TCP over IS-707[*]

Yong Bai, Andy T. Ogielski[†] and Gang Wu[‡]

WINLAB, Rutgers University
73 Brett Road, Piscataway, NJ 08854-8060
e-mail: [yongbai, ato, g-wu]@winlab.rutgers.edu

Abstract - Mobile wireless Internet access services face a problem of mitigating the time-varying, correlated frame losses in fading radio channels. The losses may be compounded by detrimental interactions between separate error recovery mechanisms in the link and transport protocol layers. In this work we investigate by simulations the performance of TCP Reno over IS-707, a radio link protocol standard for spread spectrum digital cellular systems, and study the interactions of these protocols in the high loss regime. We show the improvement of TCP goodput by a fast RLP frame error recovery scheme based on using the IS-707 idle frames. We generalize the problem of characterizing the performance of user applications accessing the Internet over wireless data services by investigating the distributions of TCP goodput at multiple time scales, that give more insight than long-time averages. Results are presented for a range of frame error rates and normalized Doppler frequencies using a standard Markov model of a fading radio channel.

I. Introduction

Recent measurements of packet traffic in a typical ISP environment showed that between about 60% to 80% of all packets and bytes are Web TCP traffic [1]. Assuming that similar traffic would be present in wireless Internet access networks poses a challenge to service providers, as TCP has been designed under the assumption that packet losses are caused almost exclusively by network congestion. TCP is a network–friendly adaptive protocol which upon loss detection initiates the congestion avoidance mechanisms [2, 3] that include rate reduction and multiplicative increase of the retransmission timeout.

In contrast, link losses in mobile wireless networks are quite high, and may exceed 10%. These losses, moreover, tend to cluster due to fade correlations and handoffs. Lower-level wireless protocols provide a number of error control methods, such as FEC (forward error correction), ARQ (automatic repeat request), and hybrid FEC/ARQ, to improve communication reliability. In practical design of wireless communication systems, such as CDMA, FEC coding can be implemented in physical layer as specified in IS–95 [4], and ARQ in Radio Link Protocol (RLP) as specified in IS–707 [5] (an extended version of IS–99 [6]). Link layer error control can improve the performance of TCP. However, if RLP and TCP operate independently of one another, their interaction may degrade the TCP goodput, especially when the link protocol employs a partial error recovery algorithm such as IS-707.

We investigate the performance of the most popular TCP variant (Reno) over IS-707 using a wireless Internet access simulation model written using a recently developed high performance Scalable Simulation Framework (SSF) [7]. We show the improvement of TCP goodput by a fast RLP frame error recovery scheme based on controlling the number of idle frames sent after a data frame is sent and the channel is idle. This use of idle frames is compatible with IS-707 (which does not clearly specify their purpose), however, the control of a number of idle frames is not a part of IS-707, and constitutes a proposed extension of the standard.

Ultimately, what matters to the customers of wireless Internet services is the behavior of user applications. To get an insight into the potential effects of radio link losses without going into details of applications, in this paper we look at the distributions of short–time TCP goodput at multiple time scales, for a range of fading severity and loss correlations, that give more information than traditionally used long–time average packet error rates.

Previous research: The performance of TCP over IS-99 RLP for circuit–mode data services was studied in [8] and [9] for a mobile–base station link and perfect feedback channel (error-free NAK transmission). The performance of TCP Tahoe over IS-707 for packet data services in a wireless Internet access scenario was studied in [10], where it was shown that with correlated losses a slow frame recovery as well as unrecovered errors cause successive TCP timeouts that trigger exponential TCP re-

---

[*] Published in the Proceedings of the 10th International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC'99), Osaka, Japan, Sept. 12-15, 1999. Research of Y. Bai and A. T. Ogielski was supported by DARPA Grant N66001-96-C-8530 and by NSF Grant NCR-9527163.

[†] Also at the DIMACS Center, Rutgers University.

[‡] Also at Communications Research Laboratory, MPT, Japan.

transmission timer back-offs. That paper proposed that probing frames can drive fast RLP frame error recovery.

