

Handover Management for Mobile Nodes in IPv6 Networks

Nicolas Montavont and Thomas Noël

LSIIT — Louis Pasteur University — CNRS, Strasbourg

ABSTRACT

In this article we analyze IPv6 handover over wireless LAN. Mobile IPv6 is designed to manage mobile nodes' movements between wireless IPv6 networks. Nevertheless, the active communications of a mobile node are interrupted until the handover completes. Therefore, several extensions to Mobile IPv6 have been proposed to reduce the handover latency and the number of lost packets. We describe two of them, Hierarchical Mobile IPv6, which manages local movements into a domain, and Fast Handover protocol, which allows the use of layer 2 triggers to anticipate the handover. We expose the specific handover algorithms proposed by all these methods. We also evaluate the handover latency over IEEE 802.11b wireless LAN. We compare the layer 2 and layer 3 handover latency in the Mobile IPv6 case in order to show the saving of time expected by using anticipation. We conclude by showing how to adapt the IEEE 802.11b control frames to set up such anticipation.

INTRODUCTION

Nowadays, wireless network access is increasingly popular since wireless communication offers interesting advantages: it allows movements during communications and network access at a fair rate among nodes. The movement of mobile nodes (MNs) between access points (APs) belonging to a common subnet is managed by the layer 2 (L2) protocol and does not involve layer 3 (L3) mechanisms. On the other hand, if an MN connects to an AP in another subnet, the IPv6 address of the MN is not topologically valid anymore. Therefore, this kind of movement has to be managed by a specific L3 protocol.

Mobile IPv6 [1] is designed to manage MNs' movements between wireless IPv6 [2] networks. The protocol provides unbroken connectivity to IPv6 MNs when they move from one wireless point to another in a different subnet, an operation known as an L3 handover.

Nevertheless, an MN cannot receive IP packets on its new point of attachment until the handover ends. This time includes the new prefix discovery on the new subnet, the new care-of address establishment, and the time needed to notify the correspondents and home agent about the new locality of the MN. This time is called *handover latency*.

Actually, handover latency can be too long for real-time multimedia applications. In most cases, the impact of handover latency strongly degrades the IPv6 stream of MNs. Therefore, there are many extensions to MIPv6 and new protocols proposed to improve the IPv6 connectivity of MNs. The aim of these proposals is to reduce the latency and the number of lost packets due to handover between one point of attachment to another [3], and the signaling load on the MIPv6 home agent and the correspondent nodes [4].

The aim of this article is to present three protocols that manage MN movements: Mobile IPv6 [1]; Hierarchical Mobile IPv6 [4], which optimizes movement in an administrative domain; and Fast Handover Protocol [3], which anticipates the movement to start the handover earlier. These protocols are studied in the next section. Then we propose an evaluation of Mobile IP over IEEE 802.11b wireless LAN [5] in a later section. We measure both L2 and L3 handover latencies. In the following section, we discuss the difference between L2 and L3 handover, particularly the saving of time expected when L3 handover is done by anticipation. Finally, we give some concluding remarks.

PROTOCOL OVERVIEW

Currently, the Internet Engineering Task Force (IETF) presents three main protocols to manage MN movements. Mobile IPv6 [1] allows the MN to acquire and register a new IPv6 address in each visited network, but the time involved in these operations can be long. Hierarchical Mobile IPv6 [4] focuses on local movements by reducing the signaling load on the network. The Fast Handover Protocol [3] provides anticipation

by using L2 events to initiate operations in advance. We present these three protocols in the following subsections. The terminology used in this article is taken from [6].

MOBILE IPv6

Mobile IPv6 [1] is designed to manage MNs' movements between wireless IPv6 networks. The architecture of a wireless access network is shown in Fig. 1. When an MN remains in its home network, it communicates like another IPv6 node with its correspondent(s). When a MN moves to a new point of attachment in another subnet, its home address is not valid anymore, and packets sent by its correspondent(s) will continue to reach its home network. Therefore, it needs to acquire a new valid address in the visiting subnet, called the care-of address, and registers it with its home agent and correspondent(s). The association made between the home address and the current care-of address of an MN is known as a *binding*. Henceforth, the home address always identifies the communication of an MN, and the care-of address locates the MN.

