

Analytic Evaluation and Experimental Validation of a Network-based IPv6 Distributed Mobility Management Solution

Fabio Giust, Carlos J. Bernardos and Antonio de la Oliva

Abstract—Mobile Internet traffic is growing steeply, mainly due to the deployment of new broadband wireless technologies and the ever increasing connectivity demand coming from new services being available to mobile users. Current mobile network architectures rely on centralized mobility protocols which intrinsically pose enormous burdens on the central anchors, both in terms of connectivity needs and user mobility management. In order to face these issues, a new paradigm, called Distributed Mobility Management, is being explored, based on flattening the network architecture by deploying multiple mobility anchors at the edge of the network.

In this article we conduct an analytic and experimental evaluation of a network-based IP distributed mobility management solution that leverages Proxy Mobile IPv6 protocol operations. We develop an analytic model of the signaling and packet delivery costs, as well as the handover latency of both Proxy Mobile IPv6 and our distributed solution. We have also implemented a Linux-based prototype of our proposal, which has been used to experimentally assess the handover latency in a real IEEE 802.11 scenario. Finally, we use the results obtained from the analytic and experimental performance to evaluate the benefits that could be achieved by deploying a distributed mobility management solution.

Index Terms—Distributed Mobility Management, IP mobility, PMIPv6, wireless systems, cellular architecture, handover mechanisms, experimental evaluation.



1 INTRODUCTION

WE are witnessing an exponential increase in the use of data services by mobile subscribers. This is motivated by a variety of different reasons, such as the affordability and wide availability of cellular and WLAN accesses (combined, coverage reaches almost 100% of dense populated areas in developed countries), the increasing popularity of High-Speed Downlink Packet Access (HSDPA) dongles for laptops and the surprising popularity of applications for smartphones and other portable devices that make use of Internet connectivity.

This huge amount of mobile data traffic is expected to grow even more with the future deployment of the Long Term Evolution (LTE) access networks, which will significantly increase the available bandwidth. In parallel, Internet services and contents are being placed closer to the users, aiming at enhancing the users' experience. All these facts are impacting the dimensioning and planning of operator networks, as current deployments are highly hierarchical and centralized, introducing serious scalability and reliability concerns. Different approaches are

currently being developed to address this aspect, such as the Local IP Access and Selected IP Traffic Offload (LIPA-SIPTO) within the 3GPP¹, and the Distributed Mobility Management (DMM) effort at the IETF². These solutions are intended to enable traffic distribution within an operator network, to alleviate congestion and to avoid traversing centralized gateways, while still providing mobility support to those applications that need it.

This article analytically evaluates a network-based IP distributed mobility management solution [1], paying special attention to the scalability and reliability problems caused by the use of current hierarchical and centralized mobility approaches. In particular, we study the signaling overhead, the data packet delivery cost and the handover latency, comparing the obtained results to those of a well known network-based centralized mobility protocol: Proxy Mobile IPv6. Moreover, we carry out the proof of concept of the design using an experimental setup with an implementation of the DMM solution. We also report on the validation and performance assessment results obtained from the evaluation of the handover latency on the implemented DMM platform.

The rest of the article is organized as follows: Section 2 briefly introduces the background and motivation for this work. Section 3 provides an overview of the proposed DMM solution, highlighting the key operations that are crucial elements for the analytic evaluation covered in Section 4. Section 5 describes the implementation

• F. Giust, C. J. Bernardos and A. de la Oliva are with the Department of Telematic Engineering, University Carlos III of Madrid, Spain.
E-mail: {fgiust, cjb, aoliva}@it.uc3m.es

The research leading to these results has received funding from the European Community's Seventh Framework Program FP7/2007-2013 under grant agreement 317941 – project iJOIN. The European Union and its agencies are not liable or otherwise responsible for the contents of this document; its content reflects the view of its authors only. The research of Antonio de la Oliva has also been partially funded from the Spanish Government, MICINN, under research grant TIN2010-20136-C03.

1. 3rd Generation Partnership Project: <http://www.3gpp.org/>
2. The Internet Engineering Task Force: <http://www.ietf.org/>

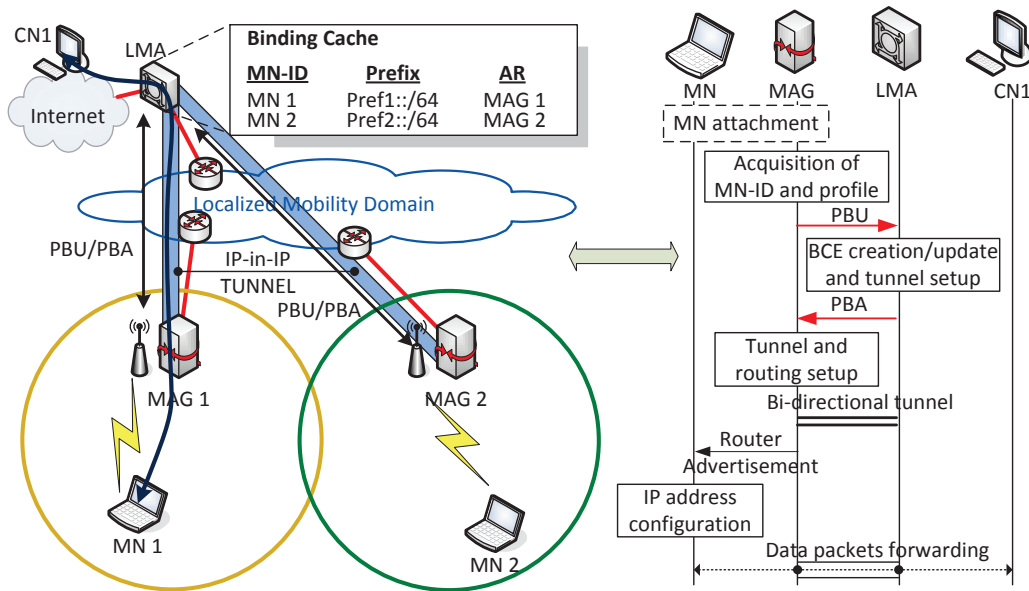


Fig. 1. Proxy Mobile IPv6 operations.

of the proposed solution, the platform deployed and the experiments conducted. Finally, Section 6 concludes the paper.

2 BACKGROUND AND MOTIVATION

Recent mobile architectures, such as the Evolved Packet System (EPS), follow an all-IP approach, triggering a real need to optimize IP mobility management protocols. Most of current solutions derive from Mobile IPv6 (MIPv6) [2], the first mobility protocol standardized by the IETF for IPv6 networks. MIPv6 enables global reachability and session continuity by introducing the Home Agent (HA), an entity located at the Home Network of the Mobile Node (MN) which anchors the permanent IP address used by the mobile node, called the Home Address (HoA). The home agent is in charge of defending the HoA's reachability when the mobile node is not at home (i.e., where the HoA is not topologically valid), and redirecting received traffic to the node's current location. When away from its home network, the MN acquires a temporal IP address from the visited network – called Care-of Address (CoA) – and informs the home agent about its current location by sending a Binding Update (BU) message. An IP bi-directional tunnel between the mobile node and the home agent is then used to redirect traffic from and to the mobile node.

While MIPv6 requires the explicit participation of the mobile node in the signaling procedures (this is referred to as client-based mobility), there is also a family of protocols that provide mobility support without the active involvement of the mobile node (the so-called network-based mobility). The effort towards network-based mobility management resulted in the standardized Proxy Mobile IPv6 (PMIPv6) [3], developed as an enhancement of MIPv6 and adopted by the 3GPP as one of

the mobility protocols within the EPS. With PMIPv6 (see Fig. 1), mobility support is offered in a portion of the network called localized mobility domain (LMD) or PMIPv6 domain. An LMD is formed by a core entity named the Local Mobility Anchor (LMA) and a set of Mobile Access Gateways (MAGs), deployed as part of the access network. PMIPv6 evolved from MIPv6 by relocating relevant functions for mobility management from the MN to the MAG. With PMIPv6, mobility is transparent for MNs: the network learns through standard terminal operation, such as router and neighbor discovery [4], about MN's changes of Point of Attachment (PoA) and coordinates routing state information using Proxy Binding Update (PBU) and Proxy Binding Acknowledgment (PBA) messages. The LMA plays the role of home agent in this architecture, being the topological anchor point for the IPv6 prefix(es) that is (are) uniquely assigned to MNs on a per user basis: the Home Network Prefix (HNP). Every time an MN roams within a PMIPv6 domain, the LMA updates the associated Binding Cache Entry (BCE), which contains – among other information – the Mobile Node Identifier (MN-ID), the HNP and the MN's location, called Proxy Care-of Address (P-CoA), which is the MAG's address where the MN is currently attached to. An IP bi-directional tunnel is maintained between the LMA and the MAG, so mobility is provided in a transparent way to the IP stack of the mobile node.

