Implementation of Hierarchical Mobile IPv6 for Linux.

Richard Nelson Greg Daley Nick Moore

Center for Telecommunications and Information Engineering,
Monash University, Melbourne, Australia

October 18, 2002

Abstract

For Mobile IPv6 (MIPv6) to be adopted for real-time flows such as Voice over IP (VoIP) applications, handover disruption must be greatly reduced. Currently, MIPv6 handovers can cause packet delay and/or packet loss for many hundreds of milliseconds. The maximum acceptable disruption for real-time voice traffic is in the tens of milliseconds.

Hierarchical Mobile IPv6 (HMIPv6) attempts to reduce the delay caused by Binding Updates (BUs) during handover. This paper examines the process of extending MIPv6 implementations to support HMIPv6, and the limitations and pitfalls of this approach. We also present some preliminary test results of our implementation in a typical network scenario.

In addition, this paper outlines some of the potential techniques for further reducing handover disruption with minimal changes to the HMIPv6 draft standard, by addressing Router Solicitation (RS) and Duplicate Address Detection (DAD) delays.

1 Introduction

Mobile IPv6 (MIPv6) [7], has been widely proposed as a universal solution for mobile devices. Designed for time-independant flows, MIPv6 is now being advocated as a transport for real-time flow dependant services such as mobile telephones and video streaming devices. MIPv6 provides route optimizations which work well for devices which are in a relatively static location away from home, but which are less effective for devices moving rapidly between access networks.

A MIPv6 Mobile Node (MN) may take many hundreds of milliseconds to handover between access networks. This delay is often unacceptable for real-time traffic – for example, the maximum acceptable disruption for real-time voice traffic is in the tens of milliseconds.

The rest of the paper is organized as follows: In Section 2, we provide as summary of the components and operational principles of HMIPv6. Section 3 outlines our implementation approach. We present an overview of our testing methodology in

Section 4, present the results of experimental data in Section 5, discuss some further improvements in Section 6 and finally we offer our conclusions.

2 HMIPv6

2.1 Route Optimization in MIPv6

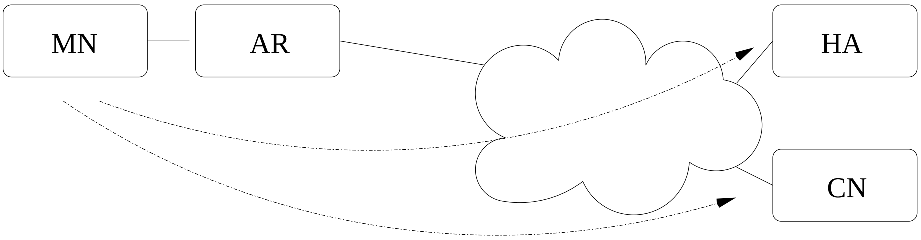

Figure 1: Handover signaling in MIPv6

In MIPv6, a Mobile Node (MN) remains reachable through its Home Address despite changing its access network. New correspondents route packets to the Home Address, and the Home Agent forwards these packets to the last known Care-of-Address (CoA) of the MN. Correspondent nodes may store the MN's current CoA in order to send the packets directly. This is called route optimization. When the MN moves, it sends a Binding Update (BU) to its Home Agent (HA) and all current Correspondent Nodes (CNs), letting them know its new location so that traffic can be routed to the new location.

The simplicity of MIPv6, and its lack of dependence on network infrastructure make it a natural choice for mobility. Route optimization prevents triangular routing while maintaining backward-compatibility for CNs which do not understand MIPv6.

However, MIPv6 handovers require the MN to signal its HA and each CN every time the MN moves. This introduces signaling overhead, and the time taken for a BU to reach the HA/CNs introduces packet loss or delay – until the BU is received by the HA/CN, packets are forwarded to the old CoA.

