

An Empirical Study of Bandwidth Predictability in Mobile Computing

Jun Yao, Salil S. Kanhere, Mahbub Hassan
School of Computer Science and Engineering
University of New South Wales
Sydney, NSW 2052, Australia
{jyao, salilk, mahbub}@cse.unsw.edu.au

ABSTRACT

While bandwidth predictability has been well studied in static environments, it remains largely unexplored in the context of mobile computing. To gain a deeper understanding of this important issue in the mobile environment, we conducted an eight-month measurement study consisting of 71 repeated trips along a 23Km route in Sydney under typical driving conditions. To account for the network diversity, we measure bandwidth from two independent cellular providers implementing the popular High-Speed Downlink Packet Access (HSDPA) technology in two different peak access rates (1.8 and 3.6 Mbps). Interestingly, we observe no significant correlation between the bandwidth signals at different points in time within a given trip. This observation eventually leads to the revelation that the popular time series models, e.g. the Autoregressive and Moving Average, typically used to predict network traffic in static environments are not as effective in capturing the regularity in mobile bandwidth. Although the bandwidth signal in a given trip appears as a random white noise, we are able to detect the existence of patterns by analyzing the distribution of the bandwidth observed during the *repeated* trips. We quantify the bandwidth predictability reflected by these patterns using tools from information theory, *entropy* in particular. The entropy analysis reveals that the bandwidth uncertainty may reduce by as much as 46% when observations from past trips are accounted for. We further demonstrate that the bandwidth in mobile computing appears more predictable when location is used as a context. All these observations are consistent across multiple independent providers offering different data transfer rates using possibly different networking hardware.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication

General Terms

Experimentation, Measurement, Performance

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.

WinTech'08, September 19, 2008, San Francisco, California, USA.

Copyright 2008 ACM 978-1-60558-187-3/08/09 ...\$5.00.

Keywords

Measurement, Bandwidth Predictability, Mobile Computing, Entropy, Information Theory, Time Series

1. INTRODUCTION

High-speed networking is no longer restricted to residential and enterprise users. Thanks to the developments in the wide area wireless networking technology, e.g. the 3G HSDPA [1], it is now possible to experience high data transfer rates inside a moving vehicle. The capacity improvement in mobile bandwidth brings the vision of pervasive computing another step closer to reality.

While the capacity improvement of wireless bandwidth is not in question, the stability or the consistency of the bandwidth is. Instead of being a deterministic entity, the interference and other scheduling issues make the wireless bandwidth behave more like a random variable that fluctuates over time and space. Sustaining the quality of service in wide area mobile computing environment therefore remains a challenging problem.

In fact, even in wired networks, bandwidth is known to fluctuate, albeit at a lesser degree, due to load fluctuations [2–4]. Consequently, bandwidth predictability has sparked significant interest in many domains, including adaptive multimedia, admission control, congestion control, and network multihoming. However, studies on network traffic predictability have hitherto focused only on static environments where the users do not move. The issue of bandwidth predictability in mobile computing, especially in the context of high-speed (vehicular) mobility, remains largely unexplored.

To gain a deeper understanding of bandwidth predictability in the mobile environment, we conducted an eight-month measurement study consisting of 71 repeated trips along a 23Km route in Sydney under typical driving conditions. The repetition of the trips is motivated by the observation that in many practical cases, we repeat the same trip over and over again. For example, one may use the same route to drive to work, or drop kids to school. The repetition is even more regular when public transport is concerned. To account for the network diversity, we measure bandwidth from two independent cellular providers. While both providers offer the same popular HSDPA mobile data service, they implement different access rates or peak bandwidths. Besides, being independent, these two providers are likely to implement their own proprietary solutions for managing their network.

From 71 vehicular trips, we collected 71 independent bandwidth traces for each provider. Initially we studied each of these traces in isolation. Interestingly, we found no evidence of correlation between the bandwidth signals at different points in time within the trace. This observation eventually

led to the revelation that the popular *time series* models, e.g., the Autoregressive and Moving Average, typically used to predict network traffic in static environments, may not be as effective in capturing the regularity in mobile bandwidth.

Failing to detect predictability with time series analysis, we explored other approaches. In particular, we investigated if information theoretic measures such as *entropy* would be useful in identifying bandwidth predictability. It is well known that, if the past behaviour of a random process is known, then entropy can help discover intrinsic patterns in the observations, which may help in predicting its future behaviour. To confirm if the above observation holds true in the context of mobile bandwidth, instead of analyzing the traces in isolation, we decided to treat the bandwidth observations from any of these 71 traces as realizations of independent *random observations*. The result was as expected. Although the bandwidth signal within a trace behaved like a random white noise, striking patterns started to emerge when the distribution of all random observations were analyzed. Our analysis confirmed that the bandwidth uncertainty is reduced drastically, by as much as 46% for one provider, when observations from past trips are accounted for. We were able to further demonstrate that, the bandwidth in mobile computing appears more predictable if location is used as a context. All these observations are consistent across multiple providers, and under different location and bandwidth quantization.

The rest of the paper is organized as follows. We review the related work in Section 2. The measurement testbed and the data collection methodology are described in Section 3. The inapplicability of time series tools in predicting mobile bandwidth is explained in Section 4. We present our entropy-based analysis in Section 5. The paper is concluded in Section 6.

