Understanding WiFi-based Connectivity from Moving Vehicles

Ratul Mahajan
Microsoft Research

John Zahorjan
University of Washington

Brian Zill
Microsoft Research

Abstract – Using measurements from VanLAN, a modest-size testbed that we have deployed, we analyze the fundamental characteristics of WiFi-based connectivity between basestations and vehicles in urban settings. Our results uncover a more complex picture than previous work which was conducted in more benign settings. The interval between a vehicle coming into and going out of range of a basestation is often marred by intermittent periods of very poor connectivity. These "gray periods" are hard to reliably predict because their arrival is not signaled by metrics such as signal strength, loss rate, speed or distance from the basestation. At the same time, they also do not consistently occur at the same spot. Our analysis suggests that gray periods are not caused by the motion of the vehicle *per se* but by the variability in the urban radio environment combined with the vehicle traversing locations that are poorly covered by the basestation. We also find that knowledge of past connectivity can be used to identify regions where gray periods are more likely to occur as well as regions where the vehicle is likely to experience good connectivity.

Categories and Subject Descriptors
C.4 [Performance of systems]: Performance attributes

General Terms
Measurement, performance

Keywords
Vehicular networks, measurement, WiFi

1. INTRODUCTION

WiFi-based networks are becoming increasingly ubiquitous, and in many cases entire cities and campuses are being covered [16, 15]. Though it may be possible to get brief periods of connectivity [3], these networks and the WiFi technology itself are not designed to provide connectivity to moving vehicles.

Like others before us [10, 7, 3], we ask if WiFi can be leveraged to provide connectivity (in areas of good coverage) from moving vehicles. Compared to cellular networks, which can enable such connectivity today,[1] 802.11 has two advantages. The key one is that it is significantly cheaper because it operates in an unlicensed band and does not require regulatory approval. It can also support higher data rates. For instance, for currently available versions, the maximum data rate with 802.11 a/g is 54 Mbps and with EVDO Rev A is 3.1 Mbps [6] (though the actual throughput is lower for both). This order of magnitude difference persists for the upcoming versions. 802.11n promises 100-600Mbps [1] and EVDO Rev B promises 9.3 Mbps [6].

However, it is not clear *a priori* that 802.11 can be co-opted for this purpose. At vehicular speeds, its short range could make it difficult to provide uninterrupted connectivity because basestations would come in and go out of range quickly. Its radio environment is also hostile [2], with high variability and interfering sources. Vehicular mobility only makes it worse: the wireless environment will change rapidly as the vehicle moves.

To investigate feasibility, we have deployed a modest-size, WiFi-based testbed, called VanLAN. It currently has eleven basestations and two clients that are mounted on vans. VanLAN has been "operational" since January 2007. More information on the testbed and some of our measurement data is available at http://research.microsoft.com/vanlan/.

In this paper, we analyze measurements from VanLAN to understand connectivity between movings vehicles and basestations. We are interested in the basic nature of WiFi-based connectivity, such as how it varies as the vehicle moves and whether it is stable across traversals of the same location. As such, unlike previous work [10, 7, 3], we consider raw connectivity rather than performance obtained by current transport protocols and do not consider the impact of overheads such as client authentication and IP address acquisition.

Our analysis reveals a challenging environment. Instead of being continuous between the times the client comes in and goes out of range of a basestation, the connectivity is often marred by intermittent "gray" periods of very poor connectivity. This is unlike the three phases of connectivity – poor quality "entry" and "exit" phases and a good quality "production" phase – reported in an earlier, smaller-scale study that was conducted in a more controlled environment [10]. We find that the occurrence of gray periods is hard to predict based on current measurements of, for instance, loss rate, signal strength, speed, or distance from the basestation. Gray periods do not consistently occur at the same location, and they are also unlikely to occur due to vehicular mobility itself. But they likely occur due to a combination of variability in the urban radio environment and the vehicle encountering locations that are poorly covered by the basestation.

