

Link-level Measurements from an 802.11b Mesh Network

Daniel Aguayo John Bicket Sanjit Biswas Glenn Judd[†] Robert Morris

M.I.T. Computer Science and Artificial Intelligence Laboratory
 {aguayo, jbicket, biswas, rtm}@csail.mit.edu

[†] Carnegie Mellon University
 glennj@cs.cmu.edu

ABSTRACT

This paper analyzes the causes of packet loss in a 38-node urban multi-hop 802.11b network. The patterns and causes of loss are important in the design of routing and error-correction protocols, as well as in network planning.

The paper makes the following observations. The distribution of inter-node loss rates is relatively uniform over the whole range of loss rates; there is no clear threshold separating “in range” and “out of range.” Most links have relatively stable loss rates from one second to the next, though a small minority have very bursty losses at that time scale. Signal-to-noise ratio and distance have little predictive value for loss rate. The large number of links with intermediate loss rates is probably due to multi-path fading rather than attenuation or interference.

The phenomena discussed here are all well-known. The contributions of this paper are an understanding of their relative importance, of how they interact, and of the implications for MAC and routing protocol design.

Categories and Subject Descriptors

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

General Terms

Measurement, Performance

Keywords

wireless, mesh, 802.11b

1. Introduction

This paper is a measurement study of the Roofnet multi-hop wireless network. Roofnet nodes are computers with

802.11b cards in apartments spread over six square kilometers of Cambridge, Massachusetts. Each node has a roof-mounted omni-directional antenna. The network’s main purpose is to provide Internet access via a few wired gateways. The initial implementation strategy was to combine existing radio, MAC, and routing technology in order to build a production-quality network as quickly as possible. This approach led to performance far less than expected, primarily due to assumptions made by MAC and routing protocols that were a poor fit to the network’s actual behavior. It is widely understood that wireless differs from simple abstract models in a number of ways [10]; the goal of this paper is to provide insight into which differences are important enough to worry about, and to draw conclusions relevant to the design of future MAC and routing protocols.

Many routing and link-layer protocols assume the validity of a “neighbor” abstraction that partitions all the pairs of nodes into pairs that can communicate directly, and pairs that cannot. This assumption justifies the use of graph-theoretic routing algorithms borrowed from wired networks, where the assumption is true. It leads to the design of MAC protocols such as 802.11 that assume that a pair of nodes will either hear each other’s control packets (e.g. RTS/CTS), or will not interfere. It justifies conservative transmit bit-rate selection algorithms that reduce the bit-rate after a few packet losses. Many existing protocols might have to be re-designed if the neighbor abstraction turned out to be a poor approximation of reality.

In principle the neighbor abstraction is supported by typical assumptions about the relationship between signal-to-noise ratio and bit error rate (S/N and BER). This relationship is typically assumed to have a rapid transition from essentially zero BER to a BER high enough to corrupt every packet. For example, the transition zone for the Intersil Prism HFA3873 baseband processor is about 3 dB, regardless of bit rate [1]. Since signal strength falls off rapidly with distance, one might expect relatively few node pairs to lie in the transition zone. As a result, one might expect almost all pairs of nodes to either be able to talk to each other with low loss, or not at all. Some empirical 802.11 measurements suggest that the neighbor abstraction usually holds [7, 10], while others do not [6, 11].

This paper starts with the observation that most Roofnet node pairs that can communicate at all have intermediate loss rates; that is, the neighbor abstraction is a poor approx-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM’04, Aug. 30–Sept. 3, 2004, Portland, Oregon, USA.
 Copyright 2004 ACM 1-58113-862-8/04/0008 ...\$5.00.

Links with intermediate loss rates are common, with no sharp transition between high and low packet loss rates.	Sec. 3
Inter-node distance is not strongly correlated with whether nodes can communicate.	Sec. 4
Most links have non-bursty loss patterns.	Sec. 5
Links with very high signal strengths are likely to have low loss rates, but in general signal strength has little predictive value.	Sec. 6
A link is likely to have a significant loss rate at its optimum 802.11b bit-rate.	Sec. 7
Multi-path fading greatly affects outdoor links and helps explain intermediate loss rates.	Sec. 9

Figure 1: Summary of major conclusions for wireless MAC and routing protocol design.



Figure 2: A map of Roofnet, with a black dot for each of the 38 nodes that participated in the experiments presented in this paper.

imation of reality. The remainder of the paper explores a series of hypotheses for the causes of packet loss in Roofnet, and for the predominance of intermediate loss rates. The hypotheses include factors that affect signal-to-noise ratio (distance and interference), choice of transmit bit rate, and multi-path fading. Figure 1 lists the paper’s main conclusions about these sources of packet loss. The conclusions in this paper should not be viewed as universal, since they are limited by the particulars of Roofnet’s configuration.

2. Experimental Methodology

Roofnet consists of 38 nodes distributed over roughly six square kilometers of Cambridge. Each consists of a PC with an 802.11b card connected to an omni-directional antenna mounted on the roof. Figure 2 shows a map of the network.

The area is dominated by tightly-packed three- and four-story houses; most antennas are mounted about two or three feet above the chimneys of these houses. There are also a number of taller buildings in the area; seven Roofnet nodes are located in such buildings. Not all nodes have roof-

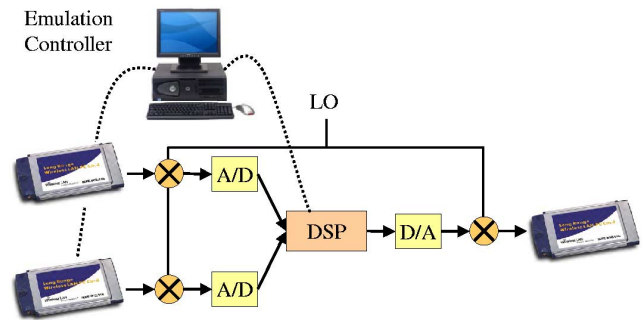


Figure 3: Architecture of the hardware channel emulator.

mounted antennas: a handful of users found it easier to place the antenna in or hanging outside of a window.

