

The Effects of Mobility on Reliable Transport Protocols

Ramón Cáceres

Matsushita Information Technology Lab
2 Research Way, Third Floor
Princeton, NJ 08540, USA
ramon@mitl.research.panasonic.com

Liviu Iftode

Department of Computer Science
Princeton University
Princeton, NJ 08544, USA
liv@cs.princeton.edu

Abstract

We explore the effects of host motion on the performance of active transport-level connections. Motion causes increased delays and packet losses while the network learns how to route data to a host's new location. Reliable transport protocols interpret these delays and losses as signs of network congestion. They consequently throttle their transmissions, further degrading performance. We quantify this degradation through measurements of protocol behavior in a wireless networking testbed. We show how current retransmission policies introduce unacceptably long pauses in communication (800 milliseconds and longer), and propose an end-to-end fast retransmission scheme that can reduce these pauses to levels more suitable for human interaction (200 milliseconds). Our work demonstrates that reliable transport protocols must be made aware of mobility, and suggests how to adapt these protocols to mobile computing environments.

1 Introduction

Reliable transport protocols have been tuned for networks composed of wired links and stationary hosts. They adapt to prevailing end-to-end delay conditions throughout the life of a connection, and interpret unexpected increases in delay as packet losses caused by congestion. In response to perceived losses, protocols like the Transmission Control Protocol (TCP) [9] aggressively slow their transmission to allow the network to recover. These congestion control policies have proven beneficial in improving the overall performance of networks like the Internet. The premise underlying these policies, that packet losses are largely due to congestion, is correct for existing networks.

Future networks, however, will include wireless links and mobile hosts. In particular, there will be local area networks composed of wireless cells of only a few meters in diameter. Such *picocellular* networks are

desirable for three important reasons: they offer high aggregate bandwidth, they require low power from mobile transceivers, and they provide accurate location information. Users in picocellular environments will often carry hosts across cell boundaries without prior warning and in the midst of data transfers.

Transport-level connections will thus encounter types of delay and loss that are unrelated to congestion. First, communication may pause while the hand-off between cells completes and packets can again be routed to and from the mobile host. Second, packets may be lost due to futile transmissions over the wireless network when a mobile host moves out of reach of other transceivers, especially in networks with little or no overlap between cells. Third, packets may be lost due to the relatively frequent transmission errors suffered by wireless links. Some performance degradation due to these delays and losses is unavoidable.

These events also trigger congestion control procedures that further degrade performance. In particular, TCP continually measures how long acknowledgements take to return. It maintains a running average of this delay and an estimate of the expected deviation in delay from the average. If the current delay is longer than the average by more than twice the expected deviation, TCP assumes that the packet was lost. In response, TCP retransmits the lost packet and initiates congestion control procedures to give the network a chance to recover [7]. First, TCP drops the transmission window size to reduce the amount of data in transit through the network. Second, it activates the slow-start algorithm to restrict the rate at which the window grows to previous levels. Third, it resets the retransmission timer to a backoff interval that doubles with each consecutive timeout.

When motion is mistaken for congestion, these procedures result in significant throughput losses and unacceptable interactive delays for active connections. The degradation is readily apparent, for example,

to users of emerging ubiquitous computing environments [3].

This paper quantifies the effects of motion on throughput and delay, identifies the factors that contribute to the loss of performance, and suggests an end-to-end approach for alleviating the problem. It shows how waits for TCP's retransmission timeouts cause pauses in communication that last 0.8 seconds and longer after each cell crossing. Other researchers have called attention to the long pauses caused by TCP's exponential backoff policy [2, 6, 8], but to our knowledge this is the first systematic treatment of this problem. This paper also describes how using TCP's fast retransmission mechanism can reduce these pauses to 0.2 seconds. We focus on TCP because it is the most widely used reliable transport protocol and will be used in at least the first generation of mobile computing environments. Furthermore, lessons learned from TCP are applicable to other reliable transport protocols that must deal with both mobility and congestion.

2 Wireless networking testbed

We explore the effects of mobility through measurements of transport protocol behavior in a wireless networking testbed. The testbed consists of mobile hosts (MH), mobility support stations (MSS), and stationary hosts (SH) deployed in an ordinary office environment. Mobile hosts connect to a 2-Mbit/second WaveLAN local area network, a radio-based wireless product from NCR. Stationary hosts connect to a 10-Mbit/second Ethernet local area network. Mobility support stations connect to both networks. Figure 1 shows the minimum testbed configuration.

