



**TCP OVER CDMA2000 NETWORKS
A CROSS-LAYER MEASUREMENT STUDY**

KARIM ABDEL MAGID MATTAR

Thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Arts

**BOSTON
UNIVERSITY**

BOSTON UNIVERSITY
GRADUATE SCHOOL OF ARTS AND SCIENCES

Thesis

TCP OVER CDMA2000 NETWORKS
A CROSS-LAYER MEASUREMENT STUDY

by

KARIM ABDEL MAGID MATTAR

Bachelor of Science, Computer Systems Engineering, UMass Amherst, 2003

Submitted in partial fulfillment of the
requirements for the degree of
Master of Arts

2008

Approved by

First Reader

Abraham Matta, PhD
Associate Professor of Computer Science Department

Second Reader

Azer Bestavros, PhD
Professor of Computer Science Department

*All truth passes through three stages.
First, it is ridiculed.
Second, it is violently opposed.
Third, it is accepted as being self-evident.*
Arthur Schopenhauer

Acknowledgments

First and foremost I would like to thank my advisor, professor Abraham Matta, for his support, guidance and infinite patience. For that I am truly indebted. I would also like to thank professor Azer Bestavros, Ashwin Sridharan and Hui Zang for their very useful feedback and input. Finally, I would like to thank my parents and my brother for always being there for me.

TCP OVER CDMA2000 NETWORKS
A CROSS-LAYER MEASUREMENT STUDY
KARIM ABDEL MAGID MATTAR

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.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Contributions	2
2	The CDMA2000 1xRTT System	4
2.1	Network Architecture	4
2.2	Wireless Scheduler	6
2.3	Radio Link Protocol	7
3	Problem Statement	9
4	Experiments and Data Sets	11
5	Empirical Evaluation: Methodology and Tools	13
5.1	Time Series Correlation	13
5.1.1	Mutual Information as a Correlation Measure	13
5.1.2	Dealing with Stochastic Delays	15
5.1.3	Dealing with Real-Valued Time Series	16
5.1.4	Guidelines for Identifying Potential Correlations	17
5.1.5	Correlation Example using Mutual Information	18
5.1.6	Other Variants of Mutual Information	19
5.2	Analysis at Multiple Time Scales: Wavelets	19
6	Characterizing the CDMA2000 1xRTT System	22
6.1	Wireless Scheduler	22
6.1.1	Impact of Data Sending Rate and Buffer Backlog	22
6.1.2	Impact of Channel Quality over a Short Time Scale	25

6.1.3	Impact of Sector Load over a Short Time Scale	26
6.1.4	General Behavior	27
6.1.5	Rate-Limiting Mechanism	28
6.2	Radio Link Protocol	30
6.3	UDP Connections	31
6.3.1	Round Trip Time	32
6.3.2	Packet Loss	32
6.3.3	Packet Reordering	32
7	TCP in CDMA2000 Networks	34
7.1	Impact of TCP Window Size and Time-stamp Option	34
7.1.1	On TCP Round Trip Time	35
7.1.2	On TCP Throughput	36
7.1.3	On TCP Timeouts, Re-transmissions and Packet Losses	39
7.2	Impact of Sector Load, User Mobility and Location on TCP Throughput	41
8	Conclusions	45
	References	47

List of Figures

2-1	Network Architecture	4
4-1	Experimental Setup	11
5-1	PDF of RTT for Different TCP ARWND	16
5-2	Correlation Results Using a Synthetic Data Model	19
5-3	Wavelet Decomposition of a Sample Forward Channel Assigned Rate Signal	21
6-1	Average Throughput for UDP CBR Experiments	23
6-2	Wireless Scheduler Behavior at Short Time Scales	23
6-3	Rate Scheduler's Sensitivity to Buffer Occupancy	24
6-4	Sample E_c/I_o Signal	25
6-5	Impact of Channel Conditions on Wireless Scheduler	26
6-6	Impact of Sector Load on Wireless Scheduler	27
6-7	Properties of the Wireless Scheduler	28
6-8	Time Series of Assigned Channel Rate for UDP CBR Experiments	29
6-9	Channel Scheduler Tracking an On-Off Data Source with Peak Rate of 153.6 kbps	30
6-10	Sources of RLP Re-transmissions	31
6-11	CDF of Round Trip Time for UDP CBR Experiments	32
6-12	CDF of the Fraction of Packets Lost Per Second for UDP CBR Experiments	33
7-1	RTT Behavior as a Function of ARWND	35
7-2	Impact of RF Factors and Queuing on RTT	37
7-3	Factors Affecting Instantaneous TCP Throughput	38
7-4	Coupling Between Sender and Wireless Scheduler	39

7·5	TCP Throughput as a Function of ARWND	40
7·6	TCP timeouts, Re-transmissions and Packet Loss	41
7·7	Histogram of Re-transmissions Versus Time Delay Since Last Ack Was Received	42
7·8	Impact of Sector Load on TCP	43
7·9	Impact of User Mobility (Speed)	43
7·10	Impact of User Location	44

List of Abbreviations

AAA	Authentication, Authorization and Accounting
ARWND	Advertised Receiver Window
ARQ	Automatic Repeat Request
BTS	Base Transceiver Station
BSC	Base Station Controller
CAIT	CDMA Air Interface Tester
CBR	Constant Bit Rate
CDF	Cumulative Distribution Function
CDMA	Code Division Multiple Access
EV-DO	Evolution Data Only
FER	Frame Error Rate
FCH	Fundamental Channel
FTP	File Transfer Protocol
HA	Home Agent
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
NACK	Negative Acknowledgment
NMI	Normalized Mutual Information
PCMD	Per Call Measurement Data
PDSN	Packet Data Serving Node
PPP	Point-to-Point Protocol
RF	Radio Frequency

RLP	Radio Link Protocol
RTT	Round Trip Time
SACK	Selective Acknowledgment
SCH	Supplemental Channel
SINR	Signal to Interference Plus Noise Ratio
SLC	Single-Link Clustering
TCP	Transmission Control Protocol
UDP	User Datagram Protocol

Chapter 1

Introduction

1.1 Motivation

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 (Lee, 2006; Benko et al., 2004; Ridoux et al., 2006). 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.

1.2 Contributions

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 the radio link protocol, we show that a high frame loss (or error) rate could be due to a congested back-haul link that causes packets to get lost (or corrupted) after they are broken down into frames by the Base Station Controller.
4. 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, c) the high variability in channel rate causes frequent spurious retransmissions which can be overcome by using the time-stamp option, and d) the radio link protocol, despite employing an aggressive retransmission scheme to significantly reduce wireless losses, has a limited impact on TCP's round trip time and throughput.

5. 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.

Thesis Organization. Chapter 2 outlines the architecture of a CDMA2000 network and highlights the relevant features. Chapter 3 outlines our problem statement, along with our objectives for conducting this measurement study. Chapter 4 presents a description of the various experiments that we conducted. Chapter 5 explains our empirical evaluation methodology which is based on correlating time series that capture the evolution of various system parameters, to better understand the interdependence between different system components. Chapter 6 characterizes various aspects of the CDMA2000 network, namely, the performance of the wireless scheduler, the radio link protocol and end-to-end UDP connections. Chapter 7 presents an evaluation of TCP's performance. Chapter 8 presents our conclusions.

Chapter 2

The CDMA2000 1xRTT System

In this chapter, we illustrate the architecture of modern cellular data networks, as well as identify salient properties of CDMA2000 1xRTT, a 2.5G technology, which is widely used in these networks and was the technology available at the time we conducted our experiments. In particular, we highlight key features of the CDMA2000 network, that either directly or indirectly affect higher layer performance and motivate the need to characterize their impact in realistic environments. We believe our findings may also be applicable to current 3G networks based on the 1xEV-DO technology because they share some similar features.

2.1 Network Architecture

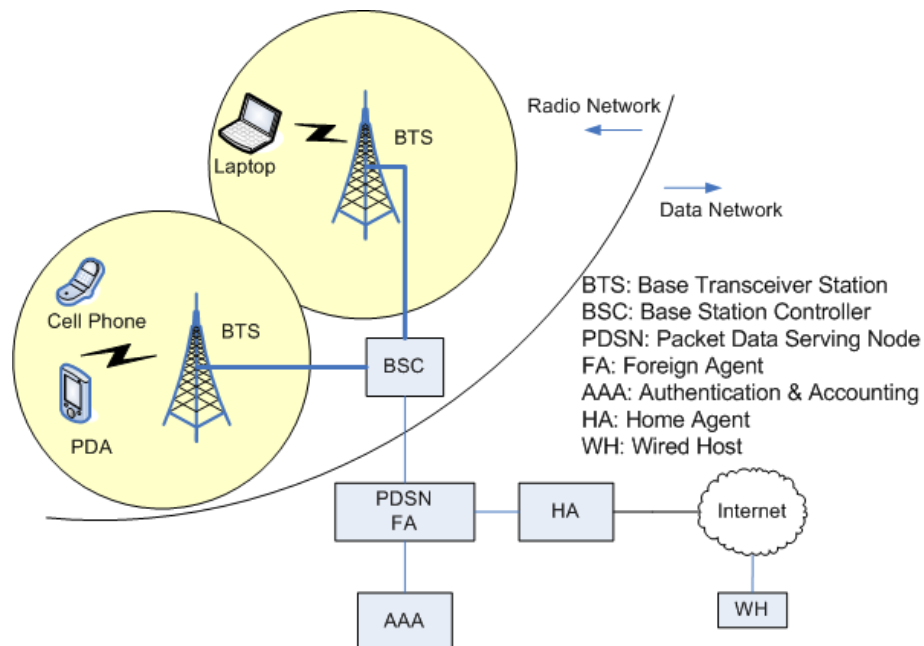


Figure 2.1: Network Architecture

Figure 2-1 sketches the architecture of a typical cellular data network. 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.

More specifically, the data network is an all-IP network comprising of the PDSN (Packet Data Serving Node), the HA (Home Agent) and the AAA (Authentication, Authorization and Accounting) server. The PDSN, residing in the data network, acts as the interface agent between the two networks. It establishes a PPP session for each cellular user and forwards traffic received from the radio network to the HA and vice versa. The HA is responsible for IP address allocation, forwarding cellular IP traffic to (and from) the Internet and more importantly, managing user mobility via Mobile IP (C. Perkins, 1996). The AAA server mainly addresses the requirements of authentication, billing, etc.

The radio network, on the other hand, which is actually the focus of this study, comprises the air interface and two basic elements: a Base Transceiver Station (BTS) and a Base Station Controller (BSC). The BTS, or simply put, the base station, is essentially a “dumb terminal” comprising only of antenna arrays to efficiently radiate RF (Radio Frequency) power to mobile users, as well as receive signals from them. Hence, it acts as the interface between the “wireline” network and the “wireless” device. Each such base station represents a “cell”. For purposes of efficient frequency re-use, the cell is typically split into three sectors by suitable alignment of the antenna profile into three geographically distinct radiation beam patterns. Users in the same cell but different sectors can operate independently, but users in the same sector must share air resources.

