Presented at European Wireless 2002, Florence, Italy, February 2002.

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

Fabio Vena, Llorenç Cerdà and Olga Casals

Polytechnic University of Catalonia - Computer Architecture Dept.
c/ Jordi Girona, 1-3, Modulo C6, E-08034 Barcelona, Spain

Ph: + 34 934016798, fax: + 34 934017055, email: `(llorenc,fvena,olga)@ac.upc.es`

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 two of these protocols namely Cellular IP and Hawaii affect the behavior of TCP and their interaction with the MAC layer.

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) [6] is the current standard for supporting mobility in an IP network. MIP is an appropriate solution to handle global IP mobility (macro-mobility) 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. Host based routing schemes such as Cellular IP (CIP) [3], HAWAII [7], etc have been proposed to provide micro-mobility management with seamless handoff, minimum packet loss and limited handoff latency. In this paper we study the behavior of TCP sources in a mobile environment using Cellular IP and HAWAII with different path setup schemes. We first provide a description of each of these protocols. 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][3][4][5] 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 forward packets. These information elements are refereed 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.

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.

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 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

Handoff in Cellular IP is always initiated by the MH. As the host approaches a new BS, it redirects its data packets from the old to the new BS. The first of these redirected packets will automatically configure a new path of Routing Cache mappings for the host, this time 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. HAWAII

HAWAII was proposed to the IETF in [2][5][7] by researchers from Lucent Bell Labs. Like in Cellular IP, HAWAII is responsible for the micro-mobility support while the macro-mobility 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, don't 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 pre-configured 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 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.

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 and all BSs within that paging area join this multicast group.

3.1. 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.

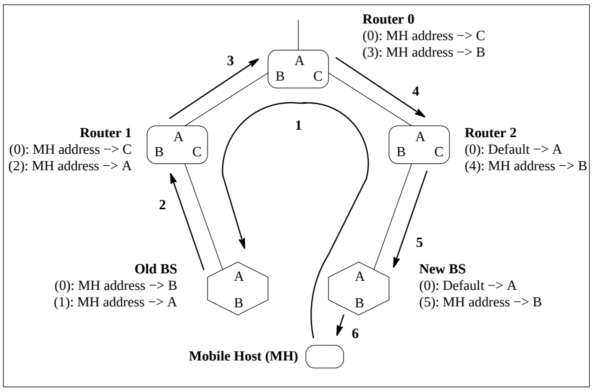

Figure 1: Hawaii MSF Path Setup Scheme.

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 1. The forwarding table entries are shown adjacent to the routers. These entries are prepended with a message number indicating which message was responsible for establishing the entry (a message number of zero indicates a pre-existing 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 1 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, Router1. 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 2 to Router 1. Router 1 performs similar action and forwards the message to Router 0. Router 0, the cross-over router in this case, adds forwarding entries 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 5 reaches the new BS that changes its forwarding entry and sends an acknowledgment of the path setup message to the MH, shown as Message 6.

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 2). 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.

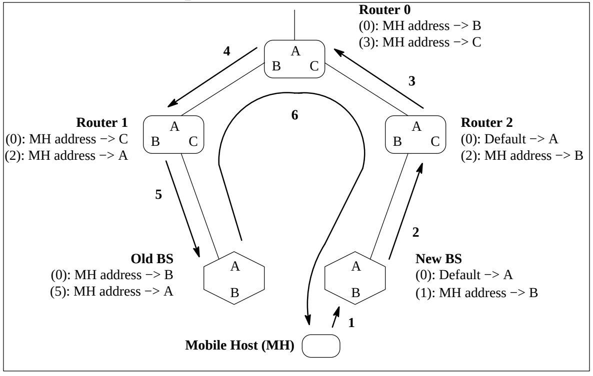

Figure 2: Hawaii UNF Path Setup Scheme.

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 2. 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.

4. SIMULATION FRAMEWORK

All simulations were conducted using the topology shown in Figure 3. This topology consists of two BSs (BS1 and BS2) and a cross-over router. The cross-over router is the Gateway for Cellular IP and the Domain Root Router for HAWAII. The topology was chosen as simple as possible, in order to avoid undesirable complexity and keep the results comprehensible.