2. RELATED WORKS

Prediction of background network traffic, which would incidentally work as a prediction of the available link bandwidth, has been the subject of intense research since the early 1990s [2–4]. In these studies, researchers collected traces of continuous traffic data from live networks for a specific period of time, and subsequently applied classical time series models, to these traces to predict the short to long term traffic fluctuations (bandwidth variations) in the network. While most of the earlier work were based on wired networking, e.g., the Ethernet, some recent initiatives involve bandwidth prediction in the popular 802.11-based wireless networks [5,6]. Wired or wireless, these studies were carried out for a fixed point of attachment to the network and hence are not quite applicable to mobile communication.

In more recent years, we have witnessed a surge of interest in measuring and characterizing the wireless bandwidth in a wide area mobile environment covered by the fast growing 3G cellular technology. In [7], the authors measured the performance of HSDPA and found that the service bit rate had a large spread during a 1-hour driving along a route that goes through a variety of radio conditions. A detailed cross-layer measurement study to evaluate TCP performance over CDMA2000 networks, and to identify factors affecting TCP performance, was reported in [8]. In that study, the authors observed high variability in TCP throughput based on the time and day of the experiments. The variability of TCP throughput was attributed to the adaptive nature of the wireless schedulers at the base stations. The authors of [9] conducted an empirical study on the variability of UMTS

data transfer capacity in the presence of voice and video calls. They have shown that available capacity can vary not only across operators, but also between sites of the same operators. Similar measurement studies were carried out for other 3G flavors, e.g., CDMA 1x EV-DO networks [10]. While these studies all confirm that 3G wireless bandwidth is highly variable, they did not investigate the predictability of the bandwidth.

Similar to our vehicular mobile measurements, there are other measurement studies reported in the literature that investigated Internet connectivity performance in driving conditions. Rodriguez et al. [11] have shown that by exploiting operator and technology diversity, data download performance can be significantly improved through smart scheduling in an onboard mobile router that travels with the vehicle. Through an intensive measurement study consisting of hundreds of hours of driving in metropolitan city areas, Bychkovsky et al. [12] have demonstrated that the unplanned installations of WiFi access points can actually provide tens of seconds of end-to-end Internet connectivity for motorists. Similar measurements conducted on German highways have provided detailed analysis of the WiFi connectivity durations and the associated TCP performance that can be achieved at high speed cruising [13]. While these studies provide valuable insight into the expected Internet performance in wide area mobile environment, they do not focus on bandwidth predictability.

In our analysis, we use information theoretic measures, entropy to be precise, to quantify the predictability in the observed bandwidth. Entropy is a well known measure with its roots in thermodynamics and signal processing [14]. Recently information theory has been used for analyzing location predictability in mobile computing [15,16]. To the best of our knowledge, use of entropy to quantify bandwidth predictability has not been reported yet.

3. MEASUREMENT METHODOLOGY

In this section we provide an overview of the measurement setup including the software and hardware components and also briefly describe the field trips.

The measurement architecture as depicted in Figure 1(a) consists of a server deployed at the University of New South Wales (UNSW) and a client, which is installed in a vehicle. The client illustrated in Figure 1(b) comprises of two Soekris Net4521 boards (133Mhz processor and 16MB memory), which are interconnected via 10Mbps Ethernet and configured to operate in the master and slave mode. Each board is equipped with one PCMCIA cellular modem. The boards are enclosed in a protective casing and housed in the boot of the OCEAN (On-board Communication Entertainment And information, a vehicular Internet project at UNSW) vehicle. To enhance the wireless signal reception, the cellular modems are connected to external antennas, which are mounted on the car windshield. The vehicle location is recorded by a GPS sensor, installed on the top of the vehicle, which is connected to one of the boards.

To account for the network diversity, we select two independent 3G providers. Both providers have deployed HSDPA services, but with different specifications. Provider A's advertised service rate ranges from 550Kbps to 1.5Mbps with the peak rate of 3.6Mbps; Provider B has a peak rate of 1.8Mbps and advertises a typical rate of 600Kbps.

In this paper, we focus on measuring the downlink mobile bandwidth, for which, we employ the popular packet-train dispersion technique [17]. This technique uses the time

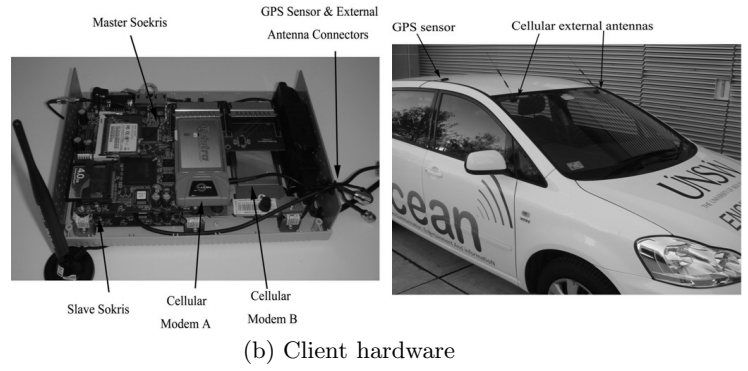
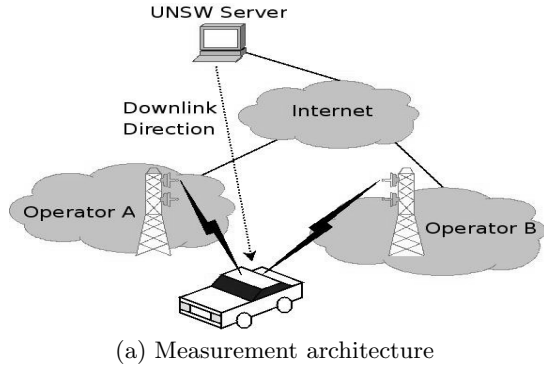


