Wireless Networks 10, 17–27, 2004

Study of the TCP Dynamics over Wireless Networks with Micromobility Support Using the ns Simulator

LLORENÇ CERDÀ *, FABIO VENA and OLGA CASALS
Polytechnic University of Catalonia, Computer Architecture Department, c/ Jordi Girona, 1-3, Modulo C6, E-08034 Barcelona, Spain

Abstract. Several IP micro-mobility protocols have been proposed to enhance the performance of Mobile IP in an environment with frequent handoffs. In this paper we make a detailed study of how some of these protocols namely Cellular IP, HAWAII and Hierarchical Mobile IP affect the behavior of TCP and their interaction with the MAC layer. The aim of the paper is to investigate the impact of handoffs on TCP by means of simulation traces that show the evolution of segments and acknowledgments during handoffs.

Keywords: IP micromobility, CIP, HAWAII, hierarchical MIP, TCP, handoff

1. Introduction

The Internet Protocol (IP) is occupying a dominant position in computer networking. The provisioning of an end-to-end IP solution for mobile users gives rise to the problem of how to route the packets to the Mobile Hosts (MH).

Mobile IP (MIP) [11] is the current standard for supporting mobility in an IP network. MIP is an appropriate solution to handle global IP mobility (macromobility) but is not optimized to handle micro-mobility management. The MH's care-of-address changes each time the user moves between neighboring Base Stations (BS), resulting in undesirable notifications to the home agent and the correspondent host (CH) on every handoff. In such an environment with frequent handoffs, low-latency handoffs are essential to avoid performance degradation and signaling overhead.

To meet such requirements a hierarchical mobility management, where Mobile IP is handling the macromobility and local mobility is handeled by more optimised micromobility protocols, seems appropriate (figure 1). In that way a home agent has not to be aware of host movements within an access network (or a region): an MH will keep the same address until it moves to another access network, in which case it will have to notify it to his home agent.

A number of micro-mobility protocols have been discussed in IETF. Two types of hierarchical mobility approach have been proposed:

(i) hierarchical tunneling,
(ii) host-based routing schemes.

Hierarchical tunneling techniques rely on a tree-like structure of foreign agents. Traffic destined to the MH is encapsulated at the home agent and then delivered to the root foreign agent. Each foreign agent in the tree decapsulates and then reencapsulates the packets as they are forwarded down the tree of foreign agents toward the MH. Location updates at the adequate points of the tree follow the MH movement between different BSs so that traffic is tunneled to the MH point of attachment. Proactive Handoff [9], Fast Handoff [5] and Hierarchical Mobile IP [7] are examples of micromobility protocols that use hierarchical tunneling.

Host-based routing schemes avoid the overhead introduced by decapsulation and reencapsulation in the previous approach. Inside the access network, packets are routed using mobile-specific routing without tunneling.

Host-based routing schemes such as Cellular IP (CIP) [1], HAWAII [13], etc. have been proposed to provide micromobility management with minimum packet loss and limited handoff latency. They also aim to minimize signaling through the introduction of paging techniques. In this paper we study the behavior of TCP sources in a mobile environment using Cellular IP, Hierarchical Mobile IP and HAWAII with different path setup schemes. These protocols have been previously evaluated using TCP in [3]. However, in this paper we give detailed traces in order to investigate the interaction between TCP and the micromobility protocols to better understand and find out the possible causes of performance degradation.

The rest of the paper is organized as follows. We first provide a description of each of the micromobility protocols studied in this paper. We then present the simulation framework using the ns simulator and show the influence of these IP micro-mobility protocols on the TCP dynamics and the interaction of the MAC protocol used.

2. Cellular IP

Cellular IP [1,2,4,6] is a proposal to the IETF made by researchers from Columbia University and Ericsson. Besides the Mobile IP protocol engine, Cellular IP MHs have to run a special Cellular IP protocol that controls the micro-mobility support.

In a Cellular IP network packets addressed to a MH are routed to its current BS on a hop-by-hop basis where each node only needs to know on which of its outgoing ports to

---

* Corresponding author.
E-mail: llorenc@ac.upc.es

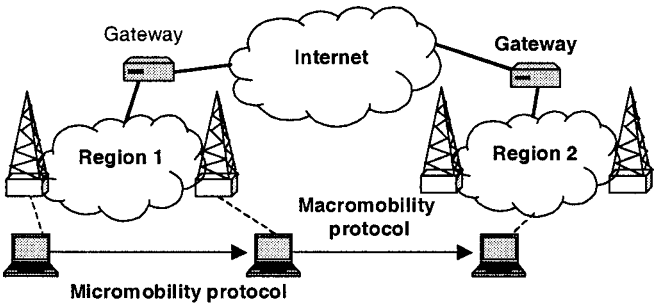

Figure 1. Macro- versus micromobility.

forward packets. These information elements are referred to as mappings because they map MH identifiers (IP addresses) to node ports. Packets transmitted by MHs create mappings. These packets travel in the access network toward the gateway router. By monitoring these packets and by mapping sender address to incoming port, nodes of the access network create a hop-by-hop reverse path for future packets addressed to the given host.

In order to minimize control messaging, mappings are not cleared in an explicit way after handoff. Rather, they are assigned timers to clear outdated mappings.

Cellular IP uses two parallel structures of mappings. Nodes maintain one set of mappings, called *Paging Caches*, for idle MHs. Independent of Paging Caches, nodes maintain another set of mappings called *Routing Caches*. These mappings are only maintained for MHs currently receiving or expecting to receive data. As the routing entries in the nodes are soft state, a MH that has no data to transmit, has to send periodically special IP packets (dummy packets) to maintain its routing path. Paging is used to route packets to MHs that have not been transmitting or receiving data for a while and whose Routing Cache entries have timed out. Paging Caches are not maintained in each node and have a longer timeout value. Cellular IP uses two different Caches because the system has two characteristic time scales. To minimize resource waste due to unused but not yet cleared mappings, the timeout should be in the order of the packet time scale. On the other hand, to avoid overloading the network with dummy packets, the system should operate at the host mobility time scale. The separation of paging and routing allows working with these two different time scales: for Routing Cache mappings the timeout is in the packet time scale and for Paging Cache mappings the timeout is in the migration frequency time scale. Usually a wireless access network will have a large number of hosts of which only a small number will be active at a time and so the Paging Caches will contain a large number of mappings which means a large data base. The separation has also the advantage that a network operator can place Paging Caches in only a small number of nodes. In what follows we explain how Paging and Routing Caches work.

Paging Cache. Idle MHs periodically generate short control packets, called paging-update packets, sending them to the nearest available BS. The paging-update packets travel in the access network toward the gateway router (GW), routed on a hop-by-hop basis. Nodes equipped with Paging Cache monitor passing paging-update packets and maintain the cache that maps MH identifiers to the port through which the paging-update packet arrived.

