

Experimental Analysis of Concurrent Packet Transmissions in Low-Power Wireless Networks *

Dongjin Son^{1,2} Bhaskar Krishnamachari¹ John Heidemann²
{dongjins, bkrishna}@usc.edu, johnh@isi.edu

¹Department of Electrical Engineering-Systems, ²Information Sciences Institute
Viterbi School of Engineering, University of Southern California

Abstract

We undertake a systematic experimental study of the effects of concurrent packet transmissions in low-power wireless networks. Our measurements, conducted with Mica2 motes equipped with CC1000 radios, confirm that guaranteeing successful packet reception with high probability in the presence of concurrent transmissions requires that the *signal-to-interference-plus-noise-ratio* (SINR) exceed a critical threshold. However, groups of radios show a wide gray region of about 6 dB. We find that this occurs because the SINR threshold can vary significantly depending on the measured signal power and radio hardware. We find that it is harder to estimate the level of interference in the presence of multiple interferers. We also find that the measured SINR threshold generally increases with the number of interferers. Our study offers a better understanding of concurrent transmissions and suggests richer interference models and useful guidelines to improve the design and analysis of wireless network protocols.

1 INTRODUCTION

There is growing awareness that realistic models of wireless links are essential for developing efficient protocols for wireless networks and evaluating them meaningfully [13]. In particular, good interference models are essential not only to improve the evaluation of existing protocols under medium-to-high traffic loads, but also to aid in the future design of novel interference-aware protocols for wireless networks.

Most research considering network interference normally assumes one of two interference models: the protocol model or the physical model [10]. In the protocol model, which is implemented by many state-of-the-art

wireless network simulators, concurrent transmissions from any node within a given range (referred to as the interference range) of a receiver will cause a collision that results in the loss of a packet from its corresponding sender. A recent study by Whitehouse *et al.* [16] has argued that this protocol model significantly overestimates packet loss during concurrent transmissions and can therefore result in the design of inefficient medium access protocols. In the physical model, a packet from the sender is lost at the receiver only if the *signal-to-interference-plus-noise-ratio* (SINR) falls below a given threshold. To our knowledge, the physical model, which is widely used in communication theory, has not been previously tested rigorously through real experiments in the context of low-power wireless networks.

Several recent empirical studies in the context of wireless sensor networks have given us an understanding of the complex non-ideal behavior of low-power wireless links [8, 14, 15, 17, 21]. However, most of these empirical studies have focused on single links, without concurrent transmissions from interfering nodes.

In this paper, we systematically study the effects of concurrent transmissions through experimental measurements with low-power Mica2 motes equipped with CC1000 radios. Our experiments involve the measurement of received signal and interference strengths as well as packet reception rates under carefully designed single-interferer and multiple-interferer scenarios. In agreement with the results in [16], we also find the simplistic interference range-based protocol model to be inadequate. Our experimental results confirm some key aspects of the SINR-based physical model, while suggesting significant ways in which it can be enhanced for applicability in real deployments.

There are several concrete findings from our experimental study that offer useful insights; these are summarized in Table 1. Our measurements, conducted with Mica2 motes, confirm that guaranteeing successful packet reception with high probability in the presence of

*This research is supported partially by the National Science Foundation (NSF) through the following grants: CNS-0435505, CNS-0347621, CCF-0430061, CNS-0325875, and CNS-0435517, and by a hardware donation from Intel Corporation.

Finding	Section
Single interferer effects	4
Capture effect is significant	4.1
SINR threshold varies due to hardware	4.2
SINR threshold does not vary with location	4.3
SINR threshold varies with measured RSS	4.4
Groups of radios show ~ 6 dB gray region	4.5
New SINR threshold model	4.6
Multiple interferer effects	5
Measured interference is not additive	5.2
Measured interference shows high variance	5.3
SINR threshold increases with more interferers	5.4

Table 1: Key findings of this paper

concurrent transmissions requires that the SINR exceed a critical threshold. However, groups of radios show a wide gray region of about 6 dB. We find that this occurs because the SINR threshold can vary significantly depending on the measured signal power and radio hardware (but not depending significantly on the location). By contrast, we find that the gray region is quite narrow for a specific hardware combination at a fixed signal strength level. We find that it is harder to estimate the level of interference in the presence of multiple (two or more) interferers for two reasons: (a) the joint interference measurements show a much higher variation when there are multiple interferers, and (b) the measured joint interference strength is not always the sum of the individual interference strengths. We also find that the measured SINR threshold generally increases with the number of interferers.

The rest of the paper is organized as follows: in Section 2, we discuss some key related empirical studies in wireless networks. We present our experimental methodology in Section 3. We discuss the results from experiments involving a single interferer in Section 4, and those involving multiple interferers in Section 5. Finally, we present our conclusions and discuss future work in Section 6.

2 RELATED WORK

In the context of wireless sensor networks, several empirical studies have given us an understanding of the complex non-ideal behavior of low-power wireless links [8, 14, 15, 17, 18, 21]. The bulk of these studies focus on wireless link quality in the absence of concurrent transmissions. Some studies (including [8, 11, 15]) do evaluate the impact of increased interference and traffic load on higher layer protocols, but they do not explain the fundamental behavior of wireless links under interference as the experiments in this paper aim to do.

One recent paper by Whitehouse *et al.* [16] does address wireless link quality in the presence of concurrent transmissions. They propose a technique to detect and recover packets from collisions taking advantage of the so-called *capture effect*, whereby a packet with the stronger signal strength can be received in spite of a collision. Their scheme works by allowing the detection of preambles even during packet reception. They study the performance of the proposed scheme through experiments with a single interferer and show that the simplistic protocol model (in which the communication range is chosen to be the interference range) significantly overestimates interference and can result in inefficient MAC design. Our study complements their work by quantifying the SINR conditions under which the capture effect can be observed (that are the conditions under which their proposed scheme shows performance gains).

We should also mention briefly that there have been several experimental studies pertaining to 802.11 radios that consider concurrent packet transmissions. Many of these are orthogonal to our work in that they pertain primarily to the evaluation of different routing metrics in the presence of multiple flows (e.g., [4, 6, 20]). Jamieson *et al.* [12] consider concurrent transmissions when they investigate MAC protocol performance by turning on and off the carrier sense functionality at different bit rates in an 802.11 testbed. They argue that a capture-aware carrier sense mechanism that considers the bit rates and SINR will improve network efficiency. Our work can provide useful guidelines for the development of similar techniques for low-power wireless networks.