The Handover Procedure — An MN detects that it has moved to a new subnet by analyzing the *router advertisement* periodically sent by the access router (AR). The MN can also request the AR to send a router advertisement by sending a *router solicitation*. The information contained in the router advertisement will allow the MN to create a new care-of address. As specified in IPv6 [2], the MN first needs to verify the uniqueness of its link-local address on the new link. The MN performs duplication address detection (DAD) on its link-local address. Then, it may use either stateless [7] or stateful [8] address autoconfiguration to form its new care-of address. Once it has obtained a new care-of address, it may perform DAD for it. However, DAD takes quite a long time with respect to the handover latency. Actually, in order to perform DAD, the MN has to send one or several *neighbor solicitation(s)* to its new address and wait for a response for at least 1 s. This implies important additional time to handover latency. For this reason, the MN should perform DAD in parallel with its communications, or choose not to perform it.

Once the new care-of address construction is done, the MN must update the binding cache in its home agent and correspondent(s) by sending a *binding update*. The MN can request an acknowledgment by setting a specific bit in its message (this bit must be set in the binding update intended for the home agent).

Handover Enhancement — MIPv6 already provides some enhancements to the handover procedure. In some cases, an MN can be reachable through multiple wireless links from physically neighboring APs. If these APs are on different subnets, the MN can configure a care-of address for each of them. One of these care-of addresses must be a primary care-of address for a default AR that will be registered in the MN home agent and correspondent(s). Then, when the default AR becomes unreachable, the

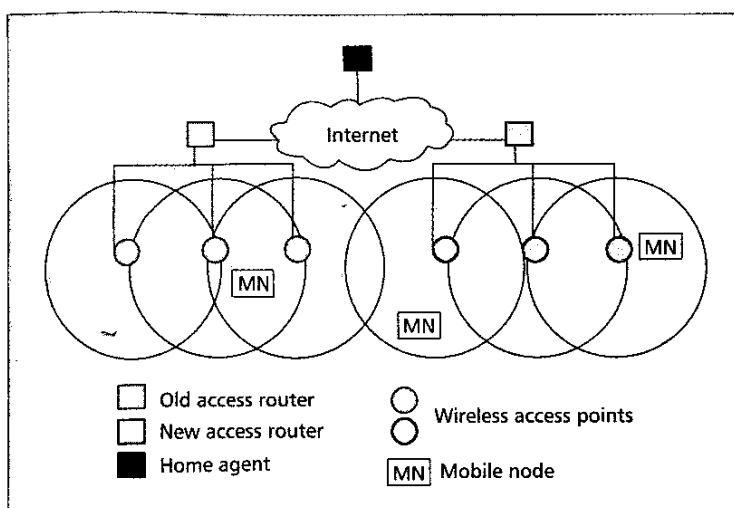


Figure 1. IPv6 wireless network architecture.

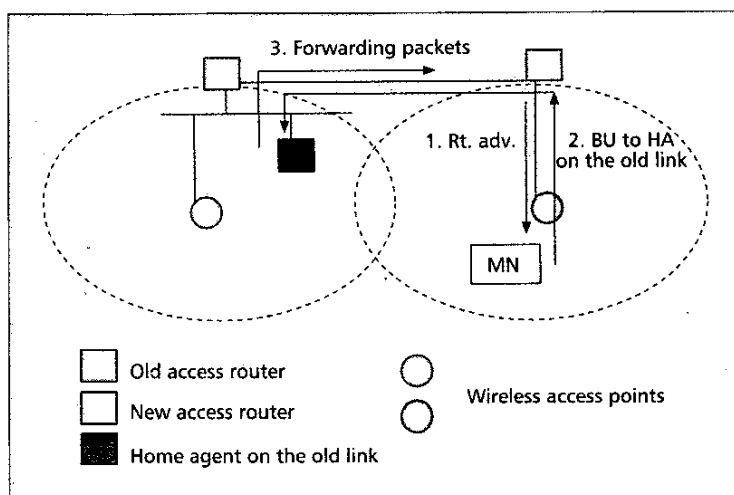


Figure 2. Forwarding from the old access network.

