Characterizing 802.11 Wireless Link Behavior

Glenn Judd
Carnegie Mellon University
Pittsburgh, PA 15213
glennj@cs.cmu.edu

Peter Steenkiste
Carnegie Mellon University
Pittsburgh, PA 15213
prs@cs.cmu.edu

ABSTRACT

Since wireless signals propagate through the ether, they are significantly affected by attenuation, fading, multipath, and interference. As a result, it is difficult to measure and understand fundamental wireless network behavior. This creates a challenge for both network researchers, who often rely on simulators to evaluate their work, and network managers, who need to deploy and optimize operational networks. Given the complexity of wireless networks, both communities often rely on simplifying rules, which frequently have not been validated using today's wireless radios. In this paper, we undertake a detailed characterization of 802.11 link-level behavior using commercial 802.11 cards. Our study uses a wireless testbed that provides signal propagation emulation, giving us complete control over the signal environment. In addition, we use our measurements to analyze the performance of an operational wireless network. Our work contributes to a more accurate understanding of link-level behavior and enables the development of more accurate wireless network simulators.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: Local and Wide-Area Networks

General Terms

Measurement

Keywords

802.11, wireless network performance

1. INTRODUCTION

Over the last decade, wireless LAN technology has been adopted at an explosive rate. As a result, wireless LANs can now be found everywhere from university campuses to airports, cafes, and private homes. The ubiquity of wireless LANs has lead to a significant amount of research on how to improve the performance of wireless networks and on new wireless applications, such as mesh and vehicular networks. Wireless research is however a challenging endeavor due to the complex nature of wireless signal propagation. Thus, while hardware-based experimentation clearly achieves the best physical layer realism, practical considerations such as ease of use, control, and repeatability have made simulation the dominant evaluation technique. Recent work [11], however, has shown that unless a great deal of care is taken, simulation can lead to incorrect results.

The problem is that a simulator must correctly model all aspects of the system, including the network protocol stack, radio, and signal propagation. This is very challenging, given how quickly wireless technology evolves. Moreover, it is not clear to what degree many simulators have been validated. For example, initial work [10] has shown that the most commonly used simulator - ns-2 - produces results that differ significantly from real-world experiments. Moreover, real-world measurements [1, 16] show that wireless networks exhibit a variety of behaviors, such as link asymmetry, that are not recreated in current simulators. This problem will become worse as researchers start to use more aggressive techniques, such as off-channel reception, to increase network capacity.

In this paper, we undertake a detailed analysis of 802.11 link-level behavior using real hardware and a physical layer wireless network emulator that gives us complete control over signal propagation. This work contributes to a better understanding of the link-level behavior of 802.11 hardware by replacing conventional assumptions and possible misconceptions with actual recorded behavior. We also discuss a number of applications of our measurement results. We discuss the implications of the results on MAC protocol design and we describe how the measurements can feed into the development and validation of more accurate wireless network simulators. Moreover, our results can assist network managers, who currently often have to rely on common wisdom, e.g. "only use channels 1, 6, and 11" or "RTS/CTS is not needed". As an example, we use our results to study the impact of hidden and exposed terminals on the performance of a deployed wireless network.

The remainder of this paper is organized as follows. In the next two sections, we summarize the capabilities of our wireless network emulator and present measurement results for clear channel reception as a baseline for later measurements. Sections 4 through 9 then present our results for the following phenomena: hidden and exposed nodes, packet capture behavior with two competing transmitters, off-channel re-

ception behavior, off-channel interference, multipath, and link asymmetry. Finally, in Section 10, we use the observed link-level behavior to analyze te performance of a production wireless network. This paper is an extended version of [7].

2. EXPERIMENTAL SETUP

Fine grained characterization of wireless link-level behavior requires tight control over signal propagation between the transmitter and receiver. This is achieved using a wireless network based on signal propagation emulation [6], which allows us to conduct network experiments using real wireless cards running in a controlled environment. The only simulated element is the *propagation* of signals between hosts. The wireless hardware, signal generation, signal reception, and software on end hosts are all real.

The operation of our emulator is illustrated in Figure 1. A number of "RF nodes" (e.g. laptops, access points) are connected to the emulator through a cable attached to the antenna port of their wireless cards. On transmit, the RF signal from each RF node is passed into a signal conversion module where it is shifted down to a lower frequency, digitized, and then forwarded in digital form into a central DSP Engine that is built around an FPGA. The DSP Engine models the effects of signal propagation (e.g. large-scale attenuation, multi-path, small-scale fading) on each signal path. Finally, for each RF node, the DSP combines the processed input signals from all the other RF nodes. The resulting signal is sent to the wireless line card of the RF node through the antenna port, after conversion into an RF signal by the signal conversion module. Our implementation supports the full 2.4 GHz ISM band.

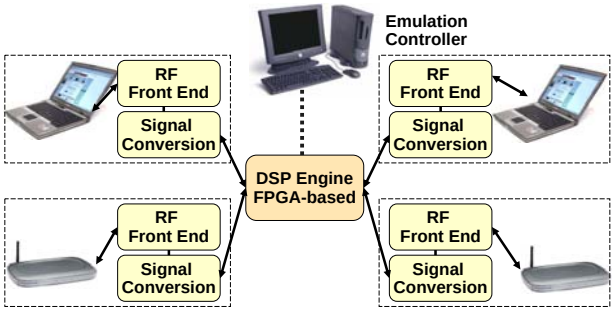

Figure 1: Emulator Implementation

The emulator simultaneously offers a high degree of realism and control. The RF nodes are shielded from each other so that **no communication occurs over the air**. Since all communication between RF nodes occurs through the emulator, we have full control over the signal propagation environment. Channels are modeled at the signal level and signals are generated and interpreted by real radios resulting in realistic system behavior. We have done extensive measurements to verify the precision of the emulator and to validate its results [5].

Emulation is controlled by an Emulation Controller PC which controls RF node movement in a modeled physical environment as well as application behavior on the end hosts. It also coordinates the modeled movement of the RF nodes with the modeling of the signal propagation environment on the emulator hardware. This is done by modifying the parameters of the channel models in real time. Different methods of channel emulation are supported, including the use of statistical models and replay of channel measurements. In the experiments discussed in this paper, the Emulation Controller directly specifies the channel characteristics - in particular path loss between devices - allowing us to construct arbitrary network topologies.

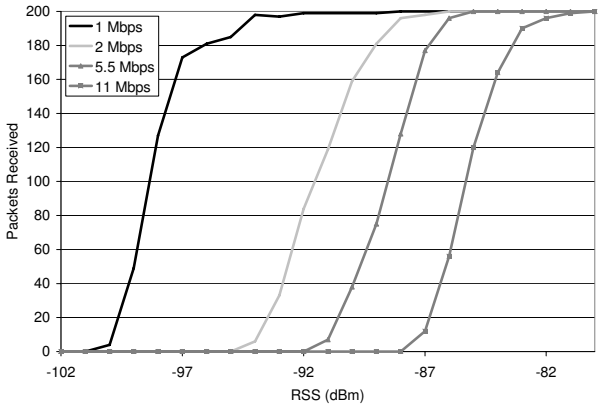

Figure 2: Clear Channel Reception

This paper characterizes wireless link behavior through a series of experiments using three wireless NICs. In some experiments, there is an implicit fourth receiver for which characterization is not necessary. All experiments use Senao 2511CD Plus Ext2 NICs. They are based on the Prism 2.5 chipset, which is one of the more popular 802.11b chipsets in both the research community and deployed networks. While the precise values we report are specific to these cards, our observations should apply to many other hardware configurations. For instance, the robustness of 802.11b's 1 Mbps spread spectrum modulation to interference is a fundamental characteristic of the standard, and all standard compliant hardware should have this feature.

3. CLEAR-CHANNEL RECEPTION

As a reference, we first consider clear-channel reception behavior. The test uses a single transmitter and a single receiver and the emulator varies the RSS (received signal strength) at the receiver from -102 dBm to -80 dBm in 1 dB increments. For each RSS value, the transmitter sends 200 broadcast packets to the receiver and the receiver records the number of successful packet receptions. As broadcast packets do not use link-level retries, this experiment allows us to measure packet delivery rate as a function of RSS. We repeated this test for each of the four 802.11b modulation rates, using the same transmitter and receiver for all tests. Our results are shown in Figure 2; note that different pairs of wireless transmitters and receivers will have results that vary slightly from the results shown in this graph (see Figure 15.) We observe that this receiver is quite sensitive and that, as expected, higher transmission rates require higher signal strengths in order to be received successfully. The noise floor and carrier sense threshold of the Senao cards was measured to be approximately -99 dBm.

