



Performance Comparison of Low Latency Mobile IP schemes

Chris Blondia, Olga Casals, Llorenç Cerdà, Nik Wijngaert van Den, Gert Willems, P. Cleyn De

► To cite this version:

Chris Blondia, Olga Casals, Llorenç Cerdà, Nik Wijngaert van Den, Gert Willems, et al.. Performance Comparison of Low Latency Mobile IP schemes. WiOpt'03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks, Mar 2003, Sophia Antipolis, France. 10 p. inria-00466577

HAL Id: inria-00466577

<https://inria.hal.science/inria-00466577>

Submitted on 24 Mar 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Performance Comparison of Low Latency Mobile IP schemes¹

C. Blondia^a, O. Casals^b, Ll. Cerdà^b, N. Van den Wijngaert^a, G. Willems^a, P. De Cleyn^a

^aUniversity of Antwerp, Dept. Mathematics and Computer Science, Universiteitsplein 1, B-2610 Antwerp – Belgium, Tel: +32 3 8202404, email: chris.blondia@ua.ac.be

^bTechnical University of Catalonia (UPC), Dept. Computer Architecture, Jordi Girona 1-3, D6, E-08034 Barcelona–Spain, Tel: + 34 934016985, email: olga@ac.upc.es

Abstract

This paper investigates the performance of two low latency handoff protocols proposed by the IETF, namely pre- and Post-Registration. These protocols use Layer 2 triggers to reduce the built-in delay components of Mobile IP. We propose a simple analytical model that allows to assess the packet loss and the delay characteristics of these protocols. Their scalability properties are investigated by means of an OPNET simulation model and their implementation using IEEE802.11 as link layer is also discussed.

1. Introduction

The growing number of portable computing devices and the requirement to provide seamless connectivity to the global Internet using end-to-end IP solutions for mobile users have stimulated the research into IP mobility protocols. Mobile IP ([12], [9]) is the current standard solution for mobility management in IP networks. It allows a Mobile Node (MN) to change its point of attachment from one access router to another across media of similar or dissimilar types. Establishment of new tunnels can introduce considerable delays in the handoff process due to the round-trip time between the Foreign Agent (FA) and the Home Agent (HA) during the registration process. Applied in an environment with frequent handoffs, this may lead to unacceptable disturbance to ongoing sessions in terms of handoff latency and packet loss.

The IETF has proposed a number of protocols that handle local movements locally. These so-called micro-mobility protocols use host-based routing schemes (such as Cellular IP [6] and HAWAII [13]), or are based on hierarchical tunneling techniques using a tree structure of FAs (such as Regional Registration [8] or Hierarchical Mobile IP [16]). These protocols are designed without any assumption regarding underlying layers. This allows the widest possible applicability and a clean separation between Layer 2 and Layer 3 of the protocol stack. However, this layer separation results in lower performance. Indeed, the MN may only communicate with a directly connected FA and therefore it can only start the registration process after completion of the L2 handoff. Moreover, the MN is unreachable during the registration process, a property that may contribute to a non-negligible handoff latency.

In this paper we analyse two mobility protocols [7] that aim at low latency Layer 3 handoff based on Layer 2 information through the use of L2 triggers. In Pre-Registration, the MN can communicate with the new FA while still being connected with the old FA. In the second type, Post-Registration, the data is delivered to the MN by the new FA, even before the actual registration process has been completed.

This paper focuses on the performance evaluation of these schemes and their comparison. We present a simple analytical model that allows a detailed analysis of the delay characteristics and the buffer requirements for a single MN involved in a handoff. In particular we investigate the influence of the timing of the L2 triggers on the delay and the packet loss of UDP streams towards an MN that is involved in a handoff. In order to investigate the scalability of the protocols, we have developed an OPNET simulation model that allows to consider more realistic systems with respect to the number of MNs and their mobility pattern. A possible implementation of the mechanisms using IEEE802.11 as link layer is also discussed.

2. Low Latency Handoff Mechanism

2.1 Reference Network

Consider a configuration depicted in Figure 1. The access network is connected to the Internet via a gateway router. The two access points (AP) belong to different subnets with access routers referred to as the old access

¹ This work was supported by the Ministry of Education of Spain under grant CICYT TIC-2001-0956-C04-01 and by the Department of Universities, Research and Information Society (Generalitat de Catalunya) under grant CIRIT 2001-SGR-00226, by the Fund for Scientific Research Flanders under project G.0315.01 and by DWTC Belgium under project IAP P5/11 MOTION (Mobile Multimedia Communication Systems and Networks)

router and the new access router. The APs may or may not be co-located with the access routers. The access routers and the gateway router are provided with mobility agents, resp. oFA, nFA and the GFA. The reference network is not hierarchical. Router 3 has been added to allow independence between the distance (oFA,nFA) and the distance (nFA,GFA). We assume that a Corresponding Node (CN) sends packets (UDP and TCP) to the Mobile Node (MN).