All nodes use identical 802.11b cards based on the Intersil Prism 2.5 chip-set. Except as noted, the cards transmit at 2.422 GHz (802.11b channel 3) with the transmission power level set to +23 dBm (200 mW). The omni-directional antennas provide 8 dBi of gain with a 20-degree -3 dB vertical beam-width. Cabling and lightning arrestors introduce an attenuation of 6 to 10 dB depending on the length of cable.

The cards can be configured to transmit at 1, 2, 5.5, or 11 Mbit/s; the experiments in this paper run with automatic bit-rate selection disabled. The cards operate in the Prism 2.5 “pseudo-IBSS” mode, which is a simplified version of the 802.11b IBSS (ad hoc) mode; use of pseudo-IBSS circumvents firmware bugs in the IBSS implementation that can cause network partition.

Nodes are located at the apartments of volunteers, who were selected with no special plan beyond basic radio connectivity. The experiments were run with Roofnet routing turned off, and thus with no Roofnet user traffic. All the experiments were executed in the early hours of the morning, so the paper’s results may underestimate the effects of non-Roofnet radio activity.

Most of the Roofnet data presented in this paper is derived from a single experiment. In this experiment, each node in turn sends 1500-byte 802.11 broadcast packets as fast as it can, while the rest of the nodes passively listen. Each sender sends for 90 seconds at each of the 802.11b bit-rates. The experiment uses 802.11 broadcast packets because they involve no link-level acknowledgments or retransmissions.

Each packet includes a unique sequence number. The sender records the time at which it sends each packet, and all the other nodes record each received packet’s sequence number, arrival time, and the “RSSI” and “silence” values that the 802.11 card reports.

The experiment was run in the early hours of June 6, 2004. Figure 16 was derived from a similar experiment on June 1, 2004, in which different power levels were tested. The authors have examined the results from many similar experiments over a period of months and verified that they are similar to the data presented in this paper.

2.1 Channel Emulator

In addition to the Roofnet experiments, this paper presents results from a wireless channel emulator [8], to which two sender laptops and a single receiver laptop are connected.




Figure 4: The distribution of link delivery probabilities for 1500-byte broadcast packets. Each point corresponds to one sender/receiver pair at a particular bit-rate. Points were restricted to pairs that managed to deliver at least one packet during the experiment. Most pairs have intermediate delivery probabilities.

The laptops use the same wireless cards used by Roofnet. Figure 3 shows the emulator’s architecture.

The outgoing signal of each source card is first attenuated and then mixed down to baseband where it is digitized and sent to the digital signal processing (DSP) unit. The DSP then independently scales the signals from each source to emulate large scale path loss. A small number of delayed copies of a signal may also be produced and independently scaled. All signals are then summed and then converted back to analog. The resulting baseband signal is then attenuated, mixed up to RF, and fed to the receiver’s 802.11 antenna input. The attenuation and delay used by the DSP are controlled by the Emulation Control Node which also controls the transmission of traffic by the source nodes.

In the emulator experiments, the receiver node operated in monitor mode and logged the headers of all frames received. These logs were then post-processed to generate the results discussed in this paper.

2.2 Signal Strength Measurements

The Prism 2.5 chip-set provides per-frame measurements called RSSI (receive signal strength indication) and “silence value.” The RSSI reflects the total power observed by the radio hardware while receiving the frame, including signal, interference, and background noise. The silence value reflects the total power observed just before the start of the frame. We found that the accuracy of the RSSI and silence readings was within 4 dB by comparison with a spectrum analyzer. This paper reports signal-to-noise ratios derived from the RSSI and silence values.

3. Distribution of Delivery Probabilities

Figure 4 shows the distribution of inter-node packet delivery probabilities on Roofnet at different 802.11 transmit rates. The graph includes only node pairs between which at least one packet was delivered, and thus reflects different numbers of pairs for different bit rates. The data for each

bit-rate is sorted separately, so the delivery probabilities for any particular x -value are not typically from the same pair of nodes.

At 1, 2, and 5.5 Mbit/s, Figure 4 shows that the distribution of loss rates is fairly uniform: there is only a slight tendency for pairs to segregate between working and not working. At 11 Mbit/s, there is a more rapid fall-off in delivery probability, but there are still many links with intermediate probabilities.

The implication of Figure 4 is that the neighbor abstraction does not apply well to Roofnet: most node pairs that can communicate have intermediate loss rates. It would be difficult to find multi-hop routes through Roofnet that did not involve one or more hops with significant loss rates. A routing protocol cannot ignore this problem by simply ignoring all but the very best links: for example, a one-hop route with 40% loss rate has better throughput than a two-hop route with loss-free links [6].

The failure of the neighbor abstraction in some real-world wireless environments has been noted before and shown to seriously reduce the performance of multi-hop routing [11, 6, 15]. The failure is perhaps surprising given that some measurements of 802.11 and 802.11-like systems suggest that nodes that can communicate at all can usually communicate with low loss [7, 10]. The rest of this paper explores the causes and implications of the prevalence of intermediate delivery probabilities, focusing on the reasons for packet loss in Roofnet and the nature of the delivery-probability distribution in Figure 4.

4. Spatial Distribution of Loss Rates

A potential explanation for the distribution of link delivery probabilities in Figure 4 is that it is determined by attenuation due to distance. Figure 5 shows three samples of how delivery probability varies with location. Each map corresponds to a different sender; the size of each node’s disk indicates the fraction of packets that node received from the sender.