4. CAPTURE UNDER DELAYED INTERFERENCE

A central question in wireless networks is understanding

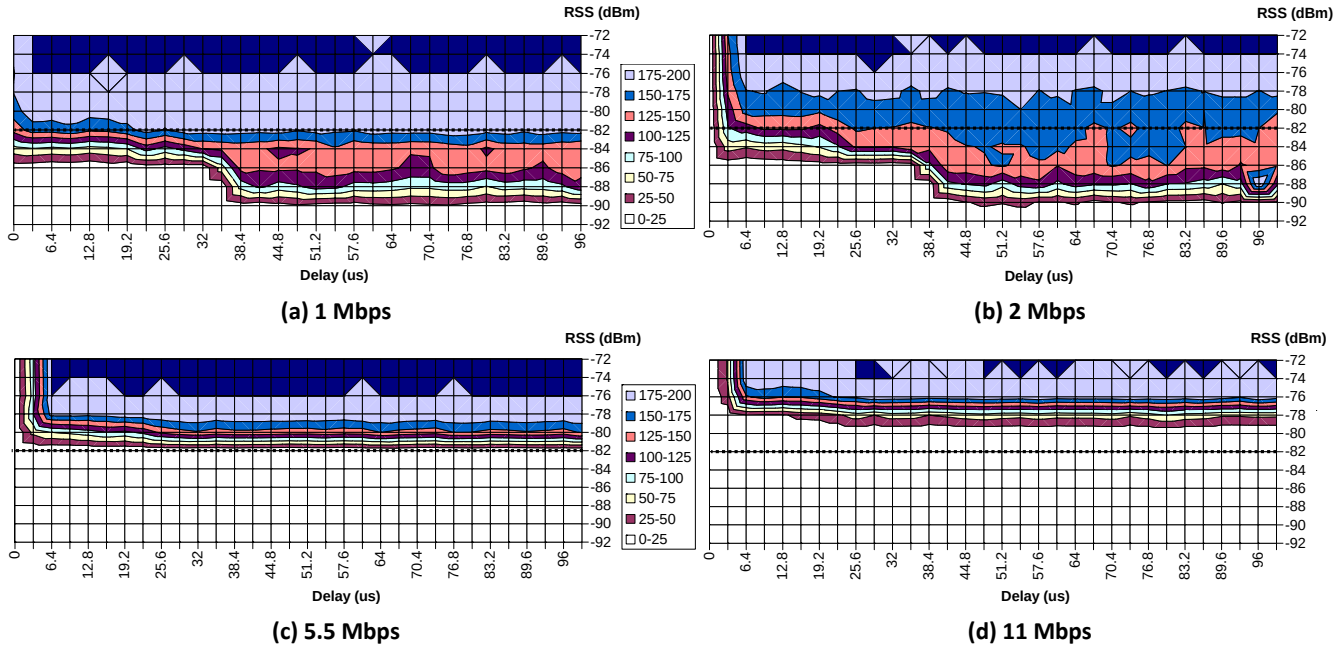

Figure 4: Capture Under Delayed Interference Results

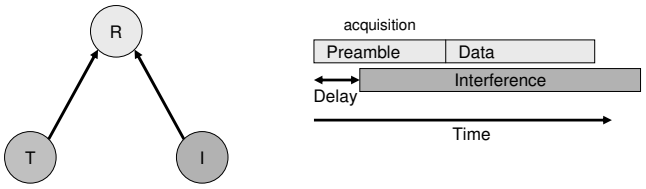

Figure 3: Setup for Capture Under Delayed Interference

what happens when two competing signals arrive at a receiver. Is a packet received and if so, which one? Is there a collision? Simulators have made contradictory assumptions, but little data exists on the behavior of actual hardware. In this section we quantify the effects of timing and received signal strength on a receiver's ability to capture a single desired signal in the presence of an undesired interfering signal. The next section will discuss the effect of received signal strength on the outcome of two competing desirable signals. We use the physical layer network emulator to construct the highly controlled signal propagation topologies necessary to examine reception behavior.

The emulator configuration for the capture experiments (Figure 3) consists of a transmitter T sending traffic to a receiver R. A second transmitter I plays the role of interferer; I constantly sends interfering 1 Mbps 1500 byte broadcast packets that are received at -82 dBm by R. T and I are hidden [22, 3] and cannot hear each other's transmissions. Moreover, we modified the code on the emulator's FPGA to allow R to only hear transmissions from I if: 1) T is actively transmitting, and 2) T's current transmission has been active for a specified delay. Note that we did not explicitly control the arrival time of packets from I. Instead, the interferer I is transmitting almost continuously and we control when I is *allowed* to interfere with T by controlling the channel between I and R.

This setup allows us to investigate the effect of interference timing and signal strength on packet reception. Figure 4 shows the results of our experiments for data rates of 1 through 11 Mbps. We show for different delay-RSS combinations, how many packets R received from T, out of a total of 200 packet sent. The x-axis shows the delay of interference from I with respect to the start of T's transmission in 3.2 microsecond increments between 0 and 96 microseconds. The y-axis shows the RSS of T at R in 1 dB increments between -72 dBm and -92 dBm.

For 1 Mbps (Figure 4(a)), we observe three performance regions, corresponding to delays of 0 microseconds, (0-37] microseconds, and $>$ 37 microseconds. As expected, reception is worst when the interference arrives at the same time as the desired transmission, although some packet reception is still possible. When the interference is delayed by at least 3.2 microseconds, we see a noticeable improvement in performance due to the fact that the receiver has begun acquisition of the desired signal. At delays greater than approximately 32 microseconds, there is a further improvement of approximately 4 dB in reception behavior. This improvement is due to the receiver having acquired the transmission from T. Of particular note is that after signal acquisition, interference can be rejected even if it is stronger than the transmission. We also noticed that when R lost the packet from T, it sometimes would switch to and receive the packet from I, similar to what was observed in [9]. At 2 Mbps (Figure 4(b)), delayed capture behavior is similar to 1 Mbps, though somewhat worse, as expected. The results at 11 Mbps (Figure 4(d)) are very different. While a longer delay in the interference still improves reception, a stronger signal from T is needed, and reception is no longer possible when the desired signal is weaker than the interfering signal. The results for 5.5 Mbps fall in between those for 2 and 11 Mbps.

Conclusion - Our results have important ramifications for MAC design. 802.11's carrier sense mechanism operates without respect to the cell in which a station resides. Not only may this cause transmitters to needlessly defer (an exposed node situation), but transmitters in different cells

(i.e. with different receivers) will tend to synchronize their attempted transmissions in order to limit the time that the medium is experiencing collisions [19]. The above results show that this may be the worst possible timing for packet capture since the very start of a frame is the most vulnerable. Avoiding needlessly synchronizing transmitters in different cells could greatly improve capture performance, and would have negligible impact on the time that the medium might experience collisions. Capture-aware MACs have been considered in different contexts [17, 14].

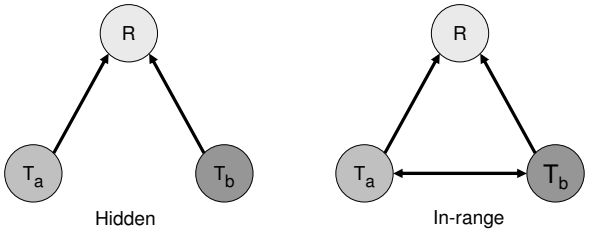

Figure 5: Setup for capture experiments

5. CAPTURE WITH COMPETING TRANSMITTERS

Another question regarding packet reception is how the signal strength of two competing *desirable* signals affects the outcome. We use the configurations of Figure 5 to determined the reception outcome for different RSS combinations from Ta and Tb at R without controlling the interference timing. This was done by having the two transmitters Ta and Tb constantly send broadcast packets to receiver R. At first, the channels are "turned off" so that no packets are actually received at R. We then simultaneously turn on the channels by setting the attenuation so that we get the desired RSS value at R from each transmitter. After a fixed time interval, we shut off the channels from Ta and Tb to R and we record how many packets R received from each transmitter. We measured all combinations of RSS values from Ta and Tb at R between -102 and -72 dBm in 1 dBm intervals and for all 802.11b transmission rates. In the "hidden" configuration, we did not allow Ta and Tb to hear each other's transmissions while in the "in-range" setup, we set the RSS from Ta to Tb at -80 dBm and vice versa, so that Ta and Tb will always hear each other's transmissions.

Figure 6 shows our results. In each of the graphs the z-axis is the number of packets received from both Ta and Tb at R. In many RSS combinations, however, packets were only received from one or the other; the regions where one source dominates are labeled on the plots.

In all in-range cases (the graphs on the left), we found that when the RSS at R from both Ta and Tb was high, CSMA did a good job of allowing the two nodes to share the medium and only a small number of collisions occurred. As a result, the throughput is close to the capacity at receiver R. Note that the region where both nodes Ta and Tb are within range shrinks as the transmit rate increases; this is consistent with the clear channel results in Figure 2. As expected, when one transmitter was out of range of R and the other was in range, the number of packets received for the in-range cases was roughly half of the channel capacity since the two transmitters defer to each other's transmissions, irrespective of the number of packets successfully received at R. In an actual network, this would only occur when the out-of-range node was sending to a receiver other than R (or broadcasting) since unicast communication requires acknowledgement of successful reception. For these "exposed node" cases, the in-range node may be needlessly deferring since the out-of-range node is not communicating with the same receiver and its signal may be too weak to interfere with the in-range node's signal.

An important question is what happens when transmissions from two nodes overlap in time at a single receiver. The "hidden node" configuration tests investigate this question. In hidden node situations, Ta and Tb send at full rate since they are out of carrier sense range. We see in the graphs on the right that when the RSS values at R from Ta and Tb are similar, R correctly receives either very few or no packets as a result of collisions. The size of the collision regions increases with higher transmit rates. For 1 Mbps, collisions only occur for a very narrow range of signal strengths where the RSS values at R from Ta and Tb are nearly identical. Even then, some packets are still received correctly. Hence, at the 1 Mbps rate, deferring transmission based on carrier sense is likely unnecessary. At higher transmission rates, the range over which collisions occur is larger, especially for the 5.5 and 11 Mbps rates. We also see that in some cases, collisions completely prevent communication. However, even at those higher rates, the collision region is relatively narrow, i.e. the RSS values from the two senders must be relatively close for collisions to prevent communication. As soon as the difference between the two RSS values increases, one of the two packets is likely to be received as a result of the capture effect, as discussed in the previous section. The trend across the different transmit rates is consistent with the measurements in Figure 4.