2.1 Limitations of centralized mobility management

Currently standardized IP mobility solutions come with the cost of handling operations at a central point – the mobility anchor – and burdening it with data forwarding and control mechanisms for a great amount of users. This central anchor point is in charge of tracking the location of the mobile node and redirecting traffic towards its current topological location. This brings several




Fig. 2. Network-based DMM solution: initial registration operations.

limitations [5]: *a)* sub-optimal routing, as traffic always traverses the central anchor, leading to paths that are, in general, longer than the direct one between the mobile node and its communication peer; *b)* scalability problems, as existing mobile networks have to be dimensioned to support all the traffic traversing the central anchors, and the anchor itself has to be powerful enough, and *c)* reliability, as the central entity is a potential single point of failure.

In order to address such issues, a new design paradigm, the so-called Distributed Mobility Management (DMM), is currently being discussed by both academy and industry [6], [7]. DMM basically develops the concept of a flatter system, in which the mobility anchors are placed closer to the user, distributing the control and data infrastructures among the entities located at the edge of the access network.

The approaches proposed so far towards distributing the mobility management functions can be divided into four different categories, namely: *i)* clean-slate approaches, proposing novel network architectures tackling the problems inherent to current IP mobility architectures from their foundation, *ii)* architecture-dependent solutions, such as the different efforts initiated in the 3GPP, *iii)* peer-to-peer approaches, distributing the mobility management functionality across several nodes in the network, and *iv)* solutions based on or extending existing IETF protocols. In this article we focus on the last category identified before, and in particular, on a partially distributed version of the IETF network-based IP mobility protocol, Proxy Mobile IPv6.

3 DESCRIPTION OF THE SOLUTION

The goal of this section is to provide enough background on how our network-based IPv6 DMM solution works, so the reader can better follow and understand the analytic and experimental evaluation performed later in

Sections 4 and 5 – which represent the main contributions of this work. For more details on the protocol operation, the reader is referred to [1] and [8].

Most of the terminology used when describing the solution is adopted from PMIPv6. Moreover, our DMM solution takes advantage of other aspects introduced by PMIPv6, as for instance the functions encompassed by the local mobility anchor and the mobile access gateway, but with the following substantial variations:

- The local mobility anchor is relieved from the data forwarding role, maintaining only the control function (i.e., mobility session management). Hence, the LMA is renamed as Central Mobility Database (CMD) to highlight its function of mobility sessions registry.
- The mobile access gateway is enriched with the LMA functionality, hence the name Mobility Anchor and Access Router (MAAR), as it acts both as connectivity provider and mobility manager. It maintains a local binding cache for the attached MNs.
- Each MAAR manages a unique set of global and locally anchored IPv6 prefixes, that are allocated by the MAAR to the MNs upon attachment. As each MAAR assigns prefixes from a separate non-overlapping pool, the definition of MN’s HNP is no longer adequate. We therefore refer in the rest of the paper to Local Network Prefix (LNP) as the prefix assigned by a MAAR to an MN attached to it³.
- The PMIPv6 signaling is extended to include some additional mobility options.

In terms of protocol operations, the MAARs leverage on the central mobility database to arrange the routing according to MNs mobility. The CMD maintains a global

3. By “local” we refer to the locally anchored nature of the prefix, not to the prefix scope, which is global to allow IPv6 communications with the whole Internet.




Fig. 3. Network-based DMM solution: handover management operation.

view of the network and keeps track of all the mobility sessions, being queried by the MAARs every time a mobile node is detected to attach/leave the mobility domain. Indeed, the MAARs are not aware of past information related to mobile nodes, so they need to contact the CMD to retrieve the required information to eventually take the appropriate actions towards providing session continuity to attached MNs.

We split the explanation of the protocol operations into two main parts: *i*) the initial registration (shown in Fig. 2), and *ii*) the handover management (Fig. 3). The left hand-side of each figure shows an example network topology, used to highlight where the MN is attached to, the main routing and mobility state maintained on each of the involved network entities (i.e., CMD and MAARs), the IP prefixes allocated and the tunnels set-up to ensure the connectivity. The right hand-side shows the corresponding message sequence charts.

3.1 Initial registration

When a mobile node (MN1) joins an access subnet, the responsible MAAR (MAAR1 in Fig. 2) retrieves the domain-wide unique MN-ID for MN1⁴. The MAAR pairs the MN-ID with the LNP (Pref1::5. This mapping is included in a local binding cache entry (BCE) created for the MN, that additionally contains other useful parameters as the session lifetime and the interface used to communicate with the MN. After the BCE has been created, a standard PBU message is sent to the CMD to create a new mobility

4. The actual procedure to generate the MN-ID can be related to the initial authentication and authorization security checks and it is out of scope in this work, but, for instance, it can be obtained from a dedicated AAA infrastructure using RADIUS.

5. For the sake of simplicity, we assume that the MN is always advertised a single prefix, arguing that extending this to multiple prefixes can be easily achieved.

session in its binding cache (a global BCE). Finally, the CMD generates back a standard PBA to the MAAR.

MAAR1, upon PBA reception, unicasts a Router Advertisement (RA) to the mobile node including the LNP allocated before. The mobile node configures an IPv6 address out of this prefix (Pref1::MN1/64), which is locally anchored at MAAR1 (i.e., the MAAR where the MN is currently attached). This MAAR is referred to as the Serving-MAAR (S-MAAR), because no encapsulation nor other special handling is performed by it to forward packets carrying the LNP, as long as the MN remains attached.

3.2 Handover management

This section describes the solution operation when a mobile node hands off from a point of attachment associated to MAAR1 to another one associated to a different MAAR (MAAR2 in the example shown in Fig. 3), which therefore becomes the new S-MAAR. Triggered by the new association, MAAR2 executes the procedure seen before: it allocates a new LNP (Pref2::

The S-MAAR receives the PBA message and advertises, in a RA message, the allocated LNP to the MN, which configures a new IP address ($\text{Pref2}::\text{MN1}/64$). The S-MAAR establishes an IP-in-IP tunnel towards the A-MAAR included in the “A-MAAR mobility option”, and sets up the routes required to handle traffic to and from $\text{Pref1}::/64$. Additionally, $\text{Pref1}::/64$ is deprecated, by advertising it with a 0-second preferred lifetime in the Router Advertisement sent to the MN. In this way, old IP flows using $\text{Pref1}::\text{MN1}/64$ are still active and anchored at MAAR1, but new flows use $\text{Pref2}::\text{MN1}/64$, which is the LNP routed by MAAR2 without any special packet handling both for uplink and downlink (see Fig. 3).

In parallel with previous operations, right after receiving the PBU from MAAR2 and updating the binding cache, the CMD notifies each of the active A-MAARs (in this example it is just MAAR1) of the new MN location through a PBU. This message is used by MAAR1 to update its local BCE for MN1 and to modify the routing and tunnel state accordingly. Indeed, MN1 is no longer reachable through a direct link, so MAAR1 establishes its tunnel endpoint with MAAR2 and sets the new required forwarding state to route packets containing $\text{Pref1}::/64$ through it (as shown on the left hand side of Fig. 3). When MAAR1’s operations are over, it sends back a PBA to the CMD to terminate the procedure. Once the PBU/PBA signaling exchanges are completed, the downlink path for $\text{Pref1}::/64$ is finally set up: packets are first received by MAAR1, encapsulated into the tunnel and forwarded to MAAR2, which de-encapsulates them and finally delivers them to their destination.

This handover procedure is repeated every time the MN changes S-MAAR. In the long run, a mobile node might have several IP flows running simultaneously, anchored at different MAARs. Flows using the LNP are routed normally (local breakout), whereas, those using a pLNP are encapsulated to the “home” (i.e., Anchor-) MAAR for that prefix, which takes care of routing the packet to the final intended destination.

The solution follows the concept known as *dynamic mobility*: only those IP flows that require session continuity will be tunneled if a handover takes place while the flows are alive. Applications that do not need mobility support from the IP layer (e.g., because the application can survive an IP address change) or that finish before a handover occurs (e.g., short lived dialogues, such as DNS queries) are tackled through standard IP routing.

4 ANALYTIC EVALUATION

In this section we analyze the costs in terms of signaling, packet delivery and handover latency of our proposed solution, and compare them with those of Proxy Mobile IPv6. This type of cost analysis has received a lot of attention in the recent past, starting during the development of cellular networks [9], [10], [11], [12], and then moving to the handover and mobility management in IP-based networks [13], [14].