Figure 1: Measurement setup

dispersion between a pair of packets or a train of back-to-back probing packets for estimating bandwidth of the bottleneck link along the end-to-end path. It is quite likely (see Figure 1(a)) that the cellular last-hop is the bottleneck link along the end-to-end path from the server to the client. Hence, the downlink bandwidth of the cellular link can be estimated using packet-train probes sent from the server to the client. Several packet-train based bandwidth estimators [17, 18] exist in literature. It is known that all these tools take a long time to converge to an estimate, which is not an issue for static networks. However, in a high mobility scenario such as ours, quick convergence is essential. Further, the probe traffic generated by these tools is quite significant (cellular bandwidth is still relatively expensive). Hence, we have designed a simplified packet-train client-server program, which converges quickly and also minimizes the amount of probe traffic. In order to evaluate the accuracy of our packet-train program, we estimated the actual capacity by completely saturating the downlink and compared the results with those from the packet-train. Each test lasted for 15 minutes and was conducted at fixed locations. Note that, there are two configurable parameters in a packet train: the number of packets in the train and the packet size [19]. We tested various combinations of these two parameters and observed that a train of 10 packets of size 1000 bytes was the most accurate. Figure 2 plots the CDF of the downlink capacity for provider B as measured by the 10 packet-1000 byte train and compares it with the CDF of the saturation test (results from one experiment). The graph shows that for the most part, the two distributions conform with each other. However, there is some deviation in the lower range of the bandwidth (0-250Kbps). Nonetheless, given that this particular range only accounts for less than 20% of the distribution, the impact on our subsequent analysis is quite minimal. In our tests, we executed an independent packet-train program for each of the cellular links for the entire duration of the trip. The frequency of generating the packet-train was adjusted such that one train was sent out for approximately every 200m segment of motion, resulting in one bandwidth sample for every 200m.

We selected a 23Km route that traverses through the Sydney metropolitan area. The starting point of the route was UNSW in the eastern suburbs and the final destination was Macquarie University (MQ) located in the north-western suburbs. The route passes through three major tunnels (including one underwater tunnel that extends across Sydney harbor), three freeways and some residential and business precincts. Figure 3 depicts the trajectory of the route. In the eight month period from Aug 7, 2007 to Apr 6, 2008, we

conducted 71 repeated trips along this route in the same direction (i.e. MQ bound). The trips were conducted during normal commuting hours including morning/evening rush hours and off-peak periods. Consequently, the trip completion time varied from 25 minutes to 50 minutes. Collectively, our measurements clocked more than 60 driving hours and covered over 1600Km. Note that, the chosen route is a fairly typical representation of a daily commute.

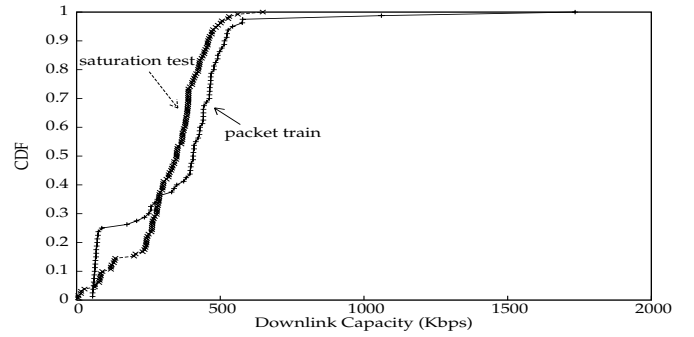


Figure 2: Estimation accuracy of the packet train




Figure 3: Route trajectory

4. TIME SERIES ANALYSIS

A time series represents an ordered sequence of values of a variable, which are measured at equally spaced time intervals. Time series analysis comprises methods that attempt

to understand the underlying forces and structure that produce the observed data (i.e. to identify trends and seasonal variations). The primary goal is to fit a model to the data, which can then be used to forecast future data points. Recent research [4] [20] has demonstrated using empirical measurements that Internet traffic exhibits strong stationary properties. The authors have captured extended network traffic traces (over few hours long), with each trace being represented as a time series. Time series models applied to the individual traces have shown that the future characteristics of the traffic (e.g., per-user traffic) can be predicted accurately. However, this analysis has exclusively focused on network traffic measurements conducted at static locations (i.e. a fixed point of attachment to the Internet). In this paper, we seek to investigate if mobile bandwidth has any inherent regularity, and is hence predictable. A logical first step is to employ time series analysis. Note that, the band-