2.2 L2 Triggers

The handoff schemes investigated in this paper make use of L2 triggers. Such a trigger is a signal related to the L2 handoff process. A first trigger that is used is an early notice of an upcoming change in the L2 point of attachment of the MN, referred to as anticipation trigger. A second trigger, the Line Down trigger (LD), indicates that the L2 link between the MN and the old AP is lost. The Line Up trigger (LU) occurs when the L2 link between the MN and the new AP is established. A trigger initiated at the old FA is referred to as a *source trigger* and a trigger initiated at the new FA is referred to as a *target trigger*.

For a detailed discussion on these L2 triggers, we refer to [7].

2.3 Pre- and Post-Registration Handoff

The Pre-Registration handoff scheme is based on an anticipated Layer 3 handoff. The L2 trigger contains the new FA's IP address and this allows the Mobile Node (MN) to communicate via the old FA with the new FA while still being connected with the old FA. The MN sends a registration request to the new FA via the old FA (if the L2 handoff is not completed) and the new FA issues a regional registration to the Gateway FA (GFA). The latter sends a registration reply message to the new FA. Until the MN actually completes the L2 handoff to the new FA and establishes the new L2 link, the new FA can receive packets for which it does not have a link layer connection. These packets may be buffered in the new FA.

In the Post-Registration scheme, the registration occurs after the L2 handoff is complete. The L2 trigger initiates the set-up of a bi-directional edge tunnel (BET) between the old FA and the new FA. When the old FA receives the L2 Line Down (LD) trigger, it starts forwarding packets destined to the MN via the BET to the new FA. When the new FA receives an L2 Line Up (LU) trigger, it delivers the packets received from the old FA to the MN. Eventually the MN performs a formal MIP registration.

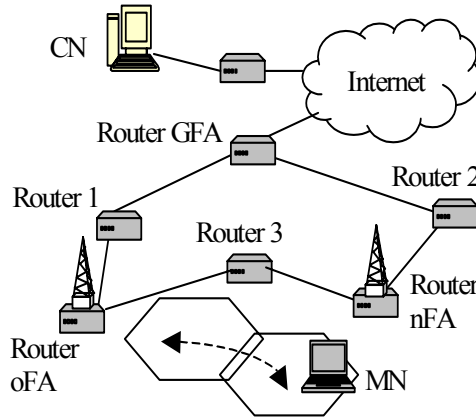


Figure 1: Reference Network

2.4 Implementation of low latency handoff mechanisms using IEEE802.11 as link layer

We describe a possible implementation in a wireless network having an 802.11 link layer [18]. Our goal is to illustrate by means of a realistic example the abstraction made in previous sections about the L2 triggers and the interaction with the L3 layer protocol.

We shall assume that the wireless network operates in *infrastructure mode*. In this mode all MN transmissions go through an *Access Point* (AP). Before an MN can send data packets, it has to be associated with an AP. The MN associates with an AP by sending an Association Request message (or Re-association if the MN is already associated), which in turn is answered with an Association (or Re-association) Reply. Before association (or re-association) the MN uses the 802.11 beacons sent by the AP in order to decide which AP would make the best connection.

For sake of simplicity, we shall assume the following: (i) The AP is an embedded entity in the FA such that the L2 triggers can be easily implemented by means of some internal interface. (ii) The MN uses the same channel with all APs (thus, during the handoff the MN can communicate with the new and the old AP while it is in

coverage with them). In this scenario the L2 link between the MN and the old AP is lost when the MN moves out of coverage of the old AP. Communication with both APs using different channels could be achieved using two transceivers [7]. (iii) The movement at L2 layer is detected upon receiving the first beacon from the new AP (nAP). Note that a timeout is needed to avoid pingponging between both APs. The traces shown in this section have been obtained with the ns simulator [19].

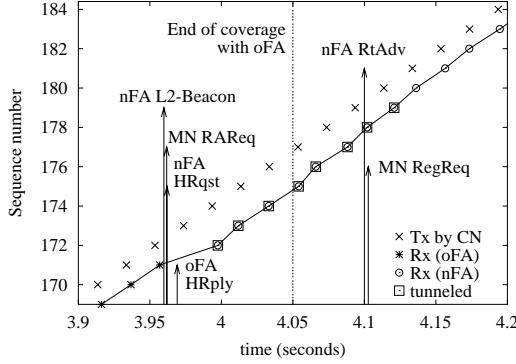


Figure 2: Target Trigger-Post Registration handoff.

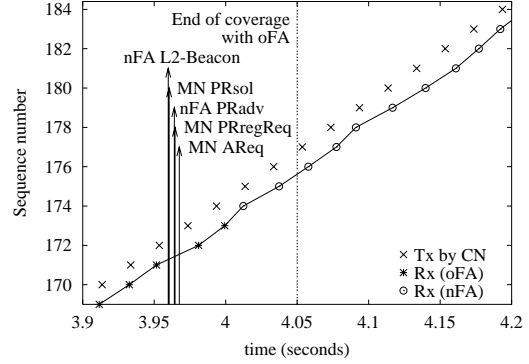


Figure 3: Target Trigger-Pre Registration handoff.

Figure 2 shows a trace with a possible handoff using the Post-Registration scheme. In this trace the CN sends a CBR stream to the MN. The figure shows: (i) The instants when the packets are sent (indicated as “Tx by CN”). (ii) The instants when the MN receives the packets. These are indicated differently according to whom they are sent by: the oFA or the nFA (indicated as “Rx (oFA)” and “Rx (NFA)” respectively). (iii) The signaling packets sent during the handoff. These are explained in the following.