Conclusion - The measurements in this section have shown that for low transmission rates, collisions occur only when the signal strengths of the competing signals at a receiver are nearly equal. Hence, packets sent at low rates, e.g. management and control packets such as beacons, RTS, CTS, and ACK, are very robust to interference. At higher rates, however, a broader range of received signal strengths will interfere. Nevertheless, even high modulation rates will very often capture packets in spite of interference. Hence, deferring transmission due to an interfering source below the capture threshold is not necessary and hurts network performance. An important implication for simulators is that realistic capture behavior cannot be recreating using the fixed threshold that is commonly used, but it requires a realistic model such as the data in Figure 6. In Section 10 we will also look at the implications of our capture results for deployed 802.11 networks.

6. OFF-CHANNEL INTERFERENCE

In the US, eleven 802.11b channels are available in 5 MHz increments from 2.412-2.462 GHz. Each 802.11b channel occupies 22 MHz so a total of three 802.11b signals can coexist - on channels 1, 6, and 11 - without interfering. Ideally, adjacent 802.11b cells would utilize non-overlapping channels. Unfortunately, it is frequently impossible to deploy an 802.11b network without placing some adjacent cells on the same frequency. For this reason, some have advocated using four channels despite the fact that there would be some signal overlap [13]. While there is some evidence to support this idea, there has not been a carefully controlled study of the impact of off-channel interference on real hardware.

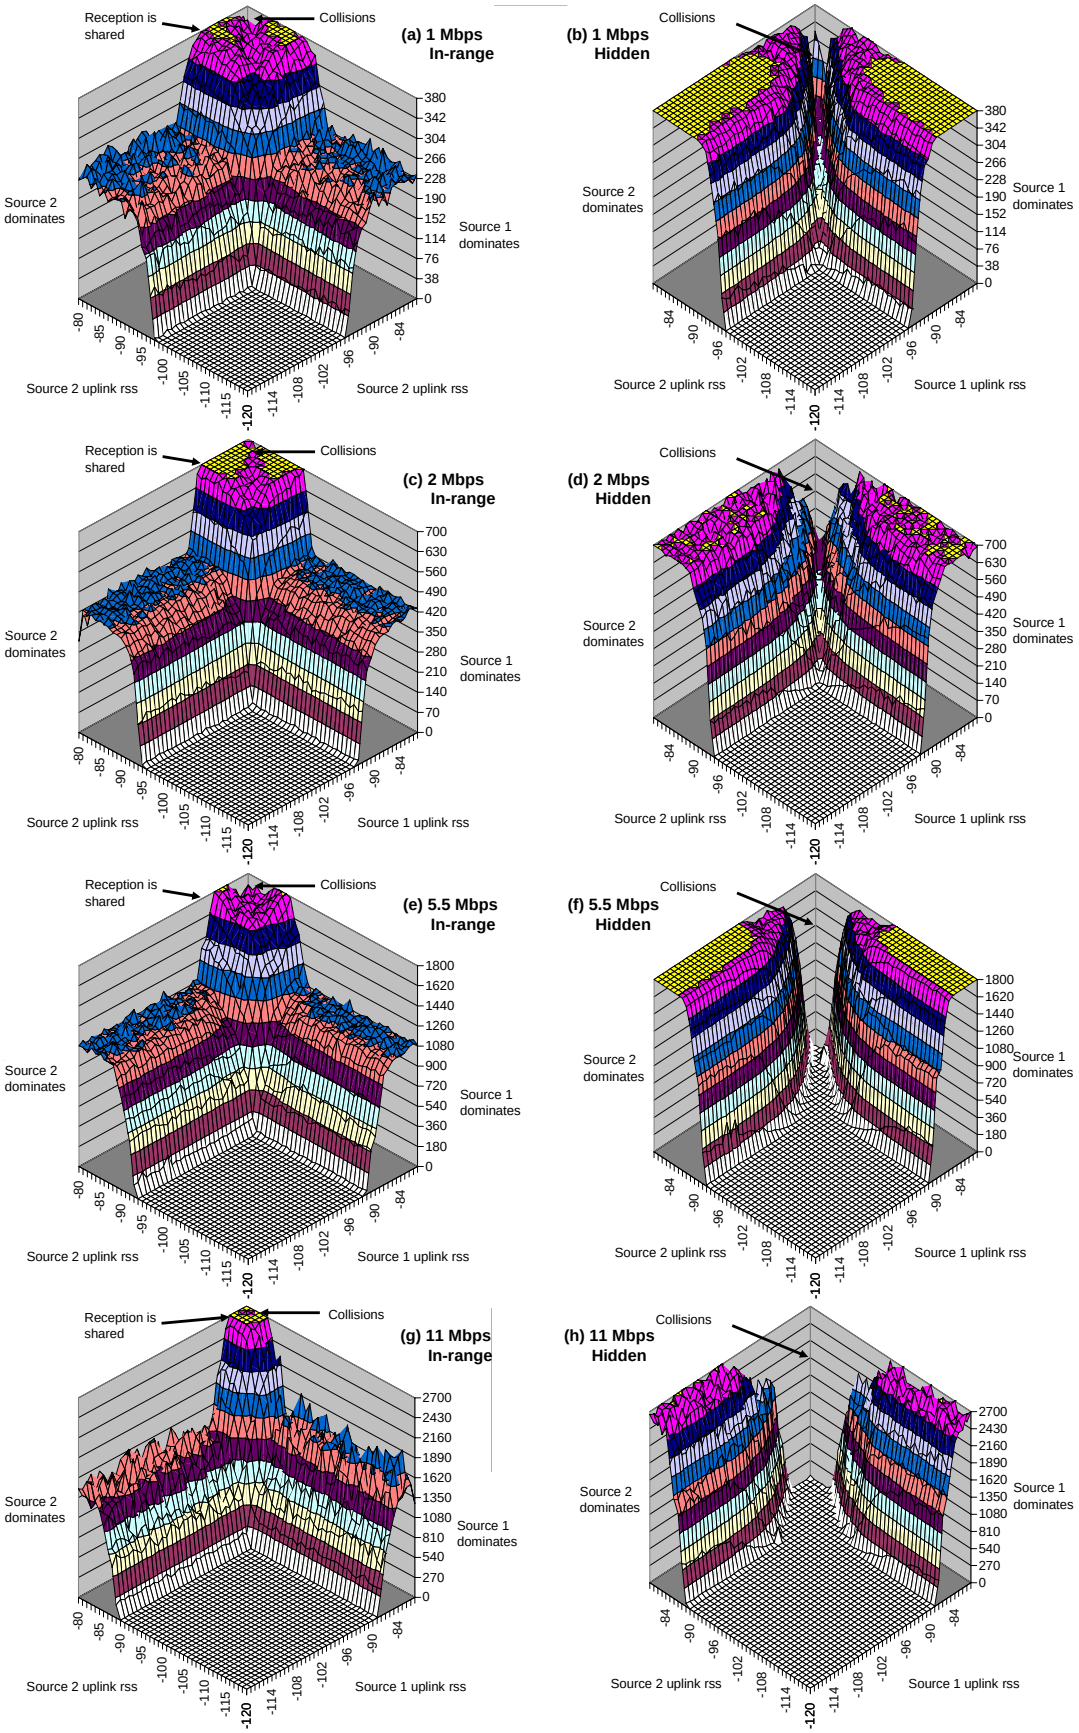

Figure 6: Packet Capture Results

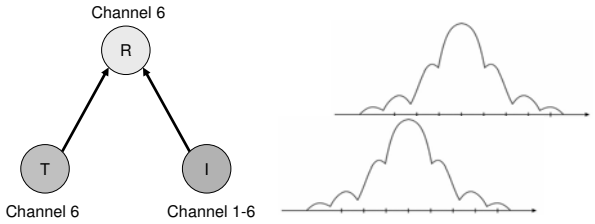

Figure 7: Off Channel Interference Setup

In order to quantify the viability of this 4-channel proposal and to understand the impact of off-channel interference on successful packet capture, we measured the impact of off-channel interference on packet reception using the setup shown in Figure 7. In this experiment we have two transmitters T and I and a single receiver R. Both T and R are on channel 6; I plays the role of an off-channel interferer on channels 1 through 6. As in the delayed capture test discussed in Section 4, the interference from I is controlled so R only hears the signal from I some specified delay after R begins to hear a packet from T. For this test, we use two delay values, 0 and 384 microseconds i.e. immediately, or well after packet acquisition. For each channel that I is placed on, the RSS at R from I is held constant at -82 dBm while the RSS at R from T is varied between -72 and -102 dBm. For each channel-RSS combination T sends a series of packets to R and R records how many were received successfully. We used broadcast packets, so no retries took place. We repeated this test for all four 802.11b modulation rates.

Our results are shown in Figure 8. For all tests where interference was prevented until well after packet acquisition (graphs on the right), we observed that the impact of interference from channels 1, 2, and 3 was low and virtually identical. Channel 4 degraded performance by approximately 4 dB, while channels 5 and 6 degraded performance more significantly. For tests where interference was allowed to occur at the start of packet reception, the interference of channels 1, 2, and 3 was still nearly identical though channel 3 was slightly worse in some cases. Interference from channels 4-6 was much more significant in this case.

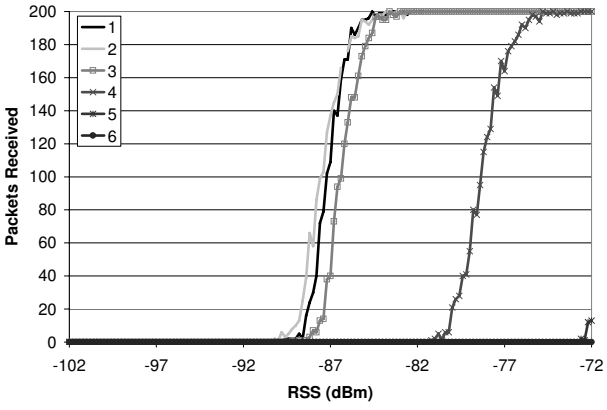

Figure 9: Off-channel Interference, 11 Mpbs, large delay, -72 dBm Interference

To investigate the effect of stronger interference, we reran the 11 Mbps large delay tests with interference of -72 dBm instead of -82 dBm. Figure 9 shows the results. When comparing with Figure 8(h), we see that the higher interference has a strong impact when the interferer is on channels 4-6. When the interferer is on channels 1-3, however, interference impact is only 2 dB stronger than it was with -82 dBm of interference.

