Understanding the Causes of Packet Delivery Success and Failure in Dense Wireless Sensor Networks

Technical Report SING-06-00

Kannan Srinivasan[†], Prabal Dutta[‡], Arsalan Tavakoli[‡], and Philip Levis[†]
[†]Department of Electrical Engineering, Stanford University
[‡]Computer Science Division, UC Berkeley

Abstract

We present empirical measurements of the packet delivery performance of the Telos and MicaZ sensor platforms. At a high level, their behavior is similar to that of earlier platforms. They exhibit link asymmetry, a reception "grey region," and temporal variations in packet loss. Looking more deeply, however, there are subtle differences, and looking deeper still, the patterns behind these complexities become clear. Packet losses are highly correlated over short time periods, but are independent over longer periods. Environmental noise (802.11b) has high spatial correlation. Packet loss occurs when a receiver operating near its noise floor experiences a small decrease in received signal strength, rather than an increase in environmental noise. These variations cause the reception "grey region." While short-term link asymmetries are not uncommon, long-term asymmetries are rare. Based on these findings, we suggest several ways in which current practices could be easily changed that would greatly improve the efficiency, performance, and lifetime of sensor networks.

1 Introduction

The complexities of low-power wireless networking are a basic challenge in sensor network research. Early studies established that, in practice, networks of the lowest power devices — "motes" — exhibit complex behaviors not easily captured by simple propagation models. CSMA-based networks can easily suffer from a very irregular form of the broadcast storm problem [28, 16], packet reception rates over distance exhibit a large and highly variable "grey region" [50] and node pairs often have asymmetric loss rates [6]. These studies established important considerations for sensor network protocols [48, 23, 13, 43, 33] and have guided design decisions and tradeoffs for a wide range of systems [44, 17, 12].

These early studies made many important observations that have since become common knowledge within the community. Nevertheless, several important unknowns remain. First, while the early studies quantified many important characteristics of packet delivery, they were unable to establish the root causes of these complexities. In many cases, the hypothesis was hardware variations, such as slight differences in receiver sensitivity [16] or oscillator calibration [48], but these effects were neither quantified nor empirically measured, leaving the hypotheses unevaluated.

Second, the experimental variables used – distance, orientation, environment – present a "human-eye" view of a network, in that they are not variables that a node can easily observe. These are excellent guides to people for designing or installing a sensornet deployment, but they provide only a partial understanding of how nodes themselves observe the network. Considering packet delivery success and failure from a "mote-eye" view, in terms of what a mote can readily observe, would give a greater sense of how protocols or systems might make decisions based on this information.

Third, since these studies, the design space of mote platforms and their radios has changed significantly. Some earlier platforms had low-level software stacks on top of a wide range of radio technologies, including on-off keying (OOK) [20], amplitude shift keying (ASK) [19], and frequency shift keying (FSK) [42], while others used well-established full protocol stacks such as Bluetooth [26, 24]. More recently, many platforms, including the micaZ [37], Telos [9], and IntelMote2 [4], have gravitated towards a single data-link protocol, 802.15.4 [38], and even a single radio chip, the ChipCon CC2420 [21]. This newer technology differs significantly from earlier radios. For example, it operates in the 2.4GHz band, rather than 900MHz (rene/mica) or 915/433MHz (mica2), and uses offset quadratic phase shift keying (OQPSK) with a direct-sequence-spread-spectrum (DSSS) encoding. These changes suggest that 802.15.4-based platforms may behave quite differently than early studies would suggest.

In this paper, we present the results of an initial set of experiments that seek to shed light on these unknowns by answering the simple question:

> "What are the observable causes of packet success and failure in modern platforms, and how can a node detect them?"

Successfully answering this question would have broad implications for sensor networks. It will lead to more accurate link estimators, which will improve the energy efficiency of almost every routing protocol and thereby increase data yields and lengthen lifetimes. It will lead to a quantitative understanding of the relationship between noise and signal strength, improving channel-sense media access protocols. It will provide insight on send queue and data-link retransmission policies, reducing the number of failed packets that waste energy. Answering this question would allow nodes to make good decisions on when, how, and with whom to communicate, which is a fundamental part of almost every sensornet application, protocol, and architecture. In short, it would improve the reliability, lifetime, and performance of mote-based sensor networks.

Our first step in answering this question is to establish what about these networks is understood, what remains uncertain, and whether or not these uncertainties present practical limitations in sensor network deployments. In Section 2, we revisit three relevant areas of prior work that provide the basis of this study: studies of the packet delivery behavior of wireless networks, the algorithms multihop and media access protocols use which consider these complexities, and observations on how these protocols perform in real world deployments in the face of unforeseen challenges. Taken together, these indicate that while we *do* understand what the complexities inherent to low-power networking are, we still *, do not* understand their root causes, how to effectively detect them, or how to deal with them.

Our second step in answering this question is to decompose it into smaller, more specific areas of inquiry. We describe our experimental methodologies in Section 3. While simple experiments can establish that at a high level current platforms have similar challenges to older ones [30], establishing root causes, however, re-

quires a broader set of experiments, as we must account for and quantify many unknowns, such as temporal trends, spatial correlation, environmental noise, channel effects, and hardware variations.

Section 4 describes our first set of experiments, which examine the relationship between packet deliveries and packet metadata such as received signal strength (RSSI) and chip correlation (CCI). Our results show that the reception "grey region" abruptly occurs when RSSI values fall below a certain threshold, and temporal variations are more likely a result of changes in signal strength than changes in the noise floor. In Section 5, we describe our second set of experiments, which examine noise traces and determine that there are variations in noise floor, that hardware is the cause, external noise is highly correlated spatially, verify the spectrum analysis that 802.11b interferes with most 802.15.4 channels, and show that noise spikes are 802.11b traffic.

Section 6 asks the next sub-question: if hardware effects account for variations in observed noise floor, how does this relate to link asymmetries? We answer this question by examining link asymmetries over long and short time scales. We find that noise floor variations are a cause of some packet reception rate (PRR) asymmetries, temporal RSSI variations lead to very few long-term but many short-term PRR asymmetries, and there are also RSSI asymmetries, which play a factor in these effects. In Section 7 we examine these temporal behaviors more closely, measuring the correlation in packet delivery failures to quantify how a lost packet might reflect a short-term change in PRR. We find that single packet failures are often a negative indicator, in that after a single packet failure the probability the next packet will be received is higher than the average PRR, but subsequent packet failures are often positive indicators, and might point to a temporary change in the PRR.

While our experiments explain many of the networking complexities of 802.15.4 mote platforms, our results are far from comprehensive. For example, while we do examine noise from 802.11 networks, we do not consider the effects of interference from concurrent transmissions within the network. While there are interference studies for the mica2 platform [34, 46], we are not aware of similar work on 802.15.4 mote platforms. Similarly, although space constraints prevent us from presenting all of our observations, our data is still only an initial investigation: greater certainty of our results is only possible after more comprehensive testing and independent corroboration.

Nevertheless, we feel the data is conclusive enough to present important findings on what the root caues of packet delivery failure and success are, as well as what indicators can be used to detect them. Overall, these findings suggest several ways in which current practices could be easily changed that would greatly improve the efficiency, performance, and lifetime of mote-based sensor networks. We defer a discussion of these suggestions to Section 8.

2 Related Studies

Experiments with early mote platforms established that low-power wireless networks have complex dynamics. These observations have guided the design and implementation of many protocols and systems. However, deployments that used these protocols and systems have demonstrated that many of them perform badly in practice, suggesting that the dynamics are not yet understood well enough. In this section, we overview these efforts, distilling a set of factors and considerations that remain uncertain or unknown, whose investigation may provide the knowledge that bridges the gap between research evaluation and practical use.

2.1 Packet Delivery

Ganesan et al [15] analyzed different protocol layers for rene motes, showing that even simple algorithms such as flooding had significant complexity at large scales. They observed that many node pairs had asymmetric packet reception rates, which they hypothesized were due to receive sensitivity differences, which Cerpa et al. [6] supported after swapping asymmetric node pairs and finding that the asymmetries were a product of the nodes and not the environment.

In order to better understand packet reception asymmetries, Woo et al [48] looked at packet reception rates (PRR) over distance for mica motes. They found that for a large range of distances, PRR and distance had no correlation and attributed this to hardware miscalibration. Zhao et al [50] confirmed the prevalence of this "grey region" but tentatively concluded that multipath effects were the probable cause, noting that further study was needed. All of these studies measured early mote platforms (e.g., rene, mica, and mica2) whose data-link stacks (e.g., encoding, CSMA, start symbol detection) resided primarily in software.

Ganesan et al. [15] showed that packet collisions, hidden terminals, link asymmetries, and the broadcast storm problem [28] make flooding a problematic approach for building trees. Whitehouse et al. demonstrated that frequency shift keying (FSK) radios, such as those on the mica2 platform, can recover from packet collisions where the stronger packet starts later by constantly looking for a start symbol [46]. Son et al. [34] took one step further and measured a precise RSSI envelope for when mica2 packets can be recovered. They showed that if the signal to interference plus noise ratio (SINR) is above a threshold, PRR is very high ($> 99.9\%$), and that this threshold varies for different nodes. These results suggest that SINR may be a good way to understand PRR more generally. If noise behaves in a simple fashion and RSSI values are stable over time, then RSSI might be a good determinant of packet delivery success or failure.

Cerpal et al. showed that PRR rates can change significantly over time, so that long-term PRR calculation can lead to very inaccurate results [7], suggesting instead that an instanteous measure of RNP – "required number of packets" – was preferable to a long-term PRR. This work also introduced the idea of using conditional probabilities in link estimation, an idea which we extend when considering the correlation between packet failures in Section 7.

Aguayo et al. [5] observed similar packet delivery behaviors in a 38-node 802.11 long haul urban mesh network, but concluded that they were most likely due to multipath effects as there was little correlation between PRR and SINR. However, their experimental methodology differs from those of the sensor network studies. For example, they consider average SINR ratios over second-long periods rather than on a per-packet basis. Nevertheless, the differences in conclusions between the efforts are interesting. Since 802.11b operates in the same ISM band (2.4 GHz) as 802.15.4 and uses a similar modulation scheme (QPSK), 802.11b transmitters could be significant sources of interference [47].

