4 Global Mobility Architecture

This chapter summarizes the *Global Mobility Architecture* (GMA) proposal. The goals of this section are: (i) to introduce GMA functional entities; (ii) to introduce the mobility support operations; and (iii) to identify the main advantages of GMA compared to MIPv6 and its main enhancements (HMIPv6 e FMIPv6).

4.1 GMA Features

Global Mobility Architecture (GMA) aims to decouple network identity from network location, usually determined from the network interface IP address configured to the terminal. This procedure enables terminal mobility without disrupting the communication sessions of the upper layers (transport and application), which are interrupted when the network interface address is changed during an ongoing communication.

In order to provide this feature, the GMA: (i) reserves a unique IPv6 Network Prefix, called *GMA Prefix* (GMAP); (ii) defines some functional entities to support mobility and (iii) defines a signaling protocol called *GMA Mobility Protocol* (GMP). The following sections present and detail these components.

4.2 GMA Functional Entities

GMA is composed of the following functional entities:

- *GMA Access Network (GAN)* – represents an access network oriented by the rules and characteristics of the GMA. The MN moves between GANs without modifying its IP address.
- *Mobile Node (MN)* – represents a host that changes its point of attachment (PoA) from one GAN to another. This change occurs without modifying

MN's IP address, which is unique, constant and formed by prepending the GMA Prefix to the interface identifier (Terminal MAC Address – TMA). [1]

- *Access Router (AR)* – represents the default access router of the GAN and provides a packet forwarding service to MNs registered at the GAN. The AR forwards packets received and addressed to MNs or sent by a MN to addresses outside the GAN. The IP address of the GAN's AR internal interface is unique, constant and reserved in all GANs. This address is represented by the GMA_DEFAULT_ROUTER constant, whose value is formed by the GMAP (GMA Prefix) appended by a suffix that does not come into conflict with any TMA – and has to be standardized.
- *Home Rendezvous Server (HRS)* – represents a server at the administrative domain of the MN that offers name and registration services. The name service is equivalent to a DNS service and the registration service enables the registering of the current location of the MN, i.e., the external interface IP address of the AR which is currently hosting the MN. The HRS IP address can be obtained by querying the DNS service at the administrative domain of the MN. The DNS Resource Record (RR) associated to the HRS will be defined in the future and is expected to be standardized.
- *Local Rendezvous Server (LRS)* – represents a server at the administrative domain of the AR that offers name and registration services. The name service is equivalent to a DNS service and the registration service enables the registration of the MN at the GAN maintained by the AR. The registration request is sent by the MN to the LRS IP address, which is unique, constant and reserved in all GANs. This address is represented by the GMA_LOCAL_RENDEZVOUS constant, whose value is formed by the GMAP (GMA Prefix) appended by a suffix that does not come into conflict with any TMA – and is expected to be standardized in the future. The registration request is intercepted by the AR which in turn registers the MN at the LRS of its domain on behalf of the MN. The AR binds MN IP address to

[1] This kind of address composition has been already proposed in [11] as a link-local address auto-configuration mechanism.

the IP address of its external interface, which is known as the *Care-of Address* (CoA) of the MN. The LRS, then, replicates this registration at the HRS of the MN.

- *Mobility Manager (MM)* – represents an entity inside the network layer of the above-mentioned entities and is responsible for processing all packets whose destination address has a GMA Prefix (GMAP). MMs use a signaling protocol called *GMA Mobility Protocol* (GMP) in order to communicate with each other.

Like MIP, hosts that communicate with the MN are known as *Correspondent Nodes* (CNs). These hosts may be mobile or stationary. Figure 7 illustrates a typical GMA scenario.

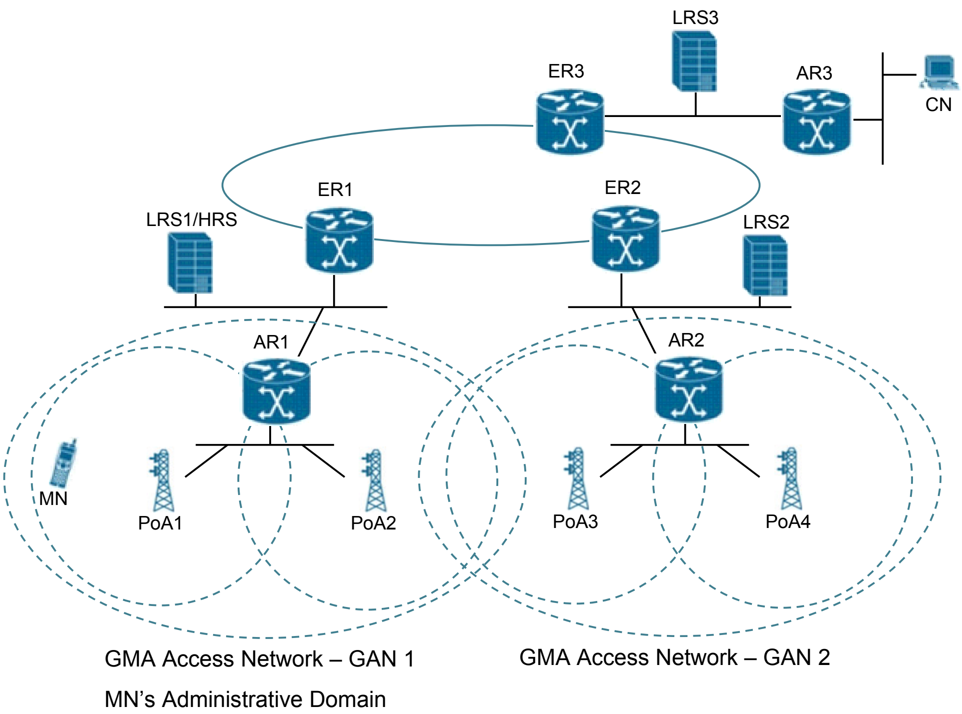

Figure 7 – GMA Scenario

A fundamental point to achieve a reliable operation of the GMA is the trustworthy relationship that should exist between the MNs and their HRSs, between the ARs and their LRSs, and between LRSs and the HRSs. Without this guarantee, GMP signaling messages could be spoofed, possibly compromising the mobility management of the MNs. This issue can be solved by the use of classic schemas of message integrity and authentication guarantees that are not covered by this GMA specification.

4.3 GMP Operation

The GMA Mobility Protocol (GMP) is the signaling protocol used between Mobility Managers (MMs) of different GMA functional entities. Three categories of messages are defined in GMP: (i) *client-network* (C2N) messages, sent by the MN to the AR of a GAN; (ii) *network-client* (N2C) messages, sent by the AR of a GAN to the MN; and (iii) *network-network* (N2N) messages, used to exchange messages between ARs of GANs, between ARs and LRSs, and between LRSs and HRSs.