Conclusion - These tests show that a well-designed receiver can cope quite well with off-channel interference that is at least three channels away. This is an important result as it demonstrates that the 802.11b five channel separation that is typically used is overly conservative. Using four channels in place of the typical three can reap nearly a 33% improvement in capacity.

7. OFF-CHANNEL RECEPTION

A recent observation that some off-channel packets can be received has lead to the proposal to leverage off-channel communication for purposes such as bridging between channel regions in multi-hop networks. The utility of this proposal, however, clearly relies on the efficacy of off-channel communication, which to our knowledge has not been analyzed in a controlled manner. To fill this void, we designed an experiment to characterize off-channel reception. We use a single transmitter-receiver pair with the transmitter fixed on Channel 6 while the receiver is varied from channels 1-6. Note that there is no interference at all in this test. For each receiver channel, we varied the RSS at the receiver from the transmitter between -102.0 and -72.0 dBm. For each channel-RSS pair, we sent 200 broadcast packets from the transmitter to the receiver and measured how many were received. We repeated this test for all 802.11b transmission rates.

Figure 10(a) shows the results of this test for 1 Mbps. At 1 Mbps, off-channel communication appears to work, although we observe increasing isolation as the channel separation increases. At 2 Mbps, however, this scheme begins to break down as shown in Figure 10(b). For this modulation, reception is still possible, but only when the signal is strong and even then we never achieve a packet delivery rate higher than 10%. Also, the fact that packet delivery rate is not monotonically increasing with RSS suggests that signal distortion may be occurring. At 5.5 and 11 Mbps, things are even worse. Up through -72.0 dBm, we received no off-channel packets as shown in Figures 10(c) and 10(d).

The trouble with off-channel reception likely lies with several features of the receiver. For instance, when the receiver filter is applied off-center with respect to the modulated signal's center frequency, the signal is distorted in time. Also, the receiver's acquisition circuitry may not be able to acquire the signal. 1 Mbps uses BPSK modulation which is somewhat robust but all other bit-rates use QPSK modulation which is much more susceptible to these effects.

Conclusion - Our results show that the opportunities for off-channel reception are limited. In particular, off-channel reception is only effective at the lowest transmission rate of 1 Mbps and only when the received signal is extremely strong. Thus, while this technique may prove useful in some unique circumstances, it is unlikely to be broadly applicable.

8. MULTIPATH CHARACTERIZATION

To evaluate the impact of multi-path on wireless link behavior we ran a sequence of experiments on the wireless emulator using the signal propagation environment shown in Figure 11. It uses two nodes connected by a wireless