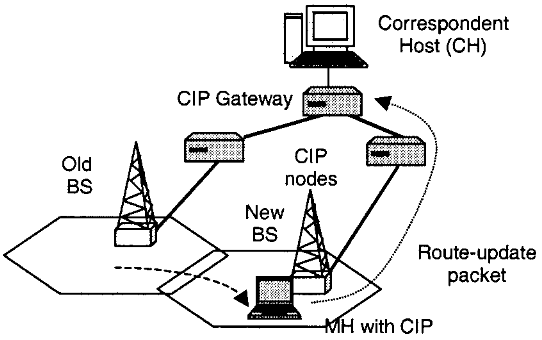

Figure 2. Handoff using CIP.

As the idle host moves, it keeps sending its paging-update packets to the nearest BS, forcing Paging Caches to have up-to-date mappings. Outdated mappings are cleared after a system-specific timeout.

When IP packets arrive at the gateway router, addressed to a MH for which no up-to-date routing information is available, the Paging Caches are used to find the host. The gateway queues the arrived IP packets and generates a control packet, called a paging packet, which contains the identifier of the MH being searched for. The paging packet is routed in the access network by Paging Caches that simply reverse the route taken by recent paging-update packets. If all nodes have Paging Caches, a full hop-by-hop route is available to the host's current location. If some nodes do not have a Paging Cache, then they will forward the paging packet to all outgoing ports.

Routing Cache. Route-update packets transmitted by the MH are routed to the GW on a hop-by-hop basis. Nodes that contain a Routing Cache monitor these packets and use them to create a mapping of host identifiers to port numbers. Data packets are used to refresh these mappings. Packets addressed to the MH are routed along the reverse path, hop-by-hop, by these Routing Caches and are broadcast where no routing information is available.

2.1. *Handoff*

As the host approaches a new BS, it sends a route-update packet to the new BS and redirects its data packets from the old to the new BS (figure 2). The route-update packet will configure a new path of Routing Cache mappings for the host to the new BS. For a time equal to the timeout of Routing Cache mappings, packets addressed to the MH will be delivered at both the old and new BSs. This guarantees that if the host's radio device is capable of listening to two logical channels, the handoff will be soft. If the host cannot listen to both BSs at the same time then the performance of hard handoff will depend on the radio device. After a while, the path to the old BS will time out and clear, while packets will continue to be delivered to the host at its current location via the new BS.

3. Hierarchical Mobile IP

Hierachical Mobile IP (HMIP) [7] is an IETF proposal made by researchers from Ericsson and Nokia that uses a hierarchy of foreign agents. IP registration messages sent by MHs to update their location information establish tunnels between neighboring foreign agents along the path from the MH to a gateway foreign agent. The inclusion of paging extensions in [8] allows idle MHs to operate in a power saving mode. A packet sent to an idle MH will be received by the home agent, tunneled to the paging foreign agent that will page the MH to re-establish a path to its current point of attachment.

4. HAWAII

HAWAII [6,12,13] was proposed to the IETF by researchers from Lucent Bell Labs. Like in Cellular IP, HAWAII is responsible for the micro-mobility support while the macromobility is handled by Mobile IP.

In HAWAII a hierarchy based on domains is used. The gateway into each domain is called domain root router. A HAWAII domain comprises several routers and BSs running the HAWAII protocol, as well as MHs. There are three types of HAWAII path setup messages: (i) power-up, (ii) update and (iii) refresh.

On power up a MH sends a Mobile IP registration request message to the corresponding BS. The BS then sends a HAWAII path setup power-up message to the domain root router, which is processed in a hop-by-hop manner. On all routers on its way to the domain root router this power-up message adds a routing entry for the concerned MH. The domain root router finally acknowledges this path setup power-up message to the BS, which finally notifies the MH with a Mobile IP registration reply.

The routing entries in the routers are soft state, i.e. they have to be refreshed periodically by path setup refresh messages, which are sent independently by each network node and which can be aggregated.

Routers, not passed by a path setup message related to a MH, do not have any knowledge about its whereabouts. Whenever a router receives a packet for such an unknown MH, e.g., from another MH within the domain, it uses a preconfigured default interface pointing towards the domain root router. This packet will be forwarded in this direction until it will arrive at a router knowing a route to the addressed host. In the worst case this will be the domain root router.

Similarly to Cellular IP, a paging mechanism is foreseen for standby MHs. Mobile hosts in standby state only have to notify the network on a change of *paging area* and not on each BS handoff. When a packet arrives for a MH in standby state, the network has to page it before it delivers the packet. This paging induces the MH to switch to active state immediately. For using HAWAII'S paging support, it is necessary to have link-layer paging functionality on the wireless link which means that the MH is able to identify its paging area and to detect paging requests.

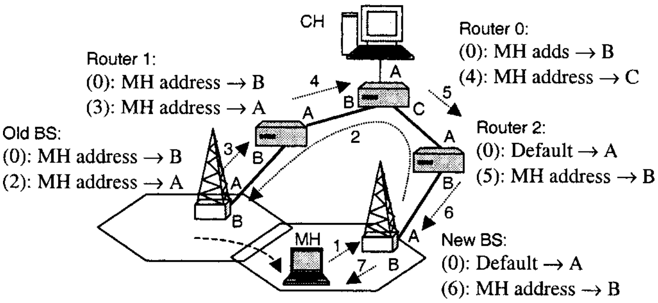

Figure 3. HAWAII MSF path setup scheme.

The network has to maintain paging information for each MH and has to deliver paging requests for these hosts up to the BSs from where on link-layer paging mechanisms are responsible. To achieve this HAWAII relies on the IP multicast routing protocol: each paging area is assigned a multicast group address all BSs within that paging area join this multicast group.

4.1. HAWAII path setup schemes

We now describe the operations of four path setup schemes used to establish path state when the MH moves from one BS to another. The four path setup schemes can be classified into two types based on the way packets are delivered to the MH during a handoff.

We define the cross-over router as the router closest to the MH that is at the intersection of two paths, one between the domain root router and the old BS, and the second between the old BS and the new BS.

4.1.1. Forwarding path setup schemes

In these path setup schemes, packets are first forwarded from the old BS to the new BS before they are diverted at the cross-over router. Two variants of forwarding schemes in HAWAII are proposed, one that works with standard IP routing tables to update the host-based entries and, another scheme where the IP routing table is extended to accommodate interface-based information. These schemes are known as Multiple Stream Forwarding (MSF) and Single Stream Forwarding (SSF). In the following the MSF scheme analyzed in this paper is described.

