Characterization by Measurement of a CDMA 1x EVDO Network

Mark Claypool, Robert Kinicki, William Lee, Mingzhe Li, and Gregory Ratner
Computer Science Department
Worcester Polytechnic Institute
100 Institute Rd, Worcester, MA 01609, USA
{claypool|rek}@cs.wpi.edu

ABSTRACT

Cellular networks are increasingly carrying Internet traffic for Email and Web browsers, VoIP and online games and streaming video. While there have been some studies characterizing the throughput of cellular data networks, there has been few measurement studies focusing on the latest Evolution-Data Optimized (EVDO) cellular protocol, considering characteristics for throughput-sensitive, latency-sensitive and streaming applications. This paper provides a measurement-based characterization of an EVDO network, using measures of TCP throughput, round-trip times and loss rates, and bandwidth estimation in order to gauge the ability of EVDO to support the QoS requirements of a range of applications. A week's worth of active measurement data reveals throughputs satisfactory for Email and Web browsing, latencies too high for interactive applications, and consistent results for bandwidth estimations, suitable for streaming.

Categories and Subject Descriptors

H.2 [**Computer-Communication Networks**]: Miscellaneous

Keywords

Measurement, Wireless, EVDO

1. INTRODUCTION

Cellular networks are increasingly used to carry Internet traffic, allowing Internet users to stay connected wherever they are, without network wires or even the need of a wireless local area network "hotspot". The advent of the third generation (3G) networks has provided potential bandwidths sufficient for many typical Internet applications, including Email, Web browsing and even streaming audio and video directly to the cell phones of mobile users. Cell phone providers have recognized the potentially lucrative market in streaming video content, such as sports highlights, news programs or even entire television shows, directly to customers' mobile phones [2].

The intent to stream video over cellular phone networks presents new networking challenges. Video streaming can be bandwidth intensive, and many cellular networks have limited bitrate capacities. A typical video stored on the Web is encoded at around 200 Kb/s [18], while in the United States the widespread single-carrier radio transmission technology (1xRTT)[1] provides data rates of only 60-100 Kb/s. Moreover, just as connection quality (commonly referred to as "bars") varies with physical proximity from a cell tower, so too does the bitrate capacity on a cellular link.

An attempt to provide more capacity for bandwidth-hungry video is 1x Evolution-Data Optimized (EVDO), a 3G Code Division Multiple Access (CDMA) wireless data standard adopted by many mobile phone service providers in North and South America, Japan and Korea, Israel, and Australia and New Zealand [24]. EVDO purports to support bandwidths ranging from 500 Kb/s up to 2.4 Mb/s, suitable for streaming video content with good quality.

Video streaming can be quite elastic, adjusting the quality of the video to fit the available network capacity, typically by decreasing the frame rate or by reducing the resolution of the frames. In order to begin streaming, a video server must estimate the bandwidth from the server to the client. Typical commercial streaming systems estimate bandwidth with packet-pair techniques [19], where the dispersion of two large packets sent back-to-back can provide an estimate of the capacity on the bottleneck link.

Previous studies of CDMA 1x EVDO networks have used theoretical analysis or simulations with various link parameters [4, 8, 16, 11]. Although there have been some measurement studies of 3G networks [3, 7, 23] including an EVDO network [17], these studies have primarily focused on application throughput and have not characterized latency for interactive applications and bandwidth estimation for streaming.

The paper provides a characterization of an EVDO network with particular focus on the impact the network characteristics have on three classes of applications: throughput-sensitive (Web, Email, file downloads), delay-sensitive (Voice over IP, online games) and streaming (streaming video). Network characterization is done via ICMP and TCP downloads, and bandwidth estimation through packet-pair trains sent over UDP. The experiments gather one week of data, with samples taken every 30 minutes to provide time of day and day of week correlations on the network characteristics.

[1] Also known as "National Access".

WiCon'06 August 2-5, 2006, Boston, MA, United States

In general, the throughputs measured (about 350 Kb/s) over the EVDO network are sufficient for Web browsing, Email, file downloading and other similar applications. Loss rates are similarly encouraging, being effectively zero at all times of the day and week. However, round-trip times (nearly 200 milliseconds) are still prohibitive for highly interactive applications, such as Voice over IP and many online games. Bandwidth estimations (about 500 Kb/s) are consistent, suggesting packet-pair techniques are effective for streaming video and other similar applications.

The rest of this paper is laid out as follows: Section 2 provides background on EVDO networks and bandwidth estimation; Section 3 describes the experimental methodology used to characterize the EVDO networks; Section 4 presents the results and analysis; and Section 5 summarizes our conclusions and possible future work.

2. BACKGROUND

2.1 CDMA and EVDO

Code division multiple access (CDMA) is a means of wireless multiplexing, not by dividing the channel by time or frequency, but instead by encoding data with a special code for each channel that uses constructive interference properties to perform the multiplexing [25]. The CDMA protocol called 1xRTT[2] is the standard used by most CDMA providers in the U.S. today. The 1xRTT protocol specifies a maximum theoretical channel capacity of 153 Kb/s on both the downward and upward link [22].

