Extending Mobile-IPv6 with Multicast to Support Mobile Networks in IPv6

Thierry Ernst[†‡], Claude Castelluccia[†] and Hong-Yon Lach[‡]

[†]INRIA Rhône-Alpes
ZIRST - 655, Avenue de l'Europe
38330 Montbonnot Saint Martin - France

[‡]MOTOROLA LABS Paris
Espace Technologique St-Aubin
91193 Gif-sur-Yvette Cedex - France

E-mail:{Thierry.Ernst | Claude.Castelluccia}@inrialpes.fr ; Hong-Yon.Lach@crm.mot.com

Abstract : This paper addresses the problems of routing datagrams to nodes located in an IPv6 *mobile network*. A *mobile network* is a network that is changing its point of attachment dynamically such as a network deployed in an aircraft, a boat, or a car. The IETF Mobile-IPv6 protocol that has been developed to support *mobile nodes* is unable to support *mobile networks* efficiently. In fact, we show that Mobile-IPv6 would not scale and would introduce some authentication problems. This paper proposes a solution that combines multicast routing with Mobile-IPv6 to support *mobile networks* in the Internet.

Keywords : Mobile networks, IPv6, Mobile-IP, Multicast, Routing

1. Introduction

Mobile-IPv4 [13] and Mobile-IPv6 [4] have introduced mobility support for IPv4 and IPv6 [3] nodes respectively. The purpose of mobility support is to provide continuous Internet connectivity to *mobile nodes*. Mobile-IP is a solution to support *mobile nodes* but does not handle *mobile networks*.

There are situations where an entire network might move and attach to different places in the Internet topology. In this paper, we refer to a *network* as a set of nodes that share the same IP prefix and that are attached to the Internet through a single border router. We refer to a *mobile network* as a *network* whose border router is dynamically changing its point of attachment to the Internet and thus its reachability in the IP topology. The internal architecture of a *mobile network* is preserved while it is roaming. As such, nodes in the *mobile network* do not move with respect to the others and shouldn't take part in mobility management.

Applications of *mobile networks* include networks attached to people (Personal Area Network or PANs) and networks of sensors deployed in aircrafts, boats, cars, trains, etc. As an example of a *mobile network*, we could imagine that an airways company provides permanent on-board Internet connectivity. This allows all passengers to use their laptops to connect to remote hosts, download music or video from any provider, or browse the web. The Internet could also be used to exchange information between the aircraft and air traffic control stations. This scenario has already been investigated by Eurocontrol (European Organization for the Safety of Air Navigation [14]). During the flight, the aircraft changes its point of attachment to the Internet and is reachable by distinct IP addresses owned by distinct Internet service providers. This scenario justifies that *mobile networks* may be of a big size, containing hundreds of hosts and several routers and may attach to very distant parts of the Internet topology. Moreover, it shows that we face two distinct levels of mobility, node mobility and network mobility, since laptops owned by passengers are themselves *mobile nodes* visiting the aircraft *mobile network*.

A *mobile network* attaches to the rest of the Internet through its border router which we refer to as the *mobile network gateway* (*MNG*). Points of attachment are called *foreign gateways* (*FGs*). We also call *mobile network node* (*MNN*) any host or router located within the *mobile network* and *correspondent node* (*CN*) any external node corresponding with some *MNN* of this *mobile network* (see figure 1 for the terminology). Then, all datagrams sent between *CNs* and *MNNs* necessarily transit through the *MNG*.

Although the designers of Mobile-IPv4 claim that it could support *mobile networks* equally as *mobile nodes* ([13] section 4.5, [12] section 5.12, [15] section 11.2), we argue that this is not true for Mobile-IPv6, which therefore requires some changes in the specification. Indeed,

