Some Implications of Low Power Wireless to IP Networking

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

Abstract

We examine and outline challenges in IPv6 routing over low-power wireless personal area networks (PANs). We present empirical measurements and analysis of an increasingly popular PAN link layer, 802.15.4. We show that over short periods 802.15.4 exhibits bimodal connectivity, but over longer periods has many intermediate links. We quantify how synchonous acknowledgments affect common low-power routing metrics, such as ETX. We identify metrics for detecting modal changes in link quality. We explore how these behaviors affect IP routing and IPv6 requirements, such as route selection and maintenance, sub-IP fragmentation and assembly, and packet scheduling.

1 Introduction

Low-power wireless is increasingly important to computer networking. Until recently, network edge devices were for the most part wired servers and desktops. As Moore's Law has pushed the price and form factor of computers down, however, networks have expanded to include large numbers of wireless desktops, laptops, palmtops and cellphones. This trend towards smaller, lower power, and more numerous devices has led to new wireless physical and data-link standards to support them, such as Bluetooth [4], 802.15.4 and 802.15.4b [17], which are designed for short range personal area networks (PANs). Economies of scale may make PAN devices more numerous than any other class of wireless node. In order to maximize lifetime, PAN devices aggressively conserve energy.

Wireless sensor networks (sensornets) are one heavily studied subclass of PANs [8]. Composed of collections of tiny, battery-limited devices with a few kB of RAM, a few MHz of CPU, and sub-1% duty cycles, sensornets impose novel and unique network requirements. Research protocol architectures [10, 7] as well as industrial standards [1] have discarded IP, arguing that it is not suitable due to addressing, network dynamics, discovery, and power. Instead, research protocols have focused on data-centric approaches, while standards such as Zigbee have defined monolithic stacks that stretch from the data-link to the application layer.

Not everyone agrees that IP is inappropriate. The IETF has recently formed a working group – 6lowpan – to standardize how to run IPv6 on low-power PAN protocols [12]. The group believes that the expected number of devices calls for a very large address space, making IPv6 better suited than IPv4. There are many reasons why IP is an attractive solution, including interoperability, a huge library of tools and utilities, and decades of research towards understanding its behavior. History has shown IP to be flexible enough to optimize for many different kinds of networks and usage patterns, working well, or at least well enough, in many domains for which it was never initially intended.

This debate raises two closely related questions. First, how do low-power wireless networks behave? Second, what implications do these behaviors have for IP? The first question has been an important part of sensornet research. Several studies have experimentally quantified low-power wireless radio performance and behavior by exploring the effects of environments, encoding, frequencies, and disambiguating causes of loss [9, 6, 5]. At this point it is clear that low-power wireless protocols have differences from the many media traditionally considered when discussing IP networking, such as bandwidth utilization, energy minimization, and packet sizes. As much of academia and research has dismissed IP, however, there has been very little thought or investigation into the second question, of how these results would affect IP-based networking. Quantifying *how* low-power wireless is different and the implications of those differences is an important first step towards understanding the challenges in bringing IP to these devices.

This paper presents measurements of the long- and short-term behavior of the dominant PAN layer 2 protocol, 802.15.4. It shows ways in which it differs significantly from higher-power protocols in the same spectrum (e.g., 802.11b) as well as the low-power radios measured in early PAN/sensornet studies. It presents some implications of these behaviors to IPv6 networking. Table 1 summarizes the contributions of this paper as its experimental observations and their implications.

2 Background

The IEEE 802.15 working group focuses on wireless PAN protocols. More recently, the 802.15.4 task group was chartered "to investigate a low data rate solution with multi-month to multi-year battery life and very low complexity." 802.15.4 uses a beacon-based scheme, which conserves power by scheduling communication without requiring an association protocol. 802.15.4 uses CSMA for media access. In terms of raw bandwidth per joule, 802.11b is cheaper; what makes 802.15.4 more attractive to PANs is its simpler electronics, which leads to lower cost, faster wakeup, and lower sleep currents.