2.2 Route Optimization in HMIPv6

HMIPv6 extends the route optimization of MIPv6, adding a Mobility Anchor Point (MAP) to assist the MN. When the MN moves into MAP coverage, it asks the MAP for a Regional Care-of-Address (RCoA) from which the MAP will forward packets to the MN. Once the MN has obtained an RCoA, it uses this address in its BUs to HA and CNs.

A MAP covers a number of access networks. While a mobile node moves within the coverage of a MAP it need not re-bind its HA or CNs, thus signalling is reduced. Since the MN is generally closer to the MAP than its correspondents, signaling latency is also reduced.

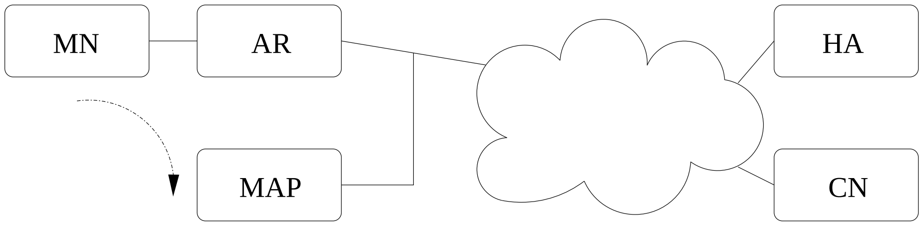

Figure 2: Handover signaling in HMIPv6

Configuration and deployment of MAPs adds infrastructure back into a previously low-infrastructure mobility system.

3 Our HMIPv6 Implementation

MIPL[9], developed at Helsinki University of Technology (HUT), is an Open Source project to extend the Linux Kernel and user-space tools to support Mobile IPv6. The current implementation is conformant to IETF draft 15 [7].

Our implementation of HMIP further extends MIPL to add support for HMIPv6 draft 6[12] Basic Mode.

Extended Mode has minimal benefits over Basic Mode[3][4], and is not essential to prove the concept of HMIP workable. We have not attempted to implement Extended Mode, nor have we implemented Authentication whose specification is still changing in the Mobile IPv6 draft[7][5].

3.1 Mobile Node

The Mobile Node (MN) functionality has been extended to handle HMIPv6 by implementing the following functions:

- Neighbour Discovery – The MN understands the new MAP option to determine when it has moved into MAP coverage.
- MAP Binding – The MN signals to the MAP and uses an RCoA on the MAP's subnet
- MAP Hierarchy – The MN keeps a list of available MAPs, and chooses an appropriate MAP for each CN.
- Listening – The MN accepts packets for its Home Address and for potentially multiple LCoAs and RCoAs. It must be able to decapsulate packets tunnelled to it by the MAP.

3.2 Mobility Access Point

The MAP is a device first introduced in HMIPv6, but with a great deal of similarity to a MIPv6 Home Agent. Our code extends the HA to add:

- Binding type information for HMIPv6 Basic and Extended Modes, Correspondent Nodes and Home Agents is kept.
- MAPs participate in Dynamic MAP Discovery (section 3.3).

3.3 Access Network Routers

Router advertisement (RA) is performed by `radvd` on each of the Access Routers (ARs). We have added features to advertise the MAP information option, and to propagate received MAP options for dynamic MAP Discovery.

4 Testing HMIPv6

For Intra-MAP handovers, HMIPv6 theoretically reduces handover delay compared to unaided MIPv6, where the time to notify the MAP is less than that required to notify the HA or CNs. Additionally, the signalling load should be reduced.

Indications of how well this works in realistic implementations are limited to that undertaken before the standardization of HMIPv6[2][1], and do not consider the requirements of other IPv6 protocols such as neighbor discovery[10] and stateless address autoconfiguration[13], each of which introduce delays before HMIPv6 or Mobile IPv6 signalling may be initiated.

