

Edge Handovers

Q4 Report

Nick Moore Y. Ahmet Şekercioglu Johnny M. Lai

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

1 Introduction

1.1 Mobile IPv6

Mobile IPv6 [1] (MIPv6) builds upon the IPv6 standards to allow hosts to maintain a *Home Address* (HAddr) through which they are always reachable, while using a changable *Care-of Address* (CoA) to connect to the Internet. This is done using a *Home Agent* (HA), which intercepts traffic for the HAddr and tunnels it to the CoA. A host using this service is referred to as a *Mobile Node*. The process of moving from one CoA to the next is known as a *handover*.

MIPv6 performance is generally not sufficient for applications like real-time voice which demand low latency handovers. Our approach to reducing handover latency has been to break down the handover latency into a number of separate delays, and to find solutions to each of them. One cause of delay is the time taken for a Binding Update to reach the HA, and for traffic redirected by the HA to reach the Mobile Node. This is called the *Binding Update Round Trip Time*.

1.2 Hierarchical Mobile IPv6

Hierarchical Mobile IPv6 [2] (HMIPv6) is an extension to MIPv6 which adds the *Mobility Anchor Point* (MAP) to the mobile network infrastructure. While the Mobile Node moves within the coverage of a MAP, it can use a stable *Regional Care-of Address* (RCoA) issued by that MAP, and the MAP tunnels packets for the RCoA to the Mobile Nodes changing *Local Care-of Address* (LCoA). The MAP covers a number of ARs, and while the Mobile

Node moves within the coverage of a MAP it only needs to inform the MAP of its movements rather than the HA.

HMIPv6 effectively reduces the Binding Update delay to the round trip time between mobile node and MAP. Binding Updates are only sent to the HA when the Mobile Node moves beyond the coverage of its MAP and must handover to a new MAP. However, it is difficult to optimize the location and number of MAPs in a network: on the one hand, a MAP should be close to the edge of the network in order to minimize the Mobile Node to MAP round trip time. On the other hand, the MAP should be far from the edge in order to minimize the number of MAP to MAP handovers. Adding a MAP into the network also goes against the ‘smart edge’ design principle which tries to keep the complexity of the Internet at the edge.

1.3 Edge Handovers for Mobile IPv6

Edge Handovers (EH) [3] attempts to solve these problems by moving all MAP functionality to the Access Routers which form the very edge of the mobile network, and by streamlining handovers between these ARs.

An *Edge Handovers enabled Access Router* (EH-AR) is capable of providing local mobility support to Mobile Nodes in much the same way that an HMIPv6 MAP does. We call the EH-AR which is currently providing this service to an Mobile Node the *Bound Access Router* (BAR).

Figure 1 illustrates the steps taken when a Mobile Node arrives on an Edge Handovers capable network, and Figure 2 shows it moving within the network.

A Mobile Node can keep the same BAR indefinitely. However, in order to minimize edge network traffic, it periodically changes its BAR to a closer access router. This does not have to be done at the time of handover. This process is illustrated in Figure 3.

There are many possible heuristics for choosing when to change BAR, for this paper we have only considered a very simple example: every ten seconds the Mobile Node tries to change its BAR to its current AR.

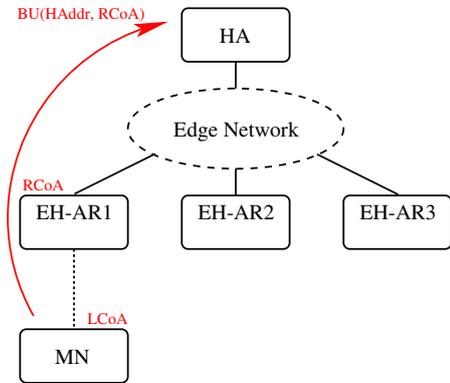
2 Simulation Scenario

For these simulations, we set up a ‘Campus Network’ topology as shown in Figure 4. In this scenario 50 Access Routers (AR) are connected in a minimal spanning tree. Each AR is directly connected to an 802.11b Access Point (AP). A gateway router at the centre of the topology connects this network to a distant Home Agent (HA) and Correspondent Node (CN). For data

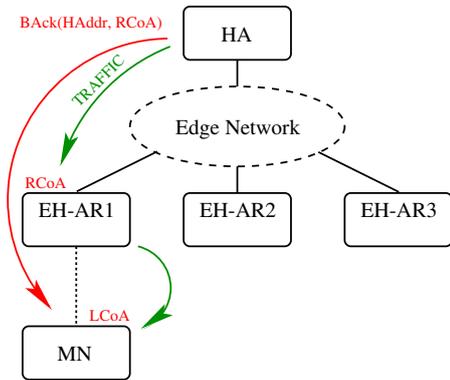
collection, we simulated 50 Mobile Nodes (Mobile Node) travelling around the campus using a random waypoint model, attaching to Access Points as available.

