

# Characterization of a 3G EV-DO Network - a Measurement Study

Zhe Zhou, Mark Claypool and Robert Kinicki  
Computer Science Department at Worcester Polytechnic Institute  
Worcester, MA USA  
{jeffz,claypool,rek}@cs.wpi.edu

**Abstract**—The growth of mobile streaming applications has increased the importance of a better understanding of 3G EV-DO performance. This paper provides performance measurements for a series of experiments that offer an in-depth, cross-layer understanding of EV-DO. Insights are provided for data rate control, packet encapsulation/combination, early completion, packet interarrival time, packet loss pattern and queue management. Key results are: data rates for non-mobile EV-DO users are more stable than previously reported with little variance over short time scales and more random behavior over long time scales; low-level packet encapsulation/combination happens frequently; and custom service provider configurations (e.g., advanced queue management), and operating system differences (e.g., clock granularity for device drivers) can obfuscate network behavior relative to EV-DO specifications.

**Index Terms**—3G, EV-DO, bandwidth estimation

## I. INTRODUCTION

The emergence of smart phones as a popular Internet mobile platform and the increased use of cellular access cards on laptops has fueled the pressure for cellular network providers to develop technologies that enable higher capacity to support the plethora of new streaming applications designed for mobile wireless clients.

Effective mobile video streaming requires video servers that adapt quickly and efficiently to the variability in wireless transmission conditions, particularly the transmission rates. Bandwidth estimation has been proposed and investigated as a technique to optimally adjust streamed video transmission rates to match the available bandwidth over the Internet path. However, historically, bandwidth estimation explorations have only considered cases where the last hop is either a wired LAN or an IEEE 802.11 wireless access point. Bandwidth estimation research for cellular networks is challenging, with only sparse published measurements, due to: the volatile nature of mobile wireless transmissions, limited access by external researchers to proprietary cellular network infrastructures, difficulties in managing controlled measurements over cellular network infrastructures, and no proven, publicly-available cellular network simulators [1], [2].

This paper presents measurement studies of the widely deployed 3G EV-DO Rev. A cellular network technology. By investigating low-level EV-DO characteristics, this research assists in the ultimate goal of developing effective bandwidth estimation techniques for mobile streaming applications. Important EV-DO details uncovered include EV-DO packet

overhead, frequency and duration of transmission data rates, EV-DO packet fragmentation and encapsulation, EV-DO and higher layer loss characteristics, and observed signal strengths. While providing a general understanding of a deployed EV-DO network, these low-level characteristics facilitate the development of an EV-DO simulator or they can be used by end hosts to improve application performance over EV-DO.

Our experiments show high variability in network conditions when the client location changes. However, measured stationary EV-DO client reception characteristics have less variability than previously reported [3]. Over time scales of a few minutes, throughput at EV-DO clients shows little variance, while over interval of tens of minutes, EV-DO client throughput changes unpredictably. Only relatively short sampling periods (less than 1 second) are needed for a consistent estimate of the standard deviation of network bandwidth. Operating systems provide timing granularities that are smaller (1 millisecond in Windows) and larger (increments of 4 milliseconds in Linux) than those provided by EV-DO packet transmissions (1.667 milliseconds). EV-DO packet errors are infrequent but when they do occur, they are bursty. The effects of upper layer queue management (e.g., RED [4]) often obfuscate EV-DO network behavior beyond what can be inferred in the specification.

The paper is organized as follows: Section II provides an overview of previous related research; Section III introduces EV-DO networks; Section IV details our measurement methodology and experiment design; Section V presents and analyzes experimental results and Section VI concludes by summarizing our main findings.

## II. RELATED WORK

To the best of our knowledge, there are no bandwidth estimation techniques designed specifically for 3G EV-DO networks. Due to the fundamental differences between EV-DO networks and wireless 802.11 or wireline networks, existing bandwidth estimation tools are likely to be inaccurate in measuring the bandwidth for EV-DO networks. The fundamental challenges for bandwidth estimation over EV-DO include possible packet fragmentation and multiple packet encapsulation, the proportional fair scheduler, a time-division multiplexing slotted forward link, and highly variable transmission rates.

A handful of literature explores the feasibility of directly applying existing bandwidth estimation techniques designed for wireline and wireless IEEE 802.11 to EV-DO networks.

For example, Castellanos *et al.* [2] compare three bandwidth estimation tools (TOPP, SLoPS and pathChirp) in mobile networks. The authors do not find any of these estimation tools to be more outstanding than the others, but the results show that pathChirp performs the best. However, the experiments only use NS-2 based simulations, and the NS-2 configurations are not presented. The paper does not provide details about the mobile network studied (for example, whether it is 1xRTT, 1xEV, or UMTS).

Claypool *et al.* [5] compare TCP throughput with the bandwidth estimation results from a tool called *npath* [6]. The estimation results remain steady at about 525 kbps, but the TCP throughput varies from 300 to 350 kbps. While the authors do not comment on the difference, the result shows *npath* does not estimate the achievable throughput accurately.

Koutsonikolas and Hu [1] use WBest [7] to estimate the achievable throughput of an EV-DO terminal. Their experimental results show that WBest fails to measure the achievable throughput defined either by the maximum TCP throughput or the maximum UDP throughput. The study concludes that bandwidth estimation in EV-DO networks is not feasible because of the high variability in transmission rates exhibited at small time scales.

