The Effect of Link Churn on Wireless Routing

Stephen Dawson-Haggerty
Jorge Ortiz
Xiaofan Fred Jiang
David E. Culler

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2008-109
http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-109.html

August 29, 2008

The Effect of Link Churn on Wireless Routing

Stephen Dawson-Haggerty, Jorge Ortiz, Xiaofan Jiang and David Culler
Computer Science Division
University of California, Berkeley
{stevedh,jortiz,xjiang,culler}@cs.berkeley.edu

Abstract

In this paper, we examine the spatiotemporal dynamics of wireless links and evaluate their effect on routing. We ran several experiments using a testbed consisting of 57 MicaZ motes, and collected data on link behavior over one entire day. We use this data to observe the overall network connectivity over time and space. We are able to examine in detail the choice of neighbors and routes using several link-selection mechanisms, both statically and over time.

We are able to verify the hairy-edge hypothesis, which states that the most important links for routing are the most difficult to predict. In order to do so, we develop precise definitions of important and unpredictable links. We also find it possible to remove these intermediate links from consideration and still have a very rich set of links to route over, while suffering from fewer difficult-to-predict links. Also, we explore the tradeoff between statically defining routes as opposed to a dynamic protocol. We find that while it is not possible to remove all temporal variations from the network, their impact can be significantly reduced through the use of local redundancy. Finally, we present a survey of how several existing routing protocols fit into the framework developed in the body of the paper.

1 Introduction

A fundamental difference between routing for wireless sensor networks and routing in conventional settings is the absence of an *a priori* connectivity graph to define "the network". Instead, logical links between nodes are discovered by a process of observation, termed topology formation[9]. In its simplest form, nodes broadcast a packet and collect responses to build a neighbor table, which is essentially a distributed adjacency list representation of the network graph. Or, in many ad-hoc routing protocols, a flood is initiated from the source to build a tree and upon reaching the destination links are reversed to form a route. Of course, experience has shown that the receiver set for a particular transmitter varies from packet to packet due to fading, attenuation, obstructions, interference, and multipath effect[16, 30, 32]. In addition, responses and retransmissions from neighbors may collide, introducing further variation in the observed connectivity[11].Thus, in most robust wireless sensor network routing protocols topology formation is a continuous process with active or passive probes that feed into an estimator. A threshold determines whether or not a link is present between a pair of nodes at points in time. The connectivity graph may be directed, to express link asymmetry, and may contain weights, to express link quality characteristics such as packet reception rate (PRR) or received signal strength indicators (RSSI). Routing is the process of determining which paths in the connectivity graph are used when forwarding packets from their source to their destination. The routing subgraph may be represented by routing tables, "parent" lists, and the like.

Because topology formation is observational, rather than a fixed property of the network, the underlying link dynamics may have a critical impact on what routing algorithms have to work with. Conversely, the role of a potential link in the routing subgraph has tremendous import on how well its dynamics need to be captured. This inter-relationship of link dynamics and routing is all the more important because routing is fundamentally used to traverse space. The placement of nodes is determined primarily by the physical sense points of interest, rather than the need to achieve a particular network topology, possibly with additional nodes to improve connectivity. Packets are forwarded across a sequence of links because range is limited and because of obstructions and attenuation in the environment.

It has been asserted that the nodes that are most valuable to route through are the ones right on edge of the connectivity threshold and hence are precisely the ones that are hardest to estimate [3]. Closer nodes are likely to have strong, stable links, but they do not take the packet as far toward the destination, so more hops are required. Nodes much farther away are easy to estimate because connectivity is so poor. In the region that is close to the edge, links may come and go due to external factors. However, estimates are based on historical behavior, so transmissions across a link determined to be above threshold may fail, while transmissions may succeed where no link is judged to be present.

While many link estimators and routing protocols have been developed, our goal is to characterize the underlying impact of link dynamics on routing protocol design in practice. We begin the study by trying to established the validity of the "hairy edge" hypothesis described above. This leads to an analysis methodology that allows us to answer a set of basic questions about topology formation and routing. If the link quality threshold is more stringent, are links more or less stable? Are the resulting shortest path routes more or less stable? If we allow suboptimal routes do we increase stability? How much routing redundancy permits local repair, rather than rerouting. To what extent does topology formation need to be a continuous process, or can it be done off-line?

To answer these questions, we first obtain a detailed record of link behavior throughout a large testbed over an extended period. A summary of our findings from a close analysis of this data follows. To greater or lesser extents, all of these conclusions reflect either common sense or the anecdotal evidence of practitioners in the field. However, we believe this is the first study to subject these ideas to a rigorous analysis using a real data set.

- We experimentally measure the size of the "grey" region of the CC2420 radio, and find it to be approximately 25% of the size of the "good" region.
- We propose a single-number metric for wireless sensor networks, the fraction of possible links which are active in the network. Our study is conducted on a 30% network.
- The hairy edge hypothesis, that intermediate links are important for routing, is substantially borne out when the topology formation criteria is very lenient.
- Considering slightly suboptimal routes provides routing with a rich set of reasonable candidates. These routes are only slightly longer then the best possible routes.
- The amount of dynamic activity (or churn) in link tables of routing protocols can be reduced through the use of local redundancy to permit local repair; however, it can not be entirely eliminated.

