

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.