All hosts and support stations are IBM PC-compatible machines equipped with 50-MHz i486 processors, 330-Mbyte hard disks, 16 Mbytes of memory, and the necessary network interface hardware. They run the 4.3BSD-Tahoe version of TCP from the University of California at Berkeley, Mobile IP software from Columbia University [6], and the Mach 3.0 microkernel and Unix server (MK77/UX37) from Carnegie Mellon University [1].

2.1 Cellular handoff procedures

Each MSS defines one cell and is responsible for the MHs in its cell. It acts as the default gateway for those MHs, routing packets that originate in an MH from the wireless to the wired part of the network. Similarly, it forwards packets destined to an MH from the wired to the wireless part of the network.

MHs and MSSs collaborate to perform handoffs

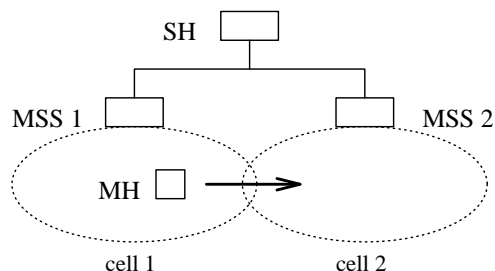


Figure 1: Wireless networking testbed

between cells. MSSs make their presence known by broadcasting periodic beacons over the wireless network. An MH decides to switch cells when it receives a beacon from a new MSS with a stronger wireless signal than the beacon from the old MSS, or when it receives the first beacon from a new MSS after failing to receive beacons from the old MSS.

To switch cells the MH sends a greeting packet to the new MSS, and changes its own routing tables to make the new MSS its default gateway. It also notifies the new MSS of the identity of the old MSS. The new MSS acknowledges the greeting to the MH, adds the MH to the list of MHs for which the new MSS is responsible, and begins to route the MH's packets accordingly. The new MSS also informs the old MSS that the host has moved and can be reached through the new MSS. The old MSS then adjusts its routing tables in order to forward to the new MSS any packets that arrive for the MH, and acknowledges the handoff to the new MSS. Finally, the new MSS acknowledges the completion of the handoff to the MH. Further details of this protocol are found in [6].

2.2 Methodology

In our experiments, we initiate a reliable data transfer over a TCP connection between an MH and an SH, we cause the MH to cross cell boundaries while the connection is active, and we measure the performance of the connection.

We simulate motion across cell boundaries in software. The MH in our testbed is always in range of both MSSs, but we modified the Mobile IP software on the MH to ignore beacons from all but one MSS. After the MH spends a specified number of beaconing periods in that MSS's cell, the modified software listens for a beacon from the other MSS in order to initiate handoff procedures with the new MSS.

An important benefit of simulating motion in software is that it lets us study networks with overlapping cells as well as networks with non-overlapping cells. When adjacent cells overlap and an MH is in the re-

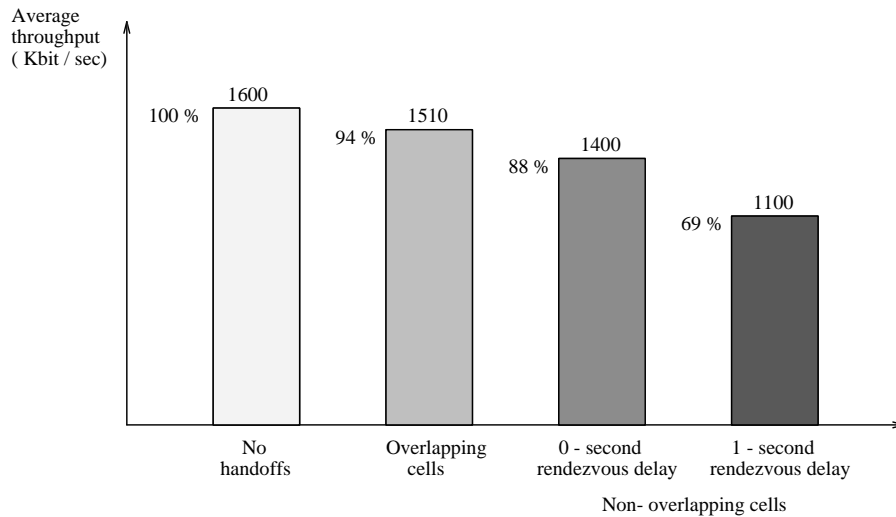


Figure 2: Loss of throughput due to host motion

gion of overlap, packets can continue to flow between the MH and the old MSS while the handoff to the new MSS is in progress. When cells do not overlap, there is an unavoidable pause in network-level communication while the MH is out of reach from the old MSS and the handoff to the new MSS has not yet completed. The testbed allows us to explore the full range of handoff scenarios, from the case when the MH is in contact with both MSSs throughout the handoff, to the case when the MH cannot communicate with any MSS for an arbitrary interval of time after it leaves the old cell.

