

Understanding Link-level 802.11 Behavior: Replacing Convention With Measurement

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, and interference. As a result, it is often 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 often have not been validated using today's wireless radios. 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. We replace conventional assumptions and possible misconceptions with actual recorded behavior. Additionally, we analyze the impact of our observations on commonly deployed networks. 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 most physical layer realism, practical considerations such as ease of use, control, and repeatability have made simulation the dominant evaluation technique. Recent work [10], however, has shown that unless a great deal of care is taken, simulation can lead to incorrect results.

To obtain accurate results, one must carefully consider both the simulation setup and the accuracy of the simulator. A simulator must correctly model all aspects of the system, including the networking protocol stack, signal transmission, propagation, and reception. Unfortunately, relatively little work has been done on validating the accuracy of many simulators. For example, initial work [9] has shown that the most commonly used simulator - ns-2 - produces results that differ significantly from real-world experiments. Moreover, real-world measurements [1, 15] 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. We first summarize the capabilities of our wireless network emulator and we present measurement results for clear channel reception as a baseline for later measurements. We then present our results for the following phenomena: hidden and exposed nodes, packet capture behavior with two competing

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

WICON 2007 October 22-24, 2007, Austin, Texas

Copyright 2007 ACM 987-963-9799-04-2/07/10 ...\$5.00.

transmitters, off-channel reception behavior, off-channel interference, and link asymmetry. Finally, in Section 9, we use the observed link-level behavior to analyze the performance of a production wireless network.

2. EXPERIMENTAL SETUP

Fine grained characterization of wireless link-level behavior requires tight control over signal propagation between the transmitters and receivers. This is achieved using physical layer wireless network 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, or *any* wireless device in the supported frequency range) 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.

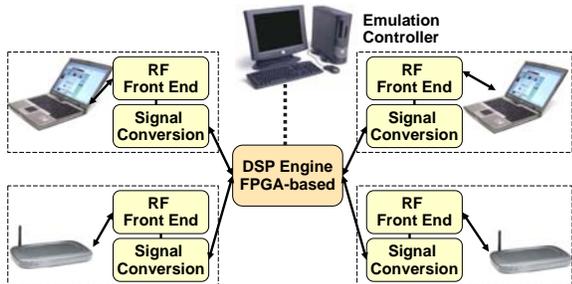


Figure 1: Emulator Implementation

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 [5] and to validate its results.

Emulation is controlled by an Emulation Controller PC which models the physical environment and coordinates the movement of RF nodes in the modeled physical environment with the modeling of the signal propagation environment with the emulator hardware. Different methods of channel emulation are supported, including the use of statistical models and replay of channel measurements. In the experiments

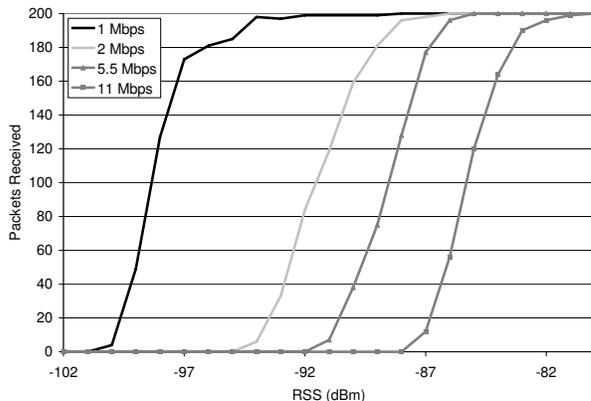


Figure 2: Clear Channel Reception

discussed in this paper, the Emulation Controller directly specifies the channel characteristics - in particular attenuation or path loss between devices - allowing us to construct arbitrary network topologies.

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. In this test, we use a single transmitter and a single receiver and the emulator to 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 sent 200 broadcast packets to the receiver. The receiver then recorded the number of successful packets. As broadcast packets do not use link-level retries, this experiment allowed 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 12.) The noise floor and carrier sense of the Senao cards was measured to be approximately -99 dBm.