II. IS–707 RLP Frame Error Recovery

TCP/IP packets are fragmented by IS-707 into 24–byte frames (19 bytes of payload) on 9,600 bps links. In the non-transparent mode, IS-707 uses a NAK (negative acknowledgement) selective repeat ARQ protocol to retransmit lost data frames. The receiver does not acknowledge correct RLP data frames. In case of a data frame loss, RLP performs a partial error recovery through a small number of frame retransmissions, and if retransmissions fail, further error recovery is left to higher protocol layers. IS-707 RLP maintains one sequence number (SN) counter $V(S)$ for sender, and two SN counters for receiver: the expected frame SN counter $V(R)$, and the oldest missing frame SN counter $V(N)$.

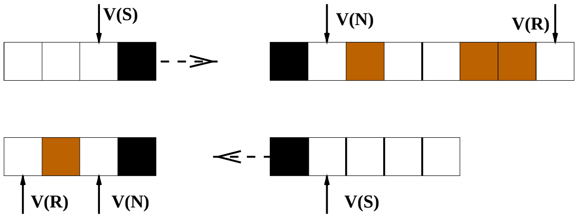

Figure 1: IS-707 sender and receiver SN counters.

The receiver can notice that a frame is missing in either of three ways: at the arrival of a valid data frame (if $SN \geq V(R)$), an idle frame (with $SN = V(S)$), or a NAK control frame (with $SN = V(S)$). The idle frames, without data payload, are transmitted at 1/8 rate of primary traffic RLP frames. At the arrival of a valid data frame or an idle frame, the receiver checks the resequencing buffer for missing frames, and may advance the NAK retransmission timer.

The arrivals of data or NAK frames may be separated by fairly long time intervals. This happens, e.g., when TCP packet arrivals are sporadic as in Web sessions. The dependence on another side to discover a frame loss thus may cause a long delay of frame loss recovery that leads to the TCP retransmission timer expirations. This suggests the use of idle frames for forcing fast frame error recovery. In the IS-95-B and cdma2000 proposals, the radio link is released to other users if the channel is idle for a certain time, so it is desirable to send as few idle frames as possible to achieve the required performance improvement. The appropriate number of idle frames under various channel conditions can be obtained from the simulation results shown in Section III.

III. Simulation Results

The simulation model is shown in Fig. 2, which is similar to the TeD simulation model in [10]. The model has been implemented in SSF [7] and executed using Cooperating Systems Corp. C++ implementation of SSF. In the model, the `Mobile Host` MH is the receiver of a continuous TCP data flow generated by the `Data Source` at Remote Host RH. We use a TCP Reno variant ported to SSF from the `ns-2` simulator [11]. The TCP maximum segment size (MSS) is set to 536 bytes. Each TCP packet from `TCP Src` is delivered to the `Interface` at `Internet Access Point` IAP and fragmented into about 30 IS-707 RLP frames for delivery over a 9600 bps IS-95 [4] physical channel. At the `TCP Sink` side, the `Interface` reassembles the incoming frames passed from the RLP. The packet size of TCP ACKs from the `TCP Sink` is 40 bytes. `Radio Channel` is modeled as a first-order binary Markov process [12]. By choosing different values of frame error rate P_e and normalized Doppler frequency $f_d T$, where f_d is the Doppler frequency and T is the frame length (20 ms for IS-707), we can model fading radio channels with different degree of correlation in the fading process. When $f_d T$ is small, the fading process has long-time correlations (long bursts of frame errors); while for large values of $f_d T$ the successive samples of the radio channel are approximately independent. We essentially ignore the wired transmission details in the `Wireline Link` entity, by choosing the bandwidth of 1.5 Mbs and delay of 200ms.

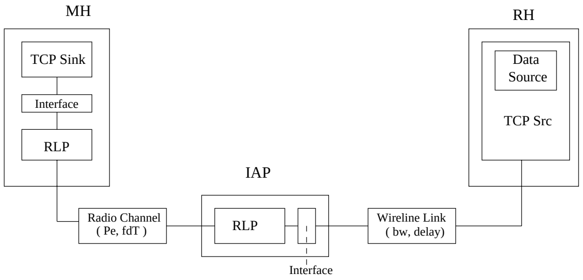

Figure 2: The simulation model.

We define the *average TCP goodput* as the average TCP data throughput normalized to the maximum net link throughput; i.e., the average fraction of maximum TCP payload delivery rate. In IS-707, the maximum net link throughput is 7600 bps at the raw data rate of 9600 bps, and the goodput is upper–bounded by $1 - P_e$.