On the positive side, information on past performance at a location is valuable in this environment. We find that it can be used to predict regions where gray periods are more likely as well as regions where connectivity is likely to be good. Combined with the constrained, and therefore predictable, paths in vehicular networks, this may be used to inform and prepare applications for future conditions and alleviate the impact of the uneven connectivity characteristics of this environment.

[1]WiMax is another possibility. Because it is not widely deployed yet, we defer its investigation to future work.

IMC'07, October 24-26, 2007, San Diego, California, USA.

2. RELATED WORK

There are many studies that characterize 802.11 behavior [2, 9, 4, 8, 11] but only a few analyze it in the face of vehicular mobility [10, 7, 3]. Ott and Kutscher study the performance of UDP and TCP traffic in a carefully planned setup with two BSs along a highway (with few other WiFi users in the vicinity) [10]. They find that the connection between the car and a BS can be roughly divided into three phases. In the "entry" and "exit" phases, when the car is far from the BS, the throughput is low. In the "production" phase, when the car is close to the BS, the throughput is high. Gass *et al.* use a single BS besides a highway in a desert to study the impact of speed and backhaul network properties on application performance [7]. In this interference-free environment, they find that the wireless link is not the main bottleneck. Instead, the likely high delays in the backhaul network, on the order of 100 ms, interact poorly with applications that require many round trip exchanges.

In contrast, we consider a larger-scale and a more realistic urban setting which contains interference from 802.11 and non-802.11 sources and obstacles such as trees and buildings in the path between BSs and vehicles. As a result, we find a significantly more challenging wireless environment which cannot be simply divided into three phases [10].

The CarTel study quantifies the performance of uploading data from a moving car using urban WiFi APs that happen to be open [3]. It finds that reasonable performance can be had – the median connection duration is 13s and median upload bandwidth is 30 KBps. It also finds that overheads due to association, authentication, and IP address acquisition are significant. In contrast, we focus on the fundamental characteristics of WiFi-based connectivity.

Like VanLAN, DieselNet [5] employs a testbed of WiFi-equipped buses as well. However, it has a different technical thrust. While we focus on vehicle-to-infrastructure connectivity for today's applications, it focuses on delay tolerant networks formed using vehicle-to-vehicle connectivity.

3. THE VANLAN TESTBED

To study the performance of 802.11 in vehicular settings, we have deployed a modest-size testbed of BSs and vehicles. We use the testbed nodes to generate and log probe traffic which we then analyze to understand the characteristics of connectivity. We describe our testbed in this section.

VanLAN currently consists of eleven BSs and two mobile clients. The BSs are spread across five office buildings on the Microsoft campus in Redmond, WA. Their geographic placement is shown in Figure 1. The box bounds the region in which at least one packet is received by mobile clients from any BS, though the vast majority of the packets are received in the right half of the box. The network of BSs is connected, but not all pairs of BSs can hear each other.

The mobile clients are vans that provide a shuttle service on and around the campus during the day. They visit the part of the campus where the BSs are present roughly ten times a day. The roads in this part are similar to urban neighborhood streets with a speed limit of around 40 Kmph.

Both BSs and clients are small form factor desktops. BSs are placed on top floors of the buildings, but their antennae are mounted on the roofs. Low-loss coaxial cables connect the radios (inside the desktops) and antennae. Similarly, the clients are placed inside the vans and their antennae are mounted on the roof. The client desktops are powered by a dedicated battery that is different from the van's main battery. This battery charges when the van is on and powers the clients for about four hours after the van is switched off. This time is used for software updates through a wireless connection with another computer located near the vans' overnight parking space. (Basestations have Ethernet connections for this purpose.)

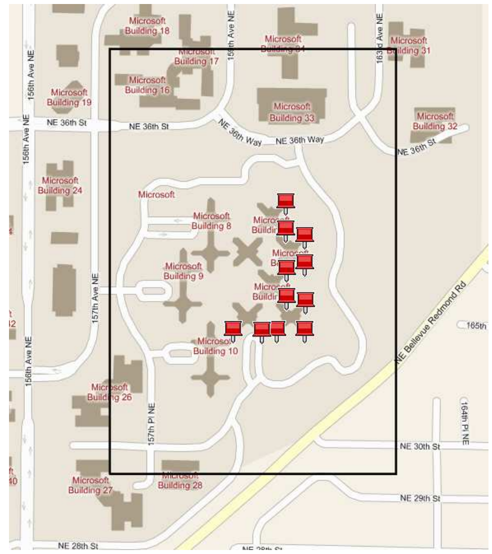

Figure 1: The placement of VanLAN BSs. The box (559m × 828m) bounds the region in which at least one packet is received by vans from any BS.