TABLE 1
Notation.

$f_X(t), f_X^*(s)$	prob. density function of X , its Laplace transform
T_{SN}	subnet residence time
μ_{SN}	subnet (i.e., cell) crossing rate
T_{PR}	active prefix lifetime
\overline{T}_{PR}	mean value of T_{PR}
T_{PR}^H	active prefix lifetime while the MN is attached to the prefix home network (prefix used as LNP)
T_{PR}^F	active prefix lifetime while the MN is visiting a foreign network (the prefix is used as pLNP)
$1/\lambda_{PR}^F$	mean value of T_{PR}^F
\overline{N}_{PR}	average number of used active prefixes
\overline{N}_{pLNP}	average number of active anchored prefixes
$\alpha(K)$	probability that K handovers occur during the interval T_{PR}^F
τ	cost per packet for tunnel transmission
ω	cost per packet for wireless transmission
λ_p	packet transmission rate per active prefix
A_{SN}, L_{SN}	area and perimeter of a subnet
R_{SN}	radius of a circular subnet
\overline{v}	average speed of mobile node
T_{BCE}	BCE lifetime
R_{BCE}	rate of BCE refresh operations

In the following, we introduce the mobility and traffic model used in our analysis, to then derive the signaling, packet delivery and handover latency cost functions for both Proxy Mobile IPv6 and our DMM solution. The notation used throughout this section is summarized in Table 1.

4.1 User Mobility and Traffic Models

A location update takes place when the MN passes from one subnet to another, hence the first step is to characterize users’ mobility. A common practice is to model the node mobility as the time T_i spent in a cell i before moving to the next one (this is also referred to as *cell residence time*), and its related probability density function (PDF) f_{T_i} . In the following, we consider the worst case in which each cell is a different IP-subnet (SN). We assume that the wireless network is composed of statistically identical cells, so that T_i can be denoted by T_{SN} , with PDF: $f_{T_{SN}} = f_{T_i}, \forall i$. One simple method that has been extensively used in the past to model users mobility is known as *Fluid Flow model* [9], [15]. With such model⁶, the MN travels within a subnet with a direction uniformly distributed in $[0, 2\pi)$ and average speed \overline{v} . If the subnet has an area A_{SN} and a perimeter L_{SN} , then

6. Note that the analytic evaluation performed in this section is not limited to using the Fluid Flow Model, but it is also valid as long as statistically identical cells are assumed.

the subnet border crossing rate μ_{SN} is given by:

$$\mu_{SN} = \bar{v} \frac{A_{SN}}{\pi L_{SN}} = \frac{2\bar{v}}{\pi R_{SN}}, \quad (1)$$

where the second equation is obtained assuming circular subnets with radius R_{SN} .

Consequently, the mean subnet residence time is given by the inverse ratio of the border crossing rate:

$$E[T_{SN}] = 1/\mu_{SN}. \quad (2)$$

The MN configures a new address at each visited subnet. This behavior is opposite from Proxy Mobile IPv6, where typically the same IP address is used by the mobile node in all cells. However, in our DMM-based solution, the mobile node might be simultaneously using a subset of all the IP addresses configured. In the following, a prefix is considered “active” if the address derived from it is being used by at least one IP flow. Among the active prefixes, one is configured from the LNP allocated by the current Serving-MAAR, and the others from the pLNPs anchored at previously visited MAARs, i.e., the A-MAARs. The prefixes not involved in active IP flows are eventually de-registered and no longer influence the mobility updates. Therefore, a key aspect of our cost analysis is the lifetime of a prefix. Indeed, the number of active prefixes determines the number of involved A-MAARs, hence impacting on the signaling overhead, as well as on the traffic tunneled between the MAARs.

In the following we evaluate the average number of active prefixes at a handover event. According to our protocol operations, a prefix is always maintained at least for the first handover, and then it may expire depending on the user activity. Hence a prefix lifetime consists of a whole subnet residence time, because it is used as LNP, plus a trailing interval in which it is used as pLNP. We refer to the former as *home prefix lifetime* (T_{PR}^H), and to the latter as *foreign prefix lifetime* (T_{PR}^F). If T_{PR} denotes the time while a prefix is active (we name it *active prefix lifetime*), then we obtain:

$$T_{PR} = T_{PR}^H + T_{PR}^F, \quad (3)$$

where $T_{PR}^H = T_{SN}$, and T_{PR}^F is the random decay interval since the MN leaves the “home” MAAR until the prefix expires. By denoting $1/\lambda_{PR}^F \triangleq E[T_{PR}^F]$, we can write:

$$\bar{T}_{PR} = E[T_{PR}] = \frac{1}{\mu_{SN}} + \frac{1}{\lambda_{PR}^F}. \quad (4)$$

At a handover event, the average number of active prefixes, \bar{N}_{PR} , is given by one (the LNP) plus the average number of active pLNPs, \bar{N}_{pLNP} :

$$\bar{N}_{PR} = 1 + \bar{N}_{pLNP}. \quad (5)$$

\bar{N}_{pLNP} is the product of the average number of handovers that occur during the foreign prefix lifetime, times the prefix generation rate r , which in our case is simply $r = 1$ prefix per handover. We refer to $\alpha(K)$ as the

probability that K handovers occur during the interval T_{PR}^F . Therefore, a pLNP remains active on average for $E[\alpha(K)]$ subnet residence intervals and from the above considerations, we obtain:

$$\bar{N}_{pLNP} = r E[\alpha(K)] = E[\alpha(K)]. \quad (6)$$

We next follow the methodology devised in [12] to compute $\alpha(K)$ and $E[\alpha(K)]$ for any distribution of the subnet residence time and the foreign prefix lifetime. Let $f_{T_{SN}}^*(s)$ and $f_{T_{PR}^F}^*(s)$ be the Laplace transforms of $f_{T_{SN}}(t)$ and $f_{T_{PR}^F}(t)$ respectively. By applying Theorem 1 of [12], we obtain:

$$\alpha(K) = \begin{cases} - \sum_{p \in \sigma_{T_{PR}^F}} \text{Res}_{s=p} \frac{1 - f_{T_{SN}}^*(s)}{s} f_{T_{PR}^F}^*(-s), & \text{for } K = 0, \\ - \sum_{p \in \sigma_{T_{PR}^F}} \text{Res}_{s=p} \frac{[1 - f_{T_{SN}}^*(s)][f_{T_{SN}}^*(s)]^K}{s} f_{T_{PR}^F}^*(-s), & \text{for } K > 0, \end{cases} \quad (7)$$

where $\sigma_{T_{PR}^F}$ is the set of poles of $f_{T_{PR}^F}^*(-s)$ in the right half complex plane, and $\text{Res}_{s=p}$ is the residue at poles $s = p$.

In a similar way, from Theorem 2 of [12], we derive the expected value $E[\alpha(K)]$:

$$E[\alpha(K)] = - \sum_{p \in \sigma_{T_{PR}^F}} \text{Res}_{s=p} \frac{f_{T_{SN}}^*(s)}{s[1 - f_{T_{SN}}^*(s)]} f_{T_{PR}^F}^*(-s), \quad (8)$$

with the same meaning of $\sigma_{T_{PR}^F}$ and $\text{Res}_{s=p}$.

Eqs. (7) and (8) hold for any generic distribution $f_{T_{SN}}(t)$ and $f_{T_{PR}^F}(t)$, but they may yield to complicated results if the corresponding Laplace transforms are not easy to manipulate, and numeric methods and simulations may be required. For this reason, approximations are commonly introduced in order to carry out closed and tractable expressions. In this sense, the generalized Gamma distribution is often used because it has a simple Laplace transform, and more important, it does not have a specific shape, so it can fit an arbitrary distribution by choosing appropriate parameters [11]. For instance, let $f_{T_{SN}}(t)$ follow a Gamma distribution with mean $1/\mu_{SN}$ and variance $1/(\gamma\mu_{SN}^2)$, and $f_{T_{PR}^F}(t)$ follow an exponential distribution with mean $1/\lambda_{PR}^F$ (an Exponential is a special case of Gamma). We obtain:

$$f_{T_{SN}}(t) = \frac{\gamma\mu_{SN}^\gamma t^{\gamma-1}}{\Gamma(\gamma)} e^{-\gamma\mu_{SN}t}, \quad f_{T_{SN}}^*(s) = \left(\frac{\gamma\mu_{SN}}{s + \gamma\mu_{SN}} \right)^\gamma \quad (9a)$$

$$f_{T_{PR}^F}(t) = \lambda_{PR}^F e^{-\lambda_{PR}^F t}, \quad f_{T_{PR}^F}^*(s) = \frac{\lambda_{PR}^F}{s + \lambda_{PR}^F}, \quad (9b)$$

and therefore, applying these expressions to (5) and (6),

we get:

$$\begin{aligned} \bar{N}_{pLNP} &= E[\alpha(K)] = \frac{f_{T_{SN}}^*(\lambda_{PR}^F)}{1 - f_{T_{SN}}^*(\lambda_{PR}^F)} \\ &= \frac{(\gamma\mu_{SN})^\gamma}{(\lambda_{PR}^F + \gamma\mu_{SN})^\gamma - (\gamma\mu_{SN})^\gamma}; \end{aligned} \quad (10a)$$

$$\bar{N}_{PR} = \frac{(\lambda_{PR}^F + \gamma\mu_{SN})^\gamma}{(\lambda_{PR}^F + \gamma\mu_{SN})^\gamma - (\gamma\mu_{SN})^\gamma}. \quad (10b)$$