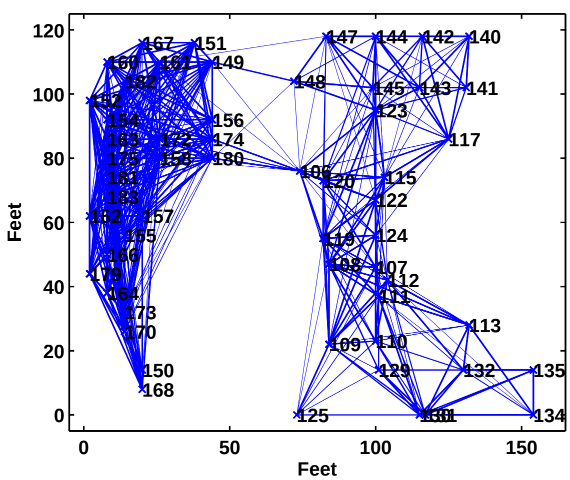

Figure 1. Positions and coarse connectivity of 57 MicaZ motes in the testbed distributed over a 120x150 foot region

2 Experimental Methodology

The experimental methodology we employ to answer these questions is a two stage process, as shown in figure 2. First, we collect detailed traces of connectivity between all pairs of nodes in a large testbed in a realistic setting over an extended period of time. Then we analyze these traces by emulating potential topology formation and routing algorithms upon this body of physical data.

2.1 Connectivity data collection

The connectivity data is obtained using 57 MicaZ motes distributed over a 120x150 feet region of an actively used office building, as shown in Figure 1. The building contains numerous sources of potential interference, including computing devices and multiple WiFi networks. Many of the nodes are attached to the ceiling. They span offices, corridors, open spaces, and laboratories. [1]

Ideally, we want to know the connectivity between every pair of nodes at every instant over a 24 hour period. However, the only way to determine if there is connectivity at a point in time is to send a packet and see whether it is received. Inferring connectivity from a single transmission is problematic because the interference and environmental factors at the instant may not be representative, so multiple such probes are required. If more than one node sends at a time, the observed connectivity is influenced by the simultaneous transmissions.

In light of these concerns, we collect connectivity data in "slices". During each slice, each node transmits a burst of 100 broadcast packets with 20*ms* between consecutive packets. Nodes are scheduled to transmit one at a time in round-robin so that transmissions do not overlap. While a node is not transmitting, it listens and logs any received probes to its flash memory, along with signal strength indicators reported by the CC2420 radio. Once every node has transmitted its packet burst, the network is switched into a download mode where `MultiHopLQI` is enabled and data is retrieved from flash and sent to the base station. This process forms a slice and takes approximately 1/2 hour. Once the download from each node has completed, we repeat the process for the next slice. The results in this paper are developed from running this sample–download cycle for one day, generating 48 slices approximately $1/2$ hour apart and containing approximately $300,000$ measurements. This process gets very close to a snapshot of the network connectivity at any given point in time.

To provide richness of topology and reasonable network diameter the radio power used when sending probe messages is set to -7dBm. Obviously, increasing the power increases the connectivity and reduces the depth, while reducing the power makes connectivity sparser. At this setting, the network is reasonably well connected and the most distant nodes are 5 to 7 hops apart.

Following data collection, the raw traces are filtered and analyzed. First, the traces are parsed and processed into a connectivity graph G, a directed graph where edge weights are the observed packet reception rates (PRRs). This can be done either for a single slice of data, or in aggregate for the entire experiment. To get a basic sense of the network we study, Figure 1 shows the most basic topology formation analysis over the entire lifetime of the connectivity trace. An edge is present if there is at least one packet received in each direction over the life of the trace. This very basic analysis with essentially a zero quality threshold is overlaid on the spatial relationship of the nodes. The region on the left has

[1] Upon publication, our data set will become freely available for download by any interested parties.

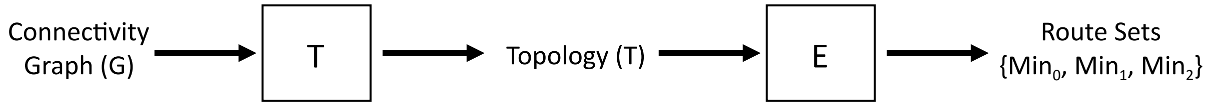

Figure 2. Graph Formation filter graph.

dense connectivity in a large open lab area; the region on the right more uniform connectivity throughout office areas, and there is a braid of a modest number of links connecting these.

In the studies below, this connectivity graph is further filtered, for example by applying a threshold $\mathbb{T}$, where links with $PRR < \mathbb{T}$ are removed, or taking various temporal slices to isolate important aspects of link dynamics and to show how they interact with routing algorithms.