All nodes have two radios. One radio is configured to Channel 1 of 802.11g and the other to Channel 11. To reduce interference, the two antennae are separated by at least one foot [12]. By comparing the cases where only one radio is active to where both are active, we have confirmed that any residual interference is minimal.

Our radios operate in *ad hoc* (IBSS) mode using a locally modified device driver. One modification forces the use of a fixed BSSID instead of a randomly generated one. This prevents (temporary) network partitions when nodes end up with different BSSIDs. It also means that a BS and a client that come into range can start communicating immediately, without waiting for their BSSIDs to be reconciled. Yet another modification lets us log every received frame along with a hardware timestamp and PHY layer information such as RSSI while communicating normally (i.e., the radio is not put in "monitor" mode).

VanLAN uses the following hardware. EnGenius' EMP-8602 modules, which are based on the Atheros 5213 chipset, are used as radios. Their output power is 400 mW at 1 Mbps and lower at higher transmission rates. HyperLink's HG2403MGU antennae are used for the vans and HGV-2404U antennae are used for the basestations. Both types are omnidirectional in the horizontal plane but radiate less energy directly above and below.

The clients also have an externally mounted GPS unit, so we know their locations. We use GlobalSat's BU-353 GPS unit which is based on the SiRF Star III chipset and outputs data once per second. The uncertainty in the location estimate of this chipset is under three meters 95% of the time.

4. CONNECTIVITY SESSIONS

We begin our investigation by studying the basic characteristics of 802.11-based connectivity between a fixed BS and a moving vehicle. The period of connectivity between a BS and the client can be divided into sessions, where a session is a contiguous period of time in which the client can communicate with the BS. To study the properties of these sessions, we leverage beacons that 802.11 BSs send roughly every 100 ms. Using beacons instead of custom traffic lets us consider connectivity sessions with not only VanLAN

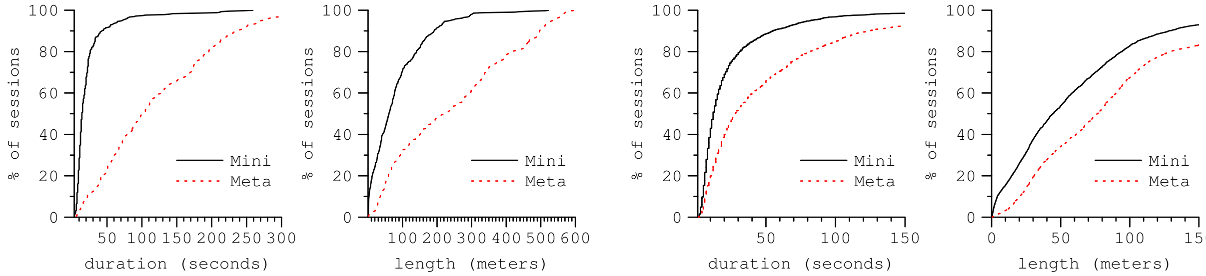

(a) VanLAN BSs (b) Other BSs in the environment

Figure 2: Mini sessions are much shorter than meta sessions, implying that meta sessions contain intermittent disruptions. The graphs plot the CDFs of duration and length of meta and mini sessions. The x-axis ranges differ for the two types of BSs.

BSs but also other BSs encountered by our vans. This helps verify whether the properties of the BSs in our testbed are similar to those of other BSs in the environment. However, beacons let us study connectivity in only one direction – from BSs to the client. We have confirmed using VanLAN BSs that connectivity in the other direction behaves similarly. The results below are based on a four-week trace of beacons logged on Channel 11. Results for Channel 1 are qualitatively similar, except that Channel 11 provides overall better connectivity in our testbed.