An interesting phenomena is that the measured TCP throughput in [5] and [1] exhibit different properties. For example, [5] observes TCP throughput being in the 300 to 350 kbps range and reasonably stable over a 24 hour period, while [1] sees TCP throughput changes much more drastically (e.g., for one specific location, TCP throughput is measured at 700 kbps at 4am and about 350 kbps at 10am). A possible explanation of this difference is the year the two sets of measurement are taken, 2006 for [5] and 2009 for [1]. The worldwide number of EV-DO subscribers had grown from 65 million in 2006 [8] to 127 million in 2009 [9]. Although service providers may have built additional network infrastructure to accommodate the increased number of users, the increased density of the EV-DO user population may have caused the throughput variability observed in [1].

Despite the difference between observed achievable throughput, both [5] and [1] find that the achievable throughput measurement by sending TCP/UDP flows stays relatively steady for several minutes at a time, suggesting streaming applications may benefit from using appropriate bandwidth estimators. For instance, a bandwidth estimate valid for several minutes can help choose a suitable transmission rate for a YouTube streaming video since about 98% of YouTube videos are less than 10 minutes long [10]. More generally, an accurate bandwidth estimation can allow for initial configuration of application network parameters (e.g., initial video encoding rate). Subsequent adjustment of bandwidth estimation can be inferred from the statistics of actual application packets received, or by rerunning the bandwidth estimation.

### III. BACKGROUND

This section briefly introduces fundamental EV-DO concepts and terminology. In this paper, *Packet* used alone refers

to an EV-DO packet, whereas protocol together with packet refers to packets at other network layers (e.g., IP packet).

Figure 1 illustrates a typical EV-DO network. *Access Terminals* (ATs) connect to an *Access Point* (AP), also known as a *Radio Node* (RN). APs are managed by *Radio Control Nodes* (RNCs) that offer services including resource and mobility management. *Packet Data Serving Nodes* (PDSNs) connect RNCs into the Internet and manage Point-to-Point Protocol (PPP) sessions for the ATs.

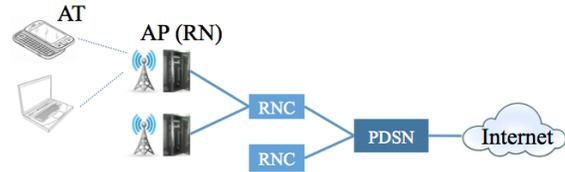


Fig. 1. Typical EV-DO Network Architecture

An EV-DO Rev. A AP transmits data to ATs at a variable rate that depends upon the received DRC (Data Rate Control) index. Every EV-DO time slot (1.667 milliseconds), each AT estimates the highest possible DRC<sup>1</sup> for individual packet reception based on its measured *Signal to Interference and Noise Ratio* (SINR) and sends the DRC index to the AP. Better wireless channel conditions yield higher DRCs where DRC 14 corresponds to the highest data rate and DRC 0 corresponds to the lowest data rate.

The AP maps an AT's DRC into a specific EV-DO MAC layer *transmission format* that consists of a packet size (in bits), a duration (in time slots) and a preamble length (in chips). Generally, a higher DRC implies a larger packet size and/or a shorter duration, which is intended to produce a higher data rate.

The duration, measured in time slots, can be understood as the number of time slots a packet is scheduled to be transmitted. If the SINR at reception time is good, the AT can receive and decode the packet in the first time slot the packet is transmitted, even when the packet could have two time slots scheduled. This event, known as *early completion*, helps compensate for an AT's inaccurate DRC estimation and typical fluctuations in channel signal strengths.

An EV-DO network maintains PPP sessions for each AT. Internet IP packets are encapsulated into PPP packets and sent through several EV-DO layers to the appropriate AT. The formation of PPP packets does not preserve IP packet boundaries, meaning that IP packets can be combined or fragmented by the EV-DO network layers.

EV-DO APs use *Proportional Fairness* (PF) scheduling when transferring data to multiple ATs. The PF scheduler essentially uses a round-robin mechanism across ATs with an adjustment that gives priority to ATs that report to the AP a current DRC that is better than its previous weighted average. While a PF scheduler tends to improve overall system

<sup>1</sup>With an objective of no more than a 1% packet error rate

throughput [11], it increases the variability of IP packet response times.

*Packet error rate (PER)* is the error rate of EV-DO packets. An EV-DO packet can fail to be received after the scheduled number of transmission attempts (as opposed to early completion) and all retransmissions attempted by the EV-DO *Radio Link Protocol*. Higher PERs increase delays in transmissions and possibly packet loss at the IP layer.

#### IV. METHODOLOGY

Our EV-DO experiments utilized a laptop with a 3G EV-DO Rev. A card as a mobile wireless client to measure cellular traffic captured by an EV-DO network sniffer and Wireshark.<sup>2</sup> By carefully controlling the IP packet size and packet rate of downstream traffic from a well-connected server (see Figure 2), the investigation correlates information across network layers to infer EV-DO behavior.

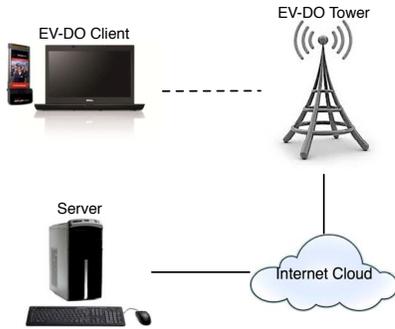


Fig. 2. Experiment Topology

Table I provides hardware details for our setup that sends traffic from a WPI desktop to an EV-DO enabled laptop to study EV-DO characteristics relevant for bandwidth estimation of streaming video down to a mobile device. Similarly, Table II lists the software tools employed in the EV-DO measurements.