4.1 Testing environment and assumptions

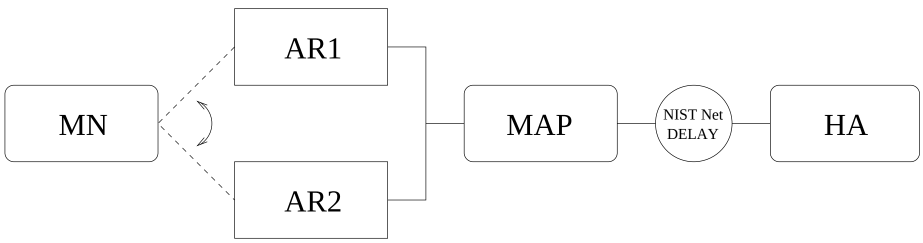

Figure 3: Test network for intra-MAP handovers

Our test network 3 has a simple bidirectional 100ms delay between the access sites (also incorporating the MAP) and the home and correspondent networks.

In order to measure the performance improvements of HMIPv6 over MIPv6, we assembled a test network with the following components:

Two Access Routers, each with one access network and each connected to a third router, which is also a MAP. This MAP connects via an IPv6 over IPv4 tunnel to a Home Network Router. One IPv4 hop exists between the MAP and Home Network Router. This IPv4 node runs NIST Net[11] in order to provide a simple network emulation of packet delay.

Each of these routers is a Pentium 166 MHz, with 32MB Ram, running the Linux 2.4.16 kernel, with MIPL release 0.9 (MIPv6 draft-14) and HMIPv6 patches (HMIPv6 Draft 5)[1]. Access routers were equipped with radvd 0.6.2pl4 patched to advertise MAPs. For the purposes of the test, only statically defined MAP information was advertised.

The Mobile Node is a Pentium 3 1GHz with 256 MB of RAM, running the same kernel as the MAP and ARs.

The Home Agent, another Linux server on the home network was a 450 MHz Pentium 3 with 256 MB of RAM. The kernel on this machine is Linux 2.4.16, with MIPL release 0.9.1 (mipv6 draft-15). No Authentication was turned on for any of the MIPv6 implementations and all physical connections are by 10/100BaseT Ethernet.

The test network in Figure 3 was set up to enable comparison of HMIPv6 and MIPv6. Handovers were forced by physically moving the MN from one AR to the other by disconnecting and reconnecting its network cable. In the HMIPv6 case, both ARs were configured to advertise the same MAP, and their own address Prefix. In the MIPv6 case, the same Router Advertising tool was used, but MAP advertisements were switched off.

Within the testing scenario, the MN has no explicit knowledge of Layer 2 events (such as link failure), and relies upon timers and reception of new Router Advertisements to determine movement to a new link. In the situation where no RA has been received and a timeout has occurred the MN sends a Router Solicitation to determine its attachment.

Since the applicability of HMIPv6 is to the portion of handover after L2 Handover (Section 2.1) HMIPv6 may not be initialised until movement is either detected or suspected. Therefore, the packet traces displayed start with either the RA or RS which cause (or indicate) the MN's awareness of movement.

5 Results

The results tables show handovers with MIPv6 only (Table 1) and HMIPv6 (Table 2). Columns show timing, packet source and destination, with a description of the packet. Each horizontal line in the table indicates a handover and its subsequent signalling.

In each table, initialization is performed in the first handover sub-table. The second sub-table represents movement to a new, previously unseen, network. The third sub-table is return to a previously visited network. Only HMIPv6 has a fourth sub-table, which shows rebinding the home address. This is performed periodically in both HMIPv6 and MIPv6.

[1] except for the IPv4 only node, which runs Linux 2.4.16 with NIST Net 2.0.10.

