An Empirical Characterization of Wireless Network Behavior

Charles Reis

creis@cs.washington.edu

Abstract

Existing work on understanding 802.11 wireless network behavior has been largely unsatisfactory for practical settings. Widely used simulators rely on unrealistic assumptions about signal propagation, while more detailed radio models are too complex to configure for predicting performance of an actual wireless system. To gain a more accurate understanding of wireless behavior in practice, we use experimentation on a wireless testbed and in controlled settings to effectively characterize packet delivery. We contribute a simple measurement-based model of wireless behavior, supported by empirical observations of relevant physical effects. Our model and observations can be used directly for designing and improving wireless protocols and systems in practice.

1 Introduction

Wireless networks present a daunting challenge for system and protocol designers, who must cope with a degree of complexity far beyond that seen in wired networks. While wired networks can often be modeled with simple undirected graphs representing connected nodes, the behavior of wireless networks is fundamentally tied to far more complex radio environments that affect the probability nodes will receive packets, often in highly irregular or unpredictable ways, and with high variability. A basic understanding of the impact of such effects is necessary for wireless system building to be feasible and effective.

Unfortunately, while 802.11 wireless networks have become explosively popular over the past several years, networking researchers and administrators continue to rely on many flawed assumptions about the physical properties of such networks, leading to ineffective protocol designs and evaluation techniques. Wireless simulators are frequently used to evaluate new proposals, yet they often rely on simplistic models of signal propagation (often based on distance alone) that cannot easily capture the complexities of real environments, leading to inaccurate results. On the other hand, more detailed physical models of radio signals are often far too low-level to be useful for system or protocol design. For example, sufficient information about the materials between a transmitter and receiver in a real environment can be nearly impossible to collect, ruling out accurate techniques for predicting signal propagation.

Meanwhile, testbeds of wireless machines provide an alternative, allowing observations of actual behavior in representative environments. Unfortunately, the management overhead, difficulties of uncontrolled environments, and the common need to test a large number configurations can make such an approach equally problematic.

We seek an improved understanding of wireless networking, as it is relevant for system and protocol design and analysis. To accomplish this goal, we use testbed experimentation to provide an empirical analysis of packet reception, and we contribute a new measurement-based model of the physical aspects of wireless networking that incorporates our findings. We aim for a simple, usable, and realistic model: it should be easy to understand and use, rely only on inputs that are readily available, and provide sufficiently accurate results to be effective as an evaluation tool for new systems and protocols.

This paper discusses our observations of the key aspects of wireless network behavior, and it describes our resulting model and its evaluation. Our methodology is presented in Section 2, while Section 3 attempts to characterize wireless networks by investigating basic packet reception, variability, packet loss, and asymmetry. We describe our measurement-based model in Section 4, and we evaluate its performance in Section 5. Finally, we discuss related attempts to improve understanding of wireless networking in Section 6, and we conclude in Section 7.

2 Methodology

To reach an improved understanding of the behavior of wireless networks in practice, we must rely on empirical observations from deployed environments, rather than on previous analytic models or simulations. While some previous work, such as Divert [9] and ExOR [5], has

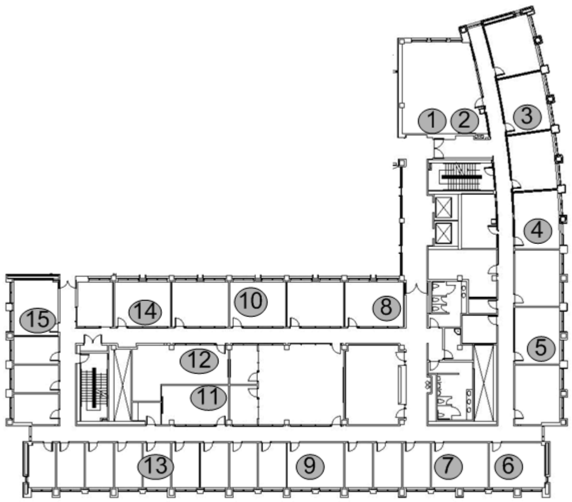

Figure 1: Our wireless testbed, consisting of fifteen 802.11 a/b/g nodes. The node locations are marked with circles. Horizontally, the building is 184 ft. long.

focused on improving wireless protocols based on isolated observations, we seek a broader approach to identify which aspects of wireless networking overall are most important for designing new protocols. We accomplish this using a combination of both controlled and uncontrolled experiments with commodity wireless networking cards, guided by the desire to construct a simple model for packet reception.

2.1 Testbed Experiments

The primary source of our experimental measurements is an in-building testbed of fifteen stationary PCs with commodity wireless cards, located on the third floor of our department building. The testbed is designed to mimic a relatively dense deployment of wireless nodes in a realistic office scenario, if wireless networking were used to augment or supplant wired connectivity. We focus on static node placement rather than mobility, in an effort to simplify the characterization problem. Even without addressing mobility, we feel that an evaluation of wireless network behavior can be valuable for many important scenarios.

Our testbed is depicted in Figure 1, showing the placement of the nodes in offices throughout the third floor of our department's building. Each node has a single 802.11 a/b/g Atheros card, forming a connected network. We employ the Click modular router [8] and the "stripped" Madwifi driver [4] from MIT to control low level aspects of the card's behavior and record the contents and detailed statistics for all packets decoded by the card. These statistics include signal strength measurements (in the form of receive signal strength indicators (*RSSI*)), as well as packet length, reception time, etc. Specifically, RSSI refers to the energy observed in the air during the PLCP header before a packet arrives, and it is thus an imperfect measure of a packet's actual signal strength [3]. Meanwhile, noise is only sampled once per second with a special-purpose script rather than per-packet, and we have limited ability to control transmit power. For the testbed, we have developed an extensive software infrastructure to deploy experiments, gather data, and analyze results, enabling us to easily test hypotheses about wireless behavior.

Our experiments are conducted using 802.11b. Because our focus is on physical level behavior, we attempt to minimize impact from MAC-level features, such as acknowledgments and request-to-send messages. This involves operating in Click's pseudo-IBSS mode (which minimizes the use of management packets) and sending broadcast packets only, such that acknowledgments and retransmissions are suppressed. Thus, we can study pure reception rates from one node to another, or while multiple transmitting nodes compete with each other. It is worth noting, though, that the use of broadcast packets limits our experiments to a bitrate of 1 Mbps, rather than allowing a variety of bitrates.

Unfortunately, while the testbed is able to provide realistic measurements, it also operates in a noisy environment, with many active people and energy sources, including an official department-wide 802.11b wireless network. We operate on a separate but non-orthogonal channel from this network, allowing departmental traffic to leak across as noise to a small degree.[1]

2.2 Attentuator Experiments

To address the high variability present in testbed environments, we also conduct a series of more controlled experiments using shielded cables and variable attenuators. This approach is valuable for identifying the key characteristics of the wireless cards' behavior, while limiting the impact of competing energy sources and changing environments.