2.2 Sensor Networking

The conclusions of these experimental studies have greatly influenced sensor network protocol and system design. The grey region and link asymmetries have led some routing protocols to incorporate link estimation algorithms that maintain tables of candidate next hops. For example, because initial studies suggested that RSSI may not be well correlated with packet delivery success or failure, Woo et al. used packet sequence numbers to directly estimate PRR [48]. The expense of doing so, however, has led several more recent protocols, such as TinyOS's Drain and MultiHopLQI [39], as well as Moteiv Corporation's Boomerang [8], to use single samples of the chip correlation indicator (CCI) of the CC2420 radio as a measure of link quality; to the best of our knowledge, there are no evaluations of this approach in the literature. In one contrasting example, the collection layer of the SNMS management system uses a sum of per-hop RSSI values to select parents [40].

Mote data link layers generally have a CSMA MAC and use a constant randomized backoff policy. If the link layer detects an active channel, it selects a backoff timer from a uniform distribution over $[a,b]$, where b is usually under a packet time. Each successive busy channel detection resets the timer in the range $[a,b]$. Some link layers, such as the mica2, will backoff indefinitely, while others, such as the CC2420, will give up after a number (e.g., 8) tries.

Many TinyOS ad-hoc routing protocols use transmission queues to absorb bursts of forwarding traffic. TinyDiffusion [14], for example, maintains a 12-packet queue [1], beacon vector routing [13] maintains a 32-packet queue [2], and the TinyOS AODV [1] implementation [3] uses a 10-packet queue. All of these queues use an *immediate retransmission* policy. Once the link-layer finishes sending a packet, the queue immediately submits the next packet for transmission, where it enters the data-link layer backoff. There is no mechanism to modify queue entries once they are submitted. Together, this means that routing layers often enqueue several packets to a single destination, then transmit them back-to-back very quickly. This approach may be appropriate if packet losses are independent and identically distributed, but is not appropriate if packet losses are correlated.

2.3 Deployment Experiences

While empirical measurements have influenced several aspects of sensornet protocol design, such as demonstrating the need for link estimation, there are many more which have never been evaluated. This has meant that, in practice, the protocols often operate poorly when deployed.

Szewczyk et al. [35] presented network data from a deployment on an island off the coast of Maine. The design of the network assumed significant end-to-end packet losses would occur and so oversampled the environment. They measured packet delivery performance for a single-hop and a multihop network which used Woo et al.'s algorithms. PRR was initially satisfactory, but the multihop network deteriorated over time, with some networks delivering under 30% of its packets, some of which was due to significant base station outages. They note that while only 15% of the links that the routing algorithm selected were stable and long-lived, those links were responsible for 80% of the packets delivered.

Tolle et al. [41] reported similarly low yields from a network designed to monitor the microclimate of redwood trees, although in this case much of the network was unable to form a routing topology. Furthermore, approximately 15% of the nodes in the deployment died one week into the deployment by exhausting their batteries due to a problem in the time synchronization component of the routing protocol.

These results suggest that a gap exists between research algorithms and their performance in real deployments. While studies have quantified many of the difficulties in low-power wireless that make developing efficient and robust protocols difficult, none have established how to detect them. Without such information, protocols and systems will suffer from always being tuned to the environment in which they were developed. The next section describes a set of experiments whose purpose is to take the next step and provide insight on not only packet delivery performance, but more importantly, the *causes* of packet delivery success and failure.

3 Experimental Methodology

This section describes the radio, platforms, and testbeds we used in our experiments, as well as the experiments themselves.

[1] `tinyos-1.x/contrib/tinydiff/tos/lib/TXMan_TM.nc`

[2] `tinyos-1.x/contrib/bvr/tos/commstack/BVRQueuedSendM.nc`

[3] `tinyos-1.x/contrib/hsn/tos/lib/SimpleQueueM.nc`

3.1 CC2420

The CC2420 is a single-chip 2.4GHz band transceiver that is 802.15.4 compliant. 802.15.4 uses OQPSK modulation to send chips at 2MHz. Each 4-bit symbol is encoded in to a pseudorandom sequence of 32 chips. The CC2420 decodes a received symbol by correlating it with all the 16 different possible symbols. 802.15.4 can operate in 16 different channels (in the ISM 2.4 GHz band), which are spaced 5 MHz apart and occupy frequencies 2405 MHz - 2480 MHz. Some of these frequencies overlap with 802.11b frequencies, as shown in Figure 7.

The CC2420 provides two pieces of metadata about received packets. The first is its received signal strength indicator (RSSI), which is the strength in dBm of the RF signal received over the first eight symbols after the start of a packet frame. RSSI can also be sampled at other times, to detect the ambient RF energy. The second is the chip correlation indicator (CCI), which is an unsigned integer in the range of 50-110. The CC2420 calculates CCI over the same eight symbols of a packet as RSSI. Roughly speaking, CCI represents how close the received chip sequences were to the decided symbols: a 110 means that they were very close, while a 50 means there were lots of wrong chips.

3.2 Platforms

We used the Telos rev B mote [36] and the MicaZ mote [37] as our two primary experimental platforms, both of which have a CC2420 radio. The platforms have two principal differences. The first is the microcontroller: Telos motes have a 16-bit Texas Instruments MSP430 F1611 microcontroller, while MicaZs have an 8-bit Atmel ATmega128L. The second is their RF engineering: Telos motes have an integrated planar inverted F-style antenna (PIFA) printed directly on the circuit board, while the MicaZs have a detachable, quarter wave, monopole antenna connected to an MMCX jack on the MicaZ circuit board. Additionally, they have some different passive components, like the oscillator, their components are placed differently, and the Telos has an RF guard ring whereas the MicaZ does not. For outdoor experiments, we used Telos rev B motes.

To measure the interference effects of 802.11b on 802.15.4, we used a Dell Optiplex SX280 (a small form-factor PC) with a USB 802.11b card attached and a Sony VAIO with integrated 802.11b.

3.3 Testbeds

Most of the experiments use the Intel *Mirage* [22] testbed, which has 100 MicaZ motes, and a *university* testbed in our computer science department building, which has approximately 30 Telos nodes spread over approximately 2500 square feet. The nodes in both testbeds are on the ceiling. We used the university testbed for two reasons. First, it allowed us to repeat all of the Mirage experiments with Telos motes, albeit with a different physical layout and possibly different RF characteristics. Second, because Mirage is a public resource, we could not ask its maintainers to swap nodes or otherwise alter the environment.

In addition to the two wired testbeds, we used several other ad hoc setups to explore specific questions which emerged from our iterative analysis. The *Outdoor* testbed consisted of 11 Telos Rev. B nodes arranged in line with 2 meter spacing placed on 6 cm high plastic boxes in a grassy field. The *Interference* testbed consisted of the 802.11b devices placed four feet apart with a Telos Rev. B mote sitting half-way between the two laptops. The *Comparison* testbed had a transmitter mote (MicaZ or Telos) placed approximately 4.5 m away from a MicaZ and Telos mote which were once 5 cm apart. The nodes in the comparison testbed were on cluttered desks, approximately two feet above ground level, and during some experiments people walking by broke line-of-sight (LOS) between the nodes.

3.4 Experiments

With these platforms and testbeds, we performed six basic experiments. In all experiments we gathered the RSSI and CCI of every received packet, and kept track of all missed packets through sequence numbers.

Burst Broadcast: Each node broadcast a burst of between 100 and 2000 packets, separated by a delay (variable between experiments), and all other nodes that received this packet forwarded it over the UART to the base station. A PC controlled when a node started a burst and what inter-packet timing a node used.

Burst Unicast: Identical to burst broadcast, except that each node sent a separate burst of packets to every other node.

Round Robin: A PC sent packets with a predefined rate, cycling through the nodes as transmitters: if the PC sent k packets with an inter-packet interval of t and there were n nodes, then each node sent $\lfloor \frac{k}{n} \rfloor$ packets and sent a packet every nt. All packets were broadcasts. Nodes forwarded all packets received to the PC.

Background Noise: Each node sampled the noise level from the RSSI register at 4Hz and sent this data back to the base station. Each experiment collected data for 24 hours in order to examine the long-term noise floor.

Correlated Noise: Similar to the background noise experiment, except that a radio beacon synchronized the start of sampling. Sampling was at 128Hz rather than 4Hz, and each experiment ran for only a few minutes, limiting the effect of clock drift on time synchronization.

RSSI Comparison. A single transmitter sent packets at 16Hz, cycling through all of the available transmit power strengths (-30 dBm to 0 dBm), so that the interval between two packets of the same strength was approximately 2 seconds. This experiment was only performed on the comparison testbed.

Interference: We placed a Telos node between the two 802.11b devices set in ad-hoc mode. The PC ran an FTP server and the laptop downladed a large file. The Telos node sampled RSSI and sent it to a PC for collection. The mote continued sampling after the transfer completed in order to compute the effect of the 802.11b transmissions.

We repeated some experiments with different parameters. We performed the Burst and Round Robin experiments at 6 different power levels, ranging from -29 to 0 dBm, and three different 802.15.4 channels (11, 20, and 26). We performed the Burst Unicast experiment on the university testbed with inter-packet times of 10ms, 100ms, 1s, and 5s on channels 11 and 26. We collected Background Noise data on channels 11, 20, and 26 in indoor environments using the Mirage, university, and outdoor testbeds. We performed the Interference experiment with every combination of 802.11b and 802.15.4 channel.

4 Packet Reception

Prior studies have shown that wireless nodes have links with intermediate loss rates, and these loss rates can vary over time. We verified that these observations are true for the mote platforms in our study by running round-robin experiments on the Mirage and university testbeds. Figure 1 shows a packet delivery trace from a 4 hour period of a single node in the Mirage testbed (node 4). This trace, which is representative of the rest of the data, shows that most links have either near-perfect or zero PRRs, but a small number have the expected intermediate behavior. In particular, the PRR from node 30 to node 4 is approximately 40%, varies from 0% at 10-30 minutes to 75% at 40-80 minutes.

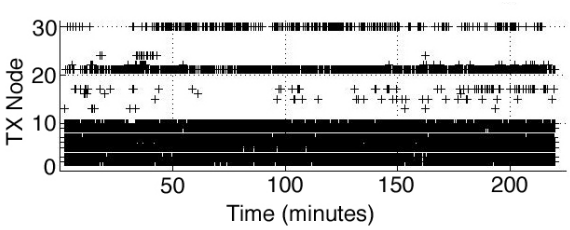

Figure 1. Packet reception over time (x axis) at node 4 from all other 29 nodes (y axis). Every packet received is marked by a '+'. Some nodes such as 2 and 3 have almost consistent packet reception while node 30 has packet reception that varies over time.

