

Adaptive Semi-Soft Handoff for Cellular IP Networks

Eriko Nurvitadhi[‡], Ben Lee[‡], Chansu Yu[†], and Myungchul Kim^{*}

[‡]*School of Electrical Engineering & Computer Science
Oregon State University
Corvallis, OR 97331
{nurviter, benl}@eecs.orst.edu*

[†]*Department of Electrical & Computer Engineering
Cleveland State University
Cleveland, OH 44115
c.yu91@csuohio.edu*

^{*}*School of Engineering
Information & Communications University
Daejeon, Korea
mckim@icu.ac.kr*

Abstract

Rapid advances in wireless networking have led to more mobile phones, PDAs, and other digital mobile devices becoming ubiquitously connected to the Internet. As the popularity of multimedia applications for these portable devices increases, providing seamless connectivity to wireless networks becomes a critical issue. For this reason, a number of micro-mobility protocols, such as Cellular IP, have been proposed to complement the Mobile IP protocol. However, providing fast and reliable handoff is still a major obstacle to enabling seamless micro-mobility in wireless access networks. Cellular IP semi-soft handoff has been proposed to address such challenge. Evaluations have been performed which show that semi-soft handoff yields better performance than the conventional hard handoff. However, these studies are based on symmetrical network topologies and loads. In practice, network topology varies and the network load fluctuates depending on numerous parameters (e.g. number of mobile nodes, amount of traffic in the network, etc.). The semi-soft handoff uses fixed delay device and semi-soft delay values for stream synchronization and mobile host's tune-in timing. Such scheme may work well for the evaluated symmetrical setup. However, this will not be the case with unbalanced and dynamically changing networks, as what are typically found in real life. This paper describes a novel adaptive protocol (Adaptive-SS), which is proposed as an extension to the current Cellular IP semi-soft handoff protocol to address such issue by assigning delay device and semi-soft delay values dynamically based on the present network condition. The simulation results show that Adaptive-SS significantly reduces network traffic and packet losses and duplications during handoff, while still minimizing handoff latency.

Key words: *Fast handoff, micro-mobility protocols, Cellular IP, multimedia streaming.*

1. Introduction

Rapid advances in wireless networking have led to the prolific availability of commercial mobile devices and services. More mobile phones, PDAs, and other digital mobile devices have become ubiquitously connected to the Internet. As the popularity of multimedia applications for these portable devices increases, providing seamless connectivity to wireless networks becomes a critical issue. Mobile IP [1] has been proposed and used as the protocol for facilitating global mobile Internet capability. Mobile IP handoff requires the *mobile host* (MH) to register with its home agent by providing the care-of-address of the new foreign agent. This enables the home agent to keep track of the exact location of the MH. This mechanism would work well at a macro-level, where the coverage area and the cells are relatively large. However, at the micro-level, the cells are smaller to support more users, and thus lead to more frequently handoffs. For such a local mobility management, conventional Mobile IP handoff would suffer from increased delay, packet loss, and signaling, caused by registration, if the home agent of the MH is at a distant location. Such overhead would invariably disrupt the smooth delivery of multimedia content. A number of micro-mobility protocols have been proposed to mitigate this problem [2, 3, 4]. One of the main improvements introduced by these protocols is the fast handoff mechanism. Handoff mechanisms in micro-mobility protocols eliminate the need for registration with the home agent as long as the MH is still within the same domain.

Cellular IP is one of the proposed micro-mobility protocols [5, 6, 7, 8], which handles two types of handoffs—hard and semi-soft handoff. In hard handoff, a MH performs a handoff by tuning in its radio to the new *base station* (BS) and sending a *route-update packet* to establish a new path. Hard handoff is simple and minimizes network traffic during handoff, but induces certain amount of delay and packet loss during switching from the old BS to the new BS. In semi-soft handoff, a MH tunes into the new BS while still preserving its connection to the old BS. A route-update packet is sent to the new BS to create a path to it. Then, the MH tunes into the new BS and the handoff process is completed. By maintaining connections to both old and new BSs, a more seamless handoff can be achieved.

Cellular IP semi-soft handoff mechanism uses a fixed *delay device* to synchronize the delays of the old and new paths. This may work well with relatively small path delay differences. However, for path delays that fluctuate, as in real networks, using a fixed value for the delay device is not sufficient. An inaccurate delay value would lead to packet losses or duplications. Therefore, the delay device value needs to be determined dynamically depending on the network condition at the time. In addition, a fixed value is also used for the *semi-soft delay*. This parameter determines how long the MH has to wait before tuning into the new BS. An inappropriate semi-soft delay value will also lead to packet losses or duplications.

Existing evaluations of semi-soft handoff, however, were conducted using a simple and symmetrical network configuration. By simple we mean only one MH is involved in the evaluation and only the network traffic generated by the MH is considered during the handoff. Symmetrical network means that the old and new paths involved in the handoff have the same number of network hops. The impact of such setup is that both old and new paths experience similar network delays. In practice, however, network loads vary dynamically, and the topology may differ depending on the location. Such conditions would lead to fluctuations in network delays between the old and new paths involved in the handoff.

This paper proposes a novel *adaptive semi-soft handoff* protocol (Adaptive-SS) that adapts to dynamic network conditions by assigning delay device and semi-soft delay values based on the network condition at the time of the handoff. Furthermore, Adaptive-SS can also be applied in a network where MH cannot communicate to both BSs simultaneously during handoff, and thus could be used as an alternative to indirect-handoff [17]. The proposed approach was implemented and tested on ns-2 extended with the Cellular IP model by the Columbia IP Micro-mobility group [11]. The simulation results show that Adaptive-SS significantly improves handoff performance in terms of network traffic and packet losses/duplications during handoff.

