Measurement-based Characterization of 802.11 in a Hotspot Setting

Maya Rodrig Charles Reis Ratul Mahajan David Wetherall John Zahorjan

University of Washington

{rodrig,creis,ratul,djw,zahorjan}@cs.washington.edu

Abstract – We analyze wireless measurements taken during the SIGCOMM 2004 conference to understand how well 802.11 operates in real deployments. We find that the overhead of 802.11 is high, with only 40% of the transmission time spent in sending original data. Most of the remaining time is consumed by retransmissions due to packet losses that are caused by both contention and transmission errors. Our analysis also shows that wireless nodes adapt their transmission rates with an extremely high frequency. We comment on the difficulties and opportunities of working with wireless traces, rather than the wired traces of wireless activity that are presently more common.

Categories and Subject Descriptors:
C.2.m [Computer-Communication Networks]: Miscellaneous

General Terms: Measurement

Keywords: Wireless networks, 802.11, measurement

1. INTRODUCTION

WiFi networks based on the 802.11a/b/g are rapidly becoming pervasive. This makes it important to understand how well the basic wireless protocols work in real settings that are replete with complexities such as asymmetric connectivity and client diversity. Surprisingly, there is very little information available to make such determinations. The vast majority of earlier trace studies focus on the characteristics of user sessions, such as duration and mobility, rather than on the workings of 802.11. Moreover, these studies are overwhelmingly based on traces taken on wired segments adjacent to APs and periodic SNMP queries of AP MIBs. They do not record what was observed "in the air", and necessarily omit (or gather at a coarse granularity) key 802.11 PHY and MAC information , *e.g.*, transmission rates, signal strengths, and 802.11 retransmissions, as well as control and management traffic.

In this paper, we report on a preliminary study of wireless traces gathered during the SIGCOMM 2004 conference. We captured a comprehensive record of network activity, totaling 70 GB of full packet wireless traces, divided over five days, five locations, and three channels. Each channel was recorded at all locations to provide a view of spatial diversity, and our traces include PHY and MAC information. We also captured traffic from the wired segment between the APs and the Internet. We are in the process of anonymizing our traces to make them publicly available.

Our goal is to understand how well the 802.11 protocols operate in this setting. An analysis of a subset of our traces shows that:

1. The overhead of 802.11 is high. Only 40% of the overall transmission time is spent sending original data packets. Most of the remaining transmission time is spent on retransmissions (35%), acknowledgements (15%), and management traffic (10%).

2. Retransmissions are common, accounting for 28% of all data transmissions and 46% of data transmission time. We find evidence that these occur due to losses caused by both contention (competing transmissions) and wireless transmission errors (signal strength).

3. Switches in client transmission rates are the common case rather than the exception: in most cases only one or two frames are sent between rate switches. Most of the transmission time is spent sending bits at 1 Mbps, the lowest rate. Contention losses have interesting consequences for 802.11 transmission rate adaptation: switching to a lower rate is not only unnecessary but will also make the medium busier.

The value of working with wireless traces is evident in that none of the analyses that support the findings above could have been performed with the combination of wired traces and SNMP statistics. Nonetheless, we believe our analysis to date is the "tip of the iceberg" in terms of what information can be extracted from these traces. For instance, in the near future, we intend to investigate spatial diversity and reuse.

A key difficulty in working with wireless traces is that they are inherently incomplete and provide a view that differs from that of the clients and APs. Nonetheless, we are able to draw useful inferences from them, e.g., estimating the utilization at an AP from the packets seen by our nearby monitor. We believe that there is significant potential for a richer set of wireless inference techniques.

In the following sections, we describe our trace environment and then present initial analyses of the utilization, overhead, retransmissions and rate adaptation. We then contrast our work with related efforts and conclude with a discussion of our experiences working with wireless traces.