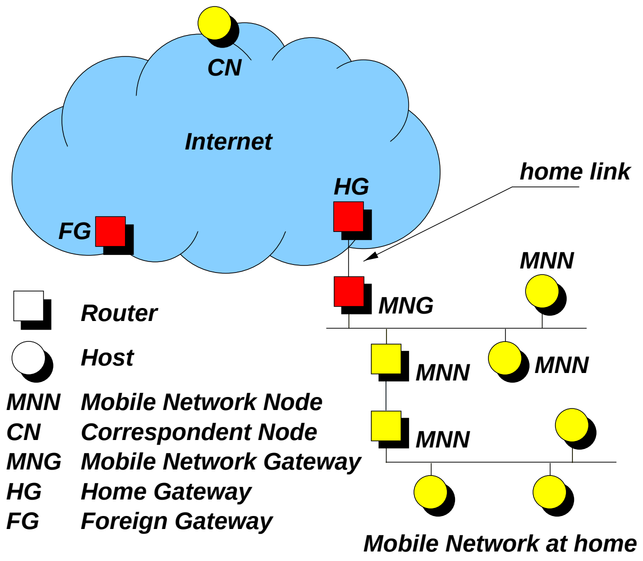

Figure 1. Terminology for Mobile Networks

we have carefully studied the adequacy of Mobile-IPv6 for supporting *mobile networks* and we came to the conclusion that some minor changes would allow Mobile-IPv6 to support them, but not optimally. As we will show in section 3, scalability becomes a more important issue for *mobile networks* compared to *mobile nodes*. As the *mobile network* may contain hundreds of nodes, each communicating with several peers, the questions of locating, optimal routing and signaling overload are significantly more important. Slightly enhancing the Mobile-IPv6 specification for supporting *mobile networks* would not provide optimal routing without overloading the network with signaling. Moreover, Mobile-IPv6 wouldn't handle authentication correctly. We therefore propose some important enhancements and modifications to the current specification. Those enhancements do not question the intrinsic features of Mobile-IPv6. Our Mobile-IPv6 extensions are based on IPv6 multicast capabilities. The current *careof address* of the *mobile network* is delivered to a multicast group formed of all *CNs* interested in getting the *Binding Updates*. We do not describe the mechanisms for joining, leaving and sending data to a multicast group. We will rely on IPv6 multicast features currently being developed in the Internet Engineering Task Force (IETF).

This paper is structured as follows: section 2 reviews the Mobile-IPv6 protocol for *mobile nodes*. Section 3 explains why Mobile-IPv6 is unable to support *mobile networks*. Section 4 describes our approach for limiting the signaling bandwidth consumption and also discusses support of *mobile nodes* visiting a *mobile network*. We then evaluate our proposal in section 5 and finally, section 6 concludes this paper.

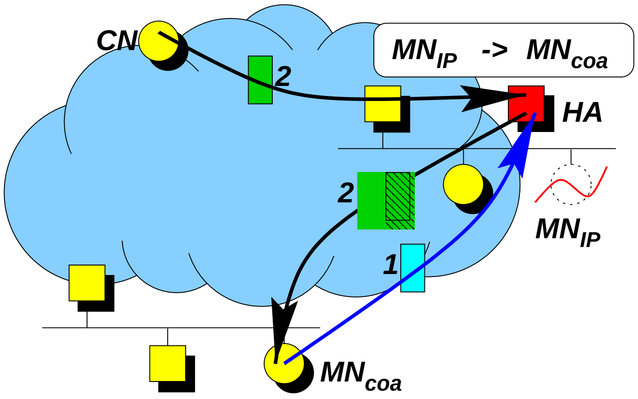

Figure 2. (1) Registration with the home agent and (2) first datagrams between correspondent nodes and mobile nodes in Mobile-IPv6

2. Review of Mobile-IPv6 for mobile nodes

Mobile-IPv6 [4] is a work in progress in the IETF, offering support for IPv6 *mobile nodes*. Although it is not yet standardized, every IPv6 node is required to implement Mobile-IPv6, which means that mobility is ought to be widely supported in IPv6.

Mobile-IP relies on a two-tier addressing scheme. The *mobile node* (*MN*) uses a permanent *home address* MN_{ip}, as an identifier, and a *careof address* MN_{coa}, as a routing directive. When roaming, the *MN* detects its movement. Then, it obtains a new *careof address* MN_{coa} on each subsequent foreign link it visits using either stateless [16] or stateful DHCPv6 Address Autoconfiguration [1]. The *mobile node* may own several MN_{coa} at anytime, one of which is selected as the primary MN_{coa}. The *MN* has to register the binding between its *home address* MN_{ip} and the primary MN_{coa} with the *home agent* (*HA*) which is a special router on the *home link* able to intercept and redirect datagrams intended to the *MN* (figure 2). This is performed by means of a *Binding Update*. The *Binding Update* is a datagram containing a *Binding Update Option* which carries the MN_{coa} (unless it is already recorded in the IP header source address field), and a *Home Address option* which specifies the MN_{ip} used as a permanent identifier of this *MN*. The *Binding Update* and the *Home Address options* are contained in an *IPv6 Destination extension header*. All datagrams carrying a *Binding Update option* must also contain an *AH* [5] or an *ESP* [6] *extension header* used for authentication. Receiving the *Binding Update*, the *HA* uses gratuitous Neighbor Advertisement messages [10] to intercept all datagrams intended for the *MN* and then encapsulates them to the current MN_{coa} (figure 3).

At this moment, the *MN* may want to provide its primary MN_{coa} to its *correspondent nodes* (*CNs*) to avoid triangular routing via the *HA* between the *CN* and the *MN*. This is done by means of periodic *Binding Updates*, which might be

