Contents lists available at ScienceDirect

Computer Networks

fP2P–HN: A P2P-based route optimization architecture for mobile IP-based community networks

Ruben Cuevas [b,*], Albert Cabellos-Aparicio [a,*], Angel Cuevas [b], Jordi Domingo-Pascual [a], Arturo Azcorra [b,c]

[a] Universitat Politècnica de Catalunya, D. d'Arquitectura de Computadors, c/ Jordi Girona 1-3, D6-118, 08034 Barcelona, Spain
[b] Universidad Carlos III de Madrid, Telematic Department, Avenida de la Universidad, 30, 28911 Leganes, Madrid, Spain
[c] IMDEA Networks, Avenida del Mar Mediterráneo, 22, 28918 Leganes, Madrid, Spain

ARTICLE INFO

ABSTRACT

Wireless technologies are rapidly evolving and the users are demanding the possibility of changing their point of attachment to the Internet (i.e. Access Routers) without breaking the IP communications. This can be achieved by using Mobile IP or NEMO. However, mobile clients must forward their data packets through its Home Agent (HA) to communicate with its peers. This sub-optimal route (lack of route optimization) considerably reduces the communications performance, increases the delay and the infrastructure load. In this paper, we present fP2P–HN, a Peer-to-Peer-based architecture that allows deploying several HAs throughout the Internet. With this architecture, a Mobile Node (MN) or a *Mobile Community Network* (i.e. a NEMO) can select a closer HA to its topological position in order to reduce the delay of the paths towards its peers. fP2P–HN uses a Peer-to-Peer network to signal the location of the different HAs. Additionally, it uses flexible HAs that significantly reduce the amount of packets processed by the HA itself. The main advantages of the fP2P–HN over the existing ones are that it is scalable, it reduces the communications delay and the load at the HAs. Since one of the main concerns in mobility is security, our solution provides authentication between the HAs and the MNs. We evaluate the performance of the fP2P–HN by simulation. Our results show that the fP2P–HN is scalable since the amount of signalling messages per HA does not increase, even if the number of deployed HAs increases. We also show that the average reduction of the communication's delay compared to Mobile IP/NEMO is 23% (with a minimum deployment) and the reduction of the load at the HA is at least 54%.

1. Introduction

Wireless technologies have rapidly evolved in recent years. IEEE 802.11 is one of the most used wireless technologies and it provides up to 54 Mbps of bandwidth in an easy and affordable way. In the current Internet status, a user can be connected through a wireless link but he cannot move (i.e. change his access router) without breaking the IP communications. That's why IETF designed Mobile IP [29], which provides mobility to the Internet. With "mobility", a user can move and change his point of attachment to the Internet without losing his network connections.

In Mobile IP, a Mobile Node (MN) has two IP addresses. The first one identifies the MN's identity (Home Address, HoA) while the second one identifies the MN's current location (Care-of Address, CoA). The MN will always be reachable through its HoA while it will change its CoA according to its movements. A special entity called Home Agent (HA), placed at the MN's home network will maintain bindings between the MN's HoA and CoA addresses.

---

* Corresponding authors. Tel.: +934054063.
E-mail address: acabello@ac.upc.edu (A. Cabellos-Aparicio).

The main limitation of Mobile IP is that communications between the MN and its peers are routed through the HA. Unfortunately, packets routed through the HA follow a sub-optimal path. This reduces considerably the communications' performance, increases the delay and the infrastructure load. In addition, since a single HA may be serving several MNs and forwarding several connections, the HA itself may become the bottleneck of the whole system and represents a single point of failure in Mobile IP-based networks [1].

Mobile IPv6 [30] solves this limitation by allowing MNs to communicate with their peers directly (route optimization) by exploiting special IPv6 extension headers. However, the NEMO protocol (NEMOv4 [2] and NEMOv6 [32]), which provides mobility to networks instead of nodes, does not support route optimization, even in IPv6. That is why we believe that route optimization is an issue in the current Internet status (IPv4) and even in the future (IPv6). Note that a NEMO (NEtwork that MOves) can be seen as a *Mobile Community Network*. From the Internet infrastructure's point of view, a *Community Network* is a set of nodes located in the same geographical area. The nodes belonging to the *Community Network* are equipped with at least one wireless interface and can share information directly using an ad-hoc protocol. Regarding the connection with the Internet, the nodes belonging to the *Community Network* share a common point of attachment. This common point can be seen as the NEMO's mobile router. This router is equipped with two interfaces: an "external" long-range wireless interface intended to attach to the Internet and an "internal" interface intended to provide connectivity to the nodes belonging to the *Community Network*.

Solving the route optimization problem has attracted the attention of the research community and several solutions have been proposed [3–6]. The main idea behind these proposals is deploying multiple HAs in different Autonomous Systems (ASes). Then, a MN may pick the best HA according to its topological position thus, reducing the delay of the paths towards its peers. The main challenge of this approach is signalling the location of the different HAs throughout the Internet. Some of authors use the exterior Border Gateway Protocol (eBGP) protocol [3,5,6] while others [4] use Anycast routing. The main issue of these proposals is the scalability. On the one hand, using the exterior BGP protocol means increasing the load in the already oversized global routing table [7]. On the other hand, anycast's defiance of hierarchical aggregation makes the service hard to scale [8]. In addition, these solutions force the MNs to send the data packets through the HAs, increasing the load on these devices that may become the bottleneck of the whole system [1].

In this paper, we propose a scalable architecture, named fP2P–HN (flexible P2P Home agent Network) that solves the route optimization issue for Mobile IP and *Mobile Community Networks* (NEMO). We propose using an overlay Peer-to-Peer (P2P) network to signal the location of the different HAs [17]. When a MN detects that its current HA is too distant it queries its *Original HA* (the one serving the MN's Home Network) that belongs to the fP2P–HN network for a closer HA. Then, the fP2P–HN network uses BGP information to locate a HA that reduces the delay of the paths between the MN and its peers, for instance by choosing a HA located in the same AS as the MN. Since security is one of the main concerns in mobility, we also present an architecture that provides trustworthiness to the HAs belonging to the P2P network and that allows that the MNs can be authenticated by the HAs (and vice versa).

Our solution allows deploying multiple HAs at different ASes without impacting the exterior BGP global routing table or requiring anycast routing; however, the HAs are still responsible of forwarding all the MN's data packets. In order to alleviate their load, we propose to deploy flexible HAs (fHA) [18]. The main idea behind the fHAs is that a registration from a MN to a HA can be viewed as an internal route from the network's point of view. That is, when a MN registers a new location into its HA, it is actually installing a new route (Home Address → Care-of Address). We believe that this route can be announced throughout the network using the interior BGP (IBGP [31]) protocol to each of the AS' Border Routers. Then, the Border Routers are aware of the current location of the MN and will decapsulate and forward any packets addressed to/from the MN directly, just as regular packets. Thus, MN's data packets are not forwarded by the HAs but by the Border Routers. It is worth to note that HAs are not necessarily devices designed for routing purpose whereas routers are routing-dedicated devices.

