

A Hierarchical Mobile IPv6 Proposal

Claude Castelluccia

N° 0226

November 1998

THÈME 1



*Rapport
technique*

A Hierarchical Mobile IPv6 Proposal

Claude Castelluccia*

Thème 1 — Réseaux et systèmes
Projet SIRAC

Rapport technique n ° 0226 — November 1998 — 25 pages

Abstract:

The IETF Mobile IPv6 protocol has been developed to manage global (macro) mobility. It is not adapted to local (micro) mobility since it does not support any kind of hierarchy.

This report presents a hierarchical protocol, built on top of Mobile IPv6, that separates local mobility (within a site) from global mobility (across sites) management. Local handoffs are managed locally and transparently to a mobile node's correspondent hosts while global mobility is managed with Mobile IPv6. Our scheme is flexible (several levels can be used), scalable, interworks with Mobile IPv6 and can be deployed gradually.

Key-words: Mobility, Mobile IP, IPv6, Internet

(Résumé : tsvp)

* Email : claude.castelluccia@inrialpes.fr

Mobile IPv6 Hiérarchique

Résumé :

Le protocole Mobile IPv6 de l'IETF permet de gérer la mobilité globale des utilisateurs dans l'Internet. Ce protocole est mal adapté à la gestion de la micro-mobilité car génère une signalisation importante. Ce rapport présente un protocole de gestion de la mobilité hiérarchique, basé sur Mobile IPv6, pour l'Internet de demain. Ce protocole sépare la gestion de la mobilité locale (à l'intérieur d'un site) de celle de la mobilité globale (entre sites). La mobilité locale est gérée localement alors que la mobilité globale est gérée par Mobile IPv6. Notre approche est très flexible, compatible avec la solution Mobile IPv6 de l'IETF et peut être déployée dans l'Internet progressivement.

Mots-clé : Mobilité, Mobile IP, IPv6, Internet

1 Introduction

Internet Mobile users require special support to maintain connectivity as they change their point-of-attachment. This support should provide performance transparency to mobile users and should be scalable. Providing performance transparency means that higher level protocols should be unaffected by the addition of mobility support. Issues that may affect performance transparency are optimum routing of packets to and from mobile nodes and efficient network transition procedures [5]. The mobility support should be scalable in the sense that it should keep providing good performance to mobile users and should keep the network load low as the network grows and the number of mobile node increases. This scalability issue is a very important one in the context of a still growing worldwide network such as the Internet. The IETF Mobile IPv6 proposal, which provides a mobility management scheme for the Internet, does not completely meet these design goals. Whereas it provides performance transparency, we argue that Mobile IPv6 is not scalable. In Mobile IPv6, a mobile node sends a location update to each of its correspondent nodes periodically and at any time it changes its point-of-attachment. The resulting signaling and processing load may become very significant as the number of mobile nodes increases. This limitation is the result of the lack of hierarchy in the mobility management procedures of Mobile IPv6. In fact, Mobile IPv6 handles global area mobility and local area mobility identically. Since 69% of a user's mobility is local [4]¹, we believe that a hierarchical scheme that separates micro-mobility from macro-mobility is preferable.

In this paper, we present a hierarchical mobility management architecture for IPv6. The proposed scheme, which is fully compatible with the IETF solution, differentiates global (inter-site) mobility management from local (intra-site) mobility management. Correspondent hosts are only aware of inter-site moves of mobile hosts. Local mobility, i.e. within a site, is managed locally and transparently to the site's external hosts.

This paper is structured as follows. Section 2 presents Mobile IPv6 very briefly. Section 3 describes our hierarchical mobility management proposal. It first provides an overview of the proposed scheme and then goes into more detail of the protocol. Section 4 compares and analyses the performance of Mobile IPv6 and our proposal. Routing, transition and scalability performances are considered. Section 5 presents the related work. Section 6 concludes the paper.

2 The IETF Mobile IPv6

The Mobile IPv6 protocol is currently being specified by the IETF IP Routing for Wireless/Mobile hosts working group [3]. With Mobile IPv6, each time the mobile host (MH) moves from one subnet to another, it gets a new *care-of address* (CoA). It then registers its *Binding* (association between a mobile node's home address and its care-of address) with a router in its home subnet, requesting this router to act as the *home agent* (HA) for the mobile host. This router records this binding in its Binding Cache. At this point, the HA serves as a proxy for the MH until the MH's binding entry expires. The HA intercepts any packets addressed to the MH's home address and tunnels them to the MH's care-of address using IPv6 encapsulation. The MH sends also a Binding Update (BU) to

¹This study examined the mobility patterns of professionals regardless of whether they were equipped with portable devices or not.

its correspondent hosts (CHs), which can then send packets directly to the MH. While this protocol optimizes the routing of packets to MHs, it is not scalable. As the number of MHs increases in the Internet, the number of BUs increases proportionally and adds a significant extra load to the network.

3 A Hierarchical Mobility Management Architecture

Mobile IPv6 handles local mobility of a host (i.e. within a site or a network) in the same way as it handles global mobility (inter-site or inter-network mobility). In Mobile IP, a mobile host sends binding updates to its home agent and its correspondent nodes each time it changes its point-of-attachment regardless of the locality and amplitude of its movement. As a consequence, the same level of signaling load is introduced in the Internet independently of the user's mobility pattern. We argue that this approach is not scalable since the generated signaling load can become quite overwhelming as the number of mobile hosts increases in the Internet.

