv6 engine and the logic to classify/manage the IP flows.

4.2 Experimental Evaluation

This section provides an experimental analysis of the mechanisms designed to enable flow mobility in PMIPv6 domains. Different tests were performed to validate the feasibility of the proposed approach. We consider two main situations in our experimental evaluation:

1. QoS triggered flow mobility. The movement of a flow (or set of flows) from one interface to another is triggered by QoS reasons. For example, the access network to which an interface is attached might not be able to cope with all the traffic, so the operator decides to offload a flow (or set of flows) to an interface connected to a less congested access network. This type of mobility is typically proactive.
2. Interface outage triggered flow mobility. A completely different situation appears when all the flows bound to a given interface have to be moved because the interface has just gone down. This might happen because the user has just manually switched down an interface (e.g., to save some battery life or money) or because of radio coverage. This type of mobility is typically reactive.

As explained in Sect. 3, two different types of mobile nodes are supported by our solution, following different paradigms: the logical interface and the weak host model. Although from a conceptual viewpoint our solution should behave quite similarly with both approaches, due to the particular implementations that we use for the experiments, there are some limitations that have an impact on the type and number of the tests that can be performed:

- The logical interface based MN is implemented by using the Linux Bonding Driver. This driver¹⁰ is designed for physical Ethernet interfaces only. Although other Ethernet-based technologies, such as WLAN, are also supported, it is not possible to bond (i.e., group under the same logical interface) 3G interfaces, as a logical PPP interface is brought up when 3G is enabled¹¹ and the bonding module does not support non-physical interfaces.
- The weak host model does not allow the prefixes assigned to an interface to survive if the interface is shut down, as they are bound to the physical interface. Because of this limitation, we do not perform tests with the weak host MN in which an interface is completely turned down (this actually would correspond to a complete handover). Note that with some support from the terminal, this limitation might be overcome by not fully shutting down the interface, but just turning the radio off.

It is important to note that the main goal of this section is to experimentally validate the design of our solution, by conducting different experiments with a real implementation.

¹⁰ <http://www.linuxfoundation.org/collaborate/workgroups/networking/bonding/>

¹¹ The Point to Point Protocol (PPP) is used between the MN and the GGSN when the PDP context is setup. A PPP interface is configured on the MN and used as default one to reach the Internet.

4.2.1 QoS Triggered Flow Mobility Handovers

This section analyzes the behavior of the flow mobility procedures when the Flow Manager (located at the LMA) receives QoS related triggers. We first proceed to analyze the WLAN to WLAN scenario for the bonding MN and then compare the obtained results with the WLAN to WLAN scenario for the weak host MN. The goal is to show that there is no difference from a flow management point of view. We then proceed to analyze the more compelling WLAN to 3G flow mobility scenario. It should be noted that the latter scenario is the baseline for any optimization algorithm aiming at offloading the 3G network.

Flow mobility triggered by QoS changes for WLAN-WLAN scenario

These experiments are performed using an MN which operates through two identical WLAN interfaces. It is worth noticing, in order to understand the experiment, that the delay between the LMA and each interface of the MN is the same, without adding any artificial delay between both entities. As TCP is the predominant type of traffic in the Internet nowadays, we use TCP flows in the tests. During this experiment we simulate a degradation of the link used by the flow under inspection, triggering a handover due to an increase in the number of packet losses. In order to do so, we use the `tc` (traffic control) properties of the Linux kernel. By using the traffic shaping module (through the `tc qdisc` interface) we are able to decrease the capacity of the tunnel between the LMA and the MAG, leading to a handover once the packet loss reaches a given threshold.

