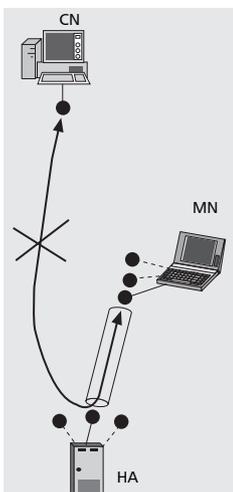


IPv6 MULTIHOMING SUPPORT IN THE MOBILE INTERNET

MARCELO BAGNULO, ALBERTO GARCIA-MARTINEZ, AND ARTURO AZCORRA, UNIVERSIDAD CARLOS III DE MADRID



The authors describe an architecture for IPv6 mobile host multihoming that enables transport layer survivability through multiple failure modes. The proposed approach relies on the cooperation between the MIPv6 and SHIM6 protocols.

ABSTRACT

Fourth-generation mobile devices incorporate multiple interfaces with diverse access technologies. The current Mobile IPv6 protocol fails to support the enhanced fault tolerance capabilities that are enabled by the availability of multiple interfaces. In particular, established MIPv6 communications cannot be preserved through outages affecting the home address. In this article, we describe an architecture for IPv6 mobile host multihoming that enables transport layer survivability through multiple failure modes. The proposed approach relies on the cooperation between the MIPv6 and the SHIM6 protocols.

INTRODUCTION

The integration of fourth-generation (4G) mobile devices to the Internet imposes the adoption of new mechanisms to fully support their multihoming features. The availability of multiple physical interfaces with different technologies in a single device greatly extends their roaming capabilities, enabling a mobile node to preserve the established communications as it moves through areas served by dissimilar access networks. Moreover, the possibility of having multiple paths associated to different technologies enables increased fault tolerance, including the preservation of established communications through different types of outages. In addition, when multiple access technologies are simultaneously available, the mobile node may choose to course different flows through different interfaces, based on cost, quality, or other preferences. However, currently available mobility protocols fail to support the aforementioned features, and specific mechanisms to provide mobile host multihoming support are needed.

In this article we present a mobile host multihoming solution for IPv6 based on the SHIM6

architecture [1] developed by the Internet Engineering Task Force (IETF). The proposed solution consists of end-to-end mechanisms that interact with the available mobile IPv6 protocol [2], enabling the use of multiple addresses (home address and/or care-of address) during the lifetime of an established communication. The end-to-end nature of the proposed solution implies that each mobile device manages its own addresses without relying on any centralized infrastructure. Moreover, transparent support to existing transport protocols, and consequently, existing applications, is guaranteed due to the network-layer nature of the SHIM6-based approach. For example, Voice over Internet Protocol (VoIP) applications layered on top of User Datagram Protocol (UDP), or even requiring both TCP and UDP, can benefit from the extended fault tolerance capabilities.

The remainder of this article is organized as follows. We provide essential background about the Mobile IPv6 (MIPv6) protocol. Next, we present the SHIM6 architecture for IPv6 multihoming, and we illustrate its use through an example. Then, we identify possible mobile host multihoming configurations and the limitations of the MIPv6 protocol to support them. We present the proposed solution that integrates both MIPv6 and SHIM6. We finish by presenting our conclusions.

MOBILE IPv6 FOR 4G MOBILE DEVICES ABOUT MOBILE IPv6

Mobile IPv6 (MIPv6) [2] enables a mobile node to change its attachment point to the Internet while preserving established communications. The main components involved in MIPv6 operation are: the mobile node (MN), originally located in the home network that roams through different visited networks; the home agent (HA) located in the home network; and the correspondent node (CN). The MN has at least one stable address, called the home address (HoA), which is topologically meaningful as long as the MN is located in the home network. When the MN moves away to a visited network, it acquires at

This work was partially supported by Projects RiNG (IST-2005-035167) and IMPROVISA (TSI2005-07384-C03-02).

least one topologically meaningful address at its new location, the care-of address (CoA). However, independent of the MN location, packets addressed to the HoA are routed to the home network. As soon as the MN has left the home network, the MN uses a MIPv6 message called binding update (BU) to inform the HA about its current location, that is, its current CoA. When the HA is aware of the MN location, it tunnels the packets addressed to the HoA to the MN at its present location, that is, the CoA, preserving the communication.

The MIPv6 protocol has two operation modes: the bidirectional tunnel (BT) mode and the route optimization (RO) mode, as depicted in Fig. 1.