Of particular relevance to this work is the study by Aguayo *et al.* [2], who perform link measurements to study the causes of packet loss in a 802.11 mesh network (*Roofnet*). They experimentally study several packet loss related factors such as SINR (which they refer to as S/N ratio), transmit bit-rate, interference, and multi-path fading. Their experimental results show a wide (greater than 3 dB) gray region of SINR with intermediate values of packet delivery probability even for the same receiver. They argue that, for this reason, SINR cannot be used as a reliable predictor of delivery rate in 802.11 networks. Our study confirms that this observation also holds for the low-power mote radios, and we explore more systematically the impact of hardware and measured signal strengths on packet reception rate as a function of SINR.

3 EXPERIMENTAL METHODOLOGY

In this section, we discuss some key aspects of our experimental methodology. In Section 3.1, we discuss the hardware and software used. We describe our experimental design for carrying out synchronized measurements in Section 3.2. We conclude this section by dis-

curring the regression model we use for mapping SINR to packet reception rates in Section 3.3.

3.1 Hardware and Software

Our study is based on systematic experiments on a PC104 [1] testbed running Linux. The experiments are conducted in a controlled indoor office environment where surrounding objects are static, with minimal time-varying changes in the wireless channel due to multipath fading effects. Any code that can be used commonly by all PC104 nodes is accessed on a central computer through an NFS-mounted directory. We use Mica2 motes, with the Chipcon CC1000 [3] radio operating at 433 MHz, as an RF transceiver on the PC104 node. We use the Linux-based Emstar software framework to take advantage of its interactive interface with sensor nodes in the testbed [9].

We use the S-MAC protocol [19], configured in fully-active mode without sleep cycles. To study collisions in a controlled manner we intentionally disable carrier sense and random backoff in the MAC. This allows us to freely transmit concurrent packets even when there is ongoing packet transmission in the same wireless channel. We also omit the MAC-level RTS/CTS/DATA/ACK sequence by sending packets as broadcasts, avoiding the complications of ARQ. We thus disable much of the MAC functionality in order to focus on the fundamental behavior of wireless links in the presence of interference.

There are several other important wireless platforms, including IEEE 802.11 and IEEE 802.15.4. As an experimental study, we can only affirm that our results apply to the CC1000 radio. However, hardware variation and large gray regions have been previously observed for 802.11 radios [2] and it is likely that low power 802.15.4 radios will show similar results.

3.2 Measurement Design

Our study requires a careful configuration to synchronize both packet transmissions as well as measurements of signal strength and packet loss. Figure 1 shows our experimental configuration. Each experiment involves four types of nodes: a sender, a receiver, one or more interferers, and a special *synchronizer* node. The synchronizer broadcasts a *sync packet* just before each single or concurrent packet transmission. This serves to synchronize the clock of every node in the testbed. The sync packet is a kind of *reference broadcast* [7]. Each transmitting node (sender or interferer) sets its packet transmission time and the receiver sets the received signal strength measurement time based on this reference time.

In our controlled experiments the hardware identity and locations of the sender, interferer, and receiver is

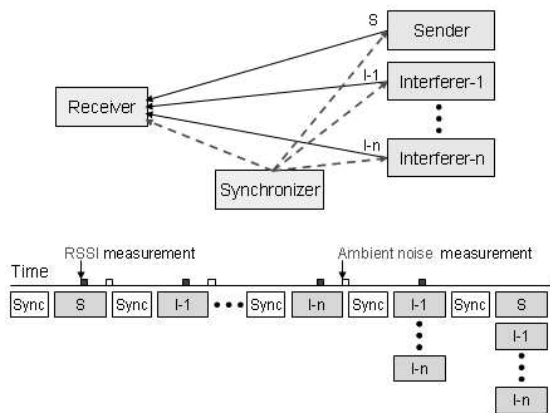


Figure 1: Overview of the testbed with experimental methodology used for time synchronization, signal strength and PRR measurement

fixed, but we vary the transmit power of the sender and interferers over some range. For each specific combination of transmit power settings, there is a series of packet transmission epochs. In each epoch, there is the following sequence of transmissions, each interleaved with a sync packet (see Figure 1): (i) the sender transmits alone; (ii) each interferer in turn transmits alone; (iii) all interferers transmit concurrently; (iv) the sender transmits concurrently with all interferers. The receiver measures signal strength in the middle of each single or concurrent transmission, except the final one, which is used to record whether the packet was received successfully or not. If a total of n packet transmission epochs are used for a particular transmit power combination, the *packet reception rate (PRR)* for that combination is calculated as the total number of packets received successfully divided by n . We typically use 75 epochs to estimate PRR with a precision of about 1.3%. In addition, ambient noise measurements at the receiver are taken at the end of reception of each of the single packet transmissions.

Due to jitter in the testbed system, transmission start times vary with a mean of 3 ms. Further, obtaining reliable signal strength measurements can take up to 7 ms (this is not a controllable parameter in the CC1000 radios [3]). Hence the signal strength measurement times need to be carefully chosen at the receiver to ensure any intended collision occurs. We take measurements in the middle of long packet transmission periods. With 230 byte packets, packet transmission time is about 97 ms and so we can tolerate substantial jitter.

As second potential timing problem can occur depending on when packets transmissions begin. When the sender and all interferers are transmit concurrently, variation in the transmission starting times can cause the sender packet to arrive 8 ms or later than the first inter-

ferring packet. In such cases we observe that the packet is never recognized at the receiver, even if its signal is strong enough to overwhelm the interferer. This problem occurs because our implementation of the radio’s physical layer requires that packet data immediately follow the start symbol of the packet. It will refuse to shift to a later, stronger packet once it has read the start symbol of the earlier packet. The 8 ms period corresponds to the transmission time required for the 18 byte preamble and 2 byte start words. This problem was identified by Whitehouse *et al.* [16]; they solved it by modifying the MAC software to retrain when it encounters subsequent start symbols of higher power. We became aware of this approach mid-way through our work. To keep a consistent methodology, rather than modify our MAC to retrain, we detect and filter out cases when the strongest packet arrives later than 8 ms. To do this we add two timestamps to each packet, recording transmission start and completion times. Fortunately, because timing error is normally distributed with a mean of 3 ms, few packets arrive late. From timestamps in logs, about 3% of epochs must be discarded due to late arrival of the strongest packet. By removing these packets, we should get loss rates comparable to a MAC that can retrain on later packets as proposed by Whitehouse *et al.*

Signal strength measurements are used to estimate the *received signal strength (RSS)* and *received interference strength (RIS)* for the concurrent packet transmissions at the end of the epoch. These include the strength of the transmission and any ambient noise. Received signal strength measurements are taken in ADC counts and converted to dBm following the manufacturer’s documentation [3, 5]. This documentation also indicates that signal strength measurements are inaccurate when they exceed -55 dBm. We confirmed this claim with tests and therefore drop measurements beyond this threshold.