We believe that a hierarchical scheme that differentiates local mobility from global mobility is more appropriate to the Internet. Using such a hierarchical approach has at least two advantages. First, it improves handoff performance, since local handoffs are performed locally. This increases the handoff speed and minimizes the loss of packets that may occur during transitions. Second, it significantly reduces the mobility management signaling load on the Internet since the signaling messages corresponding to local moves do not cross the whole Internet but stay confined to the site. This hierarchy is furthermore motivated by the significant geographic locality in user mobility patterns. According to the study presented in [4], 69% of a user's mobility is within its home site (within its building and campus). We believe that this result can be extrapolated by stating that most of a user's mobility is local i.e. within its home site or the foreign site it is visiting. It is therefore important to design a mobility management architecture that optimizes local mobility.

We propose a hierarchical architecture that separates local mobility (within a site) from global mobility.

3.1 Protocol Overview

Our proposal differentiates the intra-site mobility from the inter-site mobility. As a result, a host communicating with a mobile host is only aware of its inter-site mobility. The mobile host's intra-site mobility is completely hidden. In this paper, we define a *site* as the the highest level of our hierarchical architecture. A site is actually an arbitrary structure. It can be an ISP network, a campus network, a company network, a set of LANs or even a single LAN. A site is connected to the rest of the Internet via one or several interconnection routers that we call *Border Routers* (BR) in the rest of this document.

Our proposal is based on the deployment of *Mobility Networks* (MN). A MN of a site is a LAN that defines an address space for the mobile hosts roaming within this site. A Mobility Network contains one or several *Mobility Agents*. A Mobility Agent ² (MA) is a router of the Mobility Network that maintains a binding per mobile hosts currently visiting the site. Note that there is no

²The Mobility Agent concept is very similar to the Home Agent concept.




Figure 1: Inter-Site Mobility




Figure 2: Intra-site Mobility

constraint on the physical location of the Mobility Network. However for efficiency reasons, it is preferable to connect it to the *border router* of the network that it is serving. The mobility Network can actually be any sub-network of the site. It does not have to be dedicated to mobile hosts but instead can support ordinary (fixed) hosts.

Deploying a Mobility Agent in a separate Mobility Network instead of implementing it on the BR has two main advantages. First, it does not require any modification to the routers and is therefore easier to deploy. Second, it is more scalable since (1) it does not add additional processing constraints on the BR and (2), as we will describe in Section 3.3, several MAs could be deployed for scalability and/or robustness motivations. However the MA can be implemented within the BR if this is desirable.

The main operations of the proposed protocol are the following (all abbreviations used in this description are recalled in table 1 for clarity.):

- *Inter-site mobility*: When a mobile host enters into a new site, it gets two CoAs: a Private (or Physical) Care-of Address (PCoA), which is a CoA on the link it is attached to, and a Virtual Care-of Address (VCoA), which is a CoA in the Mobility Network of the site (note that in Mobile IPv6 only the PCoA is required.).

The mobile host then sends some BUs. It sends:

- a BU³ that specifies the binding between its *VCoA* and its *PCoA* to the site MA. Upon reception of this BU, the MA performs admission control such as authentication and charging. If the request is accepted, an acknowledgement is sent back to the MH. The issues of authentication and billing are beyond the scope of this report.
- a BU that specifies the binding between its *home address* and its *VCoA* to its HA and each of its external CHs (i.e. CHs that are outside of the site).
- a BU that specifies the binding between its home address and its *PCoA* to each of its local CHs (i.e. CHs that are within the site).

As a result,

- An external host that sends packets to the mobile host uses its *VCoA*. Packets are then routed to the Mobility Network of the visited site, intercepted by the Mobility Agent and forwarded (tunneled) to the current *PCoA* of the MH.
 - A local host that sends packets to the mobile host uses its *PCoA*. Packets are then directly delivered to the mobile host.
- *Intra-site mobility*: When a mobile host moves within the site, it gets a new *PCoA* on its new point-of attachment. The *VCoA* remains constant as long as the mobile host is roaming locally. The mobile host then sends the following BU:
 - a BU that specifies the binding between its home address and its new *PCoA* to each of its local CHs (i.e. CHs that are within the site).

³The 'Acknowledge' bit ('A' bit) must be set.

<i>MH</i>	Mobile Host
<i>CH</i>	Correspondent Host
<i>MA</i>	Mobility Agent
<i>BU</i>	Binding Update
<i>CoA</i>	Care of Address
<i>VCoA</i>	Virtual Care of Address
<i>PCoA</i>	Physical Care of Address
<i>MN</i>	Mobility Network

Table 1: Abbreviations

- a BU that specifies the binding between its *VCoA* and its new *PCoA* to the site Mobility Agent.

Note that during intra-site mobility, no BU is sent on the Internet and that transitions are performed locally. Figures 1 and 2 illustrate the Inter-site and Intra-site mobility operations.

3.2 Protocol Details

Our proposal can actually be extended to use several levels of hierarchy per site. In fact, if necessary, a site can be split into several levels of hierarchy which may themselves be sub-divided and so on. A *Mobility Network* is then deployed in each level of the hierarchy. These Mobility Networks are structured following the same hierarchy of the site, i.e. the Mobility Network that manages the whole site is at the top of the tree, followed by the Mobility Networks of the lower hierarchy level and so on until the lowest levels.

A mobile host roaming within a Site gets a *VCoA* in each MN of the MNs' tree from its current point-of attachment to the top MN. Each MA of these MNs keeps a binding between the *VCoA* of the MH in its MN and the MH's *VCoA* in the next-lower MN. When an external CH sends a packet to the MH, it uses its *VCoA* in the top MN, i.e. *VCoA₁*. The packet is then delivered to the top MN, intercepted by the MA, encapsulated to the MH's *VCoA₂*. The MA of the second node MN intercepts the packet, decapsulates it to the MH's *VCoA₃*...