Time (s)	Source	Dest.	Contents
182.218935	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1
182.219530	LCoA-1	HomeAgent	BU (Ack,HReg) HomeAddr → LCoA-1 (seq #0)
182.299978	::	ff02::1:ff31:d5e5	Neighbour Solicitation from MN (DAD for Link Address)
182.420993	AR1-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN
183.414044	AR1-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN
183.414097	LCoA-1	AR1-link	Neighbour Advertisement from MN
183.414318	HomeAgent	LCoA-1	BAck: Binding Update accepted
183.720176	LCoA-1	HomeAgent	**BU (Ack,HReg) HomeAddr → LCoA-1 (seq #1)**
183.921499	HomeAgent	LCoA-1	**BAck: Binding Update accepted (seq #1)**
184.778704	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1
281.711546	AR2-link	all-nodes-multicast	ICMPv6 (Router advertisement) (HA) for AR2
281.712218	LCoA-2	HomeAgent	**BU (Ack,HReg) HomeAddr → LCoA-2 (seq #2)**
281.914012	AR2-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN
281.914067	LCoA-2	AR2-link	Neighbour advertisement from MN
281.914278	HomeAgent	LCoA-2	**BAck: Binding Update accepted (seq #2)**
283.731359	AR2-link	all-nodes-multicast	ICMPv6 (Router advertisement) (HA) for AR2
370.718186	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1
370.718838	LCoA-1	HomeAgent	**BU (Ack,HReg) HomeAddr → LCoA-1 (seq #3)**
370.719100	LCoA-1	AR2	BU (HReg) LCoA-2 → LCoA-1 (seq #4)
370.719998	AR1-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN
370.720028	LCoA-1	AR1-link	Neighbour Advertisement from MN
370.720226	AR2	LCoA-1	BAck: Binding Update was rejected (seq #4) - Home registration not supported
370.920609	HomeAgent	LCoA-1	**BAck: Binding Update accepted (seq #3)**
371.467993	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1

Table 1: Capture of MIPv6 only (Non-Hierarchical) movement

With MIPv6 only (see Table 1) these provisional results show an average 202 ms round-trip delay between the BU being sent by the MN and the BAck being received by the BU. This is consistent with the 200ms round-trip delay betwen the ARs and the HA.

Notable in the third sub-table is the presence of a second binding update to a previous Access Router. This is an optional part of MIPv6, which is still (optionally) being performed in HMIPv6 as well.

Conversely, with HMIPv6 the handover (Table 2) speed is limited by the DAD time required to construct LCoAs and the RCoA. This is illustrated in the second and third sub-tables.

In a handover to a new Access Network (around t=84), a binding update is sent as soon as the RA is received, but the MN has to wait for 1 second for DAD to complete. This greatly impacts the usefulness of HMIPv6, given that a mobile device may very rarely revisit access networks within the advertised prefix's timeout.

In the handover back to a previously visited access network (around t=509), when DAD is not an issue, we see a delay of only 1.1ms between transmission of the BU and reception of the BAck. This is consistent with the proximity of the MAP to both ARs. Binding is not delayed by the 200ms round-trip delay between the ARs and the HA.