The MSF scheme is illustrated in figure 3. The forwarding table entries are shown adjacent to the routers. For example, *MH address* → B means that the packets with destination the *MH address* are routed through interface B. These entries are prepended with a message number indicating which message was responsible for establishing the entry (a message number of zero indicates a preexisting entry). The letters denote the different interfaces. When the MH initiates the handoff it connects to the new BS (and thus, no more packets can be received from the old BS through the air interface). Then the MH sends a path setup message (Message 1) to the old BS along the new BS (Message 2). Message 2 contains the new BS address. The old BS performs a routing table lookup for the new BS and determines the interface, interface A, and next hop router, Router 1. The old BS then adds a forwarding entry for the MH's IP address with the outgoing interface set to interface A. It then forwards Message 3 to Router 1. Router 1 performs similar action and forwards the message to Router 0. Router 0, the cross-over router in this case, changes the forwarding entry that result in new packets being diverted to the MH at the new BS. It then forwards the message towards the new BS. Eventually Message 6 reaches the new BS that changes its forwarding entry and sends an acknowledgment of the path setup message to the MH, shown as Message 7.

Note that this order of updating the routers can lead to the creation of multiple streams of disordered packets arriving at the MH. For example, during transient periods newer packets forwarded by Router 0 may arrive at the MH before older packets forwarded by Router 1 which might in turn arrive before even more older packets forwarded by the old BS. This scheme can also result in the creation of transient routing loops (for example, after old BS has changed its entry to forward packets but before the Router 1 processes Message 3). However, note that the disordered streams and routing loops exist for short periods of time. The main benefit of this scheme is that it is simple and results in no loss.

The BSs use a *forwarding buffer* for each MH in order to store the packets to be forwarded in the handoff procedure. All packets addressed to a MH are stored in the buffer (even after being transmitted to the MH). This allows that packets sent to the MH but lost because the MH moved out of coverage, will have the opportunity to reach the MH when forwarded to the new BS. Furthermore, the forwarding buffer is provided with a time out mechanism such that the buffer holds a packet only for a limited time period. When the path setup update message arrives at the old BS, all packets outstanding in the buffer for which the time out is not expired are forwarded to the new BS.

4.1.2. Non-forwarding path setup schemes

In these path setup schemes, as the path setup message travels from the new BS to the old BS, data packets are diverted at the cross-over router to the new BS, resulting in no forwarding of packets from the old BS.

There are two variants of the Non-Forwarding scheme, motivated by two types of wireless networks. The Unicast Non-Forwarding (UNF) scheme is optimized for networks where the MH is able to listen/transmit to two or more BSs simultaneously for a short duration, as in the case of a WaveLAN or Code Division Multiple Access (CDMA) network. The Multicast Non-Forwarding (MNF) scheme is optimized for networks where the MH is able to listen/transmit to only one BS as in the case of a Time Division Multiple Access (TDMA) network. In the following the UNF scheme analyzed in this paper is described.

The UNF scheme is illustrated in figure 4. In this case, when the new BS receives the path setup message, it adds a forwarding entry for the MH's IP address with the outgoing interface set to the interface on which it received this message. It then performs a routing table lookup for the old BS then forwards Message 2 to Router 2. This router performs similar actions and forwards Message 3 to Router 0. At Router 0, the cross-over router in this case, forwarding entries are added such that new packets are diverted directly to the MH at the new BS. Eventually Message 5 reaches the old BS that then changes its forwarding entry and sends an acknowledgment, Message 6, back to the MH.

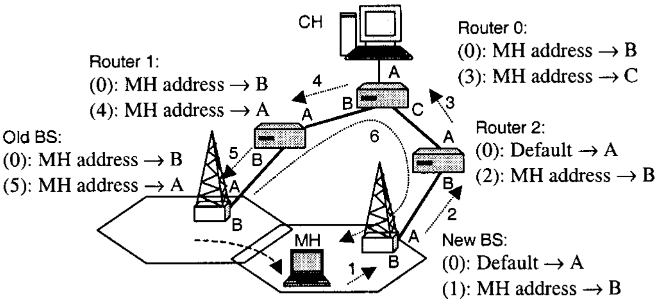

Figure 4. HAWAII UNF path setup scheme.

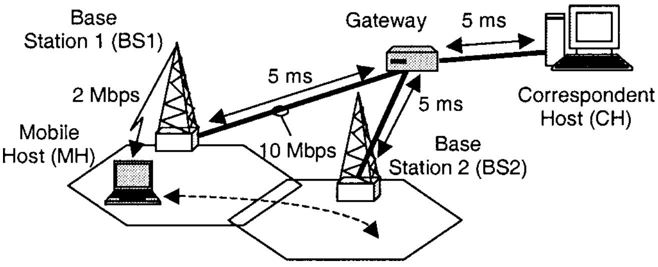

Figure 5. Simulation testbed.

5. Simulation framework

The simulations carried out in this paper have been done using the *Network Simulator* (ns) [14]. Ns is a discrete event simulator targeted at networking research. This simulator has been extensively used for evaluating TCP, MAC protocols, Satellite communications, etc.

All simulations were conducted using the topology shown in figure 5. This topology consists of two BSs (BS1 and BS2) and a Gateway router. The topology was chosen as simple as possible, in order to avoid undesirable complexity and keep the results comprehensible.

The Mobile Host moves between BS1 and BS2 such that handoffs are periodically produced. We will refer to the time between two consecutive handoffs as the *handoff period*. We assume that the MH has a speed such that the cell overlapping is 1 second. BSs send MIP Agent Advertisements every 1 second (this is the recommended maximum rate given in [11]). MHs use these Agent Advertisements as beacon signals to recognize the migration from one cell to another, and thus, initiate the handoffs.

In all the simulations a TCP download is simulated. The TCP-Reno implementation is used with a packet size of 1500 bytes and a maximum window size of 20 segments.

Several simulations have been carried out combining the following scenarios.

Radio links. Two kinds of radio links have been used (i) a shared media with an implementation of IEEE 802.11 that works like the 914 MHz Lucent WaveLAN DSSS radio interface, and (ii) an "ideal" wireless interface that consists of a fictitious non-shared media, that is, as if each sender were using a different channel at full link rate (with no collisions). We shall refer to these radio links as the *802.11 MAC* and *Ideal MAC*, respectively. In both cases the bit rate of the radio link was set to 2 Mbps.

The motivation of using the *Ideal MAC* was to eliminate the effect of a shared media access protocol as the 802.11. This allows to better observe the impact of the micromobility protocols.

Micromobility protocols. We have tested the following micromobility protocols using the ns-2 simulator implementation of [10][1]:

(i) Cellular IP with Hard Handoff and Soft Handoff.

(ii) HAWAII with MSF and UNF Path Setup Schemes: in the HAWAII-MSF a forwarding buffer with capacity for 20 packets and a time out of 400 ms was used.

(iii) Hierarchical Mobile IP with one level of hierarchy: all foreign agents in the base stations are connected to the gateway foreign agent. Therefore, direct tunnels connect the gateway foreign agent to the foreign agents located at the base stations.

