Performance Analysis of Optimized Smooth Handoff in Mobile IP

C. Blondia, N. Van den Wijngaert, G. Willems
University of Antwerp,
Dept. Mathematics and Computer Science
Universiteitsplein 1,
B-2610 Antwerp - Belgium
Tel: +32 3 8202404
chris.blondia@ua.ac.be

O. Casals
Technical University of Catalonia (UPC),
Dept. Computer Architecture
Jordi Girona 1-3, Mòdul. D6,
E-08034 Barcelona – Spain
Tel: + 34 93 4016985
olga@ac.upc.es

ABSTRACT
Mobile IP allows node mobility involving changes of point-of-attachment to the Internet. In order to reduce the impact on the performance and the signaling overhead, hierarchical mobility management schemes have been introduced. These schemes define protocols that allow movements within a domain to be handled locally, without involvement of the mobile node's home network. In order to reduce more the packet losses during handoff, new schemes have been defined, such as smooth handoff. By storing packets temporarily in the access point after the mobile host has left and forwarding them to the new access point as soon as the mobile has connected to it, it is possible to reduce significantly the packet loss. In this paper we develop an analytical model and a simulation program using OPNET Modeler to evaluate the packet loss and packet delay for UDP streams and the throughput for TCP streams that are involved in a handoff. We show that, in spite of the buffering capabilities of the previous access point, packets may still get lost. The reason for this loss is identified and solutions to this problem are proposed.

Categories and Subject Descriptors
C.2.2 [**Computer-Communication Networks**]: Network Protocols – *routing protocols.*

General Terms
Performance

Keywords
Mobile IP, Smooth Handoff, Performance Analysis, OPNET, Analytical Modelling, Micro Mobility Management.

1. INTRODUCTION
With a growing number of portable computing devices like laptops and PDAs, the need for seamless connectivity to the global Internet is driving the acceptance of different mobility solutions.

Mobile IP (MIP) [6] allows node mobility across media of similar or dissimilar types and allows a Mobile Node (MN) to communicate using only its *home address* while changing its point-of-attachment to the Internet.

When a MN is visiting a foreign network it obtains a *care-of address* (CoA) that will generally change every time the MN moves from one foreign network to another. MIP uses *tunneling*, when the MN is connected to a foreign network, to deliver packets that are destined to the MN's home address. The purpose of a tunnel is to encapsulate an original packet destined to the MN's home address within a packet destined to the care-of address. At the care-of address, the original packet is extracted from the tunnel and is delivered to the MN.

Tunneling requires a MN to report its current care-of address to its *Home Agent* (HA), in a procedure called *Registration*. This requires a procedure called *Agent Discovery* that enables a MN to determine its current location and obtain a care-of address on a foreign network.

MIP is an appropriate solution to handle global IP mobility (*macro-mobility*) but is not optimized to handle *micro-mobility* management. The MN's care-of address changes each time the user moves between neighboring Base Stations, resulting in undesirable notifications to the Home Agent and the Correspondent Node (CN) on every handoff. In such an environment with frequent handoffs, low-latency handoffs are essential to avoid performance degradation and signaling overhead.

The basic MIP protocol forces all packets for a MN to be routed through its HA. Therefore, packets to the MN are often routed along paths that are significantly longer than optimal. The IETF Draft on *route optimization* [7] defines extensions to the operation of the basic MIP protocol so that packets can be routed from a Correspondent Node (CN) to a MN without going first to the HA. The *Foreign Agent Smooth Handoff* proposal described in [7] and [8] provides a means for the MN's previous foreign agent (pFA) to be notified of the MN's new mobility binding, allowing packets in flight to the MN's pFA to be forwarded directly to its new care-of address.

Hierarchical mobility management has been introduced to reduce the impact of mobility by handling local movements locally and hiding them from HAs. In that case, the MN's address known by a HA represents the address of a gateway common to a number of network access points. When a MN moves from one access point to another – reachable through the same gateway – the HA does not need to be informed. The role of micro-mobility protocols [4] is to ensure that packets arriving at the gateway are forwarded to the appropriate access point.

Two well known proposals based on mobile-specific routing, namely Cellular IP [3] and HAWAII [9] have been analysed in [1], [2], [12]. In [10] a comparison of several micro-mobility protocols is presented, showing how the main Mobile IP issues are addressed by each protocol.

In our paper we use an analytical model based on the ones developed to model Cellular IP and HAWAII ([1], [2]) to evaluate the performance of smooth handoff ([7], [8]) obtaining the packet loss as a function of the forwarding buffer and the link delays. The results are compared with an OPNET simulation.

We first describe the protocol and illustrate its operation by means of a trace using an OPNET simulation in section 2. In section 3 we present the analytical model and in section 4 the simulation model is described. Section 5 shows the results obtained using the analytical model for UDP traffic and the behaviour of a TCP connection involved in a handoff using the OPNET simulation. Finally section 6 is devoted to the conclusions.

2. SMOOTH HANDOFF IN MOBILE IP

For the remainder of this paper, we make the following assumptions. From the IETF draft [3] used as a starting point for this analysis, only the part on "FA smooth handoff" is discussed. The route optimization problem (i.e. the triangle routing problem) is assumed solved adequately. Furthermore, in what follows, we assume a MIPv4 network with a hierarchical Foreign Agent (FA) architecture. The hierarchical tunneling approach [5] relies on a tree-like structure of foreign agents. The packets are encapsulated at the HA and delivered to the root FA. Each FA of the tree decapsulates and then re-encapsulates the packets as they are forwarded down the tree of FAs toward the MN's point-of-attachment. As the MN moves from one point of attachment to another, the location updates are made at the optimal point of the tree, tunneling packets to the new access point. We consider a hierarchical scheme with only one level, i.e. where the cross-over FA is the Gateway FA (GFA), although the approach applies to hierarchies with multiple levels as well.

