An Empirical Study of Low Power Wireless

Kannan Srinivasan*, Prabal Dutta+, Arsalan Tavakoli+ and Philip Levis*
* Stanford University, + University of California, Berkeley

1. ABSTRACT

We present empirical measurements of the packet delivery performance of the latest sensor platforms: Micaz and Telos motes. In this paper, we present observations that have implications to a set of common assumptions protocol designers make while designing sensornet protocols – specifically – the MAC and network layer protocols. We first distill these common assumptions in to a conceptual model and show how our observations support or dispute these assumptions. Understanding the implications of these observations to the conceptual model can improve future protocol designs.

2. INTRODUCTION

The complexity of the real world forces us to design protocols and systems in terms of simplifying abstractions. A wired network can be thought of as a graph; an operating system process acts as if it has sole control of a CPU; a packet either arrives or it doesn't. To achieve this simplicity, these abstractions make assumptions about how the real artifact behaves. In the case of examples such as wired networks, processes, and packets, these assumptions are reasonable except in some extreme cases, and so allow us to imagine, design, and build more complex systems that then work in practice.

One of the greatest challenges in wireless research has been that its conceptual models have been unable to give this safety of abstraction. Wireless networks are not graphs, because links are not edges: two transmissions to different destinations can easily collide. Communication does not follow a unit disc model, as real RF propagation is uneven and intereference is dependent not only on the strength of the signals but also who transmits first. Historically, this mismatch between abstraction and reality has been a tremendous impediment to protocol design.

The first step towards a good conceptual model of wireless packet communication is to experimentally determine what underlying assumptions in current models can cause protocols, when deployed, to behave differently than expected. Towards this end, this paper summarizes a 3 year study of 802.15.4 networks, placed in the context of a conceptual model typically used when designing protocols today. Early sensor network platforms used many different radios such as the TDA5250 [31], TR1000 [56], and CC1000 [10]. More recently, many mote designs have settled on a common link layer, 802.15.4, and this stability seems long-lived. While some of the observations made of earlier platforms are applicable to wireless systems in general, such as the broadcast storm problem [48, 22], other observations such as a large grey region [73] might be an artifact of the radio. These earlier platforms had simple modulation schemes such as on-off keying (OOK), amplitude shift keying (ASK), and frequency shift keying (FSK). In contrast, 802.15.4 uses much more advanced orthogonal quadrature phase shift keying (OQPSK) and direct sequence spread spectrum (DSSS). This means that such observations might be different for 802.15.4. Therefore, revisiting the earlier studies in the context of 802.15.4 might lead to different conclusions.

This study builds on those of earlier platforms in two ways. First, in addition to observation, it discusses their implications to common protocol design principles. To do this, it presents a commonly used conceptual model of a wireless sensor network, and describes four assumptions in the model which many protocol designs make. Each section in the paper examines how its experimental observations support or refute the model and its assumptions.

Second, the experimental variables most commonly used in prior studies – distance, orientation, environment – present a "human-eye" view of a network, in that they are variables a node cannot 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 and measure, gives better insight into how protocols or systems might make decisions. For example, the 802.15.4 radio chip we study provides a received signal strength indicator (RSSI). In addition to received packet strength, a node can use RSSI to compute the noise floor when there are no transmissions. As nodes can directly measure RSSI, examining it allows us to understand the network characteristics and how software can possibly adapt. For example, Section 9 shows that the noise floor differences at nodes can cause long term link quality asymmetries.

This paper makes three research contributions. First, it presents the first comprehensive, in-depth 802.15.4 link layer measurement study. It presents results and observations on two different platforms and multiple testbeds. This study uses measurements that are visible to nodes and makes observations on the temporal trends, channel effects, and spatial correlation of 802.15.4 links. Second, it addresses the implications of our observations to protocol design. It presents four major assumptions made while designing sensornet protocols as a conceptual model:

1. stability: link quality (reception ratio) changes slowly compared to the data rate,

Observation	Section	Implications to the conceptual model
Over short periods, links exhibit either 0% or 100% packet reception ratio (PRR). Short periods have few links with PRR between 10% and 90% i.e. intermediate links. The portion of intermediate links increases with time.	Sec. 5	Estimates from infrequent beacons need not be applicable to frequent data transmissions.
The reception ratio a link observes depends on the channel.	Sec. 5	Protocol performance can vary over different channels.
Links have temporally correlated reception.	Sec. 6	Assuming independent reception over time is not always valid.
External interference from 802.11 can cause losses at multiple nodes.	Sec. 8	Assuming no spatial correlation of losses is not always valid.
Per-node received signal strength indicator (RSSI) and noise floor variations cause difference in long term packet reception ratios between the forward and reverse links (asymmetric links).	Sec. 9	Assuming that forward and reverse link can have different PRRs (directed link) is valid.
Acknowledgement reception ratio (ARR) is usually greater than the packet reception ratio (PRR).	Sec. 10	Using PRR in the place of ARR is not valid and can lead to inaccurate estimates of link quality.

Table 1: The key observations of this paper.

2. channel: link quality is the same on all channels,
3. spatial: losses on different links are independent, and
4. ack: acknowledgement and packet delivery ratios are the same.

The paper discusses the implications of each of its experimental observations to these four assumptions. It shows that these assumptions are not always valid. Third, the paper presents the underlying phenomena that cause many of the observations. It shows that measurable quantities, such as signal strength, noise floor, and external noise can explain the observed link behaviors. Measuring these quantities for networks can help us reason how protocols may perform on a given network. Therefore, these observations may be used to make nodes able to adapt their protocols to achieve better performance.

Table 1 summarizes our findings and briefly discusses their implications to the conceptual model.

This paper synthesizes and extends our prior work in this area, distilling our three year study down to a set of the most important findings. One earlier workshop paper based on this work examined the prediction value of two hardware indicators (signal strength and chip correlation) on intermediate links [64]. Another workshop paper examined how 802.15.4 might affect IP routing [60]. Technical reports [61, 63] presenting some of the content in this paper have been widely referred to in the community [51, 33, 20, 47, 17]. This paper collects all of these results from peer-reviewed workshop papers and technical reports. It also adds new results and observations, digging deepeer into the underlying phenomena that affect packet delivery (Sections 7 and 8). It also presents new results on the spatial correlation of packet delivery (Sectiond 8). Finally, it discusses the implications of all of these experimental observations to low power wireless protocol design, providing guidance and insight for future protocol research.

We believe that understanding the link layer is critical for high-quality sensor network protocols and systems. Understanding the link layer can improve the energy efficiency of almost every protocol, increase data yields, and lengthen network lifetimes. Understanding the link layer can improve media access protocols as well as provide insight on send queue and data-link retransmission policies. Understanding the link layer will allow systems to make good decisions on when, how, and with whom to communicate, a fundamental consideration in almost every sensornet application, protocol, and architecture. In short, a better understanding of the link layer will improve the reliability, robustness, lifetime, and performance of low power wireless networks.

3. HARDWARE

This section describes the radio, platforms, and testbeds of our experimental study.

3.1 CC2420

The CC2420 is an 802.15.4 compliant transceiver. Although there are other 802.15.4 chips available in the market such as the RF230 chip [4], the CC2420 is the most widely used chip in sensor network research: Telos [45], Micaz [15] and Imote2 [2] platforms use CC2420. In this study, we focus on motes which have a CC2420.

802.15.4 has 16 non-overlapping channels, spaced 5 MHz apart. They occupy frequencies 2405-2480 MHz. 802.15.4 uses a direct sequence spread spectrum OQPSK modulation to send chips at 2MHz. 32 chips encode a 4-bit symbol, providing a physical layer bandwidth of 250kbps. The CC2420 uses soft chip decision: rather than convert chips to bits and match against the encodings, it decodes by choosing the symbol which maximizes chip correlation. 802.15.4 shares the same band as 802.11b [28] and Bluetooth [29].

The CC2420 attaches two pieces of metadata to every received packet, RSSI and CCI (chip correlation indicator)[1]. It measures both over the first eight symbols (32 bits, 125μs)

[1]CCI is sometimes refered to as LQI (link quality indicator), but LQI is actually a link quality metric defined in the 802.15.4 standard: the CC2420 datasheet [11] suggests using RSSI to calculate LQI. To avoid confusion we do not use the term LQI and use RSSI and CCI without any conversion.

Testbed	Platform	Number	Environment
Mirage	Micaz	30-100	6000 square feet indoor office.
University	Telos	30	2500 square feet indoor office.
Lake	Telos	20	Outdoor dry lake with people moving around.

Table 2: The Experimental Testbeds.

of a received packet. The RSSI (received signal strength indicator) is the RF signal strength, in dBm. The CCI ranges from 50 to 110 – high is good. CCI represents the correlation between the received symbol and the symbol to which it is mapped after the radio does soft decoding. The CC2420 only calculates CCI on received packets, but continuously calculates RSSI. So, software can read RSSI at any time to measure ambient RF energy.

3.2 Platforms

The Telos revB mote [16] and the MicaZ mote [15] are our two primary experimental platforms. For the sake of this study, the principal difference of the platforms 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 different passive components, like the oscillator. They have their components placed differently and the Telos has an RF guard ring whereas the MicaZ does not. Despite the differences in the RF engineering between the two platforms, the observations presented in this paper apply to both.

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. To measure the interference effects of Bluetooth, we used the laptop/PC pair with USB Bluetooth adapters.

3.3 Testbeds

We studied 3 different testbeds and several ad-hoc setups. Table 2 lists the 3 testbeds. Most of our experiments use the Intel Mirage testbed [32] of 100 MicaZ nodes spread over 6000 square feet and a University testbed of 30 Telos nodes spread over approximately 2500 square feet in UC Berkeley's Soda Hall. Both the Mirage and the University testbeds have their nodes on the ceiling. The Mirage testbed represents an ad-hoc network spread over an entire floor of a large office building, while the University testbed's size more closely resembles a small home. Unlike Mirage, which is a public resource, the University testbed was under our control. This allowed us to swap nodes for experiments and otherwise alter the environment. Both of the indoor testbeds have wired backchannels for controlling and communicating with their nodes. Nodes in the Mirage testbed transmitted at full power (0 dBm) while nodes in the University testbed transmitted at -25 dBm due to its small area. The third testbed, the Lake testbed, consisted of 20 Telos nodes with 4 feet spacing between each other in a dry lakebed on Stanford University campus. The nodes had clear line of sight and were arranged in a linear topology. There were people moving around the Lake testbed.

In addition to the three testbeds, we used several short term setups to explore specific questions which emerged from our iterative analysis. We describe these setups as they arise in the paper. All experiments use the standard TinyOS 2.0 CSMA MAC layer. This paper studies how links behave in the absence of concurrent transmissions: all experimental setups were designed to prevent collisions between nodes under the control of the experiment. External RF sources could (and in many cases do, as we explore) cause packet losses.

4. CONCEPTUAL MODEL

Because wireless, or networks in general, are very complex, protocol designers use conceptual models to explain and reason about a protocol. These conceptual models make simplifying assumptions about how a network behaves. For example, many MAC protocols retransmit immediately after a packet failure assuming that the losses on a link are independent over time. Such a protocol would work well if the losses are independent. However, when the losses are bursty then such a protocol can waste energy. Thus, evaluating the validity of these asumptions can greatly aid protocol design.

This section distills a set of assumptions that the majority of protocols make today and presents them as a conceptual model of a wireless network. The sections to come explore the validity of this conceptual model and how this impacts protocol design.

The majority of the contemporary sensornet protocols depend on at least one of the following assumptions:

1. stability: link quality (reception ratio) changes slowly compared to the data rate,
2. channel: protocol performance is the same on all channels,
3. spatial: losses on different links are independent, and
4. acknowledgment: acknowledgement and packet delivery ratios are the same.

Figure 1 shows a conceptual model incorporating all the four assumptions above. It shows how protocol designers conceptually model a wireless network: a directed graph with different packet reception ratios for different links. The graph shows observed reception ratio to be stable, assumes that these reception ratios remain the same for all channels of operation. The graph assumes that the losses on different links are independent and assumes same reception ratios for acknowledgement and data packets.

4.1 Stability Assumption

Link estimation protocols like MultihopLQI [46], CentRoute [65] and MintRoute [43] assume that the observed reception ratios do not change much over time. They use

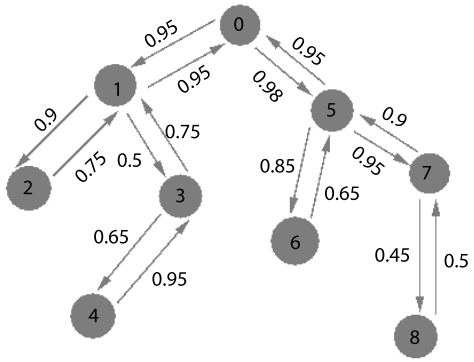

Figure 1: The Conceptual Model. The graph does not show temporal and spatial correlation observed reception ratios. It does not show if the channel of operation is relevant to protocol performance. It assumes that the observed reception ratio is the same for both data and acknowledgement packets.

estimates computed from periodic control beacons for sending data packets. Specifically, MultihopLQI can have beacons as infrequent as every 32s while the data can be as frequent as every 10ms. Such infrequent beacons are also common in high power wireless technologies. Srcr, designed for 802.11 networks, sends a beacon every second, computes the link quality over 10 such beacons and uses this estimate for sending more frequent data packets [6]. Beacon vector routing (BVR) [19], S4 [42] and optimized link state routing (OLSR) [12] also send infrequent control packets to manage network neighbor table. The assumption that these protocols make is that the link qualities change slowly over time. This is the *stability* assumption in the conceptual model.