The BSC is the main element in the radio network and is responsible for maintaining the *radio session* with the mobile device. The BSC controls hundreds of base stations, through a low-latency back-haul network. The importance 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 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 frame losses. We will outline the key features of these two components in more detail next.

2.2 Wireless Scheduler

The function of the wireless scheduler is to assign a wireless channel rate, from up to *six* discrete rates depending on the Radio Configuration Type (Vanghi et al., 2004) 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.

The network utilized for our experiments supports five channel rates. The smallest assignable rate, denoted by the *Fundamental Channel* (FCH) is 9.6 kbps. This is the standard channel assigned to all voice users and initially to a data user upon joining the network. If a user requires higher data rates, the scheduler can assign it a *Supplemental Channel* (SCH) in bursts of short durations. The Supplemental Channel can take rates from the set {**19.2, 38.4, 76.8, 153.6**} kbps.

Though a shorter Walsh Code increases the data rate, it has two drawbacks. First, the reduced code length degrades “orthogonality”, which makes the signal more susceptible to interference from other users. To overcome this problem, the scheduler employs two techniques:

1. The signal strength is boosted for users assigned a higher rate channel to overcome increased interference.
2. When a user is assigned a higher rate channel, fewer users are allowed to simultaneously transmit at high rates to reduce interference. The higher the rate is, the fewer the number of users that can be assigned this rate simultaneously. In the extreme case, only one user can be assigned the highest rate, 153.6 kbps, at any point in time. Consequently, in order to provide fairness, high rate channels are allocated only in short bursts.

The wireless scheduler may assign a supplemental channel with the appropriate rate to a user based on the following potential *secondary*¹ factors:

- *Buffer backlog*: The CDMA2000 1xRTT network deploys a per-user buffer at the BSC which is routinely monitored by the scheduler. A large data backlog is more likely to trigger assignment of a high rate Supplemental Channel for the user.
- *Channel conditions*: Each base station transmits a continuous *Pilot Signal* which is received by all users in the cell. The user then determines the channel condition by computing the Pilot SINR E_c/I_o , where E_c represents the strength of the Pilot Signal received and I_o the interference due to other users and thermal noise. A low value indicates poor channel conditions (or high loss) and vice versa². This value is fed back to the BSC which utilizes this information in deciding what rate to assign. A poor channel may result in a reduction in the assigned rate to minimize channel losses.
- *Sector load* (in terms of number of users): As mentioned earlier, shorter Walsh Codes (at the same power) experience higher interference and also cause more interference. Hence, whenever the BSC transmits a high rate SCH burst to a user, it may prevent other users from transmitting at high rates. Consequently, the BSC must take into consideration the number of other active users in a sector before determining what rate to assign.

2.3 Radio Link Protocol

Apart from wireless channel rate allocation, the other feature of the BSC that can directly affect higher layer performance is the Radio Link Protocol (RLP). The RLP is a NACK-based ARQ re-transmission mechanism developed in order to minimize the losses perceived by higher layers. The motivation for such a mechanism is the high latency on wireless links which can induce large delays before end-to-end recovery mechanisms sense and recover from a packet loss.

¹They indirectly affect application performance by either affecting the behavior of the wireless scheduler or the radio link protocol.

²Typical values for good channels are around -3 to -7 dB, while values less than -11 dB indicate a poor channel.

The BSC maintains an RLP session with each mobile user. The BSC breaks incoming IP packets from the PDSN into radio frames which are then transmitted to the mobile via the BTS. The mobile, on detecting missing (or corrupted) RLP frames requests re-transmission of the corresponding RLP frames.

Chapter 3

Problem Statement

To the best of our knowledge, almost no research has been conducted to quantify the impact of dynamic wireless channel rate allocation mechanisms on TCP performance *in practice* on commercial networks. The only experimental study we are aware of is the one presented in (Chaponniere et al., 2003) which evaluated the impact of bandwidth variation in CDMA2000 networks on the TCP timeout mechanism in a lab environment. The reader is referred to (Altman et al., 2004; Chan and Ramjee, 2002) for relevant *theoretical* TCP models. Even less research exists on quantifying the extent each of the secondary factors influences the channel assigned rate.

It has been traditionally assumed that the RLP re-transmission rate is closely related to the Frame Error Rate (FER) of the channel. In this context, the impact of the Radio Link Protocol on TCP has been extensively researched theoretically (Chaskar et al., 1999; Liu et al., 2002; Barakat and Altman, 2002; Barman et al., 2004) from the perspective of trade-off between reduced error probability and increased latency to maximize throughput. The link layer increases the reliability seen by higher layers through re-transmissions or stronger error correcting codes. Both mechanisms attempt to reduce the likelihood of TCP throttling its sending rate due to packet losses. On the other hand, these mechanisms increase latency since packets are retained longer by the link layer for successful transmission, which in turn can degrade throughput.

Also, to the best of our knowledge, *neither* the RLP re-transmission rate and its dependence on the channel FER, *nor* the impact of the link layer on TCP dynamics have been quantified *in practice* on commercial networks.

As discussed in Section 2.2, the fact that the *data sending rate* of the higher layer protocol could affect the assigned channel rate is crucial because, TCP, the most widely used transport layer protocol, is a *reactive protocol* which adjusts its rate based on feedback from the receiver.

Hence, the system becomes a closed-loop system where both TCP and the wireless scheduler vary their rate based on feedback from the other. This may result in unexpected interactions between the two control regimes, possibly leading to performance degradation.

The objective of this study is to precisely characterize these issues. Through a series of experiments, we evaluate which secondary RF layer factors affect the assigned channel rate and the Radio Link Protocol the most. We also study the impact of both RLP and the wireless scheduler on TCP dynamics.

Chapter 4

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. 4-1 to illustrate the data path¹, as well as measurement points.

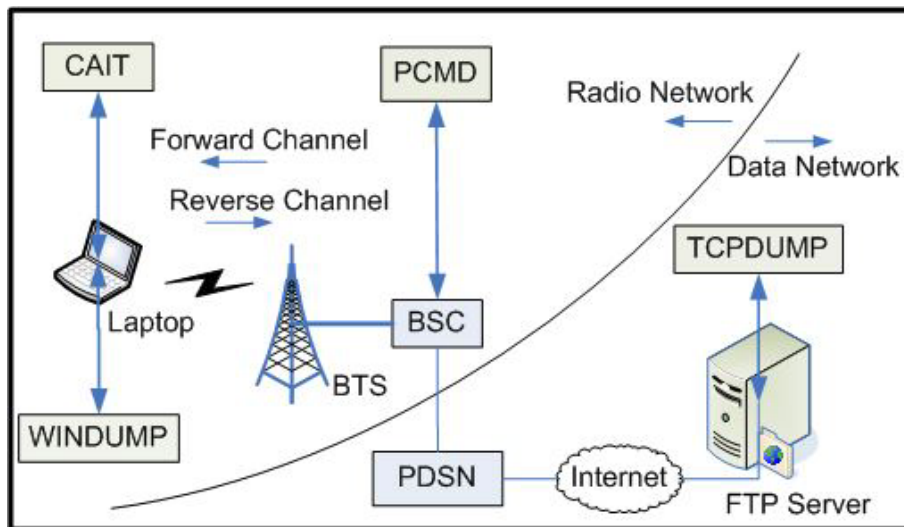


Figure 4-1: Experimental Setup

The experiments can be categorized into two classes. The first class consisted of sending UDP traffic to characterize the wireless scheduler and the Radio Link Protocol. 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

¹The 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%.

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.

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, retransmissions and the assigned wireless channel rate were collected at the laptops using an air-interface analysis tool called CAIT (Qualcomm, 2007). 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 calls², 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).

²A single sector, covered by a BTS, typically had 8-9 (and a maximum of 30) active calls.

Chapter 5

Empirical Evaluation: Methodology and Tools

In this chapter we explain the methodology and tools we used to analyze the measurement data that we collected from our active experiments in a commercial cellular network.

5.1 Time Series Correlation

Recall that our goals are two-fold: one, quantitatively characterize the impact of the various secondary RF factors on the assigned channel rate and RLP and two, perform a similar characterization of the impact of these two primary factors on TCP.

In order to achieve these goals, we must be able to measure the effect of different performance metrics and parameters on one another. For some of the objectives we can rely on standard statistical metrics like the expected mean. However, a large portion of our analysis involves quantifying the correlation between performance metrics and parameters that come in the form of time series, capturing the evolution of different aspects of our system. To tackle this aspect of our study, we require a robust technique to evaluate the correlation between time series.

We chose *normalized mutual information* (NMI) as the correlation measure to accomplish this task. Section 5.1.1 introduces the metric and motivates this choice. We were faced with some technicalities when applying this measure. They, as well as the relevant solutions are discussed in Sections 5.1.2, 5.1.3, and 5.1.4.

5.1.1 Mutual Information as a Correlation Measure

There are numerous correlation measures that have been extensively used in the literature. The most commonly used ones being Pearson's correlation coefficient and covariance. These techniques are limited, however, to only being able to measure linear dependencies. *Mutual informa-*

tion, on the other hand, is a correlation measure that can be generalized to all kinds of probability distributions and is able to detect non-linear dependencies between variables. Consequently, since it was unknown whether the system under consideration was linear or not, we use mutual information in our work to find the correlation between time series.

Mutual information can be thought of as the reduction in uncertainty (entropy) of one variable due to the knowledge of the other. It is mathematically defined as follows. Let X denote a discrete random variable that takes a value $x \in \mathcal{X}$ with probability $p(x)$. The *entropy* of X is given by the well-known definition (Cover and Thomas, 1991):

$$H(X) = - \sum_{x_i \in \mathcal{X}} p(x_i) \log p(x_i) \quad (5.1)$$

The mutual information between two random variables X and Y is then given by:

$$I(X; Y) = H(X) + H(Y) - H(X, Y) \quad (5.2)$$

$$= H(X) - H(X|Y) \quad (5.3)$$

where $H(X, Y)$ represents the joint entropy of random variables X and Y , and $H(X|Y)$ represents the conditional entropy of X given Y .