4. CAPTURE UNDER DELAYED INTERFERENCE

An essential element in understanding and modeling wireless packet reception is understanding 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

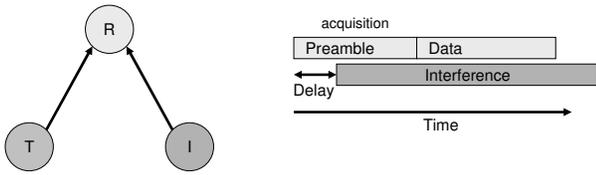


Figure 3: Capture Under Delayed Interference

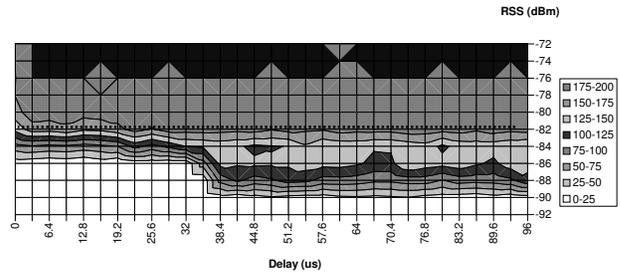
receiver’s ability to capture a single desired signal in the presence of an undesired interfering signal. The following section will discuss the effects of received signal strength on the outcome of two competing desirable signals.

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 the interferer; I constantly sends interfering 1 Mbps 1500 byte broadcast packets that are received at -82 dBm. T and I are hidden [21, 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 was actively transmitting, and 2) T’s current transmission had been active for a specified delay. Note that we did not explicitly control the arrival time of packets from I - rather we control when I is *allowed* to interfere with T.

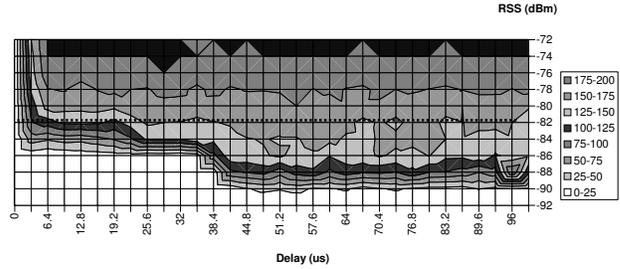
This setup allows us to investigate the effect of interference timing and signal strength on packet reception. Figure 4 shows the results of the experiments for data rates of 1 through 11 Mbps. We show for different delay-RSS combinations, how many packets R received from T, out of a total of 200 packet sent. The x-axis shows the delay of interference from I with respect to 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, we can 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. 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 [8]. The results at 11 Mbps (Figure 4(d)) are very different. While a longer delay in the interference still improves reception, stronger transmission is needed, and reception is no longer possible when the signal is weaker than the interfering signal. The results for 2 and 5.5 Mbps fall in between those for 1 and 11 Mbps.

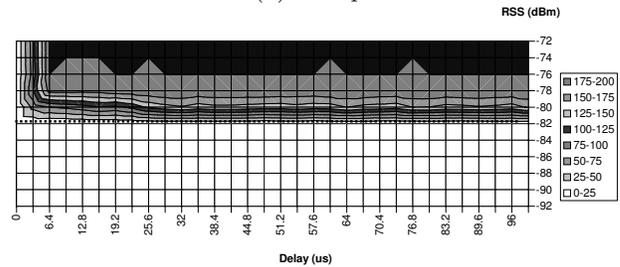
Conclusion - Our results have important ramifications for MAC design. 802.11’s carrier sense mechanism operates without respect to the cell in which a station resides. Not only may this cause transmitters to needlessly defer (an exposed node situation), but transmitters in different cells (i.e. with different receivers) will tend to synchronize their attempted transmissions in order to limit the time that the



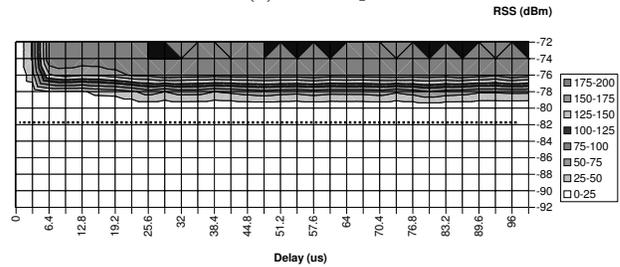
(a) 1 Mbps



(b) 2 Mbps



(c) 5.5 Mbps



(d) 11 Mbps

Figure 4: Capture Under Delayed Interference

medium is experiencing collisions [18]. 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 [16, 13].

