Performance Comparison of 3G and Metro-Scale WiFi for Vehicular Network Access

Pralhad Deshpande, Xiaoxiao Hou, and Samir R. Das
Computer Science Department, Stony Brook University
Stony Brook, NY 11794, USA
Email: {pralhad, xhou, samir}@cs.stonybrook.edu

ABSTRACT

We perform a head-to-head comparison of the performance characteristics of a 3G network operated by a nation-wide provider and a metro-scale WiFi network operated by a commercial ISP, from the perspective of vehicular network access. Our experience shows that over a wide geographic region and under vehicular mobility, these networks exhibit very different throughput and coverage characteristics. WiFi has frequent disconnections even in a commercially operated, metro-scale deployment; but when connected, indeed delivers high throughouts even in a mobile scenario. The 3G network offers similar or lower throughputs in general, but provides excellent coverage and less throughput variability. The two network characteristics are often complementary. It is conceivable that these properties can be judiciously exploited for a hybrid network design where 3G data can be offloaded to WiFi for better performance and to reduce 3G network congestion and to lower costs.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communications—*Vehicular Communications*; C.4 [**Performance of Systems**]: Measurement techniques.

General Terms

Performance, Measurement, Experimentation.

Keywords

Vehicular Internet Access, WiFi, 3G.

1. INTRODUCTION

3G cellular data and WiFi are the two wireless broadband access technologies widely available today. The two networks have significant differences. The former uses licensed bands and macrocells with large coverage areas. The base station and associated radio access network setup have significant capital and operational costs [25]. On the other hand, WiFi uses unlicensed spectrum. The access point (AP) coverage is relatively smaller and typically capital and operational costs are lower. Also, in comparison to 3G, most users expect that WiFi access would be free or inexpensive, and would provide a significantly higher bit rate (in the order of 10s of Mbps). However, they also expect that the coverage is available only in the local area, while 3G coverage is often ubiquitous.

Recent experiences with outdoor and vehicular access with WiFi [7, 11] have started to blur some of the differences between 3G and WiFi. Many metro-scale WiFi deployments [20] have also succeeded in making WiFi ubiquitous providing very good coverage in urban spaces. All these now make it viable to access WiFi in outdoors even in a mobile scenario at vehicular speeds. This makes us ask this question: *Can WiFi be used effectively in outdoors and mobile scenarios to reduce the load on the expensive 3G networks?* Demonstrating technical viability of this will certainly draw interest in developing networking protocols and applications that eventually lead to more efficient and cost effective wireless broadband access. We note that using WiFi to reduce the load on 3G networks is not new, and has been promoted by both networking [10, 27, 4] and policy researchers [16, 13]. Also, some carriers are known to incentivize such switching (e.g., TMobile @Home) to reduce load on their cellular networks. However, in spite of the general interest, systematic head-to-head comparisons of commercially operated metro-scale WiFi and 3G networks are still missing (except [4] where mostly open WiFi APs in the wild have been used).

Our paper makes the following contributions.

1. Provide an in-depth evaluation of WiFi access with vehicular mobility using a commercially operated metro-scale WiFi network. This is in contrast to prior studies that only considered open APs in the wild [11, 7, 4] or a limited WiFi deployment [12, 5, 9].

2. Provide a head-to-head comparison of WiFi and 3G access under vehicular mobility. This lays down the salient features of the two access networks and provide pointers as to how future hybrid network access could be designed by exploiting the better properties of the two networks.

The rest of the paper is organized as follows. In Section 2, we describe the related work. In Section 3, we describe the measurement setup including the networks considered, the testbed and the driving scenarios. The analysis of the mea-

IMC'10, November 1–3, 2010, Melbourne, Australia.

surement results are presented in Section 4. We conclude in Section 5.

2. RELATED WORK

2.1 Vehicular WiFi Access

Several experimental studies have explored the potential of using intermittently available WiFi connectivity from moving vehicles for data transfers. In the Drive-thru Internet project [23, 24] controlled experiments are done with a single car driving past a single access point to measure range and connectivity in an intermittent network. More recently, the CarTel project [7] has focused on upload performance while using open APs in the wild. In the ViFi project [5] the link layer performance is improved by exploiting macrodiversity (using multiple APs simultaneously), and opportunistic receptions by nearby APs. By the nature of the work, both CarTel and ViFi focus on upload.