2. TRACE ENVIRONMENT

Our trace environment is the open wireless network provided to roughly 550 participants who attended the SIGCOMM conference in Portland, Oregon from 8/30/04 to 9/03/04. We view this as characteristic of a large and busy hotspot setting. The conference took place in a hotel with a layout depicted in Figure 1. A wireless network comprised of 5 APs was set up for the conference, operating

SIGCOMM'05 Workshops, August 22–26, 2005, Philadelphia, PA, USA.

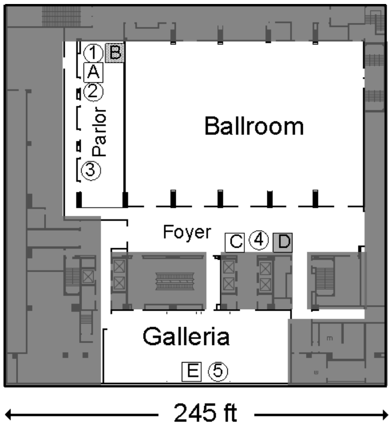

Figure 1: The wireless network and monitor machines at the conference. Ballroom hosted the main conference and had only limited wireless access on the back and left. Parlor acted as the terminal room and was most active. Galleria hosted workshops and poster sessions. The boxes represent the rough locations of the five APs and the circles represent the rough locations of the five wireless monitors. Of the five APs, A, C, and E operated on channel 1, B on channel 11 and D on channel 8. All of our monitors listened passively on all three channels.

in 802.11b mode[1] on channels 1, 8 and 11 (even though channels 8 and 11 are not orthogonal). Internet connectivity was provided by four separate DSL access lines.

The analysis presented in this paper is based on only a small portion of the overall trace: the headers of correctly received 802.11 frames gathered by Monitor 1 on Channel 1 for the active part of Tuesday, August 31, 2004.

2.1 Trace Collection

We monitored the activity on the wireless network in two ways. The first method was wireless monitors. We placed five PCs, each with three Netgear WAG311 wireless adapters, near the APs. Each monitor listened on all three active channels, using external antennas which we placed at least a foot apart from each other to avoid interference.[2] Multiple monitors were within the range of each AP to provide information on the spatial diversity of transmission and reception.

We logged all observed activity using *tcpdump*, capturing 802.11 control and management packets as well as data packets and writing complete packets to disk. The PHY information includes a receive signal strength indicator (RSSI) and the transmission rate. The MAC information is the entire 802.11 frame header and CRC. We also customized our MADWiFi driver to log reception errors that provide information on periods when transmissions were detected but not correctly decoded.

The second monitoring method was traditional wired traces, collected using tcpdump on the network segment connected to the APs. This provides a view of packets exchanged between the APs and the Internet.

[1] Some of the APs were capable of operating in 802.11g mode, but we primarily observed 802.11b rates and saw only 1, 2, 5.5 and 11 Mbps rates listed on beacon frames.

[2] We took care here as researchers have reported interference across orthogonal channels due to hardware implementation strategies [8].

Total Frames	12 million
Total Bits	2.0 GB
Total Airtime	2.4 hours (20.9% of period)
Number of distinct clients	377
Period of trace	11.5 hours

Table 1: Summary of trace on Monitor 1, Channel 1, for the active part (8:00am–7:30pm) of Aug. 31, 2004.

2.2 Trace Limitations

One limitation that was beyond our control is that the SIGCOMM network was hampered by intermittent usability problems. We understand that these problems stemmed from the DHCP and DNS configurations and caused Internet connectivity to become unavailable to some clients. For this reason, we do not focus on characteristics of client and user sessions, as they may be artificial. We believe that these problems do not affect 802.11 behavior, *e.g.*, the use of various transmission rates and signal strengths, other than lowering the load on the network. Thus, expect a pool of users of comparable size to place slightly greater demands on a smoothly operating wireless hotspot.

A second limitation is that the different APs were assigned to logically different networks (with different SSIDs), such that users or their operating systems selected one AP for connectivity. Thus we cannot study client strategies for switching to the best AP within a network.

