Wireless Networks 8, 289–299, 2002

Optimizing the End-to-End Performance of Reliable Flows over Wireless Links

REINER LUDWIG
Ericsson Research, Herzogenrath, Germany

ALMUDENA KONRAD, ANTHONY D. JOSEPH and RANDY H. KATZ
Computer Science Division, University of California at Berkeley, USA

Abstract. Pure end-to-end error recovery fails as a general solution to optimize throughput when wireless links form parts of the end-to-end path. It can lead to decreased end-to-end throughput, an unfair load on best-effort networks, and a waste of valuable radio resources. Link layer error recovery over wireless links is essential for reliable flows to avoid these problems. We demonstrate this through an analysis of a large set of block erasure traces measured in different real-world radio environments, with both stationary and mobile hosts. Our analysis is based on a case study of the circuit-switched data service implemented in GSM. We show that the throughput on this wireless channel can be increased by using a larger (fixed) frame size for the reliable link layer protocol. This yields an improvement of up to 25% when the channel quality is good and 18% even under poor radio conditions. Our results suggest that adaptive frame length control could further increase the channel throughput. Finally, we discuss link and transport layer error control mechanisms and their interactions with end-to-end congestion control schemes. For reliable flows, we argue in favor of highly persistent error recovery and lossless handover schemes implemented at the link layer.

Keywords: TCP, wireless, flow differentiation, link ARQ persistency

1. Introduction

The Internet is evolving to become *the* communication medium of the future. It will not be long before virtually all people-to-people, people-to-machine, and machine-to-machine communication are carried end-to-end in IP (Internet Protocol) [41] packets. The recent tremendous growth of the Internet in terms of connected hosts is only matched by the similar tremendous growth of cellular telephone subscribers. While most hosts on today's Internet are still wired, the next *big* wave of hosts has yet to hit the Internet. We believe that the predominant Internet access of the future will be wireless. Not only every cellular phone, but every *thing* that communicates will have: (1) an IP protocol stack and (2) a wireless network interface.

It is well known that the performance of reliable transport protocols such as TCP (Transmission Control Protocol) [42] may degrade when the end-to-end path includes wireless links. This is due to non-congestion related packet losses on the wireless link, causing a TCP sender to underestimate its share of bandwidth along the path. However, related work has mostly focused on the problem that wireless links cause for the congestion control scheme used in most implementations of TCP [1]. Employing a link layer error recovery scheme over the wireless link removes this problem. Furthermore, for TCP our previous work shows that problems resulting from competition between end-to-end and link layer error recovery are rare. In [31] we show this for the wireless network examined in this paper. For other wireless networks, related work [5,14,36] comes to the same conclusion. In [32] we provide a general analysis of this subject, and in [34] we reveal reasons for the overly high degree of conservativeness implemented in TCP's retransmission timer [51].

Motivated by this result, we study the impact of link layer frame sizes on application layer throughput and the consumption of radio resources. We then quantify the benefit of link layer error recovery by comparing it against the performance of pure end-to-end error recovery.

The key premise for our analysis is that we assume a model of a network-limited bulk data transfer based on a reliable flow (e.g., based on TCP). In addition, we assume a *flow-adaptive* link layer implementation [30,32]. A flow-adaptive link layer error recovery scheme distinguishes among unreliable and reliable flows to control its retransmission persistency. This ensures that link layer error recovery does not interfere with delay-sensitive (usually unreliable) flows (e.g., based on UDP (User Datagram Protocol) [40]). The advantages of this solution over approaches that require access to transport layer headers in the network (e.g., [2–4,9,13,21,28]), are

- its independence from transport (or higher) layer protocol semantics making it a "non-TCP-specific" solution,
- the possibility of coexistence with network layer encryption as proposed in [27], and
- the fact that no per-flow state needs to be maintained in the network making this solution more scalable.

The concept of flow-adaptive link layer implementations was first introduced in [30]. There we proposed to use the protocol identifier field in the IP header for the purpose of