The simulation data is analyzed as follows: First, we show the long–time average TCP goodput without idle frames in RLP, and illustrate the mechanism of excessive performance degradation at higher frame error rates. Then we demonstrate TCP goodput improvement with idle frames, and show its dependence on the number of idle frames transmitted. Long–time average goodput tells only a small part of the story when the application performance is of concern, as for most applications the short–time variability of TCP goodput may be a more important performance measure. This is illustrated with histograms of short–time TCP goodput measured over time windows of duration of 10, 100, and 1,000 seconds.

Figures 3 and 4 compare the long–time average TCP goodput without idle frames, and with idle frames transmitted during periods when RLP has no data to send, respectively, as a function of average frame error rate (FER) P_e for different values of normalized Doppler f_dT. It is seen that the TCP goodput improves with the use of idle frames, especially in correlated high FER regime.

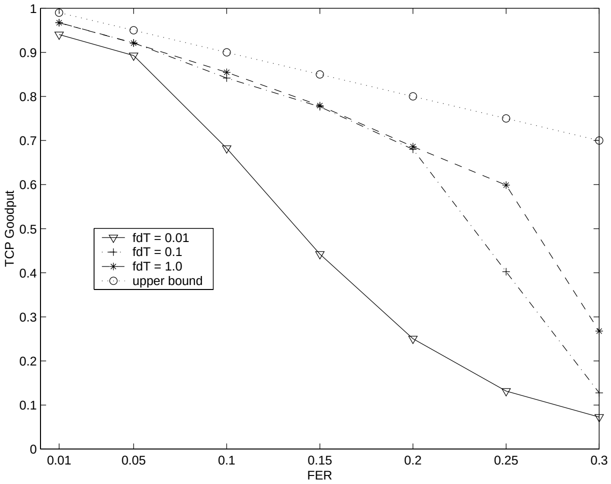

Figure 3: Long–time TCP goodput without idle frames.

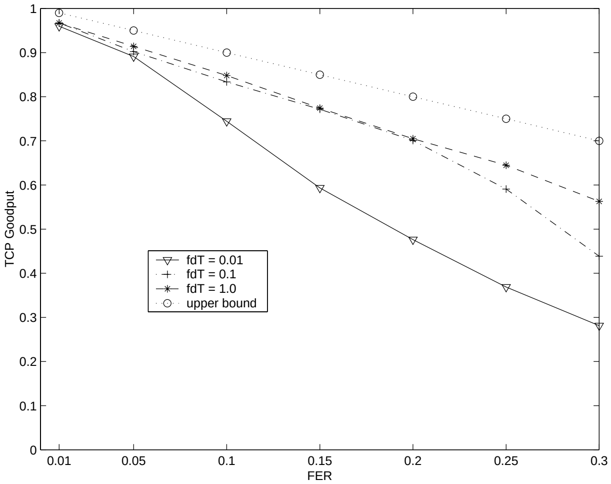

Figure 4: Long–time TCP goodput with idle frames.

The mechanism of the observed TCP goodput improvement with the use of idle frames is illustrated using simulation traces as follows. Figures 5 and 6 show a typical sample path of our simulation process in the correlated high–loss regime $f_dT = 0.01, P_e = 0.3$.

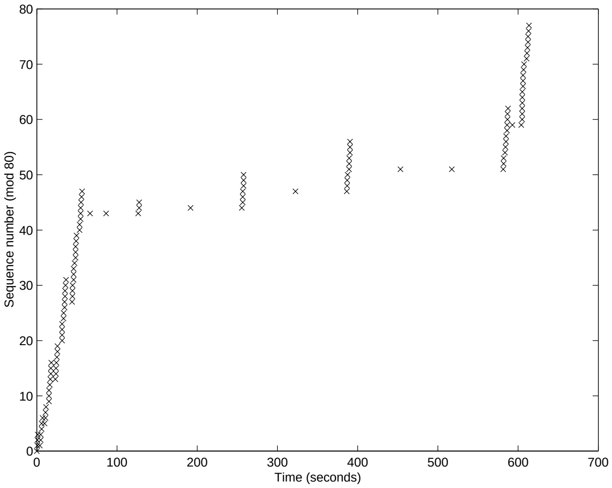

Figure 5: TCP packet sequence number in a simulation sample path without idle frames ($f_dT = 0.01, P_e = 0.3$)