In order to further simplify the problem formulation, let's assume that also $f_{T_{SN}}(t)$ follows an exponential distribution with parameter μ_{SN} . In this scenario, the computation of \bar{N}_{PR} does not require to use the Laplace transform. In fact, the assignment of a new prefix is a Poisson process with parameter μ_{SN} . Let $N_{PR}(t)$ denote the number of active prefixes at time t . It follows a Poisson distribution with mean:

$$\bar{N}_{PR}(t) = \mu_{SN} \int_0^t [1 - F_{T_{PR}}(u)] du, \quad (11)$$

where $F_{T_{PR}}(u)$ is the cumulative distribution function (CDF) of T_{PR} . By letting $t \rightarrow \infty$ in Eq. (11), we obtain the long run behavior of $\bar{N}_{PR}(t)$, which represents the average number of active prefixes:

$$\begin{aligned} \bar{N}_{PR} &= \mu_{SN} \int_0^\infty [1 - F_{T_{PR}}(u)] du \\ &= \mu_{SN} \bar{T}_{PR} = 1 + \frac{\mu_{SN}}{\lambda_{PR}^F}. \end{aligned} \quad (12)$$

Similarly, the evaluation of $\alpha(K)$ is straightforward:

$$\begin{aligned} \alpha(0) &= \Pr[T_{PR}^F \leq T_{SN}] \\ &= \int_0^\infty \Pr[T_{PR}^F \leq u] f_{T_{SN}}(u) du \\ &= \frac{\lambda_{PR}^F}{\lambda_{PR}^F + \mu_{SN}}; \end{aligned} \quad (13a)$$

$$\alpha(K > 0) = 1 - \alpha(0) = \frac{\mu_{SN}}{\lambda_{PR}^F + \mu_{SN}} \triangleq P_{PR}. \quad (13b)$$

Iterating this computation and after simplifying the expression, we obtain:

$$\alpha(K \geq k) = P_{PR}^k, \quad k \geq 1, \quad (14)$$

which finally yields to the probability $\alpha(K)$ and its expected value:

$$\alpha(K) = P_{PR}^K - P_{PR}^{K+1} = P_{PR}^K(1 - P_{PR}), \quad (15a)$$

$$E[\alpha(K)] = \frac{P_{PR}}{1 - P_{PR}} = \frac{\mu_{SN}}{\lambda_{PR}^F}. \quad (15b)$$

By merging (5), (6) and (15b), we finally have:

$$\bar{N}_{PR} = 1 + \frac{\mu_{SN}}{\lambda_{PR}^F}. \quad (16)$$

We now anticipate a definition that it is used later in the paper. Let D_i denote the number of MAAR-MAAR hops between the S-MAAR and the A-MAAR anchoring

prefix i . If K handovers take place during the foreign prefix lifetime (T_{PR}^F), then:

$$D_i = K + 1. \quad (17)$$

Intuitively, this is explained by the fact that each prefix is maintained active at least for the time that the MN stays in the subnet.

4.2 Total signaling cost

A key operation for an IP mobility protocol is maintaining the MN's mobility session up to date. As we have described in Sections 2 and 3, such operation requires dedicated signaling, thus an important performance metric is the cost associated to it.

In the following, we refer to the total signaling cost, C_{sig} , as the sum of three main components: *i*) the cost for the binding update after a handover (C_{update}), *ii*) the cost for terminating a prefix that is no longer active (C_{de-reg}), and *iii*) the cost required to periodically refresh the bindings ($C_{refresh}$) [13], [14]. With this approach, we consider the steady-state or long run behavior, that is, we omit the initial registration phase and the de-registration when the MN definitely leaves the domain.

These operations are performed for each new visited access network. Since it takes place at a rate μ_{SN} , we have:

$$C_{sig} = \mu_{SN}(C_{update} + C_{refresh} + C_{de-reg}). \quad (18)$$

Table 2 lists the expressions of the cost components for both Proxy Mobile IPv6 and our DMM solution. The notation used is as follows: C_{X-Y} represents the symmetric cost to transfer a packet⁷ between network nodes X and Y; PC_X is the cost of processing a received packet by node X, and \bar{N}_{PR} is defined in Eq. (5). We assume a network topology in which $C_{LMA-MAG}$ is the same for all MAGs (i.e., the differences in cost are negligible) and similarly is $C_{CMD-MAAR}$ for all MAARs. Additionally, and in order to be fair in any subsequent cost comparison, we consider $C_{LMA-MAG} = C_{CMD-MAAR}$. Moreover, $C_{MAAR-MAAR}$ is the cost for any pair of adjacent MAARs, and, as mentioned in Section 4.1, we assume that the MN moves away from the home MAAR increasing the cost in a linear way. This assumption has no impact for the PMIPv6 case, whereas it is the worst-case for DMM.

4.2.1 PMIPv6

Every time a mobile node changes subnet, its location has to be updated, incurring in a cost C_{update}^{PMIP} . At each MN movement, there is a PBU/PBA exchange between the mobile access gateway (MAG) and the local mobility anchor (LMA), hence 2 messages on the LMA-MAG path, plus the corresponding processing costs.

Even if a node does not move, there is some periodic signaling in place to refresh the mobile node binding at

7. In order to simplify the analysis, we are not considering different costs for different packet sizes. We also recall that a signaling message is accounted as one packet.

TABLE 2
Binding signaling costs.

C_{update}^{PMIP}	$= 2C_{LMA-MAG} + PC_{LMA} + PC_{MAG}$
$C_{refresh}^{PMIP}$	$= R_{BCE} (2C_{LMA-MAG} + PC_{LMA} + PC_{MAG})$
C_{de-reg}^{PMIP}	$= 2C_{LMA-MAG} + PC_{LMA} + PC_{MAG}$
C_{update}^{DMM}	$= (\bar{N}_{PR} + 1)(2C_{CMD-MAAR} + PC_{CMD} + PC_{MAAR})$
$C_{refresh}^{DMM}$	$= R_{BCE} (2C_{CMD-MAAR} + PC_{CMD} + PC_{MAAR})$
C_{de-reg}^{DMM}	$= 4C_{CMD-MAAR} + 2PC_{CMD} + 2PC_{MAAR}$

the LMA. Each binding is stored as part of a mobility session with an associated lifetime. Before this lifetime expires, the MAG can check if the mobile node is still connected (via neighbor discovery) and proceed to refresh or de-register the session depending on the outcome. The costs of this procedure are denoted by $C_{refresh}^{PMIP}$ which is performed on average $R_{BCE} = \lfloor 1/(\mu_{SN}T_{BCE}) \rfloor$ times for each subnet sojourn, being T_{BCE} the binding cache entry lifetime.

Additionally, the old MAG performs a de-registration (i.e., a zero lifetime PBU/PBA exchange with the LMA) when it detects that the MN is no longer connected (C_{de-reg}^{PMIP}). If layer-2 triggers are available, the de-registration phase is synchronized with the MN's movement outside the subnet, whereas without such mechanisms, the de-registration happens when the session at the old MAG expires (depending on the mobility session lifetime, it may happen that the de-registration is triggered some time after the MN had left the subnet).

4.2.2 DMM

In our DMM solution, C_{update}^{DMM} consists of a PBU/PBA registration handshake between the new S-MAAR and the CMD to register the new assigned LNP, plus the PBU/PBA messages exchanged between the CMD and each active A-MAAR to maintain updated the information about active pLNPs. Since an MN has on average \bar{N}_{PR} active prefixes prior to handover, this is also the number of A-MAARs that are updated after the handover.