4.2 Channel Assumption

Simulators like EmStar [18], TOSSIM [40], NS-2 [1] and GlomoSIM [72] do not consider the channel of operation as a simulation parameter. Protocol designers also present protocol performance results from running their protocols on a single channel in testbeds. The inherent assumption is that the protocol performs similarly on all the channels. Many multi-channel MAC protocols like Y-MAC [35], McMAC [25] and TFMAC [34] assume that the observed reception ratio is similar in all the channels. Y-MAC makes the potential senders and the receiver hop channels in a sequence so as to reduce collisions in the commonly used (base) channel and to improve the overall network energy efficiency. It implicitly assumes that the protocol performance is the same across all the channels. For example, if the packet delivery is near zero in all the channels except the base channel then by switching from the base channel, the nodes may spend more energy on failed transmissions than they may spend on collisions had they used only the base channel. This is the *channel* assumption in the conceptual model.

4.3 Spatial Assumption

Protocols that use random linear codes to encode packets work best when the chance a receiver loses a packet is independent of the chance that any other receiver loses that packet. Rateless Deluge, a dissemination protocol, encodes n data packets using linear codes and transmits m encoded packets, where $m > n$ such that upon receiving any n linearly independent encoded packets, a receiver can decode the original n packets [24]. In the event that different single hop receivers do not receive different packets, the transmitter can send only one additional coded packet so that all the receivers can decode all the n packets unlike Deluge that has to send every packet missing at every receiver. MORE, a routing protocol for 802.11, uses a similar linear coding technique to send multicast packets [9]. The assumption is that the losses over multiple links are not correlated. If the links were to have spatially correlated delivery then linear coding of packets may be cost inefficient as receivers are likely to lose the same packet in which case just sending that packet, not coded, would suffice. Simulators like TOSSIM, EmStar and ns-2 also assume independence of reception across different links. This is the *spatial* assumption in the conceptual model.

4.4 Ack Assumption

Existing energy-based route selection metrics such as ETX (the expected number of transmissions per packet [71]) and its derivatives [36] use the product of forward and reverse packet reception ratios. They assume that the acknowledgment reception ratio is the same as the packet recption ratio in the reverse direction. Some protocols like Srcr calculate reception ratio of small packets, comparable in size to ACK, to estimate acknowledgement reception ratio. However, actual acknowledgements immediately follow a successful packet which, due to neighboring nodes backing-off, keeps the channel free from other transmissions and therefore can have a better reception ratio [13]. This is the *ack* assumption in the conceptual model.

4.5 Summary

Many protocols work well on testbeds but do not act as we expect in the actual deployments [37, 68, 27, 5, 66]. We believe that studying link layer behavior is the way to explore the validity of the conceptual model underlying these designs. Understanding when these assumptions will and will not hold can help us understand how a protocol will work in different environments.

The rest of the paper presents observations from an empirical study of 802.15.4 links and discusses the implications of these observations to the conceptual model presented above. Section 5 shows that the observed reception ratio of a link differs for traffic with different inter-packet intervals. This shows that the *stability* assumption in the conceptual model is not valid. Section 6 looks at the temporal correlation of delivery and shows that losses and successes are highly correlated for small inter-packet intervals (of about 10ms). It shows that links are bursty at the timescale of about hundreds of milliseconds. This further shows that the *stability* assumption in the conceptual model is not valid. Sections 7 and 8 explore the underlying causes of the temporal behavior of reception and show the wireless channel – signal and noise power – variation as the possible cause. Section 8 shows that external high power noise sources can affect packet delivery over multiple receivers and can thus cause spatially correlated losses across nodes. It shows that the *spatial* assumption in the conceptual model is not valid. Section 9 looks at the asymmetries in the reception ratios of forward and backward links between pairs of

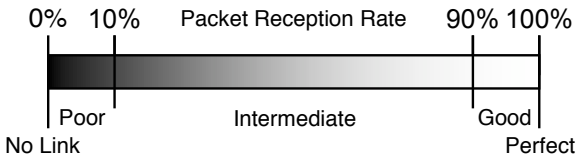

Figure 2: Link Definition. A link is dead if it has a packet reception ratio (PRR) of 0%. A link is poor if the PRR is less than 10%, is intermediate if the PRR is between 10% and 90%, is good if the PRR is between 90% and 100%, and is perfect if the PRR is 100%.

nodes. It shows that there are many short term asymmetries but very few prevail over long timescales. Finally, Section 10 shows that link layer acknowledgement and packet reception ratios are different showing that the *ack* assumption in the conceptual model is not valid.

5. DISTRIBUTION OF DELIVERY PROBABILITIES

Early mote platforms and 802.11 networks observe a wide range of packet reception ratios [21, 73, 7, 59, 3]. To see whether these observations hold in 802.15.4, we start our study with the packet reception ratio distribution on the University, Mirage and Lake testbeds. Prior studies have shown that reception ratio of 802.15.4 links can vary significantly over time [41]. Correspondingly, we also study how time scales affect reception ratio by varying the time interval between packet transmissions over four orders of magnitude, from 10 milliseconds to 15 seconds. Varying this inter-packet interval (IPI) allows us to see behavior at different time scales. In the rest of this paper, we use the terms "interval" and "IPI" interchangably.

Our results show that the reception ratio of a link depends on how often we send packets to measure them: 5% of all links in the Mirage testbed are intermediate when the IPI is 10ms while that percentage increases to 19% when the IPI increases to 1s.

5.1 Experimental Methodology

Nodes send packets at a fixed interval. The interval varies from one experiment to another. When the inter-packet interval is small, nodes, one at a time, transmit all packets as bursts. For large IPIs, the nodes take turns to send every packet such that consecutive packets from a single transmitter is interleaved with packets from other transmitters. This reduces the total experiment time. In the wired testbeds, a PC controls all transmissions. In the outdoor testbeds, nodes self-schedule and log all data to flash. Logging into flash prevents outdoor nodes from having an inter-packet interval below 50ms. We repeat the experiments for different number of packets: 100, 200 and 2000, and on different channels: 16 and 26. In all the experiments only one node transmits at a time, so there are no collisions between experimental transmitters.

Figure 2 shows how we desribe links based on their reception ratio. A node pair which receives no packets has no link. A pair which recives $\leq 10\%$ of the packets is a poor link. A pair which receives $10-90\%$ is an intermediate link. A pair which receives $> 90\%$ is a good link, and a pair which receives 100% of the packets transmitted is a perfect link.

5.2 Results

Figure 3 shows the cumulative distribution function of reception ratio of every node pair in each testbed. These plots show many subtle complexities.

Figure 3(a) shows the distribution of link qualities in the Mirage, University, and Lake testbeds on channel 26 (2.495 GHz). About 55% of all node pairs in the Mirage and University testbeds can communicate, while 90% of the pairs in the Lake testbed can communicate. This difference is due to the different setups; the Lake testbed is much denser than Mirage, yet uses the same transmit power. Of these communicating links, 19% of the Mirage links are intermediate, 5% of the University links are intermediate, and 14% of the Lake links are intermediate.

The first important observation is that 802.15.4 has far fewer intermediate links than what studies of other link layer have reported: a study of an early mote platform reported 50% [73] intermediate links and the study of an outdoor 802.11 mesh network reported 58% [3].

Figures 3(b) and 3(c) show how the interval between packet transmissions affects the observed reception ratio. The percentage of intermediate links for the University testbed increases from 5% at 10ms to 19% at 1s, and percentage of intermediate links in Mirage testbed increases from 19% at 10ms to 23% at 15s. As the inter-packet interval increases, a network observes more intermediate links. As we calculate each link reception ratio over 200 packets, the packet interval determines the time period of packet reception measurement. For example, measuring the reception ratio for a packet interval of 10ms takes 2s. In contrast, measuring for a packet interval of 15s takes 50 minutes.

Figure 3(d) shows how channel selection changes the reception ratio distribution in Mirage. Channel 16 has far fewer perfect links than channel 26: 60% in channel 26 and only 12% in channel 16. Correspondingly, 35% of the communicating channel 16 links are intermediate, compared to 17% of channel 26 links. Figure 4 shows how channel affects the observed reception ratio of a link. Figure 4(a) for transmitter shows that the link to receiver 12 is perfect on channels 15, 21, 24 and 26, while it is dead or intermediate on the others. Similarly, Figure 4(b) for transmitter 17 shows that the link to receiver 5 is perfect on channels 16 and 23, while it is dead or intermediate on others.

5.3 Observations

The high level observations from this section are:

1. 802.15.4 has fewer intermediate links than link layers measured in prior studies,

2. as the measurement duration increases, the number of intermediate links increases and the number of perfect links decreases, and

3. channel selection can affect the observed reception ratio.

5.4 Implications

The first observation does not contradict the conceptual model. Nevertheless, it means that protocols need not operate under the assumption that most links are intermediate or poor. Instead, deployments with reasonable density can

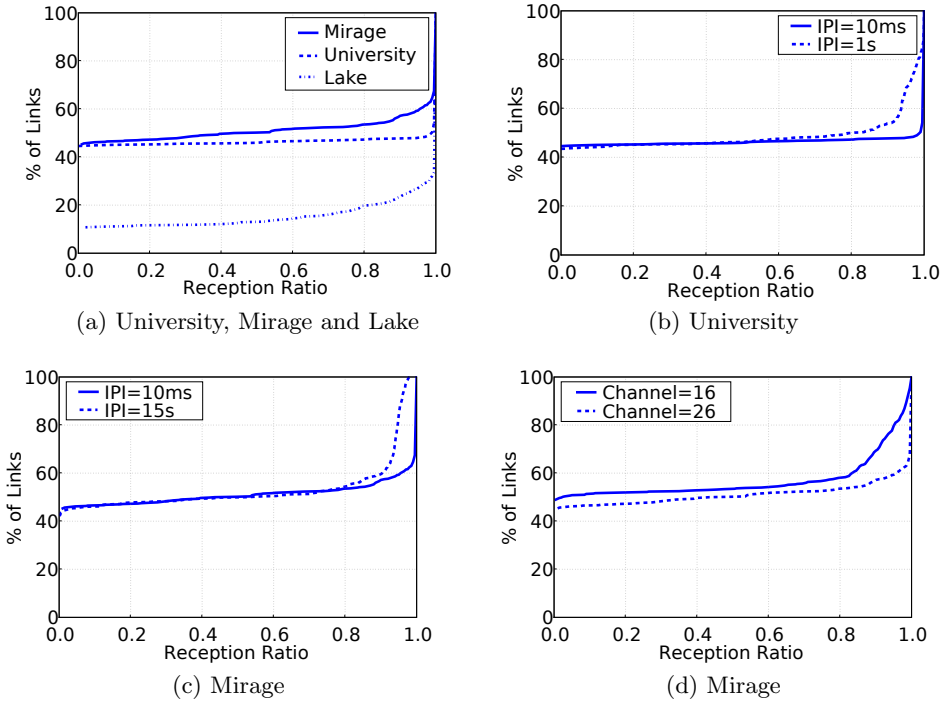

(a) University, Mirage and Lake (b) University

(c) Mirage (d) Mirage

Figure 3: (a-c) show cummulative distribution functions (CDFs) of link qualities in the three testbeds for channel 26 at different IPIs. The percentage of intermediate links is small compared to good and bad links. It increases as the inter-packet interval increases. (d) shows CDFs of link qualities in Mirage on channels 16 and 26. Channel affects the percentage of perfect links and the percentage of intermediate links.

take advantage of links which have low loss rates in the absence of interfering transmissions.

The second observation contradicts the *stability* assumption in the conceptual model. The observed reception ratio can be different for traffic with different inter-packet intervals. The percentage of intermediate links increases by 15% as the IPI increases from 10ms to 1s as seen in Figure 3(b) for the University testbed. A protocol ignoring the changes of link quality at this time scale may be trading-off delivery ratio and goodput for network tree stability. The IPIs used in this study are very similar to what Srcr uses for data and control beacons [6]. Our result for the University testbed shows that the estimates that Srcr gets from control beacons may be invalid for the more frequent data packets.

The third observation contradicts the *channel* assumption in the conceptual model. Protocol designers usually test their protocols on a single channel. A protocol may perform well on a given channel but not on some other channel. Therefore, generalizing such results to all the channels may not be applicable. We defer an exploration of this phenomenon to Section 8.

5.5 SNR Hypothesis

The second observation contradicts a widely adopted practise of computing and using link estimates to forward packets. Understanding what causes links to observe different reception ratios for traffic at different rates can guide us to design better future protocols.

Because the CC2420 handles almost all stages of packet reception in hardware, software timing issues or errors are an insignificant cause of packet loses in the above experiments. Therefore, packet delivery is a result of the underlying RF behavior, in particular, the signal-to-noise ratio (SNR). The signal-to-noise ratio describes how strong an intended RF signal is in comparison to additive white gaussian noise. The fact that hardware generates this class of noise is what has enabled the RF community to apply rigorous analysis tools to understading RF behavior.