MN can use a new default AR for which it already has a care-of address.

In addition, the packets sent by the correspondent nodes are lost until they receive the binding update indicating the new care-of address of the MN. To reduce the number of lost packets during this time, the MN can request the old AR to forward all its incoming packets to the new AR. To do so, the MN has to send a binding update to a home agent on its old link indicating its new care-of address, but with its old care-of address instead of the home address. Then, the home agent on the old link intercepts the packets intended to the old care-of address of the MN and forwards them to the current localization of the MN (Fig. 2).

Depending on the MN's movements, an MN can switch between two ARs several times (this kind of movement is usually called *ping-ponging*). In this case, Mobile IPv6 requires that the MN create and register a new care-of address after each movement. Bicasting allows the MN to simultaneously register with several ARs. All the packets intended for the MN are then duplicated

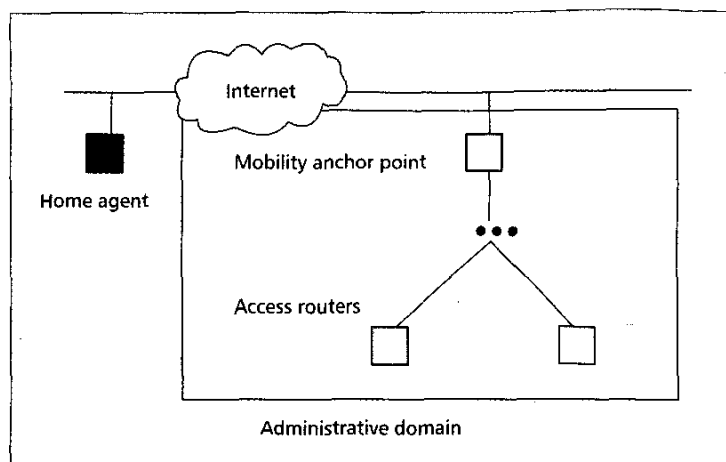


Figure 3. Hierarchical architecture.

in several potential localizations. This solution is very interesting, particularly if the multiple associations could be set up by anticipation (see a later section). However, the bicasting performed by the home agent is not scalable and generates lots of traffic on both the wired and wireless links. We will see in the next subsection how it is possible to implement local bicasting to scale this solution.

HIERARCHICAL MOBILE IPV6

Flat Mobile IPv6 requires that the MN send a binding update to each of its correspondents. According to their localization, the time to reach them and the signaling load generated can be very important. Hierarchical Mobile IPv6 [4] is designed to minimize the amount of signaling to correspondent(s) and to the home agent by allowing the MN to locally register in a domain.

The global Internet is divided in regions defining local area mobility [6]. These domains are independent from subnets and are generally managed by a unique administrative authority (e.g., a campus). Each domain is connected to the rest of the Internet by a *mobility anchor point*, which acts like an anchor point for the MN. The mobility anchor point is an AR with a publicly routable IP address at the top of several ARs (Fig. 3).

When the MN first enters a domain, it needs to make a *regional registration* to advertise to its home agent and correspondent(s) its new raw localization. It indicates a global care-of address for the domain (see next paragraph). Later, after each movement between ARs in the same domain, the MN needs to send a *local registration* to the mobility anchor point to update its localization into the domain (on-link care-of address). Thus, all MN movements within the domain are hidden from the home agent and correspondent(s) since the global care-of address of the MN does not change.

The Two Modes of Hierarchical MIPv6 — The mobility anchor point is announced in the *agent advertisement* messages sent by the AR of the domain. When an MN enters a visited

domain for the first time, it must perform a home registration. Next, when it moves within this domain, the MN can choose between *basic mode* and *extended mode*. In basic mode, the MN has two addresses: a regional care-of address based on the mobility anchor point prefix and an on-link care-of address based on the current AR prefix. In this scheme, the mobility anchor point acts as a home agent: it intercepts the packets destined to a regional care-of address and tunnels them to the corresponding on-link care-of address. These operations are totally transparent to the MN home agent, which does not need any modification.