2.1. Foreign Agent discovery

There are two ways to discover a new FA: an agent is required to advertise its service on regular instants (= *agent advertisements*); alternatively, a Mobile Node (MN) may request an advertisement by broadcasting an *agent solicitation* message. Agent advertisements are also used by the MN to check if it has entered a new subnetwork.

As soon as the MN has obtained its new regional CoA, it will register this address with its GFA (note that due to the hierarchical structure, the HA need not to be informed about this change of CoA). This is achieved by sending a *Registration Request Message* to the new FA, who on his turn sends a registration message to the GFA. As part of this registration procedure, the MN may add to the registration request message a *Previous Foreign Agent Notification extension*. The new FA will then send a *Binding Update Message* to the previous FA, with the request to reply with an acknowledgement. This ack is tunneled to the new FA, who should forward it to the MN (this may involve unauthorized traffic to the MN). In this way the MN is notified that the previous FA has the new CoA of the MN. The reason for this binding update is explained in what follows.

2.2. Smooth Handoff: Forwarding with Buffering

As long as the GFA does not receive the registration message with the new CoA from the new FA, it continues to transmit packets to the previous FA. The binding update the previous FA received from the new FA, allows the following forwarding mechanism (see [8]). When a packet arrives in the previous FA, the binding cache is checked and the packet is tunneled to the new FA, who delivers it to the MN. However, packets arriving at the previous FA after the MN left and before the binding update message from the new FA is received are lost. In order to avoid this packet loss, FAs are provided with a circular buffer referred to as the *Forwarding Buffer*. When a tunneled packet arrives at the previous FA, it is decapsulated, delivered to the MN (if possible) and copied into a buffer. When the previous FA receives a binding update originating from a previous foreign agent notification, these buffered packets are re-tunneled to the new FA and all packets arriving at the previous FA with destination the MN are immediately tunneled to the new FA. In order to avoid duplicate packets, the MN includes the pair of source address and datagram identification of the most recent received packets in the registration request that is sent to the previous FA, who uses this pair to drop the buffered packets that have been received by the MH.

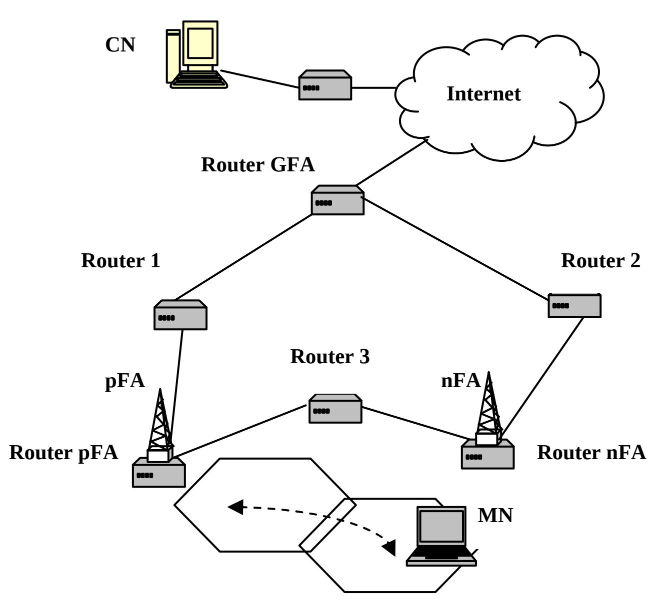

Figure 1: Network architecture

As soon as the new FA receives the registration reply from the GFA, it is allowed to forward the tunneled packets coming from the previous FA to the MN. Packets arriving at the previous FA before the registration reply has arrived are considered as unauthorized traffic and therefore are discarded by the new FA.

Hence, two sources of packet loss can be identified: firstly, packets may be pushed out of the previous FA's circular forwarding buffer due to overflow and secondly, packets tunneled by the previous FA to the new FA that arrive before the registration reply are dropped.

2.3. Events in Smooth Handoff

From the above discussion, it follows that smooth handoff in a hierarchical MIP network involves the following sequence of events:

- The MN moves from subnetwork A to subnetwork B.
- After a time Δ_{FA} the MH obtains a new CoA.
- The MN registers this new CoA with its GFA.
    i. A registration message is sent to the new FA.
    ii. The new FA sends this registration request to the GFA.
    iii. The MN may add a previous FA notification extension.
- The new FA sends a binding update to the previous FA.
- Upon receipt of this binding update, the previous FA adds an entry in its binding cache, and forwards all packets that are in his buffer to the new FA, which have not been received yet by the MN.
- After the registration message is processed by the GFA, a registration reply is sent back to the MN (via the new FA).
- As long as the new FA does not receive this reply, it drops all forwarded packets (these packets are considered as unauthorized as long as no registration reply is received from the GFA).
- After the registration reply is received, the new FA transmits the forwarded packets to the MN.