For any given RF modulation scheme, a given signal-to-noise ratio has an expected bit error rate. If one assumes fixed length packets, this bit error rate can be extrapolated to a packet reception ratio. SNR is typically measured in decibels: an SNR of 10dB means the signal is ten times stronger than the noise, while 20dB means it is one hundred times stronger. Hardware noise power varies with temperature, and so is typically quite stable over time periods of seconds or even minutes. Signal, however, is a product of many environmental surrounding transcievers, and so can vary much more quickly.

Barring software errors, we hypothesize that changes in the packet reception ratio must be due to changes in the signal-to-noise ratio.[2] This means that the increase of intermediate links as the IPI increases is due to channel variations. This means that intermediate links may be links which move between good and poor channel conditions. We explore this hypothesis in the next section. The fact that the University testbed observed fewer perfect links on channel 11 than on channel 26 (Figure 3(d)) suggests that there is

[2]To the RF theory community, this hypothesis is hardly such: it is mathematical truth. However, experimental studies, such as Roofnet [3] have suggested otherwise.

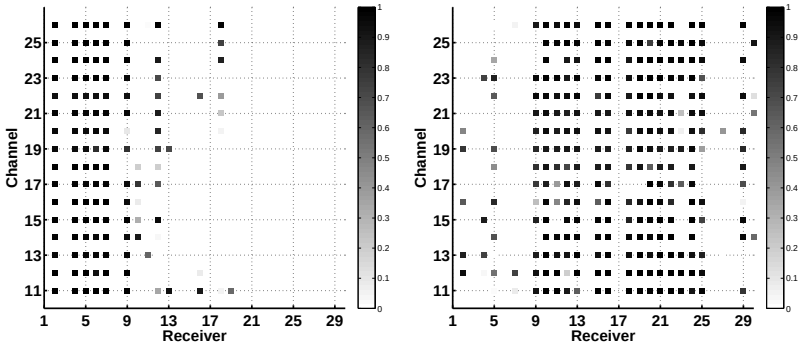

(a) Transmitter 3 (b) Transmitter 17

Figure 4: Reception ratios for different receivers over different channels for packets from the same transmitter in Mirage. The grayscale represents the observed reception ratio. Observed reception ratio is different on different channels. Some links are perfect on some channels and dead on others.

an additional phenomenon in play, which Section 8 explores.

Knowing if intermediate links see correlated delivery over time has implications to protocol design: if delivery is correlated then protocols can send bursts of packets when a link becomes available, amortizing the cost of finding when a link is available. The following section explores if links are temporally correlated. Sections 7 and 8 examine the above hypothesis that temporal correlation is due to signal and/or noise variations.

6. TEMPORAL CORRELATION OF RECEPTION

Section 5 showed that over short time periods most links are good or poor and the percentage of intermediate links increases over longer periods. This section explores this observation more deeply. We find that most of the intermediate links see highly correlated delivery when the IPI is small (about 10ms), showing intermediate links are caused by shifts between good and poor link quality.

6.1 The Conditional Packet Delivery Function (CPDF)

We have found that an effective way to see if reception is temporally correlated is to look at the probability of a packet succeeding given the fate of the previous k consecutive packets. The conditional packet delivery function (CPDF) is a function which plots this conditional probability for past k successes ($k > 0$) and failures ($k < 0$). A flat CPDF means packet events are independent: the fate of a packet in the past does not influence the fate of the current packet. For example, Figure 6 shows the CPDF of a synthetic link with a reception ratio of 90% with independent reception: the CPDF is almost 0.9 for all possible values of k. The horizontal line shows the PRR calculated over the entire trace. The CPDF shows values only between -2 and 75 because they were the maximum run lengths of failures and successes on the synthetic link.

6.2 Alternative Metrics

We use CPDFs rather than autocorrelation of the reception trace because autocorrelation does not distinguish failure-failure correlation from success-success correlation.

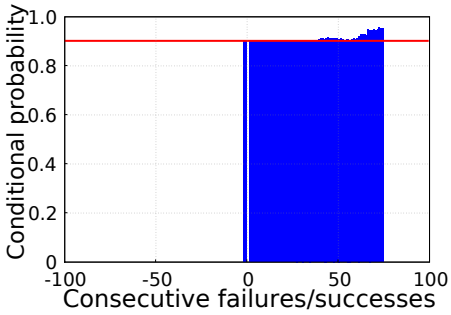

Figure 6: CPDF of an hypothetical link with PRR=90% with independent losses. Positive x-axis corresponds to consecutive successes and negative x-axis corresponds to consecutive failures. The y-axis gives the success probability of a packet given the fates of the immediate k packets before it. This CPDF is always close to 0.9 showing that the fate history of previous packets does not change the fate of the current packet – an independent link. The horizontal line shows the packet reception ratio over the entire trace of the link.

As these two correlations can be different, CPDFs are more suitable than autocorrelation.

Allan deviation calculation as proposed by Aguayo et al. is another alternative to CPDFs [3]. They propose to calculate Allan deviations of packet reception ratios computed over non-overlapping intervals for different intervals. They showed that the Allan deviation will be high for interval lengths near the characteristic burst length and will be small at smaller and longer intervals. We did not observe this pattern in the Allan deviation plots for the links from our testbeds. This led us to explore CPDF.

6.3 Experimental Methodology

To compute CPDFs, we ran an experiment on the Mirage testbed on channel 26. Thirty nodes chosen from the entire testbed, took turns, one at a time, to broadcast 100,000 packets at an inter-packet interval of 10ms. A server controlled all transmissions so there were no collisions. The

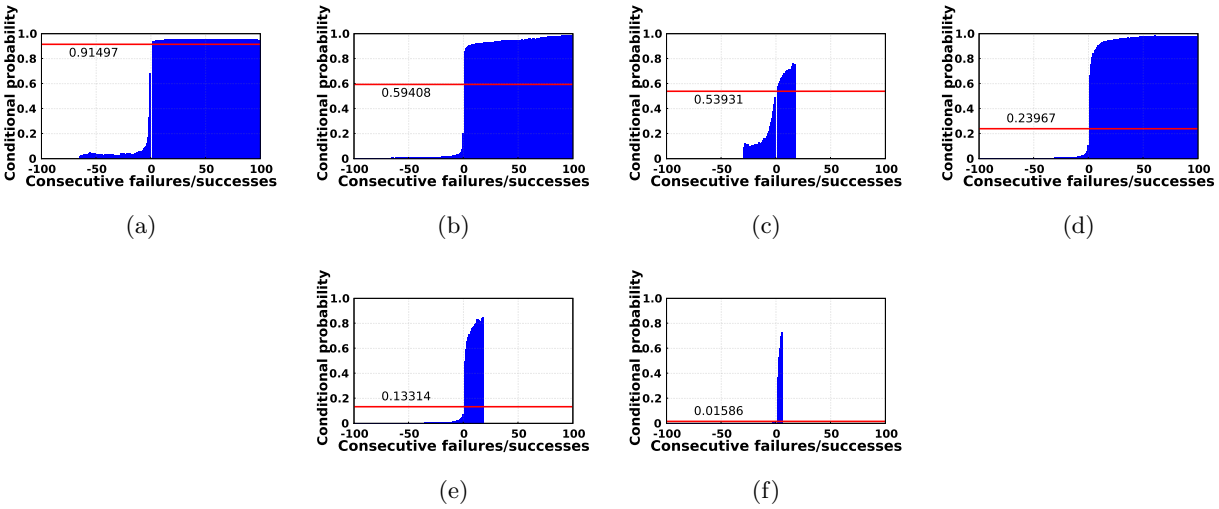

Figure 5: Conditional packet delivery functions (CPDFs) of links with different packet reception ratios (marked by a horizontal line). All of these CPDFs are from intermediate links observed in the Mirage testbed.

receivers noted down the sequence number, RSSI and CCI of received packets and sent them to the server over the Ethernet. Sending packets at an inter-packet interval of 10ms allows us to look at the temporal correlation of 802.15.4 links at a reasonably small time scale.

6.4 Results

Figure 5 shows the CPDFs of a few intermediate links when the inter-packet interval is 10ms. The CPDF includes only those conditional probabilities for which there are at least a 100 datapoints in the data trace.[3] CPDFs are not flat: packet losses are not independent. Irrespective of the reception ratio, overall, the plots show that as more packets are received, the probability of the next packet succeeding increases and as more packets are lost the probability of the next packet succeeding decreases. The CPDFs shown in Figure 5 are representative of all the intermediate links from the experiment. Packet successes are followed by a burst of successes and losses are followed by a burst of losses: links have clustered successes and failures. 802.15.4 intermediate links have temporally correlated reception when the time between packets is as small as 10ms.

Figure 7 shows the CPDFs of two links when the IPI increases from 10ms to 500ms and 2s. Figure 7 shows the CPDFs zoomed in on first few elements on the negative and positive axes to illustrate that the conditional probabilities are approaching the average PRR of the link (horizontal line) as the IPI increases. As the IPI increases, the conditional probabilities in the CPDFs of the two links look more similar to the independent case (marked by the horizontal line): the packet events on links become less correlated as the interval between packet events increases. We observed this pattern for 97% of all the intermediate links. The remaining 3% of the intermediate links consistently had highly correlated packet delivery across all the IPIs.

[3]100 data points gives a 95% confidence interval of [p-0.1, p+0.1], where p is the conditional probability.

6.5 Observations

The high level observation from this section is that the 802.15.4 links have temporally correlated reception when the packets are close in time (about 10ms apart). This means that links go through periods of perfect and zero reception. This explains why Figure 3(c) shows about 80% of all the links are either perfect or dead. The experiment with the inter-packet interval of 10ms involved only 200 packets on every link and so the measurement time was only 2s over which, due to temporal correlation of reception, most of the links either received all or no packets. However, as the interval increases, the temporal correlation between packets decreases, leading to more intermediate links.

This burstiness is a well-known phenomena in wireless at the physical layer, but this study shows it occurs at much longer time scales that the physical layer usually considers. Furthermore, it shows that intermediate links are predominantly a result of this burstiness, rather than links having a signal-to-noise ratio which causes intermediate packet reception ratios.

6.6 Implications

The temporal correlation of reception in 802.15.4 links further contradicts the *stability* assumption in the conceptual model. Many MAC protocols like CSMA are designed assuming independent link delivery: losses over time are independent of each other. The MAC for 802.15.4 can let a node retransmit immediately after a failure [30]. If the link has correlated losses then immediate retransmissions after a failure are likely to fail. However, if the losses are independent then immediate retransmissions can be useful.

Routing metrics such as ETX assume losses over time are independent. ETX calculates the average number of transmissions needed for a packet as the inverse of reception ratios calculated over an interval (usually from control beacons). This assumes independence of losses [8]. Using a metric like required number of packets (RNP) that calculates the average number of losses before a success, does not make the independent losses assumption [8]. Cerpa et al. propose

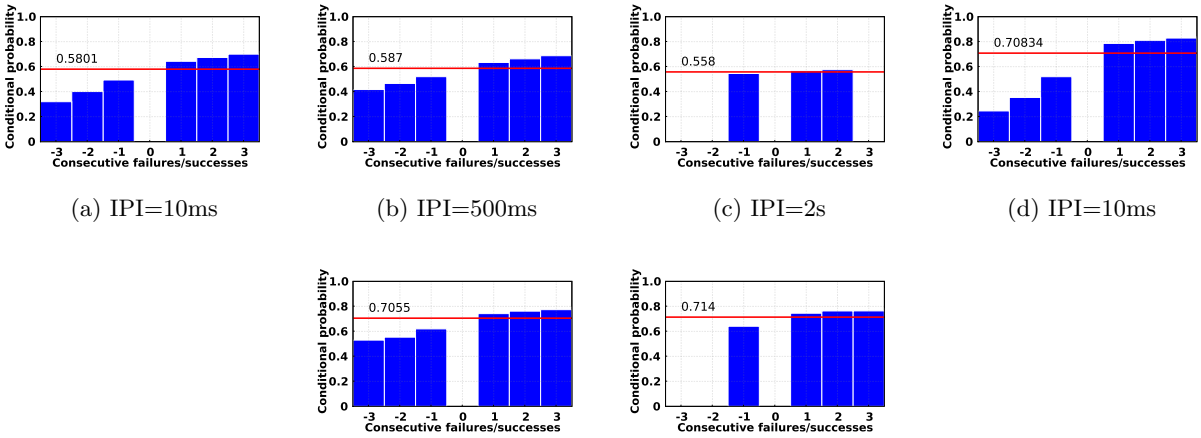

(a) IPI=10ms (b) IPI=500ms (c) IPI=2s (d) IPI=10ms

(e) IPI=500ms (f) IPI=2s

Figure 7: Conditional packet delivery functions (CPDFs) of links for different IPIs. The CPDFs look more similar to the independent case as the IPI increases: the conditional probabilities approach the average reception ratio of the link marked by the horizontal red line.

calculating RNP on control packets sent every second. Figure 3(b), however, shows that the links can change on the timescale of hundreds of milliseconds: calculating RNP over packets sent every second may yield inaccurate estimates for data packets. However, sending control packets every few milliseconds will lead to significant energy overhead. Protocols can not only use the control packets but also include the actual data packets sent over a link to estimate link quality. Four-bit wireless link estimator [20] is such a protocol.