Two aspects of the 802.15.4 MAC layer are particularly important to IPv6 networking. The first is synchronous layer 2 acknowledgments. When a node sends a unicast packet, it can request an acknowledgment from the receiver, which the receiver sends approximately 180μs later. An acknowledgment packet is 5 bytes long, containing only the format header (2 bytes), a CRC (2 bytes), and the sequence number (1 byte) of the received packet: it contains neither a source nor a destination address. The second is that the maximum 802.15.4 packet size is 127 bytes. The 128th byte is used by the physical layer to denote the size of the packet. This is important because IPv6 requires that data-link layers whose MTU is smaller than 1280 octets provide a sub-IP fragmentation and assembly layer. On one hand, the expectation is that few PAN packets will be large, but this functionality is a requirement for IPv6 interoperability, and the 6lowpan working group has proposed an approach which incorporates header compression and the ability to use short 16-bit node addresses [13]. As PAN devices are energy-constrained, using techniques to increase single-hop delivery rates are valuable, as they can significantly improve end-to-end reliability and therefore reduce the number of network-level retransmissions.

The commonly used 802.15.4 physical layer occupies the same 2.4GHz spectrum as 802.11b. Because of their different data rates, their channels occupy different spectrum widths. Figure 1 shows

Observation	Section	Implications
Over short packet bursts, links qualities are largely bimodal.	Sec. 3	Fragments may be sent in small bursts when a link is good (greedy link select). Need sub-IP acknowledgment scheme to handle fragment flushes.
Low rate traffic encounters intermediate links, which are due to SNR variations or proximity to the reception threshold.	Sec. 4	Routing low utilization traffic requires continuous link estimation or route probing/discovery. The network layer may benefit from physical-layer information such as signal strength and noise measurements.
ETX asymmetries exist and are more common in low rate than burst traffic.	Sec. 5.	Route discovery cannot assume bidirectional communication. Routes require periodic refreshing or probing.
Packet ACK failures are correlated.	Sec. 5	Naive retransmissions waste energy. Need feedback between retransmissions and route selection. Need retransmission and duplicate suppression techniques.

Table 1. Summary of observations and their implications to IPv6 routing.

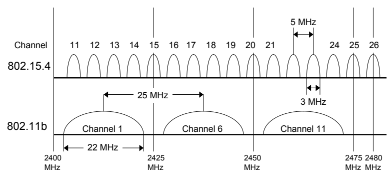

Figure 1. 802.11b and 802.15.4 spectrum utilization.

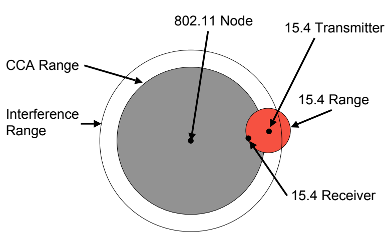

Figure 2. 802.11 hidden terminals are uncommon occurances to 15.4 nodes because of power disparities. For an 802.11 to be a hidden terminal, the 15.4 transmitter must hear a signal weak enough for clear channel assessment, the 15.4 receiver must be within interference range, and the receiver be at the edge of the transmitters range.

their overlap and how 802.15.4 networks can experience interference from 802.11 networks. Unlike 15.4, which is assumed to have low utilization, a copresent 802.11 network might be very busy. While 802.11 might experience interference from 15.4, there is a 100-fold difference in output power: 802.11 chipsets have an output power of 15-23dBm, while 15.4 chipsets are typically -3-0dBm. Figure 2 shows how in practice, the disparity lowers the chances of an 802.11 node being a hidden terminal, as it will only occur if the the signal strength at the 15.4 transmitter is below the clear channel threshold for CSMA, yet the signal strength at the 15.4 receiver is strong enough that it will corrupt the 15.4 packet.

However, as 802.15.4 has such a lower output power and is narrowband interference, 802.11 networks are not likely to respond to their transmissions when performing CSMA. 802.11 packets can be much briefer than 15.4 packets: a 300 byte packet at 11Mbps is approximately 200μs, while a 30 byte packet at 256kbps is 1ms. While most 802.11 data packets are 1500 bytes, acknowledgements and other control traffic often has smaller payloads. Therefore, a 15.4 node can detect a clear channel, start sending a packet, and receive a corrupting burst of mid-packet 802.11 interference.