Specifically, using two desktop PCs with wireless cards, we disconnect the antennas from the cards and attach shielded SMA cables, along with both fixed and variable attenuators to control the signal strength at the receiver. Our configuration follows sensitivity experiment guide-

[1] The department network operates on channels 1, 6, and 11, while our testbed operates on channel 3.

lines provided by Intersil [1]. We perform experiments both with the same Atheros cards used in the testbed, as well as with Prism wireless cards from another manufacturer, which do not have internal antennas. Because the Atheros cards do have some ability to decode wireless packets even without an external antenna, we must surround the receiver with RF shielding foam to prevent influences from the local environment.

The attenuator experiments allow us to observe card behavior with greater precision, in a variety of simulated environments. Combined with our observations of testbed experiments, we can accurately characterize wireless network behavior for actual deployed scenarios.

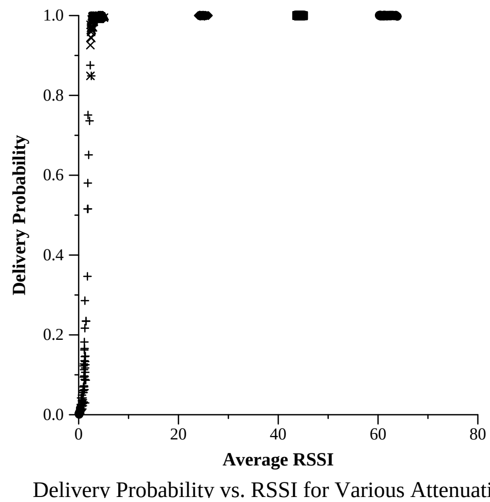

Figure 2: *Probability of successful delivery for Atheros cards as a function of received signal strength, in attenuator experiments. There is a sharp signal strength threshold indicating when packets will be received.*

3 Characterizing Wireless

In this section, we establish a more realistic characterization of wireless network behavior than distance-based models used frequently in practice. We place an emphasis on using only widely available measurement statistics to ensure our resulting model will be useful in practice, while being guided by a classical understanding of wireless packet reception. To achieve this, we conduct a series of experiments to investigate key factors in determining a node's ability to receive packets, and to evaluate the use of imperfect information for modeling.

Rather than assume packet reception depends on the distance between the transmitter and receiver, we adopt a more classical (and complex) view that it depends on the ratio between signal strength and noise plus interfering energy. If the "signal to noise plus interference ratio" (*SNIR*) is above a certain threshold during the transmission of a given packet, then the radio on the receiving wireless card will be able to successfully decode it. While this may seem relatively simple, it is traditionally difficult to use in simulation and practical models, because SNIR is not an easily observable quantity in practice. In general, the physical characteristics of an environment cannot be adequately described to model signal propagation. Additionally, while observations in a given environment can provide some information, wireless cards report only an imperfect measure of energy in the air (*i.e.*, the received signal strength indicator, or RSSI) for each packet, and not any measure of competing interference. Measurements of the thermal noise within the card may also be available, although not always on a per-packet basis.

As a result, we explore how well RSSI can be used as a proxy for SNIR for certain purposes, to determine if we can accurately model packet reception using only available information. As part of this exploration, we investigate basic packet reception in both controlled and uncontrolled settings, the impact and causes of variability between experiments, the amount of burstiness and independence between packet loss events, and potential causes for asymmetric reception abilities. Together, these aspects help to characterize wireless network behavior and reveal that RSSI can be a useful quantity in modeling, under certain assumptions.

3.1 Packet Delivery using Attenuators

We begin with a discussion of the ability of wireless cards to receive packets in a controlled environment, using the attenuator experiments described in Section 2.2. We seek to confirm the predictive power of SNIR for successful delivery probability, in a setting with minimal external interference where RSSI should be a good approximation of SNIR. We also look at the degree of variation in signal strength and noise measurements, to compare against measurements from the testbed.

For our experiments, we send a flood of large (about 1500 bytes) broadcast packets from one machine to another, using Atheros cards connected by shielded cables and attenuators. We repeat five minute floods of packets using a representative series of different signal attenuations, from zero attenuation to a level that prevents any packets from being received. The entire process is repeated for three separate trials, with the results summarized in Figures 2 and 3.

Figure 2 shows the probability of delivery as a function of signal strength (RSSI). (For our experiments, RSSI is re-

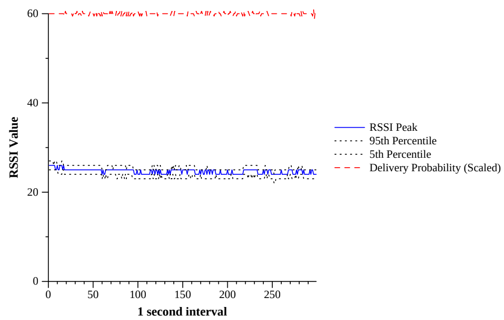

Figure 3: *Variation in signal strength and noise distributions over time, for attenuator experiments using Atheros cards.*

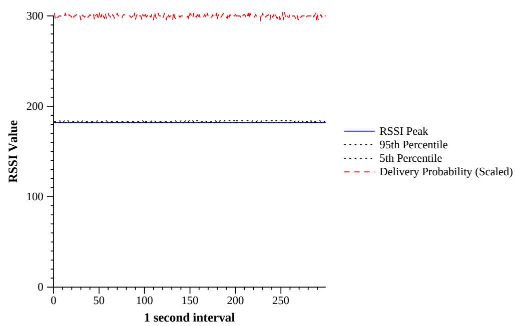

Figure 4: *Similar attenuator experiments on Prism cards without internal antennas show even less variability.*

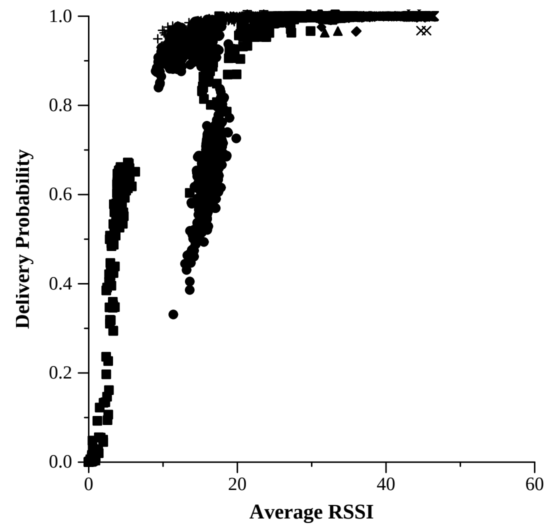

Figure 5: *Probability of successful delivery as a function of received signal strength, in overnight single sender experiments. There are multiple thresholds corresponding to different receivers, and there is a wider range of observed delivery probabilities due to impact from the RF environment.*

ported on an Atheros-specific scale, which is simply dBm plus 95.) Each point represents the average RSSI over a given five minute trial. As expected, there is a sharp threshold, above which packets are received with certain probability, and below which essentially no packets are received. There is a very small window where the delivery probability drops off rapidly, indicating the rather clear cutoff for successful packet delivery.

