Reliable Communication over Wireless Links

Anna Ewerlid
Signals and Systems, Uppsala University
ae@signal.uu.se, http://www.signal.uu.se

ABSTRACT

Within the PCC Wireless IP project, we study how to optimize TCP/IP connections over fading wireless links. An important aspect here is to evaluate transport protocols for reliable (error-free) end-to end communication.

The dominating transport protocol on the Internet is TCP/IP. This protocol has been tuned since 1977 for the primary problem in a global network, congestion and high but static latencies. The Internet is changing to utilize more wireless network segments. Wireless networks introduce new problems such as stochastic packet loss and varying latencies, which create problems for existing TCP flow control algorithms.

In this paper, we describe and compare schemes designed to improve the performance of TCP in such networks and propose our own solution for the problem.

INTRODUCTION

The Transmission Control Protocol (TCP) defines a key service provided by the Internet, namely, reliable stream delivery. TCP provides a full duplex connection between two machines, allowing them to exchange large volumes of data efficiently.

But TCP has been tuned for traditional networks comprising wired links and stationary hosts. It assumes congestion in the network to be the primary cause for packet losses and unusual delays. TCP performs well over such networks by adapting to end-to-end delays and congestion losses. The TCP sender uses the cumulative acknowledgments it receives to determine which packets have reached the receiver, and provides reliability by retransmitting lost packets. The sender identifies the loss of a packet either by the arrival of several duplicate cumulative acknowledgments or by the absence of an acknowledgment for the packet within a timeout interval. To accommodate the varying delays encountered in an Internet environment TCP uses an adaptive retransmission algorithm that monitors delays on each connection and adjusts its timeout parameters accordingly.

TCP reacts to packet losses by initiating congestion control or avoidance mechanisms (e.g., slow start) and by backing off its retransmission timer (Karn s Algorithm [9]). These measures result in a reduction in the load on the intermediate links, thereby controlling the congestion in the network.

Unfortunately, when packets are lost in networks for reasons other than congestion, these measures result in an unnecessary reduction in end-to-end throughput and hence, in suboptimal performance. Communication over wireless links is often characterized by sporadic high bit error rates, and intermittent connectivity due to handoffs. TCP performance in such networks suffers from significant throughput degradation and very high interactive delays [1] [18].

In the present project, we investigate how the performance and quality of service offered by basic TCP connections can be improved. Since TCP is not at all optimized for wireless links, there exist numerous possibilities for improvements of both throughput and latency. However, since TCP is an accepted standard which works well on fiber-optical networks, such modifications should be evaluated critically. In particular,

- Proposed modification might require extensive changes of standards, and widespread distribution throughout the existing Internet. This would delay their acceptance.
- Modifications that improve performance over wireless links could negatively affect fixed links, and perhaps even endanger the stability of the Internet.
- Modifications that improve performance by breaking the end-to-end TCP connections could possibly affect applications and will certainly affect security mechanisms.

When evaluating these and other issues, we need assumptions about the wireless links. Within the Wireless IP project [17], we assume multiple users who share fading wideband links in a coordinated way. The resource sharing and optimization is based on predictions of the channel quality [10] and uses adaptive modulation and scheduling [11], as outlined briefly in the next section.

Thus, information on the link quality is produced and used at lower layers, and it could be used also by higher layers, if desired. Our long-term aim will be to co-optimize the adaptive modulation, scheduling and protocols, so that quality of service and robustness is maintained, while waste of bandwidth is minimized.

Different known approaches to reliable communication over wireless links are outlined and discussed below. We then propose a double split connection scheme that will be the main focus of our investigation. The last section outlines the aims of our comparative simulation study, the first results of which will be presented at the NRS conference.

THE ASSUMED LINK AND MAC LAYER

In the wireless IP project we develop a system that exploits the short-term fading of channels to mobile users. The idea is that the differing fading for different users will enable them to share the available bandwidth. The task is then to optimize both quality of service and system throughput. Our proposed scheme is based on adaptive modulation and scheduling of the IP traffic that adapts to the short-term fading. We use predictions of the future channel quality for all active mobile terminals and develop predictor algorithms for this purpose. The traffic to different users is scheduled for short time intervals ahead, so that their total satisfaction is maximized. Transmissions that require low delay, such as speech, will of course be given high priority [13].

Channels of desired quality will not always be available and channel prediction will sometimes fail. To increase robustness, coding is used in efficient way. We investigate hybrid type-II ARQ schemes in combination with predictive scheduling [11] . Hybrid type-II ARQ will first perform an uncoded transmission and then transmit additional redundant symbols if the previous transmission was unsuccessful.

