

On the Feasibility of Bandwidth Estimation in 1x EVDO Networks

Dimitrios Koutsonikolas, Y. Charlie Hu
Purdue University, West Lafayette, IN, USA
{dkoutson, ychu}@purdue.edu

ABSTRACT

The problem of bandwidth estimation has been extensively studied in the wired Internet and recently, in 802.11 wireless networks, however, no tool has been developed so far for cellular networks. Although cellular networks share many common characteristics with 802.11 WLANs, they also have many differences which make it unclear whether tools developed for 802.11 WLANs can also be used in cellular networks. This paper presents the first study of the feasibility of fast bandwidth estimation in 3G networks. We study the applicability of a state-of-the-art probe-based tool for 802.11 WLANs on a commercial 1x EVDO network. We find that estimating bandwidth with a probe-based tool is very difficult in EVDO networks, due to the short-scale dynamics in this type of networks.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design—*Wireless Communication*

General Terms

Measurement

1. INTRODUCTION

Cellular networks are increasingly being used for providing Internet access. The advent of 3G technology has provided data rates sufficient to support many Internet applications, such as Email, Web browsing, bulk data transfers, and media streaming, making cellular networks a strong competitor to 802.11 WLANs as a broadband access technology.

CDMA 1x Evolution-Data Optimized (EVDO) [1] is one of the 3G technologies adopted by many cellular providers in several countries. The more recent 1x EVDO Revision A system offers peak data rates of 3.1Mbps on the downlink, from the base station (BS) to the user, and 1.8Mbps on the uplink. The 1x EVDO downlink is a TDMA slotted system. The BS selects one user in each slot and transmits to it with full power at the requested data rate. Central to this selection is the use of a Proportional Fair (PF) algorithm, an

opportunistic scheduling scheme that tries to increase channel capacity while providing fairness to all users in the same sector. The main idea is that not all users will have poor channel quality all the time; hence, the overall system throughput can be improved if each user is only served when its channel is good.

Bandwidth estimation refers to the measurement of a bandwidth-related metric over a network path, performed only by the end hosts, without access to the intermediate routers. The most obvious way to perform the measurement is via a direct download from the source to the destination of the path of interest. To avoid the intrusion of a direct download and reduce the measurement time, researchers have developed a series of tools that try to estimate a bandwidth-related metric by sending only a few probe packets and exploiting techniques such as packet pair or packet train dispersion [9, 6]. Bandwidth estimation tools will become important for 3G cellular networks, for applications such as video streaming (as in other types of networks), but also potentially for allowing devices such as smartphones, with multiple interfaces, to dynamically select the best technology (802.11 or 3G) to connect to.

The problem of bandwidth estimation has been extensively studied in the wired Internet and recently, in 802.11 wireless networks, however, no tool has been developed so far for cellular networks. Although cellular networks share many characteristics with 802.11 WLANs, they also have many differences which make it unclear whether tools developed for 802.11 WLANs can also be used in a cellular network. In particular, the rapid channel rate adaptation (at timescales of 1 millisecond [18]) and the use of the PF scheduler at the BS make it unclear whether bandwidth estimation is even feasible in such networks.

This paper presents the first study of the feasibility of fast bandwidth estimation in 3G cellular networks. Using a state-of-the-art bandwidth estimation tool for 802.11 WLANs [17], we conducted measurements in a commercial 1x EVDO Revision A network at different locations and different times of the day with the goal of evaluating the applicability of the tool on a 3G network. We discuss the challenges associated with bandwidth measurements in commercial 3G networks in contrast to in controlled lab environments. Our results show that fast bandwidth estimation through a few probe packets is not feasible in EVDO networks. We also analyze our measurement data to reveal the reasons for this negative result.

2. RELATED WORK

There have been many tools proposed for capacity and available bandwidth estimation in the wired Internet [13, 6, 7, 19, 9, 20, 25, 25, 24, 8].

Recently, researchers have shown that these tools cannot be used in wireless networks, due to the different characteristics of these

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.

