Measured TCP Performance in CDMA 1x EV-DO Network*

Youngseok Lee
Email: lee@cnu.ac.kr

School of Computer Science & Engineering,
Chungnam National University
Daejon, Korea, 305-764

Abstract. This paper investigates the long-lived TCP bulk throughput over the CDMA 1x EV-DO service that provides high-speed "always on" Internet connectivity in a wide-area mobile environment. Although the peak rates of downlink/uplink are specified as 2.4 Mbps/153 Kbps, the user-experienced application-layer throughput has not been much reported and analyzed. In our experiment, it was shown that average TCP throughputs over downlink/uplink are 572.5/94.7Kbps and the average packet loss rates of 1x EV-DO downlink/uplink are 0.2/4.7%. The average end-to-end round-trip delay was 417.4ms with the variance of 14,995ms. Although the packet loss rate is low, bursty packet losses frequently occur because of packet corruption with TCP checksum failures, which result in TCP performance degradation by the retransmission timeout. Our study showed that this TCP checksum errors are related with the TCP/IP header compression algorithm at link layer protocols such as PPP. Our measurement-based analysis of TCP performance could be used for the correct model of the 3G wireless link characteristic and for the real-world simulation of TCP behavior over the 3G wireless network.
Keywords: Measurement, TCP, CDMA 1x EV-DO, PPP, bursty loss.

1 Introduction

Code Division Multiplexing Access (CDMA) 1x EVolution-Data Only (EV-DO) [1], which has been finalized by the 3G Partnership Project 2 (3GPP2) [2] and has been published by the Telecommunication Industry Association (TIA) as Interim Standard (IS)-856, provides high-speed data service (2.4Mbps/153Kbps for downlink/uplink) with wide coverage and mobility. The CDMA 1x EV-DO service with high-speed "always on" connectivity in a wide-area mobile environment

* This research was supported by the MIC (Ministry of Information and Communication), Korea, under the ITRC (Information Technology Research Center) support program supervised by the IITA (Institute of Information Technology Assessment). All the traffic traces in this paper are available at http://networks.cnu.ac.kr/measurement/cdma-1x-evdo/.

is being deployed throughout the world[1]. The CDMA 1x EV-DO system adds a high-speed data solution for existing IS-95 or CDMA 1x service providers while it maintains the compatibility with the frequency and RF modules. While the specification of 1x EV-DO is based on the High Data Rate (HDR) proposal from Qualcomm, Inc., it includes new features such as Incremental Redundancy (IR) Hybrid ARQ for improving system performance against the fast fading condition. The 1x EV-DO system provides large service coverage as well as handover with high bandwidth. In addition, it employs a time-shared downlink which serves only one user at any instant in time-multiplexed manner. Therefore, the mobile terminal calculates its Signal-to-Interference-plus-Noise Ratio (SINR) and determines the highest data rates among available 11 data rates with the calculated SINR at every slot. Based on the periodic report of the data rates every 1.67ms (1 slot duration), the base station schedules slot allocation.

The characteristics and the performance of cellular links have been widely studied in 2/3G networks (i.e., General Packet Radio Service (GPRS), IS-95A, IS-95B, CDMA2000, and CDMA 1x EV-DO). It is well known that Forward Error Correction (FEC) and link-layer retransmissions have been implemented to defeat the challenging radio propagation environments. When a packet is transmitted, the channel has to be allocated to each packet, which also causes the variable delay. In CDMA 1x EV-DO networks, the peak rate of downlink is specified as 2.4Mbps and that of uplink as 153Kbps per user, and the maximum number of data users is 16. Recently, 1x EVolution Data and Video (EV-DV), that is being standardized, is expected to provide 3Mbps/1.5Mbps downlink/uplink data rates by integrating voice and data channels. Since the cellular network is equipped with wide coverage, mobility, and high data rates, the number of subscribers to the data service of the cellular network grows rapidly[2]. Mobile Internet users with CDMA 1x EV-DO terminals use Internet applications such as web browsing, multimedia streaming, or email. Therefore, the efficient transport protocol is important to achieve the maximum throughput over the error-prone wireless link with fluctuating bandwidth, large delay, and jitter.

