A Simulation Study on TCP Performance with Vertical Handoffs in WLAN-GPRS/UMTS Environment
Technical Report,
University of Helsinki, Department of Computer Science,
June 2005.

Laila Daniel and Markku Kojo

Department of Computer Science, University of Helsinki, Finland.
{ldaniel, kojo}@cs.helsinki.fi

Abstract In this report we examine the performance of TCP with vertical handoff in WLAN - GPRS environment.

1 Introduction

With the increasing deployment of WLAN hot-spots and the universal presence of GPRS / UMTS, it is now possible to provide continuous Internet connectivity to mobile users using a variety of devices such as laptops, PDAs and mobile phones. In order to provide this connectivity, mobile devices have several wireless network interfaces. Wireless networks are often organized in an overlay structure in which a faster network with less coverage is contained in a slower network with a larger coverage [13]. Vertical handoff introduced in [13] refers to the process where the interface to the wireless network changes from one wireless network to another wireless network which uses a different link layer technology. Thus vertical handoff enables a mobile user to have the best possible connectivity over a wide geographical area in a transparent manner.

As wireless technologies differ widely in their characteristics such as bandwidth, latency, coverage and media access techniques, there is a need to devise means to make handoff smooth. Though Mobile IP / IPv6 [11,7] can handle vertical handoff without breaking any ongoing transport connection, the change in the link characteristics due to handoff can affect the performance of TCP [3,4,5,6]. The change in bandwidth and propagation delay during the vertical handoff may cause packet reordering and delay the arrival of acknowledgments (ACKs). This in turn may trigger unwanted TCP congestion control actions thereby degrading the performance of TCP as TCP congestion control is based on the assumption that the end-to-end path of a connection remains largely unchanged after a connection is established [1].

In this report we present a systematic simulation study of the performance of TCP in various vertical handoff scenarios. The different scenarios include soft / hard handoff between WLAN and GPRS with handoff occurring at different times during a TCP connection. We also conducted simulations with handoff delay. From the study we can see that handoff affects short TCP transfers more significantly than large TCP transfers. The inherent problems of TCP with slow links such as GPRS is more significant than the delays caused by the handoff in the case of large file transfers.

The rest of the report is organized as follows. Section 2 outlines the related work in this area. Section 3 describes vertical handoff and section 4 discusses how vertical handoff affects the TCP performance. Section 5 describes the proposed modification to the TCP Sack sender algorithm. In section 6 we evaluate the algorithm by simulation. The last section presents the conclusions of the study.

2 Related Work

Many recent papers discuss the problems with TCP associated with vertical handoff in wireless networks [3,4,5,6].

A practical study on the performance of TCP with vertical handoff between GPRS and WLAN is presented in [3]. It points out that the delay in handoff often causes TCP to timeout and contributes the degradation in TCP performance. The high buffering in GPRS aggravates the performance of TCP as it inflates the round trip time (RTT) and retransmission timeout (RTO) values. This paper also presents some network layer optimizations to reduce the handoff delay.

A comparative study of the effect of vertical handoff on transport protocols such as TCP and TCP-friendly rate control (TFRC) is presented in [4]. Using measurements and simulation it shows that TFRC has significant difficulties in adapting to the network after vertical handoff because of the change in bandwidth delay product (BDP) of the link. This paper suggests over-buffering to reduce the problem due to the differences in BDP. A drawback of this scheme is that the knowledge of the link BDP is important to set the proper size of the buffer at the sender.

A dynamic value of the dupThresh parameter of TCP is proposed in [5] to reduce the effect of packet reordering during vertical handoff from a low bandwidth-high latency network to a high bandwidth-low latency network. Since dupThresh value decides when TCP acts on duplicate ACKs changes to the value of dupThresh make TCP robust to packet reordering.

The Internet draft [14] proposes a lightweight mobility detection and response (LMDR) algorithm for TCP which makes TCP aware of the path change during vertical handoff. This draft also outlines the congestion control behavior that should take place during the vertical handoff.