MICNET'09, September 21, 2009, Beijing, China.

Copyright 2009 ACM 978-1-60558-753-0/09/09 ...\$10.00.

networks which invalidate many of the assumptions made for the wired Internet [14, 10, 16].

This observation has lead to new tools, which take into account the different characteristics of wireless networks, and also provide a faster estimation than Internet tools, to deal with changes in the bandwidth due to channel variability [14, 10, 17]. One common characteristic of these tools is that they have only been evaluated in controlled environments, under controlled settings and interferers, and only for 802.11 CSMA-based WLANs. In contrast, there is almost no work on bandwidth estimation for TDMA- or CDMA-based 3G cellular networks. Only [21] proposes a method for available bandwidth estimation both for WLANs and for cellular networks. However, this method is integrated with the RTP protocol for video streaming, and cannot be used as a generic tool for fast bandwidth estimation. Also, [3] compares three bandwidth estimation tools proposed for the Internet in cellular networks but only using simulations. To our best knowledge, this is the first work that studies the feasibility of fast bandwidth estimation, using a probe-based tool, in a commercial 3G network.

Recently, there have been several measurement studies of 3G networks, both 1x EVDO [4, 15, 18] and UMTS [26, 22, 23, 12, 2, 11, 5]. The majority of these works [4, 15, 18, 22, 23, 12] study TCP performance in 3G networks by performing direct downloads. A few other works focus on delay-related metrics [2, 5]. Finally, [26, 11] try to estimate the capacity of 3G networks for data and video calls.

Our work is different from all these works in that we are interested in *estimating* the bandwidth of an EVDO downlink using a fast probe-based tool, rather than *directly measuring* it through a download. The only work that also uses a similar approach is [4], which, in addition to TCP downloads, also estimates the end-to-end capacity of an 1x EVDO Revision 0 downlink using the packet pair technique. As we discuss in Section 3, applications are mostly interested in the achievable throughput rather than the raw capacity. Another feature that distinguishes our work from [4] is that our data were obtained from measurements conducted in different locations. In contrast, the data in [4], as well as in most of the existing works (with the exception of [18, 26]), were collected from a single location.

3. MEASUREMENT METHODOLOGY

3.1 What is the metric of interest?

Two of the bandwidth-related metrics that have been traditionally used in the wired Internet are the end-to-end *capacity* and the end-to-end *available bandwidth*. The end-to-end capacity is the maximum rate a path can provide to a flow when there is no other flow in that path. In contrast, the end-to-end available bandwidth refers to the maximum rate at which a new flow can transmit over a path without reducing the rate of the existing flows in that path. Another metric of interest for a user is the *achievable throughput* of a new flow, i.e., the actual fraction of the bandwidth that will be allocated to a new flow over a path. Note that the achievable throughput can actually be higher than the available bandwidth. Consider a single-link path with capacity C saturated by a single TCP connection. The available bandwidth in this path is zero, but the achievable throughput of a new TCP flow will be $C/2$ if it has the same RTT as the existing flow.¹

The importance of achievable throughput becomes more prominent for clients that connect to the Internet through one of the wire-

less broadband access technologies (e.g., 802.11 WLANs or 3G networks). In contrast to Internet links that are shared by many flows, the wireless broadband access link only carries traffic for one user. However, due to the broadcast property of the wireless medium, a flow has to contend with flows traversing other wireless links in the neighborhood for a share of the medium. Thus, in addition to the Internet-type of cross-traffic, in the form of two flows crossing the same link and competing for the link's bandwidth, in a wireless network we have two additional cases of competition between flows: (i) flows with the same source but different destinations (e.g., from an AP to two different clients) and (ii) flows with different sources and different destinations (e.g., from AP1 to client C1 and from AP2 to client C2).² In all these cases, the available bandwidth may be zero for a new flow, but the new flow will still obtain a fraction of the channel's bandwidth.