2.1 Related Work

Experiments with early sensor platforms established that low-power wireless networks have complex dynamics. Ganesan et al [9] analyzed different protocol layers for rene motes, showing that even simple algorithms such as flooding had significant complexity at large scales. They observed that many node pairs had asymmetric packet reception rates (PRRs), which they hypothesized were due

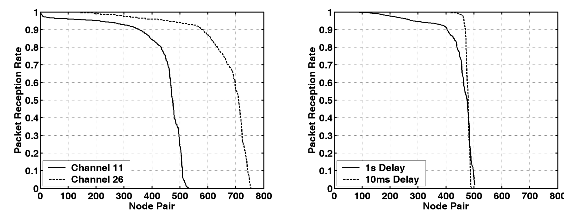

(a) Low-rate round-robin traffic. (b) Short packet bursts on channel 26.

Figure 3. PRR distributions for a 28 node indoor testbed where nodes are on the ceiling. Reception rates are generally bimodal, and the commonality of intermediate links increases with inter-packet delays.

to reception sensitivity differences, which Cerpa et al. [5] supported after swapping asymmetric node pairs and finding that the asymmetries were a product of the nodes and not the environment. While the affects of link asymmetry have been studied in TCP traffic [3] and are applicable to PANs, small packet sizes of low-power wireless and temporal variations raise separate issues to IP networking, which to the best of our knowledge have not yet been addressed.

Cerpal et al. showed that PRR rates can change significantly over time, so that long-term PRR calculation can lead to very inaccurate results [6]. They suggested instead that an instanteous measure of RNP – "required number of packets" – was preferable to a long-term PRR. This work also introduced using conditional probabilities in link estimation, an idea which we extend when considering the correlation between packet failures in Section 5.

Aguayo et al. [2] observed similar packet delivery behaviors in a 38-node 802.11 long haul urban mesh network, but concluded that they were most likely due to multipath effects as there was little correlation between PRR and signal to interference plus noise ratio (SINR). Their experimental methodology differs from those of the sensor network studies. For example, they consider average SINR ratios over second-long periods rather than on a per-packet basis. Nevertheless, the differences in conclusions between the efforts are interesting. Since 802.11b operates in the same ISM band as 802.15.4 and uses a similar modulation scheme (QPSK), its transmitters could be significant sources of interference [18].

3 Distribution of Packet Reception Rates

We start our study by examining the distribution of packet reception rates (PRRs) of 802.15.4 nodes on an indoor testbed. Figures 3(a) and 3(b) show these distributions. Each point corresponds to a single, unidirectional link over which at least one packet was delivered. There were a total of 28 nodes, giving 756 potential links. The figures specify the number of packets over which the reception rate is computed. Reception rates are sorted in descending order and the data for each line comes from a different experiment: each y-value for a given x-value is not expected to be from the same node pair. In the testbed nodes were pinned to the ceiling, people moved freely through the space and there were 802.11 access points.

We measured PRR by having each node transmit 200 broadcasts under two different traffic patterns. In the first, *round-robin*, each node took turns transmitting a single packet, and transmissions were 500 ms apart. With 28 nodes, the inter-packet time for each node was 14 seconds, and for 200 packets the entire experiment took 47 minutes. In the second, *burst*, each node transmitted its 200 packets without interruption. With inter-packet times of 10 and 100 ms, the experiments took 56 seconds and 9 minutes.

Figure 3(a) shows that Channel 11 has 40% fewer high quality ($>$ 90% PRR) links than Channel 26. There are at least three possible explanations for this data. First, 802.15.4 channel 11 shares spectrum with 802.11b while 802.15.4 channel 26 does not. Therefore, 802.11 traffic may interfere with 802.15.4 traffic on channel 11. This explanation, however, is not entirely satisfying. It seems that channel 11 should have a longer right tail since at least a few packets might have been received during the 47 minutes experiment on the 200 or more links seen on channel 26 but absent on channel 11. Second, since our experiments were carried out at different times, it is possible that the RF environment changed appreciably between the two trials. However, this explanation appears unlikely since repeating experiment at different times results in essentially the same distribution of reception rates versus node pairs. Third, the RF circuitry combined with the antenna on the mote may greater attenuate signals on channel 11 than channel 26. This, if true, can increase the communication range of a node and thus increase the number of neighbors for a node in channel 26.