Transmission Control Protocol (TCP) in a wireless network experiences several challenges [4–10]. One of the issues is how to deal with the spurious timeout caused by the abruptly increased delay, which triggers unnecessary retransmission and congestion control. It is known that the link-layer error recovery scheme, the channel scheduling algorithm, and handover often make the link latency very high. The *Eifel* algorithm has been proposed to detect the spurious timeout and to recover by restoring the connection state saved before the timeout. Although the packet loss rate of the wireless link has been reduced due to link-layer retransmission and FEC, losses still exist because of the poor radio condition and

[1] In South Korea, since 2002, two carriers (KTF and SK Telecom) have deployed CDMA2000-based 1x EV-DO service which enables Video on Demand (VOD) and Multimedia Message Service (MMS).

[2] The number of CDMA2000 1x EV-DO subscribers was 7.7 million at the end of June 2004 according to the Ministry of Information and Communication in Korea (http://www.mic.go.kr).

mobility. Therefore, non-congestion errors could decrease the TCP sending rate. Packet reordering at the TCP layer may be caused by link-layer retransmission, which also calls for unnecessary retransmission and congestion control. Bandwidth of the wireless link often fluctuates because the wireless channel scheduler assigns a channel for a limited time to a user. Thus, the variance of inter-packet arrival time becomes high, which may result in spurious timeout. In the wireless network, in general, bandwidth and latency at uplink and at downlink directions are different. Hence, the throughput over downlink may be decreased because of ACK congestion at uplink [11].

Previous studies on the throughput over the CDMA 1x EV-DO network have been focused on theoretical analysis or simulation-based performance evaluation of sector throughput as the function of the number of users and the given radio link parameters [12–16]. Although measurement results of TCP performance in the 2.5G GPRS network have been reported in several works [17–19], user's experienced application-level throughput in the real-world 3G CDMA 1x EV-DO environment has not been much studied and analyzed. Moreover, there have been insufficient measurement results of cellular data traffic because of user privacy and high cost of performing measurements. The measurement-based study on TCP performance over CDMA 1x EV-DO is important to derive the high-speed cellular link model and to simulate TCP behavior correctly [20].

In this paper, we investigate the steady-state TCP performance over CDMA 1x EV-DO downlink/uplink with the measurement data of long-lived TCP connections in the commercial CDMA 1x EV-DO network. Our analysis shows that the average TCP throughput over downlink/uplink is 572.5/94.7Kbps, and that the average end-to-end RTT is about 417.4ms with the variance of 14,995ms. The average loss rate at downlink/uplink was 0.2/4.7%. Inspite of the low packet loss rate, multiple packet losses frequently observed.

The remaining paper is organized as follows. Section 2 describes the measurement environment, and Section 3 explains the measurement result of TCP performance. In Section 4, we conclude this paper.

2 Measurement Environments

As shown in Fig. 1, we used a laptop with CDMA 1x EV-DO USB modem and a Linux or FreeBSD machine for our experiments[3]. The laptop was connected to the CDMA 1x EV-DO carrier's network through the point-to-point protocol (PPP). TCP performance has been tested with the *Iperf* [21] tool which generates long-lived bulk traffic between the laptop and the Linux/FreeBSD machine. In our experiments, our Linux machine was located at the research networks called "KOREN" which is connected to the CDMA 1x EV-DO carrier's network via high-speed Korea Internet Exchange Points (KIX, KT-IX). Although the

[3] The operating systems of the laptop for the experiment is Microsoft Windows XP home edition. The operating systems of Unix machines for the experiments are Linux kernel 2.4 and FreeBSD 4.9 release.