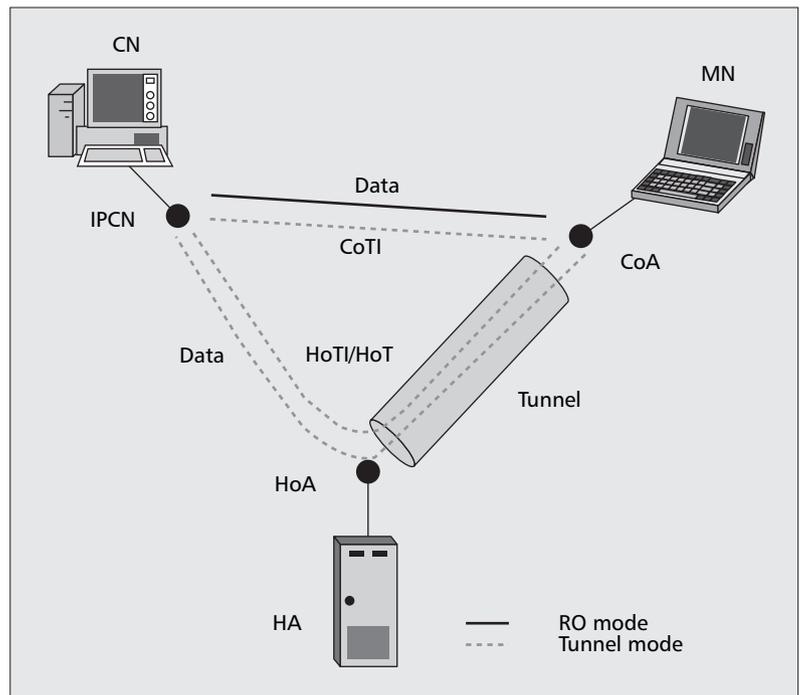
In the BT mode, packets are routed through the HA as long as the MN is away from home, as described previously.

In the RO mode, the MN also informs the CN about its current location, sending it a BU message containing its current CoA. The result is that packets are exchanged directly between the MN and the CN without HA intervention. To protect these BU messages, a security mechanism called a return routability (RR) check is used. The RR procedure consists of the CN exchanging with the MN two different nonces, one through the HoA (using a message exchange called HoTI/HoT) and another one through the CoA (using a message exchange called CoTI/CoT). If the MN can show that it has received both nonces, it can prove that the claimed HoA is co-located with the claimed CoA. To limit the scope of the time of man-in-the-middle attacks, the bindings between a HoA and a CoA that are validated through the RR procedure have a maximum lifetime of seven minutes. After this period, the RR procedure must be executed again to extend the lifetime of the binding.

WHY IPV6?

A legitimate question to ask is why would IPv6 be the right protocol for a mobile host multihoming architecture. In particular, why not use IPv4, especially considering that IPv4 is currently the most widely deployed network layer protocol? The reason why IPv6 must be the protocol used to integrate the new generation of mobile devices in the Internet is two-fold.

On the one hand, only IPv6 can provide the required scalability features, and this conclusion can be obtained by performing a simple calculation. According to International Telecommunication Union (ITU) statistics,¹ there were about 1.7 billion mobile subscribers in the world in the year 2004. This means that a block of 2^{31} addresses is needed to accommodate all of these mobile nodes in the Internet, assuming an address utilization efficiency superior to 80 percent, which is deemed quite hard to achieve [3]. In addition, there are about 2^{30} addresses still unallocated in the Internet Assigned Numbers Authority (IANA) address pool.² The conclusion is that the number of addresses required to integrate the number of mobile devices available in the year 2004 is higher than the number of currently available IPv4 addresses. One could argue that this limitation only applies to public IPv4



■ Figure 1. MIPv6 operation.

addresses and that private IPv4 addresses do not suffer from this shortage. However, to use IPv4 private addresses, network address translation (NAT) traversal techniques are required. It is documented that such techniques impose a signaling overhead that drains battery resources [4], which are critical for mobile hosts.

On the other hand, MIPv6 is more suitable than its equivalent for IPv4 (MIPv4) for the future multihoming mobile Internet because of some key features it provides. In particular, it provides any-to-any route optimization support without requiring the availability of pre-existent security architectures, and it allows direct communications between the MN and CN without passing through the home agent. Additional reasons to prefer MIPv6 over MIPv4 can be found in [5].