However, not every MN can acquire an individual regional care-of address because of scalability or a network operator policy. In extended mode, the regional care-of address is (one of) the mobility anchor point address(es). The mobility anchor point keeps a binding table with the current on-link care-of address of an MN matched with the MN home address. When it receives the packets destined to an MN, it detunnels and retunnels them to the on-link care-of address. This implies that each packet must contain the MN home address.

Bicasting in Hierarchical Architecture

The bicasting done by the home agent, presented in an earlier section, is not scalable and can generate too much delay in packet delivery. The hierarchical model allows bicasting from the mobility anchor point. When an MN moves within a domain, it can request bicasting in its local registrations. This request is forwarded to the mobility anchor point, which adds a new entry for the MN (simultaneous bindings). Then the mobility anchor point forwards the same traffic to the old and new MN localizations. When bicasting is performed in this way, the packets are only duplicated within the domain.

However, the problem of scalability is not resolved if the mobility anchor point handles too many MNs. Considering several mobility anchor points per domain that are at the same level could resolve the scalability problem, since these mobility anchor points could share the number of MNs. However, this method is still under discussion since it causes some problems: discovery of the other mobility anchor point(s), selection of one mobility anchor point by the MN, and load balancing among multiple mobility anchor points.

FAST HANDOVER

The Fast Handover Protocol [3] is an extension of Mobile IPv6 that allows an AR to offer services to an MN in order to anticipate the L3 handover. The movement anticipation is based on the L2 triggers [9]. An L2 trigger is information based on the link layer protocol, below the IPv6 protocol, in order to begin the L3 handover before the L2 handover ends. An L2 trigger contains information on the MN L2 connection and on the link layer identification of the different entities (e.g., the link layer address). The main L2 triggers used are the following:

- **Link Up:** indicating that the MN has established a connection with an access point

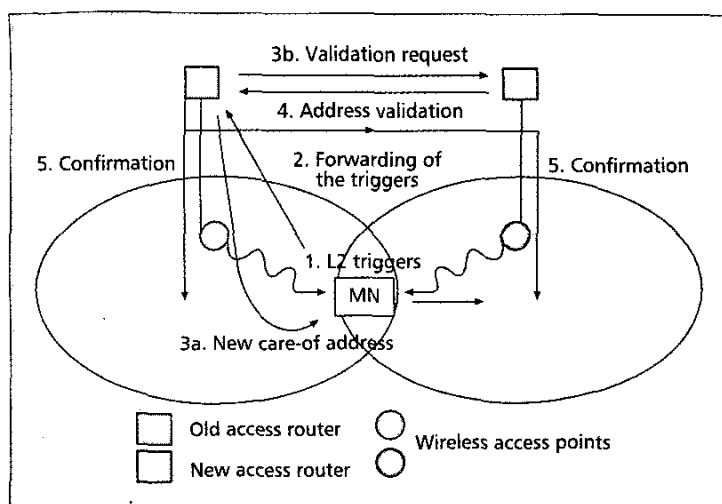
- **Link Down:** indicating that the MN has lost a connection with an access point
- **L2 Handover Start:** indicating that the MN starts an L2 handover to attach to a new access point

When an AR receives an L2 trigger, it must be capable of matching entity identification to an IP address. For example, when it receives access point identification, it must know to which subnet this access point belongs. To do so, the neighboring ARs have to exchange information to discover each other [10]. The information exchanged can be a network prefix or a list of the access points operating in an AR subnet.

Anticipated Handover — Fast Handover uses these L2 triggers to optimize the MN movements in two methods: anticipated handover and tunnel-based handover. In anticipated handover, the MN or the current AR (when L3 handover is controlled by the network) receives an L2 trigger indicating that the MN is about to perform an L2 handover (steps 1 and 2, Fig. 4). This trigger must contain information allowing the target AR identification (e.g., its IPv6 address). If the MN receives the L2 trigger, it must initiate the handover and request fast handover to its AR. The current AR then sends a valid IPv6 address for the new subnet to both the MN (step 3a) and the target AR for validation (step 3b). Then the target AR controls if the address is unique in its subnet [7] and sends the validation result to the current AR (step 4). If the address is valid, the current AR forwards the authorization (to use this address in the target subnet) to the MN in both subnets (step 5). Then when the MN establishes the connection with the new access point, it can immediately use the new care-of address as the source address in the outgoing packets and send a binding update to the home agent and correspondent(s). To minimize the loss of packets, the old AR forwards all the packets intended to the MN to the new AR.