Analogously to PMIPv6, the S-MAAR refreshes through a PBU/PBA exchange the binding at the CMD. In this case it is not necessary that the CMD forwards the refresh message to the A-MAARs, so in total there is one handshake. For each permanence time in a subnet, there are on average R_{BCE} refreshments required, as we use T_{BCE} also as the DMM binding refresh lifetime.

An A-MAAR eventually de-registers the pLNP upon detection that the prefix is no longer in use, incurring in a cost C_{de-reg}^{DMM} , which is twice the PMIPv6 one, as a message exchange is performed between the A-MAAR and the CMD, and another with the CMD and the S-MAAR.

4.3 Packet Delivery Cost

One of the key drivers of the development of flatter network architectures exhibiting distributed mobility management is the reduction of the costs caused by all the traffic traversing a centralized anchor. The use of a DMM solution allows reducing the communication delay between endpoints bypassing the core network, and reducing the costs of using very powerful network nodes (e.g., the packet data network gateway, P-GW, in a 3GPP network) and links that need to be dimensioned to transport all mobile nodes' data to the core. Indeed, with PMIPv6, user data always traverses the centralized local mobility anchor, unless the two communication endpoints reside in the same domain and the local routing feature [16] is enabled.

A DMM mechanism should mitigate the problems of mobile operators when coping with the foreseen increase in users' traffic (e.g., when next generation access technologies such as LTE are deployed). Together with this traffic demand increase, operators also expect that most of the user data sessions will be terminated in the same region where the originating peer is located [17], e.g., voice calls, which usually are established between geographically close users, instant messaging services or the penetration of content distribution networks (CDNs) which bring content closer to the users.

In the following, we develop a framework to evaluate to which extent a flat architecture brings advantages over a centralized one in terms of packet delivery cost, and under which mobility scenarios. In our analysis, we consider that the transfer of a packet without encapsulation has a unitary cost, whilst the delivery through a tunnel incurs in a penalty $\tau > 1$; similarly, a wireless link has an extra cost $\omega > 1$. We consider that packets are transferred to and from the MN at rate $\lambda = \lambda_p \cdot \bar{N}_{PR}$ packets/second, where λ_p is the packet rate per active prefix. Note that for the case of PMIPv6, only one prefix is active, but the total packet rate λ would be the same.

With PMIPv6, the path followed by user data traffic between a correspondent node (CN) and a mobile node can be divided in three different segments: *i*) between the CN and the LMA, *ii*) tunneled between the LMA and the MAG where the mobile node is attached to, and *iii*) between the MAG and the MN, via a wireless link. Therefore, the PMIPv6 packet delivery cost is given by:

$$C_{PD}^{PMIP} = \lambda(C_{CN-LMA} + \tau C_{LMA-MAG} + \omega C_{MAG-MN}). \quad (19)$$

With DMM, the path followed by user data traffic depends on the prefix used by the mobile node. This path can also be divided in different segments: *i*) between the CN and the MAAR anchoring the prefix, *ii*) if the mobile node is not directly connected to the Anchor-MAAR, traffic is encapsulated between the A-MAAR and the S-MAAR, and *iii*) between the Serving-MAAR and the mobile node, via a wireless link. Note that if the mobile node is connected to the MAAR anchoring the used prefix, only segments *i*) and *iii*) are traversed.

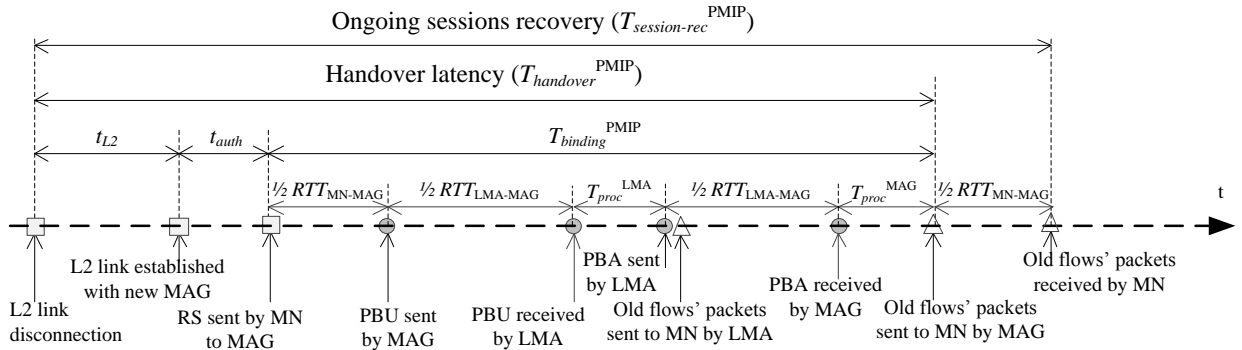


Fig. 4. Handover latency timeline in PMIPv6.

Hence packets are never encapsulated in the segments *i*) and *iii*), leading to the primary cost C_1^{DMM} , while some of them incur in the additional tunnel cost C_2^{DMM} . The packet delivery cost is thus obtained by the sum of the two components:

$$C_{PD}^{\text{DMM}} = C_1^{\text{DMM}} + C_2^{\text{DMM}}. \quad (20)$$

For the former, we simply have:

$$C_1^{\text{DMM}} = \lambda(C_{CN-\text{MAAR}} + \omega C_{\text{MAAR}-\text{MN}}). \quad (21)$$

For the latter, C_2^{DMM} , we need to consider that the tunnel cost is proportional to $D_i \tau C_{\text{MAAR}-\text{MAAR}}$, being D_i introduced in Eq. (17). Hence, in general, the additional tunnel cost for a prefix *i* not anchored at the current MAAR is given by the sum $\forall D_i > 0$ of the cost to transfer the packet for D_i MAAR-MAAR hops ($D_i \tau C_{\text{MAAR}-\text{MAAR}}$), weighted by the probability of the prefix to experience K handovers. Using Eqs. (5), (6) and (17), we have:

$$\begin{aligned} C_{tunnel}^{\text{DMM}} &= \sum_{K=0}^{\infty} D_i \alpha(K) \tau C_{\text{MAAR}-\text{MAAR}} \\ &= \sum_{K=0}^{\infty} (1+K) \alpha(K) \tau C_{\text{MAAR}-\text{MAAR}} \\ &= (1 + E[\alpha(K)]) \tau C_{\text{MAAR}-\text{MAAR}} \\ &= \bar{N}_{PR} \tau C_{\text{MAAR}-\text{MAAR}}. \end{aligned} \quad (22)$$

Finally, considering that tunneled packets belong to $\bar{N}_{PR} - 1$ prefixes, we obtain:

$$C_2^{\text{DMM}} = \lambda_p (\bar{N}_{PR} - 1) C_{tunnel}^{\text{DMM}}. \quad (23)$$

4.4 Handoff latency and Packet loss

We define the *handover latency*, $T_{handover}$, as the time interval in which an MN does not have global IP connectivity as a result of a handover. Besides restoring the IP link, the mobility operations ensure that the ongoing sessions are not disrupted and lost, even if, in general, packet losses may occur during the handover. For the sake of simplicity, we consider the session in the downstream direction, thus we refer to the *ongoing sessions recovery time*, $T_{session-rec}$, as the interval from the last data packet

of a given session received by the terminal before the handover to the first packet received after the handover. Therefore, and since packet buffering is not considered, the packet loss during the MN's movement, PL , is proportional to the length of such interval and to the compound packet transmission rate, λ , of the IP flows:

$$PL = \lambda T_{session-rec}. \quad (24)$$

We next study separately PMIPv6 and DMM.

4.4.1 PMIPv6

According to the PMIPv6 protocol, the handover operations can be divided into three phases, each of them requiring some time to be executed (see Fig. 4): *i*) the new layer-2 link establishment (L2 switch), leading to the time t_{L2} , *ii*) the authentication and authorization security checks, associated to t_{auth} , and, *iii*) the mobility binding phase, necessary to configure the correct routing between the LMA and new MAG, taking the time $t_{binding}^{\text{PMIPv6}}$. These stages take place sequentially, thus the handover latency is obtained by the sum of the three:

$$T_{handover}^{\text{PMIPv6}} = t_{L2} + t_{auth} + t_{binding}^{\text{PMIPv6}}. \quad (25)$$

However, the operations leading to t_{L2} and t_{auth} do not actually depend on the layer-3 mobility protocol. Indeed, the former element is related to the wireless technology deployed, whereas the latter is bound to the security mechanisms used, which are not necessarily tightly coupled with the mobility ones. Moreover, we can safely assume that such components would be identical if another mobility protocol were used, including DMM solutions, hence they are omitted in the subsequent analysis.