The rest of the paper is organized as follows. Section 2 overviews the issues with hard and semi-soft handoff. Section 3 presents the proposed adaptive semi-soft handoff approach. Section 4 presents the evaluation methodology and results. Section 5 discusses related work on fast handoff mechanisms and Cellular IP extensions. Finally, Section 6 concludes the paper and discusses future work.




Figure 1: A simple handoff scenario.

2. Background

Figure 1 shows a simple handoff scenario where a MH is communicating with a *corresponding host* (CH), which could be anywhere on the Internet. The Cellular IP *gateway* (GW) is the gateway of the Cellular IP network with two cells, cell1 and cell2. The *base stations* (BS1 and BS2) are the access points for cell1 and cell2, respectively. The MH moves from cell1 to cell2 and performs a handoff. CN is a *cross-over node* where the old path and the new path intersect. Thus, CN is responsible for rerouting the packets for MH to the new BS (BS2).

A Cellular IP node (GW, BS, or any other intermediate nodes) maintains a routing cache containing soft-state mappings. Each mapping consists of the address of a MH and the address of the next node that has to be followed to reach the MH. The mappings are created and updated by uplink or route-update packets sent from the MH. When a node receives a route-update packet, it records in the routing cache the IP address of the source MH and the neighboring node from which the packet came from. A node will look up for a valid mapping from its routing cache when it receives a network packet intended for a particular MH. If a valid mapping is found, the packet is forwarded using the next node address entry found in the mapping. An entry is deleted after a certain timeout value if it is not refreshed by either an uplink or a route-update packet sent from the MH.

A MH listens to beacon signals sent by the BSs as it moves, and initiates a handoff based on their signal strength measurement. The MH establishes a path by sending a route-update packet to the BS. The route-update packet travels through the BS to the GW (via all the nodes in between the BS and the GW). Each node visited by the route-update packet updates or creates a soft-state mapping in its routing cache for the MH. Thus, a new path to the new BS is established when the CN creates a soft-state mapping in its routing cache. Following this, any packets arriving at the GW for the MH will be sent to the MH based on the mappings on each of these nodes.

The MH performs a hard handoff by tuning in its radio to the new BS (i.e., BS2) right after its signal strength is detected to be stronger than the old base BS's (i.e., BS1). MH then sends a route-update packet to BS2 to establish the new path. The pitfall of this approach is that packet losses and handoff delay (i.e., the time period between when the last packet is received from the old BS and when the first packet is received from the new BS) can occur during the time MH waits for the establishment of the new path. If a packet destined for MH arrives at CN prior to the arrival of the route-update-packet at CN, CN will forward the packet through the old path. Since MH has already tuned into the BS2 by this time, the packet will be lost as it is sent via BS1. The handoff delay occurs because the MH has to wait for the route-update packet to reach and update the routing cache at CN before it can receive any packets from BS2.

Semi-soft handoff reduces handoff delay and packet losses. In semi-soft handoff, the MH first establishes the new path by sending a route-update packet to BS2 upon signal strength detection. Then, it tunes back to BS1 and waits for the path establishment for a certain semi-soft delay. After such a delay, it then tunes back to BS2 (presumably the new path is established by this time). Since the new path is already established by the time MH tunes into the new BS, the MH can start receiving packets right away. Thus, delay and packet losses can be minimized and more seamless handoff can be achieved. However, unlike the hard handoff, the semi-soft handoff does not minimize signaling. During the path establishment, semi-soft handoff requires CN to forward (bi-cast) packets destined

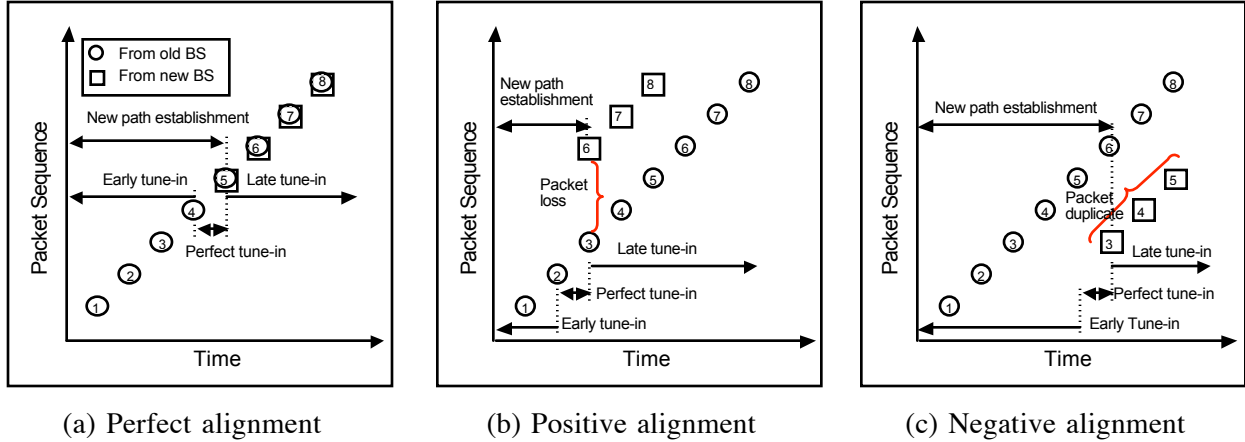


Figure 2: Stream alignment and tune-in issues.

for MH via both old and new BSs. Thus, there will be more traffic during semi-soft handoff than hard handoff. Studies have shown that semi-soft handoff provides improved UDP and TCP performance over hard handoff in terms of packet losses [6, 7, 8, 9].