Downloads and intermittent connectivity have been studied in a vehicular environment in a fleet of taxis in the Cabernet project [11] with an improved handoff scheme and a new transport protocol. In our prior work in this space [9], we also have focused on downloads. The emphasis there has been predictive methods for improving handoffs and use of prefetching in the APs for better download performance.

Improving handoffs with vehicular mobility has been an important area of research. Other than our work in [9], optimized handoffs techniques have also been developed and evaluated in the Cabernet project [11] and also in [12]. One or more of these strategies can nicely complement our work.

Applications that require maintaining a session face problems in vehicular WiFi access because of intermittent connectivity. Some papers address this issue by creating a transport layer protocol that maintains sessions transparently to changing IP addresses [11, 24].

Several other papers focus on sundry issues on vehicular WiFi access. For example, link and transport layer problems are examined in detail in [14]. Web applications are examined in presence of intermittent connectivity in [6]. Link layer measurements are reported in [18]. Use of directional antenna for better connectivity has been explored in [21].

2.2 Characterizing 3G Access

There has been less excitement about characterizing 3G access in mobile or stationary environment. This is likely because the 3G network is a closed system and under tight operator control. Thus, much of the work is measurement-based.

The important works in this space include cross-layer measurement study to evaluate TCP performance over 1xRTT networks in [19] and over EVDO networks in [15], capacity evaluation of UMTS networks in [26]. Mobile experiments are done in EVDO networks characterizing cross-layer aspects with TCP in [17]. Bandwidth predictability is evaluated for HSDPA networks in [28]. A measurement system for city-wide measurement of WiFi and EVDO networks was developed in [22].

2.3 Augmenting Mobile 3G using WiFi

A recent paper [4] has addressed the issue of augmenting mobile 3G by offloading data on WiFi whenever possible. However they mostly use open APs in the wild (more than 70% of times) and hence the median WiFi throughput numbers reported are much lower (about half) than in our 500 mile long drive case. Our work analyzes the performance of a single commercially operated WiFi network.

3. MEASUREMENT SETUP

3.1 Network

For our study we use a metro-scale WiFi deployment in the Long Island area in New York (population roughly 7 million). This service is called 'Optimum WiFi' [2] and is provided by Cablevision, a local cable TV provider and ISP. The WiFi network extends to most of populous areas of Long Island where our study is conducted. It also extends to parts of New York City, Pennsylvania, Connecticut and New Jersey, where Cablevision has service. The entire network has roughly 18,000 APs. WiFi access is provided free to all subscribers of Cablevision's TV or Internet services. We note that there are several hundreds of such metro-scale WiFi deployments in USA alone [20] While we expect that our general observations will be repeatable in another metro area deployment, we do note that more specific quantitative observations could be strongly tied to deployment density, radio characteristics of the APs and any handoff control on the APs.

For 3G access, we use Verizon's EVDO Rev. A service using the USB760 USB-based air interface available from Verizon. EVDO Rev. A service is capable of a maximum bit rate of 3 Mbps. Verizon's website [1] claims a maximum of 1.5 Mbps TCP throughout. In our experience we occasionally got a higher TCP throughput (about 2 Mbps).

3.2 Testbed

We use a Dell Latitude laptop running Linux as the client machine to be carried in the car. The original miniPCI WiFi interface from the laptop is removed and replaced by a carrier-grade interface (Ubiquity XR2 [3]) with transmit power set to 25 dBm. The interface uses Atheros chipset supported by the latest madwifi driver that we used. The WiFi card is connected to a high-gain (12 dBi) omni-directional antenna. We experimented with various other antennas (5 and 7 dBi) and transmit power choices (within FCC limits). However, the 12 dBi antenna provided the best connectivity. Note that due to the high gain, the beamwidth for this antenna on the vertical plane is relatively small (about 9°). Thus, assuming that the APs are deployed at the height of poletops, this antenna is likely to connect better to a distant AP relative to a very close AP. This is because the distant AP is much closer to the horizon relative to the car. The antenna is long – about 4ft. The antenna is carried inside the car with its top part sticking out from the sunroof.