Given the RSS, we define JRIS as the measured *joint received interference strength* when all interferers transmit concurrently. If N is the average ambient noise level measured at the receiver, we can then calculate the *signal-to-interference-plus-noise-ratio (SINR)* as:

$$\text{SINR}_{dB} = 10 \log_{10} \frac{10^{\text{RSS}_{dBm}/10} - 10^{N_{dBm}/10}}{10^{\text{JRIS}_{dBm}/10}} \quad (1)$$

Note that we base our SINR values from measurements taken directly at the receiver. This approach is central to the experimental nature of our work. Alternatives such as measuring transmit power at the sender would require the use of theoretical models of path loss and ambient noise, neither of which we know for our environment. While our approach avoids inaccurate signal strength estimation due to mismatches between model and environment, we do not claim that the measured signal strength values represent “true” signal strengths, since that would

require a calibrated comparison with a highly accurate RF measurement device. Instead, we claim that they represent signal strengths as measured by actual radios. Our results may not directly apply to future radios with more accurate measurements of signal strength, however we believe our findings have great utility with regard to practical protocols which must depend on similar measurements in real deployments.

3.3 A Regression Model Mapping SINR to PRR

Interference affects the relationship between interference and packet reception. While all of our findings are based on raw measurements, we add regression lines in some of the graphs to clarify the SINR-to-PRR relationship. The link layer model presented by Zuniga and Krishnamachari [22], especially SNR to PRR conversion formula, is the basis for our regression model.

$$\text{PRR} = \left(1 - \frac{1}{2} \exp^{-\beta_0 \text{SINR} + \beta_1}\right)^{8(2f-l)} \quad (2)$$

This regression model is intended for non-coherent FSK modulation and Manchester encoding that is used in Mica2 motes. We introduce the parameters β_0 and β_1 to fit the experimental dataset to the regression model. The β_0 value controls the shape of the regression curve and β_1 induces horizontal shifts of the curve. Based on repeated experiments, we determined that a constant β_0 value provides excellent fits (see, for example, Table 4); find the optimal β_0 for each experiment improved our R^2 values by at most 0.01. We therefore hold β_0 constant at 2.6 in all our single-interferer figures. The parameter f is the frame size (230 bytes for our experiments) of the packet and l is the preamble size in bytes (20 bytes).

4 EXPERIMENTAL STUDY OF SINGLE INTERFERERS

In this section, we describe our systematic experiments to understand how concurrent packet transmissions affect packet reception when there is a single sender and a single interferer. In Section 4.1, we begin by studying how different transmit powers cause different regions of reception, from good to noisy to bad (or white to gray to black). We then define the signal-to-interference-plus-noise-ratio (SINR) threshold for good reception and show that it varies with hardware combinations (Section 4.2) and signal strength (Section 4.4), and does not vary strongly due to location (Section 4.3). Finally, in Section 4.5, we complement our detailed studies based on small numbers of nodes with a larger 12-node experiment. Finally, from these results we propose a realistic simulation model in Section 4.6.


4.1 Interference and Black-Gray-White Regions

It is well known that stronger packets can be received even in the face of weaker, concurrent transmissions, and this result has recently been confirmed and exploited experimentally [16]. We begin our study with experiments to carefully quantify this *capture effect* as a function of the measured signal strengths from concurrent packet transmitters over a wide range of transmission powers.


In these experiments we consider two transmitting nodes, SRC1 and SRC2. By definition, we call the stronger signal source the *sender* and the weaker signal source the *interferer*. From this definition these roles change with the varying transmission powers. To study how these roles change, we vary transmission powers as both sources send 230-byte packets and calculate packet reception rate (PRR), here over 60 epochs.

Figure 2(a) presents the packet reception ratio (PRR) of SRC1 and SRC2 as the transmit power of SRC1 varies. Here we fix the transmission power level of SRC2 at -4 dBm and vary the output power of SRC1 from -17 dBm to 2 dBm. Without interference, either source has reliable communications with the destination. However, the experiment shows that three distinct regions occur as SRC1's transmit power varies. Beginning at the left of the graph, when SRC1 is less than -10 dBm, SRC2's transmissions are always received. In the middle of the graph, when SRC1 transmits at powers between -7 and -5 dBm, packets from *neither* of the senders are recognized at the receiver. At the right of the graph, with SRC1 at -1 dBm or more, SRC1 is always successful. This experiment shows two clear regions of packet capture, for SRC2 at the left, and SRC1 at the right. We call these regions the *white* regions, where one source is assured reception even in the face of a concurrent transmission. These regions can be compared to the *black* region in the middle where neither transmission is received. Finally, we observe two *gray* regions at intermediate power levels (from -10 to -7 dBm and -4 to -1 dBm), where packets reception is intermittent. We define the gray region as any combination of sender and interferer transmit power levels that result in PRRs between 10% and 90%. Our definition was inspired by the notion of the gray area described by Zhao and Govindan [21]. As with their definition, our gray region corresponds to high variation in packet reception. However, the gray area defined in their work refers to a spatial distance range, and is not related to power levels.

To measure the level of interference in the channel we directly measure the received signal strength (RSS) in Figure 2(b). Recall that we measure RSS values at the receiver, first taking separate measurements for each transmitter and then during the concurrent transmission. We



(a) Transmission power level to PRR



(b) Transmission power level to RSS

Figure 2: Transmission power for SRC1 is varied between -17 to 2 dBm while the transmission power of SRC2 is fixed at -4 dBm. Ambient noise level at the receiver is shown together. Error bars show 95% confidence intervals

align the x-axes of Figures 2(a) and 2(b) to relate RSS to PRR. We observe that when the RSS of both sources become similar (within 0.6 dBm, when SRC1 is around -6 dBm), packet reception for both transmitters is zero as the transmissions corrupt each other. Further from this point, more packet receptions are observed as the received signal strength difference between two transmitters increases.

Table 2 reproduces the PRR, RSSI, and transmit power values from Figure 2 and adds calculated signal-to-interference-plus-noise-ratio (SINR) values. SINR represents the difference between the sender (by definition, the strongest transmitter) and the interferer. We categorize each SINR value based on the corresponding PRR as

Tx Pwr of SRC1	RSS of SRC1 (dBm)	SINR (dB)	PRR	Region
-17	-76.55	9.51	1	white (SRC1)
-14	-74.07	7.08	1	
-12	-72.59	5.87	1	
-10	-71.09	4.21	0.98	
-8	-69.76	3.00	0.72	gray (SRC1)
-7	-68.22	1.56	0	black (neither)
-6	-66.33	0.58	0	
-5	-65.78	1.73	0	
-4	-63.99	2.98	0.03	
-3	-63.01	3.98	0.22	gray (SRC2)
-2	-61.96	5.02	0.82	white (SRC2)
-1	-60.36	6.54	0.98	
0	-59.64	7.08	1	
1	-58.13	8.75	1	
2	-36.85	9.93	1	

Table 2: SINR-to-PRR mapping with region distinction. RSS of SRC2 is static around -66.8 dBm and ambient noise is around -94.6 dBm




Figure 3: Packet reception rate at different RSS combination from SRC1 and SRC2. Black-gray-white regions are marked with cross, triangle, and circle respectively

being in a *black*, *gray*, or *white* region for the dominant source.