2.2 Link quality distribution

To establish the baseline connectivity of the testbed, consider the PRR on each of the 57x56 potential links over the entire lifetime of the trace. Figure 3 shows a scatter plot of PRR vs. distance for all of these links in the manner of Woo[30] and Zhao[32]. To our knowledge this is the first published record of the 'three regions' behavior for IEEE 802.154 radios with its OQPSK modulation, and the analysis is performed in a realistic setting. Prior studies utilized RFM radios with a proprietary ASK modulation and CC1000 radios with FSK modulation, but which are extremely narrow band, and the nodes were placed in a line. Here we have a more or less arbitrary arrangement of links with a wide variety of distances, orientations, and node pairs.

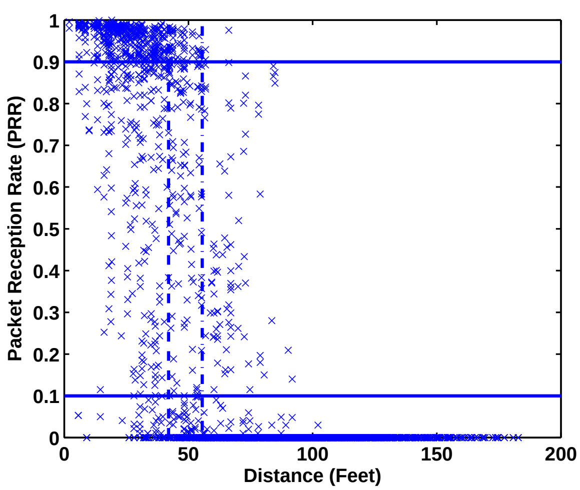

Figure 3. Distance vs. PRR. The observed grey region is smaller then what was previously observed when using the CC1000 radio; it is only 25% the size of the good region.

As expected, we see that the vast majority of short distance links are of high quality and for links beyond a certain distance there is no connectivity whatsoever. As compared to previous studies, we see much more variation in the quality of short distance links. This is consistent with the presence of obstacles, such as walls and equipment, in the actively used indoor setting. If we define the end of the good region to be the distance at which 90% of the candidate links are above 90% PRR, the good region ends at 42 feet. It is indicated by upper-left perpendicular line crossing in Figure 3. Similarly, if we define the start of the bad region to be the distance at which 90% of the candidate links are below 10% PRR, the bad region begins at 55.5 feet. The intermediate "transitional" or "grey" region shows wide variation in link quality, as in previous studies. However, this region is much narrower than in prior studies with their more primitive coding schemes. It is roughly 25% of the width of the good region, where prior results showed a value close to 100%.

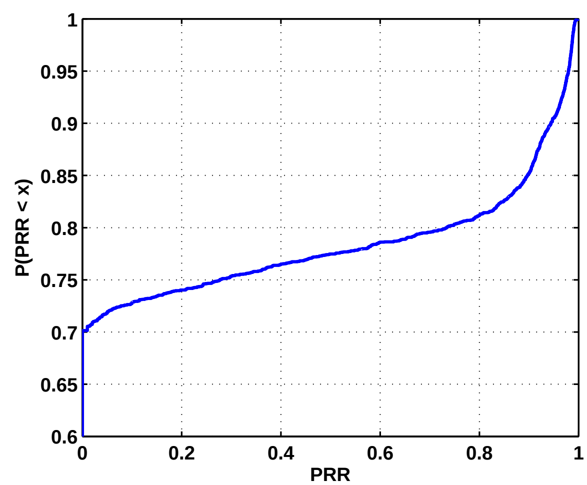

Figure 4. CDF of link quality in our testbed. Only 30% possible links in the network are connected. The set of candidate links in the network is dominated by very good links with PRR greater then 0.85.

The scatter plot in Figure 3 provides a good sense of the variation in behavior with distance, but it is difficult to see how prevalent links from the three regions are. Figure 4 shows the cumulative density function of PRR for links in the network. 70% of the potential links show no connectivity. This figure is essentially a reflection of the physical extent of the network relative to the radio range. If one enlarged the physical extent, adding nodes but retaining the same node density and radio range this value would rise. Similarly, this value would rise at the same factor as node density if the range were reduced, for instance by reducing the transmit power. However, the network becomes disconnected if that is pushed very far. Alternatively, increasing the physical density of the nodes at the same range or increasing the range at the same density decreases this value. Thus, we have useful single number figure of merit for wireless sensor network organization. Our study is conducted on a 30% connected network.

The sharp transition to good links reflects the point where the signal-to-noise ratio at the receiver exceeds the receiver sensitivity. With the rich symbol coding of IEEE 802.15.4 radios, once the signal can be obtained the packets come through clearly, Less than 15% of the potential links fall between PRR = 0.10 and PRR = 0.90. This would appear to be good news for routing protocols. It suggests that the connectivity graph is largely insensitive to the threshold used to determine where a link is present or absent. Even though connectivity is a continuous time-varying property that is influenced by a wide variety of physical effects, the conventional notion of routing by selecting paths in a derived connectivity graph is reasonable. However, caution should be exercised, as we are drawing this conclusion from the PRR over a full day. There may be substantial temporal variation at the timescales that link estimators and routing protocols operate.