MULTIHOMING SUPPORT IN IPV6

SHIM6 ARCHITECTURE

To preserve global routing system scalability, the IPv6 community is advocating the adoption of Provider Aggregatable (PA) addressing. Such an approach forces multihomed sites, that is, sites connecting to the Internet through multiple providers, to obtain multiple provider aggregative prefixes, one from each of their provider's address blocks. Moreover, since Internet service providers (ISPs) only announce their own prefix block to the global routing system, a multihomed host is reachable at a given address only through the corresponding ISP. Consequently, to be reachable through all the available ISPs, a host within the multihomed site must configure as many addresses as prefixes are available in the multihomed site.

Although this set up guarantees the scalability of the multihoming solution, such multi-

¹ ITU, *Mobile Cellular Subscribers Statistics*, <http://www.itu.int/ITU-D/ict/statistics/>.

² Internet Assigned Numbers Authority (IANA), *Internet Protocol v4 Address Space Report*, <http://www.iana.org/assignments/ipv4-address-space>.

The locator agility capability introduced by the SHIM6 protocol requires proper security measures to protect against redirection attacks, where the attacker redirects an established communication between two peers to an alternative locator of its choice.

addressed configuration presents additional difficulties when attempting to provide fault tolerance capabilities. In particular, the preservation of established communication when an outage affects the provider through which the communication is flowing becomes challenging, because to re-home the communication to another ISP, an alternative address must be used to exchange packets. Furthermore, such change of the addresses used during the lifetime of the communication must be performed in a transparent fashion with respect to transport and application layers, to actually preserve the established communication. This is so because current applications and transport layers, such as TCP and UDP, identify the endpoints of a communication through the IP addresses of the nodes involved, implying that the IP addresses selected at the communication establishment time must remain invariant through the lifetime of the communication.

To preserve established communication through outages, a multihoming mechanism located in a SHIM6 layer within the IP layer is proposed [1]. The multihoming mechanism of the SHIM6 layer translates the address used for exchanging packets according to the available providers, while always presenting a constant address to the upper layers of the stack. The result is that the SHIM6 layer performs a mapping between the *identifier* presented to the upper layers and the *locator* actually used to exchange packets on the wire. It should be noted that both nodes involved in the communication must support the mechanism to present a coherent view of the addresses involved in the communication. Both ends use the SHIM6 protocol to exchange the information about the upper layer identifiers and their alternative locator sets, which is stored in each peer in a SHIM6 context.

SHIM6 SECURITY BASED ON CGAs

The locator agility capability introduced by the SHIM6 protocol requires proper security measures to protect against redirection attacks, where the attacker redirects an established communication between two peers to an alternative locator of its choice. In the SHIM6 architecture, protection against redirection attacks can be achieved through the use of cryptographically generated addresses (CGA) [6]. The CGAs are regular unicast IPv6 addresses that incorporate into the 64-bit interface identifier a cryptographic one-way hash of a public key, the prefix of the address, along with other parameters not relevant for our discussion. This structure enables the holder of the associated private key to prove ownership over the claimed CGA address.

To secure the SHIM6 protocol, the addresses that are used as identifiers are generated as CGAs. When a context is established using the SHIM6 protocol, the alternative locator set of this identifier is conveyed to the peer protected by a signature generated with the private key associated with the CGA. The result is that the peer can verify that the owner of the CGA has authorized the use of the alternative locator set.

SHIM6 PROTOCOL WALKTHROUGH

In this section we describe the behavior of the SHIM6 multihoming solution in a common scenario. Consider two SHIM6 hosts, namely host X, holding N different addresses and host Y being configured with M different addresses. As described in the previous section, these addresses are generated as CGAs.

Consider the case in which X starts a communication with Y. Typically, an application in host X issues a domain name system (DNS) request for a name associated to host Y, obtaining in the request some subset of the addresses assigned to host Y. The regular address selection process for IPv6, specified by RFC 3484, is used by host X to select one of the addresses of host Y as the destination address and one of its own addresses as the source address. These addresses selected at the beginning of the communication also are used as identifiers for transport and application layers when required.

After the communication has been established, the SHIM6 protocol is used to create SHIM6 contexts in the peers. The SHIM6 context establishment phase is a four-way exchange through which hosts convey information about the identifiers, the alternative locators available, and related security information. Besides the locators included in the context establishment phase, any of the peers can add new addresses to the session at any time.

After the context has been established, the failure detection mechanism described in the Reachability Protocol (REAP) [7] is used to verify that the currently used path is working (note that alternative locator pairs are not tested). The failure detection mechanism relies on the periodic exchange of packets between the peers. The packet exchange rate is guaranteed by sending SHIM6 keepalive packets only when data is scarce. A failure is detected when one of the peers involved in an active communication stops receiving packets for a certain period of time. Note that it is possible that both peers detect a failure simultaneously.

