

Experiences in a 3G Network: Interplay between the Wireless Channel and Applications

Xin Liu
liux@ccs.neu.edu
College of Computer and
Information Science
Northeastern University,
Boston

Ashwin Sridharan
ashwin.sridharan@sprint.com
Sprint Research
Burlingame, California

Sridhar Machiraju
machiraju@sprint.com
Sprint Research
Burlingame, California

Mukund Seshadri
mukund.seshadri@sprint.com
Sprint Research
Burlingame, California

Hui Zang
hui.zang@sprint.com
Sprint Research
Burlingame, California

ABSTRACT

We present an experimental characterization of the physical and MAC layers in CDMA 1xEV-DO and their impact on transport layer performance. The 1xEV-DO network is currently the fastest mobile broadband cellular network, offering data rates of up to 3.1 Mbps for both stationary and mobile users. These rates are achieved by using novel capacity enhancement techniques at the lower layers. Specifically, 1xEV-DO incorporates rapid channel rate adaptation in response to signal conditions, and opportunistic scheduling to exploit channel fluctuations. Although shown to perform well in isolation, there is no comprehensive literature that examines the impact of these features on transport layer and application performance in real networks.

We take the first step in addressing this issue through a large set of experiments conducted on a commercial 1xEV-DO network. Our evaluation includes both stationary and mobile scenarios wherein we transferred data using four popular transport protocols: TCP-Reno, TCP-Vegas, TCP-Westwood, and TCP-Cubic, and logged detailed measurements about wireless channel level characteristics as well as transport layer performance. We analyzed data from several days of experiments and inferred the properties of the physical, MAC and transport layers, as well as potential interactions between them. We find that the wireless channel data rate shows significant variability over long time scales on the order of hours, but retains high memory and predictability over small time scales on the order of milliseconds. We also find that loss-based TCP variants are largely unaffected by channel variations due to the presence of large buffers, and hence able to achieve in excess of 80% of the system capacity.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Proto-

cols; C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks; C.4 [Performance of Systems]: Measurement techniques

General Terms

Measurement, Performance, Experimentation

Keywords

3G, Cellular, CDMA, 1xEV-DO, TCP, SINR, Mobility, DRC, Proportional Fair (PF), Measurement, Cross-layer

1. INTRODUCTION

Over the last three years, mobile high-speed networking, in the form of CDMA 1xEV-DO [5], has taken a quantum leap from field trials to nationwide availability in the United States. Compared to peak rates of around 100 kbps that were offered on GSM EDGE and CDMA 1xRTT networks a few years ago, 1xEV-DO offers peak rates of more than 3 Mbps on the downlink and 1.8 Mbps on the uplink even when users are traveling at high speeds. It is worthwhile noting that these rates compare favorably to those offered by current-day DSL with the added incentives of mobility and significantly larger areas of coverage.

1xEV-DO was designed to meet stringent objectives of high-speed data, wide geographical range of coverage, and mobility. This was achieved not through a single technological innovation but rather with the help of several physical and MAC layer optimizations that are unique in commercial wireless networks. Noteworthy in the context of this paper is the *rapid channel rate-adaptation* (at time-scales of around one millisecond) at the physical layer and the *opportunistic scheduling* at the MAC layer to exploit wireless channel fluctuations. These technologies are relatively well understood in isolation and have been shown to provide impressive throughput gains [4] in simulations and controlled trials.

However, a decade of research in 802.11 wireless networks (see for example [19]) has shown that physical and MAC layer optimizations do not always translate into desired higher throughputs in practice. This is because the *transport layer* plays an important role in determining application performance. Indeed, this is an overarching fact in all types of networks, and has motivated researchers to develop and evaluate the performance of a number of transport layer protocols tuned to the properties of specific networks. Furthermore, throughput gains observed in isolation for any particular

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom '08, September 14–19, 2008, San Francisco, California, USA.
Copyright 2008 ACM 978-1-60558-096-8/08/09 ...\$5.00.

optimization during simulations and trials often disappear in practice due to complex real-world interactions that are not easily modeled.

The above arguments drive the core objective of this paper. The 1xEV-DO network has been commercially available only in the last three years. The novel lower layer enhancements have been shown to optimize MAC layer data throughputs. However, they produce a unique environment of rapidly changing channel rates and transmission times that current implementations of transport layer protocols do not specifically cater to. In particular, TCP protocols developed in the past for wireless links [7, 18] assumed *channel errors* as the primary source of performance degradation. With powerful Turbo Codes and adaptive rates, channel error is negligible in 3G networks (verified in our study) and replaced with high delay as well as rate variability, as more likely causes of performance degradation. Although certain aspects of TCP degradation over 3G networks have been observed in simulations, it is unclear how the combination of link layer optimizations and present-day transport protocols perform in an actual 3G network.

Motivated by this, we carried out a measurement-based characterization of the physical, MAC and transport layer performance in a commercially deployed 3G network. Our goals were, broadly: a) to understand the *behavior* of these layers in a real environment and more importantly, b) to determine if the physical and MAC layer enhancements aid in transport layer performance. More specifically, we identified the following key questions and sought to answer them:

1. How does the channel rate in 1xEV-DO vary? Does it change rapidly, is it different for mobile users, and can we predict it?
2. Does the MAC layer opportunistic scheduler provide any throughput gains in practice?
3. How well do current state-of-the-art transport protocols (specifically, different TCP variants) perform on 1xEV-DO? Can they cope with a varying channel?

To answer these questions, we conducted extensive experiments that involved data transfer over 1xEV-DO in the downlink direction, in multiple locations as well as mobile environments. We initiated data transfers using TCP-Reno, TCP-Vegas [6], TCP-Westwood [7] or TCP-Cubic [21] as transport layer protocols. We chose these as our candidate protocols since they capture a variety of transport layer mechanisms and have readily available implementations. In each experiment, we collected detailed wireless channel-related information including instantaneous channel rate, SINR, and packet loss. We also collected information related to the behavior of the transport layer protocols. We cross-analyzed these two data sets to study dynamics at *each layer*. We summarize our findings below:

1. The observed channel data rate varied significantly over *long* time scales on the order of an hour or a day, as well as with location. This is largely in agreement with prior research [23].
2. *Short-term* behavior of channel data rate (within 150 ms) was surprisingly predictable and exhibited significant memory for both stationary *and* mobile users. In particular, we found that a one-step Markov chain can adequately model channel evolution over short-time scales. This indicates that short-term predictors utilized in transport protocols can be effective.
3. Stationary users typically retained one value of channel data rate for more than 40% of the time and transitioned to a

different rate only after a few hundred milliseconds. Mobile channels were much more variable, often changing rate within hundred milliseconds and experiencing sector hand-offs¹ every 20–30 secs on average when the user was traveling at speeds of 50–60 mph. However, rather surprisingly, we found that the *average* channel rates for stationary and mobile users was comparable.