4.1 CCI, RSSI, and PRR Correlation

Radio propagation theory [31, 25] states that a low SNR threshold causes packet delivery failure, and interference studies on some mote platforms validate that this is a very good indicator on FSK radios [34]. Sampling the ambient and packet signal strengths can be a fast and inexpensive operation, as it may not even require powering up modulation hardware. However, initial studies argued that signal strength might not be a good indicator of PRR [50]. Furthermore, initial experiments on the Telos platform suggested that CCI might be a more accurate indicator of PRR.

One way to resolve these conflicting claims is to determine whether RSSI and CCI are good indicators of PRR. Variations in PRR are ultimately the result of variations of a channel's bit error rate, which a node can estimate with the CC2420's CCI. This raises the additional question of which is a better indicator of PRR: RSSI or CCI?

To answer this question, we examined packet traces from burst and round robin experiments on the Mirage and university testbeds, as well as a burst experiment on the outdoor testbed, plotting PRR versus RSSI and PRR vs CCI. As each burst has a short duration of a few seconds, the burst experiment results are more likely to be affected by brief noise spikes, but less likely to be affected by environmental changes (such as people walking around) that affect signal attenuation.

Figures 2(a), 2(c), and 2(e) show the RSSI results from the burst experiments. Over a packet burst, RSSI values are very stable: for most links the standard deviation is $\leq$ 1dBm. There is a strong correlation between RSSI and PRR, in that there is a sharp cliff around -87dBm: if a link is -87dBm or stronger, it is almost but not completely certain to have a PRR $\geq$ 99%. The fact that the variation between readings is small means that a single packet at -86dBm is usually sufficient to determine whether a link is good for the duration of a burst. Each of the plots has a small number of outliers. We examine the causes of some of these in depth later, in Section 4.3.

Figures 2(b), 2(d), and 2(f) show the CCI results. Over a packet burst, CCI values vary widely: standard deviations of 10 are common. The mean values over all of the packets, however, show a clear trend that corresponds to PRR. Combining the results from Figure 2, this suggests that while RSSI can provide a quick and accurate estimate of whether an incoming link is in the grey region or not, averaging CCI values over a number of packets can provide an estimate of where in the gray region a link is.

Given CCI's statistical nature, this variance makes sense. Since the CC2420 calculates CCI over eight separate symbols of 32 bits each, even if the chip errors are independent there will be significant statistical variation. The CC2420 data sheet does not describe exactly how it computes a CCI value, otherwise we could apply statistical models to try to determine whether chip errors are inde-

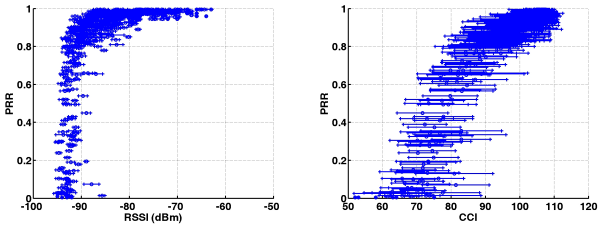

(a) PRR vs RSSI, Indoor Telos (b) PRR vs CCI, Indoor Telos

(c) PRR vs RSSI, Indoor MicaZ (d) PRR vs CCI, Indoor MicaZ

(e) PRR vs RSSI, Outdoor Telos (f) PRR vs CCI, Outdoor Telos

Figure 2. PRR versus RSSI and PRR versus CCI in a short packet burst. Each data point is a directional node pair. The error bars show one standard deviation of the measured RSSI/CCI values. The two platforms have similar distributions, and Telos behaves similarly indoors and outdoors.

pendent or correlated.

4.2 Long-Term Behavior

Figure 2 shows that RSSI is very stable over a short burst of packets. Environmental changes, however, can cause significant changes in propagation gain. RSSI is a very good predictor of whether a short-term link is high quality or not. This raises a follow-up question: is RSSI a good long-term predictor of PRR?

Figures 3(a), and 3(c) show the answer to this question. RSSI and PRR show the same relationship as in the burst experiment: there is a sharp cliff around -87dBm. RSSI values, however, have much greater variance, and there are more outliers. Figures 3(b), and 3(d) show the CCI plot of a long-term experiment. These CCI plots is very similar to that of a burst experiment, but just as with RSSI has more noise and variation.

Looking closer at the outliers in Figure 3 (the five or so data points with an average RSSI around -85dBm and PRRs between 60% and 20%) show that their RSSI variation is higher than most of the other links. These links could represent data points that moved between the plateau of good connectivity and the cliff of variable connectivity: because the plot is not a straight line, they appear as outliers below the knee of the curve.

4.3 Node 4

The plots in Figures 2- 3 lead to several hypotheses. The first is that temporal variations in PRR are more likely due to variations

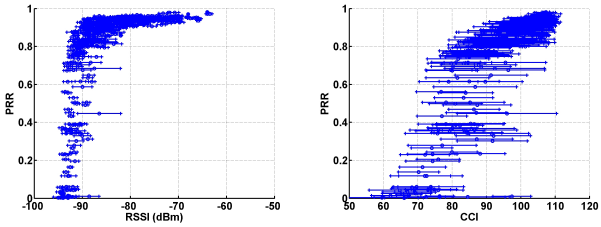

(a) PRR vs. RSSI, Indoor Telos (b) PRR vs. CCI, Indoor Telos

(c) PRR vs. RSSI, Indoor MicaZ (d) PRR vs. CCI, Indoor MicaZ

Figure 3. PRR versus RSSI and PRR versus CC2420 CCI in an 8 hour trace. Each data point is a directional node pair. The error bars show one standard deviation of the measured RSSI/CCI values.

in received signal strength than noise. This is because if the noise floor variations were to be large we would see a wider cliff in the plots of RSSI vs PRR. The second is that outliers on the curve are caused by a link moving between the plateau of good connectivity and the shelf of highly variable connectivity.

As an initial test of these hypotheses, we looked in depth at the behavior of a single node over a 4-hour trace, node 4 from the Mirage round-robin experiment at channel 11. We examined four data series: the packet reception over time, the instantaneous ambient signal strength sampled at 10Hz, a smoothed noise plot where each data point represents the average over 40s, and the RSSI of packets received from one particular node whose connectivity showed temporal variation (node 30).

The first trace is the same as Figure 1, rotated 90 degrees. The middle two traces show the signal strength values observed at node 4. The left of these two traces shows instantaneous signal strength measurements. The plot shows that there are many large spikes of approximately 35dBm. However, the time scale of the plot obscures the fact that these spikes comprise only a small percentage of the samples, and makes it difficult to see the actual noise floor. The right-center plot is therefore a coarser grained plot where each data point is the average over 400 samples (40s). The fact that the averages are between -93 and -92.5dBm shows that the noise spikes are rare. Furthermore, there does not seem to be a strong correlation between noise values and the PRR from node 30; while the period between ten and thirty minutes had some noise spikes and a very low PRR, the period of highest noise (around the 200 minute mark) had a burst of packet successes. This suggests that noise is not the principal factor behind temporal variations in PRR.

The far right plot of Figure 4 is an expansion of the column of node 30 in the far left plot: it shows the RSSI of packets received from node 30. Node 4 received one packet with an RSSI of -93 dBm, two with -92dBm, many at -91 and -90dBm, and a few at -89dBm. Making conclusive statements on the relationship between RSSI and PRR requires sampling the signal from node 30 for every packet, received or not. As this trace only shows received packets, we can only make reasonable hypotheses at what the causes are. For

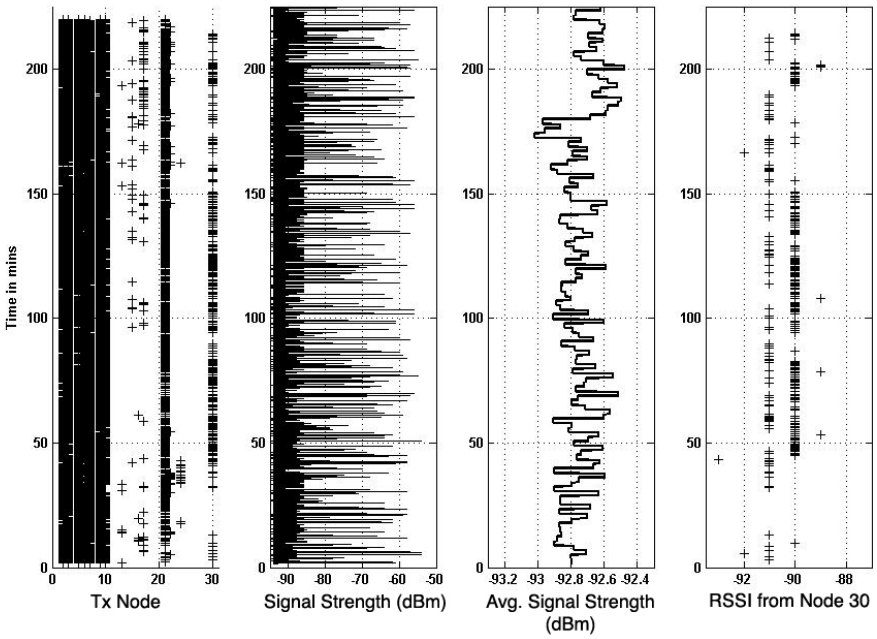

Figure 4. Observed behavior at node 4. The first plot on the left shows packet loss over time. The second plot shows the measured signal strength of the channel, which shows very, short-lived spikes. The third plot shows the singnal strength of the channel averaged over 400 samples (40s), which shows that there are not significant variations. The last plot, on the right, shows the RSSI distribution of packets received from node 30 over time.

example, while node 4 received packets at both -91 and -90dBm, the receptions at -90dBm form dense clumps, while the receptions at -91dBm are sparsely scattered. This suggests that the dense clumps of packet receptions – the periods of higher PRR – are correlated with periods of higher received signal strength.

The threshold for reasonable packet reception at node 4 seems to be approximately -91dBm. This raises the question of whether this threshold is due to environmental noise or other factors. The CC2420 datasheet [21] says that its receive sensitivity is on average -94dBm, with a worst-case sensitivity of -90dBm. The mode of its measured noise values is -94dBm, and the CC2420 has a 3dB sensitivity.