The laptop also carries Verizon's USB-based USB760 EVDO Rev. A Modem as the second network interface. A USB-based Garmin GPS receiver is connected to the logging laptop. At startup the GPS receiver is initialized and the GPS location is logged every second.

A server program runs on a lab machine with a public IP address. The server accepts incoming connections from the WiFi/3G networks and transmits 1500 byte packets continuously to the client over TCP. as selected by the client. This enables download throughput tests to be reported in the next section.

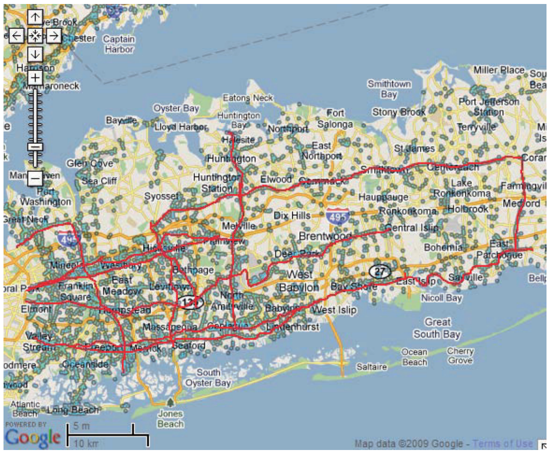

Figure 1: Map of the middle of Long Island – the area used in the long drive experiments. The route is shown in red. Approximate WiFi coverages are shown (from [2]).

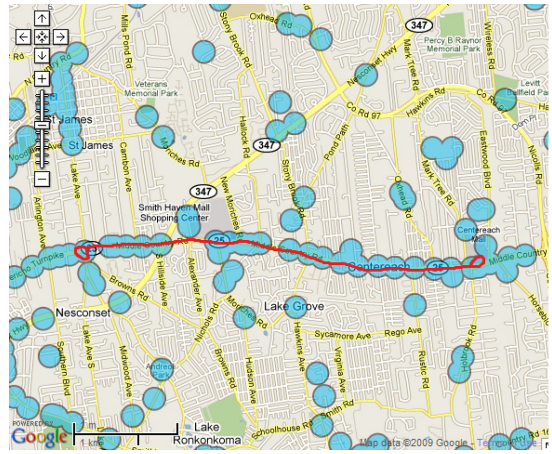

Figure 2: Map of the road stretch used for short repeated driving experiments (part of Route 25), along with the route shown in red. Approximate WiFi coverages are shown (from [2]).

We set the WiFi card's AP selection mode to automatic, i.e. the madwifi driver decides on the best AP to connect to. Auto rate control is disabled and a fixed bit rate of 11 Mbps is used. This is based on the observations in [11]. A few empirical tests indeed verified that this was the best policy in our setup.

DHCP is used to obtain the IP address. The Optimum WiFi network allows clients to retain the IP address between associations; thus DHCP delay is incurred only once at the beginning. This also allows us to retain the same TCP socket across associations. Optimum WiFi uses a browser-based authentication for the initial network access. This is done manually. Again this step is needed only once. Authentication is retained across associations.

Finally, a word about TCP. While using the stock TCP implementation on Linux we faced problems during our initial experiments when the WiFi link loses association for more than a few seconds. This sometimes makes the retransmission timeout really long (the timeout doubles at every step). This is a well-known issue for TCP in a mobile environment [8]. To tackle this problem, we simply set the maximum retransmission timer value on the server side TCP to 1 sec. We do note that this is not a general solution and more sophisticated client-side approaches are possible. However for the scope of this measurement study this solution worked reasonably well.

3.3 Driving Scenarios

Two driving scenarios are used for our evaluations.

- *Long Drive:* This is approximately a 500 miles drive shown in Figure 1 driven in a continuous fashion (close to 5 hours). The vehicle speed varied over a wide range depending on the road traffic. Each road is driven in both direction as the APs and their signal strengths visible from one side of the road are likely to be different from the other side of the road. This drive was performed only once. This long drive provides a reasonable sample of the quality of WiFi access from moving vehicles in a metro-scale deployment scenario. The association records for the long drive show over 900 unique APs.

- *Short Repeated Drives:* This is approximately a 9 mile drive on a selected stretch where the quality of the AP coverage is reasonably good. See Figure 2. Once again this drive constitutes a round-trip for the same reason. This drive was repeated 10 times to get a statistical confidence in the experimental results.