Another benefit of simulating motion in software is that it gives us precise control over the instant when handoffs begin. Under normal circumstances, handoffs begin at undetermined times based on the time remaining in a cell's beaconing period when a host enters the cell, or on the relative strengths of two wireless signals. Our testbed makes this process deterministic and therefore allows us to reliably reproduce test conditions. Finally, simulating motion in software eliminates the need to physically move test machines.

3 The effects of motion

We ran a number of experiments in the manner described above. We found that throughput dropped significantly in the presence of motion. We then analyzed the problem in more detail to determine the causes of the performance loss. We tracked the TCP sequence number and window size over the lifetime of a connection to determine how TCP behaved during handoffs. We also traced TCP and Mobile IP packets during the course of each handoff to determine if any

packets were lost and why. This section presents our results.

Due to space limitations, we only present results for the case where data packets flow from the MH to the SH and acknowledgement packets flow from the SH to the MH. However, we also ran our experiments for the opposite case, with very similar results. We summarize our results for both cases at the end of Section 4.

3.1 Loss of throughput

Figure 2 shows the average application-level throughput achieved when transferring 4 Mbytes of data between an MH and a SH. From left to right, the vertical bars represent the throughput obtained under four scenarios:

- The MH does not move.
- The MH moves between overlapping cells.
- The MH moves between non-overlapping cells and receives a beacon from the new MSS at the instant it leaves the old cell (0-second rendezvous delay).
- The MH moves between non-overlapping cells and receives a beacon from the new MSS one second after leaving the old cell (1-second rendezvous delay).

In the scenarios that involve motion, the beaconing period is 1 second and the MH switches cells every 8 beaconing periods. These parameters were chosen to allow TCP connections to attain maximum throughput between handoffs while also allowing us to observe multiple handoffs during a single data transfer.

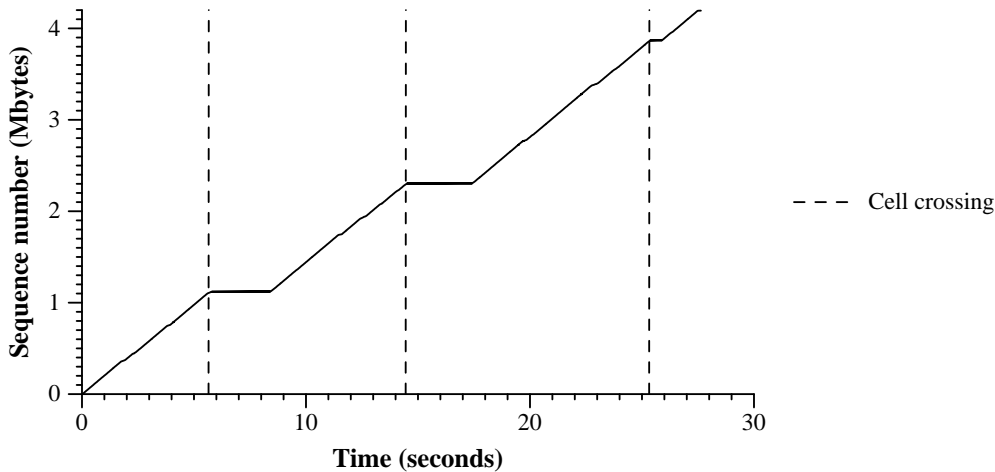


Figure 3: Behavior of TCP sequence number in response to cell boundary crossings

We believe these four scenarios show a complete and fair picture of the problems introduced by host motion. We use the no-motion scenario as a base for comparison. The motion scenario with overlapping cells represents the best handoff performance possible with our hardware and software. It is realizable in a real network only if overlap regions are large enough and hosts move slowly enough for handoff operations to complete while a moving host is still in the overlap region. The scenario with zero rendezvous delay represents the minimum network-level interruption introduced by non-overlapping cell handoffs. It is realizable only if the MH does not have to wait for a beacon before it can communicate with the new MSS, for example in a network where MSSs announce their presence by means of a continuous signal. Finally, the scenario with a 1-second rendezvous delay shows what happens as the length of network-level interruptions increases. It is the most realistic scenario when a periodic beaconing scheme is used, since an MH will in general wait up to a full beaconing period before it receives a beacon from the new MSS.