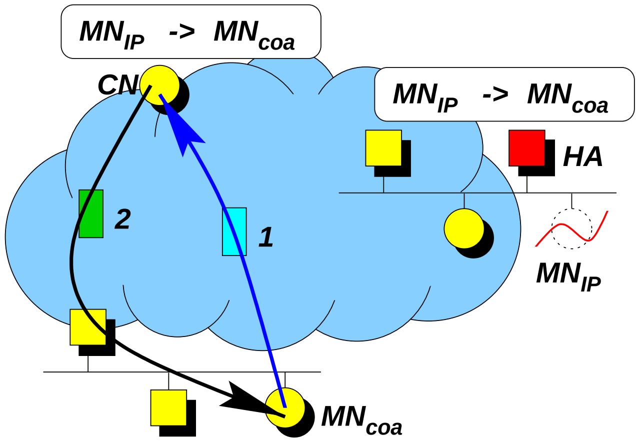

Figure 3. Optimal Routing to mobile nodes in Mobile-IPv6

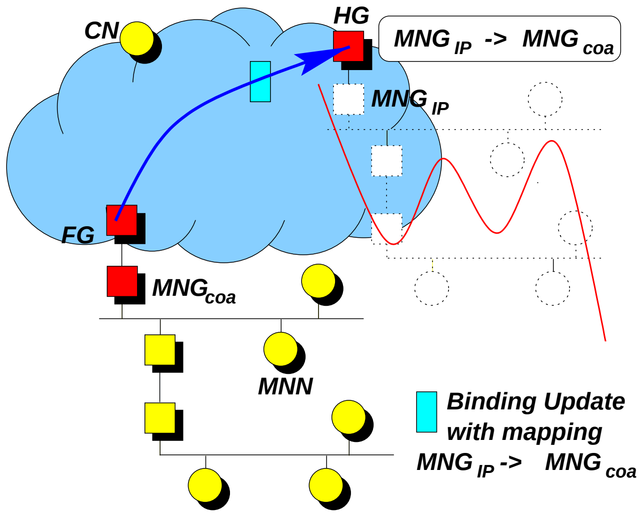
CN
HG
MNG_{IP} -> MNG_{coa}
MNG_{IP}
FG
MNG_{coa}
MNN
Binding Update with mapping
MNG_{IP} -> MNG_{coa}

Figure 4. Mobile Networks in Mobile-IPv4 - registration

piggybacked in payload datagrams or sent alone in separate packets containing no payload. The datagram must include a *Binding Update option*, a *Home Address option* and an *AH* or an *ESP extension header*. At reception, the *CN* authenticates the datagram using the *AH or ESP header* and then sends the forthcoming datagrams directly to the MN_{coa} using an *IPv6 Routing extension header* containing the MN_{ip}.

In order to bypass ingress filtering and to be identified by upper layer protocols at the destination, the source address of payload datagrams sent by the *MN* is set to the MN_{coa} while the MN_{ip} is inserted in a *Home Address Option* of the *Destination extension header*.

3. Mobile IP and Mobile Networks

From the routing perspective, distinction between nodes is only necessary for routing inside the network. From outside, a network can be virtually perceived as a single node (the border router *MNG*) with one address (or prefix) and n interfaces attached to it. According to this observation, the Mobile-IPv4 specification proposes to support *mobile networks* as standard *mobile nodes* (see [13] section 4.5). The *mobile node* is the border router *MNG* of the *mobile network*. It has a permanent *home address* MNG_{ip} and gets a new *careof address* MNG_{coa} at each subsequent point of attachment and sends a *Binding Update* to its *home agent HG*[1] (figure 4) to instruct it to intercept all datagrams intended for its *MNNs* which necessarily transit through the *MNG* (figure 5).

Mobile-IPv6 and Mobile-IPv4 with Routing Optimization [11] could in theory support *mobile networks* similarly as in Mobile-IPv4. However, although mentioned in the Mobile-IPv4 specification, the current specifications of Mobile-IPv4 with Routing Optimization and Mobile-IPv6 doesn't mention them anymore.

Obtaining a new MNG_{coa} and requesting the *HG* to redirect on the way datagrams intended for *MNNs* doesn't require modifications in the Mobile-IPv6 specification as this could be done independently for a host or for a router. Since datagrams intended for a *MNN* necessarily transit through the *MNG*, the *HG* easily claims to be the *MNG* and redirects them to the MNG_{coa} using encapsulation.

However, the emission of *Binding Updates* to *CNs* does not allow Mobile-IPv6 to support *mobile networks* as easily as in Mobile-IPv4. This is particularly true because *Binding Updates* are sent by the *mobile node* itself, and not by the *HA* as in Mobile-IPv4 with Routing Optimization. Regarding the emission of *Binding Updates* to *CNs*, it makes sense that the node which is assigned the MNG_{coa} also sends *Binding Updates*. As a result, the *MNG* would send a *Binding Update* on behalf of its *MNNs*, which has the benefit of hiding mobility of the network to the *MNNs* and frees them from any mobility management. Doing so while keeping *MNNs* out of any mobility management would require that the *MNG* tracks the *CNs* of the *MNNs* and sends them *Binding Updates*. As we have seen in section 2, *Binding Updates* could be piggybacked in datagrams sent by *MNNs*, or sent in special purpose datagrams. In any case, those datagrams require authentication. Piggybacking cannot be done by the *MNG* without rewriting the AH or ESP header which may be present. Sending *Binding Updates* in separate datagrams requires that the *MNG* uses the same security association as the *MNN* so that the *CN* accepts to send datagrams to the *MNN* via the MNG_{coa}. Both scenarios do not comply with IPv6 recommendations since no headers but the Routing extension (under some particular conditions) can be rewritten by routers along the path. We could think of mechanisms that would allow the *MNG* to authenticate itself as

[1] In order to avoid confusion with the *home agent* of a *mobile node*, subsequent paragraphs refer to the *home gateway* (*HG*) as the *home agent* of a *mobile network*.