TABLE I  
HARDWARE

Hardware	Specification
Laptop (Client)	Dell Precision M4500, Intel Core i5-M560 2.67GHz, 4G memory
Desktop (Server)	Dell Precision, Intel Core i5-660 3.33GHz, 4G memory, Intel 82578DM gigabit Ethernet NIC over 100 Mbps Ethernet connection
EV-DO Aircard	Serria Wireless AC595
EV-DO Network	Verizon EV-DO Rev. A
WPI Egress	500 Mbps over gigabit link

The *EV-DO sniffer* runs on Windows providing DRC and SINR statistics per EV-DO time slot and DRC, SINR, and packet error rate (PER) statistics per second. *QCAT* exports sniffer traces to text files for packet format analysis. Since the sniffer receives no information when the Aircard goes dormant after a 5-second idle interval, our custom *keep-alive tool* sends

TABLE II  
SOFTWARE TOOLS

Tool	Author	Purpose
EV-DO Sniffer	Airvana	Record EV-DO traffic trace for DRC, SINR and PER
QCAT	Qualcomm	Parse EV-DO traffic trace for DRCs and packet transmission formats (every second)
Keep-Alive Tool	WPI	Send dummy traffic periodically to keep aircard connected
Traffic Generation Tool	WPI	Send from server to client with controlled packet rates and sizes

a UDP packet to the server every 3 seconds. Our custom *traffic generating tool* sends UDP traffic from the server to the client using a 10-byte application header that indicates sequence numbers.

While this toolset examines traffic across multiple network layers, there are several capabilities not supported by the tools. Without EV-DO packet-level granularity, these experiments only observe packet encapsulation, overhead and interarrival times indirectly. As EV-DO packet transmission formats are aggregated per second by the tools, per slot-time formats are not available and there is no certainty that client-requested DRCs match the DRCs of the subsequent EV-DO packet sent. The lack of automated support by the sniffer (e.g., command-line arguments) meant experimental runs could not be scheduled and human intervention was required to repeat runs.

Additionally, there is evidence of EV-DO packet loss at high loads. Namely, an observed mismatch in byte count between the application layer and the EV-DO layer is attributed to the aircard missing some EV-DO traffic in the sniffer-recorded summary packet.

#### V. RESULTS

This section provides experimental findings for EV-DO network experiments run between May and November 2011. The three locations and average recorded DRC's at these locations are shown in Table III. In the rest of this paper we refer to these locations by their labels.

TABLE III  
EXPERIMENT LOCATIONS

Label	Location	DRC mode
Location A	A single-family home in Worcester, MA	13
Location B	An academic building in WPI (Worcester Polytechnic Institute)	8
Location C	A high-rise apartment in Cambridge, MA	6

Our EV-DO end device (AT) is stationary for all experiments discussed except for a mobile experiment where the laptop gathers measurements from a car moving about 30 miles per hour. For the stationary experiments, the laptop is placed in the exact same physical spot at each location.

<sup>2</sup><http://www.wireshark.org>

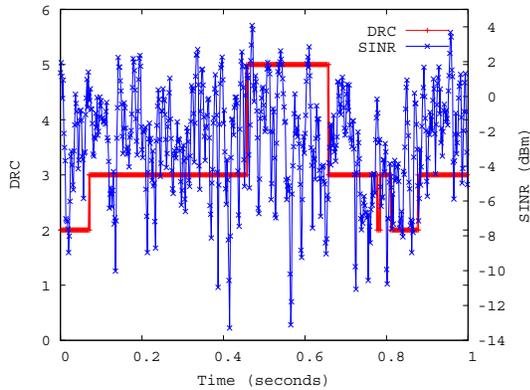


Fig. 3. DRC and SINR correlation over 1 second intervals (mobile)

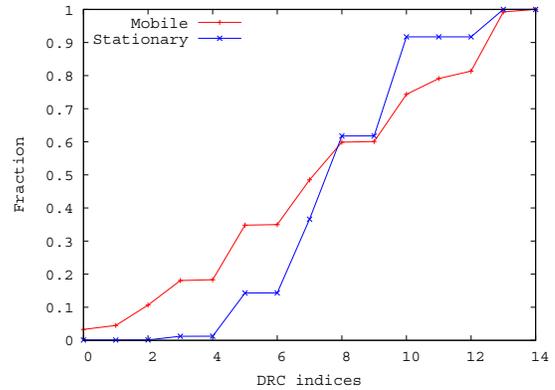


Fig. 4. CDF of DRCs

We are aware that the limited scale of our experiments may not yield ubiquitously representative results. However, previous mobile network data studied in large scale has been presented mostly only from the vantage of the network provider, likely due to the manual process needed to gather data at the AT. Some of our experiments (e.g., throughput versus sample time) are run at locations with different DRC ranges providing some breadth to the representation. Other experimental results (e.g., different packet interarrival times caused by the operating system) are generally applicable. Overall, the results offer a practical instance of an AT-side perspective of a commercial EV-DO network, while highlighting network aspects that could be studied at a larger scale in future work.

#### A. DRC Analysis

DRC indices are forward-link rate requests made per EV-DO time slot from an AT to the AP. Our experiments examine DRC distributions from both stationary and mobile EV-DO endpoints. Previous work studying 3G networks has shown DRCs to be a key attribute in determining EV-DO end-device performance [3].