GMP is based on the availability of Link Events (L2 Events) that should be provided by the link layer to the network layer (possibly by means of the MIES). Three L2 Events should be provided: (i) L2-GoingDown, indicating that a L2 handover is eminent and informs the identifications of the candidate PoAs (Point of Attachments); (ii) L2-Down, indicating that link connectivity is lost; and (iii) L2-Up, indicating that link connectivity is established. These L2 Events should be triggered by the Mobility Manager (MM) of each MN. The L2 triggers are fired when the L2 events take place at the link layer.

4.3.1 Registering the MN at the GAN

Right after the link connectivity is established (immediately after the L2-UP trigger is fired), the MN should request its registration at the GAN if there were no previous service PoA before this event. To do so, the MN should send the *C2N_RegistrationRequest* message to the GMA_LOCAL_RENDEZVOUS address. This message carries information about the identification of the selected PoA (selected PoA-ID), the HRS address (HRS_Addr) of the MN and its GMA address (GMAP_MN). This message is intercepted by the MM of the GAN's AR that checks whether the selected PoA-ID belongs to the GAN. If the checking fails, the AR sends the *N2C_Registration Response* message to the MN denying the registering at the GAN (statusFlag= DENIED_BY_AR). Otherwise, it sends the *N2N_RegistrationRequest* message on behalf of the MN to the LRS. This message carries the information related to the MN and the IP address of the external interface of the AR that will be bound to the MN as the CoA.

Next, the LRS checks whether or not it is the HRS of the MN's domain. If so, it just executes the requested registration and sends the *N2N_RegistrationResponse* message to the AR to confirm this operation (statusFlag=APPROVED). Otherwise, it decides whether or not to continue the registering of the MN at the GAN (possibly, it checks if the HRS is a *partner rendezvous*, which is an HRS that belongs to a domain of a business partner). Upon an unfavorable decision, the LRS sends the *N2N_RegistrationResponse* message to the AR denying the registering of the MN at the GAN (statusFlag=DENIED_BY_LRS). In both cases, the AR sends the corresponding *C2N_RegistrationResponse* message to the MN to inform the result of the operation.

Upon a favorable decision, the LRS sends the *N2N_RegistrationRequest* message to the HRS informing the new CoA of the MN. Then, the HRS decides whether or not to accept the registration request (possibly, it checks if the MN is an authorized mobile node). If so, the HRS registers the new CoA of the MN at its local database and sends the *N2N_RegistrationResponse* message to the LRS to confirm this operation (statusFlag=APPROVED). Otherwise, it sends the same message to the LRS denying the registering of the MN at the GAN (statusFlag=DENIED_BY_HRS). In both cases, the responses are forward to the AR, which, in turn, sends the corresponding *C2N_RegistrationResponse* message to the MN.

Upon a successful registration, the binding between the MN and the CoA must be stored in the *mobility binding cache* of the MM of the GAN's AR. This binding is volatile and expires (expiring time defaults to DEFAULT_ EXPIRING_TIME [2] seconds, which is a common value used for binding cache entries), and the same applies to the binding information stored in the LRS and HRS databases. Hence, in order to refresh this information (reset the expiring time

2 In our experiment, the DEFAULT_EXPIRING_TIME was set to 1800 seconds.

to the desired value), the MN should periodically execute the *refresh cycle*, as described in section 4.3.4 .

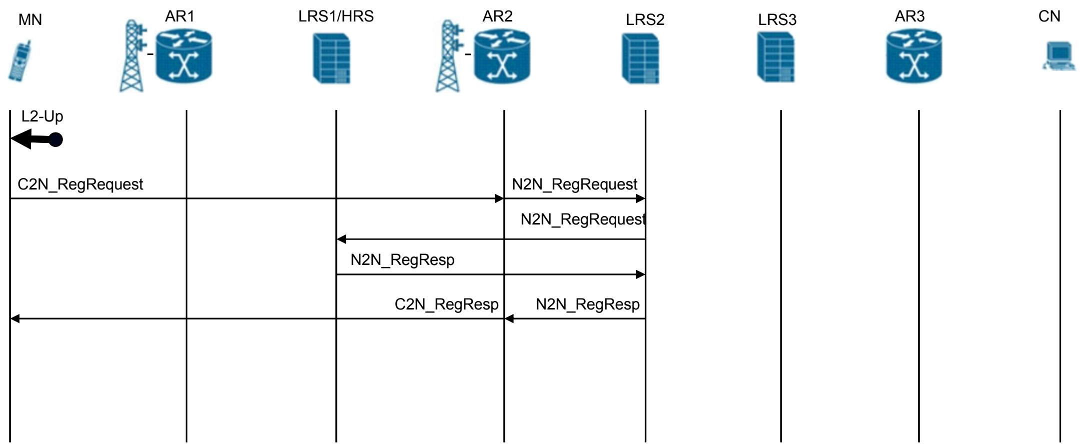

Figure 8 – Registering of the MN at the GAN of a different administrative domain

Figure 8 illustrates the registering of the MN (which address is GMAP_MN) at the GAN of a different administrative domain (HRS is different from the LRS). Table 1 presents the mobility binding cache of the AR2 (which address is ADDR_AR2) after the registration of the MN at the GAN.

Table 1 – Mobility binding cache of the AR2 after registering the MN at the GAN

Node Address	Care-of Address	HRS Address	LRS Address	Validity (s)
GMAP_MN	ADDR_AR2	ADDR_LRS1	ADDR_LRS2	1800

4.3.2 Discovering the IP address of the CN

Before initiating a communication with the CN, the MN must first submit a DNS query to the name service of the GAN using the GMA_LOCAL_ RENDEZVOUS address in order to obtain the IP address of the CN. This query is intercepted by the AR of the GAN, which, in turn, relays the query to the LRS in behalf of the MN. The LRS acts as a common DNS service when processing the name resolution related to the FQDN (Full Qualified Domain Name) queried. If the resolved address does not have a GMAP, then the CN is considered to be a stationary node. Otherwise, the CN is considered a mobile node and, then, the LRS queries the CoA of the CN registered at the HRS of the CN. As a possible optimization, the answer provided by the HRS to the first DNS query may also carry the CoA, if the CN is a mobile node.