Since the three senders are close to each other, one might expect the three reception patterns to be similar. This is true to the extent that very close nodes have high delivery probabilities for all three senders. Other than that, however, the three reception patterns are quite different. The differences are likely caused by obstacles in the environment, different antenna heights, and multi-path fading, implying that up to a certain point reception is dominated by obstacles and geometry rather than by free space path loss.

Figure 6 shows the relationship between distance and delivery probability for all Roofnet node pairs, for 1 and 11 Mbit/s. Both bit-rates exhibit a cluster of short links with high delivery probabilities, a few remarkably long links, and a significant set of links with no discernible relationship between distance and delivery probability.

5. Time Variation of Loss Rate

The significance of intermediate loss rates depends on the time scale at which loss and delivery alternate. One way in which a link might exhibit a 50% loss rate would be to deliver or drop each packet in alternation. At another extreme, a 50% link might alternate 10-second periods of total loss and total delivery. Different route selection and error correction strategies are appropriate in the two different situations.


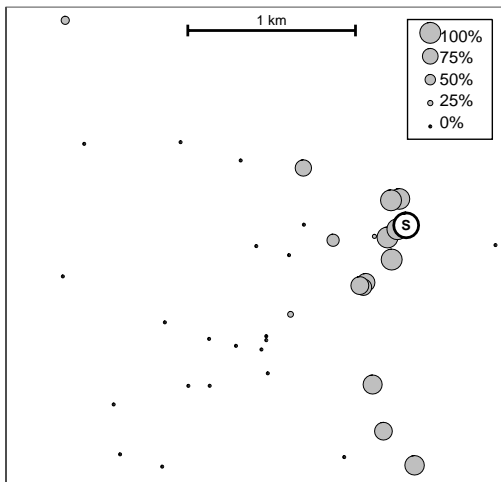



Figure 5: These maps show the delivery probabilities from three senders to all other nodes. The sender is marked S, and each receiver is indicated by a circle with radius proportional to the fraction of packets it received. There is a correlation to distance but it is not always consistent.

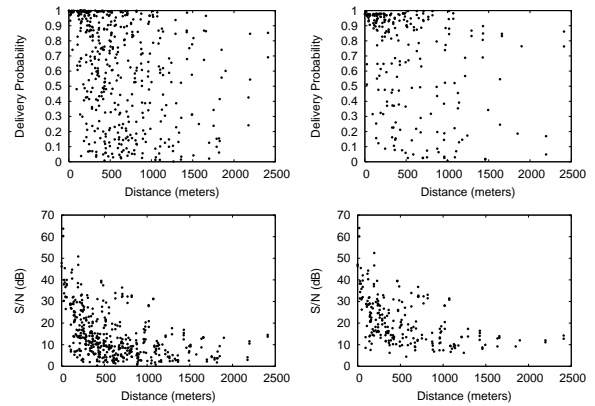


Figure 6: Scatter plots showing the relationship of distance versus delivery probability (top) and distance versus S/N (bottom). The left two graphs are for 1 Mbit/s, the right two graphs are for 11 Mbit/s. There is one point per pair of Roofnet nodes that are able to communicate at the given bit-rate.




Figure 7: Delivery probability over time (in seconds) for four 1 Mbit/s links, all with about 50% average loss rate. The send rate is about 80 1500-byte packets/second. Each point is an average over 200 milliseconds. The top graph shows one of Roofnet's most bursty links, the bottom one of the least bursty.




Figure 8: Delivery probability over time (in seconds) for four 11 Mbit/s links, all with about 50% average loss rate. The send rate is about 600 1500-byte packets/second. Each point is an average over 200 milliseconds. The top graph shows one of Roofnet’s most bursty links, the bottom one of the least bursty.

Figure 7 shows delivery probability over time for four Roofnet links running at 1 Mbit/s with 1500-byte packets. The lines indicate averages over successive 200-millisecond intervals. The four links are chosen from among the set of links with delivery probabilities near 50%: the top graph shows the link from that set with the highest short-term variation in delivery probability, the bottom graph shows the link with the lowest variation, and the other two show representative links with intermediate amounts of variation. Figure 8 shows similar data for 11 Mbit/s. These graphs suggest that there is considerable difference from link to link in the burstiness of the delivery probability.

Figures 7 and 8 show how delivery probability changes at a time scale of 200 milliseconds, but it would also be useful to know how much fluctuation there is over other intervals. Such an analysis would reveal any characteristic burst size of packet losses. One way to summarize changes at different time scales is to plot the Allan deviation [2] at each time scale. Allan deviation differs from standard deviation in that it uses the differences between successive samples, rather than the difference between each sample and the long-term mean. In this case, the samples are the fraction of packets delivered in successive intervals of a particular length. The Allan deviation is appropriate for data sets in which the data has persistent fluctuations away from the mean. The formula for the Allan deviation of a sequence of samples x_i is:

$$\text{Allan deviation} = \sqrt{\frac{1}{2n} \sum_{i=2}^n (x_i - x_{i-1})^2} \quad (1)$$

The Allan deviation will be high for interval lengths near




Figure 9: Allan deviations of loss rate for various time intervals, for the four 1 Mbit/s links in Figure 7. 1, 2, 3, and 4 refer to the graphs in Figure 7 from top to bottom. The bold line is what you would see if losses were independent. The lines begin at the time for a single packet transmission, as this is the smallest value for which the Allan deviation can be computed.

the characteristic burst length. At smaller intervals, adjacent samples will change slowly, and the Allan deviation will be low. At longer intervals, each sample will tend towards the long-term average, and the Allan deviation will also be small.

