The β-factor: Improving Bimodal Wireless Networks

Technical Report SING-07-01

Kannan Srinivasan, Maria A. Kazandjieva, Saatvik Agarwal, and Philip Levis
Computer Systems Lab
Stanford University
Stanford, CA 94131
{srikank,mariakaz,legolas}@stanford.edu, pal@cs.stanford.edu.

Abstract

We explore how wireless networks can take advantage of bursty links. Measuring 802.15.4 in three testbeds, we find that most intermediate links are bimodal: they oscillate between poor and good delivery. We present a metric to quantify link bimodality and name this value β. We propose an algorithm that can boost the observed reception ratio of high β links by trading off latency for efficiency. We find this policy can reduce end-to-end route costs by up to 80%, with 20% of routes improving more than 20%. We examine 802.15.4 physical-layer testbed data as well as traces from 802.11b networks and find β has broader relevance than the testbeds we measure. Based on these results, we show how changing a single constant in a standard sensor network data collection protocol can reduce transmission counts by 15%.

1. INTRODUCTION

Many deployment experiences have found that high performance in wireless meshes requires continually measuring the imperfect and time-varying wireless channel. Srcr [2], COPE [17], MORE [7], and ExOR [3] measure the estimated transmission count (ETX) [9] using periodic broadcasts. ENT and mETX consider how retransmissions affect losses to higher-level protocols such as TCP [19], while ETT (expected transmission time) [10] extends ETX to consider multi-rate link layers. All of these protocols have shown that using run-time link measurements improves throughput, as nodes can dynamically respond to channel changes.

In this paper, we consider an analogous but subtly different problem, improving transmission efficiency. Rather than try to optimize the number of packets received, we seek to improve the probability any given packet is received. Optimizing throughput makes sense for latency-sensitive, interactive protocols and applications. But for networks that are constrained by energy, rather than capacity, reducing the number of packets needed to deliver data is more important than throughput.

Wireless sensor networks [27, 29] are one extreme example, where long lifetimes require spreading communication over long periods of deep sleep. Many sensornets can therefore trade off latency for efficiency. In addition to choosing which node they transmit to, protocols can also decide *when* they transmit. This increased flexibility can improve protocol efficiency and lengthen node lifetimes.

Deployment studies have found packet delivery in wireless networks is often bursty [9, 6, 20]. Section 2 of this paper presents experimental data we have gathered on 802.15.4, the dominant low-power wireless link layer. Measuring three testbeds, we find that the distribution of observed packet reception ratios depends on how long one takes to measure them. Longer measurement periods see more intermediate-quality links than shorter ones do.

Looking at these results more closely, we find that many intermediate links are bimodal: they oscillate between states of high and low reception. This bimodality causes a high correlation between packet successes as well as a high correlation between packet failures. Section 3 introduces a bimodality measure, whose value we call β. A β of 1 indicates a high correlation between packet events, while β=0 indicates packet events are independent. In one testbed, we find that 40% of intermediate links have a β value above 0.9, and that bimodality decays over time. As the interval between packet sends increases, β shrinks.

These results suggest that controlling when a protocol transmits can boost link quality. Section 4 measures what happens when a protocol pauses after a packet delivery failure. As long it delivers packets successfully, a protocol sends as quickly as possible, opportunely taking advantage of correlated successes. Pausing after a lost packet breaks the correlation between failures. β distributions show

500ms to be a good tradeoff point in the networks we measured. Using packet traces, we find that opportune transmissions improve network efficiency slightly (2-4%), but improve some routes by as much as 42%, and in some extreme cases by over 80%.

These results are promising, but are they specific to the testbeds we used, or are they more generally applicable to wireless meshes? Section 5 examines the testbed packet traces to find the root causes of bimodality in 802.15.4. Bimodality in 802.15.4 is mostly due to small shifts of 1-2dB in received signal strength, propagation changes that are typical in real-world networks. Examining data from several prior publications on 802.11b [9, 23], we find in Section 6 that 802.11b observes bimodality similar to 802.15.4 in both indoor and outdoor environments (albeit probably for different reasons than those Section 5 describes), but a 1Mbps outdoor mesh exhibits very low β values.

All of these measurements suggesting improvements come from post-facto analysis of packet traces. Section 7 evaluates whether transmitting opportunely improves protocol efficiency in real networks. It presents results from modifying a standard sensornet protocol's transmit timing. Changing a single protocol constant decreases the overall network delivery cost by 15%. Furthermore, as transmitting opportunely naturally sends packets in bursts, it can reduce energy costs by amortizing wakeup packet costs. Section 8 discusses the relevance of this paper in respect to related work and concludes.

This paper makes three research contributions. First, it shows that using periodic control packets to estimate link quality introduces an inherent bias and inaccuracy when measuring bimodal links. Second, it describes an algorithm to quantify link bimodality, and finds this value to have high predictive value. Finally, it proposes a way to improve efficiency in networks with bimodal links, and exhaustively evaluates this approach using traces from five networks as well as a protocol experiment on a testbed. Together, these suggest we need to rethink current approaches to wireless protocol design and analysis: understanding how protocols perform requires measuring fine-grained temporal properties, and measuring links with special control packets has fundamental limitations.

2. 802.15.4 PACKET DELIVERY

This section introduces 802.15.4 and its packet delivery behavior. It begins with background information on 802.15, describes the testbeds we used, and defines the terminology we use to describe links. It presents packet delivery results from three testbeds.

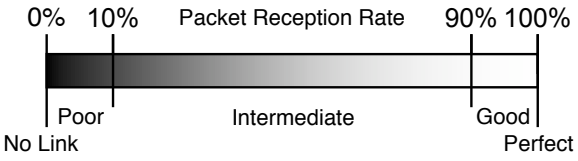

Figure 1: Terminology used to describe links based on PRR. Poor links have a PRR $< 10\%$, intermediate links are between 10% and 90%, and good links are $> 90\%$. A PRR of 100% is a perfect link. A link that receives one or more packets is a communicating link.

We observe that packet reception ratio distributions depend on how they are measured. Measurements over short intervals observe fewer intermediate links than measurements over longer periods do. The rest of the paper explores the causes and implications of this phenomenon.

2.1 802.15.4 and Testbeds

802.15.4 is a IEEE PHY-MAC standard for low power, low data rate networks. It has a data rate of 250kbps and a range of approximately one hundred meters. It provides 16 channels, numbered 11-26. The channels are 5 MHz apart in the 2.4 GHz band (2405 MHz - 2480 MHz), overlapping with 802.11b and 802.15.1 (Bluetooth).

We measured 802.15.4 using three wireless sensornet testbeds. Most experiments use the 100 node Intel Mirage testbed [15]. We also present results from the 30 node university testbed in UC Berkeley's Soda Hall. The Mirage and university nodes are on the ceiling. Finally, we examine an outdoor 20 node lake testbed, whose nodes were spaced 4 feet apart in the dry Lake Lagunita lake bed on Stanford campus. The lake nodes were arranged in a line and all had clear line of sight. All nodes in these experiments ran TinyOS [14] and use the CC2420 802.15.4 chip [8], which provides variable transmit power control from 0dBm to -20dBm.

