

TCP over CDMA2000 Networks: A Cross-Layer Measurement Study

Karim Mattar¹ Ashwin Sridharan² Hui Zang² Ibrahim Matta¹ Azer Bestavros¹

¹ Department of Computer Science at Boston University
{kmattar, matta, best}@cs.bu.edu

² Sprint Advanced Technology Labs
{ashwin.sridharan, hui.zang}@sprint.com

Abstract. Modern cellular channels in 3G networks incorporate sophisticated power control and dynamic rate adaptation which can have a significant impact on adaptive transport layer protocols, such as TCP. Though there exists studies that have evaluated the performance of TCP over such networks, they are based solely on observations at the transport layer and hence have no visibility into the impact of lower layer dynamics, which are a key characteristic of these networks. In this work, we present a detailed characterization of TCP behavior based on cross-layer measurement of transport, as well as *RF* and *MAC* layer parameters. In particular, through a series of active TCP/UDP experiments and measurement of the relevant variables at all three layers, we characterize both, the wireless scheduler in a commercial CDMA2000 network and its impact on TCP dynamics. Somewhat surprisingly, our findings indicate that the wireless scheduler is mostly insensitive to channel quality and sector load over short timescales and is mainly affected by the transport layer data rate. Furthermore, we empirically demonstrate the impact of the wireless scheduler on various TCP parameters such as the round trip time, throughput and packet loss rate.

1 Introduction

With advances in error-correction coding, processing power and cellular technology, the wireless channel need no longer be viewed as an error-prone channel with low bandwidth. Instead, modern 3G cellular networks (e.g CDMA2000 1xRTT, EV-DO, HSDPA/UMTS) deploy ARQ mechanisms for fast error recovery, as well as sophisticated wireless schedulers that can perform “on-the-fly” rate adaptation. The latter feature allows the network to adapt to diverse conditions such as channel quality, sector load and more importantly, as we show in this work, data backlog.

The *dynamic rate* adaptation of modern cellular channels implies that a source will typically experience variable bandwidth and delay, which may be caused by the scheduler’s dependency on buffer backlog. Since TCP, the dominant transport protocol in the Internet, utilizes feedback from the channel to control its transmission rate (indirectly the buffer backlog), this creates a situation where two controllers, the wireless scheduler and TCP, share a single control variable.

There are several interesting studies that have considered the performance of TCP over cellular networks [5, 2, 9]. However, they mostly rely on measurement of TCP dynamics at the transport layer and have no visibility into the underlying MAC nor the dynamics of the radio channel. In this work, we measure relevant information at all

three layers in a commercial CDMA2000 network to identify the dominant factors that affect TCP. To the best of our knowledge, this is the first study that looks at cross-layer measurements in a wireless network. Our contributions can be summarized as follows:

1. We conducted extensive active measurements in a commercial CDMA2000 cellular network to characterize the behavior of the wireless scheduler, and evaluate TCP's performance. One of our objectives was to identify the impact of various network factors on both the wireless scheduler and TCP. Towards this end, we develop a simple Information Theoretic framework that allows us to *quantify* how factors such as channel quality, sector load, *etc.*, affect the wireless scheduler, and how the scheduler in turn affects TCP.
2. In terms of the wireless scheduler, we exposed the different mechanisms that govern its operation and identified the characteristics that influence its performance. We concluded that over short timescales (1 second), the wireless scheduler: a) is highly dependent on buffer backlog, b) is surprisingly insensitive to variations in channel quality or sector load, and c) has a rate limiting mechanism to maintain fairness by throttling connections that are being persistently greedy. Over long timescales (20 minutes), however, the scheduler reduces allocated rate in response to persistently bad channel conditions or high sector load, and is unable to maintain fairness among concurrent TCP sessions.
3. In terms of TCP, we concluded that: a) there is a tight coupling between the TCP sending rate and the scheduler. This implies that rate variations, seen by TCP, on the CDMA channel are not random, b) most of the packet losses seen by TCP are congestion related, and c) the high variability in channel rate causes frequent spurious re-transmissions which can be overcome by using the time-stamp option.
4. Finally, as a general observation, we found high variability in TCP throughput based on the time and day of the experiment. We hypothesize this to be due to rapid cell dimensioning by network operators.