Figures 9 and 10 show the Allan deviations of loss rate at various intervals for 1 and 11 Mbit/s. The bold line shows the deviation for a synthetic link with independent packet loss: the deviation starts at a maximum when the interval is equal to the 1500-byte packet transmission time, then decreases because averaging over longer time intervals rapidly smooths out fluctuations. The data from Roofnet also start at a peak at one packet time, but they decrease less quickly than the synthetic independent data; this suggests that actual loss is bursty on some links.

For both 1 and 11 Mbit/s, the Allan deviation shows that loss behaves as if it were independent for time intervals less than about 0.1 seconds. For longer intervals, some of the links show bursty losses, and some do not. The bursty links all show correlation out to at least 1 second.

To illustrate what fraction of links exhibit these bursty loss patterns, Figure 11 shows the Allan deviation of loss rate for all links, for one-second intervals. The graph shows that most links vary in loss rate by only a few percent from one second to the next, but that there are a small minority of links that vary by 10% or more. That is, the relatively smooth bottom graphs in Figures 7 and 8 are the most common types.

The predominance of non-bursty links suggests that most of the links with intermediate loss rates in Figure 4 have more or less independent packet loss. That is, the links are not really alternating between “up” and “down.” One consequence of this is that, for most links, measuring a link’s loss rate over intervals as short as a few seconds will provide




Figure 10: Allan deviations of loss rate for various time intervals, for the four 11 Mbit/s links in Figure 8. The bold line is what you would see if losses were independent.

an average useful in predicting the near-term future. On the other hand, a significant minority of links (at the right in Figure 11) varies substantially in loss rate from one second to the next.

6. Effect of Signal-to-Noise Ratio

One reason that many links in Figure 4 have intermediate loss rates might be that many links have marginal signal-to-noise ratios. The Prism 2.5 specification [1] suggests that the range of S/N values for which the packet error rate would be between 10% and 90% is only 3 dB wide, assuming additive white Gaussian noise (AWGN).

Figure 12 shows the results of an emulator experiment in which the sender is essentially connected to the receiver with a cable and a variable attenuator. The x -axis shows the S/N reported by the receiver's card at each level of attenuation, and the y -axis shows the delivery probability. This experiment confirms the manufacturer's specification: most S/N values result in either very high or very low loss rates; the intermediate range is only a few dB wide.

In order for marginal S/N with AWGN to explain why so many Roofnet links have intermediate loss rates, the majority of Roofnet links would have to have S/N ratios in a narrow 3 dB range. Figure 13 shows that is not the case. The range of S/N values is much greater than 3 dB, even though most Roofnet links have intermediate loss rates. Figure 14 shows scatter plots of each Roofnet link's average delivery probability and average S/N. While high S/N values correspond to high delivery probabilities, the range of S/N values for intermediate loss rates is much wider than 3 dB.

It is possible that variations in receive sensitivity across the nodes could be responsible for the spread of S/N values in Figure 14, but that individual nodes might have the expected relationship between S/N and delivery probability. Figure 15 shows per-receiver versions of the 1 Mbit/s plot from Figure 14. These plots show a better correlation between S/N and delivery probability. However, the range of S/N values corresponding to intermediate loss rates is still

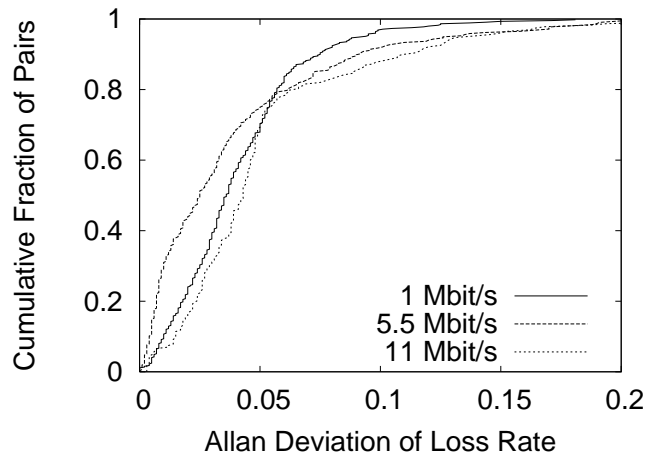


Figure 11: Allan deviation of loss rate over one-second intervals for all links at 11, 5.5 and 1 Mbit/s.

much larger than 3 dB for three of the four receivers, suggesting that S/N is not the only factor determining delivery probability.

Figure 16 shows the effect on delivery probability of varying the transmit power level, and thus the received signal strength. The data come from an experiment in which each sender transmitted at three different power levels. The three curves show the delivery probabilities between the node pairs at 10, 40, and 200 milliwatts. The power levels were verified by cabling an 802.11 card directly to a spectrum analyzer.

Figure 16 is not entirely inconsistent with simple models: assuming that signal strength falls off with the cube of distance, quadrupling the power should increase the radius of any given signal level by 1.6, and the area covered by 2.6. This is somewhat higher than the increase in the number of nodes covered when increasing the power from 10 to 40 milliwatts.

A practical conclusion from the data in this section is that although S/N does affect delivery probability, one cannot expect to use S/N as a predictive tool.

7. Effect of Transmit Bit-Rate

Figure 4 implies that the 802.11b transmit bit-rates differ in robustness; for example, there are about three times as many links at 1 Mbit/s as at 11. This section explores the effect of transmit bit-rate on losses, and particularly on net throughput, in more detail.

Figure 17 shows, for each pair of nodes, the throughput in 1500-byte packets/second at the different bit-rates. The pairs are sorted by the throughput at 11 Mbit/s. The graph is truncated so that it is missing the low-quality pairs.