2.2 Packet Delivery

Prior studies of wireless networks have observed that links have a wide range of packet reception ratios (PRR) which can vary significantly over time [1, 5, 23, 21]. To determine whether 802.15.4 behaves similarly, we measured reception ratios in the university, Mirage and lake testbeds. In the rest of this paper, we describe links as poor, intermediate, good, or perfect using the definitions shown in Figure 1 and use the terms link quality and packet reception ratio interchangeably. As prior studies have shown 802.15.4 links can vary significantly over time [21], we measured reception ratios over differ-

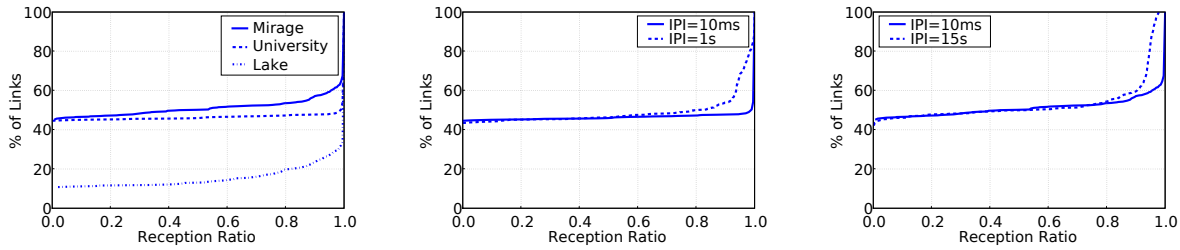

(a) University and Mirage IPI=10ms, Lake IPI=50ms (b) University (c) Mirage

Figure 2: Reception ratio and the CDF of proportion of links in the three testbeds for channel 26. The percentage of intermediate links is small compared to good and bad links, and it increases as the inter-packet interval increases.

ent time scales by sending 200 broadcasts with varying inter-packet intervals (IPI, the time between packet transmissions). We used inter-packet intervals ranging from 10ms up to 15 seconds. All packets used the standard TinyOS CSMA layer, and we controlled transmission timing so there would be no collisions. The lack of a wired backchannel prevented lake nodes from having an inter-packet interval below 50ms.

Figure 2(a) shows the reception ratio distribution in the three testbeds on channel 26 with small inter-packet intervals. 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. Of these communicating links, 19% in Mirage are intermediate, 14% in the lake are intermediate, and 5% in the University are intermediate. These are all quite smaller than what has been observed in other networks. Even the 19% in Mirage is much less than the 50% reported for earlier sensor platforms and the 58% reported for Roofnet. Compared to these other networks, 802.15.4 has a much sharper reception distribution. While there are enough intermediate links that protocols cannot ignore them, they do not dominate.

2.3 Time and Frequency Effects

Figures 2(b) and 2(c) show how the time interval between packets affects the reception ratio distribution. Increasing the IPI from 10ms to 1 second increases the percentage of communicating links in the university testbed that are intermediate from 5% to 19%. Mirage increases from 19% to 23% as IPI increases from 10ms to 15s. As the reception ratio is calculated over 200 packets, the packet interval determines the measurement time: an interval of 10ms takes 2 seconds, 1s takes 3 minutes, and an interval of 15s takes 50 minutes.

Timing is not the only factor that affects link distributions. Figure 3 shows how channel selection

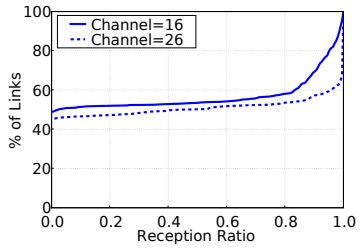

Figure 3: CDFs of link qualities in Mirage on Channels 16 and 26. The proportion of perfect links is more in channel 26 than in channel 16: 60% in 26 and 12% in 16 of all the communicating links.

changes the PRR distribution in Mirage. Channel 16 has far fewer perfect links than channel 26: of the communicating links, 60% are perfect in channel 26 while only 12% are perfect in channel 16. Correspondingly, 35% of the communicating channel 16 links are intermediate, compared to 17% of the channel 26 links.

These results lead to two major observations. The first is that frequency affects link distributions. While this is not surprising, learning why is an important step to better understanding wireless behavior. We defer this question to Section 5.

Second, and perhaps more interestingly, the percentage of intermediate links depends on the time scale over which a protocol measures them. Over shorter periods, links have a higher chance of being perfect or non-existent. Over longer periods, their chance of being intermediate increases. In the next section, we examine this behavior more closely, finding it is due to links on the edge of reception sensitivity moving between poor and good states. As the measurement period increases, so does the chance of observing a transition. While this is a simple observation, it has deep implications for wireless protocol design: it means that the data plane may observe different link qualities than a control plane which sends link measurement packets.

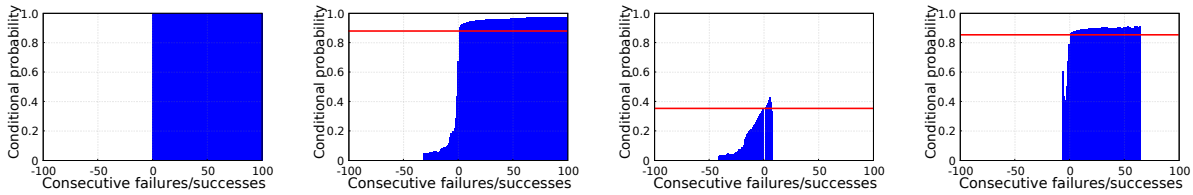

(a) KW= $0, \beta = 1$. (b) KW= $0.06, \beta = 0.8$. (c) KW= $0.20, \beta = 0.5$. (d) KW= $0.14, \beta = 0.3$.

Figure 4: CPDFs with their corresponding KW-distances and β value. The ideal bimodal link is shown in (a), while (b)-(d) are links observed in the Mirage testbed. The red horizontal line shows the overall reception ratio of the link. Positive X-axis values are consecutive failures while negative X-axis values are consecutive failures.

3. BIMODAL RECEPTION

In this section, we examine intermediate links in greater detail. Many intermediate links experience exhibit bimodal reception: they go through periods of good and poor link quality. As the inter-packet interval increases, so does the chance of observing a mode transition and therefore an intermediate link. We define β, a metric to quantify this bimodal behavior. We measure how β changes as the inter-packet interval increases. We see many Mirage links have high β values over short time scales, denoting correlated delivery successes and failures. As the time between packets increases, β decreases and packet deliveries appear more independent.

3.1 Conditional Delivery

Before we examine how links behave, we need a way to describe their behavior. Conditional packet delivery functions (CPDFs) provide a succinct way to describe the durations of packet delivery correlation [20]. The conditional packet delivery function $C(x)$ is the probability the next packet will succeed given x consecutive packet successes (for $x > 0$) or failures (for $x < 0$). For example, $C(5) = 83\%$ means that the probability a packet will arrive after five successful deliveries is 83%, while $C(-7) = 18\%$ means that the probability after seven consecutive losses is 18%.