4. The opportunistic channel-aware scheduler typically yields higher gains for mobile users ($\approx 20\%$) as opposed to stationary users ($\approx 5\%$) when compared to a simple Round Robin scheduler.
5. TCP performance was, in general, not significantly influenced by wireless channel characteristics primarily due to the presence of large buffers. TCP-Cubic, TCP-Reno and TCP-Westwood could generally utilize $\approx 85\%$ of channel capacity under good channel conditions and low propagation delays by creating large backlogs. TCP-Vegas, a delay-based variant that controls queue size fared poorer, with a channel utilization of around 75%. The large buffer however incurs a penalty in that delay-sensitive applications were found to fare more poorly in the presence of TCP-Cubic, TCP-Reno or TCP-Westwood as compared to TCP-Vegas. Somewhat surprisingly, we found that under excellent channel conditions utilization dropped to around 80% due to reasons that are as yet unclear.

Our hope with this study is to shed light on the performance of the 1xEV-DO network in the context of commonly used applications, and spur further research on this topic.

The rest of this paper is organized as follows. Section 2 gives the reader a quick primer on the 1xEV-DO system. Section 3 surveys prior work related to this paper's area of study. Section 4 describes our measurement methodology in detail. Section 5 studies the 1xEV-DO physical layer characteristics and details our findings thereof. Section 6 describes our study and findings on the performance of the PF scheduler, while Section 7 investigates on how TCP performance is affected by the wireless channel. Section 8 concludes the paper.

2. 1xEV-DO

The 1xEV-DO mobile broadband cellular network offers peak speeds of 3.1 Mbps on the downlink (base-station to user) and 1.8 Mbps on the uplink (user to base-station) for both mobile and stationary users. The network incorporates several interesting features at the physical and MAC layer in order to achieve these speeds. Of these, we primarily focus on two relatively unique and novel aspects, which have the potential to impact transport layer performance: 1) The variable transmission rates on the downlink, and 2) The opportunistic Proportional Fair (PF) Scheduler. Note that the 1xEV-DO network also incorporates uplink optimizations. But, our primary focus in this paper is on the characterization of the downlink channel and performance.

The *downlink* channel from a base station to the user is a TDMA slotted system, with each slot duration equal to 1.67 ms. At the beginning of each slot, the mobile device computes the perceived signal-to-noise ratio (SINR) with the help of a pilot signal transmitted by the base-station. It then maps this SINR to a channel

¹A hand-off is a phenomenon where a user with decreasing signal quality (for example, due to mobility) shifts from the currently used base-station to another base-station with better signal strength.

data rate (called Data Rate Cover or DRC) based on a certain target Frame Error Rate (usually 1%). The DRC is sent to the base-station, indicating the rate at which the device wishes to receive data in the current slot.

In 1xEV-DO, the channel rate requested by the device, *i.e.*, DRC, can take 15 potential values (or *states*), each indicating a modulation and coding scheme, that translates into a specific desired channel rate in that slot. These rates range from 38.4 kbps up to 3.072 Mbps. We refer the reader to [5] for more details on the exact rates. As mentioned above, in each slot, the computed SINR is mapped to the highest DRC that achieves the target Frame Error Rate. This mapping relation is *dynamic* in that it may be changed on-the-fly based on actual observed FER. The mapping also introduces quantization since it maps a continuous variable (SINR) to a discrete DRC (which can take one of 15 values).

At the base-station *one* user is selected for transmission in the current slot based on the received DRCs. The base-station then transmits to the user in the slot with *full* power at the requested rate. Given the potential dynamic nature of the channel in each slot, the base station utilizes the Proportional Fair (PF) algorithm, an opportunistic scheduling scheme, to decide which user to serve in a particular slot. The PF scheduling decision is based on the past history of users as well as current requested DRCs and tries to exploit the variability of the channel to increase capacity. The PF scheduler is explained in more detail in Section 6.

In order to handle mobility, 1xEV-DO utilizes *fast cell-switching* or 'hand-off'. Each device monitors several sectors and if the serving sector falls below a threshold, it can rapidly request a switch to the next strongest one.

3. RELATED WORK

Today, wireless data networks fall predominantly into two categories: IEEE 802.11 standard based Wi-Fi WLANs, which represents the significant majority and the recent 3G mobile broadband networks that are rapidly proliferating. There is a very large body of work in literature characterizing the channel and performance of the former type of networks (for example, [1, 20]).

However the two networks differ significantly enough in objectives, and, consequently, design principles as well as features at the physical and MAC layer. 802.11 networks aim to offer high data rates, but within limited coverage for only reasonably stationary users. The 3G network is geared to offer high data rates over large areas and support mobility. To achieve this, as outlined in the previous section, it incorporates several unique features that are absent in 802.11 Wi-Fi. The significant differences between the two networks and the relative novelty of 3G networks, creates a clear rationale for the need to characterize and understand the efficacy of these unique features in commercial 3G networks.

A few recent studies have addressed various issues regarding 3G cellular networks. [9, 10] have studied the performance of TCP over a variable rate wireless channel via simulations and proposed ack and buffer management schemes at the MAC layer in order to overcome potential limitations of TCP. [13, 14] proposed new TCP-aware scheduling mechanisms to replace the Proportional Fair scheduler. Simulations were used to show that these new mechanisms improve TCP throughput. It is however unclear as to what extent, if any, the performance degradation of TCP observed in simulations actually occurs in a commercial 1xEV-DO network.

Measurements of performance in actual 3G networks have recently appeared in [11, 17, 23, 8]. The authors of [11] conducted several TCP downloads and probing experiments to characterize the latency, TCP throughput and stability of the 1xEV-DO channel. They observed that TCP throughputs are reasonably satisfactory

(compared to posted peak speeds), and the channel relatively stable but with high latencies. [17] conducted similar studies. [23] conducted TCP and video tests to determine the capacity of a 3G network (they do not specify the type/technology). Their conclusions included the presence of high unpredictability in such networks due to customized engineering of each cell site. [8] was mostly focused on 2.5G networks (such as GPRS and CDMA2000) in terms of application performance with a light-weight evaluation of UMTS-based 3G technology. The authors evaluated various optimization techniques and designed and implemented performance-optimizing proxies.

Our work differs from these previous works in that none of them characterize the underlying channel or *mobility*. In this paper, we undertake a detailed study of the physical channel for stationary and mobile users and try to gauge the extent of its impact on transport layer performance. Furthermore, our work also takes a closer look at the performance of popular transport protocols on a high-speed variable-rate channel in stationary and mobile scenarios.

4. EXPERIMENTAL SET-UP

We carried out our evaluations on a commercial 1xEV-DO network as well as a fully functional test-bed. All our experiments were carried out using Lenovo T-60 Thinkpad laptops running Windows XP and equipped with Sierra Wireless 1xEV-DO data cards as clients and, when required, Dell Edge Servers running Linux Kernel 2.6.21 as TCP and UDP servers. The servers were dedicated to our experiments and had high-bandwidth network connectivity.

Given the objectives of this work, our measurement process involved collection of detailed channel information as well as transport layer information. For the former, in each experiment we collected the SINR and DRC values in *each slot* as well as other relevant radio information, *e.g.*, packet error rate and hand-offs, using a proprietary CDMA measurement tool running on the clients. This allowed us to generate time-series that traced various metrics such as DRC, SINR, and packet loss. At the transport layer, we collected standard Windump and Tcpdump logs at the client and server respectively. The logs were parsed and packet/ack pairs matched at each end in order to generate packet loss and round trip time series at the transport layer.