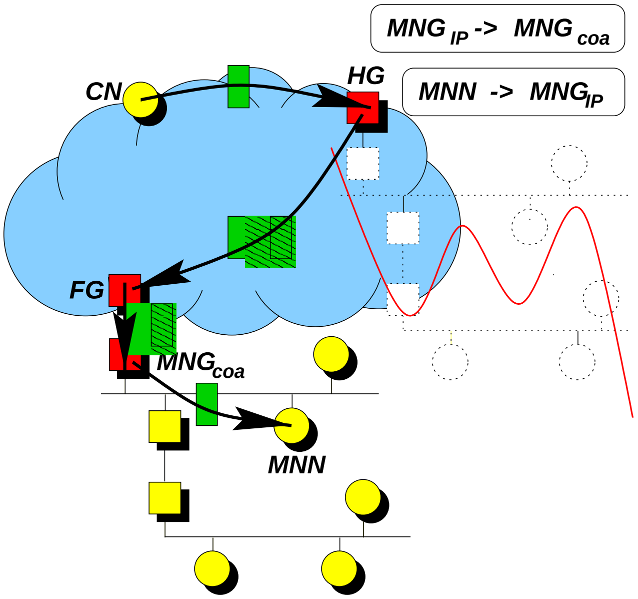

Figure 5. Mobile Networks in Mobile-IPv4 - data transmission

one of its *MNNs*. But, although the *MNG* and its *MNNs* are likely to trust each other and adopt the same administrative policy, it is not desirable to mislead the recipients since, as stated in Mobile-IPv6, no node is authorized to send *Binding Updates* on behalf of a mobile node [4, section 10.8].

To avoid these constraints, we may nevertheless consider that *MNNs* directly send *Binding Updates* to their *CNs*. This solution requires a mechanism to distribute the MNG_{coa} to all *MNNs*. They would consequently take part in mobility management. This approach is quite advantageous since the process of sending and authenticating *Binding Updates* is left unchanged. *Binding Updates* could be piggybacked or sent alone and are authenticated as coming from the *MNNs*. On the other hand, it requires changes of the Mobile-IPv6 *mobile node* operation as *MNNs* do not need to perform the tasks of obtaining a *careof address* and registering it with some *HA* and previous attachment point.

As we can see, both approaches have drawbacks and necessitate changes in the Mobile-IPv6 specification. More importantly, periodic *Binding Updates* are sent to each *CN*. As *mobile networks* may contain hundreds of nodes, each communicating with several peers, the number of *CNs* is growing with the size of the *mobile network* and is likely to be very large. In this situation, the emission of *Binding Updates* to a large number of *CNs* would cause a *Binding Update* explosion as shown on figure 6. We also note that *CNs* are misled by the origin of *Binding Updates*. As they may be communicating with several *MNNs* in the same *mobile network*, they would redundantly record a binding containing the same MNG_{coa} for each *MNN*.

Since the issue of optimal routing between *CNs* and *MNNs* cannot be left aside for *mobile networks*, we see that the intrinsic mechanisms of the current Mobile-IPv6 specification are not able to support *mobile networks* and would potentially cause an important waste of bandwidth resources and processing power. Simply updating the specification for supporting *mobile network* would lead to a *Binding Update explosion* and a hack in the authentication of these *Binding Updates*. We therefore conclude that major changes are required in the specification for efficiently supporting *mobile networks*.

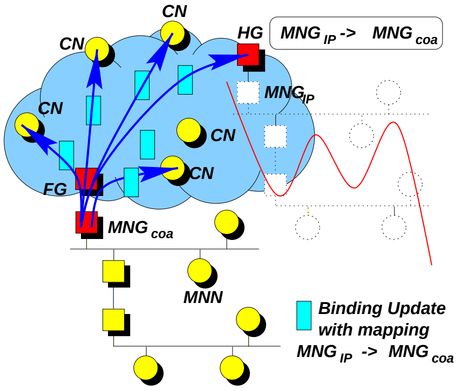

Figure 6. Binding Updates explosion

Before going any further, we should identify some constraints for a workable solution. First, routing from *CNs* to *MNNs* should be optimal and be maintained with minimal signaling overload while ensuring authentication of control messages. Second, *MNNs* shouldn't be concerned with mobility of their network. Although they may encounter variable delays of transmissions and loss with their *correspondent nodes* as the network is moving, mobility management should be transparent to them. Thus, *MNNs* should have no responsibility in the periodic delivery of *Binding Updates*. Third, Mobile-IPv6 has to work for *mobile nodes* visiting a *mobile network*. We also make the assumption that the *mobile network* has only one *MNG* and is not multihomed. At last and more importantly, the solution must scale to the size of the *mobile network* and the number of its *CNs*. It also has to scale to an important number of *mobile networks*.

4. Binding Updates multicasting

We have seen in section 3 that, following some of the Mobile-IPv6 mechanisms, the *MNG* would get a new *careof address* MNG_{coa} at each subsequent point of attachment. Whether it is performed by the *MNG* or by its *MNNs*, the *careof address* would then be sent periodically to each *CN* corresponding with *MNNs*. As we observe, each *CN* would receive exactly the same MNG_{coa}. Some *CNs* might even receive duplicate *Binding Updates* carrying the same MNG_{coa} in case they are corresponding with several *MNNs* residing in the *mobile network*. It would then be wise that *CNs* use a unique entry for their corresponding *MNNs*.