Figure 4 shows four sample CPDFs. A link with independent losses will have a flat CPDF: the probability of reception is independent of any history. In contrast, a perfectly bimodal link – one which between perfect and non-existent – will have a CPDF that looks like Figure 4(a).[1]

We calculated CPDFs by programming nodes on the Mirage testbed to broadcast 100,000 packets with an inter-packet interval of 10ms. A node did not start its broadcasts until the one before it had completed all 100,000. We used 100,000 packets to provide reasonable confidence intervals to the CPDF values. Each value in a CPDF has a minimum of 100 data points.[2] Figures 4(b)-4(d) show the CPDFs of three sample links.

Unlike autocorrelation, which quantifies the correlation of a signal with a phase-shifted version of itself, CPDFs quantify the conditional probability of an event. As the reception ratio is typically neither 0% or 100% in bursts, modality shifts do not have a regular periodicity and better resemble a random process. Therefore, links in our experiments do not have significant autocorrelation, and Allan deviation analyses are inconclusive. We defer a more detailed discussion of the relationship to existing metrics and statistical models to Section 8.

3.2 Kantorovich-Wasserstein

CPDFs distill a long vector of packet delivery successes and failures into a concise representation of burstiness. While CPDFs can give a good visual intuition of link behavior, we need to distill bimodality down to a scalar value. To do so, we borrow an approach from our prior work [20] and use the Kantorovich-Wasserstein (KW) distance [24] to measure how close a CPDF is to the CPDF of the ideal bimodal link (Figure 4(a)). In the rest of this paper, When we refer to the distance of a link, we mean the Kantorovich-Wasserstein distance to the ideal bimodal link. A lower distance therefore means a link is more bimodal.

The captions in Figure 4 include the distances of the three example Mirage links. A CPDF that is further from the ideal bimodal link has a larger distance. Figure 4(d) has a larger distance than Figure 4(b) because the small number of data points with negative x values have reasonably high recep-

[1]There is an inherent timescale assumption in this description: the modal shifts must be at an interval longer than the CPDF x-axis range. Bimodal shifts are fast enough to occur within the CPDF range make a link's losses look more independent.

[2]100 data points gives a worst case 95% confidence interval of [p-0.1,p+0.1], where p is the empirical conditional probability

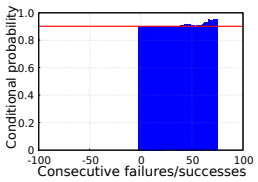

(a) KW=0.1, $\beta = 0$

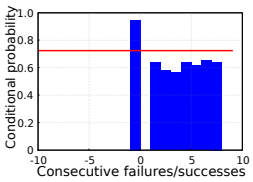

(b) KW=0.45, $\beta = -0.36$ (note smaller x range)

Figure 5: Two link edge cases. Independent links with high or low reception ratios have a low distance but a β close to 0. Links with negative correlation have a $\beta < 0$.

tion ratios.

3.3 β: The Bimodality Metric

While distance is a useful measure, it is not sufficient to quantify bimodality. CPDFs often do not have the same number of elements with positive and negative x values. For example, consider a link with a reception ratio of 90% whose packet deliveries are independent. Figure 5(a) shows the CPDF of such a link (synthetically generated with a random process), where the conditional probability is for the most part constant for all x. This hypothetical link, however, has a low distance from the ideal bimodal link: 0.1.

This low distance is because the link has many data points for $x > 0$ and few for $x < 0$. The values for $x > 0$ are very close to the bimodal link's 100% and there are many of them. The values for $x < 0$ are very far from the bimodal link's 0%, but with a packet reception ratio of 90%, the chances of many consecutive losses is small.

As bimodal and independent reception are two two ends of a spectrum, we quantify bimodality as distance from the ideal bimodal link compared to an independent link with the same PRR. For brevity, we call this bimodality metric β:

$$\beta = \frac{KW(I) - KW(C)}{KW(I)}$$

where $KW()$ is distance from the ideal bimodal link, C is the CPDF of the link in question, and I is the CPDF of an independent link with the same PRR.

A perfectly bimodal link has a β=1, while a link with independent deliveries has a β=0. To get a sense of what β looks like, the captions in Figure 4 include the corresponding β values of the four example links.

It is possible for a link to have a negative β value. This happens when there is a negative correlation in packet reception: as more packets are received the next packet is more likely to fail and as more packets are lost the next packet is more likely to be received. Figure 5(b) shows the CPDF of link we measured where the probability of next packet succeeding decreases as more packets are received. This can happen, for example, if bursts are so short that a few packet receptions mean the end of the burst is more likely to be near.

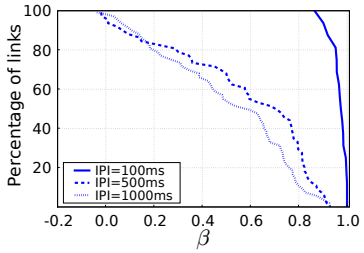

Figure 6: CCDF of intermediate link β values on the Mirage testbed for different inter-packet intervals (up and to the right means more bimodal). Larger inter-packet intervals observe many more intermediate links, and these links are less bimodal: at 10ms, the minimum intermediate link β is ≈0.85.

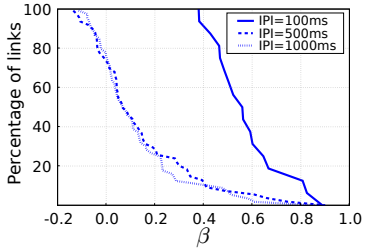

Figure 7: CCDF of intermediate link β values in Mirage on channel 16, which overlaps with a nearby 802.11 network. Fewer than 15% of the links have β above 0.8 for any inter-packet interval, even at an inter-packet interval of 10ms. WiFi decreases bimodality.

3.4 β Distributions

Now that we can quantify link bimodality, we can examine the distribution of β values of intermediate links to better understand how bimodal they are. Figure 6 shows the complementary CDF (CCDF) of β for intermediate Mirage links. This plot is from the same data as in Figure 4. By subsampling each 100,000 (16 minute) packet trace, we can calculate β for different inter-packet intervals.

Figure 6 shows that as the inter-packet interval increases, β decreases. At an IPI of 10ms, 40% of intermediate links have a β above 0.9. At 500ms and 1s, however, less than 5% of the links have a β this high. Furthermore, the percentage of links that have a β close to zero increases as the inter-packet interval increases: no links at 10 ms, 10% at 100ms

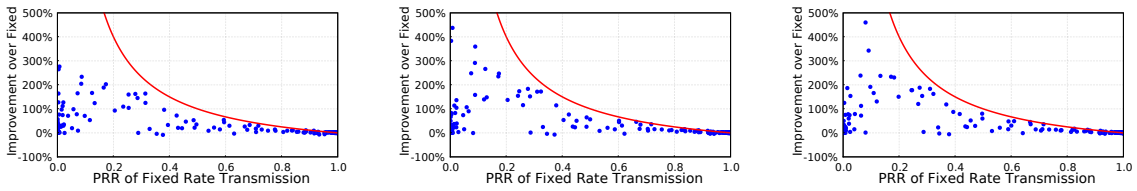

(a) IPI=100ms (b) IPI=500ms (c) IPI=1000ms

Figure 8: Link-level reception ratio improvement of opportunistic transmissions over fixed rate transmissions in the Mirage traces. The red curve shows the maximum possible improvement (a PRR of 100%). Increasing the probe interval to 500ms improves greatly over 100ms, but 1 second does not improve much further