We illustrate the operation of smooth handoff by means of a trace obtained by a simulation using the OPNET modeler. Figure 2 shows the different signaling events together with the end-to-end delay packets experience. A handoff occurs at time instant 0 and the simulation starts 100 ms earlier (indicated on the X-axis). The end-to-end delays (between the GFA and the MN) of the packets are depicted on the Y-axis. We assume the network topology as depicted in
Figure 1 and use the following parameters. The propagation delays are given as follows: on the links connecting the gateway and the previous FA packets experience a propagation delay of 5 ms, on the links connecting the Gateway and the new FA a propagation delay of 10 ms and finally on the links connecting the previous FA and the new FA a propagation delay of 5 ms. From the figure, we see that packets 1 – 8 are directly sent from the pFA to the MN and experience a delay of around 25 ms (2 link propagation delays of 5 ms and 3 router delays of 5 ms). Packet 9 and following arrive when the MN has no layer 2 connection with the previous FA any longer. Packets 9, 10 and 11 are stored in the forwarding buffer of the previous FA. They will be forwarded to the new FA as soon as the binding update arrives at the previous FA. This batch of packets arrives at the new FA before the registration reply has arrived and is therefore lost. Also packet 12, who is arriving at the previous FA after the binding update has arrived, is forwarded immediately to the new FA. As it arrives before the registration reply it is also lost. Packet 13 and 14 are reaching the new FA via the previous FA, but after the registration reply has arrived, and therefore they are forwarded to the MN. Remark the extra delay due to the forwarding mechanism. Packet 15 and the following packets arrive at the GFA after the registration request has arrived and therefore they are routed directly to the new FA. They experience a delay of 35 ms (2 link propagation delays of 10 ms and 3 router delays of 5 ms). Remark that due to the forwarding mechanism packet 15 arrives before packet 14 at the MN.

Remark: Smooth Handoff when no Hierarchy of FA is used

When no hierarchy of FAs is used, the Corresponding Node (CN) needs to be informed about the change of point of attachment of the MN in order to avoid triangle routing. The CN maintains a binding cache containing the current CoA of the MN it is communicating with. When the MN leaves subnetwork A, the CN obtains the new CoA using binding update messages. As long as the CN has not received this binding update message, it continues to transmit packets to the previous FA. Here again the above forwarding mechanism may be applied.

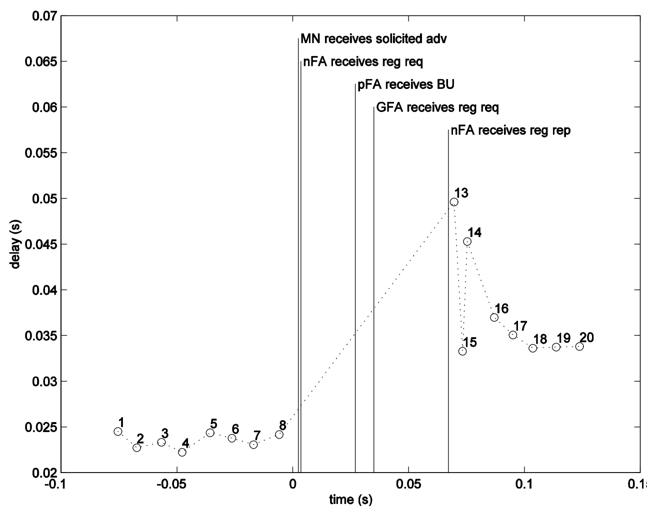

Figure 2: Smooth handoff - an example

3. THE ANALYTICAL MODEL

We now present a mathematical model for the smooth handoff scheme based on a queuing network as in [1] and [2]. We assume a network architecture as depicted in
Figure 1. The following assumption is essential for computational tractability reasons. All routers are modeled as simple M/M/1 queues. The exponentially distributed service time of a packet includes both the processing time and the transmission time. Denote the service rate of Router i (i= 1, 2, 3, GFA, pFA and nFA) by μ, and the load by ρ, then its response time random variable R_i is exponentially distributed with rate $\mu(1-\rho)$.
In order to model the handoff procedure, let us define the following time instants.

- t_0: the time instant the MN leaves subnetwork A (and hence has no layer 2 connection any longer with it) and enters subnetwork B.

- t_1: the time instant the binding update message containing the new CoA of the MN, sent by the new FA, reaches the previous FA;
- t_2: the time instant the regional registration request reaches the GFA and is processed by the GFA
- t_3: the time instant the regional registration reply, originating from the GFA, reaches the new FA.

In our M/M/1 queuing model, the instances t_1, t_2 and t_3 are random variables distributed as sums of exponentially distributed random variables and constants (conditioned on a fixed value of Δ_{FA}):

- $t_1 = t_0 + \Delta_{FA} + R_{nFA} + R_3 + R_{pFA} +$ fixed link delays
- $t_2 = t_0 + \Delta_{FA} + R_{nFA} + R_2 + R_{GFA} +$ fixed link delays
- $t_3 = t_2 + R_{GFA} + R_2 + R_{nFA} +$ fixed link delays

Each packet of a stream belongs to exactly one of the following classes or subclasses:

- Class 0 : packets routed via the previous FA and directly forwarded to the MN
- Class 1: packets routed via the previous FA and buffered before being forwarded to the new FA.
  - Subclass (a): packets forwarded but lost because they arrive at the new FA before the Registration Reply.
  - Subclass (b): packets forwarded and arriving at the new FA after the Registration Reply.
  - Subclass (c): packets that are lost due to buffer overflow at the previous FA.
- Class 2: packets routed via the previous FA and directly forwarded to the new FA.
  - Subclass (a): packets lost because they arrive at the new FA before the Registration Reply.
  - Subclass (b): packets arriving at the new FA after the Registration Reply.
- Class 3: packets routed via the new FA.

Remark that subclasses can be empty.