In order to obtain a consistent interpretation of the correlation measure across different experiments, we utilize the *normalized mutual information* (NMI), defined as:

$$I_N(X; Y) = \frac{I(X; Y)}{H(X)} = 1 - \frac{H(X|Y)}{H(X)} \quad (5.4)$$

To illustrate the intuition behind $I_N(X; Y)$, assume Y completely determines X (i.e., Y captures all the information in X), then $H(X|Y)$ would be close to 0 and $I_N(X; Y)$ would be close to 1. On the other hand, if Y contains no information about X , then $H(X|Y)$ would be close to $H(X)$ and $I_N(X; Y)$ would be close to 0. The closer $I_N(X; Y)$ is to 1, the larger the amount of information that Y carries about X . Note that $I_N(X; Y)$ is asymmetric. Eqn. 5.4 computes the *relative* amount of information that Y contains about X given the entropy of X . If we simply wanted

to compute the normalized mutual information irrespective of direction, we could divide $I(X; Y)$ by $\min(H(X), H(Y))$. In our work, we are more interested in the amount of information that one variable has about another and therefore chose to use Eqn. 5.4.

Next, we address two key issues that we faced in utilizing NMI as our correlation measure: 1) accounting for delays between time series when performing time series correlation, and 2) the discretization of time series to compute the joint and marginal probability distributions necessary for evaluating NMI.

5.1.2 Dealing with Stochastic Delays

In general, when correlating two time series capturing the evolution of two processes, one must consider possible delays between them because of the potential time lag between when a state change in one process actually affects the other. For example, if we were to correlate the instantaneous data sending rate (measured at the sender) with the instantaneous data receiving rate (measured at the receiver), we need to consider the one-way delay between the sender and the receiver (including any possible queuing delays in the network). To overcome this problem, we compute the normalized mutual information between each pair of time series over a wide range of possible time shifts. The mutual information is now defined as:

$$I(X; Y; d) = H(X) + H(Y) - H(X, Y_d) \quad (5.5)$$

where $H(X, Y_d)$ denotes the joint entropy of the random variable X and a time-delayed version of Y if $d > 0$ or a time-advanced version of Y if $d < 0$. The NMI is then defined as:

$$I_N(X; Y; d) = \frac{I(X; Y; d)}{H(X)} \quad (5.6)$$

It is important to note that the time shift between two time series being correlated could be stochastic in nature. This is true, for example, when correlating the data sending rate and the data receiving rate time series. In order to capture any potential correlation, one needs to account for the one-way propagation delay in the network. If the network's round trip time is uniformly distributed then one cannot expect to capture any correlation by considering constant time shifts

between the two time series. In all our experiments, however, the round trip time distribution did have a dominant peak at some fixed constant value as shown in Fig. 5-1. NMI will simply capture any potential correlation between the two time series at that dominant time shift which is sufficient for the purposes of our study.

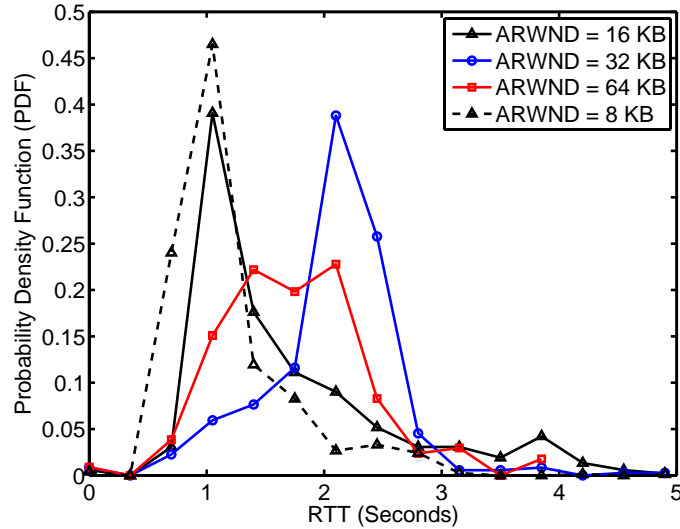


Figure 5-1: PDF of RTT for Different TCP ARWND

5.1.3 Dealing with Real-Valued Time Series

Observe from Eqn. 5.3 that in order to use mutual information to quantify the correlation between two time series, we need to estimate the marginal and joint probability distributions of both time series. Since time series like RTT are real-valued, they must be discretized for this purpose. Towards this end we utilized two techniques for discretization¹.

The first technique we used was proposed by Dimitrova et al. (Dimitrova et al., 2005) which seeks to *bin* any real-valued time series data into a finite number of discrete values. The algorithm assumes no knowledge about the distribution, range or discretization thresholds of the data. It is based on the single-link clustering (SLC) algorithm and aims to minimize the information loss (measured by the entropy), which is inherent to any discretization. The algorithm has also been

¹Most of the time series we were correlating were discrete in nature thus allowing us to model them as *discrete* random variables.

shown to maintain prior correlation between the original time series, which was one of our main criteria for the selection of this technique.

Although we found the technique to be quite effective, the range of time series behavior in our experiments is quite large and there were cases where discretization by this technique fails to capture important properties. This was especially significant in cases with slowly varying signals with sudden variations. Hence, we also utilized standard binning with equidistant bin sizes as an alternate discretization technique for such cases. The bin sizes we chose were: 10kbps for rate time series, and 500ms for RTT time series. For purposes of verification, we compared equidistant binning with the discretization technique by Dimitrova et al. in cases where the latter worked. The NMI values were found to be the same thus confirming the suitability of these bin sizes.

5.1.4 Guidelines for Identifying Potential Correlations

It finally remains to discuss how the Normalized Mutual Information, $I_N(X; Y; d)$, was used to identify potential correlation. Clearly in order to identify any potential correlation, the value of $I_N(X, Y; d)$ needs to be sufficiently *large*. However, like most measures, the NMI is a continuous metric. Hence, it only measures the *strength* of correlation and it is up to the interpreter to judge whether the strength is sufficiently large. In general, it is very hard to justify choice of a specific threshold to categorize the strength of correlation as large or small. This difficulty is compounded even further in our case because: a) discretization of the time series can introduce noise and b) strong correlation between two time series requires the presence of a characteristic delay over which they interact. However, the stochastic nature of the system we study can *spread* the delay range over which the time series are correlated. Both factors can not only *lower* the intrinsic NMI values, but also magnify the potential range of NMI values that a correlated pair of time series may take.

To circumvent this issue, we exploit two aspects in our study. One, when studying any feature like TCP RTT or assigned channel rate, we are only interested in the *relative impact* of various factors on this feature. Hence, we need only focus on the *relative* NMI values. Two, if strong correlation exists between two time series, the NMI values should peak at some characteristic delay

despite the distortion in delay due to stochastic perturbation. Put another way, a sharp NMI peak at a particular time shift² indicates the presence of correlation large enough to overcome potential delay perturbations. This is also corroborated by our experiments, where NMI values between two time series that did not exhibit any strong peak invariably had very low values (compared to those that exhibited peaks).

We found these two guidelines to be quite useful in analyzing the various features on a case-by-case basis without having to resort to choosing a specific threshold value for strong correlation. Specifically, when studying the impact of various factors on a given feature, as a first step the sharpness of the NMI curves as a function of the time shifts helps narrow the potential correlations. The peak value of the NMI curve for each factor is then used to rank the relative strength of correlation.

5.1.5 Correlation Example using Mutual Information

To illustrate the operation of NMI as a correlation measure, we used a simple synthetic data model, shown in Fig. 5.2(a) that is identical to the one in (Gillblad and Holst, 2001). In this model we take a completely random signal $X(n)$ and create two distinct signals from it. The first signal is $X(n)$ delayed by 3 units of time and multiplied by a constant factor of 5.0. The second signal is $X(n)$ delayed by 7 units of time and multiplied by a constant factor of 2.0. We then sum these two signals along with another random signal (noise) denoted by $Y(n)$ to obtain the output signal $Z(n)$. We are interested in finding the correlation between signals $X(n)$ and $Z(n)$. Clearly there should be a correlation between these two signals, one at a time delay of 7 units and the other at a time delay of 3 units. We would also expect the component delayed by 3 units of time to be the dominant one since it is amplified by a larger constant factor of 5.0, as opposed to 2.0. Our results, as shown in Fig. 5.2(b), verify our expectations.

²We use the terms *delay* and *time shift* interchangeably.

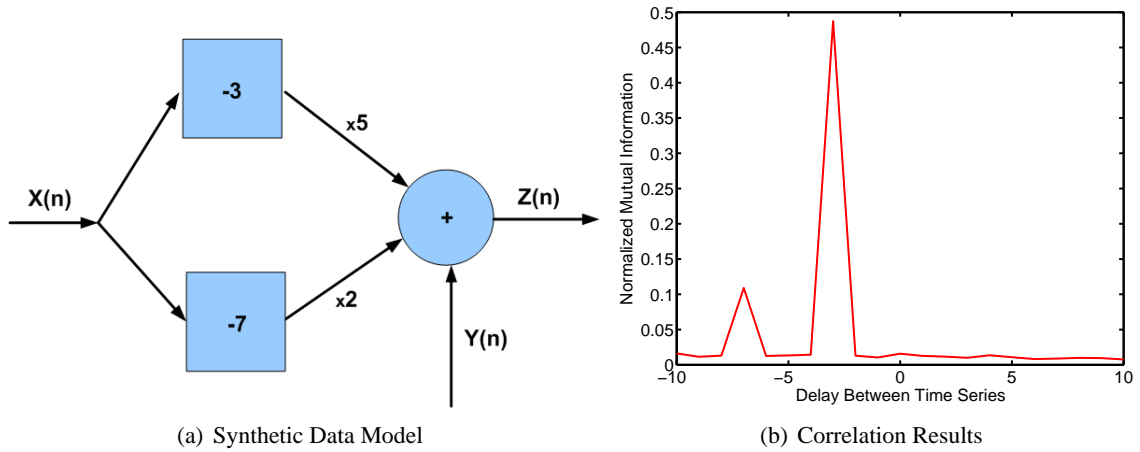


Figure 5-2: Correlation Results Using a Synthetic Data Model

5.1.6 Other Variants of Mutual Information

We evaluated and experimented with several other variants of mutual information: 1) *mutual information of state transitions* where each sample in a time series represents a state and we are interested in capturing dependencies in the *transitions* between these states as opposed to the states themselves, 2) *mutual information of magnitude variations* where we are interested in capturing dependencies in the magnitude changes between consecutive samples, and 3) *mutual information rate* proposed by Gillblad et al. in (Gillblad and Holst, 2001), which is more suited for correlating time series than mutual information but requires making assumptions about the probability distributions of the variables being correlated to obtain meaningful results. The use of these correlation measures only supported the conclusions we made based on normalized mutual information and the results were therefore omitted.

5.2 Analysis at Multiple Time Scales: Wavelets

The last aspect of evaluation that we wish to touch upon is the time scale of different events. Specifically, some RF factors like channel rate, RLP re-transmissions and channel conditions vary over a very small time scale³ while others like the sector load change more slowly. Similarly, TCP

³In our experiments the observable variations are lower bounded by 1 sec.

reacts at the time scale of round trip times, which for wireless links, we show can be in the order of seconds.