There are two critical requirements for establishing seamless handoff with semi-soft handoff: (1) *Perfect alignment of streams* and (2) *efficient tune-in time*. Figure 2 illustrates the alignment and tune-in time issues. Stream alignment refers to the skewness of the packet sequences arriving at MH from the old and new BSs during handoff. Figure 2(a) shows a perfect alignment where the last packet received from the old BS and the first packet received from the new BS is in sequence. Misaligned streams result in packet losses or duplications. If the new stream (i.e., packets received from the new BS) is ahead the old stream, the first packet received by the MH upon tuning into the new BS would be a few packet sequences ahead of the last packet received from the old BS, and thus some packets are lost. We refer to this as *positive misalignment* (Figure 2(b)). A *negative misalignment* occurs when the old stream is ahead of the new stream and packets are duplicated (Figure 2(c)).

To have a perfect alignment of streams, the path delays between the MH and the CN via the old BS and the new BS must be the same. However, streams can be misaligned due to fluctuations in various delay factors (e.g., node delay, propagation delay, and MAC contention) that causes erratic path delays between MH and its CN. Positive misalignment occurs when the new path delay is smaller than the old path delay. Negative misalignment occurs when the new path delay is longer than the old path delay.

Semi-soft handoff mitigates the positive misalignment problem by having the CN induce a delay, using its delay device, to packets forwarded to the new path to increase its delay to the same level as the old path delay. Thus, the delay device value (t_{dd}) directly impacts stream alignment. An appropriate t_{dd} has to be selected so that both path delays are equal and thus a perfect stream alignment can be achieved. Semi-soft handoff does not provide any mechanism for handling negative stream misalignment.

Tune-in time refers to the time the MH tunes into the new BS. An efficient tune-in time is the one that minimizes handoff delay, packet losses/duplications, and handoff traffic. As can be seen from Figure 2, tuning-in too early may cause packet losses and increased handoff delay. On the other hand, late tune-in induces more network traffic than needed. In Cellular IP semi-soft handoff, semi-soft delay (t_{ss}) is used as a guide to when MH tunes into the new BS. For example, after sending a route-update packet to the new BS to establish the new path and tuning back into the old BS, the MH waits for t_{ss} before it tunes into the new BS. Thus, the t_{ss} value impacts the tune-in time directly.

The Cellular IP semi-soft handoff uses fixed values for both t_{dd} and t_{ss} , which is not sufficient for handling dynamic network behavior and unsymmetrical network topologies.

3. Adaptive Semi-soft Handoff Mechanism

The proposed *Adaptive Semi-Soft* (Adaptive-SS) handoff scheme is a simple, yet effective method to dynamically determine the delay device (t_{dd}) and semi-soft delay values (t_{ss}) and thus improve handoff performance. Adaptive-SS determines t_{dd} and t_{ss} based on the network condition at the time of the

Adaptive-SS mechanism:

1. MH sends semi-soft route-update packets, p_{old} and p_{new} to old and new BSs, and tunes back to old BS.
2. CN records arrival times, t_{old} and t_{new} :
 - (a) If $t_{old} > t_{new}$, then set $t_{dd} = t_{old} - t_{new}$
 - (b) Else, if $t_{ss-old} \leq t_{ss-new}$, then set $t_{dd} = 0$
3. CN sends ACK to MH via old BS upon receipt of both route-update packets.
4. (a) If ACK received, MH tunes into the new BS.
 (b) Else, MH waits t_{ss} and tunes into the new BS.

Figure 3: The adaptive semi-soft handoff mechanism.

handoff, and as such, aims to cope with the dynamic network situations and differing network topologies. Furthermore, it also minimizes handoff traffic compared to the Cellular IP semi-soft handoff mechanism. Adaptive-SS uses route-update packets to probe the delays of each of the paths involved in the handoff. From the timestamps of the route-update packets traversing the different paths, the CN calculates the total delay of these paths and injects an appropriate amount of delay to the delay device. Then, the CN sends an ACK to the MH to indicate that the new path has been established and it should tune into the new BS.

The proposed Adaptive-SS scheme works as follows. Just before a MH initiates a semi-soft handoff, it sends semi-soft route-update packets towards both the new and old BSs, p_{old} and p_{new} , and tunes its radio back to the old BS. There are two ways to do this. If the adjacent cells operate using the same channel, route update packets can be broadcasted to both BSs. On the other hand, if the adjacent cells operate under different channels, the MH can first send a route-update packet to the new BS and then send the other route-update packet to the old BS. The CN can use the timestamps of these route-update packets to take into account the time delay between the transmissions of these packets in the calculation of t_{dd} .

When the CN receives p_{new} , it records the arrival time of the packet and checks whether p_{old} has already arrived. If not, this means that the delay of the old path is greater than the delay of the new path, i.e., $t_{old} > t_{new}$. When p_{old} arrives from the old link, its arrival time is subtracted by the arrival time of the packet from the new link. The delay device value, $t_{dd} = t_{old} - t_{new}$, is then adjusted accordingly. If the CN has already received p_{old} (i.e. $t_{new} > t_{old}$), t_{dd} is set to 0 since there is no need to inject any delay to the new path. This way (in conjunction with the use of ACK packet that will be discussed shortly), packet duplications due to negative stream misalignment can be minimized. For the ideal case, where t_{old} is equal to t_{new} , the delay device is set to 0.