In the example of Figure 3, the site is decomposed into 2 sub-sites S_2 and S_3 . MN_1 manages local mobility of hosts between S_2 and S_3 . MN_2 manages local mobility within S_2 while MN_3 manages local mobility within S_3 .

In the rest of this section, we detail the proposed mobility management protocol with several levels of hierarchy. This description is divided into three parts: (1) the *registration*, (2) the *MN and Mobility Agents Discovery*, and (3) the *packet delivery phases*.




Figure 3: Site Hierarchy

3.2.1 Registration Phase

In our proposal, a MH gets several (V)CoAs (instead of one single CoA as in Mobile IP) and registers each of them with its mobility agents, and possibly with its CHs and HA⁴. This registration phase differs in local (intra-site) and global (inter-site) mobility.

- **Intra-site Mobility**

When the mobile host moves locally (i.e. within the site), it needs to find out the lowest MN in the branch from its current location to the top MN that has changed. This is performed by comparing each MN of the branch connecting the top MN to its previous point-of attachment with the MNs advertised in the Mobility Agent Information Option of the router advertisements. If l is the rank of the lowest node (the rank of the top MN is one) and N the number of MNs on the new branch, the Mobile host performs the following operations:

- it gets a new VCoA in each MN from MN_l to MN_N (we note MN_i , the MN of rank i in the branch from the top MN to the mobile host's point-of attachment with MN_1 being the top MN),
- it gets a new PCoA on the link,

⁴As with the regular Mobile IPv6, a mobile host requires the service of a home agent in its home network. This HA intercepts packets addressed to the MH and forwards them toward the MH's current $VCoA_1$.

- it registers the $(VCoA_{i-1}, VCoA_i)$ binding with MA_{i-1} , for i going from l to N (we note MA_j the mobility agent of MN_j and $VCoA_j$ the (Virtual) Care-of Address of the MH in MN_j),
- it registers the $(VCoA_N, PCoA)$ binding with MA_N ,
- it registers the $(HomeAddress, PCoA)$ binding with its local CHs⁵.

The Mobility Agents must, as the Home Agent in Mobile IPv6, acknowledge the reception of the Bindings coming from the MH. Consequently, the BUs sent by the MHs to the MAs must have the 'acknowledge' bit set to 1.

• Inter-site Mobility

When the mobile host moves globally (i.e. it enters into a new site), the Mobile host performs the following operations:

- it gets a new VCoA in each MN from MN_1 to MN_N ,
- it gets a new PCoA,
- it registers the $(VCoA_{i-1}, VCoA_i)$ binding with MA_{i-1} for i going from 1 to N ,
- it registers the $(VCoA_N, PCoA)$ binding with MA_N ,
- it registers the $(HomeAddress, PCoA)$ binding with its local CHs (i.e. within the site),
- it registers the $(HomeAddress, VCoA_1)$ binding with its distant CHs (i.e. external to the site) and Home Agent.

Note that Binding Updates are only sent outside of the site (to the Home Agent and distant Correspondent Hosts), when the mobile host moves from one site to another. As a result, the local signaling load (i.e. within the site) is reduced since BUs are only sent locally when a MH is roaming within a site.

3.2.2 Mobility Networks and Mobility Agents Discovery

To perform the previous registration operations, a mobile host gets the following information:

- *the prefix of the site* (this information is used by the mobile host to define the site boundary),
- *the depth of the hierarchy* i.e. the number of Mobility Networks on the branch from its current point-of attachment to the top MN,
- and for each MN on the branch to the top MN, *its network prefix and the IP address of the mobility agent.*

⁵The MH uses the *Site Prefix* field in the new Mobility Information Option to differentiate the local CHs from the distant ones (i.e. within the site)

This information is advertised by a new option used in the Router Advertisement messages of the IPv6 Neighbor Discovery [6]. The format of this option, called the *Mobility Information Option*, is presented in Appendix A.

A mechanism similar to the *dynamic home agent address Discovery* mechanism of Mobile IPv6 could be defined instead. In this case, the Mobile Host would send a Binding request to the anycast address of the MN and get back the address of the mobility agent.

3.2.3 Packet Delivery

When a distant correspondent host sends a packet to a mobile host, it uses its $VCoA_1$. Packets are then delivered to the MN of the level 1 hierarchy, intercepted by the Mobile host mobility server and encapsulated to the MH's $VCoA_2$. The mobility agent of the MH in the level 2's MN intercepts the packet, decapsulates it and encapsulates to $VCoA_3$. The packet is then forwarded down until the current PCoA of the mobile host⁶.

When a local CH sends a packets to a MH, it uses its $PCoA$. Packets are directly delivered to the mobile host.

When sending a packet, a mobile host sets the source field of the IP header to its $PCoA$ regardless whether its correspondent host is local, site-local or distant and includes an *Home Address Option* (as in Mobile IP) specifying its Home Address. The use of the $VCoA$ is avoided to bypass ingress filtering.

3.3 Deploying Several Mobility Agents per MN

The problem with hierarchical schemes [7, 2] is that they usually use a tree-based structure. In these proposals, the mobility agent of the site must keep one entry per mobile host roaming locally. We believe that this structure is not scalable and that this mobility agent can become a performance bottleneck as the site grows and/or the number of mobile hosts increases.

