"I just want to say one word to you—just one word—'plastics.'"

— From "The Graduate," 1967.

Chapter 8

Conclusions and Future Work

We conclude this dissertation with a summary of our contributions and directions for future work.

8.1 Summary

This thesis identifies three fundamental challenges that degrade the performance of TCP in heterogeneous wireless networks:

1. The preponderance of packet losses due to wireless bit-errors and user mobility.

   TCP performance in many wireless networks suffers because bit-error-induced packet losses, which occur in bursts because of the nature of the wireless channel, are misinterpreted by the TCP sender. TCP attributes these losses to network congestion because of the implicit assumptions made by its congestion control algorithms today. This causes it to reduce its transmission window in response and often causes long timeouts during loss recovery that keep the connection idle. Thus, bit-error losses lead to degraded end-to-end performance. In addition, packet losses that occur due to user mobility cause the TCP sender to remain idle for long periods of time even after the handoff is completed, resulting in unacceptably low throughput.

2. Asymmetric effects and latency variability.

TCP performance in bandwidth-asymmetric networks like wireless cable modem networks suffers because of the low bandwidth ACK channel. TCP over packet radio networks comprised of half-duplex radios does not perform well because of large latency variation and ACK queueing that occurs in the network. In both these networks, the root cause of problems is that TCP performance degrades if the ACK feedback is imperfect, slow or variable.

3. Small transmission windows due to low bandwidths and short connections.

   The bit rate commonly available between distant hosts in the Internet today is often limited and time-varying. This is especially true in many wide-area wireless networks where maximum bandwidths are often orders of magnitude lower than their wired counterparts. In addition, today's dominant TCP application, the World Wide Web, uses short TCP connections because most Web transfers are short. When the available bandwidth is low, so is the transmission window. Small TCP transmission windows prevent the sender from recovering from losses without incurring expensive timeouts that keep the connection idle for long periods of time. As a result, end-to-end throughputs are significantly lower than the already-low maximum channel bandwidth: TCP transfers in these situations yield throughputs that are only a fraction of what is achievable.

While TCP adapts well to network congestion, it does not adequately handle these problems of wireless media. This dissertation analyzed the problems posed by the above challenges, and solved them using a combination of link-layer techniques and modifications and enhancements to TCP at the sender and receiver.

- **The Berkeley Snoop protocol:** This protocol combines a *transport-aware link protocol* and a *link-aware transport protocol,* which improves performance in networks composed of wired and single-hop cellular wireless links. The key idea here is the deployment of a soft-state agent at the base station to perform local loss recovery and shield the TCP sender from the vagaries of the wireless link. The details of the protocol are discussed in Chapter 4.
- **Explicit Loss Notification (ELN):** ELN is a general *framework* by which receivers, base stations, or other elements in the network can inform the TCP sender of losses that occur for reasons other than congestion. When combined with algorithms to distinguish congestion losses

from corruption, this framework provides a powerful way by which TCP senders can separate congestion control from loss recovery and recover from non-congestion-related losses without invoking congestion control. The details of ELN are discussed in Chapters 4 and 5.

- **Asymmetric TCP enhancements:** We first formally defined asymmetry with respect to TCP performance in terms of how the characteristics of one direction affect performance in the other and classified different types of asymmetry based on bandwidth, latency, media-access, and loss-rate asymmetries. We then presented a suite of end-to-end and link-layer algorithms to overcome the adverse effects of asymmetry. These involve techniques to reduce the frequency of ACKs (ACK filtering) and techniques to handle reduced ACK feedback (TCP sender adaptation and ACK reconstruction), and improve performance in both packet radio networks that exhibit variable latencies and ACK queueing, and in bandwidth-asymmetric networks with low-bandwidth ACK paths. These details are discussed in Chapter 6.

- **Enhanced TCP loss recovery:** This is an enhancement to TCP's data-driven loss recovery mechanism that improves performance by reducing the number of timeouts when transmission windows are small. These details are discussed in Chapter 7.