Despite the absolute differences in the distributions shown in Figure 3(a), both curves exhibit similar numbers of intermediate links with reception rates between 10% and 90%. Both channel 11 and channel 26 have approximately 150 intermediate links indicating that over timeframes of about an hour, approximately 20% to 40% of the links had intermediate reception rates. In contrast, Figure 3(b) shows that over the much shorter timeframe of one minute (10 ms delay), packet reception is sharply bimodal. Approximately 85% of links exhibit a 100% reception rate, 10% of links have between 90% and 99% reception, while fewer than 5% of the links have less than 90% reception rate. As the timescale of the experiment is increased by a factor of ten to just over nine minutes (1 s delay), fewer than 20% of the links exhibit a 100% reception rate, 60% of the links exhibit between 90% and 99% reception rate, and 20% of the links exhibit a reception rate below 90%.

Overall, the data indicate that distribution of PRRs in our indoor testbed are largely bimodal. The vast majority of links exhibit either greater than 90% or zero reception rate over short periods of time. The fraction of intermediate links over these timeframes is also small, as indicated by the pronounced knee and sharp fall-off in reception rate shown in Figures 3(a) and 3(b). These observations contrast with the work of Aguayo et al. [2] which showed that in Roofnet – an outdoor 802.11b mesh network – between 50% to 70% of links have intermediate PRRs over a 90 second interval.

Our experiments did not include concurrent transmitters.[1] In the presence of hidden terminals, concurrent transmissions can lower packet reception rate due to collisions at the receiver. However, Aguayo et al. concluded that in their experiments, it seemed unlikely that interfering traffic caused the observed losses [2], so we can factor out foreign traffic as a source of significant differences in both cases. Even with 802.11b-induced interference, the distributions from the two experiments are considerably different.

[1] Therefore our results may exhibit *fewer* intermediate links than a network with concurrent transmitters would. For low-rate PANs or TDMA-based networks, concurrent transmitters are be rare, so our results suggest an upper bound for MAC protocols.

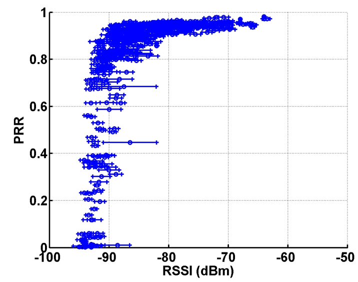

Figure 4. Packet Reception vs. Received Signal Strength for a long-term trace.

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

Table 2. Distribution of the mode of noise readings across 26 802.15.4 nodes.

These results show that over very short time scales, links are for the most part bimodal, having either a PRR of 0% or greater than 95%. As the time scale increases, the chances of link qualities changing increases, leading to larger proportion of intermediate links. Once a node detects a good link, that link is likely to be good for a burst of packets, such as a large IPv6 datagram. If a node has only a single burst to send, as soon as it finds a good link, greedily choosing that link may be a good strategy. In 802.15.4, a "good" link can still have a 5% packet loss rate. Succesfully transmitting a large IPv6 packet (10 fragments) therefore requires a sub-IP acknowledgment layer. Link-layer acknowledgments can provide one part of this mechanism, if a system follows the 6lowpan requirement that an overlapping fragment flush all other fragments, then imperfect duplicate suppression may cause a receiver to flush fragments that were acknowledged at the data link layer.

4 Intermediate Links

The previous section outlined the bimodal nature of 802.15.4 connectivity. Over longer time periods the proportion of intermediate links increases. This section explores the reasons behind those observations and their implications to IPv6.

One of the theories laid out in prior studies of 802.15.4 [16] is that the Signal-to-Noise Ratio (SNR) is the main factor determining packet reception success. We ran long term round-robin (8+ hours) experimental traces across different platforms, in varying environments. The results from all of these experiments were similar to those shown in Figure 4. When the RSSI is greater than some lower bound (-87dBm in this particular experiment), the PRR is high (greater than 80%) with a high likelihood. Otherwise the link falls into a grey area in which the PRR is difficult to predict.

Table 2 shows the varied range of noise floors calculated as the mode of samples of the signal strength indicator (SSI). Note that SSI is not same as RSSI. RSSI is the signal strength of successfully received packets while SSI is the signal strength sampled periodically. SSI is a good measure of the background noise at a node. For the same RSSI, different nodes will see different SNRs due to differences in their noise floors and thus will observe different PRRs.