6. Simulation results

All results discussed in this section have been obtained using the simulation topology described in section 5. The traffic is evaluated for down-link (data is sent from CH to MH).

The section is organized as follows: First traces obtained using the 802.11-MAC are shown in order to discuss the impact of the radio link access mechanisms. Then, the dynamics of TCP using CIP with hard and soft handoff, and HAWAII with MSF and UNF Path Setup Schemes are presented using the Ideal-MAC and the 802.11 MAC. We do not give traces of HMIP because they are very similar to the traces obtained using CIP. Finally, the goodput obtained with CIP, HAWAII and HMIP protocols is shown for comparison.

6.1. Impact of the radio link access method

Figures 6(a)–(c) show traces obtained in a down-link transmission (from CH to MH) using Cellular IP with soft handoff and a 802.11-MAC radio access link.

Figure 6(a) shows the sequence number of each transmitted segment, the instants when the MH loses the coverage (end of overlapping) with the old BS, and the evolution of the congestion window used by TCP (cwnd multiplied by 10 to see it better) at the CH. Figure 6(b) shows the queue length of TCP segments built up at the BS and the queue length of acks at the MH. Figure 6(c) is a zoom of figure 6(a). This figure shows:

(i) Instants at which the TCP segments are transmitted by the CH (indicated as "tcp Tx by CH" in the figure).

(ii) Instants at which TCP segments arrive at the MH (indicated as "tcp Rx (BS1)/(BS2)"). These instants are marked differently according to the route taken to reach the MH (through BS1 or BS2). Note that these are also the instants when the acks are generated. Transmissions instants of lost acks are marked with a square.

(iii) Instants at which acks arrive at the CH (indicated as "ack Rx (BS1)/(BS2)"). Again, these instants are marked differently according to the BS used to send them (BS1 or BS2).

(iv) The instant when the new BS generates the beacon causing the handoff and the instant when the MH generates the route update message (indicated as *MH sig lost*).

The queue of acks built up at the MH shown in figure 6(b) seems to be counterintuitive since the 10 Mbps links connecting the CH with the BS are much faster than the 2 Mbps radio link. Therefore, we would expect a queue of TCP segments at the BS, but not the queue of acks at the MH. The reason of this effect is that the acks sent by the MH tend to find the wireless medium occupied by the TCP segments, and their back-off makes that the number of acks per time unit that the MH is able to send into the shared media is lower than the number or TCP segments per time unit sent by the BS. The queue of acks is responsible of the long delay that occurs between the transmission of the ack by the MH and the reception at the CH, as shown in figure 6(c).

The beacon shown in figure 6(c) is the first one received from the new BS (BS2). This beacon causes the MH to initiate the handoff. Therefore, the MH switches the radio connection from the old BS (BS1) to the new BS and sends the route update message through the new BS to the gateway router. As indicated in the figure, this route update message and the following 11 acks sent to the new BS are lost. These losses are caused by the address resolution protocol (ARP) and are motivated by the ack queue built up at the radio link driver of the MH, as explained in the following. The IP module at the MH issues an ARP request to the new BS when the route update message is to be sent. The ARP packet is stored at the driver queue and the route update message is kept while waiting for the address resolution. Since the following acks sent by the MH are also addressed to the new BS, and the ARP module keeps only one packet waiting for the resolution of an address, each ack pushes out the previous packet waiting for the resolution of the new BS address. Only when the ARP query leaves the queue and the address is solved, the following acks are sent to the new BS. The first ack sent along the

---

[1] We introduced some modifications in the source code given in [10] in order to solve some bugs found during the simulations shown in this paper.