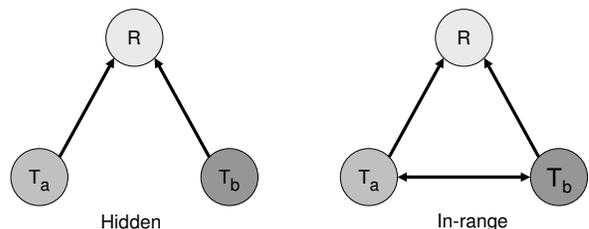


Figure 5: Setup for capture experiments

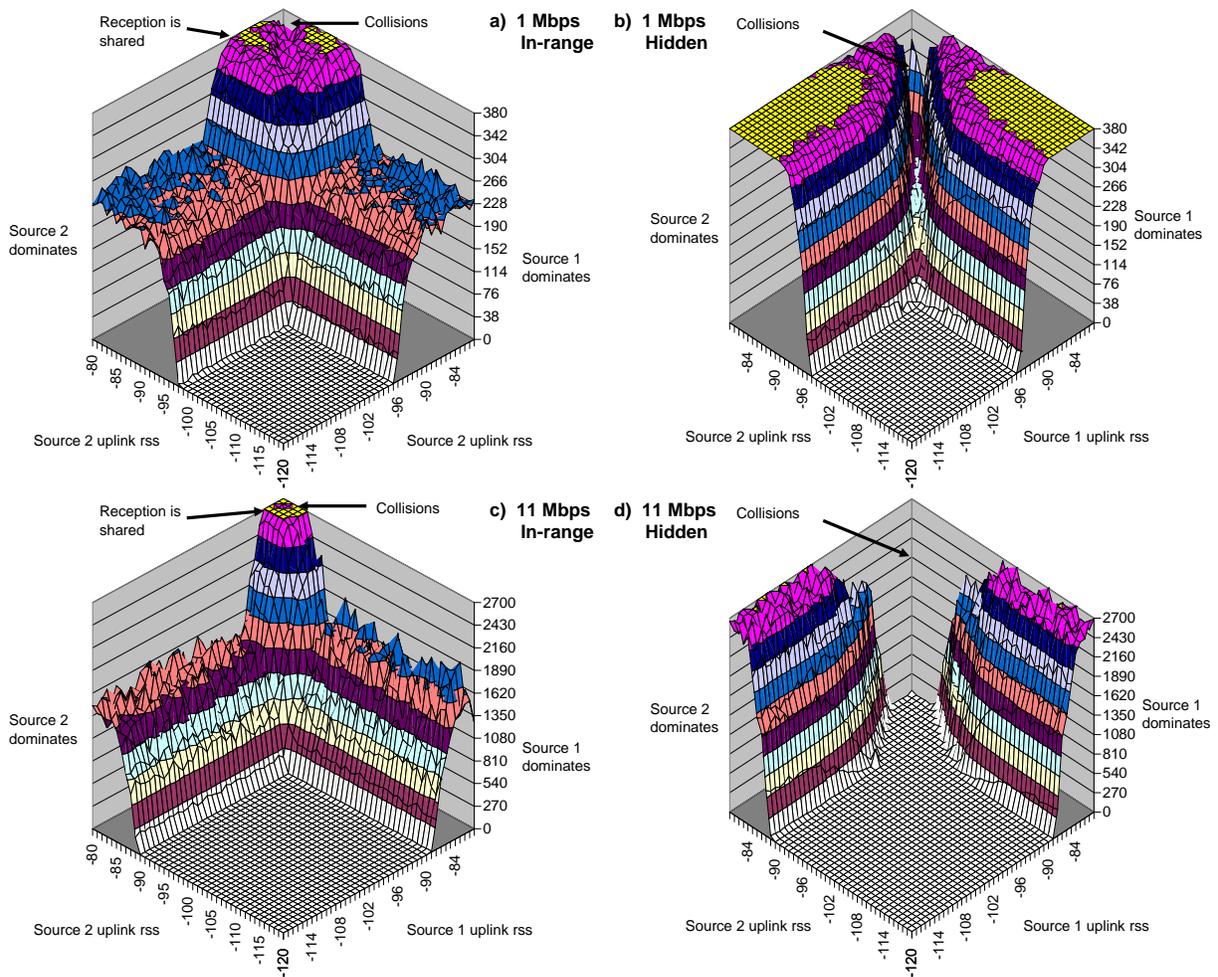


Figure 6: Packet Capture Results

5. CAPTURE WITH COMPETING TRANSMITTERS

We now discuss the effects of received signal strength on the outcome of two competing *desirable* signals. Using the configurations shown in Figure 5, we determined the reception outcome for different RSS combinations from T_a and T_b at R without controlling the interference timing. This was done by having the two transmitters T_a and T_b 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 T_a and T_b to R and we record how many packets R received from each transmitter. We measured all combinations of RSS values from T_a and T_b 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 T_a and T_b to hear each other’s transmissions while in the “in-range” setup, we set the RSS from T_a to T_b at -80 dBm and vice versa, so that T_a and T_b 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 T_a and

T_b 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. The results for 2 Mbps and 5.5 Mbps are again omitted but they are similar to the 1 Mbps and 11 Mbps cases.

In all in-range cases, we found that when the RSS at R from both T_a and T_b 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 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 reception between T_a and T_b is still good, and they defer for 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 isn’t communicating with the same receiver.