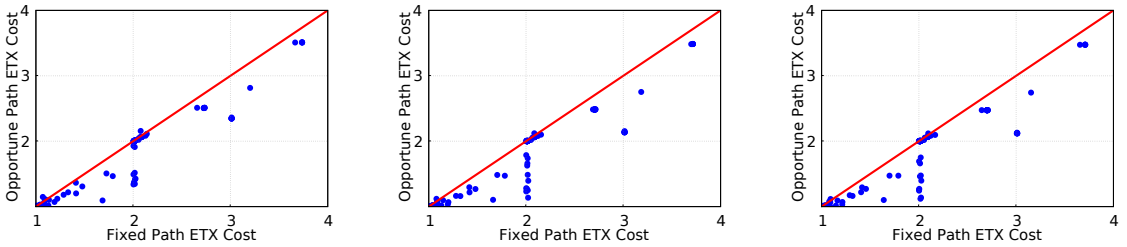

(a) IPI=100ms (b) IPI=500ms (c) IPI=1000ms

Figure 9: Opportune transmissions improve the path ETX in Mirage (lower is better).

and 25% at 1 second.

Section 2 noted that PRR distributions in the Mirage testbed differ based on channel. Figure 7 shows the cumulative distribution of the β-factor in channel 16. The percentage of links with $\beta >$ 0.8 is less than one third than that in the channel 26 experiment (Figure 6). In the Mirage testbed, links on channel 16 are less bimodal than those on channel 26. Section 5 investigates why the choice of 15.4 channel affects link bimodality.

3.5 Observations

This inverse relationship between the inter-packet interval and β means that, sending packets further apart decreases the correlation in their fate. This correlation is in terms of both successful and failed transmissions. In networks which exhibit high β values, a failed packet transmission means that the chances that an immediate retransmission will succeed is low. However, if a node waits long enough, it will experience a small β, and the probability of delivery will be independent of the prior loss.

A protocol that understands β can play the odds of delivery, trading off latency to improve its communication efficiency. The results in Figure 6 suggest that waiting for 500ms represents the knee of the efficiency benefit curve in the Mirage network. The following section shows one way protocols can improve their efficiency with this knowledge.

4. IMPROVING BIMODAL NETWORKS

This section presents an algorithm to increase the reception ratio of bimodal links, thereby improving network energy efficiency. A node sends packets as quickly as possible until one is lost. When a packet delivery fails, a node waits. This back-off gives the next packet an independent chance of delivery. While pausing breaks the packet loss correlation, sending packets back to back preserves the correlation between packet successes, thus improving link quality. We call this approach *opportune transmissions*, because a node takes advantage of opportune moments of high reception. T

This section presents a trace-based evaluation of our approach on the Mirage network, finding it can elevate many intermediate links to be good, thereby reducing end-to-end route costs of some routes by up to 45%. Supporting the observations on β in the prior section, we find the best back-off time for Mirage to be 500ms.

4.1 Reception Ratio Improvement

As a first step in quantifying what effect opportune transmissions have on a network, we examine how they change link reception ratios. For each link in the 100,000 packet Mirage traces, we measure the fixed period PRR of an inter-packet interval as if a node tried to periodically transmit. We also measure the PRR if a node uses opportune transmis-

sions. On a packet delivery success, a node sends back-to-back packets until a failure occurs. On a failure, it waits until the next time the fixed period algorithm would transmit. We use this variable waiting period to a fixed point rather than waiting for a fixed period, because the latter introduces significant noise due to phase shifts. Both cases send the same total number of packets.

Comparing fixed and opportune reception ratios shows whether a node can improve its efficiency by taking advantage of correlated successes. Figure 8 shows how opportune transmissions affect the observed reception ratio at different inter-packet intervals. Many links see improvements, some to near-optimal levels, and a small number see small degradations. For example, in Figure 8(b) two links with fixed reception ratios of about 0.67 and 0.73 reach a ratio very close to 1. Increasing the probe interval garners larger reception ratio improvements because it decreases the correlation of packet delivery failures. As the results in Section 3 indicated, the knee of the curve in breaking correlations occurs at 500ms, after which the improvements level off. The leveling off is hard to see in Figure 8 due to the density of points with high fixed reception ratios, but later results show it more clearly.

4.2 End-to-End Path Improvement

Link improvements do not guarantee improvements in end-to-end routes. It may be that the improved links are irrelevant for low-cost routes. Additionally, large improvements do not necessarily lead to large gains: a 100% improvement from 0.1 to 0.2 is still a reception ratio of 0.2. We therefore examine whether the link improvements observed in Figure 8 translate to lower end-to-end costs.

Figure 9 compares the minimum-cost paths between all node pairs for fixed and opportune transmissions, where a link cost is its expected transmissions per delivery (ETX), or $\frac{1}{PRR}$. For a probe interval of 500ms, the mean improvement over all paths is about 4.5% with the maximum reaching 45%. Overall, using opportunistic transmissions decreases the least-cost path for nearly all node pairs, thus improving the network's transmission efficiency. While a 500ms pause observes significant improvements over 100ms (note the larger number of links with a fixed ETX of 2 and their downward shift), increasing the value to 1 second does not see significant further improvements.

4.3 Lower Transmit Power

The end-to-end results so far have come from one Mirage dataset where nodes transmitted at 0dBm.

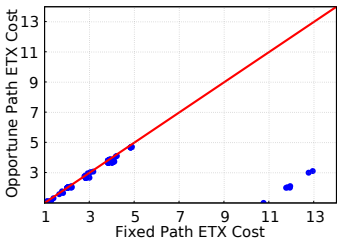

Figure 10: ETX improvements at -15dBm and a packet interval of 500ms. A number of paths all experience a constant shift ETX reduction of approximately 90%. This suggests that they all shared a single link which opportune transmissions improved greatly.

This gave a well connected network with a maximum path ETX of only 4. As a simple sanity check of whether our observations have overfit to one particular network trace, we examine what happens in a less connected network.

Figure 10 shows the CDF of end-to-end ETX reductions when nodes transmitted at -15dBm. For most paths, the reduction in ETX is small. For several paths, however, the reduction is dramatic – up to 90%! A linear shift of the data points in the lower right of the figure suggests that a number of paths have improved due to one link's improvement. This suggests that mesh protocols that seek to minimize route cost can see significant benefits from even a small number of link improvements.

4.4 Three Caveats

The end-to-end study carried out in this section provides evidence that opportune transmissions can reduce end-to-end costs, in some cases by over 90%. While these results are inspiring, three aspects of our methodology so far prevent us from generalizing these results to real-world wireless protocols:

- While Figure 2 showed that 2 other 802.15.4 testbeds see similar PRR distribution shifts to Mirage, generalizing our results to 802.15.4 networks as a whole requires understanding the cause of bimodality and whether it is common.

- Even if bimodality is a general phenomenon in 802.15.4, there are many other link layers. It could be β is only relevant in 802.15.4. Understanding which link layers have high β values and which do not can provide a quantitative basis for when opportune transmissions are effective. Furthermore, if the applicability is not uniform, this introduces challenges in cleanly abstracting wireless to higher-layer protocols.