3. UTILIZATION

We begin by analyzing how heavily the wireless network is used. Table 1 gives summary statistics for the portion of the trace we focus on in this paper. In the table, Frames counts correctly decoded 802.11 transmissions of all kinds, Bits includes 802.11 MAC headers and higher layers, and Airtime covers PHY transmissions (preamble and PLCP) as well as Bits. The trace records usage by the majority of the conference participants, judging by the number of distinct clients. It captures 1.6 GB of downstream traffic (from the APs to clients) and 0.4 GB of upstream traffic. This corresponds to a long-term average data rate over the wireless link of roughly 380 Kbps. There was also a very small amount (25 MB) of non-SIGCOMM wireless traffic.

A key problem in working with wireless traces is that they are inherently incomplete, missing packets that were sent by clients or APs but not correctly received by our monitor, even though it is adjacent to the AP. In fact, our monitor generated 3.9 million reception error events compared to 12 million correctly decoded frames. To account for this potential loss of information, we developed simple techniques based on 802.11 protocol behavior to analyze our trace for monitor loss. The 802.11 MAC sequence numbers on packet sent from the AP are normally consecutive. By observing gaps in these sequence numbers, we estimate that we observe 96% of AP to client transmissions. Similarly, (non-broadcast) packets sent to the AP by clients are ACKed by the AP. By matching ACKs to transmissions, we estimate that we observe 75% of the client transmissions out of those that were correctly decoded by the AP. This implies that the statistics above already represent roughly 90% of the transmitted Frames and Bits.

To better understand network load, we show how it varies over time in Figure 2. Here, we selected Airtime as the most suitable measure for conveying load. This is because Airtime allows transmissions at different rates to be combined in a meaningful way (including low rate headers and high rate data), whereas Frames and Bits do not. (We will see considerable rate variation in later sec-

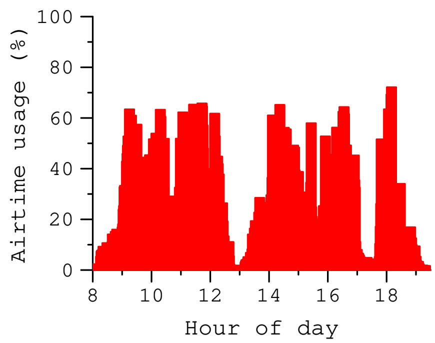

Figure 2: Airtime utilization over time. The binning interval is one minute.

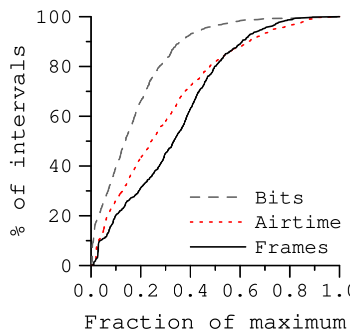

Figure 3: Cumulative distribution of activity per minute.

tions.) Note that this graph gives a lower bound on the time that the medium around the monitor was busy. It does not include frames lost by our monitor and omits some small PHY layer overheads, such as forced delay due to the contention window, which we could not readily consider.

We see both heavily and lightly loaded periods, even with a one minute averaging interval, and significant variability throughout the day. This is consistent with the results of Yeo *et al.* [11], which show high variability in activity in an academic department's wireless network. The patterns evident in usage tend to correspond to events in the conference program. For example, there is a noticeable drop over lunch, though activity remained high during the keynote and talks despite the lack of APs in the ballroom where the talks were given. We also observed high night-time activity (until around 1 am) on some nights.

To see the differences between Frames, Bits and Airtime as measures of utilization, we computed their cumulative distributions. This is shown in Figure 3. To fit all three measures on one scale, the x-axis values have been normalized to the fraction of the maximum activity observed in an interval. In all cases a small portion of the trace time contains the top third of the loads. Bits is the most skewed measure, with half of the trace minutes having low loads of around 10% or less of the maximum, and less than 10% of the trace minutes having loads above 40% of the maximum. Frames is a much flatter distribution, with more than 90% of the minutes roughly uniformly spread up to 60% of the maximum load. Airtime falls between the other measures. These results suggest that Frames, Bits and Airtime are not interchangeable measures of load even when averaged over small intervals such as a minute.