Tunnel Based Handover — In tunnel-based handover [3], the MN delays the new care-of address establishment when it moves to a new AR. Therefore, it only performs an L2 handover and continues to use its old care-of address in the new subnet. Moreover, the MN does not need to exchange any packets: the two ARs set up a bidirectional tunnel from the L2 triggers without interacting with the MN. The packets intended for the MN reach the old subnet where they are captured and forwarded to the new AR by the old AR. The outgoing packets of the MN take the reverse path from the new AR to the old AR, which forwards them in the Internet. Later, the MN will create and register a new care-of address in parallel with its communications. Otherwise, if the MN moves quite fast, the tunnel would be extended to a third AR (*handover to a third*).

The use of L2 triggers allows the AR to detect MN movement without the need to send any packets. This is very interesting since the cost to send a packet on a wireless interface is more expensive than on the wired interface.



■ Figure 4. Anticipated handover.

MOBILE IPV6 EVALUATION OVER IEEE 802.11B

In this section we evaluate Mobile IPv6 over wireless LAN. We do not consider Hierarchical Mobile IPv6 [4] in our measurements since we focus on interaction with the L2 protocol. We also only consider a single local correspondent communicating with our MN. Therefore, in this testbed, the handover latency in Hierarchical Mobile IPv6 would be the same as in Mobile IPv6.

IEEE 802.11b [5] aims to manage wireless communications. This is the most used norm in wireless LAN, and several products are already available. To ensure interoperability among different product vendors, the specification defines a radio propagation model interface, an encoding and modulation method, and a medium access control (MAC) layer.

We mainly focus on IEEE 802.11b in order to have an overview of the real MN possibilities over wireless LAN. IEEE 802.11b access points allow communications within their cover area at a configurable rate of 1, 2, 5.5, or 11 Mb/s. We measure the L2 and L3 handover latencies for each bandwidth.

THE TESTBED

We consider a single MN ping-ponging between two access points in two different subnets: the home network and a visited network. Whole measurements are taken in an optimized case, where the MN is the only user, and in a more realistic case with four additional static active users. Each time we measure L2 handover latency (the disruption time to establish the new connection after disconnecting from the old access point) and L3 handover (L2 handover plus the time needed to acquire and register a new valid care-of address). Especially for L3 handover, we take into account the time for a local correspondent to redirect the traffic to the new localization of the MN. The results are presented in Fig. 5. The ARs are configured to send a router advertisement every second.

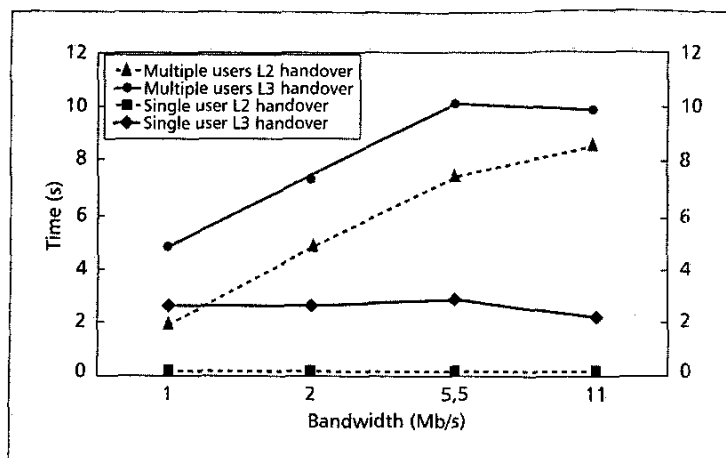


Figure 5. L2 and L3 handover latencies.

ANALYSIS

The averages of our measurements are shown in Fig. 5. The red curves represent the measurements for a single MN, and the blue curves represent the multiuser case. For each case, the dotted line is L2 handover latency and the plain line is L3 handover latency.