DRC requests are based on the SINR computed at the AT. Figure 3 shows an example of the correlation between DRC and SINR over a one-second period for a mobile device. Although the DRC fluctuates among 2, 3, and 5, there is no visual correlation between SINR and DRC over this time scale. Thus, we provide no further analysis of performance based on SINR.

Figure 4 depicts the CDF of DRCs reported over two 90-minute experiments. The mobile experiment is run in a car moving at 30 mile per hour as aforementioned and the stationary experiment is run at Location B. Both cases yield nearly a full-range of DRC values, with DRC values of 1-13 seen clearly. The DRC distribution is slightly broader during the mobile experiment. There are few DRC 14 values reported in the mobile case, and none in the stationary case.<sup>3</sup>

<sup>3</sup>Reported DRC 0 values appear to be mode changes by the AT and not actual requests for data transmission rates.

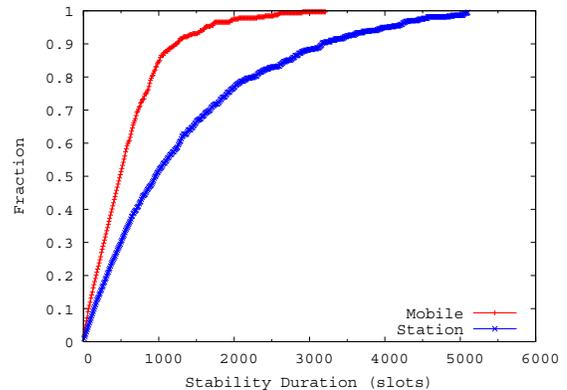


Fig. 5. CDF of DRC steadiness

DRC fluctuation has a direct impact on AP transmission rate variability, which in turn impacts streaming applications and bandwidth estimation. Generally, a more stable DRC distribution should yield steadier streaming data rates, more accurate bandwidth estimation and reduced estimation times.

To examine this variation, AT DRC requests made every EV-DO time slot (1.667 milliseconds) are recorded over a period of 800 seconds. From this data, the size of continuous bursts of non-changing DRC values (referred to as "steady region" in the rest of the text) are extracted and then computed as a fraction of the total number of DRC values. Figure 5 graphs the CDF of this DRC steadiness for the same stationary and mobile experiments.

In the stationary experiment, DRC values remain constant over relatively long time periods. For the stationary AT, over half the time the DRC stays in steady regions of 1000+ time slots (approximately 1.5 seconds), exhibiting more steadiness than reported previously [3]. Almost 10% of the time, the DRCs belong to steady regions of 3000+ time slots (more than 5 seconds). Even in the mobile AT case, DRCs remain in steady regions of 600+ time slots (more than a second) over half the time.

### B. Throughput versus Sample Time

Determining the sample size needed for accurate measurements is important, both for experimentation and for tools such as bandwidth estimation that make throughput predictions from a minimal number of samples. To this end, three traces collected at all three locations (representing three different AT throughputs) are selected for comparison. The three traces last for 200 seconds during which time the server sends saturating UDP traffic to the AT. The average throughputs are: 285 kbps for Location C, 441 kbps for Location B, and 813 kbps for Location A. The packet size used for Location B and C is 1500 bytes and the packet size used for Location A is 33 bytes.

Figure 6 graphs the standard deviation of the average throughput versus the sampling interval in seconds. As expected, the standard deviation stabilizes with increased sample interval. However, the standard deviation is quite low even for very few samples. This suggests bandwidth estimation techniques may be able to obtain fairly accurate predictions with relatively few measurements. Note, for all three traces, the coefficient of variation (standard deviation divided by the mean) is approximately 1.

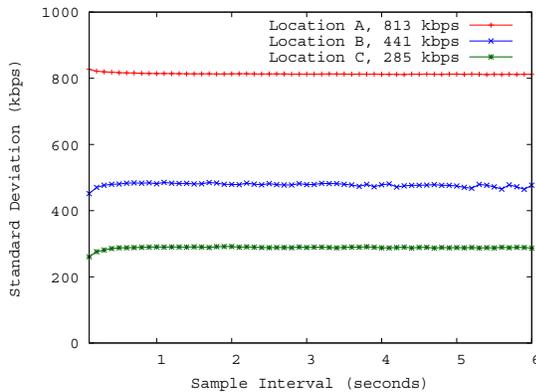


Fig. 6. Standard deviation versus sampling interval

### C. EV-DO Packet Analysis

IP packets can be fragmented during transfer to EV-DO ATs. To better understand how and when packet fragmentation occurs within EV-DO, experiments run at Location A that elicited predominantly a 14 DRC are examined over a range of IP packet sizes. UDP packets are sent from the server to the AT in ten-packet batches. All packets in a batch are the same size sent at a rate of 1 packet per 3 seconds. After sending one batch, the server pauses 10 seconds, increases the packet size by 1 byte, and sends the next batch. The 3 second interval is chosen since it is long enough to distinguish packet(s) sent at the EV-DO layer for each UDP packet, but short enough to avoid the AT going dormant and disconnecting. If there is no data transmission in 5 seconds the AT will enter a dormant state which affects the statistics of the next packet.