Now consider a UDP stream originating from a CN destined to the MN. The handoff does not affect the path of the stream until it reaches the GFA, therefore we adopt the point of view of packets arriving at the GFA. We assume that every T ms a packet arrives at the GFA (the jitter introduced by the network between CN and GFA is not taken into account). Let us denote the time of arrival in the GFA by t_{GFA}.

Packets are lost if they belong to subclasses 1(a), 1(c) or 2(a). So, the probability that a packet will be lost equals

$$P[\text{packet lost}] = P[1(\text{a})] + P[1(\text{c})] + P[2(\text{a})]$$

The different probabilities of the right hand side are obtained as follows:

$$\begin{aligned}P[1(a)] &= P[(t_{GFA} + c < t_2) \text{ AND } (t_0 < X < t_1) \text{ AND } (X' > t_1)\\ &\quad \text{AND } (t_1 + Y' < t_3)] + P[(t_{GFA} < t_2 < t_{GFA} + c)\\ &\quad \text{AND } (t_0 < X < t_1) \text{ AND } (t_1 + Y' < t_3)]\end{aligned}$$

$$P[1(c)] = P[t_{GFA} + c < t_2 \text{ AND } (t_0 < X < t_1) \text{ AND } (X' < t_1)]$$

$$P[2(a)] = P[t_{GFA} < t_2 \text{ AND } (X > t_1) \text{ AND } (X + Y < t_3)]$$

where the random variables X, X', Y and Z, and c are given by

$$X = \{t_{GFA} + R_{GFA} + R_1 + R_{pFA} + \text{fixed delays}\}$$

$$X' = \{t_{GFA} + c + R_{GFA} + R_1 + R_{pFA} + \text{fixed delays}\}$$

$$Y = \{R_{pFA} + R_3 + R_{nFA} + \text{fixed delays}\}$$

$$Y' = \{\text{burst delay} + R_{pFA} + R_3 + R_{nFA} + \text{fixed delays}\}$$

$$c = BS \times T\ ,$$

where BS denotes the size of the forwarding buffer.

The burst delay is determined as the expected number of packets in the queue in front of the current packet, and counting one extra service time per packet per router.

Remark that all these random variables are the sums of three independent exponential variables with rate $\mu\ (1\text{-}\rho)$ and some constants, hence the computation of $P[1(\text{a})]$, $P[1(\text{c})]$ and $P[2(\text{a})]$ is fairly straightforward.

As for the delay distribution we have that

$$P[\text{delay} > t] = P[\text{packet lost}] + \sum_{class \in A} P[\text{class AND delay} > t]$$

where $A = \{0, 1(\text{b}), 2(\text{b}), 3\}$.

Here we have that

$$\begin{aligned}P[0 \text{ AND delay} > t] &= P[(t_{GFA} < t_2) \text{ AND } (X < t_0)\\ &\quad \text{AND } (\text{delay} > t)]\end{aligned}$$

$$\begin{aligned}P[1(\text{b}) \text{ AND delay} > t] &= P[(t_{GFA} + c < t_2) \text{ AND } (t_0 < X < t_1)\\ &\quad \text{AND } (X' > t_1) \text{ AND } (t_1 + Y') > t_3) \text{ AND } (\text{delay} > t)]\\ &\quad + P[(t_{GFA} < t_2 < t_{GFA} + c) \text{ AND } (t_0 < X < t_1)\\ &\quad \text{AND } (t_1 + Y' > t_3) \text{ AND } (\text{delay} > t)]\end{aligned}$$

$$\begin{aligned}P[2(\text{b}) \text{ AND delay} > t] &= P[(t_{GFA} < t_2) \text{ AND } (X > t_1)\\ &\quad \text{AND } (X + Y > t_3) \text{ AND } (\text{delay} > t)]\end{aligned}$$

$$P[3 \text{ AND delay} > t] = P[(t_{GFA} > t_2) \text{ AND } (\text{delay} > t)]\ .$$

The delay is a random variable that takes on different forms according to the class:

- 0 : delay $= X - t_{GFA}$
- 1(b) : delay $= t_1 + Y' - t_{GFA}$
- 2(b) : delay $= X + Y - t_{GFA}$
- 3 : delay $= Z - t_{GFA}$,

where $Z = \{t_{GFA} + R_{GFA} + R_2 + R_{nFA} + \text{fixed delays}\}$ and the other variables are defined as before. Thus the delay is the time of arrival in the current FA minus the departure time t_{GFA}.

When we are interested in the total end-to-end delay CN-MN, we approximate this by adding the expected end-to-end delay CN-GFA, which is the same for every packet.

The M/M/1 assumption allows us to compute each of these probabilities in a fairly straightforward way.

In order to compute the expected number of lost packets due to the handoff, we can proceed as follows. If we set the instance of handoff $t_0 = 0$, then we can compute the loss probability for a number of N consecutive packets, starting sufficiently before the handoff, say $t_{GFA} = -100$, and ending sufficiently after the handoff. The expected number of lost packets for such a stream is then given by the sum of the individual probabilities:

$$E[\text{number of lost packets}] = \sum_{k=1}^{N} P[\text{lost}, t_{GFA} = -100 + (k-1) \times T]$$

The result does not depend on the length of the stream that is considered here, as long as the first and last packet considered have negligible loss probabilities.