Consequently, it is of interest to study the correlation between time series of various parameters at different time scales. For example, an important case that we study is whether changes in TCP sending rate at small time scales are correlated to the rapid variations in the wireless channel rate.

Wavelet is an ideal tool for the purpose of multi-time-scale analysis. We employed the wavelet decomposition strategy outlined by the authors of (Barford and Plonka, 2001) to decompose each time series into low and high frequency signals. The low frequency signal extracts the general slow-varying trend of the original signal. The high frequency signal captures the fine-grained details of the original signal, such as spontaneous variations. Continuing with our example, correlation between the high frequency time series of the TCP sending rate and the wireless channel rate would indicate that the former is affected by (or causes) rapid changes in the latter. Similarly, correlation of the low frequency signals would indicate that the TCP sending rate tracks the assigned channel rate over long time scales (or vice versa). In Fig. 5-3 we show the decomposition of a sample forward channel assigned rate signal. In all our wavelet decompositions, the slow-varying component captured variations over a 32-second duration (which is roughly 10 times the average RTT for many of our connections) and the fast-varying component captured variations over a 2-second duration (which is less than the average RTT for many of our connections).




Figure 5.3: Wavelet Decomposition of a Sample Forward Channel Assigned Rate Signal

Chapter 6

Characterizing the CDMA2000 1xRTT System

6.1 Wireless Scheduler

In this section we present an empirical evaluation of the various factors that affect the behavior of the wireless scheduler.

Recall from Section 2.2, that the wireless scheduler's decisions can be affected by three factors:

- the data sending rate and buffer backlog,
- the channel conditions, and
- the sector load

We quantify the impact of each of these factors in this chapter.

6.1.1 Impact of Data Sending Rate and Buffer Backlog

We performed numerous UDP experiments using constant bit rate (CBR), as well as on-off traffic sources where the on and off durations, as well as the peak data rate were varied. The CBR traffic source allows us to determine if the scheduler tracks the user's sending rate over a long time scale. An on-off traffic source, on the other hand allows us to probe the channel with a bursty-like data source to see if the channel rate scheduler is able to track the data source over a short time scale.

Figure 6.1 plots the average throughput obtained by a UDP connection over the entire experiment's duration (1200secs) versus the different data sending rates for the CBR experiments. Observe that for data rates up to about 60 kbps, the channel tracks the source closely and is able to honor the requested rate. Beyond that, however, the throughput drops significantly. We hypothesize that this is because of a "rate-limiting" mechanism built in the scheduler which preempts users that are persistently greedy. We will discuss this aspect further in Section 6.1.5.

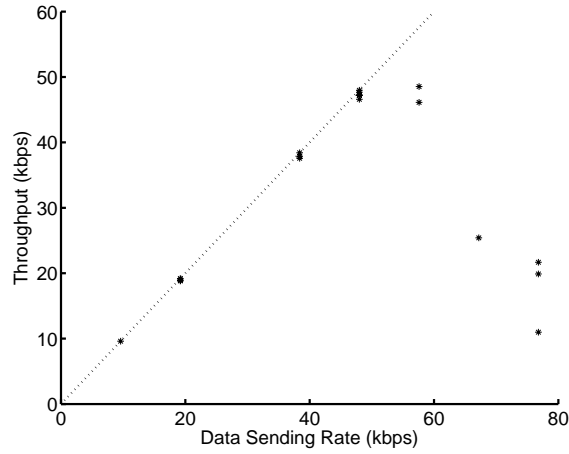


Figure 6-1: Average Throughput for UDP CBR Experiments

Turning next to shorter time scales, in Fig. 6-2 we plot the time series of both the assigned channel rate and the data sending rate for two on-off experiments that had on-durations of 1 second and off-durations of 5 seconds. It clearly shows that the channel rate scheduler is highly sensitive to the data rate even over short time scales since it is able to assign the appropriate rates for every burst of data transmitted. This indicates that the scheduler checks the user’s buffer occupancy at time scales less than 1 second before assigning an appropriate rate.

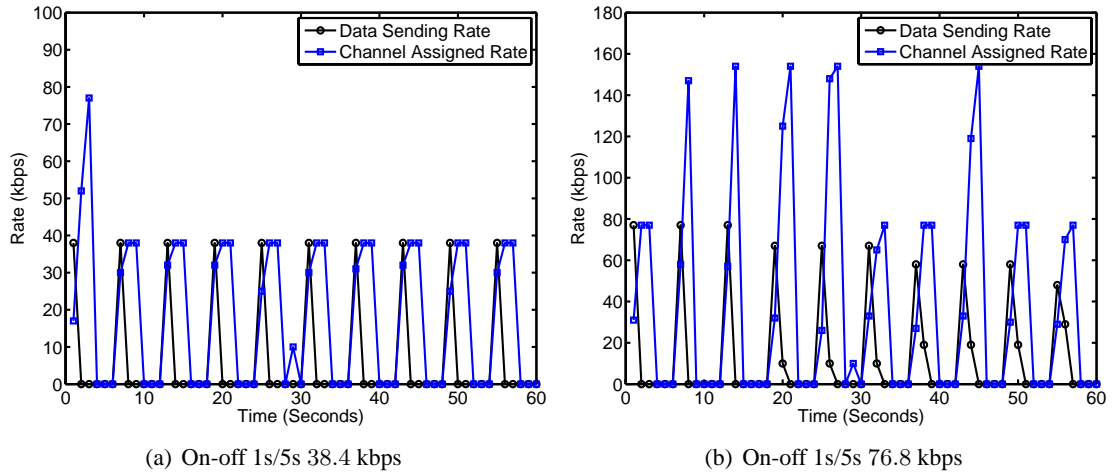


Figure 6-2: Wireless Scheduler Behavior at Short Time Scales

To further illustrate the strong dependency of the assigned channel rate on buffer occupancy

as well as showcase the operation of NMI, we plot the NMI values between the two time series in Fig. 6-3 for various on-off experiments at different time shifts¹. The figure shows several interesting features. Experiments that had low peak rates or short burst (on) periods show high NMI peaks (0.2 and higher) indicating that the scheduler assignment and the data sending rate are strongly correlated, i.e., the wireless scheduler is able to closely track the sending rate. Furthermore, the NMI values have multiple peaks at more than one characteristic delay since the data rate (and hence the channel rate also) are periodic signals, as can be seen in Fig. 6-2.

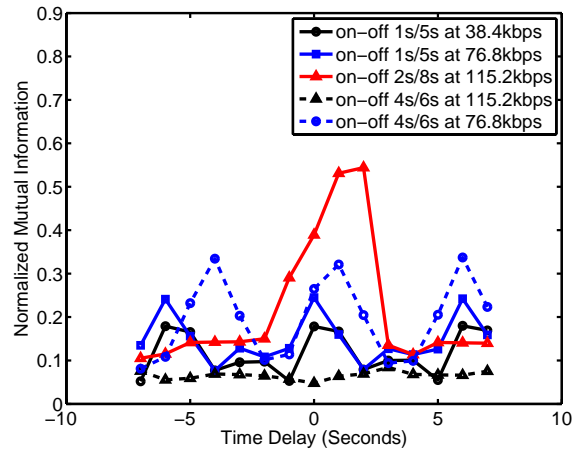


Figure 6-3: Rate Scheduler's Sensitivity to Buffer Occupancy

The NMI correlation also reflects the impact of the scheduler's fairness mechanism. Bursts with high duty cycles and high on-rates cause the wireless scheduler to deny resource allocation in which case the channel assigned rate stops tracking the data source resulting in a drop in correlation. As an example, the experiment with an on-rate of 115.2 kbps but a high duty cycle of 4/6 yields very low NMI values compared to the experiment with the same on-rate but a lower duty cycle of 2/8 which results in high NMI peaks.

Another interesting observation we make regarding the scheduler is that due to the small discrete set of supplemental channel rates, the rate scheduler may assign a much higher rate than the one requested, as shown in Fig. 6.2(b), which could have implications on the stability of the

¹Recall from Section 5.1.2, that when correlating time series we compute the NMI for different delays between the time series.

system.

6.1.2 Impact of Channel Quality over a Short Time Scale

It is well known that channel conditions can introduce significant signal distortion. However modern technologies like CDMA2000 incorporate techniques like rate control (through adaptive coding, modulation, Walsh Code length) as well as power control that allow them to either vary the rate or increase the power to adapt to channel conditions without sacrificing packet integrity. We wish to quantify the role of the former factor, i.e., adaptive changes in the channel rate assigned to the mobile in response to channel conditions, since it can directly affect higher layer protocols.

As explained in Section 2.2, the channel quality in CDMA networks is estimated using the metric E_c/I_o , where E_c is the pilot strength, and I_o , the overall interference. This metric was logged in our experiments by CAIT at a granularity of 1 second. Figure 6-4 shows a sample E_c/I_o signal. In order to quantify the correlation between the assigned channel rate by the wireless scheduler and E_c/I_o , we utilized the UDP CBR experiments, since the data sending rate is constant, independent from the channel rate, and hence should have a minimal influence on the scheduler's behavior.




Figure 6-4: Sample E_c/I_o Signal

Figure 6-5 shows the *maximum* NMI between the assigned channel rate and E_c/I_o for *all* the CBR experiments that we conducted. For each experiment, the maximum NMI was obtained over all time shifts. One can see that compared to the NMI values obtained when quantifying the impact

of the data sending rate, the NMI of E_c/I_o is much smaller (by two orders of magnitude) across all time shifts. Hence, the empirical evidence indicates that in our experiments, the channel condition did not have a significant impact on the assigned channel rate. While the lack of correlation between channel conditions and the assigned rate may be surprising, we believe that this is because of the availability of sufficient sector power, which allows the CDMA network to temporarily boost the strength of the signal to combat adverse channel conditions. In other words, the network adapts to channel conditions via power control rather than explicit rate control.




Figure 6.5: Impact of Channel Conditions on Wireless Scheduler

6.1.3 Impact of Sector Load over a Short Time Scale

The last factor that we analyze is the sector load. The sector load time series represents the number of active voice and data calls that originate and/or terminate in the same sector as our client. We also computed the *maximum* NMI value between the sector load time series and the assigned channel rate for all the CBR experiments. Figure 6-6 shows the *maximum* NMI value between the sector load time series and the assigned channel rate for all the CBR experiments. As with the channel quality, the sector load does not have a significant impact on the assigned channel rate. In most cases, this happened due to low sector load conditions. We also hypothesize that another contributing factor could be the small number of concurrent active data sessions and consequently little cross-traffic from other data users.




Figure 6-6: Impact of Sector Load on Wireless Scheduler

The set of experiments above clearly show the dominant influence of the data sending rate on the wireless scheduler. We further explore characteristics of the wireless scheduler as a function of the higher layer traffic pattern since they will be useful in analyzing the impact of the scheduler on TCP.