Our solution fP2P–HN is simple, scalable and secure. Moreover it does not require deploying any new entities on the Internet. At the Inter-domain level, we signal the location of the HA using a P2P network instead of using eBGP or anycast. At the Intra-domain level we signal the location of the MN using IBGP, in this way the Border Routers are aware of the location of the MN and the load of the HA is significantly reduced. As we will see later, we evaluate the performance of our proposal through simulation. Our results show that the fP2P–HN is scalable since the amount of signalling messages per HA does not increase, even if the number of deployed HA increases. This amount of signalling, in the worst case, is around 20 kbps per HA. We also show that the average reduction of the communication's delay compared to Mobile IP/NEMO grows from 23% (with a minimum deployment) up to 80% (with large deployments). Whereas the reduction of the load at the HA varies between 54% (in the worst case) and nearly 100% (in the best case).

In our previous work, we presented a P2P Home Agent network that signals the location of different HAs throughout the Internet [17]. In [18], we presented the flexible HAs, that reduce significantly the traffic load. The main contributions of this paper are three: the first contribution is the novel architecture fP2P–HN (Section 2) which is based on both solutions. The second contribution is the evaluation of the solution (Section 3). Finally, the third contribution, is a security architecture (Section 2.7) that provides authentication to the nodes belonging to the network.

2. Flexible P2P home agent network

In this section, we detail the fP2P–HN architecture. Please note that an fHA (flexible HA) is a Home Agent that

belongs to the architecture and that has special features. In this paper, we will refer to a HA or an fHA indistinctively.

2.1. Overview

The main goals of the fP2P–HN architecture are to reduce the delay of the communications of the MNs and the load at the fHAs. Fig. 1 shows an overview of the architecture.

When a Mobile IP or NEMO client changes its point of attachment to the Internet it establishes a new tunnel with its HA to communicate. Depending on the MN's topological position, this new path may have a large delay. We propose to deploy several HAs throughout the Internet in order to reduce this delay. When the MN detects that the new path to its currently assigned HA has an unacceptable performance (e.g. RTT ⩾ a given threshold) it queries its *Original* HA (the HA serving the MN's Home Network at the MN's administrative domain) for a closer one (i.e. an HA located in the MN's current AS). Our architecture is flexible and allows using any metric to trigger the discovery of a closer HA. In this paper, we use the RTT because it is a simple metric able to capture the performance of a path. It is worth noting here that any other metric can be used.

Our proposal requires deploying several HAs throughout the Internet and has four differentiated phases. The HAs organize themselves in a P2P network which stores the information regarding their addresses and their topological position (HA's AS number). This P2P network is formed during the *P2P Setup phase*. The MNs are always bound to a HA belonging to this P2P network. Thus, when the MN detects that the RTT to its current HA is unacceptable it triggers the *fHA Discovery phase* and queries the P2P network for a closer HA. Once the MN has the IP address of this closer HA it sends a registration message (Binding Update) and obtains a new HoA (*fHA Registration phase*). The MN keeps using this HoA while the RTT remains below a given threshold.

All the HAs deployed in the fP2P–HN architecture are in fact flexible HAs. This means that they belong to the IBGP

domain of its AS. When their assigned MNs are attached directly to their AS they act as a regular HA. However, when the MNs are outside their AS, they announce the location of the MNs (Care-of Address) through IBGP to the AS' Border Routers (BR). This announcement is just a new route: To reach the MN (Home Address) packets must be addressed to its topological position (Care-of Address). This way, packets addressed from/to the MN are directly processed by the BR and thus, the load at the HA is considerably reduced. This is the last phase of the proposal known as *Data Packet Forwarding*.

2.2. P2P Setup phase

This subsection details how the P2P network is created. The P2P network is used to store the location of the fHAs (AS number) and their IP addresses. This information is used by MNs to locate a closer fHA to its topological position.

fHAs organize themselves forming a structured P2P overlay (also known as DHT-based P2P overlay). The fP2P–HN is fully flexible and can be deployed using any of the proposed structured P2P schemes [13]. In the remainder of the paper, we will consider Chord [14] as the P2P scheme, thus, the overlay's structure is a ring.

In the fP2P–HN, the search key is the *AS-key* that is computed as *hash(AS number)*. When a new fHA joins the fP2P–HN it chooses an identifier (Peer-ID). In our case, this is the *hash(fHA's IP Address)*. The fHA's position in the ring is determined by its *Peer-ID*: the fHA is placed between the two overlay nodes with the immediately higher and lower *Peer-ID* to its own id. Each overlay node has direct references to its two neighbours and also to other overlay nodes (crossing the ring) thus making the routing within the fP2P–HN faster. These nodes are named *fingers*. Each overlay node uses these fingers to create its fP2P–HN routing table.

Finally, each fHA must register its AS number within the fP2P–HN. The fHA obtains the AS-key by computing the *hash(AS number)*. Then, it looks for the overlay node with the immediately higher *Peer-ID* to the *AS-key*, named *successor*, and sends to this node the *AS-key*, its IP address and its AS number. Moreover, the fHA sends some security information (See Section 2.7 for more details). The *successor* stores an entry with all this information.

2.3. fHA Discovery phase (inter-domain)

This subsection details (Fig. 2) how a MN can use the fP2P–HN to discover a closer fHA. An MN connected to fHA_1 eventually detects (after a handover) that the RTT to fHA_1 is above a given threshold. Then, it triggers the procedure to discover a closer HA. The MN sends to its *Original fHA* a special BU soliciting the IP address of a closer fHA. At this point, the *Orginal fHA* discovers (using BGP) the AS number associated to the MN's CoA. Afterwards, it obtains the *AS-key* by computing the *hash(AS number)*.

The search method within the fP2P–HN is as follows. The *Original fHA* sends a query with the AS-key. The search query is routed in the overlay towards the AS-key's *Successor*. This fHA (e.g. fHA_2) is responsible of storing the


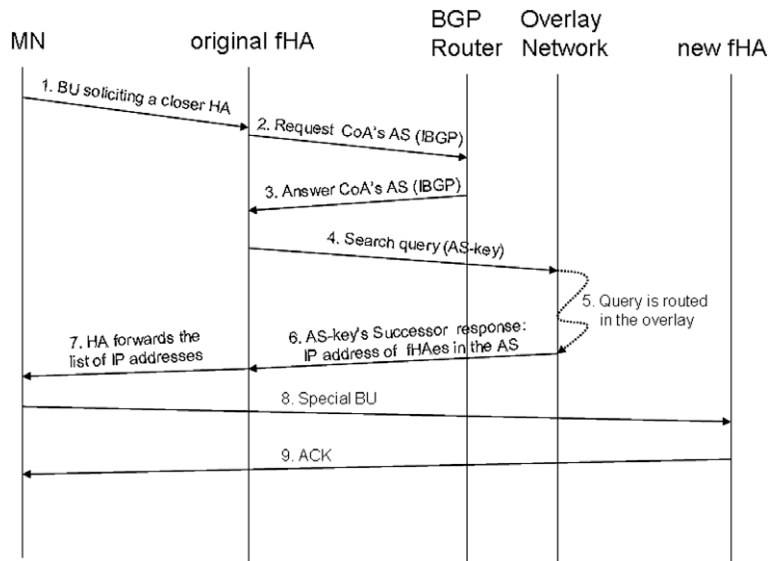**Fig. 1.** Overview of the fP2P–HN architecture.