In order to study the effect of wireless channel characteristics on transport layer protocols, we compared the performance of different variants of TCP. Four Linux servers were configured to each use TCP-Reno, TCP-Vegas, TCP-Westwood or TCP-Cubic. We downloaded large files from either a single server or all four servers concurrently, depending on the scenario. Our experiments can be broadly classified as stationary or mobile.

For stationary experiments, in order to study the physical layer, we passively monitored the 1xEV-DO channel at three different locations (periodic 3 second pings were sent to ensure the traffic channel was not relinquished). At each location, 24 *contiguous* hours worth of channel rate and SINR (broken into hour-long traces for reporting) were logged at the granularity of 1 slot (1.67 ms). Due to equipment constraints, the logging took place on different days. Our three locations covered a large metropolitan region tens of kilometers in diameter and henceforth are referred to as Locations 1, 2, and 3. The distance from the nearest cell tower was about 400, 1000 and 600 meters for Location 1, 2 and 3 respectively.

In addition, we collected traces for 24 contiguous hours from three co-located laptops at Location 1 so as to infer channel behavior of closely spaced devices. We also collected 30-minute wireless traces in the morning and evening at Location 1 over a period of 25 non-contiguous (due to holidays) week-days to study long-term trends.




Figure 1: Average DRC (Left) and Modal Fraction (Middle), over 24 hours at Locations 1,2, and 3; average DRC over 25 days at Location 1 (Right).

When studying transport layer behavior on stationary devices, clients placed in one of the above locations downloaded 250 Mbyte files from each Linux server configured with a particular variant. We conducted experiments to study each TCP variant in isolation as well as concurrent downloads by co-located laptops at the same location to study how the variants share base-station resources. Our experiments were performed on the commercial 1xEV-DO network and a laboratory testbed. The latter comprised of a 1xEV-DO base-station and the four linux servers connected to an internal 100 Mbps network which had no commercial traffic and a minimum round trip time (RTT) of 40 ms. The per-user buffer at the base station was configured to be 64 Kbytes, the same as the maximum default TCP window size in Windows XP. In all testbed experiments, the clients were placed at Location 1.

The mobile experiments were conducted in the San Francisco Bay Area. The typical vehicular speed was [50 – 60] mph and the drive-test experiments lasted at least 30 minutes (it could vary due to traffic conditions). During each experiment, the client would initiate download of a 1 Gbyte file and simultaneously log transport and physical layer information. Similar to stationary tests, we conducted tests where a single laptop connected to a particular server, which was running one of the TCP variants, as well as experiments involving multiple laptops in the vehicle, each connecting concurrently to a server running a different TCP variant. The entire period of mobile trace collection lasted over 2 months, and each scenario was repeated several times to get dependable results. More specific details about the data-sets are covered in the appropriate sections that analyze the data.

5. CHANNEL CHARACTERISTICS

This section presents our findings regarding the nature of the wireless channel observed in 1xEV-DO networks. The objective of the characterization is threefold: first, to explore the differences in channel behavior at different locations; second, to study the impact of mobility on the wireless channel; and third, to characterize the temporal behavior of the channel.

5.1 Location

We first consider the impact of location on the wireless channel characteristics of stationary users. Though the effect of location on the *long-term* signal quality (SINR) of stationary users has been comprehensively documented in literature [23], it is the DRC that ultimately dictates the channel rate in 1xEV-DO networks. Given the dynamic mapping of SINR to DRC, and the associated quantization, it is useful to characterize the behavior of the DRC time-series at different locations, and examine whether it indeed tracks

the behavior of the SINR. Towards this end, we conducted experiments to address the following two questions:

- How does location impact DRC over long time scales? In this context, we consider measures of DRC aggregated over periods of 30 minutes or 60 minutes, and study the evolution of such measures.
- How does location affect the DRC time series at short time scales? In this case, we consider the values of DRC measured during each slot (recall that a slot is 1.67 milliseconds).

5.1.1 Long time-scale analysis

We now consider the behavior of DRC when averaged over an hour, and observed for 24 hours, at Locations 1, 2 and 3; this is plotted in Fig. 1(Left). The impact of location is clearly evident on perceived channel rate. Location 1 enjoyed far higher mean DRC channel rates (2 – 3 Mbps) compared to Locations 2 and 3 both of which experienced average channel rates of about 1 Mbps. We also verified that at this time scale, in spite of dynamic mapping and quantization, SINR and DRC were strongly correlated, with a cross-correlation coefficient above 0.98 at each location.

The spatial variation of mean DRC across our locations is not unexpected, given that these locations are separated by several kilometers. However, we notice from Fig. 1(Left) that even when considering one particular location, the mean DRC varies significantly over time. For example, the mean DRC at Location 1 varies from 2 Mbps, at the beginning of the 24-hour measurement period, to 3 Mbps at the end of the measurement period. This property is further illustrated in Fig. 1(Right), which plots the DRC (averaged over 30 minutes), over 25 days, for Location 1. The mean DRC varies significantly over time, ranging from 500 kbps to 3 Mbps, even at a single location.

5.1.2 Short time-scale analysis

We now investigate the nature of DRC behavior over shorter time scales, specifically, at every slot? We start our analysis by considering the fraction of time spent by the channel in different DRC states. In particular, we identify the most frequently observed DRC state, and focus on the fraction of time spent in that state - we term this quantity the *Modal Fraction*. Fig. 1(Middle) shows the Modal Fraction over each hour, for 24 hours at Locations 1, 2, and 3.

Several interesting observations can be made from Fig. 1(Middle). At all locations, the channel retained a particular DRC state greater than 40% of the time, indicating the presence of a dominant DRC value. To further study this property, we plot the Cumulative Distribution Function (CDF) of the Modal Fraction values for all our




Figure 2: CDF of Modal Fraction for all experiments (Left); Average DRC (Middle) and Modal Fraction (Right), for co-located laptops over 24 hours.

experiments, in Fig. 2(Left). The curve for stationary users shows that the Modal Fraction was at least 60% in more than 60% of our experiments. This further confirms that stationary users have a DRC state that is frequently experienced.

Prior work has shown the unpredictability of SINR in wireless channels over long time scales [23]. Our findings indicate that this is true of DRC as well. In contrast, our observation of a dominant DRC state is an indication that DRC evolution might indeed be predictable *at short time scales*. We explore this further in Section 5.3. We note, however, that this does not give us any evidence of long-term predictability. The actual value of the dominant DRC state also varies with location. And as we shall see next, even at similar locations, there can be significant differences over time.

We noticed earlier, from Fig. 1(Left), that Locations 2 and 3 had roughly the same DRC values, when aggregated over longer time periods. However, Fig. 1(Middle) indicates that Location 2 had a much more ‘stable’ DRC, with the most frequent DRC state being retained between 60 – 80% of the time. The wireless channel at Location 3 was far more variable, visiting any DRC state only 40 – 60% of the time. Thus, even if two locations exhibit similar long time scale behavior, short time scale analysis can identify significant differences between them.