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, T_a and T_b send at full rate

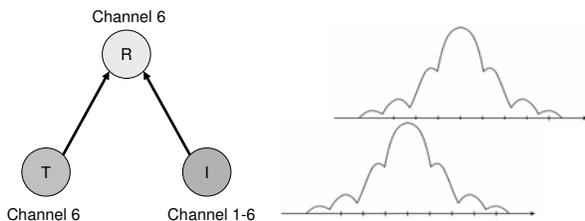


Figure 7: Off Channel Interference Setup

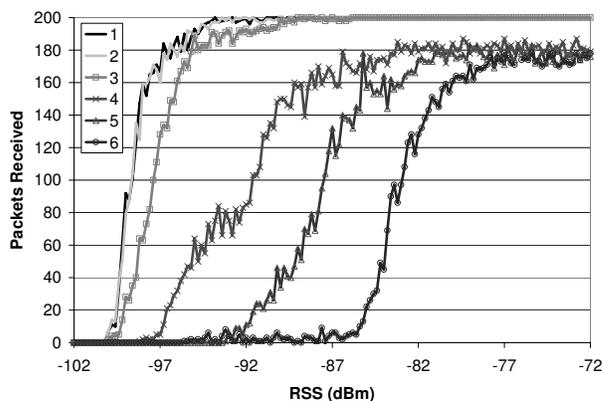
since they are out of carrier sense range. Looking at the 1 Mbps results, we see collisions only occur for a very narrow range of signal strengths where the RSS at R is nearly identical from both Ta and Tb. Hence, in the exposed node situation, waiting is likely unnecessary if the transmission of the in-range node is 1 Mbps. At higher transmission rates, the range over which collisions occur grows especially for 5.5 and 11 Mbps rates. Hence, higher transmission rates require a larger signal to interference and noise ratio (SINR) in order to be captured successfully.

Conclusion - These tests 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.

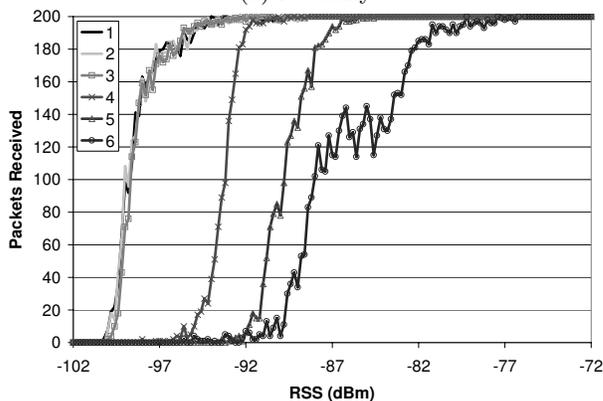
6. OFF-CHANNEL INTERFERENCE

In the US, eleven 802.11b channels are available in 5 MHz increments from 2.412-2.462 GHz. Each 802.11b channel is designed to have 22 MHz occupied bandwidth which implies that a total of three 802.11b signals may 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 [12] despite the fact that there would be some signal overlap. While there is some evidence to support this idea, there has not been a tightly controlled measurement of the impact on real hardware.