The paper [6] proposes three network layer schemes to the problem of packet reordering, BDP mismatch and increased RTT in soft handoff scenarios. These schemes called fast response, slow response and acknowledgment delaying, reduce the difference in RTT between the old link and the new link. The fast response scheme requires using the old link for a short period to be used for sending the ACKs after the handoff. In the slow response scheme, a few ACKs are sent through the slow link just before handoff. These two schemes may have practical problems if the old link is not available just after handoff and if the new link cannot be used just before handoff. In the acknowledgment delaying scheme, the ACKs of the last few fast packets are delayed at the IP layer of the sender before passing them to the TCP sender. Here a proper selection of how much the packets are to be delayed at the IP layer is needed.

3 Vertical Handoff and TCP performance

The term *Vertical handoff* refers to the process of switching the network interface between wireless networks which employ different link layer technologies. Switching between WLAN to GPRS network is an example of vertical handoff.

Handoff can be classified as hard handoff and soft handoff. *Hard handoff* is set to occur if the mobile device can be associated with only one type of access point at a time. This is a break-before-make type of connection. When the mobile device undergoes a hard handoff, the network interface has to decouple with its old access point and then connect to the new one. This will result in packet losses and disruption in connectivity.

The *soft handoff* occurs when the mobile node has more than one access points at a time. If the mobile device has more than one network interface, it can simultaneously connect to two access points of two wireless networks. This can be viewed as a make-before-break type of connection. The mobile node ends its connection to the old network only after establishing the connection to the new network. Thus soft handoff can prevent packet losses due to handoff. An example of soft handoff is the scenario where a GPRS connectivity is always present and when a WLAN hot-spot becomes available (as noticed from network scan) a handoff taking place from GPRS to WLAN.

Mobility management protocols like Mobile IP / IPV6 [11,7] make the transport layer protocols such as TCP unaware of the change in the network interface. Since vertical handoff switches the mobile node between wireless networks of different link characteristics such as bandwidth and latency, the handoff will affect the performance of TCP.

In this section we discuss in detail the effect of vertical handoff on TCP performance. We present the handoff from WLAN to GPRS which is a representative of a handoff from a high bandwidth, low latency link (called a fast link) to a low bandwidth, high latency link (called a slow link). We also present the behavior of TCP when the handoff occurs from GPRS to WLAN, i.e., from a slow link to a fast link.

3.1 WLAN-GPRS Handoff

In the case of WLAN to GPRS handoff, TCP performance is affected by the change in BDP and also by the increase in RTT.

Usually the size of the link buffer is set to the BDP of the link. If the congestion window (cwnd) matches the link BDP, TCP can highly utilize the link without any buffer overflow. If there is a mismatch between the cwnd and the link BDP, there can be either buffer overflow or under utilization of the link.

When a handoff occurs from a high BDP network to a low BDP network, some data already buffered can be lost due to insufficient buffer space. On the other hand, during the handoff from a low BDP network to a high BDP network, the number of buffered packets may not be enough to utilize the new link. In the case of vertical handoff from WLAN to GPRS, handoff can occur at a point where TCP has increased the cwnd to a large value due to high transmission rate of WLAN. TCP uses the same cwnd after handoff and sends a burst of packets to GPRS which cannot handle this burst due to small buffer size in GPRS and many packets will be dropped at the link.

Another problem which arises due to handoff from WLAN to GPRS is the erroneous calculation of TCP timeout. TCP calculates the RTO value based on Jacobson's algorithm [**?**]. Upon receiving an acknowledgment and obtaining a new SampleRTT, TCP updates the EstimatedRTT. RTO is calculated using the following formulas

$$EstimatedRTT(i) = (1 - \alpha) * EstimatedRTT(i - 1) + \alpha * sampleRTT(i)$$

$$RTO(i) = EstimatedRTT(i) + \beta * Deviation(i)$$

$$Deviation(i) = (1 - \alpha) * Deviation(i - 1) + \alpha * |SampleRTT(i) - EstimatedRTT(i - 1)|$$

RTO is calculated based on the EstimatedRTT and Deviation where Deviation is an estimate of how much SampleRTT deviates from EstimatedRTT. From the above equations it can be seen that RTO is influenced by the history of the observed RTT and the amount of RTT deviation. RTO increases as the SampleRTT increases and also increases if the difference between RTTs is large.