Finally, dynamic mapping and quantization were found to have a more pronounced effect at smaller time-scales. The correlation coefficient between SINR and DRC within each hour of measurement was much lower, albeit covering a wide range: [0.05 – 0.711] - indicating that quantization hides away small SINR variations from higher layers.

5.1.3 Co-located Users

Another unusual observation from our measurements was that *even* co-located laptops can have different channel characteristics. Fig. 2(Middle) plots the average DRC, measured at 1 hour intervals, over a 24-hour period for three *co-located* laptops separated by less than 50 centimeters at Location 1. All three laptops were verified to be connected to the same network sector. As is clearly evident, all three laptops have markedly different average DRC evolution. This difference also extends to the short time scale behavior in Fig. 2(Right), which plots the Modal Fraction on an hourly basis.

This indicates that even nominal separation between laptops is sufficient to provide significantly different data rates. Note, however, that all three laptops always show a significantly large Modal Fraction, reinforcing our earlier observation that stationary users have a dominant DRC state.

5.2 Mobility

A key feature of cellular networks is their ability to support mobility, which is achieved at the physical layer by accounting for Doppler shifts, and at the MAC layer through fast hand-offs. Intuitively, one expects this to significantly affect channel characteristics. For example, as a user’s position relative to the cell tower changes, the perceived signal quality could change. Our goal in this section is to understand this effect; we compare and contrast the properties of the wireless channel for mobile users versus stationary users. In this subsection, we focus mainly on DRC; in later sections, we report results pertaining to the TCP downloads done during these experiments.

Fig. 3(Left) plots the average DRC achieved over *each* mobile experiment and compares it against 24-hour traces from the stationary experiments at Locations 1 and 3. The mobile laptop typically experienced channel rates around 1 Mbps, similar to Locations 2 and 3, but less than the 2 – 3 Mbps obtained at Location 1. Overall, our experience with mobile experiments indicates that average DRC channel rates were quite reasonable compared to stationary users, in contrast to common perception that channels for stationary users are necessarily better. In retrospect, the cellular network design plays a key role in mitigating the difference between the two, at least over long time scales.

We next look at how variable the channel was, considering short time scales (on the order of a few slots). Fig. 3(Middle) plots the standard deviation of the channel while Fig. 3(Right) plots the Modal Fraction for each experiment. Both metrics capture the variability of the channel; unlike average channel rates, they clearly highlight differences between stationary and mobile users. The mobile channel had a far higher standard deviation of DRC rate, more than 600 kbps, when compared to the stationary channels’ standard deviation, which was always below 400 kbps. In addition, the mobile channel occupied a particular DRC state less than 20% of the time. This is in significant contrast to our observation (in Section 5.1.2) that stationary channels’ Modal Fraction was above 40% and often higher. We explore the differences between stationary and mobile channels in more detail in the next section.

5.3 Channel Variability

The previous two sub-sections compared long term average DRC channel rates that arise when users are stationary or mobile and provided an initial look at short-term behavior. The main observation regarding the latter aspect was that mobile users have a more variable channel. There remain several pertinent questions about the *nature* of the variability: what is the range of variability? how fast does the channel vary? and how much memory does it retain?




Figure 3: Mean(Left), Standard Deviation(Middle) and Modal Fraction (Right) of DRC for Stationary and Mobile Traces.




Figure 4: CDF of Entropy values for all experiments (Left); Average Time (in slots) spent in any state (Middle); RMSE between observed Sojourn Times and analytically computed Sojourn Times with a one-step Markov Model (Right).

While interesting on their own account, these questions also have implications for transport layer performance. For example, VoIP applications require low jitter and analytical models of TCP indicate that channel variability and rate of variation affects throughput [2, 3]). In this subsection, we answer these questions by exploring the temporal properties of the channel in more detail.

To begin with, we study the entropy of the DRC distribution: this provides a simple summary of the range and frequency of channel rates observed by a user. Larger entropy values indicate that the channel experiences a larger set of DRC values more frequently. The entropy for the DRC distribution of each trace is computed empirically as:

$$H(X) = - \sum_{i \in DRC} \frac{F(i)}{N_{samples}} \log_2 \frac{F(i)}{N_{samples}} \quad (1)$$

where i represents one of the potential 15 DRC values and $F(i)$ the frequency with which DRC value i was observed. Fig. 4(Left) plots the cumulative distribution function (CDF) of entropy values for all mobile and stationary experiments. We see that most of the stationary user experiments exhibited entropy values less than one, whereas all the mobile user experiments showed entropies greater than two. This shows, as expected, that a mobile user experiences a larger range of channel rates compared to a stationary user.

To provide a more complete picture, we next characterize the *rate* at which users transition between the various DRC values and the *time* spent in each state. Fig. 4(Middle) plots the average number of contiguous slots spent in any state across both stationary and mobile traces, in other words the sojourn time in a state. Observe that the stationary traces have large sojourn times, on the order of a few hundred or more slots. The mobile channel, clearly more

rapidly changing, has a typical sojourn time on the order of a few tens of slots (the Y axis is plotted on a logarithmic scale).

5.3.1 Memory in the Wireless Channel

We now quantify the amount of *memory* in the 1xEV-DO wireless channel. Intuitively, we wish to determine the correlation between DRCs at different time-slots in the wireless channel. This is useful in developing models of data rates and determining their efficacy of prediction.

More formally, we attempt to model the observed data rates as a discrete-time Markov chain [22]. A Markov chain with depth n has the property that the evolution of state at any time slot is independent of the entire past, *given* the state of the past n time slots. Let X_i denote the random variable determining the DRC value (ranging from 0 to 14) in time slot i , and $P(X_i)$ the probability of the DRC being a particular value at time slot i . Then the Markov property implies that:

$$P(X_{i+1}|X_i, \dots, X_0) = P(X_{i+1}|X_i, \dots, X_{i-n}). \quad (2)$$

To model our observations as a Markov chain, we need to determine how many past slots (n) are required to satisfy the Markov property. For this purpose, we follow the approach used in [15], and use conditional entropy as our primary metric. The conditional entropy of a random variable Y as a function of a given random variable X is defined as:

$$H(Y|X) = \sum_{x \in DRC} p(x) \sum_{y \in DRC} p(y|x) \log_2 p(y|x) \quad (3)$$

where $p(x)$ is the probability that X takes the value x . These probabilities are computed empirically from our observations. Intuitively, the conditional entropy quantifies the amount of informa-




Figure 5: Reduction in normalized conditional entropy as a function of past n slots: Stationary Hosts (Left) and Mobile Hosts (Middle); and Normalized Conditional Entropy as function of lag d (Right).

tion the random variable X provides about the random variable Y . If they are highly correlated $H(Y|X) \approx 0$. If they are uncorrelated then $H(Y|X) \approx H(Y)$. Normalizing the conditional entropy by $H(Y)$, i.e., $\frac{H(Y|X)}{H(Y)}$ allows us to verify if the two variables are strongly correlated (close to zero) or uncorrelated (close to one).