The newest CDMA protocol, *1x Evolution-Data Optimized* (or just *EVDO*), allows for a maximum capacity ranging from 500 Kb/s up to 2.4 Mb/s on the downlink and a maximum capacity of 153 Kb/s on the uplink. EVDO networks are optimized for data transfer on the downlink through the use of packet-based TDM (time division multiplexing) and by scheduling packets based on a user's channel quality, rather than utilizing traditional round-robin scheduling. In contrast, CDMA is used on the uplink, but with adaptive rate control, which allows an increase in uplink throughput that can be higher than 1xRTT networks [1].

Second and third generation networks, including CDMA, combat link-layer errors by using forward error correction (FEC) and link-layer retransmissions to repair frames damaged during radio propagation.

2.2 Bandwidth Estimation

Bandwidth estimation usually refers to the end-to-end measurement of bandwidth-related metrics performed by the end hosts of a path without requiring administrative access to intermediate routers along the path [20]. In this paper, the term *bandwidth* is used to mean the maximum possible throughput that a link or end-to-end path can deliver. The end-to-end bandwidth B is defined as:

$$B = \min_{i=1,\ldots,H} B_i \tag{1}$$

where B_i is the capacity of i-th hop, and H is the number of hops in the end-to-end path. The hop with the minimum bandwidth is the *narrow link* on the path. For most network configurations with a wide-area wireless link (such as EVDO), the last hop to the end-host is the narrow link.

A popular bandwidth estimation technique is packet dispersion, first introduced in [5, 12, 13] and improved via various research approaches in tools such as *bprobe/cprobe* [6], *nettimer* [14, 15], *sprobe* [21] and *pathrate* [9, 10].

Packet dispersion sends two packets (a packet-pair) of the same size back-to-back into the network. After the packets traverse the narrow link, the time dispersion between the two packets is linearly related to the narrow link capacity. Figure 1 [20] illustrates the basic concept of packet dispersion. When packets of size L with initial dispersion Δ_{in} go through the link of with bandwidth B_i, the dispersion after the link Δ_{out} becomes [20]:

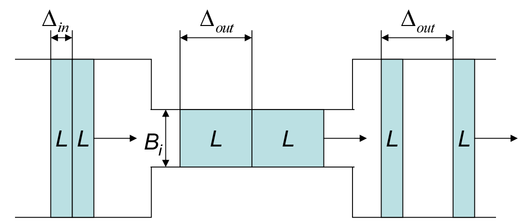

Figure 1: Packet Dispersion

$$\Delta_{out} = \max(\Delta_{in}, \frac{L}{B_i}) \tag{2}$$

After packets go through each link along an H hops end-to-end path, the final dispersion Δ_R at the receiver is:

$$\Delta_R = \max_{i=0,\ldots,H}(\frac{L}{B_i}) = \frac{L}{\min_{i=0,\ldots,H} B_i} = \frac{L}{B} \tag{3}$$

where B is the end-to-end bandwidth. Therefore, the end-to-end path bandwidth can be estimated from $B = L/\Delta_R$.

Larger packets result in a wider dispersion that is more sensitive to queuing delay in intermediate hops and produce a more accurate estimation of end-to-end capacity. However, if packets are too large, such as packets larger than the maximum transmission unit (MTU) of a network, additional delay can be introduced as packets are fragmented and reassembled from server to client. Small packets are not ideal either, since they result in proportionally shorter dispersion times which can lead to overestimation of end-to-end capacity [10].

Packet train dispersion extends packet-pair dispersion by using multiple back-to-back probing packets. However, the concepts for a packet train are similar to that of a single packet-pair.

Compared to other bandwidth estimation techniques, packet dispersion techniques usually have a faster measurement time and induce less load on the network path. Packet dispersion techniques have also been used in some commercial streaming video applications. For instance, Windows video streaming uses a train of three packets (two packet-pairs) to estimate the end-to-end bandwidth before beginning the streaming from server to client.

[2]The RTT in this context stands for "Radio Transmission Technology" and *not* "Round-Trip Time".

3. METHODOLOGY

The objective was to characterize the EVDO channel in terms of throughput, round-trip time, loss rate and bandwidth estimation. Figure 2 shows a high-level overview of the network setup used in the experimentation.

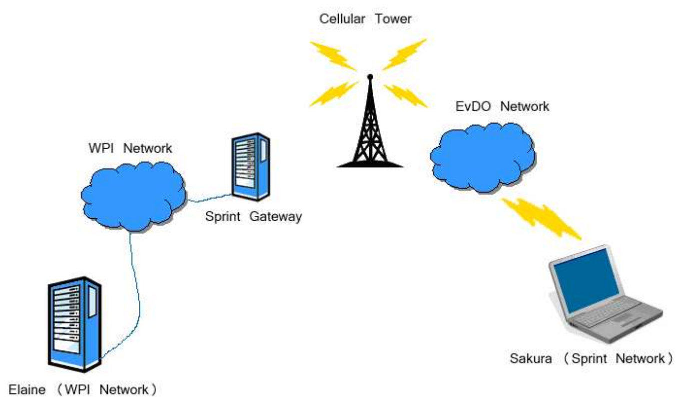

Figure 2: Experimental Network

3.1 Equipment Testbed

The experimental traffic for this study was sent between two computers, one with a wired Internet connection and the other a laptop on the end of wireless EVDO Internet connection. The server, labeled Elaine in Figure 2 sent data downstream to the client for bandwidth estimation and TCP throughput experiments. Elaine has a Celeron 733 MHz processor with 256 MB RAM, and is connected to a 100 Mb/s LAN on the Worcester Polytechnic Institute (WPI) campus,[3] which in turn has a 50 Mb/s uplink to the Internet. Elaine runs the GNU Linux operating system with the 2.6.11 kernel and a minimal amount of services necessary for basic network operation.