5 Interference and Noise

The prior section showed that long-term changes in PRR are the result of changes in the RSSI of a borderline link dips slightly. However, the RSSI grey region (where RSSI values that can have intermediate PRRs) in Figures 2- 3 is approximately 6dBm wide. This is significantly greater than the expected receive sensitivity of 3dBm. While the noise floor at a given node is fairly stable, it is possible that the RSSI grey region is due to inter-node differences in the noise floor. Additionally, the noise traces had large spikes of 35dBm or more. If these spikes are not spatially correlated – that two nodes close to each other do not both observe them – then this would be a significant problem for CSMA algorithms. Furthermore, understanding the source of these spikes and whether or not they can be avoided through channel selection could improve the throughput and energy efficiency of protocols.

5.1 Noise Floor Differences

We ran a background noise experiment on the univeristy testbed, using 802.15.4 channel 11, collecting approximately 160,000 samples from each mote. Figure 5 shows a two-minute subset of the

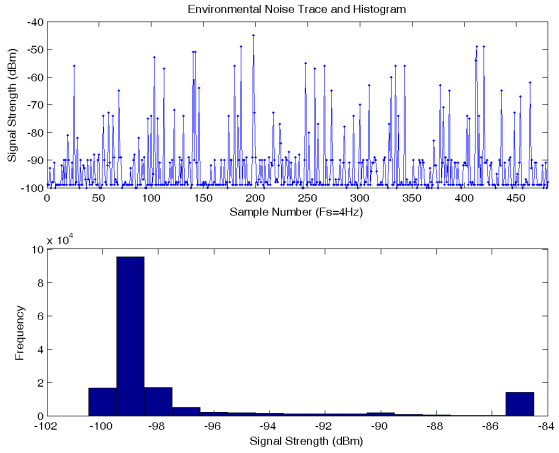

Figure 5. Sampled signal strength trace and histogram. (a) Sampled Signal Strength (dBm) measured over a two minute period at a single mote; (b) A histogram of the sampled signal strength over 11 hours measured at the same mote.

values measured at a single node as well as a histogram of the values over an 11 hour period. 59.4% of the samples have a value of -99 dBm, 10.3% have a value of -100 dBm, and 0.002% have a value less than or equal to -101 dBm. This indicates that the mode of the samples is a good estimator of a mote's noise floor. The distribution of samples is right-tailed, with more than 8.6% of the samples having a value greater than -85 dBm.

Table 1. Distribution of estimated noise floor across 25 motes.

SSI (dBm)	-99	-98	-97	-96	-95	-94
# Nodes	7	6	11	0	0	1

Figure 5 is representative of the other motes with one exception. Table 1 shows the distribution of the noise floor values across the motes in the university testbed. The noise floor at all but one mote falls in the range -99 dBm to -97 dBm. The outlier had a noise floor of -94 dBm. To determine whether the higher noise floor indicates an enviromental difference at the mote's location or a hardware anomaly, we swapped the positions of the outlier and a quieter node and re-ran the experiment. The noise floors showed that the outlier values were due to the node and not the environment.

5.2 Correlated Noise

Our first background noise experiment collected signal strength statistics at 4 Hz, which is sufficient to measure the noise floor and obtain the overall distribution of noise values. One of the key observations of the background noise experiment is that the distribution of noise values across nodes are quite similar. The traces also show noise spikes of varying amplitude ranges. This raises the question of whether the noise at different motes is merely statistically similar or actually well correlated. Since background noise experiments do not synchronize sampling across motes, answering this question requires running a correlated noise experiment.

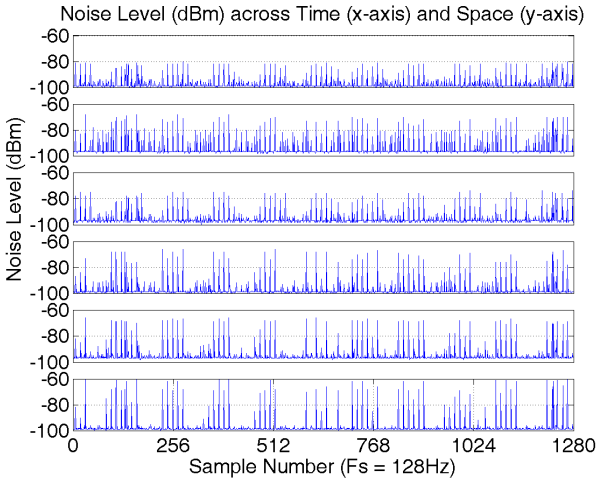

Figure 6. Sampled signal strength over 10 seconds and across six motes.

We ran a correlated noise experiment on the university testbed. To avoid picking up other TinyOS users, we measured channel 20 (the default TinyOS channel is 11). Figure 6 shows the noise level over 10 seconds at six different motes that are located at the five convex corners of the testbed and the one centrally-located concave corner and Table 2 shows their correlation coefficients. The data appear well correlated both visually and statistically. The correlation coefficients range from 0.80 to 0.97, which when squared gives a range of 0.64 to 0.94, and describes the proportion of variance in common between the noise levels at different motes.

Table 2. Correlation coefficients of the noise measurements.

Node	1	6	13	19	26	28
1	1.00	0.85	0.95	0.80	0.82	0.80
6	0.85	1.00	0.83	0.91	0.88	0.84
13	0.95	0.83	1.00	0.78	0.80	0.77
19	0.80	0.91	0.78	1.00	0.97	0.94
26	0.82	0.88	0.80	0.97	1.00	0.97
28	0.80	0.84	0.77	0.94	0.97	1.00

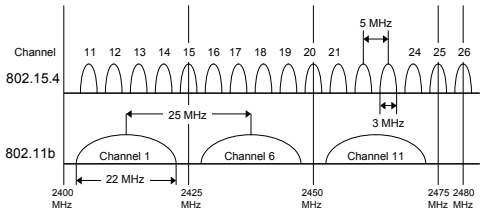

(a) 802.11b and 802.15.4 spectrum usage.

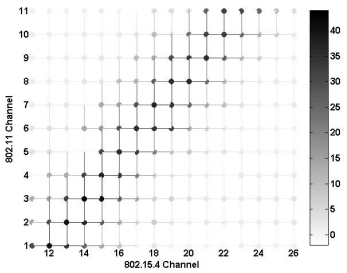

(b) Measured 802.11b interference with 802.15.4 (dBm).

Figure 7. Interactions between 802.11b and 802.15.4.

5.3 Interference from 802.11b

The plots in Figure 6 show that the periodic noise spikes are correlated over space. While it is possible that there were other 802.15.4 devices using channel 20, there is a more likely cause: 802.11b. The standard beacon period for an 802.11b access point is 100 Kμsec (0.1024 sec or 9.7656 Hz), which is consistent with the data, and there was an access point located close to mote 28.

802.11b and 802.15.4 share the 2.4GHz spectrum band. Because of their different data rates, their channels occupy different spectrum widths. Figure 7 shows how they overlap. Because 802.11b channels 1, 6, and 11 do not overlap significantly with each other, they are what most wireless APs use. To better measure the level of interference that 802.11b causes in the 802.15.4 band, we performed an interference experiment for every combination of 802.15.4 and 802.11b channels. Figure 7 shows the *difference* in received signal strength between the presence and absence of 802.11b traffic. The data indicate an inopportune choice of 802.15.4 channel can result in even 35 dBm to 45 dBm of interference from 802.11b traffic. The data also shows that only 802.15.4 channel 26 is largely immune from 802.11b interference. This suggests that the selection of the 802.15.4 should consider the channel of a nearby 802.11 network as 802.11b interferes with 802.15.4 in a way that reduces packet reception rates [10, 47].

One plausible conclusion from these experiments is that the default 802.15.4 channel should be set to 26 to minimize 802.11b in-

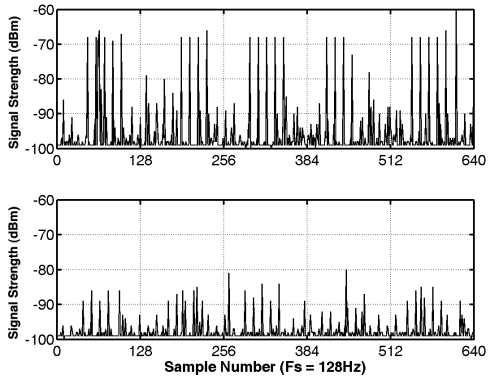

Figure 8. The effect of an 802.11 access point on an 802.15.4 node on channel 20. The access point and node are separated by 2 m and a wall. (a) Sampled signal strength with the access point operating. (b) Sampled signal strength with the access point unplugged. Difference between two indicates spikes are interference from 802.11b beacons.

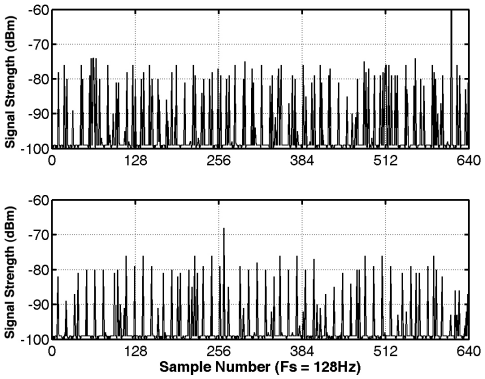

Figure 9. The effect of an 802.11 access point on an 802.15.4 node on channel 20. The access point and node are separated by 7 m and a wall. (a) Sampled signal strength with the access point operating. (b) Sampled signal strength with the access point unplugged. Smaller spikes correspond with greater distance from access point. More frequent spikes in (a) and remaining spikes in (b) indicate interference from two access points.

terference. While this is a perfectly reasonable conclusion given our results, it is important to recognize that multi-channel radios are often optimized for the center channel, so selecting by default a channel at one extreme may lower the radio performance. In addition, channel 26 may not always be free of 802.11b interference: while the Figure 7 shows simple, distinct 802.11 channel widths of 22MHz, in reality the specification is a bit more complex.