The majority of the bandwidth estimation tools for the Internet calculate the capacity or the available bandwidth. The knowledge of these two metrics is valuable for ISPs and network operators, for operations such as capacity planning, detection of congested or underutilized links, routing and traffic engineering, admission control, and QoS provisioning, which require a global view of the network. However, applications ran by individual users do not see a global view of the network. Hence, from a single client's perspective, achievable throughput seems to be the most important metric, and this is particularly true in wireless networks, due to the contention on a per-node basis for the medium's bandwidth. In the example with the two APs mentioned above, client C2 does not care if the channel is saturated by client C1 that belongs to another network and it should not be prevented from sending or receiving its own traffic due to zero available bandwidth. For these reasons, we will use the achievable throughput in this paper as the metric of interest.

3.2 Overview of the probe-based tool

The WBest tool [17] calculates the effective capacity, the achievable throughput, and the available bandwidth, on a network path where the last hop (which is also the bottleneck link) is over an 802.11 wireless network.

WBest has been shown in [17] to offer much more accurate bandwidth estimations compared to three popular tools proposed for the Internet. Compared to other tools for wireless networks, it is very fast, using only one train of 30 packets to estimate the achievable throughput and the available bandwidth (in contrast, e.g., DietTopp [10] sends a series of packet trains, each one at a different rate), and it calculates all three bandwidth-related metrics (in contrast, e.g., ProbeGap [14] only estimates available bandwidth but it requires a priori knowledge of the capacity), being the only tool to measure achievable throughput. Finally, the code for WBest is publicly available. For these reasons, we selected WBest for our study.

WBest uses a two-stage algorithm. In the first stage, n packet pairs are sent to estimate the capacity C_e . To mitigate the effect of cross and contending traffic, C_e is calculated as the *median* instead of the *mean* of the n dispersion samples from the n packet pairs. In the second stage, a train of m probe packets is sent at rate C_e . The achievable throughput A_t is equal to the average dispersion rate of the train [17]. The available bandwidth is then calculated as a function of C_e and A_t .

The measurement of achievable throughput becomes more challenging in an EVDO network. In contrast to an 802.11 AP or the interface of an Internet router, that simply transmit packets in FIFO order, the PF scheduler at the BS implements a type of QoS, guar-

¹The achievable throughput of a TCP flow is actually another metric in the Internet, known as *Bulk Transfer Capacity (BTC)*.

²This is referred to as *contending traffic* in [16], as opposed to cross-traffic.

anteeing that all the arriving flows will receive a (proportionally) fair share. This TDMA-based scheduling may insert random intervals between packet transmissions, when the BS decides to serve some other user. This makes it unclear whether it is even feasible to estimate the achievable throughput using a probe-based tool in an EVDO network. We are investigating the feasibility of this approach in the rest of the paper.

3.3 The challenge of determining the ground truth

The evaluation of the accuracy of any bandwidth estimation tool always faces the challenge of correctly determining the *ground truth*, i.e., the actual value of the metric estimated by the tool. This is more challenging in wireless networks, since the capacity of wireless links is affected by link layer mechanisms such as rate adaptation or ARQ, and the available bandwidth or achievable throughput over a given link is affected not only by cross traffic over that link but also by the presence of other clients. To deal with this challenge, previous tools for 802.11 networks [14, 10, 17] were evaluated in a controlled environment, with additional transmitters used to create the desired amount of interference, as well as the desired amount and type of contending traffic.

In our study, we are interested in evaluating the accuracy of a bandwidth estimation tool in a commercial 3G network, with no control over the environment and no knowledge about the amount or type of traffic from other users of the network during our experiments. Under these conditions, the definition of ground truth is infeasible for the available bandwidth and not clear even for the achievable throughput.