After the CN assigns the delay device value, it sends an *acknowledgement* (ACK) packet back to the MH via the old link. In addition, the CN also stops forwarding packets to the old path. Upon receiving this ACK, the MH immediately tunes its radio to the new BS and the handoff process is complete. With the use of such an ACK packet, efficient tune-in time could be achieved as MH can be sure that the routing for the new path has been established prior to tuning in to the new BS. In addition, this ACK mechanism minimizes the downlink traffic imposed by bi-casting, since the packet forwarding to the old BS is immediately stopped instead of waiting for its soft-state mapping to time out. Also, the ACK packet acts as a “mark” for the last packet to be received by the MH via the old BS, similar to the Last Packet Marking approach described in [19]. The differences between LPM and Adaptive-SS approach will be discussed in Section 5.

Note that it is possible for the ACK packet to be not received by the MH (e.g., lost on its way down to MH, MH is out of reach of old BS, dropped due to contention, etc). As a contingency measure, if such a case happens, the original t_{ss} parameter set in the Cellular IP is used. The MH assumes that the ACK packet is lost if the packet does not arrive after it waits for t_{ss} . Thus, when Adaptive-SS is used, t_{ss} should not be set too small, since it would diminish the effectiveness of the mechanism. That is, the MH would not detect ACK packet as being lost, even though it might not be the case. We recommend setting t_{ss} to the upper bound of the handoff latency based on observing the handoffs in a particular network of interest. Naturally, this would mean t_{ss} for Adaptive-SS is larger than that of the original semi-soft handoff. The Adaptive-SS handoff scheme is summarized in Figure 3.

4. Evaluation

4.1. Simulation Framework

The Adaptive-SS hand-off scheme was implemented in ns-2 extended with the Cellular IP model from the Columbia IP Micro-mobility group [10]. To evaluate the effectiveness of Adaptive-SS in anticipating stream misalignment and in formulating efficient tune-in time, the topology shown in Figure 1 was simulated with varying path delay differences and amount of network traffic. Such a simple topology was chosen because it allowed us to accurately control the path delay difference and focus on observing its effect on handoff performance. Traces were obtained and analyzed from the MAC layer of the MH.

Note that for all the simulations discussed in this paper, the speed of MH is set so that the time the MH spends in the overlapping coverage of the cells (i.e., handoff area) is longer than the round-trip time of the paths from MH to CN. Such setup was used so that complete handoff behavior can be observed. The issue of reduction in handoff performance due to lack of time the MH spends in the handoff area is beyond the scope of this paper.

4.2. Effect of the path delay difference on handoff performance

In this simulation, a single MH receives a UDP stream at 4 ms inter-transmission interval ($t_{inter-trans}$) from the CH. A 4 ms packet inter-transmission rate was observed to provide sufficient pressure to the network while not giving too much traffic to it. All the wired link bandwidths were set to 10 Mbps, which was more than sufficient for the one network stream simulated. Because there is only one MH in the simulation, there is no RTS/CTS contention due to other network devices trying to use the wireless medium. Thus, the propagation delay of the wireless medium is fairly constant. Also, it is fairly small relative to the link propagation delays (i.e. L3 and L4). In addition, for a constant packet rate that was used in this simulation, the processing time of the nodes (i.e., GW, CN, BSs, and MH) is negligible. Therefore, the link propagation delays are comparable to the total path delays. Based on this, *Path Delay Ratio* (PDR), which is the ratio of the old path delay relative to the new path delay, i.e., L3:L4, is varied to simulate the various path delay differences. The PDR is varied from 30:1 to 1:30 by varying the propagation delay of L3 and L4 from 2 ms to 60 ms (e.g., 30:1 = 60ms:2ms; 1:1 = 2ms:2ms; 1:30 = 2ms:60ms). A PDR of 30:1 to 2:1 simulates negative misalignments, a PDR of 1:1 perfect alignment, and a PDR of 1:2 to 1:30 simulates positive misalignments.

The Adaptive-SS scheme was compared against hard handoff (Cip-Hard), semi-soft handoff with t_{dd} of 0 ms (Cip-SS), and semi-soft handoff with t_{dd} of 50 ms (Cip-SS-50). The delay value of 50 ms was chosen because it has been shown to provide the best performance in a previous Cellular IP study [9]. Figure 4 shows a comparison of the four handoff schemes for PDR of 5:1. The x-axis represents packet arrival time at the MH, while the y-axis shows the sequence number of the packet (added to the UDP packets for simulation purpose). The packets received from the old BS are shown in a different color than those received from the new BS. When packets are received at a constant rate, the data points would shape as a straight line with a positive slope since the interval between packet arrivals would be relatively constant and the packet sequence number would increase by one for each new packet received.

With Cip-SS, all the downlink packets received from the new BS arrive unsynchronized, which creates a positive stream misalignment. Thus, several packets are lost at the time the MH tunes into the new BS. On the other hand, with Cip-SS-50, the delay device puts too much delay to the packets forwarded to the new BS, which results in a negative stream misalignment. This causes the MH to receive several duplicate packets. For Cip-Hard, several packets are lost during the period MH waits for the establishment of the new routing path. These packets are routed to the old path and are not received since the MH has already tuned into the new BS by the time these packets are sent out by the old BS. As can be seen by the figure, the Adaptive-SS scheme works the best. Packets coming from the new BS arrive perfectly aligned with the packets from the old BS. Furthermore, the ACK mechanism allows the MH to know the efficient tune-in time (i.e., right after the packets are synchronized).