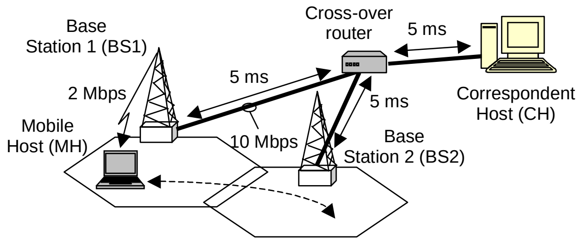

Figure 3: Simulation testbed.

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*. There is a cell overlapping of 1 second. BSs send beacon signals every 1 second. The MH uses the 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 1460 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.

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 micro-mobility protocols. In both cases the bit rate of the radio link was set to 2 Mbps.

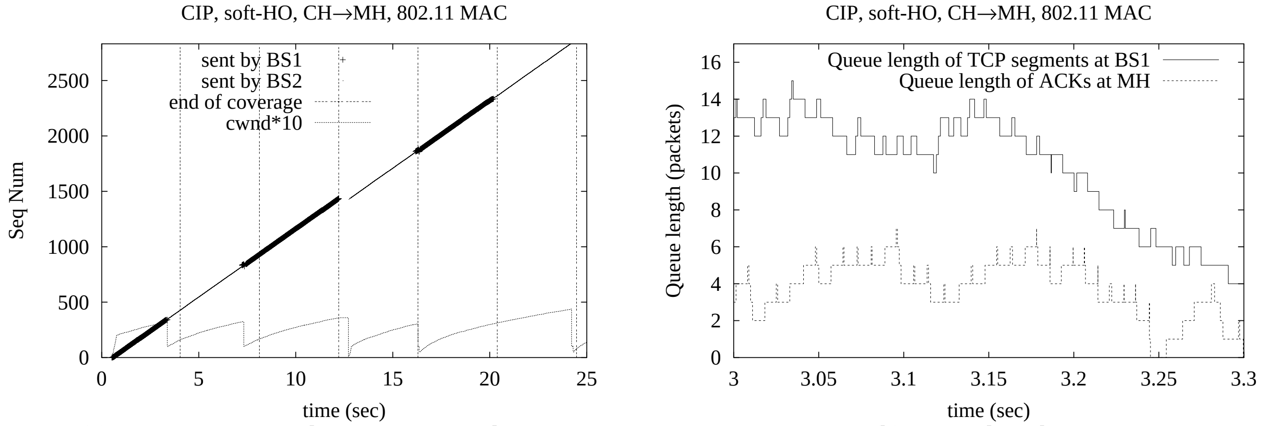

(a) TCP segments and congestion window.

(b) Queue length.

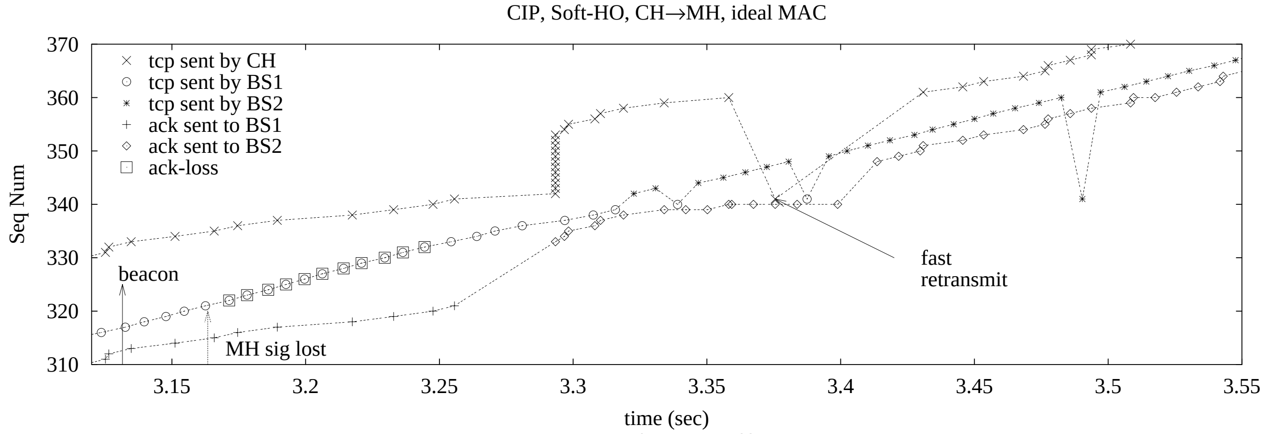

(c) Zoom of a handoff.