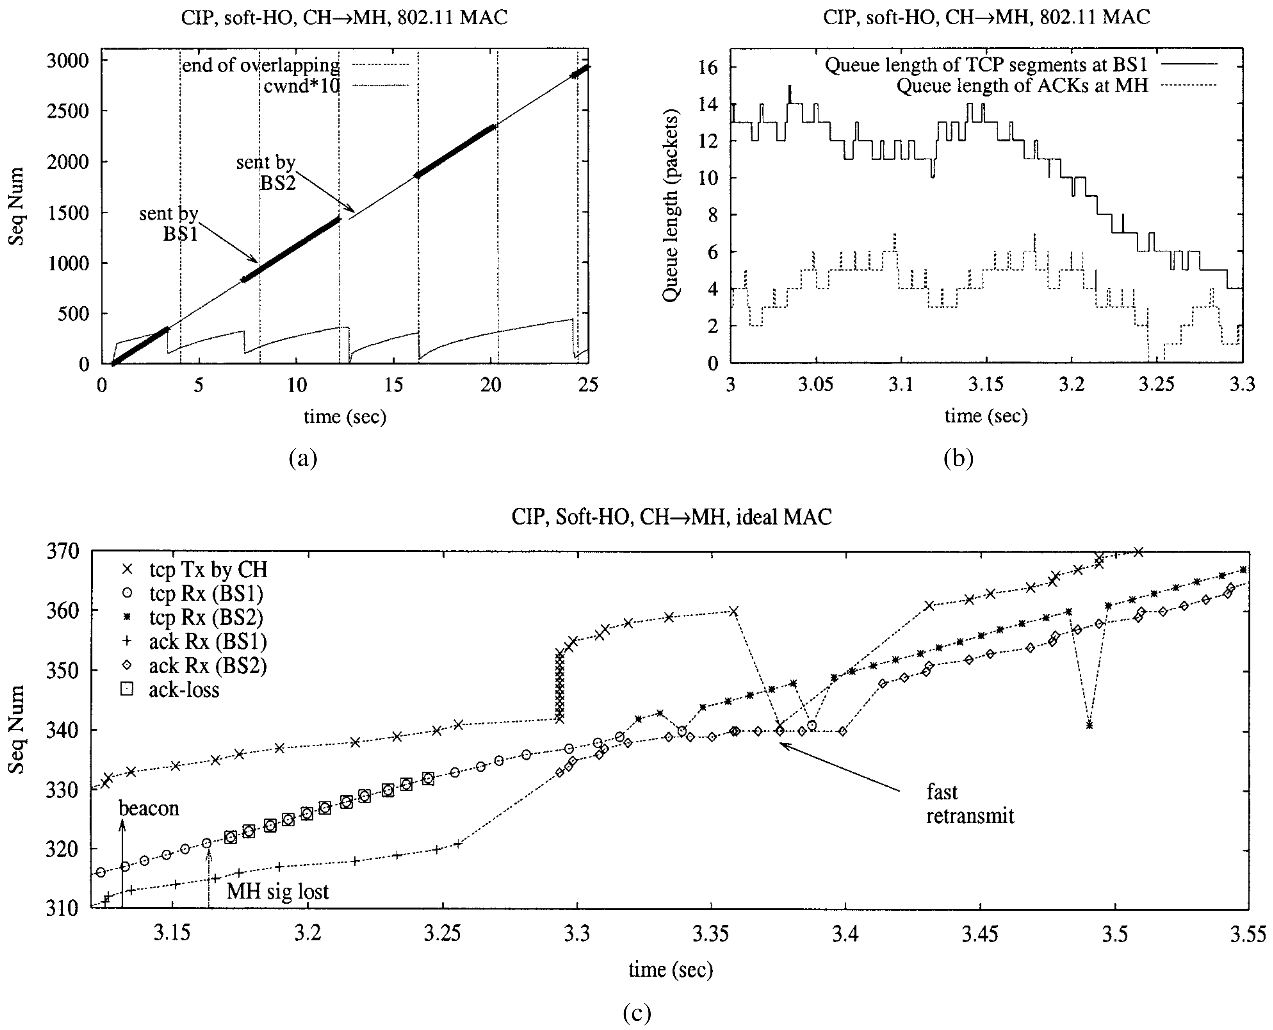

Figure 6. CIP with soft handoff traces using a 802.11-MAC. (a) TCP segments and congestion window. (b) Queue length. (c) Zoom of a handoff.

new path changes the Routing Cache to point to the new BS[2]. During this time the TCP packets are still able to reach the MH through the old BS. These packets are not lost because the MH is able to simultaneously listen to both BSs. However, in our simulations we only had these losses due to the ARP resolution in the first handoff, since the ARP table was already initialized in subsequent handoffs.

Figure 6(a) shows that although no TCP segments are lost during this first handoff, the TCP sender reduces the congestion window. This is because some of the packets arriving from the old BS reach the MH later than packets reaching the MH through the new BS (as shown in figure 6(c)). These packets arrive out of order causing the MH to send duplicated acks that trigger the fast retransmit mechanism of TCP. In this case, the fast retransmit of TCP unnecessarily retransmits a packet and reduces the congestion window.

In order to better understand the dynamics of the micro-mobility protocols and the influence of the MAC, in the following evaluations we shall use both an Ideal and a 802.11 MAC.

6.2. *CIP soft handoff*

Figure 7 shows the trace obtained with CIP using soft handoff using an Ideal MAC. When the MH receives a beacon signal from the new BS, a handoff is initiated while the connection with the old BS is maintained. The MH listens to both BSs during the overlapping of their cell coverage.

The trace shows the transmission instants of the data segments at the CH, and the transmission instants of the data segments at the old BS (BS2) and the new BS (BS1). When the new BS starts transmitting TCP segments, the old BS remains transmitting the segments that are enqueued. Since the queue at the new BS is empty, the delay of the segments that go along the path to the new BS is smaller than the ones that go along the path to the old BS. As a result of this, every time a segment transmitted by the new BS arrives at the MH, the MH sends a duplicated ack since the segment is out of order. Only when the last expected packet is transmitted by the old BS, all segments arrived out of order are acknowledged at once and the TCP source sends a whole window of new segments.

---

[2] The current version of CIP says that only route-update packets change the Routing Caches (as explained in section 2). However, the first version of CIP said that data or route-update packets could do so, as it is done in the ns implementation used in this paper.

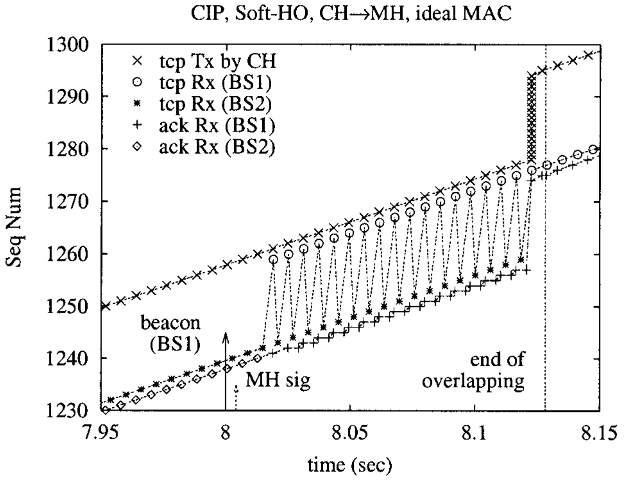
CIP, Soft-HO, CH→MH, ideal MAC

Figure 7. CIP with soft handoff (Ideal MAC).

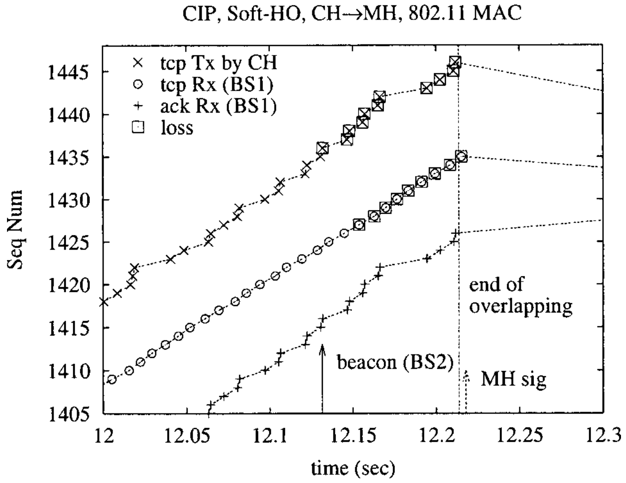
CIP, Soft-HO, CH→MH, 802.11 MAC

Figure 9. Losses using CIP with soft handoff (802.11 MAC).

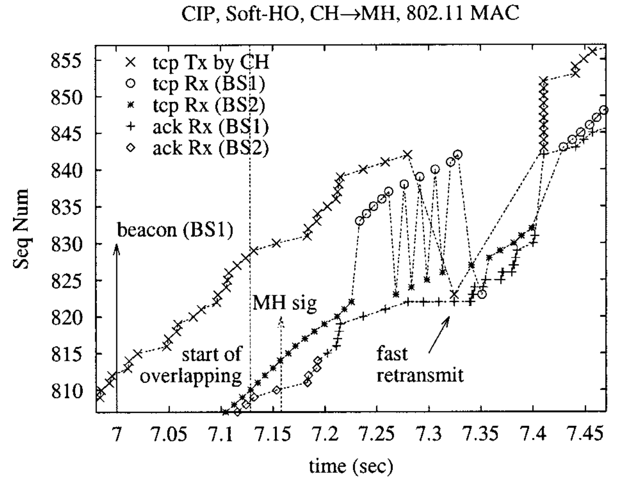
CIP, Soft-HO, CH→MH, 802.11 MAC

Figure 8. CIP with soft handoff (802.11 MAC).