Figure 5 shows the efficient t_{ss} interval for Cip-SS and Cip-SS-50 for each of the simulated PDRs. Utilizing t_{ss} value that lies within this interval would yield efficient tune-in time as illustrated in Figure 2. These intervals are determined based on observing the handoff trace of each of the simulation runs for differing PDR values. As the figure shows, the efficient t_{ss} interval varies with differing PDR values. This indicates a need for dynamic determination of t_{ss} value that could fit in the efficient t_{ss} interval given a particular network condition. A fixed t_{ss} value would not be able to keep up with

the changing efficient t_{ss} interval requirement to provide efficient and seamless handoffs. Lastly, the figure also shows that higher t_{ss} value is needed with larger path delay difference (t_{ss} values for PDR of 5:1 to 30:1 of Cip-SS-50 do not increase because of the impact of adding the 50 ms to the new path).

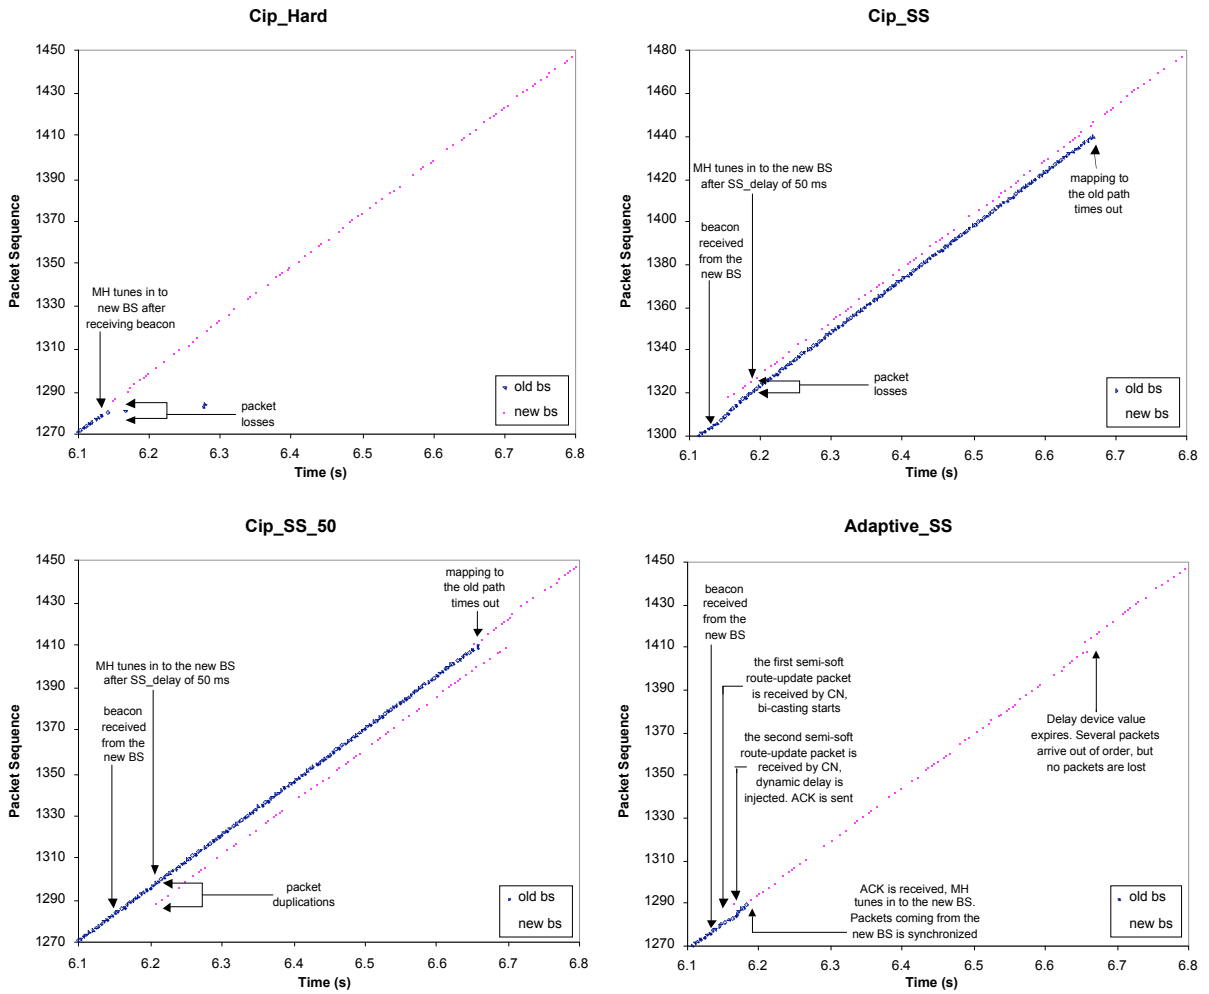


Figure 4: The phase of the simulated handoff schemes (PDR = 5:1).




Figure 5: Efficient tune-in times.

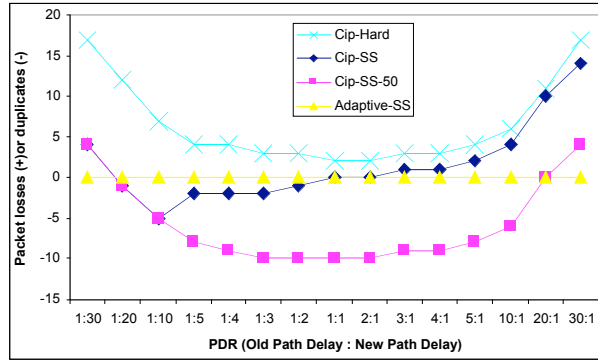


Figure 6: Handoff performance for various PDRs.




Figure 7: Handoff traffic for various PDRs.