The links between the directly connected Access Routers incur a 10ms delay in either direction, representing the fast edge network, and the links between the Core Router and the HA and CN each incur a 100ms delay, representing the slow Internet.

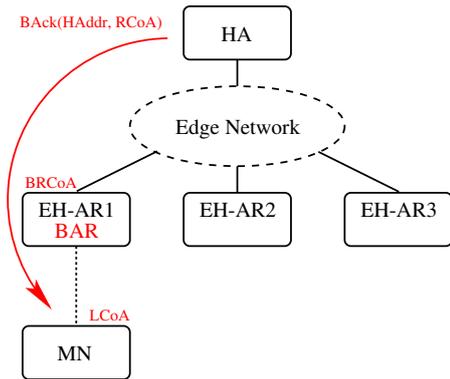
In the HMIPv6 tests, MAP functionality is implemented on the gateway router. In the EH tests, EH-AR functionality is implemented on each of the ARs.



When the Mobile Node arrives on a link which supports EH, it configures an LCoA and an RCoA, and sends a BU to its HA to informing it of the new RCoA

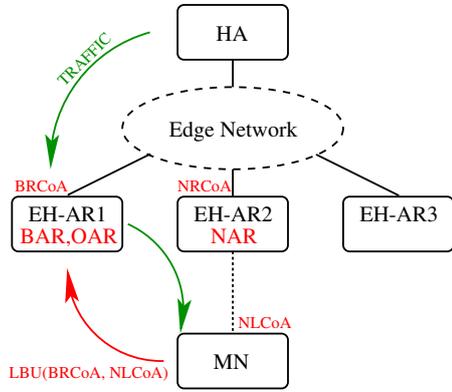


The HA sends a BAck to the Mobile Node, and can now forward traffic to the Mobile Node via the RCoA.

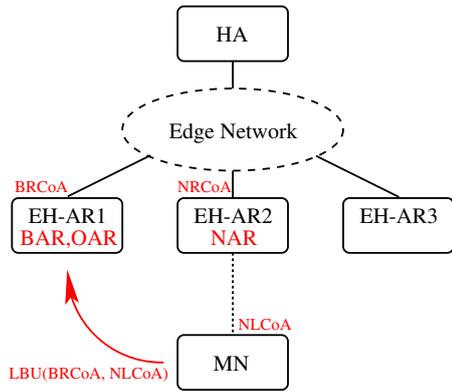


When the Mobile Node receives the BAck, marks that RCoA as its *Bound Regional Care-of Address* (BRCOA) and the the AR which provided the BRCOA as its *Bound Access Router* (BAR).

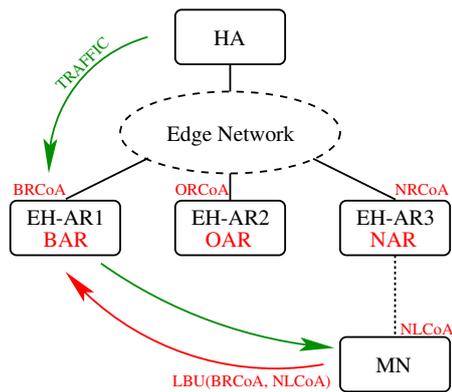
Figure 1: Binding an Access Router



When the Mobile Node moves to a new AR, it obtains a new LCoA and sends a *Local Binding Update* (LBU) to inform the BAR of its new LCoA.

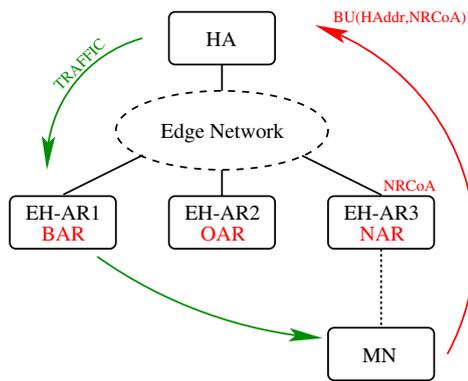


The HA continues to forward traffic to the BRCoA, and the BAR forwards this traffic to the new LCoA.

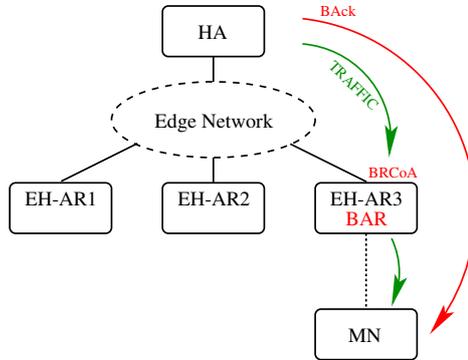


The Mobile Node moves again, and informs the BAR of its new location. The BAR forwards traffic to the new LCoA.