4. THE SIMULATION MODEL

OPNET Modeler is a discrete-time event simulator with an extensive and very detailed model library. In order to preserve all the features in the application models, TCP/IP stack and WLAN models, Mobile IP (MIP) [6] has been integrated into this standard OPNET-supplied framework. The two modes of operation of MIP have been implemented (foreign agent and co-located care-of address), including some extensions: Smooth Handoff [7], Optimized Smooth Handoff [8] and Regional Tunnel Management (Hierarchical MIP) [5].

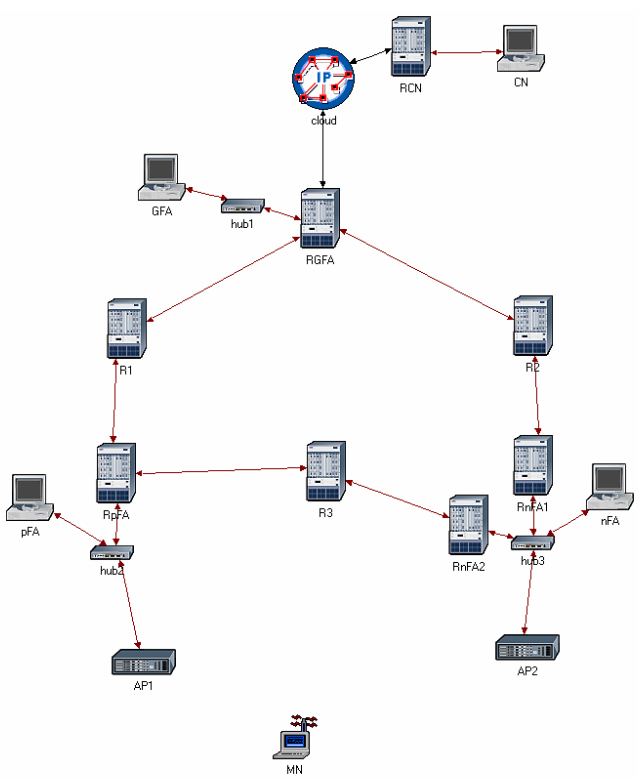

Figure 3: OPNET network set-up

The work on this implementation has been two-fold. Firstly, the standard IP routing logic had to be modified in several places, e.g. a foreign agent must be able to forward traffic for a registered mobile node at all times (i.e. act as a default gateway for the MN), home addresses of MNs should be considered local on the link, etc... The ARP model had to be modified as well to allow for techniques like Proxy ARP and Gratuitous ARP. Secondly, several new processes needed to be developed in order to allow correct protocol functionality: IP-in-IP tunneling, agent advertisement and solicitation support, handling of registration messages. These models are in strict accordance with their respective RFC or internet draft, and this in combination with the fact that this is all integrated with OPNET's code gives us a rich framework in which to run simulations in.

The network used for the purpose of this paper is shown in Figure 3.

All the links in the network are 10Mb/s Ethernet connections, with the connections between the different routers in the branch under the GFA having a fixed delay of 5ms. The links in the path from the GFA to the nFA are changed to achieve different path lengths in the branches of the tree. All the other links have a fixed delay of 0ms. The routers in the network have a background load of 80% on the central CPU in charge of the IP forwarding (set to 1000 packets/s). Due to the specific OPNET implementation of these background-loaded routers, the topology in the network of the nFA has been changed somewhat to avoid unnecessary delays when two packets are submitted to it at the same time. This would lead to discrepancies with the analytical model behaviour as the registration request to the GFA and the binding update to the pFA are sent at the same time during handover. Access to the network from the mobile node is done through 802.11b WLAN interfaces, with AP1 and AP2 being WLAN-Ethernet bridges.

If the mobile node's handover time is at t seconds, the correspondent node (CN) starts sending UDP packets at t-0.1 until t+0.1 seconds, generating at 500bit packet every 10ms.

5. NUMERICAL RESULTS

In this section we validate the analytical model by means of a simulation performed using the OPNET modeler. Once validated, the models are used to illustrate the influence of different parameters (distance between routers, buffer capacity) on the system performance (packet loss, delay and throughput during handoff).

The analytical model is used to evaluate both the packet loss and the induced delay on the packets of a UDP stream during handoff. A second class of results are obtained using the OPNET simulation program for a TCP stream involved in a handoff, for which the throughput is evaluated.

5.1. Performance Measures for a UDP Stream using the Analytical Model

First we investigate the loss probability of packets in the forwarding buffer of the previous FA and the loss probability of forwarded packets in the new FA. At the same time, these results are validated using the OPNET simulation described in Section 4.

Consider the network depicted in
Figure 1 with the following system parameters. Each router is loaded up to 0.8, the propagation delay between routers is $\tau = 5$ ms, the average processing time of a packet in a router is 1ms.

In Figure 4 and Figure 5 the expected number of lost packets is shown as a function of the buffer size at the previous FA. The analytical results (A) are compared to simulation results (S). The expected packet loss due to buffer overflow is given by the dashed line, while the solid line represents the additional loss at the new FA, due to early arrival. Figure 4 shows the results for link delays equal to 5ms on every link, while for Figure 5, the 2 links on the nFA-GFA path are increased to 10ms each.