For simplicity, Figure 2 varied only one source’s transmission power while holding the other constant. By contrast, Figure 3 shows measured results when the transmit powers of *both* sources are varied. This extensive set of experiments confirms that the results of Figure 2 hold regardless of which transmitter is varied or what power levels are considered. A horizontal or vertical slice through this figure would show white regions for either SRC1 or




Figure 4: Effect of different packet sender and interferer hardware on SINR-to-PRR relationship

SRC2, a black region in the middle, and gray regions on the border. We also observe that the edge of the gray region is not strictly linear as power varies. We will study this issue in more detail in Section 4.4.

Figures 2 and 3 show that concurrently transmitted packets are all corrupted when they have nearly equivalent signal strength at the receiver. However, there is a significant range of transmission powers in which the capture effect occurs and the stronger packet is received successfully. These results lend further evidence to show that the simplistic protocol interference model can be highly misleading. Capture-aware MAC schemes are indeed likely to provide significant improvements in efficiency.

These observations motivate us to analyze various factors that impact relationship between SINR and PRR. We define the *SINR threshold* as the minimum SINR which guarantees a reliable packet communication with $PRR \geq 0.9$. In the following sections, we examine the impact of hardware combinations, node locations, and signal strength variations on the measured SINR threshold. In particular, we seek to know whether there is a constant SINR threshold for all scenarios.

4.2 SINR Threshold and Transmitter Hardware

Section 4.1 demonstrated the packet capture effect and defined the SINR threshold. We next study SINR threshold to see if it is affected by transmitter hardware.

We consider two pairs of nodes, SRC1-SRC2 and SRC1-SRC3. As in Section 4.1, we hold one transmitter’s received signal strength constant at -66 dBm and vary the others from -66 to -77 dBm. We then measure the SINR threshold.

Figure 4 presents these experimental results. On the

left side of the graphs, SRC1 is the sender and SRC2 or SRC3 is the weaker interferer. On the right side, the opposite holds, with SRC1 being weaker. The x-axis shows the SINR (the negative signs on the left hand side should be ignored as an artifact of the presentation). In addition, the solid and dotted lines fit our regression model (defined in Section 3.3) to the experimental data.

First, we compare the experiment results from SRC1-SRC2 pair, shown as the solid line model and asterisk points. The SINR threshold values are different for each transmitter; SRC1 has an SINR threshold of 3.4 dB and SRC2 has an SINR threshold of 5.3 dB. There is a nearly 2 dB difference between these thresholds. When we compare the experiment results with different pairs of hardware (i.e., between the solid and dotted regression lines), we can see that SRC1 requires a stronger signal strength to reach the same level of PRR at the same receiver when the interferer is changed from SRC2 to SRC3. SRC1’s regression line (shown in the left side of the figure) moves about 1 dB to the left with interferer SRC3 and SRC3 requires about 1.7 dB higher SINR threshold compared to SRC2 when the same node SRC1 is the interferer. These results indicate strongly that the specific hardware combination of sender and interferer change the measured SINR threshold. (We rule out location differences as an alternative explanation in Section 4.3.)

Note that since our SINR calculations in all cases are based on measurements at the same receiver, we can rule out differences that have to do with transmit-side calibration settings, receiver sensitivities, or differences in the magnitude of the path loss from different transmitter locations. We speculate that the hardware-combination-specific variations in the SINR-threshold result from distorted signals due to non-linear effects in the radio transmitters. Even at the same measured signal strength at the receiver, the signals from different sources may have different levels of distortion, in turn affecting the packet reception differently.

4.3 Effects of Location on PRR and SINR

Multipath reflections are a major source of interference and are strongly dependent on location. One possible explanation for the variations in hardware shown in Section 4.2 could be that the nodes were in different locations. We therefore next study the effect of location on the SINR-to-PRR relationship.

To study the possible effect of packet sender and interferer location on the SINR-to-PRR relationship, we swap the location of SRC1 and SRC2 and performed the same experiments as in Section 4.2. Swapping the sender locations changes the channels observed between the two transmitters and the receiver. Figure 5 compares the ex-

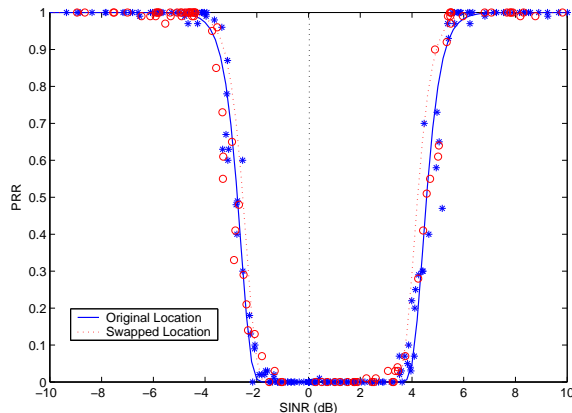


Figure 5: Effect of different packet sender and interferer location on SINR-to-PRR relationship

Location	Source	β_1 (95% confidence)
Original	SRC1	-0.914 (± 0.108)
	SRC2	3.802 (± 0.127)
Swapped	SRC1	-0.587 (± 0.157)
	SRC2	3.774 (± 0.147)

Table 3: Parameter β_1 and 95% confidence intervals for two different locations

periment results from new, swapped location with previous experiment results at the original node location. There is no noticeable difference in SINR-to-PRR relationship between these two set of experiment results. When we compare the parameter β_1 value used for each regression model (presented in table 3), β_1 values are very close for the same sender, not for the same location. But, SRC1 β_1 value is still located a little bit outside of 95% confidence interval of β_1 value used for switched location. This difference is from the effect of location change but it is minor compared to the hardware effect, as can be observed from the corresponding curves in figure 5.

From this comparison, we can verify that the main difference in SINR threshold between two nodes is from the transmitter hardware (or signal distortion level) difference rather than their location difference. We have run similar experiments with a two additional pairs of nodes, as well as with different locations for the nodes used above. We consistently find that location change does not make distinguishable difference in SINR threshold. However, all our experiments have been carried out in an office environment. An area of future work is to study if these results apply in other environments, both indoors and outdoors.




Figure 6: Experiments with wide range of sender and interferer signal strength. Sender: SRC2, Interferer: SRC1

4.4 Effect of Sender Signal Strength on the SINR Threshold

Our studies with two senders showed that the edge of the white region does not exhibit a linear relationship with unit slope (see Figure 3), which would be expected if the SINR threshold remained a constant regardless of the measured signal strength. In Section 4.2, we showed that different transmitter hardware results in different SINR thresholds. We next study more carefully how the measured sender signal strength affects the SINR threshold.