Figure 3 shows how the signal strength values vary over time, during a particular trial. The graph shows the peak RSSI value for the distribution of values observed during each 1 second interval, along with the 95th and 5th percentiles to show the spread of values within a given interval. While there is some spread in observed RSSIs (which we attribute to imperfect RF shielding for the receiver's internal antenna), the peaks are quite stable within a trial.

To verify that the spread in RSSI values is likely a result of poor RF shielding, we also conducted experiments with similar Prism cards that do not have internal antennas. The RSSI values are reported on a different scale, but the distributions of observed values shown in Figure 4 have virtually no spread during each 1 second interval, such that the peak, 95th, and 5th percentile all appear the same. Thus, RSSI appears to be a good proxy for SNIR in predicting delivery probabilities in controlled settings.

3.2 Packet Delivery over the Air

Next, we investigate the difference between controlled and uncontrolled environments by running experiments deployed on the testbed in our building. We look at delivery probabilities for all nodes in the testbed, while either a single sender or a pair of senders transmits a flood of broadcast packets.

Single Sender For the single sender experiments, we allow a single node to transmit large packets (about 1500 bytes) for an hour at a time. We conduct two trials late at night to minimize impact from regular department activity. The first trial started at 10:30 pm, and the second started just before 1:00 am.

Similar to Figure 2, Figure 5 shows the probability of delivery as a function of average signal strength. The graph includes points for all receivers, with each point corresponding to a single 10 second interval during the trials. Unlike in the attenuator experiment, there appear to be multiple thresholds for successful delivery in this graph. The different thresholds actually correspond to different receivers, indicating local differences in either

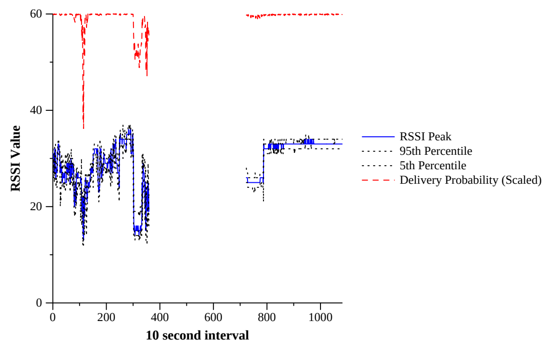

RSSI Distribution over Time for Node 12 -> Node 10

Figure 6: *Variation in signal strength and noise distributions over time, for a particular receiver in the overnight single sender experiments. Delivery probability is overlaid with a dashed line, scaled so that 100% is at the top of the graph.*

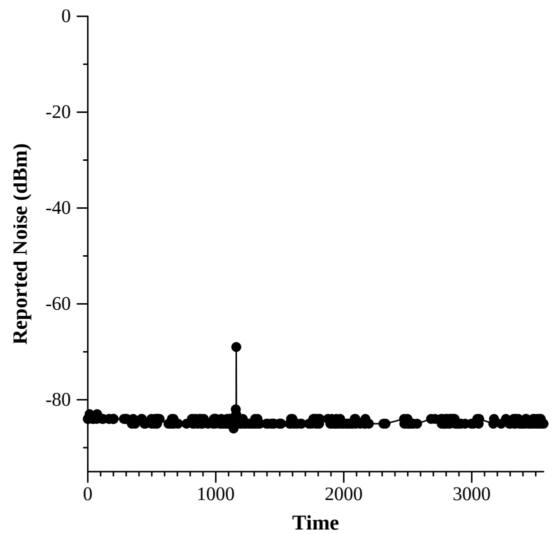

Noise values over time for Node 12 -> Node 10

Figure 7: *Variation in noise observations over time, for the first trial depicted in Figure 6. Noise is notably more stable than RSSI.*

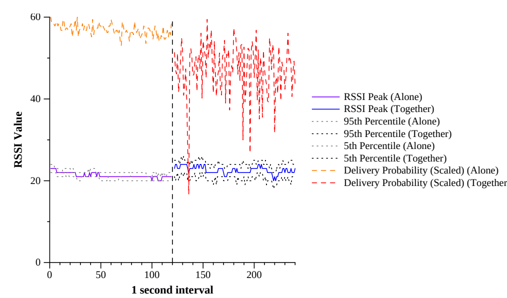

RSSI Distribution over Time for Node 15|Node 9 -> Node 11

Figure 8: *Variation in signal strength and noise distributions over time, for sender A to receiver C. The first half shows A sending to C in isolation, and the second half shows A sending to C, given that B is also transmitting.*

receiver sensitivity or environment that impact the probability of delivery. It is also noteworthy that the thresholds have shifted right and become more gradual relative to Figure 2, suggesting higher signal strength values are required for successful delivery in a realistic environment. Additionally, this graph shows evidence of non-binary connectivity, where nodes can consistently receive some given fraction of the transmitted packets.

To see the variation in signal strength over the course of the experiment, Figure 6 shows the distributions of RSSI values for each 10 second interval, for a typical receiver in range of the sender. The graph shows results from each of the 2 one-hour trials. The corresponding delivery probability for the interval has also been overlaid with a dashed line, scaled such that 100% is at the top of the graph. We note that variations in delivery probability tend to coincide with large variations in RSSI values (apparent in the first trial), while more stable RSSI values in the second trial correspond to stable delivery probabilities. This suggests the usefulness of an RSSI threshold as a predictor of delivery probability in cases with low interfering energy.

For comparison, we can see the variation of reported noise measurements over time, as sampled once a second from the card with a special purpose script. The results are shown in Figure 7, which covers the first of the two trials. Note that despite the wide variation in signal strength measurements, the noise reported by the card remains quite stable, with only one outlier. This finding is consistent across different receivers, showing that thermal noise is not likely to have a visible effect on variability.

Two Senders To explicitly observe the impact of competing signal energy, we also ran experiments with pairs of senders transmitting simultaneously. We handpicked the pairs to consider primarily cases where the senders were out of range of each other (such that they would transmit simultaneously without deferring), and we looked at receivers that could hear both senders well in isolation. We conducted three trials of two minute long experiments, allowing us to look at a large number of configurations.