When a peer detects a failure, it initiates an exploratory phase in which it sends probe packets with different source and destination locators to discover working alternative locator pairs. If the peer had not yet detected the failure, it then starts its own exploratory phase upon the reception of the first probe packet. When the first reply to the probe packets is received, the host selects the associated locator pair as the new working path and diverts the communication through it, preserving the established communication.

After the communication is diverted to an alternative locator pair, a SHIM6 extension header is included in the data packets to incorporate a *context tag*. This context tag allows the receiver to identify the SHIM6 context to be used for restoring the original identifiers.

HOST MULTIHOMING SUPPORT IN MOBILE IPV6

In this section, we first present different multihoming configurations, and then we identify the limitations of the MIPv6 protocol for multihoming fault tolerance support.

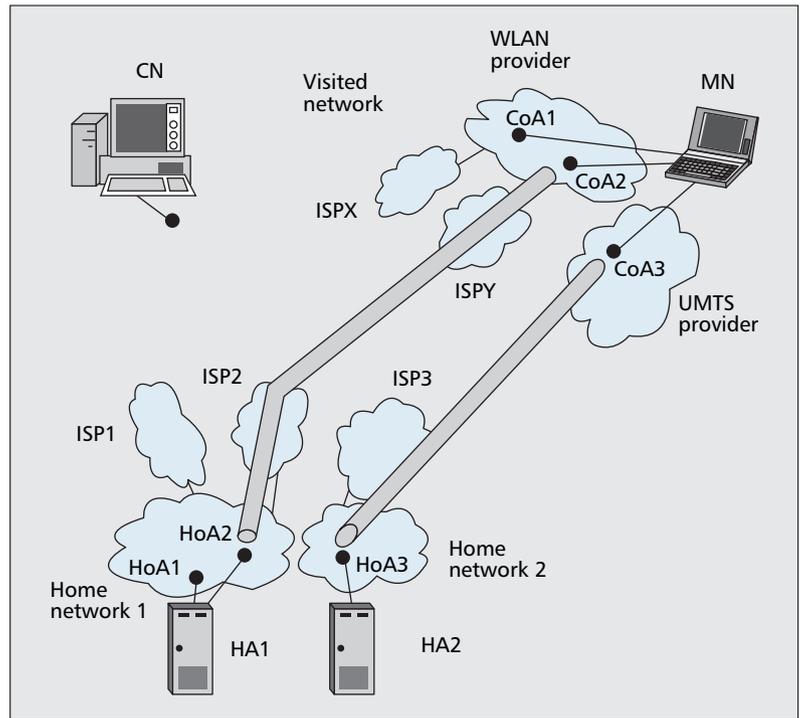
We consider that an MN is multihomed when it has more than one HoA and/or more than one CoA. We can identify the following scenarios where an MN is multihomed, as illustrated in Fig. 2:

- A 4G MN that has multiple physical interfaces, presumably with different access technologies. In this case, the MN may have multiple HoAs, because each interface may have a different home network; and it also may have different CoAs, because each physical interface may be located at a different visited network.
- An MN with a multihomed home network. In this case, the home network is connected to multiple ISPs, each of which delegates a prefix, resulting in multiple HoAs, one per prefix.
- An MN that is roaming in a multihomed visited network. As in the previous case, when a visited network is multihomed, multiple prefixes are available. Therefore, an MN visiting the multihomed network has the possibility of configuring multiple CoAs.

A feature that is common to all of the identified multihoming configurations is the availability of multiple paths between the MN and the CN. The existence of multiple paths enables extended fault tolerance, because in the case of a failure, the communication can be preserved by using an alternative path. However, as we describe next, current MIPv6 protocol fails to provide full fault tolerance capabilities, because failures may affect ongoing communications even though alternative working paths are available.

In the case of a multihomed MIPv6 node with multiple HoAs and multiple CoAs that is communicating in BT mode (Fig. 3a), it is trivial to see that a failure affecting the reachability to the HoA would break the established communication. This is true even in the case that other reachable HoAs are available, because MIPv6 does not provide support for changing the HoA used for an established communication. In summary, in BT mode, a failure in the path between the CN and the MN through the HA affects any communication established using the corresponding HoA, even if there are other working HoAs available.