Figure 17 has a number of implications for 802.11b bit-rate selection algorithms. First, an algorithm should wait until a high bit-rate is performing very badly (i.e. delivering only half the packets) before it reduces the bit-rate. Second, 11 Mbit/s often provides higher throughput than 5.5 Mbit/s even when the loss rate at 11 Mbit/s is higher than 50%. Third, performance at a low bit-rate is not a good predictor of performance at higher rates: for exam-

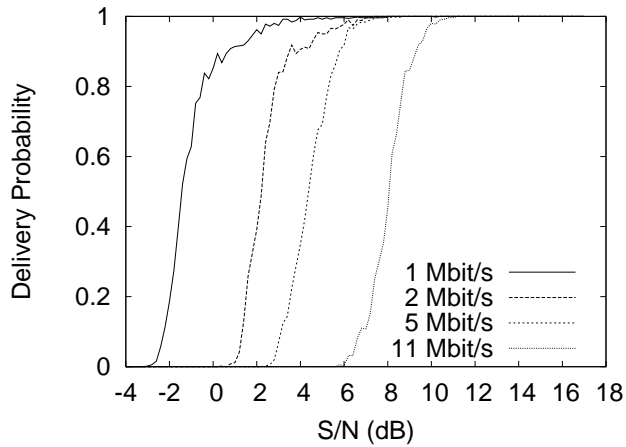


Figure 12: Delivery probability versus S/N, measured using the emulator and two Prism 802.11b cards. The S/N values are derived from values reported by the receiving card.




Figure 13: Distribution of Roofnet links' average S/N values, for the same experiments as Figure 4. There is one value in the CDF per sender/receiver pair.




Figure 14: Delivery probability at 1, 2, 5.5, and 11 Mbit/s versus the average S/N. Each data point represents an individual sender-receiver pair.




Figure 15: A scatter plot of average S/N vs average delivery probability at 1 Mbit/s. Each graph corresponds to a different receiver, each point shape corresponds to a different sender, and there is one point per one-second interval per sender.

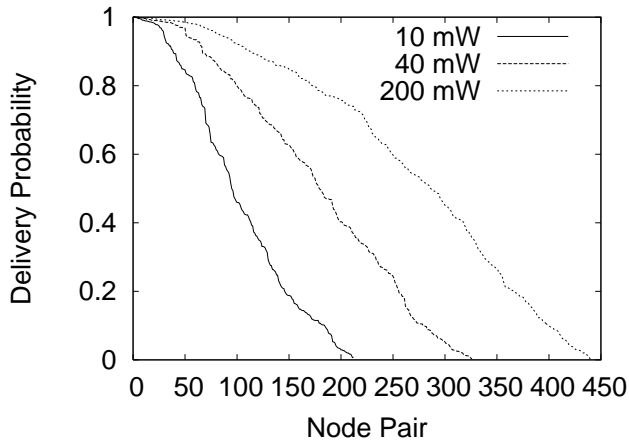


Figure 16: The effect of varying the transmit power level on the delivery probability, for 1 Mbit/s. For example, raising the power level from 10 to 40 milliwatts almost doubles the number of nodes that have delivery probabilities of 40% or more.

ple, there are many links with high loss rates at 1 Mbit/s that would have a higher throughput at 11 Mbit/s. These observations imply that bit-rate selection must be based on explicit measurements of throughput at the different rates, rather than on indirect prediction.

8. Interference from 802.11 Sources

Another possible reason for the links with intermediate delivery probabilities could be interference from other 802.11 activity. Packets could be lost due to interference from other 802.11 senders on the same channel, or from overlapping channels.

These packets might be data, or they might be periodic 802.11 beacons. Data traffic would probably be bursty, while beacons would likely maintain a relatively steady rate. Roofnet itself generates no 802.11 beacons because it runs in pseudo-IBSS mode.

Table 1 shows the number of packets per second received on each channel, averaged over all the Roofnet nodes. These numbers were acquired with the Prism 2.5 “monitor” mode, which reports all packets on a given channel, not just packets in the same BSSID as the node. During the experiment, all the Roofnet nodes were placed in monitor mode simultaneously. In this mode Roofnet itself produces no packets, so all packets in the table are from non-Roofnet sources. The measurements were taken just after midnight.

Table 1 shows that, when our experiments run, most intelligible 802.11 traffic takes the form of beacons rather than data packets. On channel 3, which the experiments in this paper use, each Roofnet node received an average of 46 packets per second.

Are the numbers of packets received from non-Roofnet sources consistent with the quantity of losses observed by Roofnet receivers? Figure 18 shows a scatter plot, with one point per Roofnet pair, relating the number of losses per second during a broadcast experiment to the number of packets per second observed by the receiver on the same channel during an immediately preceding monitor experiment.

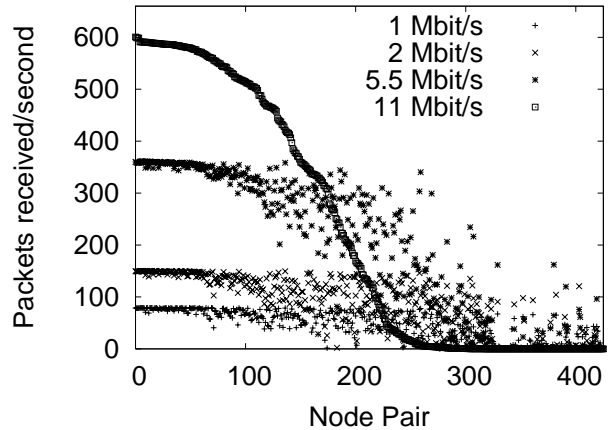


Figure 17: Throughput for each link at each 802.11b transmit bit-rate. The throughput values are in units of received 1500-byte packets per second. The node pairs are sorted by their throughput at 11 Mbit/s.