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 has 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



(a) No delay



(b) Large delay

Figure 8: Off-channel Interference, 1Mbps

of packets to R and R records how many were received successfully. Packets were broadcast, so no retries took place. We repeated this test for all four 802.11b modulation rates.

Our results are shown in Figures 8 and 9. For all tests where interference was prevented until well after packet acquisition, we observed that the impact of interference from channels 1, 2, and 3 was low and virtually identical. Channel 4 degraded performance 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. (We omit the “no delay” results for 2, 5.5 and 11 Mbps in the interest of space.)

To investigate the effect of stronger interference, we reran the 11 Mbps large delay tests with interference of -72 dBm. We found that the larger interference had 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

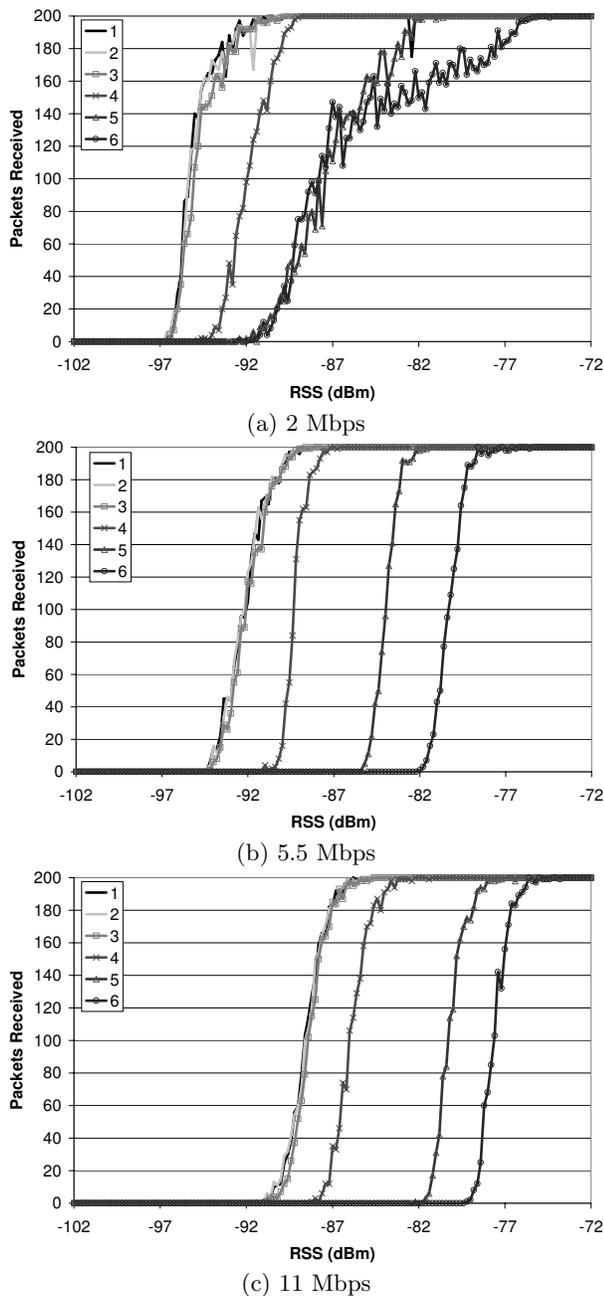


Figure 9: Off-channel Interference, large delay

Recently the observation that some off-channel packets can be received has led 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 characterized 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 as anticipated by providing 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 possible, but only when the signal is strong and even then it is imperfect. 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 or 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.

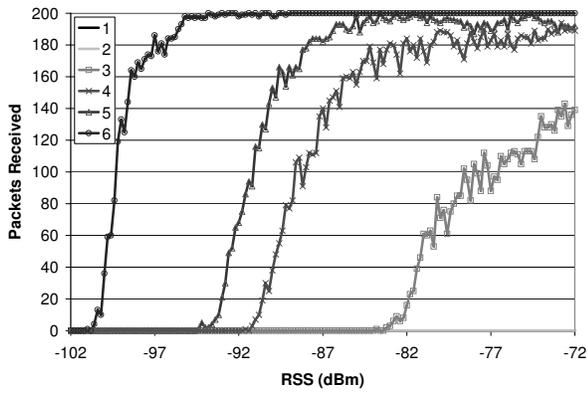
8. LINK ASYMMETRY

Several research groups [15, 11, 9] have independently observed asymmetric wireless link behavior. In particular, several instances of asymmetric packet delivery rate have been observed, i.e. the packet delivery rate from node A to node B is not the same as the packet delivery rate 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. We quantify two possible causes, transmit power and receive noise floor variations, and we also discuss two site-specific factors, antenna diversity and interference variations.