In our proposal, several MAs can be deployed in a MN transparently to the CHs or the higher MNs in the tree hierarchy. When a packet addressed to a MH's VCoA gets to the MN, the packet is intercepted by the MH's MA. The actual MA identity is not revealed to the source of the packet. As a result, MAs can dynamically be duplicated or exchanged transparently to the CHs. An administrator wishing to reduce the MA processing load of a MN can also deploy several MAs in this MN. Each of these MAs would then be in charge of some of the lower networks in the MN hierarchy based, for example, on a geographical partitioning of the site. These MAs would then be advertised through the new Mobility Information Option in the lower networks...

The duplication of MAs is very useful to share the load at the mobility agents (BU processing, packet forwarding and bindings' storage). This technique is also useful to improve the robustness of the system (if one mobility server fails, only one part of the site will become unreachable).

⁶Note that instead of encapsulating and decapsulating packets, mobility agents (except for the first one) can merely change the source and destination IP addresses of the encapsulating IP header.

3.4 The Multiple Border Routers Case

We propose in our scheme to deploy a Mobility Network per site and connect it directly (if possible) to the Border Router. If a site is connected to the Internet through several Border Routers then several Mobility Networks should be deployed otherwise the routing of the packets reaching the site through a border router that is far from the mobility network would be suboptimal. Figure 4 illustrates this problem. All packets that reach the site through BR_2 are first routed to MN and then redirected to the MH's current PCoA. We suggest the following algorithm:

- The site deploys one Mobility Network per Border Router.
- Each of these MNs, MN_y is defined by two network prefixes: P and P_y . The prefix P is common to all MNs and P_y is specific to each of them.
- A mobile host roaming within the site configures its $VCoA$ using the prefix P and registers its Binding ($PCoA, VCoA$) with a Mobility agent of each mobility network (the addresses of the mobility agents, which prefix is P_y , are obtained via the Mobility Information Option which has to be extended for this purpose. See Appendix B for more detail about this new option.).
- Each border router is configured to route packets with a destination address belonging to the sub-network defined by the prefix P to the MN directly attached.

As a result of this algorithm, when a packet addressed to a mobile host with address $VCoA$ reaches a border router of the site, it is routed to the closest Mobility Network, intercepted by the mobility agent and forwarded to the current $PCoA$ of the mobile host (see Figure 5).

Note that if a site contains several levels of hierarchy which each has several border routers, a mobile host roaming in the site must register a Binding Update with border routers of each hierarchy level from its current point of attachment to the hierarchy top level. This has two consequences for highly connected sites that contain many levels of hierarchy: (1) the size of the router advertisement messages which contains the list of MAs the mobile host must register with can be large (the size of an Extended Information Option, S , is $(28 + (12 + 16*n)*m)$ bytes where n is the number of BR per level and n the number of levels. If $n=2$ and $m=2$ then $S=116$ bytes, if $n=4$ and $m=10$ then $S=788$ bytes), and (2) the local (i.e. within the site) signaling load, generated by the emission of the BUs can be significant. As a result, we propose to make this multi-BR registration extension optional for these large sites. Routers only advertise one MA per level using the regular Mobility Information Option. However if necessary a mobile host may obtain the complete list of MAs by sending a solicitation message to the local router. Upon reception of this solicitation message, the router returns a router advertisement with an Extended Mobility Information option containing the list of all MAs the MH should register with.

4 Comparison and Evaluation

In this section we compare the performance of our proposal and Mobile IPv6. When comparing the performance of different mobility management schemes, several factors have to be taken into




Figure 4: The Multiple Border Routers Case: the problem

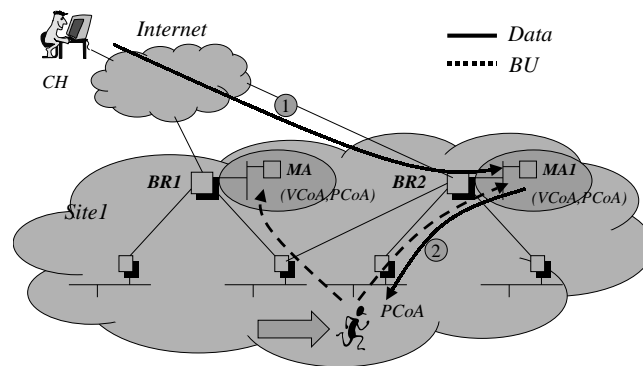


Figure 5: The Multiple Border Routers Case: the solution

consideration. Among these factors, three are particularly important [5]: (1) The routing performance of the schemes, i.e. what is the extra latency introduced by each of the schemes. (2) The transition performance of the schemes, i.e. how fast are the transition phases performed. (3) The scalability property of the schemes, i.e. how do the schemes behave as the network grows and the number of mobile hosts increases.

4.1 Routing and Transition Performance

The routing and transition performances of both schemes are quite similar.

With mobile IP, the routing is optimum, i.e. packets follow the shortest path from the CHs to the MH, except for the first packets which have to go through the mobile host's home agent. With our hierarchical Mobile IP, an extra indirection through the MA is required. We believe that the cost of this indirection is small especially if the mobility agent is close to the border router as suggested.

Handoffs are performed locally in both proposals. In our proposal, local handoffs are managed within the site. In Mobile IPv6, while location updates have to cross the whole Internet to reach the mobile host correspondent nodes, a mechanism is provided to smooth out transitions. After switching to a new default router, a mobile node may send a Binding Update to its previous default router, asking him to redirect all incoming packets to its new Care-of Address.