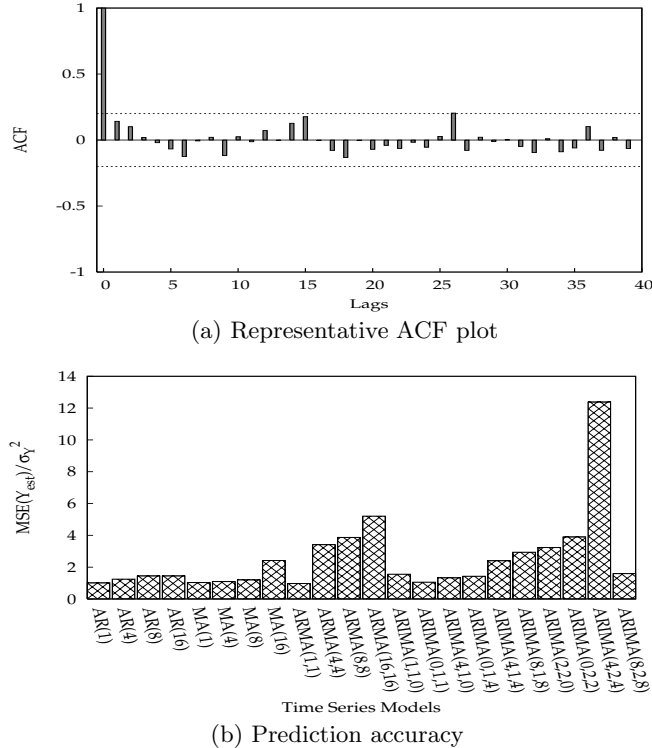


Figure 4: Time series analysis

width samples from each trace in our measurements can be represented as a time series, since the data points recorded in a given trip are measured at successive times. However, measurements across different trips cannot be combined, since there is a time delay between successive trips. This implies that the traces from each trip represent different time series. Since, the analysis requires a discrete time signal, the usual practice is to bin the samples into non-overlapping bins of a fixed size and then average the samples grouped in the same bin to obtain the representative bin sample. We have used a range of values for the bin intervals (10 to 40 seconds) and observed similar results. We only present results for an interval of 20 seconds.

To identify any repeating patterns that may exist in the signal, we first compute the Autocorrelation Function (ACF - i.e., the correlation between the signal at different points in time) of the signal. An ACF plot that has a slow decaying shape indicates the existence of patterns in the underlying signal [4]. Figure 4(a) presents an ACF plot for the trace

collected from one trip for provider A. Observe that the ACF coefficients for lags greater than 0 are not significant. This implies that there is no significant correlation between the bandwidth signals at different points in time. Hence, the best predictor is probably the mean of the signal. We observe similar behavior with the traces from different trips for both providers (results are not presented for reasons of brevity). The initial observations from the ACF analysis seem to suggest that the signal is actually close to random white noise and hence, time series model might not be suitable for analyzing individual traces.

To confirm this, we evaluate the accuracy of a number of linear time series models, AR, MA, ARMA and ARIMA, to be precise, in predicting the future data points (i.e. bandwidth) for the traces from each trip. The usual practice in this analysis [3] is to divide the time series into two equal halves. The first half of the data is used to train the time series model, The trained model is then used to predict the values of the second half. The error for each data point can be computed as the difference between the predicted and actual values. The metric used to evaluate the accuracy of the predictions is the ratio of the prediction mean square error to the variance of the actual values, i.e. $MSE(\hat{Y})/Var(Y)$. The smaller the ratio (i.e. < 1), the better is the predictability of the model. Note that, if we just use the mean value of the first half of the signal as a predictor for future values, the prediction accuracy is close to one. Thus, for time series models to be effective, the ratio should be lower than 1. We follow the aforementioned steps and compute the prediction accuracy achieved by a number of time-series models. Figure 4(b) presents the results for the bandwidth samples collected from the provider A network during one trip. The number in the parenthesis indicates the order of the model. The plot shows that the best performing models are the simple AR (autoregressive) and MA (moving average) with order *one*, which achieve similar performance (i.e., $MSE(\hat{Y})/Var(Y) \approx 1$) as using the mean value of the training set as a predictor. Observe that higher order and non-stationary models perform worse than these simplistic models. The general trend of the results is similar for the traces from different trips and providers.

The results presented in this section suggest that time series models when applied to individual traces, are unable to reveal the existence of any regular patterns in the downlink mobile bandwidth. One possible reason for this could be the relatively short duration of the trips (between 25 to 50 minutes). This time period is probably not long enough for time series models to gather sufficient knowledge required for identifying any underlying patterns in the bandwidth variations. However, note that, typical commuting times in urban environments are not significantly longer than those encountered in our field trips. It should also be noted that traces from the repeated trips cannot be combined together, since there is a time lag between subsequent trips.

5. INFORMATION-THEORETIC ANALYSIS

Instead of analyzing each trace in isolation, in this section, we treat the bandwidth observations from the repeated trips as realizations of independent random observations. Further, we investigate if information-theoretic measures, entropy in particular, can provide insights into the predictability of mobile bandwidth.

The outcome of a stochastic process usually appears to be highly unpredictable due to the associated randomness.