After gathering proper information about the CN, the LRS of the MN's GAN should notify the LRS related to the CN. If the CN is a stationary node, the *N2N_BindingNotification* message is sent to the LRS of the CN's domain. Otherwise, this message is sent to the LRS of the CoA's domain of the mobile CN. In its turn, the LRS should relay the *N2N_BindingNotification* message to the AR of the CN. Thus, the MM of the CN's AR creates a record in its mobility binding cache to bind the MN to the CoA-HRS-LRS triplet informed.

By the end of the notifying operation, the LRS of the domain the MN is visiting must send the DNS query response to the AR of the GAN the MN is registered at. If the CN is a mobile node, the MM of the AR creates a record in its mobility binding cache to bind the CN to the CoA-HRS-LRS triplet informed. Then, only the CN's IP address is sent to the MN as a DNS query response. The CoA of a mobile CN is never informed to the MN.

Figure 9 illustrates the discovering of the IP address of the CN when the MN is registered at the GAN of the AR2.

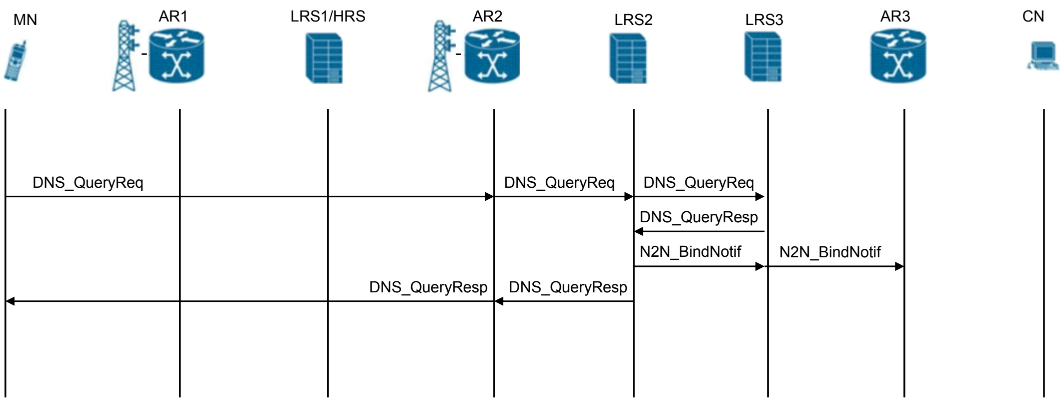

Figure 9 – Discovering the IP address of the CN

Table 2 presents the mobility binding cache of the AR2 (whose address is ADDR_AR2) after discovering the IP address of the mobile CN (whose address is GMAP_CN). Note that the first entry was recorded in the table during the registration performed by the MN, so its corresponding Validity field has already been decremented for 700 seconds since the last refresh. The second entry is a consequence of the DNS query response and has just been recorded by the AR2, so its corresponding Validity field is set to DEFAULT_EXPIRING_TIME seconds.

Table 2 - Mobility binding cache of the AR2 after discovering the IP address of the mobile CN

Node Address	Care-of Address	HRS Address	LRS Address	Validity (s)
GMAP_MN	ADDR_AR2	ADDR_LRS1	ADDR_LRS2	1100
GMAP_CN	ADDR_AR3	ADDR_LRS1	ADDR_LRS3	1800

Table 3 presents the mobility binding cache of the AR3 (whose address is ADDR_AR3) after receiving the *N2N_BindingNotification* message from the LRS3 (whose address is ADDR_LRS3). Note that the first entry was recorded in the table during the resgistration performed by the CN, so its corresponding Validity field has already been decremented for 400 seconds since the last refresh. The second entry is a consequence of the binding notification and has just been recorded by the AR3, so its corresponding Validity field is set to DEFAULT_EXPIRING_TIME seconds.

Table 3 - Mobility binding cache of the AR3 after receiving the *N2N_BindingNotification* message

Node Address	Care-of Address	HRS Address	LRS Address	Validity (s)
GMAP_CN	ADDR_AR3	ADDR_LRS1	ADDR_LRS3	1300
GMAP_MN	ADDR_AR2	ADDR_LRS1	ADDR_LRS2	1800

4.3.3 Discovering the IP address of the MN

Before initiating a communication with the MN, the CN must first submit a DNS query to a name service. If the CN is a mobile node, the process works the same way described in section 4.3.2 . Otherwise, the stationary CN must use the name service of the LRS of its domain. The LRS acts as a common DNS service when processing the name resolution related to the FQDN queried. As the resolved address of a MN always has a GMAP, the LRS should query the CoA of the MN at the HRS of the MN. As a possible optimization, the answer provided by the HRS to the first DNS query may also carry the CoA of the MN.

After gathering proper information about the MN, the LRS of the CN's domain should send the *N2N_BindingNotification* message to the AR of the CN. Thus, the MM of the AR creates a record in its mobility binding cache to bind the MN to the CoA-HRS-LRS triplet informed. Then, the MN's IP address is sent by the LRS to the CN as a DNS query response. The CoA of a MN is never informed to the CN.

Figure 10 illustrates the discovering of the MN's IP address by a stationary CN. Table 4 presents the mobility binding cache of the AR3 (which address is ADDR_AR3) after discovering the IP address of the MN (which address is GMAP_MN).

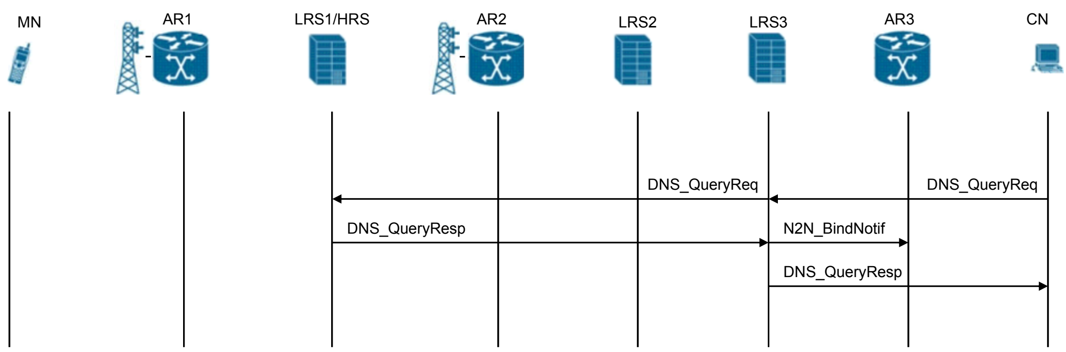

Figure 10 – Discovering the IP address of the MN

Table 4 - Mobility binding cache of the AR3 after discovering the IP address of the MN

Node Address	Care-of Address	HRS Address	LRS Address	Validity (s)
GMAP_MN	ADDR_AR2	ADDR_LRS1	ADDR_LRS2	1800

4.3.4 Refreshing the record of the MN

The binding between the MN and its CoA is stored in the *mobility binding cache* of the MM of each of the ARs that needs to forward packets addressed to/from the MN. This binding is volatile (expiring time defaults to DEFAULT_EXPIRING_TIME seconds), and the same applies to the binding information stored in the LRS and HRS databases. Hence, in order to refresh this information (reset the expiring time to the desired value), the MN should periodically execute the *refresh cycle*.