- Data traces are not representative of real net-

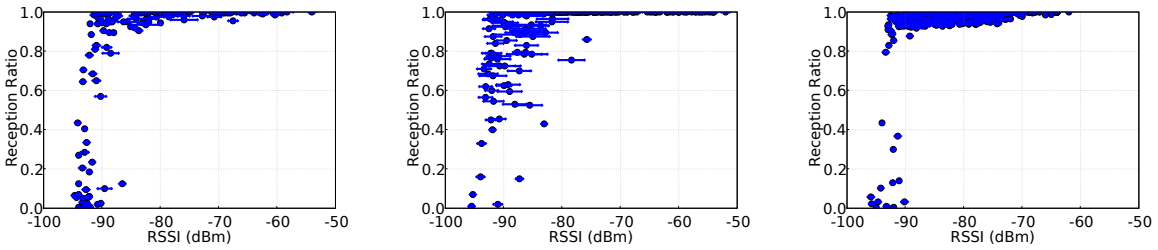

(a) Mirage. (b) Lake. (c) University.

Figure 11: Packet reception ratio versus received signal strength on channel 26 with IPI=50ms. Each data point is for a directional node pair. The average RSSI is marked by circles and the error bars show one standard deviation. Overall, RSSI and reception ratio correlation is similar across different testbeds with some outliers.

work traffic. In the traces of 100,000 broadcasts, for example, each set of broadcasts (and therefore link measurements) occurred 15 minutes apart, yet when we compute minimum-cost paths we assume they were taken at the same time. Furthermore, the analysis so far assume a protocol has perfect knowledge of whether a packet was delivered; in real networks a protocol must rely on possibly lossy link-layer acknowledgements. Concluding whether opportune transmissions are beneficial requires measuring a protocol on a real network.

The next 3 sections address these caveats in turn.

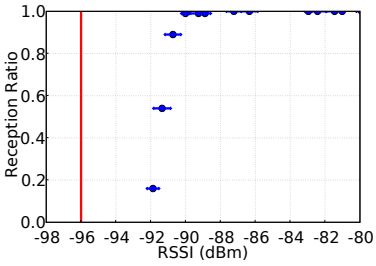

Figure 12: PRR versus RSSI plots for two nodes connected through a variable attenuator. The red line shows the noise floor of the receiving node. Intermediate PRRs are within 1.5dB range.

5. CAUSES OF BIMODALITY

This section examines what causes the bimodality we observe in the 802.15.4 testbeds. We find that bimodal links are often links on the edge of reception sensitivity, such that small 1-2dB swings in signal strength significantly change the observed packet reception ratio. Returning to the observation that channels 16 and 26 behave differently, we find that low-rate 802.11b traffic can appear as independent packet losses and decrease β.

5.1 Received Signal Strength (RSSI)

Figure 11 plots the signal strength of received packets of links against their reception ratio. In all the three testbeds, there is a general trend with a few outliers: if the mean received signal strength (RSSI) is above -80dBm then the link is almost always good. The two exceptions to this rule occur in the lake testbed, where there were people actively moving between the nodes. Below this value, there is a grey region of good, intermediate, and poor links [26].

To better understand this grey region, we followed Aguayo et al.'s methodology [1], wiring two nodes together through a variable attenuator via shielded SMA cables. At each attenuation level (1 to 64dB) one node transmitted 100 packets with an inter-packet interval of 50ms: the receiver logged the RSSI and sequence number of received packets to flash (which limited the IPI to $\geq$50ms). We measured background (hardware/AWGN) noise by sampling the RSSI register of the CC2420 of both nodes when there was no traffic. We calculated the noise floor as the mode of these samples.

Figure 12 shows that there is a crisp RSSI/PRR curve. The small error bars show that the RSSI at each attenuation level was very stable. The receiver does not receive any packets below a signal to noise ratio of 4dB. All of the intermediate PRRs are within a 1.5 dBm range of -92 to -90.5dBm. If the signal strength is close to the noise floor, a 1.5dB shift can change it from a good link to a bad link and vice versa. We repeated this experiment for two separate node pairs, and both near-identical 4dB signal-to-noise thresholds and intermediate link windows of 1.5 dB.

But not all nodes have the same noise floor. In the attenuator experiment, one node had a noise floor of -98dBm and observed a good link at -92dBm;

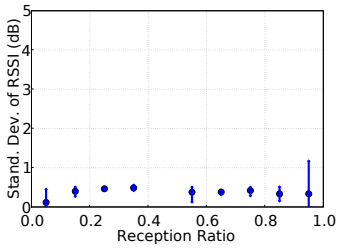

(a) Mirage, IPI = 10ms, Channel 26.

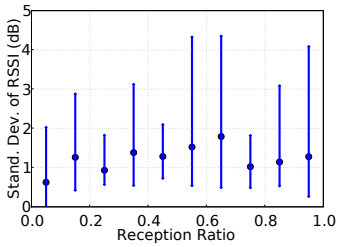

(b) Mirage, IPI = 14s, Channel 26.

Figure 13: The plot shows mean, max and min RSSI observed at different reception ratios. Each data point is a bin of all links within a 10% range: the fourth data point, for example, is 30-40%. RSSI is stable for short term traffic and varies more over longer time periods.

the node in Figure 12 had a noise floor of -96dBm and observed a good link at -90dBm. Across a large network, this variation can be large. In the Mirage testbed experiment in Figure 11(a), 5 nodes had noise floors at -98dBm, 8 at -97dBm, 4 at -96dBm, 3 at -95dBm, 2 at -94dBm, 3 at -93dBm and 1 at -92dBm. Therefore, even if every node observes a crisp signal-to-noise/packet reception curve, the threshold values for these curves are spread across 6dB. While the broad spread is partially due to the fact that these are very inexpensive, low-power radios, examining datasets from the Roofnet 802.11b mesh [1] we saw a 6dB spread in noise floors, albeit with 80% of the nodes having the same value.

Noise floor variations only partly explain the grey region. The grey region's range (-85dBm to -96dBm: 11dB) is greater than the range of the noise floors (-98dBm to -92dBm: 6dB). We believe this additional 5dB of range is due to an inherent measurement bias common to all such studies: nodes only measure the signal strength of received packets. If there are 6-7dB changes in signal strength, then link may transition from good to non-existent, yet the receiver will see only the RSSI of a good link. The comparatively short range and low bit rate of 802.15.4 means it does not observe the same multipath inter-symbol self-interference observed in Roofnet's 802.11b [1].

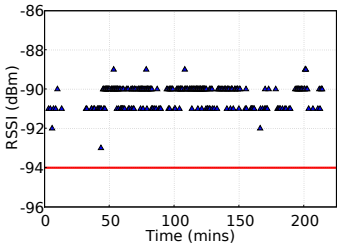

Figure 14: RSSI variation over time on a single link. The noise floor at the receiver node was -94dBm. The receptions at -90dBm form dense clumps, while the receptions at -91dBm are sparsely scattered.

Nevertheless, we cannot definitively explain the entire width of the grey region, and leave such investigations to future work with software radios.