In addition to solving the specific problems in wireless networks discussed above, we also derive a set of general principles and lessons to design transport protocols for heterogeneous networks. These include:

- **Cross-layer protocol optimizations:** In our work, we develop and refine the idea of *cross-layer protocol optimizations* as a design methodology by which information is exposed across multiple layers of the conventional protocol stack to optimize end-to-end performance. For example, the Snoop protocol involves a transport-aware link agent, which takes advantage of the reliability mechanisms of the transport layer and does not generate its own independent messages. It also suppresses certain duplicate TCP ACKs from the sender. Another example is the design of the ACK filter and ACK reconstructor agents to combat the adverse effects of asymmetry; these agents perform TCP-aware connection-specific actions to achieve better performance.

  We believe that this is a powerful technique that can be generalized to many other situations where protocols are being designed with good network performance in mind. This design principle is orthogonal to Integrated Layer Processing (ILP) [Clark90], which proposes that multiple layers be coalesced to achieve efficient end-host implementations; in contrast, cross-layer interactions preserve the principle of layering, but expose certain critical parts of different layers to each other for improved network performance. Chapter 5 discusses this principle in greater detail.

- **Soft-state:** A second general design principle that we use in the design of many of the protocols presented here is the notion of soft-state agents in the network. To enhance performance, agents are often deployed in the network and have state associated with them. However, this state is soft [31], i.e., it is not essential for the correct end-to-end functioning of the connection. It is also periodically refreshed by the arrival of successive packets and ACKs, which implies that even if it gets corrupted, the connection itself is not compromised. This is in direct contrast to protocols that deploy hard-state agents in the network, such as split-connection methods for wireless TCP. Chapters 4, 5, and 6 discussed this principle in detail in the context of specific network agents, such as the Snoop protocol, ELN-enhanced TCP, and link-layer solutions for asymmetric networks.

  In our work, the design of soft-state agents has been guided by the end-to-end argument in systems design [127]. These agents must be viewed as performance optimizers, are not essential for correctness, and do not compromise semantics. Another guiding principle has been the conformance to the IP service model; in particular, we use the fact that packets can be dropped or duplicated in an IP network to great advantage, by suppressing certain types of ACKs, generating local retransmissions, etc.

- **Data-driven loss recovery:** This applies in general to the design of network protocols that attempt to achieve reliability using retransmissions. There are two methods by which reliable transport protocols decide that packets are lost and then retransmit them: using timers (*timer-driven retransmissions*) and using information from other data packets that have reached, usually later ones (*data-driven retransmissions*). Because it is important for transport protocols not to retransmit packets that might still be in transit, retransmission timers are necessarily conservative. Thus, to achieve low recovery latencies and avoid long idle periods that timeouts impose, it is important to make loss recovery as data-driven as possible. We achieve this in the

Snoop protocol by taking advantage of the information provided by TCP cumulative ACKs (and SACK, when available) as well as in the enhanced TCP loss recovery for small window connections.

8.2 Availability

All the protocol and software implementations described in this dissertation are available on-line in source and binary forms.

The Berkeley Snoop protocol is available for BSD/OS 2.1 and BSD/OS 3.0 from

`http://www.cs.berkeley.edu/~hari/software/snoop.html`

Our TCP Selective Acknowledgments implementation for BSD/OS 2.1 and 3.0 is available from

`http://www.cs.berkeley.edu/~hari/software/tcpsack.html`

Implementations of the protocols presented in this thesis are also available in the *ns* network simulator. This can be downloaded from

`http://www-mash.cs.berkeley.edu/ns/`

The `netperf` network performance utility is available at

`ftp://daedalus.cs.berkeley.edu/netperf/np.tar.gz`

8.3 Future Directions

There are several interesting directions for future work based on the work described in this dissertation. Some of these are extensions of our work, while some others are motivated by the more general problem of ubiquitous information access over wireless technologies.

Transport over other wireless networks: In this dissertation, we investigated transport performance over in-room and in-building wireless LANs, the Ricochet packet radio network, and the asymmetric-bandwidth Hybrid wireless cable network. Several other wireless technologies have not been addressed in this thesis, including Cellular Digital Packet Data (CDPD), Global System for Mobility (GSM), and satellite networks. The CDPD network piggybacks packet data on the existing analog cellular network and provides of less than 19.2 Kbps. In addition to wireless bit-errors that are overcome using a link-layer protocol, this network displays asymmetry in media-access. Base stations gain easy access to the channel, while end-stations have to contend with each other to send packets. This leads to variations in latency and degrades performance. In the GSM network, the current link-layer error control protocol (RLP) interacts adversely with TCP, as mentioned in Section 2.5.2.