The Figures 1 and 2 show the approximate WiFi coverage taken from the provider [2]. Verizon coverage map for the same road stretch shows complete uninterrupted EVDO Rev. A coverage.

4. MEASUREMENT RESULTS

The laptop logs per second TCP throughputs (called instaneous throughputs) on both the connections along with GPS location and vehicle speed. The logs are postprocessed to develop the following analysis.

4.1 Quality of WiFi Coverage

To determine the quality of WiFi coverage we plot the CDF of run lengths (consecutive 1 sec segments) with zero and non-zero throughputs seen on WiFi. See Figure 3. Note that the 90-percentile zero throughput run length is 25 sec and 45 sec, for the short and long drives respectively. The same non-zero throughput numbers are 90 sec and 30 sec, respectively. The median numbers for non-zero throughputs are very short, however, just a few seconds. This reinforces the general experience that there are frequent disconnections on WiFi with vehicular mobility. However, we will note in the next subsection that we are nevertheless able to demonstrate impressive performance with respect to 3G.

It is important to note that this evaluation does not reflect true AP visibility; rather it shows whether there is actual connectivity and non-zero throughput on TCP. True AP visibility periods are likely to be longer than non-zero TCP throughput periods as we do not use any optimized

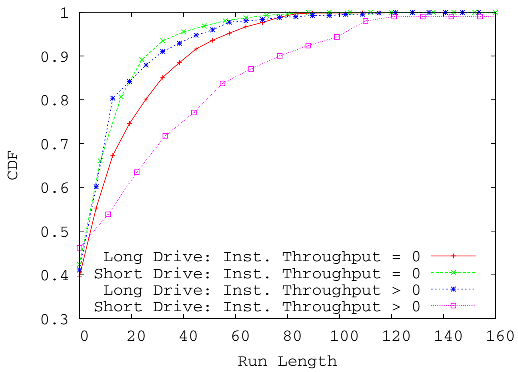

Figure 3: CDF of run lengths (consecutive 1 sec segments) with zero and non-zero throughputs seen on WiFi. Note that the short drive has zero throughput 25% of the times and the long drive 42% of the times.

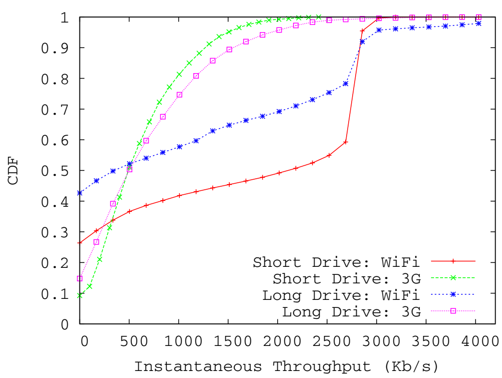

Figure 4: CDF of instantaneous TCP throughputs for WiFi and 3G.

handoff technique here. The TCP is also not optimized either except limiting the timeout period. Optimizing both of these will likely produce a much better WiFi experience [11]. Thus, our experience could be viewed as a lower bound for the WiFi performance.

4.2 Comparing WiFi and 3G Throughputs

In Figure 4 we compare the CDFs of instantaneous TCP throughputs on WiFi and 3G links. The long and short drives are shown separately. Note that WiFi throughputs are generally better in the short drive because of better coverage and slightly slower average driving speed experienced. However, 3G throughputs are very similar. While WiFi provides substantially better median throughput in the short drive (roughly 2400 Kbps vs. 500 Kbps) the median throughputs are similar in the long drive. However, WiFi has zero throughput on more occasions (roughly 25% vs. 10% for the long drive and 42% vs. 15% for the short drive). As noted before, this is expected as no special optimized handoff mechanism has been used. Also, WiFi coverage holes do exist (See Figures 1 and 2).