First, we notice an important difference between the optimized case and the most realistic case at the L2 level: when the moving MN is alone, L2 handover is quite constant around 0.158 s, while when there are other active users, first the L2 handover increases with the bandwidth, and on the other hand the values are greater: from 1.754 s at 1 Mb/s up to 8.618 s at 11 Mb/s. This is because the data frames and control frames share the same single channel of the access point cover area. Then when there are several active users, collisions are more frequent, and the time to access the channel is longer. Therefore, the MN needs more time to synchronize with the target access point. We also observe that at faster bandwidths the error rate is higher, since we saw more erroneous frames in the trace files.

Second, L3 handover seems to be independent of the access rate, except at 11 Mb/s. We observe that L3 handover is the addition of L2 handover latency and a certain constant. The additional time encompasses the new prefix discovery (contained in the router advertisement), and the time to create, validate, and register a new care-of address. Although this additional time strongly varies from one measurement to another (the variance is significant), the average of all the measurements is closed for the different rates.

When the moving MN is alone, total L3 handover latency is around 2.54 s. Otherwise, when there are several active users, total L3 handover latency ranges from 4.772 s at 1 Mb/s up to 9.731 s at 11 Mb/s. In the most realistic case, L2 handover latency strongly contributes to L3 handover latency, especially for highest bandwidths. However, Mobile IPv6 does not anticipate actions before L2 handover ends, while L2 handover can take up to 8 s. If L3 handover could be initiated during L2 handover, L3 handover latency could

be strongly reduced. This kind of anticipation is proposed in the Fast Handover Protocol exposed earlier with the L2 triggers. The feasibility of L2 triggers over IEEE 802.11 is studied in the next section, as well as the expected results on L3 handover latency.

APPLICATION OF ANTICIPATION

As we explained earlier, L3 handover can be initiated in advance of L2 connection with a new access point by the L2 triggers. The L2 triggers we considered are *link up*, *link down*, and *L2 handover start*. We present here how the IEEE 802.11b control frames can be mapped to these L2 triggers. We then conclude on the realization of the Fast Handover [3] and approximate the saving of time.

IEEE 802.11b ROAMING

When the signal strength between an MN and its current AP drops under a predefined threshold, the MN starts an L2 handover. Thus, it sends an IEEE 802.11b control frame called a *probe request* to an L2 broadcast address. All the APs that hear this message reply with a *probe response*. Then the MN chooses one depending on the characteristics contained in the reply. Subsequently the MN and the target AP continue to exchange probe requests and probe responses to synchronize. Finally, the MN sends an authentication and requests the association.

The first exchange of probe request and probe response indicates that the MN may establish a new connection. These messages can be considered an *L2 handover start*. Moreover, since an MN cannot communicate once it starts an L2 handover, these messages can also be considered a *link down*. On the other hand, the information contained in these control frames is not sufficient for the requirements of L2 triggers. Actually, the L2 trigger must contain an identification of the MN and an identification of the two APs. To do so, the two MAC addresses of the old and new APs must be added in the probe request sent by the MN. This information allows the receiving AP to forward the L2 trigger to its AR.

When an AR receives the L2 trigger, it must be capable of identifying the subnet of the other AP. This identification can be achieved by a beforehand exchange between the neighboring ARs [10]. The neighboring ARs exchange characteristics on their own subnets, including the list of the MAC addresses of the APs operating in their subnet.

MOBILE IPV6 LIMITS

As we explained earlier, Mobile IPv6 only initiates L3 handover after the MN connects to the new AR. Nevertheless, we see in the Fig. 5 that the L2 handover can be long, especially when the moving MN is not alone. This L2 handover time lets lots of time to initiate L3 operations in advance. The Tunnel-Based Handover presented in an earlier section seems to be very suitable since there is sufficient time to perform (or at least initiate) the exchange between the two ARs. Moreover, the realization of anticipation

from the IEEE 802.11b control frames seems to be possible (see above). However, the anticipated L3 handover must be controlled by the network, since the MN cannot send an L3 packet once it has started an L2 handover.