In the United States, there are 11 legal 802.11b channels numbered from channel 1, which has a center frequency of 2.412 GHz, to channel 11, which has a center frequency of 2.462 GHz. There are 16 available 802.15.4 channels numbered from 11, which has a center frequency of 2.405 GHz, to channel 26, which has a center frequency of 2.480 GHz. Both the 802.11b and 802.15.4 center frequencies are separated by 5 MHz but the 802.11b channels have a much greater bandwidth. The 802.11b standard specifies the center frequency of each channel as well as a spectral mask for that channel. The spectral mask requires the signal be attenuated by 30 dB or more from its peak energy at $\pm$ 11 MHz from the center frequency, and attenuated by at least 50 dB from its peak energy at $\pm$ 22 MHz from the center frequency.[4] Therefore, even though our experiments show little interference from 802.11b on 802.15.4 channel 26, for a sufficiently powerful transmitter, there could still be considerable interference.

Figure 6 demonstrates how this can manifest. The simplistic depiction in Figure 7 would suggest that 802.15.4 channel 20 would not experience interference from 802.11 channels 6 and 11 that access points use, but the noise spikes in Figure 6 suggest otherwise. To verify that the spikes were coming from the access point, we performed two additional background noise experiments across multiple 802.15.4 channels. Figure 8 shows the signal strength at the node closest to the intefering 802.11 access point. A decrease of 15 dBm to 20 dBm in peak signal strength is observed when the access point is turned off. Figure 9 shows the signal strength at a node further away from the access point and the less pronounced decrease in signal strength. The key difference appears to be a reduction in the *density* but not peak signal strength of interfering traffic, suggesting that the second node is situated between two access points and continues to sense traffic from the other access point.

In the previous section we saw that packet reception rates and RSSI showed a strong correlation, and long-term temporal variations in PRR could be attributed to variations in signal strength rather than noise. However, there is an RSSI grey region of approximately 6dBm in which PRR is highly variable. In this section we examined this grey region more closely, and showed that nodes can have a range of noise floors, which represent one form of hardware variation. We also examined the sources of noise spikes and established that they are from 802.11b traffic, they have strong spatial correlation, and can be for the most part avoided by selecting the right 802.15.4 channel.

6 Packet Reception Asymmetry

Prior sections showed that RSSI can effectively predict packet loss, that RSSI varies over long time periods, and noise floors differ. These conclusions give us the basis from which we can address two major unanswered questions: why are there asymmetric links, and how can a node detect them? For this study, we define a bidirectional link between a node m and a node n to be asymmetric if $|PRR_m - PRR_n| > 0.4$.[5] But before we address these questions, we first quantify how common asymmetric links are in 802.15.4, as there are interesting temporal complexities to consider.

6.1 Temporal Characteristics of Asymmetry

Using a unicast burst experiment on 30 nodes of the Mirage testbed on channel 11, we measured the PRR between every node pair and from this computed link asymmetries. This experiment has two interesting qualities: first, the two directions of a node pair may have been measured several minutes apart, and the PRR is calculated over a very brief interval. We repeated this experiment twice, once on a Wednesday evening and once on a Saturday afternoon. Figure 10 shows the results. There are many asymmetric links. However, looking more closely, only two of the 16 or so asymmetrics links occur in both experiments. This suggests that there may be significant temporal effects in link asymmetry.

[4] In contrast, 802.15.4 channels have a 3 MHz "bandwidth."

[5] While ETX might be a better metric, as it reflects the true cost of a link, ETX has interesting temporal properties due to its assuming a retransmission policy. As Section 7 will show, this means that ETX cannot be considered in a precise fashion without making constraining assumptions that might reduce the applicability of the conclusions.

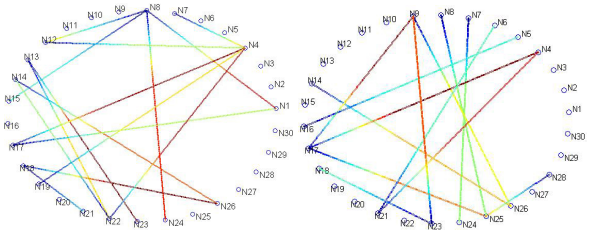

(a) Wednesday (b) Saturday

Figure 10. PRR Asymmetry for 30 nodes in the Mirage testbed over a 200 packet unicast burst for two different trials. The 30 nodes were a subset of the testbed nodes that covered the entire lab area. While they are shown in a circle solely for visualization purposes, nodes close to each other on the circle were close to each other in the testbed. Nodes having asymmetry are connected using a colored line, where the red end of the line is the node that had trouble receiving packets. A larger gradient on the line indicates higher asymmetry. While each trial had a significant number of asymmetric links, there are only two (N14-N26 and N17-N4) present in both.

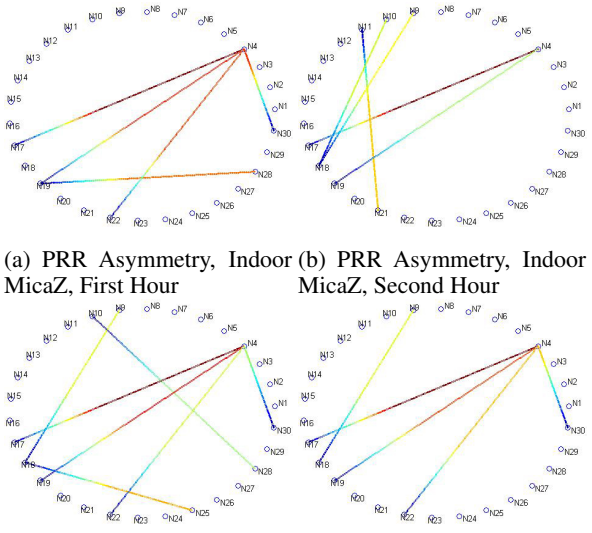

(a) PRR Asymmetry, Indoor MicaZ, First Hour (b) PRR Asymmetry, Indoor MicaZ, Second Hour

(c) PRR Asymmetry, Indoor MicaZ, Third Hour (d) PRR Asymmetry, Indoor MicaZ, Forth Hour

Figure 11. Hour-by-hour asymmetry plots for a four hour round-robin experiment on the Mirage testbed. The visualization methodology is the same as in Figure 10. A small number of links such as N17→N4 are consistently asymmetric and there are also transiently asymmetric links such as N18→N10. Node 4 also seems to be a "bad node," in that many of the stable asymmetric links have it as a bad receiver.

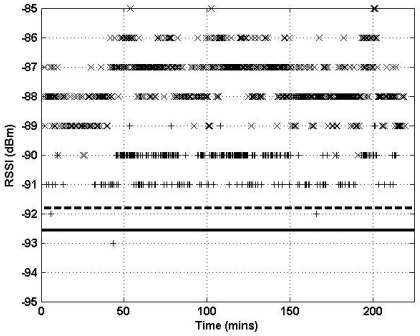

Figure 12. RSSI variation over time at nodes 30 (+) and at node 4 (×) for packets received from each other. Both observe variation of a few dB, and 4 observes an RSSI of approximately 2dBm lower than 30 does.

To determine the time scale of variations in PRR asymmetry, we examined the data from the round-robin experiment used in Figure 4 and calculated link asymmetry over four separate one hour periods. Figure 11 shows the results. It is clear that a few links such as N17→N4 are consistently asymmetric while some such as N18→N10 are not. Furthermore, the number of asymmetric links is much smaller. These results suggest that there are significant differences between long-term and short-term link behavior.

Since Section 4 showed that temporal variations in RSSI were the cause of temporal changes in PRR, it is a reasonable hypothesis that it is the cause here as well. Figure 4 supports this hypothesis: in the second hour, node 4 is able to receive packets from node 30 because the RSSI increased to be mostly -90dBm readings rather than -91dBm, as also seen in Figure 11(b).

Table 3. Distribution of estimated noise floor across the motes in the Mirage round robin experiment. Only 26 of the 30 nodes reported noise data.

SSI (dBm)	-98	-97	-96	-95	-94	-93	-92
# Nodes	5	8	4	3	2	3	1

6.2 Causes of Asymmetry

Together, the RSSI temporal variation results from Section 4 and the inter-node noise variation results from Section 5 present a picture of what causes asymmetric links and why they might change over time. Table 3 shows the distribution of noise floors in the Mirage round-robin experiment. Just as with the university testbed, there are significant inter-node variations, which would affect SNR and therefore lead to PRR asymmetry. However, in comparison to the university testbed (Table 1), they have a larger range and are also several dBm higher. This could be due either to environmental variations or differences between the platforms. These minor differences aside, round-robing experiments on the university testbed had similar asymmetry results.

Figure 13 ties all of these results together. It shows node 4's view of its communication. It receives no packets below its noise floor (mode). Figure 3(c) showed that the distributions of RSSI values are fairly small: most had standard deviations of 1dB or less. This plot, however, shows that there are several extreme outliers. Node 4 receives no packets below its noise floor (-93dBm), and very few below one standard deviation above that (-90dBm), and its noise floor is one of the highest, as shown by the distribution in Table 3.

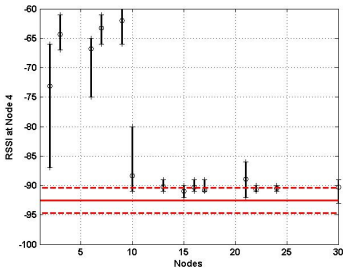

Figure 13. Average noise at node 4 and RSSI of packets received from all nodes for round robin experiment on Mirage. The circle on each vertical line marks the average RSSI while the ends of each line correspond to the minimum and maximum RSSI of packets received from that node. It receives no packets below its noise floor (-93dBm), and very few below one standard deviation above that (-90dBm).

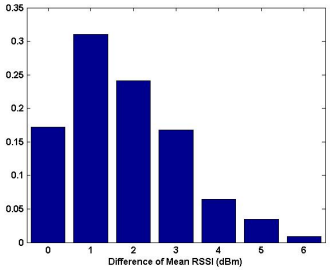

Figure 14. Disribution of RSSI asymmetries in the Mirage round-robin experiment. 50% of all communicating node pairs have pairwise RSSI differences of 2dBm or below.

Returning to Figure 11, node 4 had four asymmetric links, with nodes 17, 19, 22, and 30. In each of these asymmetric links, node 4 was the bad receiver. Examining the RSSI of these links, we can see that all of them are on the edge of receivable RSSI. One, in particular, 19, has no successfully delivered packets. Examining the reverse direction, node 19's noise floor was -98dBm, and the average RSSI of received packets from node 4 was -93dBm; unless there were significant RSSI asymmetry in node 4's favor, it is unlikely to receive any packets.