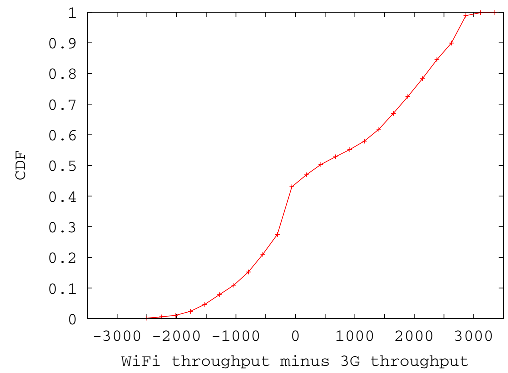

Figure 5: CDF of relative difference of instantaneous throughputs (in Kbps) between WiFi and 3G (i.e., WiFi minus 3G). Plot for the long drive only.

Also, notable is the fact that the *3G throughputs are well distributed in its entire range while WiFi demonstrates approximately tri-modal distribution* – zero or very low throughput, 0-2700 Kbps, higher upto about 3300 Kbps. Depending on the scenario each of these three modes persists for about 25%-40% of the times. The latter high throughput regions are specifically interesting for WiFi in our context.

Now we turn our attention to relative performance differentials between WiFi and 3G in the same time instants. See Figure 5 that plots the CDF of the algebraic difference between WiFi and 3G throughputs in the same time instants. This plot is done only for the long drive. Note that the fraction of times 3G performs better (45%) is only slightly higher than the times when WiFi has zero throughput (42%). In other words, when WiFi does have network layer connectivity, it is very likely that WiFi demonstrates better throughput. This again points to optimizing handoff making a serious impact.

Also, when WiFi outperforms 3G, it does it overwhelmingly so. The median difference is over 1500 Kbps. On the other hand, when 3G outperforms WiFi, the median difference is roughly 500 Kbps.

We have evaluated the correlation between WiFi and 3G throughputs for the same 1 sec intervals. *The correlation is very poor.* For the short drives it is 0.03 and for the long drive it is -0.03. This shows that these two networks can complement each other quite well.

4.3 Correlation with Vehicle Speed

It is also interesting to find out whether high throughputs are specifically correlated to slow vehicular speeds. In Figure 6 we present speedwise throughput statistics by breaking up the throughput data in the previous plot in three categories of speed – slow (0-20 Km/h), medium (20-40 Km/h) and high (40 Km/h and up). Note that the 3G throughput plots at different speeds are roughly similar. They also provide similar median throughputs. But the same cannot be claimed for WiFi, where slow speed has a clear advantage. This is expected as the WiFi physical layer is not built to support vehicular mobility, and the stock WiFi handoff implementations are not optimized for speed. The differences are much more for the short drive. This is likely because

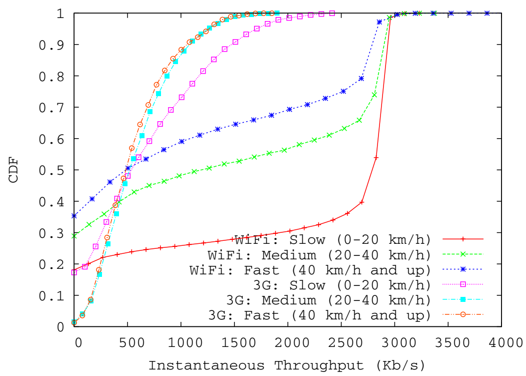

(a) Short drive.

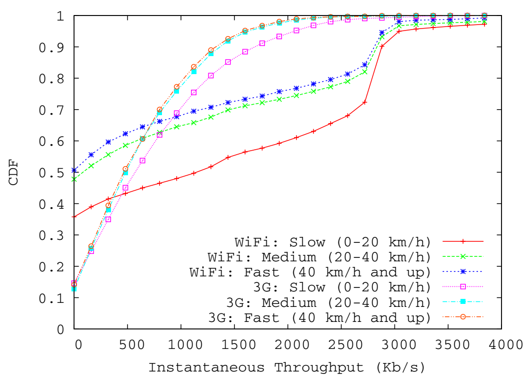

(b) Long drive.

Figure 6: CDF of Instantaneous TCP throughputs at different speeds for WiFi and 3G.

a better coverage better exposes the relationship between performance and speed.

Still at the fast speed, the median WiFi throughput is comparable to 3G for short drive. At slow speeds, it roughly doubles (for the long drive) or quadruples (for the short drive). WiFi is indeed impressive at slow speeds – providing over 2500 Kbps for over 60% or 35% of the times for short and long drives respectively. Thus, it is likely that WiFi would be much preferred in urban roadways in rush hour traffic, relative to high speed drives on rural highways. Incidentally, the former scenarios are likely to see metro-scale WiFi deployments.