4.2 Scalability Performance

The main performance difference between the compared approaches resides in their scalability property. The scalability property of a protocol can be evaluated in terms of its overhead growth on the Internet with the size of the Internet, the number of mobile hosts and the number of correspondent nodes.

One of the most important criteria that affects the scalability property of a mobility management scheme is its signaling load, i.e. the bandwidth used by the control messages, such as the Binding Updates, to support mobility.

In this section, we compare the signaling load of Mobile IPv6 with the signaling load introduced with our proposal on the Internet backbone (we do not consider the local signaling load since they are comparable for both schemes and we argue that local resource is not the most critical).

We evaluate, for each of these schemes, the aggregated signaling load bandwidth consumed on the Internet. This aggregated bandwidth is independent of the number of nodes that the Binding Updates have to cross until their destinations, but rather corresponds to the signaling bandwidth on one link. In this evaluation, we differentiate three types of mobility: (1) local mobility of a host within its home site, (2) local mobility of a host within a foreign site, and (3) inter-site mobility of a host. We then evaluate the average signaling load over these three mobility patterns.

Binding Update Emission Frequency

The signaling load of a scheme depends directly on the Binding Update Emission Frequency. According to [3], a mobile host sends a Binding Update to:




Figure 6: f_{HA} and f_{CH} for $\#CH = 2$ and $f_{REF} = 1/60$

- its *Home Agent*, each time it changes its point-of attachment (the HA must acknowledge this BU). We denote f_{HA} the Binding Update emission frequency from the mobile host to its Home Agent.
- each of its *correspondent hosts*, each time it changes its point-of attachment and then periodically to refresh the corresponding cache entries. After sending M consecutive Binding Updates at a frequency of f_B to a particular node with the same care-of address, the mobile node should reduce its frequency of sending Binding Updates to that node to f_{REF} . We denote f_{CH} the average Binding Update emission frequency from the mobile host to its Correspondent Hosts.

The emission frequencies of a Binding Update, f_{HA} and f_{CH} , are dependent on the mobility frequency of a host, f_{MOV} , and the refresh frequencies f_{REF} and f_B . They are defined as follows:

$$f_{HA} = \begin{cases} (\lceil f_{REF}/f_{MOV} \rceil + 1) \times f_{MOV} & \text{if } f_{REF} > f_{MOV} \\ 2 \times f_{MOV} & \text{if } f_{MOV} \geq f_{REF} \end{cases}$$

$$f_{CH} = \begin{cases} (\lceil f_{REF}/f_{MOV} \rceil + (M - 1)) \times f_{MOV} & \text{if } f_{REF} > f_{MOV} \\ M \times f_{MOV} & \text{if } 1/M \times f_B \geq f_{MOV} \geq f_{REF} \\ (\lceil f_B/f_{MOV} \rceil) \times f_{MOV} & \text{if } (f_{MOV} \geq 1/M \times f_B \geq f_{REF}) \end{cases}$$

Figure 6 displays f_{HA} and f_{CH} as a function of f_{MOV} with $f_B = 1$, $f_{REF} = 1/60$ and $M = 5$.

Local Mobility within the Home Site

When a mobile host, using Mobile IPv6, is moving within its home site, it sends a Binding Update to each of its external correspondent nodes at a frequency of f_{CH} . If our hierarchical proposal is used, two cases are possible:

1. *the Virtual Care-of Address of the MH is advertised in the Domain Name Server (instead of its home address).* As a result, no binding has to be sent on the Internet as long as the mobile home is roaming within its home site.
2. *the home address of the mobile host is advertised in the DNS (as in Mobile IP).* As a result, the mobile host has to send a Binding Update to each of its external correspondent nodes at a frequency of f_{REF} .

We recommend using the first solution since it is more scalable and has the nice property of hiding mobility of users that are roaming within their home site. We consider this solution in the rest of this analysis.

The signaling bandwidth respectively generated by Mobile IP on the Internet, $BW_{SIG_MIP,home}$ and by our proposal, $BW_{SIG_HMIP,home}$, when a MH is roaming within its home site, are defined as follows:

$$BW_{SIG_MIP,home} = Size_{BU} \times f_{CH} \times \#CH$$

$$BW_{SIG_HMIP,home} = 0.$$

where $Size_{BU}$ ⁷ is the size of a Binding Update and $\#CH$ is the number of correspondent hosts that are not in the home site.

Local Mobility within a Foreign Site

When a mobile host, using Mobile IPv6, is moving within a foreign site, it sends a Binding Update to each of its correspondent nodes and to its home agent at a frequency equal to f_{CH} and f_{HA} . If our proposal is used, the mobile host only sends a Binding Update to each of its correspondent nodes and to its home agent at a frequency respectively equal to the refresh frequency, f_{reqREF} .

As a result, $BW_{SIG_MIP,foreign}$ and $BW_{SIG_HMIP,foreign}$ are defined as follows:

$$BW_{SIG_MIP,foreign} = Size_{BU} \times (f_{CH} \times (\#CH + 1) + f_{HA})$$

$$BW_{SIG_HMIP,foreign} = Size_{BU} \times f_{REF} \times (\#CH + 1)$$

⁷The size of a BU is equal to the size of an IPv6 header (40 bytes) + the size of a Binding Update Extension Header (28 bytes), so 68 bytes. A Binding Update can sometimes be appended to an outgoing packet. The size of the BU is then reduced to the size of a Binding Update Extension Header.