To identify the best value of depth to use in our Markov chain model of DRC, we evaluated the conditional entropy of the state in slot i , given history of past n slots $X_{i-1}, X_{i-2}, \dots, X_{i-n}$, or in other words, $H(X_i|X_{i-1}, X_{i-2}, \dots, X_{i-n})$. By increasing n , we increase the amount of past history incorporated. Calculating the reduction in conditional entropy as a function of n allows us to determine the significance of this increase in history, and therefore the channel memory. If the reduction in conditional entropy (or normalized conditional entropy) is significant when we increase n from k to $k+1$, then a Markov chain of depth $k+1$ is a significantly better model than one of depth k ; the converse is also true.

Fig. 5(Left) traces the *normalized* conditional entropy when considering Markov chain models of depth $n = 1, 2, \dots, 4$ slots, for traces collected from stationary hosts. Fig. 5(Middle) plots the same for traces collected from mobile hosts. Note that we have added some randomness to the X axis values to make it easier to distinguish between the large number of points, many of which would be nearly coincident otherwise.

Our first observation, from both plots, is a low value of normalized conditional entropy (less than 20%), for all values of n . This indicates that the history of past states (of even just 1 state) significantly determines the current state. Next, we observe in both plots that the reduction in normalized conditional entropy as we increase n is negligible. For example, the maximum percentage reduction of normalized conditional entropy from $n = 1$ to $n = 4$ is around 10%.

These results indicate that one-step Markov models are typically sufficient to characterize 1xEV-DO data rates for both mobile and stationary users. We verify these results by analytically computing the average duration spent in any state using our Markov model, and comparing it to the empirically derived results from the trace. The analytical method assumes a one-step Markov model and utilizes a transition matrix with 15 rows and columns. We then compute the relative mean squared error (RMSE) between these two methods, for each trace T , as:

$$RMSE_T = \frac{1}{15} \sum_{i=1}^{15} \sqrt{\left[\frac{S_{obs}^i - S_{an}^i}{S_{an}^i} \right]^2} \quad (4)$$

where S_{obs}^i was the empirically computed state duration for state i and S_{an}^i obtained analytically. Low values of RMSE imply that the analytical model is in close agreement with the observed data.

Fig. 4(Right) plots the values of RMSE between the analytically derived and empirically computed state durations, across all traces. We observe low values of RMSE (typically $< 10^{-2}$), providing further evidence in support of a one-step Markov model.

We now examine the correlation between DRC values at time slots separated by much larger time lags, for example, tens or hundreds of milliseconds. This is useful in a practical context, since real world applications typically make observations (of network characteristics) separated by such time lags. Therefore, we ask the question: how correlated is the DRC value in slot i (say X_i) with the DRC in a slot at lag d (X_{i+d}), as a function of the lag d ? Again, we utilize the normalized conditional entropy $\frac{H(X_{i+d}|X_i)}{H(X_{i+d})}$ to quantify this correlation. Fig. 5(Right) plots the normalized condition entropy for both stationary and mobile traces, as a function of lag. For each scenario, we average the normalized conditional entropy across all traces. In both cases, the plot shows that the conditional entropy is less than 50% of the original entropy up to a lag of almost 100 slots. This indicates that slots separated by a wide gap can still retain significant correlation. Note that this does not contradict the one-step Markovian nature demonstrated previously.

5.4 Summary

We now summarize our main findings regarding the 1xEV-DO wireless data rates. We first note that mobile users experience bandwidths comparable to stationary users, although the variability is higher. Similar to prior studies, we found that there is significant variability of channel conditions over long time scales (on the order of hours), depending on location and time, and even for co-located laptops.

In contrast, short term behavior (on the order of a few slots) was much more predictable. In particular:

- All our experiments with stationary laptops showed that a large fraction (more than 40%) of time was spent in one DRC state, indicating the presence of a dominant channel condition.
- The time spent in any particular DRC state is on the order of ten to a hundred slots (on average) and, as a result, channel conditions are highly correlated across time lags of tens of slots. This is favorable news for techniques like channel prediction and rate-estimation that are utilized in transport-layer protocols.

- We find that the short term evolution of the DRC time series can be effectively modeled by a one-step Markov chain, for both stationary and mobile users.

Having characterized the physical layer, we now proceed to study the impact of the 1xEV-DO scheduling mechanism on end user performance.

6. PROPORTIONAL FAIR SCHEDULER

The Proportional Fair (PF) scheduler is an opportunistic scheduling mechanism that aims to increase channel capacity by leveraging channel fluctuations and at the same time be ‘fair’ to all users in the same sector. Broadly speaking, it is based on the philosophy that in an environment with variable channels, not all users will have bad channels all the time. Hence, if delay is not critical (as is true for data), then one can improve system throughput by serving a user only when their channel is ‘relatively’ good.

We make this notion more precise below. Let us assume in a given sector, there are \mathcal{I} users. For a particular $i \in \mathcal{I}$, let $R_i[n]$ be the DRC observed in slot n . Let $A_i[n]$ be its current exponentially averaged throughput which is computed in the following fashion:

$$A_i[n] = (1 - \alpha)A_i[n - 1] + \alpha I_i[n - 1]R_i[n - 1] \quad (5)$$

where $I_i[n - 1] \in \{0, 1\}$ depending on whether user i was served (1) or not (0) in slot $n - 1$. In slot n , the PF scheduler picks the user i^* such that

$$i^* = \operatorname{argmax}_{i \in \mathcal{I}} \frac{R_i[n]}{A_i[n]}. \quad (6)$$

The PF scheduler possesses some useful properties: it is simple to implement, as the name suggests it shares the wireless channel among users in a proportionally fair manner, and it is shown to maximize the log utility function $\lim_{n \rightarrow \infty} \sum_{i \in \mathcal{I}} \log A_i[n]$, under fairly general conditions [16]. However, these properties have only been explored in simulations.

To the best of our knowledge, it is unclear as to the extent to which the PF scheduler actually provides throughput benefits in realistic environments compared to simple channel-unaware schedulers. In this section, we quantify potential gains that can be obtained with the PF scheduler when compared to the simplest blind mechanism, Round Robin, with traces of data rates collected from the 1xEV-DO network. For purposes of simplicity, in this particular comparison, we assumed that the user data queue always has ‘data’ to send, *i.e.*, presence of a *perfect* transport protocol.

Our comparison was carried out using up to four *co-located* laptops at Location 1 as well as when mobile. Note that in the latter case, users do not actually *stay* in a single sector. From that perspective, this comparison is biased in favor of the PF scheduler since the PF requires a certain amount of time to converge to fair sharing and the maxima of the utility function. We shall address this aspect in more detail at the end of this section.