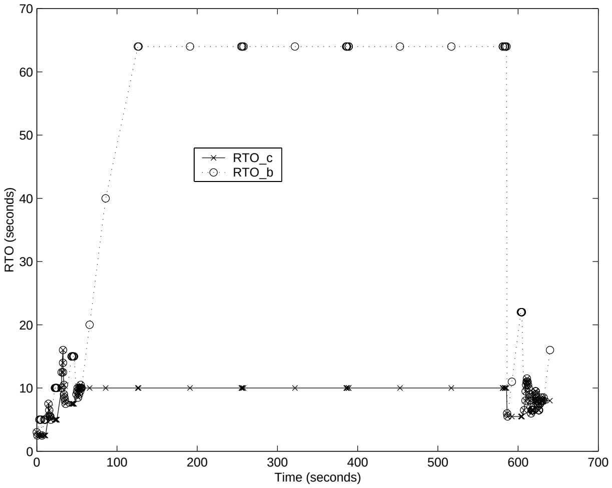

Figure 6: TCP RTO values for the sample path shown in Fig. 5

Consider the history of the TCP RTO updates (shown in Figure 6) occuring when an ACK for a non-retransmitted TCP packet arrives at the remote host. RTO_c is the RTO value calculated from the smoothed RTT A and the smoothed mean deviation D ($RTO_c = A + 4D$), and RTO_b is the RTO value after RTO_c is backed off and eventually upper bounded at RTO_{max} by the Karn's algorithm. When a TCP timeout occurs, TCP congestion avoidance is invoked and congestion window shrinks to one segment. At each retransmission, the TCP RTO increases multiplicatively with the back–off factor β = 2. After several consecutive retransmissions the RTO_b stays at its upper bound value until a valid ACK for non–retransmitted TCP packet arrives During this period, TCP waits a long timeout interval (64 seconds) for each packet loss. From these observations, we can see that in the high FER regime a long idle waiting time significantly impacts performance.

Next, consider an analogous sample path analysis obtained in a simulation with idle frames, shown in Figures 7 and 8, in the same correlated high–loss regime $f_d T = 0.01$, $P_e = 0.3$ for comparison with Figures 5 and 6. The addition of idle frames considerably reduces the timeout intervals, and the TCP data transmission tends to progress more smoothly.

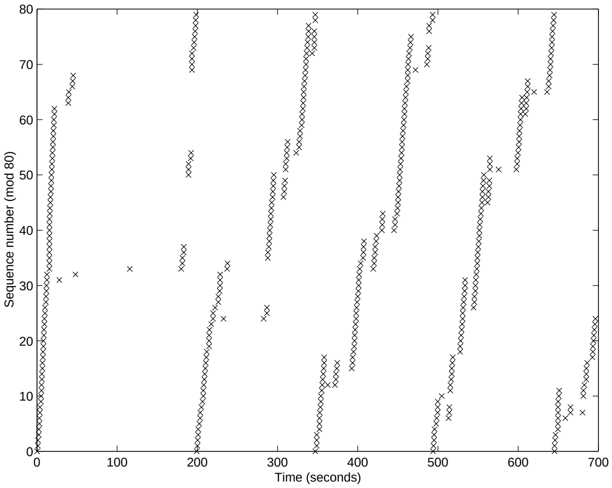

Figure 7: TCP packet sequence number in a simulation sample path with idle frames ($f_d T = 0.01, P_e = 0.3$)