We decided to use the throughput of TCP and UDP downloads as ground truth – we call this “best effort” ground truth. We used both TCP and UDP, since different smartphone applications have different requirements; the goal is to evaluate how accurately a probe-based tool such as WBest can predict UDP or TCP download throughput through the estimated A_t metric. The duration of the download is critical for the definition of the ground truth. We selected a 30-sec interval since it is long enough to provide a stable, average behavior, as we show in Section 4, and at the same time short enough to avoid the impact of user dynamics (e.g., due to mobility). In addition, this interval is sufficient for many smartphone applications (e.g., webpage downloads, youtube videos, etc).

3.4 Measurement setup

We used a laptop with Intel Core(TM)2 Duo 1.8GHz CPU and 2GB RAM running Windows Vista as a client and a PC with P4 2GHz CPU and 512MB RAM running Mandrake Linux 10.1 (kernel 2.6.8-12) as a server. The client was moved to different locations and used a Verizon HTC Touch Pro smartphone to connect to the commercial Verizon 1x EVDO Revision A network. The smartphone ran Windows Mobile 6.1, and it operated as a modem, connected to the client via USB. The server is located in our lab and it is connected to the Internet through a 100Mbps LAN. In all our measurements, we used the default values of WBest ($n = 30$ packet pairs, $m = 30$ packets, 1460-byte probe size).

We repeated our measurements in three different locations in West Lafayette, denoted as LOC1, LOC2, LOC3. In each location, the experiments were repeated at two or three different times of the same day (4:00-5:00 AM, 10:00 AM-11:00 AM, and 9:00-10:00 PM).

Each experiment, at a given location and time of the day, denoted as (location, time), was performed as follows. First, the server sent TCP traffic to the client for 30 sec, and the client recorded the number of bytes received in each second. Then the server sent



(a) UDP throughputs.

(b) TCP throughputs.

Figure 1: Throughputs of 12 30-sec UDP and TCP transfers taken every 5 min within 1 hour for each (location, time). The standard deviation varies between 0.09Mbps and 0.23Mbps for UDP and between 0.08Mbps and 0.39Mbps for TCP.

1460-byte UDP packets to the client at the maximum possible rate for 30 sec, and the client again recorded the number of bytes received in each second. Finally, WBest was run 15 times to obtain 15 estimates of the downstream achievable throughput. This process was repeated every 5 min within 1 hour, i.e., we got 12 sets of measurements for each (location, time).

4. “BEST EFFORT” GROUND TRUTH

In this section we examine whether the throughput of a 30-sec UDP or TCP transfer preceding the WBest measurements can be used as the ground truth to evaluate the accuracy of the tool.

Figures 1(a), 1(b) plot the throughputs of the 12 30-sec UDP and TCP transfers conducted within 1 hour for each (location, time). Two consecutive points in each curve correspond to two measurements taken 5 min apart from each other. We also calculated the standard deviation over the 12 measurements for each (location, time).

These two figures show a picture very different from the one presented in [4]. In [4] the authors observed that TCP throughputs were very stable throughout a 24-hour period, ranging from 300-400 Kbps. In contrast, we observe that for the same location, both UDP and TCP throughputs can be very different at different times of the day, and they cover a much larger range, of 120-1176 Kbps for UDP, and 77-1474 Kbps for TCP.

In spite of large variations at different times of day, we observe that throughputs remain relatively stable for smaller time intervals (several minutes) for the majority of locations and times of day, in particular for UDP. The standard deviations of the 12 measurements for UDP in locations LOC1 and LOC2 are 93-141 Kbps for all 3 times. In LOC3, they are slightly higher, equal to 209 Kbps for the 9PM experiment and 236 Kbps for the 10AM experiment, implying that throughput is less stable in that location. For TCP, the standard deviations generally are larger, up to 390 Kbps. Again, (LOC3, 10AM) is the most unstable (location, time) pair. However, for the remaining of the locations and times, throughputs remain relatively stable for intervals of several minutes.

Overall, Figures 1(a), 1(b) show that throughputs of both UDP and TCP 30-sec transfers remain relatively stable within 1 hour, and most importantly, within 5 min intervals, i.e., for consecutive measurements. Hence, we will use them as the ground truth in Section 5. The goal is to evaluate how accurately the A_t metric estimated by WBest can approximate these values, when the tool is run directly after the 30 sec transfers.