Nevertheless, after the MN has successfully completed the first two operations, it obtains again the local IP connectivity, but it is not able to send nor receive packets to and from the Internet. Indeed, the MN's address has not changed, nor the MN's default gateway, but from the network's perspective, the routing is not yet aligned with the actual MN's location. Thus, the binding phase is necessary to accomplish such requirement and at the same time to recover ongoing IP flows. In PMIPv6, $t_{binding}^{\text{PMIPv6}}$

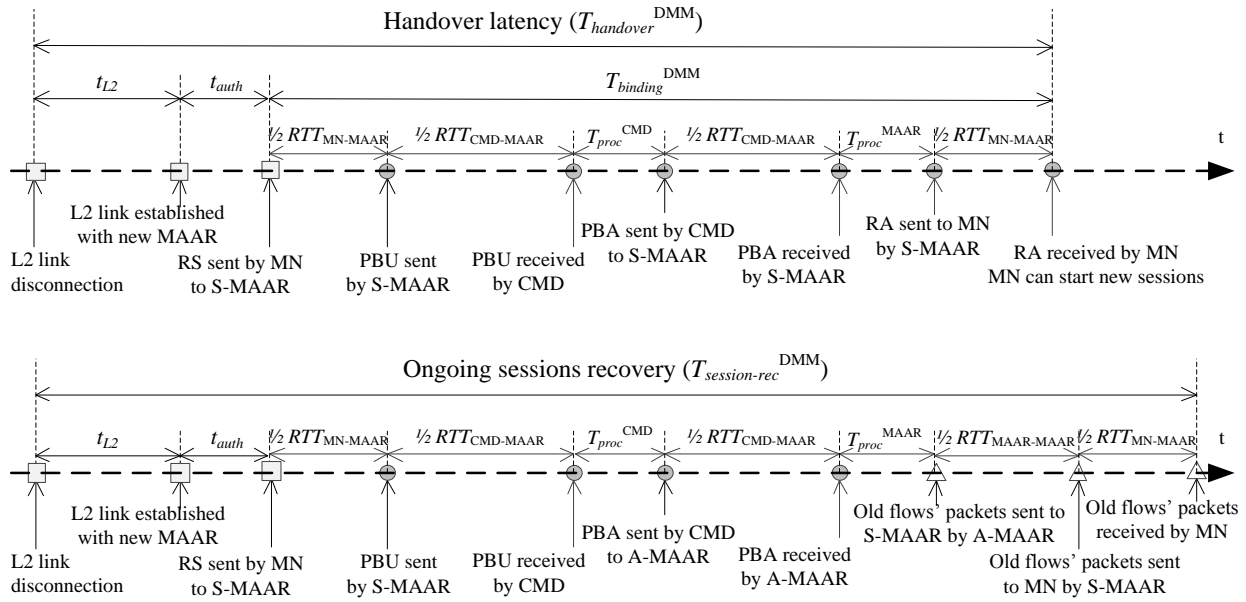


Fig. 5. Handover latency timeline in DMM.

is given by:

$$t_{binding}^{PMIP} = \frac{1}{2}RTT_{MN-MAG} + RTT_{LMA-MAG} + T_{Proc}^{MAG} + T_{Proc}^{LMA}, \quad (26)$$

because it is the time required to send the router solicitation (RS) message⁸ from the MN to the MAG, which can be approximated as half RTT between the MN and the MAG, plus the PBU/PBA message exchange, which accounts for one RTT between the MAG and the LMA, $RTT_{LMA-MAG}$, plus the processing time of the two nodes, T_{Proc}^{MAG} and T_{Proc}^{LMA} . We assume that $RTT_{LMA-MAG}$ is the same for all MAGs.

From the above considerations we can conclude that, in PMIPv6, the ongoing sessions recovery time coincides with the handover latency plus the time for the data packet to travel from the MAG to the MN, that is $1/2 RTT_{MN-MAG}$:

$$T_{session-rec}^{PMIP} = T_{handover}^{PMIP} + \frac{1}{2}RTT_{MN-MAG}. \quad (27)$$

Thus, combining Eqs. (24) and (27), the associated packet loss is:

$$PL^{PMIP} = \lambda(t_{L2} + t_{auth} + t_{binding}^{PMIP} + \frac{1}{2}RTT_{MN-MAG}), \quad (28)$$

where $\lambda = \lambda_p \bar{N}_{PR}$ is the packet transmission rate for the MN.

4.4.2 DMM

In DMM, the MN obtains global IP connectivity after a handover through the address configured from the new local network prefix, while the session recovery is tied

8. In this description we do not consider any layer-2 based attachment detection techniques hence the RS message is mandatory.

to old addresses. So, the handover latency in DMM can be expressed similarly to PMIPv6:

$$T_{handover}^{DMM} = t_{L2} + t_{auth} + t_{binding}^{DMM}, \quad (29)$$

but in this case the term $t_{binding}^{DMM}$ accounts also for the time required to send the RA message, because it conveys the prefix for the new address (see the upper part of Fig. 5). The expression for $t_{binding}^{DMM}$ is thus:

$$t_{binding}^{DMM} = RTT_{MN-MAAR} + RTT_{CMD-MAAR} + T_{Proc}^{MAAR} + T_{Proc}^{CMD}. \quad (30)$$

Again, we assume the CMD is at the same distance from all the MAARs.

However, as described in Section 3, configuring the new address guarantees the possibility for the MN to immediately start new communications, but ongoing sessions are recovered when the corresponding A-MAARs are updated by the CMD with the PBU/PBA exchange. Then old packets are received by the MN after traversing the path between the A-MAAR and the S-MAAR (see the lower diagram of Fig. 5). Hence every session is affected by a different interruption time in relation to the prefix they are using.

In the following we assume for simplicity that the S-MAAR and A-MAARs are updated simultaneously, and the extra time required for old packets to flow to the MN is given by the average distance between the S-MAAR and the A-MAARs plus half the RTT between the MN and the S-MAAR. With respect to the expressions in Eqs. (29) and (30) we thus obtain:

$$T_{session-rec}^{DMM} = T_{handover}^{DMM} + \frac{1}{2}\bar{N}_{PR}RTT_{MAAR-MAAR}, \quad (31)$$

and therefore, substituting in Eq. (24):

$$PL^{DMM} = \lambda(t_{L2} + t_{auth} + t_{binding}^{DMM} + \frac{1}{2}\bar{N}_{PR}RTT_{MAAR-MAAR}). \quad (32)$$




Fig. 6. DMM vs. PMIPv6: Signaling cost.

4.5 Performance analysis

This section builds on top of the results presented above and aims at providing insights about the advantages and disadvantages of deploying a network-based DMM solution compared to using PMIPv6.

We start by studying the signaling operations as a result of applying Eq. (18) with the PMIPv6 and DMM cost expressions presented in Table 2. For a fair comparison, we assume that $C_{LMA-MAG} = C_{CMD-MAAR}$, and also that $PC_{LMA} = PC_{CMD}$ and $PC_{MAG} = PC_{MAAR}$. Thus we have:

$$\frac{C_{sig}^{DMM}}{C_{sig}^{PMIPv6}} = 1 + \frac{1 + \bar{N}_{PR}}{2 + \lfloor 1/\mu_{SN} T_{BCE} \rfloor}. \quad (33)$$

Fig. 6 plots the above equation for $T_{BCE} = 300s^9$, and \bar{N}_{PR} computed from Eq. (11), that is, assuming a Gamma distribution for the subnet residence time.

Fig. 6 illustrates the effects of the subnet residence time, $1/\mu_{SN}$, and the prefix lifetime, $1/\lambda_{PR}^F$, on the signaling cost for various values of the parameter γ (note that for $\gamma = 1$, the subnet residence time becomes an exponential distribution with parameter μ_{SN}). The leftmost part of the graph (values below 100 seconds) indicates scenarios where the residence time of the user in the subnet is very short. As expected, the signaling overhead of DMM is higher than PMIPv6 for these scenarios (since several MAARs must be updated), especially when the prefix lifetime is long, while it converges to the signaling cost of PMIPv6 for the scenarios with comparable prefix life and residence times.

Regarding the packet delivery cost, we aim at understanding the trade-offs of the DMM deployment compared with PMIPv6, hence we focus only on the delivery cost of the packets within the mobility domain, leaving apart the cost related to reaching the domain¹⁰ and the

9. PMIPv6 Configuration Guide, Cisco IOS XE Release 3S. http://www.cisco.com/en/US/docs/ios-xml/ios/mob_pmip6/configuration/xe-3s/imo-pmip6-mag-support-xe.html

10. Note that this represents a worst-case scenario as we are neglecting the gains of our DMM solution when the CN is closer to the A-MAAR than to the LMA for the PMIPv6 case.