These algorithms provide links with a stable and prespecified frame error rate and throughput over time intervals characterized by the short-term fading. However, long-term fading and shadow fading will cause variations in these parameters, which have to be countered by the higher-level protocols studied here.

OVERVIEW OF EXISTING APPROACHES

The various existing approaches to improve TCP/IP performance over wireless links can be divided into two groups. The first group of approaches tries to hide all non-congestion related losses from the TCP sender. The idea behind these approaches is that since the problem is local for the wireless part of the link, it must be solved locally (e.g. by split connections, snoop protocol, or by different link-layer solutions). At the other side of the solution spectrum are end-to-end approaches, which are based on making the sender of packets aware of existing wireless hops (Explicit Loss Notification, Selective Acknowledgments) and different propositions for new versions of TCP/IP (e.g. TCP Westwood).

End-to-end solutions

The end-to-end protocols attempt to make the TCP sender handle losses through the use of two techniques. First, some form of selective acknowledgments (SACKs) [15] can be used to allow the sender to recover from multiple packet losses in a window, without resorting to a coarse timeout. Furthermore, they attempt to have the sender distinguish between congestion and other forms of losses by using an Explicit Loss Notification (ELN) [5] mechanism. The main drawback of these solutions is that they require TCP-stack modifications at all endpoints. They therefore require standardization of modifications in TCP followed by widespread acceptance of these changes.

TCP Westwood

A new version of the TCP protocol - TCP Westwood [16] was proposed recently. TCP Westwood enhances the performance of the TCP window congestion control by using an end-to-end measurements of the available bandwidth as feedback. The available bandwidth is estimated at the TCP source by measuring and low-pass filtering the returning rate of acknowledgments. The estimated bandwidth is then used to properly set the congestion window and the slow start threshold after a congestion episode (a timeout or 3 duplicate acknowledgments). The advantage of this approach is that the TCP sender recovers faster after losses, especially over connections with large round trip times. It also improves the performance over wireless links where sporadic losses are due to unreliable links rather than congestion.

Split connections

Split-connection approaches [3] [4] completely hide the wireless link from the sender by terminating the TCP connection at the base station. Such schemes use a separate reliable connection between the base station and the destination host. The second connection can use techniques such as negative or selective acknowledgments, rather than just standard TCP, or some specially designed protocol, to perform well over the wireless link. Advantages of this approach are:

- Each of the connections can be optimized independently;
- Packet losses due to congestion and transmissions errors can be distinguished;
- It allows earlier deployment of enhancements to TCP over wireless links;
- It allows exploitation of different application level enhancements.

However the standard split-connection has several drawbacks. Some of objections that may be raised are:

- It breaks TCP end-to-end semantics. This disables end-to-end usage of IP layer security mechanisms;

- Crashes of the intermediate node (containing TCP state machines in the radio access network) become irrecoverable;
- It introduces extra overhead when moving TCP states between the intermediate nodes at handoff between radio access points.

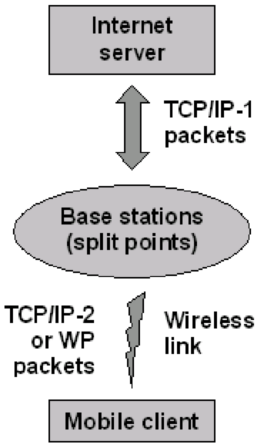

Figure 1: Split-connections

Snoop protocol

The snoop protocol [7] introduces a module, called the snoop agent, at the base station. The agent monitors every packet that passes through the TCP connection in both directions and maintains a cache of the TCP segments sent across the link that have not yet been acknowledged by the receiver. A packet loss is detected by the arrival of a small number of duplicate acknowledgments from the receiver or by a local timeout. The snoop agent retransmits the lost packet if it has it cached and suppresses duplicate acknowledgments. In our classification of the protocols, the snoop protocol is a link-layer protocol that uses the knowledge of the higher-layer transport protocol (TCP).

The main advantage of this approach is that it suppresses duplicate acknowledgments for TCP segments that are lost and retransmitted locally. It thereby avoids unnecessary fast retransmissions and congestion control invocations by the sender. The per-connection state maintained by the snoop agent at the base station is soft, and correct transmission can be maintained without it [6].

Two main concerns about the snoop protocol are that it is not beneficial when the radio link protocol provides in sequence delivery of frames and that it creates problems with IP security.

Link-layer solutions

There have been several proposals for reliable link-layer protocols [8]. The two main classes of techniques employed by these protocols are: error correction, using techniques such as forward error correction (FEC), and retransmission of lost packets in response to automatic repeat request (ARQ) messages. The link-layer protocols for the digital cellular systems in the U.S., both CDMA and TDMA, primarily use ARQ techniques. The TDMA protocol guarantees reliable, in-order delivery of link-layer frames. The CDMA protocol only makes a limited attempt and leaves eventual error recovery to the (reliable) transport layer. Other protocols like the AIRMAIL [2] protocol employ a combination of FEC and ARQ techniques for loss recovery.