Frame type and subtype	**Airtime** (secs)	**Bits** (MB)	**Frames** (1000s)	**Avg. Rate** (Mbps)
Data	6802	1884	5540	6.46
Originals	3616	1276	3988	7.30
Retransmits	3185	608	1552	4.31
Control	1418	74	5442	1.89
Ack.	1332	69	5135	1.90
RTS	42	3	142	1.69
CTS	40	2	155	1.75
PS poll	2	0	10	1.60
Management	878	82	1098	1.12
Assoc. Req.	1	0	2	1.42
Assoc. Res.	1	0	3	1.08
Authentication	6	0	13	1.13
Beacon frame	412	39	428	1.00
Deauth.	0	0	0	1.30
Dissassoc.	6	0.40	13794	1.00
Probe Req.	177	16.07	333707	1.35
Probe Res.	270	25.44	296250	1.00
Reassoc. Req.	0	0.03	2727	1.00
Reassoc. Res.	0	0.03	621	1.00
Totals	9098	2040	12080	3.92

Table 2: Breakdown by frame type and subtype. (Originals and Retransmits are not 802.11 frame subtypes; we list them here for ease of exposition.)

4. OVERHEADS

We now consider the various overheads involved in data transmission that reduce its effectiveness compared to an ideal setting. These include management and control frames, 802.11 retransmissions, and PHY and MAC headers. Table 2 presents a breakdown of transmissions by frame type and subtype. For each, it shows the three usage measures and the average transmission rate.

We see that the vast majority of Bits (92%) are for data frames. However, management frames (mostly beacons and probes) make up 10% of the frames, and acknowledgements make up almost half the frames (as they are roughly one-for-one with data frames). Moreover, these frames are transmitted at a lower average rate than data frames. Combined with PHY+PLCP headers, which are sent at 1 Mbps, this means that they occupy more of the medium than suggested by their Frame counts – data frames obtain only 75% of the Airtime.

Looking at the data frame category, it is surprising to see that 28% of them are retransmissions (i.e., data frames with Retry bit turned on). The impact of these retransmissions is heightened because they occur at a lower average rate than original frames. Overall, only slightly over half (53%) of the data frame Airtime is used by original transmissions. This leads us to investigate retransmissions and rate adaptation in the next sections.

As a final source of overhead, we note that a further 21% of the original transmission Airtime is consumed by 802.11 PHY and MAC overheads. The cumulative effect of control and management traffic, retransmissions, and PHY and MAC overhead is that only 31% of the Airtime is being used to transfer original data (IP or higher layer). We were somewhat surprised at this relatively low overall efficiency as we had expected large data frames to amortize the overheads of short frames.

5. RETRANSMISSIONS

In this section, we explore the retransmission behavior observed in our traces. We also investigate whether signal strength and contention correlate closely with high retransmission rates. We leave investigating other potential causes of retransmissions, such as packet size, and quantifying their relative impact for future work.

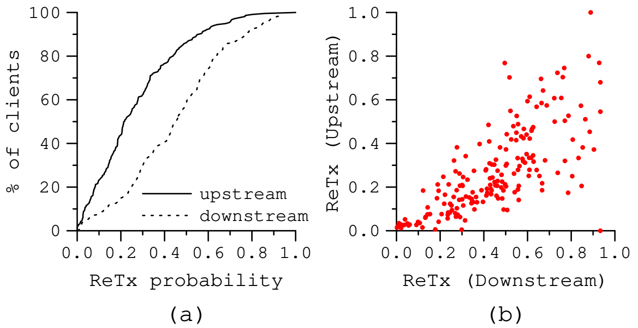

Figure 4: (a) Retransmission probability (ReTx) as a cumulative fraction of clients. (b) Per-client upstream and downstream ReTx.