Here we vary the transmission power level of both packet sender and interferer over a wide range so that the received signal strength range varies from -91 to -52 dBm at the intended receiver. Figure 6 shows these experimental results, where SRC1 is an interferer and SRC2 is a packet sender.

This figure shows a gray region that is about 4.2 dB wide from SINR values of just above 2 to above 6 dB. This wide range applies even though locations and hardware are both constant—the only difference we have made for this experiment was to vary the transmit signal strength of the sender.

To better understand the data in Figure 6, we collected the RSS values into 1.5 dB intervals (10 raw ADC counts) and then fit our regression model to each set of experimental data. Table 4 shows the RSS ranges and corresponding model parameters (β_1) and SINR thresholds, along with goodness-of-fit (R^2) data. (We use a constant 2.6 of β_0 based on the analysis from the experimental data set as described in Section 3.3.) This table shows that our model provides an excellent fit to the data, even with a constant value for β_0 , since the worst case R^2 fit value is 0.963. We therefore conclude that our regression model can accurately summarize the experimental data. We also observe that the model parameter




Figure 7: SINR-to-PRR relationship categorized for different received signal strength levels. Experiment results in each category are represented with a regression line. Sender: SRC2, Interferer: SRC1

RSS range (dBm)	β_1	$SINR_{\theta}$ (dB)	R^2
-55.2 – -56.7	0.425	3.99	0.998
-56.7 – -58.2	3.827	5.30	0.982
-58.2 – -59.7	6.894	6.48	0.993
-59.7 – -61.2	7.183	6.59	0.992
-61.2 – -62.7	6.873	6.47	0.987
-62.7 – -64.2	6.373	6.28	0.995
-64.2 – -65.7	3.856	5.31	0.963
-65.7 – -67.2	3.802	5.29	0.979
-67.2 – -68.7	2.589	4.82	0.997
-68.7 – -70.2	1.232	4.30	0.997
-70.2 – -71.7	0.223	3.91	0.992

Table 4: β_1 , SINR threshold ($SINR_{\theta}$), and R^2 (goodness-of-fit) value for sender SRC2 for SRC1-SRC2 pair experiments when β_0 is set to 2.6

β_1 varies non-linearly over these measured RSS values. This variation in β_1 shows that the SINR threshold also varies with measured RSS in some non-linear manner, even when hardware and location are unchanged.

To investigate how the SINR value relates to transmission power we plot the regression models in Figure 7. These show that the SINR threshold is highest at medium measured RSS values and lowest when the measured RSS value is strong or weak. For example, in Figure 7 the fitted model shifts to the right (higher SINRs) as the RSS shrinks from -56.0 to -60.5 dBm (see arrows 1, 2, 3, and 4), then shifts back to the left as RSS reduces further to the lowest observed values of -71.0 (arrows 5 through 11).

To confirm that this experimental result was not peculiar to our hardware or location, we repeated similar experiments with several other pairs of nodes. Due to

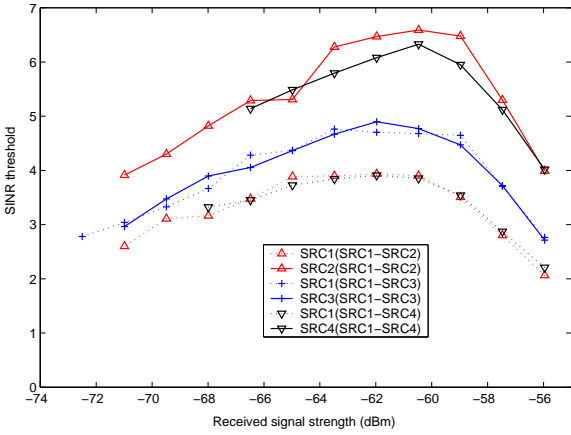


Figure 8: SINR threshold for 0.9 PRR change at different received signal strength level

space limitations, we do not reproduce the raw PRR-SINR graphs, but instead fit a model to each experiment and compute the SINR threshold. Figure 8 shows how the SINR threshold value (for 0.9 PRR) changes over different levels of sender signal strength for three different pairs of node experiments: SRC1 with each of SRC2, SRC3, and SRC4. For each pair of nodes, the figure shows two lines, one line each for when one of the transmitters behaves as a packet sender while the other behaves as an interferer.

All six SINR thresholds in Figure 8, show maximum values when the sender’s signal strength (measured at the receiver) is around -61 dBm. In this region, the SINR threshold, the β_1 parameter value, and the width of the black region are all highest. This result suggest that MAC protocols designed to exploit the capture effect and simulations designed to realistically model wireless collisions both must consider the magnitude of the signal strengths in addition to the ratio of signal and interference powers. We believe that curves such as those plotted in Figure 8 can be used as the basis for realistic simulations.

An important open question is understanding what physical phenomena causes this variation in SINR threshold. One possibility is that the radio transfer function exhibits nonlinear effects that affect signals with high and low signal strengths; another is that the RSSI measurement process itself is skewed at these extremes. A more detailed understanding of the causes of this RSS-SINR-PRR relationship is an area of future work.

4.5 Testbed Experiments

To confirm that our hardware and signal strength effects on SINR apply generally, we performed testbed experiments that consider a wider range of hardware combina-




Figure 9: Testbed experiments with 12 neighbor nodes

tions and RSS levels. We randomly deployed 12 PC104 nodes in two large rooms where the distance between the intended receiver and the farthest node in the testbed was around 18 meters. We selected an intended receiver node and a time synchronizer (using the procedure described in Section 3.2) and performed pairwise experiments with the remaining 10 nodes in the testbed. For each pair, one node is the sender (with stronger RSS) and the other node behaves as an interferer.

We set the interferer’s transmission power constant at -8 dBm so that it has a constant received interference strength (RIS) at the receiver. Measured RIS values from different interferers range from -81 to -63 dBm, but we observe a change of up to 1 dBm RIS from the same interferer at different times, presumably due to time-varying changes in the environment. We then vary the transmission power of the sender from strength equal to the interferer’s RIS value until a power level where the RSS is strong enough to provide reliable (close to 100%) packet reception.

We calculate SINR values based on the measured RSS and RIS pair information as well as the measured ambient noise and plot the SINR-to-observed-PRR relationship in Figure 9. Experimental results show a large variation in the SINR-to-PRR relationship (or in SINR threshold values). This is because different interferers in the testbed generate signals with different distortion levels and different RISs at the intended receiver. Also, different senders have different SINR thresholds for the same interferer.