After further investigation, we observed that not only do the unstable links have average SNRs that are on the edge of the "good link" threshold, but that the RSSI value of packets received from the same node can fluctuate by a few dBm over time.

Figure 5 shows packet reception, noise, and RSSI data for a single node over a round-robin trace. The left graph shows packet reception over time for a single node (node 4) from all the other nodes in the experiment. During periods with a high PRR, the RSSI of the packets received from node 30 are generally -90dBm. The RSSI

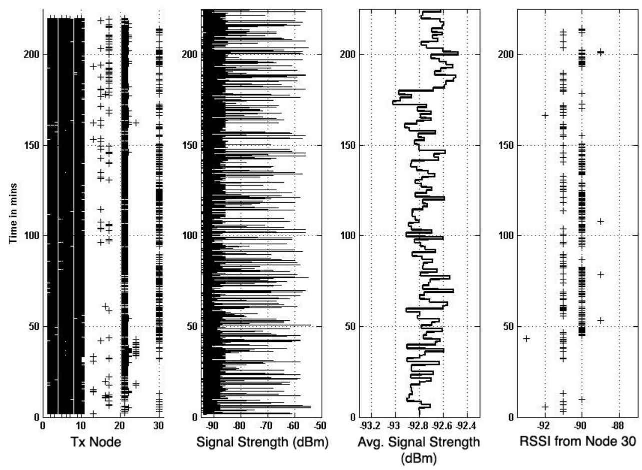

Figure 5. Observed behavior at a single node during a round-robin packet trace. The first plot on the left shows packet loss over time. The second plot shows the measured signal strength of the channel, which shows very short-lived spikes. The third plot shows the noise value averaged over 400 samples (40s), which shows that there are not significant long-term variations. The last plot, on the right, shows the RSSI distribution of packets received over time. The long-term PRR is correlated with RSSI variations.

during the poor link periods, from the packets that were received, are -91dBm or -92dBm, indicating that this slight drop in the received signal strength caused a drop in the PRR. This shows how nodes whose SNR is in the edge of receive sensitivity experience temporal variations.

While these observations shed some light on the behavior of intermediate links, generalizing them to all environments is inappropriate. In these traces, for example, there were little correlation between noise spikes from 802.11b and PRR. This is possibly due to low traffic in the 802.11 network. However, if an 802.11 network is very busy it can cause packet losses, especially as 802.11 might not consider 15.4 traffic a busy channel. This can also lead to long term intermediate PRR links. This case is based on the general hypothesis that the intermediate links are due to variations in SNR. The SNR variations could be when the RSSI (S in SNR) varies close to a stable noise floor, when the noise floor (N in SNR) varies (due to 802.11 and other sources such as microwave) but RSSI is stable, or when RSSI and noise floor both are changing independently.

This hypothesis suggests ways to identify intermediate links based on physical layer information (RSSI and noise floor). An IPv6 router, after identifying possibly intermediate links, may discard them to avoid unstable or time-varying routes. However, avoiding such links may result in a sparsely connected network and cause bottlenecks. Furthermore, discarding these links prevents greedy link selection that can make use of them during periods of good quality. Determining whether a borderline link is good at a particular time, requires either periodic link maintenance or explicit probing. The results in Section 3 suggest that once a link is discovered to be good, it is likely to be good for a packet bursts.

5 Acknowledgements

In this section, we examine the performance of 802.15.4 link-layer acknowledgments and how they affect link quality estimates used for routing. PAN devices often have very limited RAM in order to minimize cost and energy consumption. This constraint makes 802.15.4's synchronous acknowledgments very valuable, as they have a bounded latency and so define how long a retransmission layer must hold onto a packet. However, losing link layer acks (false negatives) leads to unnecessary retransmissions and duplication of a packet within the network. Packet duplication, in turn, requires duplicate suppression techniques, which can increase the complexity of higher layers. In the 6lowpan sub-IP assembly layer, for example, duplicates can lead to flushing received fragments.