Figure 13 is a plot for the Mirage testbed showing the minimum, mean and the maximum standard deviations of RSSI of different reception ratios for inter-packet intervals of 10ms and 14s. With 10ms intervals, the average standard deviation is below 1dB across all PRRs for all links and the maximum is 1.5dB. With 14 second intervals, the average standard deviation is more than 1dB for all links with a reception ratio above 0.1 and the maximum is as high as 4.2dB.

The stability of RSSI over short time spans and its variation over longer time spans suggests that the channel variations may be the cause of bimodality. Figure 14 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 above -91dBm, a few are as low as -93dBm. If a link is near the cusp of reception sensitivity, then slight variations can cause packet losses and make the link intermediate. Clustered receptions such as these are common in high β links, and are a dominant cause of link bimodality. RSSI shifts of this kind are typical of all but the most controlled environments, either due to environmental effects [21] or simple multipath fading.

5.2 Bimodality and Interference

Figure 3 and Figure 6 showed that Mirage had more intermediate links on channel 16 than 26. The university testbed observed a similar shift, while we found channel choice did not affect the link distribution in the lake testbed. In Mirage, this change in link distribution is accompanied by much lower β values on channel 16 than 26. This change in β is due to interfering 802.11 transmissions. Channel 16 is in the middle of 802.11b channel 6, while channel 26 is outside the spectrum used by 802.11b [25].

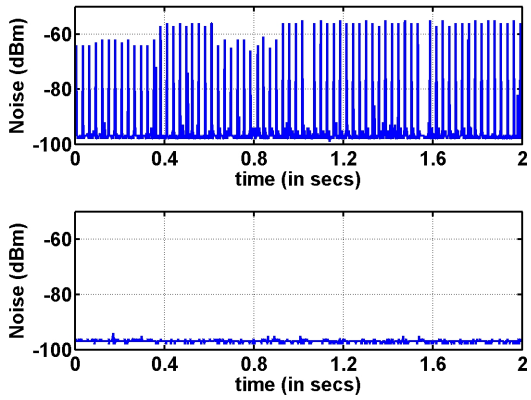

Figure 15: High frequency noise samples on channels 16 (top) and 26 (bottom) in Mirage. No 802.15.4 nodes were transmitting. Channel 16 shows large spikes while channel 26 is quiet.

Figure 15 shows 1kHz RSSI samples at a single node on channels 16 and 26 in Mirage. There were no 802.15.4 transmissions during these measurements. While channel 16 shows large spikes, channel 26 shows none. Because these spikes of 802.11 traffic are independent of 802.15.4 transmissions, they appear to 802.15.4 nodes as independent packet losses. The university testbed also has nearby 802.11b, and therefore observes observes similar external interference. The lake testbed, in contrast, has very little 802.11 interference and so bimodality is independent of the channel used.

Section 3 found Mirage nodes on channel 16 saw fewer links with $\beta > 0.8$ than on channel 26. While channel 16 observes similar RSSI shifts to channel 26, external interference decreases bimodality.

6. OTHER LINK LAYERS

Sections 4 showed that opportune transmissions improve link and end-to-end path quality in the Mirage datasets. None of these results indicate whether high β links occur in other link layers. This section analyzes 802.11b packet traces from recent SIGCOMM publications, one from the Roofnet project at MIT [1] and the other from Charles Reis at the University of Washington [23]. The Roofnet dataset is a large-scale outdoor 802.11b mesh, while the Washington dataset is a small indoor testbed. We measure β using these traces and compute the link as well as end-to-end efficiency improvements.

6.1 Roofnet: Outdoor 802.11b

Figure 16 shows the complementary CDF of β values from the 1Mbps and 11Mbps Roofnet SIGCOMM packet traces. At 11Mbps, about 20% of

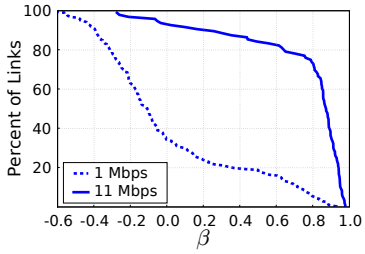

Figure 16: Complementary CDFs of β for links in the 11Mbps and 1Mbps Roofnet data (up and to the right is more bimodal). 11Mbps links are highly bimodal, while 1Mbps links show very low β values. This suggests that 1Mbps data will not benefit from opportune transmissions and that 11Mbps data will.

the links have a β of 0.8 or higher. At 1Mbps, on the other hand, less than 2% of the links have β above 0.8. Furthermore, 40% of the 1Mbps links observe a negative β value, indicating a negative correlation between past and future packet events. While high β values are not unique to 802.15.4, not all networks exhibit them.

To test whether opportune transmissions reduce path cost in Roofnet, we ran the same link quality experiment from Section 4 on the Roofnet data. Figures 17(a) and 17(b) show how opportune transmissions affect link quality. The 11Mbps data, having many links with high β values, benefits from opportune transmissions. The 1Mbps data set, in contrast, sees many links whose quality degrades with opportune transmissions. This is due to negative β values: bursts reduce the probability of delivery.

Figures 18 and 19 show results for end-to-end ETX effects. 1Mbps is consistent with the link improvement results and shows that few paths improve while most degrade due to the decrease of link packet reception ratio. On the other hand, the 11Mbps paths shows average improvement of about 6% with a maximum improvement of 54%: this results are very similar to those in Mirage.

6.2 Washington: Indoor 802.11b

The Washington testbed is interesting because it is located inside a building where there are many sources of interference (e.g people, microwave ovens, and building-wide 802.11). Figure 17(c) shows the link quality improvement for the 8 1Mbps links with an IPI of 450ms. One of the intermediate links experiences a 200% increase in reception ratio (0.2 to 0.6). The small number links and single-hop nature of the experiment precludes computing end-to-end costs or β distributions. However, this data suggests that if more nodes are deployed and each

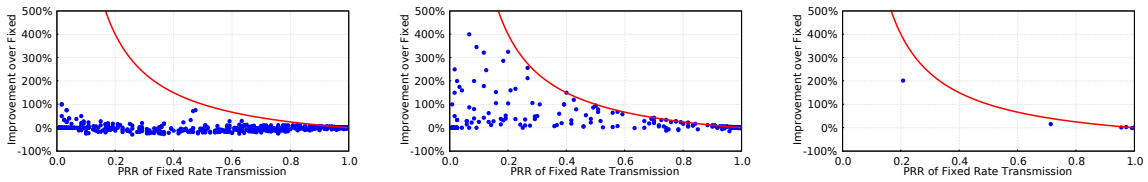

(a) 1Mbps Roofnet (b) 11Mbps Roofnet (c) 1Mbps Washington

Figure 17: While 1Mbps Roofnet sees many links degrade with opportune transmissions, 11Mbps Roofnet and Washington both observe significant link improvements. The Roofnet results are consistent with the low β values in Figure 16.

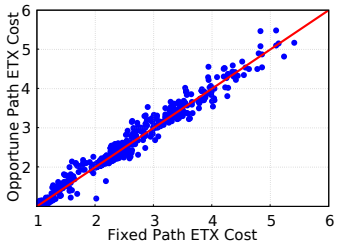

Figure 18: Opportune transmissions do not decrease end-to-end ETX in Roofnet at 1Mbps.