The rest of the paper is organized as follows: Section 2 outlines the architecture of a CDMA2000 network and highlights the relevant features. Section 3 presents a description of the various experiments that we conducted. Section 4 characterizes the wireless scheduler and quantifies the relative impact of various factors on it. Section 5 presents an evaluation of TCP's performance. Section 6 presents our conclusions.

Due to lack of space, some characterization details were omitted from this paper. The reader is referred to [7] for all the details.

2 The CDMA2000 1xRTT System

Figure 1(a) depicts the architecture of a typical cellular data network in order to illustrate its salient features. The network consists of two main components: a) the *data network* which is responsible for operations like managing PPP sessions, IP mobility and billing, and b) the *radio network* which manages radio resources. The focus of this work is on the latter.

The main element of the radio network is the Base Station Controller (BSC), which is responsible for maintaining the *radio session* with the mobile device. The BSC controls hundreds of Base Transceiver Stations (BTS), which are essentially the air interfaces to the mobile devices, through a low-latency back-haul network. The importance

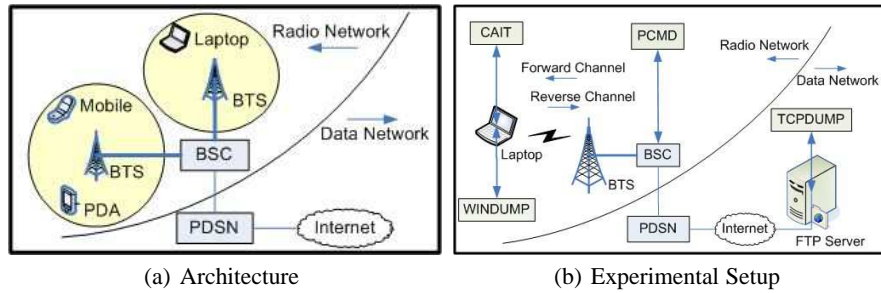


Fig. 1. CDMA2000 Network Architecture and Experimental Setup

of the BSC arises from the fact that it hosts two critical components which can directly affect higher layer performance: a) the *wireless scheduler* that dynamically controls the wireless channel rate assigned to the mobile device, and b) the *Radio Link Protocol* (RLP) that is responsible for fast MAC layer error recovery through the re-transmission of radio frames to recover from losses either over the low-latency back-haul connecting the BSC to the BTS or the wireless channel.

The function of the wireless scheduler is to assign a wireless channel rate (from up to six discrete rates) to a mobile device *on-the-fly*. This objective is primarily achieved by controlling the CDMA code length and channel power. Since higher rates require more power and resource reservation, the decision on when to allocate higher rates and to which user, must be made judiciously. In practice, the scheduler's decision could be influenced by three factors, which we investigate in detail in Section 4:

1. The queue length at the BSC (each user is assigned a separate buffer)
2. The channel conditions experienced by the mobile device, which is defined as the ratio between the received pilot signal strength E_c and the ambient noise and interference I_0 , *i.e.*, E_c/I_0
3. The number of active users in the same sector (or sector load)

3 Experiments and Data sets

Our primary focus is on the *downlink*. We performed end-to-end experiments which involved data transfer via either UDP or TCP SACK from a RedHat Linux server on the Internet to one or more laptops running Windows XP that were connected to the cellular data network via CDMA2000 1xRTT air-cards. A typical experimental setup is shown in Fig. 1(b) to illustrate the data path³, as well as measurement points.

The experiments can be categorized into two classes. The first class consisted of sending UDP traffic to characterize the wireless scheduler. UDP was chosen to remove any transport layer feedback so that the wireless scheduler could be characterized in isolation. The second class comprised of downloading files via TCP in order to characterize long term TCP behavior, as well as its dependency on RF factors. These experiments were conducted under different TCP-specific and wireless configurations to evaluate their relative impact and obtain a better understanding of the system.

³ The end-to-end path on average had a: a) propagation delay of 450-550ms, b) 25-35KB bottleneck buffer at BSC, c) 70KB-120KB wireless channel rate, and d) packet loss rate of 0.9%.

Each experiment, under every configuration, was run 10 times at various times during the day to obtain a reasonably large set of samples for statistical characterization. All plots include error bars denoting the 90% confidence interval around the mean. For TCP downloads, we used a single file size of 5MB since we are interested in long-term TCP behavior. The typical duration of both TCP and UDP flows was 15-20 minutes.