Figure 2: Moving between Access Routers



When the Mobile Node decides to get a new BAR, it configures a new RCoA and sends a BU to its HA, informing the HA of its new RCoA.



The HA sends a Back and starts forwarding traffic to the new RCoA. When the Mobile Node receives the BAck, it records the AR as its BAR and the RCoA as its BRCoA.

Figure 3: Changing Bound Access Router

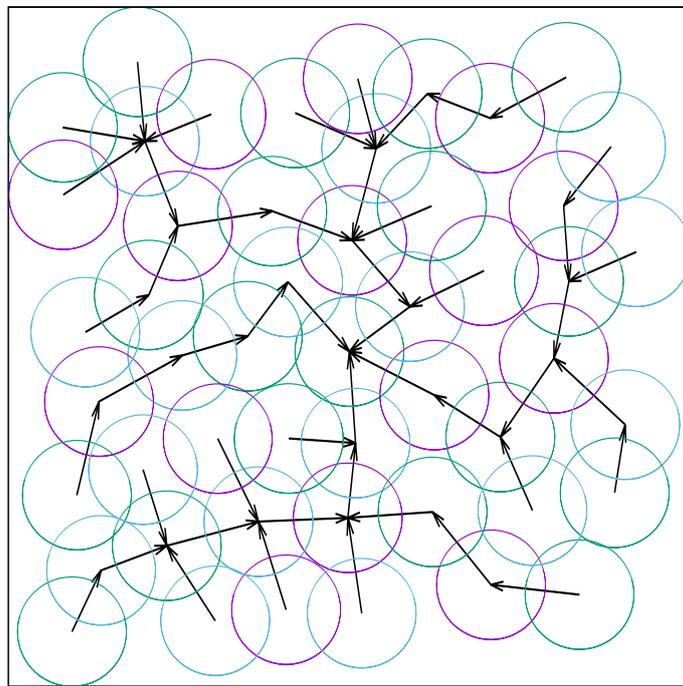


Figure 4: Simulated campus network topology

PROTOCOL	HANDOVER LATENCY (MS)					
	MEAN	5%	25%	MEDIAN	75%	95%
MIPv6	622	292	313	351	391	3263
HMIPv6	431	92	114	151	191	3065
EH	249	63	73	91	152	1544

Table 1: Statistics for L3 handover latency for MIPv6, HMIPv6 and EH

3 Results

Comparison of the distribution of L3 handover latencies for MIPv6, HMIPv6 and EH are presented in Figure 5 and Table 1. Also, a more detailed illustration of the distribution of handover latencies can be found in Figure 6. Each L3 handover latency is measured by taking the time between the L2 Up indication and the reception of a BAcK (in the case of MIPv6) or an LBAcK from the MAP (in the case of HMIPv6) or an LBAcK from the BAR (in the case of EH). These signals indicate the start and completion of L3 handover. EH and HMIPv6 both show a significant advantage over MIPv6, as L3 handover can occur without signalling across the Internet.

4 Conclusions

EH provides a significantly better L3 handover delay performance as compared to HMIPv6. The simulation results reveal that, EH provides 42% and 40% shorter handover latency as compared to HMIPv6 as measured in average and median terms respectively.

Unlike HMIPv6, EH does not need optimization and performs well even with a simple handover heuristic.