We therefore propose to deliver *Binding Updates* con-

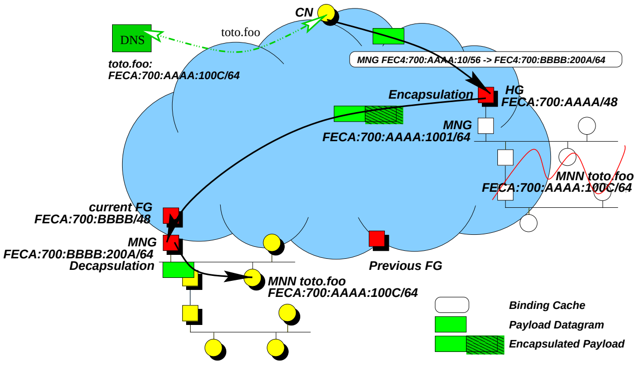

Figure 7. First datagrams transit through the home agent HG

taining network scope bindings, i.e. the binding associates the *careof address* MNG_{coa} with the network prefix MNG_{prefix} of the *mobile network* instead of the full 128-bits IPv6 *home address* MNG_{ip} as in the existing Mobile-IPv6. The MNG_{prefix} is a bit string that consists of some number of initial bits of the *home address* MNG_{ip} which identifies the *home link* within the internet topology. All *MNNs* share the same IP prefix.

Therefrom, it makes sense that the *MNG* sends the *Binding Update* to a multicast group to which *correspondent nodes* would have subscribed. At the same time, it alleviates the need for the *MNG* to know the list of *CNs* since the source of a multicast datagram does not need to know who the particular group members are. At last, authentication is made easier since *Binding Updates* containing the *careof address* are always sent by the sole *MNG* to the same multicast address. It might also reduces signaling and memory requirements when *CNs* have several correspondents in the same *mobile network*.

We therefore propose a solution based on multicast routing protocols for delivering network scope *Binding Updates*. The *mobile network* has a permanent multicast address which the *MNG* registers in the DNS [7, 8]. The *MNG* sends periodic *Binding Updates* containing a binding between its MNG_{prefix} and its MNG_{coa} to the multicast address. *CNs* join the multicast group using IPv6 multicast mechanisms. The *Binding Update* instructs *CNs* to add an entry in their binding cache. Before sending a datagram, the *CN* checks if the prefix of the destination address matches the MNG_{prefix} recorded in the binding cache. If so, datagrams are sent via the MNG_{coa} using a Routing extension header.

As detailed below, our proposition makes use of most of the existing Mobile-IPv6 facilities with some extensions. We define a new mobile entity, the *MNG*, which performs most of the existing Mobile-IPv6 features and we redefine the *CN* operation. We also define a new DNS Resource Record and a new *IPv6 Destination option*.

4.1. Extensions

A new IPv6 Destination Option The *Network Update option* is a new *IPv6 Destination option* that we propose for instructing a *CN* to redirect datagrams intended for all its communication peers sharing the same address prefix. The *Network Update option* contains the MNG_{prefix} used as a netmask and the MNG_{coa}. This option is inserted into the *Binding Update* datagram.

DNS extensions The *Mobile Network DNS Resource Record* is a new entry in the Domain Name Server. It records the multicast address used by *MNGs* for *Binding Update* delivery. Queries for the IP address of a *MNN* is processed in the following manner: the DNS understands that the domain name provided in the query corresponds to a *MNN*. It therefore returns the relevant IP address together with the multicast address used by the corresponding *MNG*.

Mobile-IPv6 extensions Our solution requires extensions to the Mobile-IPv6 specification. The *MNG* is a new entity very similar to the *mobile node* as described in [4]. Thus, the *MNG* is a *mobile node* enhanced with the ability to obtain a multicast address and to register it in the DNS. To the contrary of a standard *mobile node*, it includes the *Network Update option* in the *Binding Update* instead of the *Binding Update option* and sends the *Binding Update* to a multicast address instead of to individual *CNs*. No changes are required to the standard *MN* operation. On the other hand, the *CN* operation is extended to process the multicast address contained in the reply of the DNS. The *CN* is also enhanced to process the *Network Update option* and to transmit via the MNG_{coa} all datagrams bearing a destination address matching the same prefix as the MNG_{prefix}. The *HG* corresponds to the existing *home agent*. Our solution only requires that the *HG* is able to redirect datagrams intended to *MNNs* and not only to the *MNG*. At last, and as a standard IPv6 router, there is no need to define a *FG* entity in the Mobile-IPv6 specification.

4.2. Protocol operation

Initialization Prior to its first movement, the *mobile network* gets a multicast address which defines the *Binding Update recipient group* for this *mobile network*. The *MNG* registers it with the DNS by means of DNS Dynamic Update [17] and using a newly defined DNS Resource Record. Upon reception of this dynamic update, the name server serving the *mobile network* understands that it must return the multicast address to all nodes enquiring for the IP address of any *MNN* of the network.