As shown in Figure 2, throughput degrades substantially in the presence of motion across non-overlapping cells. In the overlapping cell scenario, throughput degrades only slightly, by 6%. In the non-overlapping cell scenario with zero rendezvous delay, throughput drops by 12% even though only 3 handoffs occur in the roughly 24-second lifetime of the connection. Throughput drops much further with a 1-second rendezvous delay, by 31% with 3 handoffs in roughly 29 seconds.

In the rest of this section we study the causes of this performance degradation with increasing detail. We concentrate on single handoffs to eliminate from our results any dependencies on the parameters of the throughput test discussed above (4 Mbytes of data with handoffs every 8 seconds). Our results will thus apply to all cell handoffs in each motion scenario.

3.2 Pauses in communication

Figure 3 shows how the TCP sequence number behaves over the life of a connection. In this example, the MH moves between non-overlapping cells with a 1-second rendezvous delay. As shown, the sequence number ceases to advance for 2.8 seconds after the first two cell crossings, and for 0.8 seconds after the last crossing. A 2.8-second pause is typical of a 1-second rendezvous delay, while a 0.8-second pause is more typical of a 0-second rendezvous delay. During these pauses, TCP transmits no new data and transport-level communication comes to a halt.

The effect is also visible in Figure 4, which graphs the TCP congestion window over the life of the same connection. The congestion window is an upper bound on the transmission window, which in turn controls how much unacknowledged data a TCP connection can have in transit over the network. As shown, the congestion window stops growing with every cell crossing. Some time after the crossing, the window shrinks to its minimum value and eventually begins to grow again. The intervals between when the window stops growing and when it begins to grow again correspond to the 2.8-second and 0.8-second pauses in communication noted above.

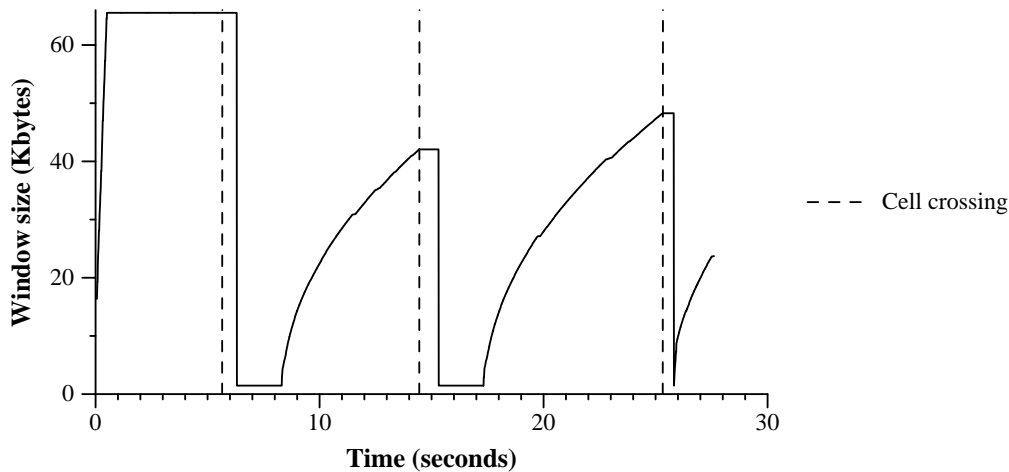


Figure 4: Behavior of TCP congestion window in response to cell boundary crossings

3.3 Packet losses

The long pauses in communication are caused by TCP's response to packet losses. Losses occur due to routing inconsistencies during non-overlapping cell handoffs. Consider the route from the MH to the SH. When the MH leaves a cell without warning, its routing tables continue to point to the old MSS as the default gateway. The MH cannot change its routing tables until a beacon arrives from the new MSS. Until then, the MH continues to send packets destined for the SH directly to the old MSS. These packets are lost because the MH can no longer reach the old MSS through the wireless interface.

Inconsistencies persist longer with the route from the SH to the MH. The old MSS does not know that the MH has left the cell until an explicit notification arrives from the new MSS, which cannot send the notification before it receives a greeting from the MH. Until the old MSS learns of the MH's motion, it continues to route packets directly to the MH. These packets are also lost because the old MSS can no longer reach the MH. Any other parts of the network involved in the handoff also wait for explicit notification that the MH has moved before they change their routing tables to point away from the old MSS to the new MSS.