When the MN approaches the nFA, the L2 beacons sent by the nAP trigger the L2 handoff at the MN, which sends a *Re-association Request* (RAReq in the figure) to the nAP. Upon receiving this frame, there is a *target trigger* at the nFA (LU trigger), which sends the Handoff Request (HRqst in the figure) to the oFA. Upon receiving the HRqst, the oFA sends the Handoff Reply (HRply in the figure) and establishes a tunnel with the nFA. In this way, the packets can reach the MN via the nFA after the coverage with the oFA has been lost. The oFA would start sending the packets addressed to the MN through the tunnel immediately after it is established. This instant corresponds with the LD trigger introduced in section 2.3. Finally, when the nFA sends the Router Advertisement, the MN makes a Registration with the nFA.

Figure 3 depicts a trace analogous to Figure 2 but using the Pre-Registration mechanism. The signaling packets sent during the handoff are the following: the L2 beacons sent by nAP produce an L2 trigger at the MN indicating the imminent movement to this AP. Then the MN sends the Proxy Router Solicitation (indicated as “MN PRsol”) to nFA. Upon reception of this packet, the nFA sends a Proxy Router Advertisement (indicated as “nFA PRadv”) to the MN. Upon reception of this packet, the MN sends a Registration Request to the oFA with destination address the nFA (indicated as “MN PregReq”). Then, the MN sends the Re-association to the nAP.

A more detailed description of possible implementations of Pre- and Post-Registration handoff mechanisms using IEEE 802.11 can be found in [4] and [5].

3. Performance Evaluation of Low Latency Handoff Mechanisms

In this part we analyse and compare the performance of the two mechanisms. An analytical model is introduced and used to compare the delay characteristics and the buffer requirements for a MN that switches from one access point to another. A similar model has been used to investigate other protocols (Cellular IP [1], HAWAII [2], Optimized Smooth Handoff [3] and Post-Registration [5]).

3.1 An analytical model for Pre- and Post-Registration Handoff

In this section we propose a simple analytical model that enables us to compute and compare performance characteristics of the Pre- and Post-Registration handoff procedure. In particular, we are interested in the delay distribution of packets involved in a handoff, and in the buffer requirements of the relevant FAs.

Consider the network architecture as depicted in Figure 1. We model all routers in the network as ordinary M/M/1 queues. Hence, each packet passing through some router has an exponentially distributed random service time, which is assumed to both include the processing time in the router and the transmission time. Moreover the response time of a packet is also exponentially distributed.

Now consider a MN moving from the oFA to the nFA, and suppose an overlapping area between the two subnetworks. We assume that a handoff is initiated when the anticipation trigger occurs (this is about at the moment the MN enters this area) and we denote this instant by t_0 . In this section we furthermore assume that we

have a source-triggered handoff, but the model can also be applied for target-triggered handoff. The timing of the relevant L2 triggers (L2 LD and L2 LU) is considered to be constant. We define D_{LD} and D_{LU} as the time between t_0 and the respective triggers and we have that $0 < D_{LD} < D_{LU}$.

The performance of the handoff procedures strongly depends on the timing of the triggers. Additionally, there are some important random time instants for both methods.

- For Pre-Registration, notably important is the moment the regional registration request arrives at the GFA. From that moment on packets will be directed to the nFA instead of the oFA. This moment ($t_{1,pre}$) is random, and is distributed as the sum of several exponential variables (the routers on the path from oFA via nFA to GFA), and several constant values (the propagation delays on the links between all routers).
- For Post-Registration, an essential random value is the moment the BET is established. Here too, this instant ($t_{1,post}$) is distributed as the sum of exponential variables and constants.

We now look at a UDP stream (i.e. a constant bit rate packet stream) originating from a CN with destination the MN. Suppose that every T ms a packet arrives at the GFA. We can then observe the stream starting from a packet arriving some time before t_0 , and compute the distribution of the end-to-end delay (GFA to MN) of this packet and each subsequent packet involved in the handoff.

The delay is again a random variable made up from exponential variables and constants, and its specific form depends on the path the packet follows. The latter is in turn dependent on certain stochastic events. As an example consider the case of Pre-Registration (for Post-Registration we refer to [5]). Denote by t_{GFA}^1 the chosen arrival instant of the first observed packet at the GFA. The arrival of the k -th packet then equals $t_{GFA}^k = t_{GFA}^1 + (k-1)T$. The arrival of this packet at the oFA or nFA is denoted by t_{oFA}^k or t_{nFA}^k respectively, and both are random variables. The following stochastic events determine the path of the packet and thus the delay distribution:

1. $t_{GFA}^k < t_{1,pre} \rightarrow k$ -th packet is routed via the oFA
 - 1a. $t_{oFA}^k < t_0 + D_{LD} \rightarrow k$ -th packet is forwarded from the oFA to the MN
 - 1b. $t_{oFA}^k > t_0 + D_{LD} \rightarrow k$ -th packet is lost
2. $t_{GFA}^k > t_{1,pre} \rightarrow k$ -th packet is routed via the nFA
 - 2a. $t_{nFA}^k < t_0 + D_{LU} \rightarrow k$ -th packet needs to be buffered, in which case it will be forwarded to the MN at $t_0 + D_{LU}$.
 - 2b. $t_{nFA}^k > t_0 + D_{LU} \rightarrow k$ -th packet is forwarded from the nFA to the MN

The M/M/1 assumption allows us to compute the overall delay distribution of the k -th packet in a fairly straightforward way. The case of Post-Registration is similar but slightly more complicated.