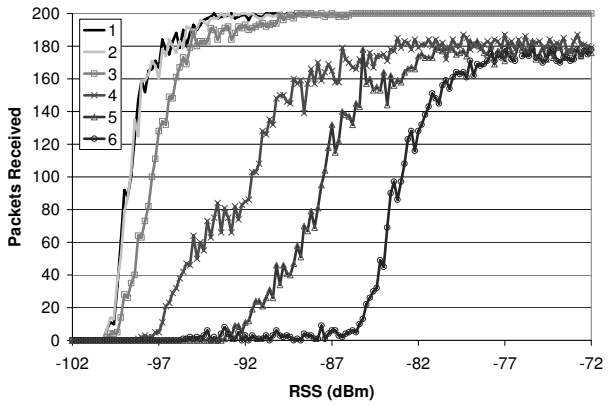

(a) 1 Mbps - no delay

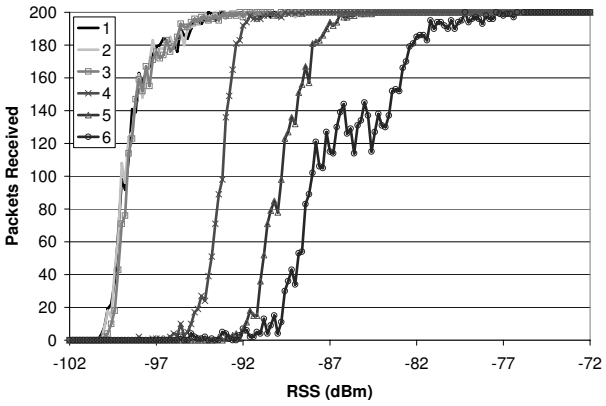

(b) 1 Mbps - large delay

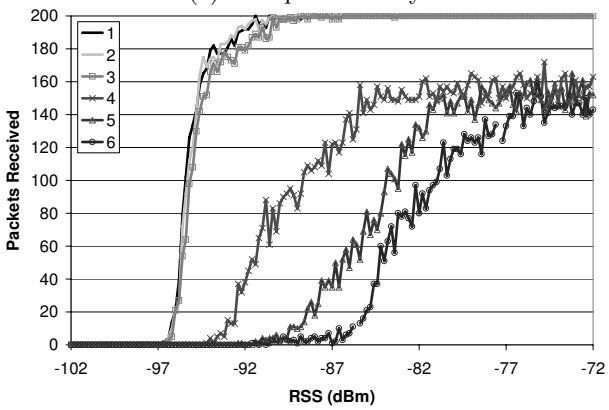

(c) 2 Mbps - no delay

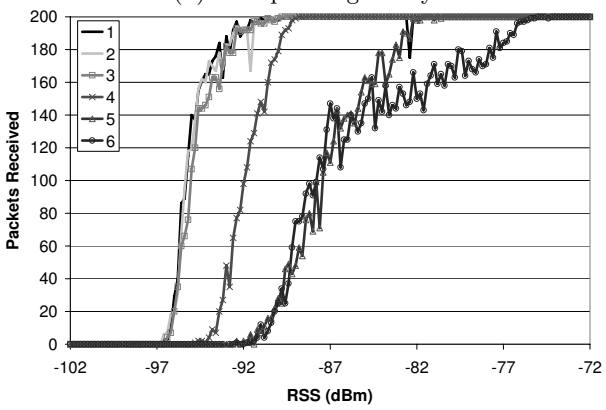

(d) 2 Mbps - large delay

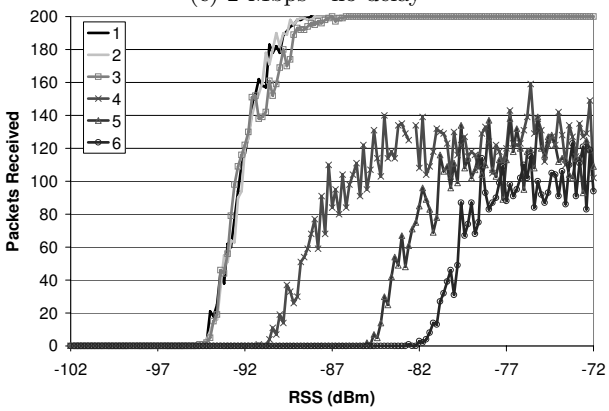

(e) 5.5 Mbps - no delay

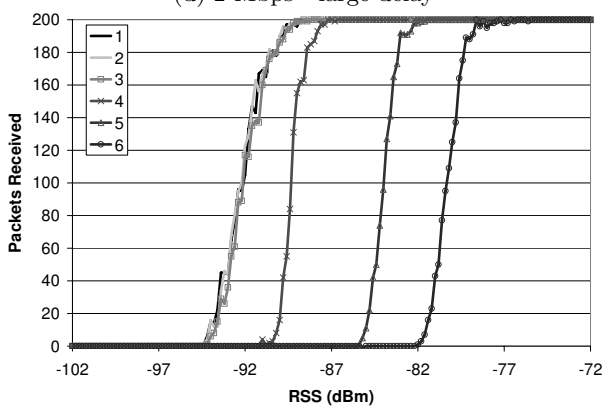

(f) 5.5 Mbps - large delay

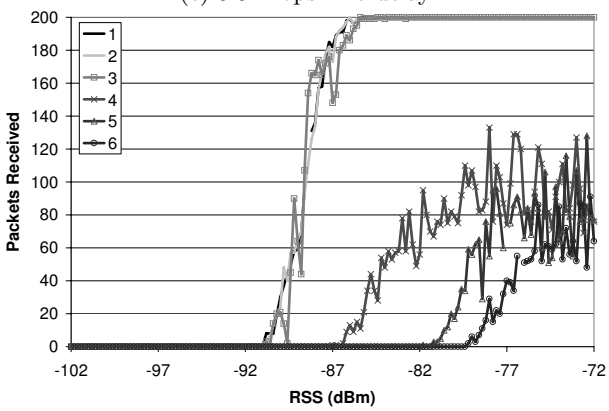

(g) 11 Mbps - no delay

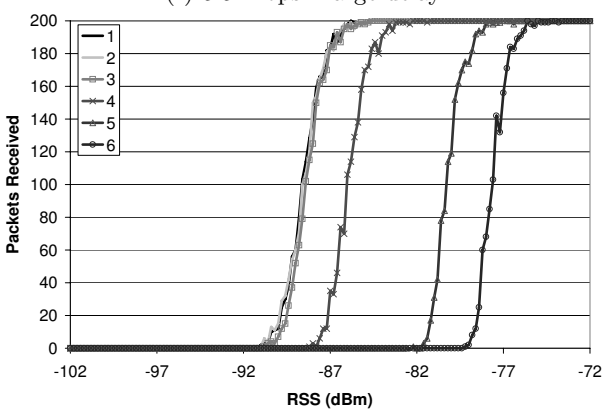

(h) 11 Mbps - large delay

Figure 8: Off-channel Interference Measurement Results

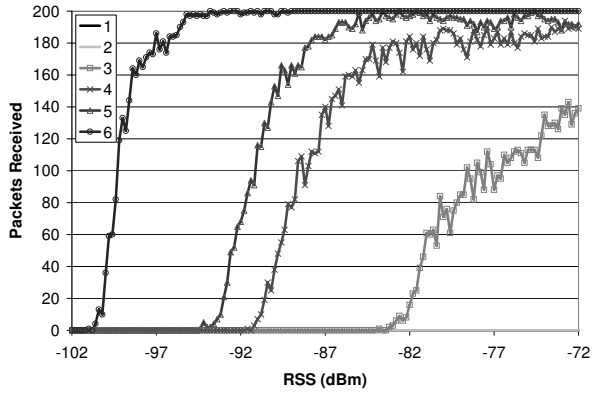

(a) 1 Mbps

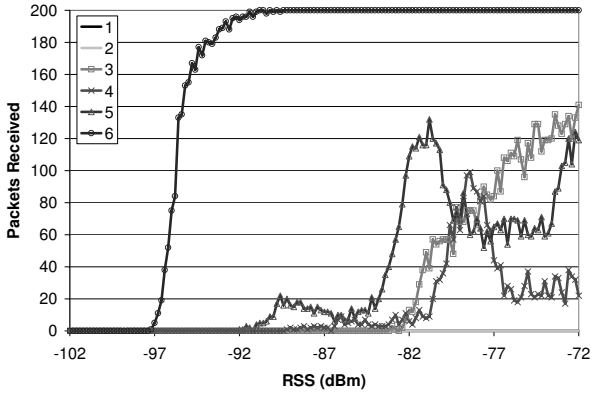

(b) 2 Mbps

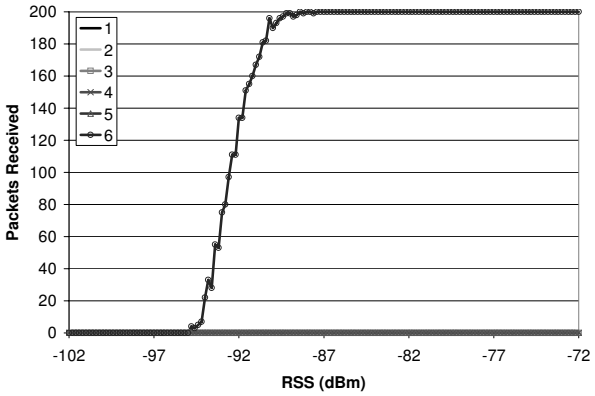

(c) 5.5 Mbps

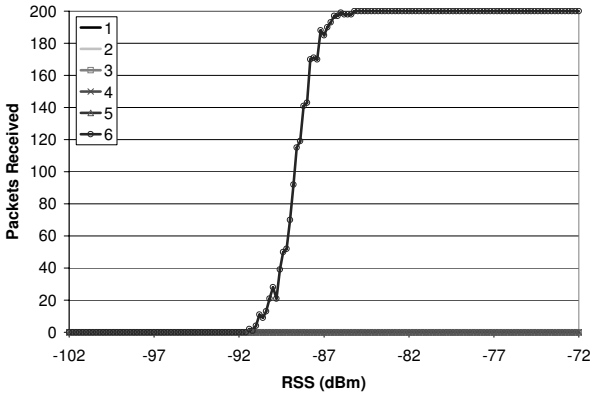

(d) 11 Mbps

Figure 10: Off-channel Reception

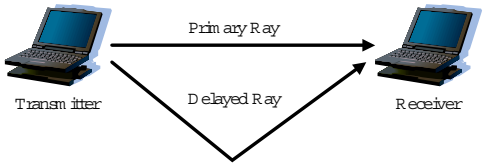

Figure 11: Setup Multi-path Experiments

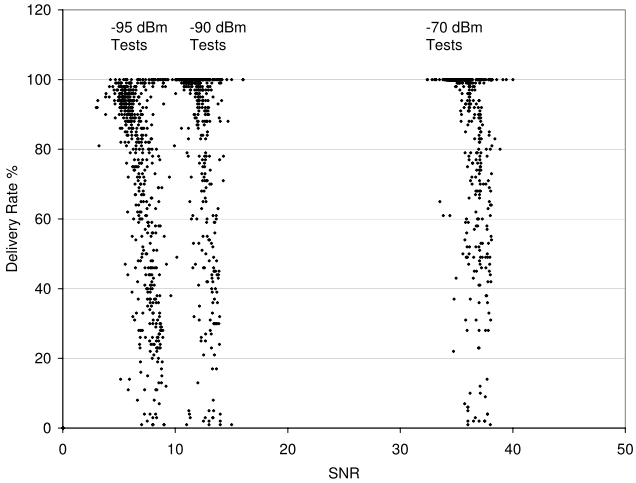

Figure 12: Impact of Multi-Path on Packet Delivery Rate

channel that consists of two paths. We can control both the path loss of the two channels and the relative delay between the two paths. For three different primary signal strengths (-70 dBm, -90 dBm, and -95 dBm) we ran experiments for different attenuations and delays of the secondary ray and measured success rate [6]. Figure 12 shows the results of the multipath experiment: each point represents the delivery rate for one combination of primary signal strength, delayed signal strength, and delay spread. We see that the delivery rates as a function of SNR exhibit a large variation. In fact, although we use only three different values for the primary signal strength (corresponding to three discrete path loss values), the packet success rates are fairly random.

This experiment was motivated by observations made in the Roofnet testbed, an outdoor mesh network in Cambridge, MA. Measurements in Roofnet showed that there was virtually no correlation between the packet success rate of a wireless link and the RSSI measured by the receiver [1]. After a number of other possible causes for this result had been eliminated, the wireless network emulator was used to evaluate whether multipath effects could be responsible. The similarity between the Roofnet measurements [1] and the results in Figure 12 suggests that multi-path is indeed the cause of the behavior observed in Roofnet. Our results also suggest that received signal strength, and by implication RSSI, is a poor indicator of packet delivery rate when significant multipath is present.

To better understand the results in Figure 12, we later used the same scenario (Figure 11) to more carefully study the impact of delay. For a fixed primary signal strength, we varied the attenuation of the secondary path relative to the primary path between 0 and 10 dB in steps of 0.2 dB. We also changed the relative delay between the two paths between 0 and 2.22 us in 0.0185 us increments. Figure 13 shows the results for all four transmit rates of 802.11b (1, 2, 5.5, and 11