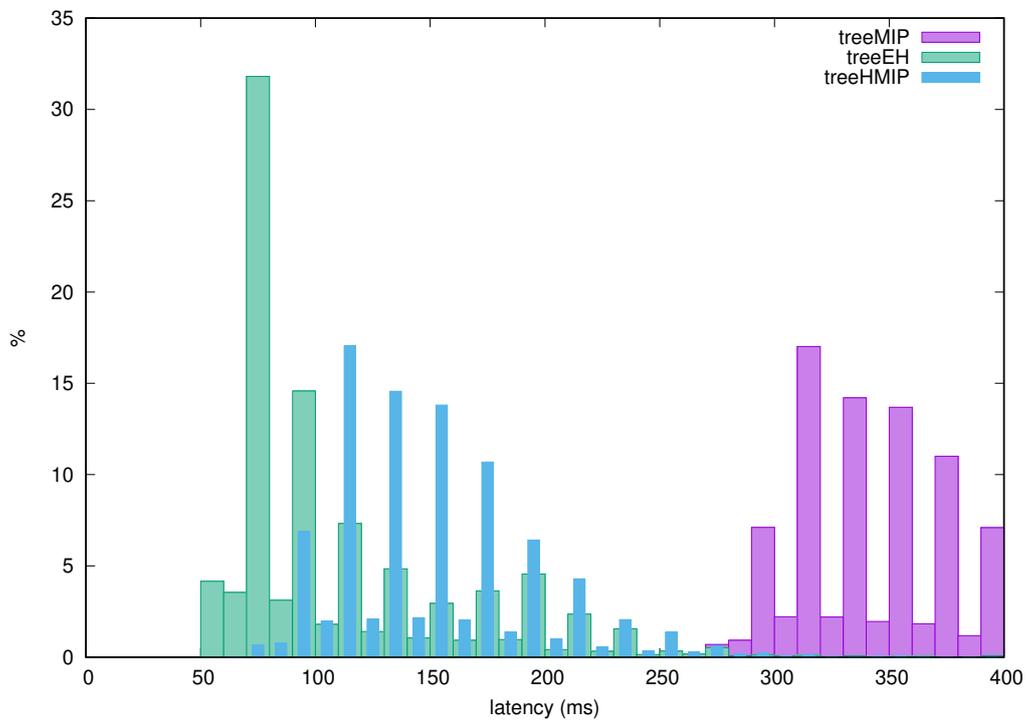


Figure 5: Histogram of L3 handover latency for MIPv6, HMIPv6 and EH

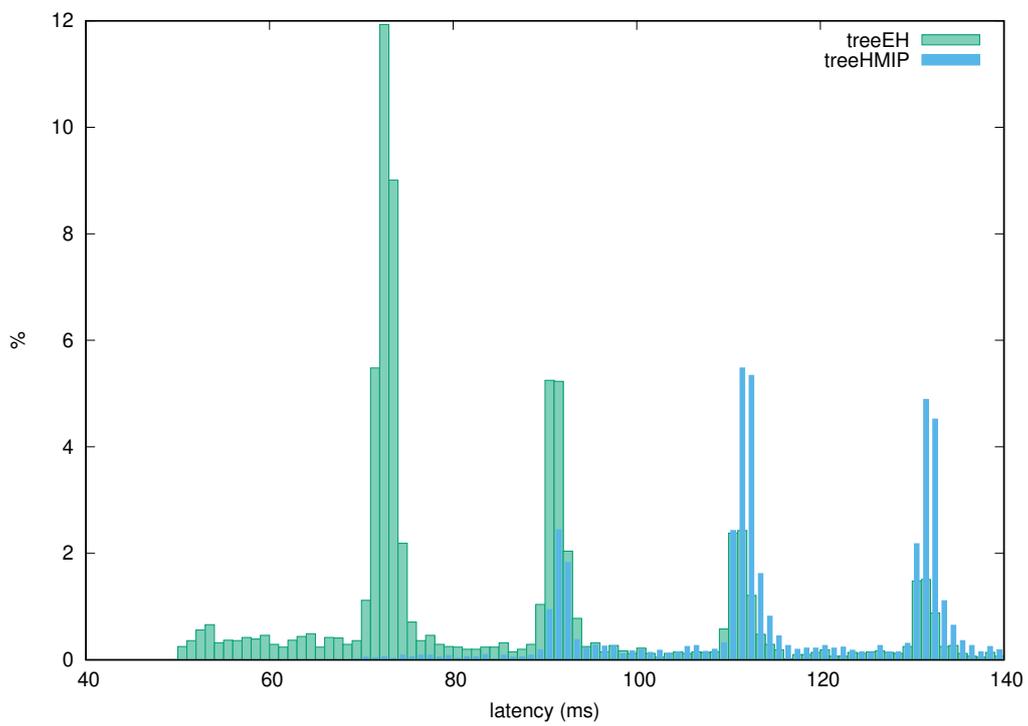


Figure 6: More detailed histogram of L3 handover latency for EH and HMIPv6

References

- [1] D. Johnson, C. Perkins, and J. Arkko. RFC3775: Mobility Support in IPv6 (proposed standard). URL reference: <http://www.ietf.org/rfc/rfc3775.txt>, June 2004.
- [2] H. Soliman, C. Castelluccia, K. El Malki, and L. Bellier. RFC4140: Hierarchical Mobile IPv6 Mobility Management (HMIPv6) (experimental). URL reference: <http://www.ietf.org/rfc/rfc4140.txt>, August 2005.
- [3] N. Moore, JH. Choi, and B. Pentland. Edge Handovers for Mobile IPv6. URL reference: <http://www.ctie.monash.edu.au/ipv6/fastho/draft-moore-mobopts-edge-handovers-01.txt>, July 2004.