Remark that we can compute the delay distribution either on the assumption of available buffer size at the FAs or on the assumption of absence of buffers. In the latter case we can e.g. compute the expected number of lost packets due to the handoff.

To determine the size of the buffer that should be installed in the FAs in order to avoid or minimize packet loss, we can proceed as follows. We focus on buffering at the nFA, although for Post-Registration, a buffer at the oFA could also be considered. At the nFA the buffer is needed for packets that arrive before the LU trigger occurs. The required buffer size can be determined by computing the distribution of N_b , the number of packets that would be lost in the absence of a buffer. If we denote by $p_{loss}(M)$ the probability that at least 1 packet will be lost at the nFA with buffer capacity M , then we have

$$p_{loss}(M) = P(N_b > M)$$

A specific (random) time interval can be determined for which it holds that a packet is lost if and only if it arrives in that interval at the nFA. The length, denoted by I_b , of that interval has a distribution composed of sums and differences of exponential variables, as follows from our M/M/1 model. Conditioned on I_b , it is straightforward to approximate the distribution of N_b , so we can use

$$P_{loss}(M) \approx \sum_{i=1}^N P(N_b > M | I_{b,i}) P(I_{b,i})$$

where $I_{b,i}; i = 1, K, N$ is some discretization for I_b .

We can then determine the required buffer size by

$$\min M : P_{loss}(M) < 10^{-\alpha}, \text{ where e.g. } \alpha = 5.$$

In the next section we will show some numerical results concerning M , as well as results for the delay distribution of the stream of packets.

3.2 Numerical results

We assume a network topology as depicted in Figure 1. The results that will be shown are all obtained with the following network characteristics. The propagation delays on the links connecting the GFA and the oFA, as well as the links connecting the nFA are all set to $\tau_1 = 5\text{ms}$, while the links connecting the oFA and the nFA have a propagation delay of $\tau_2 = 3\text{ms}$. The service rate μ in each router is set to 1 packet per ms and all routers have a load of $\rho = 0.8$. This leads to an exponentially distributed response time for each router with rate equal to $\mu(1-\rho) = 0.2/\text{ms}$. Furthermore we assume that the CN transmits a packet every $T = 10\text{ms}$ (UDP stream) destined for the MN. These characteristics are arbitrarily chosen and are not essential since the model is used mainly for comparison between Pre- and Post-Registration.

First we present some results for the delay distribution of a stream of packets involved in a handoff. The start of the handoff is set to $t_0 = 0$, and we observe a stream of 30 packets, the first of which arrives at the GFA at $t_{GFA}^1 = -80\text{ms}$. We define the playout time as the maximum allowed end-to-end delay from GFA to MN: if a packet's end-to-end delay exceeds the playout time it will be dropped.

For Figure 4 we assumed there was no buffer capacity available, which means that packets can get lost due to the handoff. The curves depict the expected number of packets from the stream that are dropped due to the expiration of the play-out time or the absence of a buffer. Both the results for Pre-Registration and Post-Registration are shown, for two different values of the time between the LD and the LU trigger. The timing of the LD trigger is set to 60ms since t_0 . Note that these curves tend to the expected number of lost packets due to the absence of buffer capacity when the playout time tends to infinity. It can be seen that Pre-Registration implies more losses than Post-Registration, while the average delay for packets that are not lost is slightly larger for the Post-Registration scheme. The latter follows from the fact that packets using the BET have a longer delay. When the time between the LD and the LU trigger increases, more packets are lost, so more buffer capacity would be needed to avoid losses.

In Figure 5 we show the impact of installing infinite buffers at the FAs. For Post-Registration the expected number of dropped or lost packets in case of infinite buffers is compared with the case of the absence of buffers. Note that the packets that would be lost when there is no buffer available, will obviously not be lost when buffers are installed, but they will experience a larger average end-to-end delay.