The following section explores if our observations are specific to the testbeds we used or if they are applicable to 802.15.4 networks in general. Specifically, it explores if the temporal variations in delivery are due to channel variations, the SNR hypothesis we put forth in Section 5.

7. PHYSICAL LAYER PARAMETERS: RSSI AND CCI

Section 6 showed that links have temporally correlated delivery. In Section 5, we hypothesized that this temporal correlation is due to the channel stability over short time spans. This section explores this hypothesis by looking at the physical layer parameters that the nodes can measure. Showing that a general phenomenon such as the channel variation is the cause of our observations will extend the applicability of our observations to other 802.15.4 networks.

We find that small variations in the received signal strength, as small as 2dB, can make a good link poor and vice versa. We find that the signal to noise power ratio (SNR) is a good indicator of the reception ratio and that the slow variation of SNR is a cause of the temporal behavior of our links.

7.1 Controlled Attenuation

Borrowing Aguayo et al.'s methodology, we connected two shielded micaZ motes through a variable attenuator via shielded SMA cables [3]. At each attenuation level (1 to 64 dB), one node transmitted 100 packets with a small inter-packet interval of 50ms: logging data in to flash limited the maximum packet rate. The receiver logged RSSI, CCI

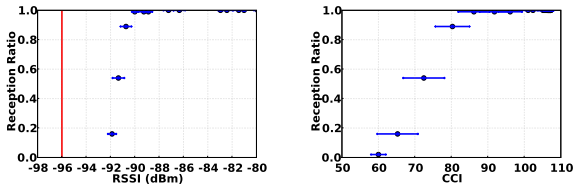

(a) RSSI vs PRR, Traffic with IPI = 50ms, Attenuator Experiment, Channel 26.

(b) CCI vs PRR, Traffic with IPI = 50ms, Attenuator Experiment, Channel 26.

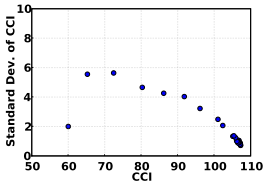

(c) Standard Deviation of CCI vs Average CCI, Traffic with IPI = 50ms, Attenuator Experiment, Channel 26.

Figure 8: PRR versus RSSI and CCI plots for a chosen pair of nodes connected through a variable attenuator. Each data point is for an attenuation level. The error bars show one standard deviation of the measured values. Intermediate PRRs are within 1.5dB range. CCI is a statistical value. CCI varies least on perfect links.

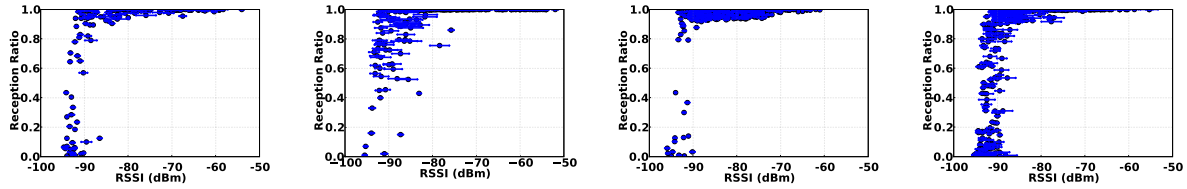

(a) Mirage, IPI = 50ms, Channel 26. (b) Lake, IPI = 50ms, Channel 26. (c) University, IPI = 50ms, Channel 26. (d) Mirage, IPI = 15s, Channel 26.

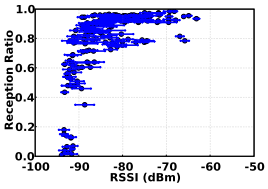

(e) University, IPI = 10s, Channel 26.

Figure 9: PRR versus RSSI in Mirage, university and lake testbeds. Each data point is for a directional node pair. The error bars show one standard deviation of the measured RSSI values.

and sequence number of every successfully received packet. Both the nodes sampled the RSSI register of CC2420 when there were no 802.15.4 traffic, to measure background (hardware/additive white Gaussian) noise. We calculated a node's noise floor as the mode of these samples. Throughout this paper, we refer to the background thermal noise in nodes as noise and the transmissions from external systems other than 802.15.4 such as 802.11 as external noise.

Figure 8(a) shows a plot of RSSI vs. PRR from different attenuation levels. The plot shows a strong correlation between RSSI and the PRR just as wireless communication theory asserts. The error bars show the one standard deviation from the measured average value. The tiny bars show that the RSSI at each attenuation level was very stable. The small variations observed in RSSI, close to the noise floor, is due to small variations in the noise floor: as RSSI is the sum of signal and the local receiver noise, small variations in the noise floor causes small changes in the RSSI. The solid vertical line at -96 dBm is the noise floor at the receiver. The receiver does not receive any packets below a signal to noise ratio of 4dB. The intermediate PRRs are within a 1.5 dB range of -92 to -90.5 dBm. For links operating close to the noise floor, this indicates that a small variation in RSSI of only 1.5dB can cause a good link to a poor link link and vice versa.

Figure 8(b) shows that CCI of intermediate links has high variation even when signal to noise ratio is stable. This supports the observation in our earlier work that CCI is a probabilistic quantity and so a single CCI value should not be used as a link quality indicator for intermediate links [64]. Figure 8(c) plots the standard deviation across different average CCIs. For links with very high CCI (close to 110), the variation is minimum (standard deviation of 0.5), suggesting that a single value might be sufficient to identify perfect and good links. As the signal to noise ratio has strong correlation with the reception ratio, we focus the rest of our study on it.

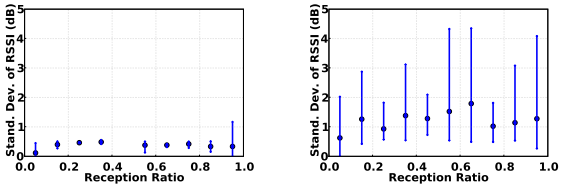

(a) Mirage, IPI = 10ms, Channel 26. (b) Mirage, IPI = 15s, Channel 26.

Figure 10: Plot of standard deviation of RSSI for different reception ratios for Mirage testbed for two inter-packet intervals:10ms and 15s. The plot shows mean (o), and max (+) and min (+) calculated for every bin of 10% of reception ratio. RSSI is stable for the short term traffic and varies for the longer term traffic.

7.2 Received Signal Strength Indicator (RSSI)

Now, we examine if real channel conditions have the same correlation between RSSI and reception ratio that controlled attenuation does. The data used in this section is from the same experiments as in Section 5.

7.2.1 Results

Figure 9 plots RSSI against the reception ratio for two intervals: 10ms and 15s. The three orders of magnitude difference in interval causes a similar difference in the time period over which the measurement was taken. At a 10 millisecond interval, each link is measured over 2 seconds, while at a 15 second interval each link is measured over 50 minutes.

The plots show that RSSI is generally stable over a short time span (2s) in all testbeds [64]. This is consistent with the notion of coherence time of a channel in wireless communications [52], which is the time over which the channel state remains highly correlated with itself. The RSSI over longer

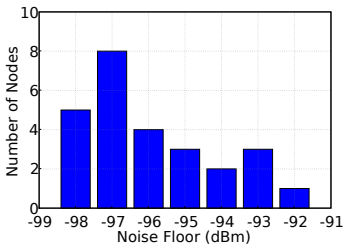

Figure 11: Distribution of estimated noise floor across 26 motes on Mirage testbed.

time spans (50mins) shows greater variance (Figures 9(d) and 9(e)). Figure 10 is a plot for the Mirage testbed showing the minimum, the mean and the maximum standard deviations of RSSI for different reception ratios for 10ms and 15s intervals. The plots show that over short term, the average standard deviation is lower than 1 dB across all PRRs for all links and the maximum is 1.5 dB. However, over longer term, the average standard deviation is more than 1 dB for all links with reception ratio greater than 10% and the maximum is as high as 4.2 dB in some links. RSSI is stable over short time spans and varies as the inter-packet interval increases.

7.3 Noise

The plots in Figure 9 show that links with an average RSSI above -87dBm are good links [64]. Below this threshold, however, there is a "grey region" of many different reception ratios, with no clear correlation to RSSI. The attenuator experiment showed that RSSI forms a very precise and smooth curve. Figure 11 shows the distribution of noise floors of 26 nodes in the Mirage testbed. The noise floor varies across nodes between -98 to -92dBm. For a given RSSI across all nodes, this difference in noise floor at different nodes means different signal to noise ratios and thus different reception ratios. Unlike the RSSI vs. reception plot for the attenuator experiment in Figure 8(a), Figure 9 plots not just one but several links. The difference in the noise floor across nodes will therefore skew the correlation between RSSI and reception ratio.

Figure 12(b) shows the plot of SNR plotted against PRR and shows the range of intermediate links. We use this range as a crude measure of the grey region. This range was only 2 dB for the attenuator experiment. The corresponding RSSI vs. PRR plot in Figure 12(a) has a wider range for the intermediate links: the SNR plot has a range of ∼5 dB and the RSSI plot has ∼7 dB. While taking noise floor differences across nodes in to account reduces the grey region, it is still greater than 2dB. This means that there are factors other than noise floor that affect the RSSI-PRR relationship.

As with other similar studies, there is an inherent bias in our measurements: averaging RSSI only over received packets biases it upwards. For example, let RSSI at a receiver be stable at -91 dBm for 50 packets and then let it drop below the receiver's noise floor to -95 dBm for the next 50 packets. As the RSSI is below the noise floor for the later 50 packets, they are not received. Averaging over only successful packets gives an average RSSI of -91 dBm. However, the actual average is -93 dBm. Thus, different links going through different levels of channel variations will have different biases. We could not confirm this as our nodes do not give RSSI values for unsuccessful packets. Son et al. have shown that bad transmitters that distort signals can introduce RSSI asymmetries [59]. This could be another factor that skews the RSSI-PRR relationship.

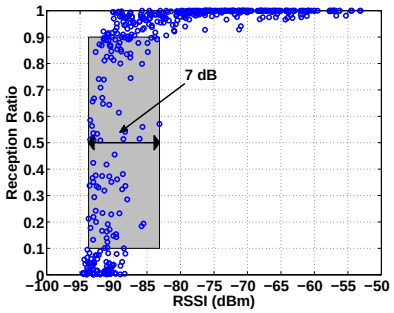

(a) RSSI vs PRR.

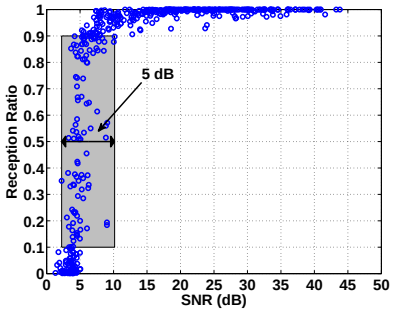

(b) SNR vs PRR.

Figure 12: Plots of average RSSI and average SNR vs. PRR. The plots show the range of intermediate qualities. The width for the SNR plot is smaller than for the RSSI plot.

Figure 13 shows a detailed look at the RSSI for a sequence of received packets for an intermediate link. While most of the received packets have an RSSI of at least -94 dBm, a few are as low as -95 dBm. If a link is near the cusp of reception sensitivity, then slight variations can cause packet losses and make the link intermediate. Figure 13 shows a dip in the reception ratio just after 7s. This dip happens after weak packets with RSSI close to -95dBm. There are no reception ratio dips when strong packets are received. This shows that the losses were not due to external noise. RSSI shifts of this kind are typical of all but the most controlled environments, either due to environmental effects [41] or simple multipath fading.

7.4 Observations

The high level observations from this section are:

(i) the signal to noise ratio has good correlation with the reception ratio,

(ii) RSSI is stable over short time spans and varies over long time spans,

(iii) a variation in RSSI as small as 1.5dB can change a good link to a bad one if the link is operating near the noise floor, and

(iv) single CCI values can identify good links but not intermediate links.

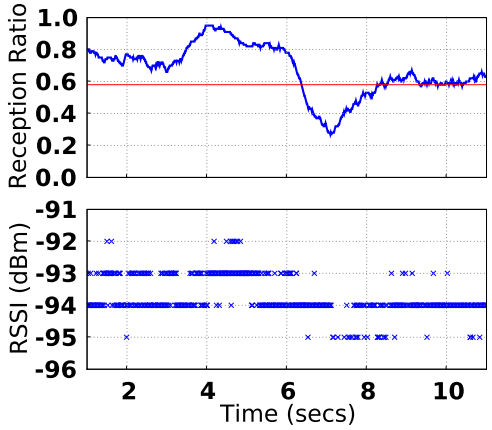

Figure 13: RSSI and PRR variations over time on a single link. The PRR over time is from a sliding window of size 100 packets (1s). Red horizontal line in the PRR plot shows the overall average PRR of the link. The RSSI oscillates with scattered reception close to -95 dBm and denser reception above. The clustered reception and losses show the burstiness of the link.

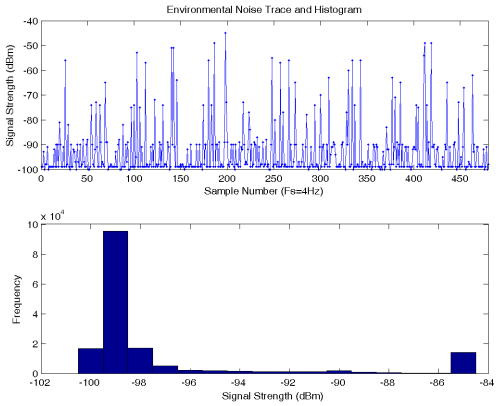

Figure 14: Sampled signal strength trace and histogram. (a) Sampled signal strength (dBm) measured over a two minute period at a single node; (b) A histogram of the sampled signal strength over 11 hours measured at the same node.