Fig. 2. fHA Discovery phase in the fP2P–HN architecture.

information regarding the AS-key. Thus, it stores the IP addresses of all the fHAs located in the AS where the MN is currently attached. Then, fHA_2 sends these IP addresses to the *Original fHA* which in turn forwards them to the MN. Finally, the MN selects one of them and sends a special BU message to the new fHA in order to obtain a new HoA.

Although the fHAs are expected to be very stable entities, the fP2P–HN includes the mechanisms to make the solution dynamic and adaptive. For this purpose, every fHA periodically checks whether its neighbours and fingers are still reachable and running. If necessary, the fHA reconfigures its fP2P–HN routing table and establishes new neighbours or fingers.

Moreover, to make the solution more robust, reliable and load-balanced we use redundancy. Each AS-key is stored for several *successors* instead of just one. Then, in case of failure of a *successor* the others are still available and can reply to the queries. In addition, each MN has the list of the fHAs obtained during the last fHA discovery phase. Thus, if its current fHA fails, the MN can re-connect to one placed on the same AS.

2.4. fHA Registration phase (intra-domain)

This subsection details the registration phase of a MN into a new fHA. At the Intra-Domain level, each MN selects a given fHA through the above-mentioned mechanism. Our fHA has the same functionalities as a regular HA but it uses IBGP to signal the location of the MNs to reduce the load. The fHA acts just as a regular HA when the MN is directly attached to its network.

When the MN is not directly attached to its AS, the fHA has to announce the new location of the MN (CoA) to the AS' BRs. To distribute this type of information we use the Interior Border Gateway Protocol (IBGP). In the fP2P–HN, the fHAs and the BRs create an IBGP domain. This IBGP domain may be an already existing one or a separate one. The

routes announced through this IBGP domain always have the longest prefix (/32) and never affect regular BGP routes. It should be noted that the routes announced by the fHAs will *never* be distributed outside the AS. Finally, the entities participating in the IBGP domain have pre-configured keys to provide confidentiality, integrity and authentication for the communications.

For each received registration message (Binding Update) from outside the AS, the fHAs send an IBGP UPDATE message to the BRs. We introduce new options in the IBGP UPDATE message. The UPDATE message sent to the BRs includes the following information: ⟨*Home Address, Care-of Address, Lifetime*⟩. Upon reception of this message, the BRs setup a tunnel endpoint with the MN. The tunnel source address is the one of the BR's address while the destination address is the Care-of Address. In addition, each BR adds the following route to its routing table: *HomeAddress\32 → Tunnel*. The tunnel and the route are automatically deleted after "Lifetime" seconds. Finally the fHA will reply to the MN informing that the registration was successful and with the list of addresses of the BRs; this way the MN can address its tunnelled packets towards the BRs (see section below for details).

Once the MN is assigned to a new fHA or returns home it sends a registration message to the previous fHA. Upon reception, the fHA sends an IBGP WITHDRAWAL message to the BRs to immediately remove all the routes and tunnels related to the MN's Home Address.

Finally, since several fHAs can be deployed on the same AS, the MNs will receive a list of the available fHAs and will choose one based on any criteria (load balancing, RTT,...).

2.5. Data packet forwarding phase (intra-domain)

This subsection details how an MN's data packets are forwarded. If the MN is connected to the fHA's AS, then packets are forwarded just as in Mobile IP or NEMO.

However, when the MN is attached to a foreign AS, then the MN should forward the packets through its HA. However, since the HA is an fHA, the MN encapsulate its data packets towards the BRs (Fig. 3). Since the fHA has previously configured (using IBGP) a new tunnel (*HomeAddress\32 → Tunnel*) in the BRs, packets sent by the MNs are automatically de-capsulated and forwarded towards the packet's destination address (the MN's peer address). If the exit point of the MN's peer address is another BR, then the packet traverses the network as a transit packet.

Regarding the packets addressed towards the MN's HoA they will reach the fHA's AS. The BRs have learned the location (CoA) of the MN through IBGP and will automatically encapsulate and forward the packet directly towards the MN.

2.6. Flexible home agent location

In the previous sections, we have assumed – for clarity – that each fHA belongs to an IBGP domain with the AS's BRs. However, our solution is flexible and allows that multiple sets of fHAs can be deployed in different networks of the AS. Then, each set of fHAs belongs to an IBGP domain with its network's *Exit Routers*. Fig. 4 presents an example.

In this example, the AS has two different networks (A and B). Two different sets of fHAs are deployed in network A and B. Thus, only routers labelled in black must belong to the IBGP domain with the fHAs of their network. The only requirement that these *Exit Routers* have to fulfill is being in the path of the packets addressed to the HoA delegated by the fHA.

A MN attached to this AS (A or B) is assigned to a given fHA; let's say one located at the A network. Then, it will receive a Home Address that belongs to the prefix of the network A. Thus, all the packets sent towards the MN will be received by the A's Exit Routers and forwarded directly to the MN. As noted previously, the MN encapsulates its data packets towards the A's Exit Routers that, in turn, de-cap-

sulate and forward towards the packet's destination address (the MN's peer).

2.7. Security considerations

In Mobile IP and NEMO, the mobile clients and the Home Agents are under the same administrative domain. That is why they are equipped with pre-configured keys. These keys provide, among others, two essential security properties to the mobile communications, trustworthiness and confidentiality. This means that the MNs and the HA can trust each other since they are authenticated. Additionally, ciphering techniques can protect the communications.

However, the MNs of the fP2P–HN may connect to different fHAs that, may or may not be under the same administrative domain. This section addresses the security at the fP2P–HN. Our goal is to achieve the same level of security as in Mobile IP and NEMO, that is: trustworthiness and confidentiality. In addition, we also provide mechanisms to achieve a third security property, non-repudiation, but only when it is required.

It must be considered that security solutions are highly dependent on the application scenario. In this section, we analyze security in two potential fP2P–HN scenarios: (i) the fP2P–HN is deployed by an unique organization and (ii) the fP2P–HN is formed by fHAs belonging to different organizations, typically Internet Service Providers (ISPs). In both scenarios, we address the security of the two types of communications present in the proposed solution: fHA-fHA and fHA-MN communications.

2.7.1. Scenario I: fP2P–HN deployed by an unique organization

In the first scenario, all the fHAs are deployed by the same organization. Several approaches can be used in order to provide fHA-fHA trustworthiness. For instance, all the fHAs own a X.509 certificate [25] provided by the organization that authorizes them to use the fP2P–HN services. This certificate provides trustworthiness, because any fHA can require another fHA's certificate in order to validate this second one as a legitimate entity. After being trusted, the fHAs involved in a communication can negotiate a shared key to provide confidentiality. This can be done by negotiating a session key based on Public/Private keys pair generated by each fHA (A public key could be also included along with the certificate provided by the organization). Finally, non-repudiation is obtained if each fHA is required to sign every data packet with its private key.

For fHA-MN communication, MNs are granted with a credential from the organization in charge of the fP2P–HN. This credential allows unique identification of a MN in the system and could be provided in different ways: hardware device, SIM card, a user/password pair, a certificate, etc. Thus, in order to achieve trustworthiness, the MN obtains the fHA's certificate and the fHA requests the credential from the MN. Again, confidentiality is obtained by negotiating a session key between the MN and the fHA. Finally, if non-repudiation is required, it is achieved if fHAs sign the data messages using their private keys and MNs include their credentials within the messages.


Fig. 3. Data packet forwarding.

