

***Improving the Performance of Reliable
Transport Protocols in Mobile Computing
Environments***

Ramon Caceres and Liviu Iftode

Presented by Venkata N. Padmanabhan (padmanab@cs)

CS294-7 Lecture

March 22, 1996

Overview

- Explored the effect of mobility (handoffs across cells) on TCP performance
 - one of the early investigations of this problem
- Testbed
 - base stations and stationary hosts connected to Ethernet
 - 2 Mbps WaveLAN connecting mobile hosts to the BSs
 - 50 MHz, 486-based hosts running Mach 3.0 + Unix server
 - 4.3 BSD-Tahoe version of TCP, Columbia Mobile IP
- Methodology
 - data transfer between MH and SH with periodic handoffs
 - overlapping/non-overlapping cells
 - zero/non-zero rendezvous delay
 - beaconing frequency: once per second

Performance Effects of Handoffs

- Metrics
 - throughput for bulk-transfer
 - interactive response latency (just after a handoff)
- Throughput degradation
 - overlapping cells: 6%
 - non-overlapping cells + zero rendezvous delay: 12%
 - non-overlapping cells + 1 sec. rendezvous delay: 31%
- Reason:
 - long time to stabilize routes: 0.05 sec. to update MH's route + 0.10 sec. to update old BS's route = 0.15 sec.
 - packets get dropped (up to a full window's worth)
 - wait till TCP sender's retransmission timer expires: 0.8 sec
 - slow start kicks in throttling transmissions

Alleviating the Effects of Handoffs

- Buffering and retransmission of packets after handoff
 - can not guarantee no loss
- More accurate TCP retransmission timers
 - might lead to excessive retransmissions/backoffs
- Fast retransmissions
 - modern TCP senders retransmit a packet after the third repetition of the corresponding ack
 - deliberately invoke fast retransmissions after handoff
 - advantage: no support needed from base stations
 - latency: 0.2 sec. instead of 0.8 sec.
 - throughput degradation:
 - zero rendezvous delay: 7% instead of 12%
 - 1 sec rendezvous delay: 14% instead of 31%

Critique

- Long latency for making routing changes
 - an artifact of microkernel OS?
 - context switches instead of system calls
- Side-effect of fast retransmissions: window size is halved
 - might reduce throughput significantly when the connection traverses a WAN
 - but gently probing network in new cell might be the right thing to do
- Need to invoke fast retransmission on all open TCP connections
 - difficult to determine which ones suffered packet losses