We focus on one such configuration in Figures 8 and 9, which show the RSSI distribution over time for each sender at a particular receiver, at 1 second intervals. The first half of each graph shows the results of the sender transmitting in isolation. The second half corresponds to the time period when both senders transmitted simultane-

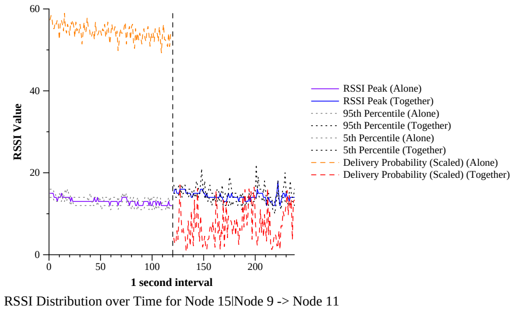

Figure 9: *Variation in signal strength and noise distributions over time, for sender B to receiver C. The first half shows B sending to C in isolation, and the second half shows B sending to C, given that A is also transmitting.*

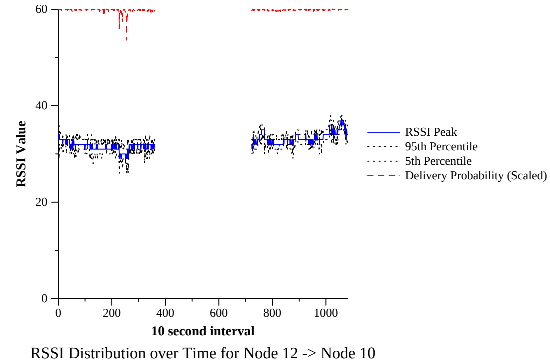

Figure 10: *Variation in signal strength and noise distributions over time, for experiments starting at noon and 3 pm.*

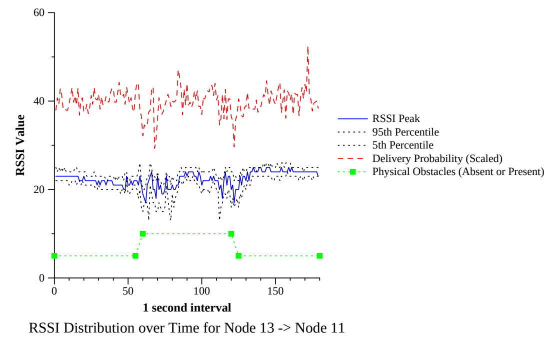

Figure 11: *Variation in signal strength and delivery probability, as impacted by shadowing effects. The dotted line indicates when physical obstacles are introduced and removed from the receiver's local environment.*

ously.

Note that the receiver was able to hear both senders well in isolation, but it exhibits a “capture effect” when both transmit concurrently. That is, it effectively only hears packets from A, since A has the higher signal strength. Very few packets from B are ever received, despite the fact that the receiver has a high delivery probability from B when A is not transmitting.

It is important to note that the signal strength observed during the concurrent half is higher than during the first half, for both senders. This illustrates the problem of using RSSI as a proxy for SNR, since the signal strength does not predict the drop in delivery probability from sender B.

Overall, we can see that RSSI values tend to be predictive of delivery probability, but only when there are not notable source of interfering energy.

3.3 Variability

Given the differences apparent between the attenuator and testbed experiments, we set out to find distinct causes of variability, to determine their relative impact and to better understand our observations. We look primarily at how the distribution of RSSI values over time are affected by various conditions, including time of day, introduction of obstacles, competing energy sources, and changes in position. We find a number of effects with different implications for modeling wireless networks with RSSI measurements.

Time of Day In Figure 10, we see the RSSI distribution for the same configuration as Figure 6, but for experiments conducted at noon and 3:00 pm instead of late at night. While there are not as many large variations as in the 10:30 pm trial, it is still more diverse than the 1:00 am trial. We can note the extreme difference between the 10:30 pm and 1:00 am trials within Figure 6, which may be explained by a greater chance that students remained working in the building before the lights switched off at midnight. These graphs suggest that variability in the results can be minimized by conducting experiments late at night, after most people have left the building.

Shadowing Effects In previous figures, there are several instances where delivery probability varies jointly with signal strength. We hypothesize that such events are the results of signal obstructions, or “shadowing” of the receiver. That is, if the signal does not arrive at the receiver with full strength, the probability of delivery will drop.

To test this hypothesis, we attempted to introduce obstacles between a sender and receiver during a short experiment. The experiment was somewhat limited, consisting of closing an office door and positioning two chairs and a person near the transmitter. As seen in Figure 11, the

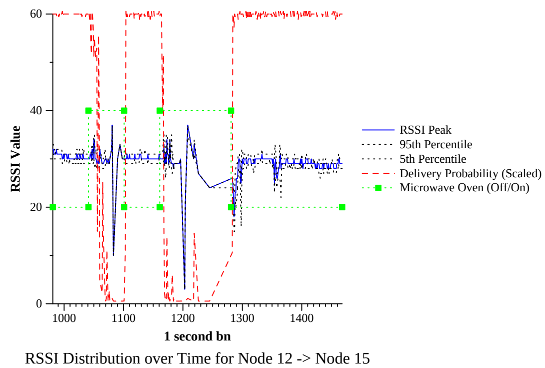

Figure 12: *Variation in signal strength and delivery probability, as impacted by a nearby microwave. The dotted line indicates when the microwave was turned on and off.*

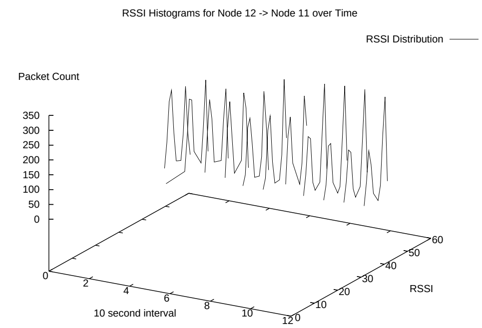

Figure 13: *Signal strength distribution at a particular receiver, exhibiting two distinct peaks.*

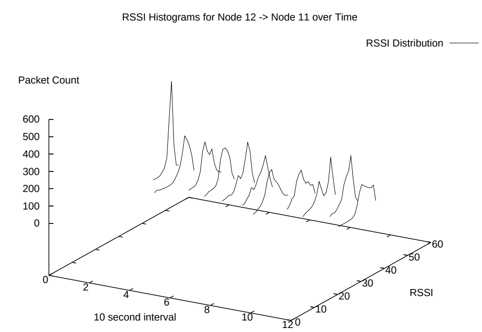

Figure 14: *Signal strength distribution at the same receiver after being moved by one foot, no longer exhibiting two distinct peaks.*

obstacles do not have an extremely large impact on the results, though there do appear to be higher variations in both signal strength and delivery probability during the introduction and removal of the obstacles. This seems to indicate that the ad hoc introduction of obstacles did affect the strength of the signal, but not always in detrimental ways.

It is worth noting that such variations in signal strength are valuable for using RSSI as an approximation of SNIR, since they provide different delivery probabilities at different signal strength values, allowing us to build up a predictive function.