This cycle should be executed at each DEFAULT_REFRESH_CYCLE[3] seconds, half of the default expiring value of a record. To do so, the MN's time handler should trigger the system timer to be waken-up every DEFAULT_ REFRESH_CYCLE seconds. Then, at each cycle, the MN sends the *C2N_RegistrationRefresh* message to the GMA_LOCAL_RENDEZVOUS address. This message carries the current PoA-ID, the HRS address of the MN and its GMA address, the address list of the CNs communicating with the MN and the expiring value of the record (default is DEFAULT_EXPIRING_TIME seconds).

The *C2N_RegistrationRefresh* message is intercepted by the MM of the GAN's AR, which, then, refreshes the record of the MN in its binding cache and sends the *N2N_RegistrationRefresh* message in behalf of the MN to the LRS. This message carries the same information delivered by the *C2N_RegistrationRefresh* message plus the CoA of the MN. To continue the refresh cycle, the LRS sends the *N2N_RegistrationRefresh* message to the HRS of the MN and sends the *N2N_BindingNotification* message to the LRSs of all the CNs the MN is communicating with. And, to complete the cycle, each LRS of each CN sends the *N2N_BindingNotification* message to the AR of each CN. These messages do not require acknowledgement of the receiver and do not interfere in the packet flow between the MN and the CNs.

Figure 11 illustrates the refresh cycle of the MN currently registered at the GAN of the AR2.

3 In our experiment, the DEFAULT_REFRESH_CYCLE was set to 900 seconds.

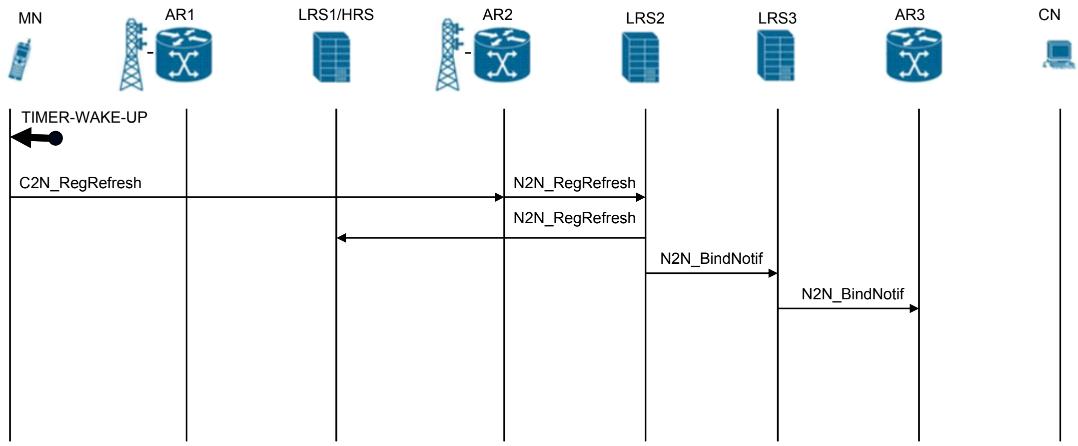

Figure 11 – Refreshing the record of the MN

4.3.5 Communication from the MN to the CN

From the point of view of the MN, its communication with the CN is not different from a traditional IP communication. That is, packets addressed to CNs outside the GAN are always forwarded to the AR acting as the default router.

When these packets arrive at the AR of the GAN, they are delivered to the MM module because there is a GMAP in the source address field of the IP header (MN's IP address). The MM checks if the MN is registered at the GAN of the AR by looking for the MN's record in the mobility binding cache. If so, the MM fills in the source address field of the IP header with the CoA of the MN and keeps the address of the CN in the destination address field. Next, the MM adds a new *IPv6 Home Address Destination* (HAD) [7] option to carry the IP address of the MN. Finally, the packet is forwarded through the network backbone until it reaches the AR of the CN. This schema is similar to the *route optimization* described in the MIPv6 specification.

If the MM of the AR finds no record associated to the MN in its mobility binding cache, one of four situations may have occurred: (i) the MN was not successfully registered at the GAN of the AR; (ii) the MN did not execute the refresh cycle; (iii) the MN moved to a different GAN and its record expired or was removed; and (iv) the packet was improperly forwarded to the AR (probably, it was spoofed). Upon any of these situations, the packet should be discarded.

When a packet arrives at the AR of the CN carrying the HAD option in the IP header, the packet should be delivered to the MM. Then, the MM must check whether or not the source address of the packet is the CoA of the address present in the HAD of the packet. If so, the MN's address is filled in the source address field of the packet, the HAD option is removed from the IPv6 header and the packet is forwarded to the CN. Otherwise, the packet should be discarded.

Figure 12 illustrates a packet sent by the MN to the CN when the CoA of the MN is registered at the MM of the CN's AR.

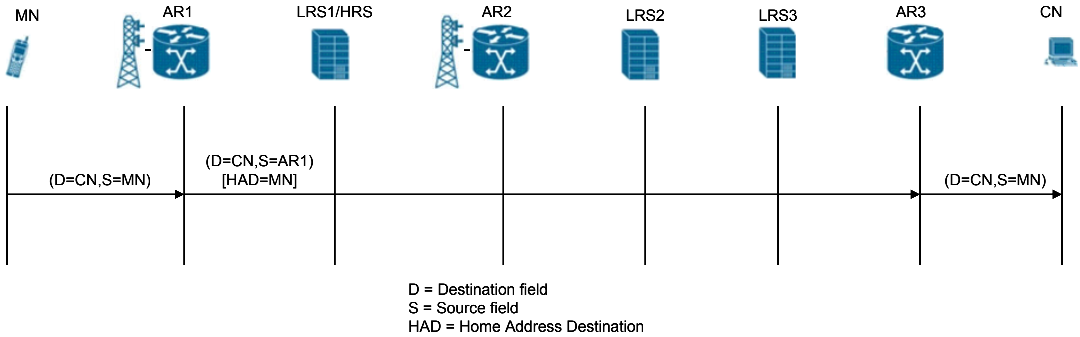

Figure 12 – Packet sent by the MN to the CN (MN's CoA registered at the CN's AR)