6.1.4 General Behavior

Intuitively, when characterizing the rates assigned by the scheduler, we are interested in how *bursty* the scheduler is, for example, whether or not it oscillates rapidly between different rates, and if so, how many *different* rates it cycles through.

To answer these questions, we introduce two metrics: the *burstiness* (denoted by β) and *entropy* (denoted by μ) of the assigned channel rates². In order to compute β , we utilize the wavelet decomposition of the original signal. Specifically, the burstiness is defined as the ratio of the average energy in the fast-varying component to the average energy in the slow-varying component of the original signal. This captures the magnitude of fast variations in the channel rate *relative* to the average rate.

An increase in μ implies an increase in the number of rates being allocated and an increase in β implies an increase in the *rate* of variations. Hence, we can think of μ and β as measures of the

²The channel assigned rate signal, which is a time-average of the FCH and SCH rates over 1 second bins, was discretized using a bin-width of 10 kbps to compute the marginal probability distribution.

scheduler’s stability. If the rate scheduler was tracking the CBR data source perfectly, then both β and μ should be close to zero.

We plot μ and β as functions of the data sending rate for the various UDP CBR experiments in Fig. 6.7. Figure 6.7(b) clearly shows that as the data sending rate increases, the burstiness in the assigned channel rate generally increases. Figure 6.7(a) shows that the entropy of the channel initially increases with the sending rate, till the rate-limiting mechanism kicks in at high data rates (around 60 kbps). Beyond that point, the allocated rate frequently drops to zero, which results in a drop in the entropy of the rates assigned by the wireless scheduler.

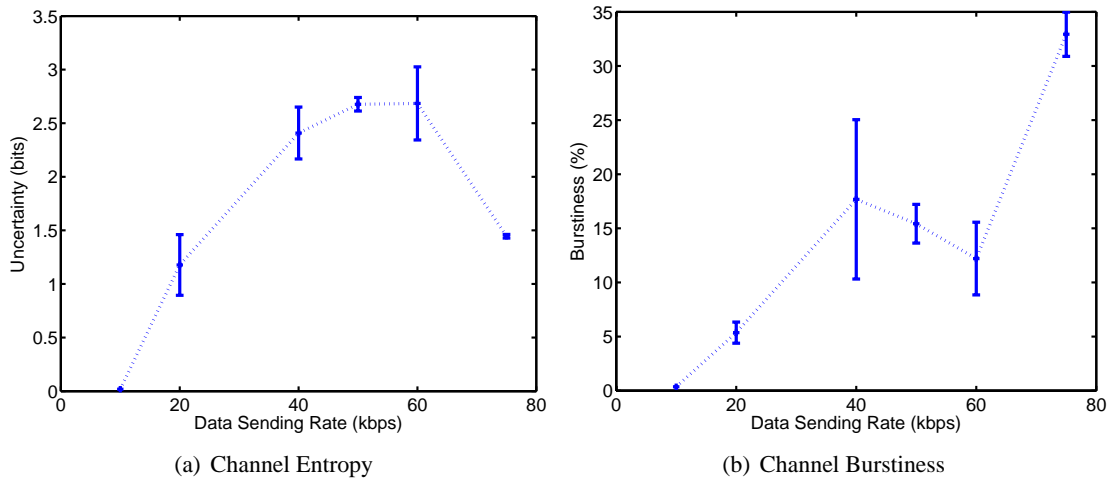


Figure 6.7: Properties of the Wireless Scheduler

To further illustrate the scheduler’s behavior, we plot the time series of the assigned channel rate for some representative UDP CBR experiments in Fig. 6.8. They clearly highlight the increase in burstiness of the channel assigned rates as the data sending rate increases. We next discuss the rate-limiting mechanism as a potential contributing factor to this burstiness.

6.1.5 Rate-Limiting Mechanism

In the previous paragraph, we showed that as the user’s data rate increases, the channel becomes bursty, i.e., it is unable to closely track the sending rate. Furthermore, from Fig. 6.1, we note that at high data rates, the throughput actually *drops* implying that the scheduler stops honoring these rate requests. These observations indicate the presence of some kind of “rate-limiting” mechanism

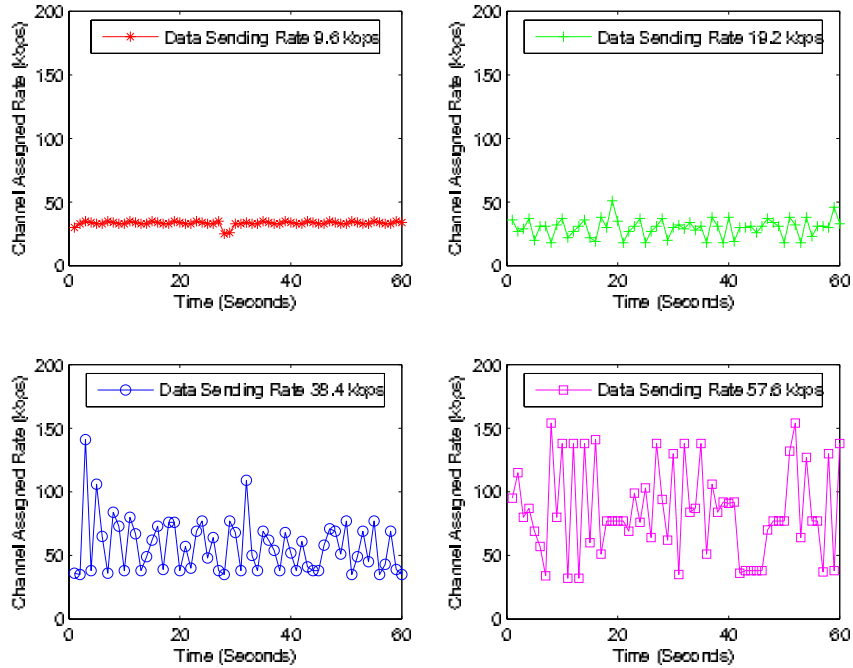


Figure 6-8: Time Series of Assigned Channel Rate for UDP CBR Experiments

in the wireless scheduler to potentially maintain fairness among all the connections being serviced.

Although a precise inference of the mechanism is difficult to achieve solely through experiments, we highlight some specific features of its operation based on our observations. The scheduler periodically monitors the rates assigned to all connections. Connections that persistently request high channel rates (i.e., are continuously backlogged) are likely to be denied rate requests (or are assigned lower rates) for a period of time. Our experiments indicate that the likelihood of denial increases with the intensity of the previous assigned rate as well as the duration of the assignment. This was seen to occur over long time scales in Fig. 6-1 and over short time scale in Fig. 6-3. Specifically, Fig. 6-3 shows that for the experiments with a peak rate of 115.2 kbps, the wireless scheduler stops tracking the sending rate when the duty cycle increases from $2s/8s$ to $4s/6s$, even though the peak rate remains the same.

To further verify the existence of such a mechanism, we probed the channel using an on-off traffic source with a peak rate of 153.6 kbps. The experiment was conducted at 6am on a weekend, to eliminate the effect of sector load (3 users on average), and from a location that is geographically

close to the BTS to guarantee a good channel (E_c/I_o between -3dB and -4dB all the time). Fig. 6-9 clearly shows that after a duration of around 40 seconds, the wireless scheduler stops tracking the user's sending rate and the assigned rate drops to zero, even though the BSC is backlogged.




Figure 6-9: Channel Scheduler Tracking an On-Off Data Source with Peak Rate of 153.6 kbps

6.2 Radio Link Protocol

The Radio Link Protocol is designed for fast recovery of link losses in wireless networks. Traditionally, these losses have been assumed to primarily arise from channel errors that can corrupt the radio frames. Hence, one can expect that the RLP re-transmission rate is highly correlated with the *Frame Error Rate* (FER). To verify this aspect, we analyzed the correlation between the RLP re-transmission rate and the FER. In most of our experiments, the FER was typically very low (zero) indicating strong error-correction and accurate power control. In experiments where there were instances of high FER on the channel, we did find a corresponding increase in the RLP re-transmission rate. However, quite surprisingly, we also observed several experiments where even though the FER was at or near zero, there were significant RLP re-transmissions. In Fig. 6-10 we plot the FER and RLP re-transmission rate from two experiments that highlight both scenarios. Figure 6.10(a), for a low rate UDP experiment, shows that spikes in the FER (upper graph) result in corresponding jumps in the re-transmission rate (lower graph). However Fig. 6.10(b),

for a high rate experiment (153.6 kbps) shows that even in the absence of frame errors, the RLP re-transmission rate is often very high.

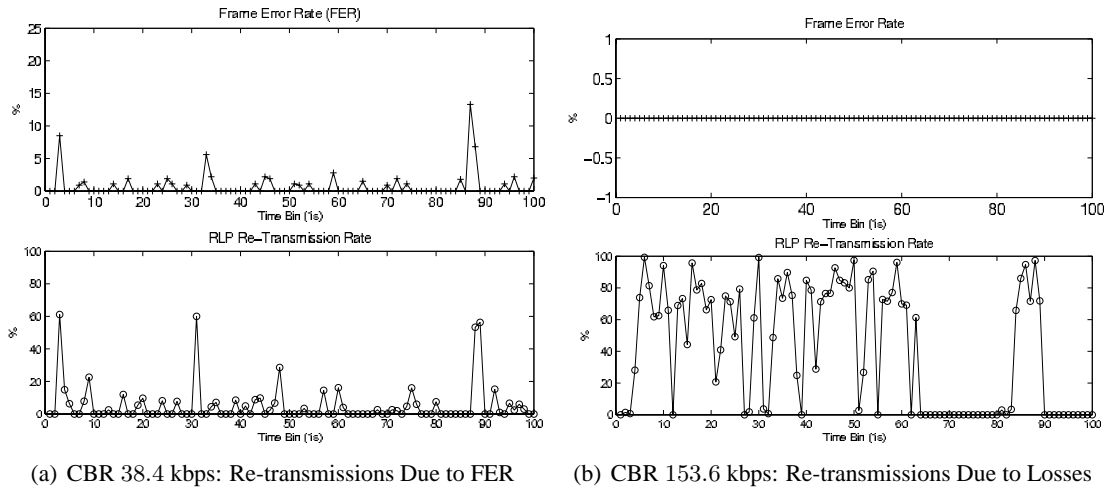


Figure 6-10: Sources of RLP Re-transmissions

The presence of significant RLP re-transmissions, even in the absence of FER suggests that, apart from wireless channel errors, packet losses may be occurring in the back-haul network between the BSC and BTS *after* IP packets are converted into RLP frames, potentially due to congestion from cross-traffic.

Having characterized both the wireless scheduler and the radio link protocol, a general characterization of the CBR UDP connections, characterizing the end-to-end paths, is provided next for completeness.