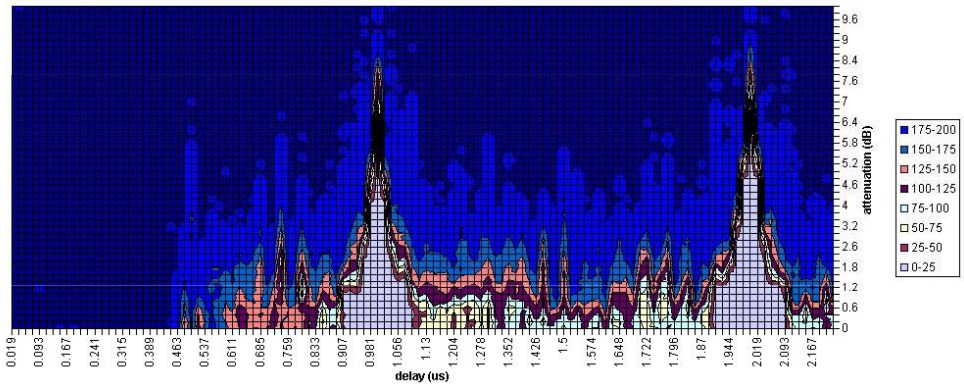

(a) Result 1 Mbps

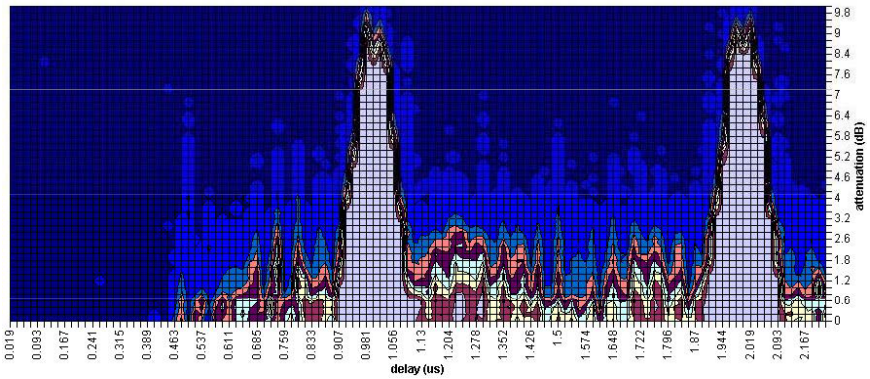

(b) Result 2 Mbps

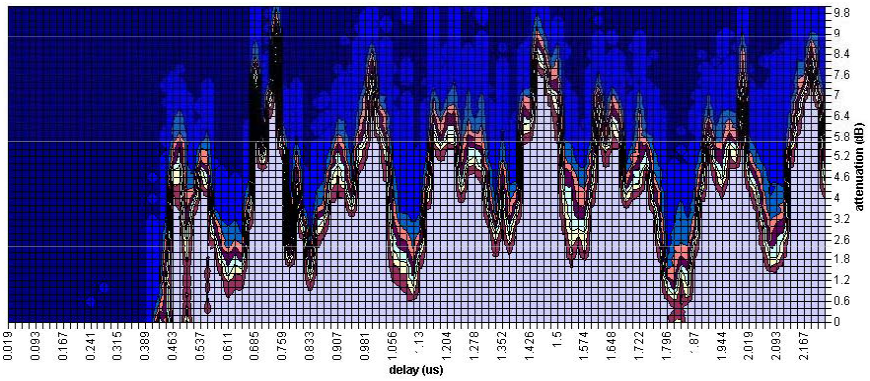

(c) Result 5.5 Mbps

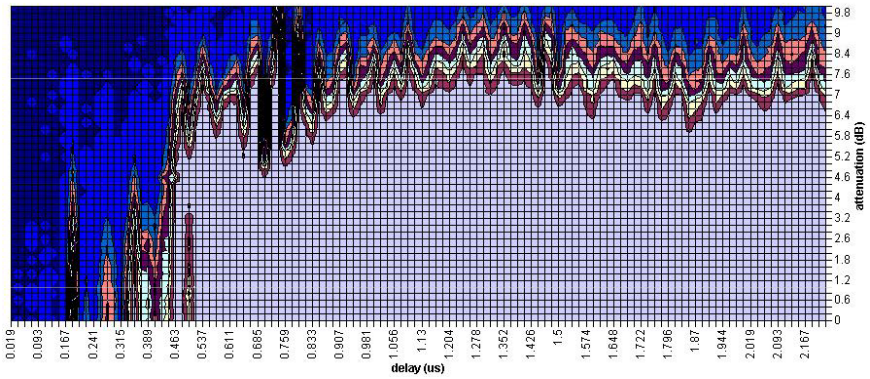

(d) Result 11 Mbps

Figure 13: Multi-path Results as a Function of Relative Path Signal Strength and Delay

Mpbs). It shows the number of packets received as a function of path delay (x-axis) and secondary path attenuation (y-axis). We observe, not surprisingly, that packet success rates go up as the transmit rate is reduced. We also see that a larger attenuation of the secondary path general improves packet success rates as the channel becomes more similar to a clear channel. Finally, we note that for the 1 and 2 Mbps rates, performance is very poor for a delay of 1 and 2 microseconds, which is when the symbols (sent at a rate of 1 Mbps) on the two paths are offset by 1 or 2.

A interesting point is that for relative path delays of less than about 0.45 microseconds, the packet success rates are uniformly high, i.e. multi-path appears to have almost no effect on performance. The reasons turns out to be simple: the wireless network cards have dynamic equalizers to fight multipath [18]. The ability to compensate for multipath is however limited by the depth of the their pipeline, in this case 0.45 microseconds or a difference in path distance of about 150 meters. The manufacturer confirmed that this was the case: since the cards were designed for indoor use, this distance was considered sufficient for normal operating conditions. However, Roofnet is an outdoor testbed with links that cover much longer distances, hence the degradation of performance due to multipath.

Conclusion - We can draw two important conclusions. First, when used in an indoor environment, well design wireless cards deal well with multipath. Second, in such environments, RSSI is indeed a reasonable predictor for received signal strength and it will generally be a good indicator of the likely packet success rate.

9. LINK ASYMMETRY

Several research groups [16, 12, 10] have independently observed asymmetric wireless link behavior. In particular, people have observed that the packet delivery rate from node A to node B may not the same as from B to A. Nevertheless, there has not been an investigation into the source of link asymmetry. In this section we present a controlled analysis of the possible causes of link asymmetry. Let us first however mention a "non-cause" for link asymmetry: **asymmetric signal propagation**. Asymmetric signal propagation is physically impossible according to the reciprocity theorem [21], which states that *if the role of the transmitter and the receiver are interchanged, the instantaneous signal transfer function between the two remains unchanged.* Nevertheless asymmetric signal propagation is sometimes posited as an explanation for link asymmetry.

Transmit power variation - Using asymmetric transmit power on a link can cause asymmetric packet delivery rates due to the disparity in received signal strength. As real wireless networks are typically composed of a heterogeneous mix of devices, real networks will likely have asymmetric links due to asymmetric transmit power. Link asymmetry has been observed, however, even when the same model card is used on both nodes. To assess transmit power variability, we measured the transmit power of 11 different Senao cards using a spectrum analyzer. We added 0.5 dB to the measurements to account for pigtail loss (an estimate).

Figure 14 shows the average of 23 individual measurements for each card and the computed 95% confidence intervals. We observed that the cards fell into two distinct sets A and B. The cards in set A had an averaged transmit power close to 23 dBm with very little variation. In contrast,

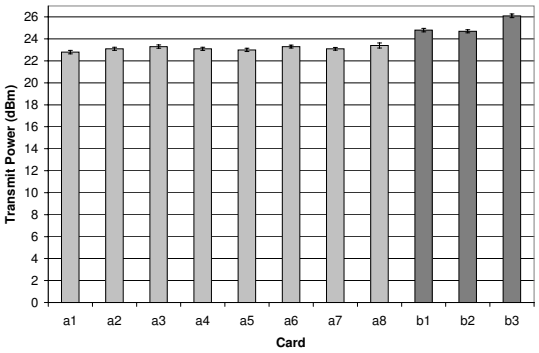

Figure 14: Senao Card Power

cards in set B had a higher transmit power and exhibited more variation. While the 11 cards were marketed, sold, and labeled as identical they were purchased at different times from different vendors, and they have MAC addresses that fall into two distinct ranges corresponding to sets A and B. Clearly even cards that appear outwardly to be identical, may actually be different and have different transmit power. We conclude that transmit power variation is one likely contributor to link asymmetry.

Receiver noise floor variation and quality variations in the transmitter and receiver - The noise floor of the receiver is determined largely by the performance of the low noise amplifier (LNA). LNAs are designed to amplify the weak signal received at the antenna into a stronger signal that can be processed without introducing much noise. However, anything that touches a signal adds some degree of noise. The figure of merit for LNAs is their "noise figure", measured in dB. It quantifies how much noise they introduce into the signal. Moreover, quality variations in transmit modulation and in the receiver, including factors such as linearity, can affect fidelity of the transmitted signal and signal acquisition on the receiver.

Since we cannot separate these factors without dissecting the radio hardware, we use a single experiment to quantify the combined effects of these three factors on link asymmetry. Specifically, we measured the pairwise packet delivery rate between all possible pairs of four wireless cards using 2 Mbps broadcast packets. In this case, we used coaxial cable and a variable attenuator to vary the transmit power between these nodes. We corrected for transmit power variation in order to isolate the desired effects. We varied the received signal strength between -80 and -98 dBm. Figure 15 shows the results. We observed approximately 3 dB of variation over all of the links that we measured.

Antenna Diversity - Some degree of link asymmetry could arise when different or multiple transmit and/or receive antennas are being used on one or both ends of the link. The degree of asymmetry will depend both on the algorithms used to exploit antenna diversity and the channel conditions. Note however that asymmetry has been observed even in cases where no antenna diversity exists.

Interference variation - A final potential contributor to link asymmetry is interference level variation. Interference variation is likely to contribute to asymmetry in a way that is highly site-specific and variable over time and evaluating its impact requires a careful study of interference at a specific site and under specific conditions. An important distinction of interference compared with the previous factors

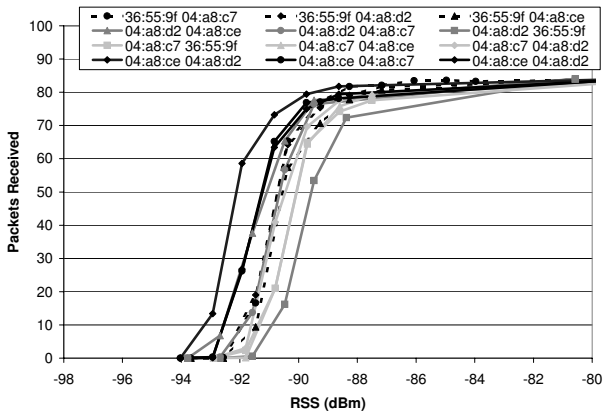

Figure 15: Packet Delivery Rate Variation

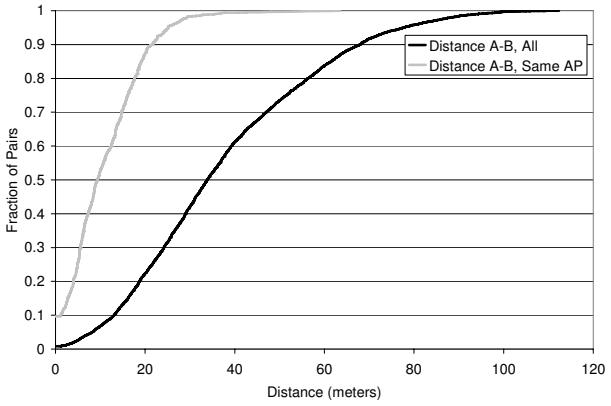

Figure 16: Distance CDF for Operational WLAN

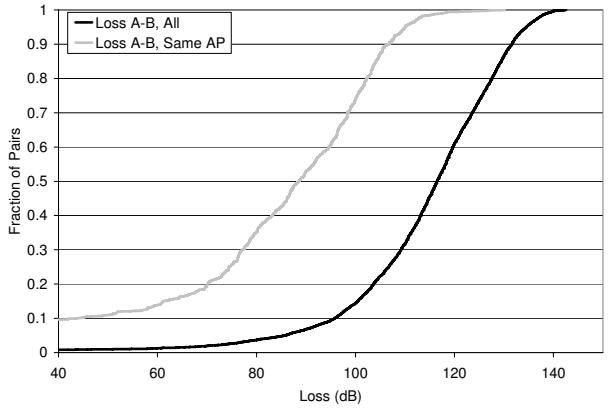

Figure 17: Path Loss CDF for Operational WLAN

is that it is typically not constant. Most sources of interference e.g. competing 802.11 traffic, non-802.11 data traffic, cordless phones, microwave ovens, etc. are bursty on some time scale. Several researchers have observed asymmetric links that have a fairly consistent constant bias. Thus, in at least some cases it is unlikely that bursty interference is the cause of link asymmetry.