Acknowledgements also affect route selection metrics. Existing energy-based metrics such as ETX (the expected transmission count including retranmissions[19]) and its derivatives, such as mETX and ENT [11] use the product of the packet reception rate of the forward and reverse links between a pair of nodes. This approach assumes that the acknowledgment loss rate is the same as the packet loss rate in the reverse direction, and gives each direction of a link an identical ETX. There are two reasons why this assumption may not hold. First, 802.15.4 acknowledgment packets are very small, so for any given bit error rate they are less likely to be lost than a data packet. Second, CSMA causes a data packet transmission to suppress other nodes around it. As acknowledgments are very soon (tens of microseconds) after the data packet, the channel conditions around a transmitter are different than those at an arbitrary receiver.

If the acknowledgment reception rate (ARR) can be significantly different from the reverse PRR, then it is possible that the two directions of a link have different ETX values. In practice, the ETX from A to B (ETX_{AB}) is $\frac{1}{PRR_{AB} \cdot ARR_{BA}}$. In cases of asymmetry, even if two neighbors can communicate reasonably well, optimizing routes for energy might cause directions to differ.

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 seek to minimize ETX, and so choosing very expensive routes is unlikely. Figure 6 plots ETX asymmetries for burst and round-robin traffic on channels 11 and 26. Burst 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). Round-robin traffic observed many more asymmetries. On channel 11, many of these asymmetries were severe, while on channel 26 they were for the most part slight.

There are several possible causes to the larger number of severe asymmetries in channel 11. One possibility is that the signal strength characteristics of channels 11 and 26 are different. This could be due to multipath effects or RF impedance. However, if

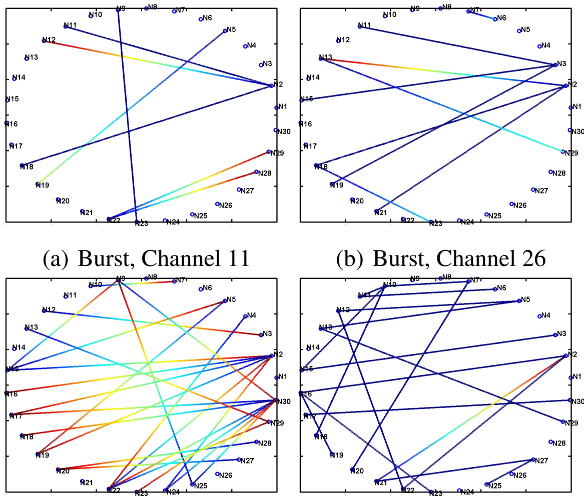

(a) Burst, Channel 11 (b) Burst, Channel 26

(c) Round-Robin, Channel 11 (d) Round-Robin, Channel 26

Figure 6. ETX asymmetries in burst and round-robin traffic for 30 nodes in an indoor testbed. The nodes are in a circle solely for visualization purposes, but nodes close to each other on the circle were close to each other physically. Nodes having asymmetry have a colored line, where the red end of the line is the node that had a higher ETX. A larger gradient indicates higher asymmetry.

this were the case then channel 11 would have many more asymmetric links during bursty traffic, or at least a greater distribution. Another possibility is the fact that channel 11 overlaps with 802.11, while 26 does not. As the data points for ETX measurements occur at different times, it is possible that when one node transmitted there was conflicting 802.11 traffic at the receiver but not vice-versa. This could explain both the larger number of severe asymmetries and the high asymmetry rate observed in round-robin traffic. Regardless of the cause, however, Figure 6 shows that significant ETX asymmetries can exist, they are more pronounced over low-rate than bursty traffic, and channel choice affects the severity.

As ETX asymmetries exist, ARR must differ from PRR. Figure 7 shows the relationship between PRR and ARR. As burst traffic observes predominantly bimodal links, its values are clustered at high reception rates. In contrast, round-robin 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.

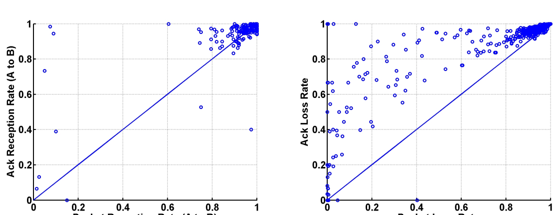

(a) PRR vs ARR, Channel 11 Burst (b) PRR vs ARR, Channel 11 Round-robin

Figure 7. PRR for A→B link vs ARR for A→B. Burst traffic shows bimodal reception rates, while round-robin 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. Differences in intermediate values lead to more ETX asymmetries. Similar plots observed for channel 26 and are not shown for brevity.