Satellite networks pose interesting problems because of large communication latencies. In addition, several satellite networks like the Direct Broadcast Satellite (DBS) system display bandwidth asymmetry. While we expect some of our techniques from Chapter 6 to apply to these networks, there are other important problems to overcome, such as unfairness issues and error control. This is an area of active research (e.g., [140], [53]).

Automatic detection of heterogeneity: Throughout our work, we have assumed that we know in advance what networks are present on our path, thus being able to deploy our solutions. For example, by explicitly knowing that the network is bandwidth asymmetric, we know to deploy schemes like ACK filtering for that connection. Ideally, we should deduce that a connection traverses an asymmetric path from end-to-end measurements and observations, and automatically configure the appropriate protocols.

Impact of short-term unfairness in wireless media-access protocols: Common wireless media-access protocols like CSMA/CA provide long-term (asymptotic) fairness, but are unfair in the short-term. That is, over short time-scales, multiple contending stations do not gain equal access to the wireless medium. This affects the performance of higher-layer protocols and applications in many ways. For example, constant-rate audio streams have large jitter values, increasing playback buffer requirements at end hosts. For TCP, ACKs arrive in bursts rather than interleaved with segments in a window, leading to ACK compression and burst transmissions at the sender. We need to develop a framework to quantify short-term unfairness and design distributed contention resolution protocols (e.g., [19]) that have better fairness properties than CSMA/CA.

Networks of Devices (NOD): Until now, research in the areas of mobile computing and wireless networking has focused on providing access to information “anytime” and “anywhere.” This includes our work too. There are several applications, such as remote sensing, surveying, monitoring, and control applications, that are enabled by providing network access to unconventional devices. Examples of these devices include video cassette recorders, temperature sensors, climate monitors, household appliances, remote controls, automobiles, printers, copiers and cameras. Now, in addition to "anytime" and "anywhere" access, these devices enable access "*from anywhere*." Easy access to information located at remote and inconvenient sites is a key ingredient in the vision of ubiquitous information access.

However, there are numerous difficult technical challenges that must be overcome if this vision is to become a reality. We must handle issues of scale, reliability, power consumption, unconventional and harsh communication channels (e.g., wireless links and power lines), and the possible scarcity of computational resources at end devices. Scaling a NOD to numbers on the order of several thousand or more devices is important in many contexts, such as in wide-area sensor networks. Many wireless protocols are notorious for poor performance when the communication channel is shared by more than a small number of hosts. Improving media-access protocols and designing scalable control protocols between NOD entities will result in a system that works well as the number of devices grows.

As an example of the rich class of applications enabled by NODs, consider a large cooperative network of video (or still-image) cameras surveying a remote geographical area. The goal of this system is to provide a reliable and complete a coverage of the remote area, disseminating data in real-time to a remote station. Each camera can change its orientation based on messages sent to it and may also move. Clearly, robust wireless communication is essential for this to work, which implies the need for protocols to combat channel errors and other problems posed by this harsh environment. Most transport protocols today optimize throughput or latency, without considering power consumption in end devices. In this environment, we have to explicitly consider this issue in protocol design. More importantly, we need to design scalable, multicast-based, application-specific control protocols between devices, making sure that message implosion is avoided. Finally, such large amorphous networks must be designed to *self-organize* without any prior structure or hierarchy being imposed on the system.

A fruitful approach to this research goal would be to build some applications like the one described above, and then extract a toolkit that other new applications could use to perform tasks related to scalable control messages, self-organization, and error-free data transport.

Home networking: Another example of a new genre of networks is the class of networked consumer devices in homes and offices. The challenges here are similar to those in NODs (e.g., self-organization), but system security and service location become especially important problems. Furthermore, there are likely to be several types of communication elements in networks at home (wireless, telephone, cable, power lines), making it important to understand and solve performance problems.

Location-dependent applications and services: Little attention has been paid so far to developing novel location-dependent applications for wireless mobile systems, especially on a large scale involving hundreds to thousands of devices. Developing a general infrastructure for applications to exploit knowledge about the geographical or physical locations of users, devices and services in likely to be an area of fertile research in the near future, with the promise of significant commercial impact. We believe that the deployment of a general purpose system infrastructure for location-dependent applications will foster the development of new, compelling applications and accelerate the deployment and use of wireless technologies.