For each experiment, we collected data from the higher layer protocols through standard UDP/TCP logs at the client (windump) and server (tcpdump), as well as RF layer information. The RF statistics were collected from two observation points. Messages related to instantaneous channel quality, frame errors, re-transmissions and the assigned wireless channel rate were collected at the laptops using an air-interface analysis tool called CAIT [8]. These messages were aggregated to generate a time-series tracking the value of the above RF variables at a time-granularity of 1 second. The second source of measurement was the *Per Call Measurement Data* (PCMD) obtained from the BSC. PCMD contains details for all voice and data calls⁴, such as source/destination phone numbers, cell/sector ID's, call duration, number of transmitted bytes, and call success/failure information. We used the PCMD logs to infer the number of active sessions in a sector (*i.e.*, sector load).

4 Wireless Scheduler Characterization

In this section we present an empirical evaluation of the factors that affect the behavior of the wireless scheduler. The exact implementation details of a commercial 1xRTT wireless scheduler is proprietary and hence we have to infer its behavior indirectly.

As mentioned in Section 2, channel rates assigned by the wireless scheduler can be influenced by the user's queue length, channel conditions and sector load. To understand the extent to which each of these factors affects the scheduler's decisions, we performed numerous UDP experiments with constant bit-rate (CBR) and 'on-off' traffic sources.

We begin by examining the impact of the application data rate on the wireless scheduler. Figure 2 plots the average throughput of a connection as a function of the data sending rate of a UDP CBR traffic source. The figure indicates that the achieved throughput tracks the sending rate up to a rate of 50 kbps after which it *decreases* sharply.

We next study *how* the assigned channel rate tracks the source's data rate and empirically show that the wireless scheduler assigns channel rates by tracking the user's *buffer backlog* over *short timescales*. Figure 4 plots the source's data rate and the assigned channel rate time-series (measured by CAIT) when the source utilizes an 'on-off' traffic pattern. The two figures show the data sending rate for two different on-rates (38.4kbps and 76.8kbps) and the same duty cycle (on for 1-second and off for 5-seconds). For the wireless scheduler to assign the correct channel rates for every 1-second burst of data transmitted by the sender, it must inspect the buffer backlog (or some function of it) at least once every second. This implies that the scheduler's decisions are very sensitive to the data sending rate determined by the transport and application layers.

Even though Fig. 4 indicates the presence of a strong correlation between the scheduler and the buffer backlog, visual identification is not always reliable or more importantly quantifiable, which is necessary for comparison. Consequently, we have developed an Information-Theoretic methodology based on a metric that we refer to as the

⁴ A single sector, covered by a BTS, typically had 8-9 (and a maximum of 30) active calls.




Fig. 2. Throughput for UDP CBR experiments

Variable	NMI Value
Buffer Backlog	0.57
Channel Quality (E_c/I_0)	0.15
Sector Load	0.11

Fig. 3. Factors affecting channel rate

Normalized Mutual Information to quantify the correlation between two time-series. We explain the main idea briefly. The Normalized Mutual Information (NMI) measures the amount of information [4] a time-series X can provide about another time-series Y (taking into consideration time-shifts between the two sequences). NMI lies between 0 and 1 and the larger the peak value, the more the two sequences are dependent on each other. NMI is basically a time-shifted correlation measure that can capture both linear and non-linear dependencies between time-series. For completeness, NMI is defined as:

$$I_N(X; Y; d) = (H(X) + H(Y) - H(X, Y_d)) / H(X) \quad (1)$$

where $H(X)$ denotes the entropy of X , and $H(X, Y_d)$ denotes the joint entropy of X and a version of sequence Y that is time delayed by d time units⁵.

Armed with this tool, we evaluated the impact of both the *channel conditions* and the *sector load* on the rate decisions made by the wireless scheduler, over short timescales, when sending UDP CBR traffic. We also computed the NMI between the buffer backlog and the channel assigned rate time-series for the ‘on-off’ traffic source. Figure 3 presents the peak NMI values indicating the *relative* impact that all three factors (averaged across all experiments) have on the channel assigned rates.