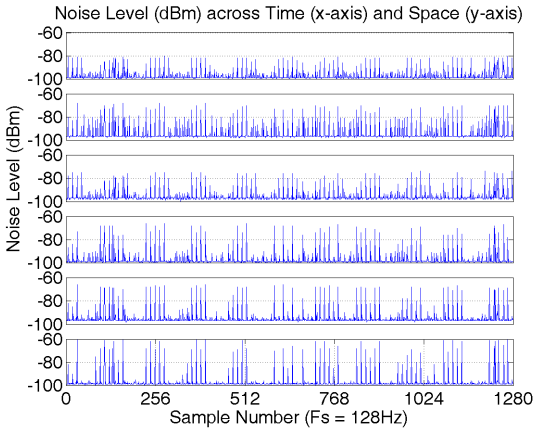

Figure 15: Sampled signal strength over 10s and across six nodes.

7.5 Implications

The first observation contradicts numerous earlier studies. For example, the Roofnet measurement study [3] argued that the signal-to-noise ratio was a poor predictor of packet reception. We defer a more detailed examination of this topic to Section 12, but believe this conclusion was due to two overlooked issues in the experimental methodology.

The second observation is consistent with the hypothesis proposed in Section 5 that over short time spans the channel is stable. The second and the third observations, together, suggest that the reception ratios measured from traffic with different intervals may be different. As the channel varies over time, based on the timing, each packet may observe a different channel state than its predecessor and so may have a different fate. This supports our observation in Section 5 that the observed reception ratio depends on the rate at which we send packets. This is, once again, showing that the *stability* assumption is not valid.

The second and the third observations also suggest that the chances of a link being an intermediate link increases as the measurement timespan increases. This supports our observation from Section 5 that the number of intermediate links increases as the inter-packet interval increases.

The fourth observation has implications to link estimators. Link estimators should not use single value CCI as the estimate for intermediate links. Protocols such as MulithopLQI, which use CCI to compute link costs, do so by scaling the value such that the protocol greatly favors only perfect links: a link with a CCI of 90 has eight times the cost of a link with a CCI of 110, even though Figure 8(b) shows links with a CCI of 90 are good to near-perfect. Correspondingly, while MulithopLQI produces stable and robust topologies, they are sub-optimal [20] and waste energy.

Figure 3(d) showed that the observed reception ratio could be different on different channels. The observations from this section do not explain this, suggesting there may be an additional factor at play. External noise sources can also introduce packet losses in a network. The following section studies effects of external noise from other wireless systems that share the spectrum with 802.15.4.

8. EXTERNAL NOISE

802.15.4 operates in the same 2.4GHz ISM band as Bluetooth and the far more powerful 802.11b. Since all of these systems co-exist in the same wireless spectrum, we focus our study on 802.11 and Bluetooth. We find that most of the 802.15.4 channels are vulnerable to interference from cohabiting 802.11 networks.

8.1 Experimental Methodology

The best way to observe noise is to sample the signal strength when there are no transmissions in our network. Nodes read raw RSSI samples from the CCC2420 radio on channel 11 at 4Hz for 11 hours on the University testbed. No nodes transmitted 802.15.4 packets during this time.

Figure 14 shows a two-minute subset of the values measured at a single node as well as a histogram of the values

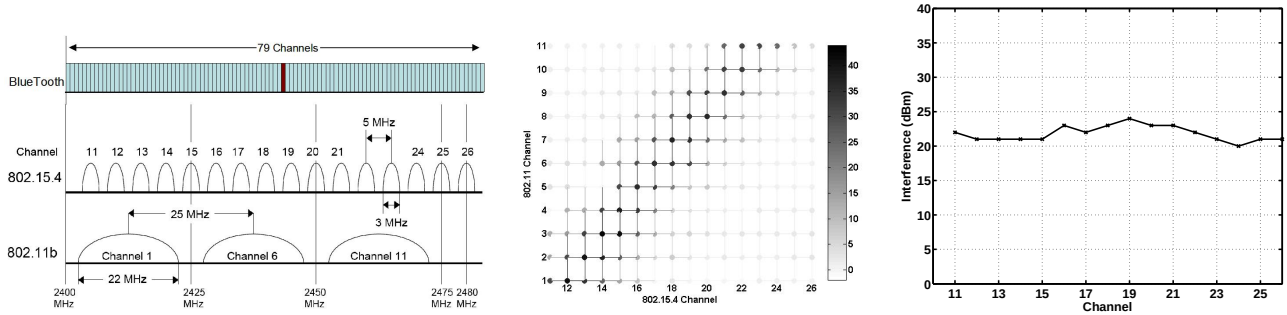

(a) 802.15.1 (Bluetooth), 802.11b and 802.15.4 spectrum usage.

(b) Increase in average noise for all channel pairs observed by an 802.15.4 receiver during 802.11b transmissions.

(c) Increase in average noise for all 802.15.4 channels for an 802.15.4 receiver during 802.15.1 transmissions.

Figure 16: Interactions between Bluetooth, 802.11b and 802.15.4. Bluetooth adds approximately 20dBm to the noise floor while 802.11 adds between 0 and 45dBm to the noise floor.

over a 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. The distribution of samples is right-tailed, with more than 8.6% of the samples having a value greater than -85 dBm. If these spikes were to be from within the node then this can be an issue for the CSMA/CA protocols as it could make the node think that the channel is busy while the spike is due to the local noise. To verify if these spikes are external we ran an experiment in which node synchronized their noise samples. A high correlation between spikes observed by the nodes would suggest external interference as the cause.

Figure 15 shows 128Hz noise samples from six synchronized nodes. The minimum calculated correlation coefficient between the traces is 0.77. This indicates that the noise spikes are highly correlated and are likely external to the nodes. We confirmed that the source of the spikes were due to co-located 802.11b network access points by shielding the near-by 802.11b access point (AP). Shielding the AP reduced the peak noise by approximately 15dB.

8.2 802.11b

Although 802.11b and 802.15.4 share the same spectrum, their channels occupy different bandwidths. Figure 16(a) illustrates how the channels nominally overlap. We note that most wireless 802.11b access points use channels 1, 6, and 11 because these three channels are mutually non-interfering. We now measure the interference 802.11 nodes may cause to 802.15.4 nodes by sampling the signal strength at the 802.15.4 nodes when the 802.11 nodes are busy.

8.2.1 Experimental Methodology

A Telos node placed between two 802.11b devices samples the CC2420 RSSI register at 4Hz. The 802.11b devices are in ad-hoc mode and transfer a large file using FTP. We measure the noise observed at the node for all combinations of 802.15.4 and 802.11b channels.

Figure 16(b) shows the *difference* in received signal strength caused by the presence of 802.11b traffic. An inopportune choice of 802.15.4 channel can result in up to 45 dBm of interference from 802.11b traffic under our test conditions. The data also indicates that only 802.15.4 channel 26 is largely immune to 802.11b interference.

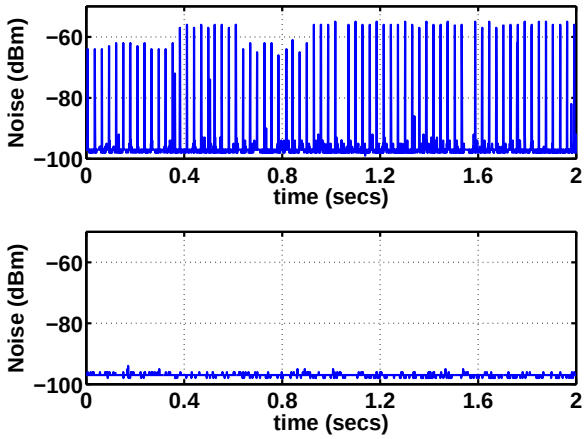

Figure 17: High frequency noise samples on channels 16 and 26 in Mirage testbed when no 802.15.4 transmissions are present. Channel 16 shows large spikes while channel 26 shows none. Channel 16 shares the spectrum with a cohabiting 802.11 network while channel 26 does not.

This data suggests that when selecting an 802.15.4 operating channel, one should avoid channels of coexisting 802.11b networks to minimize interference and loss [14, 70]. Packet losses are because 802.11b nodes usually do not defer transmission when an 802.15.4 packet transmission is in progress. This is because of the difference in transmission power between the two technologies. 802.11b transmission power is larger than that of 802.15.4 by a factor of 100. In contrast, 802.11 transmissions can prevent clear channel assessment at 802.15.4 nodes and increase latencies.

Figure 17 shows the noise samples observed by a single node on channels 16 and 26, when no 802.15.4 transmissions were present. While channel 16 shows large spikes from 802.11, channel 26 shows no such spikes. Returning to Figure 3(d) in Section 5, we observed that channel 16 had fewer perfect links than channel 26. External interference from 802.11b explain why their distributions are different.

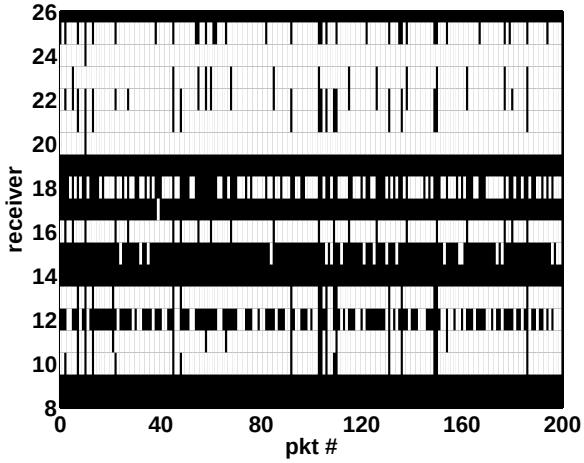

Figure 18: Reception at different nodes (y-axis) of broadcast packets from a single transmitter over time (x-axis) on a channel (802.11 Channel 17) where 802.11 was active in the Mirage testbed. Black bands correspond to losses at different nodes. This shows that losses have spatial correlation.

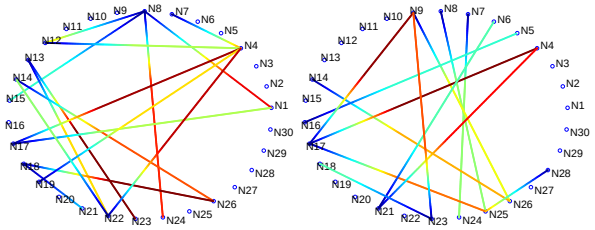

(a) Traffic with IPI = 10ms, Mirage, Channel 11, Wednesday

(b) Traffic with IPI = 10ms, Mirage, Channel 11, Saturday

Figure 19: The nodes 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.

8.3 Bluetooth

Unlike 802.11b and 802.15.4, Bluetooth is based on frequency hopping spread spectrum (FHSS) technology. Bluetooth uses 79 different 1MHz channels. Figure 16(a) shows the overlap between Bluetooth and 802.15.4 channels. We investigate the effect of coexisting 802.15.4 and Bluetooth networks using an approach similar to that used for 802.11.

8.3.1 Experimental Methodology

A Telos node placed between two Bluetooth devices samples the RSSI register of the CC2420 at 4Hz. The Bluetooth devices transfer a large file using the Bluetooth file transfer protocol. We measure the noise observed at the node for every 802.15.4 channel.

Figure 16(c) shows the difference in received signal strength between the presence and absence of Bluetooth traffic measured in different 802.15.4 channels. Interference is as high as 25 dBm. We suspect that since the interference magnitude is on the same order as our observations for the 802.11 networks, similar implications hold. However, since Bluetooth does frequency hopping every 625 microsecs the packet losses may not be as bad as we saw from 802.11 transmissions.

8.4 Spatial Correlation of Packet Loss

As the 802.11 systems are high power external noise sources for 802.15.4, they can cause packet losses at multiple nodes. Thus, such external noise sources can introduce spatial correlation of losses in a network.

Figure 18 shows reception of broadcast packets at different nodes in the Mirage testbed. It shows that losses (black bands) at different nodes happen simultaneously: losses are spatially correlated.

8.5 Observations

The high level observations from this section are:

(i) 15.4 is vulnerable to interference from 802.11 and Bluetooth systems, and

(ii) high power external noise source such as 802.11 can cause spatially correlated losses.

8.6 Implications

The first observation shows that the *channel* assumption in the conceptual model is not valid when there are cohabiting external interference sources. Many protocol studies in 802.15.4 use channel 26 for evaluation – a channel that is most immune to 802.11 interference. Motivated by our work, the TinyOS community changed its default channel from 11 to 26. Our observations from this section suggest that protocol designers should evaluate their protocols on multiple channels to fully understand their behavior and report the channel used in their evaluations.

The second observation shows that protocols should not always assume reception on different links to be independent, the *spatial* assumption in the conceptual model. In the presence of high spatial correlation of losses, Rateless Deluge and MORE may not benefit from linear coding techniques as the packets lost at different neighbors are very likely to be the same. In fact, in the presence of high spatial correlation of losses, linear encoding might result in inefficient use of the bandwidth due to transmissions of additional packets due to encoding.