The change in RIS level causes a similar effect as the RSS level change (presented in Section 4.4). This change is because different interference levels require different RSS levels to provide the same level of link reliability. For one pair of sender and interferer, we intentionally change the default transmission power of the interferer (which results in the RIS between -74.2 and -60.5 dBm)

to see the effect of RIS change on the SINR threshold apart from the hardware effect. Figure 9 marks these results with triangles. This RIS level change causes a change in SINR threshold similar to our previous observations, with a 1.9 dB gray region.

In the figure, the circles represent experiment results corresponding to having different sender hardware for a given fixed interferer. This sender hardware change results in about 3.1 dB gray region. The width of gray region varies between 1.6 and 3.6 dB for different individual interferers with 9 different senders. Overall, we observe a 6.1 dB wide gray region in the testbed experiments.

Thus, the testbed experiments confirm the two identified causes of SINR threshold variation (hardware combination and measured signal strength). These causes can explain the high variation in SINR-to-PRR mapping observed in previous experimental studies [2], and strongly suggest that constant SINR-to-PRR mappings will not model all realistic situations. Upper layer protocols designed based on the constant SINR threshold assumption may therefore be inefficient or work incorrectly.

4.6 Modeling the SINR Threshold

Now that we have identified that hardware and signal strength each affect the SINR threshold, we propose a simple simulation model for single interferer scenarios that considers these effects.

Based on the collected data in the testbed, we model the RSS and SINR threshold relationship with a quadratic function. We then allow hardware choice to shift this model with a normal distribution around our observed mean. We have verified that a quadratic fits signal strength reasonably well, but confirming the normal distribution of hardware is an area of future work. (We do not have enough hardware combinations to confirm normality at this time.) The model for SINR threshold ($SINR_\theta$) for sender S at a given RSS is therefore:

$$SINR_\theta(S, RSS) = \alpha_2 RSS^2 + \alpha_1 RSS + \alpha_0 + \zeta_S \quad (3)$$

where $\zeta_S \sim \mathcal{N}(0, \sigma^2)$

Where we set $\alpha_2 = -0.0305$, $\alpha_1 = -3.855$, $\alpha_0 = -116.91$. The hardware effect is modelled as a zero-mean Gaussian random variable ζ_S with a variance of $\sigma^2 = 1.33$, that moves the curve up and down.

This model represents one application of our experiments to modeling reception of real radio in simulation. An important area of future work is to model not only how packet reception is observed by radios, but also measurements taken with an accurate RF spectrum analyzer to provide ground truth.

5 EXPERIMENTAL STUDY OF MULTIPLE INTERFERERS

In this section, we consider concurrent packet transmissions involving more than two transmitting nodes (i.e., involving two or more interferers). In Section 5.1, we define how we empirically measure the joint interference as well as a conventional estimator assuming additive interference strengths. We then show that the measured joint interference generally does not match the additive assumption (Section 5.2). We then show in Section 5.3 that it is difficult to estimate the joint interference in the presence of more than one interfere, because measurements show high variance. Finally, we investigate the impact of multiple interferers on the SINR threshold in Section 5.4.

5.1 Joint Interference Estimator

When there is a single interferer (IFR) (i.e., a concurrent packet transmitter), we can estimate the interference strength from this interferer based on the individually measured *received interference strength (RIS)*. We now consider how joint interference may be estimated in the presence of multiple interferers.

The following two metrics are estimators of joint interference, with n interferers and k measurements from a given setup:

$$JRIS(e) = \sum_{i=1}^n RIS_{IFRi}$$

$$JRIS(m) = \frac{\sum_{i=1}^k JRIS_i}{k} \quad (4)$$

$JRIS(e)$ is the estimation based on the summation of individual RIS measurements from each interferer where RIS measurement for each interferer is taken without any interference in the same channel. $JRIS(e)$ is a conventional way to calculate the interference from multi-sources in theoretical studies.

$JRIS(m)$ uses the mean of multiple JRIS measurements as the estimator of joint interference. $JRIS(m)$ is a more practical method to estimate the joint interference from multiple interferers using the collected, actual JRIS measurements in real systems.

5.2 Additive Signal Strength Assumption

We first investigate the following question: “is the additive signal strength assumption valid in the measurements with low-power RF radios?”. Here, our aim is to examine the validity of using the measurement-based $JRIS(m)$ as an interference estimator in practice.

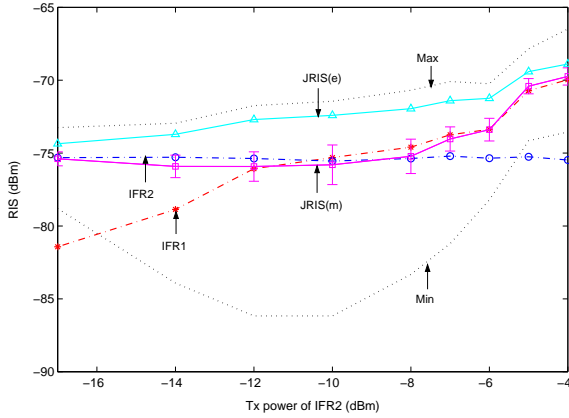


Figure 10: Two node experiments IFR2 at -75 dBm and IFR1 between -82– -70 dBm RIS at the receiver




Figure 11: Experiment results with two interferers (IFR1 and IFR2) causing equivalent RIS at the receiver

5.2.1 Two interferer experiments

We carefully design experiments (as described in Section 3.2) to measure the $JRIS$ at the intended receiver. First, we run some experiments with two concurrent interferers (IFR1 and IFR2) to see the effect of combined interference on the $JRIS$ values. IFR2 uses constant transmission power and the RIS from IFR2 is around -75 dBm at the receiver. IFR1 varies its transmission power between -17 to -4 dBm and this power adjustment results in the RIS between -82 to -70 dBm at the receiver.

Figure 10 presents the following information: (1) IFR1 and IFR2: mean RIS at the receiver from each interferer (IFR1 and IFR2) measured individually without any interference (2) $JRIS(e)$: joint interference estimation based on the additive signal strength assumption (3) $JRIS(m)$: mean measured $JRIS$ from both interferers (4) Min-Max: minimum to maximum value range of $JRIS$ measurements in two dotted lines. Each data point represents a mean measurement value over 100 experiments with 230B packets. Error bars show 95% confidence intervals for $JRIS(m)$ values.

While it is intuitive to see the dominance of stronger interference signal over the weaker interference due to the logarithmic nature of dBm unit, we still expect to measure a higher $JRIS(m)$ value from the intensified joint interference than single RIS when both interferers have equivalent RIS at the receiver, as with the $JRIS(e)$ estimates. However, the $JRIS(m)$ value follows the single stronger RIS value within the 95% confidence interval even at the point where both interferers have about the same individual RISs at the receiver (e.g. when transmission power of IFR1 is -10 dBm in Figure 10).