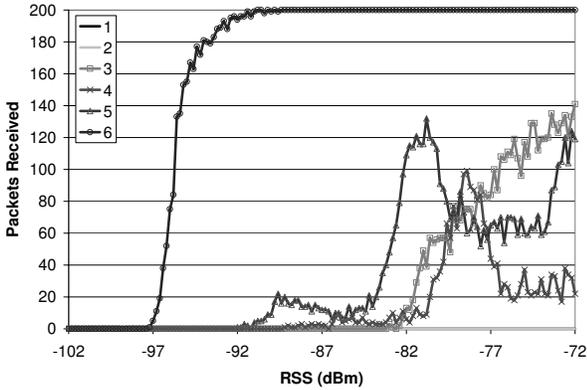
However, let us first mention a "non-cause": **asymmetric signal propagation**. Asymmetric signal propagation is physically impossible according to the reciprocity theorem [20], which states that *if the role of the transmitter and the receiver are interchanged, the instantaneous signal transfer function 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. Link asymmetry has been observed, however, even when the same model card is used. 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).

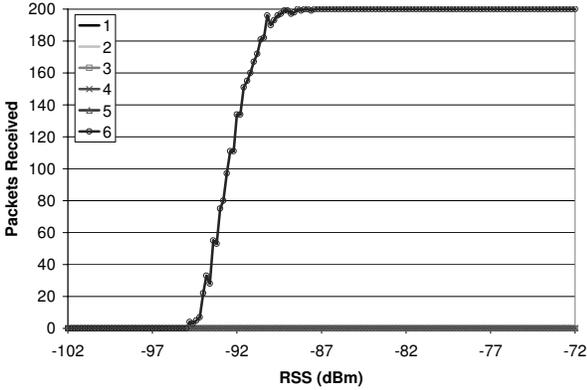
Figure 11 shows the average of 23 individual measurements for each card and the computed 95% confidence intervals. We observed that the cards fell into two distinct sets A and B. The cards in set A had an averaged transmit power close to 23 dBm with very little variation. In contrast, cards in set B had a higher transmit power and exhibited



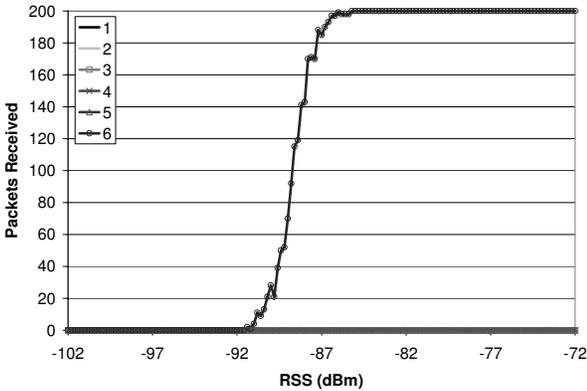
(a) 1 Mbps



(b) 2 Mbps



(c) 5.5 Mbps



(d) 11 Mbps

Figure 10: Off-channel Reception

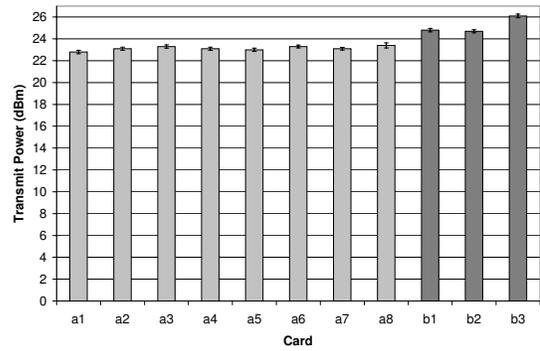


Figure 11: Senao Card Power

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.

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). Moreover, quality variations in transmit modulation and 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 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 12 shows the results. We observed approximately 3 dB of variation over all of the links that we measured.