However, by understanding the past behavior of the process, one may be able to discover intrinsic regularities in the underlying forces that govern the process, which may help in predicting its future behavior. Information entropy [21] is a well-known metric that measures the level of uncertainty associated with a random process. It quantifies the information contained in a message, usually in *bits/symbol*. The entropy of a discrete random variable X , is defined as,

$$H(X) = \sum_{x \in X} p(x) \log_2 p(x) \quad (1)$$

where $p(x)$ is the probability mass function, $0 \leq p(x) \leq 1$.

Since entropy informs us about the uncertainty associated with a process, it can implicitly provide information about its predictability. Note that, when the entropy is 0, the outcome of the process is completely deterministic and hence completely predictable. On the other hand, when the process is completely random, $p(x)$ takes on a uniform distribution, and the corresponding upper bound on the entropy can be calculated using Equation 1. In general, the lower the entropy, the lower is the information uncertainty associated with the process, which in turn implies that the future outcomes can be predicted with greater ease.

The measurements from each trip in our study provide us with a trace of discrete samples of the cellular downlink bandwidth at different locations along the route. Each repeated trip generates a distinct trace. Instead of analyzing each trace in isolation (as in Section 4), we can treat the samples from the 71 traces as realizations of independent random observations of the cellular bandwidth (which is a random variable). In this section, we investigate if the past knowledge can be leveraged to identify any particular trends that may exist and thus, lead to better predictability. In the first part of our analysis, we assimilate all samples to form a collective sample set. We then use entropy to quantify the associated uncertainty. Next, we investigate if the bandwidth predictability can be improved by using location as a context. For this, we evaluate the *location entropy* i.e. the entropy of the bandwidth samples at each location along the route and observe if the associated uncertainty reduces.

5.1 Entropy Analysis of the Collated Samples

From our traces we observed that the downlink bandwidth varies from 0 to 3Mbps and 0 to 600Kbps for providers A and B, respectively. This is line with the typical advertised capacity for both providers. For mathematical tractability, we quantize this continuous range of values to a small set of discrete symbols. Choosing an appropriate number of symbols is important. With too few symbols, each symbol would represent a large bandwidth range. In this case, there would be little use of correctly predicting the symbol since the corresponding range of values is quite large. On the contrary, with too many symbols it may be hard to capture any patterns that may exist in the underlying process. To strike a balance, we use 7 symbols for Provider A (i.e. a quantization interval of 500Kbps for the higher bandwidth provider) and 6 symbols for Provider B (i.e., a quantization interval of 100Kbps for the lower bandwidth provider) as shown in Table 1.

In our measurements we have collected one bandwidth sample for approximately every 200m section of the route. Thus, the data samples collected over one trip represent a space-ordered sequence (i.e. ordered by locations along the route). On occasions the probes used in estimating the bandwidth are lost, leading to a few missing samples. Fur-

ther, it is desirable to bin consecutive samples together and use the average value as a representative sample as opposed to using their instantaneous values. We have used different sizes of location segments ranging from 500m to 1500m as the bin size. Our results are consistent for different bin sizes. Due to space limitations, we only present results for a segment size of 500m for the most part. The route used in our measurements is approximately 23.5Km long. With a resolution of 500m each trip provides us with 47 samples. Since we have collected data for 71 trips, the total sample set is $71 \times 47 = 3337$.

Table 1: Symbol Definition


Symbol	Bandwidth Range	
	Provider A	Provider B
A	$\leq 0.5Mbps$	$\leq 0.1Mbps$
B	$> 0.5Mbps \ \& \ \leq 1Mbps$	$> 0.1Mbps \ \& \ \leq 0.2Mbps$
C	$> 1Mbps \ \& \ \leq 1.5Mbps$	$> 0.2Mbps \ \& \ \leq 0.3Mbps$
D	$> 1.5Mbps \ \& \ \leq 2Mbps$	$> 0.3Mbps \ \& \ \leq 0.4Mbps$
E	$> 2Mbps \ \& \ \leq 2.5Mbps$	$> 0.4Mbps \ \& \ \leq 0.5Mbps$
F	$> 2.5Mbps \ \& \ \leq 3Mbps$	$> 0.5Mbps$
G	$> 3Mbps$	-

We first compute the entropy of the bandwidth assuming that we do not have any information of the past behavior. In this case all symbols in Table 1 are equiprobable. The resulting entropy is $\log_2 7 = 2.81$ and $\log_2 6 = 2.58$ for providers A and B, respectively. Note that, this will be the maximum possible value (upper bound) for the entropy. Next, we evaluate if knowledge of the past behavior as represented by our measurement sample set can reduce the uncertainty. By counting the frequency of occurrence of each symbol in our sample set, we first estimate the PDF of the bandwidth for both providers. Then, using Equation 1, the entropy for providers A and B are calculated to be 1.52 and 1.92, respectively. The minor difference in the entropy of the two providers is due to the different quantization intervals used. Observe that the bandwidth entropy for both providers has reduced significantly from their corresponding upper bounds (46% and 26% drop for providers A and B, respectively) by utilizing information about the past behavior. This also implies that the bandwidth predictability improves considerably. To further elaborate this, we plot the estimated PDFs for the downlink bandwidth of both providers in Figure 5(a) and Figure 5(b). Observe that there are two dominant symbols for each provider, i.e. *C* and *D* for provider A, and *D* and *E* for provider B. Further, the most dominant symbol (*D* for provider A and *E* for provider B) accounts for about 50% of the entire distribution. This implies that even a simple prediction algorithm such as *Best Guess* [21], which always picks the dominant state can achieve reasonable accuracy.