Figure 7 shows an experimental segment where the UDP packet size varies from 612 bytes to 620 bytes (shown by

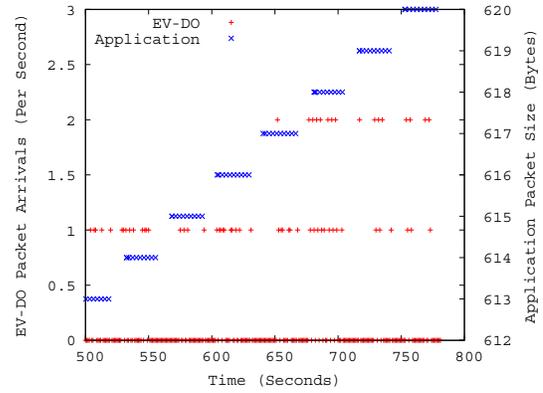


Fig. 7. EV-DO packet arrivals with increasing IP packet sizes over time

the right-hand y-axis) and EV-DO packet arrivals per second (shown by the left-hand y-axis). At approximately 675 seconds, some of the EV-DO packet arrival rates jump from 1 per second to 2 per second. This suggests that at 617 bytes the IP packet size is too big to fit in one EV-DO packet and must be fragmented into two EV-DO packets. In addition to confirming the presence of EV-DO packet fragmentation, this data also provides enough detail to compute EV-DO header sizes. At DRC 14, EV-DO sends 640 bytes per slot and Figure 7 implies the EV-DO header is 24 bytes.

To ascertain EV-DO overhead for other DRC values, experiments with three distinct IP packet sizes (50 bytes, 400 bytes, and 750 bytes) are conducted at Location B and C. Each experiment has the server send packets to the AT at a fixed packet rate for 10 seconds, followed by a three-second quiet period where no packets are sent, then doubling the packet rate and repeating transmissions up to a saturating rate (higher than the EV-DO maximum). Each experiment is repeated 5 times for each packet size and location combination. Figure 8 shows the EV-DO packets arrivals for one such experiment (packet size 50 bytes, Location C). For example, at around 90 seconds, the previous batch of IP packets are received (as EV-DO packets), a 3 seconds pause is executed, and the next batch of IP packets (sent at double the previous packet rate) start to be received. In correspondence to the doubling of IP packets sending rate, the EV-DO packet arrival rate increases every 13 seconds up to the saturation point, with no noticeable drop-off in packet arrival rates once saturation is reached. EV-DO packet arrival rate seems to increase less when packet rate is higher (e.g., for the two batches at time 128 and 256), suggesting EV-DO is combining IP packets.

Figure 9 presents the ratio of the number of bytes of the EV-DO traffic to that of IP layer traffic (including IP header) for all experimental runs. Some experiments finished earlier than others because those experiments use the larger packet sizes and reach the saturation faster. The EV-DO overhead converges from about 15:1 at a low packet per second rate, to close to 1 at a high packet per second rate. The high overhead for small IP packets at low sending rate is because EV-DO

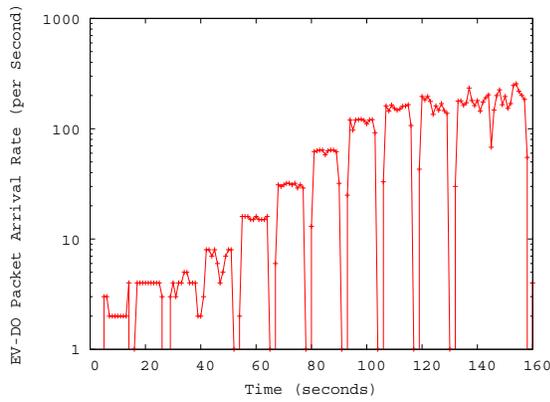


Fig. 8. EV-DO packets over time

does not have a backlog from which the IP packets can be combined. At high sending rate, EV-DO combines IP packet to help keep overhead low.

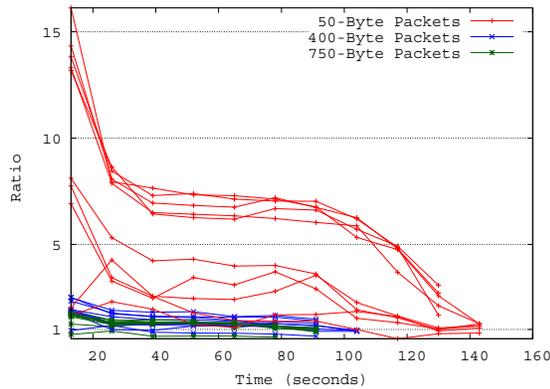


Fig. 9. Ratio of EV-DO traffic to IP traffic

#### D. Application Packet Analysis

Some bandwidth estimation techniques use packet dispersion to infer available bandwidth. These techniques have been proven to work well in wired networks [12] and IEEE 802.11 networks [7], but to the best of our knowledge, they have not been thoroughly studied in mobile phone networks. The next experiments offer a preliminary view of packet dispersion in EV-DO networks.

Figure 10 shows the CDF of UDP packet interarrival times gathered at the receiving laptop’s application layer of the EV-DO AT for two experiments. The two experiments are run at Location A with the same setup – the only difference being the operating system, Microsoft Windows (top) or Linux (bottom). The packet size is set to 50 bytes to reduce the possibility of packet fragmentation.

The graphs demonstrate that the packet interarrival times are significantly affected by the operating system. The Windows interarrival times cluster in multiples of 1 millisecond and the Linux interarrival times cluster in multiples of 4 milliseconds.