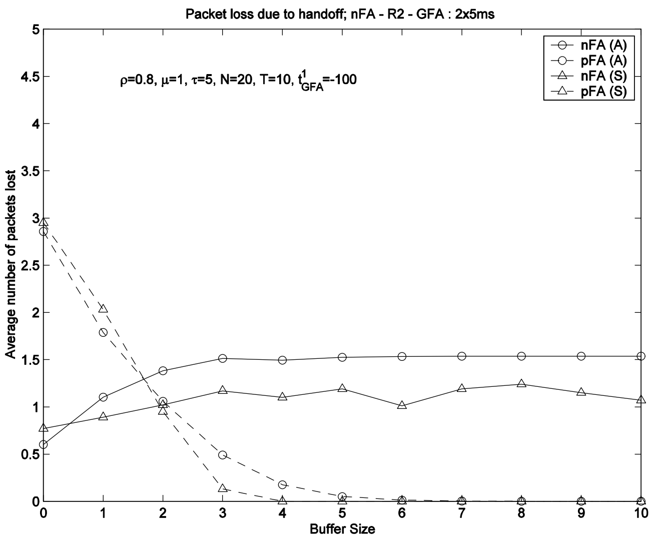

Figure 4: Packet loss as a function of the buffer size

Obviously, the loss in the buffer at the previous FA diminishes when the buffer size is increased. The packets that are lost in case of very small buffer size do go through the buffer when this buffer size increases. But in the latter case they possibly contribute to the number of lost packets at the new FA, hence the rise of the solid curve for small buffer sizes. This is especially true for the case of 10ms link delay on the nFA-GFA path, because then the whole buffer is likely to arrive too early at the new FA (i.e. before the registration reply message from the GFA has arrived). In other words, in the case of long delay on the nFA-GFA path, if a packet is not dropped at the previous FA, it will most likely be lost at the new FA.

In order to avoid packet loss at the previous FA, the forwarding buffer need to be dimensioned such that it can store packets of the order of the product bit rate of the stream times delay (MN -new FA – previous FA). The loss at the new FA on the other hand depends on the difference between the distance (new FA – GFA) and (new FA – previous FA). If the latter is smaller than the former, then packets may get lost. A possible solution would be to provide the new FA with a buffer to store temporarily unauthorized traffic until the registration reply from the GFA arrives at the new FA. Another solution consists of sending the binding update message from the new FA to the previous FA via the GFA in order to allow the registration reply message to arrive before the first forwarded packet. A similar solution has been applied in the Multiple Stream Forwarding scheme of the HAWAII protocol (see [9]).

In the next figures we show the delay experienced by individual packets. We consider a stream of 20 packets, the first packet of which has t_{GFA} = -100, and T = 10ms. We set ρ = 0.8 and μ = 1 in every router. The fixed link delays are set at 5ms each.

Figure 6 and Figure 7 depict the delay distributions of the specified packets of this stream, for the capacity of the forwarding buffer given by BS=0 and BS=6 respectively. The delay that is depicted is the end-to-end delay from GFA to the MN. As noted earlier, we consider the CN-GFA delay to be the same for each packet and do not take it into account here.

Remark that, as t grows, the curves tend to the loss probability of each packet.

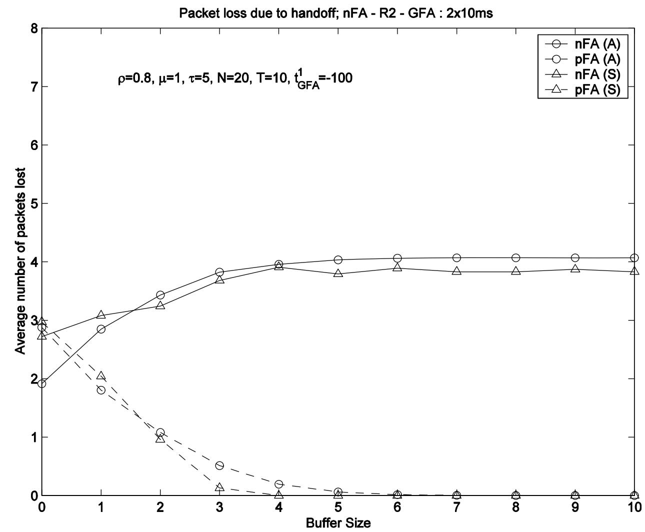

Figure 5: Packet loss as a function of the buffer size

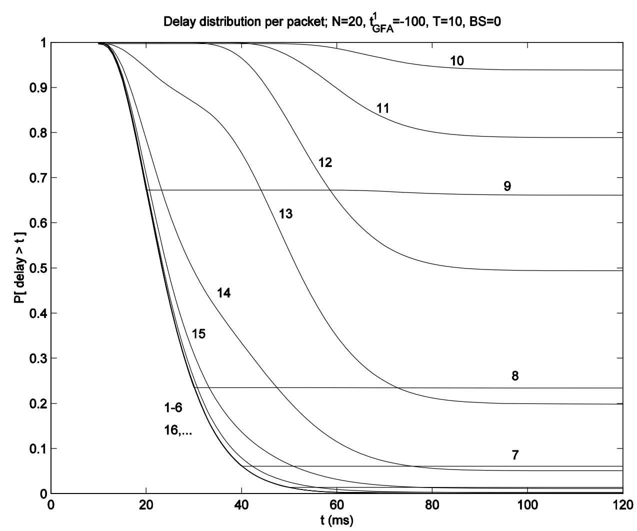

Figure 6: Delay distribution of individual packets (without forwarding buffer)

As can be seen, the first few packets as well as the last ones are not really involved in the handoff. Their delay distribution is the one of a packet traveling directly from the GFA to the MN via either the previous FA or the new FA.

From packet number 7 on, the packets have a significant probability to experience delay due to the handoff. Apparently the 10^{th} packet of the stream is the one that has the most extra delay to expect. It certainly has the highest loss probability. The 10^{th} packet clearly is also the first packet that has zero probability of belonging to class 0, i.e. traveling directly from the previous FA to the MN.