RSSI asymmetry also contributes to PRR assymetry. While RF theory (and the laws of physics) state that the two directions of RF propagation have identical attenuation, in practice this is not the case. Figure 14 shows a distribution of the RSSI asymmetries in the Mirage round-robin experiment. The largest asymmetry is 6dBm.

7 Packet Loss Correlation

Examining link asymmetries showed that there are significant short-term variations in PRR, but that over longer time periods PRRs are stable. This suggests that in addition to being able to gauge long-term PRR values, motes may need to be able to quickly adapt to abrupt and transient changes. Independent and identically distributed (i.i.d.) events are a common assumption in routing protocols and the formulation of wireless networking. In this model, the success or failure of each packet transmitted is independent and uniform. The results in Section 6 directly contradict this abstraction. However, if packet delivery is not i.i.d., what is it?

If packet delivery failures are not independent, they are correlated. More precisely, consider the conditional probability $F(n)$, which is the probability that the next packet will fail if the past n failed. The temporal behavior observed in Section 6 suggests that a more expressive form might be useful in practice, $F(n,t)$, where t denotes the time interval between the packets. If $F(n)$ is uniform for all n, then packet delivery failures are i.i.d. By definition, if $F(n,t)$ is uniform, then $F(n,kt)$ for any integer k must be uniform as well. If packet delivery failures are not independent, then $F(n,t)$ will not be uniform.

7.1 Correlation Periods

We ran a series of broadcast burst experiments on the university testbed using channel 11 and 26 to compute $F(n,t)$ for a range of t values. The differences in time scale of the short-term and long-term experiments in Section 6 were over three orders of magnitude (seconds versus hours). By examining packet failure correlations at a wide range of time scales, we can hopefully understand the degree of temporal dynamics in PRR variation.

Figure 15 shows $F(n,t)$ for channel 26 2000 packet bursts with t values of 5ms, 10ms, 100ms, and 1s. Each graph represents the results from a single chosen unidirectional link; each experiment had over 700 links. Each link selected experienced an overall packet reception rate between 40% and 60%[6], and was chosen because it is representative of the common behavior observed. The plots show that packet losses are not i.i.d. Instead, the conditional probability of packet delivery failure fluctuates greatly, often with a value greater then 90%, even though the packet reception rate is around 50%.

$F(n,t)$ describes the conditional probability of loss, but it does easily represent how common strings of losses are. The probability that there will be k consecutive losses with a packet timing of t is $P(k,t)$, which is defined as

$$P(c,t) = (1 - F(c,t)) \cdot \prod_{n=1}^{c-1} F(n,t)$$

Figures 15(e)-15(h) show the corresponding $P(c,t)$ for the $F(n,t)$ plots. We see that in all four graphs, single packet loss bursts make up over 50% of loss sequences. Looking at the top 4 graphs, we see that, given the fact that only one packet has been lost since the last successful transmission, the probability that the next packet is lost is approximately the expected *1 - PRR*. Coupled with the bottom four graphs, this suggests that a good portion of the losses that occur after a successful transmission are due to a short-lived condition, and consequently are independent of the result of the next transmission. However, by the same token, we notice that beginning with 2 consecutive losses, the probability of losing the next packet increases rapidly. Therefore, this suggests that if at least two consecutive packets are lost, then there exists a correlation between successive packet transmissions. Since the test was performed on channel 26, 802.11b interference should have been minimal. Only one node broadcasted at a given time, to the best of our knowledge no other devices were operating on channel 26 concurrently, and as Figure 8 showed, channel 26 sees little if any interference from other devices. We therefore assume that packet losses due to interference are unlikely.

However, as packet losses are not i.i.d,, some process is creating a dependency. Returning to the results of Section 4, the temporal variations in RSSI suggest a possible cause. Because there is temporal correlation in RSSI, down-swings cause periods of lower PRR, which in turn create a statistical dependency in packet loss.

[6]The link selected for the 1s burst had an overall packet reception rate of 81%, as there existed no links that had an overall packet reception rate between 10% and 80%.

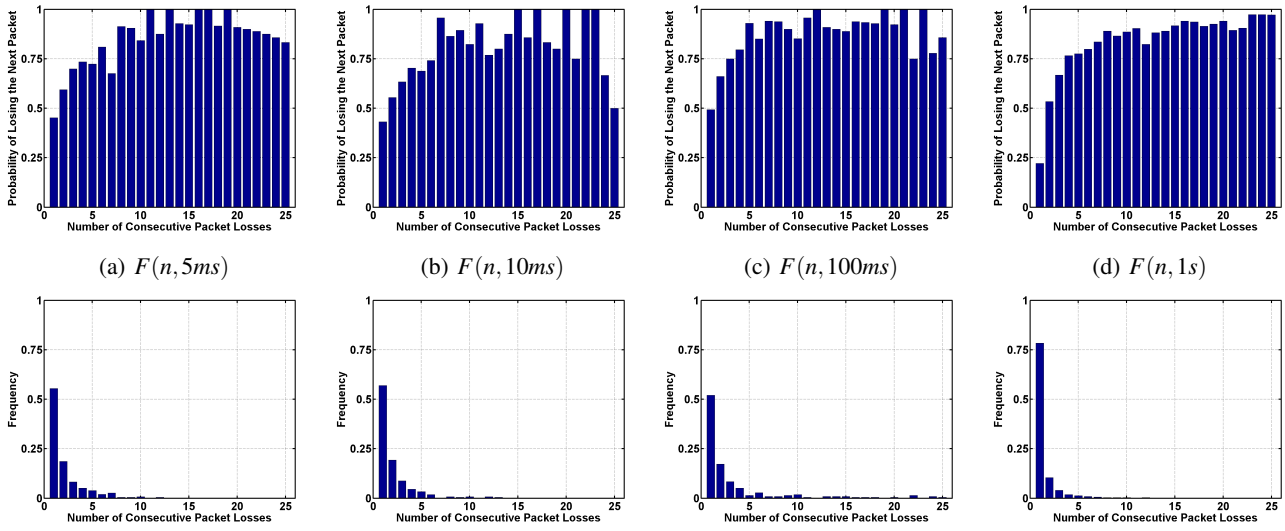

(a) $F(n,5ms)$ (b) $F(n,10ms)$ (c) $F(n,100ms)$ (d) $F(n,1s)$

(e) $P(c,5ms)$ (f) $P(c,10ms)$ (g) $P(c,100ms)$ (h) $P(c,1s)$

Figure 15. $F(n,t)$ and $P(n,t)$ for 4 values of t using channel 26 on the university testbed. The 5ms, 10ms, and 100ms plots are 2000 packets, while the 1s plot is 60000 packets. Each F,P pair is from a single node pair with an intermediate PRR, but link PRR variations mean that each t value is a different node pair.

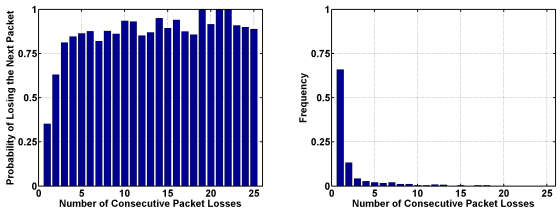

(a) 500ms Round Robin (b) 500ms Round Robin

Figure 16. $F(n,t)$ and $P(c,t)$ plots from a representative pair in a round-robin experiment on the university testbed using 802.15.4 channel 26.

7.2 Long Term Correlation

Consequently, the cause of this correlation can not be determined from this data set alone, and requires additional experiments. Since this data was for packets separated by delays of one second or less, the next natural question is whether this correlation exists for greater amounts of delay.

Given that packet losses are not i.i.d. due to temporal RSSI correlation, the next natural question is whether these dependencies exist for greater inter-packet timings. Because burst experiments run serially, running a long-term experiment for an entire testbed is prohibitively long, long enough that there could be significant environmental changes. Therefore, to examine longer-term behavior, we ran a round-robin experiment on the university testbed, also using channel 26. In both cases, packet transmissions are timed so as not to interfere: the round robin experiment can be thought of as n burst experiments running in parallel.

Although this used a different experimental methodology, the data can be examined in the same way. Only one node transmits at a given time, and using a 500ms delay between node transmissions, the effective delay is 14s (500ms * 28 nodes per round). The trace lasted for 12 hours, during which each node broadcast approximately 3000 packets. Figure 16 displays the results of this experiment. Initially, it seems no different from the short-term burst results, but a closer look reveals some subtle differences. Just as in the short experiments, given a successful transmission followed by a packet loss, the probability that the next packet is lost is approximately *1 - PRR*. However, from there, the probabilities increase more rapidly than in the burst experiments, and remain reasonably stable.

Further and deeper exploration into this phenomena is needed in order to be able to state the cause definitively, but the data seems to point to an interesting result. The nearly flat portion of the graph seems to indicates that at that point, successive packet losses are independent of each other. Recall that if losses are independent, the probability of the next packet being lost is *1 - PRR*. Looking at the graph, this would imply a PRR of approximately 15%, while we know that the overall PRR for that link is closer to 60%. The explanation for this could be that with packets spread so far apart, two or three consecutive losses could indicate that the quality of the link has temporarily deteriorated, and this deterioration is fairly stable and long-lived.

More interestingly, the shape of the curve suggests an interesting relationship between different PRR states. A single loss indicates a relatively low probability of additional losses (25%), and 65% of all losses are single loss events. This indicates that most of the time, the PRR is above 75%. The long tail in Figure 16(b) shows that there are a small number of long strings of packet losses, beyond what would be expected statistically from a uniform PRR. Within these long tails, $F(n)$ is reasonably stable, around 85%. This suggests a predimoniantly bimodal behavior; drawing a parallel to the observations of node 4 in a Mirage round-robin experiment (Figure 4), this suggests movements between two different principal signal strengths.

8 Implications and Future Work

As with any study, some of the observations in Sections 4- 7 are conclusive, while others are preliminary or speculative. The former establish a set of guidelines for protocol and system design, many of which contradict currently held beliefs and practices, while the lat-

ter provide future areas of investigation and experimental methodologies.