Fig. 7. DMM vs. PMIPv6: Packet delivery cost.

packet delivery on the last wireless hop. Following this reasoning, Fig. 7 depicts the ratio between Eq. (20) and (19) obtained setting $C_{CN-LMA} = C_{CN-MAAR} = 0$ and also $C_{MAG-MN} = C_{MAAR-MN} = 0$:

$$\frac{C_{PD}^{DMM}}{C_{PD}^{PMIPv6}} = (\bar{N}_{PR} - 1) \frac{C_{MAAR-MAAR}}{C_{LMA-MAG}}. \quad (34)$$

In order to show the different deployment choices that an operator may face when designing its DMM architecture, in Fig. 7 we present several curves according to the ratio $\frac{C_{MAAR-MAAR}}{C_{LMA-MAG}}$. This ratio expresses the distance in terms of delay between each entity. In our case, we have studied the following cases; *i*) the distance between MAARs is similar to the distance between LMA and MAGs ($\frac{C_{MAAR-MAAR}}{C_{LMA-MAG}} = 1$), *ii*) the distance between MAARs is half of the distance between LMA and MAGs ($\frac{C_{MAAR-MAAR}}{C_{LMA-MAG}} = 0.5$), *iii*) the difference in distances corresponds to a local vs regional delay case ($\frac{C_{MAAR-MAAR}}{C_{LMA-MAG}} = 0.2$, e.g., 10ms for local and 50ms for regional delays respectively), and *iv*) the distance between the LMA and the MAG is 10 times higher than the MAAR to MAAR one ($\frac{C_{MAAR-MAAR}}{C_{LMA-MAG}} = 0.1$). We can conclude from Fig. 7 that the benefits obtained from the deployment of DMM highly depend on the ratio $\frac{C_{MAAR-MAAR}}{C_{LMA-MAG}}$. In the case of large operator networks, where the distance between the LMA and the MAGs might be considerably large, the deployment of a DMM solution (provided that there is a shorter path between MAARs available) can reduce significantly the cost of providing mobility to the users, hence reducing the OPEX of the network.

Finally, we study the packet loss for PMIPv6 and DMM in a common scenario dimensioned as follows:

$$\begin{aligned} RTT_{MN-MAG} + T_{Proc}^{MAG} + T_{Proc}^{LMA} &= \\ &= RTT_{MN-MAAR} + T_{Proc}^{MAAR} + T_{Proc}^{CMD} \end{aligned} \quad (35a)$$

$$\begin{aligned} &\triangleq T_{Common} \\ RTT_{LMA-MAG} &= RTT_{CMD-MAAR}. \end{aligned} \quad (35b)$$




Fig. 8. DMM vs. PMIPv6: Packet loss.

From Eqs. (28) and (32) we obtain the ratio:

$$\frac{PL^{DMM}}{PL^{PMIP}} = \frac{T_{session-rec}^{DMM}}{T_{session-rec}^{PMIP}} = 1 + \frac{\overline{N}_{PR}}{2} \frac{RTT_{MAAR-MAAR}}{T_{session-rec}^{PMIP}} \quad (36)$$

Fig. 8 plots Eq. (36) for different values of the $RTT_{LMA-MAG}$. Also, we have supposed that $RTT_{MAAR-MAAR} = 10ms$, and $t_{L2} + t_{auth} + T_{Common} = 50ms$. Note that these values are aligned with the results obtained in the experimental evaluation, presented later in this article. In all cases, the packet loss probability of DMM is higher than the one of PMIPv6, due to the time required for packets to flow from the Anchor-MAAR to the Serving-MAAR through the tunnel.

Through the performed analysis we have clearly identified the trade-offs involved in the design of DMM solutions compared with PMIPv6 deployments. Our analysis indicates that the key points to consider for a DMM deployment are *i*) the size of the network in terms of delay between the LMA and MAG, *ii*) the connectivity between the MAARs and, *iii*) the ratio between subnet residence time versus the sessions lifetime. In summary, our conclusions show that DMM approaches are particularly suited for low to medium mobility scenarios, where the active session lifetime does not largely exceed the cell residence time. These results are very useful input when designing how to deploy the DMM solution, namely in which access entity the MAAR function should be co-located. To exemplify this analysis, let's consider a very demanding application, such as mobile video, where the average session duration is of 8.5 minutes¹¹, and a user roaming in a large operator network (with $RTT_{LMA-MAG}$ in the order of one hundred milliseconds). For this case, the DMM design outperforms PMIPv6, in terms of packet delivery cost, if the user performs less than one handover per minute. Considering current cellular

11. http://www.bytemobile.com/news-events/mobile_analytics_report.html

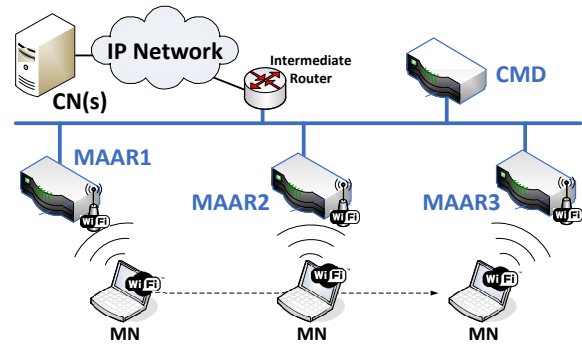


Fig. 9. Testbed deployed for the experiments.

and WiFi hotspot deployments, we argue that this case is quite common for pedestrian and low/medium mobility users, hence deploying a DMM architecture will benefit the operators.

Moreover, DMM enables the scenario where IP sessions are not topologically bound to a gateway deployed in the operator's core network, like the LMA, but are instead anchored at the MAAR where the session started. In practical terms, if a mobile terminal is not moving and the communication peer is close to the MN, the communication latency can be dramatically reduced and the throughput increased. As already mentioned in Section 4.3 and in [17], we can devise multiple examples where this scenario takes place, for instance VoIP telephony or the use of CDN platforms.

5 EXPERIMENTAL EVALUATION

This section reports on the experimental evaluation conducted using a real implementation of our solution¹².

Our prototype is written in C and runs in Linux-based machines. The testbed deployed to perform the experiments is depicted in Fig. 9. It comprises five Linux Ubuntu 12.04 boxes (running a Linux-3.2.37 kernel): four regular PCs playing the role of three MAARs and one CMD, and one laptop playing the role of mobile node. The wireless access is provided by using IEEE802.11a/b/g WLAN cards.

Our experimental evaluation focuses on the handover latency. We use Wireshark¹³ at the MN to extract the events produced when repeating the following sequence: the MN attaches to MAAR1 and requests a UDP stream from a remote server outside the DMM domain; the MN then visits MAAR2 and MAAR3 before coming back to MAAR1, and then this handover sequence is repeated until the stream is over. We iterated the experiment obtaining more than 500 handovers. Note that MAAR1 is always the anchor for the data flow. Moreover, the CMD is attached to the same local subnet in our testbed, and therefore the obtained values do not clearly show the impact of the delay between the MAARs and the

12. MAD-PMIPv6: <http://www.odmm.net/>

13. <http://www.wireshark.org/>

TABLE 3
DMM Handover latency measurements.

	Mean (ms)	Std. Dev. (ms)
t_{L2}	29.49	9.92
$t_{binding}^{DMM}$	13.20	3.38
$T_{handover}^{DMM}$	42.69	10.58
$T_{session-rec}^{DMM}$	61.01	21.52

CMD, which might be non-negligible in large mobile network deployments. The tests consisted in measuring the different components of the handover latency (see Fig. 5) while the mobile node roams among the three deployed MAARs:

- t_{L2} : time required to perform a layer-2 switch from one MAAR to another (in our testbed using IEEE 802.11 technology). We measured it as the time interval between the L2 De-authentication and the Association Response messages.
- $t_{binding}^{DMM}$: time elapsed since the layer-2 handover has been completed until the MN receives the router advertisement message for the LNP. After this stage an MN configures an address from the IPv6 prefix locally anchored at the current MAAR and can use this address for new communications. We consider the *handover latency* as the sum of this and the previous term, $T_{handover}^{DMM} = t_{L2} + t_{binding}^{DMM}$.
- $T_{session-rec}^{DMM}$ (*ongoing sessions recovery time*): time elapsed since the mobile node starts moving to a new MAAR until it resumes reachability and connectivity of an old prefix anchored at a different MAAR from the serving one. It is measured as the interval between the last UDP packet before a handover and the first one after the same handover.