Fig. 4. Example of location of the fHAs.

2.7.2. Scenario II: fP2P–HN deployed by several organizations

This second scenario requires more complex security mechanisms because many different organizations are involved in the fP2P–HN deployment. Again, the most important requirements for the proposed solution are trustworthiness and confidentiality, but also non-repudiation is analyzed.

We propose using a trusted third party (TTP) in order to achieve these goals. This TTP is trusted by all the organizations participating within the fP2P–HN and thus, by all the fHAs belonging to these organizations.

In this scenario, the organizations that offer mobility services are typically the ISPs. In addition, an ISP is (usually) an AS within the Internet architecture. Thus, we assume that all the fHAs belonging to an AS are managed by a single ISP.

In this architecture, each ISP participating in the fP2P–HN is granted with a X.509 certificate obtained from the TTP. This certificate contains, among other elements: the AS Number, the AS public key (AS_pu_key) and the valid period. It must be taken into account that each ISP has an AS private key (AS_pr_key) paired with the AS_pu_key. Then, all the fHAs deployed in a given AS use that certificate within the fP2P–HN. Only fHAs belonging to an ISP participating in the fP2P–HN are provided with such certificate. Therefore, based on this approach, we are able to provide the required security properties in the fHA–fHA communications.

Trustworthiness is achieved because only fHAs owning such a certificate (provided by the TTP) are trusted by the rest of fHAs within the fP2P–HN. Therefore, at any time a given fHA, fHA_1, could request from another fHA, fHA_2, its certificate to check whether fHA_2 is an authorized entity or not.

After both fHAs trust each other, they negotiate a shared key in order to provide confidentiality to the fHA-fHA communication. Several approaches could be applied at this point. For instance, the fHA_1 can provide a $nonce_1$[1]

---
[1] A nonce is a long random number.

encrypted with the AS_pu_key_2 to the fHA_2, and similarly fHA_2. Therefore, both peers create a shared key using the nonces as input parameters to a given function. For instance, $Shared\ Key = f(nonce_1, nonce_2) = nonce_1\ XOR\ nonce_2$.

In order to secure, the fHA-MN communications, we propose a similar approach to that used in GSM [26–28] that validates users owning a SIM card using a credential. In GSM, when an user is attached to a foreign operator (roaming), it has to present its credentials to the new operator. Then, the new operator contacts the home operator and uses the received credentials to validate the user.

Following this approach, in the fP2P–HN the home AS (an ISP with the certificate provided by the TTP) provides credentials to its MN clients. This credential could be: a certificate, an unique ID like in GSM networks, etc. Therefore, once a MN selects a new fHA from a different ISP, it presents its credential and its home AS number to the new fHA. In turn, the new fHA validates the MN by sending to one of the fHA in the MN's home AS the credential. Then, based on the received credential, the fHA in the home AS checks if the credential's owner is an authorized user and returns the validation result to the new fHA. If the validation is successful the new fHA can trust the MN.

Finally, each MN has a permanent trusted connection with its *Original fHA*. Thus, the MN also trusts the new fHA because it has been authenticated by its *Original fHA*. This means that the new fHA is trusted by the *Original fHA* and also by the MN. Therefore, trustworthiness is achieved in both directions. After that, a shared key could be negotiated between the fHA and the MN in order to provide confidentiality for the communications. Non-repudiation is achieved (if required) by applying the same mechanism introduced in the previous scenario.

2.8. Final remarks

In this subsection, we discuss the final considerations of the fP2P–HN. First, changing the MN's HoA may break the existing connections. In order to solve this issue, we propose that these connections are forwarded through the

previous fHA while new connections are forwarded through the new fHA. A MN changes its HoA only when it is outside of its currently assigned fHA's AS and the RTT is above a given threshold. ASes usually provide connectivity to very large geographical areas, thus, this will occur rarely. In addition, 98% of the connections last less than 15 minutes [16], this means that very few connections may be affected. Regarding the inbound connections, the MN may still use its original HoA (the one from its Home Network). In fact, as we have seen in Section 2.7, the MNs use this connection for authentication purposes, therefore, MNs are always reachable through their regular Home Addresses. It is worth to note that MNs are clients (not servers) and with the current deployment of firewalls and NATs inbound connections are almost non-existent.

Second, the regular Mobile IP or NEMO handovers (i.e. changing the access router) are not affected by the fP2P–HN. That is, the procedural operations of the regular handovers are exactly as defined in the Mobile IP and NEMO standards. Therefore, the latency of these handovers is the same in our approach as in Mobile IP or NEMO. Furthermore, the fP2P–HN adds a second handover type that occurs when the MN changes its HA and its HoA. Then the handover latency is higher than in the regular one because it includes the search process in the P2P network. However, since the existing connections are being forwarded through the previous HA, this extra handover latency does not affect the communications. We can conclude that although our solution introduces a new type of handover that suffers from a higher latency, this does not impact the performance of the communications.

Finally, the architecture requires minor modifications in the MNs and HAs. Obviously, the HAs must include an implementation of the fHA and the P2P algorithms. Regarding the MNs, they must include a triggering mechanism to discover a closer HA. As noted previously, this mechanism can use any metric, in our paper, we have used the RTT. In addition, the MNs must support multiples HoAs, this is already under standardization by the MEXT WG [19]. The signalling between the MNs and the fP2P–HN can be accommodated into the Mobile IP signalling by exploiting the *Extensions* field present in the Binding Update messages (see [29] for details). Finally, the rest of the entities participating in the solution (CNs and routers) do not need to be modified. Since Mobile IP has not been deployed yet, we believe that the deployment cost of Mobile IP enhanced with the fP2P–HN would not increase.

3. Evaluation

The fP2P–HN architecture introduces two major improvements on Mobile IP and NEMO which are: the reduction in the delay of the communications and the reduction in the load at the HAs. However, these improvements increase the signalling load in both, Intra (IBGP) and Inter-domain (P2P) levels. In order to evaluate the advantages (*reduction in the communication's delay* and *reduction in the load at the fHAs*) and the costs (*Inter-Domain Signalling* and *Intra-Domain Signalling*), we have implemented the fP2P–HN in a simulator.

3.1. Simulation setup

In order to simulate the proposed solution, we have used Internet-like topologies generated with the last version (3.0) of *Inet* [21]. An earlier version of this random topology generator was presented in [20]. We have chosen *Inet* as the topology generator because it has been designed based on the analysis of public NLANR (National Laboratory for Applied Network Research) data-traces [22]. These traces, well-known by the passive measurements research community, have been collected from a variety of links at different networks. This means that *Inet* does not produce synthetic topologies, but realistic topologies based on real data-traces. In addition, *Inet* fulfils the requirements since it is intended to model AS-level connectivity instead of router-level connectivity. Regarding the mobility model, we have used the Random Waypoint Mobility simulator [15]. This simulator implements the well-known Random Trip Model [23] that was proposed as a generic mobility model. We refer the reader to [15,21] for further details.

Node-level simulators such as NS-2 or OMNET do not scale when simulating a large number of ASes. On the other hand AS-level simulators such as C-BGP or simBGP are not intended to include end-host mobility. That is why we have developed an ad-hoc simulator. We have implemented our simulator using Perl [33], the topology is generated using the *Inet* topology generator and the Random Waypoint Mobility model has been implemented into the simulator. The AS topology is stored as a graph using CPAN's *Graph* library and, for each MN, and after each movement, the shortest path to it fHA is computed using the Floyd–Warshall algorithm [34].