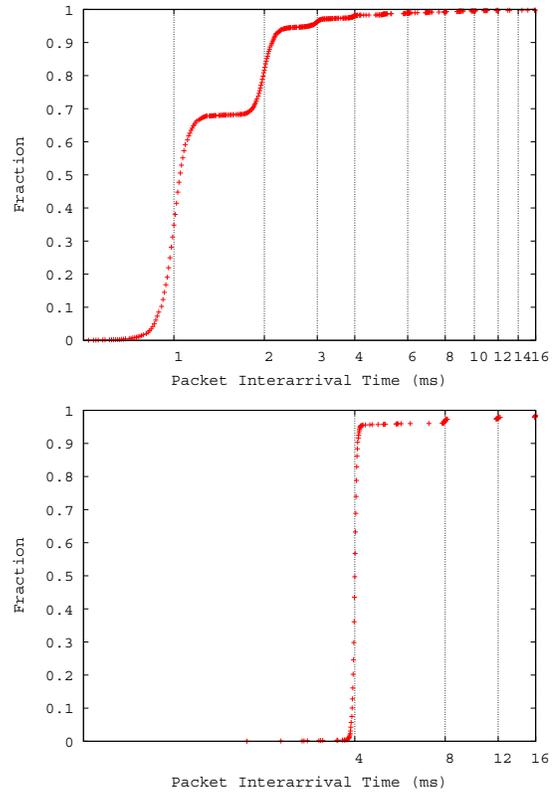


Fig. 10. Interarrival time of application packets (Windows top, Linux bottom)

IP Packets that arrive 1 millisecond apart are likely sent in the same EV-DO packet (EV-DO packets must be separated by at least one EV-DO time slot). The timing clusters observed are likely due to distinct buffering schemes in each OS stack or from different timing mechanisms in the device driver. Thus, any application-level packet dispersion techniques that infer bandwidth based on packet timings on the order of a few milliseconds may not be able to ascertain EV-DO packet arrivals accurately.

#### E. PER and Out-of-Order Packets

When an AT requests a DRC from the AP, the DRC is chosen to achieve less than a 1% EV-DO packet error rate (PER) given the measured SINR. For a 500 second interval, the PER is recorded each second at Location B. Over the 500 seconds, only 15 seconds are observed with non-zero PER. Figure 11 plots the CDF and CCDF of PER observed when there are errors. From the figure, 90% of the time the PER is between 0.15 to 0.40, with the bulk of PERs around 0.20. The CCDF also includes seconds when there is no packet error to give a more complete view of the data.

Figure 12 shows the IP packet loss for experiments that are run in Location A with a slowly increasing data rate to observe loss behavior at link saturation. With small, 33-byte IP packets (top), the rate increases steadily until link saturation is reached and there is significant loss. There is a noticeable

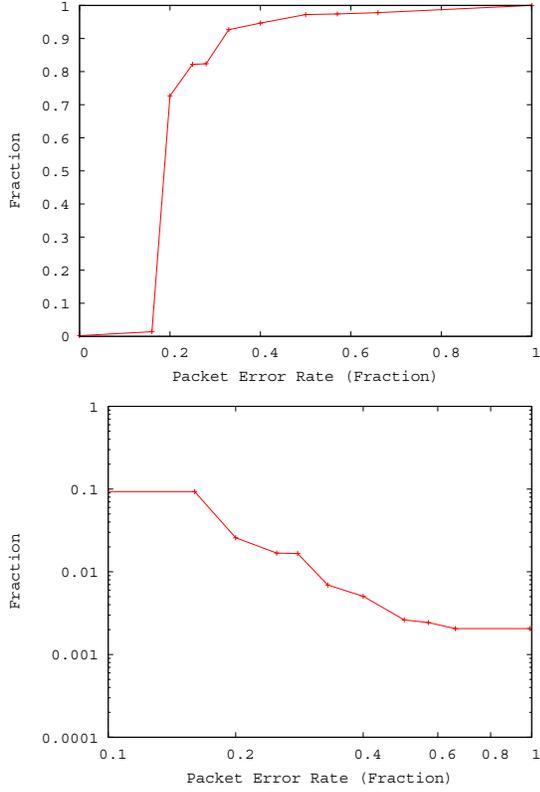


Fig. 11. CDF of PER (top), CCDF of PER (bottom)

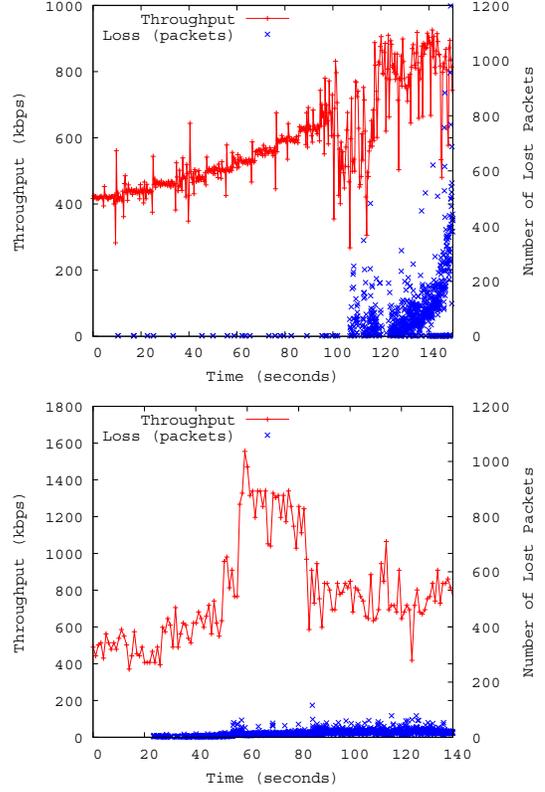


Fig. 12. Packet loss with increasing throughput (top: 33 byte packets, bottom: 1500 byte packets)

dip in throughput around 100 seconds. With large, 1500-byte IP packets (bottom), there is a similar increase in rate until saturation, but a noticeable increase in throughput around 60 seconds before the rate decreases to the saturation point.