Table 3 shows the mean values and standard deviations for the previously listed components, whereas in Fig. 10, we plot the cumulative distribution functions of t_{L2} , $T_{handover}^{DMM}$ and $T_{session-rec}^{DMM}$. The results depicted in Fig. 10 show that the binding latency ($t_{binding}^{DMM}$) can be approximated to a constant deterministic component, that, added to the WLAN switching time, results into a handover latency of less than 50 ms for more than 90% of the handovers.

By observing the latency incurred to recover the IP flow ($T_{session-rec}^{DMM}$), it is interesting to note that some values are very close to the handover latency ($T_{handover}^{DMM}$), while others are higher, reaching almost 90ms. This high variability is due to the dependency between the handover latency and the distance between the Anchor-MAAR and serving MAAR targeted by the handover: if the mobile node is moving back to the Anchor-MAAR, the stream is recovered quicker than when the node is moving to a MAAR farther away from the Anchor-MAAR.

The results obtained from the experiments described in this section show that our proposed DMM approach




Fig. 10. CDFs of the DMM handover latencies.

is feasible and can be easily implemented. In Proxy Mobile IPv6, the handover latency depends on the latency between LMA and MAG (see Section 4.4), which can be considered approximately constant in the domain. In the case of our DMM solution, there are two different types of “handover latencies”, one representing the time elapsed since the MN detaches until a new prefix is available and ready for use at the new point of attachment, and one representing the time elapsed since the MN left the previous MAAR until the reachability of a previously assigned and still active prefix is re-gained. While the former mainly depends on the distance between CMD and MAARs which can be considered as approximately constant for all the MAARs, the latter depends on such distance between the CMD and the MAARs, and also between the A-MAAR and the S-MAAR. This means that keeping the reachability of a prefix anchored at a MAAR far away from the Serving-MAAR should be avoided.

Nevertheless, the advantage of DMM over PMIPv6 is the dynamic anchoring of data flows: in PMIPv6 all the traffic is encapsulated and has to traverse the operator’s network core. With DMM, mobility support is selectively offered only for those applications that cannot survive an IP address change, while the rest of application can always benefit from a better performance provided by the best available path from the S-MAAR to destination, and from not encapsulating the packets. This behavior is achieved by selecting the most appropriate source IP address, among those configured by the terminal, depending on the application’s needs. Modern portable devices already run a software acting as a connection manager to handle the data connections and the radio interfaces, hence the function of an advanced source address selection algorithm could be provided by this software.

6 CONCLUSIONS

Triggered by the explosion of mobile users' traffic demand, operators are looking at new network architectures where mobility and connectivity management are no longer centralized, so user traffic can be routed towards its destination at local breakout points at the edge of the network without traversing the core network. The Distributed Mobility Management (DMM) paradigm proposes to relocate anchors closer to the users, flattening the network.

In this article, we have conducted an analytic and experimental evaluation of a network-based DMM approach. We have derived analytic expressions of the signaling and packet delivery costs, the handover latency and the packet loss, of our solution, and compared them with the ones of a centralized network-based solution: Proxy Mobile IPv6. Additionally, we have performed experiments using a Linux-based prototype, which validate the feasibility of our solution, as well as have been used to experimentally assess the achievable handover latencies. The platform also served to show the DMM architecture during public exhibitions such in [18] and during the 87th IETF meeting in Berlin, Germany. To the best of our knowledge, this is the first work performing an analytic and experimental study of a DMM solution, showing the trade-offs to be considered when deploying this kind of architecture.

Obtained results show that the use of a DMM approach allows to save resources in the network in some situations, due to a reduced packet delivery cost. However, there are also some scenarios in which DMM incurs in higher costs, for example when the mobile node is running long-lasting applications and/or the cell residence time is short. In these situations, the use of a centralized mobility management solution would be preferred. Based on this observation, we can conclude that future mobile network architectures will exhibit an hybrid centralized-DMM behavior in which the mobility management of some traffic will be kept centralized, while other will be distributed.

REFERENCES

- [1] F. Giust, A. de la Oliva, C. Bernardos, and R. Da Costa, "A network-based localized mobility solution for Distributed Mobility Management," in *Wireless Personal Multimedia Communications, 2011 14th International Symposium on*, Oct. 2011, pp. 1–5.
- [2] C. Perkins, D. Johnson, and J. Arkko, "Mobility Support in IPv6," RFC 6275, July 2011.
- [3] S. Gundavelli, K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil, "Proxy Mobile IPv6," RFC 5213, Aug. 2008.
- [4] T. Narten, E. Nordmark, W. Simpson, and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)," RFC 4861, Sept. 2007.
- [5] P. Bertin, S. Bonjour, and J. Bonnin, "Distributed or centralized mobility?" in *Global Telecommunications Conference, 2009. GLOBECOM 2009. IEEE*. IEEE, 2009, pp. 1–6.
- [6] H. Chan, "Requirements for Distributed Mobility Management," Internet-Draft (work in progress), draft-ietf-dmm-requirements-12.txt, Dec. 2013.
- [7] D. Liu, J. C. Z. an P. Seite, H. Chan, and C. J. Bernardos, "Distributed Mobility Management: Current practices and gap analysis," Internet-Draft (work in progress), draft-ietf-dmm-best-practices-gap-analysis-02.txt, Oct. 2013.

- [8] C. Bernardos, A. de la Oliva, and F. Giust, "A PMIPv6-based solution for Distributed Mobility Management," Internet-Draft (work in progress), draft-bernardos-dmm-pmip-03, Jan. 2014.
- [9] F. Baumann and I. Niemegeers, "An evaluation of location management procedures," in *Universal Personal Communications, 1994. Record., 1994 Third Annual International Conference on*, Sept. 1994, pp. 359–364.
- [10] I. F. Akyildiz, J. S. M. Ho, and Y.-B. Lin, "Movement-based location update and selective paging for PCS networks," *IEEE/ACM Trans. Netw.*, vol. 4, no. 4, pp. 629–638, Aug. 1996. [Online]. Available: <http://dx.doi.org/10.1109/90.532871>
- [11] Y.-B. Lin, "Reducing location update cost in a PCS network," *IEEE/ACM Trans. Netw.*, vol. 5, no. 1, pp. 25–33, Feb. 1997.
- [12] Y. Fang, "Movement-based mobility management and trade off analysis for wireless mobile networks," *Computers, IEEE Transactions on*, vol. 52, no. 6, pp. 791–803, 2003.
- [13] S. Pack and Y. Choi, "A Study on Performance of Hierarchical Mobile IPv6 in IP-Based Cellular Networks," *IEICE transactions on communications*, vol. 87, no. 3, pp. 462–469, 2004-03-01.
- [14] C. Makaya and S. Pierre, "An Analytical Framework for Performance Evaluation of IPv6-Based mobility Management Protocols," *Wireless Communications, IEEE Transactions on*, vol. 7, no. 3, pp. 972–983, Mar. 2008.
- [15] I. Akyildiz and W. Wang, "A dynamic location management scheme for next-generation multitier PCS systems," *Wireless Communications, IEEE Transactions on*, vol. 1, no. 1, pp. 178–189, Jan. 2002.
- [16] S. Krishnan, R. Koodli, P. Loureiro, Q. Wu, and A. Dutta, "Localized Routing for Proxy Mobile IPv6," RFC 6705, 2012.
- [17] E. Demaria and L. Marchetti, "Dimensioning considerations for distributed mobility architecture," Internet-Draft (work in progress), draft-demaria-dmm-dimensioning-considerations-00, Mar. 2012.
- [18] F. Giust, A. de la Oliva, and C. Bernardos, "Mobility Management in Next Generation Mobile Networks [demo]," in *IEEE WoWMoM 2013*, June 2013.



Fabio Giust received his Bachelor's and Master's degree in Telecommunications Engineering at University of Padova, Italy. After an internship at Alcatel-Lucent Bell Labs in France, he undertook a Master in Telematics Engineering at University Carlos III of Madrid (UC3M), Spain. Currently he is working at UC3M, where he is also pursuing his Ph.D. His research interests cover IP mobility and wireless mobile networks, on which he has published several papers in international conferences and journals.



Carlos J. Bernardos received a Telecommunication Engineering degree in 2003, and a PhD in Telematics in 2006, both from 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 mobility in heterogeneous wireless networks. He has published over 50 scientific papers in international journals and conferences, and he is an active contributor to the IETF. He has served as guest editor of IEEE Network.



Antonio de la Oliva received a Telecommunication Engineering degree, and a PhD in Telematics in 2008, both from UC3M. In the last years, he has served as Vice-chair of the IEEE 802.21b task group and Technical Editor of IEEE 802.21d, contributing significantly to the development of the IEEE 802 standards for Media Independent Handover Services. Currently, he works as Visiting Professor at the Telematics Engineering department of UC3M, where he is performing lecturing and research activities.