The client computer, labeled Sakura in Figure 2, received data from the server for throughput and bandwidth estimation and ran `ping` experiments to the Sprint gateway for round-trip time and loss rate measurements. The client is a notebook with a Sierra Wireless Aircard 580 PCMCIA EVDO network card connected to Sprint PCS's EVDO network in Worcester, Massachusetts. The notebook is equipped with a Pentium M 1.8 GHz CPU, 1 GB RAM, and operates on the Linux 2.6.11 kernel. All additional services not required for network connections were turned off during the experiments.

3.2 Tools

The tools used are open-source networking tools and a customized bandwidth estimation tool, *npath*.

3.2.1 Open-source Tools

TCPdump[4] captures all traffic information on a specific network interface device. For our experiments, TCPdump was used on Sakura to capture activity on the EVDO network card. *TCPTrace*[5] is a tool to parse and process *TCPdump* files, and was used to calculate the average TCP throughput.

GNU *wget*[6] is a tool that is capable of retrieving content from servers using various TCP/IP-based protocols, such as FTP (hypertext transport protocol). Wget was used to transfer a 20 MB file via FTP from Elaine to Sakura.

The standard Linux/Unix *ping* utility was used to estimate round-trip times and packet loss rates between hosts by sending ICMP "echo request" packets every second to the server from the client over the EVDO network. The time between sending the ICMP packet and receiving a response is the round-trip time for that packet. If a response is not received within one second, the packet loss is recorded. One-hundred pings were sent during each test run.

3.2.2 Custom Tool – npath

Loosely based on *Pathrate*, a bandwidth estimation tool discussed in [10], *npath* is a bandwidth estimation tool created specifically for this project.

npath includes a server and a client. The server runs on the server (Elaine) and waits for a connection from the client program (run on Sakura). The client initiates a TCP connection with the server, signaling the server to begin the bandwidth estimation experiment.

The server then generates a train of packets (providing a series of packet-pairs) and sends them to the client over UDP. Each packet contains the train number and the packet number. As the client receives each packet, it checks the train number and packet number. If either is not the expected number, the client discards the train and sends a Negative Acknowledgment (NACK) packet to the server to begin retransmission of the train. If the entire train is received without error, the experiment continues. The server pauses for one second between each packet train and ten seconds between every ten trains. After thirty packet pair trains, the server changes packet sizes and begins anew. Packet dispersion information is recorded at the client.

3.3 Experimental Design

Tests were run every thirty minutes for seven days with all output logged on Sakura for later analysis. Each test lasted approximately twelve minutes and consisted of (in-order): 1) measure wireless signal strength; 2) transfer a 20 MByte file via TCP; 3) ping the server 100 times; and 4) estimate bandwidth with four different packet sizes, in sequence.

3.3.1 Signal Strength Baseline

The received signal strength of the EVDO card was measured prior to each run by sending a special serial command to the card. Since the intent was not to characterize the performance of the EVDO network under alternate wireless conditions (this is left as future work), for all tests the signal strength ranged from -90 to -83 dBm which is equivalent to excellent signal strength relative to cell phone signal strength metrics.

3.3.2 TCP and ICMP Traffic Performance Analysis

To analyze the performance of TCP traffic, a 20 MB file was transferred using wget via a TCP connection from Elaine to Sakura. TCPdump was used to generate and log an in-depth trace of all activity on the EVDO network interface device during the file transfer. After the TCP transfer

[3]WPI is physically located in Worcester, Massachusetts, USA.
[4]http://www.tcpdump.org/
[5]http://jarok.cs.ohiou.edu/software/tcptrace/
[6]http://www.gnu.org/software/wget/

was completed, 100 ICMP pings were sent from Sakura to measure round-trip time and loss rates performance. Since WPI's network infrastructure drops all incoming ICMP traffic, it was necessary to send the ICMP packets to a node within Sprint's network. Round-trip time and packet loss information was timestamped and logged for each ping on Sakura for later analysis.

3.3.3 *Bandwidth Estimation*

To estimate the available bandwidth of the EVDO network, npath sent packet trains from Elaine to Sakura on the EVDO network. The npath tool sent 30 trains, each with 25 packets. First the 30 trains all used 1472 byte packets. Then, 30 trains were run with 1200 byte, 1000 byte, and finally 800 byte packets. The different packet sizes were designed to ascertain the impact of packet size on bandwidth estimation results.

The length of the trains (25 packets) was initially determined by using an estimation tool that would send packet trains of increasing length until packet loss was seen. At this train length, it was assumed that longer trains overflow a packet queue (most likely from the high-bitrate, wired connection to the lower-bitrate, wireless EVDO connection) from the server to the client. Using this technique, the longest train that could be sent without error was determined to be near 30 packets, so a train length of 25 was used for each packet size to reduce the chance of queue overflow.