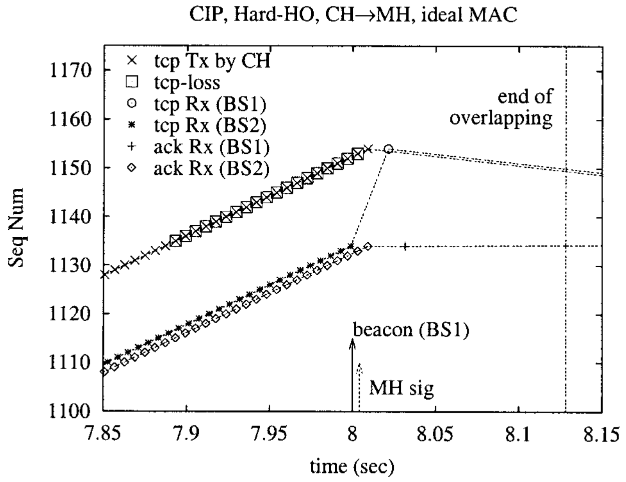
CIP, Hard-HO, CH→MH, ideal MAC

Figure 10. CIP with hard handoff (Ideal MAC).

Figure 8 shows the same handoff using the 802.11 MAC. Note that the beacon triggering the handoff in this figure (generated at 7 seconds) is the beacon generated one second before than the beacon triggering the handoff in figure 7. This is because beacons suffer a longer delay using the 802.11 MAC. The ack packets sent by MH cause such delays. Note that the BS is likely to find the shared media occupied by theses packets when it is going to send a beacon. This makes that the beacon generated at 7 seconds reaches the MH when being in the cell coverage of the new BS. The variable delay introduced by the 802.11 MAC can be observed in figure 8. This variable delay produces that several consecutive TCP segments arrive out-of-order at the MH during the handoff. As shown in the figure, these segments trigger the fast retransmit mechanism, and thus, the TCP sender unnecessarily retransmits segments and reduces the transmission window.

Finally, figure 9 shows that even using soft handoff losses can occur. Remember that the time between beacons are generated is 1 second, and the duration of the overlapping area is also 1 second. Therefore, it can occur that the MH receives the beacon from the new BS triggering the handoff after crossing the overlapping area, as shown in figure 9. In this case, after the MH sends the route update packet establishing the path along the new BS, no more TCP segments arrive at the MH because the coverage with the old BS is lost. Therefore, the transmissions stalls until the retransmission time-out is triggered and the TCP segments are transmitted through the new BS.

6.3. *CIP hard handoff*

Figure 10 shows a trace capturing a handoff obtained with CIP using hard handoff and an Ideal MAC. This figure shows that the hard handoff procedure has an important impact on the TCP dynamics. At each handoff, packets get lost when the MH switches from the old BS (BS2) to the new BS (BS1). Packets waiting at the old BS when the connection is switched cannot reach the MH and are lost. This burst of lost TCP segments causes the TCP sender to wait for the retransmission time out and start with a slow start phase. This produces a goodput degradation.

Figure 11 shows the same handoff using the 802.11 MAC. This trace shows a similar behavior than using the Ideal MAC. Now, however, 7 TCP segments arrive to the MH through the

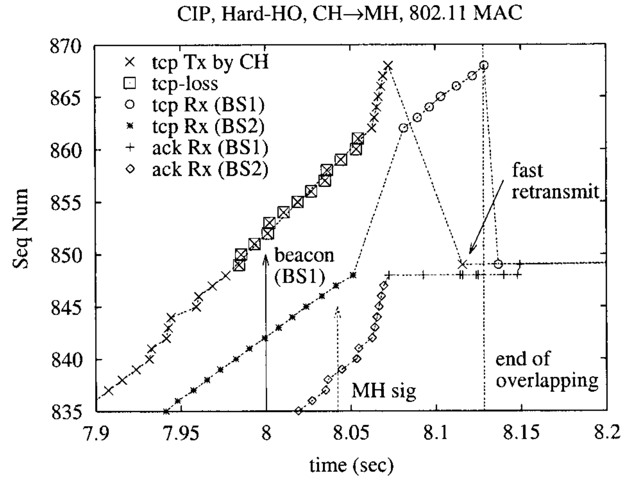

Figure 11. CIP with hard handoff (802.11 MAC).

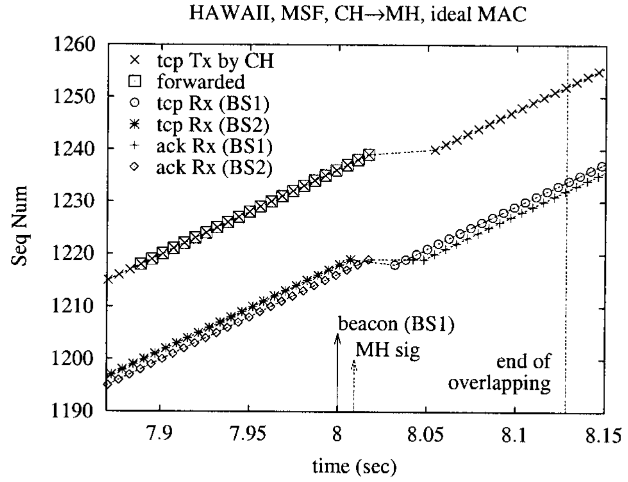

Figure 12. HAWAII with MSF (Ideal MAC).

new BS, producing a burst of duplicated acks that triggers the fast retransmit mechanism. Since multiple segments are lost in the same window, TCP-Reno is not able to recover the losses and it has to wait for the retransmission time out. Note however that these losses could be recovered using TCP-SACK and thus the impact of the handoff would be lower.

Figure 11 shows that the last 6 acks sent by BS2 have a shorter interarrival time. This is due to the fact that the MH has already disconnected from BS1. Therefore, the TCP segments dropped by BS1 are not sent to the shared media and the acks sent by the MH do not have to compete with them.

6.4. HAWAII with MSF/UNF

Figure 12 shows the traces obtained using HAWAII with the MSF path setup scheme and an Ideal MAC. In this scenario the MH is able to maintain an ongoing connection with only one BS. Remember that this protocol tries to avoid losing the packets outstanding at the old BS when the connection is switched to the new BS. This is accomplished by forwarding these outstanding packets to the new BS. The trace shows how these forwarded packets effectively avoid TCP packets to be lost, thus solving the goodput degradation that occurs in CIP with hard handoffs. Figure 13 shows the same handoff using the 802.11 MAC. This figure shows that one packet is received through the new BS before the forwarded segments. Furthermore, two of the forwarded packets were already received through the old BS. These segments generate three duplicated acks that trigger the fast retransmit mechanism.