Interference We are also aware of periods where delivery probability drops without any corresponding drop in RSSI. This likely corresponds to interfering energy from competing signals or other sources, as seen in Figure 9. To further confirm this hypothesis, we measured the impact of another source of competing energy— a microwave oven. As Figure 12 shows, turning on a microwave between the sender and receiver dramatically affects the delivery probability, while the RSSI for any packets that are received remains high.

Unlike shadowing effects, such interference is detrimental the ability to model wireless networks using RSSI, since it causes RSSI to diverge from SNIR and thus is not predictive of delivery probability. Thus, we would like to minimize the amount of interference during any experimental observations for RSSI to remain a useful quantity.

Positional Effects Finally, we made an interesting observation that one of the receiver nodes near the sender consistently displayed two separate peaks in its distribution of RSSI values, regardless of the time scale or other experimental conditions. This effect can be seen in Figure 13. Since this was happening consistently at a very fine granularity, we hypothesized that it could be a result of RF fading and multi-path effects, which can account for a wider spread in observed RSSI values for a given time interval. (The presence of fading and multi-path effects can already be noted in the difference between the testbed and attenuator experiments, as shown by the 95th and 5th percentile lines for the signal strength variability graphs.)

To test our hypothesis, we moved the node approximately one foot further from the sender and repeated the experiment, and the two peaks flattened, as seen in Figure 14. Thus, it is clear that only small changes in position can have notable effects on the distribution of observed signal strength values, and thus the node's ability to receive packets. This is consistent with fine-grained changes affecting RF fading in practice.

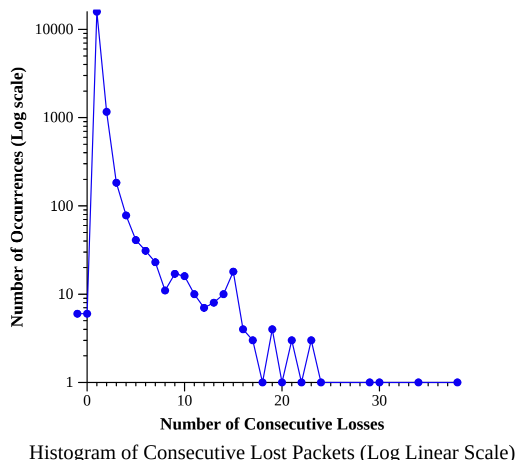

Figure 15: *Histogram of how often bursts of consecutive losses of different lengths occur. Most loss burst events tend to be less than 20 packets long.*

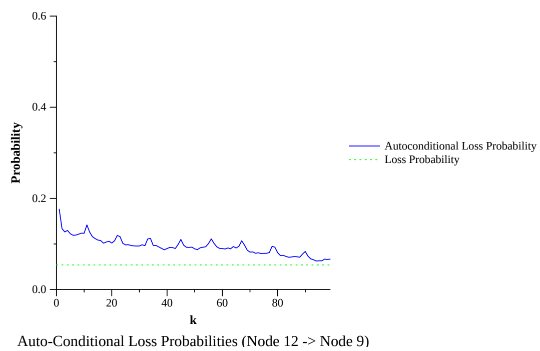

Figure 16: *Auto-conditional loss probability, compared to overall loss probability. Loss events appear to be independent for larger sized bursts.*

Variability Summary Through a series of experiments, we have confirmed several sources of variability in the behavior of deployed wireless networks. Additionally, we can note that shadowing effects increase the quality of RSSI observations for modeling purposes by adding valuable diversity, while interfering energy decreases the accuracy of using RSSI in place of SNIR.

3.4 Loss

We next investigate the nature of packet losses, and whether losses occur independently over time or in short or long bursts. This information will reveal how well time intervals of different sizes can be used to characterize the behavior of a node.

Figure 15 shows how often we observe bursts of consecutive losses of different lengths, for a given sender and receiver during a one hour experiment at night. The graph is on a log linear scale and shows an almost geometric drop-off in observed bursts as the burst length increases. This suggests that loss events are largely independent, and do not frequently occur in large bursts. Indeed, while we do observe some cases of 20 or more consecutive losses, these events are relatively rare, and most losses tend to occur in much shorter bursts.

An alternative view is shown in Figure 16, which covers data from the same sender and receiver as above. The dotted line shows the overall loss probability for the trial, indicating that this receiver had a 5.4% chance of loss from the given sender. The solid line shows the "auto-conditional loss probability," as inspired by similar graphs in the Divert paper [9]. That is, given that packet $i-k$ is lost (*i.e.*, A_{i-k}), what is the probability that packet i will also be lost? If losses are fully independent, we would expect $P(A_i|A_{i-k}) = P(A_i)$, and thus the solid line should be the same as the dotted line. However, we can see that for small values of k, the auto-conditional loss probability is slightly greater, indicating that losses are bursty at small intervals. As k increases, the lines converge, showing that losses become independent for larger burst sizes. Similar to the authors of [9], we hypothesize that the periodic increases in auto-conditional loss probability correspond to periodic beacon frame interference from the department network.

Overall, we can see that losses tend to occur in small bursts (*i.e.*, less than a few dozen packets), but can be treated as independent for larger time intervals of several seconds or minutes.

3.5 Asymmetry

We have observed cases where delivery probabilities between two nodes A and B are dramatically asymmetric, such that A receives packets reasonably well from B, but B receives almost nothing from A, as seen in Figures 17(a) and (b). Under the assumption that the local RF environments in different parts of the third floor are at least relatively uniform, one possible explanation for such asymmetry is differences in receiver sensitivity, as found by Judd in [7].

To test this hypothesis, we swapped the wireless cards in use by nodes A and B and repeated the experiment. However, despite a slight shift in signal strength, the predominant asymmetry still holds, as shown in Figures 18(a) and (b). We did note substantially higher noise floor readings in the node with poor reception, in both experiments. This suggested an unlikely possibility that some aspect of the PC itself (perhaps a power supply) was the cause of the noise and thus the asymmetry.