4. ANALYSIS

Data was gathered every 30 minutes for an entire week, from November 28, 2005 to December 4, 2005. From the data, performance was seen to be similar from one weekday to the next (Monday through Friday) and from one weekend day to the next (Saturday and Sunday),[7] thus the detailed analysis concentrates on two specific days: Wednesday November 30, 2005, and Saturday December 3, 2005, hereafter called the *weekday* and the *weekend day*, respectively.

4.1 TCP Throughput

TCP throughput was measured by downloading a 20 MB file via FTP from the wired server to the EVDO client.

Figure 3 shows a cumulative distribution function (CDF) for the average TCP throughput for both the weekday and the weekend day, obtained from TCPdump and TCPtrace packet traces of the 20 MByte wget transfer. The x-axis is the throughput in kilobits per second, and the y-axis is the cumulative distribution. The average throughput is obtained by dividing total number of bytes transferred during a given run (including retransmissions) by the total time in microseconds between the first and last packets.

The relatively steep shape with a small horizontal spread in throughput in Figure 3 indicates that the TCP bitrates are fairly consistent throughout the day. Ninety percent of throughput values for the weekday were below 375 Kb/s and for the weekend day ninety percent were below 384 Kb/s. Figure 3 also shows that the extremes of the TCP throughput, namely those outside of the top and bottom five percent, are slightly higher for the weekend day then

[7]The Appendix contains graphs of throughput and bandwidth estimates for all days of the week, visually showing consistency from day to day.

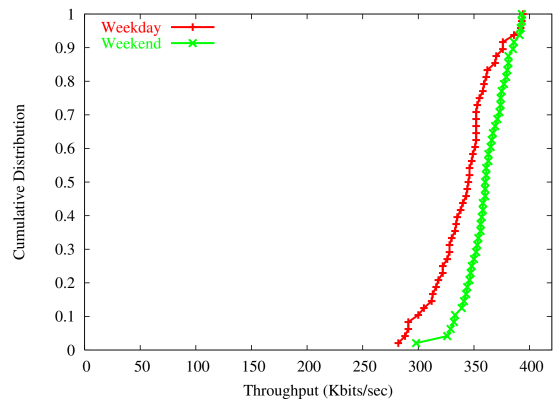

Figure 3: Cumulative Distribution of Average Throughput for a Single Day

for the weekday. This suggests that the network effectively has more available capacity for users during a weekend than during a weekday, most likely because there are fewer users on the EVDO network during the weekend.

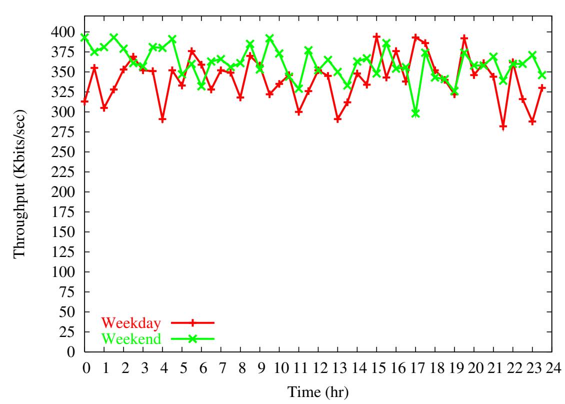

Figure 4: Average Throughput over Time for a Single Day

Figure 4 provides the average throughput of the 20 Mbyte wget TCP transfer for every half hour run throughout the weekday compared to the weekend day. The x-axis is the time of day in 24-hour military time, and the y-axis is the average throughput in Kb/s.

There is no visual-apparent throughput pattern over the course of the day, and no visual correlation between the average throughput and the time of day. For example, even during the early hours of the weekday an average throughput of 376 Kb/s was reached at 5:30, but was followed by a drop in throughput to 360 Kb/s at 6:00 and down to 328 Kb/s at 6:30. An equivalent weekday result can be found after the afternoon peak at 15:00. The average throughput drops to 344 Kb/s just a half hour later, then increases again to 376 Kb/s by 16:00, only to drop down to 340 Kb/s in another half hour.

While similar to the weekday data trends, throughout the day the weekend has slightly higher average TCP throughputs than does the weekday. Peaks of around 392 Kb/s can be seen even in the very early hours of 12:00, 1:30, and at 4:30, that are higher than their counterparts during the weekday. This is in contrast to the earliest hours of the weekday observations, where only 1:00 had an average throughput of around 304 Kb/s, for example. Yet several times during the weekend, average throughputs were below their weekday counterpart, such as a significant drop from 16:30 - 17:30.

Impact. In general, throughputs are somewhat lower than can theoretically be achieved by EVDO networks. However, the average throughput of just over 300 Kb/s is still more than five times faster than typical throughputs over a dialup modem and suitable for typical Email, Web browsing and file transfer applications.

4.2 Round-Trip Time

To measure round-trip times, 100 ICMP pings were sent every half hour from the laptop on the EVDO network to a gateway within the Sprint network. Similar to the TCP analysis, weekday versus a weekend day performance is considered.