Figure 2 shows a scatter plot of TCP vs. UDP throughputs taken in all 8 (location, time) experiments, i.e., a total of $8 \times 12 = 96$





Figure 2: Scatter plot of UDP and TCP throughputs in all 8 (location, time) experiments.



(a) Relative error for UDP. (b) Relative error for TCP.

Figure 3: CDFs of relative error for UDP and TCP for different locations and times. UDP and TCP throughputs are calculated as averages over a 30 sec period.

points. We make two observations. First, in almost all cases, TCP throughputs are higher than UDP throughputs. We will discuss more about this rather unexpected finding in Section 5. Second, for a given UDP throughput value, the corresponding TCP throughput (measured only 30 sec later) can be very different, and vice versa. This implies that a single WBest estimate will be able to correctly approximate either UDP or TCP throughput but not both in most cases.


5. RESULTS

Overall Result. We define the relative error for UDP as $Error_{UDP} = \frac{WBest - UDP}{UDP} \times 100\%$ and the relative error for TCP as $Error_{TCP} = \frac{WBest - TCP}{TCP} \times 100\%$, where UDP and TCP denote the average UDP and TCP throughput, respectively, over 30 sec, and WBest is the A_t metric estimated by a single run of the tool.

Figures 3(a), 3(b) plot the CDF of the relative UDP and TCP error for each (location, time). We used only the first 3 runs of WBest, which always span a time interval of less than 30 sec, since the channel may change in a longer period of time. This means that each of the 12 UDP or TCP throughput values for a given (location, time) is compared against 3 A_t values, which gives a total of 36 error values for each CDF. Negative values in Figures 3(a), 3(b) mean that the tool underestimated UDP or TCP throughput, respectively; positive values mean that it overestimated them.

From Figures 3(a), 3(b), we make two observations. First, the fraction of negative values for the UDP relative error ranges from 14% to 36% for different CDFs; the same fraction for the TCP relative error ranges from 47% to 86%. Hence, the A_t estimates obtained by WBest tend to overestimate UDP throughput and underestimate TCP throughput. Note, however, that for TCP some CDFs have heavy right tails.

Second, the achievable throughput metric estimated by WBest



(a) UDP

(b) TCP

Figure 4: Average UDP and TCP throughput for the first T sec ($1sec \leq T \leq 30sec$).

in many cases fails to accurately approximate either UDP or TCP throughput. The fraction of relative error values that lie outside $[-15\%, +15\%]$ (i.e., the fraction of inaccurate approximations) typically ranges from 62% – for (LOC1, 10AM) up to 78% – for (LOC3, 10AM), for UDP, and from 65% – for (LOC3, 9PM) up to 71% – for (LOC3, 10AM), for TCP. The only exceptions are (LOC1, 4AM) for the UDP error and (LOC1, 10AM), (LOC1, 9PM) for the TCP error, with only 31%, 44% and 40% of the error values, respectively, lying outside $[-15\%, +15\%]$.


Overall, we observe that in most cases WBest cannot give an accurate approximation of either UDP or TCP throughput measured over a 30 sec period. One reason for this could be that WBest, which runs in a 3-5 sec interval, cannot capture the dynamics of a much longer period.

UDP vs. TCP. To verify this hypothesis, we plot in Figures 4(a), 4(b) the UDP and TCP throughput, respectively, calculated over the first T sec ($1sec \leq T \leq 30sec$), for each (location, time). Each point on a given curve corresponds to the average of the 12 measurements taken over a 1 hour period at a given (location, time).