6.3 UDP Connections

Here we give some preliminary insight into higher layer metrics, important for TCP performance, such as RTT, packet loss, and packet reordering that were experienced by the *CBR UDP connections*. This will aid in our discussions regarding TCP in Chapter 7.

The client in our UDP application responded to received packets with acknowledgments (the server’s transmission rate is not influenced by this). The data packets and acknowledgments had sequence numbers that allowed us to compute the RTT, identify the packets that were lost and infer

any potential reordering of packets.

6.3.1 Round Trip Time

In Fig. 6-11 we plot the Cumulative Distribution Function (CDF) of the RTT for the different data sending rates. The RTT increases by a factor of 8 as the data sending rate increases from 19.2 kbps to 76.8 kbps which indicates the existence of a large buffer at the BSC, causing the queuing delay to increase significantly.




Figure 6-11: CDF of Round Trip Time for UDP CBR Experiments

6.3.2 Packet Loss

In Fig. 6-12 we plot the CDF of the fraction of transmitted packets lost per second for a few different data sending rates. As the data source increases its sending rate the fraction of packets lost per second increases significantly. This indicates that the channel rate scheduler is unable to support high data rates. More importantly, it also indicates that packet losses in the cellular network could be due to congestion rather than wireless losses, even when the data sending rate is below the advertised system capacity (153.6 kbps).

6.3.3 Packet Reordering

In general, packets can be reordered due to the traversal of a flow on multiple paths as a consequence of load balancing or RLP re-transmissions. In practice, however, routers typically utilize

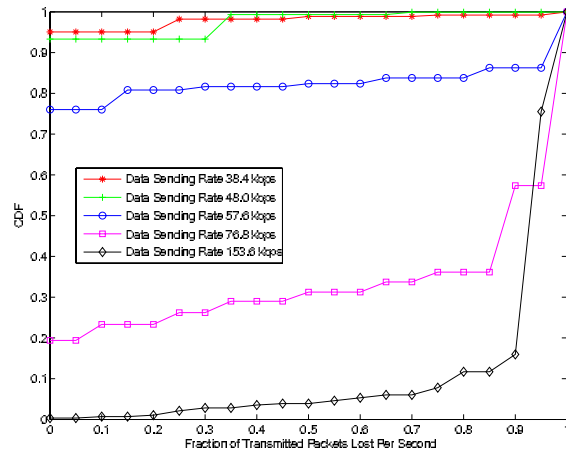


Figure 6-12: CDF of the Fraction of Packets Lost Per Second for UDP CBR Experiments

per-destination load balancing to avoid packet reordering. As a result, packets belonging to the same flow³ are not routed on different paths. The RLP re-transmission mechanism can also reorder packets in the following fashion. Assume packet p_1 was transmitted before p_2 . In the cellular network, both are converted into radio frames before transmission on the wireless channel. Assume they are transmitted back to back. It may then happen that some frames belonging to packet p_1 are delayed due to channel/congestion loss causing repeated re-transmissions. In such a scenario, packet p_2 would be reconstructed first and sent up to the higher layer before p_1 . Analysis of *all* our UDP CBR experiments shows that packets were never reordered in the network.

³A flow is typically identified using the source and destination IP addresses where every unique pair of addresses constitutes a flow.

Chapter 7

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, where the following parameters were varied, namely: a) TCP's advertised receiver window size (ARWND), b) TCP's time-stamp option was enabled/disabled, c) number of users in the sector, and d) user mobility (speed) and location. 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 achieved throughput varied over a large range of values (more than 10%), 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 have to do with dynamic cell-dimensioning of neighboring cells by the network operator which is the focus of our future work.

7.1 Impact of TCP Window Size and Time-stamp Option

The first two configurations we study involve TCP's behavior as a function of the advertised receiver window, both when the time-stamp option (V. Jacobson, 1992) was disabled/enabled. The window size was varied from 8 KB to 128 KB.

7.1.1 On TCP Round Trip Time

We first start with a general characterization of the round trip time (RTT) as a function of ARWND. We infer the instantaneous RTT observed by the sender from the tcpdump log. Every packet that is transmitted is time stamped. For every explicit (cumulative) acknowledgement or sack that is received for a transmitted packet, we compute the RTT. We only compute RTT estimates for packets that were sent out once to avoid ack ambiguity. If duplicate acknowledgements are received we only consider the first one. Each RTT estimate is time stamped with the time at which the packet was *sent*.

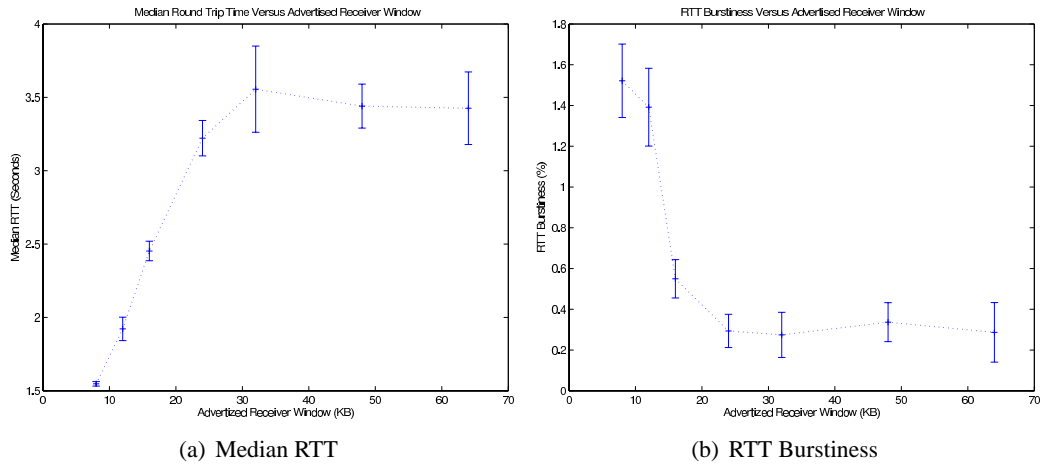


Figure 7-1: RTT Behavior as a Function of ARWND

In Fig. 7-1 we plot the median and burstiness of the RTT as a function of ARWND. The burstiness was computed in the same fashion as in Section 6.1.4. The values reported are in the 90% confidence intervals.

Fig. 7.1(a) shows that as ARWND increases, the RTT increases significantly, indicating high queuing delays, and possibly a high correlation between RTT and the amount of data in flight. We note that the authors of (Biaz and Vaidya, 2003) have shown that in the Internet, the correlation between the RTT and the amount of data in flight (indirectly the receiver's window size) is quite weak. This discrepancy is not entirely unexpected since wireless links have typically far less bandwidth.

Fig. 7.1(b) shows that the RTT burstiness, on the other hand, decreases as ARWND increases since the average RTT increases significantly causing the *relative* burstiness to decrease. This could potentially decrease in the number of spurious timeouts caused by inaccurate RTO estimates (since the RTT estimates are *relatively* smoother).

We next focus on evaluating the impact of RF factors on RTT. The NMI metric is used to quantify the relative impact of the RLP re-transmission rate (which affects time spent at the link layer), wireless channel rate (which affects both queuing delay and transmission time), as well as the buffer occupancy, approximated by the number of outstanding packets (which affects the queuing delay). The impact of the three factors was found to vary as a function of ARWND. As Fig. 7.2 indicates, for small and medium window sizes (8 KB, 16 KB) the wireless channel transmission 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 even when ARWND is very small, and b) in the absence of queuing, the channel rate (and hence transmission time) determines RTT. In other words, any fluctuations in the channel rate have a direct impact on the RTT observed by a TCP connection.

7.1.2 On TCP Throughput

Since RTT directly impacts throughput, we can expect TCP's throughput to be strongly dependent on the assigned channel rates. This is indeed true as shown in Fig. 7.3 which presents the amount of information (NMI) that the channel rate and RLP re-transmission rate have about TCP's throughput. We showed in Section 6.1.1 that the channel rate is influenced by the transport data rate. This implies a strong coupling between the scheduler and TCP. More importantly, it indicates that the rate variations in the wireless channel are in fact *not* completely random, as is commonly assumed in models (Altman et al., 2004; Chan and Ramjee, 2002). Instead, it is *strongly* 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, potentially causing timeouts.

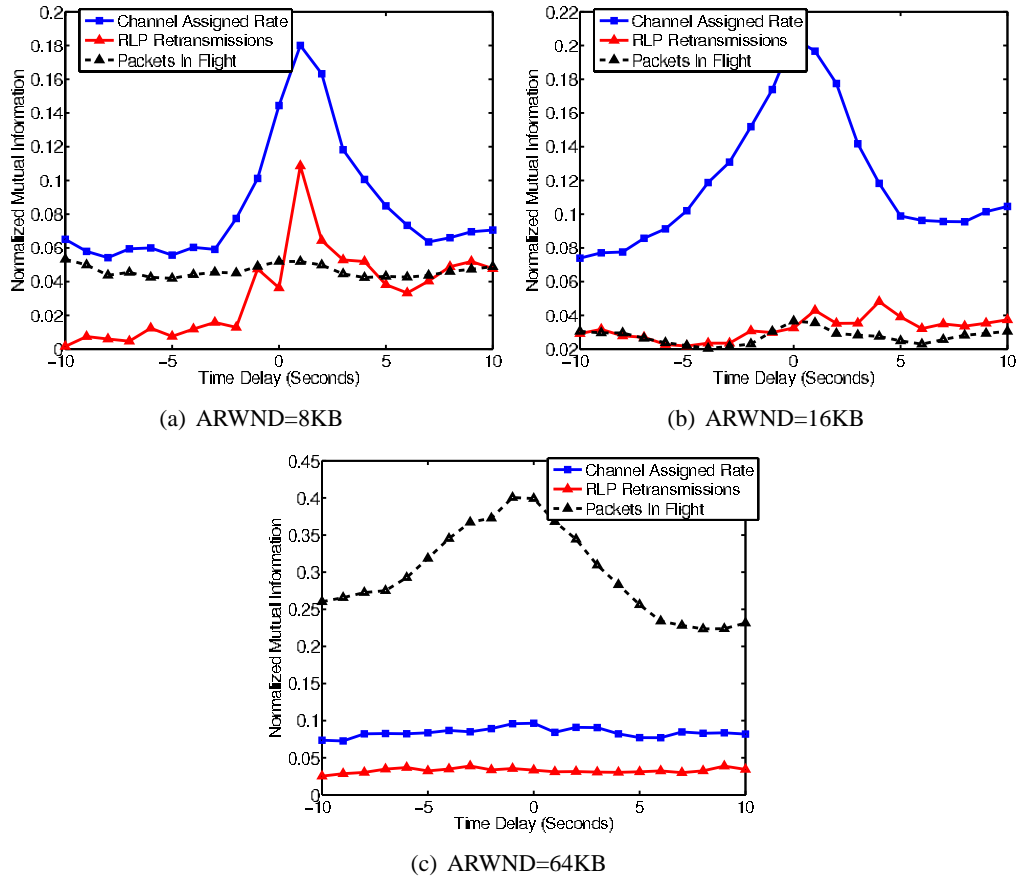


Figure 7-2: Impact of RF Factors and Queuing on RTT

We study the latter two issues in more detail. Figure 7-5 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, 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 computed from CAIT traces¹. 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).