Quite surprisingly, as evident from the table, we found that the *short-term* channel quality and sector load have a very limited impact on the scheduler. We believe this can be attributed to two reasons: a) fast power control⁶ deployed in CDMA networks combats short-term channel fluctuations thus eliminating its impact on rate assignment, and b) the sector load varies too slowly to have a significant *short-term* impact on the scheduler’s behavior. In Section 5 we show that these two factors affect *long-term* average scheduler behavior. The table also supports our initial observation that the *buffer backlog* is the most dominant factor influencing the scheduler’s decisions.

5 TCP in CDMA2000 Networks

This section is devoted to results related to TCP behavior in experiments that were conducted under four different configurations on a commercial CDMA2000 network,

⁵ When computing the peak NMI we consider all possible delay shifts d .

⁶ BTS boosts the transmitted signal’s strength to increase the signal-to-noise ratio at the receiver.




Fig. 4. Behavior of the wireless scheduler over short timescales. On-off 1s-5s at 38.4kbps (Left). On-off 1s-5s at 76.8 kbps (Right).

where the following parameters were varied, namely: a) TCP’s advertised receiver window size (ARWND), b) TCP’s time-stamp option, c) number of active users in the sector, and d) user mobility (*i.e.*, speed) and location (*i.e.*, channel quality). The first two configurations are specific to TCP, while the latter two are more characteristic of wireless users. Clearly, there are several other variations (e.g. the TCP congestion mechanism) that are possible, however, we believe that the configurations we focused on are the most fundamental ones. Specifically, almost all operating systems have support for modifying TCP’s ARWND and time-stamp options. Similarly, mobility and user location are the main characteristics of *any* wireless network.

Before discussing the experimental results, it is worthwhile making a general observation regarding our results. In almost all configurations, we found that the amount of variation in attained throughput is at least 10% and at times higher than 100%, even across consecutive runs of the same experiment, depending on location and time. Our measurements indicate that these variations in capacity were not caused by the channel quality or sector load. Instead, we believe they may be due to dynamic cell-dimensioning⁷ performed by the network operator which is the focus of our future work.

5.1 TCP Window Size and Time-stamp Option

The first two configurations we study involve TCP’s behavior as a function of ARWND, both when the time-stamp option [10] was disabled/enabled. The window size was varied from 8 KB to 128 KB to control the number of packets in-flight in the network and hence the bottleneck queue size (equivalently queuing delay). This allowed us to subject the wireless scheduler to different queuing loads. Also, setting small advertised receiver windows (particularly ones that are smaller than the bottleneck buffer at the BSC) allowed us to emulate environments where wireless losses are more prevalent (*i.e.*, no congestion losses, due to buffer overflow, could ever occur). The time-stamp option, on the other hand, primarily allows the TCP sender to obtain more accurate RTT/RTO estimates which aid in detecting and avoiding spurious re-transmissions. Enabling/disabling this option allowed us to evaluate its impact on a connection’s performance.

⁷ Adapting a BTS’s maximum transmission power allows it to vary its coverage area.

We begin by evaluating the impact of RF factors on a particular aspect of TCP, the round trip time (RTT). The NMI metric is used to quantify the relative impact of the RLP re-transmission rate, wireless channel rate, as well as the buffer occupancy (approximated by the number of un-acknowledged packets in-flight). Table 1 presents the peak NMI values between the RTT time-series and each of the three factors for a few different ARWND values. The table clearly indicates that for small and medium window sizes (8 KB and 16 KB) the wireless channel rate has the strongest influence on RTT, while at large window sizes (64 KB), buffer occupancy is high, and hence queuing delay becomes the dominant component in RTT.

These observations have several implications: a) RLP re-transmissions do not add significant latency, and b) in the absence of queuing, the channel rate determines RTT. Since RTT directly impacts throughput, we can expect TCP's throughput to be highly dependent on the assigned channel rates. This is indeed true as shown in Table 2 which presents the amount of information (NMI) that the channel rate and RLP re-transmission rate have about TCP's throughput. We showed previously in Section 4 that the channel rate is influenced by the transport data rate which implies that there is a strong coupling between the wireless scheduler and TCP. More importantly, it indicates that the rate variations in the wireless channel are *not* completely random, as is commonly assumed in models [1, 3]. Instead, the channel's rate is *highly* correlated with TCP's state and must be taken into consideration in future models, and c) since the two controllers (TCP and wireless scheduler) take each other as input, this can lead to oscillations resulting in highly variable RTT, causing spurious re-transmissions.