Connectivity sessions start when a beacon is first received from the BS and expire if no beacon is received for a fixed time threshold. To highlight intermittent disruptions, we use two time thresholds to expire current sessions. A high threshold of 60s is intended to capture *meta* sessions, i.e., the entire period between the client coming and going out of range of the BS. A low threshold of 2s is intended to capture *mini* sessions that are contained within a meta session. Consecutive mini sessions are separated by periods of very poor or no connectivity. The 2s limit is chosen so that enough (at least 20) beacons are transmitted to reliably identify these periods. If disruptions are common, mini sessions will be much shorter than meta sessions.

Figure 2 shows the distribution of the duration and length of meta and mini sessions for VanLAN BSs and for other BSs in the environment. It excludes sessions in which less than 10 beacons are received to ignore trivial sessions with distant BSs. Length of a session is distance covered by the van between the starting and ending positions of the session.

Meta sessions are much longer than mini sessions. For VanLAN, the median durations are 100s for meta sessions and 14s for mini sessions; for other BSs, the corresponding durations are 30s and 12s. This implies that intermittent disruptions often break meta sessions into multiple mini sessions. We refer to these disruptions as "gray" periods. They do not necessarily represent a complete loss of connectivity because only a finite number of beacons are transmitted during this period. They do, however, represent a period of very poor connectivity in which none of the 20 or more transmitted beacons are received. Such disruptions can significantly hurt interactive applications such as voice and even those that use TCP if connections time out.

The differences in VanLAN and other BSs in Figure 2 contrast their connectivity sessions. Both have gray periods, but VanLAN BSs have longer meta sessions, likely because they are closer to nearby roads whereas other BSs could be deep inside buildings. Interestingly, however, due to gray periods these longer meta sessions do not translate to longer mini sessions. Both types of BSs have comparable mini session durations.

We now illustrate the behavior of connectivity sessions in more detail by showing individual examples. Figure 3 shows the experience of the van in one round with respect to the circled BS. We see that the connectivity is not continuous but contains several silent periods of at least 2s. Except for the one big silent period at the bottom of the path, where the van is completely out of range, the disruptions are short and represent gray periods. Observe that gray periods are not limited to regions where the client is far from the BS and sometimes occur close to the BS.

Figure 4 provides a detailed view of an example meta session. It plots the beacon reception ratio (BRR) over 1-second intervals. This meta session has two major mini sessions, separated by a gray period at 80s. Even within a mini session, the connection quality varies significantly due to vehicular mobility and changes in the wireless environment.

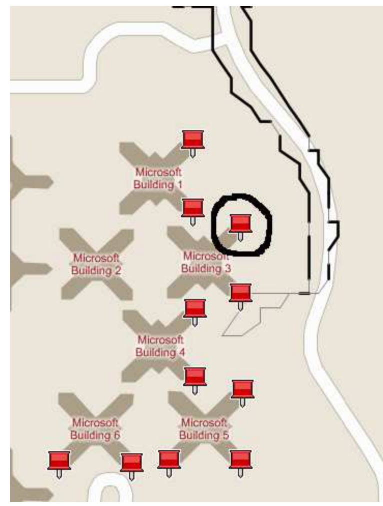

Figure 3: Gray periods can occur even close to the BS. The figure shows the experience of the van in an example round with respect to the circled BS. Thick lines represent regions where beacons were received without a silent period of at least 2s. Thin lines represents the complement.

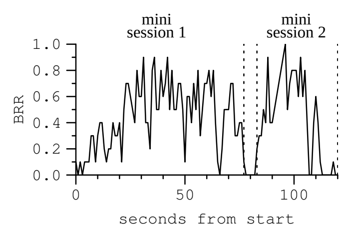

Figure 4: Connection quality within a meta session varies significantly. The graph plots the BRR over 1-second intervals in an example meta session.