Figure 6 summarizes the performance of the various handoff schemes in terms of packet loss (+) or duplicate (-). For both Cip-SS approaches, $t_{sw}=50$ ms is used as the default value. As expected, the hard handoff results in most packet losses. Moreover, employing a fix delay device value does not work effectively in eliminating packet losses or duplicates. This is evident from the result for Cip-SS-50, which has virtually the same increasing trend as the result of Cip-SS, except that it is shifted down. This is because of the extra 50 ms delay only shifted the stream alignment point. The figure shows that there is an increasing trend with higher path delay difference. The increasing trend in the handoff performance for PDR values of 1:2 to 1:30 (and 2:1 to 30:1) is due to earlier MH tune-in time caused by having a fixed t_{sw} value. As can be seen, Adaptive-SS performs the best out of all the schemes. For all the differing PDRs, packet losses and duplications are entirely eliminated.

Figure 7 shows the downlink traffic of the old path during handoff in terms of the number of packets (e.g., from when the semi-soft route-update packet is sent to the new BS until when there is no more packet sent via the old path). The amount of handoff traffic increases with increasing path delay difference. Cip-SS and Cip-SS-50 generate significantly more traffic compared to the Cip-Hard and Adaptive-SS. This is due to the bi-casting technique used in these schemes, which requires CN to continue forwarding packets to MH via the old BS until the routing table timeout. Both Cip-SS and Cip-SS-50 induce similar amount of traffic because they use the same timeout interval of 1.5 second. As anticipated, Cip-Hard generates the least amount of traffic. Nevertheless, the traffic generated by Adaptive-SS is only slightly higher than that induced by the Cip-Hard. For delay ratios of 1:10 to 1:30 (10:1 to 30:1), the handoff traffic of Adaptive-SS increases at a higher rate compared to delay ratios of 1:5 to 5:1. This is because Adaptive-SS offsets the higher rate of increase in path delay differences by assigning higher values to the delay device and effectively inducing more downlink traffic via the old path.




Figure 8: Handoff performance for various cell densities.

4.3. Effect of other network streams on handoff performance

The previous simulation results are based on path delay differences that are relatively constant. In realistic situations, however, handoffs are likely to occur in the presence of other network streams, which induce fluctuations in the path delays. Therefore, such a situation was evaluated based on two MH migration scenarios: (1) From high to low cell density (H→L), and (2) from low to high cell density (L→H). Cell density refers to the number of active MHs in a cell. A high density cell is simulated by having two cell-resident MHs in the cell, in addition to the moving non-resident MH. These cell-resident MHs stay in the cell and do not perform handoffs. Each of them receives a downlink UDP stream from the CH at the same rate as the moving MH. Low cell density refers to a cell without the cell-resident MHs. In a high-density cell, there is more network traffic due to the cell-resident MHs, and thus delay variability is higher due to contention among the MHs in the cell.

In order to capture only the behavior of fluctuating delays that comes from having multiple streams, the base path delay difference (i.e., $|L3-L4|$) is eliminated by setting the L3 and L4 propagation delays to the same value of 2 ms. The amount of network traffic is varied by increasing the packet transmission interval of all the three streams from 4 ms to 8 ms in steps of 1 ms. This range is selected to represent the reasonably slow to busy network condition. Our experiment with lower packet inter-transmission time (< 4 ms) resulted in an extremely congested network where packets were dropped in the high-density cell even when there were no handoffs occurring. On the other hand, for packet transmission interval higher than 8 ms, the network is light and the observed behavior is very similar to the network traffic with 8 ms packet transmission interval (e.g., very good handoff performance). The same four handoff schemes as in the previous PDR analysis are evaluated using the two scenarios.

In general, the lower the packet transmission interval (i.e., higher packet rate), the more congested the network is, and thus, the higher the path delay variation is expected to be. Thus, the handoff performance is expected to be worse with lower transmission interval. Furthermore, since the link propagation delay is set to be the same value, the path delay fluctuations are caused by other delay factors. Based on observing the simulation data, the main factor for path delay variation is the MAC protocol, which is based on CDMA with RTS/CTS control packets. In the high-density cell, the time required to gain access to the wireless medium via RTS/CTS varies greatly due to contention among the MHs residing in the cell. The variability due to RTS/CTS will be discussed in the later part of this section.

Figure 8 shows the handoff performance for the various packet-transmission intervals simulated for both L→H and H→L scenarios. As expected, Cip-Hard resulted in the largest number of packet losses. The Cip-SS handoff schemes resulted in less packet losses than Cip-Hard and adding the fixed t_{dd} value of 50 ms only offsets the handoff performance to yield either decreased packet losses or increased packet duplicates in most cases. For the L→H scenario with intervals of 5 and 4




Figure 9: BS processing times of the high-density cell for Adaptive-SS.

ms, the impact of the delay device is not noticeable because the high density of the cell the MH is moving to results in increased delay in the packet queue time at the new BS, which in turn absorbs the impact of t_{dd} . As can be seen, Adaptive-SS yielded the best overall performance in terms of absolute number of packet losses or duplicates. On average, Adaptive-SS performed significantly better than the other approaches for the L→H scenario. For the H→L scenario, Adaptive-SS performed moderately better than the other approaches, but also the performance of all the schemes simulated were better than the L→H scenario.

In general, long inter-transmission intervals of 8 ~ 6 ms resulted in very few or no packet losses for both L→H and H→L scenarios. However, the handoff performance for all the schemes degrade significantly for packet transmission intervals of 5 ~ 4 ms. At these intervals, there are too many packets sent and the network is handling these packets at its limit. Analyses of the simulation traces indicate that the main source of variance in the total path delay of the high-density cell comes from the time packets spend in the BS. Other delay parameters, such as the wired link and air propagation delays and CN and GW processing times, remain relatively constant. BS processing time, however, varies quite significantly for packet transmission interval of 5 ms or less. This is due to the CTS/RTS mechanism used to multiplex the access to the wireless medium. As multiple MHs compete for the wireless channel, packets experience longer delays waiting to be sent.