Generally, the MM of the CN's AR will always find the record of the MN in its mobility binding cache. But, eventually, the binding record may expire before the packet reaches the AR of the CN. Under this circumstance, the AR of the CN must send the *N2N_BindingCheckRequest* message to the LRS of its domain to confirm the binding between the address of the source field of the packet (supposed to be the CoA) and the address of the HAD option (supposed to be the address of the MN). In its turn, this LRS relays the *N2N_BindingCkeckRequest* message to the LRS of the domain of the supposed CoA. Then, this last LRS checks the binding requested and sends the *N2N_BindingCheckResponse* message to the LRS of the CN's domain, which in turn, relays the *N2N_BindingCheckResponse* message to the AR of the CN. If the binding is confirmed, the MM of the AR updates its mobility binding cache and prepares the packet to be forwarded to the CN. Otherwise, the packet should be discarded.

Figure 13 illustrates a packet sent by the MN to the CN when the CoA of the MN is not registered at the MM of the CN's AR.

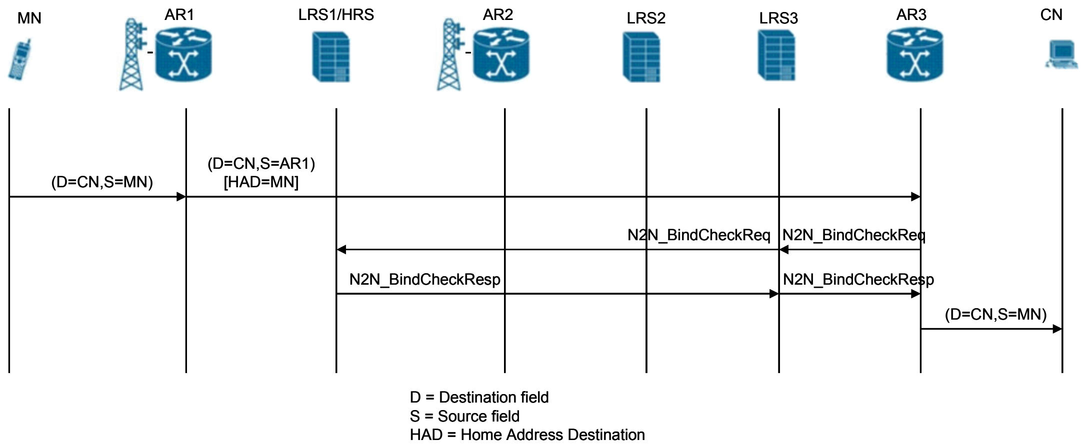

Figure 13 - Packet sent by the MN to the CN (MN's CoA not registered at the CN's AR)

4.3.6 Communication from the CN to the MN

From the point of view of the CN, the communication from the CN to the MN does not differentiate from a traditional IP communication. That is, packets addressed to MNs outside the GAN are always forwarded to the AR acting as the default router.

When these packets arrive at the AR of the CN, they are delivered to the MM because of the GMAP presence in the destination address field of the IP header (MN's IP address). Then, the MM looks for the record of the MN in the mobility binding cache. If the record is found, the MM fills in the destination address field of the IP header with the CoA of the MN and keeps the address of the CN in the source address field. Next, the MM adds a new type of IPv6 routing header called *Type 2 Routing Header* (T2RH) [7] to route the packet to the MN by way of the CoA indicated in the mobility binding cache. At last, the packet is forwarded through the network backbone until it reaches the AR of the MN. This schema is also part of the *route optimization* described in the MIPv6 specification.

As described in section 4.3.5 , whenever the CoA of the MN fails to be validated for any reason, that is, neither the mobility binding cache record nor the LRS binding check response confirms the binding between the CoA and the MN address, the packet should be discarded.

When a packet arrives at the AR of the MN carrying the T2RH in the IP header, the packet should be delivered to the MM. Then, the MM must check whether or not the destination address of the packet is the CoA of the address present in the T2RH of the packet. If so, the address of the MN is filled in the destination address field of the packet, the T2RH is removed from the IPv6 header and the packet is forwarded to the MN. Otherwise, the packet should be discarded and the Destination Host Unreachable ICMP message should be sent to the CN.

Figure 14 illustrates a packet sent by the CN to the MN when the CoA of the MN is registered at the MM of the CN's AR.

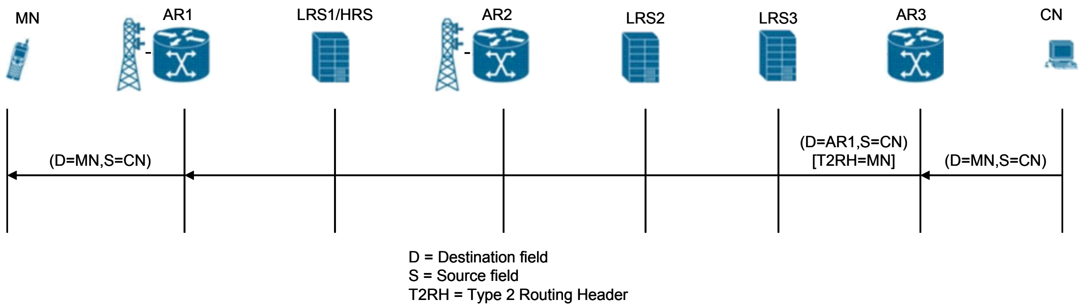

Figure 14 - Packet sent by the CN to the MN (MN's CoA registered at the CN's AR)

4.3.7 Communication between MNs

The communication between MNs is very simplified in the GMA if they are within the same GAN. That's because both of them are configured with valid local addresses of the GAN, and, so, packets can be directly delivered.

When MNs are located in different GANs, the same routing process used in sections 4.3.5 and 4.3.6 should be used. Thus, if the MN1 at the GAN1 transmits a packet to the MN2 at the GAN2, the packet is captured by the AR1 of the GAN1 and forwarded to the AR2 of the GAN2. And, then, the AR2 forwards the packet to the MN2.

Briefly, the AR1 forwards the packet to the AR2 immediately after its MM prepares the packet header, as follows: (i) the destination address field is replaced by the CoA2 of the MN2, while the source address filed is replaced by the CoA1 of the MN1; (ii) the HAD option should be added and filled in with the address of the MN1; and (iii) the T2RH should be added and filled in with the address of the MN2.