G_{home}	1.0
$G_{foreign}$	$(\#CH \times (f_{CH} - f_{REF}) + (f_{HA} - f_{REF})) / (f_{CH} \times \#CH + f_{HA})$
$G_{transit}$	0
G_{AV}	$0.69 + 0.31\alpha \times (\#CH \times (f_{CH} - f_{REF}) + (f_{HA} - f_{REF})) / (f_{CH} \times \#CH + f_{HA})$

Table 2: Gains of HMIP over MIP

Inter-Site Mobility

The signaling bandwidth introduced on the Internet when a mobile node is transiting from one site to another is the same in both schemes. The mobile host sends a Binding Update to its home agent (and receives an acknowledgment) and M Binding Updates to each of its external correspondent hosts. Therefore, $BW_{SIG,transit}$ is defined as follows:

$$BW_{SIG,transit} = Size_{BU} \times (M \times \#CH + 2)$$

where $\#CH$ is the number of external correspondent hosts of the mobile host.

Analysis of the Results

In this section, we evaluate, for each of the mobility patterns, the gain achieved by our proposal over Mobile IPv6. We note G_{home} the gain when the host is roaming within its home site, $G_{foreign}$ the gain when the host is roaming within a foreign site, and $G_{transit}$ the gain when the host is transiting from one site to another. G_Y (with Y = home or foreign), and $G_{transit}$ are defined as follows:

$$G_Y = (BW_{SIG_MIP,Y} - BW_{SIG_HMIP,Y}) / BW_{SIG_MIP,Y}$$

$$G_{transit} = (BW_{SIG_MIP,transit} - BW_{SIG_HMIP,transit}) / BW_{SIG_MIP,transit}$$

We also evaluate G_{AV} the pondered average of G_{home} , $G_{foreign}$ and $G_{transit}$. By making use of the results established in [4] that 69% of a host's mobility is local, G_{AV} is defined as follows:

$$G_{AV} = 0.69 \times G_{home} + 0.31 \times (\alpha \times G_{foreign} + \beta \times G_{transit})$$

where $\alpha + \beta = 1$, $\alpha = (N - 1)/N$ and $\beta = 1/N$, N being the average number of different points-of attachment that a mobile host gets within a site before moving to another site.

α and β characterizes the mobility pattern of a mobile host outside of its home site. α defines the intra-site versus inter-site mobility ratio of a mobile host. A large β means that the host is frequently changing sites. A large α means that the host is mainly roaming within a site and barely changes sites. For example, an α of 0.9 means that the mobile host changes, in average, 10 times its point-of attachment within a site before moving to another site.

The gains computed from the previous results are presented in the table 2. These results show that:

- The gain of our hierarchical Mobile IP over Mobile IP when a Mobile Host is roaming within its home site is 100% since our proposal does not send any BU over the Internet.
- The gain of our scheme is 0% during inter-site mobility. In fact, during inter-site mobility our proposal behaves exactly as Mobile IP.
- The gain of our proposal during local mobility within a foreign site is a function of $\#CH$, f_{CH} , f_{HA} and f_{REF} . Figure 8 plots the gain $G_{foreign}$ as f_{MOV} varies from 0 to 0.4 moves/second for several values of f_{REF} (1/600, 1/60 and 1/10). These plots show that:
 - the gain, $G_{foreign}$, gets larger as f_{reqMOV} increases (it actually converges to 100% as f_{MOV} gets close to 100%). Our proposal avoids the emission of BU on the Internet when the mobile host is roaming locally but does not avoid the emission of the refresh BUs sent periodically. As a result, when a mobile host is not moving frequently (f_{reqMOV} is low), most of the signaling load is generated by the refresh binding updates and the gain of our proposal is low. As f_{reqMOV} increases, the number of BUs generated by the MH's mobility increases and consequently the gain achieved by our solution gets larger.
 - the gain, $G_{foreign}$, is larger for smaller values of f_{REF} . The gain increases quickly for f_{MOV} between 0 and f_{REF} and then stabilizes around 1.0. This behavior is better explained by analyzing the signaling bandwidth generated by Mobile IP (see Figure 7). The bandwidth increases with a rate of $\lceil f_{REF}/f_{MOV} \rceil$ when f_{MOV} is smaller than f_{REF} . f_{REF} being fixed, this rate gets larger as f_{MOV} gets closer to 0. When f_{MOV} is larger than f_{REF} but smaller than $1/M \times f_B$ (0.2 packets/sec.), the bandwidth increases linearly with a rate of M (5 packets). Finally when f_{MOV} is larger than $1/M \times f_B$, the bandwidth is more or less constant and equals to f_B (1 packet/sec.).
- The average gain is always larger than 69% and converges to 1.0 as f_{MOV} gets larger. Figure 9 displays G_{AV} for several values of α (1.0, 0.5, 0.1). When $\alpha = 1.0$, the gain converges to 100% since there is no cost due to intra-site mobility. When $\alpha = 0.5$, the gain converges to 85% while when $\alpha = 0.1$, the gain converges to 72%. The average gain is larger for larger α . In fact for large α , the relative cost of the inter-site mobility is small compared to the gain achieved during local mobility.