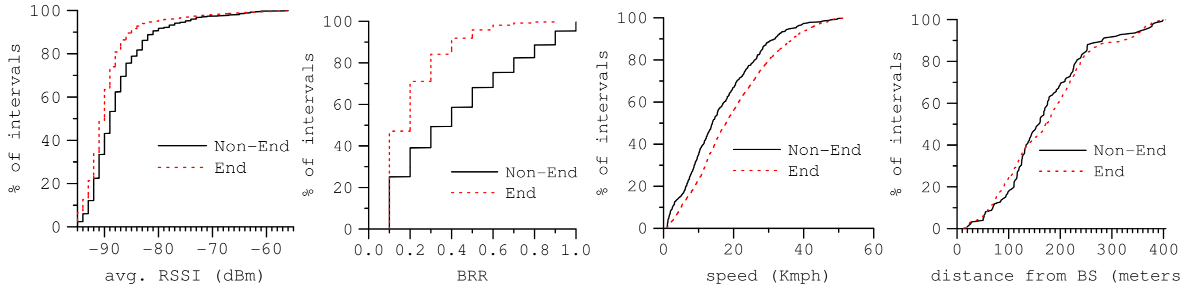

Figure 6: Gray periods are hard to predict based on current measurements. The graphs plot the CDFs of four measures in 1-second intervals immediately preceding a gray period ("End") and in other 1-second intervals ("Non-End").

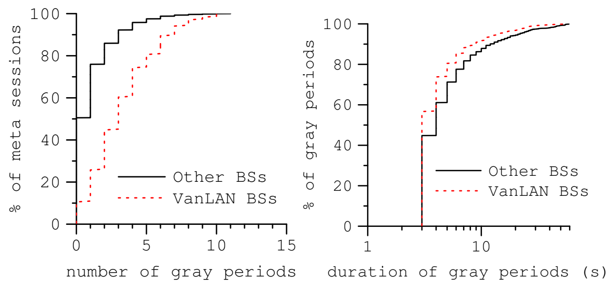

(a) (b)

Figure 5: *(a):* Gray periods occur frequently. The graph plots the CDF of the number of gray periods in a meta session. *(b):* Many gray periods are short-lived but some are long. The graph plots the CDF of the duration of gray periods.

Thus, the nature of connectivity between a BS and a mobile client that we find is quite complex. It is different from that observed in previous WiFi-based studies of controlled environments [10, 7]; well-defined phases are absent, and poor connectivity periods can arise even close to the BS. It is, however, similar to studies of cellular networks in urban settings [13].

5. UNDERSTANDING GRAY PERIODS

Gray periods can pose a significant challenge to providing uninterrupted connectivity to applications. In this section, we first study their frequency and duration and then investigate whether their occurrence can be predicted. We also speculate on the factors that lead to gray periods.

5.1 Frequency and Length

Figure 5(a) shows the distribution of the number of gray periods in a meta session. It excludes gray periods that separate mini sessions with less than 10 packets, to ignore trivial mini sessions. We see that gray periods are common: roughly 80% of meta sessions for VanLAN BSs and 50% of the meta sessions for other BSs contain at least one gray period. (Recall that VanLAN meta sessions tend to be longer than those for other BSs.)

Whether an application is robust to gray periods depends on the duration of poor connectivity that it can withstand. Figure 5(b) shows the duration of gray periods. The x-axis is in log scale. The minimum is 3s because we expire a mini session only when the silence period lasts for more than 2s. Most gray periods are short-lived, though some last for more than 10s. Given typical RTTs of less than 200 ms [14], our results suggest that TCP-based applications are likely to suffer under these conditions because gray periods tend to last for more than a few RTTs.

5.2 Prediction using Current Measurements

We now investigate if the occurrence of a gray period can be predicted. This prediction ability would enable applications to take steps to blunt their impact. We consider two classes of prediction techniques: $i)$ using current measurements; and $ii)$ using longer-term history of performance at the current location.

To evaluate the effectiveness of current measurements in predicting gray periods, we consider measurements of RSSI, BRR, speed, and distance from the BS. Many wireless clients today use the first two measures to determine if the connection to the BS is about to falter. The last two may be able to predict the onset of a gray period if gray periods commonly occur, for instance, at high speeds or far away from the BS.