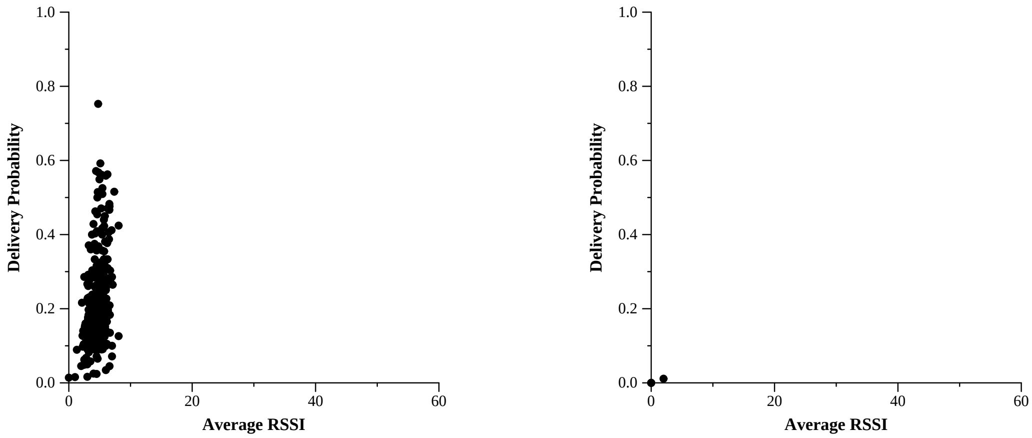

(a) Delivery Probability vs. RSSI for Node 10 -> Node 13

(b) Delivery Probability vs. RSSI for Node 13 -> Node 10

Figure 17: *Asymmetric delivery probabilities for two nodes A and B.*

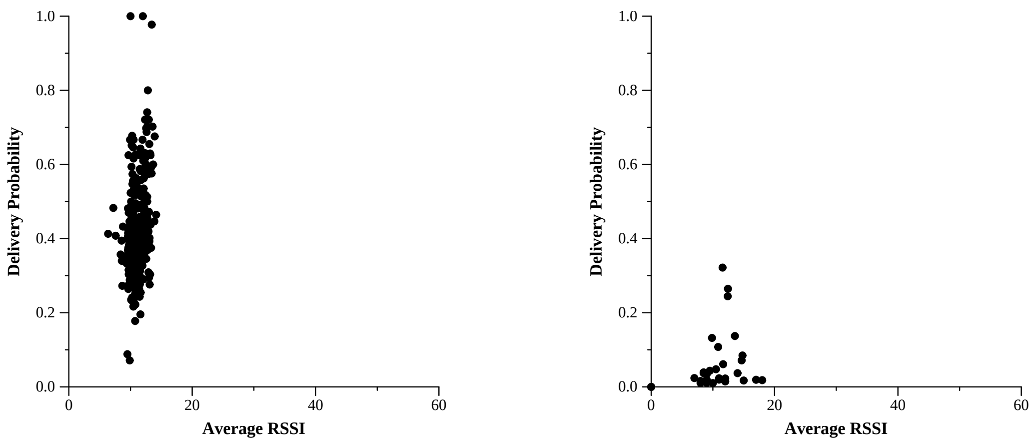

(a) Delivery Probability vs. RSSI for Node 10 -> Node 13

(b) Delivery Probability vs. RSSI for Node 13 -> Node 10

Figure 18: *Asymmetric delivery probabilities for nodes A and B, after swapping the wireless cards. The asymmetric relationship remains unchanged.*

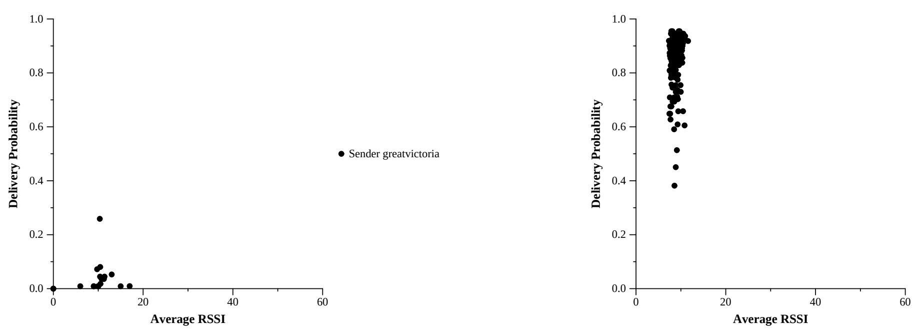

(a) Delivery Probability vs. RSSI for Node 10 -> Node 13 (swapped locations)

(b) Delivery Probability vs. RSSI for Node 13 -> Node 10 (swapped locations)

Figure 19: *Asymmetric delivery probabilities for nodes A and B, after swapping the locations of the PCs themselves. The asymmetric relationship is tied to the location of the PC, and not the PC itself.*

Thus, we repeated the experiment after swapping the entire PCs, as shown in Figures 19(a) and (b). The more successful delivery probability remains tied to the location, and not the PC or card. We can thus conclude that notable differences in local environments exist even in relatively small testbeds, and these differences can lead to largely asymmetric behavior. As a result, we are forced to consider each receiver's reported RSSI values in isolation, since they cannot be compared across locations.

3.6 Summary

Overall, we can conclude several things about the characteristics of wireless packet reception, and about the interpretation of RSSI values reported by wireless cards. There is surprisingly high variability between local environments around the receivers, and observed RSSI values are specific to a location and cannot easily be compared. Additionally, most variability over time comes from the environment and not thermal noise reported by the cards. RSSI values do tend to vary notably at fine time scales and links exhibit short bursts of losses, but our measurements tend to stay near stable long term averages, with larger time intervals independent of each other. Finally, in terms of predicting delivery probability, RSSI values can vary in useful ways due to shadowing, and in detrimental ways due to external interference. In general, these observations can be used to help employ RSSI as a proxy for SNIR, as we will discuss in the next section.

4 Measurement-Based Model

Given the set of observations and hypotheses about wireless packet reception covered in Section 3, we can now take a step toward a better understanding of the physical behavior of wireless networks. We describe the construction of a *measurement-based model* of packet reception, designed to explain and predict behavior of real wireless networks in practice. This model addresses the ability of nodes to receive packets in a physical sense, and it can thus be viewed as a low layer model upon which separate models of protocol or workload behavior could be built.

The key idea is to exploit a set of empirical measurements to capture most of the complexity of the environment, allowing the model itself to be quite simple. This is accomplished using simple "calibration" measurements on the network of interest, with each machine sending packets in isolation while all other machines record what they receive. No measurements of combinations of transmitters are necessary, as our model will predict such scenarios using computed SNIR values.

To ensure that the resulting model is usable in practice, we will limit ourselves to using observed RSSI, rather than requiring SNIR values. To use RSSI as an approximation of SNIR, we must rely on our previous observations as guidelines for how to interpret and use the values. As such, we know that we must treat RSSIs from separate nodes as incomparable, and we must rely on time intervals around a minute or longer to obtain relatively stable results.

The model itself can be broken into separate components for describing receiver behavior and transmitter behavior, to simplify each component. As mentioned above, the complex nature of signal propagation from transmitter to receiver is incorporated into the calibration measurements, which allows us to avoid any explicit (and inevitably inaccurate) modeling of such behavior.

Finally, we can consider different metrics for the model to address. At a basic level, we wish to predict delivery probabilities for a receiver in various scenarios, with arbitrary combinations of senders transmitting packets. (Here, we will limit ourselves to pairs of senders, with the intent for the model to generalize to multiple senders.) Throughput can be viewed as a more useful metric in practice, however, as it conveys how much data will be transferred in a given period of time. We design the model such that both metrics can be computed. In settings where we are able to experiment with various bitrates for transmissions, a stronger notion of airtime utilization may be useful, but as we are limited to 1 Mbps for broadcast packets, we will ignore this metric for now.