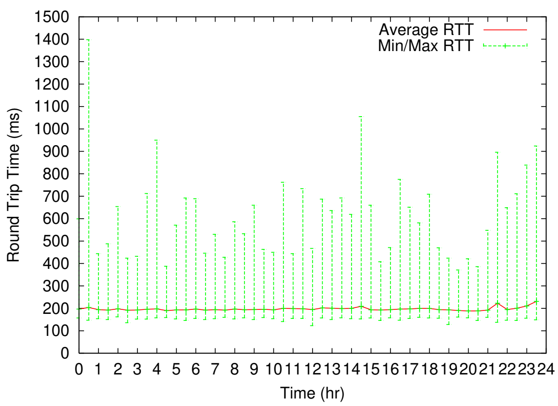

Figure 5: Round-Trip Times for Weekday

Figure 5 depicts the results for the weekday. The x-axis is time of day, and the y-axis is the round-trip time in milliseconds. The data sets show components of the round-trip time information – the vertical bars show the minimum and maximum round-trip times for the 100 ICMP packets sent at the indicated time, and the horizontal line connects the round-trip times averaged for all 100 packets during each run.

Similar to the results for TCP throughput, Figure 5 shows no visual correlation between time of day and round-trip times. The maximum round-trip times vary greatly throughout the day, with a maximum of about 1400 milliseconds at 12:30am and a minimum of about 370 milliseconds at 7:30pm. However, the minimum round-trip times exhibit little variability throughout the day, oscillating around the 160 millisecond range. Despite the variation in the maximum round-trip times, the average round-trip time remains consistent at around 200 milliseconds throughout the day.

Figure 6 provides data similar to Figure 5, except for a weekend day. The axes, vertical bars, and line represent the same metrics as in Figure 5. For the weekend day as well, there is no visual correlation between time of day and the round-rip time. Round-trip time minimums stay around 180 milliseconds, round-trip time maximums are generally around 700 milliseconds but spike as high as 1000 milliseconds, and the average round-trip time stays around 200 milliseconds at all times of the day.

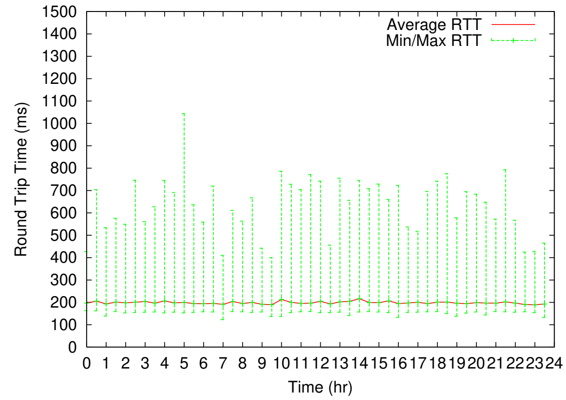

Figure 6: Round-Trip Times for Weekend Day

Comparing Figures 5 and 6, there is no significant difference in round-trip time performance between weekday and weekend.

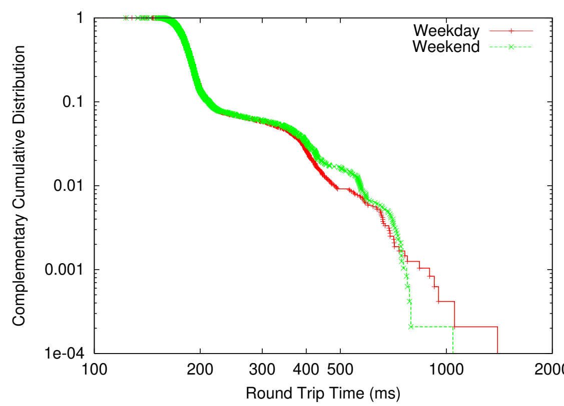

Figure 7: CCDF for Round-Trip Times

In order to better depict the maximum round-trip times, Figure 7 provides the complementary CDF (CCDF) of all the round-trip time measurements gathered on the weekday and the weekend day. The x-axis is the round-trip time in milliseconds, and y-axis is the complimentary cumulative distribution (1 - the cumulative distribution), shown on a logarithmic scale for both axes. The two data sets are the round-trip times for the weekday and the round-trip times for the weekend day. Note that the median of the weekday and weekend round-trip time measurements are the same at about 180 milliseconds, and only about ten percent of the round-trip times are above 220 milliseconds throughout either day. The distributions overlap through about 97% of the round-trip time values and only significantly diverge at the very end of the tail, 0.1% of the distribution, with the largest weekend round-trip times being around 800 milliseconds while the largest weekday round-trip times tend towards 1000 milliseconds.

Impact. Round-trip times have the largest impact on interactive applications that require low-latency to be responsive to user input. In general, the median latencies tending towards 200 milliseconds are too high for interactive applications that require latencies of 150 milliseconds or lower. Additionally, maximum round-trip times of around 1 second have a seriously detrimental effect on interactive application performance.

4.3 Bandwidth Estimation

Bandwidth estimates were made by sending 30 trains of 25 packets (each train thus providing 24 packet-pair measurements) from the wired server to the EVDO client. The results shown are all from the 1200-byte packets because the largest packet size used in our measurements, 1472 bytes, gave inconsistent results, possibly due to fragmentation, and smaller packet sizes generally provide less accurate estimates of bandwidth than do larger packets.