Figure 4: CIP with soft handoff traces using a 802.11-MAC.

Micro-mobility protocols: we have tested the following micro-mobility protocols using the ns-2 simulator implementation of [1]: (i) Cellular IP with Hard Handoff and Soft Handoff and (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.

5. NUMERICAL RESULTS

All results discussed in this section have been obtained using the simulation topology described in section 3. The traffic is evaluated both for down-link (data is sent from CH to MH) and up-link (data is sent from MH to CH).

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 each of the micro-mobility protocols (CIP with hard and soft handoff, and HAWAII with MSF and UNF Path Setup Schemes) are presented using the Ideal-MAC. Finally, the goodput obtained with both protocols is shown for comparison.

5.1. Impact of the radio link access method

Figure 4 shows 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.

Trace (a) shows the sequence number of each transmitted segment, the instants when the MH loses the coverage (end of coverage) 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. Trace (b) shows the queue length of TCP segments built up at the BS and the queue length of acks at the MH. Trace (c) is a zoom of trace (a). This trace shows: (i) Instants at which the TCP segments are transmitted by the CH (indicated as “tcp sent by CH” in the figure). (ii) Instants at which TCP segments arrive at the MH (indicated as “tcp sent by 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 sent to BS1/BS2”). Again, these instants are marked differently according to the BS used to send them (BS1 or BS2). (iv) The transmission instant of the beacon from the new BS causing the handoff and the transmission instant of the route update message sent by the MH (indicated as *MH sig lost*).

The queue of acks built up at the MH shown in Figure 4.(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

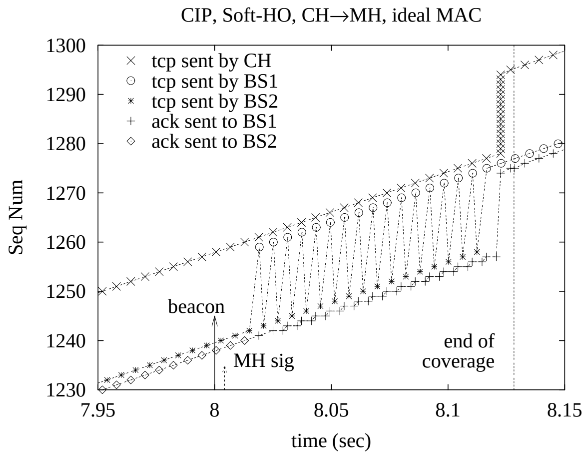

Figure 5: CIP with soft handoff.

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 CN, as shown in trace (c).

The beacon shown in Figure 4.(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 loses are caused by an address resolution failure motivated by the ack queue built up at the radio link driver of the MH. 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 push 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 reaching the cross over router changes the routing cache to point to the new BS. 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.

Figure 4.(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 trace (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.

To avoid the described problems caused by the 802.11 MAC influencing the performance of the Micro-Mobility protocols, in the following evaluations we shall use an Ideal MAC.