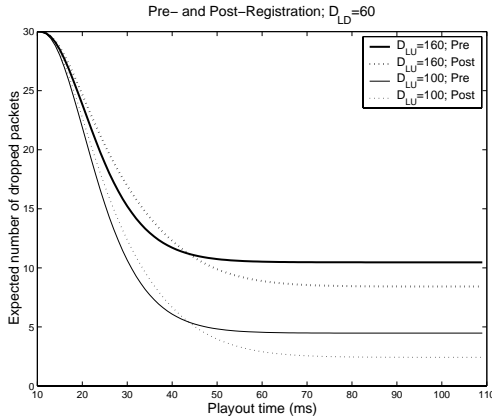


Figure 4: Post- and Pre-Registration

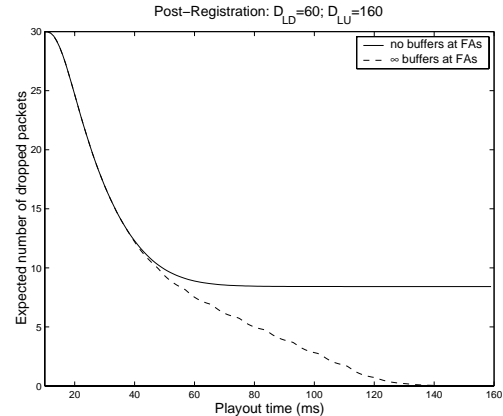


Figure 5: Post-Registration

Figure 6 and Figure 7 depict the delay distribution of each individual packet of the observed stream. In particular the curves present the probability that the GFA-MN end-to-end delay of the k -th packet exceeds t . The timing of the LD and LU trigger is fixed at 80ms and 120ms respectively. Here we assume infinite buffer size available. The curves for the first few packets and the last few packets are all alike, since these packets are not influenced by the handoff. It can then be seen that for Post-Registration the influence of the handoff on the stream's delay starts later as well as lasts longer. This corresponds to e.g. the fact that Post-Registration (optimally) keeps using the link between the oFA and the MN, while Pre-Registration possibly already redirects packets before this link is down. Such packets (e.g. $k=12$ and $k=13$) then have a longer delay since they have to wait for the LU trigger. The disadvantage of Post-Registration is reflected by the fact that it lasts until packet 25 for the influence of the

handoff procedure to have disappeared. Note that packets 19 to 22 have the same delay distribution. These packets, according to our model, all travel with a high probability through the BET, straight to the MN without having to be buffered. In summary, these figures show that Post-Registration is better adjusted to the timing of the triggers, but more packets have a slightly larger delay. A Pre-Registration handoff affects fewer packets, but some of these can have unnecessary high delays.

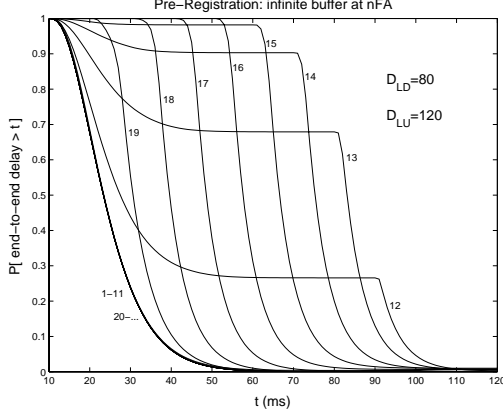


Figure 6: Individual delays – Pre

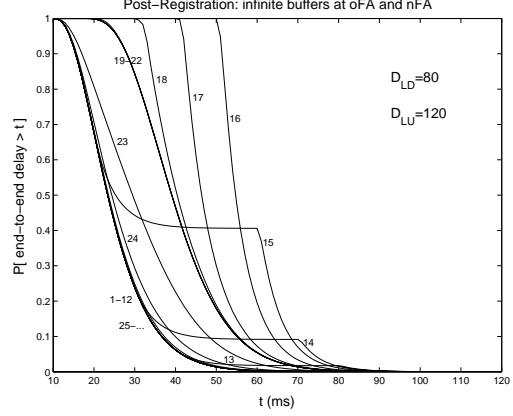


Figure 7: Individual delays – Post

Finally we show some results on the required buffer capacity as a function of the timing of the L2 triggers. Figure 8 and Figure 9 each depict both the expected number of packets that would be lost if there would be no buffer available at the nFA (thick line), as well as the minimum buffer size M to ensure that $P_{loss}(M) < 10^{-5}$ (thin line). Each figure compares Pre-Registration (solid line) to Post-Registration (dotted line). In Figure 8 we fixed the timing of the LU trigger at $D_{LU} = 200$ ms, and we varied D_{LD} . In case of Pre-Registration, only the timing of the LU trigger affects the situation at the nFA, and not the timing of the LD trigger, hence the horizontal curves. As the time between the LD and LU trigger decreases, the number of packets that need to be buffered at the nFA in case of Post-Registration decreases as well. As can be seen from the figure, Post-Registration offers considerable gain over Pre-Registration when the time between the triggers is limited.

In Figure 9 we fix the time between LD and LU and let both vary. The required buffer capacity for Pre-Registration is obviously linearly related to the timing of the LU trigger, while for Post-Registration it is essentially the difference (which is fixed here) between the two triggers that determines the buffer requirements. Concerning Post-Registration, the fact that for smaller values of D_{LU} fewer packets need to be buffered, is explained by noting that in these cases fewer packets are transmitted through the BET because the BET might not yet be established at the time of D_{LD} . These last two figures indicate that Pre-Registration in general, but not necessarily, requires more buffer capacity than Post-Registration.