Figure 5 shows what happens during one handoff in the case of zero rendezvous delay. Although the beacon from the new MSS arrives concurrently with the cell crossing, the MH's routing tables do not point to the new MSS until 0.05 seconds after the cell crossing. Similarly, the old MSS's routing tables do not point to the new MSS until 0.15 seconds after the cell crossing. Although the system overhead implicit in these figures can be reduced through careful implementa-

tion, handoff latency cannot be altogether eliminated because at least two packet exchanges are needed to notify both the new MSS and the old MSS that the MH has changed cells. Because these packets incur unavoidable propagation delays, there will always be a window of opportunity during which both data and acknowledgement packets can be routed to unreachable wireless transceivers.

An active TCP connection thus loses up to a full window's worth of packets and related acknowledgements during each handoff. Once the transmission window fills, communication stops until the retransmission timer expires. When the timeout occurs, TCP retransmits the earliest unacknowledged packet, doubles the retransmission interval, and resets the timer. If the handoff is not yet complete when the timeout occurs, the retransmitted packet is also lost and TCP waits for yet another timeout before retransmitting. A single timeout is typical of zero rendezvous delay, as shown on Figure 5. Two consecutive timeouts are typical of a 1-second rendezvous delay, as shown on Figure 6.

It is evident how waits for retransmission timeouts freeze transport-level communication for 0.8 seconds or more with each cell crossing across non-overlapping cells, and are responsible for a large part of the throughput losses reported earlier. In contrast, handoffs between overlapping cells do not cause the same long pauses in communication because the implementation of overlapping cells in our testbed insures that no packets are lost during those handoffs. The slight throughput losses reported earlier for the overlapping cell scenario are due only to encapsulation and forwarding delays during handoffs.

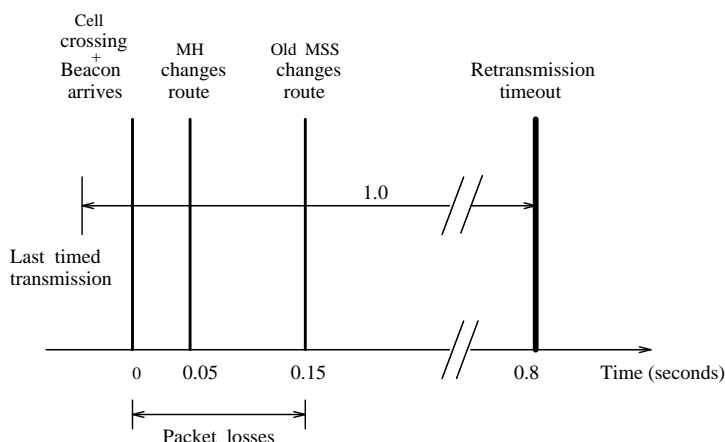


Figure 5: Handoff latency and related packet losses with a 0-second rendezvous delay

3.4 Slow recovery

As shown in Figure 4, the congestion window does not immediately return to its previous level after a cell crossing. Rather, it rises gradually once transport-level communication resumes. TCP's slow-start algorithm [7] is responsible for this behavior. The slow recovery after each handoff contributes to the loss of throughput discussed earlier, but only moderately. Our measurements show that the algorithm throttles transmissions for approximately 1 second after communication resumes. At that point the connection again reaches the maximum throughput supported by the wireless network (1.6 Mbit/second), and the congestion window ceases to affect performance.

3.5 Unacceptable interactive response

Interactive delays are a concern in addition to throughput. Studies of human factors indicate that people perceive interactive response to be "bad" if it takes longer than 100-200 milliseconds [11]. As discussed above and shown in Figures 3, 4, 5, and 6, transport-level communication comes to a halt for 800 milliseconds or longer after non-overlapping cell crossings. Furthermore, these pauses grow exponentially with growing rendezvous delays due to TCP's exponential retransmission backoff policy. In interactive applications that use TCP for reliable data transport, user inputs and their responses will be unable to travel between mobile hosts and remote servers during these pauses.

Although users may not always want to interact with their computers while moving, there will certainly be times when they will want to do so soon after stopping. Our results show that pauses will persist from 650 milliseconds to several seconds after a

host enters a new cell and the handoff completes. Motion will thus lead to unacceptable interactive response unless we solve the problems presented in this section.