Armed with a topology generator and a mobility model we have developed an ad-hoc simulator. Unless noted otherwise, we have simulated an average number of 100 mobile clients per fHA. The MNs are distributed randomly (uniformly) among the fHAs, this means that the fHAs do not necessarily serve the same number of MNs. Each MN is assigned to a given Home Network (uniformly); the location if this Home Network is assigned randomly. For each handover, the MN has a 10% of probability of remaining in the same AS and, after a handover it remains attached to the same access router during a random amount of time distributed as (Gaussian) $N(5,1)$ s. When the MN remains in the same AS, it means that it is changing its access router (CoA). Obviously, these values produce highly mobile nodes compared to the movements in real environments, however, we aim to evaluate our solution in a stressful scenario. Regarding the delays of the links, we consider that each link has a constant delay uniformly distributed as $U[10,25]$ ms. Finally, each MN sends 1 unit of bandwidth per second towards its Home Agent (for Mobile IP) and 1 unit towards its flexible Home Agent (for fP2P–HN). Since we aim to compare the load of both proposals a CBR data stream suffices. The MN's threshold to trigger the fHA discovery procedure is set to 75 ms.

We run each simulation during 1000 s (simulation time) running fP2P–HN and Mobile IP/NEMO. We consider the following deployment scenarios {0.01, 0.1, 0.3, 0.6, 0.75, 0.9}. These numbers represent the probability of deploying one fHA for each AS. In the case of Mobile

IP/NEMO, we consider the same number of HAs and the same number of MNs. Finally, we repeat the simulation of each deployment scenario 50 times with a different topology of 3500 ASes. The different topologies are generated using *Inet* (different seeds). In total, we have run 300 simulations. With this setup we simulate a wide range of scenarios, and we obtain the needed statistical information to assure the accuracy of the results. This accuracy is represented by the 90% Confidence Intervals included in every table and figure.[2] In order to run this huge amount of simulations, we have used a cluster of 70 machines (Intel Xeon, 16Gb RAM) that uses Sun's N1 Grid Engine [24].

The graphics included in this section represent the Cumulative Distribution Function,[3] (CDF) of the different evaluated aspects and also provides the Confidence Intervals of the calculated CDF. In order to obtain the CDF, first we compute the discrete probability density function (pdf) of the data. That is, we calculate the data distribution histogram. The histogram resolution (i.e. the width of the histogram intervals) was selected small enough to avoid information losses. Once we had the histogram, the CDF is the result of computing the histogram's cumulative sum. This process was repeated for each one of the 50 simulation samples. Thus, once we had the 50 CDFs we estimated the Confidence Interval for each one of the CDF points (that is, for each one of the histogram intervals). Since the histogram resolution is very high, the Confidence Intervals are not represented for every point since the figure would not be understandable.

3.2. Simulation results

3.2.1. Reduction of the communication delay

Firstly, we focus on the analysis of the communication delay since this is the main issue of Mobile IP and NEMO. Fig. 5 shows the delay of the communications in the path between the MN and its current HA, both for Mobile IP and for the fP2P–HN. The figure presents the CDF of the average delay suffered by each MN. The results show that, for a very low deployment (1%), the fP2P–HN slightly outperforms Mobile IP/NEMO. However, increasing the deployment up to 10%, the reduction of the delay achieved by the proposed solution is around 30%. This confirms, that even in the case of low deployments, our solution clearly outperforms Mobile IP or NEMO. Moreover, if we analyze the cases of higher deployments, fP2P–HN reduces the communication delay up to 6 times compared to Mobile IP or NEMO.

Table 1 summarizes the results of Fig. 5. It shows the mean MN-HA communication delay for both fP2P–HN and Mobile IP/NEMO.

Thus, we can conclude that in terms of delay, fP2P–HN introduces a major improvement compared to the Mobile IP or NEMO solutions.

3.2.2. Reduction of the load at the fHAs

In addition to the Route Optimization problem, the fP2P–HN addresses the reduction of the data traffic load at the HA as well. For this purpose, we have introduced the concept of fHA. Fig. 6 depicts the Complementary CDF (CCDF) of the percentage of saved traffic at the fHA compared to the regular Mobile IP's HA. The obtained results show that fP2P–HN introduces a major reduction of the load at the HA. The percentage of load reduction decreases along with the deployment. In the case of 1% of deployment, we find that around half of the fHAs are free of data traffic load. This means that they delegate the forwarding task to the Exit Routers. Even considering large deployments ($d = 0.9$), 80% of the fHAs experience a load reduction larger than 50%.

Table 2 shows the mean values. It must be noted that even in the worst case ($d = 0.9$) the mean load reduction with the fP2P–HN is 54.56%.

The reader may wonder why the percentage of saved traffic decreases as the deployment increases. This is because the fHAs delegates the forwarding of traffic from/to the MN when this is not directly attached to the fHA's AS. Whereas, if the MN is attached to its fHA's AS, then the fHA is responsible for forwarding the traffic from/to the MN. Hence, if we consider a large deployment of fHAs, it is more likely that the MNs are attached to its current fHA's AS so that the fHA suffers from higher load. On the other hand, in case of low deployments, the probability that the MN finds an fHA in its current AS is lower. Then, the MN maintains the connection to the fHA located in a different AS which delegates the forwarding task to the Border Routers. Thus, the fHA's load is lower with low deployments.

In a nutshell, the higher the deployment, the higher the probability that a MN uses an fHA placed at its current AS; thus more data traffic is forwarded by the fHAs.

3.2.3. Inter-domain signalling

As has been explained above, existing solutions addressing the problem of Route Optimization for Mobile IP and NEMO are not scalable. However, the fP2P–HN uses P2P (an scalable technology) in order to signal the location of the HAs. In this section, we evaluate the number of Inter-domain (P2P) signalling messages required to run the fP2P–HN.

Fig. 7 shows the inter-domain (P2P) signalling generated by the fP2P–HN to signal the location of the different fHAs. This figure depicts the CDF of the number of inter-domain signalling messages per second (sent + received) that a fHA has to support in the fP2P–HN. We can observe that the signalling overload introduced by the fP2P–HN remains between 50 and 100 messages/s for all the analyzed deployments. Therefore, the fP2P–HN requires a low number of Inter-domain signalling messages. Moreover it must be considered that these messages are usually short messages; thus the bandwidth consumption is negligible. For instance if we consider the worst case of the figure (50 sent + 50 received messages per second) and we assume that each message has 50 bytes (a Mobile IPv4's Binding Update message has 44 bytes, see [29]); then the amount of signalling traffic that an fHA has to support in the fP2P–HN is 20 kbps (both uplink and downlink).

---

[2] In some figures, the Confidence Intervals are so narrow they appear as a point in the figure or are smaller than the symbol representing the point.
[3] In case of Fig. 6. the Complementary CDF is represented instead of the CDF.