2.3 Degree and diameter

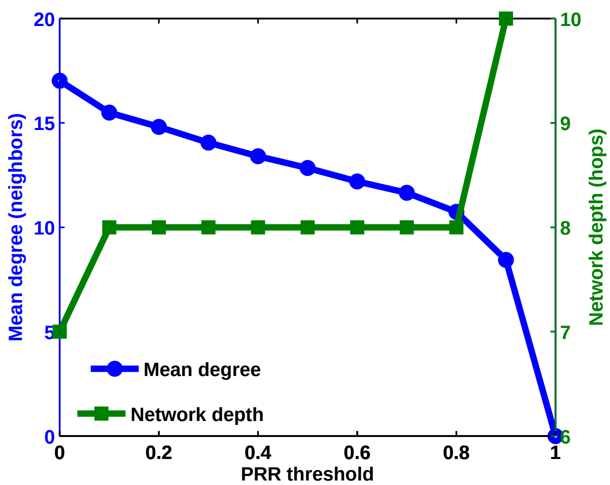

Figure 5. The connectivity of the network varies as the PRR threshold used in topology formation is changed.

Having established the connectivity graph can be derived by applying a threshold to the PRR over a long interval, we can analyze properties of this connectivity graph, such as its degree and its diameter. The degree is essentially the size of the neighbor set while the diameter is a bound on the number of routing hops. Figure 5 shows the effect of the PRR threshold, $\mathbb{T}$, on the mean degree and the maximum depth for the network. If the link quality threshold is extremely lenient, i.e., essentially zero, some very long links appear, substantially reducing the diameter. While one or more of these serendipitous links are likely to show up in a topology-free flooding algorithm, they are too unreliable to use in routing. Fortunately, even a very modest threshold filters them out. The value of the threshold hardly matters until it becomes close to the mean value in the good region, whereupon the graph becomes disconnected and the diameter explodes.

The threshold has a more continuous effect on the mean node degree. As it is increases, the connectivity graph becomes sparser and sparser. As nodes have varying links to various neighbors, some drop out as the threshold increases.

From this simple analysis of overall all connectivity, we have our first important results on the interaction of topology formation and routing. The more stringent the criteria for determining that a link exists between a pair of nodes the sparser the connectivity graph that is available for routing. The threshold choice has little impact on the network diameter, unless it is very extreme, but it is likely to have a substantial impact on the number of choices at each hop, and on the storage complexity of the neighbor table. As the threshold approaches the mean PRR of good links, the connectivity graph collapses and routing becomes impractical.

3 The Hairy Edge Hypothesis

We turn now to a more subtle interaction of variations in link quality and routing. The figures above suggest that link estimation for the purpose of topology formation should be easy. The vast majority of potential links are either nonexistent or of high quality. However, it has been asserted that the minority of intermediate links are likely to be the ones that are most important to use for routing. They are not just of intermediate quality, but they may be highly variable and possible hard to estimate.

To make the "hairy edge hypothesis" more precise, we can formulate the assertion in terms of basic communication theory. The signal strength from a transmitter fades with distance due to increasing spatial coverage, absorption, attenuation, obstructions, and multipath effects. Typically it is neither radially symmetric nor simple along the normal because of antenna anisotropy, device mechanical construction, and variations in the spatial environment. A packet is received, with high probability, if the resulting signal-to-noise ratio at the receiver sufficiently exceeds the receiver sensitivity. However, the noise is also a varying function of space and time, due to varying sources of interference. Therefore, the points in space where the SNR is adequate also move around. So, the amoeboid "cells" of coverage that have been witnessed in empirical studies are not surprising [16].

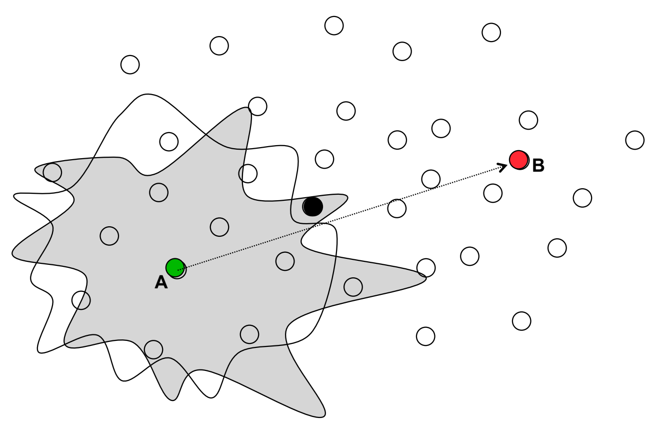

Figure 6. Illustration of the important of nodes at the edge of connectivity in wireless routing

This variation is especially important for wireless routing, as illustrated in Figure 6. If multihop communication is used to route from a source A to a destination B, which node should A choose for the first hop? Nodes that are in close proximity tend to have very high quality links, but they do not move the packer much closer to its destination, so many more hops may be required. Nodes far along towards the destination but outside the range of connectivity are useless. The node that offers a good link and reduces the number of remaining hops is likely to be at the very edge of the cell of connectivity. However, this suggests that the quality of that link may vary with time above and below the selection threshold. In other words, it is likely that topology formation may alternatively accept and reject this important link. In addition, if the selection threshold is raised, a different node on the edge of the new selection criteria will take on the same role. Links that are not in the connectivity graph are more likely to receive packets, but the set of known good links that are also good for routing may be just as unstable, or even more so.