CONCLUSION

This article proposes an overview of current mobility management in the IETF for wireless IPv6 networks. We described Mobile IPv6 and two extensions, Hierarchical Mobile IPv6 and Fast Handover Protocol. Then we evaluated L2 handover and Mobile IPv6 handover over IEEE 802.11b wireless LAN. We discussed the results and showed how anticipation can be achieved with IEEE 802.11b control frames.

Mobile IPv6 allows an MN to register a new care-of address acquired in a visited network with its correspondent(s) and its home agent. The role of the home agent is to intercept the packets intended to a distant MN and to redirect them to the current localization of the MN. However, the time required for the MN to acquire a new care-of address and receive the redirected traffic can be too long, especially for real-time applications.

Hierarchical Mobile IPv6 proposes to manage the MN movement within an administrative domain. As long as an MN remains in a domain, it performs local registration and its movements are transparent to the rest of the Internet. Otherwise, Fast Handover uses anticipation to set up services for MNs before their connection with a new AP. The L2 triggers, which are messages informing about a change in the MN L2 connection, allow L3 handover to be initiated with or without interaction with the MN.

The tests we performed over IEEE 802.11b show that L2 handover could take a long time, especially if there are several active users. While basic Mobile IPv6 can take up to 8 s, we expect that this time could be considerably reduced with anticipation. Moreover, the IEEE 802.11b control frames seem to be well

suited to the L2 triggers, and thus allow L3 handover to be initiated as soon as L2 handover starts.

ACKNOWLEDGMENTS

Thanks are given to Christophe Jelger for technical review and encouragement.

REFERENCES

- [1] D. Johnson and C. Perkins, "Mobility Support in IPv6," IETF draft, draft-ietf-mobileip-ipv6-15.txt, July 2001.
- [2] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," IETF RFC 2460, Dec. 1998.
- [3] G. Dommety et al., "Fast Handovers for Mobile IPv6," IETF draft, draft-ietf-mobileip-fast-mip6-03.txt, July 2001.
- [4] H. Soliman et al., "Hierarchical MIPv6 Mobility Management," IETF draft, draft-ietf-mobileip-hmip6-05.txt, July 2001.
- [5] IEEE 802.11b, IEEE Std. 802.11-1999.
- [6] J. Manner et al., "Mobility Related Terminology," IETF draft, draft-manner-seamoby-terms-01.txt, Mar. 2001.
- [7] S. Thomson and T. Narten, "IPv6 Stateless Address Autoconfiguration," IETF RFC 2462, Dec. 1998.
- [8] J. Bound et al., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," IETF draft, draft-ietf-dhc-dhcpv6-20.txt, Oct. 2001.
- [9] J. Kempf et al., "Supported Optimized Handover for IP Mobility - Requirements for Underlying Systems," Internet Engineering Task Force (IETF) draft-manyfolks-l2-mobilereq-01.txt, Nov. 2001.
- [10] D. Trossen et al., "Issues in Candidate Access Router Discovery for Seamless IP Handoffs," IETF draft, draft-ietf-seamoby-cardiscovery-issues-00.txt, July 2001.

BIOGRAPHIES

NICOLAS MONTAVONT (montavont@dpt-info.u-strasbg.fr) received his B.Sc. (1998) and M.Sc. (2001) in computer science from the Universities of Lyon and Strasbourg, France, respectively. Currently, he is studying toward a Ph.D. within the Network Research Team at LSIIT laboratory (ULP-CNRS). His research work has included IP mobility, multiple interfaces management, and wireless LAN architectures.

THOMAS NOEL (Thomas.Noel@dpt-info.u-strasbg.fr) is assistant professor at Louis Pasteur University, Strasbourg, France. He is a member of the network research team of the LSIIT Laboratory (ULP-CNRS). He received his M.Sc. in 1995 and Ph.D. in computer science in 1998. His research interests include network architectures, protocols, and especially wireless IP networks and multicast routing protocols for mobile IP nodes.

The tests we performed over IEEE 802.11b show that L2 handover could take a long time, especially if there are several active users. While a basic Mobile IPv6 can take up to 8 s, we expect that this time could be considerably reduced with anticipation.