Figure 4 presents the plot of TCP sequence number and throughput versus time for the scenario explained before and using a bonding MN. It can be observed how the sequence number graph presents six step regions, starting in 37, 89, 130, 224, 257 and 314 s. These step regions correspond to the packets losses due to the effect of the traffic shaping. Once the flow is moved appropriately, the TCP sequence number starts increasing again since in the new path no losses occur. The same effect can also be appreciated in the throughput. At the same time intervals when the sequence number graph reduces its slope, the average throughput depicted in the figure decreases, since packets are lost at the receiver, and data segments are retransmitted. A close-up of one of the step regions is also presented in Fig. 4 for better understanding. It shows that the step region is not continuously flat as packets are being dropped by the traffic shaper progressively. Note that the mechanisms used to emulate congestion and to detect packet losses are not perfect. Some packets need to be lost before detecting the congestion of a particular path, and then triggering the subsequent flow mobility handover. This has an impact on the performance experienced by the user, which could be reduced by deploying more intelligent network congestion mechanisms. In a real operator's network there are more complex tools available that could be used to help triggering flow mobility in a more effective way (i.e., shorter —close to zero—service disruption times). The main goal of these experiments is not to fully characterize experimentally the performance of the solution, but to validate its feasibility.

In order to compare the weak host model and bonding interface concepts regarding the flow mobility due to QoS constraints, we perform the same experiment using the weak host MN (results are shown in Fig. 5). Comparing Fig. 4 and Fig. 5, it can be concluded that there are no significant differences between the observed behavior, which supports the idea that the performance of our solution is not affected by the type of MN (weak host or bonding one).




Fig. 4 Bonding MN, QoS scenario, TCP sequence number and average throughput versus time




Fig. 5 Weak host MN, QoS scenario, TCP sequence number and average throughput versus time

Flow mobility triggered by QoS changes for WLAN-3G scenario

This experiment explores the inter-technology flow mobility due to QoS changes. The experiment setup is similar to the one previously depicted, but herein we focus on the relevant aspects of the handover between two different technologies. The experiment consists in the streaming—using TCP—of a video to an MN connected to two different MAGs through WLAN and 3G. As in the previous tests, the quality of the links between the LMA and MAG




Fig. 6 Weak host MN, WLAN-3G QoS scenario, TCP sequence number and average throughput versus time

is affected by the use of the traffic shaping characteristics of the Linux Kernel, through the `tc qdisc` command. Figure 6 presents the results obtained.

Figure 6 shows the sequence of the different handovers, triggered by the packet loss ratio crossing a configured threshold. Again, we should note that in a real operator's scenario, the network would be able to predictively trigger flow mobility handovers, without needing to wait for a reaction upon packet losses. The experiment starts with the MN attached to the 3G network, since this is the interface defined as default. A total of eight handovers are performed in this test, each one moving the flow from the congested access network to the one without QoS constraints. As in the WLAN to WLAN experiment, the sequence number graph does not remain completely flat during the retransmissions, since the interface is affected by losses, but it never goes completely down. The instants where a flow is moved from one interface to another can be easily identified due to the fact of the average throughput decreases during the handover (this would not be the case for handovers triggered predictively by the network). Once the handover is performed, we can see an abrupt increment in the sequence number graph caused by the TCP retransmissions.

Finally, from Fig. 5 and Fig. 6, we can conclude that the designed solution is feasible and works in a real environment. Therefore, our approach could be used by network operators to provide seamless inter-technology flow mobility, fulfilling operators desires while not impacting the final user's experience. Note that while in our experimental validation handovers have been triggered upon reaction to packet losses, in a real QoS-enabled mobile operator scenario, the network would be able to predict path congestion, and therefore react accordingly to solve this by issuing flow mobility.

4.2.2 Interface Outage Triggered Flow Mobility

This section describes the flow mobility procedures when the LMA receives Proxy Binding Update messages with a lifetime value set to zero (in terms of protocol operations it means




Fig. 7 Bonding MN, outage scenario, TCP sequence number and average throughput versus time

that an MN has disconnected from the sending MAG). Due to the limitations explained before, we first test the scenario for the WLAN to WLAN case using the bonding MN. We argue however that from a protocol operation point of view the same considerations apply to weak host terminals. We finally relate an out of coverage scenario to one in which a weak host MN performs a WLAN to 3G flow handover triggered manually. It should be noted that there is no impact on the protocol operation (only the trigger changes).