4.4 Correlation with Location

Here, we evaluate how throughput is correlated with location. Previous research has indicated that WiFi signal strength in the same location has a reasonable degree stability and this property has been utilized to develop location-based handoff techniques [9]. However, throughput stability has not been evaluated. In regards to 3G, throughput stability has been evaluated to conclude that the entropy of the throughput distribution reduces significantly when conditioned on location [28]. However, no direct comparison between WiFi and 3G exists in this aspect.

For our analysis here, we use the general approach reported in [28]. The idea is to use the notion of *information entropy* that measures the level of uncertainty associated

	3G		WiFi	
Grid Size	Total H	Loc. H	Total H	Loc. H
10x10	1.56	1.19	2.26	1.44
20x20	1.52	1.16	2.25	1.56
30x30	1.51	1.12	2.23	1.71

Table 1: Comparison of total and location entropies for 3G and WiFi networks.

with a random process. The information entropy of a discrete random variable X is defined as,

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

where $p(x)$ is the probability mass function, $0 \leq p(x) \leq 1$. The lower the entropy, the lower the uncertainty associated with the process. Hence, lower entropy indicates more predictability. Here, the random variable is the throughput. To address the influence of location on throughput we use the concept of *location entropy* as in [28]. Location entropy $H(X|l_i)$ is the entropy of throughput for a specific location l_i. This is defined as follows,

$$H(X|l_i) = \sum_{x \epsilon X} p(x|l_i) \log_2 p(x|l_i), \quad (2)$$

where $p(x|l_i)$ is the probability mass function of throughput at location l_i. We use a similar technique as in [28] to calculate the entropy of instantaneous throughput measurements from the data corresponding to the short drives. Both total entropy (as in Equation 1) and location entropy (as in Equation 2) are computed.

The entire geographic space containing the drives is checkered with grids of different sizes: 10m×10m, 20m×20m and 30m×30m, with each grid square being indicative of a location. The grids simply discretize the space. The three different sizes are chosen to estimate what level of discretization is appropriate. The average instantaneous throughput per drive per grid square is treated as the realization of single independent random observation. Hence, each drive can produce at the most one observation per grid square. The total entropy is computed over all instantaneous throughput measurements for 3G and WiFi separately. The location entropy is computed similarly for each grid square and then averaged over all squares. The results are presented in Table 1. We can see that the total entropy of 3G is much lower than the total entropy of WiFi showing better predictability of 3G throughputs. Note that the size of grid squares makes little difference. We also see that the location entropy of both 3G and WiFi is lower than their respective total entropies. This result suggests that when conditioned on location, instantaneous throughput becomes more predictable in both networks. But the degree of improvement is somewhat similar.

4.5 Temporal Correlation

We also evaluate the temporal correlation of instantaneous throughput values. Figure 7 shows the autocorrelation for the long drive. Note the excellent temporal stability of 3G relative to WiFi. This is somewhat in contrast with measurements in [28], where the authors have found that the autocorrelation in HSDPA throughputs was quite small. Also,

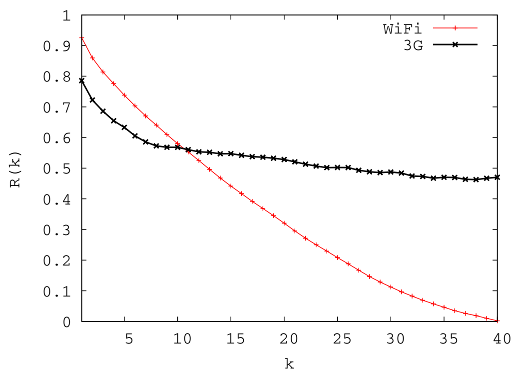

Figure 7: Autocorrelation $R(k)$ of the instantaneous throughputs measured in 1 sec intervals, where k denotes the time lag in seconds.

AR(k) Model	MSE(Y_{est})/VAR(Y)
AR(1)	0.120
AR(2)	0.119
AR(4)	0.119
AR(16)	0.117

Table 2: Prediction accuracy using AR(k) model.

interestingly for lower values of lag (k), the autocorrelation in WiFi throughputs is higher than that of 3G, though it decays promptly.