5 Related Work

Caceres and al. have proposed a hierarchical mobility scheme based on Mobile IPv4 that separates three cases : local mobility, mobility within an administrative domain and global mobility [2]. This proposal has been made in the context of Mobile IPv4 which uses foreign agents; agents that mobile hosts connect to when they visit a foreign network. [2] defines a hierarchy of foreign agents. In this proposal, each subnet that a mobile node could visit has one or more subnet foreign agents, which manage local mobility. On top of those subnet foreign agents, a domain foreign agent manages mobility across the different subnets of an administrative domain. The mobile node's home agent




Figure 7: $BW_{SIG_MIP,foreign}$ and $BW_{SIG_HMIP,foreign}$ for $\#CH = 2$ and 10 and $f_{REF} = 1/60$





Figure 8: $G_{foreign}$ for different values of f_{reqREF}

Figure 9: G_{AV} for different values of α

only keeps track of the movement of the mobile node across administrative domain boundaries. As a result, the mobile node's motion within an administrative domain is transparent to the home agent and its correspondent nodes. Charles Perkins defines an architecture that uses a hierarchy of Foreign Agents to reduce the signaling load [7]. This proposal is very similar to the Caceres's one but the author goes into much more detail in the protocol specification. In this solution, FAs are arranged hierarchically, as a tree, in the site topology, and the mobile node is then allowed to move from one local area of the site topology to another one without requiring approval by or re-binding at the home agent (or correspondent hosts). A site is decomposed in sub-networks that may themselves be decomposed and so on. When a mobile node moves to a new point of attachment, it searches the lowest level of the hierarchy in the new list of FAs (this list is advertised by the lowest FA through Agent Advertisements), which has a different care-of address of its previous list of FAs, and then it notifies the foreign agent at the next-higher level of the hierarchy about the different care-of address. This is done by the new Registration Request message, called the Regional Registration Request message. This request is then forwarded to the next-higher level of hierarchy and a Registration Reply is returned to the MH. When the foreign agent receives the Request from the mobile node, it must pass the Request along to its next nearest ancestor in the hierarchy along the way to the agent listed as the Home Agent. In this way, each foreign agent in the hierarchy between the mobile node and the home agent will be able to maintain a binding for the mobile node. Similarly, Site Registration Replies are passed down from one level of the hierarchy to the next along the way to the mobile node, so that each foreign agent can determine the status of the corresponding mobile node. Packets arriving at the top of the hierarchy will be forwarded down to the current location of the mobile node.

Our proposal has a lot of similarities with the solutions described above. However our proposal takes advantage of the IPv6 new functionalities to provide a solution that is more robust, scalable and flexible. The Caceres's and Perkins's approaches use the agent advertisement at the lowest level to advertise the FA hierarchy to the mobile host. This imposes that a FA be deployed in each subnet that hosts mobile nodes. We believe this is a very strong design constraint. By using the Neighbor Discovery mechanisms, we do not impose any constraint on the location of the Mobility Agent. We argue that our proposal is:

- *easier to deploy* : The proposed solution can be implemented with the current Mobile IPv6 protocol without modifications. It only requires the definition of a new router advertisement option and some minor modifications at the mobile hosts.
- *more flexible* : A mobile host can decide to bypass some levels of hierarchy if appropriate. For example, a mobile host that does not move too frequently and/or wants to save bandwidth on the last hop (that may be wireless) by limiting the number of emitted BUs may only register with the top mobility agent. As a result, this mobile host will not suffer from the cost of the indirections and intermediary mobility agent processing.
- *more scalable* : Caceres's and Perkins's proposals impose that the FAs be arranged as a tree. The FA that is at the top of the tree must keep one entry for each mobile host in the region. This can become a problem as the number of mobile hosts increases. In contrast, in our proposal, several MAs can be deployed at any level of the hierarchy resulting in a distribution of the Mobility Agent processing load .

6 Conclusions and Future Work

This paper presents a solution that makes use of IPv6 new functionalities, such as a large address space and the Neighbor Discovery mechanism, to propose a mobility management scheme that is hierarchical, flexible and scalable. We propose a hierarchical architecture that separates local mobility (within a site) from global mobility (across sites). Local handoffs are managed locally and transparently to a mobile node's correspondent hosts. Our scheme reduces the signaling bandwidth on the backbone from 69% to 100% by hiding local mobility while still providing optimal routing and fast transition performance. A solution that hides local mobility to correspondent hosts provides several benefits. First, it reduces the signaling load since less Binding Updates are sent over the Internet. As a result, the global load on the Internet, the BUs' losses and consequently the mobile hosts' connectivity losses are reduced. Second, it improves partially mobility confidentiality since the correspondent hosts do not know the exact location of mobile hosts.

Our proposal is built on top of and is fully compatible with the IETF Mobile IPv6 protocol. It does not require installation everywhere to be useful but instead can be deployed gradually.

Our current work include:

- *Simulation*: We are currently working on the simulation of our proposal using NS [1] and are hoping to get some simulation results soon.

- *An Adaptative Hierarchical Mobile IPv6.* We are also studying an adaptive version of our hierarchical Mobile IPv6. As illustrated in this report, our proposal is mainly advantageous for fast moving hosts communicating with distant hosts. For these mobile hosts, our solution reduces the signaling load significantly, but also improves their handoff performance and reduces the Binding Updates' losses. However, these benefits do not come for free. There is a slight latency cost, due to the mobility agent indirection, and an additional complexity to the architecture to be paid. We believe that this cost is acceptable especially when a mobile host is moving quickly and communicating with a host which is far from its location. However for slow moving hosts and local communications, a hierarchical scheme is probably not necessary. An Adaptive Hierarchical Mobile IPv6 which uses different levels of hierarchies according to the mobility patterns and the needs of each host seems preferable.
- *Header Compression:* Within a site, the packets are encapsulated from the MA to the MH's current PCoA. This encapsulation has a cost in term of bandwidth utilization. We are studying how header compression techniques could be used to reduce this overhead.
- *DNS Dynamic Updates:* The signaling load reduction of our proposal could be further improved by using the DNS dynamic update mechanisms [8]. In fact, if the mobile host could dynamically update the address of its DNS entry with its current top-level *VCoA*, no mobility management protocol would be necessary for communications initiated and terminated while the mobile host is roaming within the same site. Updating the DNS entries with the top-level *VCoAs* seems reasonable since these *CoAs*, as opposed to the Mobile IPv6 *CoAs*, change very rarely. We are currently studying this extension.