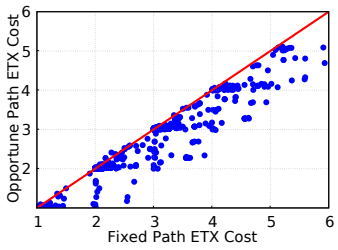

Figure 19: Opportune transmissions decrease end-to-end ETX in Roofnet at 11Mbps. Some routes see 50% reductions in cost.

node were a sender, we would see end-to-end path ETX reduction in the opportune transmissions approach. At 1Mbps, Roofnet sees no benefit from transmitting opportunely, but the 1Mbps Washington testbed has noticeable improvements.

Overall, we can conclude that high β values are not unique to 802.15.4, and therefore opportune transmissions may have broader applicability. However, as not all link layers observe high β values, experimental studies can aide in deciding whether opportune transmissions bring improved results (11Mbps Roofnet, Washington) or hamper the network's efficiency (1Mbps Roofnet). So far, results show that opportune transmission can provide decrease transmission cost in a variety of network scenarios. The next section presents a real-world evaluation.

7. PROTOCOL IMPROVEMENTS

Taking advantage of high-β links can reduce path costs and improve energy efficiency. So far, our results have been based on packet traces. This section examines how a protocol can use opportune transmissions to improve its energy efficiency. Changing a single constant in TinyOS 2.0's Collection Tree Protocol (CTP) [11] implementation reduces its end-to-end delivery costs by 15%.

7.1 CTP

The Collection Tree Protocol (CTP) is the standard TinyOS 2.0 data collection protocol. CTP provides an anycast service to data sinks using a distance vector protocol which chooses the minimum cost (transmission) route. While CTP seeks to provide high reliability using a high retransmission counts and an agile link estimator, it provides neither end-to-end reliability nor flow control.

At the single-hop level, CTP uses a timer to control when it sends packets. The timer institutes a waiting period before the next transmission. CTP uses two timer values to regulate data rates. The first, the success interval, is how long CTP waits after it sends a packet and it receives the link-layer acknowledgement. The second, the no-ack interval, is how long CTP waits after it sends an unacknowledged packet. For the 802.15.4 MAC in TinyOS, both intervals are 16-31ms, which is approximately 1-2 packet times.

To examine whether opportune transmissions can improve CTP's efficiency, we modified CTP's no-ack interval to be a fixed 500ms. We made this change to CTP from TinyOS release 2.0.2.

7.2 Results

To measure how the effect of this change, we ran CTP on 80 nodes in the Intel Mirage testbed at two transmission power levels. Each node generated a data packet every 10 seconds. CTP had a single collection root, located at one corner of the network. Each node sent 128 packets, for a total run time of 21 minutes and 10,240 packets. Because CTP

	-7dBm	-15dBm
Immediate	4.73	6.71
Opportune	4.02	5.65
Reduction	15%	15%

Figure 20: Affects that opportune transmissions have on the average routing cost (transmissions/delivery) for CTP at two transmit power levels. Opportune transmissions reduce the average cost by 15%.

discovers topology quickly, we measured receptions once every node had delivered a packet to the root: in every case the counts begin at the fifth packet.

We measured routing cost by logging every packet reception and transmission to Mirage's wired backchannel. By counting the number of transmissions and dividing by the number of unique packets the root receives, we can measure the average number of transmissions per delivery. We count unique packets because multiple copies of a packet can arrive at the root due to duplication in the network (this happens to approximately 1% of packets).

Figure 20 shows the results. By waiting 500 milliseconds after an unacknowledged packet, CTP's average delivery cost drops by 15%. This is larger than the results in Section 4, which calculated average improvements over an entire network to be small (2-4%). This difference stems from the time scale of the experiments. The trace-based results in Section 4 measured the minimum cost path based on the PRR of an entire trace, while CTP may change its next hop as often as every five data packets. As high β values cause links to come and go, CTP can dynamically take advantage of active good links.

To measure the distribution of path effects, we measured the difference in single-hop and route ETX for each node. A node's single-hop ETX is the number of packets it transmitted divided by the number that were acknowledged [3]. A node's path ETX is the transmission count across the entire network for packets it originated, divided by the number of unique packets received at the collection root.

Figure 21 shows the results. Overall, the link improvements are negligible. While 20% of the −15dBm nodes observe improvements of 10% or more, just as many -7dBm nodes observe a 10% degradation. The average improvement at -15dBm is 1.2%, while at -7dBm the average is a 0.5% degradation.

Despite these anemic link-layer results, both power levels see significant end-to-end improvements. A small number of nodes see large (≈100%) degradations because they have 2-hop, rather than 1-hop, routes. Because the two experiments ran at different times, it is hard to determine if these changes were due to opportune transmissions or the underlying channel. Nevertheless, these outliers are more than made up for by nodes whose routes shorten or become more efficient: the 80th percentile sees route cost reductions of 34-42%. This causes an overall improvement of 15%.

[3]As false positive acknowledgements are < 0.1% of packets and uniformly distributed, we ignore them for simplicity.

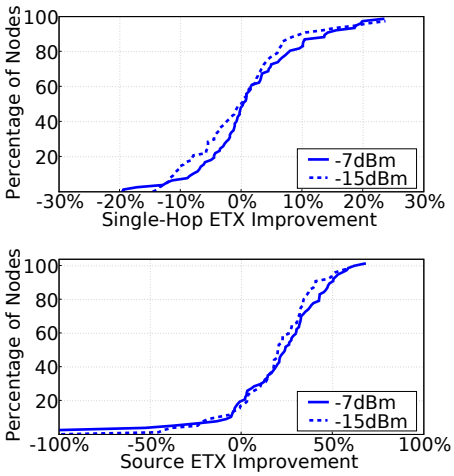

Figure 21: CDF of ETX improvements at the first hop and end-to-end. The average link improvement is ±1%, but the average end-to-end improvement is ≈15%. The large end-to-end degradations reflect a few nodes having different length routes in the two experiments.

For the high-power CTP experiment, the maximum observed end-to-end packet latency when transmitting opportunely was 4 seconds; in the low-power experiment it was 25 seconds. All packets arrived in under 1 second with immediate transmissions.

CTP's route selection explains the seemingly contrary link layer and end-to-end results. While the average link-layer improvement is close to zero, nodes do not have a uniform transmission load. In seeking the minimum-cost path, CTP automatically selects nodes with improved links. Even though opportune transmissions only improve 40% of the links in the network, higher layer protocols then preferentially use these links, leading to significant overall improvements. This result echoes what was observed in Figure 10, where a single link improvement reduced a number of routes by 80% or more.

Because CTP has no rate control or feedback, the bursty pattern of opportune transmissions may cause more more queue overflows. Contrary to intuition, opportune transmissions *reduced* the number of queue drops by 60–77%. Looking at the code, we found that every queue drop was due to a mem-

ory leak in CTP, where it would leak a forwarding packet when it suppressed a duplicate. Pausing after failures causes CTP to detect fewer duplicates, so it leaks fewer packets. Examining the traces, we calculated that without this memory leak neither version would have dropped any packets due to queue overflows. The end-to-end reception ratios in all experiments were 96-97%.