4 Alleviating the effects of motion

Our results demonstrate that we must improve the performance of reliable transport communication in mobile computing environments. We clearly should strive for smooth handoffs in order to minimize packet losses during cell crossings. As we have shown with our testbed, engineering enough overlap between cells to insure that handoffs complete before the MH loses contact with the old MSS is one way to achieve this goal. However, there are compelling reasons to build networks with little or no overlap between small cells:

- They offer high aggregate bandwidth because they can use the same portions of the electromagnetic spectrum in nearby cells. Bandwidth is scarce in wireless networks.
- They support low-powered mobile transceivers because signals need only reach short distances. Mobile computers have stringent power consumption requirements.
- They provide accurate location information because cells are small and sharply-defined. Location information adds important functionality to distributed systems.

It is therefore worthwhile to investigate ways to alleviate the effects of motion across non-overlapping cells.

4.1 Fast retransmissions

An attractive end-to-end solution [10] is for the transport protocol to resume communication as soon

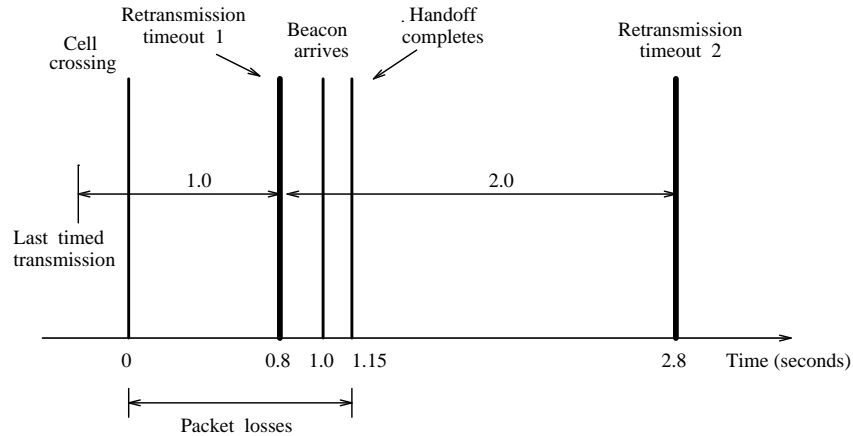


Figure 6: Handoff latency and related packet losses with a 1-second rendezvous delay

as handoffs complete, thereby avoiding the long waits for retransmission timeouts. Modern TCP implementations already perform similar *fast retransmissions* when a transmitter receives duplicate acknowledgments from a remote receiver. When activated, the fast retransmission procedure immediately retransmits the earliest unacknowledged packet. It also drops the transmission window and initiates the slow-start algorithm.

We made modest changes to the TCP and Mobile IP software in our testbed to invoke the existing fast retransmission procedure as soon as routes become consistent following a cell crossing. First, the Mobile IP software on the MH signals the TCP software on the MH when a greeting acknowledgment arrives from the new MSS. Second, the TCP transmitter on the MH invokes the fast retransmission procedure when it receives such a signal. The signal is delivered through normal shared-memory channels between co-located TCP and IP software.

An additional communication step is necessary to inform the TCP software on the SH of the events occurring at the other end of the connection. First, the Mobile IP software on the MH notifies the TCP software on the MH of the completion of the handoff, as described above. Second, the TCP software on the MH sends to the SH an acknowledgement packet containing the sequence number of the last data packet successfully received by the MH. Third, the TCP software on the SH invokes the fast retransmission procedure when it receives such an acknowledgement. The acknowledgement travels from the MH to the SH through normal IP routes as an ordinary TCP packet.

The fast retransmission approach has three important advantages:

- It requires only minimal changes to software on the end hosts. It changes Mobile IP only to propagate available handoff information one layer up in the protocol hierarchy. It changes TCP only to invoke the existing fast retransmission procedure under two new conditions.
- It does not depend on special support from the network, including mobility support stations or other intermediate routers. It therefore does not depend on any one mobile networking environment and will work over an internetwork.
- It follows established congestion control policies by closing the transmission window and using the slow-start algorithm after the initial retransmission. It thus avoids congesting the cell the MH has just entered as well as the rest of the network.

4.2 Summary of results

We measured the effectiveness of the fast retransmission approach in our testbed. Figure 7 shows the pauses in transport-level communication caused by motion across non-overlapping cell boundaries, together with the improvements gained by applying the fast retransmission procedure. As shown, when the transmitter resides on the MH, fast retransmissions reduce these pauses from 0.8 to 0.2 seconds for a 0-second rendezvous delay, and from 2.8 to 1.2 seconds for a 1-second rendezvous delay.