In the case of WLAN-GPRS handoff, TCP has a relatively small RTO due to the small value of RTT in the WLAN link. The packets transmitted just before handoff and during handoff will get their ACKs through the slow GPRS link even though the packets transmitted just before handoff are received at the receiver. The retransmission timer will expire as the current RTO value is based on the estimated timeout in the fast WLAN link but not on the slow GPRS link to which the mobile device is switching to. Because of the RTO expiry the TCP sender retransmits the packets sent just before handoff and during handoff unnecessarily and also enters slow-start phase by reducing cwnd to 1. In this case the packets are not lost and the ACKs are coming through the slow link, the RTO does not reflect the current state of the network and the congestion control actions are erroneous. Retransmissions in this case is a waste of network resources and resetting cwnd underutilizes the link.

Figure 1 illustrates the TCP behavior during WLAN-GPRS handoff. Here the handoff from WLAN to GPRS occurs at 0.6 sec for a 200 KB TCP flow. The buffer overflow occurs at the time of handoff and 12 packets are dropped. Since the ACKs are delayed, a timeout happens at 0.89. There are 29 unnecessary retransmissions. The transfer of 200KB data by WLAN takes only 0.65 sec but due to the handoff to GPRS at 0.6 sec, the transfer completes only at 5.2 secs.

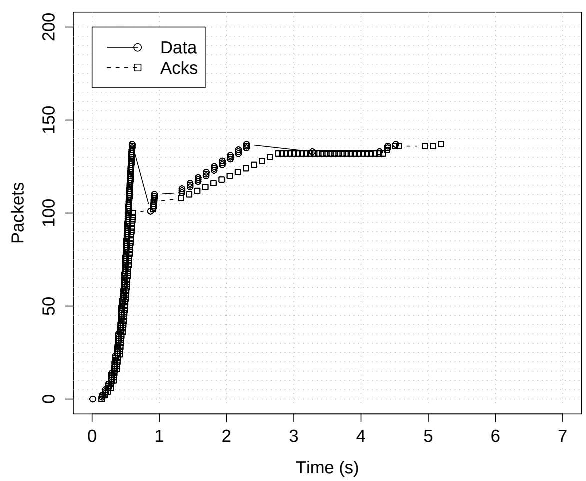

Figure 1. WLAN ⇒ GPRS, handoff at 0.6 for 200KB, the effect of delayed acks

3.2 GPRS-WLAN Handoff

Packet reordering is the problem that TCP faces when handoff takes place from a slow link to a fast link, in our case from GPRS to WLAN. The packets sent just before handoff will arrive through the slow GPRS link whereas packets sent after handoff traverse the fast WLAN link and reach the destination earlier. This reordering causes the TCP sender to receive multiple duplicate ACKs (dupacks). When the TCP sender receives a dupThresh number of duplicate ACKs, it takes congestion control measures such as fast retransmit and recovery and halves the cwnd. In current TCP implementations a typical value of dupThresh is 3 [1].

Figure 2 shows the packet reordering problem when handoff occurs from GPRS to WLAN. The handoff occurs for 200KB TCP flow at 6.2 secs and dupacks start coming at 6.32 secs. Even though this creates unnecessary cwnd reduction, the transfer completes at 6.8 secs as WLAN is a fast link.

We can use either Eifel [9,8] detection algorithm or F-RTO [12] algorithm to detect the unnecessary retransmissions. Using the response part of these algorithms to restore the cwnd and ssthresh values may not be always correct as it may restore the ssthresh and cwnd values of the old link.

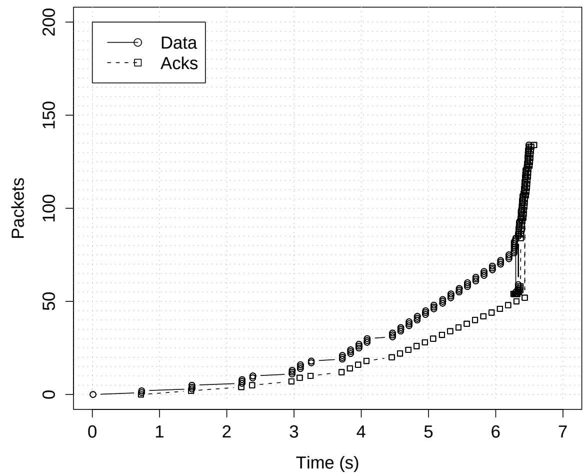

Figure 2. GPRS ⇒ WLAN, handoff at 6 for 200KB, the effect of dupacks

4 Simulation Environment for Vertical Handoff

The main objective of the simulation study presented here is to explore the effect of handoffs on TCP performance when handoff takes place between WLAN and GPRS / UMTS network and vice versa. During handoff link characteristics such as bandwidth and delay change and link connectivity may drop.