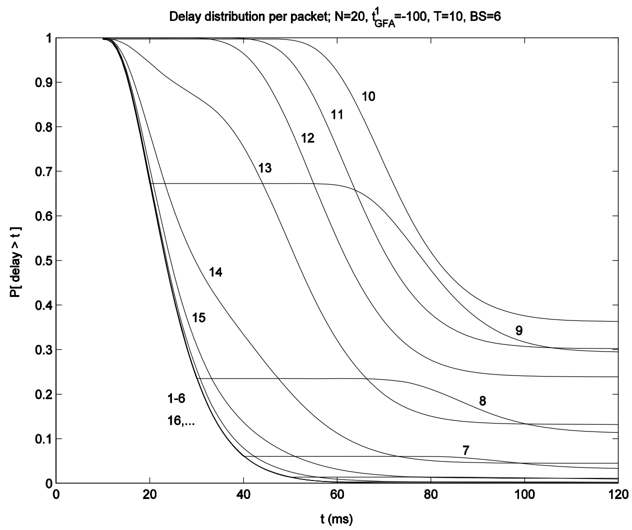

Figure 7: Delay distribution of individual packets (with forwarding buffer)

A comparison of Figure 6 and Figure 7 reveals that packets 8 to 12 are the ones that take most advantage from the presence of the forwarding buffer at the previous FA. The other packets have a low probability of belonging to the class of packets that need to be buffered.

We now consider the same stream of 20 packets, a forwarding buffer size of 3 packets and we vary the value of the router processing rate μ. Figure 8 shows the expected number of packets of the stream that will have a delay larger than *t*, for μ =1, 2 and 5.

As *t* grows, the curves indicate the average number of lost packets.

As can be expected, the number of late or lost packets decreases when the service rate in the routers increases. Indeed, the paths are shorter and so the handoff will be completed sooner.

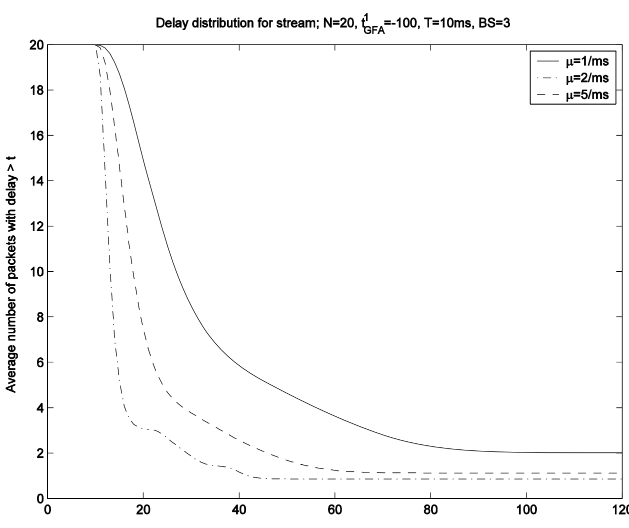

Figure 8: Delay of UDP stream for variable router processing rate

5.2. Performance Measures for a TCP Stream Using the OPNET Simulation

Consider a similar network as depicted in Figure 3. Instead of considering a tagged UDP packet stream from CN to MN, we consider a TCP Reno connection, and we are interested in the received goodput. Remark that possible duplicate packets are eliminated as explained in [8] and Section 2.2.

As was shown in the previous section, packets may get lost during handoff (in the previous FA or in the new FA). This leads to a throughput degradation for TCP traffic. The TCP trace depicted in Figure 9 illustrates this decrease of throughput. The figure on the left hand side shows the different segment arrival times at the MN for a handoff that occurs at time instant 95. The size of the forwarding buffer is assumed to be 1 and the propagation delay on the path new FA – R2 – GFA is given by 2 x 2ms (short enough to allow the registration reply to arrive before the forwarded packets). Hence, packet loss is due to forwarding buffer overflow. From the figure it is clear that one packet gets lost and is retransmitted by the CN. In the right hand side figure, the buffer size is 20 packets, but the path new FA – R2 – GFA is given by 2 x 20ms. This implies that losses do not occur in the forwarding buffer, but at the new FA. In this example, three packets get lost and are retransmitted by the CN.

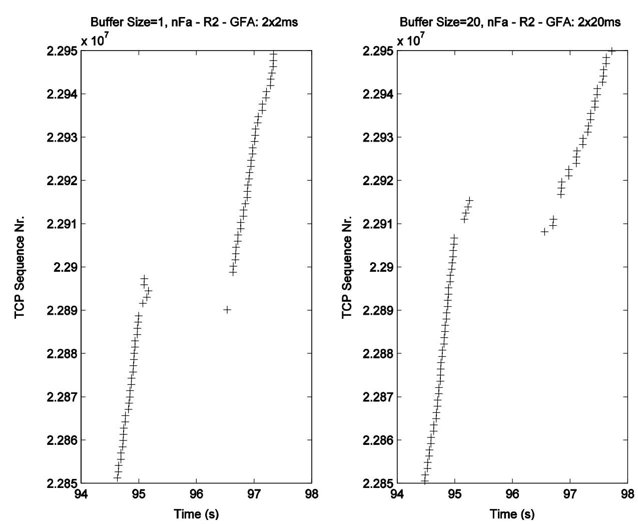

Figure 9: TCP trace during handoff

Now let us illustrate the impact of packet loss on the TCP goodput. The MN initiates an FTP transfer of 1,000,000 bytes from the CN to the MN. As soon as the request is initiated, the MN starts switching between the two networks every three seconds.