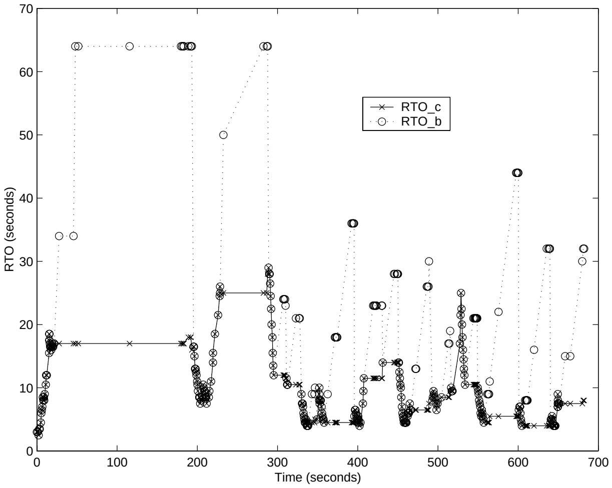

Figure 8: TCP RTO values for the sample path shown in Fig. 7

We also evaluate the effect of the number of idle frames on the TCP goodput. Long–term average TCP goodput for different number of idle frames, at a high error rate (FER $P_e = 0.3$), is shown in Fig. 9. A small number of idle frames makes a big difference in the weakly correlated radio link error regime, and about 5 idle frames can give a very significant performance improvement at the higher $f_d T$ values (0.1 and 1.0), but a further increase of the number of idle frames has no practical effect. In contrast, in the case of strongly correlated link errors ($f_d T = 0.01$) the rate of goodput improvement with increasing number of idle frames is very slow.

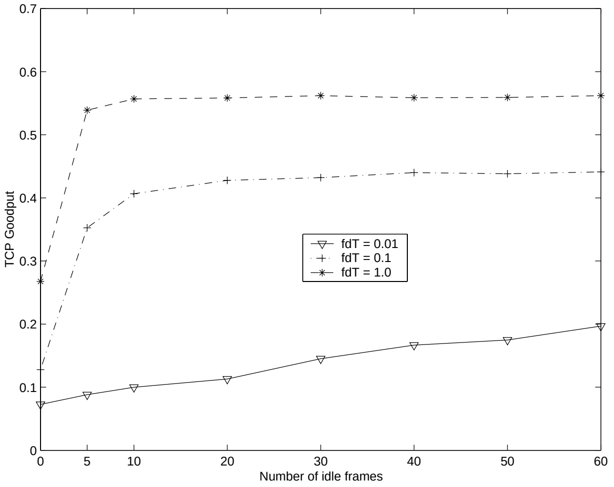

Figure 9: Long-term average TCP goodput as a function of the number of idle frames in high FER regime, $P_e = 0.3$.

Finally, we touch on the issue of variability of TCP goodput on the shorter time scales, that are more relevant to application–level performance than long–time averages. Figures 10, 11, and 12 show a few histograms of short–time TCP goodputs (with unlimited idle frames) in the high frame loss regime, and varying degree of loss correlations. From top to bottom, the goodput is measured over a sequence of 10,000 non-overlapping time intervals of length 10, 100, and 1000 seconds, respectively. The goodput values shown on the x–axis correspond to the number of correctly received (in sequence) TCP segment bytes in the time interval, normalized to the maximum rate. This explains the pronounced quantization effect at the scale of 10 seconds, where maximum of 17 segments can be received.

IV. Summary

In this paper, the performance and interactions of TCP Reno and IS-707 Radio Link Protocol in a low-speed wireless Internet access was investigated by computer simulation. We demonstrated how the use of idle frames in the high correlated frame loss regime can improve the useful throughput (goodput) of TCP data transfers. To get an insight into the potential effects of radio link losses on user applications, we provided examples of the statistics of TCP goodput at short time scales of 10, 100 and 1,000 seconds.

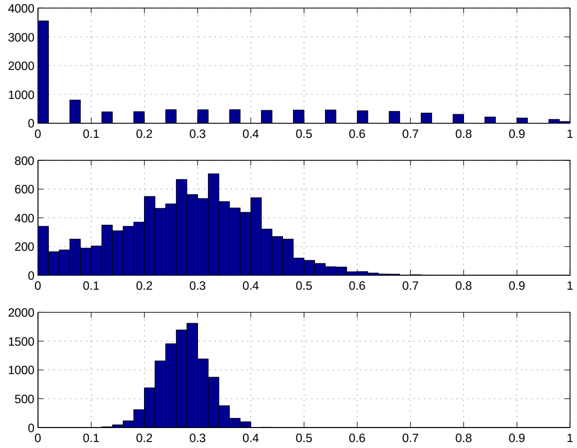

Figure 10: Histogram of TCP goodputs at $P_e = 0.3$, $f_d T = 0.01$, in intervals of 10, 100 and 1,000 seconds.