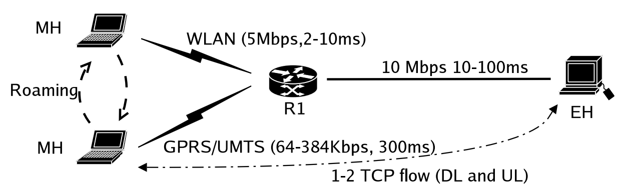

Figure 3. Network topology for handoff tests

4.1 Link Characteristics

The network topology for the handoff scenario is given in Figure 3. The Mobile host (MH) moves from WLAN to GPRS / UMTS and vice versa. The router connects the WLAN and GPRS network to the End host (EH) which is a part of the fixed network.In our simulation model we take the following parameters. The GPRS links have bandwidth of 200 kbps and the WLAN has bandwidth of 5 Mbps. The MTUs used in the tests is 1460 bytes for TCP traffic. The buffer sizes at the sender and the receiver are adjusted to the delay-bandwidth product of the link. We are doing the simulation in network simulator ns-2 [10]. We are using TCP Sack [2] for our simulations. Table 1 describes the parameters of the link configuration.

Characteristics	GPRS/UMTS	WLAN
Propagation delay	300 ms	2-10ms
Transmission rate	64 kbps & 384 kbps	5Mbps
MTU	1460 bytes	1460 bytes
Router queue size	Bandwidth*RTT	

Table 1. Link characteristics for Handoff Tests

4.2 Workload Models

The workload in the handoff scenario is TCP only. There is a single TCP flow from the mobile host to the end host and vice versa. The TCP data transfer is long enough to study the handoff between the two networks. As a first step, we carry out experiments with 20KB, 200 KB and 20MB data transfers in the downlink direction.

5 Simulation Models for Vertical Handoff

We have the following simulation models.

5.1 Ideal Vertical handoff

By ideal vertical handoff we mean that there is no handoff delay. The handoff delay refers to the total duration of the handoff at the IP layer. In the ideal vertical handoff we model the handoff by just changing the link characteristics. The link bandwidth, latency and buffer size can be changed instantly in ns-2. We model both soft and hard handoff with the ideal vertical handoff scheme

– Ideal vertical handoff with soft handoff: There are two kinds of behaviors in ns-2 when the buffer size is changed. The default behavior is not to discard packets if the new link has a buffer smaller than the old link. This models the soft handoff. In the other model, if the new link has less buffer compared to the old link, the excess packets will be dropped. This can be done in ns-2 using the 'shrink_queue' procedure which will shrink the queue if the buffer after handoff is smaller than the buffer before handoff.
– Ideal vertical handoff with hard handoff: To simulate this handoff model, we make the buffer size to for a very short period of time, say 1 msec when the handoff takes place. Using the shrink_queue procedure all the packets are dropped since the new buffer has zero size. Then make the buffer size the value of the buffer of the link after handoff. We can simulate the situation where all packets are dropped or a percentage of the total packets are lost.

5.2 Vertical handoff with handoff delay

In this model, we include a handoff delay. We have taken three sets of handoff delays. The handoff delay from GPRS to WLAN as 1, 3 and 5 as the minimum, typical and maximum values respectively. The corresponding values for the handoff delay from WLAN to GPRS as 4, 7 and 12. The modeling of delay is done using the delayer in ns-2 and blocking it at the handoff time and unblocking it after the delay time.

– Vertical handoff with delay and with soft handoff: The soft handoff is as described in the case of ideal vertical handoff model.

– Vertical handoff with delay and with soft handoff: The hard handoff is as described in the case of ideal vertical handoff model.

The simulation models used in this study are summarized in Table 2.