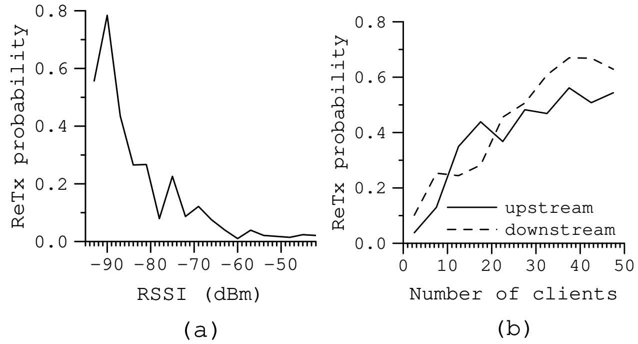

Figure 5: (a) Retransmission probability as a function of RSSI (upstream traffic only), and (b) number of clients.

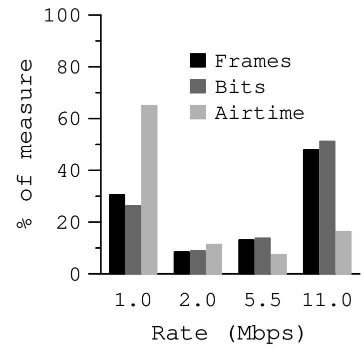

Figure 6: Relative prevalence of transmission rates.

The analysis in this section is based on the *Retry* bit in the 802.11 header. This bit is turned on when a node retransmits a packet because the acknowledgement for the previous transmission was not received. It enables us to distinguish the original transmission from the subsequent ones (but we cannot distinguish among different retries). We quantify retransmissions using *retransmission probability*, defined as the ratio of retransmitted data frames with to the total number of data frames. This is a measure of the quality of the link between the sender and receiver; higher probability implies that more transmissions or their acknowledgements are lost.

In Figure 4a, we show the distribution of retransmission probabilities for upstream (from clients to the AP) and downstream (from the AP to the clients) traffic. The results indicate that retransmission is much more likely in the downstream than in the upstream direction. They also show that there is significant variation in retransmission probabilities across clients, in both directions. Figure 4b correlates the upstream and downstream retransmission probabilities for all clients. While there is a general trend along downstream retransmission probability being twice upstream, there are many outliers. For all but a very few clients, though, downstream probabilities are higher than upstream.

We now study how retransmission probability varies with signal strength and contention level. The analysis below assumes that these two factors are independent of each other.

Signal strength The first factor we study is the strength of the signal from the client to the AP (RSSI). Since we cannot measure the signal strength at the AP itself, we assume that the relative signal strengths of different clients measured at a monitor near the AP are reasonable approximations to the relative strengths seen by the AP itself. Because we have no way of approximating the RSSI values seen by the clients, we do not consider downstream traffic in this analysis.

Figure 5a shows retransmission probability as a function of signal strength. Not surprisingly, we see a strong correlation between the two. This implies that the signal strength has an important factor influence on retransmission probability, even if it is not an accurate predictor on its own [1].

Contention level The second factor we study is the effect of contention in the network, i.e., the number of nodes in the network with data to send. Because we cannot determine this measure from our traces, we approximate this using the number of clients active in a short time interval. A client is considered active in a given interval if we see at least one packet from it in that interval.

Figure 5b shows the variation of the retransmission probability with the number of active clients. The interval size in this analysis is one minute. The retransmission probability increases with the number of active clients in the interval.

That the retransmissions increase with increased contention has consequences for rate adaptation. Most adaptation algorithms reduce their transmission rate in the face of losses. But if many losses are caused by contention, rate reduction is unlikely to help. In fact, rate reduction is exactly the wrong thing to do as it increases contention by occupying the media for a longer time. For this reason, rate adaptation algorithms should either be driven by throughput [4] or try to distinguish between the various causes that lead to loss.