Application sequence numbers can be used to determine lost packets, but the analysis is made difficult by out-of-order arrivals. To infer out-of-order packets caused by the EV-DO network, two experiments are run. Using 33 byte packets, the first experiment sends data at 100 kbps from the server to the AT (achieved rate of 85 kbps) and the second experiment sends data at 750 kbps (achieved rate of 585 kbps). Both experiments run for three minutes. In the first experiment, there are no out-of-order packets, while in the second experiment there are 167 out-of-order packets, accounting for 0.04% of all 379198 packets that are received.

Examining traceroute output before and after the experiments indicates routing changes occurred in both the first and second experiments. These route changes, and not the EV-DO link, could be the cause of the out-of-order packets observed in the second experiment.

#### F. Early completion

EV-DO early completion can be observed in Table IV. In the table, each DRC implies one or more combinations of packet sizes (in bits) and time slots scheduled to send those packets. This experiment runs for 278 seconds, with the server sending 33-byte packets at a rate of 1 Mbps to the EV-DO end device

TABLE IV  
PACKET TRANSMISSION FORMAT REPORT

DRC	Size	Slots	slot1	slot2	slot3	slot4	PER
0	1024	16	0	7	0	0	0
2	128	8	12	0	0	0	0
2	512	8	5	0	0	0	0
2	1024	8	14	12	0	0	0
3	128	4	4	0	0	0	0
3	512	4	5	0	0	0	0
3	1024	4	<b>113</b>	<b>17</b>	0	0	0
4	1024	2	33	1			0
5	2048	4	14	937	60	0	<b>0.4%</b>
6	1024	1	31				0
7	2048	2	1475	18373			<b>0.3%</b>
8	3072	2	0	26774			<b>0.7%</b>
9	2048	1	4				0
10	4096	2	<b>0</b>	<b>4984</b>			<b>1.4%</b>
13	5120	2	0	94			0

at Location A (the receiving rate is 377 kbps). Transmission formats which have no packets transmitted are omitted from the table. From the table, early completion is prevalent for low DRCs.<sup>4</sup> For example, DRC 3 with 1024 bit EV-DO packets has 113 packets complete in slot 1 and 17 in slot 2 and no packets in the subsequent slots. However, for higher DRCs, all DRC slots are needed for transmission. For example, DRC 10 has all EV-DO packet completions during slot 2.

<sup>4</sup>DRC 0 and 2 have more than 4 slots but no EV-DO packets are sent in these slots.

In most cases the PER is less than 1% [13], though at DRC 10 the PER is a little higher. The PER exhibits an increasing trend with the DRC going up, implying that packets are more likely to be lost at higher DRC for the same location.

### G. Queue management

Saturating traffic is sent from our server to the EV-DO end device, with the packet loss pattern is plotted in Figure 13. The experiment is conducted at Location A, with a fixed packet size of 33 bytes. All losses are in bursts of 3 or more packets, with most losses in bursts 400-1000 packets in a row. Figure 14 shows the corresponding CDF of the time between two consecutive packet losses. Loss intervals are mostly between 30 to 100 milliseconds. The bursty nature of the packet drops suggests queue management, likely above the EV-DO layer, perhaps employing an algorithm such as Random Early Detection (RED) [4].

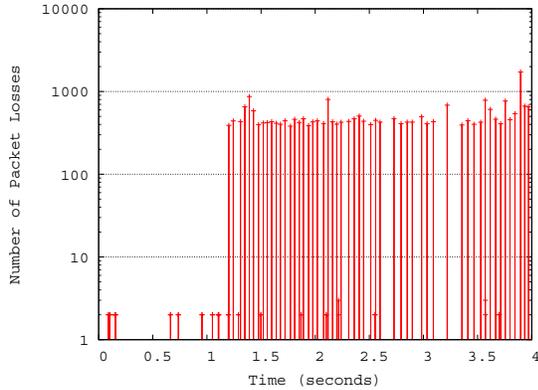


Fig. 13. Packet loss over time

There is no significant packet loss in Figure 13 in the initial 1.2 seconds, suggesting a downlink queue at the AP. Using the method in [14], we calculate the queue size to be about 235 bytes. This size does not change significantly with the packet size which implies the queue is a per-byte queue, rather than a per-packet queue.

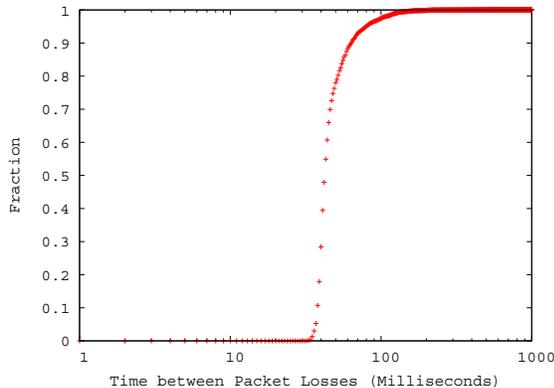


Fig. 14. CDF of time between packet losses

## VI. CONCLUSIONS

Though wide-spread in deployment, low-level 3G EV-DO analysis suitable for simulation and bandwidth estimation techniques for streaming video are not well-established. Through careful experimentation and measurements with an EV-DO sniffer, Wireshark and application-level tools, this paper provides insights into EV-DO performance not previously reported or demonstrated. The results should be useful for those building EV-DO simulators or building bandwidth estimation tools for streaming applications over EV-DO – both areas of our future work.