From Figures 4(a), 4(b) we observe that the two protocols exhibit very different trends. For each curve, TCP throughput shows an increasing trend at the beginning, and it is stabilized only after about 10 sec due to the slow start phase. On the contrary, for most of the UDP curves, throughput shows a decreasing trend at the beginning, and it is stabilized after 5-10 sec. We are not sure why this behavior is observed. One reason could be due to traffic policies at the BS, that restrict high volume UDP traffic or give priority to TCP traffic. For each (location, time), UDP throughput in the first 5 sec is much higher than the corresponding TCP throughput. But after the first 10 sec, UDP throughput is always lower than the corresponding TCP throughput by 100-200 Kbps.

We now plot again in Figures 5(a), 5(b) the relative error of WBest with respect to UDP and TCP, as in Figures 3(a), 3(b), but this time UDP and TCP throughput (i.e., the ground truth) are calculated as the overage over only the first 3 sec. The goal is to see if WBest can approximate UDP or TCP throughput in this small timescale, similar to the duration required for one run of the tool.

From Figure 5(a) we observe that the accuracy of WBest for UDP is generally not improved when we consider throughput of the first 3 sec only as ground truth. Compared to Figure 3(a) for the 30 sec ground truth, we observe the fraction of error values that lie outside $[-15\%, +15\%]$ has significantly decreased only for (LOC1, 9PM), and slightly decreased for (LOC2, 9PM) and (LOC3, 10AM), but it has increased for the remaining 5 curves. We also observe that 7 out of 8 curves have been shifted towards left compared to Figure 3(a); this implies that the tool underestimates the throughput of short UDP transfers more often than for longer transfers.



(a) Relative error for UDP (b) Relative error for TCP

Figure 5: CDFs of relative error when throughput only over the first 3 sec is used as ground truth.

On the other hand, when we compare Figures 5(b) and 3(b) for TCP, we observe that accuracy drops significantly for all 8 (location, time) experiments when the ground truth is based on the first 3 sec, instead of the total 30 sec. We also observe a significant shift to the right for all 8 curves, compared to Figure 3(b), showing that the tool highly overestimates throughput for short TCP transfers. This was expected, since in Figure 4(b) we saw that TCP throughput is always much lower in the first 3 sec.


Overall, WBest fails to accurately approximate throughput of short transfers.³ For TCP, this is rather expected, since TCP throughputs for short transfers are much lower than for longer transfers. However, for UDP, the reason is not clear. WBest calculates A_t sending a train of 30 UDP packets at a rate equal to the estimated capacity of the 1x EVDO downlink. Since we use the same packet size for the tool and the UDP transfer, and one run of the tool (C_e estimation plus A_t estimation) takes between 3 and 5 sec, one would expect the A_t estimates to approximate at least throughputs of short UDP transfers.

Short-scale dynamics. To understand why this is not the case, we look at the short-scale dynamics of UDP and TCP, i.e., how throughput changes in each second; we also look at the results of consecutive runs of WBest. We show in Figure 6 one example for (LOC2, 10AM). Figure 6(a) plots the per second UDP throughput of 3 consecutive UDP transfers, initiated at 10:20AM, 10:25AM, 10:30AM. Figure 6(b) plots the per second TCP throughput for the same 3 measurements (each TCP transfer followed the corresponding UDP transfer). Finally, Figure 6(c) plots the A_t estimates of 15 runs of WBest following each TCP transfer.


In Figures 6(a), 6(b), we observe that both UDP and TCP per-second throughputs exhibit very large variations, and this is true even when the average throughput (over 30 sec) between 2 consecutive measurements does not change significantly. For example, the average UDP throughput is 260 Kbps for the 4-th measurement and 259 Kbps for the 5-th measurement; however, the per second throughputs vary from 115-503 Kbps, and 8-541 Kbps, respectively.

From Figure 6(c) we observe that the A_t estimates obtained by 15 WBest runs, following the UDP and TCP transfers, also exhibit very large variations covering a range of 159-572 Kbps for the 4-th measurement, 178-743 Kbps for the 5-th measurement, and 103-332 Kbps for the 6-th measurement. This result is very different from the result in [4], in which individual packet pair measurements gave consistently the same estimate of capacity (520 Kbps) over a 24-hour period. Remember that in our previous results we

³We also looked at the relative error when throughputs are calculated over the first 1, 5, and 8 sec and the results were similar to the 3 sec case.