In the case of an MN with multiple HoAs and multiple CoAs that is communicating with a CN in RO mode (Fig. 3b), it also is clear that a failure in the path between the CN and the CoA used for the communication would affect the ongoing communication. The MIPv6 protocol could provide means to survive this outage if it could detect it and try to use an alternative CoA, or fall back to the path through the HoA. Because MIPv6 does not have any mechanism to detect this type of outage, the communication will be interrupted in this case. In addition, the established communication is not only vulnerable to outages in the path used to exchange data packets, but it is also vulnerable to failures in the path between the CN and the MN through the HA (Fig. 3c). This is so because the path through the HA is used to periodically exchange HoT/HoTI packets. In case an outage affects this path, the HoT/HoTI packet exchange would be interrupted. The result is that the binding



■ **Figure 2.** Mobile node multihoming scenarios.

between the HoA and the CoA in the CN will expire [8], and the communication will fall back to the path through the HA, which is not working. So, in RO mode, ongoing communications are not only vulnerable to failures in the path between the CN and the currently used CoA, but they are also vulnerable to outages in the path between the CN and the MN through the HA.

PROPOSED MOBILE HOST MULTIHOMING SUPPORT ARCHITECTURE

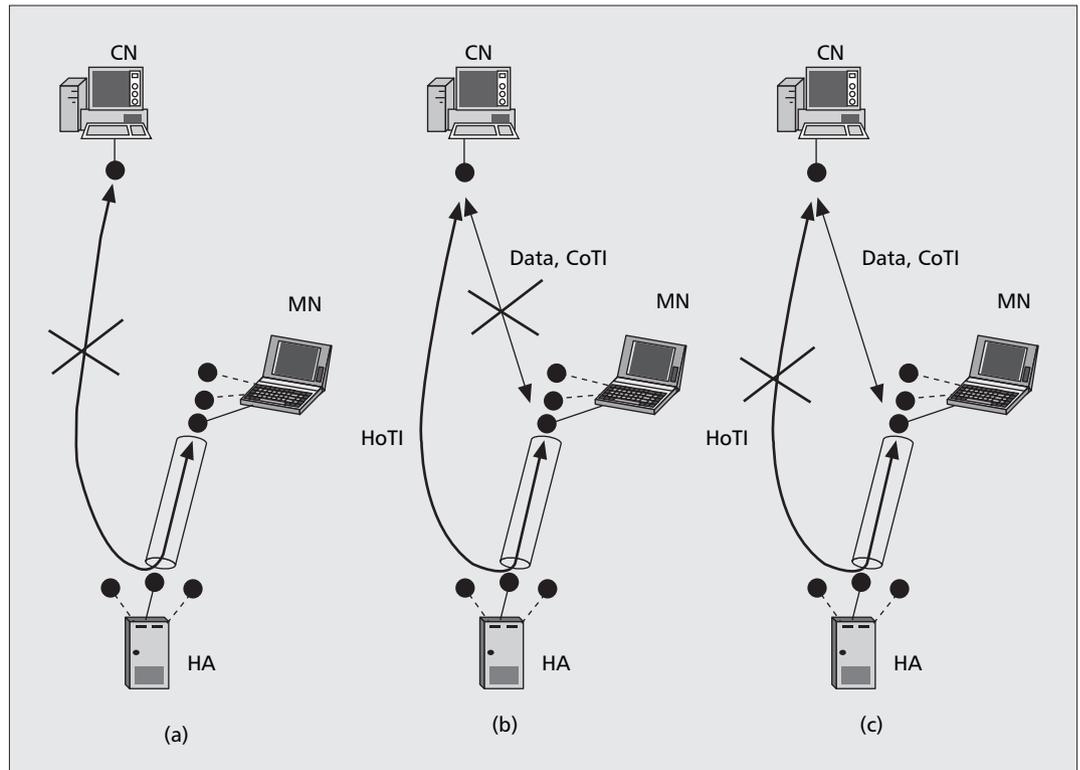
PROPOSED ARCHITECTURE

In this section, we describe how to integrate the MIPv6 and SHIM6 protocols, without requiring any modification to the protocol messages in any of them, to provide fault tolerance capabilities to multihomed mobile devices.

In this configuration, both the MN and the CN include in their stack a SHIM6 and a MIPv6 module (Fig. 4). In both the MN and the CN, the SHIM6 layer is located below the IP endpoint sublayer, that is, the sublayer within IP that performs end-to-end functions like fragmentation. The MIPv6 mechanisms are placed underneath the SHIM6 layer and on top of the IP forwarding sublayer, that is, where the forwarding functions of the IP layer are situated.