The transmission of a packet in the opposite direction, that is, from MN2 to MN1, should use the same procedure. When the AR2 captures the packet, its MM should prepare the packet header, as follows: (i) the destination address field is replaced by the CoA1 of the MN1, while the source address field is replaced by the CoA2 of the MN2; (ii) the HAD option should be added and filled in with the address of the MN2; and (iii) the T2RH should be added and filled in with the address of the MN1.

When the packet arrives at the AR1, the MM of the AR1 should remove the HAD option and the T2RH from the packet header and, then, it should fill in the destination address field with the address of the MN1 and the source address field with the address of the MN2. Next, the AR1 should forward the packet to the MN1.

Figure 15 illustrates the communication between MNs at different GANs.

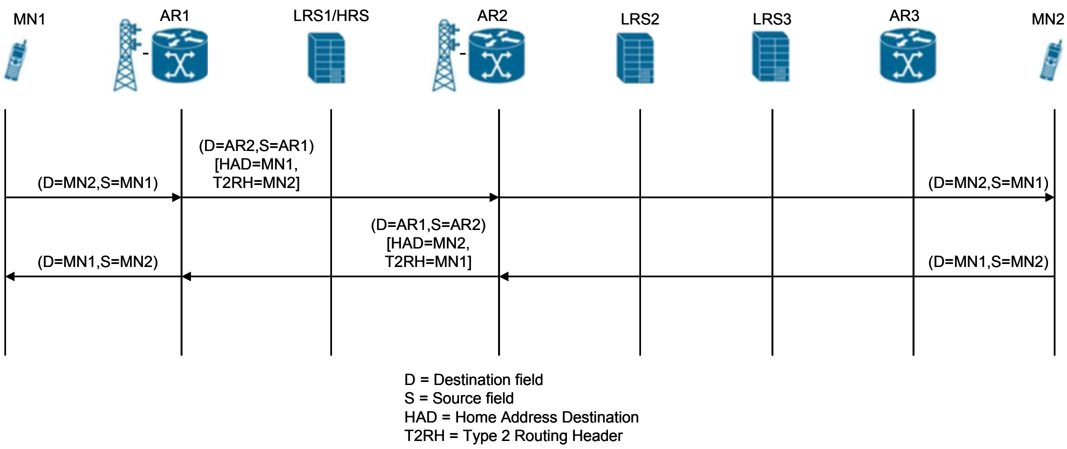

Figure 15 – Communication between MNs at different GANs

4.3.8 Handover handling in the GMA

The handover handling requires monitoring of at least three triggers of the link layer: (i) L2-GoingDown, indicates that a L2 handover is eminent and informs the identifications of the candidate PoAs; (ii) L2-Down, indicates that link connectivity is lost; and (iii) L2-Up, indicates that link connectivity is established. Based on these triggers, the network layer handover can be successfully coordinated by the MM of the MN, what maximizes the chance of a seamless handover to take place.

However, as there are different network wireless technologies, there are no guarantees about uniformity of the link layer triggers, which can make L2 handover monitoring unfeasible. In order to solve this issue, as presented in section 3.2 , the IEEE is standardizing the MIH to help detecting and initializing the L2 handover. The MIH Function (MIHF), located between layer 2 and 3, offers an event service (Media Independent Event Service – MIES) to notify the network layer about link events, and, among them, the triggers required by GMA.

The handover of the network layer (L3-handover) begins as soon as the L2-LinkGoingDown trigger is fired by the link layer. After this event, a very short period of time elapses until the L2-Down trigger is fired and link connectivity is lost. So, there is no guarantee that the MN will be able to send the required information to the AR of the GAN in order to accomplish a seamless handover. Thus, the GMP provides two possible modes of operation during L3-handover: (i) *anticipated mode of operation*; and (ii) *reactive mode of operation*. The problem addressed here is the same problem addressed by the FMIP. These modes of operation are detailed in the following sections.

4.3.8.1 Anticipated Mode of Operation

In the anticipated mode of operation, the MN is informed about the candidate PoAs by the L2-GoingDown trigger. Immediately after this event, the MN should select one of the candidate PoAs and send the *C2N_HandoverInitiate* message to the AR of the GAN. This message carries the current PoA identification (oldPoA-ID), the selected PoA identification (newPoA-ID) and the IP addresses of all the CNs which are communicating with the MN.

When the *C2N_HandoverInitiate* message arrives at the AR of the GAN, the MM of the AR checks whether or not the selected PoA-ID is part of the current GAN of the MN. If this is the case, the AR stops the L3-handover signaling because it will continue to be the AR of the GAN after the L2-handover. Otherwise, the AR, which will be the future previous router (PAR) of the MN, continues the L3 handover in behalf of the MN.

Next, the AR of the current GAN prepares the *handover context* of the MN to be forwarded to the NAR (next access router). The handover context of the MN is composed of the following information: (i) the identification of the selected PoA; (ii) the triplet composed of the MN's address, CoA and HRS; (iii) the addresses of the stationary CNs; and (iv) all the records associated to the mobile CNs found in the mobility binding cache of the AR.

Then, the AR sends the *N2N_HandoverIndication* message to the LRS with the selected PoA-ID and the handover context of the MN. In turn, the LRS relays the message to the LRS of the NAR's domain. To do so, the LRS must implement a method capable of mapping the selected PoA-ID to the IP address of the LRS of the NAR's domain. Here, we suppose that the LRS works based on maps that are exchanged by LRSs. The technical aspects of this operation are out of the scope of this work.

Finally, the *N2N_HandoverIndication* message is relayed once more and reaches the NAR. Here, we suppose that the LRS of a domain knows the PoA-IDs of all the PoAs connected to the ARs within a domain. To conclude this phase, the NAR should update its mobility binding cache based on the handover context received and send the *N2N_HandoverAcknowledge* message (statusFlag=OK) to the PAR.

Next, the PAR sends the *N2N_HandoverForward* message to the NAR to request the initialization of the buffering mechanism. Immediately after, the NAR sends the *N2N_HandoverAcknowledge* message (statusFlag=FORWARD) to confirm the initialization. As a consequence, the PAR starts the forwarding mechanism and forwards all in-flight packets of the downstream flow addressed to the MN. These packets are buffered at the NAR until the MN finishes the handover process.

Two of the main features of the anticipated mode of operation of the GMP are: (i) the anticipated registering of the MN at the new GAN; and (ii) the anticipated updating of the mobility binding cache of the ARs of the CNs the MN is communicating with. Both operations take place in parallel to the L2-handover of the MN. These features are accomplished as follows:

1. The NAR sends the *N2N_RegistrationRequest* message to the LRS of the new GAN in behalf of the MN to anticipate the regis-

tering of the MN. The registering procedure here is the same as presented in section 4.3.1 .