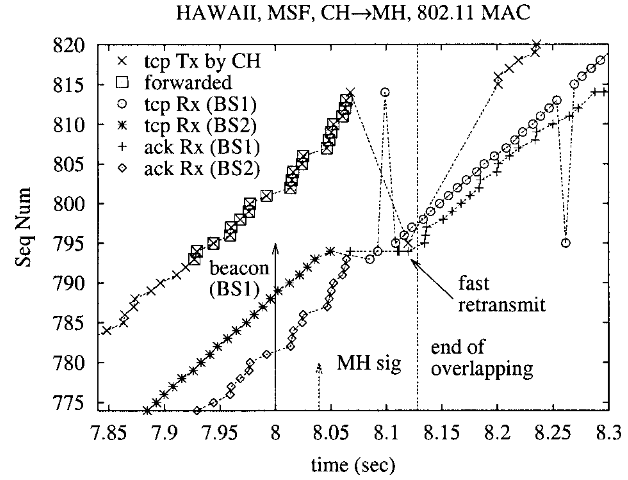

Figure 13. HAWAII with MSF (802.11 MAC).

When using HAWAII with the UNF Path Setup scheme the MH is able to listen to both the new and old BS. This scheme and CIP with soft handoffs have only small differences in the handoff procedures that are not relevant for their performance evaluation, showing similar behavior.

6.5. Goodput results

Figure 14 depicts the average goodput obtained with an Ideal-MAC and a 802.11-MAC using CIP, HAWAII and HMIP. The figures show how the handoff period length affects differently each protocol. In a real situation the handoff period will depend of many factors like MH behavior, cell size, etc. Therefore, the aim of these figures is not to reflect a real situation, but compare the protocols under study in terms of their handoff seamless. The goodput was computed as the sequence number increment of the TCP segments sent multiplied by the payload in bits $(1460 \cdot 8)$ divided by the simulation time. The measures were taken doing a simulation of 40 trips between BSs. The first four handoffs are not considered in the measures in order to avoid the impact of the losses due to the ARP resolution (as explained in section 6.1). The x-axis shows the handoff period, i.e. the time elapsed between handoffs. Figures 14(a) and (c) depict results obtained using the micromobility implementations where the MH is able to listen only to one BS (CIP and HMIP with hard handoff and HAWAII with MSF). Figures 14(b) and (d) depict results obtained using the micromobility implementations where the MH is able to listen to two BSs simultaneously (CIP and HMIP with soft handoff and HAWAII with UNF).

The graphs show that similar results are obtained using the Ideal and the 802.11 MAC. Goodputs are lower using

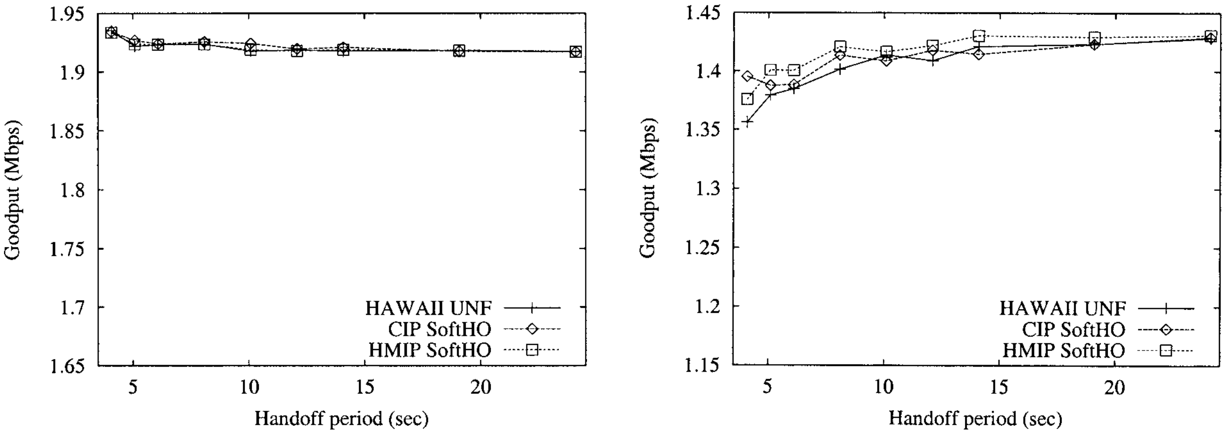

(a) Ideal MAC (b) 802.11 MAC

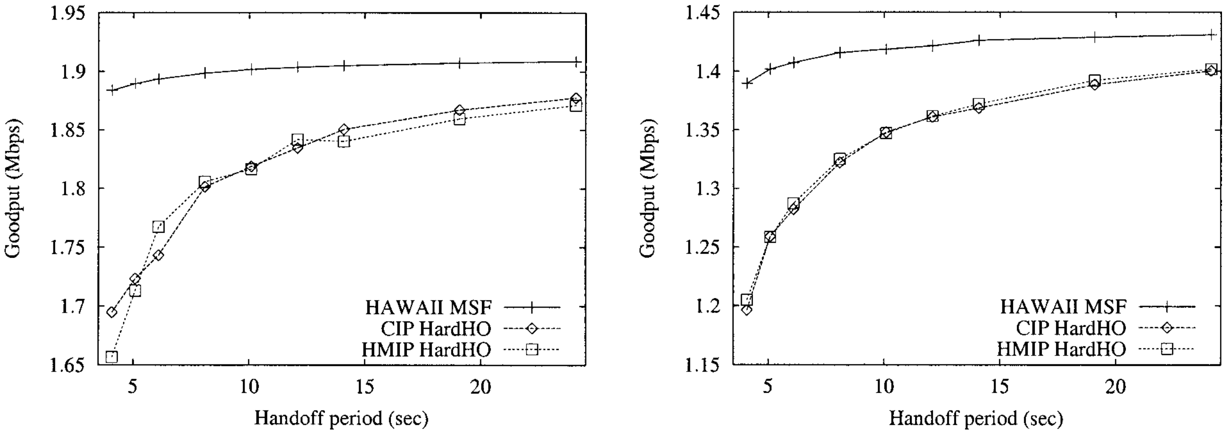

(c) Ideal MAC (d) 802.11 MAC

Figure 14. Goodput obtained with CIP and HAWAII using an Ideal-MAC ((a) and (c)) and 802.11-MAC ((b) and (d)), for down-link and up-link traffic.

the 802.11-MAC because the media is not shared as with the Ideal-MAC.

The results obtained with CIP and HMIP are very similar. The protocols differ in the routing at the IP level, but have a similar behavior when used with hard and soft handoffs. A better performance of CIP could be observed using a more complex topology than the one considered in this paper (see [3]). This is because when a handoff occurs using HMIP, the packets start only to be routed toward the new BS once the registration message reaches the Gateway (and the tunnel is switched toward the new BS). Using host routing (as CIP), routing tables are updated while the registration message travels toward the Gateway. This may produce a shorter delay to establish a path from the Gateway to the new BS.

In case of a MH being able to listen to only one BS simultaneously, graphs (c), (d) show that HAWAII with MSF achieves higher goodput than with CIP and HMIP with hard hand-offs. This demonstrates that forwarding the packets from the old to the new BS when a handoff occurs is better than simply discarding them. Note also that the goodput degradation of CIP and HMIP with hard handoff only occurs in the downlink download. The TCP segments transmitted by the MH are not lost when the connection is switched from the old to the new BS. Only the acks outstanding at the old BS will be lost, but since each ack confirms the previous ones, this loss does not decrease the goodput.