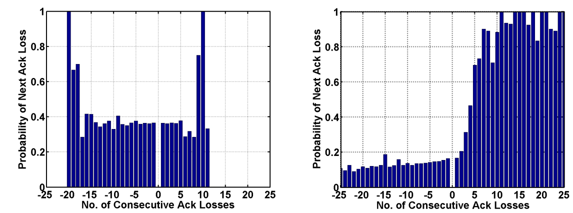

(a) $A \rightarrow B$, Channel 11 (b) $B \rightarrow A$, Channel 11

Figure 8. Conditional probability of a packet not being acknowledged given n consecutive prior failures. The motes sent 100000 packets seperated by 10 msecs to each other in a burst. Negative numbers indicate n consecutive delivery successes. Acknowledgment packet losses are not independent. Channel 26 shows similar behavior but is not shown for brevity.

While greater than the PRR, ARR is usually less than 100%. A protocol that uses link-layer retransmissions may deliver a packet more than once. As routing protocols adjust next hop selections in response to changes in link quality, these duplicates might also be accompanied by route changes, which can lead to multiple copies of an IP packet being in flight.

In Section 3, we showed that over long time periods links can have intermediate PRRs due to transitions between high and low short-term loss rates. An IP routing layer can easily interact with either common case. The difficult case is when a link transitions in the middle of a packet or stream of packets.

Figure 8 shows what transitions look like to a routing layer. It shows the *conditional* probabilities of a successful data transmission and acknowledgement based on prior packets. This plot was generated from 100,000 transmissions between a single node pair with an intermediate loss rate. If failures are independent, then loss probabilities will be constant. Figure 8 shows conditional deliveries for each direction of a single node pair. Figure 8(a) shows failures that follow this pattern. The two edges of 100% loss represent rare cases. For example, there were 0 cases of 10, 3 cases of 11, and one case of 12, leading to values of 100% and 33%. Figure 8(b) shows a very different pattern, where packet losses are not independent: there are two cases, of approximately 10% loss and 80% loss. If a B does not hear acknowledgments from A for several consecutive packets, then the probability of hearing future acknowledgments (whether due to data or ack failure) drops significantly.

The traces from the two experiments show a significant difference which explains these distributions. Approximately halfway through the burst from B to A, packet RSSI values increased for a long period, reaching an average of 5 dBm higher. This increase in RSSI similarly increased the packet delivery rate. The link underwent an RSSI shift, which transitioned it from the low quality to the high quality mode, producing an intermediate link. This is in contrast to the link from A to B, which during its burst happened to be on exactly the edge of receive sensitivity.

Link-level asymmetries preclude broadcast-based route selection techniques, such as those used in AODV [15]. Similarly, ETX asymmetries mean that the two directions of an IP 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. Packet loss correlation suggests that the sub-IP retransmission layer can provide useful feedback to IP route selection, telling it that a link has failed and choosing a different one will save energy.

6 Implications

Our experiments have four major observations.

1. Links are predominantly bimodal for short packet bursts.
2. Sporadic traffic observes intermediate links, which are due to SNR variations.
3. There are ETX asymmetries, which are larger over longer time intervals.
4. Acknowledgement failures are correlated.

The first and second observations indicate that once a node detects a good link, sending IPv6 fragments as quick bursts on that link is effective. The bimodal delivery behavior means that there will be few reassembly failures at the receiver. However, as a link may transition from a good to a bad link during a transmission, a sender needs to maintain all fragments until a single recipient acknowledges all of the fragments.

The second and third observations together indicate that routing small, sproadic IPv6 packets requires different approaches than bursts of traffic, as link qualities may have changed. Continuously probing links (e.g., DSDV [14]) or establishing routes (e.g., AODV [15]) can easily consume more energy than transmitting the data. For latency-sensitive PANs, such as a lighting control system, this cost may be unavoidable. For less stringent PANs, however, such as a lawn monitoring or heating system, nodes can amortize route discovery costs by buffering packets into bursts. Alternatively, with physical layer knowledge a router can choose links with strong signal strengths, which are less likely to have temporal variations. As a single packet is sufficient for detecting a change in signal strength, this is an inexpensive measurement.