Time (s)	Source	Dest.	Contents
6.296977	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1
6.297659	LCoA-1	HomeAgent	**BU (Ack,HReg) HomeAddr $\rightarrow$ LCoA-1 (seq #0)**
6.499560	AR1-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN-link
6.799977	::	ff02::1:ff31:d5e5	Neighbour Solicitation from MN-link
			(DAD for Link Address)
7.492118	AR1-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN-link
7.800196	LCoA-1	HomeAgent	**BU (Ack,HReg) HomeAddr $\rightarrow$ LCoA-1 (seq #1)**
8.010461	LCoA-1	MAP	*BU (Ack,DAD,MAP) RCoA $\rightarrow$ LCoA-1 (seq #2)*
8.010760	LCoA-1	HomeAgent	**BU (Ack,HReg) HomeAddr $\rightarrow$ RCoA (seq #3)**
			(using AltCoA)
8.492006	AR1-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN-link
8.492053	LCoA-1	AR1-link	Neighbour Advertisement
8.492259	HomeAgent	LCoA-1	**BAck: Binding Update accepted (seq #1)**
8.492284	MAP	LCoA-1	*BAck: Binding Update accepted (seq #2)*
8.492317	HomeAgent	RCoA	**BAck: Binding Update accepted (seq #3)**
			(via MAP tunnel)
8.736879	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1
84.889257	AR2-link	all-nodes-multicast	ICMPv6 (Router advertisement) (HA) for AR2
84.890065	LCoA-2	MAP	*BU (Ack,DAD,MAP) RCoA $\rightarrow$ LCoA-2 (seq #4)*
84.890907	AR2-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN-link
			(AR wants to deliver BAck, but MN DAD incomplete)
85.669978	::	ff02::1:ff31:d5e5	Neighbour Solicitation from MN-link
			(DAD for Link Address)
85.884924	AR2-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN-link
			(AR wants to deliver BAck, but MN DAD incomplete)
86.390196	LCoA-2	MAP	*BU (Ack,DAD,MAP) RCoA $\rightarrow$ LCoA-2 (seq #5)*
86.884803	AR2-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN-link
			(AR wants to deliver MAP's response)
86.884853	LCoA-2	AR2-link	Neighbour Advertisement
86.885044	MAP	LCoA-2	*BAck: Binding Update accepted (seq #4)*
86.885045	MAP	LCoA-2	*BAck: Binding Update accepted (seq #5)*
87.479026	AR2-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR2
509.362213	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1
509.363019	LCoA-1	MAP	*BU (Ack,DAD,MAP) RCoA $\rightarrow$ LCoA-1 (seq #6)*
509.363323	LCoA-1	AR2	BU (HReg) LCoA-2 $\rightarrow$ RCoA (seq #7)
			(Tell old Access Router RCoA)
509.363914	AR1-link	ff02::1:ff31:d5e5	Neighbor Solicitation to MN-link
509.363942	LCoA-1	AR1-link	Neighbor Advertisement from MN-link
509.364141	MAP	LCoA-1	*BAck: Binding Update accepted (seq #6)*
509.364326	AR2	LCoA-1	BAck: Binding Update was rejected
			- Home registration not supported (seq #7)
512.111232	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1
808.778405	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1
809.000207	LCoA-1	HomeAgent	**BU (Ack,HReg) HomeAddr $\rightarrow$ RCoA (seq #8)**
			(periodic rebinding of Home agent at 80% of lifetime)
809.202182	AR1-link	ff02::1:ff31:d5e5	Neighbour Solicitation to MN-link
809.202245	LCoA-1	AR1-link	Neighbour Advertisement from MN-link
809.202462	HomeAgent	LCoA-1	**BAck: Binding Update accepted (seq #8)**
			(via MAP tunnel)
809.538453	AR1-link	all-nodes-multicast	ICMPv6 (Router advertisement) for AR1

Table 2: Capture of movement within a HMIPv6 hierarchy

6 Extensions to HMIPv6

While HMIPv6 reduces disruption caused by binding delay, there are a number of other potential improvements which can be made to the handover process. Major improvements should be possible using Fast Handovers, Fast Router Advertisements and Dummy CoA Networks.

6.1 Fast Handovers for IPv6

Fast Handovers, as documented in draft-ietf-mobileip-fast-mipv6-04 [6], allow a MN to immediately detect movement rather than waiting Neighbour Discovery.

Additional specification of edge tunnelling in this specification temporarily mitigates signalling requirements to CN's and the HA.

Combinations of HMIPv6 with Fast Handovers[6] have been proposed, although this has not yet been implemented. It is expected that the combination of predictive handovers with HMIPv6 will significantly impact the required handover time to signal the MAP, as well as minimizing the extra-network signalling provided by HMIPv6[12].

6.2 Fast Router Advertisement

While Fast Handover techniques allow a MN to send a Router Solicitation immediately upon entering a network, RFC 2461[10] requires a router to delay its response to a Router Solicitation by a small random amount, in order to prevent routers from simultaneously sending Router Advertisements.

Fast Router Advertisement[8] allows at most one router on a network to respond immediately to an RS, thus eliminating this delay. This improvement is available to both MIPv6 and HMIPv6.