The knowledge of autocorrelation is useful is estimate how well the recent throughput experiences can predict the throughputs to be seen in near future. This could be useful in packet scheduling in a hybrid access scheme.

The prediction model can use one of several standard techniques such as autoregressive (AR) model, moving average (MA) model, autoregressive moving average (ARMA) model, etc. In [28] the authors have found that even the simplest of the models (AR) produces a very good estimate of the current throughput based on throughput measurements in the recent past. Thus, we have kept our analysis limited to AR models only. We have evaluated the performance of the AR(k) model for several values of k to see the quality of estimation. Table 2 shows the mean-squared estimation error divided by the sample variance (MSE/s^2). This is based on a standard AR(k) fitting method used in the mathematical software Octave, when applied on the time series WiFi throughput data.[1]

5. SUMMARY AND CONCLUSIONS

We now summarize the experiences we gathered from the measurement study.

- There is upto a factor of 4 improvement of median throughput on WiFi relative to 3G. This is somewhat drive specific.
- Instantaneous thoughputs on WiFi can be zero occasionally (very roughly one-third of the times). On the other hand, very roughly one-third of the times it can deliver significantly high throughput (over 2.5 Mbps).
- When WiFi is available, it is very likely (roughly 90%) that it outperforms 3G. Such outperformance is significant (a median difference of 1500 Kbps).
- There is more variability in WiFi throughputs relative to 3G. But variability reduces when conditioned on location. This is true for both networks.
- There is a good temporal correlation for the instantaneous throughputs on both networks, at least for short lags (more for WiFi).
- Slow speed means higher throughput on WiFi. But 3G is much less sensitive to speed.
- Correlation between WiFi and 3G throughputs are poor signifying that diversity techniques will be successful.

In conclusion, we have shown by experimental measurements that a metro-scale WiFi and 3G networks exhibit very different characteristics under vehicular mobility. WiFi has frequent disconnections even in a commercially operated, metro-scale deployment; but when connected indeed delivers high throughout even in a mobile scenario. The 3G network offers much lower throughputs, but provides much better coverage and less throughput variability. Our general experience indicates that a hybrid design that exploits the best properties of the two networks opportunistically can be very successful [4]. This can benefit the consumer with better throughput and as well as lower the cost for the provider by moving expensive 3G bits onto WiFi networks.

We do note that while our experiences are indeed provider specific, the general observations are not surprising and very likely would be repeated in other deployments. We expect that our work will encourage similar measurement studies. We also note that lack of access to the provider network lets us measure end-to-end throughput only, instead of the throughput on the wireless hop alone. The latter would certainly be more interesting to study. Our future work includes (i) understanding and improving handoff characteristics and TCP dynamics for vehicular WiFi access, specifically for metro-scale WiFi deployments, and (ii) designing hybrid access methods that exploit the best of 3G and WiFi, particularly targeting interactive use.

[1]If x_t is the average throughput at the t-th sec, the AR(k) model estimates x_t as $\hat{x_t} = c + \sum_{i=1}^{k} \phi_i x_{t-i} + \epsilon_t$, where c is a constant, ϕ_i's are the parameters of the model, and ϵ_t is white noise.

6. REFERENCES

[1] Coverage Locator – Verizon Wireless. http://www.verizonwireless.com/b2c/CoverageLocatorController?requesttype=NEWREQUEST.

[2] Optimum WiFi. http://www.optimum.net/MyServices/WiFi/.

[3] Ubiquity Networks, Inc. http://www.ubnt.com.

[4] A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting Mobile 3G Using WiFi. In *Proc. ACM MobiSys Conference*, 2010.

[5] Aruna Balasubramanian, Ratul Mahajan, Arun Venkataramani, Brian Neil Levine, and John Zahorjan. Interactive WiFi connectivity for moving vehicles. *ACM SIGCOMM Comput. Commun. Rev.*, 38(4):427–438, 2008.

[6] Aruna Balasubramanian, Yun Zhou, W. Bruce Croft, Brian Neil Levine, and Aruna Venkataramani. Web search from a bus. In *Proc. ACM CHANTS Workshop*, 2007.