Even though $JRIS(e)$ value is normally considered as an estimator of joint interference, our experiments show that the measured $JRIS(m)$ values are generally always lower than the estimated $JRIS(e)$ values.

Even though $JRIS(e)$ value is normally considered as an estimator of joint interference, our experiments show that the measured $JRIS(m)$ values are generally always lower than the estimated $JRIS(e)$ values.

5.2.2 Additivity and RIS levels

To investigate if the observation from -75 dBm individual RIS level holds at different interference strength levels, we perform further experiments with two interferers at multiple RIS levels between -76 and -59 dBm. Figure 11 shows the experiment results for the cases when both interferers generate equivalent RIS at the receiver at different interference strength levels. While the $JRIS(m)$ value normally follows the stronger RIS value when the RIS values are not equal as well as at extreme values of RIS when they are equal for all interferers, in this experiment we find some intermediate RIS levels (around -68 dBm) where the $JRIS(m)$ value is larger than the stronger value. However, it is still the case that the $JRIS(m)$ is smaller than the $JRIS(e)$ value.

5.2.3 Additivity with additional interferers

To see the effect of additional interferers on $JRIS(m)$ and $JRIS(e)$, we have performed experiments with one to four interferers each with equivalent individual RIS levels. We incorporate the change in $JRIS(m)$ value at different RIS levels (identified in previous section) by repeating the same experiments at the following two RIS levels: -73.0 dBm (where $JRIS(m)$ is close to the single strongest RIS) and -68.8 dBm (where $JRIS(m)$ is higher than the single strongest RIS). We have individually measured RIS values from each interferer and the $JRIS$ value from different number of concurrent interferers over the 75 packet experiments for each setup.

When we compare the $JRIS(e)$ and $JRIS(m)$ at two different RIS levels in Table 5, there are smaller differences between the two interference estimators at -68.8 dBm individual RIS. This is in agreement with our previous re-

# of IFRs	Individual RISs (dBm)	JRIS(e) (dBm)	JRIS(m) (dBm)
1	-72.9 — — —	-72.9	-72.9
2	-72.9 -73.4 — —	-70.1	-72.7
3	-73.0 -73.5 -73.3 —	-68.5	-70.4
4	-72.9 -73.5 -73.5 -73.0	-67.2	-68.9

(a) RIS from each interferer around -73 dBm

# of IFRs	Individual RISs (dBm)	JRIS(e) (dBm)	JRIS(m) (dBm)
1	-68.8 — — —	-68.8	-68.8
2	-69.0 -68.7 — —	-65.8	-67.1
3	-69.1 -68.6 -68.7 —	-64.0	-64.2
4	-68.9 -69.0 -68.8 -68.2	-62.7	-63.7

(b) RIS from each interferer around -68.8 dBm

Table 5: Comparison of JRIS(e) and JRIS(m) metric for JRIS estimation at two different individual RIS levels




Figure 12: Frequency distribution of JRIS measurement values for two interferer experiments

sults that shows higher signal strength additivity at -68.8 dBm than at -73 dBm (presented in Figure 11). These results with multiple interferers also confirm our previous observation that the JRIS(e) estimates stronger interference than measured by JRIS(m).

5.3 Variation in JRIS Measurements

When we look at the each JRIS measurement value rather than the mean value (i.e., JRIS(m)), there is significant variation in the JRIS measurements especially when IFR1 and IFR2 have close interference strength at the receiver. The wide minimum to maximum JRIS value range (in Figure 10 and 11) clearly represents a significant variation in JRIS measurements. The standard deviation of the JRIS measurements is around 3 dBm (2.75

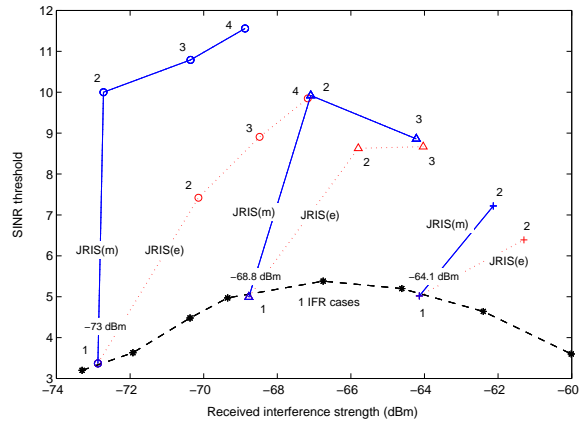


Figure 13: SINR threshold changes with different number of interferers which changes the received interference strength

to 3.65 dBm) over the experiments with different levels of two equivalent interference strength (shown in Figure 11). And the minimum-to-maximum JRIS range is consistently very wide throughout the experimented signal strength levels.

Figure 12 shows one example of the frequency histogram from the 300 JRIS and RIS measurements from two interferer experiments. While RIS measurements from each interferer (RIS1 and RIS2) are clustered together near the mean value (-68.2 and -68.5 dBm respectively), the JRIS values are widely distributed around its mean value (-66.2 dBm). This histogram clearly shows the wide variation from the multiple interferers in the JRIS measurements (where the standard deviation is 3.02 dBm) compared to the single interference cases (where the standard deviation is 0.30 and 0.37 respectively) and some additive behavior (about 2 dBm increase in JRIS(m)) from multiple interference at the given individual RIS level. The JRIS values are still normally distributed. Similar frequency distributions are observed from the experiments with two to four interferers.

In wireless communication protocols, collecting the received signal strength indication (RSSI) is a natural way to estimate the current interference plus noise level. However, single RSSI measurement (which we call RIS for interference measurement) cannot be an appropriate estimator of current interference if there is any possibility of having multiple interferers, due to the significant variance in the measurement values.

5.4 Effects of Joint Interference

By comparing JRIS(m) and JRIS(e), we have evaluated how measured joint interference levels from multiple interferers compare to estimated joint interference. We

next relate this back to the SINR threshold for reliable packet reception.

To evaluate the SINR threshold with multiple interferers, we vary both the number of interferers and the individual RIS levels. We consider from 1 to 4 interferers, and RIS levels of -73, -68.8, and -64.1 dBm, matching the experiments in Table 5 and adding the -64.1 dBm level.

Figure 13 shows the experiment results, comparing the SINR threshold against the received interference strength (RIS). We mark each data point with the number of interferers in each experiment and also indicate the method of joint interference estimation (either JRIS(e) or JRIS(m)) for each branch. The experiments in the same branch use the same individual RIS level. As indicated in Section 5.1, JRIS(e) values are predicted from individually measured RIS values, while JRIS(m) are joint measurements.