One of the most remarkable aspects of this architecture is related to the management of the diverse name spaces involved. According to what we have presented earlier, the transport and application protocols located on top of the SHIM6 layer use identifiers to name the communicating peers. Those identifiers are IPv6 addresses that are selected by the applications to initiate the communication. The SHIM6 layer

If we put all of this together, we see that the upper layer protocols standing above the IP layer use a given address to identify the parties involved in the communication. This address is used by the SHIM6 layer as an identifier.



■ **Figure 3.** Failure scenarios for a multihomed MN: a) tunnel mode; b) route optimization node (data path failure); c) rout optimization mode (control path failure).

creates a context state that stores alternative locators that can be used to reach the identifier of this context.

As currently defined, the MIPv6 protocol creates a single binding cache entry (BCE) to a particular CoA for a given HoA. As available CoAs change, the BCEs are changed accordingly. In this case, the MIPv6 layer translates the HoA used by the protocol located above (the SHIM6 protocol) to the associated CoA.

In the particular case of the SHIM6 running on a MIPv6 MN, the local identifiers available to the upper layer protocols are likely to be the HoAs, because they are stable addresses and susceptible to be published in the DNS, and the alternative locators are likely to be the available HoAs and CoAs.

So, if we put all of this together, we see that the upper layer protocols standing above the IP layer use a given address to identify the parties involved in the communication. This address is used by the SHIM6 layer as an identifier. The SHIM6 layer may be required to translate this identifier to an alternative locator if the identifier is not working as a locator (e.g., because of a failure). The locator selected by the SHIM6 layer may be a HoA or a CoA. When the locator is a CoA, it is not processed by the MIPv6 layer, and it is directly included in the actual IPv6 address field. When the locator selected by the SHIM6 layer is a HoA, the MIPv6 layer performs a mapping between the selected HoA and the CoA currently associated to the HoA. This CoA is the address that is included in the address field of the forwarded packet.

RESULTING BEHAVIOR

To illustrate the operation of the proposed approach, we consider the following scenario:

- An MN with multiple CoAs (CoA_1, \dots, CoA_n) and multiple HoAs (HoA_1, \dots, HoA_m)
- A CN with a single address IPCN

In this scenario, the MN establishes a communication with the CN. This communication uses one of the HoAs available in the MN (e.g., HoA_i) as the upper layer protocol identifier of the MN and the address of the CN (IPCN) as the identifier of the CN.

After the communication has been established, the SHIM6 layer decides through some heuristics (such as elapsed time of communication or number of packets flowing) to create a SHIM6 context to protect that communication. The SHIM6 context is established between the MN and the CN. In this case, the SHIM6 layer uses the selected upper layer identifier (e.g., HoA_i) as the SHIM6 identifier of the MN and includes the alternative addresses, HoAs and CoAs, as alternative locators for this identifier. For the CN, the SHIM6 identifier is IPCN and the only available locator is IPCN itself. After the context is established, the SHIM6 layer uses the REAP protocol to detect possible outages.

In addition, as soon as the MN leaves the home network, the MIPv6 layer of the MN creates a binding between HoA_i and one of the CoAs, available at the visited network (e.g., CoA_p). The MN notifies about the binding through a BU message to the HA (and to the CN in the case of RO mode). It also may update the binding for other HoAs with the available CoAs.

The resulting state is the following:

- Upper layer protocols: a communication is established between the HoA_i and IPCN.
- The SHIM6 context has the following information:
 - IPCN as the CN identifier and HoA_i as the MN identifier.
 - IPCN as the only available locator for IPCN and (HoA₁, ..., HoA_m) and (CoA₁, ..., CoA_n) as the available locators for HoA_i.
- The MIPv6 layer has a BCE binding HoA_i to CoA_p. It also may have other BCE bindings for other HoAs.

So during the whole lifetime of the communication, the application uses HoA_i and IPCN as identifiers. As long as there is no outage, the SHIM6 layer does not perform any transformation, and the MIPv6 layer uses the CoA to reach the MN, that is, it will transform the HoA_i to CoA_p (note that in BT mode the CoA_p is included in the address field of the outer header of the tunnel, and in RO mode, the CoA_p is included in the address field of data packets that also carry information about the HoA_i in the destination option).