Figure 7 also shows our results for the case when the TCP transmitter resides on the SH, where pauses drop to from 0.6 to 0.3 and from 2.6 to 1.3 seconds for 0- and 1-second rendezvous delays, respectively. Pauses before the improvements are shorter when the transmitter is on the SH (e.g., 0.6 vs. 0.8 seconds)

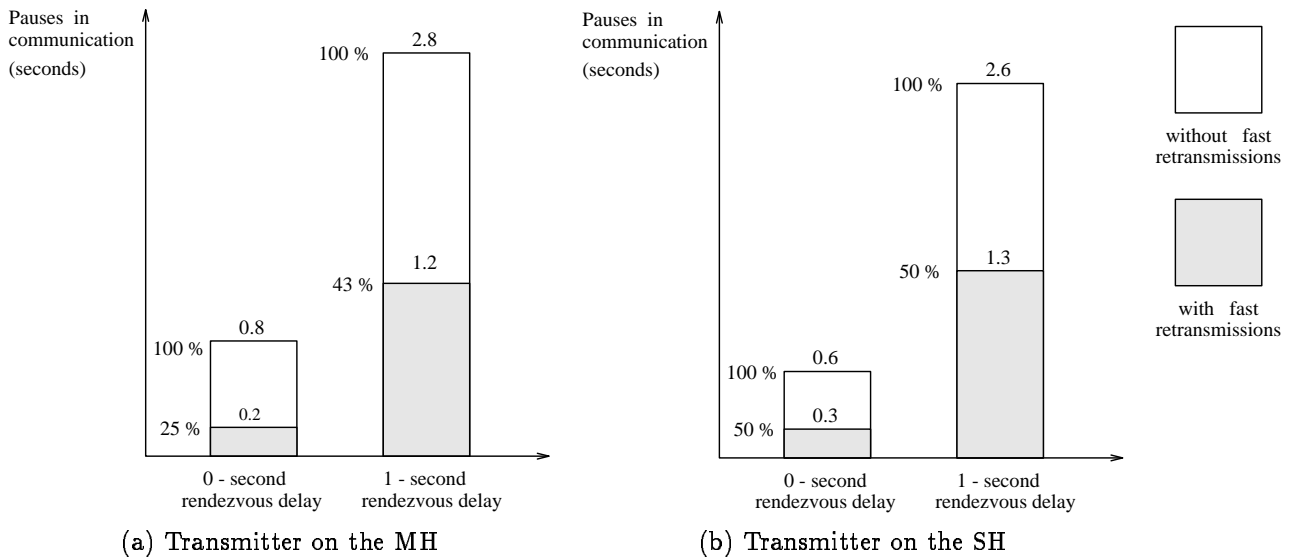


Figure 7: Improvements in latency due to fast retransmissions

because data packets incur added propagation delay before they are lost. Effectively, lost packets are sent earlier before the cell crossing, and thus retransmission timeouts occur earlier after the crossing. Pauses after the improvements are longer when the transmitter is on the SH (e.g., 0.3 vs. 0.2 seconds) because the fast retransmission must wait for an acknowledgement packet to travel between the MH and the SH after the handoff completes.

The fast retransmission scheme thus succeeds in reducing interactive delays to 200-300 milliseconds beyond the rendezvous. Reducing handoff latency through careful implementation would further reduce this remaining delay. The Mobile IP software in our testbed is an early example of support for mobile inter-networking. It provides important functionality, but was not written with fast handoffs in mind. For example, it incurs substantial system overhead by employing user-level daemons to process beacons and other control packets during handoffs. A careful implementation of handoffs combined with the fast retransmission strategy should in all cases bring the pauses in communication to 200 milliseconds or less after the rendezvous. If users do not attempt to interact with their mobile computers until they stop moving across cell boundaries, interactive delays will then drop to acceptable levels.

We also measured improvements in throughput due to the fast retransmission scheme. For example, for the test described in Section 3.1, throughput improves from 1400 to 1490 and from 1100 to 1380 Kbit/second for 0- and 1-second rendezvous delays, respectively.

Some throughput losses remain because fast retransmissions do not entirely eliminate pauses in communication, and because the slow-start algorithm continues to throttle the connection after cell handoffs.