Simulation Models		
WLAN ⇒ GPRS	vhoideal	soft handoff hard handoff
	vhodelay	soft handoff hard handoff
GPRS ⇒ WLAN	vhoideal	soft handoff hard handoff
	vhodelay	soft handoff hard handoff

Table 2. Simulation Models for Vertical Handoff

We have also carried out simulations with the same amount of buffers for both WLAN and GPRS and also with different amount of buffer space for WLAN and GPRS.Since we want to study the effect of handoff on TCP, we make the handoff happen at different instants during the TCP connection (at the beginning of the connection, just before slow-start overshoot, after slow-start overshoot, in the congestion avoidance phase). We carry out the simulations for all these cases as handoff instants vary from 0 to 30 secs. We also randomize the starting point of the handoff instant. The staring point for each handoff is uniformly distributed in a 1 second interval. For each handoff time, 32 replications are done to observe the average behavior in the neighborhood of the chosen point. We repeat the experiments with TCP Sack and modified TCP Sack algorithms and then compare the results. In our experiments we assume that handoff occurs only once in the lifetime of a TCP connection.

6 Simulation Results: Vertical Handoff in Ideal Environment

6.1 GPRS-WLAN Handoff

In the case GPRS-WLAN soft handoff, as discussed in section 3.2 dupacks are generated due to reordering that takes place as the mobile host moves from GPRS to WLAN. The generated dupacks cause TCP to take unnecessary congestion control measures and the congestion window (cwnd) reduces. The effect due to reordering is more prominent when the handoff happens during the early phase of the TCP connection because window reduction due to handoff at the early phase affects the growth of the window and thus increases the transfer time.

If the buffers in the GPRS and the WLAN are of same size, there are no packet losses due to handoff. The only packet losses are due to buffer overflow and slow start overshoot. In our experiments with different buffer sizes we have taken the buffer size of GPRS is smaller than that of WLAN, no packet losses occur when mobile host moves from GPRS to WLAN.

In our experiments we select different points in time at which handoff occurs. Tables 3 and 4 show the transfer times for 200KB. 2MB and 20 MB data transfer. We can see that the transfer time is reduced if the handoff happens at the beginning of the connection. This is obvious as the WLAN has a high bandwidth and low propagation delay compared to GPRS and if the handoff happens during the early phase, WLAN can transfer the data more efficiently.

In the case of GPRS-WLAN hard handoff, no reordering of packets will take place as all the packets in the buffer before handoff are lost. So all these packets are to be retransmitted and the transfer time increases, and it is slightly more than the corresponding soft handoff.

6.2 WLAN-GPRS Handoff

As discussed in section 3.1 we can see that in the case of WLAN-GPRS handoff there is an early expiry of retransmit timer and the cwnd reduces. many packets are retransmitted unnecessarily.

If the handoff occurs during the slow start overshoot or at the point of buffer overflow, many packets will be lost. The loss recovery done through the slow wireless links increases the transfer time. In our experiments with different handoff times, we can see that for a particular handoff time there can be 2- 5 % difference in the transfer time among the 32 replications of a particular handoff time. This is mosly due to the buffer over flow.

In the case of WLAN-GPRS handoff, the effect of handoff is more prominent when handoff happens during the early phase of the connection. The slow start overshoot in WLAN occurs around 0.7 and if the handoff happens during this time, a number of packets will be lost and the recovery by GPRS will take a long time. Since GPRS buffer size is smaller compared to WLAN, packet losses occur if the WLAN buffer is full. Transfer time reduces as the handoff happens later during the TCP connection.

In the case of WLAN-GPRS hard handoff all the packets in the WLAN buffer just before handoff will be lost. The transmitter timeouts and all those packets are retransmitted. Since loss recovery happens through the GPRS link, we can see that there is a 1- 5 % increase in the transfer time compared to that of soft handoff.

7 Simulation Results: Vertical Handoff with delay

We run both GPRS-WLAN handoffs and also WLAN-GPRS handoffs with different handoff delay as given in section 5.2.