We study the latter two issues in more detail. Figure 5 (Left) plots the cumulative TCP throughput as a function of ARWND. As a general observation, we note that TCP throughput increases as ARWND is increased, which is to be expected. First, let's consider the coupling between the scheduler and TCP. To highlight how this is different from an arbitrary random channel with some mean rate, we plot throughput obtained from an *ns-2* simulation that has the same parameters as the active experiments except that the channel rate was set to the *average* assigned channel rate inferred from the CAIT logs. One can clearly see that the simulation predicts a far higher throughput than that obtained from the experiments (the lowest curve being the one with the time-stamps option disabled). Quantifying the exact relative impact of random channel variations (considering the timescale and magnitude of the variations) and coupling between scheduler and TCP (considering sensitivity of assigned channel rate to buffer backlog and vice versa) is part of our future work.

The second aspect we mentioned was that oscillations produced by such a coupling could result in highly variable RTT, causing spurious re-transmissions, including spurious timeouts [6]. To test this hypothesis, we ran experiments where the time-stamp option was enabled, and thus RTT variability is more accurately captured. Indeed, with time-stamps, the attained throughput is much higher as shown in Fig. 5 (Left).

Finally, we briefly look at packet losses as a function of window size. Figure 5 (Right) shows that packet losses increase with ARWND. When the window size is less than the bottleneck buffer at the BSC (25-35KB), all losses are due to the wireless channel. For window sizes larger than 35KB, *congestion* becomes the dominant cause of packet loss as the packet loss rate curve flattens out.

(ARWND) Advertised Receiver Window Size	Peak NMI Values		
	Channel Rate	RLP	Packets in Flight
8 Kbytes	0.18	0.11	0.05
16 Kbytes	0.20	0.04	0.038
64 Kbytes	0.08	0.03	0.41

Table 1. Impact of various factors on RTT

(ARWND) Advertised Receiver Window Size	Peak NMI Values	
	Channel Rate	RLP
8 Kbytes	0.16	0.06
64 Kbytes	0.26	0.03

Table 2. Impact of RF factors on instantaneous TCP throughput




Fig. 5. TCP throughput (Left) and packet loss (Right) as a function of ARWND.

5.2 Sector Load, User Mobility and User Location

Our final two configurations incorporate characteristic wireless behavior. For these experiments TCP’s ARWND was set to the default (64 KB) with no time-stamp option.

We first varied the number of active TCP sessions (*i.e.*, data calls) within a single sector to study how TCP throughput changes with sector load, as well as evaluate the wireless scheduler’s fairness. Towards this end, we simultaneously downloaded files from up to 4 co-located laptops. The experiments were conducted during off-peak hours to ensure that the only users in the sector were the experiment laptops. In Fig. 6 we plot the cumulative TCP throughput (Left), as well as the Fairness Gap (Right) as a function of active users. A perfectly fair scheduler would result in a Fairness Gap of 0. The larger the gap, the more unfair the scheduler. For any given set of throughput values (y_1, y_2, \dots, y_n) , the Fairness Gap is defined as:

$$f_{\text{gap}}(y_1, y_2, \dots, y_n) = \frac{\max(y_1, y_2, \dots, y_n) - \min(y_1, y_2, \dots, y_n)}{\min(y_1, y_2, \dots, y_n)} \quad (2)$$

As expected, the average throughput achieved per user decreases as the number of active connections increases. However, we note that the fairness of the scheduler degrades with the number of active connections, as reflected by a larger Fairness Gap. Indeed, manual inspection of our experiments indicate that the throughput achieved by concurrent connections can be highly disparate with typically one user dominating.