Figure 6 shows that none of the measures above can reliably predict an impending gray period. It plots the distribution of these measures in the 1-second interval before a gray period ("End," of a mini session) and in other 1-second intervals. ("Non-End"). The distance graph is based on VanLAN BSs alone because we do not know the location of other BSs; the other graphs include data from all BSs. Out of the four, RSSI, speed, and distance measurements for periods right before a gray period are almost indistinguishable from measurements during other times. For BRR, the intervals before a gray period usually have lower BRR, which suggests that gray periods often follow times of poor connectivity. Nevertheless, there is no threshold a client can use to reliably predict the onset of a gray period; any threshold will have many false positive or false negatives. We also considered variations of these measures, such as combining them and using exponential averages, but find that they too are ineffective at predicting an impending gray period.

5.3 Prediction using Longer-term History

Another potential method for predicting gray periods is using the history of the connectivity experienced at a location. This would be effective if most gray periods consistently occur at the same location, for instance, due to permanent obstructions (e.g., trees).

By comparing where gray periods occur in individual traversals to average performance of those locations across multiple traversals, we find that even history cannot reliably predict the occurrence of a gray period. We illustrate our point by showing the average behaviors of the paths in Figures 3 and 4. To compute the average behavior at a location, we first map each location to a 10m x 7m grid, by rounding raw coordinate degree values to four decimal places. We do this not only because of potential imprecision in GPS output but also because the van does not traverse the exact same locations on each traversal of a path. We then compute the average over all locations that map to the same grid.

Figure 7 shows the performance of the paths we showed previously when averaged over the entire day. Observe that almost all intervening locations of poor connectivity disappear upon averaging, which implies that most gray periods do not consistently occur at the same spot, and thus are not caused by physical obstructions. The variability in BRR for the average view is much less than that

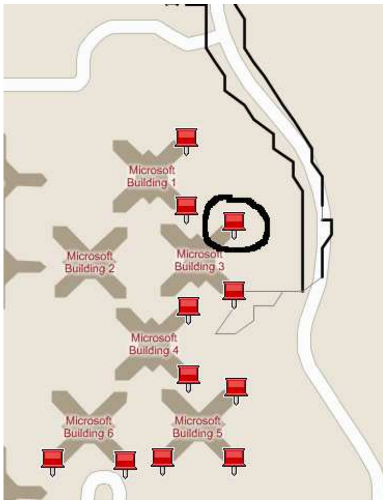

(a) Average view of the path in Figure 3

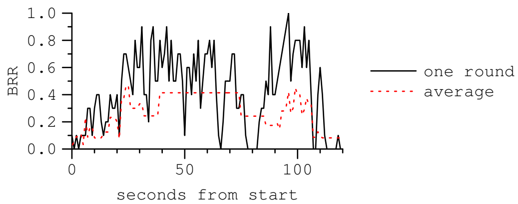

(b) Average view of the session in Figure 4

Figure 7: Gray periods do not consistently occur at the same spot. *(a)* The average (over a day) performance at each location along the path. Thin lines connect locations with very poor average connectivity (less than 10% BRR). Thick lines represents the complement. *(b)* The single-round and average (over a day) BRR for locations along the session.

for the single-round view in Figure 7(b). This suggests that interference and wireless variability in the urban environment are often responsible for gray periods.

To show this more directly, we consider for each location its mean BRR and the probability of observing a gray period there. The mean BRR is computed across traversals in one day of that location and the probability is the fraction of traversals for which a gray period is observed. Figure 8 shows a scatterplot of the two measures. While the probability of observing a gray period is non-zero even for locations with high mean BRR, it is higher for locations with lower mean BRR. Thus, the client is more likely to experience a gray period when traversing locations with poorer average connectivity. The wireless variability in a given traversal is more likely to cause gray periods at those locations.

While this behavior does not allow us to deterministically predict a future gray period, can it identify regions where the client is likely to encounter a future gray period? Figure 9 shows that this is true in our environment. It plots the average BRR observed across a day versus the probability of encountering a gray period during the next day. We see that there is enough stability in the environment such that past performance can be used to predict the probability of encountering a gray period in the future.