7 Acknowledgement

The author is very grateful to Mary Baker, Thierry Ernst, Charles Perkins, Thierry Turetletti and Xinhua Zhao for their careful readings of earlier versions of this paper and insightful discussions.

References

- [1] UCB/LBNL/VINT Network Simulator - ns. In <http://www-mash.CS.Berkeley.EDU/ns/>, 1998.
- [2] Ramon Caceres and Venkata N. Padmanabhan. Fast and Scalable Handoffs for Wireless Inter-networks. In *Proc. 2st Annual International Conference on Mobile Computing and Networking*, 1996.
- [3] David B. Johnson and Charles E. Perkins. Mobility Support in IPv6. In *Internet Draft*, 1998.
- [4] G. Kirby. Locating the User. In *Communication International*, 1995.
- [5] Andrew Myles and David Skellern. Comparing Four IP Based Mobile Node Protocols. In *Proc. 4th Joint European Networking Conference*, 1993.

- [6] T. Narten, E. Nordmark, and W. Simpson. Neighbor Discovery for IP Version (IPv6). In *RFC 1970*, 1996.
- [7] C. Perkins. Mobile-IP Local Registration with Hierarchical Foreign Agents. In *Internet Draft*, 1996.
- [8] P. Vixie and al. Dynamic Updates in the Domain Name System (DNS UPDATE). In *RFC 2136*, 1997.

A Mobility Information Option Format

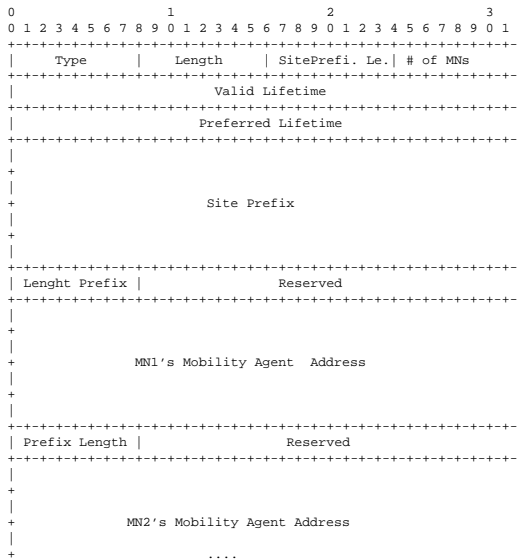


Figure 10: Mobility Information Option

The format of the new Mobility Information Option used by routers in their advertisement messages is displayed in Figure 10. This new option is used by the mobile hosts to get the information about the site hierarchy such as the active MNs and mobility agents.

The fields of this option are the following:

- *Type* : ?
- *Length* : length of the option in units of 8 octets = $3.5 + 2.5 * \text{number of levels}$

- *Site-Prefix Length*: number of leading bits of the *Site Prefix* field that defines the site. This information is used by the MHs to differentiate the local CHs from the distant ones.
- *# of MNs* : number of MNs in the hierarchy from the MH to the top MN.
- *Site Prefix* : prefix of the Site IP address. Used by the MHs to identify the site's boundary.
- *Prefix Length* : number of leading bits of the MN's Mobility Agent Address to define the MN prefix. This MN prefix is used by a MH to get a VCoA in the MN address space.
- *MN_i Mobility Agent Address*: The IP address of the *MA* of *MN_i*

B Extended Mobility Information Option Format




Figure 11: Extended Mobility Information Option

The format of the new Extended Mobility Information Option used by routers in their advertisement messages is displayed in Figure 11. In addition to the information contained in the Mobility Information Option it also contains, for each level the hierarchy, the common Prefix, to be used by MHs to configure their VCoA and the list of MAs of each border router.

The fields of this option are the following:

- *Type* : ?
- *Length* : length of the option in units of 8 octets = $3.5 + (1.5 + 2 * \#MAs) * \#levels$
- *Site-Prefix Length*: number of the leading bits of the *Site Prefix* field that defines the site. This information is used by the MHs to differentiate the local CHs from the distant ones.
- *# of levels* : number of levels in the hierarchy from the MH to the top MN.
- *Site Prefix* : prefix of the Site IP address. Used by the MHs to identify the site's boundary.
- *level #* : level in the hierarchy
- *#MAs* : number of MAs in this level
- *Commun Prefix*: Prefix used by the MHs to configure their VCoA
- *MA_i Address*: The IP address of *MA_i*

Table of Contents

1	Introduction	3
2	The IETF Mobile IPv6	3
3	A Hierarchical Mobility Management Architecture	4
3.1	Protocol Overview	4
3.2	Protocol Details	7
3.2.1	Registration Phase	8
3.2.2	Mobility Networks and Mobility Agents Discovery	9
3.2.3	Packet Delivery	10
3.3	Deploying Several Mobility Agents per MN	10
3.4	The Multiple Border Routers Case	11
4	Comparison and Evaluation	11
4.1	Routing and Transition Performance	13
4.2	Scalability Performance	13
5	Related Work	17
6	Conclusions and Future Work	20
7	Acknowledgement	21
A	Mobility Information Option Format	22
B	Extended Mobility Information Option Format	23



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399