Transmitting opportunely has an additional energy benefit besides reducing the total number of packet transmissions. To operate at a low duty cycle, wireless sensors keep their radio off most of the time, periodically turning it on long enough to hear channel activity. [22]. To wake up a receiver, a transmitter sends a packet long enough for a receiver to sample. As opportune transmissions are more likely to send bursts of packets, they can automatically amortize this wakeup cost over multiple packets. Measuring the possible benefit of this effect is future work.

8. RELATED WORK AND CONCLUSION

Starting from the observation that measurement durations affect link PRR distributions, this paper traces the causes of this behavior and whether it is commonly observed in wireless networks. It proposes transmitting opportunely as one way to take advantage of this behavior. By simply adjusting when it sends a packet, a node can improve link reception ratios by up to 80%, thereby reducing end-to-end delivery costs by 10%, with some nodes observing reductions over 40%.

Early work on wireless sensornet included many experimental studies on the behavior of low-power radios in the 433MHz and 915MHz bands [30, 12]. This work led to two major conclusions that have since become principal considerations in low-power protocol design. First, most node pairs that can communicate have intermediate links [30]; second, link asymmetry is common. While 802.15.4 exhibits both phenomena, Section 2 showed far fewer intermediate links and our prior studies showed that large asymmetries are uncommon [25].

Cerpa et al. found that early sensornet 916MHz radios [6] had much stronger correlations between packet successes than failures, an observation which our 802.15.4 results do not agree with. As Section 2 showed, modern low-power radios exhibit very different behavior than these early platforms and therefore call for different solutions. Some MAC layers resemble opportune transmissions in their behavior. For example, 802.11b's exponential backoff in response to dropped acknowledgements can been seen as a logarithmic search for a good pause interval. But MAC layers operate on time scales that are orders of magnitude smaller than β: neither 802.15.4 nor 802.11b backs off 500ms.

Our observation that wireless links have a strong temporal component is not new: Roofnet and many of the above sensornet studies observed similar behavior. The challenges caused by the combination of correlated loses like those in Section 2, as well as independent losses like those observed on the SNR/PRR knee in Figure 11 , led Noble et al. to propose a suite of link estimators, some of which flip-flop between different estimation time scales [18]. Some sensornet link estimators take similar approaches [28]. There is also work in modeling packet burstiness in wired LANs, for example, with packet trains [16].

While there is a long history of modeling wireless links as n-state or n-stage Markov models [4, 13], the variety and complexity of CPDFs we observe show that these models can capture some, but not all of the important behaviors in wireless networks. Furthermore, we do not go so far as to model communication based on CPDFs. At the very least, we have yet to look into how bimodal shifts across different node pairs are correlated (something the Markov models typically ignore), an area of future work we are very interested in.

This paper differs from all of this prior work in one critical way. Rather than assume that a link has a packet reception ratio, which must be measured, it explores how a protocol can adjust its packet timing in order to influence the reception ratio it observes. The PRR depends on how and when a protocol measures it. As few real protocols send packets at uniform and fixed rates, basing decisions on such measurements is problematic. Instead, data traffic itself should be used to measure links. Furthermore, by controlling data traffic timing through opportune transmissions, protocols can trade off latency for efficiency, boosting their PRR.

As opportune transmissions can introduce large swings in end-to-end latency, they will disrupt some protocols which seek to measure it, such as TCP. This problem is not new: ExOR notes that its packet batches create the same challenge, and advocates edge protocol proxies to smooth the traffic [3]. Using opportune transmissions, however, a protocol can avoid these latency by changing the next hop destination after a single loss; this provides another mechanism to trade off latency and efficiency. A better understanding of the spatial correlation of link modal shifts will enable us to gauge the tradeoff between opportune receptions and opportune transmissions; the effectiveness of ExOR certainly suggests that the correlation is not strong.

These results suggest we should rethink how we think about, model, and design protocols for wireless networks. Protocols do not only need to decide where to send a packet, they also need to decide when to send a packet: the packet loss rate between two nodes is not independent of packet timing. Barring opportune reception schemes, lost packets are wastes of the wireless channel: if the link layer wishes to maximize its throughput, it may want to use opportune transmissions when scheduling packets. Rather than retry a failed packet immediately, a node can send packets with other destinations for a suitable pause interval. While these timing decisions are probably best handled at the MAC layer, network layers may wish to choose different destinations if there may be a large latency. "Cross-layer design" is often praised as an important technique for improving wireless networks: β provides strong evidence that simple information flow on packet timing has significant benefits.

9. REFERENCES

[1] 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.

[2] 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.

[3] S. Biswas and R. Morris. Exor: opportunistic multi-hop routing for wireless networks. In *SIGCOMM '05: Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications*, 2005.

[4] A. Cerpa. Low-power wireless links properties: Modeling and applications in sensor networks. In *PhD thesis, University of California, Los Angeles, California*, 2005.

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

[6] 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*, 2005.

[7] 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.

[8] ChipCon Inc. CC2420 Data Sheet. `http://www.chipcon.com/files/CC2420_Data_Sheet_1_4.pdf`, 2006.

[9] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*, 2003.

[10] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking*, pages 114–128, New York, NY, USA, 2004. ACM Press.

[11] R. Fonseca, O. Gnawali, K. Jamieson, S. K. an d Philip Levis, and A. Woo. TEP 123: Collection Tree Protocol. http://www.tinyos.net/tinyos-2.x/doc/.

[12] 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.

[13] E. N. Gilbert. Capacity of a burst-noise channel. *Bell System Tecnical Journal*, 39:1253–1266, 1960.

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

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

[16] R. Jain and S. Routhier. In *IEEE Journal on Selected Areas in coomunications*, volume 4, September 1986.

[17] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Médard, and J. Crowcroft. Xors in the air: practical wireless network coding. In *SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, 2006.

[18] M. Kim and B. Noble. Mobile network estimation. In *MobiCom '01: Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 298–309, New York, NY, USA, 2001. ACM Press.

[19] C. E. Koksal and H. Balakrishnan. Quality-aware routing metrics for time-varying wireless mesh networks. In *IEEE: Journal on Selected Areas in Communications*, volume 24, pages 1984–1994, 2006.

[20] H. Lee, A. Cerpa, and P. Levis. Improving wireless simulation through noise modeling. In *IPSN '07: Proceedings of the 6th international conference on Information processing in sensor networks*, 2007.

[21] 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.

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

[23] 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.

[24] Y. Rubner, C. Tomasi, and L. J. Guibas. A metric for distributions with applications to image databases. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 59–66, 1998.

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

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

[27] R. Szewczyk, A. Mainwaring, J. Polastre, J. Anderson, and D. Culler. An analysis of a large scale habitat monitoring application. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 214–226, New York, NY, USA, 2004. ACM Press.

[28] A. Woo and D. Culler. Evaluation of efficient link reliability estimators for low-power wireless networks. Technical report, UC-Berkeley, November 2003.

[29] P. Zhang, C. M. Sadler, S. A. Lyon, and M. Martonosi. Hardware design experiences in zebranet. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 227–238, New York, NY, USA, 2004. ACM Press.

[30] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems*, Los Angeles, CA, Nov. 2003.