Flow mobility triggered by interface outage for WLAN-WLAN scenario

As in the previous experiment, herein an MN with two identical WLAN interfaces is considered and no artificial delay is added to any of the paths between the LMA and the MN. This experiment analyzes the flow mobility when triggered by an out of coverage scenario of the interface serving the flow. When the MN's currently active interface is switched off, the flow is automatically moved to the remaining active interface (thanks to the Layer-2 attachment/detachment code, which allows the MAG quickly detect the MN detachment). We then move back and forth the flow by alternating the active interface.

Figure 7 presents the TCP sequence number and Average throughput versus time graphs. As in the scenario presented in the previous experiment, four step regions can be identified in the sequence number versus time graph. These step regions start at 26, 51, 76 and 100s respectively. If we analyze the close-up of the figure, it can be seen how in this case the region is completely flat, in contrast with the results shown in the previous experiments (QoS triggered flow mobility handovers). There is no progressive loss of packets, since the interface is abruptly turned down. The different step regions are for all cases shorter than the ones presented in the previous experiments, because in this case the interruptions correspond to the time required by the network to detect and signal the interface disconnection, and then to re-route appropriately the affected flow.

It is worth noticing that we only perform this experiment for the bonding MN, for the reasons highlighted at the beginning of this section regarding the weak host MN. In the case




Fig. 8 Weak host MN, WLAN-3G handover scenario, TCP sequence number and average throughput versus time

of the bonding terminal, the IP prefix is delegated to the unique logical IP interface, instead to each individual IP (physical) interface as in the case of the weak host MN. This difference yields to a strange behavior of the weak host terminal when the interface is turned down, removing the IP prefix from the shutdown interface. Hence the outage experiment cannot not be conducted using the weak host model node.

Flow mobility triggered by interface outage for WLAN-3G scenario

This experiment considers an MN which has an IEEE 802.11a/b/g card as one of its interfaces, while the second interface is a standard 3G modem. Herein we focus on the evaluation of a handover case emulating an out of coverage scenario. The MN starts a TCP video flow in the 3G interface and this flow is manually switched to the WLAN and 3G back and forth. Figure 8 presents the results of this test. As shown in the figure, the bandwidth requirements of the video are quite low, hence the video does not suffer from congestion while being transmitted/received at any of the interfaces. We select this scenario since we want to assess the impact of changing the underlying technology to a standard traffic without QoS constraints. Observed results show that the handover between both technologies is almost transparent from the viewpoint of the flow performance. In the case of WLAN to 3G handover, we find that for each handover, some retransmissions occur, as the bandwidth of the 3G interface is lower than the WLAN one, and its delay is higher. This decrease in the performance would be hardly noticeable due to the low requirements of the traffic being used and the fact that the TCP pace is recovered quickly. For the case of the 3G to WLAN handover we find the inverse behavior, observing an increase in the speed of the sequence number growth. Observed results show that our design does not impose any penalty in the performance of the flow apart from the effect of changing the characteristics of the underlying technology, which is known to affect the TCP performance. Nevertheless, the flow handover itself is seamless and transparent for the involved communications peers.

5 Comparison with Previous Work

The concept of flow mobility has been extensively analyzed for client-based mobility protocols, and there already exist standardized solutions, such as the flow bindings extensions for Mobile IPv6 [22]. The use of this kind of client-based solution has been proposed as a mechanism to enable mobile operators to offload data from their 3G networks [23], and there even exist approaches based on the IP Multimedia Subsystem (IMS) framework [24]. We argue that client-based solutions have several disadvantages, since they require to modify the users' devices to include an IP mobility stack, which also has to be provisioned with proper configuration and security credentials (in addition to those required to access the operator's network). This additional requirements might limit the usability of a solution due to the difficulties involved in its deployment.