The final configuration involved evaluating the impact of user mobility and location on the connection’s performance. The mobility experiments were conducted on a 30-mile stretch of highway (RT 101) between San Francisco and Palo Alto, during non-peak hours (*i.e.*, at night). Connections lasted 10-15 minutes which is the time it takes to




Fig. 6. Impact of sector load: throughput (Left) and fairness gap (Right)

download a 5MB file. A BTS's coverage area is 2-3 miles causing a 10-15 minute connection to experience at least two hard hand-offs⁸, assuming the car's speed is 45mph. Due to mobility, specifically path loss and shadowing, the mobile device experiences large variations in channel quality causing E_c/I_0 to fluctuate between 0dB (best possible channel) and -16dB (very poor channel). Figure 7 (Left) shows the achieved TCP throughput for three different average speeds of 45, 60 and 75 mph, respectively. Surprisingly, user speed had little impact on TCP throughput indicating that the cellular network is well engineered for fast hand-offs. We note that mobility is a major concern in 802.11 networks which are not *a priori* designed to handle fast transitions.

The last set of experiments were conducted to investigate the impact of *average (long-term)* channel conditions. In Section 4 we showed that the *short-term* scheduler behavior was not affected by *instantaneous* variations in channel conditions. However, it is unclear whether this observation carries over to *longer timescales*. To investigate this, we performed two sets of experiments, where the laptop was placed in locations with either *consistently good* or *bad* channels. The average throughput for each location is plotted in Fig. 7 (Right)⁹. One can clearly see that the throughput in locations with better channel conditions (*i.e.*, higher E_c/I_0) is much higher. This indicates that the long-term scheduler behavior is indeed affected by *average* channel conditions and not overcome by power control.




Fig. 7. Impact of user mobility (Left) and user location (Right)

⁸ A hard hand-off occurs when the BSC completely switches data from one BTS to another.

⁹ The RTT, general path characteristics and variation in E_c/I_0 , for both locations, were very similar.

6 Conclusions

We conducted a detailed cross-layer measurement study to evaluate TCP behavior over CDMA2000 networks. The study was conducted under various configurations that involved simple variations of TCP, as well as, changing user mobility and sector load. By conducting measurements across all three (*i.e.*, transport, MAC and RF) layers, we were able to evaluate the system comprehensively. Our main findings were:

1. The RLP layer allows fast error recovery that almost eliminates packet loss observed at the transport layer, while having a minimal impact on TCP's RTT.
2. The wireless scheduler in CDMA2000 networks is unaffected by channel conditions or sector load over short timescales. Instead, the decisions are highly dependent on the transport data rate. However, the long-term scheduler rate allocation is indeed affected by *average* channel conditions and sector load. Furthermore, increasing sector load deteriorates the fairness of the scheduler.
3. The wireless scheduler and TCP are strongly coupled which can result in highly variable RTT. Apart from modeling implications since the rate variations are not completely random, it motivates the need for robust RTT estimation to prevent spurious re-transmissions.
4. Mobility is well supported in the CDMA2000 network and hence had no major impact on TCP throughput.

References

1. E. Altman, C. Barakat, and V. M. R. Ramos. Analysis of AIMD protocols over paths with variable delay. In *Proc. IEEE INFOCOM*, Hong Kong, March 2004.
2. P. Benko, G. Malicsko, and A. Veres. A Large-scale Passive Analysis of End-to-End TCP Performance over GPRS. In *Proc. IEEE INFOCOM*, Hong Kong, 2004.
3. M. C. Chan and R. Ramjee. TCP/IP Performance over 3G Wireless Links with Rate and Delay Variation. In *Proc. ACM MOBICOM*, pages 71–82, 2002.
4. T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 1991.
5. Y. Lee. Measured TCP Performance in CDMA 1xEV-DO Network. In *PAM*, Adelaide, Australia, 2006.
6. R. Ludwig and R. H. Katz. The Eifel Algorithm: Making TCP Robust Against Spurious Retransmissions. *SIGCOMM*, 30(1):30–36, 2000.
7. K. Mattar, A. Sridharan, H. Zang, I. Matta, and A. Bestavros. TCP Over CDMA2000 Networks : A Cross-Layer Measurement Study. Available at <http://research.sprintlabs.com>, Sprint ATL, October 2006.
8. Qualcomm. CDMA Air Interface Tester. www.cdmatech.com/products/cait.jsp.
9. J. Ridoux, A. Nucci, and D. Veitch. Seeing the difference in IP Traffic: Wireless versus Wireline. In *Proc. IEEE INFOCOM*, Barcelona, Spain, 2006.
10. D. Borman V. Jacobson, R. Braden. RFC 1323: TCP Extensions for High Performance.