2. The NAR sends the *N2N_BindingNotification* message to the LRS. This message carries the address of the MN, the address of the new CoA, the address of the HRS, and the address list of all the CNs the MN is communicating with. For mobile CNs, the address of the respective LRSs and CoA are also informed. All of these are taken from the handover context of the MN or from the mobility binding cache of the NAR.

3. The LRS sends the *N2N_BindingNotification* message to all LRSs of all CNs of the MN to request the update of the mobility binding caches of the ARs of the CNs. So, the MN is bound to the new CoA in these ARs.

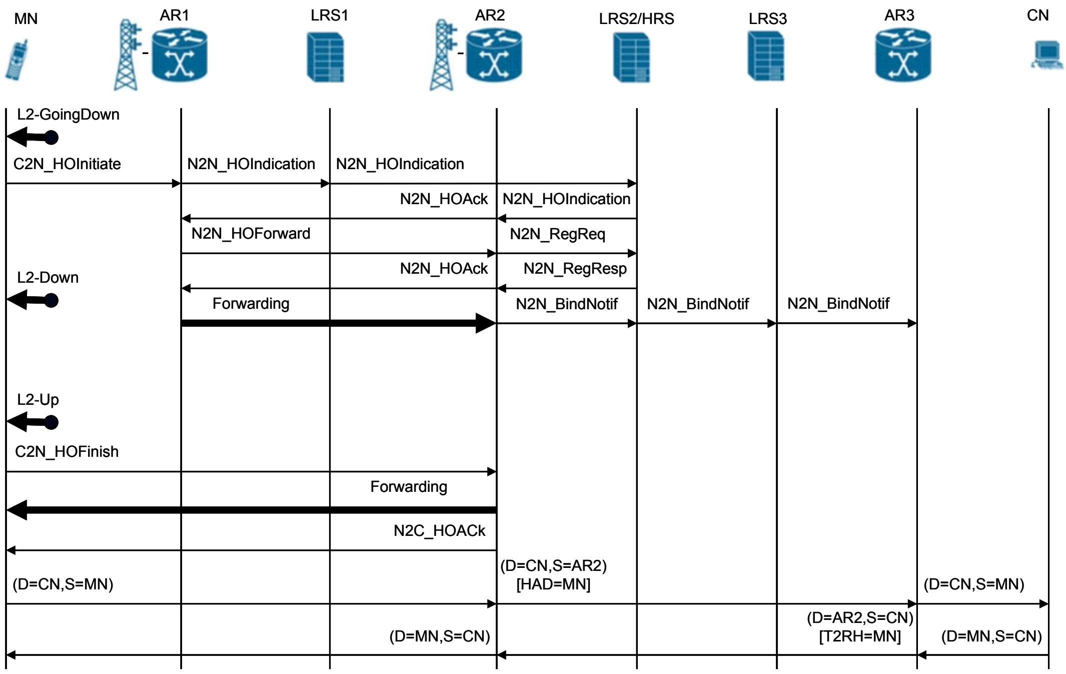

Figure 16 – Handover at the GMA - anticipated mode of operation of the GMP

Finally, the L2 handover of the MN finishes and the L2-Up trigger is fired by the link layer. After this event, the MN requests the conclusion of the L3 handover by sending the *C2N_HandoverFinish* message to the NAR. This message works as a trigger to start the delivering of the buffered packets addressed to the

MN. The end of the handover process is successfully confirmed by the NAR by sending the *N2C_Handover Acknowledge* message (statusFlag=OK) to the MN.

Figure 16 illustrates the anticipated mode of operation of the GMP during L2 handover.

4.3.8.2 Reactive Mode of Operation

The reactive mode of operation of the GMP takes place whenever the handover context of the MN is not sent by the PAR to the NAR before the end of the L2 handover. In this case, some important mobility optimization features are not executed in parallel to the L2 handover, as follows: (i) the buffering and forwarding mechanisms are not started; (ii) the mobility binding cache of the NAR is not updated with information about the CNs the MN is communicating with; and (iii) the mobility binding cache of the ARs of the CNs the MN is communicating with is also not updated.

So, the reactive mode of operation aims to restore these features after the L2 handover is completed. This mode starts when the unexpected MN requests the NAR to finish the L3 handover process. To do so, the MN sends the C2N_HandoverFinish message to the NAR. In this case, this message works as a trigger to start the reactive mode of operation.

In this mode, the NAR first checks whether or not the newPoA-ID belongs to the GAN. If so, there is nothing to do because the NAR is the PAR and the handover context of the MN doesn't change in this situation. Otherwise, the NAR should request the handover context of the MN. So, it sends the N*2N_HandoverIndication* message to the LRS of the GAN with the following: (i) the oldPoA-ID of the old GAN; (ii) the address of the MN; and (iii) the address list of all the CNs which are communicating with the MN.

Once again, we suppose that the LRS of the GAN is able to map a PoA-ID to the respective LRS. So, the LRS maps the oldPoA-ID informed to the LRS of the PAR's domain. Then, it relays the *N2N_HandoverIndication* message to this entity. Once more, the LRS must know the AR that is connected to a PoA in its domain. Hence, the LRS of the PAR's domain relays the message to the AR of the oldPoA-ID (PAR). In turn, the PAR prepares the handover context of the MN and sends it to the NAR. This information is carried by the *N2N_HandoverAcknowledge* (statusFlag=OK) sent by the PAR to the NAR.

After receiving the acknowledgment, the NAR should update its mobility binding cache based on the handover context received. Then, it should request the PAR to start the forwarding mechanism to forward the packets addressed to the MN. To do so, it sends the *N2N_HandoverForward* message. The forwarding starts just after the message reaches the PAR and the *N2N_HandoverAcknowledge* message (statusFlag=FORWARD) is sent by the PAR to the NAR to confirm the operation. The forwarded packets are immediately forwarded to the MN when they arrive at the NAR in order to restore the downstream flow.

To restore the upstream flow of the MN, the NAR needs first to update the mobility binding cache of all the ARs of all the CNs that are communicating with the MN and to register the MN at the new GAN. These tasks are executed in parallel to the forwarding of the packets of the downstream flow. The registering of the MN is executed by the NAR in behalf of the MN and is quiet similar to the procedure presented in section 4.3.1 . And the notification of the new binding of the MN is executed as explained above, by sending the *N2N_BindingNotification* message to the LRS of the GAN, which in turn, relays this message to the LRS of each CN. Finally, the LRS of the CN relays the new binding information to the AR of the CN.