6. PERFORMANCE PREDICTABILITY

In this section, we investigate if the ability to predict performance exists more broadly than predicting gray periods. That is, is the performance at a location across traversals stable enough for a client to be able to predict future performance?

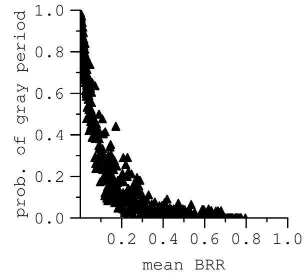

Figure 8: Gray periods are more likely to be observed at locations that have poorer connectivity on average. The graph is a scatterplot of the average BRR for a location across traversals in a day versus the probability of observing a gray period at that location during that day.

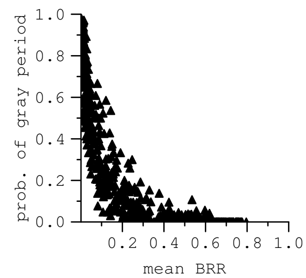

Figure 9: Past performance can help identify locations where a gray period will likely be encountered in the future. The graph is a scatterplot of the average BRR at a location across traversals in a day versus the probability of encountering a gray period at that location during the next day.

The results in this section are based on a two-week trace of broadcast probes generated by BSs and vans on Channel 11. These probes had the same frequency as 802.11 beacons, but were 500 bytes in size, to mimic the larger size of data packets.

To evaluate performance stability across time, we measure the standard deviation of the reception ratio across different traversals of the same location. In addition to the connectivity between a client and BS, we study the connectivity between two BSs. This lets us isolate the impact of vehicular mobility from the inherent variability of the outdoor wireless environment. Further, there is variability inherent in our measurement methodology because reception ratios at individual locations are computed using a small number of transmission attempts. To isolate this, we study the variability of reception ratio obtained using a simulation with fixed reception probability but the same number of (synthetic) samples.

Figure 10 shows the standard deviations for the three cases, as a function of mean reception ratio over the course of a day. Individual reception ratios are computed over a 1-second period which is comparable to the time our clients spend at each location (normalized to a grid); the means are classified into bins of size 0.1. For the client-BS case, each y-value is the standard deviations across all client-BS-location triplets with the given x-value. For the BS-BS case, each y-value is the standard deviations for all BS-BS pairs with the given x-value.

The graph shows that the variability in the Client-BS case is similar to that in the BS-BS case. This implies that vehicular mobility has at most a second order effect on variability beyond what is inherent in the outdoor environment. Thus, the speeds with which our vans move and the differences in the exact location across traversals have minimal impact. However, this does not mean that vehicular mobility does not introduce additional complexity in such

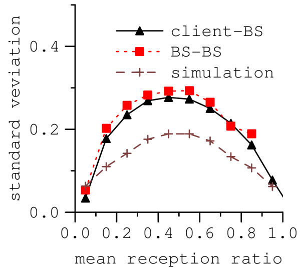

Figure 10: Vehicular mobility does not introduce additional variability in client-to-BS performance compared to BS-to-BS performance. The graph plots the standard deviation as a function of the reception ratio at the location.

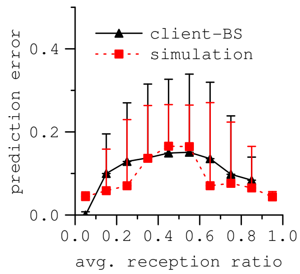

Figure 11: Predicting performance at a location based on historical data is feasible. The graph plots the prediction error as a function of the average reception ratio observed in the previous day. The lines connect the mean and the whiskers depict the 90th percentile error.

networks. Because of it, we find that the variability in performance across time is higher for client-BS connectivity than BS-BS connectivity. This will likely render protocols such as transmission rate adaptation that base future behavior on recent past less effective.

That the standard deviation of simulation in Figure 10 is lower implies that the wireless environment is more variable. However, the standard deviation for a setting where the reception ratio changes randomly in each interval is greater still (roughly 0.35). Thus, there is some consistency in performance at a location.