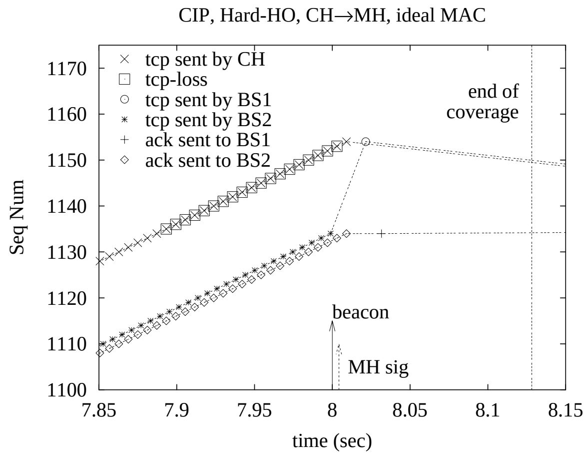

Figure 6: CIP with hard handoff

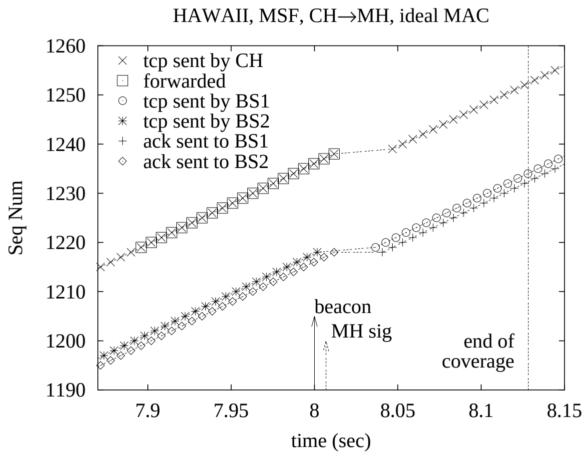

Figure 7: Hawaii with MSF.

5.2. CIP Soft Handoff

Figure 5 shows the trace obtained with CIP using Soft Handoff. 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 start 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.

5.3. CIP Hard Handoff

Figure 6 shows a trace capturing a handoff obtained with CIP using Hard Handoff. The points depicted in this figure are analogous to those of Figure 5 but now some TCP segments are lost as shown.

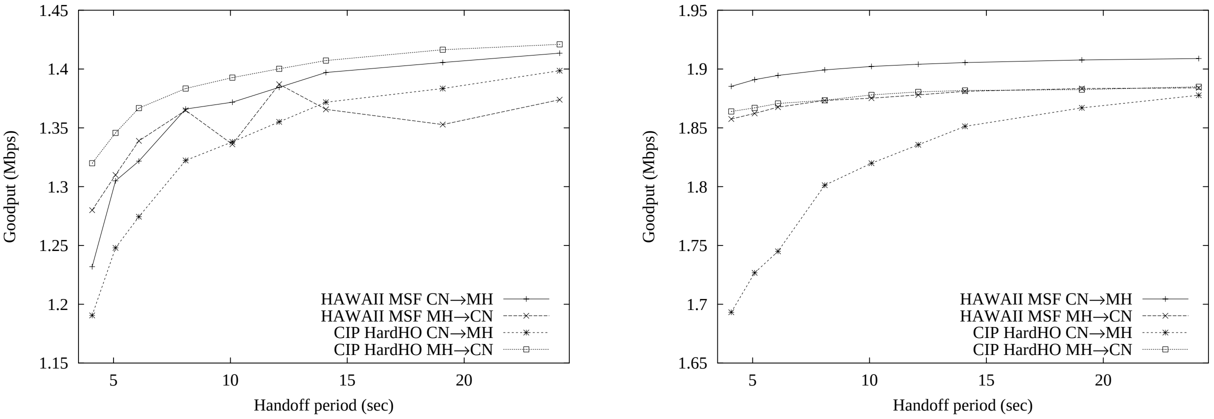

(a) 802.11 MAC. (b) Ideal MAC.