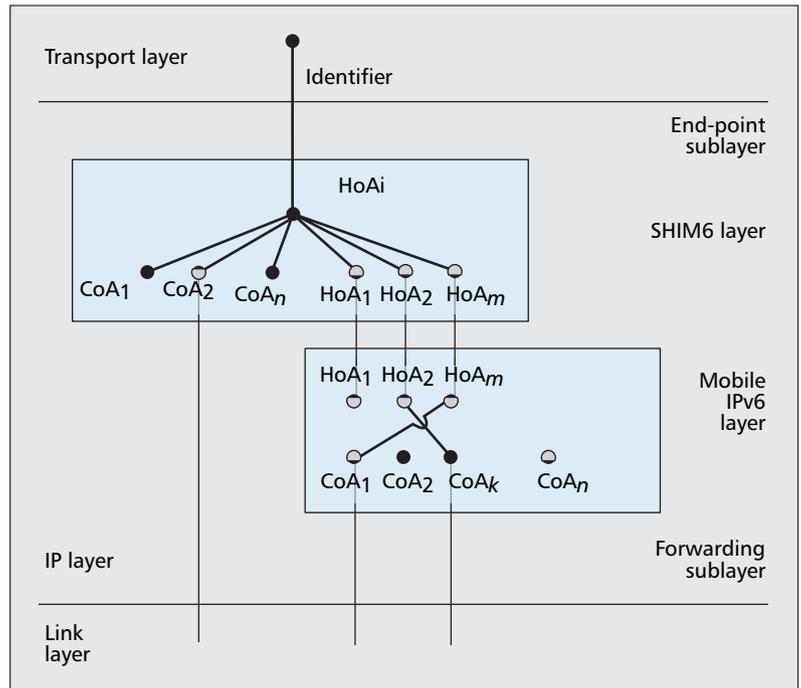
As a case study, we next consider the response of the proposed approach to a failure when RO mode is being used for the communication between the MN and the CN.

Suppose that a communication is established between the MN and the CN using HoA_i and IPCN. In addition, through MIPv6 protocol, a BCE is created in the CN associating the HoA_i with one of the CoAs, CoA_p. So, packets associated with the communication are flowing directly between the CN and the MN carrying CoA_p and IPCN in the source and destination address fields.

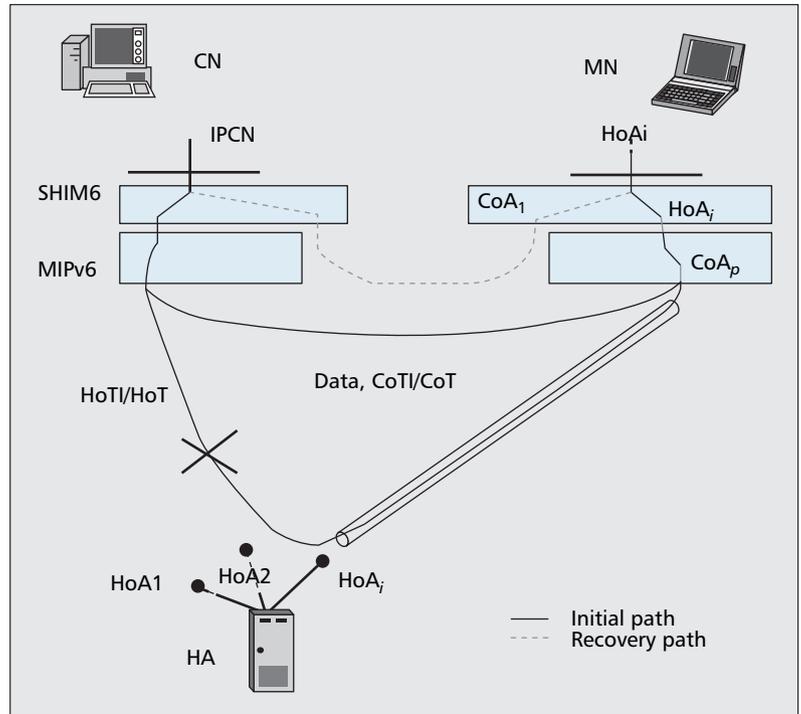
We next analyze how this configuration reacts to different failure modes.

Consider the case where the path between IPCN and CoA_p fails. The SHIM6 detects the outage and tries with alternative locators available in the SHIM6 context. If an alternative HoA is selected by the SHIM6 layer as an alternative locator, when the SHIM6 layer passes the packet with an alternative HoA to the MIPv6 layer, the MIPv6 layer will route the packets through the corresponding CoA available in the BCE associated with the new HoA, possibly falling back to BT mode, but potentially recovering the failure. If an alternative CoA is used by the SHIM6 layer as an alternative locator, the MIPv6 layer will not translate the alternative CoA (because there is no BCE for the CoA), and packets will be routed directly between the MN and the CN, in a kind of SHIM6-based RO mode.