The main advantage of employing a link-layer protocol for loss recovery is that it fits naturally into the layered structure of network protocols. The link-layer protocol operates independently of higher-layer protocols and does not maintain any per-connection state.

The main concern about link-layer protocols is that since end-to-end TCP connection passes through the lossy link, the TCP sender may not be fully shielded from wireless losses. This can happen either because of timer interactions between the two layers or more likely because TCP's duplicate acknowledgments cause the sender to fast retransmit even segments that are retransmitted locally.

DOUBLE-SPLIT CONNECTIONS

This is a variant of a split connection with two split points. The first split point should be at a base gateway; TCP packets will here be converted to some special wireless protocol (WP), designed for the wireless link. These packets are transferred via one of the wireless links and are then converted back to TCP at the other split-point, which is located in the wireless device. The sender will thus think that he is talking to the gateway, instead of the ultimate receiver and receiver will think that it receives data from the router located within the wireless terminal [see Figure 2]. Any TCP applications (as identified by port number) that can not handle the semantic breakage caused by a split will not have their packets converted to WP, but will instead just use header compression (e.g. ROHC [19]).

Advantages:

- delays that occur due to fadings over the wireless link can be hidden from the sender, if we choose to design the system that way;
- We can use our knowledge of the channel and of MAC/link layer scheduling algorithms by feeding this information to the wire/wireless gateways;
- The introduction and design of the WP will not depend on possible future changes of TCP (in all servers/clients);
- Methods based on a single split require direct changes in TCP stacks or requires re-linking of ap-

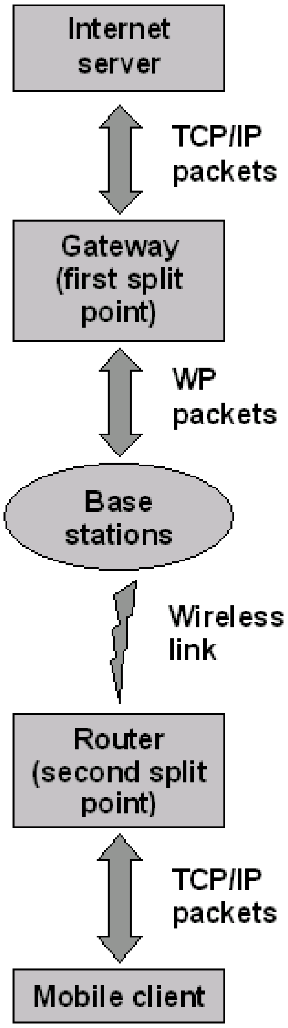

Figure 2: Double-split connections

plications. The double split proposed here is realized using two gateways and does not involve any changes to the TCP end-points.

Potential problems are due primarily to semantic breakage.

It is worth noting that a semantic breakage already exists for typical mobile devices, i.e., masquerading[1]. Masquerading can be done along with the split in the base gateway.

Problems with state migration during handoff can be reduced by locating the first split point at a gateway that controls multiple wireless access points. Migration of state information between gateways will then have to be performed less frequently.

The most serious problem with split connections, and the here proposed double split connections, would be if the gateway acknowledges the successful transmission of a packet that will then in fact never be transmitted without error over the wireless link. Such errors would not be recoverable. However, we believe that this danger can be eliminated by an appropriate design of the WP and by taking the conditions in wireless links into account when controlling the data flow through the gateway. Design of such algorithms are a central part of the present project.