(a) deployment=1%

(b) deployment=10%

(c) deployment=30%

(d) deployment=60%
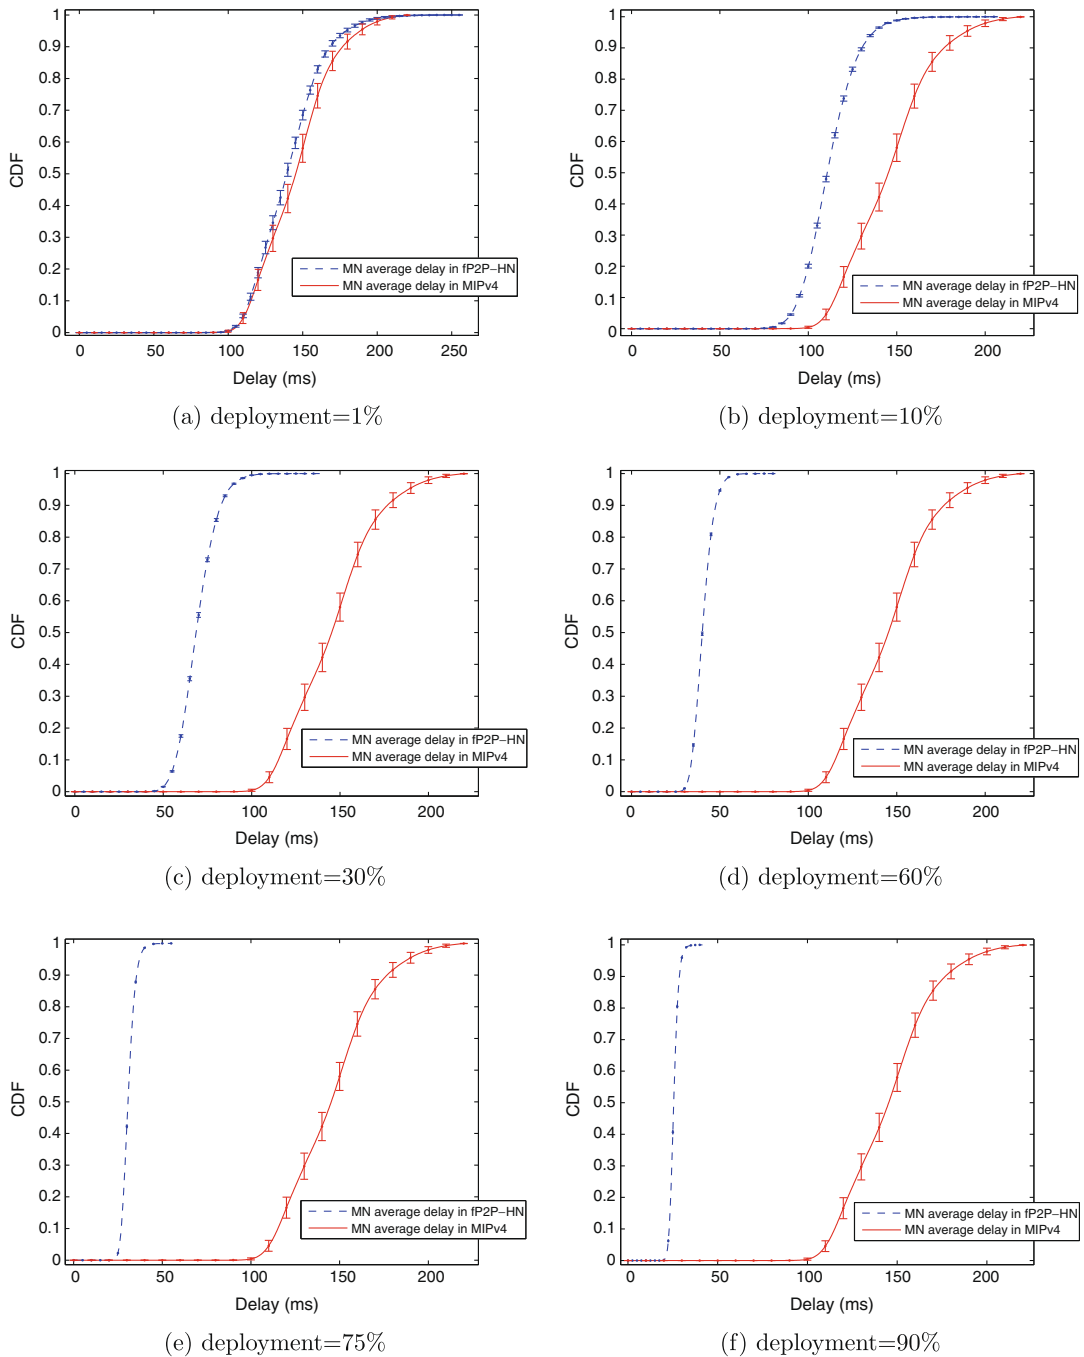

(e) deployment=75%

(f) deployment=90%

Fig. 5. Average communications delay in the MN-HA path.

Table 3 presents the mean number of total messages/s supported by the fHA.

Again it is worth analyzing the signalling overload as function of the deployment. The reader can observe that the overload increases as the deployment goes from 1% to 10%, and from this point it decreases along with the deployment increment. There are two parameters affecting the inter-domain signalling: the number of fHAs forming the fP2P–HN and the number of *special BUs* soliciting a new fHA (fHA discovery procedure). The number of fHAs has an influence since the fHA discovery procedure takes place at the overlay level and the query is routed by several fHAs within the fP2P–HN. The number of fHAs routing each query is bounded by $O(\log_2(N))$ [13] (where N is the number of fHAs forming the fP2P–HN). Thus, as deployment grows (larger N), more fHAs are involved routing each query. On the other hand the number of *special BUs* gets reduced as the deployment increases. With large

Table 1
Mean MN-HA communication delay.

| Deployment	fP2P–HN (ms)	Mobile IP (ms)	Reduction of the delay (%)
0.01	140.86 ± 0.95	145.83 ± 0.29	3.41
0.10	112.12 ± 0.31	145.83 ± 0.29	23.12
0.3	69.63 ± 0.16	145.83 ± 0.29	52.25
0.6	40.77 ± 0.07	145.83 ± 0.29	72.04
0.75	31.22 ± 0.04	145.83 ± 0.29	78.59
0.9	25.93 ± 0.03	145.83 ± 0.29	83.25

Fig. 6. Percentage of fP2P–HN's saved data traffic regarding MIPv4.

Table 2
Mean load reduction at the fHA compared to mobile IP.

| Deployment	Load reduction (%)
0.01	99.31 ± 0.02
0.10	92.72 ± 0.03
0.3	78.94 ± 0.06
0.6	64.81 ± 0.04
0.75	59.35 ± 0.02
0.9	54.56 ± 0.72

Fig. 7. fP2P–HN Inter-domain signalling traffic.

Table 3
Mean number of interdomain signalling messages/s per fHA.

Deployment	Number of fHAs	Mean number of (sent + received) messages/s
0.01	35	66.77 ± 0.14
0.10	350	94.46 ± 0.16
0.3	1050	89.23 ± 0.44
0.6	2100	75.21 ± 0.60
0.75	2625	63.32 ± 0.50
0.9	3100	67.63 ± 8.99

deployments is expected that MNs will always be connected to very close fHAs and that the fHA discovery process will be rarely unsuccessful. Therefore, both parameters compensate each other. Thus, when the deployment increases from 1% to 10%, the increment of the number of fHAs outweighs the increment of the number of *special BUs* and the signalling load grows. For larger deployments the situation is reversed resulting in a signalling load reduction.