In each experiment, we collected RF traces from all four laptops and used the DRC from one million slots for comparison. α was set to $1/1000$ which is the recommended value in practice [12]. Fig. 6 plots the system gain as a function of the number of users in the system. We ran ten simulations with random start slots for users and found negligible difference in results. Hence error-bars are not shown. The system gain is defined as T_{PF}/T_{RR} , where T_S is the total system throughput under scheduling discipline S . Interpreting the figure, gains for stationary users is minimal (less than 5%). This can be attributed to the low variability of the channel observed in practice. In the case of mobile channels which have higher variability, the opportunistic nature of PF comes into play, allowing it post




Figure 6: System Throughput Gain of PF over Round Robin as a function of number of users.

higher gains of around 20% compared to Round Robin scheduling. We can see that the gain also increases with number of users. We also evaluated the Jain’s fairness index for all scenarios and found no discernible difference between PF or RR across all scenarios. Both had values of 0.99 indicating excellent fairness.

For the mobile experiments, the above comparison implicitly assumed that the same PF scheduler was serving all users. In practice, since the mobile users experience hand-offs, this would not be the case. At each hand-off, the new serving sector starts serving the hand-off user as a new user. To explore this scenario, we determined the average time a mobile spends in a sector *between* hand-offs and re-computed the throughput gains using this timeframe. In our experiments, a mobile typically experienced about 40 hand-offs. The maximum time spent between hand-offs varied from 114 to 140 secs (around 70,000 1xEV-DO slots) while the average time was around 20 – 30 secs (around 20,000 1xEV-DO slots). Within this framework, the throughput gains were found to reduce dramatically to 5% and 8% respectively. The results indicate that during hand-off, it is beneficial to retain the user PF state ($A_i[n]$) in order to retain advantages of the PF scheduler.

7. TCP OVER 1xEV-DO

We next consider the impact of the 1xEV-DO wireless channel and network characteristics on TCP, since it is the most commonly used end-user transport protocol. Our goal is to gain insights into which factors affect transport-layer performance over 1xEV-DO. Therefore we perform experiments with a variety of different TCP variants and compare their performance. Specifically, we evaluate TCP-Reno, TCP-Cubic, TCP-Vegas and TCP-Westwood [21, 6, 7] in this study. Our choice of these four variants is motivated by the fact that they capture a variety of different TCP algorithms, optimized towards different goals - loss-based and delay-based congestion control, high-speed and shared-medium variants. Moreover, these four variants are available as free open-source Linux kernel modules and facilitate easy experimentation.

Our experiments involved download of large files for each TCP variant to obtain sufficient physical layer and transport layer statistics. Details of the experimental set up are presented in Section 4. The metrics used for judging TCP performance, are the TCP goodput, and the mean and standard deviation of the excess delay experienced by the TCP traffic. (We define excess one-way delay as the residual delay after subtracting the smallest delay sample and use it to minimize problems with clock synchronization).




Figure 7: EX_{SINGLE} : (Left) TCP goodput and mean DRC achieved with 4 different TCP variants in 3 experiments each. The total height of the bars is the mean goodput (and mean DRC). The shaded portion is the standard deviation observed. (Middle) Mean delay versus the TCP goodput (as a fraction of the mean DRC). (Right) Identifiers of base stations connected to by different mobile users on the same route, different days. The base stations were assigned identifiers in the same order that they were seen. Note the similarity of hand-offs. The difference in speeds on both days accounts for an increasing clock skew.

We performed experiments with these TCP variants in various different settings: mobile and stationary; all four variants simultaneously, or separately. We report our findings in each of these scenarios.

7.1 Preliminaries

Our first set of results focus on each TCP variant in isolation, *i.e.*, a single laptop downloading a file from a single server via the commercial 1xEV-DO network. While we performed this experiment in both stationary and mobile scenarios, our findings were similar, and we report the TCP performance only for the mobile scenario here.

We shall henceforth refer to this set of experiments as EX_{SINGLE} . Fig. 7(Left) shows the TCP goodput from three experiments run with each variant of TCP². Recall from Section 5 that the DRC, which represents a limit on TCP performance, can vary widely across these time scales. Therefore, we also plot the mean DRC for the experiments corresponding to each of these variants.

We observe that TCP-Vegas achieves a much lower throughput than the other variants. Since this could be either due to difference in wireless channel rates or lower channel utilization it is more useful to examine the TCP goodput as a fraction of the mean DRC and also as a function of the end-to-end observed excess delay. Fig. 7(Middle) presents such a scatter plot for each experiment. The experiments with TCP-Vegas lie in the lower left corner which indeed indicates that the low throughput is due to poor utilization of the channel rather than the wireless channel rates. TCP-Vegas also has the lowest packet delays in keeping with its nature : it responds to increasing delays by reducing queue sizes, a possible cause of the low utilization. The other variants generally have higher goodput fractions at the expense of higher packet delays (indicating higher queue build-ups) with TCP-Cubic having the largest latency. However, we note that across all variants, the maximum channel utilization is only 70%.

This experiment highlights two important factors which could significantly affect our TCP experiments. First, there could be cross-traffic and wired network bottlenecks inside the commercial 1xEV-DO network that varies across experiments, making it difficult to compare the TCP variants and the role of the channel. Second, the queue sizes in the 1xEV-DO network could affect the

performance of different TCP variants. We address the issue of mitigating cross-traffic effects in the subsequent sections and also discuss the role of the queue size.

In order to reduce the impact of cross-traffic, we performed mobile experiments with four laptops downloading files at the same time, each from a server configured with a different TCP variant. If all four laptops are connected to the same base station, then they share the same cross traffic effects. In order to test whether this condition would be met, we conducted experiments with different laptops driving along the same routes at different times. Fig. 7(Right) shows the result of one such trial, plotting the identifier of the base stations to which the two laptops were connected. We see that during almost all time instants, both laptops are connected to the same base station.

Therefore we proceed with the assumption that four laptops simultaneously downloading files, in close proximity or in the same vehicle, will share the same cross traffic effects. While we performed such experiments in both stationary and mobile settings on the commercial 1xEV-DO network, we only describe the mobile scenario in this paper. These experiments are described in Section 7.2. However, these experiments still do not isolate the effect of bottlenecks and congestion in the wired network from those of the wireless channel. Hence we also performed experiments on a dedicated 1xEV-DO testbed that was isolated from the commercial network. This is described in more detail in Section 7.3.

7.2 Co-located Mobile Laptops

In this section, we describe results for four mobile users simultaneously downloading a large file from servers configured with different TCP variants. Each experimental drive-test was conducted four times and henceforth shall be referred to as EX_{COLO} . We plot the results from the four experiments in Fig. 8.

Fig. 8(Left) shows the TCP goodput and average excess packet delay for each variant, as well as the mean DRC during the experiments. Note that this plot is normalized so that the maximum value of each metric shown is one. We see that the mean DRC is generally the same for all variants, but TCP-Vegas has significantly lower throughput and TCP-Cubic the highest throughput amongst the variants. The packet delays exhibit a similar relation.

Given the varying nature of the wireless channel conditions, even though the co-located laptops had the same mean DRC, we explore if differences in channel dynamics between the laptops could be responsible for the relative performance results. Therefore we

²The height of the unshaded bar represents the average and the shaded portion represents the standard deviation.