Chan	Data	Beacons
1	11.7	55.2
2	8.9	25.7
3	8.9	36.7
4	6.6	69.1
5	7.0	66.0
6	6.5	237.2
7	5.9	54.7
8	4.7	42.6
9	5.0	31.9
10	5.0	42.1
11	9.0	43.2

Table 1: Data and beacon packets per second received on each channel, averaged over all Roofnet nodes. These numbers include all frames recognized by the 802.11 hardware in “monitor” mode, including non-Roofnet traffic and damaged frames.

While the numbers of foreign packets are of the same order of magnitude as the numbers of lost packets, there does not seem to be any correlation between foreign packets received by each receiver and Roofnet packets lost by each receiver. It does not seem likely that foreign 802.11 packets on channel 3 are causing Roofnet losses.

9. Effect of Multi-path

A receiver may hear not just the signal that travels directly from the sender, but also copies of the signal that reflect from objects such as buildings. The reflected signals follow longer paths than the direct signal, so the receiver sees the combination of multiple copies of the signal at different time offsets. The Intersil HFA3873 baseband processor in the Prism 2.5 chip-set has a RAKE receiver and equalizer capable of suppressing reflected copies with delays of up to 250 nanoseconds [1]. However, studies of outdoor urban radio propagation [13, 5] find that delay spreads often exceed one microsecond. Theoretical models [4] demonstrate that




Figure 18: Each point indicates one host pair. The x -axis shows the number of foreign packets received per second by the receiver in that pair (same data as Table 1). The y -axis shows the number of 1500-byte packets lost per second at 1 Mbit/s. The data for the two axes are from experiments performed within a few minutes of each other. There is no obvious correlation.

such delay spreads significantly increase packet loss rates. While we cannot characterize the reflective paths present in Roofnet, we can evaluate the impact of longer delay spreads on packet loss with the channel emulator described in Section 2.1.

The emulator uses a two-ray channel model, in which a delayed copy of the transmitted radio signal is attenuated and mixed with the original before arriving at the receiving radio. This emulates the original signal following a line-of-sight path and a single reflective signal which followed a longer path. The parameters of the model are the delay between the two signals and their relative strengths. A real physical environment would produce many reflective rays, so the emulation results probably provide a lower bound on the losses caused by reflections.

In the experiment, the sender transmitted batches of 200 broadcast packets at 1, 2, 5.5 and 11 Mbit/s. Measurements were taken while varying both the delay and attenuation of the reflected ray in increments of 0.02 microseconds and 0.2dB respectively. The original ray was not attenuated.

Figure 19 presents the results for each bit-rate. Each bar corresponds to a different delay difference, indicated on the x -axis. The black part of each bar indicates the attenuation levels which resulted in $\geq 90\%$ loss, while the gray part of each bar indicates the attenuation levels which resulted in loss rates between 10% and 90%.

For example, in the 1 Mbit/s data, at a delay of 1 microsecond, at least 90% of the packets are lost when the reflected ray is attenuated by 9.2dB or less, and at least 10% of the packets are lost when the reflected ray is attenuated by 7dB or less. For a delay of 1.5 microseconds, 90% packet loss never occurs, even when the reflected ray and the original ray have equal signal strength; 10% of packets are lost when power of the reflected signal is within 4dB of the original. No packets are lost if the reflected signal is more than 4 dB weaker.

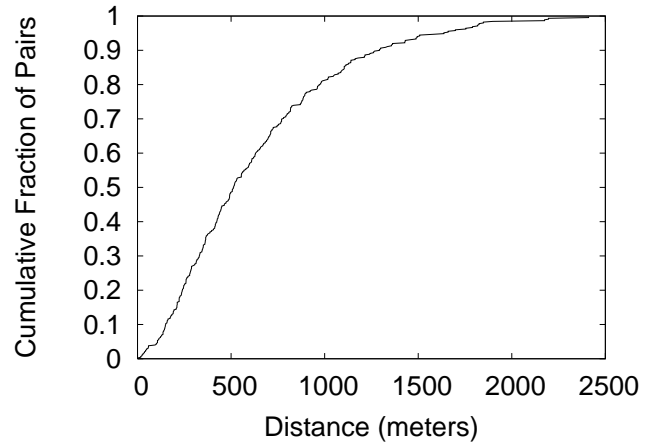


Figure 20: A CDF of the distance between all pairs of Roofnet nodes that have non-zero delivery probabilities for 1500-byte packets at 1 Mbit/s.

Figure 19 shows that a delay spread of less than a few hundred nanoseconds has little effect on packet loss, regardless of the relative strength of the reflected ray; this is consistent with Intersil's specification for the RAKE receiver. Packet loss rates increase for delays above a few hundred nanoseconds because the RAKE receiver has trouble distinguishing the original signals from the reflections.

The data show that some delays cause more loss than others; these delays are multiples of the modulation's symbol boundaries. For the phase shift keying modulation used by the 1 and 2 Mbit/s data rates, each symbol lasts one microsecond, so the data show peaks of loss at x values of one and two microseconds. At the 5.5 and 11 Mbit/s data rates, the complementary code keying modulation has symbol duration of 0.73 microseconds. At delay offsets that are not multiples of a symbol time, the delayed path's symbols look like random noise with respect to the symbols of the direct path, and do not interfere very much; at offsets that are a multiple of a symbol time, the delayed path delivers valid symbols that the receiver cannot distinguish from the direct symbols.