Another intriguing investigation is how well TCP is able to track the available network ca-

¹The CDMA200 rate allocation occurs at small time scales compared to RTT.




Figure 7.3: Factors Affecting Instantaneous TCP Throughput

capacity (*i.e.*, channel assigned rates) over both long and short time scales. To this end, wavelet decomposition was performed on both the data sending rate and channel assigned rate time series to obtain their constituent low and high frequency signals.

In Fig. 7.4, we plot the NMI between the high and low frequency signals and for completeness, the original signals as well. We also plot the low and high frequency signals in Figures 7.4(b) and 7.4(c) respectively. From the high NMI values it is clear, that the sender and the wireless scheduler are tightly coupled over both long and short time scales, with a stronger inter-dependence over long time scales. In other words, the rates assigned by the wireless scheduler are highly dependent on the data sending rate over long time scales, and vice versa.

The second aspect we mentioned was that oscillations produced by such a coupling can lead to a variable RTT and result in inaccurate RTO estimates and spurious re-transmissions (Ludwig and Katz, 2000). To test this hypothesis, we ran experiments where the time-stamp option was enabled. The time-stamp option helps obtain more accurate RTT estimates and thus reduces spurious re-transmissions. The throughput curves when the time-stamp option is enabled/disabled are plotted in Fig. 7.5. It shows that throughput obtained with the time-stamp option enabled is much higher than when it is disabled.

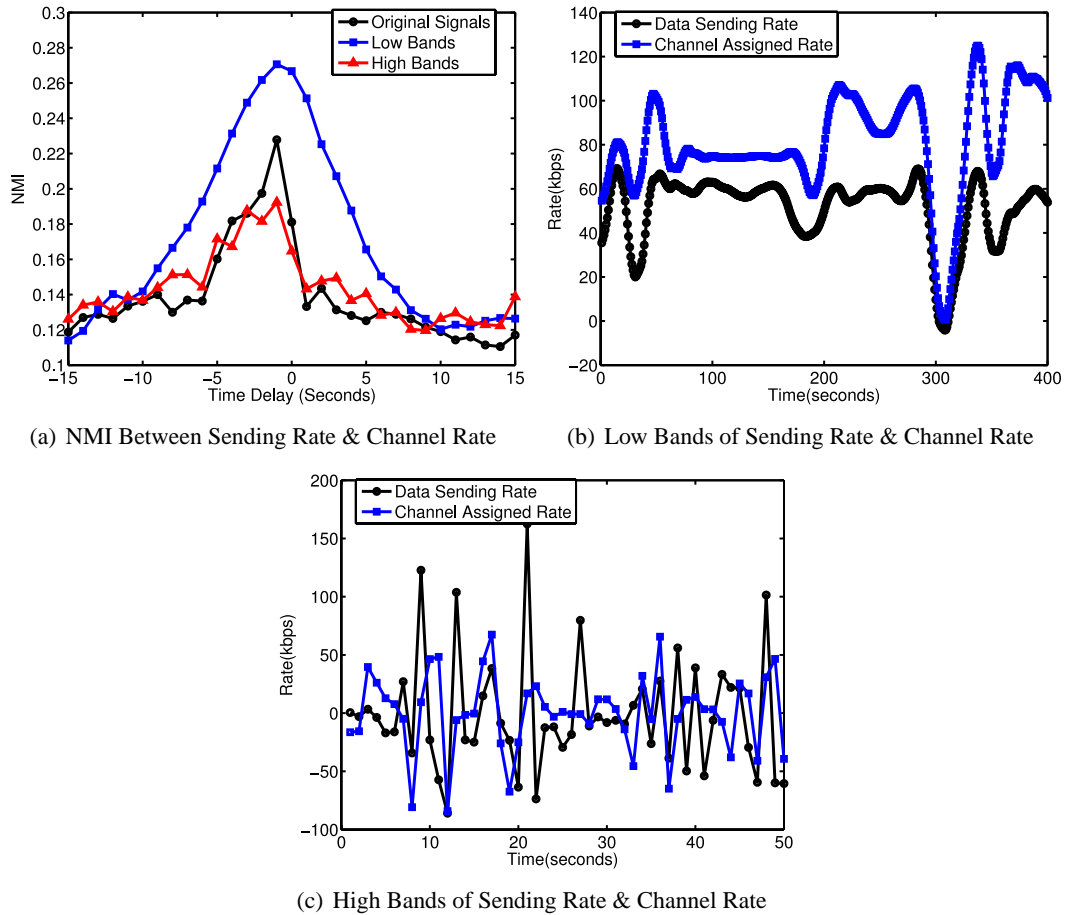


Figure 7-4: Coupling Between Sender and Wireless Scheduler

7.1.3 On TCP Timeouts, Re-transmissions and Packet Losses

Fig. 7-6 plots the number of timeouts, number of re-transmissions and number of packet losses as a function of ARWND when the time-stamp option is enabled/disabled. The number of timeouts was relatively low and the time-stamp option only resulted in significant gains when ARWND was very small. This indicates that *spurious timeouts* is not a problem. On the other hand, when the time-stamp option is disabled, both the number of re-transmissions and the number of packet losses increase significantly as shown in Fig. 7.6(c) and Fig. 7.6(b), respectively. We believe this to be the main factor that causes the degradation in attainable throughput and warrants further investigation. Finally, Fig. 7.6(c) indicates that the number of packets lost increases with ARWND, which is

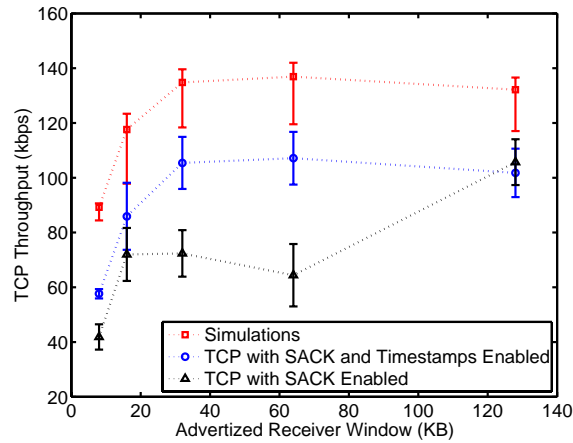


Figure 7-5: TCP Throughput as a Function of ARWND

indicative of *congestion* being the dominant cause of packet loss.

For completeness, we will describe how TCP timeouts were inferred. In general, a timeout can be detected when a packet is re-transmitted and *cwnd* drops to 1 segment² and TCP starts operating in the slow-start phase. We inferred timeouts from the tcpdump log collected at the server. In any TCP version, including the Sack Enabled Linux TCP (Sarolahti and Kuznetsov, 2002) we are using, a packet is re-transmitted either due to 1) the reception of 3 duplicate acks, 2) the reception of an ack including sacked blocks, or 3) timer expiration. In the first two cases the re-transmission occurs *shortly* after receiving an ack from the client. In the third case, on the other hand, the re-transmission *does not have to* occur after receiving an ack from the client. Our timeout inference algorithm is a threshold-based separation of cases 1 and 2 from 3. In Fig. 7-7 we show a histogram of the time delay between the occurrence of a re-transmission and the reception of the last ack from the client across *all* our TCP experiments. The two bars represent the total number of re-transmissions that occurred within *less* than 10ms and *more* than 100ms of the last ack received. There is a noticeable gap between the two bars, between 10ms and 100ms, where no re-transmissions occurred. We therefore used a threshold of 100ms to distinguish between fast re-transmissions and timeout-triggered re-transmissions.

²Some TCP implementations set the *cwnd* to 2 segments after a timeout is detected.

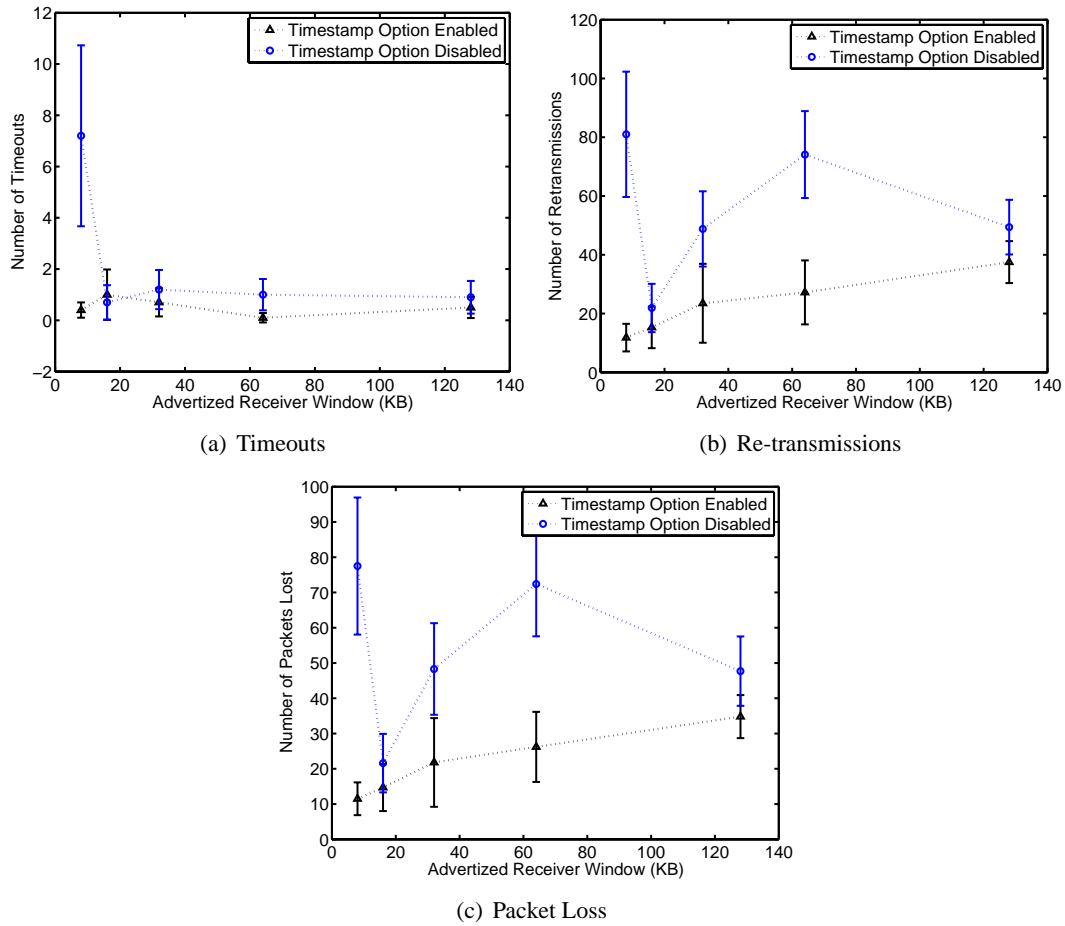


Figure 7-6: TCP timeouts, Re-transmissions and Packet Loss

7.2 Impact of Sector Load, User Mobility and Location on TCP Throughput

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. 7-8 we plot the cumulative TCP throughput, as well as the Fairness Gap as a function of the number of active users. A perfectly fair scheduler would result