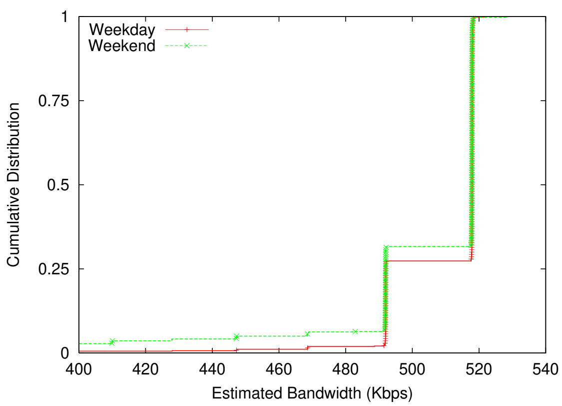

Figure 8: CDF of Bandwidth Estimation for 5:30pm

Figure 8 shows a CDF of the estimated bandwidth, obtained from each of the packet-pairs, at 5:30pm for the weekday and weekend day. The x-axis is the estimated bandwidth in kilobits per second, the y-axis is the cumulative distribution, with the data sets showing the weekday and weekend day. Again, the distributions are drawn from the bandwidth estimates obtained from 30 trains of 24 packet-pairs, so 720 individual bandwidth estimates. Similar results were obtained for different times of the day, but 5:30pm was selected as perhaps the busiest (and therefore variable) network period.

The bandwidth estimates, each the result of a single packet-pair, are remarkably consistent with approximately seventy percent (from 0.3 to 1.0) of the values being exactly 520 Kb/s (note, that the far left of the axis starts at 400 Kb/s and not at 0). 520 Kb/s is at the lower end of the purported maximum capacity for EVDO networks. There is a modest (about 20%) set of bandwidth estimates that are just under 500 Kb/s. Approximately 5% of the bandwidth estimates obtained with a single packet pair were significantly lower than 500 Kb/s, even less than 100 Kb/s. These latter values arise when there is a large dispersion between two packets in a pair, probably co-incidental with the large round-trip times reported in Figure 7.

Figure 9 visually compares the bandwidth estimation results for the weekday to the average TCP throughputs (reported in Section 4.1). The x-axis is the time of day in 24-hour military time, and the y-axis is the bandwidth estimate or throughput in Kb/s. The data sets shown are the average TCP throughput, with the fifth percentile and

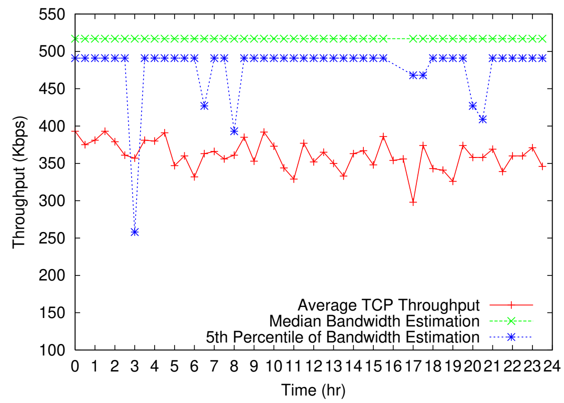

Figure 9: Bandwidth Estimation for Weekend Day

fiftieth (median) bandwidth estimations.

The median bandwidth estimates are remarkably consistent, appearing as a flat line at 520 Kb/s. The 5th percentile is typically at just under 500 Kb/s, but is a bit less consistent, with some dips as low as 400 Kb/s and even one dip (at 3am) lower than the average TCP throughput.

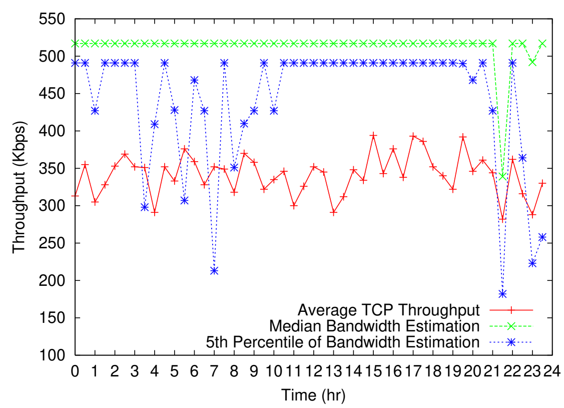

Figure 10: Bandwidth Estimation for Weekday

Figure 10 graphs similar data to Figure 9, but for the weekday. The median bandwidth estimates are still consistent, but deviate from 520 Kb/s at 9:30pm and 11:00pm. The 5th percentile estimates are also less consistent than on the weekend day, falling below the TCP throughput on 5 occasions, in the early morning and then late in the evening.

Impact. For applications such as streaming media that use packet-pairs for bandwidth estimation, the results suggest that many pairs (or long trains) are not needed in order to obtain the estimated bandwidth. In fact, based on the distribution in Figure 8, the two packet-pair schemes used by some commercial video players [19] would probably yield the 500 Kb/s bandwidth estimate.