5.2 Impact of Location on Entropy

It is well-known that the cellular bandwidth exhibits frequent variations. One common reason for these fluctuations can be attributed to location. For example, a sharp drop in bandwidth is almost always experienced when crossing an underground tunnel, or when obscured by a high-rise building [9]. In this section, we first confirm that our traces indeed exhibit a correlation between the mobile bandwidth and location. Next, we investigate how this dependency can be leveraged to improve the predictability.

As in Section 5.1, we use a granularity of 500m for each location segment, which gives us a total of 47 distinct locations along the route. Recall that, we have collected bandwidth samples for 71 repeated trips of the same route. This



(a) Provider A, $H(X)=1.5$



(b) Provider B, $H(X)=1.8$

Figure 5: PDF of downlink bandwidth

means that we have 71 bandwidth samples for each of the 47 locations. By counting the frequency of occurrence of each samples, we estimate the PDF of the bandwidth distribution at each segment. Figure 6(a) and Figure 6(b) compares the PDFs of two distinct locations for both providers. It is quite evident that the bandwidth distributions vary significantly from one location to another. Note that, similar differences exist between several locations. We have only shown one instance for the sake of brevity.

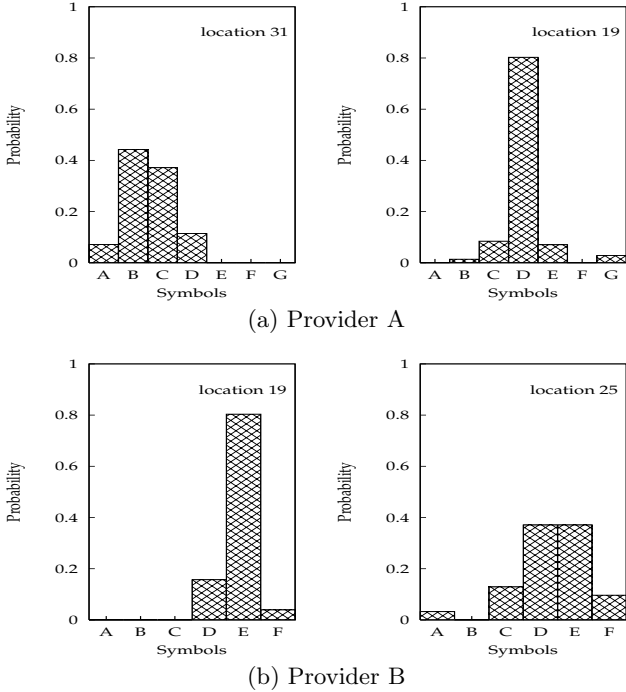


Figure 6: Comparing PDFs for different locations

To quantify this difference, we use the well known metric, L_1 distance [22]. Given the PDFs of two discrete random variables X and Y , the L_1 distance is computed as follows,

$$L_1(X, Y) = \sum_i |p(x_i) - p(y_i)| \quad (2)$$

where $p(x)$ and $p(y)$ represent the respective probability mass functions. Note that, $0 \leq L_1 \leq 2$. Using Equation 2, we compute the L_1 distance between the bandwidth distributions for all possible combinations of locations (a total of 1081 combinations). Table 2 provides a summary of the results (i.e. average, minimum and maximum value) for both providers. Observe that the maximum value is as high as 1.6. Even the average value for both providers is non-trivial,

implying that the bandwidth distributions at different locations vary significantly.

Table 2: L_1 distance

Network	Average	Max	Min
A	0.55	1.60	0.03
B	0.62	1.59	0.04

Note that, the variation in bandwidth between two locations would potentially have the most impact on the mobile application when these two locations are adjacent to each other. Figure 7(a) and 7(b) illustrate the L_1 distance between the bandwidth distribution of successive locations for both providers. Since the route comprises of 47 segments, we have a total of 46 data points. Observe that there are several sharp peaks in both the graphs. For example, the L_1 distance between location 28 and 29 for provider B is 1.4. In general, we can conclude that the bandwidth distribution of adjacent segments can vary significantly.


Finally, we investigate if the bandwidth distribution of the collated samples (independent of location as in Section 5.1) is an accurate representation of the bandwidth distribution of the individual segments. We compute the L_1 distance between the bandwidth distribution at each location and that of the collective sample set. Figure 8 illustrates that there is a significant difference between the distribution of the collective sample set and that of each location.

The results from the above analysis using L_1 distance collectively highlight the fact that there is a strong correlation between the distribution of the mobile bandwidth and location. Further, in most cases the bandwidth distribution varies significantly from location to location. The conclusions are consistent across both providers. We now proceed to evaluate if these observations can be used to improve the bandwidth predictability.