5 Wireless transmission errors

Even in the absence of motion, the WaveLAN network in our testbed suffers from relatively frequent packet losses due to physical transmission errors. A separate measurement study found that WaveLAN exhibited excellent packet capture rates (over 99%) in an indoor environment [5]. However, in our environment, packet loss frequency varies widely even across short distances and depends on such factors as the positions of antennas in a room. Such problems are common in wireless communication because wireless media are vulnerable to ambient noise and multipath interference. Commonly cited bit error rates for radio and infrared links are in the range of 10^{-6} or worse, compared to 10^{-12} or better for fiber optic links.

Wireless transmission errors will also trigger the delay and throughput problems described in Section 3. One possible solution is for the link-layer protocol that controls wireless links to retransmit packets lost on those links and thus hide the losses from higher layers. However, recent research shows that, under certain packet loss conditions, competing retransmission strategies in the link and transport layers can interact to reduce end-to-end throughput while increasing link utilization [4]. Alternative techniques such as selective retransmissions at the transport layer may prove more effective than link-layer retransmissions.

We wanted to isolate the effects of motion across cell boundaries from the effects of transmission errors. We solved the problem by positioning the WaveLAN antennas physically close together in an area relatively free from ambient radiation and multipath problems. Packet losses in the absence of cell crossings then dropped to negligible levels. We also repeated all our experiments using a second Ethernet in place of the WaveLAN and found no fundamental differences in our results. We did not treat transmission errors any further in order to concentrate on handoffs. Nevertheless, the impact of wireless transmission errors on reliable transport protocols warrants further study.

6 Conclusions

Mobility changes important assumptions on which distributed systems have been built. In particular, networks that include wireless links and mobile hosts suffer from delays and packet losses that are unrelated to congestion. Reliable transport protocols react to these delays and losses by abruptly slowing their transmissions, a response that further degrades the performance of active connections. We have identified the factors that contribute to this performance degradation and have quantified their effects in detail. We have shown how waits for retransmission timeouts cause pauses in communication at least 650 milliseconds longer than the underlying network-level interruption. These pauses are readily noticed by interactive users and significantly reduce throughput.

We have also described a fast retransmission scheme that reduces the pauses in communication to 200-300 milliseconds past the moment when network-level communication resumes. The fast retransmission approach is attractive because it calls for only minimal changes to end systems, relies on no special support from the underlying network or intermediate routers, and follows established congestion control procedures. This makes our approach applicable to a large and varied internetwork like the Internet.

Our work makes clear the need to make mobility explicit to the transport level. Our results can be used to adapt current TCP implementations to mobile computing environments. They also apply to other reliable transport protocols that must cope with both congestion and mobility.

Acknowledgements

Dan Duchamp and John Ioannidis provided the Mach 2.5 version of the Mobile IP software. Greg Minshall and the anonymous referees provided helpful comments on earlier drafts of this paper.

References

- [1] M. Accetta, R. Baron, W. Bolosky, D. Golub, R. Rashid, A. Tevanian, and M. Young. Mach: A new kernel foundation for UNIX development. *Proc. of the USENIX 1986 Summer Conference*, July 1986.
- [2] B. R. Badrinath, A. Bakre, T. Imielinski, and R. Marantz. Handling mobile clients: A case for indirect interaction. *Proc. of IEEE WWOS-IV*, October 1993.
- [3] S. Deering and M. Weiser. Private communication. Xerox PARC, October 1993.
- [4] A. DeSimone, M. C. Chuah, and O. C. Yue. Throughput performance of transport-layer protocols over wireless LANs. In *Proc. of Globecom '93*, December 1993.
- [5] D. Duchamp and N. F. Reynolds. Measured performance of a wireless LAN. In *Proc. of the 17th IEEE Conf. on Local Computer Networks*, September 1992.
- [6] J. Ioannidis and G. Q. Maquire Jr. The design and implementation of a mobile internetworking architecture. *Proc. of the USENIX 1993 Winter Conference*, January 1993.
- [7] V. Jacobson. Congestion avoidance and control. *Proc. of ACM SIGCOMM '88*, August 1988.
- [8] A. Myles and D. Skellern. Comparison of mobile host protocols for IP. *Journal of Internetworking Research and Experience*, 4(4), December 1993.
- [9] J. Postel. Transmission control protocol. *Request for Comments 793*, 1981.
- [10] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in system design. *Proc. of the 2nd International Conference on Distributed Computing Systems*, April 1981.
- [11] B. Shneiderman. *Designing the User Interface*. Addison-Wesley, 1987.