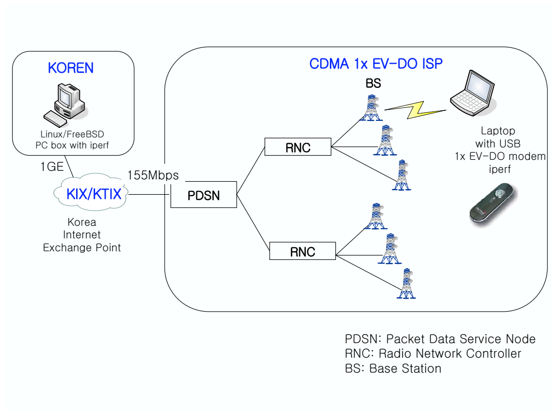

Fig. 1. The measurement environment.

location is not the perfect place to monitor the CDMA 1x EV-DO link performance, it is assumed that the average TCP throughput of the CDMA 1x EV-DO subscribers could be approximately found from many runs of experiments. The Maximum Segment Size (MSS) was set to 1460 or 1448 bytes, SACK was enabled, and the duration of the test was 150 - 300 seconds. It is believed that the test duration is long enough to observe the TCP steady state performance, since the outstanding window has become stable within 20 seconds in our measurements. The packet traces have been collected both at the server and at the client with *tcpdump* or *windump* [22]. The collected packet traces were analyzed with *tcptrace* [23].

3 Measurement Results

3.1 Overall TCP performance statistics

Table 1. Statistics of TCP performance in CDMA 1x EV-DO network

	Downlink	Uplink
No of traces	269	291
Duration per trace(sec)	150-300	150-300
Avg throughput(Kbps)	572.5	94.7
Min/Avg/Max RTT(ms)	109.2/417.4/1278.7	170.3/730.9/2097.1
Variance of avg RTT(ms)	14955.4	16628.1
Avg retransmission rate(%)	0.2	4.7

Table 1 summarizes the overall TCP performance results for stationary hosts in the CDMA 1x EV-DO network. The average TCP throughput over downlink/uplink was 572.5/94.7Kbps and the average packet retransmission rate was 0.2/4.7%. The packet retransmission rate includes the number of retransmitted packets. Therefore, the real packet loss or drop rate at the CDMA 1x EV-DO link is less than the packet retransmission rate. The cumulative distribution function (CDF) plots of average TCP throughput and the average packet retransmission rate for the traces are shown in Fig. 2. Since the average throughput over downlink is widely distributed, it is difficult to find the representative throughput value. However, the throughput over uplink is stable with the average of 94.7Kbps. Generally, it is observed that packet losses occur more often in uplink than in downlink. As shown in Fig. 2, the average packet retransmission rate at uplink is much higher than at downlink.

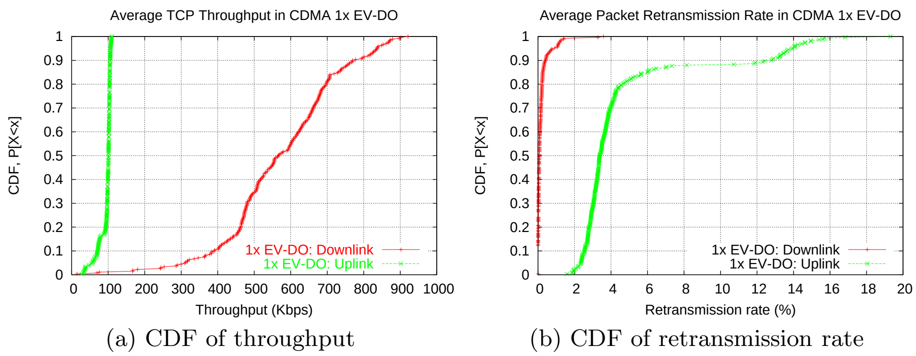

Fig. 2. CDF of TCP throughput and packet retransmission rate in CDMA 1x EV-DO downlink/uplink.