So, in essence the hairy-edge hypothesis says that the links that are most important for routing are the most unpredictable. We can attempt to validate this hypothesis using our dense, sliced connectivity data set. First, we must define precisely what it means to be important for routing. Then we must get a handle on unpredictability. To do this, we look at the population of links and various sub-populations. While the scenario in picture in Figure 6 can be expected to occur, it could possibly be rare because of other constraints in the forming of the shortest path.

3.1 Sets of important links

A link is important for routing if it is likely to be chosen by a routing protocol. Obviously, for routing to nearby nodes the direct link is used, so we are most interested in routing among fairly distant nodes over multiple hops. Most routing protocols seek to find a minimum cost path according to a cost metric. The most common are shortest path algorithms using hop count as the metric. In some cases, the expected number of retransmissions is used as a weight on each of the links. We focus first on min hop routing. With this assumption, in routing from source S to destination D a link is *important* if it is on some shortest path from S to D. For each pair of nodes, this defines a relatively narrow braid of links that are the ones that could be chosen by a shortest path algorithm on the specific connectivity graph.

Specifically, to pick out a set of important links, we choose a set of sources S, and destinations D from the topology graph TG. We then form all paths between the sources and destinations to form the set MIN_0. In some of the analysis to follow, we consider nearly optimal routing protocols as well. To capture the behavior of these algorithms, we may form the set of all links on a path of the minimum length plus one, called (MIN_1). An example of this formation is shown in Figure 7. MIN_0 is the set of solid links, while MIN_1 is the remaining dashed links.

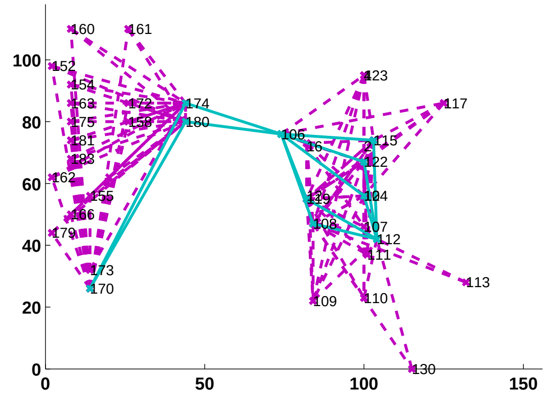

Figure 7. An example of "important" links when routing from node 170 to node 112. Solid lines are links on a shortest, 4-hop path. Dashed lines are on a 5-hop path.

3.2 Link unpredictability

The unpredictability of a link is a more subtle phenomena. Clearly, links that have very high PRR or very low PRR are very predictable, whereas links of intermediate PRR have, by definition, greater opportunity for variation. Still, it may be that a link that is good most of the time has significant epochs where it is bad, and vice versa for poor links. Similarly a link of intermediate quality may behave essentially as a set of independent random events, or it may be highly correlated and predictable. We examine several distinct notions of predictability and look at how links that are important for routing fair compared to the rest of the links.

We begin with a lenient topology formation criteria to define the connectivity graph, a threshold of essentially zero. Figure 8(a) shows in solid lines the probability distribution of links according to the physical distance between the transmitter and receiver for the entire population of links in the connectivity graph. The dashed line shows the corresponding distribution of lengths for a 'important' set constructed by taking MIN_0 for eight sources on the far left of the network and eight destinations on the far right. We see that these populations are indeed quite different. Comparatively, MIN_0 includes few short links and an abundance of fairly long ones. Although the mode is of similar length to the general population, the mean length is considerably greater. (We discuss the third curve in this graph below.)

Figure 8(b) compares the distribution of packet reception rate (PRR) between the population and the links in MIN_0. Indeed, the mean PRR for links in MIN_0, at 62% (dotted vertical line) is considerably less than the average PRR of the population (solid vertical line) at 74detail, we find that a fairly high fraction of low quality links are present in MIN_0 with this lenient criteria. This make sense, since long links are so attractive to shortest path route. Interesting, even with a link threshold near zero, the mean PRR for links used for routing is 62%. Often, very good links are utilized.