Figure 8: EX_{COLO} : (Left) Performance and network conditions (Middle) Standard deviation of the delay vs. TCP goodput. (Right) Downlink loss rate vs. TCP goodput.




Figure 9: Overall mean and a few individual per-experiment sample Frame Error Rates for EX_{COLO} (Left); and for EX_{TBED} : Mean DRC versus the TCP goodput (Middle), and Standard deviation (Right) of DRC versus fraction of the mean DRC that TCP exploits, for each TCP variant.

look at delays and losses of the individual experiments comprising EX_{COLO} . Note that these metrics are measured end-to-end between the laptop and server, since we were constrained in placement of measurement points.

Fig. 8(Middle) plots the standard deviation of average end-to-end excess delay versus the TCP goodput for each individual experiment. We observe that the standard deviation of delay is *always* much lower for TCP-Vegas, when compared to the other variants. This indicates that the delay behaviors we observe is due to different queuing properties for the different TCP variants, rather than wireless channel variations. Similarly, we plot the end-to-end loss rate for each experiment, in Fig. 8(Right), and observe little correlation between the loss rate and throughput. Further, we extracted the frame error rates (FERs) for EX_{COLO} from the CDMA measurement logs and show them in Fig. 9(Left). We observe FERs consistently less than 1%. This is well within the target FER that 1xEV-DO was designed for, implying that built-in link-layer mechanisms involving turbo-coding and hybrid ARQ [5] reduce the packet error rates to negligible values. We verified this in a separate, but extensive set of experiments which indicated that the channel loss rate in 1xEV-DO was consistently lower than 0.01%.

Thus, our investigations lead us to conclude that difference in channel conditions are *not* responsible for the relative performance of the TCP variants. Instead, this reinforces our earlier indications (from Section 7.1) that queuing effects govern the differences in the performance of the algorithms. The delay-based algorithm used by TCP-Vegas yields smaller and more stable queue sizes, but results in lower goodputs in the context of a varying wireless channel.

TCP-Reno, TCP-Westwood and TCP-Cubic have a more aggressive algorithm that can lead to larger queues with higher variance, but yields higher goodputs. Though queueing dominates, we note that the varying nature of the wireless channel has a subtle effect: it leads to larger configured limits on queue sizes and causes a wider range of delay variation. TCP-Vegas, by virtue of maintaining smaller, more stable queue sizes has a less variable delay. The other variants, TCP-Cubic in particular, obtain higher throughputs by sending traffic more aggressively, and exhibit larger, less stable queue sizes as a by-product. In passing we make a note regarding fairness. Even though the average channel rates are similar, TCP-Cubic typically gets a higher fraction ($> 30\%$) of the system bandwidth, while TCP-Vegas gets the least ($< 20\%$) with TCP-Reno and TCP-Westwood getting around 25% of bandwidth.

7.3 TCP variants on the Testbed

In our previous experiments, it is still conceivable that wired network congestion and bottlenecks affected our observations. We further mitigate this factor by performing experiments on the laboratory testbed described in Section 4. Furthermore, we ran experiments only at night to minimize any possibility of background traffic causing wired-network bottlenecks.

Our goal was to evaluate TCP efficiency, verify our earlier findings, and further examine whether the relative performance of the TCP variants was caused by difference in experienced channel conditions or algorithmic behavior, as we hypothesize. For the first objective, we conducted four file downloads with each TCP variant. Across these experiments, TCP-Reno, TCP-Cubic, TCP-Westwood




Figure 10: For EX_{TBED} : (Left) TCP Goodput, Mean DRC and the average fraction of packets that were retransmissions (height of the bars normalized w.r.t the maximum in the group). The colored portions show the magnitude of the standard deviations. (Middle) End-to-end loss rate, and (Right) Mean delay.

and TCP-Vegas were found to yield average channel utilizations (ratio of TCP throughput and mean DRC) of 0.86, 0.89, 0.85 and 0.75 respectively. Although this indicates that the former three are reasonably efficient (given protocol overhead), it is somewhat expected given the small propagation delays, large buffer and minimal cross-traffic on the testbed. A closer examination revealed that the high efficiency was present only for lower average channel rates around 1.5 Mbps, which occurred in a majority of the experiments. Surprisingly, in excellent channel conditions (average rate > 2.4 Mbps), the efficiency dropped to around 80% for all three variants. Analysis of tcpdump logs indicate the presence of numerous re-transmissions when channel rates were higher. We hypothesize that these are spurious time-outs induced due to large relative rate fluctuations for large channel rates (the large buffer mitigates congestion loss and analysis of wireless logs reveals no channel losses) that prevent TCP from efficiently utilizing the channel. We plan to conduct a more detailed analysis of this aspect in further work.

Next, in order to determine relative performance of the variants, we conducted experiments where 4 laptops connected to the testbed base station simultaneously downloaded a large file, each from a server configured with a different TCP variant. We refer to this set of experiments as EX_{TBED} and summarize the results from eight experiments below. In Fig. 9 (Middle and Right), we plot the performance for each TCP variant (in terms of the goodput as a fraction of the mean DRC) versus channel conditions experienced by each variant (mean and standard deviation of DRC). From the Middle plot we observe numerous data points for all variants with the same mean DRC but significantly different performance, and the Right plot shows that the variations in DRC have little correlation on relative performance.

Similarly, Fig. 10 plots the performance of the TCP variants (for EX_{TBED}) as a function of other possible indicators that difference in channel conditions could be responsible for performance: loss rate (Middle), and delay (Right). In both these plots, we notice no significant correlation between the plotted metric and the TCP goodput. Fig. 10 (Left) shows the average number of TCP retransmissions for the different TCP variants. Again, we notice TCP-Vegas shows few retransmissions while TCP-Cubic shows the most, indicating that TCP-Cubic is far more aggressive and builds up larger queues, whereas TCP-Vegas favors smaller queues and therefore fewer retransmissions.

All the above results from Section 7.2 and Section 7.3 lead us to hypothesize that the difference in observed performance of the congestion control algorithms is dominated by queuing rather than




Figure 11: Performance of Simultaneous VoIP and TCP sessions for different TCP variants.

the wireless channel, although the latter introduces variability. In other words, the experiments point to a scenario where the base-station possesses a large *per-user* buffer that absorbs and mitigates the impact of channel fluctuations, allowing high channel utilization with little congestion losses. This is indeed the case as noted in Section 4. The reasoning behind large buffers stems readily from the nature of the 1xEV-DO channel. Given that channel rates can vary from 38.4 kbps to 3.1 Mbps, it represents a wide range of Bandwidth-Delay Product (BDP). For example, an RTT of (say) 100 ms and 1500 byte packets translates into a BDP that ranges from 1 to 25 packets. As shown in Section 5.1.2, the channel state can fluctuate over a few tens of milli-seconds. By having a buffer larger than the peak BDP, the system potentially allows high utilization without incurring congestion loss. Analytical models in [3] have been shown to support a similar hypothesis. Though large buffers help mitigate congestion loss via channel variability, they also have negative side-effects as shown next with a case-study.