The first and second observations, together, show that simulation tools should include external noise effects to better reflect reality. A recent study [38], motivated by our observations on external noise in our previous work [60], is such an attempt.

The observations from this section and from Section 7 show that the observations we have made so far in Sections 5 and 6 are due to channel variations, a phenomenon that is common to any wireless network: our observations may not be specific to our testbeds and may be more generally applicable to other 802.15.4 networks.

Given that signal to noise ratio and the external noise affect delivery, we hypothesize that variation of these two factors over time means that link asymmetries – difference between the forward and reverse link PRRs – also come and

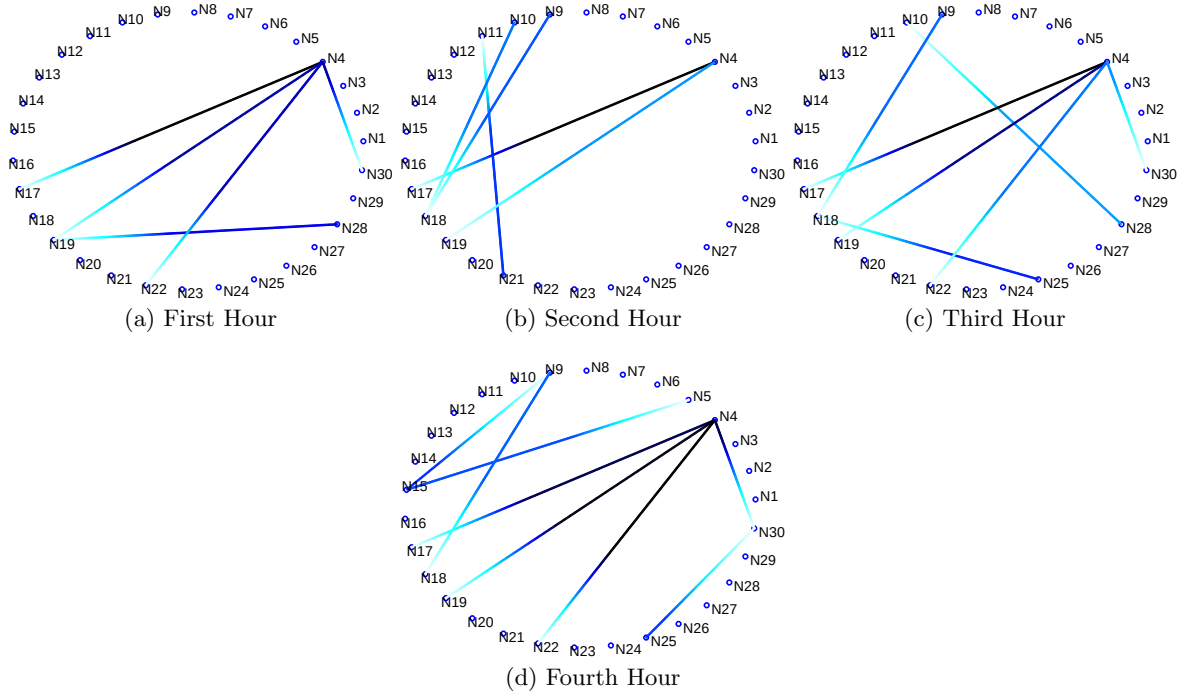

(a) First Hour (b) Second Hour (c) Third Hour

(d) Fourth Hour

Figure 20: Hour-by-hour asymmetry plots for a four hour IPI experiment on the Mirage testbed with IPI = 15s on channel 11. The visualization methodology is the same as in Figure 19. 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.

go over time. We also hypothesize that differences in the noise floor would cause some links to be consistently asymmetric. We explore these hypotheses in Section 9.

Given that RSSI is stable over short timespans and that acknowledgements are sent immediately after successful packet receptions, we hypothesize that the acknowledgement reception ratio (ARR) of a backward link is more than the packet reception ratio (PRR) of that link. We explore this hypothesis in Section 10.

9. LINK ASYMMETRIES

Link asymmetry is a well known phenomenon in wireless network studies. However, to our knowledge, temporal variation of link asymmetry has never been studied. The observations in Section 5 suggest that asymmetries might look different for traffic with different IPIs. We start with examining the commonality of asymmetric links in 802.15.4 networks and explore the causes of these asymmetries. While we find that there are many short term asymmetric links, there are fewer long term asymmetric links.

9.1 Experimental Methodology

Thirty nodes on the Mirage testbed send 200 unicast packets to each other node with an IPI of 10ms. Receivers send the RSSI, CCI and sequence number of every packet to a logging PC over a wired backchannel. This approach means the two directions of a node pair may be measured several minutes apart, while the PRR is calculated over a very brief interval of only 2s (200*10ms). We run this experiment once on a Wednesday evening and once on a Saturday morning.

We define a bidirectional link between a node m and a node n to be asymmetric if $|PRR_m - PRR_n| > 0.4$. Figure 19 shows the results from our experiment. While 802.15.4 has asymmetric links, only two of the 16 observed asymmetric links are persistent across the two experiments. Having only 2 links be persistently asymmetric in the two runs implies that link asymmetry is transient in nature.

To determine the time scale of variations in PRR asymmetry, we examine the data from an experiment with 15s IPI traffic and calculate link asymmetry over four separate one hour periods. Figure 20 shows the results. A few links such as N17→N4 are consistently asymmetric while others such as N18→N10 are not. Furthermore, the number of asymmetric links in each period is significantly fewer than what was observed in the experiments with an IPI of 10ms. These results reemphasize that there are significant differences between long-term and short-term link behavior.

Section 8 showed that a cohabiting 802.11b network can create significant interference. This lowers the SNR, leading to significant packet losses. External noise can cause asymmetry if it affects only one side of reception. Due to the nature of our experiments with low IPI where the low IPI traffic are sent as bursts, the bursts on the forward and the backward links could be several minutes apart. This is why it is possible that the external noise source causes packet losses on only one of the two links and make the links look asymmetric. However, with larger IPI traffic, as we send such traffic in a round-robin fashion, the external noise would have to consistently affect only one of the nodes over a relatively extended period of time to cause links to be asymmetric, which is unlikely. Therefore, there must be

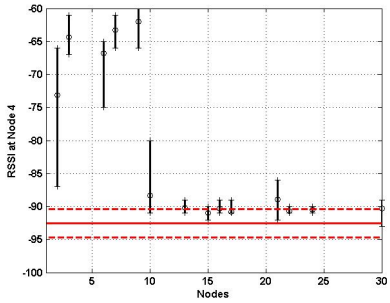

(a) Reception at node 4.

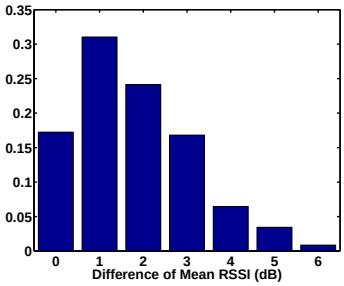

(b) RSSI asymmetry

Figure 21: (a) shows average noise at node 4 and RSSI of packets received from all nodes for traffic with IPI = 15s 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). (b) shows distribution of RSSI asymmetries on Mirage with traffic with IPI = 15 secs. 50% of all communicating node pairs have pairwise RSSI differences of 2dBm or below.

some other factor causing the long term asymmetries in Figure 20.

Figure 21(a) shows node 4's view of incoming traffic. Node 4 receives no packets below its noise floor (-93dBm), and very few below one standard deviation above that (-90dBm), consistent with the notion of an SNR threshold (Section 7). Its noise floor is also one of the highest in the network.

Looking at Figure 20, node 4 had four asymmetric links appearing in at least 3 of the 4 one hour periods with nodes 17, 19, 22, and 30. In each of these asymmetric links, node 4 was the bad receiver. Figure 21(a) shows that all of them are on the edge of receivable RSSI. In particular, node 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; without significant RSSI asymmetry in its favor, node 4 is unlikely to receive any packets.

Another factor that can contribute to PRR asymmetry is RSSI asymmetry. Figure 21(b) shows a distribution of the RSSI asymmetries in the low IPI experiment. The largest asymmetry is 6 dB. If noise floor differences and RSSI asymmetries are correlated, this would mean that there are some

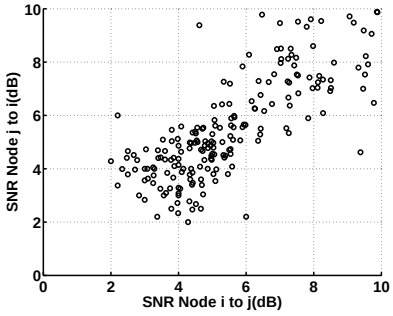

Figure 22: SNR asymmetry for traffic with IPI = 15s on Mirage.

miscalibration issues in the analog to digital converter (ADC). If noise floor differences were due to ADC miscalibration then this would invalidate our theories on the grey region in the RSSI versus PRR plot and on the causes of PRR asymmetries. Figure 22 shows the signal to noise ratio plotted for one link against that of the opposite link. If the RSSI and noise floor were correlated we would have seen all the points to be on a straight line. However, Figure 22 shows no such pattern suggesting that the asymmetries in RSSI and noise floor are not due to ADC miscalibration issues. Consequently, the cause of RSSI asymmetry still remains unanswered. Son et al. [59] present one theory, suggesting that RSSI asymmetries are due to oscillator miscalibration issues.

9.2 Observations

The high level observations from this section are:

(i) while there are many transient asymmetric links, very few links are asymmetric over long time periods of hours,

(ii) these long term asymmetric links are due to noise floor differences and RSSI asymmetries.

9.3 Implications

The first and the second observations together means that nodes may use noise floor and RSSI asymmetry information to choose neighbors and prune their neighbor tables to do efficient routing. Using routing protocols that assume bidirectionality in the presence of asymmetries might perform poorly.

The following section explores the hypothesis we put forth in Section 8 that the acknowledgement reception ratio (ARR) is more than the packet reception ratio (PRR) of a link. This is directly relevant to the *ack* assumption in the conceptual model.

10. ACKNOWLEDGEMENT PACKETS

Many routing protocols use control packets from forward and backward links between a pair of nodes to estimate the expected number of transmissions per packet (ETX). They implicitly assume that the acknowledgement and the packet reception ratios do not differ, as ETX between nodes A and B is usually calculated as $\frac{1}{PRR_{AB} \cdot PRR_{BA}}$. This section explores if this is a valid assumption. We find that the acknowledgement reception ratio (ARR) is, in general, higher than the packet reception ratio (PRR): ETX calculated from

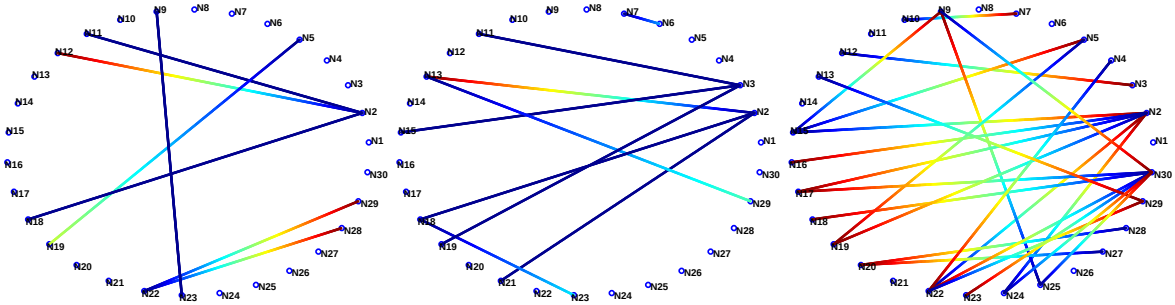

(a) Mirage, Channel 11, IPI = 10ms (b) Mirage, Channel 26, IPI = 10ms (c) Mirage, Channel 11, IPI = 14 secs

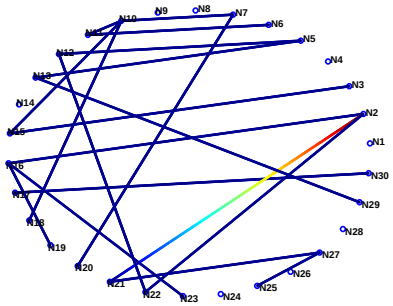

(d) Mirage, Channel 26, IPI = 14 secs

Figure 23: ETX asymmetries for low and high IPI traffic. The nodes are in a circle solely for visualization purposes. Nodes close on the circle were physically close. Asymmetry is a colored line, where the red end of the line is the node that had a higher ETX. A larger gradient indicates higher asymmetry.

packets usually gives conservative estimates of the number of transmissions needed per packet.

10.1 ARR vs PRR

If the acknowledgment reception ratio (ARR) can differ significantly from the reverse PRR, then it is possible that the two directions of a link have different ETX values, as the ETX from A to B (ETX_{AB}) is $\frac{1}{PRR_{AB} \cdot ARR_{BA}}$. There are two reasons why ARR may differ from PRR. First, 802.15.4 acknowledgment packets are very small, thus likely to be corrupted. Second, CSMA causes a data packet transmission to suppress other nodes around it. As acknowledgments are sent shortly (tens of microseconds) after the data packet, the channel conditions around a transmitter are different than those at an arbitrary receiver.