While being of intermediate quality is an indicator of variation, we also look at the actual variation in link quality from

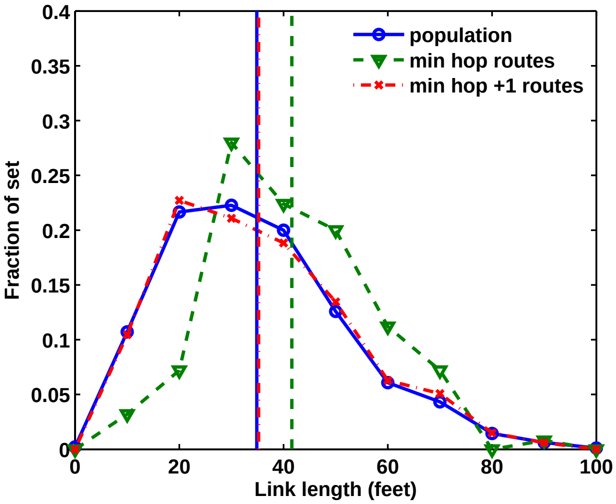

(a) Distribution of link lengths

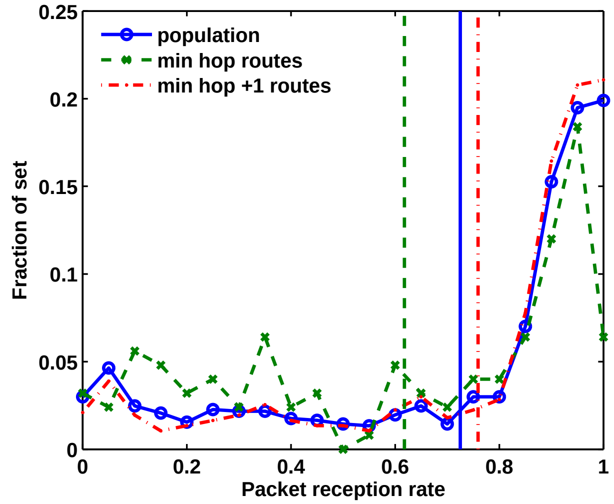

(b) Distribution of link PRR's

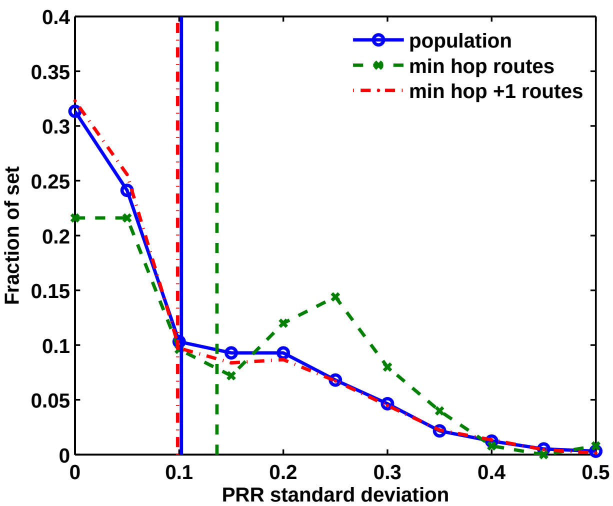

(c) Distribution of link standard deviations

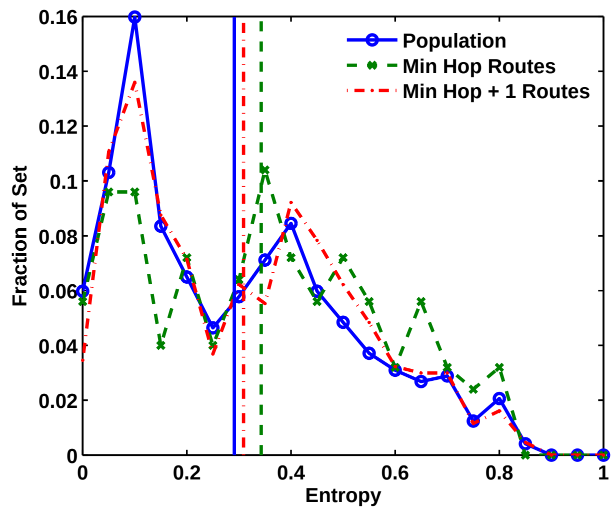

(d) Distribution of link entropies

Figure 8. Four characteristics of important links versus the general population for a lenient topology formation threshold ($\mathbb{T} = 0$).

slice to slice. This is shown for the two populations in Figure 8(c). The standard deviation of links in MIN_0 are 30% higher than the general population.

Perhaps a more fundamental notion of the unpredictability of a link is the entropy observed in the sequence of probes. We define the entropy of a link as follows.

$$H(X) = -\sum_{i=1}^{n} p(x_i) \log_2 p(x_i) \tag{1}$$

Where X is a binary random variable of packet delivery and x_i is a value of the random variable. Figure 8(d) shows that the entropy is, on average, high for minhop links than for the general population.