6. TRANSMISSION RATE ADAPTATION

Little is known about transmission rate adaptation in current hotspot environments. In this section we use our trace to investigate rate adaptation in such settings.

6.1 Summary View

We first investigate the use of different transmission rates in aggregate across all clients. Figure 6 shows the percentage of Frames and Bits transmitted at each rate, along with the percentage of Airtime utilized by that rate. The greatest fraction of frames (around 50%) are sent at the highest 802.11b rate of 11 Mbps. This is because most rate adaptation algorithms have a strong preference toward this rate. For instance, we see clients that always try to transmit a new packet at 11 Mbps irrespective of the rate at which the last transmission succeeded; such clients reduce their rate only when one or a few consecutive transmissions at 11 Mbps fails. The figure also shows that in contrast to Frames and Bits, most of the Airtime is utilized by 1 Mbps data. This is a direct consequence of 1 Mbps frames taking a lot longer than other frames. Hence, while

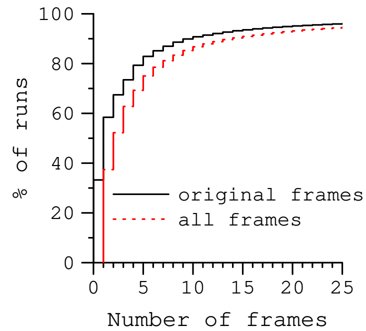

Figure 7: The number of frames clients send consecutively before switching transmission rate. The *original frames* curve excludes retransmitted frames.

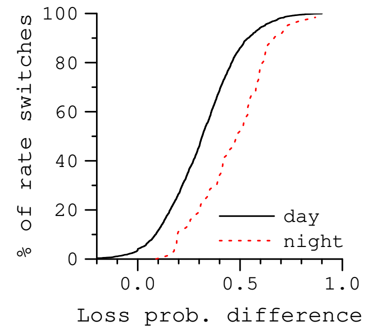

Figure 8: The difference is the loss probability at the higher rate minus that at the lower rate for the two transmission rates comprising a rate switch.

Distinct rates used	% of clients
1	8%
2	10%
3	18%
4	62%

Table 3: Percentage of clients that use a particular number of distinct transmission rates for data packets.

	1	2	5.5	11
1	0.90	0.04	0.02	0.03
2	0.12	0.77	0.06	0.05
5.5	0.03	0.04	0.84	0.09
11	0.01	0.01	0.02	0.96

Table 4: The rate transition state machine.

clients are more likely to use 11 Mbps, at any given instance the medium is more likely to be carrying a 1 Mbps frame.

We next consider whether the diversity of rates seen in aggregate results from different clients using different rates, or from each client operating at a number of rates. Table 3 shows the percentage of clients that use a given number of unique transmission rates in our trace. It considers only data packets and excludes clients that send fewer than 25 packets so that clients active for only a short period do not bias our results. We see that individual clients commonly use multiple rates: fewer than one in ten clients limit themselves to one transmission rate (which usually is 1 Mbps), while more than 60% of them use all four available rates.

6.2 Dynamics

We explore the dynamics of rate adaptation by studying how frequently nodes switch their rates and what switches are more frequent. We consider that a client has switched its rate when it sends a frame at a different rate within one second of sending the last frame. This is done to exclude spurious rate switches, those that result not from rate adaptation but from the client going temporarily idle; clients often start with a pre-determined rate after an idle period.

Figure 7 shows the distribution of the number of frames that clients send at their current rate before switching to a different rate. Surprisingly, clients change their transmission rates very frequently. Half of the time clients send only one or two frames before switching again. In the future, we will investigate whether such frequent switching hurts application performance.