7.1 GPRS-WLAN Handoff

TransferTime-200KB-GPRS-WLAN					
HTime	**Vhoideal soft handoff**	**Vhoideal hard handoff**	**Vhodelay1**	**Vhodelay3**	**Vhodelay5**
0	0.95	0.95	1.2	3.2	5.2
1	2.2	2.2	2.2	4.8	6.8
3	4.1	4.1	4.2	6.5	8.6
5	5.6	5.8	5.9	8.1	10.2
TransferTime-2MB-GPRS-WLAN					
0	3.9	3.9	4.2	6.2	8.2
1	5.2	5.2	5.1	7.9	9.9
3	7.3	7.3	7.0	9.6	11.7
5	8.7	8.8	8.9	11.1	13.5

Table 3. GPRS-WLAN Handoff. HTime refers to the time at which handoff occurs

TransferTime-20 MB-GPRS-WLAN				
HTime	**Vhoideal**	**Vhodelay1**	**Vhodelay3**	**Vhodelay5**
0	33.2	33.8	35.8	37.8
1	34.2	34.7	37.5	39.5
3	36.5	36.6	39.2	41.3
5	38.1	38.4	40.7	43.1
7	39.7	40.4	42.9	45
11	43.7	44.3	46.6	48.6
15	47.7	48.3	50.3	52.6
20	52.8	53	55.3	57.9
25	57.8	57.8	60.1	62.5
30	62.4	62.6	64.9	66.9

Table 4. GPRS-WLAN Handoff. HTime refers to the time at which handoff occurs. The table gives the time taken to transfer 20 MB

In the case of GPRS-WLAN soft handoff with a handoff delay of 1, the transfer time is not increased much compared to the handoff in the case of ideal vertical handoff. This is because the handoff delay was small compared to the TCP RTO and no timeout occurred due to the delay. Window halving takes place due to the arrival of dupacks.

When the handoff delay becomes 3 and 5, timeout takes place in both cases. That is why there is a 10 % increase in the transfer time. Once WLAN takes over, the transfer of data will happen in a faster way. The transfer times for the ideal handoff and the handoff with different delays for 200KB, 2 MB and 20 MB flows are shown in Tables 3 and 4.

There is a slight increase in the transfer time in the case of hard handoff with handoff delay.

7.2 WLAN-GPRS Handoff

In all the cases with WLAN-GPRS soft handoff with delay, timeout happens and there is an increase of about 1-2% in the transfer time. The Tables 5 and 6 give the transfer times of handoff happens in the ideal case and with delay.

TransferTime-200KB-WLAN-GPRS					
HTime	**Vhoideal soft handoff**	**Vhoideal hard handoff**	**Vhodelay4**	**Vhodelay7**	**Vhodelay12**
0	8.2	8.2	16.4	19.5	24.3
1	0.66	0.66	0.66	0.66	0.66
3	0.66	0.66	0.66	0.66	0.66
5	0.66	0.66	0.66	0.66	0.66
TransferTime-2MB WLAN-GPRS					
0	82.6	82.6	88.8	93.7	98.5
1	59.8	59.8	67.4	70.4	75.4
3	10.1	10.1	19.1	22.1	27.1
5	3.6	3.6	3.6	3.6	3.6

Table 5. WLAN-GPRS Handoff. HTime refers to the time at which handoff occurs

TransferTime-20 MB-WLAN-GPRS				
HTime	**Vhoideal**	**Vhodelay4**	**Vhodelay7**	**Vhodelay12**
0	824.1	828.6	833.4	838.3
1	807	807	810.1	815.1
3	759.3	759	762.1	767.1
5	711	711.1	714.1	719.1
7	663	663	666	671
11	567.1	567.1	570	575
15	472	471	474	479
20	351.7	351.7	354	359
25	231.7	231.7	234	239
30	111.3	111.1	114.1	119

Table 6. WLAN-GPRS Handoff. HTime refers to the time at which handoff occurs

The transfer times for the ideal handoff and the handoff with different delays for 200KB, 2 MB and 20 MB flows are shown in Tables 5 and 6. It can be seen that the effect of delay is more on 200KB and 2 MB flows than on 20MB flow. The low bandwidth and high latency of the GPRS link increases the transfer time. Also due to buffering, the RTT gets inflated and there is an increase in RTO [3]. This also increases the transfer time.

8 VHO and TCP Enhancements

It can be seen from the Table 6 that in the case of WLAN-GPRS handoff, the transfer time for 20 MB is large if handoff occurs during the early phase of the connection. The main reason for this increased transfer time is due to the low bandwidth and high latency of GPRS links. The buffer overflow occurring in the GPRS link also degrades TCP performance.