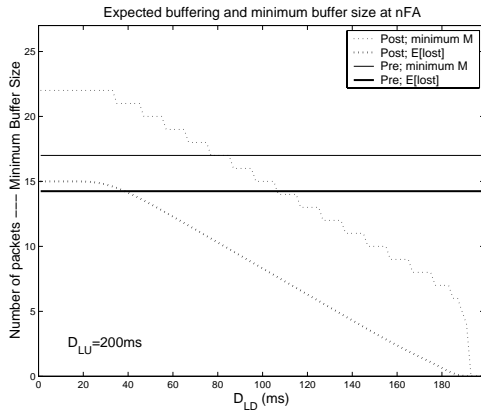


Figure 8: Buffering in nFA

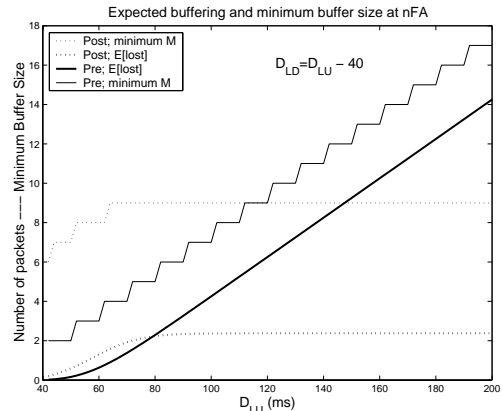


Figure 9: Buffering in nFA

4. Scalability of the Handoff Protocols

4.1 OPNET Model Description

The OPNET simulation models used in this paper are built into OPNET's standard models, extending their functionality and allowing extensive use of the stock router, link and host models. The implemented protocols are Mobile IP [12] extended with Route Optimization [13], Hierarchical MIP [8], Optimized Smooth Handoff [14] and Low Latency Handoffs [7].

Foreign agents are modeled as regular hosts (i.e. as opposed to routers/access points), which are located on the same IP (sub)network as a layer 2 bridge, which in this case is an 802.11b WiFi <-> Ethernet bridge. These hosts are extended with specific process models to handle MIP signaling and tunneling. Mobile nodes are very similar to these hosts, except that they are multi-homed: there is a layer 2 interface for each network they can roam unto (not necessarily using the same access technology on each network). For the course of the paper all of these interfaces are WiFi.

The handover protocols discussed in this paper rely on integration between IP and the underlying access technology in a way that the IP layer needs to receive certain triggers from the underlying one. When using WiFi, these triggers are not always present in the system. In order to evaluate the protocols as proposed by the IETF, 802.11 is used as the MAC layer, but the triggers are provided through other simulation means: the mobile node has a fixed mobility script that is executed and triggers/handoff events can be set relative to each other.

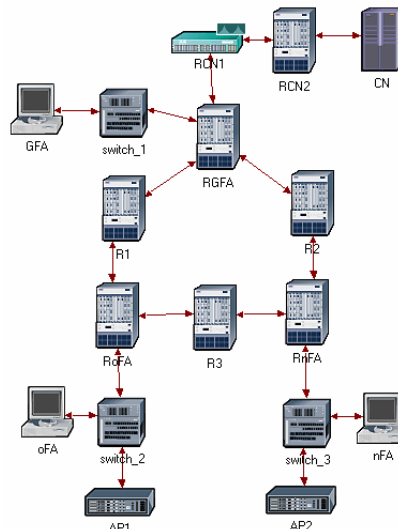


Figure 10 OPNET Network Model

Figure 10 shows a typical OPNET set-up. One hierarchical branch is shown here, along with a corresponding node. The foreign agent entities are connected with the access points through Ethernet switches. The GFA is also a simple host on an Ethernet domain, which is 10Mbps throughput. Special attention is needed when configuring this network for a large number of MNs.

4.2 Results for CBR traffic

The CBR traffic consists of constant-sized (500 bytes) UDP packets sent to all the mobile nodes every 50ms. The mobility script for each MN is as follows: it starts in one of the two domains and draws a uniformly distributed result between 0 and 5s to initiate its first move; the MNs are divided equally among the APs at the start of the simulation. Every MN then waits between 4 and 6 seconds (again uniformly distributed) in one domain before switching to the other.

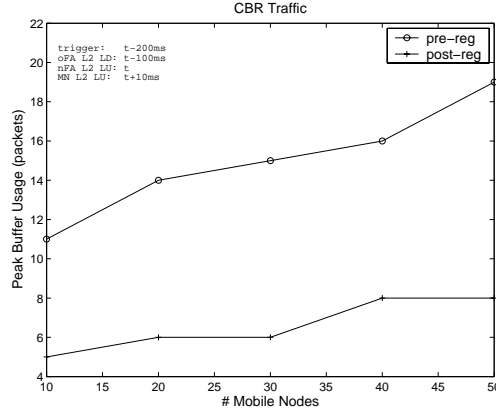


Figure 11 Peak buffer usage

In Figure 11 the peak buffer usage is measured in both foreign agents when each MN conducts 100 handovers. Post-Registration needs significantly less buffers as the time needed to buffer packets depends on the spacing of oFA LD and nFA LU: the bi-directional edge tunnel (BET) is only active between these two events. In Pre-Registration however, the anticipation trigger causes the MN to register with its nFA through the oFA. This registration goes up to the GFA who will change the routing path for the MN's traffic.