The third observation indicates that the two directions of an IP route may need to differ. The first observation implies that if the route is needed for a longer period then periodic rediscoveries may be needed, introducing a tradeoff in the cost of discovery and a route's energy efficiency. Continuously maintaining a routing table (e.g., DSDV) is also problematic, but the first observation implies that the rate at which the bidirectional quality of links need to be probed may consume a lot of energy. A novel routing protocol may combine parts of AODV and DSDV to overcome these challenges. A DSDV-like approach generates a set of candidate links, which are then probed with unicast messages to establish a route using an AODV-like approach, using seperate route requests may be needed for forward and back routes.

The fourth observation indicates that except for a small number of links which happen to be just at the reception sensitivity threshold, acknowledgments are an effective feedback mechanism for higher-layer decisions. A naive retransmission scheme will waste energy when there are several consecutive failures. A more sophisticated scheme that has an estimate of the cost-benefit tradeoff can choose to wait before retransmitting after a suitable number of failures. Alternatively, the link layer can give feedback to the routing layer that there is a latency-efficiency tradeoff, giving an opportunity to choose another link depending on the kind of traffic. Changing links introduces tradeoffs and issues in fragment caching, as a receiver may not be able to distinguish a sender that is waiting due to a period of high loss or has chosen a new destination. Given the energy cost of communication and RAM limitations, these are difficult tradeoffs, and may benefit from packet control bits that indicate what policy the transmitter will follow.

Acknowledgment losses introduce an additional wrinkle in packet assembly. As a node may receive multiple copies of the same fragment, it must have a mechanism for suppressing these duplicates. If the sub-IP fragmentation and assembly layer does not have a cumulative acknowledgment scheme, then failed suppressions can lead to unnecessary packet delivery failures.

7 Conclusion

In this paper we have presented several key observations of low power 802.15.4 nodes. We have shown their implications to IPv6 routing over low power wireless networks. While we have not clearly illustrated what these algorithms and policies have to be, we have shown which of the policies currently used for other IP over wireless networks need modifications. The exact definition of these policies remains an open research topic. However, pointing out the implications of low-power wireless to IPv6 routing is a first step to bringing IPv6 to personal area network nodes, which in the near future will be the most numerous class of networking device.

References

[1] Zigbee Alliance. `http://www.zigbee.org`.

[2] 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.

[3] H. Balakrishnan, V. N. Padmanabhan, and R. H. Katz. The effects of asymmetry on TCP performance. In *Mobile Computing and Networking*, pages 77–89, 1997.

[4] Bluetooth SIG, Inc. `http://www.bluetooth.org`.

[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*, pages 414–425, New York, NY, USA, 2005. ACM Press.

[7] D. Culler, P. Dutta, C. T. Eee, R. Fonseca, J. Hui, P. Levis, J. Polastre, S. S. r, I. Stoica, G. Tolle, , and J. Zhao. Towards a sensor network architecture: Lowering the waistline. In *Tenth Workshop on Hot Topics in Operating Systems (HotOS X)*, 2005.

[8] D. Estrin et al. *Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers*. National Acedemy Press, Washington, DC, USA, 2001.

[9] 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.

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

[11] C. E. Koksal and H. Balakrishnan. Quality-aware routing in timevarying wireless networks. In *To appear in IEEE Journal on Selected Areas of Communication Special Issue on Multi-Hop Wireless Mesh Networks)*.

[12] N. Kushalnagar and G. Montenegro. 6lowpan: Overview, assumptions, problem statement and goals. IETF Internet draft, draft-ietf-6lowpan-problem-05.txt, August 2006.

[13] G. Montenegro and N. Kushalnagar. Transmission of ipv6 packets over ieee 802.15.4 networks. IETF Internet draft, draft-ietf-6lowpan-format-04.txt, August 2006.

[14] C. Perkins and P. Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, pages 234–244, 1994.

[15] 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).

[16] 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 SING-06-00, 2006.

[17] The Institute of Electrical and Electronics Engineers, Inc. Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs), Oct. 2003.

[18] 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.

[19] 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.