For the Adaptive-SS scheme, depending on whether broadcast or unicast is used to send route-update packets, the *Address Resolution Protocol* (ARP) becomes another factor that adds to the variance in the total path delay. If broadcast is used, unlike unicast, an ARP request is not initiated since the broadcast MAC address is used instead of the MAC address of the BSs. Therefore, when the new BS receives the first packets destined for the MH, it does not have the MAC address of the MH and needs to initiate ARP. Then, the packet is put into the ARP buffer while the ARP request contends for the wireless medium. During this period, if another packet destined for the MH arrives, the packet in the ARP buffer is over written (i.e., dropped). These ARP packet drops especially degrades the performance of Adaptive-SS for the L→H scenario, as the ARP delay becomes much longer when the MH enters the high-density cell. Note that this problem does not exist if unicast is used to send route update packets.

A simple solution is to have either the new BS or the MH to send an ARP request right after the route-update packet is sent by the MH (rather than waiting until the first packet is received by the new BS). This approach, called *ARP follow-through*, allows the ARP process and the new path establishment to be done in parallel. As shown in Figure 8, the *ARP follow-through* significantly improves the performance of Adaptive-SS under L→H scenario for congested situations (i.e. 4 ms and 5 ms intervals). However, for less congested situations, there is very little or no improvement. This is because for these situations, the packet losses are not due to ARP packet drops.




Figure 10: Handoff traffic for various cell densities.

Figure 9 shows the average (with max. and min.) processing time of BS in the high-density cell during handoff for Adaptive-SS. The BS processing time for a packet is determined by the time between when BS receives the packet and when BS sends the packet to the MH. The average (and max/min) processing times are obtained from the last (H→L) or the first (H→L) 100 packets sent by the BS to MH in the high-density cell. As can be seen from the figure, there is a significant increase in the average processing time for packet transmission intervals of 6 ms to 5 ms, which was the main reason for the degradation in performance of Adaptive-SS. Furthermore, for packet transmission interval of 5 ms or less, BS average processing time increases steadily, with the range (i.e., max. – min.) increasing at a higher rate than the processing time. This means that as packet transmission interval decreases, packets arriving at the BS would experience a wider possible range of delays before they are sent to the MH.

For the H→L scenario, this means that the Adaptive-SS’ ACK packets sent from CN to the MH suffers from the high variability in MAC delay and as a result less accurately predicts the efficient tune-in time for the MH. For the L→H scenario, the MH experiences tune-in latency in the high-density cell the MH is migrating to due to the network traffic contention, increased RTS/CTS wait time, and ARP time.

Finally, Figure 10 shows the downlink traffic generated by each of the handoff schemes. For the H→L scenario, we only calculated the number of packets sent via the old path for the moving MH, since the amount of packets for the other two mobile hosts are relatively constant. As discussed previously in Subsection 4.2, the semi-soft handoff generates significantly more traffic out of all the simulated schemes. Cip-Hard yields the least amount of handoff traffic, while Adaptive-SS is only slightly higher.

5. Related Work

There are various micro-mobility protocols proposed in the literatures [2, 3, 4]. The three prevailing ones that have been extensively evaluated in the previous studies are Hawaii [11], Hierarchical MIP [12], and Cellular IP [8]. Each of these three protocols presents different fast handoff mechanisms. In general, however, fast handoff mechanisms are based on two underlying ideas: *bi-casting* and *buffering and forwarding* [9].

In *bi-casting*, the knowledge of the new BS is gathered ahead of time. Using this information, the connection to the new BS is prepared in advance. As such, this approach involves bi-casting of downlink streams to the MH via both the old and new paths during handoff. The bi-casting approach is used in Cellular IP semi-soft handoff, Hierarchical Mobile IP Fast Handoff [13] and Foreign Agent Assisted [14] handoff. In *buffering and forwarding*, no knowledge of the new BS is needed. Instead, this approach focuses on recovery after handoff. During handoff, packets are placed in a buffer in the old BS. After handoff, these packets in the buffer are forwarded to the mobile host via the BS. Hawaii MSF path setup scheme, and Buffer management [15] and Generalized IP [16] handoffs for Hierarchical Mobile IP are based on the buffering and forwarding approach.

There are two handoff extensions that have been proposed for the Cellular IP protocol: *Indirect semi-soft handoff* and *Last Packet Marking*. Indirect semi-soft handoff is included in the Cellular

IP draft [17] as an extension to the handoff protocol for such cases where the wireless technologies used do not allow MHs to simultaneously communicate with the multiple stations involved in the handoff (e.g., TDMA networks). Evaluation of the handoff extension showed that it has better performance compared to the hard-handoff but still worse than semi-soft handoff [18]. Last Packet Marking (LPM) [19] mechanism addresses the same issue as our Adaptive-SS. It aims to ameliorate packet losses and duplications due to fix delay parameters usage proposed in Cellular IP. LPM works as follows. MH sends route-update packet to the new BS to initiate a handoff. Upon receiving the route-update packet, CN sends a semi-soft reply (i.e., ACK) to the old BS and then multicasts data packets to both new and old BSs. MH tunes into the new BS after it receives the semi-soft reply. Instead of using a delay device, the CN forwards data packets immediately and the new BS buffers them in case these packets arrive before the MH tunes into the new BS. In the case where the MH tunes into the new BS before semi-soft reply arrives, data packets to the old BS before the semi-soft reply are forwarded to the new BS. If the MH tunes into the new BS after it has received several data packets after the semi-soft reply, packet duplications are eliminated at the new BS.