To study which rate switches are more common, we model the rate adaptation of clients as a state machine in which each state corresponds to a transmission rate. We then assign probabilities to state transitions based on the behavior observed in our traces. These probabilities, which depend on both the rate adaptation algorithms of our clients and our wireless environment, are shown in Table 4. An entry (x, y) denotes the probability of moving from state x to y, that is, the likelihood of using rate y given that the last packet was sent at rate x. Several interesting inferences can be made from the table. For instance, continuation of the existing rate is most likely for 11 Mbps and least likely for 2 Mbps. When decreasing rate from 11 Mbps, clients are twice as likely to move to 5.5 Mbps than 1 or 2 Mbps. But when doing so from 5.5 Mbps, clients are almost equally likely to move to 1 or 2 Mbps. Presently, we are working on using such state machines to reverse engineer the exact rate adaptation algorithms implemented by various clients.

6.3 Effectiveness

In this section we present a preliminary analysis to gain some insight into the efficacy of current rate adaptation algorithms. Our analysis considers downstream data traffic only. Analyzing upstream traffic for this purpose cannot be done reliably: because our monitors are less sensitive than the AP, if the AP misses an upstream frame our monitors will likely miss it as well, which biases our upstream sample toward successful receptions. In contrast, our monitors capture nearly all downstream traffic, and so retransmissions are reliable indicators of previously failed attempts.

We assess the effectiveness of rate switching using the difference in *loss probability* before and after the switch. Loss probability at a given rate is the fraction of packets transmitted at that rate for which we see a subsequent retransmission at any rate.[3] We use the sequence number in the 802.11 header to identify multiple instances of the same packet. Thus, if k instances of a packet are observed, we consider the first $k-1$ as lost and the last as successfully received. This ignores the possibility that the source simply gave up after too many retries. The retransmission probabilities observed in Section 5 suggest that this rarely occurs.

Figure 8 plots the distribution of the change in loss probability seen before and after a rate switch. This change is always computed as the the loss probability at the higher rate minus that at the lower

[3]The retransmissions are often at a lower rate because of the rate adaptation algorithms used. An informal analysis of our logs indicates that the AP tends to transmit a single packet at 11, 11, 5.5, 5.5, 2, 1, 1, ..., 1 Mbps.

rate, irrespective of the direction of the switch. Each sample corresponds to a session, switch type (e.g., $5.5 \rightarrow 11$, $11 \rightarrow 5.5$, $1 \rightarrow 2$) pair. A session is a sequence of downstream frames such that the gap between each two consecutive frames is less than a minute. The average difference across all instances of a switch type in a session is plotted, excluding switch types observed fewer than 10 times. *Day* and *night* correspond to busy (800–1930 hours) and mostly idle (1930–800 hours) periods. While we plot the aggregate distribution for all switch types, the distributions for individual switch types are similar.

The graph shows that in general lower transmission rates have lower loss probability. But in many cases the loss probability difference is so minor that moving to the lower rate may not be justified from a throughput perspective. For instance, during the day, the loss probability difference is less than 0.2 for 25% of the cases. Similar observations were made in a measurement study of an outdoor 802.11 mesh network [1].

Further, the difference in the night and day curves suggests that rate adaptation is less helpful when there is contention in the network. This further supports our earlier inference (Section 5) that moving to a lower rate in the face of lost packets without understanding the cause behind the loss is not a good strategy.[4] In the future, we will investigate rate adaptation strategies that distinguish between the causes behind lost packets.

7. RELATED WORK

Our study is uncommon in that it analyzes wireless traces from a production 802.11 network. Except for a concurrent publication [7], the only other work that analyzes wireless traces from a production network is that of Yeo *et al.* [11]. They advocate monitoring wireless networks for security reasons and explore how traces from different monitors can be merged to obtain a more complete view. Toward this goal, they develop and test a trace merging methodology. In contrast, our focus is to use such traces to further our understanding of how well 802.11 hotspot networks function in practice.

All of the other work that uses wireless traces of which we are aware has the flavor of experiments over a controlled rather than a production network, *e.g.*, characterization of wireless losses in Roofnet [1]. The two methods have different strengths. While controlled experiments can enable a deeper investigation into some of the observed phenomena, analyzing a production network provides valuable insight into the operation of real networks.