4.4 Summary

Table 1 summarizes average throughput, round-trip time, packet loss and bandwidth estimation for the weekday and weekend. The average throughput is slightly higher on the weekend than on the weekday, the round-trip times are consistently around 200 milliseconds at all times, there is no sig-

	Weekday		Weekend	
	Avg	StdDev	Avg	StdDev
Tput (Kb/s)	341	27	360	19
RTT (ms)	197	68	198	72
Loss (%)	0	0	0	0
	Median	5th %-tile	Median	5th %-tile
Bwidth (Kb/s)	520	491	520	470

Table 1: Summary of EVDO Measurements

nificant packet loss and the bandwidth estimates are slightly above the achievable TCP throughputs. As a comment on the packet loss, as mentioned in Section 2, EVDO utilizes forward error correction (FEC) and automatic repeat request (ARQ) to correct errors and retransmit frames at the wireless data link layer, thus hiding the packet loss from the end-hosts.

5. CONCLUSIONS

The increasing availability of high bitrate, wide-area wireless networks presents the opportunity for an expanding variety in applications to mobile users, including Email and Web browsing, Voice over IP and online games, and streaming video. Email and Web browsing are primarily concerned with throughput, while VoIP and online games care mostly about latency, and streaming video, being elastic, needs effective bandwidth estimation do determine the most appropriate streaming rate.

This paper presents a study of a 1x EVDO network, characterizing performance based on TCP throughput, ICMP loss and round-trip times, and packet-pair bandwidth estimations. Active measurements gathered every 30 minutes, 24-hours per day, for an entire week provide data for time of day and day of week inspection. Overall, throughputs are adequate for Email and Web browsing and, due to its elastic nature, enough for streaming video, as well. Average round-trip times are too high for VoIP and online games, with maximum round-trip times being especially problematic. Even minimum round-trip times are not conducive to these interactive applications. Bandwidth estimation measurements are consistent across packet-pairs, suggesting small trains of packet-pairs are sufficient to obtain accurate network capacity estimates. Neither throughput, round-trip time, nor bandwidth estimation correlates with time of day or the day of the week.

There are several areas for future work. The study was for Sprint's EVDO network, but competitors' EVDO networks, such as that of Verizon Wireless, may have somewhat different performance characteristics. In addition, a similar study could also examine GSM-based networks, such as T-Mobile's or Cingular's, and compare performance to the EVDO networks.

A revision to the EVDO standard, EVDO Revision A, currently available in Korea and Japan, is expected to be deployed by CDMA carriers in the United States in 2006 [22]. This new standard is specifically designed to decrease network latency to support latency-sensitive applications, such as VoIP (Voice over IP) and online games [1]. Measuring the performance of this new standard as compared to the results in this paper may be useful for gauging EVDO suitability for highly interactive applications.

Additionally, end-hosts in cellular data networks often have a range of connection conditions based on their distance from the cellular tower, urban interference from buildings and trees, and even from the weather. Study of EVDO performance over a range of connection conditions would provide useful data to more fully characterize EVDO performance. Moreover, the impact of mobility on network characteristics, particularly when the end-host is handed off from one cellular tower to another, may show characteristics different than those studied for the stationary end-host in this work.

6. REFERENCES

[1] Airvana, Inc. All-IP 1xEV-DO Wireless Data Networks: A Technical White Paper, Aug. 2004. [Online at: http://www.airvananet.com/files/-Airvana_1xEV_Tech_White_Paper.pdf].

[2] AP Press. Networks Rush to Offer TV Shows Online. The Boston Herald, Apr. 2006.

[3] P. Benko, G. Malicsko, and A. Veres. A Large-scale, Passive Analysis of End-to-End TCP Performance over GPRS. In *Proceedings of IEEE INFOCOM*, 2004.

[4] Q. Bi and S. Vitebsky. Performance Analysis of 3G-1X EVDO High Data Rate System. In *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC)*, Orlando, FL, USA, 2002.

[5] J.-C. Bolot. End-to-end Packet Delay and Loss Behavior in the Internet. In *Proceedings of the ACM SIGCOMM Conference*, pages 289–298, San Francisco, CA, USA, Sept. 1993.

[6] R. L. Carter and M. E. Crovella. Measuring Bottleneck Link Speed in Packet-Switched Networks. *Performance Evaluation*, 27(8):297–318, Oct. 1996.

[7] R. Chakravoty and I. Pratt. Performance Issues with General Packet Radio Service. *Journal of Communications and Networks (JCN)*, 4(2), Dec. 2002.

[8] W. Chung, H. Lee, and J. Moon. Downlink Capacity of CDMA/HDR. In *Proceedings of the IEEE Vehicular Technology Conference*, Spring 2001.

[9] C. Dovrolis, P. Ramanathan, and D. Moore. What do Packet Dispersion Techniques Measure? In *Proceedings of IEEE INFOCOM*, pages 905–914, Anchorage, AK, USA, Apr. 2001.

[10] C. Dovrolis, P. Ramanathan, and D. Moore. Packet-Dispersion Techniques and a Capacity-Estimation Methodology. *IEEE/ACM Transactions on Networking*, 12:963 – 977, Dec. 2004.

[11] E. Esteves, M. Gurelli, and M. Fan. Performance of Fixed Wireless Access with CDMA 2000 1xEV-DO. In *Proceedings of the IEEE Vehicular Technology Conference*, Fall 2003.

[12] V. Jacobson. Congestion Avoidance and Control. In *Proceedings of the ACM SIGCOMM Conference*, Stanford, CA, USA, Aug. 1988.

[13] S. Keshav. A Control-Theoretic Approach to Flow Control. In *Proceedings of the ACM SIGCOMM Conference*, Sept. 1991.