[1] Needed due to a shortage of IP-addresses in IPv4

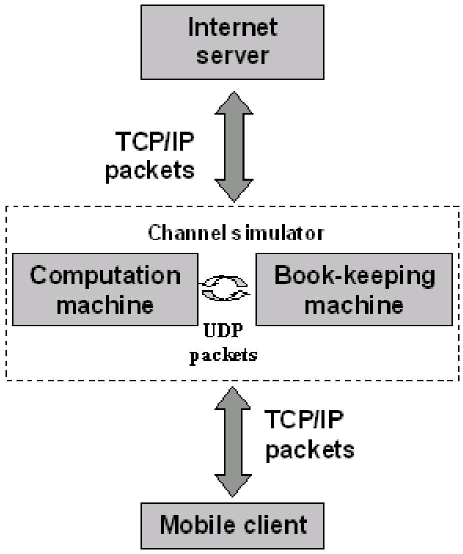

Figure 3: Simulation environment

CONCLUSIONS AND FUTURE WORK

None of the existing solutions for improving the performance of congestion control in TCP/IP over wireless is perfect. Their performance, isolated and in combination, will be evaluated by simulations, assuming usage of forward error corrections and scheduling proposed in [11]. The real-time simulation system [see Figure 3] is used. The mobile client uses a robot-application for downloading web pages from the Internet-server. Packets are trapped at the book-keeping machine, which places them in a queue and sends UDP request to the computation machine. The computation machine simulates the wireless link based on [11] and responds with a delay time (or drop request) for each packet. Within this simulation environment, all methods described above for improving the TCP/IP performance over wireless link are going to be tested. Appropriate performance measures are throughput, latency statistics and average data rates over wireless links. The performance and possible problems with double-split connection will also be investigated.

References

[1] E. Amir, H. Balakrishnan, S. Seshan, R. H. Katz, “Efficient TCP over networks with wire-

less links", *Proc. 5th HoTOS'95*, pp.35-40, 1995.

[2] E. Ayanoglu, S. Paul, T. F. LaPorta, K. K. Sabnani, and R. D. Gitlin, "AIRMAIL: A Link-Layer Protocol for Wireless Networks" *ACM/Baltzer Wireless Networks Journal*, pp.1-47, February 1995.

[3] A. Bakre, B.R. Badrinath, "I-TCP: indirect TCP for mobile hosts", *Proc. IEEE ICDCS'95*, pp. 136-143, 1995.

[4] A. Bakre and B. R. Badrinath, "Handoff and System Support for Indirect TCP/IP", In *Proc. Second Usenix Symp. on Mobile and Location-Independent Computing*, April 1995.

[5] H. Balakrishnan and R. H. Katz "Explicit Loss Notification and Wireless Web Performance", *Proc. IEEE Globecom Internet Mini-Conference*, Sydney, Australia, November 1998.

[6] H. Balakrishnan, V. Padmanabhan, S. Seshan, and R. H. Katz, "A Comparison of Mechanisms for Improving TCP Performance over Wireless Links", *IEEE/ACM Transactions on Networking*, December 1997.

[7] H. Balakrishnan, S. Seshan, E. Amir, R. H. Katz, "Improving TCP/IP performance over wireless networks", *Proc. ACM MobiCom'95*, pp.2-11, 1995.

[8] A. Chockalingam, M. Zorzi, V. Tralli, "Wireless TCP performance with link layer FEC/ARQ", *Proc. IEEE ICC'99*, 1212-1216, 1999.

[9] D. E. Comer *Internetworking with TCP/IP. Principles, protocols and architectures Volume 1*, Prentice Hall, Englewood Cliffs, NJ, 2000.

[10] T. Ekman "Prediction of Mobile Radio Channels", Licentiate Thesis, Signals and Systems, Uppsala University, December 2000.

[11] N. C. Ericsson, S. Falahati, A. Ahlén and A. Svensson, "Hybrid type-II ARQ/AMS supported by Channel Predictive Scheduling in a Multi-User Scenario" *IEEE Vehicular Technology Conference - VTC'2000-Fall*, Boston, September 24-28 2000.

[12] N. C. Ericsson "Adaptive Modulation and Scheduling over Fading Channels" *Globecom99*, Rio de Janeiro, December 5-9, 1999, pp 2668-2672.

[13] N. C. Ericsson, "Scheduling as an Optimization Problem", To appear in *PCC Workshop 2001*, Nynäshamn 2001.

[14] K. Fall, and S. Floyd, "Simulation-based Comparisons of Tahoe, Reno, and SACK TCP", *Computer Communication Review*, vol. 26 no. 3, July 1996, pp.5-21.

[15] S. Floyd, J. Mahdavi, M. Mathis, M. Podolsky, and A. Romanow, "An Extension to the Selective Acknowledgment (SACK) Option for TCP", *RFC 2883*, July 2000.

[16] S. Mascolo, C. Casetti, M. Gerla, S. S. Lee, M. Sanadini, "TCP Westwood: congestion control with faster recovery".In *UCLA CS Technical Report #200017*, 2000

[17] M. Sternad, A. Svensson, A. Ahlén, T. Ottosson, *"PCC wireless IP - optimizing throughput and QoS over fading channels"* To appear in *Nordic Radio Symposium 2001*, Nynäshamn 2001.

[18] G. Xylomenos, G. C. Polyzos, "TCP and UDP performance over a wireless LAN", *Proc. IEEE 18th InfoCom'99*, pp.439-446, 1999.

[19] *RObust Header Compression (ROHC)*, Internet drafts, http://www.dmn.tzi.org/ietf/rohc