Of course, entropy and PRR are closely related. Figure 9 shows in a scatter plot the relationship average entropy of each link over the lifetime of our experiment and the average PRR measured for that link. As expected, the most predictable links are those with high PRR and low PRR. However, it is interesting to see that links of intermediate PRR

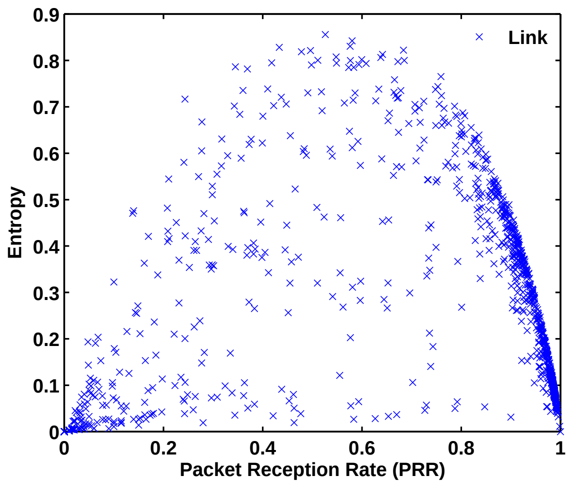

Figure 9. Above is the packet reception rate for each link and its relation to entropy. The higher the entropy, the more unpredictable the sequence (packet delivery success and failures). The most predictable links are the ends (near 0 and 1), while the medium links show the largest values and variance in predictability.

exhibit a wide range of behavior, from very low entropy to very high. Even links of 90% PRR have substantial entropy.

3.3 Stringent link threshold

So the hairy-edge hypothesis appears to be substantially borne out in a realistic network when the topology formation criteria is very lenient. This corresponds to the case where any packet arrival causes a node to be placed in the neighbor table and the associated link considered for routing. Most topology formation algorithms apply link estimation with a threshold to determine whether a link is good enough to be considered for routing. We model this behavior in our analysis by filtering out links with a PRR below the threshold. Figure 10 shows the link distributions after applying a PRR=0.5 filter.

With this more stringent criteria, the lengths of the links used for routing closely track the overall population. In Figure 10(b) we see that all of the low quality links have been eliminated; many of these were long links. Even though the threshold is set at 50%, the mean link PPR in the routing set is 0.9. Thus, the set of important links for routing are actually more reliable, on average than the general population (0.72).

Similarly, in Figure 10(c), the standard deviation of links in MIN_0 is lower than that of links in the general population, presumably as a result of choosing links with higher PRRs that are more predictable. Interestingly, the entropy of links in MIN_0 is somewhat higher than the general population.

Thus, it appears that setting even a modest threshold for link quality in topology formation is sufficient to eliminate most nodes on the edge of the cell. Minimum hop routes can be found that utilize a slightly larger number of quite reliable and quite predictable hops.

3.4 Suboptimal routing

This population perspective on the set of candidate links considered by routing protocols leads to another interesting question. What if the routing protocols considered slightly suboptimal routes, say routes that are only one hop longer than the shortest path? Would this bit of leniency in the protocol provide more robustness and ease of estimation. Does it back us away from the edge to where simple estimators are sufficient and little effort is expended in route repair?

The additional lines in in Figure 8 and Figure 10 show the length, PRR, standard deviations, and Entropy for this broader (MIN_1) set of links that are important for routing. Even with a near zero threshold, this broader set of candidates is similar to the general population in length and variation, and yet is even more reliable on average than either the general population or the strict minhop set. Raising the link threshold makes only a small difference in path lengths, mostly be eliminating the small minority of long, low quality links from consideration.

3.5 Discussion of ETX

Our analysis up to this point has been performed by measuring cost with hops. A natural question is the effect of choosing routes using Expected Transmission Count (ETX) on our conclusions. To address this, we repeated the analysis summarized by figures 8 and 10 using ETX as the metric used in path formation. The topology graph described earlier is converted to an ETX graph where edges are labeled by the number of expected transmissions, and all paths are found between a set of sources and destinations. Since the paths now are weighted on a continuous scale, we pick out the minimum cost path with cost m, and consider sets of links on paths with cost $m+k\varepsilon$, for $k = 1,2,\ldots$

Although we have omitted a lengthly discussion in consideration of space, the results are quite interesting. The result of using ETX with no topology formation threshold ($\mathbb{T}$) results in a distribution of links in a pattern very similar to that shown in figure 10. Thus, ETX functions as both a routing metric and a topology formation criterion, since it will not use links which are of too low a quality. Using ETX, we found that only 1% of the links present in our important sets had PRR < 0.55. This is not surprising, since in order to choose a link with ETX greater then 2, the link must reduce the length of the path by more then 2 hops, and this is rarely possible in this data set.

This is itself an interesting result: for this particular network size, forming paths based on ETX shows results very similar to picking paths by hop count after applying a threshold to the connectivity graph. This may provide additional flexibility to protocol designers since in some cases they may be able to forgo a complicated link estimation approach for a simpler threshold. However, there is some evidence from other studies which indicates that a static threshold leads to poor performance in the face of congestion; our study does not attempt to address the dynamics encountered due to workload.