For the purpose of this study, a link has an ETX asymmetry if the ETX for the two directions differs by 0.1 and at least one direction has an ETX below 3. The second condition is based on the observation that protocols typically minimize ETX. Figure 23 plots ETX asymmetries for both low and high IPI traffic on channels 11 and 26. Low IPI traffic on channel 11 observed 7 links with an ETX asymmetry, some of which were very asymmetric (N22-N28, N22-N29) while on channel 26 there were 9 asymmetric links, only one of which was very asymmetric (N2-N13). High IPI traffic observed many more asymmetries. On channel 11, many of these asymmetries were severe, while on channel 26 they weren't.

Figure 23 shows that significant ETX asymmetries can exist, they are more pronounced over high IPI traffic than low IPI traffic, and channel choice affects the severity. As ETX asymmetries exist, ARR and PRR must differ. Figure 24 shows the relationship between PRR and ARR. As low IPI traffic observes predominantly bimodal links, its values are clustered at high reception rates. In contrast, high IPI traffic has more intermediate links. In both cases, however, the ARR is almost always greater than the PRR. Using PRR instead of ARR (as is commonly done in current protocols) overestimates ETX. Thus, a protocol using PRR instead of ARR can overlook usable links and end up with fewer choices.

10.2 Observations

The high level observation from this section is that the acknowledgement and packet reception ratios are not equal.

10.3 Implications

The observation this section presents shows that the *ack* assumption is not valid. Link-level asymmetries preclude broadcast-based route selection techniques, such as those used in AODV [49]. Similarly, ETX asymmetries mean that the two directions of a route may differ. Just as with link quality variations, ETX asymmetries increase with time duration, and so routes require periodic probing or refreshing. As acknowledgments are imperfect and energy conservation generally calls for link-level retransmissions to improve reliability, nodes require duplicate suppression mechanisms. As the acknowledgement reception ratio is usually higher than the packet reception ratio for a link, ETX based approaches over estimate cost. This may lead a protocol to ignore some usable links and thus perform sub-optimally.

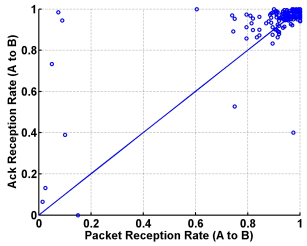

(a) IPI = 10ms

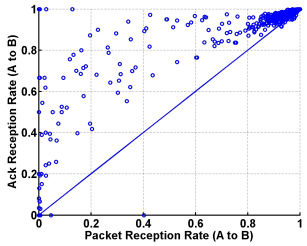

(b) IPI = 14 secs

Figure 24: PRR for A→B link vs ARR for A→B. Low IPI traffic shows bursty reception rates, while high IPI traffic shows more intermediate links. In almost all cases, ARR is higher than PRR. ARR and PRR are close at low loss rates, leading to few ETX asymmetries. Similar plots observed for channel 26 and are not shown for brevity.

11. SYNTHESIS

Section 4 presented a conceptual model of a wireless network which underlies many wirless protocol designs, noting four assumptions it makes on network behavior. Sections 5–10 examined the actual behavior of 802.15.4 networks in detail, finding that in many cases these assumptions do not hold. The roots of each of these discrepancies comes from unpredictable and uncontrolled phenomena at the physical layer: hardware assymetries, temporal variations in signal attenuation, and external interference.

1. The stability assumption does not hold due to RSSI variations at the edge of reception sensitivity.
2. The channel assumption does not hold due to external interference on some channels and RSSI variations across channels.
3. The spatial assumption does not hold due to external interference.
4. The acknowledgement assumption does not hold due to how and when acknowledgement packets are generated.

These discrepancies, however, are not fatal. Rateless Deluge, for example, still works if there are spatially correlated losses. MultihopLQI works reasonably well even though it uses only physical layer information and low-rate beacons. Instead, they lead to inefficiencies in protocol design. For example, a Deluge that estimates whether losses are independent or correlated can adapt between a rateless model, where the overhead of network coding is much less than the overhead of negative acknowledgements, or an unencoded model, where a few NACKs are sufficient for sending all missing data.

In the next section, we compare these results with existing work in wireless measurement, and in Section 13 conclude.

12. RELATED WORK

A sizeable library of literature exists that has explored low power wireless communication technologies. The results of such studies were instrumental in pointing out gaps and uncertainties in our knowledge of the space, and consequently served as important considerations when designing our own experiments. In this section, we review these prior works, pointing out the key findings and the impact this had on our work.

12.1 Prior Work with Early Motes

Experiments with early mote platforms demonstrated the complex dynamics of low-power wireless networks. The resulting observations have guided the design and implementation of numerous protocols and system stacks. However, the failure of deployments of these networks to behave as anticipated indicate that the dynamics of these systems are still a mystery to developers, and as such warrant another look. We overview previous such 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. We discuss many of these factors as part of our study.

12.1.1 Deployment Experiences

Szewczyk et al. [66] 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 [71]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. [67] 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 evaluated on small-scale testbeds 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, the underlying causes of these challenges often remain a mystery. If these root causes are left unexplored, protocols and systems will be designed with *reactive* rather than *proactive* mechanisms, resulting in a loss of efficiency. At best the resulting networks will be tuned to work well in a single deployment environment, that in which they were developed. Our work

attempts to dig into these root causes and examines the implications of our findings on sensornet system development.

12.1.2 Packet Delivery

Ganesan et al [21] analyzed different protocol layers for Rene motes, an early-generation sensor node, 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. Cerpa et al. [7] validated this theory by 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 [71] 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 [73] 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 focused on early mote platforms (e.g., rene, mica, and mica2). We explored if 802.15.4 experiences such asymmetries in detail in Section 9. We took a step forward in looking at the temporal variations of asymmetry, which was not investigated in any prior studies. We also looked at the causes of the temporal nature of asymmetries and found that noise floor differences and RSSI asymmetries can cause long term asymmetries, while short term asymmetries are due to SNR variations in the channel.

Ganesan et al. [21] showed that packet collisions, hidden terminals, link asymmetries, and the broadcast storm problem [48] 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 when the stronger packet starts later by constantly looking for a start symbol [69]. Son et al. [59] 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. Sections 7 and 8 showed that signal to noise ratio (SNR) determines reception and that it can explain all of our observations such as link burstiness and asymmetries.

Cerpa et al. showed that PRR rates can change significantly over time, so that long-term PRR calculation can lead to very inaccurate results [8], suggesting instead that a short term measure of RNP – "required number of packets (RNP)" – was preferable to a long-term PRR. Sections 5 and 6 explored temporal variations of observed packet reception ratios and showed that links can be bursty. This observation supports Cerpa et al.'s notion of RNP because the average PRR does not capture the link burstiness. However, Cerpa et al. calculate RNP based on infrequent packets (sent every second). Instead, Section 6 showed that links change in the order of hundreds of milliseconds and so RNP should be calculated over timescales of that order.

12.1.3 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 [71]. As a consequence, several more recent protocols, such as TinyOS's Drain and MultiHopLQI [46], as well as Moteiv Corporation's Boomerang [44], simply use single samples of the chip correlation indicator (CCI) of the CC2420 radio as a measure of link quality. Section 7 showed that CCI had variability even under the controlled attenuator experiment and so might not be a good indicator of link quality for intermediate links [60].

12.2 Prior Work with 802.15.4

The experimental studies on 802.15.4 primarily focus on 802.15.4's coexistence with 802.11 systems and on MAC layer performance.

Howitt et al. provided an analytic framework for analyzing the impact of 802.15.4 on 802.11 [26]. They conclude that 802.15.4 has minimal or no effect on 802.11 systems unless an 802.11 node is near a cluster of 802.15.4 nodes with high aggregate activity level.

Shin et al. explored the effects of 802.11 on 802.15.4 nodes [57]. They conclude that if the carrier frequencies of 802.11 and 802.15.4 are separated by at least 7 MHz then 802.11 has negligible effect on 802.15.4.

Petrova et al. showed that the effects of 802.11 and 802.15.4 depend on the center frequency separation between the two systems [50]. They also show that 802.15.4 may affect 802.11b but not 802.11g systems. A similar study [58] shows that 802.15.4 has minimal impact on 802.11's throughput but Bluetooth has significant impact while 802.11 can significantly degrade 802.15.4's throughput.

These studies already suggest that 802.11 systems can affect 802.15.4 links. Section 8 studied the effects of 802.11 on all the channels of 802.15.4. We showed that 802.15.4's channel 26 is the most immune to 802.11.

Zheng et al. developed 802.15.4 MAC models in NS-2 simulator and study its performance [74]. Lee et al. presented many observations on the performance of the 802.15.4 MAC layer [39]. Our study looks at understanding the more fundamental link layer behavior which, we believe, gives us insights in to the behavior at higher layers.

Jamieson et al. used a scheme that uses lower layer parameters to partially recover corrupt packet and thus improve the throughput performance of the network [33]. This suggests that using lower layer information can improve overall network performance. Sections 7 and 8 showed that physical layer parameters like signal strength and noise can give us a lot of insight about link reception, supporting Jamieson et al.'s idea of using such parameters to improve higher layer performance.

12.3 Prior Work with 802.11

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 [3, 54]. The two technologies (802.11 and 802.15.4) use the same spectrum and have similar modulation schemes (BPSK or QPSK vs. OQPSK). However, these studies and ours reach

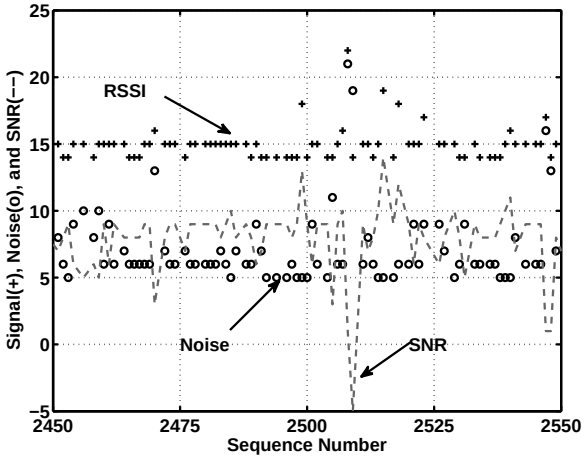

Figure 25: Signal power (+), noise power (o), and signal-to-noise ratio (- -) versus packet sequence number for a single link from the Roofnet traces. The transmit rate is approximately 78 packets per second. Units of power are reported in dBm with an unknown offset. Signal and noise both vary significantly over time.

opposite conclusions. Aguayo et al. observe very little correlation between SNR and PRR [3] and attributes this to mulitpath effects, while Reis asserts that RSSI asymmetries are a product of the environment rather than the node or wireless card [54].

To understand why our observation on SINR contradicts the findings of Aguayo et al., we examined publicly available Roofnet data. Figure 25 shows the signal and noise power plotted against sequence number of packets on an intermediate Roofnet link. As noise values are measured at the beginning of a packet transmission after a node has detected a clear channel, they represent a biased sample. Both the signal strength and noise power vary significantly over time. As the Roofnet study averaged SNR over one second periods, these variations are lost. Given the mentioned variation, and the fact that average SNR over this second-long period will lead to different averages for the same PRR, it is logical that plotting this average SNR versus PRR will show little, if any, correlation. As packets are discrete events, *averaged* SNR is a poor predictor of link quality. A recent study shows that Roofnet results are not due to multipath and instead they are due to interference from other near-by 802.11 networks [23]. This further suuports our argument that averaging SNR over 1s durations will not show good correlation with reception ratio especially when there is time varying interference.

Reis et al [55] carried out an indoor experiment with 15 802.11 nodes. They observed pair-wise links to be fairly stable over short durations. Our findings are similar to those of Reis et al. [55], displaying stability over short periods but increasing variability proportional to duration length. Reis et al. also examined link asymmetries and found them to be location-specific. Previous studies for low power wireless nodes [7, 21, 73], however, found asymmetry to be node-specific, and location independent. We ran experiments as outlined in section 9 with and without swapping nodes that saw asymmetries. We found that the asymmetries were in-fact node-specific. It is possible that external noise was at work in the experiment from Reis et al's work as external noise can cause packet losses at a node closer to it making the asymmetry location-specific (hidden terminal problem). The authors confirmed [53] that one of the nodes with the asymmetric link was under a desk with running machines and that the node had higher average noise readings than the other node.

Koksal et al. [36] used data from Roofnet to come up with routing metrics that capture short-term variability of links. They have proposed the use of modified ETX (mETX) and effective number of transmissions (ENT) as new routing metrics that consider both the link variability and the higher layer requirement. Section 10 briefly explored the implications of acknowledgements for routing metrics and showed that such metrics should not assume that the acknowledgement reception ratio is the same as the packet reception ratio of the reverse link.

Our paper differs from the prior studies presented here in one other way: we discuss the implications of our observations to the assumptions commonly made by protocol designers. The following section discusses the implications of our observations to higher layer protocols in detail and proposes a few ways to modify them to make them more efficient.