Figure 14 shows the benefits of the MH being able to simultaneously listen to two BSs. For example, if we consider the 802.11-MAC, this condition would correspond to the scenarios depicted in figures 14(a) and (b). These figures show that the goodput is higher and depends less on the handoff frequency than the scenarios depicted in figures 14(c) and (d).

Finally, note that the Ideal-MAC (figure 14(a)) shows that in the down-link scenario the goodput slightly increases when the handoff period decreases. This counterintuitive result is due to the Ideal-MAC allowing both BSs to simultaneously transmit to the MH, and thus, having double bandwidth during the handoff (while the MH is located in the overlapping region of both BS).

7. Conclusions

In this paper we have analyzed the dynamics of TCP in a cellular network using an IP micromobility protocol. We have considered three micromobility implementations where the MH is able to listen only to one BS: CIP and HMIP with hard handoff and HAWAII with MSF, and three micromobility implementations where the MH is able to listen to two BSs simultaneously: CIP and HMIP with soft handoff and HAWAII with UNF.

Detailed traces are given for each of the micromobility protocols using a TCP download. By means of these traces handoff procedures that occur in each scenario are analyzed. Furthermore, the impact of a 802.11 MAC protocol is studied.

The numerical results show that:

- The 802.11 MAC protocol may have an important impact on TCP goodput due to the various interactions of theradio link layer with the handoff schemes:
  - The acks sent by the MH tend to find the wireless medium occupied by TCP segments and a queue of acks is built up which delays a lot the reception of the acks,
  - The ack queue built up at the radio link driver causes the address resolution protocol to drop acks at the driver queue.
- Handoffs have a low impact in case of a MH being able to simultaneously listen to the old and new BS.
- In case of a MH being able to listen to only one BS, HAWAII MSF is superior to CIP and HMIP with hard handoff. In fact, goodput of CIP and HMIP with hard handoff can be considerably reduced if the handoff frequency is high.

Acronyms

Ack	Acknowledgement
ARP	Address Resolution Protocol
BS	Base Station
CDMA	Code Division Multiple Access
CH	Correspondent Host
CIP	Cellular IP
GW	Gateway Router
HAWAII	Handoff Aware Wireless Access Internet Infrastructure
HMIP	Hierarchical Mobile IP
IETF	Internet Engineering Task Force
IP	Internet Protocol
MAC	Medium Access Control
MH	Mobile Host
MIP	Mobile IP
MNF	Multicast Non-Forwarding
MSF	Multiple Stream Forwarding
SSF	Single Stream Forwarding
TCP	Transmission Control Protocol
TDMA	Time Division Multiple Access
UNF	Unicast Non-Forwarding

Acknowledgements

This work was partly supported by the Ministry of Education of Spain under grant TIC-2001-0956-C04-01 and by the European project MOEBIUS.

References

[1] A.T. Campbell, J. Gomez, C.-Y. Wan, S. Kim, Z. Turanyi and A. Valko, Cellular IP, Internet draft, `draft-ietf-mobileip-cellular-00-txt` (December 1999).

[2] A.T. Campbell, J. Gomez, C.-Y. Wan, S. Kim, Z. Turanyi and A. Valko, Cellular IP performance, Internet draft, `draft-gomez-cellularip-perf-00.txt` (October 1999).

[3] A.T. Campbell, J. Gomez, S. Kim, Z. Turanyi, C.-Y. Wan and A. Valko, Comparison of IP Micromobility Protocols, IEEE Wireless Communications Magazine 9(1) (2002).

[4] Cellular IP homepage, Columbia University, `http://comet.ctr.columbia.edu/cellularip`

[5] K. El Malki and H. Soliman, Fast hand-offs in mobile IPv4, Internet draft, `draft-elmalki-mobileip-fast-handoffs-03.txt` (September 2000).

[6] B. Gloss and Ch. Hauser, The IP micro mobility approach, in: *Proc. of the EUNICE 2000* (September 2000).

[7] E. Gustafsson, A. Jonsson and C. Perkins, Mobile IP regional registration, Internet draft, `draft-ietf-mobileippreg-tunnel-02.txt` (March 2000).

[8] H. Haverinen and J. Malinen, Mobile IP regional paging, Internet draft, `draft-haverinen-mobileip-reg-paging-00.txt` (June 2000).

[9] J. Kempf, P.R. Calhoun and C. Pairla, Foreign agent assisted handoff, Internet draft, `draft-calhoun-mobileip-proactive-fa-01.txt` (June 2000).

[10] ns with micromobility support, `http://comet.ctr.columbia.edu/micromobility`

[11] C.E. Perkins, IP mobility support, RFC 2002 (October 1996).

[12] R. Ramjee, homepage, Bell Laboratories, `http://www.bell-labs.com/user/ramjee/`

[13] R. Ramjee, T. La Porta, S. Thuel, K. Varadhan and L. Salgarelli, IP micro-mobility support using HAWAII, Internet draft, `draft-ietf-mobileip-hawaii-00.txt` (June 1999).

[14] The Network Simulator – ns-2, `http://www.isi.edu/nsnam/ns`

Llorenç Cerdà was graduated in telecommunications engineering from Polytechnic University of Catalonia (UPC) in 1993. He joined UPC in 1994, presented the Ph.D. Thesis entitled *Traffic Management of the ABR Service Category in ATM Networks* in January 2000 and currently he is Assistant Professor at the Computer Architecture Department. His areas of interest include TCP, IP micromobility, C++ programming, simulation, performance evaluation, queuing networks.
E-mail: llorenc@ac.upc.es

Fabio Vena was graduated at the Swiss Federal Institute of Technology, Zurich in 2001. He realized his master thesis at the Polytechnic University of Catalonia, Computer Architecture Department.
E-mail: fveria@ac.upc.es

Olga Casals obtained the engineering degree in telecommunication in 1983 from the Polytechnic University of Catalonia, and received the Ph.D. degree from the same University in 1986. Since the end of 1983 she has been working at the Computer Architecture Department of the Polytechnic University of Catalonia where she is head of a research group for traffic research in B-ISDN communication systems. In 1994 she became Full Professor. From May 1998 till March 2001 she was Head of the Department of

Computer Architecture. Her areas of interest include computer and communication system performance evaluation and B-ISDN traffic control mechanisms. Olga Casals is a member of IFIP Working Group 6.3 on Performance of Communication Systems. She has published a substantial number of papers in international journals and conferences on performance evaluation of computer and communication systems and on traffic control in B-ISDN networks. She has participated in many projects (RACE, ACTS, IST, COST) funded by the European Commission.
E-mail: olga@ac.upc.es