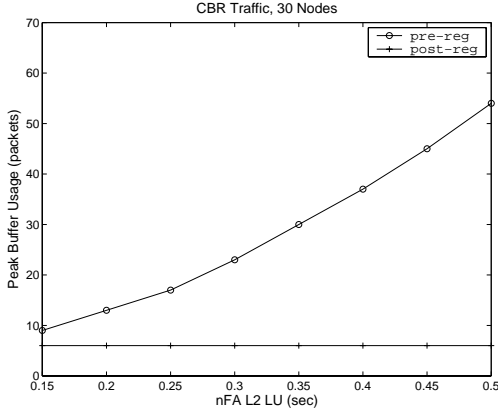


Figure 12 Peak buffer usage varying L2 LU

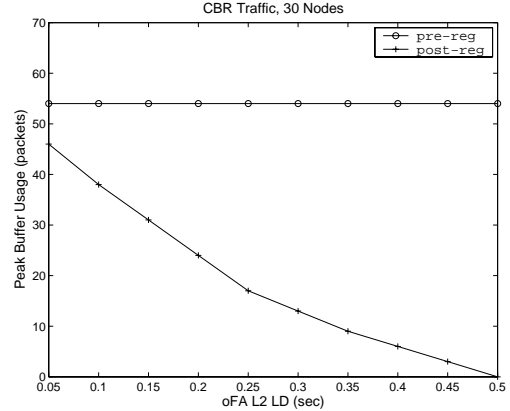


Figure 13 Peak buffer usage varying L2 LD

Figure 12 and Figure 13 show peak buffer usage measured in the foreign agents with a constant number of MNs but when varying both the nFA L2 LU and the oFA L2 LD trigger. In Figure 12 the anticipation trigger arrives at 0 secs, and oFA L2 LD precedes the LU event by 100ms. In Figure 13, oFA L2 LD varies between the anticipation trigger and nFA L2 LU at 500ms. They clearly illustrate the dependence of buffer requirements on one of the triggers: Post-Registration is not influenced by the timing of the anticipation trigger, as long as there is time enough for BET set-up (this is dictated by the protocol requirements). Similarly, Pre-Registration's buffer requirements remain unaffected when varying the oFA L2 LD trigger. Remark that these results are completely in correspondence with the conclusions drawn from the analytical results (see Figure 8 and Figure 9).

A certain amount of traffic delay is of course associated with these handovers. Apart from the delay experienced due to buffering, which will usually be higher in Pre-Registration, the Post-Registration will have an added delay due to use of the BET. This tunnel is used until the MN node chooses to register. Neither pre- nor Post-Registration induce an extra delay when the number of nodes is increased: the buffers in place are logical per-node buffers.

4.3 Results for TCP traffic

TCP traffic consists of FTP file transfers (10MB) to each of the mobile nodes. FTP response time is then measured at each of the mobile nodes (time between start and finish of request) and averaged over the total of nodes. TCP Reno is used in all the runs, but the MSS has been reduced to 1440 bytes to avoid fragmentation induced by tunneling (an extra IP header is added to tunneled packets). Ethernet links are loaded with background traffic to avoid overloading the wireless channel.

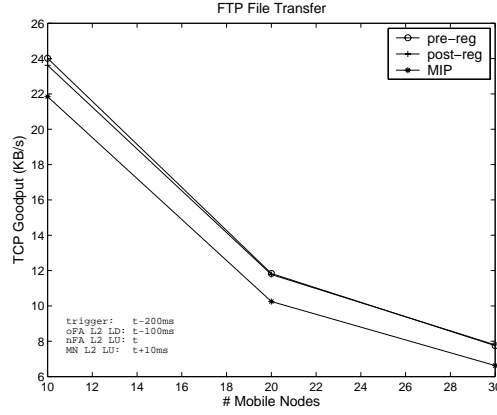


Figure 14 TCP Goodput

Figure 14 shows the advantage of using a low latency handover method as opposed to plain MIP. The smoothness of the handoffs clearly benefits TCP flows, but there is no significant difference between the two methods discussed here. Post-Registration might behave slightly worse because there is a chance of packet re-sequencing: packets following the old route through the tunnel might arrive later than packets sent directly to the nFA if the MN has just performed normal registration.

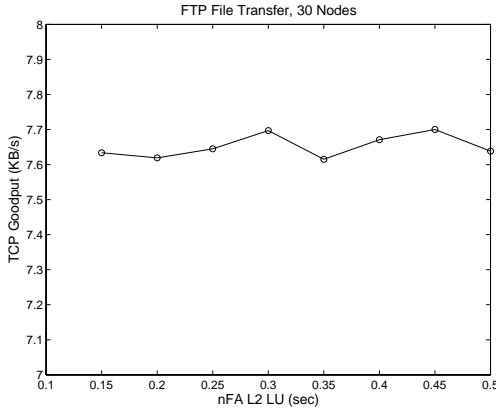


Figure 15 TCP Goodput, Pre-Reg

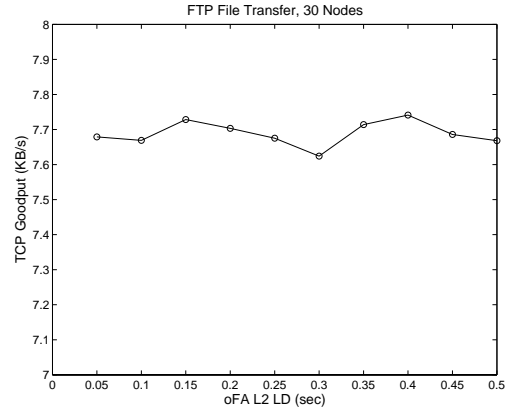


Figure 16 TCP Goodput, Post-Reg

The influence of trigger timing on TCP's performance is negligible: the time that packets spend in the FA's buffer is not large enough to cause performance degradation as the timeout values are larger: TCP commonly has an absolute bottom value of 500ms for timeout timers. This is illustrated in Figure 15 and Figure 16.