Figure 3 shows the CDF of min/avg/max RTT values for TCP connections at CDMA 1x EV-DO downlink/uplink. Although the RTT value represents the end-to-end round trip delay between hosts, it is dominated by the delay of the CDMA 1x EV-DO link. One of the causes of wide variation of RTT will be the queueing delay at the wireless link which employs link-layer error-recovery and bandwidth-adaption schemes. Especially, the low packet retransmission rate is achieved due to the link layer error control scheme with the overhead of variable delays as shown in Fig. 3.

3.2 Correlation between TCP throughput and RTT/retransmission rate

As the RTT value is a key factor to affect the TCP throughput in Eqn. (1) [24], we plot the TCP throughput in the aspect of the average RTT value. In Fig. 4 and 5, we compared the measured TCP throughput results of TCP connections in 1x and 1x EV-DO environments with the estimated TCP throughput (T in Eqn. (1)) in the aspect of the average RTT and the average packet retransmission rate.

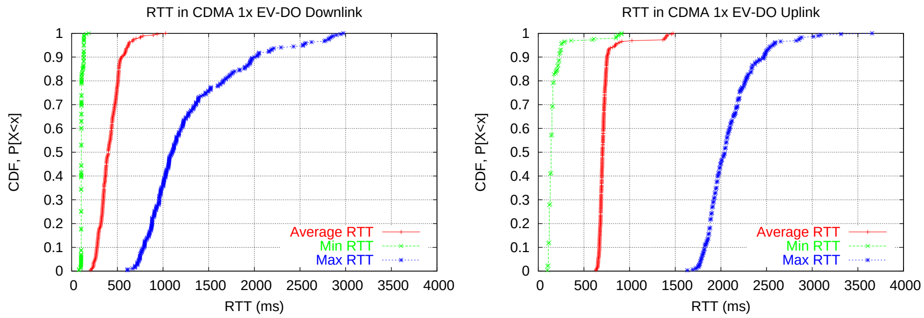

(a) CDF of RTT at downlink (b) CDF of RTT at uplink

Fig. 3. CDF of end-to-end RTT for TCP connections in the CDMA 1x EV-DO network.

$$T = \frac{C \cdot MTU}{RTT \cdot \sqrt{b \cdot p}} \tag{1}$$

, where the loss rate, p, the number of TCP segments per ACK, b, and the constant, C, and MTU are given ($b = 2$, $MTU = 1500$B). It is shown that the estimated TCP throughput is approximating to the measured one in terms of the RTT value in Fig. 4 ($C = 2$, $p = 0.002$).

However, in the aspect of the loss rate as shown in Fig. 5, it is seen that the TCP throughput over CDMA 1x EV-DO is not well described with the estimated model[4]. That is, the low packet loss rate does not always improve the TCP throughput, because the link-layer error recovery mechanism increases the queueing delay at the intermediate node, resulting in the high RTT. Thus, the TCP throughput is decreased by the amount of the increased RTT value. Similarly, the TCP throughput over uplink is decreasing as the RTT value becomes high. In CDMA 1x in Fig. 4, it is shown that the throughput is stable in spite of different RTT values.

3.3 Bursty packet losses

Besides the average packet retransmission rate, we investigated the distribution of the loss pattern which is important for the enhancement of TCP performance, because multiple packet losses will typically induce the retransmission timeout, resulting in TCP slow start. Figure 6 shows an example of multiple packet retransmissions and the CCDF (complementary CDF) plot regarding the number of retransmitted packets in a loss period, which represents bursty packet losses. For example, in Fig. 6(a), an example of multiple packet retransmissions[5] is illustrated. A loss period in our work is defined as the time between the first

[4] In Fig. 5, the packet loss rate was assumed to be the packet retransmission rate for simplicity.

[5] The bottom line represents the ACK from the receiver, and the top line is the advertised window size. Points in the middle are the sequence numbers of arrived TCP segments. Points deonted by "R" mean that these segments have been retransmitted.