[7] Vladimir Bychkovsky, Bret Hull, Allen Miu, Hari Balakrishnan, and Samuel Madden. A measurement study of vehicular Internet access using in situ Wi-Fi networks. In *Proc. ACM MobiCom Conference*, 2006.

[8] R. Caceres and L. Iftode. Improving the performance of reliable transport protocols in mobile computing environments. *IEEE Journal on Selected Areas in Communications*, 13(5):850–857, 1995.

[9] Pralhad Deshpande, Anand Kashyap, Chul Sung, and Samir R. Das. Predictive methods for improved vehicular WiFi access. In *Proc. ACM Mobisys Conference*, pages 263–276, 2009.

[10] A. Doufexi, EK Tameh, AR Nix, and A. Molina. Hotspot wireless LANs to enhance the performance of 3G and beyond cellular networks. *IEEE Communications Magazine*, 41, 2003.

[11] Jakob Eriksson, Hari Balakrishnan, and Samuel Madden. Cabernet: A WiFi-Based Vehicular Content Delivery Network. In *Proc. ACM MobiCom Conference*, 2008.

[12] Anastasios Giannoulis, Marco Fiore, and Edward W. Knightly. Supporting vehicular mobility in urban multi-hop wireless networks. In *Proc. ACM MobiSys Conference*, 2008.

[13] V. Gunasekaran and F.C. Harmantzis. Towards a Wi-Fi ecosystem: Technology integration and emerging service models. *Telecommunications Policy*, 32(3-4):163–181, 2008.

[14] David Hadaller, Srinivasan Keshav, Tim Brecht, and Shubham Agarwal. Vehicular opportunistic communication under the microscope. In *Proc. ACM MobiSys*, 2007.

[15] Y. Lee. Measured TCP performance in CDMA 1xEV-DO network. In *Proc. PAM Conference*, Adelaide, Australia, Mar. 2006.

[16] W. Lehr and L.W. McKnight. Wireless Internet access: 3G vs. WiFi? *Telecommunications Policy*, 27(5-6):351–370, 2003.

[17] 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. ACM MobiCom Conference*, pages 211–222, 2008.

[18] Ratul Mahajan, John Zahorjan, and Brian Zill. Understanding WiFi-based connectivity from moving vehicles. In *Proc. Internet Measurement Conference (IMC)*, pages 321–326, 2007.

[19] K. Mattar, A. Sridharan, H. Zang, I. Matta, and A. Bestavros. TCP over CDMA2000 networks: A cross-layer measurement study. In *in Proc. Passive and Active Measurements (PAM) Conference*, 2007.

[20] Muniwireless.com. List of cities and counties with large WiFi networks. http://www.muniwireless.com/reports/Mar-28-2009-list-of-cities.pdf.

[21] Vishnu Navda, Anand Prabhu Subramanian, Kannan Dhanasekaran, Andreas Timm-Giel, and Samir R. Das. MobiSteer: using steerable beam directional antenna for vehicular network access. In *Proc. ACM MobiSys Conference*, 2007.

[22] J. Ormont, J. Walker, S. Banerjee, A. Sridharan, M. Seshadri, and S. Machiraju. A city-wide vehicular infrastructure for wide-area wireless experimentation. In *Proc. ACM WinTech Workshop*, pages 3–10, 2008.

[23] J. Ott and D. Kutscher. Drive-thru internet: IEEE 802.11b for automobile users. In *Proc. IEEE Infocom*, 2004.

[24] J. Ott and D. Kutscher. A disconnection-tolerant transport for drive-thru internet environments. In *Proc. IEEE Infocom Conference*, 2005.

[25] Qualcomm. The economics of wireless mobile data. http://www.wirelessdevnet.com/library/WirelessMobileData.pdf.

[26] Wee Lum Tan, Fung Lam, and Wing Cheong Lau. An empirical study on 3G network capacity and performance. In *Proc. IEEE INFOCOM*, pages 1514–1522, 2007.

[27] Peng Yang, Hui Deng, and Yuanchen Ma. Seamless integration of 3G and 802.11 wireless network. In *Proc. MobiWac*, pages 60–65, 2007.

[28] Jun Yao, Salil S. Kanhere, and Mahbub Hassan. An empirical study of bandwidth predictability in mobile computing. In *Proc. ACM WiNTECH*, pages 11–18, 2008.