4.1 Receiver Component

The task of the receiver component of the model is to predict the probability that packets will be successfully delivered from a sender to a receiver, potentially in the presence of interference from a second sender. The component relies on previous information about delivery probability as a function of RSSI for the given receiver, making the assumption that future trials for a given sender will show the same characteristics as past trials. For such a model to be effective using RSSI, it is important for external interfering energy to be kept to a minimum (or at least a constant level), to the extent that it can be controlled in practice.

Single Sender In the case of a single node transmitting packets, the receiver's predicted delivery probability relies entirely on past behavior. The model simply looks up and reports the observed average RSSI and delivery probability for the given sender and receiver from the calibration data. Under the assumption that the time intervals are long enough (*i.e.*, on the order of a minute), the nodes' behavior should be sufficiently stable to yield accurate results.

Two Sender To predict delivery probabilities from two separate senders transmitting concurrently, we can compute an approximation of the SNIR that the receiver will observe, and then use this to look up a corresponding delivery probability from the calibration data. Specifically, for two senders A and B and a receiver C, we can look up $RSSI_A$ and $RSSI_B$ from the calibration data, which are the signal strengths likely to be observed from the senders in isolation. We can then compute a pseudo-

SNIR, $P = RSSI_A - RSSI_B$, as the ratio of the two signal strengths.[2] Finally, we can interpolate a point for P in the calibration data to determine the likely delivery probability for such a SNIR.

4.2 Transmitter Component

The main goal the transmitter component of the model is to facilitate prediction of throughputs between two nodes, based on how many packets a sender will transmit during a given time period, and how many of those packets will arrive concurrently with packets from another sender. This component can be seen as a slightly higher layer model than the receiver component, as it builds in knowledge of the 802.11 clear channel assessment (CCA) and deferral strategies to determine when a node will transmit packets and when it will choose to defer. In the absence of 802.11, nodes will transmit without deferring (and can thus interfere with each other), allowing the receiver model to directly compute throughputs.

Assuming the use of 802.11's CCA, a node will defer from sending a packet if it senses sufficient energy in the air. For broadcast packets, 802.11 does not employ an exponential backoff, but instead causes each node to pick a random number of slots to wait after each packet, between 1 and $cwmin$, before attempting to transmit. Thus, there is effectively a race after each packet for the next sender to begin transmitting.[3] If the two senders cannot hear each other well, neither will defer and their packets will be sent concurrently.

To compute the number of packets a sender A is expected to transmit, given that sender B is also attempting to transmit, we first determine the probability that the senders will hear each other ($h(B \rightarrow A)$ and $h(A \rightarrow B)$, respectively). We can accomplish this using the basic receiver model, under the assumption that the RSSI threshold for deferring is close to the RSSI threshold for receiving packets. Given that $cwmin$ is 31 for 802.11b, we can compute the probability that A will either win the race to transmit its packet or that it will lose but not hear B and transmit anyway:

$$P(A) = \frac{33}{64} + \frac{31}{64} * (1 - h(B \rightarrow A))$$

A similar equation holds for the probability that B will transmit its packet. Both equations can be translated into expected transmission counts if the attempted transmission rate and time frame are known.

[2]The ratio is expressed as a difference of log values, since RSSI is measured in dBm.

[3]Technically, any node that loses this race and defers will continue its backoff countdown after the next packet, rather than selecting a new number of slots to wait. We disregard this effect to treat each packet as an independent event that does not depend on previous events.

Next, we can compute the fractions of time that A and B send concurrently or in isolation. A will send alone if it wins the race and B hears it (and vice versa for B):

$$T(A) = \frac{31}{64} * h(A \rightarrow B)$$

A and B will send concurrently if they either tie or do not hear each other when the other wins the race:

$$T(AB) = \frac{1}{32} + (\frac{31}{64} * (1 - h(A \rightarrow B))) + (\frac{31}{64} * (1 - h(B \rightarrow A))$$

We can adjust our notion of delivery probability to account for these deferrals, by combining the delivery probabilities for the time periods that the senders transmit alone and concurrently. Finally, these new delivery probabilities can be multiplied by the expected transmission counts to obtain expected throughputs.

Overall, we have seen ways to use and combine existing measurements to predict future behavior in new and more complicated scenarios. The receiver model is able to compute delivery probabilities for multiple active transmitters, and the transmitter model is able to predict throughputs based on 802.11's deferral mechanisms and the relationship between the transmitters.

5 Evaluation

Given the model constructed in Section 4, we wish to judge the validity of our assumptions, as well as the quality of the model. We can evaluate this by looking at the accuracy of the predictions of delivery probability and throughput at a receiver, for pairs of simultaneous senders. If the model is able to use single sender measurements to predict these more complex scenarios successfully, then it can provide value in practice, and the assumptions about network behavior it relies upon can be seen to be valid.

Experimental Setup Our prediction experiments consisted of measurements of both single sender and sender pair scenarios, conducted on the testbed described in Section 2.1. All possible single sender scenarios were measured for the calibration data, but only a hand-picked subset of sender pair experiments were conducted, to limit the combinatorial number of possible experiments. We selected sender pairs that were within range of each other as well as pairs that were not, to ensure a variety of types of interference was evaluated.

Our model was used to predict the delivery probabilities and throughputs for the two sender experiments, given only the single sender calibration data and the knowledge of which senders are active as its inputs. To judge accuracy, we compare the results against the observed data,

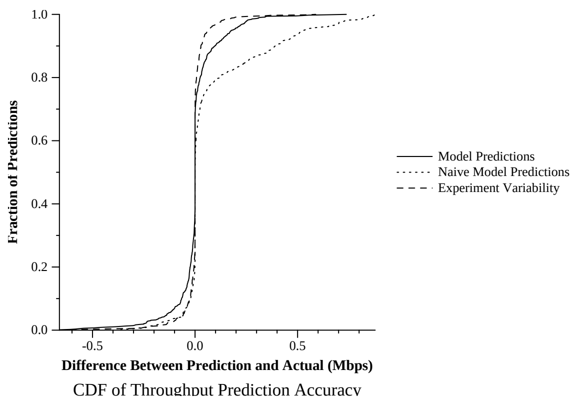

Figure 20: *Throughput prediction accuracy for pairs of senders, comparing the performance of our model and a naive model to the variability between experiments.*

and treat a prediction as correct if it is within +/- 5% of the observed value. This window allows us to gracefully deal with small amounts of variability in the experiments, while still predicting larger effects. For comparison, we also judge the accuracy of a "naive" model, which simply uses calibration data to make predictions without considering the impact of interfering senders. (It is worth noting that even this naive model may outperform many analytic existing models, as it already incorporates information about signal propagation from the calibration data.)

Results After performing 1356 individual predictions of delivery probabilities (across multiple trials with different receivers and sender pairs), the model correctly predicted 86.14% of the cases. In contrast, the naive model correctly predicted only 70.21% of the cases, indicating a substantial improvement for taking interfering signals into account.