Many studies of client behavior on wireless networks have emerged over the past five years (e.g., [10, 2, 6, 3, 9, 5]). These complement our study by using a mix of SNMP and wired network traces to analyze user and application behavior. However, this methodology cannot be used to infer the properties of the wireless media itself to make inferences such as those we make in this paper.

8. DISCUSSION AND FUTURE WORK

We conclude with a discussion of our experience working with wireless traces; our results are summarized at the start of this paper.

Traces gathered by sniffing the wireless medium at one or more points present both new opportunities and challenges compared to wired traces and SNMP logs. Of course, their key advantage is obtaining the detailed PHY and MAC information that would otherwise be lost and the new kinds of studies that this enables. For instance, none of the analyses in this paper could have been completed with wired traces and SNMP statistics.

The key disadvantage of wireless traces is that they are incomplete. The monitors log only correctly decoded frames or some error frames that were close to being decoded. This omits periods of energy that correspond to lost frames. More subtly, even "complete" wireless traces at a monitor cannot directly record whether a particular node successfully received specific frames; location and hardware capabilities are first-class features of wireless media.

We expect methods for processing such traces to become more sophisticated to mitigate these disadvantages. Yeo *et al.* take a step in this direction by merging traces from several monitors using beacons to provide a more complete view of activity in a region [11]. Our approach is different. We found that the traces usually contain enough hints (e.g., 802.11 sequence numbers, Retry bits, and Data-Ack pairs) to allow us to estimate or quantify the information that is missing. In this manner, many kinds of interesting inferences can be drawn, though they may be statistical and depend on various independence assumptions.

Finally, we believe that we have just scratched the surface when it comes to extracting information from passively collected wireless traces. Better inference procedures are a prime area for future work that may provide deeper insight into the effectiveness of 802.11. The major areas we hope to explore include: understanding the spatial diversity by comparing and merging across monitors; using the error indications to see how transmission errors and collisions affect performance; and matching Data-Ack pairs to detect hidden terminals. In the long term, we hope these techniques and insights can be leveraged by wireless nodes to improve the quality of 802.11 networks as deployed and used in practice.

9. ACKNOWLEDGEMENTS

We thank Donald Newell for assistance with setting up the monitoring infrastructure and Ed Lazowska for useful discussions throughout this work. This work was supported in part by Microsoft Research and the NSF (Grant No. ANI-0133495).

10. REFERENCES

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

[2] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan. Characterizing user behavior and network performance in a public wireless lan. In *SIGMETRICS*, 2002.

[3] M. Balazinska and P. Castro. Characterizing mobility and network usage in a corporate wireless local-area network. In *MobiSys*, 2003.

[4] J. Bicket. Bit-rate selection in wireless networks. Master's thesis, MIT, 2005.

[5] F. Chinchilla, M. Lindsey, and M. Papadopouli. Analysis of wireless information locality and association patterns in a campus. In *INFOCOM*, 2004.

[6] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *MobiCom*, 2004.

[7] A. Jardosh, K. Ramachandran, K. Almeroth, and E. Belding-Royer. Understanding link-layer behavior in highly congested IEEE 802.11b wireless networks. In *workshop on Experimental Approaches to Wireless Network Design and Analysis (E-WIND)*, 2005.

[8] J. Robinson, D. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy. Experimenting with a multi-radio mesh networking testbed. In *1st workshop on Wireless Network Measurements (WiNMee)*, 2005.

[9] D. Schwab and R. Bunt. Characterising the use of a campus wireless network. In *INFOCOM*, 2004.

[10] D. Tang and M. Baker. Analysis of a local-area wireless network. In *MobiCom*, 2000.

[11] J. Yeo, M. Youssef, and A. Agrawala. A framework for wireless lan monitoring and its applications. In *ACM workshop on Wireless Security (WiSe)*, 2004.

[4]This problem is similar to TCP reducing its congestion window in response to all types of losses, which leads to unnecessary reduction in throughput when losses are not related to congestion.