After the binding notification procedure, the NAR sends the *N2N_HandoverAcknowledge* message (statusFlag=OK) to the MN. This message works as a trigger at the MN and stops the buffering of the packets of the upstream flow which, then, are delivered to the NAR to be forwarded to the CNs.

Figure 17 illustrates the reactive mode of operation of the GMP after L2 handover.

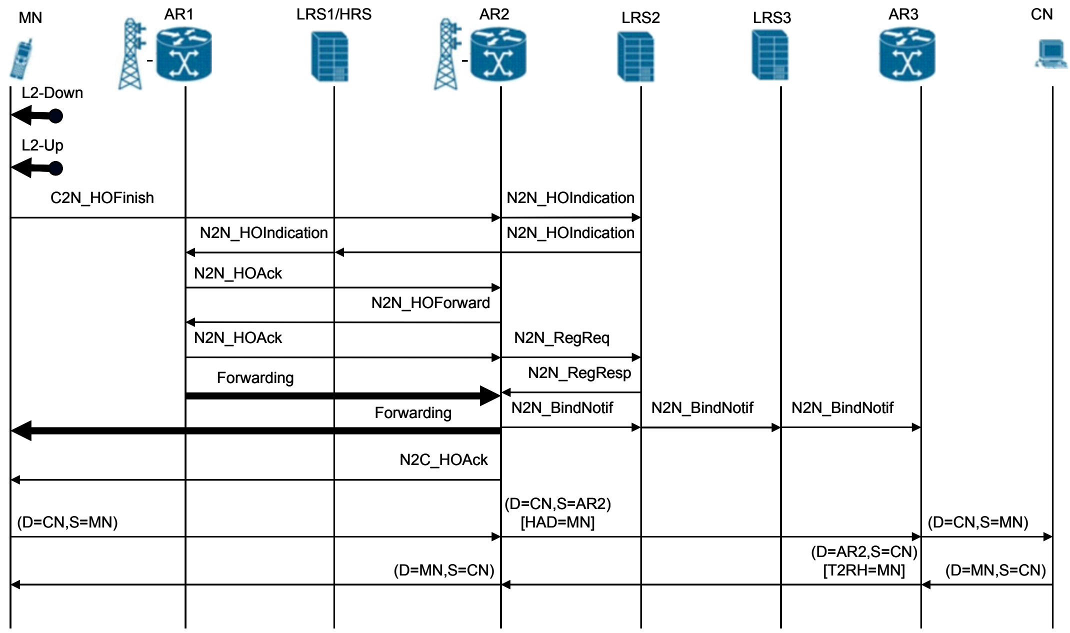

Figure 17 – Handover at the GMA - reactive mode of operation of the GMP

4.4 GMA Advantages

This section presents the main advantages of the GMA over the traditional IP mobility architecture based on MIPv6 and its optimizations (HMIPv6 and FMIPv6).

4.4.1 Simplification of the signaling protocol on the MN

Unlike the traditional IP mobility, the GMA makes a very significant simplification of the signaling protocol on the MN. Only two basic and fundamental operations are executed by the MN: (i) registration at the GMA; and (ii) indication of an eminent handover. The complexity of the mobility management of the MN is transferred to the network functional entities, that is, the AR, the LRS and the HRS. The main benefit is that, the less signaling is assigned to the MNs, the lower is the risk related to errors caused by incorrect signaling and malicious attacks that may compromise the network operation.

Also, in order to speed up the mobility management actions, the priority of the signaling messages generated by the network functional entities may be increased with a lower risk to the system (N2N messages are restricted to network functional elements and should be discarded if originated by MNs).

4.4.2 Optimization of the binding update procedure

Because of the absence of a secure relationship between the MN and the CN in the MIPv6, the MN should execute the *return routability* procedure before sending the *BindingUpdate* message to the CN after a handover. This procedure requires the MN to exchange messages with the CN via the tunnel established between the MN and the HA in order to obtain part of a new security token to be used in the *BindingUpdate* message. As the path between the HA and the MN increases in number of hops, the communication latency also increases, which may disturb real-time, interactive or delay sensitive applications.

The binding update procedure in the GMA is executed by the network functional entities and a secure relationship between the MN and the CN is not required. So, such return routability procedure is also not required in the GAN. The secure relationship is established between the network functional entities as a requisite of their operation at the GMA. Hence, the binding update procedure that takes place at the GMA has a lower overhead than the same procedure in the MIPv6.

4.4.3 Optimization of the registering procedure

The registering procedure executed by the MN with the MIPv6 requires the MN to communicate to the HA, which may be many hops away from the MN. The classical hierarchic solution of HMIPv6 differentiates local registration from regional registration in order to optimize this registration. However, while local registration of the LCoA at the MAP improves the intra-domain handover latency, the regional registration of the RCoA at the HA increases the inter-domain handover latency.

On the other hand, the registration performed by the MN using the GMP is restricted to the GAN. The CoA of the MN is registered at the HRS by the LRS of the domain the MN is located in. This registration does not impact on the MN communication after the handover. Actually, in the anticipated mode of operation, the AR of the GAN is informed about the future presence of the MN (handover context) by its LRS prior to the arrival of the MN. In this mode, the registration of the MN at the new GAN is executed by the new AR in behalf of the MN and in parallel to the L2-handover of the MN.

4.4.4 Optimization of the L3 handover

The most optimized L3 handover procedure of the MIPv6 is specified by the FMIPv6. This procedure has two operation modes, the best of them known as the *anticipated mode*. In this mode, the MN needs to exchange four messages with the AR before the beginning of the handover in order to achieve two goals: (i) to discover the network prefix of the NAR; and (ii) to request the packet forwarding from the PAR to the NAR. Because of the complexity of these operations, it is very difficult to obtain a successful result in the anticipated mode.

In the GMA schema, the network prefix of the access network is constant and there is no need to discover it. So, before the start of the handover process, the MN is required to send only one message to the AR, just to inform about the eminent handover and communicate its handover context. This message is one of the requisites to perform a seamless handover. Because of the simplicity of this procedure, the chance of achieving a successful seamless handover is increased.

All other signaling is performed by network functional entities, including the messages to: (i) transfer the MN´s handover context to the AR of the new GAN and (ii) start the packet forwarding from the AR of the old GAN to the AR of the new GAN. These operations are performed in parallel with the lost and recovery of the L2 connectivity of the MN. So, when the MN arrives at the new GAN, its handover context is expected to be already registered at the AR of the new GAN.