There are two major differences between LPM and the proposed Adaptive-SS: (1) LPM uses buffering and forwarding techniques while Adaptive-SS uses route-update packet sends. The advantage of Adaptive-SS is that it adds very little to the network bandwidth during handoff. On the other hand, packet forwarding would add extra traffic during the handoff. (2) The evaluation of LPM includes only two sets of delay parameters. Thus, it is difficult to determine how the mechanism will perform as path delay difference increases, or as more packets congest the network. On the other hand, Adaptive-SS was evaluated under many sets of delay parameters by testing for differing path delay differences and network loads. Our simulation results show that Adaptive-SS performs significantly better than semi-soft and hard handoffs in reducing packet losses/duplicates and handoff traffic.

6. Conclusion

In this paper, we show that semi-soft handoff is ineffective in handling dynamic network behavior due to the use of fix values for the delay device and semi-soft delay. The proposed Adaptive-SS scheme addresses this issue by dynamically determining the delay values based on the network condition at the time.

Our simulation results show that Adaptive-SS is very effective in anticipating base path delays, such as the link propagation delay, which resulted in relatively constant delay differences between the old and new paths involve in the handoff. For the path delay differences caused by factors that fluctuate, such as the presence of other network streams and the intensity of the streams, the adaptive approach more effective than other methods simulated. In addition, Adaptive-SS also minimizes handoff traffic up to the point that is almost comparable to the handoff traffic generated by hard handoff approach.

For future work, we plan investigate further the effectiveness of the Adaptive-SS approach for TCP/IP streams and mixed typed streams under a more dynamic and erratic experimental setup. In addition, it would also be interesting to analyze real wireless network infrastructures and characterize the various factors that contribute to the path delays during handoff and their variability. Lastly, perhaps a more aggressive approach utilizing periodic signaling packets that monitor the network condition could be devised to better anticipate the dynamic delay factors. Other optimizations for Adaptive-SS could be developed as well.

7. References

- [1] C. Perkins, Ed., "IP Mobility Support", *Internet RFC 2002*, Oct 1996.
- [2] A. T. Campbell and J. Gomez, "IP Micro-Mobility Protocols," *ACM SIGMOBILE, Mobile Computing and Communications Review*, vol. 4, no. 4, pp. 45-54, Oct. 2001.
- [3] P. Reinbold and O. Bonaventure, "A Comparison of IP Mobility Protocols," *Infonet Technical Report 2001-07*, June 2001.
- [4] K. Wong *et al.*, "Performance of IP Micro-Mobility Management Schemes using Host Based Routing," *The 4th International Symposium on Wireless Personal Multimedia Communications*, September 2001.

- [5] A. Valko, "Cellular IP: A New Approach to Internet Host Mobility," *ACM Computer Communication Review*, pp. 50-65, Jan. 1999.
- [6] A. G. Valko, J. Gomez, S. Kim, and A. T. Campbell, "On the Analysis of Cellular IP Access Networks," *IFIP Sixth International Workshop on Protocols for High Speed Networks (PfHSN'99)*, August 1999.
- [7] S. Kim *et al.*, "A Cellular IP Demonstrator," *Sixth IEEE International Workshop on Mobile Multimedia Communications (MOMUC'99)*, November 1999.
- [8] A. T. Campbell *et al.*, "Design, Implementation and Evaluation of Cellular IP," *IEEE Personal Communications, Special Issue on IP-based Mobile Telecommunications Networks*, June/July 2000.
- [9] A. Campbell *et al.*, "Comparison of IP Micromobility Protocols," *IEEE Wireless Communications Magazine*, Vol. 9, No.1, Feb. 2002.
- [10] Micromobility home page, <http://comet.columbia.edu/micromobility>
- [11] R. Ramjee *et al.*, "HAWAII: A Domain-Based Approach for Supporting Mobility in Wide-area Wireless Networks," *Proceedings of IEEE International Conference on Network Protocols*, 1999.
- [12] E. Gustafsson, A. Jonsson, and C. Perkins, "Mobile IP Regional Registration," Internet draft, draft-ietf-mobileip-reg-tunnel-03, work in progress, July 2000.
- [13] K. El-Malki and H. Soliman, "Fast Handoffs in Mobile IPv4," Internet draft, draft-elmalki-mobileip-fast-handoffs-03, work in progress, Sept 2000.
- [14] P. Calhoun *et al.*, "Foreign Agent Assisted Handoff," Internet draft, draft-calhoun-mobileip-proactive-fa-03, work in progress, Nov. 2000.
- [15] C. Perkins, G. Krishnamurthi, and R. Chalmers, "Buffer Management for Smooth Handovers in Mobile IPv6," Internet draft, draft-krishnamurthi-mobileip-buffer6-00, work in progress, July 2000.
- [16] A. O'Neill, S. Corson, and G. Tsirtsis, "Generalized IP Handoff," Internet draft, draft-oneill-craps-handoff-00, work in progress, Aug. 2000.
- [17] Z. D. Shelby *et al.*, "Cellular IPv6," Internet Draft, draft-shelby-seamoby-cellularipv6-00.txt, IETF Mobile IP Working Group Document, November 2000.
- [18] M. Ghassemian and A. H. Aghvami, "Comparing different Cellular IP with Hawaii handoff schemes," *The third international conference on Mobile Communication Technologies*, May 2002.
- [19] J. Kim, K. Kim, J. Y. Park, and C. Kim, "An Enhanced Handoff Mechanism for Cellular IP," *The Third International Network Conference*, July 2002.
- [20] D. C. Plummer, "An Ethernet Address Resolution Protocol," *Internet RFC 826*, Nov 1982.