Telecommun Syst (2010) 44: 69–78
DOI 10.1007/s11235-009-9226-1

Cross-layer TCP with bitmap error recovery scheme in wireless ad hoc networks

Rung-Shiang Cheng · Der-Jiunn Deng · Yueh-Min Huang · Lianfen Huang · Han-Chieh Chao

Published online: 19 November 2009

Abstract Current TCP is not able to distinguish corruption losses from packet loss events. Hence, high transmission errors and varying inherent latency within a wireless network would cause seriously adverse effects to TCP performance. To improve TCP in IEEE 802.11 multi-hop ad hoc wireless networks, this study proposes an error recovery mechanism based on coordination of TCP and IEEE 802.11 MAC protocols. The simulation results confirm that the proposed error recovery approach could provide a more efficient solution for frequent transmission losses, and enable TCP to distinguish between congestion errors and transmission errors, and thus, to respond with proper remedial actions.

R.-S. Cheng (✉)
Department of Computer and Communication, Kun Shan University, No. 949, Dawan Rd., Yongkang City, Tainan County 710, Taiwan, ROC
e-mail: rscheng@mail.ksu.edu.tw

D.-J. Deng
Department of Computer Science and Information Engineering, National Changhua University of Education, Bao-Shan Campus, No. 2, Shi-Da Rd., Changhua City 500, Taiwan, ROC
e-mail: djdeng@cc.ncue.edu.tw

Y.-M. Huang
Department of Engineering Science, National Cheng Kung University, No. 1, University Rd., Tainan City 701, Taiwan, ROC
e-mail: huang@mail.ncku.edu.tw

L. Huang
Department of Communications Engineering, Xiamen University, No. 422, Simingnan Rd., Xiamen City, Fujian Province, 361005, China
e-mail: lfhuang@xmu.edu.cn

H.-C. Chao
Department of Electronic Engineering, National Ilan University, No. 1, Sec. 1, Shen-Lung Rd., I-Lan City 260, Taiwan, ROC
e-mail: hcc@niu.edu.tw

Keywords Multi-hop · Cross-layer design · Bitmap · Ad hoc · IEEE 802.11

1 Introduction

Wireless technologies provide mobile access to networks and eliminate the need of fixed cable infrastructures, thus enabling cost-effective network deployment. In recent years, wireless communication networks have been extensively deployed and are generally specified in accordance with IEEE 802.11 standards. However, a wireless link is commonly characterized by an unpredictable bit-error rate and varying latencies [1], hence, wireless environments pose great challenges when attempting to provide reliable data transmission for transport protocols, such as TCP.

TCP is the most widely applied transport layer protocol for achieving reliable data transfer services over the Internet. As wireless networking gains in popularity, TCP continues to be the leading transport layer protocol. TCP guarantees end-to-end error-free data delivery services, and was originally designed for wired networks with a relatively reliable physical layer, in which packet losses arise primarily because of network congestion. When running TCP over wireless networks, however, packet losses due to network congestion rarely occur, but result primarily from repeated contentions or erratic errors in wireless links.

The current TCP format is not able to distinguish between transmission errors and network congestion. Once packet losses are detected, TCP considers the event as an indication of network congestion and invokes a congestion control mechanism that results in an undesirable reduction of transmission rate. To improve TCP performance in wireless networks, it is important to handle packet losses caused by network congestions and wireless transmission errors differently.

This paper investigates difficulties arising when utilizing the TCP protocol in wireless networks. The performance of data transfer relies on support from the end-to-end transport layer, as well as on the quality of the hop-by-hop communication link. Therefore, the present study attempts to enhance the legacy of IEEE 802.11 MAC and TCP protocols, in such a manner that TCP is rendered capable of differentiating between congestion and corruption losses, using information received from the MAC layer.

The remainder of this paper is organized as follows. Section 2 provides a brief description on the 802.11 standard. Section 3 describes major challenges arising when utilizing TCP in multi-hop wireless networks. Section 4 presents the proposed Bitmap-based error recovery scheme based on a cross-layer design. Section 5 discusses the results of evaluation simulation performances. Finally, Sect. 6 presents brief conclusions.