Conclusion - We have discussed several potential causes of link asymmetry and we have shown that several of these are contributing factors. While in some cases, one factor such as transmit power asymmetry may be the dominant factor, in many cases, we expect that link asymmetry may be the result of the additive effects of several causes. Importantly, we have shown that link asymmetry can exist even when using homogeneous hardware and when external interference does not play a role. Thus, protocol designers should consider link asymmetry even when hardware is uniform.

10. WLAN PERFORMANCE ANALYSIS

The results in this paper can be used to build more accurate models for when packets are received by commercial 802.11 cards. These models can then be used by both researchers and network managers. For example, the data collected in Sections 4 through Section 9 can be used to improve the accuracy of simulators, as has been explored by others for packet capture [9]. Alternatively, the insights provided in wireless links can be used to understand and improve the performance of operational networks. As an example, we now use the reception characterization of Section 5 to analyze the behavior of a deployed 802.11b network in the Tepper School at CMU. The network consists of 17 access points on channels 1, 6, and 11, and it covers a single large campus building. We are particularly interested in gaining insight into the issue of hidden and exposed nodes: why do WLANs seem to work well despite the fact that RTS/CTS is rarely used?

We constructed a radio map of the building by sampling received signal strength from access points throughout the building, and storing the physical location of each sample. For the sake of this analysis, we considered each node to have the same transmit power which, as discussed earlier, is only approximately correct. We then analyzed the likelihood of hidden terminals and exposed nodes as follows. We generated a random distribution of 400 clients within this building taking into account the likelihood of a particular location's occupancy, e.g., clients are much more likely to be located in lecture halls than offices. We used the actual observed access point locations and channel assignments. Each client picked a random recorded set of access point signal samples at its location in the radio map. Each node was then associated with the access point having the strongest signal. In our analysis, client to access point path loss is computed directly from radio map measurements. Between clients, however, we have no direct measurement, so we model path loss using a log distance path loss model [18] with a d0 of 1.0 meter, pld0 of 40.0 dB, and a path loss exponent "n" of 5.0. Figure 16 plots the CDF of client pair distances for all pairs and also for pairs associated with the same access point. Figure 17 plots the CDF of path loss for all client pairs and also for client pairs associated with the same access point. Both the distance and path loss results are for a single execution of our analysis, but other runs produced very similar results.

We then looked at each client (called A) in the network and analyzed its pairwise interaction with all other clients (called B) in the network to identify possibly hidden or exposed terminal scenarios. Specifically, we looked for the cases depicted in Figure 18(b), (c), and (e). If A and B are associated with the same access point then they must be able to communicate with it, and we have a hidden terminal scenario if they are out of carrier sense range (Figure 18(b)). If A and B are associated with different access points we need to consider two cases. First, if A and B are in carrier sense range of each other, but B does not interfere with A's transmissions to its access point, then we have an exposed terminal scenario (Figure 18(d)). Second, if A

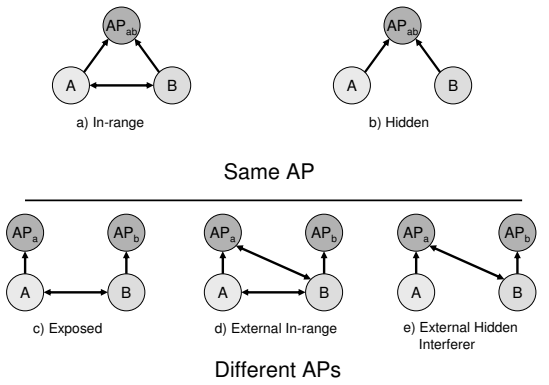

Figure 18: Infrastructure Topologies

Table 1: WLAN Performance Analysis Summary

Total Pairs	159600
Same AP Pairs	12230
Hidden Pairs	406
Exposed Pairs	11438
External Interferer Pairs	34374

and B are out of carrier sense range from each other and B's transmissions can interfere with A's transmissions to its access point (Figure 18(e)). In cases Figure 18(a) and (d), carrier sense will avoid interference occuring. Note also that we must only consider A's interaction with its access point. B's interactions are considered when it is "A".

Based on the client-AP path loss measurements and the client-client path loss estimates (Figure 17), we analyzed how often the interactions in Figure 18 were found in a single run of our analysis (other runs were quite similar). The results are summarized in Table 1. Clearly, hidden nodes were very uncommon. The reason can be found by analyzing Figures 16 and 17. The wireless network in this building is fairly dense, so nodes associated with the same AP tend to be quite close to each other. To be out of range requires a loss of 115 dB which occurred for very few pairs associated with the same access point. Exposed pairs and external interferer pairs, however, were much more common.

Next we analyzed the impact that hidden nodes might have on performance, since hidden nodes do not necessarily result in failed transmissions. For each hidden pair, we used the data obtained in Section 5 to estimate the probability that A's transmissions would be received by its access point despite the fact that it is interfered with by a transmission from B. We used the path loss measured in the radio map to compute the RSS at A's access point for both A and B and then computed the capture probability from the data in Section 5.

Figure 19 shows the result for both 1 and 11 Mbps transmissions. At 1 Mbps, very few of the hidden pairs are likely to have high collision probabilities. At 11 Mbps, however, there is a fair chance for collision. To explain this result, we show in Figure 20 the CDF of the difference in RSS from the two nodes at the AP. We see that the RSS values of nodes A and B at A's access point are often very similar. As a result, the hidden nodes will often fall in the collision regions marked in the graphs on the right in Figure 6. This region is quite small at 1 Mbps (Figure 6(b)), but is substantially larger at 11 Mbps (Figure 6(h)). In practice, we

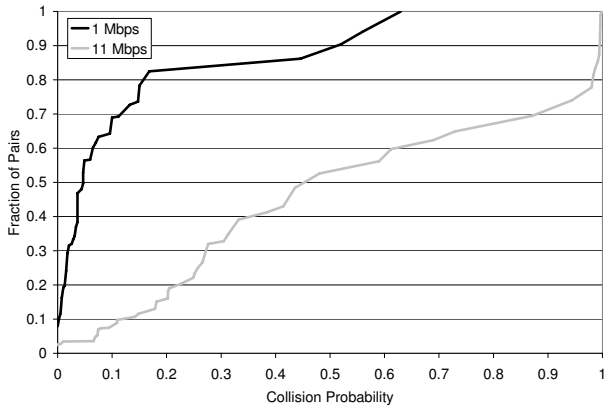

Figure 19: Hidden Node Collision Probability, 1 Mbps vs. 11 Mbps

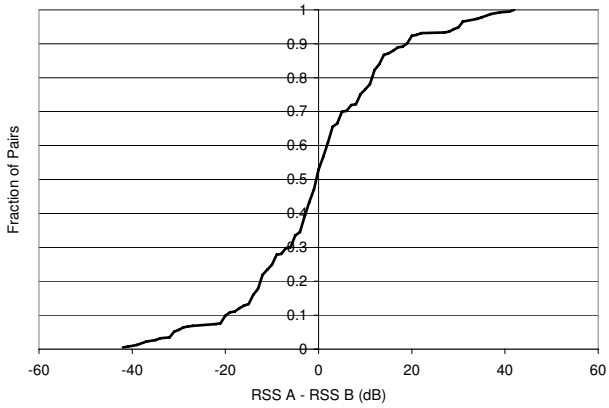

Figure 20: Hidden Node RSS Difference CDF

expect that most nodes in this network would communicate at 11 Mbps, so hidden nodes could significantly interfere with each other. Nevertheless, the relatively small number of hidden nodes indicates that they are not likely to present much of a problem.

We then performed a similar analysis for external interferer pairs, again using our capture measurements. In this case we have three possible outcomes: a collision; A's packet is captured - the desired outcome; B's packet is captured, causing A's packet to fail. The results are shown in Figure 21 for transmit rates of 1 and 11 Mbps. We see that at 1 Mbps the odds of A's packet not being captured are extremely small. While the odds of A's packet being received at 11 Mbps are somewhat worse, they are still very good for most pairs. Thus, although there are many external interferer pairs, their impact on performance is limited.

To understand why external interference has so little impact, consider Figure 22 which shows a CDF of the difference in received signal strength at A's AP from A and B. For the vast majority of external interferer pairs, A enjoys a significant advantage in signal strength over B: for more than 95% of the pairs, A's advantage is greater than 20 dB. [2] reports similar results for an operational network. Our analysis yields insight into the likely cause of the behavior seen in their data. Moreover, we found that exposed nodes - not measured in [2] - are a more significant source of network inefficiency.