13. DISCUSSION

Prior sections presented a conceptual model of assumptions often used in protocol design. They address how 802.15.4 link layer observations do not always support these assumptions. They show when the assumptions made in the model may be true and when not. For example, if there is no high power external noise in the network then the *spatial* assumption in the conceptual model may hold. Knowing which assumptions will hold when can allow protocol designers to make informed design choices. For example, in the presence of high rate, high power external noise source such as 802.11, using MAC protocols like RTS/CTS that assume sole channel ownership may not work well as the data packet following RTS/CTS can still get corrupted by the external noise. In such cases, a protocol designer may choose to use CSMA/CA-like MAC protocol instead. However, if there is no external noise and if links have bursty reception then the success of RTS and CTS means that the data packets to follow are likely to succeed as well.

As all our observations in this paper are based on parameters that nodes can measure such as RSSI and noise, nodes can easily measure network characteristics. For example, nodes can measure long term link asymmetry from noise floor differences among them. Similarly, nodes monitoring noise samples as outlined in Section 8 can identify which channels are prone to external noise sources. They may use this information to avoid channels occupied by external noise or use such channels without assuming sole ownership.

At the highest level, these observations mean that when we think about protocols, we need to think not only about to whom a node sends a packet, but also when. Links have dynamics at a time scale both much faster (hundreds of milliseconds) and slower (hours) than we typically consider in protocol design. Link layer mechanisms, such as exponential backoff, are intended to avoid collisions, and react badly when applied to links on the edge of reception. Of course, these edge links are often the most valuable, as they repre-

sent the edge of a node's transmission range.

13.1 Open Questions

This paper identifies temporal and spatial correlations of packet delivery and link asymmetry as three key characteristics of a network. We believe that measuring these network characteristics can help us understand how protocols may work in that network and that they can give insights on why some protocols work differently on different networks. However, how we measure these metrics and what the methodologies are for measuring them are open questions. Our current research is to quantify a temporal correlation metric from conditional packet delivery functions (CPDFs) presented in Section 6 [63]. Of course, there may be more metrics that we can measure that characterize the network better. What those metrics are is also another open question. We hope that such metrics will give insights into designing more efficient future protocols.

Our study illuminates a fundamental challenge in packet-based network studies: one can only measure successfully received packets. Simply disabling CRC checking does not result in meaningful data, as random RF noise can often appear to be a start symbol. Therefore, all measurements are biased. A link whose SNR varies from +20 to -20dB will appear to have a 50% PRR at +20dB, which is incorrect. While one could possibly observe RSSI itself and look for an increase that denotes a packet (even if corrupted), such an approach can reduce but not remove the problem. The measurement bias occurs when received packets are close to the hardware noise floor, and packets with signal well below it cannot be observed. This also has implications to protocols, which use physical layer information to make decisions, such as MultihopLQI: the protocol can observe packets with few bit errors but miss half of the packets [20].

Perhaps the most challenging observation in this study is that the observed packet reception ratio greatly depends on exactly how it is measured. The behavior of a link measured at 1Hz is very different from its behavior measured at 100Hz. From a scientific standpoint, this means that, to be useful, network measurements should not only report link qualities but also their temporal properties, such as how long they are stable for. In some very recent work on a metric for measuring burstiness [62], we have proposed one way to quantify these temporal properties, but it is only a first attempt at what is a very complex problem.

13.2 Looking forward

Section 11 noted that the observations in this study do not necessitate discarding the existing work in protocol design: the incorrect assumptions of protocol design lead to sub-optimal rather than fundamentally broken protocols. In the past decade, the sensor network community has grown from every packet delivery being a struggle to having reasonably robust protocols and systems. Now that networks work, our next task is to make them work well and be as efficient as they can.

14. REFERENCES

[1] The Network Simulator - ns-2. `http://www.isi.edu/nsnam/ns/`.

[2] 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.

[3] 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.

[4] Atmel. Rf230 datasheet. `http://www.atmel.com/dyn/resources/prod_documents/doc5131.pdf`.

[5] R. Beckwith, D. Teibel, and P. Bowen. Unwired wine: Sensor networks in vineyards. In *Proceedings of IEEE Sensors*, 2004.

[6] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, 2005.

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

[8] 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.

[9] S. Chachulski, M. Jennings, S. Katti, and D. Katabi. Trading structure for randomness in wireless opportunistic routing. In *SIGCOMM '07: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, 2007.

[10] ChipCon. Cc1000 data sheet. `http://focus.ti.com/lit/ds/symlink/cc1000.pdf`, 2007.

[11] ChipCon. Cc2420 data sheet. `http://focus.ti.com/docs/prod/folders/print/cc2420.html`, 2007.

[12] T. Clausen and P. Jacquet. Optimized link state routing protocol (olsr). United States, 2003. RFC Editor.

[13] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A highthroughput path metric for multihop wireless routing. In *In Proceedings of MOBICOM 2003*. ACM Press, 2003.

[14] Crossbow. Micaz-based zigbee and wifi coexistance. `http://www.xbow.com/products/Product_pdf_files/Wireless_pdf/ZigBeeandWiFiInterference.pdf`.

[15] Crossbow. Micaz datasheet. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/MICAz_Kit_Datasheet.pdf`, 2005.

[16] Crossbow. Telosb datasheet. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/TelosB_Datasheet.pdf`, 2006.

[17] P. Dutta, D. Culler, and S. Shenker. Procrastination might lead to a longer and more useful life. In *The Sixth Workshop on Hot Topics in Networks (HotNets-VI)*, Nov. 2007.

[18] J. Elson, L. Girod, and D. Estrin. Emstar: Development with high system visibility. *IEEE*

Wireless Communication Magazine, December 2004.

[19] R. Fonesca, D. Culler, S. Ratnasamy, S. Shenker, and I. Stoica. Beacon vector routing: Scalable point-to-point routing in wireless sensornets. In Proceedings of the 2nd Symposium on Networked Systems Design and Implementation (NSDI '05), 2005.

[20] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis. Four-bit wireless link estimation. In *Hotnets-VI*, Atlanta, GA, Nov. 2007.

[21] 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.

[22] 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.

[23] D. Gokhale, S. Sen, K. Chebrolu, and B. Raman. On the feasibility of the link abstraction in (rural) mesh networks. In *The 27th Conference on Computer Communications IEEE INFOCOM 2008*, 2008.

[24] A. Hagedorn, D. Starobinski, and A. Trachtenberg. Rateless deluge: Over-the-air programming of wireless sensor networks using random linear codes. In *IPSN '08: Proceedings of the 2008 International Conference on Information Processing in Sensor Networks (ipsn 2008)*, pages 457–466, Washington, DC, USA, 2008. IEEE Computer Society.

[25] J. W. Hoi-Sheung Wilson So. Mcmac: A multi-channel mac proposal for ad-hoc wireless networks. Technical Report Technical Report, 2005.

[26] I. Howitt and J. A. Gutierrez. Ieee 802.15.4 low rate - wireless personal area network coexistence issues. In *WCNC '03: Wireless Communications and Networking, 2003*. IEEE, 2003.

[27] J. W. Hui and D. Culler. The dynamic behavior of a data dissemination protocol for network programming at scale. In *Proceedings of the Second International Conferences on Embedded Network Sensor Systems (SenSys)*, 2004.

[28] IEEE802.11. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2007.

[29] IEEE802.15.1. Part 15.1: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs(tm)), 2005.

[30] IEEE802.15.4. Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs), Oct. 2003.

[31] Infineon. Tda5250 data sheet. `http://www.infineon.com`, 2007.

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

[33] K. Jamieson and H. Balakrishnan. PPR: Partial Packet Recovery for Wireless Networks. In *ACM SIGCOMM*, Kyoto, Japan, August 2007.

[34] M. D. Jovanovic1 and G. L. Djordjevic1. Tfmac: Multi-channel mac protocol for wireless sensor networks. In *8th International Conference on Telecommunications in Modern Satellite, Cable and Broadcasting Services, 2007 (TELSIKS 2007)*, Sept. 2007.

[35] Y. Kim, H. Shin, and H. Cha. Y-mac: An energy-efficient multi-channel mac protocol for dense wireless sensor networks. In *IPSN '08: Proceedings of the Seventh International Conference on Information Processing in Wireless Sensor Networks*. ACM Press, 2008.

[36] C. E. Koksal and H. Balakrishnan. Quality-aware routing in timevarying wireless networks. 24:1984–1994, 2006.

[37] K. Langendoen, A. Baggio, and O. Visser. Murphy loves potatoes: Experiences from a pilot sensor network deployment in precision agriculture. In *The Fourteenth Int. Workshop on Parallel and Distributed Real-Time Systems (WPDRTS)*, 2006.

[38] H. Lee, A. Cerpa, and P. Levis. Improving wireless simulation through noise modeling. In *IPSN '07: Proceedings of the Sixth International Conference on Information Processing in Wireless Sensor Networks*. ACM Press, 2007.

[39] J.-S. Lee. An experiment on performance study of ieee 802.15.4 wireless networks. In *IEEE international conference on emerging technologies and factory automation*, 2005.

[40] P. Levis, N. Lee, M. Welsh, and D. Culler. TOSSIM: Simulating large wireless sensor networks of tinyos motes. In *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems (SenSys 2003)*, 2003.

[41] S. Lin, J. Zhang, G. Zhou, L. Gu, J. A. Stankovic, and T. He. Atpc: adaptive transmission power control for wireless sensor networks. In *Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 223 – 236, 2006.

[42] Y. Mao, F. Wang, L. Qiu, S. S. Lam, and J. M. Smith. S4: Small state and small stretch routing protocol for large wireless sensor networks. In Proceedings of the 4th Symposium on Networked Systems Design and Implementation (NSDI '07), 2007.

[43] MintRoute. Mintroute. TinyOS repository, www.tinyos.net/tinyos-1.x/tos/lib/MintRoute/.

[44] Moteiv. Boomerang. `http://www.moteiv.com`.

[45] Moteiv. Telos datasheet. `http://www.moteiv.com/products/docs/tmote-sky-datasheet.pdf`.

[46] MultiHopLQI. `http://www.tinyos.net/tinyos-1.x/tos/lib/MultiHopLQI`, 2004.

[47] R. Musaloiu-E and A. Terzis. Minimising the effect of wifi interference in 802.15.4 wireless sensor networks. 3:43–54, 2008.

[48] 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.

[49] C. E. Perkins, E. M. Belding-Royer, and S. Das. Ad hoc on demand distance vector (AODV) routing. IETF Internet draft, draft-ietf-manet-aodv-09.txt, November 2001 (Work in Progress), 2001.

[50] M. Petrova, L. Wu, P. Mähönen, and J. Riihijärvi. Interference measurements on performance degradation between colocated ieee 802.11g/n and ieee 802.15.4 networks. In *ICN '07, Sixth International Conference on Networking (ICN 2007)*. IEEE Computer Society, 2007.

[51] D. Puccinelli and M. Haenggi. Duchy: Double cost field hybrid link estimation for low-power wireless sensor networks. In *The Fifth Workshop on Embedded Networked Sensors (Hot EmNets 2008)*, June 2008.

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

[53] C. Reis. Asymmetry in reception. `http://www.cs.washington.edu/homes/creis/wireless/asymmetry/asymmetry.html`, 2005.

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

[55] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In *SIGCOMM*, pages 51–62, 2006.

[56] RFM. Tr1000 data sheet. `http://www.rfm.com/products/data/tr1000.pdf`, 2008.

[57] S. Y. Shin, H. S. Park, S. Choi, and W. H. Kwon. Packet error rate analysis of ieee 802.15.4 under ieee 802.11b interference. In *WWIC '05, Wired/Wireless Internet Communications, 2005*, 2005.

[58] K. Shuaib, M. Alnuaimi, M. Boulmalf, I. Jawhar, F. Sallabi, and A. Lakas. Performance Evaluation of IEEE 802.15.4: Experimental and Simulation Results. *Journal of Communications*, 2(4), 2007.

[59] D. Son, B. Krishnamachari, and J. Heidemann. Experimental analysis of concurrent packet transmissions in low-power wireless networks. In *SIGCOMM*, 2006.

[60] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. Some implications of low power wireless to ip networking. In *Proceedings of The Fifth Workshop on Hot Topics in Networks (HotNets-V)*, Nov. 2006.

[61] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. Understanding the causes of packet delivery success and failure in dense wireless sensor networks. Technical Report Technical Report SING-06-00, 2006.

[62] K. Srinivasan, M. Kazandjieva, S. Agarwal, and P. Levis. The beta factor: Measuring wireless link burstiness. In *Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2008.

[63] K. Srinivasan, M. A. Kazandjieva, S. Agarwal, and P. Levis. The β-factor: Improving bimodal wireless networks. Technical Report Technical Report SING-07-01, 2007.

[64] K. Srinivasan and P. Levis. Rssi is under appreciated. In *Proceedings of the Third ACM Workshop on Embedded Networked Sensors (EmNets 2006)*, May 2006.

[65] T. Stathopoulos, L. Girod, J. Heidemann, and D. Estrin. Mote herding for tiered wireless sensor networks. Technical Report Tech Report 58, 2005.

[66] 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.

[67] 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.

[68] 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.

[69] 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.

[70] 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.

[71] 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.

[72] X. Zeng, R. Bagrodia, and M. Gerla. Glomosim: A library for parallel simulation of large-scale wireless networks. In *Workshop on Parallel and Distributed Simulation*, pages 154–161, 1998.

[73] 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.

[74] J. Zheng and M. J. Lee. A comprehensive performance study of ieee 802.15.4. IEEE Press Book, 2004.