As PMIPv6 is the standardized solution for network-based mobility management, the 3GPP and the IETF are currently working on the design of PMIPv6 extensions to enable flow mobility. The NETEXT WG of the IETF has been recently re-chartered to work on extensions to enable inter-technology handovers and flow mobility. An early version of the solution described in this paper has been presented in the IETF, being one of the first ones addressing the flow mobility issue that was presented and discussed there (even before the NETEXT group was actually re-chartered to work on flow mobility) [25]. There are other solutions which tackle the same problem, although no standard solution exists yet. We next summarize some of the most relevant existing proposals and compare them with the solution we have presented and evaluated in this paper.

Koodli and Chowdhury [26] propose in new signaling between the LMA and the MAG to enable the LMA control flow mobility. Two messages are defined: the Flow Handover Request (FHRQ)—that is sent by the LMA to the MAG set up forwarding for one or more flows to an MN—and the Flow Handover Reply (FHRP)—sent by the MAG in reply to a FHRQ message. While this signalling can be used to bind particular flows of an MN to specific MAGs, authors do not include any considerations on the mobile node behavior/support, nor provide any validation result or report on experimental tests.

Hui et al. propose a similar approach in [27] and [28], consisting on a extension of the BCE format at the LMA so the same HNP can be bound to several MAGs. The Binding Update List Entry (BULE) data structure is also modified to include the service flow information at the MAG. As opposed to [26], the handover control is on the MN and not on the LMA, and therefore it can be considered as an approach less attractive for mobile operators.

As far as the authors know there is no published work about flow mobility extensions for PMIPv6 that include validation results based on real prototype experimentation. Wakikawa et al. present in [5] an approach based on the use of the virtual interface¹² to enable inter-technology handovers in PMIPv6. The approach is validated via implementation but it does not tackle the flow mobility issue. In [29], the same authors propose for the first time the use of the virtual interface to solve the problem of inter-technology handovers and multihoming in PMIPv6, but no details on the protocol changes (i.e. signalling between the LMA and MAG) required to support flow mobility are given.

6 Conclusions

In this paper we present an end-to-end system design featuring flow mobility extensions for the PMIPv6 protocol. Starting from ongoing discussions in the 3GPP and IETF stan-

¹² The term *virtual interface* refers to a particular implementation of the *logical interface* concept.

standardization fora, we derive the required design choices covering both network components and multi-mode mobile devices. Specifically, given the expensive nature (in terms of battery consumption) of simultaneous usage of heterogeneous wireless network interfaces, we first validate our design by measuring, through experiments, the power consumption of a device equipped with WLAN and 3G interfaces. The obtained results justify our choices and the proposed end-to-end design.

We then describe the solution emphasizing the implications of flow mobility support on hand-held devices. Two different configurations (single logical IP interface and multiple IP interfaces) have been presented and validated. The tests show that flow mobility in PMIPv6 based networks is feasible for TCP based data traffic. It is worth noticing that the testbed setup features a real 3G in-house network compounded by WLAN coverage, and that experiments have been conducted with commercially available tools (e.g., 3G USB dongle). The implementation work is documented in an annex witnessing the effort in combining standard PMIPv6 routing with enhanced procedures for flow management. The reader should be comfortable in reproducing a similar setup if required.

To the best of authors knowledge this is one of the first and most complete studies on flow mobility support for the PMIPv6 protocol. The paper combines an extensive implementation effort with an up to date review of current standardization activities. The next steps include promoting these ideas [30] at the NETEXT IETF working group while evolving the platform as the standard itself will evolve.