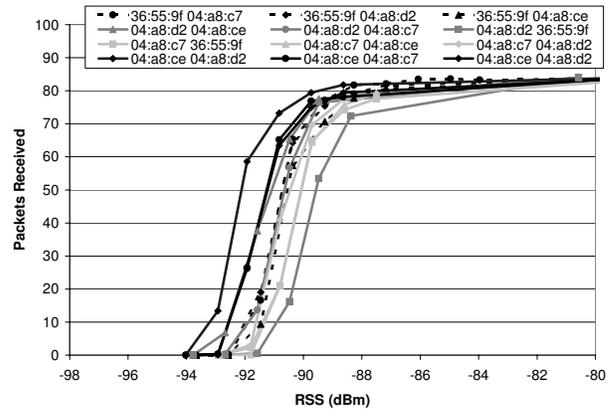


Figure 12: Packet Delivery Rate Variation

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 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 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 asymmetry may result from 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.

9. 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 the researchers and network managers. For example, the data collected in Sections 4 through Section 8 can be used to improve the accuracy of simulators, as has been explored by others for packet capture [8]. 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. 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? To answer these questions we analyzed the performance of an operational and heavily utilized 802.11b wireless 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 constructed a radio map of the building by sampling received signal strength 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 measurements, so we model path loss using a log distance path loss model [17] with a d_0 of 1.0 meter, pld_0 of

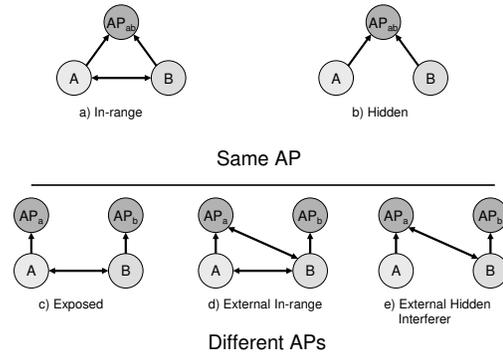


Figure 13: Infrastructure Topologies

40.0 dB, and a path loss exponent “ n ” of 5.0.

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 13(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 13(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 13(d)). Second, if A 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 13(e)). In cases Figure 13(a) and (d), carrier sense will avoid interference occurring. Note also that we must only consider A’s interaction with its access point. B’s interactions with its access point are considered when it is “A”.

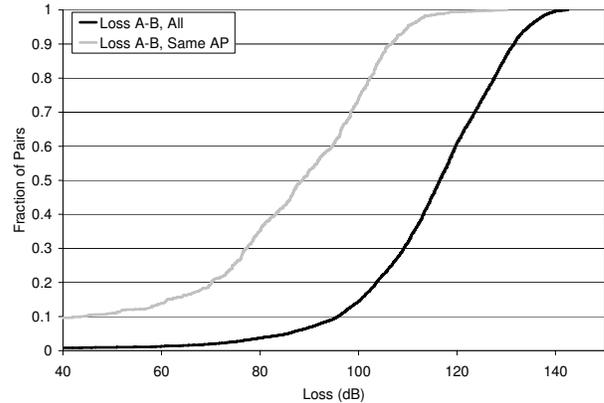


Figure 14: Path Loss CDF for Operational WLAN

Figure 14 plots the CDF of path loss for all client pairs and also for client pairs associated with the same access point. This CDF is for a single execution of our analysis (other runs produced very similar results). Based on this path loss data, Table 1 shows how often the interactions in Figure 13 were found in a single run of our analysis (other runs were quite similar). Clearly, hidden nodes were very uncommon. The reason is that 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