From Figure 10, we see that increasing the forwarding buffer leads to a better goodput, as fewer packets get lost. This result is confirming the conclusions drawn in [8]. However, as shown in the previous section, packets forwarded by the previous FA may still get lost at the new FA when they arrive before the registration reply message. This phenomenon is illustrated by means of the two curves in Figure 10: when the distance new FA – GFA increases (the propagation delay on the path new FA – R2 – GFA increases from 2 x 5ms to 2 x 10ms), the registration reply message arrives later and hence the number of lost forwarded packets increases. This implies a degradation of TCP throughput.

This is clearly illustrated in Figure 10, where a goodput degradation of more than 10% may be observed.

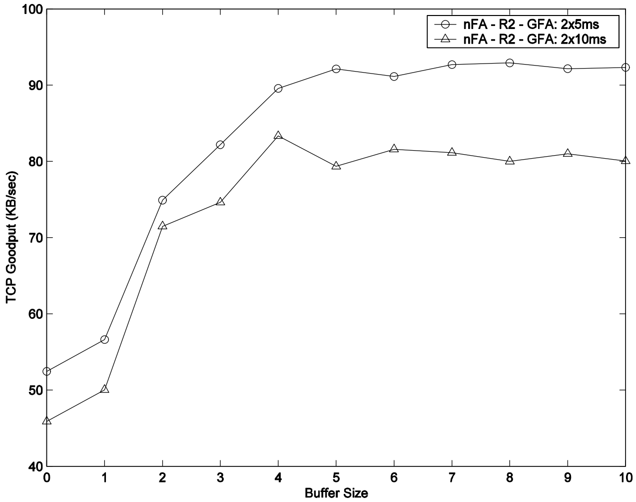

Figure 10: TCP goodput for different forwarding buffer sizes and path delays

6. CONCLUSIONS

In this paper we have presented an analytical model for smooth handoff in MIPv4. This model is used to assess the packet loss and packet delay of a UDP stream that is involved in a handoff. In addition, we have developed a simulation program for this system using the OPNET modeler. The simulation results are not only used to validate the analytical model, but they also allow for an evaluation of the throughput of a TCP connection involved in a handoff. From the results, we may draw the following conclusions. First, the validation presented in Section 5.1 shows the accuracy of the analytical model developed in Section 3. Next, we have seen that packet loss still occurs during handoff and that it may affect both the performance of UDP and TCP traffic. We have seen that the origin of packet loss is two-fold: first, packets may get lost in the previous FA when the forwarding buffer overflows and secondly, they may get lost in the new FA when upon their arrival the registration reply from the GFA has not arrived yet in the new FA. The first reason for loss may be avoided by appropriately dimensioning the forwarding buffer. As a guideline, this buffer should be able to store arriving packets at least during a time equal to the delay on the new FA - previous FA path. The second loss is more difficult to deal with. It is determined by the difference between the delays of the paths previous FA – new FA and new FA – GFA. A number of solutions are possible to solve this problem. Similar to the Multiple Stream Forwarding scheme of the HAWAII protocol, the binding update message sent by the new FA to the previous FA could be routed via the GFA in order to allow the registration reply message to arrive before the first forwarded packets. This however would increase (in some cases unnecessarily) the handoff latency. A second solution consists of storing the forwarded packets temporarily in a buffer at the new FA, until the new registration reply has arrived. This buffer could be dimensioned based on the distance between the FA and its neighboring FAs.

7. ACKNOWLEDGEMENTS

The authors are grateful to Mr. Marc Ribet (OPNET Technologies, Paris) for the valuable support during the preparation of this paper.

8. 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.

[2] C. Blondia, O. Casals, L. Cerda and G.Willems, “Performance analysis of a forwarding scheme for handoff in HAWAII”, Proceedings of Networking 2002, Pisa 2002.

[3] 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.

[4] A. T. Campbell, J. Gomez-Castellanos. “IP Micro-Mobility Protocols”, ACM SIGMOBILE Mobile Computing and Communications Review, October 2000, Vol. 4, No. 4, pp 45-54.

[5] E. Gustafsson, A. Jonsson, C. Perkins. “Mobile IP Regional Registration”, draft- ietf-mobileip-reg-tunnel-02.txt, March 2000.

[6] C. Perkins, Editor. “IP Mobility Support”, RFC 2002, October 1996.

[7] C. Perkins and D. Johnson, “Route Optimization in Mobile IP”, IETF Internet Draft, draft-ietf-mobileip-optim-11.txt

[8] C. Perkins and K-Y. Wang. “Optimized smooth handoffs in Mobile IP”, Proceedings of IEEE Symposium on Computers and Communications, Egypt, July ’99.

[9] 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.

[10] P. Reinbold and O. Bonaventure. “A Survey of IP Micro-Mobility Protocols”, submitted (http://www.infonet.fundp.ac.be).

[11] N. Van den Wijngaert and C. Blondia, “Integration of IP mobility in OPNET: modeling and simulation”, to appear in proceedings of OPNETWORK 2002, Washington D.C., August 2002.

[12] F. Vena, L. Cerda, O. Casals. “Study of the TCP Dynamics over Wireless Networks with Micomobility Support Using the ns Simulator”, European Wireless 2002, Florence, February 2002.

[13] K. El Malki, Editor. “Low Latency Handoffs in Mobile IPv4”, draft-ietf-mobileip-lowlatency-handoffs-v4-3.txt, November 2001