(a) Per second UDP throughput (b) Per second TCP throughput.



(c) Consecutive WBest A_t estimates.

Figure 6: Short-scale dynamics for 3 consecutive measurements at (LOC2, 10AM).

only used the first 3 WBest estimates in each case, which span an interval of about 30 sec. Notice in Figure 6(c) that even the first 3 estimates in each measurement vary significantly with each other.

The large throughput variations, in a very short timescale (1 sec) cause WBest estimates obtained a few seconds apart to also have similar variations, thus making it infeasible for the tool to predict either TCP or UDP throughput even in short timescales. As an illustrative example, the average UDP throughput over the first 3 sec for the 6-th measurement at (LOC2, 10AM) was equal to 138 Kbps. The first WBest run gave an estimate of 332 Kbps (Figure 6(c)), corresponding to one of the largest positive errors in the corresponding CDF of Figure 5(a), of about 140%. The second run gave a negative error of 25%. An the third run gave an estimate of about 138 Kbps, i.e., an error less than 0.05%.

6. CONCLUSIONS

In this paper, we conducted the first study of the feasibility of fast bandwidth estimation in 1x EVDO cellular networks. We examined how well a state-of-the-art probe-based tool developed for WLANs can approximate the UDP or TCP throughput of a commercial 1x EVDO Revision A downlink. Our results show that, although average UDP and TCP throughputs remain relatively stable in many cases for several minutes, they exhibit very large fluctuations in smaller timescales (in the order of 1 sec). Due to these large short-scale fluctuations, we found that bandwidth estimation is not feasible in an 1x EVDO downlink. More specifically, WBest failed to correctly approximate UDP or TCP throughputs of both short (1-8 sec) and longer (30 sec) transfers.

In future work, we plan to study the feasibility of bandwidth estimation on the uplink direction. The uplink direction violates one of the basic assumptions of the tool we used in this paper, i.e., that the bottleneck link is the last one of the network path. In addition, in contrast to the TDMA-based downlink, the 1x EVDO uplink uses CDMA, allowing more than one users to transmit simultaneously,

and posing even greater challenges to the problem of bandwidth estimation.