5. Conclusions

In this paper we have modeled, analyzed and compared two low latency handoff mechanism proposed by the IETF: the Pre- and Post-Registration Handoff method. Both protocols use Layer 2 triggers to reduce the built-in delay components of Mobile IP. We indicate how the triggers can be implemented when IEEE802.11 is used as link layer. An analytical model allows a detailed investigation of the influence of the timing of the L2 triggers used by the schemes on the delay and packet loss of UDP streams towards mobile nodes involved in a handoff. The numerical results show that with Pre-Registration, a UDP connection experiences more losses than Post-Registration, while the average delay for packets that are not lost is slightly larger for the Post-Registration scheme. A Pre-Registration handoff affects fewer packets, but some of these may experience unnecessary high delays. Finally, Post-Registration offers considerable gain over Pre-Registration when the time between triggers is limited.

The scalability of the protocols in terms of required buffers has been investigated using a simulation model. The results for a high number of MNs receiving UDP streams are in accordance with the analytical results obtained for a single MN. Investigations with TCP traffic show that the trigger timing has a no influence on the TCP goodput.

References

- [1] C. Blondia, O. Casals, P. De Cleyn and G.Willems, "Performance analysis of IP Micro-Mobility Handoff Protocols", Proceedings of Protocols for High Speed Networks 2002 (PfHSN 2002), Berlin 2002, pp 211-226.
- [2] C. Blondia, O. Casals, Ll. Cerdà and G.Willems, "Performance analysis of a forwarding scheme for handoff in HAWAII", Proceedings of Networking 2002, Pisa 2002, pp 504-514.
- [3] C. Blondia, O. Casals, N. Van den Wijngaert, G. Willems, "Performance analysis of smooth handoff in Mobile IP", Proceedings of The Fifth ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM2002), Atlanta, USA, September 2002, p.22-29
- [4] C. Blondia, O. Casals, Ll. Cerdà, N. Van den Wijngaert, G. Willems, P. De Cleyn, "Low Latency Handoff Mechanisms and their implementation in an IEEE 802.11 Network". Submitted to ITC 2003.
- [5] O. Casals, Ll. Cerdà, G. Willems, C. Blondia, N. Van den Wijngaert, "Performance evaluation of the Post-Registration method, a low latency handoff in MIPv4", to appear in Proceedings of ICC 2003.
- [6] A.T. Campbell, J. Gomez, C-Y. Wan, S. Kim, Z. Turanyi, and A. Valko, "Cellular IP", draft-ietf-mobileip-cellular-00.txt, December 1999.
- [7] K. El Malki and others, "Low Latency Handoffs in Mobile IPv4", IETF draft-ietf-monileip-lowlatency-handoffs-v4-04.txt, 2002.
- [8] E. Gustafsson, A. Jonsson, C. Perkins. "Mobile IP Regional Registration", draft-ietf-mobileip-reg-tunnel-02.txt, March 2000.
- [9] D. Johnson and C. Perkins, "Mobility Support in Ipv6", www.ietf.org/internet-drafts/draft-ietf-mobileip-ipv6-19.txt.
- [10] OPNET Modeler, www.opnet.com.
- [11] M. Ohta, "Smooth Handover over IEEE 802.11 Wireless LAN", IETF draft-ohta-smooth-handover-wlan-00.txt, 2002, June 2002.
- [12] C.E. Perkins. "IP Mobility Support", rfc 2002, October 1996.
- [13] C. Perkins and D. Johnson, "Route Optimization in Mobile IP", IETF Internet Draft, draft-ietf-mobileip-optim-11.txt.
- [14] C. Perkins and K-Y. Wang. "Optimized smooth handoffs in Mobile IP", Proceedings of IEEE Symposium on Computers and Communications, Egypt, July '99.
- [15] R. Ramjee, T. La Porta, S. Thuel, K. Varadhan, and L. Salgarelli, "IP micro-mobility support using HAWAII", draft-ietf-mobileip-hawaii-00.txt, June 1999.
- [16] H. Soliman, Editor, "Hierarchical MIPv6 mobility management", draft-ietf-mobileip-hmipv6-07.txt, October 2002
- [17] N. Van den Wijngaert and C. Blondia, "Integration of IP mobility in OPNET: modeling and simulation", proceedings of OPNETWORK 2002, Washington D.C., August 2002.
- [18] ANSI/IEEE Std 802.11, 1999 Edition.
- [19] ns with micromobility support, <http://comet.ctr.columbia.edu/micromobility>.