Recall that, in Section 5.1, we combined the bandwidth samples from all traces, independent of location, and computed the entropy of the collective sample set. We refer to this as the *collated entropy* in the rest of this discussion. Given the strong influence of location on the bandwidth distribution as observed in our analysis above, we define *location entropy* $H(X|l_i)$ as the entropy of the bandwidth for an individual location segment, l_i , as follows,

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

where X is a discrete random variable representing the bandwidth, l_i is a location segment from the set of segments L along the route and $p(x|l_i)$ is the probability mass function of the bandwidth at segment l_i .




(a) Provider A




(b) Provider B

Figure 7: L_1 distance between adjacent location



(a) Provider A



(b) Provider B

Figure 8: L_1 distance between the PDF of each location and the PDF of the collective samples

We compute the location entropy for all 47 locations of the trip and compare it with the collated entropy. Table 3 presents a comparison of the average entropy for different resolutions of the location segment. One can readily observe that by considering each location independently, the average entropy has reduced. This implies that by maintaining per-location statistics, better predictability can be achieved as opposed to collating all samples together. Further, the results are consistent for both providers and also for different sizes of the location segments. Note that, with a larger segment size, more individual bandwidth samples are averaged, thus reducing the perceived uncertainty.

Table 3: Comparison of average entropy

Network	Resolution	Location Entropy	Collated Entropy
A	500m	1.29	1.52
	750m	1.16	1.37
	1000m	1.07	1.32
	1500m	0.99	1.23
B	500m	1.65	1.92
	750m	1.58	1.80
	1000m	1.61	1.82
	1500m	1.51	1.69

However, the table only compares the averages. We plot the location entropy for all 47 locations in Figure 9. The collated entropy is represented as a straight line. These graphs indicate that a majority of the locations (over 75%) exhibit lower entropy than the collated entropy. This implies that maintaining per-location statistics can improve the chance of correctly predicting the bandwidth. Notice that a handful of locations exhibit higher entropy than the collated entropy. However, this does not mean that predictions made based on the collated samples will necessarily be accurate. Rather, it implies that these locations exhibit greater bandwidth uncertainty and are hence inherently harder to predict.

6. CONCLUSION AND FUTURE WORK

To gain a deeper understanding of bandwidth predictability in a mobile environment, we have conducted an eight-month measurement study consisting of 71 repeated trips along a 23Km route in Sydney under typical driving conditions. The route runs through locations with different radio conditions including terrestrial and underwater tunnels. We have measured bandwidth from two independent cellular providers implementing HSDPA technology in two different peak access rates (1.8 and 3.6 Mbps). We have observed no significant correlation between the bandwidth signals at different points in time within a given trip. We have found that the popular time series models, e.g. Autoregressive and Moving Average, typically used to predict network traffic in static environments are not as effective in predicting the mobile bandwidth. Although the bandwidth signal in a given trip appears as a random white noise, we were able to detect the existence of patterns by analyzing the distribution of the bandwidth observed during *repeated* trips. We quantified the bandwidth predictability reflected by these patterns using information theory, *entropy* in particular. Our entropy analysis revealed that the bandwidth uncertainty reduces drastically when observations from past trips are used to predict bandwidth. We further demonstrated that the mobile bandwidth appears more predictable when location is used as a context. Our observations are consistent across multiple independent providers offering different data rates using possibly different hardware and management tools.

The results presented in this paper are based on measurements collected from 71 repeated trips. We believe that this is a sufficiently large sample set to draw meaningful conclusions. In our on-going work, we are analyzing the relationship between the sample set size and the entropy in greater detail. Further, in this work, we solely focussed on exploring the correlation between mobile bandwidth and location. In

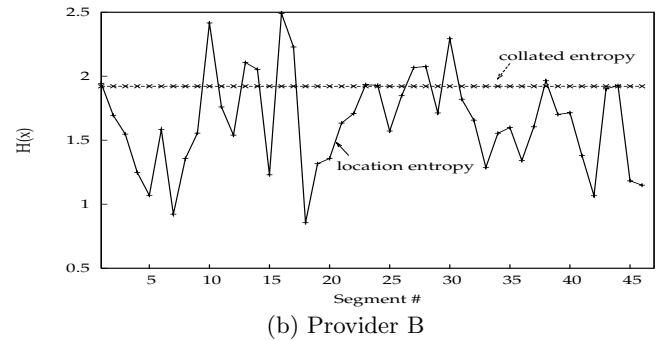
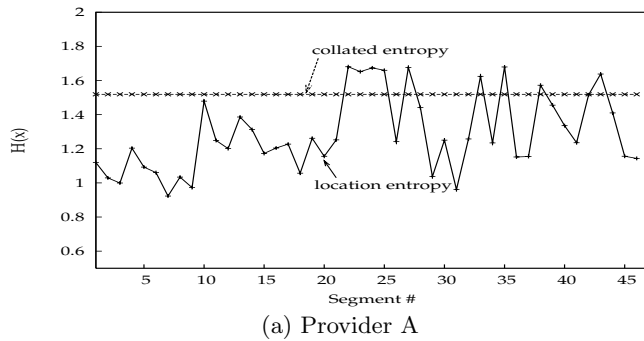


Figure 9: Comparison of collated entropy with location entropy

our future work, we intend to investigate the impact of other contextual information such as speed, time, etc. Since, the focus of this paper was on understanding the predictability of mobile bandwidth, we have purposely not discussed about the applications of these findings. In our future work, we plan to investigate how this knowledge can be used to design intelligent protocols for improving the performance of mobile Internet access systems, with a particular focus on vehicular Internet access.