MNG Operation Similarly to *mobile nodes*, the *MNG* obtains a new MNG_{coa} at each of its subsequent points of attachment using either stateless or stateful DHCPv6 address autoconfiguration. Following this, it sends a *Binding Update* datagram to its *HG* and to the multicast address

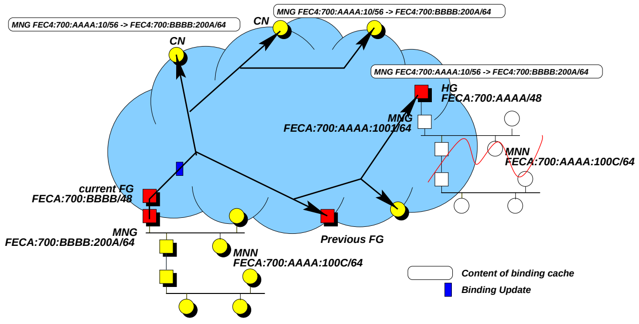

Figure 8. Binding Updates Distribution from the MNG

identifying the *Binding Update recipient group*. The *Binding Update* datagram contains a *Network Update option*, a *Home Address option*, and an *AH or ESP header*. The *Binding Update* is sent at periodic time intervals to ensure that *CNs* do not delete the binding because its lifetime has expired. The *Home Address option* is the same as for the existing Mobile-IPv6 specification. It contains the *home address* MNG_{ip} used as an identifier. Figure 8 shows *Binding Update* emission towards the multicast group of subscribed *CNs*.

CN Operation Prior to communication establishment with a node *MNN*, the *CN* calls the DNS for the IP address corresponding to the domain name of that node. The DNS supplies the IP address MNN_{ip} and also returns the multicast address where *Binding Updates* are sent since the requested node is a *mobile network node*. From the multicast address, the *CN* understands that its peer is a *mobile network node*. First datagrams are therefore sent to the *MNN*. Datagrams are intercepted by the *HG* and tunneled to the MNG_{coa}(figure 7). Meanwhile, the *CN* joins the *Binding Update recipient group* with the multicast address provided by the DNS. This could be performed by means of Multicast Listener Discovery [2], the protocol used for routers to discover neighboring hosts interested in getting multicast datagrams. Following *Binding Updates* emitted by the *MNG* to the multicast address are forwarded up to the *CN*. The *CN* verifies the authenticity of each *Binding Update* it receives and registers in its binding cache the binding between the MNG_{prefix} and the MNG_{coa} and sets the expiration timer. When sending datagrams, the *CN* checks its binding cache and understands that datagrams intended to any destination address bearing a prefix matching the MNG_{prefix} should be routed via the MNG_{coa}. Subsequent datagrams are sent to MNN_{ip} via the MNG_{coa} using an IPv6 Routing extension header (figure 9). When communication is over, the *CN* may leave the multicast group.

4.3. Mobile nodes visiting a mobile network

The existing Mobile-IPv6 can still be used by *mobile nodes* moving into and out of a *mobile network*, but not optimally. A visiting *mobile node MN* has a permanent home address MN_{ip} and obtains a *careof address* MN_{coa} from a router in the *mobile network*. The prefix of this MN_{coa} corresponds to the MNG_{prefix}. The *MN* registers its current MN_{coa} with its *home agent* and its *correspondent nodes*. As shown on figure 10, datagrams emitted by *CN* and destined to *MN* are routed to the MN_{coa} using a routing extension header. The prefix of the MN_{coa} corresponds to the prefix of the home address of the *MNG*. Datagrams are therefore routed towards the *HG*. As the *mobile network* is not at home, datagrams are intercepted by *HG* and encapsulated to MNG_{coa}. The *MNG* decapsulates and forwards them to the *MN* where the routing extension header is processed. The mobility of the visited network is therefore transparent to the *MN* which keeps sending *Binding Updates* containing the MN_{coa} to its *home agent* and its *correspondent nodes*. We note that routing is not optimal because the *home agent* and the *correspondent nodes* of the *MN* are unaware of the current location of the *mobile network* although they are aware that the *MN* is located in the *mobile network*. The solution is to provide the *correspondent nodes* of the *MN* with the MNG_{coa}. This would obviously require specific extensions to Mobile-IPv6 since *correspondent nodes* of the *MN* would have to understand that datagrams have to be sent to MNG_{coa} using a *Routing extension header* including both the MN_{coa} and the MN_{ip}.

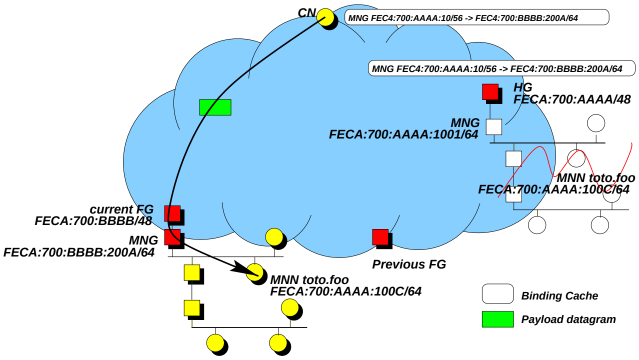

Figure 9. Optimal routing between CN and MNN

5. Evaluation

We believe that our solution is well adapted for supporting large *mobile networks*. Network scope *Binding Updates* ensure optimal routing between *CNs* and *MNNs*. By using multicast features for delivering *Binding Updates*, our solution is best designed to reduce signaling load and to scale to a large number of *CNs*, particularly when the emission rate of *Binding Updates* is significant. However, we have identified a couple of advantages and drawbacks, and issues on which to further conduct our study.