Our analysis leads us to conclude that the most serious

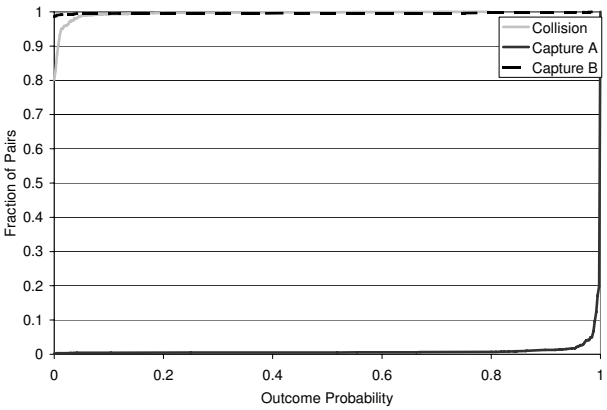

(a) 1 Mbps

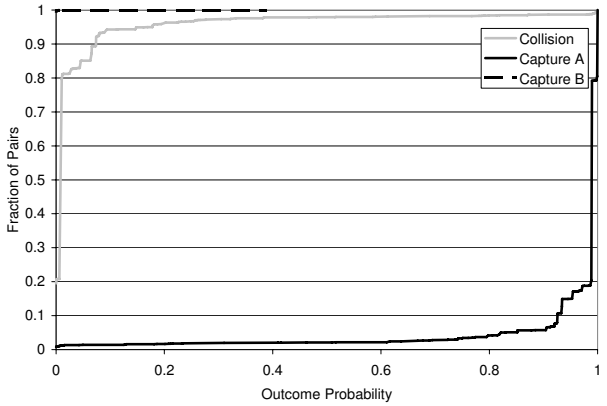

(b) 11 Mbps

Figure 21: External Interferer CDFs for 1 and 11 Mbps

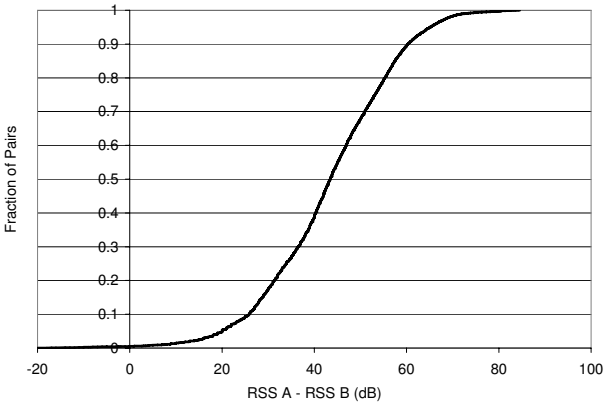

Figure 22: External Interferer CDFs for 1 and 11 Mbps

inefficiency plaguing this network is exposed nodes. For the vast majority of pairs, A enjoys a significant advantage, thus A need not defer when B is transmitting.

11. RELATED WORK

Early wireless measurement efforts focused on high level statistics such as offered load, session time, and user mobility, but more recently a number of projects have collected detailed wireless network measurements in a variety of settings. Aguayo et. al. [1] investigate the link-level behavior of an active metropolitan mesh network; in particular, they measure delivery rates across links and consider possible sources of delivery rate variation. They consider a limited set of controlled experiments to help understand the behavior observed in the actual testbed (e.g. see Section 8). Papagiannaki et. al. [16] measure link-level behavior for in-home wireless networks. Cheng et. al. [2] record and reconstruct the behavior of an enterprise wireless LAN. With the exception of [1], these studies were not performed in a controlled setting. In contrast, we perform an exhaustive study of 802.11 link-level behavior in a controlled setting that allows us to understand the behavior that is occurring, and examine behavior that cannot be easily observed in a live network. Our carefully controlled measurements complement these earlier efforts in that they provides a knowledge base that can be leveraged to understand the behavior that is observed in deployed networks.

A number of papers have presented in depth studies of a particular aspect of wireless packet reception. The capture effect has received the most attention, e.g. [9, 23, 24]. [23] studies capture models for 802.11 using experimental trace data. [9] presents a controlled set of experiments characterizing packet capture using Prism2 chipset at 2 Mbps, e.g. similar to our results in Figure 4(b). They observed that under some conditions a later, stronger packet can be received at the expense of an earlier weaker packet, which is a case that we observed but did not present results for. A number of papers have studied the impact of packet capture on throughput, delay, and/or fairness [24, 15, 8, 4]. They do not directly characterize the capture effect, but instead focus on how it affects performance. [25] looks at the capture effect in sensor networks using low-power radios. Robinson et. al [20] measure multi-radio performance in a small multihop network. They address off-channel interference. Mishra et. al. [13] propose leveraging off-channel isolation and reception. Our use of a physical layer network emulator offers a higher level of control that allows us to run more exhaustive experiments. This provides new powerful insights into issues such as capture and off-channel performance.

12. CONCLUSION

A clear understanding of wireless device performance is critical for understanding how wireless networks behave and how they might be improved. Despite this need, little data exists for modern wireless networks on important performance issues such as packet capture, collision, off-channel reception and interference and how these interplay with issues such as hidden and exposed nodes. We have conducted a large controlled study of 802.11 device behavior aimed at replacing convention and assumption with measured device behavior. We analyzed the capture effect both as a function of delay and signal strength and showed that it is quite strong, especially at lower transmit rates. Next, we measured off-channel interference and reception behavior. We found that off-channel interference rejection can perform very well, confirming that it may be possible to use four partially overlapping, instead of three non-overlapping, channels in 802.11b networks. Our results show, however, that off-channel reception behavior is quite poor and this feature should be used with great caution. We also studied the effect of multipath. We found that today's wireless cards compensate well for multipath in indoor environments, where delay differences are limited, e.g. lower then 0.45 mi-

croseconds for our cards. Higher delay differences result in packet success rates that are largely uncorrelated with the RSSI.

Our measurements can be used to improve simulators and to provide guidance to network managers. As an example, we used our data to study the performance of a deployed wireless LAN. We found that hidden nodes are uncommon in dense wireless networks and that true collisions are unlikely for low modulation rates.

13. REFERENCES

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proc. of SIGCOMM 2004*, Portland, August 2004.

[2] Y. Cheng, J. Bellardo, and P. Benko. Jigsaw: Solving the Puzzle of Enterprise 802.11 Analysis. In *Proc. of SIGCOMM 2006. Pisa, Italy*, September 2006.

[3] C. Fullmer and J. Garcia-Luna-Aceves. Solutions to Hidden Terminal Problems in Wireless Networks. In *Proc. of Sigcomm 1997*, Cannes, France, September 1997.

[4] Z. Hadzi-velkov and B. Spasenovski. Capture Effect in IEEE 802.11 Basic Service Area Under Influence of Rayleigh Fading and Near/Far Effect. In *IEEE International Symposium on Personal Indoor Communication*, 2002.

[5] G. Judd. Repeatable and realistic wireless experimentation through physical emulation, October 2006.

[6] G. Judd and P. Steenkiste. Using Emulation to Understand and Improve Wireless Networks and Applications. In *Proc. of NSDI 2005*, Boston, MA, May 2005.

[7] G. Judd and P. Steenkiste. Understanding Link-level 802.11 Behavior: Replacing Convention with Measurement. In *Wireless Internet Conference 2007 (Wicon07)*, Austin, Texas, October 2007.

[8] J. Kim and J. Kim. Capture Effects of Wireless CSMA/CA/Protocols in Rayleigh and Shadow Fading Channels. *IEEE Transactions on Vehicular Technology*, 48, July 1999.

[9] A. Kochut, A. Vasan, A. Shankar, and Agrawala. Sniffing out the correct Physical Layer Capture Model in 802.11b. In *Proceedings of the 12th IEEE International Conference on Networking Protocols*, October 2004.

[10] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *Proc. of MSWiM 2004*, Venice, Italy, October 2004.

[11] S. Kurkowski, T. Camp, and M. Colagrosso. Manet simulation studies: The incredibles. *Mobile Computing and Communications Review*, pages 50–61, October 2005.

[12] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Analyzing the MAC-level Behavior of Wireless Networks in the Wild. In *Proc. of SIGCOMM 2006. Pisa, Italy*, September 2006.

[13] A. Mishra, E. Rozner, S. Banerjee, and W. Arbaugh. Exploiting Partially Overlapping Channels in Wireless Networks: Turning a Peril into an Advantage . In *Proc. of IMC 2005*, Berkeley, CA, October 2005.

[14] K. Mutsuura, H. Okada, K. Ohtsuki, and Y. Tezuka. A New Control Scheme With Capture Effect. In *International Conference on Communications*, June 1989.

[15] C. Namislo. Analysis of mobile radio slotted aloha networks. *IEEE Transactions on Vehicular Technology*, 33(3):199–204, August 1984.

[16] K. Papagiannaki, M. Yarvis, and W. Conner. Experimental characterization of home wireless networks and design implications. In *Proc. of Infocom 2006*, Barcelona, Spain, April 2006.

[17] B. Ramamurthi, A. Saleh, and D. Gooodman. Perfect-Capture ALOHA for Local Radio Communications. *IEEE Journal on Selected Areas of Communication*, 5, June 1987.

[18] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice-Hall, Englewood Cliffs, NJ, 2002.

[19] L. Roberts. Aloha packet system with and without slots. *ARPA Network Information Center, TR ASS Note 8*, 1972.

[20] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy. Experimenting with a Multi-Radio Mesh Networking Testbed. In *Proc. of WiNMee 2005. Trento, Italy*, April 2005.

[21] C. Tai. Complementary reciprocity theorems in electromagnetic theory. *IEEE Trans. on Antennas and Propagation*, 40(6):675–681, 1992.

[22] F. Tobagi and L. Kleinrock. Packet switching in radio channels: Part ii–the hidden terminal problem in carrier sense multiple-access and the busy-tone solution. *IEEE Trans. on Comm.*, 23(12):1417–1433, 1975.

[23] C. Ware, J. Chicharo, and T. Wysocki. Modelling of Capture Behavior in IEEE 802.11 Radio Modems. In *IEEE International Conference on Telecommunications*, June 2001.

[24] C. Ware, J. Judge, J. Chicharo, and E. Dutkiewicz. Unfairness and Capture Behavior in 802.11 Adhoc Networks. In *IEEE International Conference on Communications (ICC 2000)*, June 2000.

[25] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the Capture Effect for Collision Detection and Recovery. In *The Second IEEE Workshop on Embedded Networked Sensors (EmNetS-II*, May 2005.