In the platforms we measured, there are three root causes which lead to a very complex set of higher level behaviors. The basic explanation behind these causes is the signal to noise ratio, where the noise is predominantly a hardware effect, as we did not consider concurrent transmissions. The first cause is that a node pair can have significant (as much as 6 dBm) signal strength asymmetry, but the asymmetry between 50% of pairs is within 1dBm. The second is that RSSI values change over time. This leads to temporal variations in PRR as intermediate links oscillate between high and low SNR, resulting in packet losses that are not i.i.d. The third is differences in what noise floor nodes observe, which is a product of their hardware. The fact that these causes can be easily measured, combined with their results, such as packet loss correlation, the temporal dynamics of asymmetry, and the grey region, has broad implications on the design and implementation of many systems, protocols, and algorithms, such as those we discussed in Section 2.

8.1 MAC Protocols

The interference 802.11b introduces can greatly effect MAC protocols, as the 802.11b device may not consider an 802.15.4 signal strong enough to warrant backoff, given that the latter is narrowband in comparison to the former. However, 802.11b can interfere with 802.15.4 and cause significant packet losses. This means that MAC protocols cannot assume that they are the sole users of the channel, which means that analytical results showing optimal behavior in a closed system may be very difficult to apply to real-world systems.

Furthermore, 802.11b interference may affect low power listening MAC layers [11, 29] that use RSSI as a wakeup signal. 802.11b packets can lead a protocol to wake up, wasting energy. However, given that most sensornets are heterogeneous and have 802.11b-based microservers [12] or similar devices, this behavior could also be used in a productive fashion. Because 802.11b has a longer range and these devices are not as energy-constrained, a microserver can issue a wakeup signal to most of the network.

Deployments in similar environments to those we measured could take advantage of the fact that RSSI values are stable over the short term to make intelligent MAC decisions. For example, if an RTS/CTS protocol such as SMAC [49] can assume that RSSI does not change significantly over a rapid 30 packet burst, since there is a strong correlation between RSSI and PRR, the CTS can estimate how many retransmissions may be needed in order to more tightly limit the wakeup time.

Some data-link protocols use link-level retransmissions in conjunction with explicit or implicit acknowledgements. Generally, the policy in these schemes is to retransmit immediately. Our results in packet loss correlation suggest that repeated immediate retransmissions may be ineffective. Instead, if faced with multiple consecutive losses, servicing another send request to another node may be a more efficient use of both the channel and a node's energy reserves.

Finally, the default clear channel assessment (CCA) value on the CC2420 is -72dBm. CC2420 stacks in most sensor network operating systems that support CC2420 platforms, such as TinyOS [2], SOS [18], and MOS [3], use this value. If the CC2420 less then -72dBm in the channel, it will transmit a packet. The PRR vs. RSSI plots in Figure 3 show that this can cause nodes to transmit packets that collide with ones which the nodes could readily receive since reception below 90dBm is common. Additionally, because 802.11b is a major source of interference and is highly spatially correlated and RSSI assymetries are limited to approximately 6dB, a MAC protocol could use an adaptive CCA scheme, in which it selects a CCA value based on its estimate of the RSSI asymmetry, correlation between noise values, and receive sensitivity in the presence of interference.

8.2 Link Estimation

Studies of earlier platforms argued that RSSI was not a good PRR indicator [50]. Our results show that this conclusion does not hold for more recent platforms, which has significant effects for link estimation. As RSSI values are stable and have low variance, changes can be detected very quickly. While by themselves RSSI values do not provide very much information about the PRR grey region, combined with a knowledge of the receiver's noise floor, they can. Of course, RSSI asymmetry means that a transmitter must estimate the RSSI at its destination, but nonetheless, a transmitter can at the very least quickly and accurately estimate whether it is using an intermediate or good link.

Most of our results focused on the relationship of RSSI with packet delivery behavior, but the CCI data in Figures 2 and 3 show that CCI can also be very useful. Although CCI readings are noisy, such that a single reading on an intermediate link provides very little information, the average CCI over several packets can be an accurate predictor of PRR. Combined with the RSSI results, this suggests that a hybrid estimator, one which generally uses RSSI but if needed takes advantage of CCI, might be an effective way to inexpensively detect good links yet be able to effectively consider intermediate ones.

Finally, there are often tradeoffs in link estimation between estimation cost, agility, and accuracy. The temporal effects and variations shown in Sections 4 and 6 suggest that, just as a hybrid approach for estimation itself is possible, a hybrid approach in terms of time could be worthwhile. In this model, nodes maintain long-term estimates, which short-term variations can temporarily overshadow, but which provide a longer-term assessment of a possible neighbor. The information used for these two mechanisms might be very different, as they have different resource-fidelity tradeoffs, especially given that probing a bad link can be more costly than just using a different one.

8.3 Network Protocols and Localization

Short-term PRR variations mean that network protocols must be able to either quickly (e.g., within a few packets) change routes or store packets for later forwarding. The former policy may have problematic interactions with underlying link estimators, however, as sudden shifts in traffic patterns may affect the estimation algorithm depending on when and how it gathers data.

Mote RAM constraints and the belief that many links are asymmetric made neighbor blacklisting an unattractive approach for early protocols. However, our measurements show that there are very few long-term asymmetric links; this means that revisiting this assumption and using this technique in protocols running on CC2420 networks may be useful.

Given the ubiquity of 802.11b, it is very likely that sensor networks deployed in indoor (especially office or urban) environments will be within range of a few 802.11b base stations. Section 5 suggests that nodes can use these beacons for RSSI-based distance estimation. There is a tremendous amount of literature on the subject, and while most deployed efforts have shown that it is unlikely to compute very precise locations [45], others have shown that it can produce good rough estimates [27]. As 802.11b beacons have a regular period, nodes could possibly distinguish sources without needing to decode the signal, although the details of channel overlapping (and the fact that a mote may not be able to know which 802.11 channel is being used) could introduce significant complexity.

8.4 Wireless Networking

While we believe this study is the first to closely examine many aspects of the behavior of low-power, 802.15.4-based devices, there have been in-depth studies of 802.11 [5, 32]. These two protocols use the same spectrum and have similar modulation schemes (BPSK or QPSK vs. OQPSK). However, these studies and ours reach opposite conclusions. Aguayo et al. observe very little correlation between SNR and PRR [5] and attributes this to mulitpath effects, while Creis asserts that RSSI asymmetries are a product of the environment rather than the node or wireless card [32].

The Aguayo et al. study used Roofnet, a much larger-scale network that operates on the roofs of buildings in a dense urban area. It is possible that the differences in our conclusions stem from differences in environment. Alternatively, as their study was in an urban area, there is almost certainly be a great deal of interfering traffic, a complexity that our study does not consider.

Another possibility is that the discrepancies come from differences in experimental methodology. Our results show that 802.15.4 can have significant short-term RSSI variations, which we measure on a packet-by-packet basis. Aguayo et al., undoubtedly due to the fact that 802.11b has many more packets per second, consider the SNR averaged over a second. Furthermore, their hardware enables them to measure the noise floor immediately before and after a packet, which our experimental setup did not. Just as viewing PRR vs RSSI, it may be that averaging SNR over time intervals leads to different conclusions than on a packet-by-packet basis.

8.5 Open Questions

On one hand, our data point strongly at a wide range of conclusions that have broad implications for a broad spectrum of problems in wireless sensor networking. On the other, and is often the case in these studies, they raise just as many questions as they answer. This is in part due to several limitations in our studies: we did not consider concurrent transmissions, while our data indicates that outdoor deployments behave similarly, the variety that this environment class encompasses deserves greater study, and we did not consider some of the more complex aspects of 802.15.4, such as acknowledgments.

Going forward, we believe our results raise several important questions which, if answered, could provide excellent insight on what challenges in this domain are inherent, and what challenges are temporary. Specifically,

1. 6dBm seems too great an RSSI asymmetry to be the result of CMOS electronics. What causes the observed hardware variations?

2. Are long-term RSSI trends caused by the surrounding environment or aspects of the hardware?

3. Why does 802.15.4 show very different behavior from what has been reported for 802.11? How much of this is hardware and how much of it is the physical layer?

Our belief is that for wireless sensor networks to truly become robust, long-lived, and effective tools for society, we must understand their networking deeply and fully, and hope that this work is an initial step towards this goal.

References

[1] Tinyaodv protocol. TinyOS repository, tinyos-1.x/contrib/hsn.

[2] *Tenth International Conference on Architectural Support for Programming Languages and Operating Systems*, San Jose, CA, USA, Oct. 2002. ACM Press.

[3] H. Abrach, S. Bhatti, J. Carlson, H. Dai, J. Rose, A. Sheth, B. Shucker, J. Deng, and R. Han. MANTIS: System Support for MultimodAl NeTworks of In-situ Sensors. In *2nd ACM International Workshop on Wireless Sensor Networks and Applications (WSNA)*, 2003.

[4] R. Adler, M. Flanigan, J. Huang, R. Kling, N. Kushalnagar, L. Nachman, C.-Y. Wang, and M. Yarvis. Intel mote 2: An advanced platform for demanding sensor network applications. In *Proceedings of the Second ACM Conferences on Embedded Networked Sensor Systems (SenSys)*, 2005.

[5] D. Aguayo, J. C. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *SIGCOMM*, pages 121–132, 2004.

[6] A. Cerpa, N. Busek, and D. Estrin. Scale: A tool for simple connectivity assessment in lossy environments. Technical Report 0021, Sept. 2003.

[7] A. Cerpa, J. L. Wong, M. Potkonjak, and D. Estrin. Temporal properties of low power wireless links: modeling and implications on multi-hop routing. In *MobiHoc '05: Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, pages 414–425, New York, NY, USA, 2005. ACM Press.

[8] M. Corporation. Boomerang. `http://www.moteiv.com`.

[9] M. Corporation. Telos datasheet. `http://www.moteiv.com/products/docs/tmote-sky-datasheet.pdf`.

[10] I. Crossbow. Micaz-based zigbee and wifi coexistance. www.xbow.com/products/Product_pdf_files/Wireless_pdf/ZigBeeandWiFiInterference.pdf.

[11] A. El-Hoiydi. Aloha with preamble sampling for sporadic traffic in ad hoc wireless sensor networks. In *Proceedings of IEEE International Conference on Communications*, New York, NY, Apr. 2002.

[12] J. Elson, L. Girod, and D. Estrin. Emstar: Development with high system visibility. *IEEE Wireless Communication Magazine*, December 2004.