References

1. Soliman, E. H. (2009). Mobile IPv6 support for dual stack hosts and routers. Internet engineering task force, RFC 5555, June 2009.
2. Kempf, J. (2007). Problem statement for network-based localized mobility management (NETLMM). RFC 4830, April 2007.
3. Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., & Patil, B. (2008). Proxy Mobile IPv6. Internet engineering task force, RFC 5213, August 2008.
4. Bernardos, C. J., Gramaglia, M., Contreras, L., Calderon, M., & Soto, I. (2010). Network-based localized IP mobility management: Proxy Mobile IPv6 and current trends in standardization. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 1(2/3), 16–35.
5. Wakikawa, R., Kiriya, S., & Gundavelli, S. (2009). The applicability of virtual interface for inter-technology handoffs in Proxy Mobile IPv6. *Wireless Communications and Mobile Computing*, 11(4), 500–507.
6. Yokota, H., Gundavelli, S., Trung, T., Hong, Y., & Leung, K. (2010). Virtual interface support for IP hosts. Internet engineering task force, draft-yokota-netlmm-pmipv6-mn-itho-support-03.txt (work-in-progress), March 2010.
7. Braden, R. (1989). Requirements for internet hosts—communication layers. Internet Engineering Task Force, RFC 1122, October 1989.
8. Thaler, D. (2008). Evolution of the IP model. Internet engineering task force, draft-thaler-ip-model-evolution-01.txt (work-in-progress), July 2008.
9. Johnson, D., Perkins, C., & Arkko, J. (2004). Mobility support in IPv6. Internet engineering task force, RFC 3775, June 2004.
10. Narten, T., Nordmark, E., Simpson, W., & Soliman, H. (2007). Neighbor discovery for IP version 6 (IPv6). RFC 4861, September 2007.
11. LAN/MAN Committee of the IEEE Computer Society. (2008). IEEE Std 802.21-2008, Standards for local and metropolitan area - part 21: Media independent handover services.
12. Soto, I., Bernardos, C. J., Calderon, M., & Melia, T. (2010). PMIPv6: A network-based localized mobility management solution. *The Internet Protocol Journal*, 13(3), 2–15.
13. Kiriya, S., Wakikawa, R., Xia, J., & Teraoka, F. (2010). Context reflector for Proxy Mobile IPv6. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 1(2/3), 36–51.
14. Hui, S. Y., & Yeung, K. H. (2003). Challenges in the migration to 4G mobile systems. *IEEE Communications Magazine*, 41(12), 54–59.

15. Devarapalli, V., Kant, N., Lim, H., & Vogt, C. (2009). Multiple interface support with Proxy Mobile IPv6. Internet engineering task force, draft-devarapalli-netext-multi-interface-support-00.txt (work-in-progress), March 2009.
16. Thaler, D. (2007). Multi-link subnet issues. Internet engineering task force, RFC 4903, June 2007.
17. Gundavelli, Sri, Melia, Telemaco (2010). Logical interface support for multi-mode IP hosts. Internet engineering task force, draft-ietf-netext-logical-interface-support-01.txt (work-in-progress), October 2010.
18. Serrano, P., Bernardos, C.J., de la Oliva, A., Banchs, A., Soto, I., & Zink, M. (2010). FloorNet: Deployment and evaluation of a multihop wireless 802.11 testbed. *EURASIP Journal on Wireless Communications and Networking*, 10. Article ID 153102 (17 pp).
19. Rysavy, P. (2010). Mobile broadband capacity constraints and the need for optimization, February 2010.
20. Wasserman, M. (2010). Current practices for multiple interface hosts. Internet engineering task force, draft-ietf-mif-current-practices-03.txt (work-in-progress), August 2010.
21. Templin, F., Gleeson, T., & Thaler, D. (2008). Intra-site automatic tunnel addressing protocol (ISATAP). Internet engineering task force, RFC 5214, March 2008.
22. Soliman, H., Tsirtsis, G., Montavont, N., Giaretta, G., & Kuladinithi, K. (2010). Flow bindings in mobile IPv6 and NEMO basic support. Internet engineering task force, draft-ietf-mext-flow-binding-06.txt (work-in-progress), March 2010.
23. Qualcomm. (2010). 3G/Wi-Fi seamless offload (whitepaper), March 2010.
24. Imai, N., Isomura, M., & Idoue, A. (2009). Coordination path control method to reduce traffic load on NGN access gateway. In *13th International conference on intelligence in next generation networks, 2009 (ICIN 2009)*, October 2009.
25. Bernardos, C. J., Melia, T., Seite, P., & Korhonen, J. (2010). Multihoming extensions for Proxy Mobile IPv6. Internet engineering task force, draft-bernardos-mif-pmip-02.txt (work-in-progress), March 2010.
26. Koodli, R., & Chowdhury, K. (2009). Flow handover for Proxy Mobile IPv6. Internet engineering task force, draft-koodli-netext-flow-handover-01.txt (work-in-progress), October 2009.
27. Hui, M., & Deng, H. (2010). PMIPv6 multihoming extension and synchronization in LMA and MAG. Internet engineering task force, draft-hui-netext-multihoming-01.txt (work-in-progress), March 2010.
28. Hui, M., Chen, G., & Den, H. (2010). Service flow identifier in Proxy Mobile IPv6. Internet engineering task force, draft-hui-netext-service-flow-identifier-02.txt (work-in-progress), March 2010.
29. Wakikawa, R., Kiriya, S., & Gundavelli, S. (2008). The use of virtual interface for inter-technology handoffs and multihoming in Proxy Mobile IPv6. In *Proceedings of the international conference on mobile technology, applications, and systems*. NY, USA: ACM New York.
30. Bernardos, C. J. (2010). Proxy Mobile IPv6 extensions to support flow mobility. Internet engineering task force, draft-bernardos-netext-pmipv6-flowmob-01.txt (work-in-progress), October 2010.