[14] K. Lai and M. Baker. Measuring Bandwidth. In *Proceedings of IEEE INFOCOM*, pages 235–245, New York, NY, USA, Apr. 1999.

[15] K. Lai and M. Baker. Measuring Link Bandwidths Using a Deterministic Model of Packet Delay. In *Proceedings of ACM SIGCOMM*, pages 283–294, Stockholm, Sweden, Aug. 2000.

[16] S. Lee. The Performance Improvement Principles of TCP Protocol Stack on Packet Switching High Speed Wireless DS-CDMA Links. In *Proceedings of the IEEE Vehicular Technology Conference*, Fall 2001.

[17] Y. Lee. Measured TCP Performance in CDMA 1x EV-DO Network. In *Proceedings of the Passive and Active Measurement Conference (PAM)*, Adelaide, Australia, Mar. 2006.

[18] M. Li, M. Claypool, R. Kinicki, and J. Nichols. Characteristics of Streaming Media Stored on the Web. *ACM Transactions on Internet Technology (TOIT)*, 5(4), Nov. 2005.

[19] J. Nichols, M. Claypool, R. Kinicki, and M. Li. Measurements of the Congestion Responsiveness of Windows Streaming Media. In *Proceedings of the 14th ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, June 2004.

[20] R. Prasad, M. Murray, C. Dovrolis, and K. Claffy. Bandwidth Estimation: Metrics, Measurement Techniques, and Tools. *IEEE Network*, November-December 2003.

[21] S. Saroiu, P. K. Gummadi, and S. D. Gribble. SProbe: A Fast Technique for Measuring Bottleneck Bandwidth in Uncooperative Environments, Aug. 2001. Online: http://sprobe.cs.washington.edu/.

[22] M. W. Thelander. Taking CDMA2000 into the Next Decade, Oct. 2005. [Online at: http://www.cdg.org/resources/white_papers/files/-3G_Evol_Oct05.pdf].

[23] A. Wennstrom, A. Brunstrom, J. Rendon, and J. Gustafsson. A GPRS Testbed for TCP Measurements. In *Proceedings of International Workshop on Mobile and wireless Communications Network*, 2002.

[24] Wikipedia. Code Division Multiple Access - Wikipedia, the Free Encyclopedia, 2006. [Online at: http://en.wikipedia.org/wiki/EV-DO; accessed 9-April-2006].

[25] W. C. Yee. Overview of Cellular CDMA. *IEEE Transactions on Vehicular Technolgy*, 40:291–302, May 1991.

APPENDIX

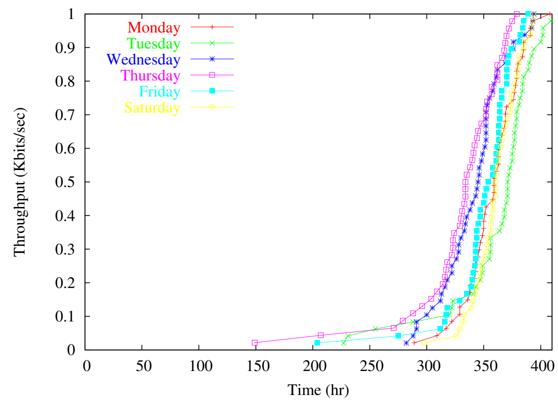

Figure 11: Cumulative Distribution of Average Throughput for Each Day (Figure 3 has the same data for Wednesday and Saturday)

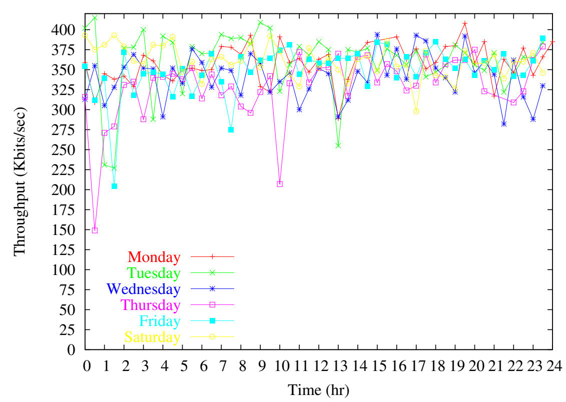

Figure 12: Average Throughput over Time for Each Day (Figure 4 has the same data for Wednesday and Saturday)

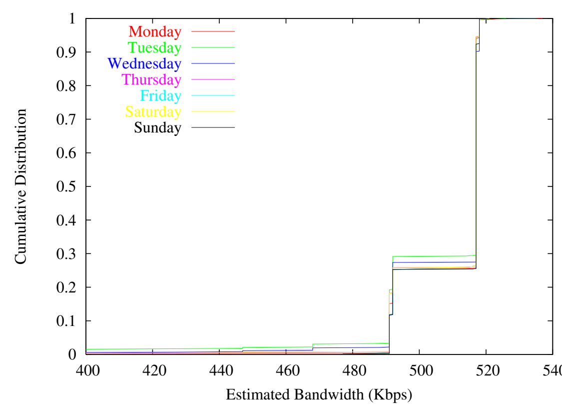

Figure 13: CDF of Bandwidth Estimation for 5:30pm for Each Day (Figure 8 has the same data for Wednesday and Saturday)