We draw three conclusions from this experiment. First, we consider how SINR varies as we add interferers at a given RIS level. We have three examples in the strings of experiments starting at -73, -68.8, and -64.1 dBm. Regardless of the estimator used (JRIS(m) or JRIS(e)), we observe that additional interferers raises the SINR level required to successfully receive a packet. This trend is clearest for the -73 dBm case where 1 to 4 interferers are considered, but it holds for all three cases.

Second, we can compare SINR threshold for two different estimators JRIS(e) and JRIS(m) (i.e., the dotted and solid lines in the figure). We find that JRIS(e) has consistently lower SINR threshold than JRIS(m). Recall from Section 5.2 that JRIS(e) has a consistently higher estimation of interference level. A lower SINR threshold with higher interference estimation sounds counter-intuitive, but this is a consequence of the way in which the SINR threshold is calculated. We have the measured received signal strength and its corresponding packet reception rate from the experiments. The only difference is in the interference level obtained by the two different estimators. We calculate the SINR threshold with this pre-identified RSS and the estimated interference with both methods, taking into account the ambient noise level. Hence the JRIS(e) estimator, which offers a higher level of interference, results in a lower SINR threshold. This illustrates the point that a careful selection of interference estimator is important because that can significantly affect the calculated SINR threshold value.

Finally, we can compare SINR threshold values as the JRIS increases. JRIS will rise either due to increase in the individual RIS in our three sets of experiments, and also due to increase in the numbers of interferers. In Section 4.4 and 4.5 we show that SINR threshold values changes at different signal strength levels. We highlight the variation in SINR threshold with a single interferer

at different RIS levels with an arched, dashed line at the bottom. We may perhaps expect multiple interferers to generally follow a similar trend. Unfortunately we do not have enough data to conclusively support or refute this trend for multiple interferers. The trend in two interferers shows a monotonically decreasing trend but this could be due to missing points at lower power levels, and we were unable to investigate three and four interferers at other power levels. Investigating this further is an area of future work.

6 CONCLUSIONS AND FUTURE WORK

In this paper we have presented experimental analysis of the effects of concurrent packet transmissions in low-power wireless link communications. We have confirmed the capture effect and the existence of the SINR threshold which ensures the successful delivery of the strongest packet under the concurrent packet communication situations. Our main contributions and findings are as follows:

- We perform the first systematic experimental study, to our best knowledge, which verifies a difference between the conventional approximation of the interference effect and the actual implication of concurrent packet transmissions in reality. Our experimental study provides new guidelines for more realistic simulation models.
- Our study shows that the SINR threshold is not a constant value, but that it depends on the transmitter hardware and the signal strength level. While the combinations of different hardware and signal strength in the testbed generate large (about 6 dB) gray region with mixed reception rate at the same SINR value, the gray region is small for a fixed hardware combination at the same signal strength level.
- Upper layer protocols that assume a constant SINR threshold can fail or be inefficient due to the significant variation in SINR threshold. Protocols designed considering capture effects and variability in SINR threshold will be more dependable and efficient.
- Single RSSI value measurement is not always a good estimator of current interference level because there is a large variation in measured signal strength in a multiple interference situation.
- The measured interference from multiple transmitters is generally less than theoretically predicted by

the assumption that interference is additive. For a given measured signal strength, therefore, the measured joint interference results in higher calculated SINR threshold values than predicted by theory.

- The SINR threshold generally increases with the number of interferers.

As future work, we plan to further study the causes of inconsistent SINR threshold with more experimental analysis and develop a realistic model for use in simulations. Understanding the fundamentals of concurrent packet transmissions in low-power wireless networks will establish a good starting point for efficient protocol designs. As a start, we hope to apply the lessons learned in this study to improve our previously proposed link quality and topology control scheme (PCBL [15]: transmission power control with blacklisting) through a more intelligent cross-layer link quality metric which considers dynamic network flows.

References

- [1] ISI laboratory for embedded networked sensor experimentation (ILENSE), PC104 based nodes, <http://www.isi.edu/ilense/testbed/pc104/>
- [2] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *ACM SIGCOMM*, aug 2004.
- [3] Chipcon. CC1000 Single chip very low power RF transceiver. http://www.chipcon.com/files/CC1000_Data_Sheet_2.2.pdf
- [4] D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *IEEE Mobicom*, sep 2003.
- [5] Crossbow. MPR/MIB user's manual 7430-0021-06, aug 2004. <http://www.xbow.com/Support/manuals.htm>
- [6] R. Draves, J. Padhye, and B. Zill. Comparison of routing metrics for static multi-hop wireless networks. In *ACM Sigcomm*, aug 2004.
- [7] J. Elson, L. Girod, and D. Estrin. Fine-grained network time synchronization using reference broadcasts. In *Fifth Symposium on Operating Systems Design and Implementation (OSDI)*, dec 2002.
- [8] D. Ganesan, D. Estrin, A. Woo, D. Culler, B. Krishnamachari, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical Report CS TR 02-0013, UCLA, 2002.
- [9] L. Girod, J. Elson, A. Cerpa, T. Stathopoulos, N. Ramanathan, and D. Estrin. Emstar: a software environment for developing and deploying wireless sensor networks. Technical Report CENS TR34, UCLA, dec 2003.
- [10] P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Transactions on information theory*, 46, 2, mar 2000.
- [11] B. Hull, K. Jamieson, and H. Balakrishnan. Mitigating congestion in wireless sensor networks. In *ACM Sensys*, nov 2004.
- [12] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. Understanding the real-world performance of carrier sense. In *SIGCOMM'05 Workshop*, aug 2005.
- [13] D. Kotz, C. Newport, and C. Elliott. The mistaken axioms of wireless-network research. Technical Report TR2003-467, Dartmouth, jul 2003.
- [14] D. Lal, A. Manjeshwar, F. Herrmann, E. Uysal-Biyikoglu, and A. Keshavarzian. Measurement and characterization of link quality metrics in energy constrained wireless sensor networks. In *IEEE Globecom*, dec 2003.
- [15] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of the effects of transmission power control and blacklisting in wireless sensor networks. In *IEEE SECON*, oct 2004.
- [16] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *IEEE Workshop on Embedded Networked Sensors (EmNetS-II)*, may 2005.
- [17] J. L. Wong, L. Kuang, M. Potkonjak, and D. Estrin. Statistical model of lossy links in wireless sensor networks. In *IPSN*, apr 2005.
- [18] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *ACM Sensys*, nov 2003.
- [19] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient mac protocol for wireless sensor networks. In *IEEE INFOCOM*, june 2002.
- [20] H. Zhang, A. Arora, and P. Sinha. Learn on the fly: Data-driven link estimation and routing in sensor network backbones. In *IEEE Infocom*, apr 2006.
- [21] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *ACM Sensys*, nov 2003.
- [22] M. Zuniga and B. Krishnamachari. Analyzing the transitional region in low power wireless links. In *IEEE SECON*, oct 2004.