Significant losses due to multi-path are only likely to occur if the inter-node distances in Roofnet are long enough that a reflected signal could be delayed on the order of a microsecond. This delay corresponds to approximately 300 meters, so Roofnet would have to have direct paths significantly longer than that. Figure 20 shows a CDF of the distances between all pairs of Roofnet nodes that have non-zero delivery probabilities at 1 Mbit/s. The median is 500 meters, and about a quarter of the links are longer than 1000 meters. Links of this length seem compatible with delay spreads of at least a few hundred nanoseconds. It is common to assume delay spreads of up to a microsecond for similar urban wireless environments [4, 13].

The emulator experiment shows that multi-path interference could cause loss rates in a way that would be hard to predict from S/N alone. For example, at 5.5 Mbit/s, Figure 19 shows that loss rate could vary widely depending on the exact length of the reflective path; if reflective path lengths are more or less uniformly distributed, then

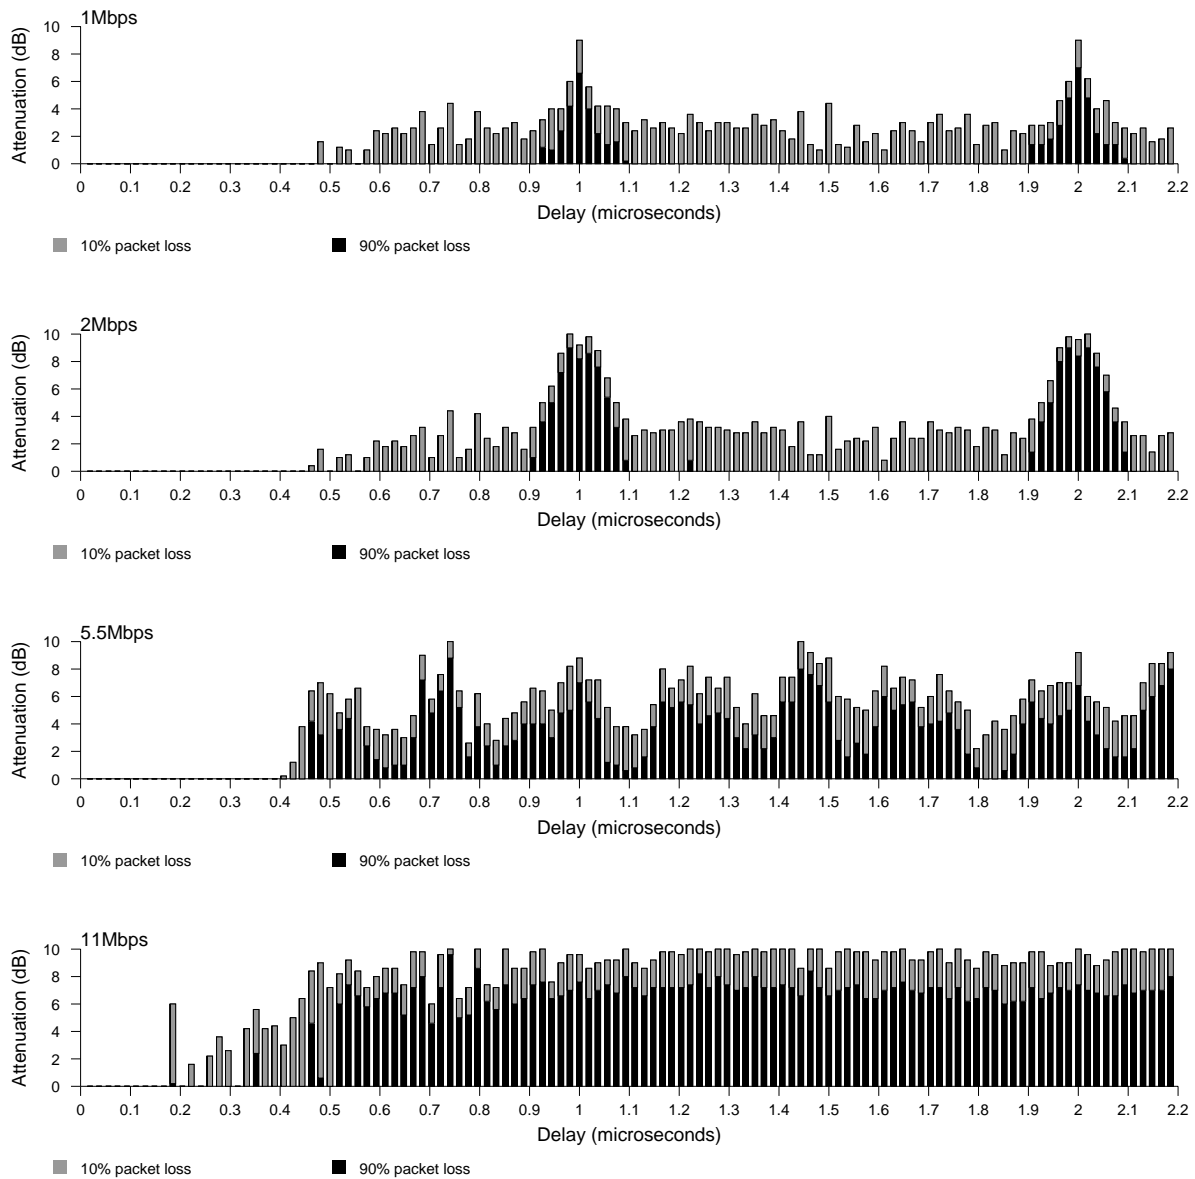


Figure 19: The effect at 1, 2, 5.5 and 11 Mbit/s of various combinations of multi-path delay and attenuation, measured using the emulator. The x -axis indicates how long the reflected ray is delayed relative to the direct ray. The y -axis indicates how much the reflected ray is attenuated relative to the direct ray. The tops of the gray and black bars at each delay indicate the attenuation levels which result in 10% and 90% packet loss respectively, so the height of the gray bar indicates the region of intermediate packet loss. Packet loss is common for delay spreads greater than a few hundred nanoseconds and occurs more often when the reflected ray is delayed by a multiple of the symbol time.

one might expect the loss rate caused by multi-path to be roughly uniformly distributed as well. This uniformity is a potential contributor to the prevalence of intermediate loss rates in Figure 4.