Our preliminary investigations show Data Rate Controls (DRCs) requested by the EV-DO ATs are centered around a few values in the case of stationary end-hosts and somewhat more scattered in case of mobile end-hosts. In all cases, DRCs observed are steadier than reported in previous work. Measured throughputs over these DRCs exhibit consistent variance, suggesting accurate bandwidth predictions may be made after relatively few samples.

EV-DO constructs forward link packets based on the maximum packet size the current DRC allows regardless of the original IP packet boundaries. IP packets are observed being combined and fragmented to best fit into each EV-DO packet. DRC 14 packets have 24-byte headers, but the exact size of the EV-DO headers appears to vary with different DRCs.

The interarrival times of application packets are multiples of a few milliseconds that differ by operating system on the EV-DO end-host. At such granularities, the interarrival times observed at the application layer are not necessarily a direct reflection of the interarrival times at the EV-DO layer.

Normal packet error rates (PERs) are less than 1%, but losses tend to be bursty. When losses do occur, per-second loss rates typically range from 20-40%.

Early completion can achieve significant gains. For DRCs that have a nominal transmission duration longer than 1 time slot, in some cases most of the packets are successfully transmitted in the first time slot. In contrast, in some cases, particularly for higher DRCs, none of the packets are received in the first time slot.

EV-DO networks may use queue management to stop accepting packets into the queue when the network is congested. When an EV-DO network receives incoming packets at a rate higher than EV-DO can forward, the packets are dropped in a regular discrete fashion rather than a continuous fashion as would be the case with a drop-tail queue.

Possible future work includes more measurements, particularly with different environmental (e.g., urban versus suburban) and end-host (e.g., mobile versus stationary) characteristics.

## REFERENCES

- [1] Dimitrios Koutsonikolas and Y. Charlie Hu, "On the Feasibility of Bandwidth Estimation in 1x EVDO Networks," in *Proceedings of ACM Workshop on Mobile Internet through Cellular Networks*, Beijing, China, Sep. 2009.
- [2] C. U. Castellanos, D. L. Villa, O. M. Teyeb, J. Elling and J. Wigard, "Comparison of Available Bandwidth Estimation Techniques in Packet-switched Mobile Networks," in *Proceedings of IEEE Personal, Indoor and Mobile Radio Comm. (PIMRC)*, Helsinki, Finland, Sep. 2006.
- [3] X. Liu, A. Sridharan, S. Machiraju, M. Seshadri, and H. Zhang, "Experiences in a 3G Network: Interplay between the Wireless Channel and Applications," in *The ACM Annual International Conference on Mobile Computing and Networking*, San Francisco, CA, USA, Sep. 2008.
- [4] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Transactions on Networking*, Aug. 1993.
- [5] M. Claypool, R. Kinicki, W. Lee, M. Li, and G. Ratner, "Characterization by Measurement of a CDMA 1x EVDO Network," in *International ICST Wireless Internet Conference*, Boston, MA, USA, Aug. 2006.
- [6] C. Dovrolis, P. Ramanathan, and D. Moore, "Packet-Dispersion Techniques and a Capacity-Estimation Methodology," *IEEE/ACM Transactions on Networking*, vol. 12, no. 6, pp. 963–977, 2004.
- [7] M. Li, M. Claypool, and R. Kinicki, "WBest: A Bandwidth Estimation Tool for IEEE 802.11 Wireless Networks," in *Proceedings of IEEE LCN Conference*, Montreal, QC, Canada, Oct. 2008.
- [8] Q. Inc., "EV-DO Rev. A and B: Wireless Broadband for the Masses," 2007, <http://www.qualcomm.com/documents/ev-do-rev-and-b-wireless-broadband-masses-whitepaper>.
- [9] CDMA Development Group, "Over 1/2 Billion CDMA Subscribers Served and 128 Million EV-DO," 2009, <http://www.wirelessandmobilenews.com/2009/09/over-12-billion-cdma-subscribers-served-and-128-million-ev-do.html>.
- [10] Xu Cheng, Dale, C. and Jiangchuan Liu, "Statistics and Social Network of YouTube Videos," in *International Workshop on Quality of Service*, Enschede, The Netherlands, Jun. 2008.
- [11] Q. Bi, "A Forward Link Performance Study of the 1xEV-DO Rev. 0 System Using Field Measurements and Simulations," *Bell Lab Technical Journal*, Mar. 2004.
- [12] M. Jain and C. Dovrolis, "End-to-End Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput," in *ACM SIGCOMM*, Pittsburgh, PA, USA, Aug. 2002, pp. 295–308.
- [13] Bhushan, N., Lott, C., Black, P., Attar, R., Yu-Cheun Jou, Fan, M., Ghosh, D., Au, J., "CDMA2000 1xEV-DO Revision A: A Physical Layer and MAC Layer Overview," *IEEE Communications Magazine*, vol. 44, no. 2, pp. 75–87, 2006.
- [14] M. Claypool, R. Kinicki, M. Li, J. Nichols, and H. Wu, "Inferring Queue Sizes in Access Networks by Active Measurement," in *Passive and Active Network Measurement (PAM)*, Antibes Juan-les-Pins, France, Apr. 2004.