Author Biographies



Telemaco Melia received his Informatics Engineering degree in 2002 from the Polytechnic of Turin, Italy, and his Ph.D. in Mobile Communications from the University of Goettingen in April 2007. From June 2002 to December 2007 he worked at NEC Europe Ltd. in Heidelberg, Germany, in the Mobile Internet Group. He worked on IPv6-based Mobile Communication focusing on IP mobility support across heterogeneous networks and resource optimization control. In September 2008 he joined Alcatel Lucent Bell Labs. He is currently working on interworking architectures spanning 3GPP, WiMAX forum, and IETF standardization bodies. His main research interests include wireless networking and next-generation networks. He is author of more than 20 publications and he actively contributes to the IETF.



Carlos J. Bernardos received a Telecommunication Engineering degree in 2003, and a Ph.D. in Telematics in 2006, both from the University Carlos III of Madrid (UC3M), where he worked as a research and teaching assistant from 2003 to 2008 and, since then, has worked as an Associate Professor. His current work focuses on vehicular networks and IP-based mobile communication protocols. He has published over 30 scientific papers in prestigious international journals and conferences, and he is also an active contributor to the Internet Engineering Task Force (IETF). He served as TPC chair of WEEDEV 2009 and as TPC co-chair of the Mobility track of NTMS 2011. He has also served as guest editor of IEEE Network.



Antonio de la Oliva received a Telecommunication Engineering degree in 2004, and a Ph.D. in Telematics in 2008, both from the University Carlos III of Madrid, where he worked as a research and teaching assistant from 2005 to 2008 and, since then, has worked as an Visiting Professor. His research is focused to IP mobility in general and specifically to Media Independent Handovers. He is an active contributor and voting member of the IEEE 802.21.



Fabio Giust received his Telecommunications Engineering degree in 2010 from University of Padova, Italy. In September 2010 he joined Institute IMDEA Networks, and he is currently attending a Ph.D. in Telematics Engineering at University Carlos III of Madrid. His research topic focuses on IPv6-based Mobile Communications and heterogeneous access, following the interest he matured while developing his Master thesis at Bell Labs.



Maria Calderon is an associate professor at the Telematics Engineering Department of University Carlos III of Madrid. She received a computer science engineering degree in 1991 and a Ph.D. degree in computer science in 1996, both from the Technical University of Madrid. She has published over 40 papers in the fields of advanced communications, reliable multicast protocols, programmable networks and IPv6 mobility. Her current work focuses on vehicular networks and IP-based mobile communication protocols.