Advantages First, *MNNs* do not have to take part in the mobility management of their network, this is entirely man-

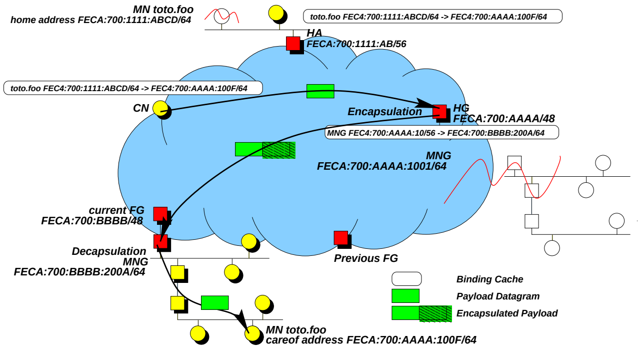
MN toto.foo
home address FECA:700:1111:ABCD/64
toto.foo FEC4:700:1111:ABCD/64 -> FEC4:700:AAAA:100F/64
HA
FECA:700:1111:AB/56
toto.foo FEC4:700:1111:ABCD/64 -> FEC4:700:AAAA:100F/64
CN
Encapsulation
HG
FECA:700:AAAA/48
MNG FEC4:700:AAAA:10/56 -> FEC4:700:BBBB:200A/64
MNG
FECA:700:AAAA:1001/64
current FG
FECA:700:BBBB/48
Decapsulation
MNG
FECA:700:BBBB:200A/64
Previous FG
Binding Cache
Payload Datagram
Encapsulated Payload
MN toto.foo
careof address FECA:700:AAAA:100F/64

Figure 10. Mobile node visiting a mobile network

aged by the *MNG*. Moreover, the *MNG* doesn't need to track the *CNs* interested in getting *Binding Updates* which saves memory and processing power. As a result of receiving a network scope *Binding Update*, *CNs* are able to optimally route datagrams to *MNNs* while registering a unique entry in their binding cache for all nodes located in the same *mobile network*. More importantly, *Binding Updates* allow optimal routing and are easily authenticated as originated from the *MNG* without misleading *CNs* about the origin of the sender. It takes advantages of existing IPv6 features and mostly reuses Mobile-IPv6 mechanisms. Moreover, by advertising the multicast address identifying the *Binding Update recipient group* in the DNS and joining the multicast group simultaneously while sending the first datagrams, the *CN* is likely to get the first *Binding Update* more quickly. Otherwise, it would have waited until payload datagrams reach the *mobile network*, which would tell the *MNG* to send the first *Binding Update*. At last, our solution doesn't require that all *correspondent nodes* subscribe to the group to enable communication with a *mobile network* although routing would not be optimal for those *correspondent nodes*.

Drawbacks Our solution requires changes at *CNs*, i.e. at every IPv6 node, but the commercial deployment of IPv6 has not started yet, which limits the impact of this drawback. Our solution also requires new DNS records and a new entity *MNG* in the Mobile-IPv6 specification, but this is not an issue either as this is only done at the *mobile network* side, which doesn't impact all implementations. The most important drawback is that *Binding Updates* cannot be piggybacked and that the *MNG* has no control over the *Binding Update recipient group* members.

Privacy Issues Since *Binding Updates* are sent to a multicast group, the *MNG* has no clue over the identity of the group members. The *MNG* is therefore unable to make discrimination between subscribed *CNs* since multicast *Binding Updates* are sent to all or none group members. As a result, the *MNG* can not hide its location to some of the *CNs* and ensure privacy to some of the *MNNs* which may wish so. To overcome this, we may further enhance our proposal to add, for instance, encryption of the multicast address recorded in the DNS. Only allowed *CNs* may decrypt the multicast address and subscribe to the *Binding Update recipient group*.

Multicast Issues The performance of our protocol may depend on the underlying multicast routing protocol. Any multicast technique such as source-based tree or core-based tree may be used to build the multicast delivery tree leading to subscribed *CNs*. However, the gain of multicasting *Binding Updates* must be balanced against the computation cost of the multicast tree, its maintenance, and the density of the group members. The computation cost is indeed not the same whether *CNs* are sparsely located or not. If a source-based tree rooted at the *MNG* is used, the source of the multicast group is mobile and the delivery tree has to be recomputed upon every new point of attachment of the *mobile network*. If a core-based tree is used, the delivery tree may not be optimal. The candidate multicast routing protocol should therefore take into account the specific characteristics of dynamic multicast groups with a unique and mobile source. Obviously, it should not be limited to intra-domain multicasting and should scale to a large number of *mobile networks*.

6. Conclusion

In this paper, we have discussed the Mobile-IPv6 ability to support *mobile networks*. As we have seen, Mobile-IPv6 cannot be used without major changes if we want to provide optimal mobility support to networks. Particularly, Mobile-IPv6 doesn't scale to the size of the *mobile network* because *Binding Updates* should be sent to each outside node corresponding with a node of the *mobile network*. As a *mobile network* may contain hundreds of nodes and as each node in the *mobile node* may communicate with several *correspondent nodes*, periodic *Binding Updates* would thus overload the backbone network. Thus, Mobile-IPv6 wouldn't scale and faces a *Binding Updates explosion*.