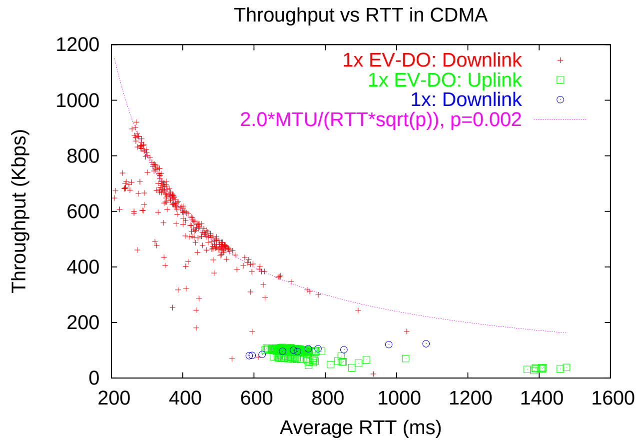

Fig. 4. Correlation between throughput and RTT of TCP connections at CDMA downlink/uplink.

retransmission and the last retransmission of multiple retransmissions without any new packet transmission. The median of retransmissions is one/three at downlink/uplink in CDMA 1x EV-DO. In downlink, the percentage of a single-packet retransmission is 58%, whereas the percentage of multiple (more than two) retransmissions is 92% in uplink. Therefore, uplink is more prone to bursty packet losses

While examining the bursty packet-loss pattern at the receiver, we found that a lot of multiple packets are retransmitted even though the same packets arrived previously. Figure 7 plots the sequence numbers of TCP segments arrived at the receiver (the middle line), ACKs generated by the receiver responding to TCP segment arrivals (the bottom line), and the advertised window value (the top line). It is observed that after a few ACK arrivals the receiver does not respond to the newly arrived TCP segments. Then, the same packets (denoted as "R") arrived again at the receiver, which were retransmitted by the sender after the TCP retransmission timeout. In our experiment, the reason of not responding the TCP segments by the receiver was the incorrect TCP checksum value. In our traces, a lot of packet corruptions with TCP checksum failures have occurred.

After investigating the reason of bursty packet losses due to TCP checksum errors, we have recognized that the TCP/IP header compression scheme (called VJ) [25] at PPP is closely related with bursty packet losses. The TCP/IP header compression mechanism of PPP will compress the TCP/IP header by sending only the changed fields of the TCP/IP header compared with those of the previous header. However, this TCP/IP header compression algorithm is prone to packet losses, because even a single packet loss could make the receiver not recover header-compressed packets that arrived after the lost packet. On the other hand, this header compression algorithm will not work under the TCP timestamp option, because each timestamped header has TCP options that differ from

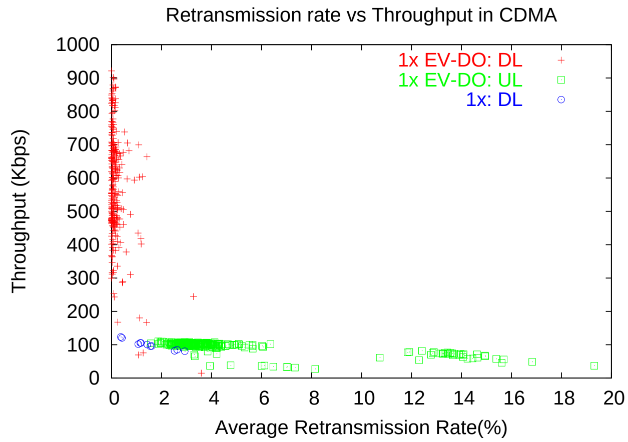

Fig. 5. Correlation between throughput and retransmission rate of TCP connections at CDMA downlink/uplink.