2 Preliminaries

The IEEE 802.11 standard [2] defines two types of services, namely a contention-free polling-based point coordination function (PCF) and a contention-based distributed coordination function (DCF). PCF is a centralized scheme, while DCF is distributed. In infrastructure-based networks, DCF can operate alone or in conjunction with PCF, however, in ad hoc networks, DCF operates alone. DCF is the basic access method for the 802.11 standard and is based on the conventional carrier sensing multiple access with a collision avoidance (CSMA/CA) scheme. DCF comprises both a basic access method and an optional channel access method based on RTS/CTS exchanges.

In 802.11, priority access to the wireless medium is controlled by the application of an inter-frame space (IFS) time between the frame transmissions. To prevent collisions, the transmitter is obliged to wait for the channel to remain free for a specified interval of time, designated as the distributed inter-frame space (DIFS), before sending a frame. If the medium is currently busy, or becomes busy during this interval, the transmitter defers the frame transmission until it detects a DIFS. At this point, the transmitter selects a random interval, referred to as the backoff time, to determine the moment to commence transmission. The backoff time is an integer number of slots, uniformly chosen from the interval (0, *CW*-1), where *CW* is the backoff window, also referred to as the contention window. The backoff number counts down slot-by-slot, and when it reaches zero, the frame is transmitted.

Due to the half-duplex nature of wireless interfaces, stations in the network are unable to detect a collision simply by listening to their own transmissions. Therefore, an immediate positive acknowledgment technique is employed to confirm the successful reception of a frame. Specifically, having received a frame, the receiver waits for a short inter-frame space (SIFS), and then transmits a positive MAC acknowledgment to the transmitter, confirming that the frame has been correctly received. The SIFS is deliberately assigned a shorter interval than the DIFS in order to assign the receiving station a higher priority than any other stations waiting to make a transmission. The ACK is only transmitted if the frame is received correctly, and hence if the transmitter does not receive an ACK, it assumes that the data frame must have been lost, and therefore, schedules a retransmission. Figure 1 illustrates the basic operations involved in 802.11 DCF.

To alleviate the hidden-station problem [3], 802.11 DCF also provides an optional channel access method using a virtual carrier sensing mechanism based on two special control frames, namely request-to-send (RTS) and clear-to-send (CTS). As shown in Fig. 2, before transmitting a frame, the transmitter transmits an RTS frame asking the receiver if the medium in its vicinity is free. Once the receiver receives this RTS frame, and assuming no transmission interference is present, it waits for the specified SIFS interval, and then sends a CTS frame to the transmitter. Both transmitter and receiver neighbors overhear these frames and consider the medium reserved for the transmission duration.

Fig. 1 Basic access mechanism in DCF

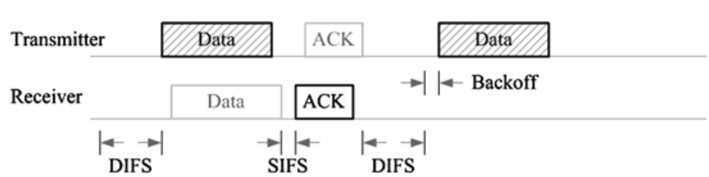

Fig. 2 Virtual carrier sensing mechanism in DCF

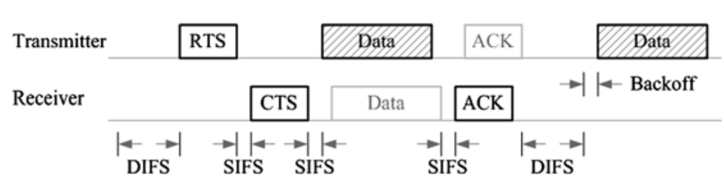

3 Transport layer challenges in wireless networks

TCP is a reliable end-to-end acknowledgment-clocking window based protocol. TCP controls the sending rate using a congestion window parameter [4]. In theory, TCP should be independent of the technology used to implement the underlying infrastructure. However, in practice, the high error rates typical of wireless networks generally cause the backoff mechanism to be inappropriately invoked. Consequently, the utilization of the network bandwidth is severely degraded.