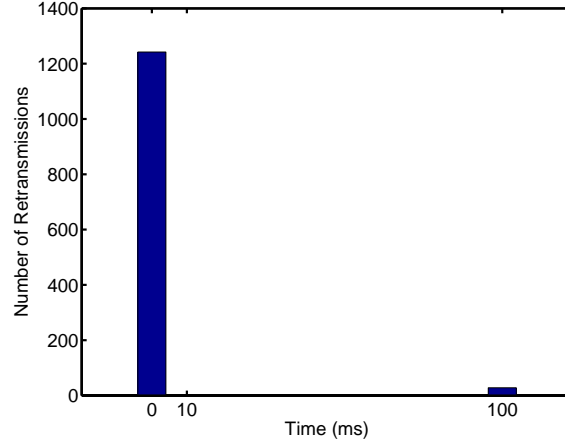


Figure 7-7: Histogram of Re-transmissions Versus Time Delay Since Last Ack Was Received

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:

$$f_{\text{gap}}(y_1, y_2, \dots, y_n) = \frac{\max(y_1, y_2, \dots, y_n) - \min(y_1, y_2, \dots, y_n)}{\min(y_1, y_2, \dots, y_n)} \quad (7.1)$$

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 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-offs³, 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

³A hard hand-off occurs when the BSC completely switches data from one BTS to another.

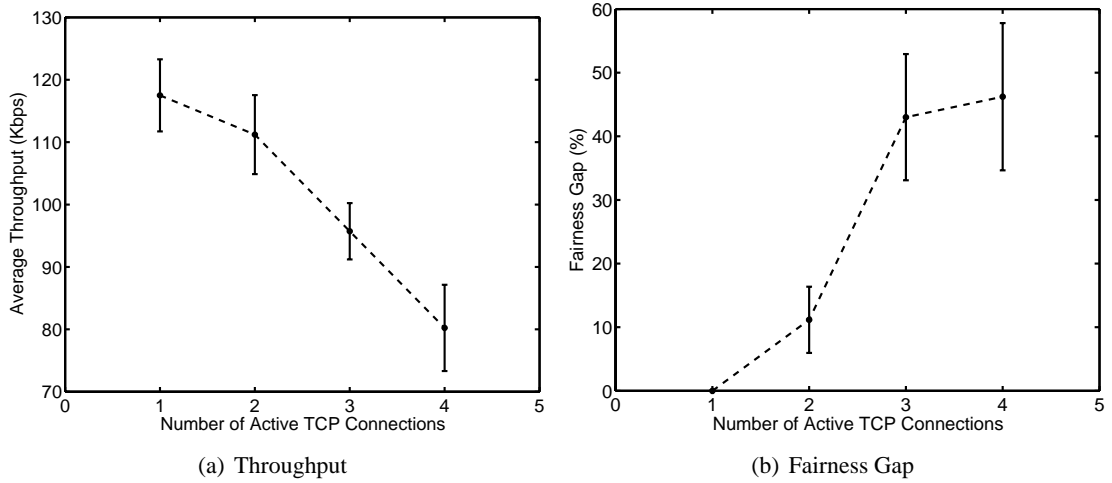


Figure 7-8: Impact of Sector Load on TCP

possible channel) and -16dB (very poor channel). Figure 7-9 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.

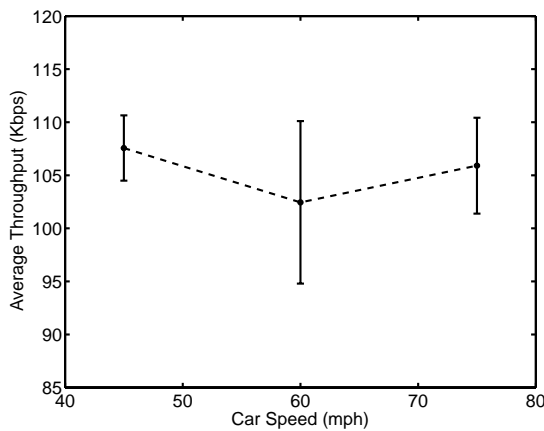


Figure 7-9: Impact of User Mobility (Speed)

The last set of experiments were conducted to investigate the impact of *average (long-term)* channel conditions. In Section 6.1.2 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 time scales*. 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. 7·10⁴. 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.

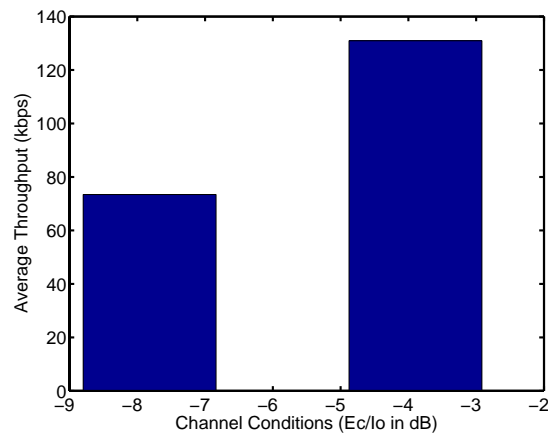


Figure 7·10: Impact of User Location

⁴The RTT, general path characteristics and variation in E_c/I_0 , for both locations, were very similar.

Chapter 8

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 modifications to TCP, as well as, changing user mobility and sector load. By conducting measurements across all three 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 due to channel errors but has minimal impact on TCP's RTT.
2. The wireless scheduler in CDMA2000 networks is, rather surprisingly, 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 because the rate variations are not completely random, it motivates the need for robust RTT estimation to prevent spurious re-transmissions.
4. Mobility in the CDMA2000 network had no major impact on TCP throughput.

Our findings indicate that current state-of-the art cellular networks have significantly different characteristics than wireless channels that were traditionally studied. In particular, robust power control and rate adaptation combined with RLP mechanisms have reduced channel error losses

significantly, much like wireline networks. However, notwithstanding these improvements, significant differences still exist that set apart such networks, namely:

1. We found a large variation in channel capacity depending on the time and location of measurement. These variations are not directly attributable to channel fluctuations. Instead we hypothesize that they are a byproduct of dynamic network dimensioning.
2. The TCP sending rate and the wireless scheduler's decisions were found to be strongly coupled indicating that rate variations seen by TCP in a CDMA network are not completely caused by channel randomness, which should be incorporated into new TCP models.

References

- Altman, E., Barakat, C., and Ramos, V. M. R. (2004). Analysis of AIMD protocols over paths with variable delay. In *Proc. IEEE INFOCOM*, Hong Kong.
- Barakat, C. and Altman, E. (2002). Bandwidth tradeoff between TCP and link-level FEC. *Computer Networks*, 39(5):133–150.
- Barford, P. and Plonka, D. (2001). Characteristics of Network Traffic Flow Anomalies. In *Proc. ACM SIGCOMM IMW'01*, San Francisco, CA, USA.
- Barman, D., Matta, I., Altman, E., and Azouzi, R. E. (2004). TCP optimization through FEC, ARQ and transmission power trade offs. In *Proc. WWIC*, Frankfurt, Germany.
- Benko, P., Malicsko, G., and Veres, A. (2004). A Large-scale Passive Analysis of End-to-End TCP Performance over GPRS. In *Proc. IEEE INFOCOM*, Hong Kong.
- Biaz, S. and Vaidya, N. H. (2003). Is the Round-Trip Time Correlated with the Number of Packets in Flight? In *Proc. Internet Measurement Conference*, Miami, Florida, USA.
- C. Perkins, E. (1996). IP Mobility Support. INTERNET-DRAFT, `draft-ietf-mpls-rfc2002.txt`. (work in progress).
- Chan, M. C. and Ramjee, R. (2002). TCP/IP Performance over 3G Wireless Links with Rate and Delay Variation. In *Proc. ACM MOBICOM*, pages 71–82.
- Chaponniere, E., Kandukuri, S., and Hamdy, W. (2003). Effect of physical layer bandwidth variation on TCP performance in CDMA2000. In *Proc. IEEE VTC Spring*, volume 1, pages 336–342.
- Chaskar, H. M., Lakshman, T. V., and Madhow, U. (1999). TCP over Wireless with link level error control: Analysis and design methodology. *IEEE/ACM Transactions on Networking*, 7:605–615.
- Cover, T. M. and Thomas, J. A. (1991). *Elements of Information Theory*. Wiley-Interscience.
- Dimitrova, E. S., McGee, J., and Laubenbacher, R. (2005). Discretization of time series data. Available at <http://polymath.vbi.vt.edu/discretization/>.
- Gillblad, D. and Holst, A. (2001). Dependency derivation in industrial process data. In *ICDM '01: Proceedings of the 2001 IEEE International Conference on Data Mining*, pages 599–602, Washington, DC, USA. IEEE Computer Society.
- Lee, Y. (2006). Measured TCP Performance in CDMA 1xEV-DO Network. In *PAM*, Adelaide, Australia.

- Liu, B., Goeckel, D. L., and Towsley, D. (2002). TCP-cognizant adaptive forward error correction in wireless networks. In *Proc. IEEE Globecom*, Taipei, Taiwan.
- Ludwig, R. and Katz, R. H. (2000). The Eifel Algorithm: Making TCP Robust Against Spurious Retransmissions. *SIGCOMM*, 30(1):30–36.
- Qualcomm (2007). CDMA Air Interface Tester. www.cdmatech.com/products/cait.jsp.
- Ridoux, J., Nucci, A., and Veitch, D. (2006). Seeing the difference in IP Traffic: Wireless versus Wireline. In *Proc. IEEE INFOCOM*, Barcelona, Spain.
- Sarolahti, P. and Kuznetsov, A. (2002). Congestion control in linux tcp. In *Proceedings of USENIX*.
- V. Jacobson, R. Braden, D. B. (1992). RFC 1323: TCP Extensions for High Performance.
- Vanghi, V., Damnjanovic, A., and Vojcic, B. (2004). *The cdma2000 System for Mobile Communications*. Prentice Hall.