In the case of GPRS-WLAN handoff, the transfer time is reduced the sooner the handoff to WLAN happens. As mentioned in section 6.1 reordering takes place and TCP congestion window reduces. We can use Eifel [9,8] detection algorithm or F-RTO [12] algorithm to detect the unnecessary retransmissions, but using the response part of these algorithms is not logical. This is because we cannot use the earlier ssthresh or cwnd value as the link has changed.

As suggested in the Internet draft [14] after the handoff the TCP connection is to be treated as a new connection and we have to use the new RTO, ssthresh and cwnd values. In the case of GPRS-WLAN handoff we are not fully utilizing the link capacity of WLAN. Even if we use an arbitrary high ssthresh value, because of reordering, TCP will switch to operate in the congestion avoidance region. So it is important that we have to find the time at which we can use the new value of ssthresh. The transfer time can be reduced further, if we are able to detect the unnecessary retransmissions and set the correct value for ssthresh and cwnd.

9 Conclusions and Future Work

In this report we present a systematic simulation study of the performance of TCP in various vertical handoff scenarios. The different scenarios include soft/ hard handoff between WLAN and GPRS with and without handoff delay and handoff occurring at different times during a TCP connection. We can see due to the change in bandwidth and propagation delay, problems such as reordering, delayed acknowledgments and buffer overflows occur due to handoff. From our study we can see that handoff affects short TCP transfers than the long TCP flows. In the case of long transfers, the inherent problems of TCP on slow links such as GPRS are more significant than the delays caused by the handoff. Another conclusion of our study is that in order to fully utilize a high bandwidth, fast link like WLAN after a GPRS-WLAN handoff, we have to use some techniques such as Quick-Start to find out the ssthresh and cwnd and also we have to also find at what time we should set these values of ssthresh and cwnd.

References

1. M. Allman, V. Paxson, and W. Stevens. TCP congestion control. RFC 2581, IETF, April 1999.
2. E. Blanton, M. Allman, K. Fall, and L. Wang. A conservative selective acknowledgment (SACK)-based loss recovery algorithm for TCP. RFC 3517, IETF, April 2003.
3. R. Chakravorty, P. Vidales, K. Subramanian, I. Pratt, and J. Crowcroft. Performance issues with vertical handovers: Experiences from gprs cellular and wlan hot-spots integration. In *Proceedings of the IEEE Pervasive Communications and Computing Conference, (IEEE PerCom 2004)*, March 2004.
4. A. Gurtov and J. Korhonen. Effect of vertical handovers on performance of tcp-friendly rate control. *ACM Mobile Computing and Communications Review*, 8(3):73–87, July 2004.
5. W. Hansmann and M. Frank. On things to happen during a tcp handover. In *Proceedings of the 28th Annual Conference on Local Computer Networks, LCN'03)*, October 2003.
6. H. Huang and J. Cai. Improving tcp performance during soft vertical handoff. In *Proceedings of the 19th International Conference on Advanced Information Networking and Applications (AINA'05)*, volume 2, pages 329–332, March 2005.
7. D. Johnson, C. Perkins, and J. Arkko. Mobility support in ipv6. RFC 3775, IETF, June 2004.
8. R. Ludwig and A. Gurtov. The eifel response algorithm for TCP. RFC 4015, IETF, February 2005.
9. R. Ludwig and M. Meyer. The eifel detection algorithm for TCP. RFC 3522, IETF, April 2003.
10. The network simulator, ns-2.
11. C. Perkins. IP mobility support for IPv4. RFC 3344, IETF, August 2002.
12. P. Sarolahti and M. Kojo. Forward rto-recovery (F-RTO): An algorithm for detecting spurious retransmission timeouts with TCP and SCTP - draft-ietf-tcpm-frto-02.txt. Internet draft, IETF, February 2005.
13. Mark Stemm and Randy H. Katz. Vertical handoffs in wireless overlay networks. *Mobile Networks and Applications*, 3(4):335–350, 1998.
14. Y. Swami, K. Le, and W. Eddy. Lightweight mobility detection and response (LMDR) algorithm for TCP-draft-swami-tcp-lmdr-05.txt. Internet draft, IETF, November 2004.