Consequently, we propose a solution to reduce the number of *Binding Updates*. The first key idea is the use of multicast mechanisms for the delivery of *Binding Updates* to *correspondent nodes*. Basically, our solution makes use of Mobile-IPv6 mechanisms with some extensions to support multicast delivery of *Binding Updates*. Actually, instead of delivering a *Binding Update* to each single *correspondent node* of the *mobile network*, *Binding Updates* are delivered to a multicast group to which *correspondent nodes* have subscribed. The second key innovation is the use of bindings with a network scope instead of a node scope: a binding is valid for all datagrams bearing a destination address matching a network prefix instead of matching a full 128-bits IPv6 address. Our paper has outlined the changes required to the Mobile-IPv6 specification and the protocol operation of the enhanced Mobile-IPv6.

We are currently performing performance simulations of our solution to show when it is beneficial to build a multicast tree over sending *Binding Updates* by unicast. Results will appear in a forthcoming paper. It is not the focus of this paper to specify which multicast protocol should be better used. This is an open research area in which we are conducting further studies.

Previous papers addressing mobility management for *mobile nodes* have already suggested the use of multicasting. For instance, in [9], *mobile nodes* are identified by a unique multicast address, which is location independent and invariant. However, the multicast infrastructure is used to route packet destined for the *mobile node*, not to deliver *Binding Updates*.

References

[1] J. Bound and C. Perkins. Dynamic Host Configuration Protocol for IPv6 (DHCPv6). Internet-Draft draft-ietf-dhc-dhcpv6-14.txt, Internet Engineering Task Force (IETF), February 1999.

[2] S. Deering, W. Fenner, and B. Haberman. Multicast Listener Discovery (MLD) for IPv6. Internet-Draft draft-ietf-ipngwg-mld-02.txt, Internet Engineering Task Force (IETF), June 1999.

[3] S. Deering and R. Hinden. Internet Protocol Version 6 (IPv6) Specification. Request For Comments 2460, Internet Engineering Task Force (IETF), December 1998.

[4] D. B. Johnson and C. Perkins. Mobility Support in IPv6. Internet Draft draft-ietf-mobileip-ipv6-12.txt, Internet Engineering Task Force (IETF), April 2000. Work in progress.

[5] S. Kent and R. Atkinson. IP Authentication Header. Request For Comments 2402, Internet Engineering Task Force (IETF), November 1998.

[6] S. Kent and R. Atkinson. IP Encapsulating Security Payload (ESP). Request For Comments 2406, Internet Engineering Task Force (IETF), November 1998.

[7] P. Mockapetris. Domain Names - Concepts and Facilities. Request for Comments 1034, Internet Engineering Task Force (IETF), November 1987.

[8] P. Mockapetris. Domain Names - Implementation and Specification. Request for Comments 1035, Internet Engineering Task Force (IETF), November 1987.

[9] J. Mysore and V. Bharghavan. A New-Multicasting-based Architecture for Internet Host Mobility. In *Proc. of the Third ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, The Palace of the Hungarian Academy of Sciences, Budapest, Hungary, September 1997. University of Illinois at Urbana-Champaign. Available on http://www.timely.crhc.uiuc.edu/publications.html.

[10] T. Narten, E. Nordmark, and W. Simpson. Neighbor Discovery for IP version 6 (IPv6). Request For Comments 2461, Internet Engineering Task Force (IETF), December 1998.

[11] C. Perkins and D. B. Johnson. Route Optimization in Mobile IP. Internet Engineering Task Force (IETF) Internet Draft draft-ietf-mobileip-optim-09.txt, Sun Microsystems and Carnegie Mellon University, February 2000. Work in progress.

[12] C. E. Perkins. *Mobile IP, Design Principles and Practices*. Wireless Communications Series. Addison-Wesley, 1998. ISBN 0-201-63469-4.

[13] C. Perkins (Editor). IP Mobility Support. Request For Comments 2002, Internet Engineering Task Force (IETF), October 1996.

[14] T. Quinot. An IPv6 architecture for Aeronautical Telecommunication Network. Master's thesis, Ecole Nationale Supérieure des Télécommunications Paris, EUROCONTROL - European Organization for the Safety of Air Navigation - ISA project (IPv6, Satellite communication and AT-Mode for ATN), 1998. http://www.eurocontrol.fr/.

[15] J. D. Solomon. *Mobile IP, The Internet Unplugged*. Prentice Hall Series in Computer Networking and Distributed Systems. Prentice Hall PTR, 1998. ISBN 0-13-856246-6.

[16] S. Thomson and T. Narten. IPv6 Stateless Address Autoconfiguration. Request For Comments 2462, Internet Engineering Task Force (IETF), December 1998.

[17] P. Vixie, S. Thomson, Y. Rekhter, and J. Bound. Dynamic Updates in the DNS. Request for Comments 2136, Internet Engineering Task Force (IETF), February 1998.