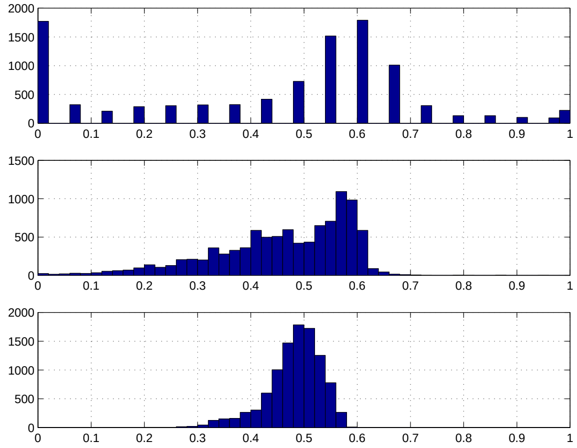

Figure 11: Histogram of TCP goodputs at $P_e = 0.3$, $f_d T = 0.1$, in intervals of 10, 100 and 1,000 seconds.

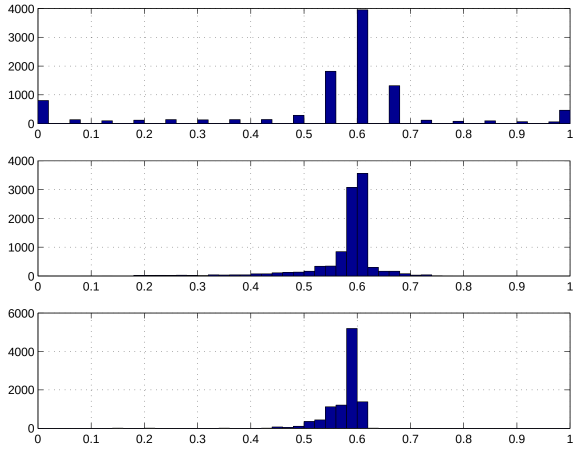

Figure 12: Histogram of TCP goodputs at $P_e = 0.3$, $f_d T = 1.0$, in intervals of 10, 100 and 1,000 seconds.

References

(1) A. Feldmann, A. C. Gilbert, and W. Willinger, “Data Networks as Cascades: Investigating the Multifractal Nature of Internet WAN Traffic,” in *Proc. of ACM SIGCOMM ’98*, pp. 25–38, 1998.

(2) V. Jacobson, “Congestion avoidance and control,” *Computer Communication Review*, vol. 18, no. 4, pp. 314–329, 1988.

(3) P. Karn and C. Partridge, “Improving round-trip time estimates in reliable transport protocols,” in *Proc. of ACM SIGCOMM ’87*, pp. 2–7, Aug. 1987.

(4) TIA/EIA/IS–95, “Data services option standard for wideband spread spectrum digital cellular system,” May 1995.

(5) TIA/EIA/IS–707, “Data services option standard for wideband spread spectrum digital cellular system,” Feb. 1998.

(6) TIA/EIA/IS–99, “Data services option standard for wideband spread spectrum digital cellular system,” July 1995.

(7) J. Cowie, D.M. Nicol, and A.T. Ogielski, “Modeling the Global Internet,” *Computing in Science & Engineering*, vol. 1, pp. 42–50, Jan./Feb. 1999.

(8) A. Chockalingam and G. Bao, “Performance of TCP/RLP Protocol Stack on the Correlated Fading DS-CDMA Wireless Links,” in *Proc. VTC’98*, (Ottawa, Canada), pp. 363–367, May 1998.

(9) G. Bao, “Performance Evaluation of TCP/RLP Protocol Stack over CDMA Wireless Link,” *Wireless Networks*, vol. 2, no. 3, pp. 229–237, 1996.

(10) Y. Bai, G. Wu, and A.T. Ogielski, “TCP/RLP Coordination and Interprotocol Signaling for Wireless Internet,” in *Proc. of VTC’99 Spring*, pp. 1945–1951, May 1999.

(11) S. Bajaj et al., “Virtual InterNetwork Testbed: Status and research agenda,” Tech. Rep. 98-678, University of Southern California, July 1998.

(12) M. Zorzi, R. R. Rao, and L. B. Milstein, “On the accuracy of a first-order Markov model for data block transmission on fading channels,” in *Proc. IEEE ICUPC’95*, pp. 211–215, Nov. 1995.