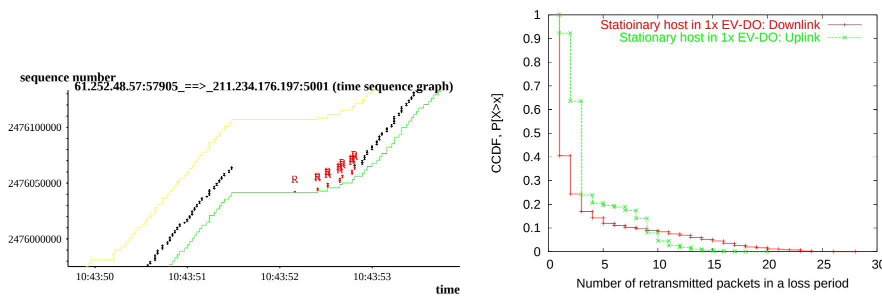

(a) Time-segment graph for an example of multiple packet retransmissions at the sender

(b) CCDF of multiple retransmissions

Fig. 6. Retransmitted packets in a loss period for the stationary host.

the previous header. Therefore, bursty packet losses due to the header compression algorithm at PPP could be reduced when the TCP timestamp option was enabled. In Fig. 8, we compared the typical packet loss patterns of TCP connections with/without the TCP timestamp option over CDMA 1x EV-DO downlink [6]. When the TCP timestamp option was disabled, the average retransmission rate and the number of multiple retransmitted packets become high. Therefore, the TCP/IP header compression algorithm should be carefully used. Otherwise, an enhanced compression algorithm such as robust header compression (ROHC) [26] is necessary.

[6] In this experiment at the CDMA 1x EV-DO downlink, the data file 20060127 is for the disabled TCP timestamp option, and the data file 20060125 for the enabled TCP timestamp option, respectively.

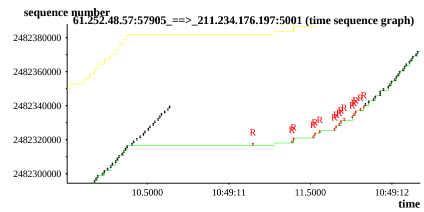

Fig. 7. Multiple packet retransmissions because of packet corruption at the receiver.

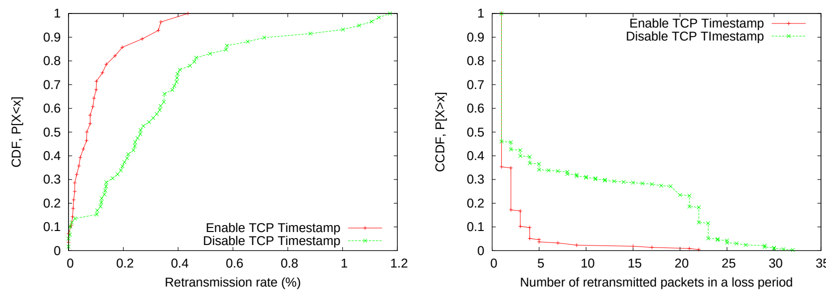

(a) CDF of retransmission rate

(b) CCDF of multiple retrasnmissions

Fig. 8. Typical packet loss patterns when the TCP timestamp option is enabled or disabled.

4 Conclusion

In this paper, TCP performance over the CDMA 1x EV-DO wireless Internet access link has been studied with the measurement data. From the results, it is shown that high throughput over downlink (average of 572.5Kbps) and the moderate throughput over uplink (average of 94.7Kbps) can be achieved. In addition, it is also seen that the large delay and delay variance are the most detrimental reasons to reduce the TCP throughput. Although the packet loss rate is not so high at downlink, multiple packet losses are frequently observed, which are related with the packet corruption. A lot of TCP checksum errors in the measurement results are caused by the TCP/IP header compression algorithm used in PPP. Hence, the TCP/IP header compression algorithm at the link layer should be carefully used to prevent the degradation of the TCP throughput.

References

1. R. Parry, "CDMA2000 1x EV-DO," IEEE Potentials, vol. 21, no. 4, pp. 10 - 13, Oct. - Nov. 2002.
2. 3GPP2, http://www.3gpp2.org
3. 3GPP C.S0024 ver. 4.0, "CDMA2000 High Rate Packet Data Air Interface Specification," 3GPP2, October 2002.