To understand whether this can be tapped to predict performance based on historical data, lets consider the task of predicting reception ratio over 1-second intervals. As our prediction for a location, we use the average reception ratio observed at that location across all traversal in the previous day. For simplicity, we ignore any time-of-day effects; taking those into account may improve prediction accuracy. We measure prediction error as the difference between the predicted and actual reception ratio.

Figure 11 plots the mean and 90th percentile prediction error. To quantify the error inherent in our measurement methodology, it also plots the prediction error in a simulation where the underlying reception probabilities are fixed. We see that while the 90th percentile error is higher, the mean error in the vehicular environment is comparable to the simulation setting. This suggests that performance prediction in this environment is feasible. The average-case accuracy would be high, though the worst-case accuracy may be poor. The graph also shows that predictions are more accurate for locations with very poor or very good connectivity on average.

7. SUMMARY

Our work uncovers a complex picture of WiFi-based connectivity between basestations and moving vehicles. The period between the vehicle coming in and going out of range of a BS is often marred by intermittent "gray" periods of very poor connectivity. Gray periods are hard to predict because their arrival cannot be reliably signaled using current measurements and because they do not consistently occur at the same location.

Our analysis suggests that gray periods are caused by variability in the urban radio environment combined with the vehicle traversing regions that are poorly covered by the basestation. This means that gray periods are likely to be part and parcel of WiFi-based vehicular access. Because its difficult to provide blanket good coverage in large, outdoor spaces, clients are bound to encounter regions of poor connectivity. Gray periods can pose a significant challenge to providing uninterrupted connectivity to applications.

We also show that history of past performance is useful in vehicular settings. The knowledge of performance experienced at a location during past traversals can be used to identify both regions where gray periods are more likely to occur and those where the vehicle is likely to experience good performance.

Acknowledgments Victoria Poncini (Microsoft IT) helped with basestation deployment. Several employees of MV Transportation Inc. helped with deployment on the vans. Our shepherd, Dina Papagianaki, helped improve the presentation in this paper. We thank them all.

8. REFERENCES

[1] 802.11n report. http://grouper.ieee.org/groups/802/11/Reports/tgn_update.htm.

[2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *SIGCOMM*, Aug. 2004.

[3] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden. A measurement study of vehicular Internet access using unplanned 802.11 networks. In *MobiCom*, Sept. 2006.

[4] Y. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the puzzle of enterprise 802.11 analysis. In *SIGCOMM*, Aug. 2006.

[5] UMass DieselNet. http://prisms.cs.umass.edu/dome/.

[6] What is EVDO. http://www.evdoinfo.com/content/view/37/61/.

[7] R. Gass, J. Scott, and C. Diot. Measurements of in-motion 802.11 networking. In *WMSCA workshop*, Apr. 2006.

[8] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding congestion in IEEE 802.11b wireless networks. In *IMC*, Oct. 2005.

[9] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Analyzing the MAC-level behavior of wireless networks in the wild. In *SIGCOMM*, Aug. 2006.

[10] J. Ott and D. Kutscher. Drive-thru Internet: IEEE 802.11b for automobile users. In *INFOCOM*, Mar. 2004.

[11] K. Papagiannaki, M. Yarvis, and W. S. Conner. Experimental characterization of home wireless networks and design implications. In *INFOCOM*, Apr. 2006.

[12] J. Robinson, D. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy. Experimenting with a multi-radio mesh networking testbed. In *WiNMee workshop*, Apr. 2005.

[13] P. Rodriguez, R. Chakravorty, I. Pratt, and S. Banerjee. MARS: A commuter router infrastructure for the mobile Internet. In *MobiSys*, June 2004.

[14] F. Smith, J. Aikat, J. Kapur, and K. Je. Variability in TCP round-trip times. In *IMC*, Oct. 2003.

[15] City-wide Wi-Fi rolls out in UK. http://news.bbc.co.uk/2/hi/technology/4578114.stm.

[16] Cities unleash free Wi-Fi. http://www.wired.com/gadgets/wireless/news/2005/10/68999.