7. REFERENCES

- [1] "High Speed Downlink Packet Access (HSDPA); Overall description; Stage 2." [Online]. Available: <http://www.3gpp.org/ftp/specs/html-info/25308.htm>
- [2] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the Self-similar Nature of Ethernet Traffic (Extended Version)," *IEEE/ACM Trans. Netw.*, vol. 2, no. 1, pp. 1–15, 1994.
- [3] Y. Qiao, J. Skicewicz, and P. Dinda, "An Empirical Study of the Multiscale Predictability of Network Traffic," in *Proc. of 13th IEEE International Symposium on High Performance Distributed Computing (HPDC-13 '04)*, Honolulu, Hawaii USA, Jun. 2004.
- [4] A. Sang and S. Q. Li, "A Predictability Analysis of Network Traffic," in *Proc. of IEEE Infocom 2000*, Tel Aviv, Israel, Mar. 2000.
- [5] J. Zhang and I. Marsic, "Link Quality and Signal-to-Noise Ratio in 802.11 WLAN with Fading: A Time-Series Analysis," in *Proc. of Vehicular Technology Conference 2006 Fall*, Montreal, Canada, Sep. 2006.
- [6] J. P. Singh, T. Alpcan, P. Agarwal, and V. Sharma, "An Optimal Flow Assignment Framework for Heterogeneous Network Access," in *Proc. IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 2007)*, Helsinki, Finland, Jun. 2007.
- [7] J. Derksen, R. Jansen, M. Maijala, and E. Westerberg, "HSDPA Performance and Evolution," *Ericsson Review*, no. 03, pp. 117–120, 2006.
- [8] K. Mattar, A. Sridharan, H. Zang, I. Matta, and A. Bestavros, "TCP Over CDMA2000 Networks : A Cross-Layer Measurement Study," in *Proc. of PAM 2007*, Louvain-la-neuve, Belgium, Apr. 2007.
- [9] W. L. Tan, F. Lam, and W. C. Lau, "An Empirical Study on 3G Network Capacity and Performance," in *Proc. of IEEE Infocom 2007*, Anchorage, Alaska, USA, May 2007.
- [10] Y. Lee, "Measured TCP Performance in CDMA1x EV-DO Network," in *Proc. of PAM 2006*, Adelaide, Australia, Mar. 2006.
- [11] P. Rodriguez, R. Chakravorty, I. Pratt, and S. Banerjee, "MARS: A Commuter Router Infrastructure for the Mobile Internet," in *Proc. of ACM MobiSys 2004*, Boston, MA, USA, Jun. 2004.
- [12] V. Bychkovsky, B. Hull, A. K. Miu, H. Balakrishnan, and S. Madden, "A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks," in *Proc. of 12th ACM MobiCom*, Los Angeles, CA, USA, Sep. 2006.
- [13] J. Ött and D. Kutscher, "Drive-thru Internet: IEEE 802.11b for Automobile Users," in *Proc. of IEEE Infocom 2004*, Hong Kong, China, Mar. 2004.
- [14] N. Gershenfeld, "Signal Entropy and the Thermodynamics of Computation," *IBM Systems Journal*, vol. 35, no. 3 & 4, pp. 577–587, 1996.
- [15] J.-M. François and G. Leduc, "Entropy-based Knowledge Spreading and Application to Mobility Prediction," in *Proc. of the 2005 ACM Conference on Emerging Network Experiment and Technology*, Toulouse, France, Oct. 2005.
- [16] A. Bhattacharya and S. K. Das, "LeZi-Update: An Information-Theoretic Framework for Personal Mobility Tracking in PCS Networks," in *Proc. of ACM MobiCom 1999*, Seattle, Washington, USA, Aug. 1999.
- [17] C. Dovrolis, P. Ramanathan, and D. Moore, "Packet-Dispersion Techniques and A Capacity-Estimation Methodology," *IEEE/ACM Trans. Netw.*, vol. 12, no. 6, pp. 963–977, 2004.
- [18] R. Kapoor, L.-J. Chen, L. Lao, M. Gerla, and M. Y. Sanadidi, "CapProbe: A Simple and Accurate Capacity Estimation Technique," in *Proc. of ACM SIGCOMM 2004*, Portland, OR, USA, Aug. 2004.
- [19] M. Claypool, R. Kinicki, W. Lee, M. Li, and G. Ratner, "Characterization by Measurement of a CDMA 1x EVDO Network," in *Proc. of the Wireless Internet Conference (WICON) 2006*, Boston, MA, USA, Aug. 2006.
- [20] N. Thompson, G. He, and H. Luo, "Flow Scheduling for End-host Multihoming," in *Proc. of IEEE Infocom 2006*, Barcelona, Catalunya, Spain, Apr. 2006.
- [21] C. E. Shannon, "Prediction and entropy of printed English," *Bell System Technical Journal*, vol. 30, pp. 50–64, 1951.
- [22] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White, "Testing That Distributions Are Close," in *Proc. of the IEEE 41st Annual Symposium on Foundations of Computer Science*, Redondo Beach, CA, USA, Nov. 2000.