In order to further study this behaviour, let's consider Table 4. This table details the probability of triggering the fHA discovery procedure for each deployment scenario (the values have been collected from the simulations). As the table shows, when the deployment is low, the MNs initiate the fHA discovery procedure more often. This is because MNs detect that the RTT is above a given threshold, ask for a closer fHA, but, since deployment is low, do not find one. Hence, the probability of triggering the fHA discovery procedure decreases as the deployment increases.

Finally, we can conclude that the fP2P–HN is scalable. Considering a highly mobile simulation scenario and 100 MNs per fHA, the number of signalling messages in the worst case is 20 kbps. On the other hand, Table 3 shows that the number of signalling messages is irrespective of the number of deployed fHAs. In fact independent of the deployment, the overload values are within the same order of magnitude (hundreds). Hence, the inter-domain cost of the proposed solution is $O(1)$.

3.2.4. Intra-domain signalling

Finally, we analyze the Intra-Domain signalling. This signalling includes the IBGP (UPDATE and WITHDRAWN) messages sent to the *Exit Routers* and the BGP queries sent to discover the MN's AS (see steps 2 and 3 in Fig. 2). This overload must be supported within each AS. Fig. 8 shows the CDF of the amount of signalling per AS (per second), considering the different deployment scenarios. As the figure shows the number of signalling messages is bounded between 0 and 70 (sent + received) messages/s. Again,

Table 4
Probability of triggering the fHA discovery procedure.

Deployment	Probability
0.01	0.73
0.10	0.64
0.3	0.47
0.6	0.35
0.75	0.29
0.9	0.27

Fig. 8. fP2P–HN Intra-domain signalling traffic.

Fig. 9. Reduction of the communication's delay and fHA's load.

considering a message size of 50 bytes, the download/up-load rate is less than 15 kbps. Additionally it has to be taken into account that this number is the total amount of signalling traffic supported inside each AS. Since the fP2P–HN allows deploying multiple fHAs within an AS (Section 2.6) each fHAs should only process a part of this signalling overload.

Regarding the mean values, Table 5 shows the results. The Intra-Domain signalling decreases as the deployment decreases. This is an expected result, since when MNs are directly attached to its fHAs no IBGP signalling is produced.

3.3. Summary of the obtained results

This section has evaluated the advantages and costs introduced by the fP2P–HN in front of the standard Mobile IP/NEMO protocols. The conclusion is that the fP2P–HN solves the main drawbacks of Mobile IP/NEMO (communication's delay and HA overload) with a low cost, some dozens of kbps in terms of extra signalling traffic. The obtained improvement depends on the deployment of the fP2P–HN. Fig. 9 summarizes in a single graphic the improvements (load reduction and communication delay reduction) introduced by the fP2P–HN as function of the deployment. This figure allows us to determine the required deployment in order to achieve a given performance. For instance if we aim to reduce both the communication delay and the load at the HA over 60% then we should have an fHA deployment between 45% and 65%. Finally, large deploy-

ments improve the communication's delays while low deployments improve the reduction of the load at the fHAs.

4. Related work

Incorporating route optimization to Mobile IP and NEMO clients is a key issue when considering the deployment of a truly mobile Internet. That's why this topic has attracted the attention of the research community and many solutions have been proposed.

First, the research community focused on solving this problem specifically for Mobile IPv4 [9] and NEMO clients [10–12]. The main idea behind these proposals is to deploy a new entity at the correspondent network that helps the MN to communicate directly with the CN. Usually this new entity authenticates the location (CoA) and the identity (HoA) of the MN. In addition, this device acts as a tunnel endpoint; this way the MN can send the packets tunnelled directly to the correspondent network. The main drawback of all these proposals is that they require deploying a new entity on each correspondent network. In the current Internet status, this would imply deploying a new entity on each network or at least, on each AS (currently there are roughly 22.000 ASes on the Internet). That's why we believe that the deployment cost of these solutions is too high.

As we mentioned in Section 1, Wakikawa presented recently a different approach [3] used by other researchers [4–6]. Since these proposals are not scalable [7,8], we propose using a P2P network that it is fully scalable and we benefit from the fHA that reduces the load at the HAs significantly.

5. Conclusions

The Mobile IP and NEMO protocols provide mobility for the Internet. Both protocols force the mobile nodes to send their data packets through a special entity (Home Agent) when communicating with their peers. This Home Agent is located at the mobile node's Home Network and forces the packet to follow a sub-optimal route. This

Table 5
Mean number of intradomain signalling messages/s per AS.

Deployment	Average number of (sent + received) messages (messages/s)
0.01	49.60 ± 0.03
0.10	45.96 ± 0.05
0.3	39.00 ± 0.09
0.6	32.57 ± 0.11
0.75	30.21 ± 0.12
0.9	29.24 ± 1.00

reduces considerably the communications' performance, increases the delay and the infrastructure load. The research community has focused on solving this issue deploying several Home Agents throughout the Internet. Then a mobile node may pick a closer one to its topological position in order to reduce the delay. Different authors use different technologies to signal the location if these Home Agents: eBGP, Anycast or a static list. Although this approach reduces the delay it is not Internet-scalable. Additionally, the Home Agents still have to forward all the mobile node's data packets and may become the bottleneck for the whole system.

In this paper, we have presented the fP2P–HN architecture that takes into account these issues. First, the architecture also deploys several Home Agents in order to reduce the delay. Second, it uses a P2P network to signal the location of these Home Agents in an scalable way. Third, the Home Agents of the architecture are in fact flexible Home Agents. These agents signal the location of the mobile nodes within a network using the IBGP. This way the network's exit routers are aware of the location of the mobile nodes an can forward the packets by themselves, thus, the load at the flexible Home Agent is significantly reduced.

It is reasonable to consider a Community Network as a NEMO in a mobile environment. Therefore, the proposed solution has a clear application in Mobile Community Networks, specifically by reducing the delay of the communications of such networks and the infrastructure load.

We have implemented the fP2P–HN in a simulator and we have evaluated the benefits and the costs of the architecture. The benefits are two: reduction of the delay and of the load at the Home Agents. The costs are the extra Inter- and Intra-domain signalling messages. We have put special attention on evaluating the Inter-domain overload since this cost must be scalable. In order to provide significant results, we have simulated the architecture using large Internet-like topologies of 3500 autonomous systems and a mean number of 100 mobile nodes per Home Agent. Additionally, each simulation has been repeated 50 times, using a different Internet-like topology, this way we can provide confidence values of the results. We tested different scenarios of deployment of the architecture, from 0.01 flexible Home Agents per Autonomous System to 0.9.

The main conclusions that can be extracted from the results are:

- The fP2P–HN effectively reduces the delay of the mobile nodes compared to Mobile IP/NEMO. Even with low deployments (0.1) the reduction is 23%. As the deployment grows so does the reduction that can be up to 83% (0.9).
- Our architecture reduces the traffic processed by each flexible Home Agent compared to that of Mobile IP/NEMO. As expected, the reduction of the traffic decreases as the deployment increases. In the worst case, the reduction of the traffic processed by a flexible Home Agent is 54% (0.9). This reduction grows further to 99% (0.01).