Table 1: WLAN Performance Analysis Summary

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

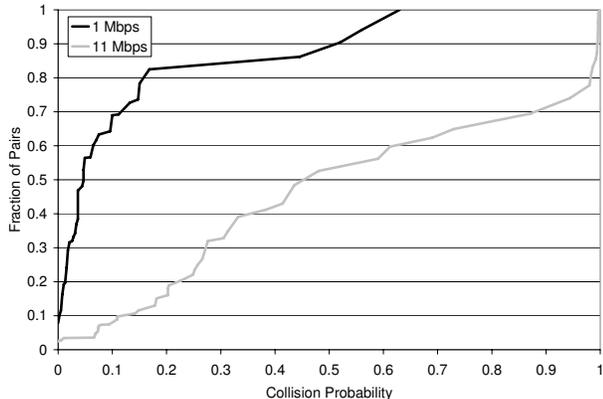


Figure 15: Hidden Node Collision Probability, 1 Mbps vs. 11 Mbps

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 if 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 15 shows the result for both 1 Mbps transmissions and for 11 Mbps transmissions. At 1 Mbps, very few of the hidden pairs are likely to have high collision probabilities. For these cases, A will receive more throughput to its access point than B will. At 11 Mbps, however, there is a fair chance for collision. In practice, we expect that most nodes in this network would communicate at 11 Mbps, so hidden nodes could significantly interfere with each other. Nevertheless, the prevalence 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 16. We found that at 1 Mbps the odds of A’s packet not being captured are extremely small; we omit this result for brevity. 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. The reason is that for the vast majority of external interferer pairs, A enjoys a significant advantage in signal strength (greater than 20 dB for more than 95% of pairs.) [2] reports similar results for an operational network. Our analysis yields insight into

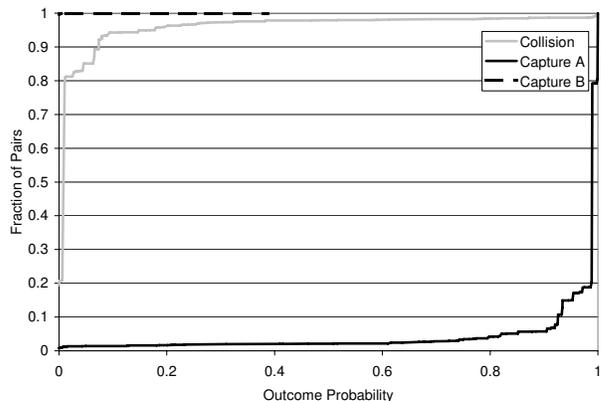


Figure 16: External Interferer CDFs, 11 Mbps

the likely cause of the behavior seen in their data. Moreover, we show that exposed nodes - not measured in [2] - are a major source of network inefficiency.

Our analysis leads us to conclude that the most serious 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.

10. RELATED WORK

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. Papagiannaki et. al. [15] 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. 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. [8, 22, 23]. [22] studies capture models for 802.11 using experimental trace data. [8] 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 [23, 14, 7, 4]. They do not directly characterize the capture effect, but instead focus on how it affects performance. [24] looks at the capture effect in sensor networks using low-power radios. Robinson et. al [19] measure multi-radio performance in a small multi-hop network. They address off-channel interference. Mishra et. al. [12] propose leveraging off-channel isolation and reception. Our use of a physical layer network emulator offer a higher level of control that allows to do more exhaustive experiments.

11. 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 quite strong, especially at lower transmit rates. We also measured off-channel interference and reception behavior. We have found that off-channel interference rejection can perform very well, confirming the potential benefits of this feature to allow more than three channels to be used 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.

The measurement 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.

12. 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] 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.
- [8] 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.
- [9] 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.
- [10] S. Kurkowski, T. Camp, and M. Colagrosso. Manet simulation studies: The incredibles. *Mobile Computing and Communications Review*, pages 50–61, October 2005.
- [11] 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.
- [12] 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.
- [13] K. Mutsuura, H. Okada, K. Ohtsuki, and Y. Tezuka. A New Control Scheme With Capture Effect. In *International Conference on Communications*, June 1989.
- [14] C. Namislo. Analysis of mobile radio slotted aloha networks. *IEEE Transactions on Vehicular Technology*, 33(3):199–204, August 1984.
- [15] 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.
- [16] B. Ramamurthi, A. Saleh, and D. Goodman. Perfect-Capture ALOHA for Local Radio Communications. *IEEE Journal on Selected Areas of Communication*, 5, June 1987.
- [17] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice-Hall, Englewood Cliffs, NJ, 2002.
- [18] L. Roberts. Aloha packet system with and without slots. *ARPA Network Information Center, TR ASS Note 8*, 1972.
- [19] J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy. Experimenting with a Multi-Radio Mesh Networking Testbed. In *Proc. of WiNMe 2005. Trento, Italy*, April 2005.
- [20] C. Tai. Complementary reciprocity theorems in electromagnetic theory. *IEEE Trans. on Antennas and Propagation*, 40(6):675–681, 1992.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.