4. H. Balakrishnan, V. Padmanabhan, S. Seshan, and R. Katz, “A Comparison of Mechanisms for Improving TCP Performance over Wireless Links,” IEEE/ACM Transactions on Networking, vol. 5, no. 6, pp. 756-769, 1997.
5. H. Inamura, G. Montenegro, R. Ludwig, A. Gurtov, and F. Khafizov, “TCP over Second (2.5G) and Third (3G) Generation Wireless Networks,” IETF RFC3481, February 2003
6. M. C. Chan and R. Ramjee, “TCP/IP Performance over 3G Wireless Links with Rate and Delay Variation,” IEEE WCNC, 2003.
7. M. C. Chan and R. Ramjee, “Improving TCP/IP Performance over Third Generation Wireless Networks,” IEEE INFOCOM, 2004.
8. R. Ludwig, and R. H. Katz, “The Eifel Algorithm: Making TCP Robust Against Spurious Retransmissions,” ACM Computer Communication Review Journal, vol. 30, no. 1, pp. 30-36, January 2000.
9. R. Ludwig, and M. Meyer, “The Eifel Detection Algorithm for TCP,” IETF RFC3522, April 2003.
10. A. Gurtov and R. Ludwig, “Responding to Spurious Timeouts in TCP,” IEEE INFOCOM, 2003.
11. H. Balakrishnan, V. N. Padmanabhan, and R. H. Katz, “The Effects of Asymmetry on TCP Performance,” ACM/IEEE MobiCom’97, 1997.
12. S. Lee, “The Performance Improvement Principles of TCP Protocol Stack on Packet Switching High Speed Wireless DS-CDMA Links,” IEEE VTC Fall, 2001.
13. W. Chung, H. Lee, and J. Moon, “Downlink Capacity of CDMA/HDR,” IEEE VTC Spring, 2001.
14. Q. Bi, “A Forward Link Performance Study of the 1x EV-DO System through Simulations and Field Measurements,” Lucent Technologies, March 2004.
15. Q. Bi and S. Vitebsky, “Performance Analysis of 3G-1X EVDO High Data Rate System,” IEEE WCNC, 2002.
16. E. Esteves, M. I. Gurelli, and M. Fan, “Performance of Fixed Wireless Access with cdma2000 1xEV-DO,” IEEE VTC Fall, 2003.
17. A. Wennstrom, A. Brunstrom, J. Rendon, J. H. Gustafsson, “A GPRS Testbed for TCP Measurements,” International Workshop on Mobile and Wireless Communications Network, 2002.
18. P. Benko, G. Malicsko, and A. Veres, “A Large-scale, Passive Analysis of End-to-End TCP Performance over GPRS,” IEEE INFOCOM, 2004.
19. R. Chakravoty and I. Pratt, “Performance Issues with General Packet Radio Service,” Journal of Communications and Networks (JCN), vol. 4, no. 2, Dec. 2002.
20. A. Gurtov and S. Floyd, “Modeling Wireless Links for Transport Protocols,” ACM Computer Communication Review, vol. 34, no. 2, pp. 85-96, April 2004.
21. *Iperf*, http://dast.nlanr.net/Projects/Iperf/
22. *Tcpdump*, http://www.tcpdump.org
23. *Tcptrace*, http://www.tcptrace.org
24. J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, “Modeling TCP Throughput: A Simple Model and its Empirical Validation,” in *Proc. ACM SIGCOMM’98*, pp. 303-314, 1998.
25. V. Jacobson, ”Compressing TCP/IP Headers for Low-Speed Serial Links,” IETF RFC1144, Feb. 1990.
26. L-E. Jonsson, ”RObust Header Compression (ROHC): Requirements on TCP/IP Header Compression ,” IETF RFC4163, Aug. 2005.