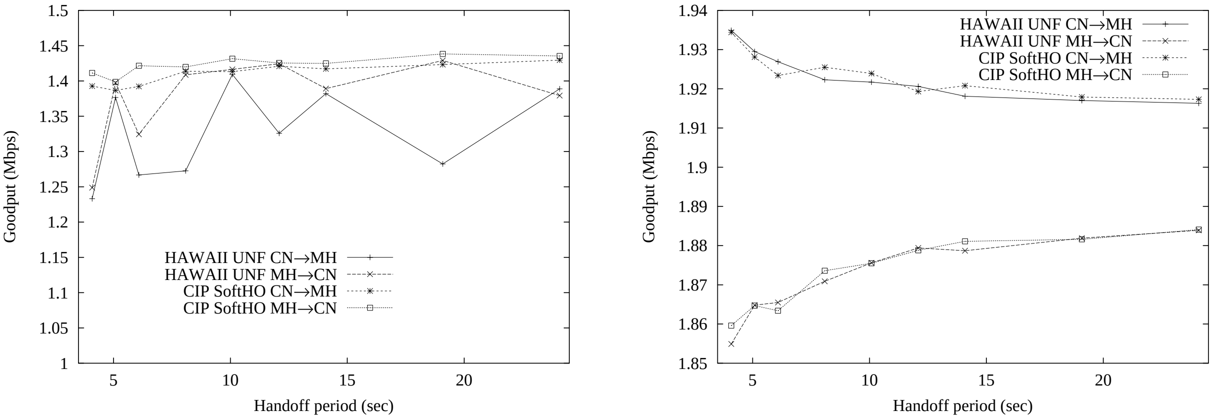

(c) 802.11 MAC. (d) Ideal MAC.

Figure 8: Goodput obtained with CIP and HAWAII using a 802.11-MAC (a), (c) and Ideal-MAC (b), (d) MAC, for down-link and up-link traffic.

The hard handoff procedure has an important impact on the TCP dynamics. At each handoff, packets get lost when the MH switches the connection 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.

5.4. Hawaii with MSF/UNF

Figure 7 shows the traces obtained using Hawaii with the MSF Path Setup Scheme. 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.

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.

5.5. Goodput results

Figure 8 depicts the average goodput obtained with a 802.11-MAC and an Ideal-MAC using the four micro-mobility scenarios described in the previous sections. The goodput was computed as the sequence number increment of the TCP segments sent multiplied by the payload in bits divided by the simulation time. The measures were taken doing a simulation of 40 trips between BSs. The x-axis shows the handoff period, i.e. the time elapsed between handoffs. Graphs (a) and (b) depict results obtained using the micro-mobility implementations where the MH is able to listen only to one BS (CIP with hard handoff and HAWAII with MSF). Graphs (c) and (d) depict results obtained using the micro-mobility implementations where the MH is able to listen to two BSs simultaneously (CIP with soft handoff and HAWAII with UNF).

Goodputs are lower when using the 802.11-MAC. The media is not shared with the Ideal-MAC and the peculiar behavior observed using the 802.11-MAC (as the ARP failure explained in section 5.1) does not occur.

In case of a MH being able to listen to only one BS simultaneously, graph (b) shows that Hawaii with MSF achieves higher goodput than with CIP with hard handoffs. This demonstrates that forwarding the packets from the old to the new BS when a handoff occurs is better than simply discarding them. This advantage is not reflected in graph (a) since the impact of the 802.11-MAC is higher that the impact caused by the handoff procedure. Note also that the goodput degradation of CIP 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 8 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 the graph (c). The graph shows that the goodput is higher and depends less on the handoff frequency than the scenarios depicted in graph (a).

Finally, note that the Ideal-MAC (graph (d)) shows that in the up-link scenario the higher is the handoff period the lower is the goodput. This counterintuitive result is due to the Ideal-MAC allowing both BS 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).

6. CONCLUSIONS

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

Detailed traces are given for each of the micro-mobility 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:

- 802.11 MAC protocol may have an important impact on goodput.
- 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 with hard handoff. In fact, goodput of CIP with hard handoff can be considerably reduced if the handoff frequency is high.

ACKNOWLEDGEMENTS

This work was partly supported by the Ministry of Education of Spain under grant TIC-1998-1115-C02-01 and by the European project MOEBIUS.

REFERENCES

[1] Cellular IP Homepage, Columbia University. http://comet.ctr.columbia.edu/cellularip

[2] R. Ramjee’s Homepage, Bell Laboratories. http://www.bell-labs.com/user/ramjee/

[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, C-Y. Wan, S. Kim, Z. Turanyi, and A. Valko. “Cellular IP Performance”. draft-gomez-cellularip-perf-00.txt, October 1999.

[5] B. Gloss and Ch. Hauser. “The IP Micro Mobility Approach”. In *Proc. of the EUNICE 2000*, September 2000.

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

[7] 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.