Packet losses in a wired network are mainly caused by buffer overflows at the bottleneck router. However, in a multi-hop wireless network, packet losses due to buffer overflows at intermediate stations rarely occur (unless the station buffer is very small), but result primarily from link-layer contentions or transmission errors [5]. Collision occurrences in a shared channel increase as in-flight packet numbers increase, and hence, a large sized TCP window leads to a higher degree of link-layer contention, due to the half-duplex nature of wireless links, and thus to a higher number of dropped packets.

As discussed in [6] and [7], channel access contentions may occur between different flows passing through the same vicinity, or between different packets within the same flow (e.g. consider the case where the forwarded TCP data competes for the channel with the backward ACK of the previous data), causing TCP transmission rate to fall as a result of frequent invocations of the congestion control mechanism. Moreover, frequent packet exchanges taking place in multi-hop wireless networks exacerbate channel contention problems, resulting in more packets dropped in the MAC layer. This exacerbation causes inappropriate invocation of the TCP back-off mechanism, further degrading network bandwidth utilization.

Several recent studies [8–11] have proposed alleviating the effects of non congestion-related losses in TCP performance over networks with wireless links, by introducing intelligence at the base station using TCP-aware SNOOP mechanisms, or by splitting the entire path into two distinct parts, namely the wired connection and the wireless connection. Under this approach, upon successful transmission of a packet in one connection, an acknowledgement message is sent to the TCP source and the packet is then relayed to the next connection. However, a major disadvantage of this split-connection approach is that the base station is required to maintain information relating to the state of both connections, as well as caching unacknowledged packets for every TCP connection passing through it, causing an overflow problem.

The Explicit Loss Notification (ELN) scheme [8] has been proposed as a means to provide TCP with the ability to differentiate between congestion and wireless losses. However, because lost packets can only be retransmitted after the round-trip time has elapsed, error recovery is slow compared to that achieved by SNOOP mechanisms. The TCP-Feedback and TCP-ELFN [12] schemes aim to improve TCP performance by sending an Explicit Link Failure Notification (ELFN) from the node immediately upstream of the link failure to every TCP connection passing through that link. Although ELFN-based approaches perform far better than conventional TCP schemes in mobile scenarios, such approaches typically provide lower throughputs in static networks [13].

The end-to-end scheme SACK [14], uses a SACK option carried by the return selective ACK to inform the sender of successfully received data. However, the SACK option is not bit-efficient in a contention-based 802.11 DCF environment due to its use of two 32-bit sequence numbers to specify a single SACK block. The resulting limitation of three maximum SACK blocks per ACK segment (when other options, e.g. Timestamp, are also used), is too restrictive for TCP over erratic wireless links [15].

In [16], Mascolo et al. introduced the TCP Westwood scheme. This scheme estimates available bandwidth based on the measured inter-arrival rate of successive ACK packets, and then uses this estimate to set appropriate values of the slow-start threshold and the congestion window size. This approach avoids the default blind-halving of the window size, as applied in conventional TCP (e.g., Reno), and enables TCP Westwood to achieve a high link utilization in the presence of random sporadic losses, typical of wireless links. However, TCP Westwood generally overestimates the available bandwidth because of ACK packet clustering. As shown in [17], in a heterogeneous scenario, TCP Westwood consistently achieves a higher throughput than its fair share. In other words, Westwood improves the performance of the TCP connection, but inevitably introduces a trade-off between its throughput and its friendliness to other TCP implementations.

4 Cross-layer TCP with bitmap error recovery scheme

Reliable data transmission requires that the standard TCP protocol employs an ACK mechanism to confirm successful packet reception. Although the transport protocol is error-tolerant, wireless losses impose additional error recovery requirements. Efficient error recovery is therefore vital for maintaining acceptable application performance.

4.1 Proposed bitmap-based error recovery procedure

The present study applies the selective negative acknowledgement (SNACK) option to improve TCP performance in wireless networks. SNACK is similar to the solution