10. Related Work

Eckhardt and Steenkiste [7] found relatively few links with intermediate loss rates in an indoor 802.11 network.

Kotz et al. [10] examine the question of whether, if two nodes can hear each other at all, they can hear each other perfectly (their Axiom 4). They conclude that the assumption is very nearly correct. Our measurements show that Roofnet behaves very differently from the network studied by Kotz, perhaps because the latter is an 802.11 access-point network in which clients are typically close to the nearest access point.

Lundgren et al. [11], Yarvis et al. [15], and De Couto et al. [6] all report much lower performance on deployments of multi-hop routing systems than predicted in simulation, and all observe in one way or another that the problem is a predominance of intermediate-quality links. They propose solutions that involve measuring link quality and carefully routing through the best links. While our work does not propose any solutions, it does contribute to an understanding of the reasons for intermediate-quality links.

Earlier studies of 802.11 links [14, 12, 3] find packet losses are bursty and require a multi-state Markov model for accurate prediction. These bursts were caused by movement in the environment, either of the receiver or obstacles, which induced slow changes in the channel due to Rayleigh fading. In the case of static outdoor 802.11 measurements [3], the highly directional antennas are susceptible to motion caused by wind, which is a possible explanation for their bursty losses. The omni-directional antennas used in Roofnet are not as easily affected by weather conditions. In general, multi-state Markov models such as the Gilbert Model and those developed in [9] are found to fit loss patterns induced by movement well, but are over-specified for static environments.

11. Conclusions

This paper is a study of packet loss on a 38-node urban 802.11b mesh network. Links with intermediate levels of loss are the common case; there is no clear distinction between “working” and “non-working” links. Link distance and S/N ratio do have an effect on loss rates, but the correlation is weak. Experiments using a hardware channel emulator suggest that an important cause of intermediate loss rates is multi-path fading due to reflections in the radio environment.

The measurements on which this paper is based will be available at <http://pdos.lcs.mit.edu/roofnet>.

Acknowledgments

This research was supported by grants from NTT Corporation under the NTT-MIT collaboration, and by MIT’s Project Oxygen. We thank the many volunteers who host Roofnet nodes for their time and patience. Ben Chambers deployed the original Roofnet, Doug De Couto wrote the original software and designed the initial protocols, and Ed-die Kohler maintains the indispensable Click.

References

- [1] *ISL3873: Wireless LAN Integrated Medium Access Controller with Baseband Processor*. Intersil Corporation, 2000. Application Note FN4868.
- [2] D. W. Allan. Time and frequency (time domain) characterization, estimation and prediction of precision clocks and oscillators. In *IEEE Trans. UFFC*, vol. 34, no. 6, November 1987.
- [3] H. Balakrishnan and R. Katz. Explicit loss notification and wireless web performance. In *IEEE Globecom Internet Mini-Conference*, October 1998.
- [4] M. V. Clark, K. K. Leung, B. McNair, and Z. Kostic. Outdoor IEEE 802.11 cellular networks: Radio link performance. In *Proc. of IEEE ICC 2002*, April 2002.
- [5] D. C. Cox. Delay Doppler characteristics of multipath propagation at 910 MHz in a suburban mobile radio environment. In *IEEE Transactions on Antennas and Propagation*, AP-20(5):625-635, September 1972.
- [6] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *Proceedings of ACM MobiCom Conference*, September 2003.
- [7] D. Eckhardt and P. Steenkiste. Measurement and analysis of the error characteristics of an in-building wireless network. In *Computer Communication Review* 26:4, pp. 243-254, SIGCOMM '96, October 1996.
- [8] G. Judd and P. Steenkiste. Repeatable and realistic wireless experimentation through physical emulation. In *HotNets-II*, Cambridge, MA, November 2003. ACM.
- [9] A. Konrad, B. Y. Zhao, A. D. Joseph, and R. Ludwig. A Markov-based channel model algorithm for wireless networks. In *Proceedings of Fourth ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2001.
- [10] D. Kotz, C. Newport, and C. Elliott. The mistaken axioms of wireless-network research. Technical report TR2003-647, Dartmouth CS Department, July 2003.
- [11] H. Lundgren, E. Nordstrom, and C. Tschudin. Coping with communication gray zones in IEEE 802.11b based ad hoc networks. In *ACM WoWMoM Workshop*, September 2002.
- [12] G. Nguyen, R. H. Katz, B. Noble, , and M. Satyanarayanan. A trace-based approach for modeling wireless channel behavior. In *Proc. Winter Simulation Conf.*, December 1996.
- [13] E. S. Sousa, V. M. Jovanovic, and C. Daigneault. Delay spread measurements for the digital cellular channel in Toronto. In *IEEE Trans. on Veh. Tech.*, vol. 43, no. 4, pp. 1-11, November 1994.
- [14] A. Willig, M. Kubisch, C. Hoene, and A. Wolisz. Measurements of a wireless link in an industrial environment using an IEEE 802.11-compliant physical layer. In *IEEE Transactions on Industrial Electronics*, vol. 43, no. 6, pp. 1265-1282, December 2002.
- [15] M. Yarvis, W. Conner, L. Krishnamurthy, J. Chhabra, B. Elliott, and A. Mainwaring. Real-world experiences with an interactive ad hoc sensor network. In *Proceedings of the International Workshop on Ad Hoc Networking*, August 2002.