[13] R. Fonesca, D. Culler, S. Ratnasamy, S. Shenker, and I. Stoica. Beacon vector routing: Scalable point-to-point routing in wireless sensornets. In submission.

[14] D. Ganesan. TinyDiffusion Application Programmer's Interface API 0.1. `http://www.isi.edu/scadds/papers/tinydiffusion-v0.1.pdf`.

[15] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks, 2002.

[16] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. An empirical study of epidemic algorithms in large scale multihop wireless networks. UCLA Computer Science Technical Report UCLA/CSD-TR 02-0013, 2002.

[17] R. Govindan, E. Kohler, D. Estrin, F. Bian, K. Chintalapudi, O. Gnawali, S. Rangwala, R. Gummadi, and T. Stathopoulos. Tenet: An architecture for tiered embedded networks, November 10 2005.

[18] C.-C. Han, R. Kumar, R. Shea, E. Kohler, and M. Srivastava. A dynamic operating system for sensor nodes. In *MobiSYS '05: Proceedings of the 3rd international conference on Mobile systems, applications, and services*, 2005.

[19] J. Hill and D. E. Culler. Mica: a wireless platform for deeply embedded networks. *IEEE Micro*, 22(6):12–24, nov/dec 2002.

[20] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. E. Culler, and K. S. J. Pister. System Architecture Directions for Networked Sensors. In *Architectural Support for Programming Languages and Operating Systems*, pages 93–104, 2000. TinyOS is available at http://webs.cs.berkeley.edu.

[21] C. Inc. Cc2420 data sheet. `http://www.chipcon.com/files/CC2420_Data_Sheet_1_0.pdf`, 2003.

[22] Intel Research Berkeley. Mirage testbed. `https://mirage.berkeley.intel-research.net/`.

[23] B. Karp and H. T. Kung. Gpsr: greedy perimeter stateless routing for wireless networks. In *MOBICOM*, pages 243–254, 2000.

[24] R. Kling. Intel research mote. `http://webs.cs.berkeley.edu/retreat-1-03/slides/imote-nest-q103-03-dist.pdf`.

[25] W. C. Y. Lee. *Mobile Cellular Telecommunications, Second Edition*. McGraw-Hill Book Company, 1995.

[26] M. Leopold, M. B. Dydensborg, and P. Bonnet. Bluetooth and sensor networks: a reality check. In *Proceedings of the Second ACM Conferences on Embedded Networked Sensor Systems (SenSys)*, pages 103–113, 2003.

[27] K. Lorincz and M. Welsh. Motetrack: A robust, decentralized approach to rf-based location tracking. In *Proceedings of the International Workshop on Location and Context-Awareness (LoCA 2005)*.

[28] S.-Y. Ni, Y.-C. Tseng, Y.-S. Chen, and J.-P. Sheu. The broadcast storm problem in a mobile ad hoc network. In *Proceedings of the fifth annual ACM/IEEE international conference on Mobile computing and networking*, pages 151–162. ACM Press, 1999.

[29] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the Second ACM Conferences on Embedded Networked Sensor Systems (SenSys)*, 2004.

[30] J. Polastre, R. Szewczyk, and D. E. Culler. Telos: enabling ultra-low power wireless research. In *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 364–369, 2005.

[31] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice-Hall, 1996.

[32] C. Reis. An empirical characterization of wireless network behavior. Quals Paper, University of Washington, http://www.cs.washington.edu/homes/creis/publications.shtml, 2005.

[33] I. Rhee, A. Warrier, M. Aia, and J. Min. Z-mac: a hybrid mac for wireless sensor networks. In *Proceedings of the Third ACM Conferences on Embedded Networked Sensor Systems (SenSys)*, pages 90–101, 2005.

[34] D. Son, B. Krishnamachari, and J. Heidemann. Experimental analysis of concurrent packet transmissions in low-power wireless networks. Technical Report

ISI-TR-2005-609, Nov. 2005.
[35] R. Szewczyk, J. Polastre, A. Mainwaring, and D. Culler. An analysis of a large scale habitat monitoring application. In *Proceedings of the Second ACM Conference on Embedded Networked Sensor Systems (SenSys 2004)*, 2004.
[36] C. Technology. Telosb datasheet. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/TelosB_Datasheet.pdf`.
[37] C. Technology. Micaz datasheet. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/MICAz_Kit_Datasheet.pdf`, 2006.
[38] The Institute of Electrical and Electronics Engineers, Inc. Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs), Oct. 2003.
[39] TinyOS. MultiHopLQI. `http://www.tinyos.net/tinyos-1.x/tos/lib/MultiHopLQI`, 2004.
[40] G. Tolle and D. Culler. Design of an application-cooperative management system for wireless sensor networks. In *Proceedings of Second European Workshop on Wireless Sensor Networks(EWSN 2005)*, 2005.
[41] G. Tolle, J. Polastre, R. Szewczyk, D. E. Culler, N. Turner, K. Tu, S. Burgess, T. Dawson, P. Buonadonna, D. Gay, and W. Hong. A macroscope in the redwoods. In *Proceedings of the Second ACM Conferences on Embedded Networked Sensor Systems (SenSys)*, pages 51–63, 2005.
[42] University of California, Berkeley. Mica2 schematics. `http://webs.cs.berkeley.edu/tos/hardware/design/ORCAD_FILES/MICA2/6310-0306-01ACLEAN.pdf`, Mar. 2003.
[43] T. van Dam and K. Langendoen. An adaptive energy-efficient mac protocol for wireless sensor networks. In *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems*, Los Angeles, CA, Nov. 2003.
[44] M. Welsh and G. Mainland. Programming sensor networks with abstract regions. In *Proceedings of the First USENIX/ACM Symposium on Network Systems Design and Implementation (NSDI)*, 2004.
[45] K. Whitehouse, C. Karlof, A. Woo, F. Jiang, and D. E. Culler. The effects of ranging noise on multihop localization: an empirical study. In *Proceedings of the Fourth International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 73–80, 2005.
[46] 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.
[47] C. Won, J.-H. Youn, H. Ali, H. Sharif, and J. Deogun. Adaptive radio channel allocation for supporting coexistence of 802.15.4 and 802.11b. In *Proceedings of the 62nd IEEE Vehicular Technology Conference (VTC2005-Fall)*, September 2005.
[48] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proceedings of the first international conference on Embedded networked sensor systems*, pages 14–27. ACM Press, 2003.
[49] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient mac protocol for wireless sensor networks. In *In Proceedings of the 21st International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, New York, NY, June 2002.
[50] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of the First International Conference on Embedded Network Sensor Systems*, 2003.

A Hardware Comparison

Our experiments use both Telos Rev B and MicaZ motes. Although both of these motes use the CC2420 radio, their different antenna, passive component, and physical layout choices raise the question of whether these differences affect signal strength and packet delivery performance. To explore this question, we performed a series of side-by-side RSSI comparison experiments on the comparison testbed. Each experiment had a transmitter mote send packets to two side-by-side receiver motes. One receiver was a micaZ and one was a Telos. We used two different receiver pairs, $\{\{MicaZ_1, Telos_1\}$ and $\{MicaZ_2, Telos_2\}$, and two different transmitters, $MicaZ_3$ and $Telos_3$, generating data traces from eight node pairs.

The two receivers in each experiment were placed side-by-side approximately 5 cm from each other and equidistant from the transmitter which was approximately 4.5 m away. The nodes were all on cluttered desks. For successive experiments in which the same transmitter was used, the transmitter was not moved between the experiments. Each transmitter sent packets on 802.15.4 channel 11 at a 16Hz rate, cycling through all of the available transmit power strengths (-30 dBm to 0 dBm) so that the interval between two packets of the same strength was approximately two seconds. Every packet transmission had a unique sequence number, allowing us to compare the RSSI of a single transmission at two receivers.

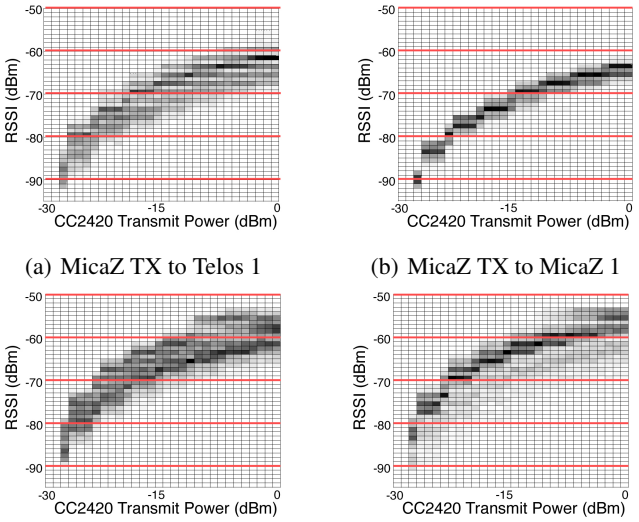

(a) MicaZ TX to Telos 1 (b) MicaZ TX to MicaZ 1

(c) Telos TX to Telos 1 (d) Telos TX to MicaZ 1

Figure 17. RSSI value distributions from experiments measuring the effect of a platform. The X axis is transmit power and the Y-axis is RSSI. Each cell corresponds to packets transmitted at a power level that were received with an RSSI value, and the darkness of the cell indicates how many were received.

We provide a cautionary note against generalizing the results reported here. First, since only channel 11 was used in this experiment, we cannot generalize about the radio performance of Telos and MicaZ motes since the radio hardware for the different motes may be optimized for a different center frequencies. Second, since different antennas have different directional gain, yet our experiments only considered one direction, we cannot generalize about omni-directional antenna gain. Third, since multipath effects can cause variations in received signal strength. Despite these caveats, this data provides a quantitative basis for comparing the results of the earlier sections.

Figure 17 shows the results. On one hand, there are what seem to be a few trends: micaz nodes, for example, tend to have lower RSSI values than their Telos counterparts, but the distribution of their RSSI values seems to be tighter. On the other, given the noise and complexity in the data, drawing concrete conclusions from this small study is problematic. At first glance, it seems that their RF behavior are more complex than a simple determination of one being "better" than the other. For example, while higher RSSI values can let a node transmit further, more stable RSSI values can make its algorithms run better, and which is preferable depends on the protocol as well as traffic pattern. While our results are in and of themselves unclear, what is clear is that this may deserve an entire study by itself.