Consider next the case where the path between the MN and the CN through the HA fails (Fig. 5). While data traffic is not routed through the HA, HoTI/HoT packets are exchanged through the HA. If the path between the MN and the CN through the HA fails, then the HoTI/HoT exchange will fail. A few minutes later, the corresponding BCE will expire, and the communication will fallback to the BT mode through the HA. However, because we are considering the case where the path through the HA is down, then the com-



■ Figure 4. Mobile node multihoming architecture.



■ Figure 5. Response to a failure in RO mode.

munication will definitely fail. At this point, SHIM6 will detect the outage and use an alternative locator pair. Analogously to the previous case, SHIM6 can try with an alternative CoA (as depicted in Fig. 5) or an alternative HoA as alternative locators for the communication. In any case, similar considerations to the ones described previously apply, and the communications will be restored, whether in BT mode (alternative HoA) or in a SHIM6-based RO mode (alternative CoA).

CONCLUSION

In this article, we presented an architecture for the provision of multihoming support to 4G mobile nodes. Such architecture enables the preservation of established communication through outages. While the preservation of established communication through failures affecting the CoA may seem quite straightforward to achieve through simple extensions to the MIPv6 protocol, the preservation of communication in case of an outage affecting the HoA is a much more complex problem, because fundamental parts of the MIPv6 protocol are built upon the underlying assumption that the HoA is always reachable. While these assumptions may hold true for single-homed mobile devices, it is not the case for multihomed mobile hosts. The proposed architecture overcomes these limitations, enabling the preservation of established communication across outages affecting the CoA and/or the HoA without any modifications to the MIPv6 protocol, but using a standard multihoming support mechanism on top of it. Such an approach substantially reduces the complexity of the resultant solution, because it minimizes the changes required to available protocols.

REFERENCES

- [1] E. Nordmark and M. Bagnulo, "SHIM6: Level 3 Multihoming Shim Protocol for IPv6," Internet draft, draft-ietf-shim6-08, May 2007.
- [2] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6," IETF RFC 3775, June 2004.
- [3] A. Durand and C. Huitema, "The Host-Density Ratio for Address Assignment Efficiency: An Update on the H Ratio," IETF RFC 3194, Nov. 2001.

- [4] "S60 Platform: Application Networking Considerations," Nokia Corp., May 2006.
- [5] S. J. Vaughan-Nichols, "Mobile IPv6 and the future of Wireless Internet Access," *IEEE Comp.*, vol. 36, no 2, Feb. 2003, pp. 18–20.
- [6] T. Aura, "Cryptographically Generated Addresses (CGA)," *6th Info. Sec. Conf.*, Bristol, U.K., Oct. 2003.
- [7] J. Arkko and I. van Beijnum, "Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming," Internet draft, draft-ietf-shim6-failure-detection-07, Dec. 2006.
- [8] G. Huston, "Architectural Approaches to Multihoming for IPv6," IETF RFC 4177, Sept. 2005.

BIOGRAPHIES

MARCELO BAGNULO (marcelo@it.uc3m.es) received his Ph.D. in telematics engineering from Universidad Carlos III de Madrid, Spain, in 2005. He is an associate professor in the Telematics Engineering Department of the same university. He is involved in the standardization process in the IETF, authoring several RFCs and working group drafts. He has participated in several international research projects on IPv6 funded by the European Union, such as RING and 6LINK.

ALBERTO GARCIA-MARTINEZ (alberto@it.uc3m.es) received his Ph.D. in telematics engineering from Universidad Politécnica de Madrid, Spain, in 1999, and was awarded the Fundacion Telefonica Award on Telecommunication Network and Services. He is an associate professor in the Telematics Engineering Department of the Universidad Carlos III de Madrid. He has participated in several research projects on IPv6 funded by the European Union such as LONG and 6LINK.

ARTURO AZCORRA [SM] (azcorra@it.uc3m.es) received his telecommunication engineering degree in 1986 and his doctorate in 1989 from the Universidad Politécnica de Madrid. He is currently a full professor at Universidad Carlos III de Madrid. He has been appointed as chief scientific advisor of the international research institute IMDEA Networks. He has participated in a large number of European research projects, having coordinated the European Networks of Excellence E-NEXT and CONTENT.