6.3 Dummy CoA Networks

A further delay in HMIPv6 is that caused by Duplicate Address Detection (DAD) while the MAP is configuring an RCoA. DummyCoA[4] is a technique to remove this delay by configuring all RCoAs on a separate prefix. This 'dummy prefix' is a proper, on-topology prefix, but allocated to a dummy interface of the MAP. In this way, the MAP can be sure that no-one but itself is using addresses on the dummy prefix, and thus eliminate DAD delay for RCoA assignments entirely.

7 Conclusions

Our results indicate that significant reductions to mobileip handover signalling delay can be achieved with HMIPv6. In particular, the delay introducted by BU signalling can be practically eliminated.

Other delays also prevent smooth handovers. In particular, our results show that delays caused by DAD dominate handover delay when moving onto a new access network.

By identifying each cause of handover delay, and devising a method to minimize it, the unacceptably long delays associated with real-time traffic over MIPv6 should be able to be brought into the acceptable range of tens of milliseconds. Further testing is necessary to reveal if this improvement is practical in the field.

Acknowledgements

- Thanks to Brett Pentland for his help with setting up the testbed environment
- Thanks to Ahmet Şekercioğlu for his help with the paper.
- This work was supported by the ATCRC Next Generation Internet project.

References

[1] C. Castelluccia. Toward a hierarchical mobile IPv6. In *Eighth IFIP Conference on High Performance Networking (HPN'98)*, Vienna, Austria, September 1998.

[2] C. Castelluccia. HMIPv6: a hierarchical mobile IPv6 proposal. *ACM Mobile Computing and Communication Review (MC2R)*, April 2000.

[3] Greg Daley. Comparison of basic and extended mode function in hierarchical mobile IPv6. Australian Telecommunications CRC (ATcrc) Internal Technical Report P1-PR1-TR-009, July 2002.

[4] Greg Daley. Optimization of address selection for HMIPv6 using dummy interfaces. Australian Telecommunications CRC (ATcrc) Internal Technical Report P1-PR1-TR-008, May 2002.

[5] Charles Perkins David B. Johnson and Jari Arkko. Mobility support in IPv6 *internet draft - work in progress*. URL reference: http://www.mipl.mediapoli.com/drafts/draft-ietf-mobileip-ipv6-17.txt, May 2002.

[6] G. Dommety (Ed) et al. Fast handovers for mobile IPv6 *internet draft - work in progress*. URL reference: http://www.ietf.org/internet-drafts/draft-ietf-mobileip-fast-mipv6-04.t%xt, March 2002.

[7] David B. Johnson and Charles Perkins. Mobility support in IPv6 *internet draft - work in progress*. URL reference: http://www.mipl.mediapoli.com/drafts/draft-ietf-mobileip-ipv6-15.txt, 2001.

[8] James Kempf, Mohamed M. Khalil, and Brett Pentland. IPv6 fast router advertisement *internet draft - work in progress*. URL reference: http://www.ietf.org/internet-drafts/draft-mkhalil-ipv6-fastra-01.txt, 2002.

[9] MIPL. Mobile IPv6 for Linux, 2002. URL reference: http://www.mipl.mediapoli.com/.

[10] T. Narten, E. Nordmark, and W. Simpson. RFC 2461 Neigbhour Discovery for IP Version 6 (IPv6), 1998. URL reference: http://www.ietf.org/rfc/rfc2461.txt.

[11] NIST. NIST Net home page, May 2002. URL reference: http://snad.ncsl.nist.gov/itg/nistnet/.

[12] H. Soliman, C. Castelluccia, K. El-Malki, and L. Bellier. Hierarchical MIPv6 mobility management (HMIPv6) *internet draft - work in progress*. URL reference: http://www.ietf.org/internet-drafts/draft-ietf-mobileip-hmipv6-06.txt, 2002.

[13] S. Thomson and T. Narten. RFC 2462 IPv6 Stateless Address Autoconfiguration, 1998. URL reference: http://www.ietf.org/rfc/rfc2462.txt.