- Our architecture is highly scalable since the amount of Inter-Domain signalling is within the same order of magnitude (hundreds) and irrespective of the number of flexible Home Agents deployed, thus, the cost is $O(1)$. Additionally, the amount of Inter-Domain signalling traffic per flexible Home Agent is around 20 kbps.
- The extra Intra-Domain signalling of the fP2P–HN is very low, around 15 kbps per Autonomous System. Since the architecture allows that multiple flexible Home Agents are deployed within an Autonomous System this overload may be shared among several entities.

References

[1] T. Clauser et al, NEMO route optimisation problem statement, RFC 4888 (2004). October.
[3] R. Wakikawa et al., Virtual mobility control domain for enhancements of mobility protocols, IEEE INFOCOM, 2006.
[4] Y.S. Yet et al., Global dynamic home agent discovery on mobile IPv6, Wireless Communications and Mobile Computing (August) (2006).
[5] Boeing Connexion Service. <http://www.connexionbyboeing.com>.
[6] Marcelo Bagnulo et al., Scalable Support for Globally Moving Networks, ISWCS (2006).
[7] G. Huston, Commentary on inter-domain routing in the internet, RFC 3221 (December) (2001).
[8] Katabi Dina et al., A framework for scalable global IP-Anycast (GIA), SIGCOMM (2000).
[9] Chun-Hsin Wu et al., Bi-directional route optimization in mobile IP over wireless LAN, Vehicular Technology Conference, September, 2002.
[10] C. Ng et al., Network Mobility Route Optimization Problem Statement, RFC 4888 (2007).
[11] M. Calderon et al., Design and experimental evaluation of a route optimization solution for NEMO, IEEE JSAC (2007).
[13] K. Lua et al, A survey and comparison of peer-to-peer overlay network schemes, IEEE Communications Surveys and Tutorials (2005).
[14] I. Stoica et al., Chord: a scalable peer-to-peer lookup service for internet applications, ACM SIGCOMM (2001).
[15] S. Pal Chaudhurri et al., Perfect simulations for random trip mobility models, 38th Simulation Symposium, 2005.
[16] N. Brownlee et al, Understanding Internet traffic streams: dragonflies and tortoises, IEEE Communications Magazine (2002).
[17] R. Cuevas, C. Guerrero, A. Cuevas, M. Caldern, C.J. Bernardos, P2P based architecture for global home agent dynamic discovery in IP mobility, 65th IEEE Vehicular Technology Conference, 2007.
[18] Albert Cabellos-Aparicio, Jordi Domingo-Pascual, A flexible and distributed home agent architecture for mobile IPv6-based networks, IFIP Networking (2007).
[19] Mobility EXTensions for IPv6 (mext), <http://www.ietf.org/html.charters/mext-charter.html>.
[20] S. Jamin et al., On the placement of internet instrumentation, Proceedings of IEEE INFOCOM, 2000.
[21] Internet Topology Generator. <http://topology.eecs.umich.edu/inet/>.
[22] Passive Measurement and Analysis (PMA). <http://pma.nlanr.net>.
[23] Amit Jardosh, M. Elizabeth, Belding-Royer, C. Kevin Almeroth, Subhash Suri, Towards realistic mobilitymodels for mobile ad-hoc networks, Proceedings of the 9th Annual International Conference on Mobile Computing and Networking, 2003.
[24] N1 Grid Engine. <www.sun.com/software/gridware/>.

[25] R. Housley et al., Internet X.509 public key infrastructure certificate and CRL profile, RFC 2459 (January) (1999).
[26] European Telecommunications Standards Institute, GSM 03.20: Security Related Network Functions, June 1993.
[27] European Telecommunications Standards Institute, GSM 02.09: Security Aspects, June 1993.
[28] European Telecommunications Standards Institute, TS 133 102: Universal Mobile Telecommunications System (UMTS); 3G Security; Security Architecture, version 3.6.0, October 2000.
[29] C. Perkins, IP Mobility support for IPv4, RFC 3344 (2002).
[30] D. Johnson et al., Mobility support in IPv6, RFC 3775 (2004).
[31] Y. Rekhter et al., A border gateway protocol 4 (BGP-4), RFC 1771 (1995).
[32] V. Devarapalli et al., Network mobility (NEMO) basic support protocol, RFC 3963 (2005).
[33] The Perl directory<www.perl.org>.
[34] Cormen, Thomas H. Leiserson, Charles E. Rivest, L. Ronald, Introduction to Algorithms, MIT Press and McGraw-Hill, 1990, ISBN 0-262-03141-8.

Ruben Cuevas Rumin got his M.Sc. in Telecommunication Engineering and M.Sc. in Telematic Engineering at Universidad Carlos III de Madrid in 2005 and 2007, respectively. Furthermore, he obtained his M.Sc. in Network Planning and Management at Aalborg University in 2006. Currently he is Ph.D. Candidate at the Department of Telematic Engineering at University Carlos III de Madrid. His research interests include Overlay and P2P Networks and Content Distribution.

Albert Cabellos-Aparicio received a B.S. (2001) and M.S. (2005) degree in Computer Science Engineering from the Technical University of Catalonia (www.upc.edu). In 2002, he joined the Advanced Broadband Communications Center (CCABA, http://www.ccaba.upc.edu) where he worked as Research Assistant. In 2004, he was awarded with a full scholarship to carry out Ph.D. studies at the Department of Computer Architecture, Technical University of Catalonia (UPC), Spain. Scholarship granted by University through the Spanish Ministry of Science and Technology. In september 2005, he became an Assistant Professor of the Computer Architecture Department. His main research interests are IP mobility, IPv6, active and passive measurements and emulated testbeds for network measurements. He has participated on several European and National funded Research Projects.

Angel Cuevas Rumin got his M.Sc. in Telecommunication Engineering and M.Sc. in Telematic Engineering at Universidad Carlos III de Madrid in 2006 and 2007, respectively. He got an Erasmus Scholarship and complete his Master Thesis at The University of Reading. Currently he is Ph.D. Candidate at the Department of Telematic Engineering at University Carlos III de Madrid. Also, he got a research Internship at SAP Labs France. His research interests include Overlay and P2P Networks and Wireless Sensor Networks.

Jordi Domingo-Pascual is Full Professor of Computer Science and Communications at the Technical University of Catalunya (UPC) in Barcelona. He is co-founder and researcher of the Advanced Broadband Communications Center of the University (CCABA). His research topics are Broadband Communications and Applications, QoS management and provision, traffic engineering, inter-domain QoS routing and TE, IP traffic analysis and characterization, QoS measurement and measurement.

Arturo Azcorra received his Telecommunication Engineering degree from the Universidad Politecnica de Madrid in 1986 and the Doctor degree in 1989 from the same University. In 1992, he graduated as Master in Business Administration from Instituto de Empresa, one of the most prestigious Business School in the world.

He is a Full Professor at Universidad Carlos III de Madrid since 1999. He is an IEEE Senior Member, an ACM-SIGCOMM member and an Internet Society member. On November 2006, he was appointed as Director of the International Research Institute IMDEA Networks.He has participated in a large number of European Research Projects in the frame of ESPRIT, RACE, COMETT, Telematics, ACTS and IST Programs, having coordinated the European Networks of Excellence E-NEXT and CONTENT, and is currently the coordinator of the European Project CARMEN. He has published over 100 scientific papers in prestigious international magazines and conferences.