7.4 Impact on Applications : A Case Study

Our finding that behavior in commercial 1xEV-DO systems is dominated by the buffer have important ramifications for end-user applications. Consider the following scenario: an end-user laptop with two simultaneous applications, one large FTP download and one Voice-over-IP session (VoIP), or some other delay-sensitive application. Now recall that all traffic from an end-user device shares the same queue at the base station. Then, if the TCP variant was aggressive in order to improve throughput (e.g., TCP-Cubic), one could expect the VoIP session to suffer degraded performance. In

comparison, TCP-Vegas would receive lower throughput, but allow the VoIP session to receive better performance.

We verified this conjecture via an experiment on the testbed, where a client downloaded a 250 Mbyte file using each TCP variant and simultaneously received a low rate 9.6 kbps stream. The end-to-end stream delay and TCP throughput with each variant are reported in Fig. 11. The stream delay is least with TCP-Vegas as would be expected, but at the expense of low TCP throughput. The other variants yield higher throughput at the expense of longer stream-delay. TCP-Cubic is at the end of this spectrum offering high throughput but significantly higher latency than TCP-Vegas. This is clearly an impact of the large buffer *and* the aggressive congestion control of TCP-Cubic. Note that TCP-Westwood, which incorporates rate estimation in its congestion control mechanism yields a reasonable trade-off in terms of a slightly increased stream delay compared to TCP-Vegas and reasonably high throughput (about 90% of TCP-Cubic). Indeed, this aspect of TCP-Westwood is noticeable across other experiments too.

8. CONCLUSIONS

We performed a detailed characterization of the physical, MAC and transport layer performance in the 1xEV-DO network. Our scope covered both stationary as well as mobile channels and also involved evaluation of various state-of-the-art transport protocols. The analysis was carried out by cross-analyzing detailed channel level information: rate, packet loss, SINR, *etc.* as well as tcpdump logs. The physical channel was found to be highly variable over long time scales of hours and days. However, at short time scales on the order of a few milli-seconds it shows significant memory. This translated into the channel retaining the same rate over the range of a few tens to few hundreds of milli-seconds. Through simulations using physical rate traces we found that the opportunistic Proportional Fair scheduling scheme is beneficial compared to the Round-Robin schedule in mobile scenarios (gains of around 20%) but has minimal gain (4 – 5%) for stationary scenarios. Somewhat surprisingly, the performance of all TCP variants was dominated by queueing effects rather than channel fluctuations. This can be attributed to the presence of large buffers that mitigate rapid channel fluctuations allowing high ($\approx 85\%$) channel utilization in moderate channel conditions and low propagation delays. They however can also induce large delay in the presence of loss-based congestion control mechanisms (TCP-Cubic) that seek to fill the queue as opposed to delay-based mechanisms that control queue size (TCP-Vegas), though the latter offer smaller throughput.

9. REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proc. ACM SIGCOMM*, August 2004.
- [2] E. Altman, C. Barakat, and V. M. R. Ramos. Analysis of AIMD protocols over paths with variable delay. In *Proc. IEEE INFOCOM*, Hong Kong, Mar. 2004.
- [3] A. Baiocchi and F. Vacirca. TCP fluid modeling with a variable capacity bottleneck link. In *Proc. IEEE INFOCOM*, May 2007.
- [4] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushyana, and S. Viterbi. CDMA/HDR: a bandwidth efficient high speed wireless data service for nomadic users. *IEEE communications magazine*, 38(7):70–77, July 2000.
- [5] N. Bhusha, C. Lott, P. B. and R. Attar, Y.-C. Jou, M. Fan, D. Ghosh, and J. Au. CDMA200 1xEV-DO Revision A: A physical layer and MAC layer overview. *IEEE Comm. Mag.*, 44(2), February 2006.
- [6] L. Brakmo, S. O’Malley, and L. Peterson. TCP Vegas : New Techniques for congestion detection and avoidance. In *Proc. ACM SIGCOMM*, pages 24–35, August 1994.
- [7] C. Casetti, M. Gerla, S. Mascolo, M. Sanadidi, and R. Wang. TCP Westwood : Bandwidth Estimation for Enhanced Transport over Wireless Links. In *Proc. ACM MOBICOM*, pages 287–297, 2001.
- [8] R. Chakravorty, S. Banerjee, P. Rodriguez, J. Chesterfield, and I. Pratt. Performance optimizations for wireless wide-area networks: comparative study and experimental evaluation. In *Proc. ACM MOBICOM ’04*, pages 159–173, 2004.
- [9] 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.
- [10] M. C. Chan and R. Ramjee. Improving TCP/IP performance over third generation wireless networks. In *Proc. IEEE INFOCOM*, Hong Kong, 2004.
- [11] M. Claypool, R. Kinicki, W. Lee, M. Li, and G. Ratner. Characterization by measurement of a CDMA 1xEVDO network. In *Proc. Wireless Internet*, 2006.
- [12] A. Jalai, R. Padovani, and R. Pankaj. Data Throughput of CDMA-HDR: A High Efficiency High Data Rate Personnel Communication Wireless System. In *Proc. IEEE Vehicular Technology Conference*, May 2000.
- [13] S. K. Kasera, R. Ramjee, S. Thuel, and X. Wang. Congestion Control Policies for IP-based CDMA Radio Access Networks. In *Proceedings of IEEE INFOCOM*, San Francisco, CA, USA, 2003.
- [14] T. E. Klein, K. K. Leung, and H. Zheng. Improved TCP Performance in Wireless IP Networks through Enhanced Opportunistic Scheduling Algorithms. In *Proc. IEEE Globecom*, 2004.
- [15] A. Konrad, B. Zhao, and A. Joseph. A Markov-Based Channel Model Algorithm for Wireless Networks. In *Wireless Networks*, volume 9, pages 189–199, 2003.
- [16] H. Kushner and P. Whiting. Convergence of proportional-fair sharing algorithms under general conditions. *IEEE Transactions on Wireless Communications*, 3:1250–1259, July 2004.
- [17] Y. Lee. Measured TCP Performance in CDMA 1xEV-DO Network. In *Proc. Passive Active Measurement Conference*, 2006.
- [18] R. Ludwig and K. Sklower. The Eifel retransmission timer. *SIGCOMM Comput. Commun. Rev.*, 30(3):17–27, 2000.
- [19] A. C. H. Ng, D. Malone, and D. J. Leith. Experimental evaluation of TCP performance and fairness in an 802.11e test-bed. In *Proc. E-WIND ’05: ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, 2005.
- [20] K. Papagiannaki, M. Yarvis, and W. Conner. Experimental Characterization of Home Wireless Networks and Design Implications . In *Proc. IEEE INFOCOM*, April 2006.
- [21] I. Rhee and L. Xu. CUBIC: A New TCP-Friendly High-Speed TCP Variant. In *Proc. of PFLDnet*, 2005.
- [22] S. Ross. *Stochastic Processes*. John Wiley and Sons, 1996.
- [23] W. Tan, F. Lam, and W. Lau. An Empirical Study on 3G Network Capacity and Performance. In *Proc. IEEE INFOCOM*, pages 1514–1522, 2007.