7. REFERENCES

- [1] N. Bhushan, C. Lott, P. Black, R. Attar, Y.-C. Jou, M. Fan, D. Gosh, and J. Au. Cdma2000 1xEV-DO revision A: A physical layer and MAC layer overview. *IEEE Communications Magazine*, 44(2), February 2006.
- [2] J. M. Cano-Garcia, E. Gonzalez-Parada, and E. Casilari. Experimental analysis and characterization of packet delay in UMTS networks. In *Proc. of LNCS NEW2AN*, 2006.
- [3] 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 *Proc. of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2006.
- [4] M. Claypool, R. Kinicki, W. Lee, M. Li, and G. Ratner. Characterization by measurement of a CDMA 1xEVDO network. In *Proc. of ACM Wireless Internet Conference (WICON)*, 2006.
- [5] R. Cosma, A. Cabellos-Aparicio, J. Domenech-Benloch, J. Gimenez-Guzman, J. Martinez-Bauset, M. Cristian, A. Fuentetaja, A. Lopez, J. Domingo-Pascual, and J. Quemada. Measurement-based analysis of the performance of several wireless technologies. In *Proc. of IEEE Workshop on Local and Metropolitan Area Networks (LANMAN)*, 2008.
- [6] C. Dovrolis, P. Ramanathan, and D. Moore. Packet-dispersion techniques and a capacity-estimation methodology. *IEEE/ACM Transactions on Networking (TON)*, 12(6):963–977, 2004.
- [7] A. Downey. Using pathchar to estimate internet link characteristics. In *Proc. of ACM SIGCOMM*, 1999.
- [8] N. Hu and P. Steenkiste. Evaluation and characterization of available bandwidth probing techniques. *IEEE Journal on Selected Areas in Communications*, 21(6), 2003.
- [9] M. Jain and C. Dovrolis. End-to-end available bandwidth: measurement methodology, dynamics, and relation with TCP throughput. In *Proc. of ACM SIGCOMM*, 2002.
- [10] A. Johnsson, B. Melander, and M. Bjorkman. Bandwidth measurement in wireless networks. In *Proc. of Mediterranean Ad Hoc Networking Workshop*, 2005.
- [11] R. Joyce, B. Graves, T. Griparis, I. Osborne, and T. Lee. Case study: The capacity of a WCDMA network – Orange UK. In *Proc. of IEEE International Conference on 3G Mobile Communication Technologies*, 2004.
- [12] M. Kohlwe, J. Riihijarvi, and P. Mahonen. Measurements of TCP performance over UMTS networks in near-ideal conditions. In *Proc. of IEEE Vehicular Technology Conference (VTC)*, 2005.
- [13] K. Lai and M. Baker. Measuring link bandwidths using a deterministic model of packet delay. In *Proc. of ACM SIGCOMM*, 2000.
- [14] K. Lakshminarayanan, V. N. Padmanabhan, and J. Padhye. Bandwidth estimation in broadband access networks. In *Proc. of ACM Internet Measurement Conference (IMC)*, 2004.
- [15] Y. Lee. Measured TCP performance in CDMA 1xEV-DO network. In *Proc. of Passive and Active Measurement Conference (PAM)*, 2006.
- [16] M. Li, M. Claypool, and R. Kinicki. Packet dispersion in IEEE 802.11 wireless networks. In *Proc. of IEEE International Workshop on Performance and Management of Wireless and Mobile Networks (P2MNet)*, 2006.
- [17] M. Li, M. Claypool, and R. Kinicki. Wbest: a bandwidth estimation tool for IEEE 802.11 wireless networks. In *Proc. of IEEE LCN*, 2008.
- [18] X. Liu, A. Sridharan, S. Machiraju, M. Seshadri, and H. Zang. Experiences in a 3G network: interplay between the wireless channel and applications. In *Proc. of ACM MOBICOM*, 2008.
- [19] B. A. Mah. pchar: A tool for measuring internet path characteristics, 1999.
- [20] B. Melander, M. Bjorkman, and P. Gunningberg. A new end-to-end probing and analysis method for estimating bandwidth bottlenecks. In *Proc. of IEEE Global Telecommunications Conference (GLOBECOM)*, 2000.
- [21] H.-M. Nam, K.-S. Shin, J.-Y. Jeong, S.-J. Bae, and S.-J. Ko. Probing based wireless channel bandwidth estimation and its application to video streaming over wireless networks. In *Proc. of IEEE International Symposium on Consumer Electronics (ISCE)*, 2008.
- [22] K. Pentikousis, M. Palola, M. Jurvansuu, and P. Perala. Active goodput measurements from a public 3G/UMTS network. *IEEE Communications Letters*, 9(9), 2005.
- [23] P. Reichl, M. Umlauf, J. Fabini, R. Lauster, and G. Pospischil. Project WISQY: A measurement-based end-to-end application-level performance comparison of 2.5G and 3G networks. In *Proc. of IEEE Wireless Telecommunications Symposium (WTS)*, 2005.
- [24] V. Ribeiro, M. Coates, R. Riedi, S. Sarvotham, B. Hendricks, and R. Baraniuk. Multifractal cross-traffic estimation. In *Proc. of ITC*, 2000.
- [25] J. Strauss, D. Katabi, and F. Kaashoek. A measurement study of available bandwidth estimation tools. In *Proc. of ACM Internet Measurement Conference (IMC)*, 2003.
- [26] W. L. Tan, F. Lam, and W. C. Lau. An empirical study on the capacity and performance of 3G networks. *IEEE Transactions on Mobile Computing*, 7(6):737–750, 2008.