If we extend the model to produce throughput predictions using the transmitter model, the overall accuracy numbers change from 86.14% to 78.24% for the model, while the naive model remains at 70.21%. This drop in accuracy is perhaps to be expected, given that multiple delivery probability predictions are incorporated into each throughput prediction. However, this still remains a high accuracy rate, suggesting the viability of this approach.

A more direct comparison is shown in Figure 20, which shows the difference between the predicted throughput and the actual throughput for both models, as a cumulative fraction of the predictions with a particular difference. If the model was perfectly accurate, the graph would show a vertical transition from 0% to 100% of the predictions at a difference on the X axis of 0. However, the graph shows that some predictions are too low (yielding a negative difference), and many predictions are too high (yielding a positive difference). This is especially the case for the naive model, which overestimates delivery probability in cases with interfering signals. Fortunately, the predictions of our model are a reasonably steep curve, indicating its accuracy across all cases.

As an additional point of comparison, this graph also shows the difference between two observed trials of the same experiment, to show the variability inherent in the measurements. No prediction technique can be expected to predict more accurately than this inherent variability. When comparing our model's accuracy to this variability between experiments, we can see that the model's performance is quite reasonable.

6 Related Work

A great deal of work has tried to address the complex behavior of wireless networks, from many perspectives. Wireless simulators such as ns-2 [10] are widely regarded as providing insufficient realism for effective evaluation of wireless networks [7]. Thus, the sources of complexity, potential techniques for addressing it, and the implications for improving wireless protocols have all been studied in some depth, as described below.

Many studies have attempted to characterize aspects of wireless behavior in practice for particular settings, to achieve a more accurate understanding of wireless networks. For example, the Roofnet project has investigated characteristics of packet loss on a city-scale wireless network [2], while the authors of Divert [9] addressed the burstiness of losses on an indoor wireless testbed. Our characterization work is complementary, attempting to reveal a broader picture of the relevant aspects of indoor wireless networking.

Given the inaccuracies of existing wireless simulation, Judd and Steenkiste have proposed an emulation approach for evaluating wireless networks, using hardware emulation of signal propagation [7]. While their technique provides notable improvements for realism and experimental control, the complexity of the construction of such emulators remains an obstacle. Additionally, though our attenuator experiments were partly inspired by emulation, our modeling approach attempts to reduce the number of physical experiments to run (comparable to emulation). In this sense, we only require collecting measurements for a simple calibration phase, and not every configuration of interest. Other work has instead attempted to improve the use of simulation, such as dealing with interference effects through the use of conflict graphs [6]. While this takes a step toward improving standalone simulation results, it does not address how to relate simulation results to actual deployed networks, and it retains many of the challenges of correctly modeling a complex environment.

Finally, many recent papers have attempted to improve wireless protocols based on empirical observations, rather than on simulation or modeling. Divert [9] attempts to reduce packet loss rates in WLAN distribution systems by rapidly selecting between APs, based on an observation of bursty loss rates. Similarly, Extremely Opportunistic Routing (ExOR) [5] attempts to leverage the unpredictable nature of wireless (and an implicit assumption of loss independence between receivers). It exploits this idea to reduce the number of transmissions in multi-hop routing, by forwarding packets in an opportunistic manner. These proposals indicate the room for improvement in wireless protocol design when considering more realistic network behavior. Our work extends this realistic view of wireless networks, and it provides a more accurate modeling approach for easily evaluating such improvements.

7 Conclusions

Through a series of experimental observations of wireless network behavior in both controlled and deployed testbed settings, we have gained an improved practical understanding of wireless packet reception. We have discussed several important implications for modeling such behavior using readily available data:

- The characteristics of local environments differ greatly, even within seemingly uniform settings. These differences result in profoundly different abilities to receive packets.
- Observed signal strength (RSSI) values must be interpreted as unique to a receiver and location, and cannot easily be compared across receivers.
- Observed RSSI values can vary widely over small time intervals, especially during periods of physical activity in the building. However, at larger time scales of a minute or longer, RSSI values tend to appear more stable.
- Thermal noise values reported by cards are quite stable at time scales as small as a second, indicating that most variability comes from the surrounding environment.
- Unknown sources of interfering energy can reduce the utility of RSSI values, suggesting that experiments be performed at times with minimal competing interference.

We have leveraged these observations to build a measurement-based model of packet reception at the physical layer, incorporating the complexity of signal propagation into measurements from actual networks. We expect this approach to be general across testbeds, allowing efficient evaluation of new scenarios with only a basic set of physical measurements.

Together, the experimental observations and measurement-based model provide notable contributions for designing and evaluating practical wireless protocols, in a manner that is simple, usable, and realistic.

8 Acknowledgments

I would like to thank David Wetherall and John Zahorjan for their substantial guidance, feedback, and suggestions throughout this project, as well as their contributions to the model and experiment design. Additionally, Ratul Mahajan and Maya Rodrig contributed to the design and evaluation of the model, and they helped to draft an earlier paper on some of these ideas. Ed Lazowska provided useful comments throughout the process, and helped to keep our group on track.

Finally, I would like to acknowledge the valuable help of two undergraduate students who have worked with our group. Joseph Lai helped develop some of the initial graphing infrastructure for analyzing our results, and Sergei Kaganovsky contributed substantially to the design and execution of the attenuator-based experiments.

References

[1] R. L. Abrahams. Intersil - measurement of wlan receiver sensitivity. http://www.demarctech.com/techsupport/rw-wireless-cards-support/wlan-re%ceiver-test.pdf, Februrary 2000.

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

[3] J. Bardwell. You believe you understand what you think i said... the truth about 802.11 signal and noise metrics. http://www.connect802.com/download/techpubs/2004/you_believe_D100201.pd%f, 2004.

[4] J. Bicket. Madwifi stripped driver. http://pdos.csail.mit.edu/~jbicket/madwifi.stripped/, 2004.

[5] S. Biswas and R. Morris. Opportunistic routing in multi-hop wireless networks. In *HotNets-II*, Nov. 2003.

[6] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and net-*

working, pages 66–80, New York, NY, USA, 2003. ACM Press.

[7] G. Judd and P. Steenkiste. Using emulation to understand and improve wireless networks and applications. In *2nd Symposium on Networked Systems Design and Implementation (NSDI 2005)*, May 2005.

[8] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The click modular router. *ACM Transactions on Computer Systems*, 18(3):263–297, August 2000.

[9] A. K. Miu, G. Tan, H. Balakrishnan, and J. Apostolopoulos. Divert: Fine-grained Path Selection for Wireless LANs. In *2nd International Conference on Mobile Systems, Applications and Services (Mobisys 2004)*, Boston, MA, June 2004.

[10] The network simulator - ns-2. http://www.isi.edu/nsnam/ns/.