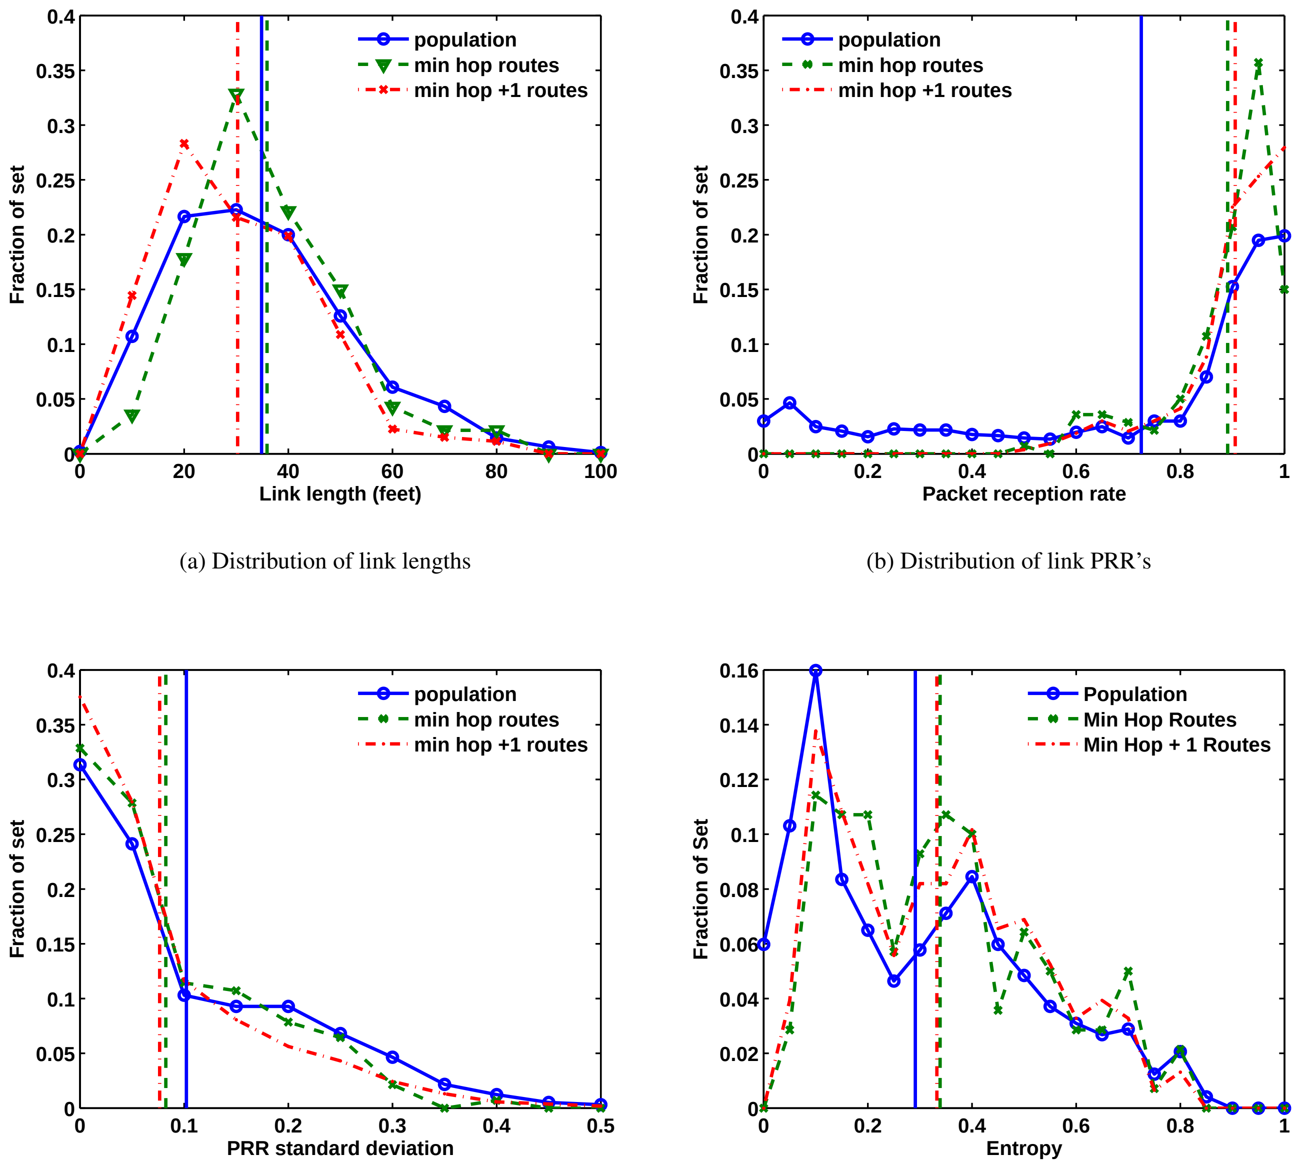

(a) Distribution of link lengths

(b) Distribution of link PRR's

(c) Distribution of link standard deviations

(d) Distribution of link entropies

Figure 10. The links which make up the shortest paths between source and destination have properties which diverge from the population. The best routes avoid links with the highest PRR's, but variances are closer to the population. $\mathbb{T} = 0.5$

3.6 Discussion

The *goodness* of a link with respect to routing reflects an inherent tension. One the one hand, we prefer links that go a long ways towards the destination because the number of hops are few, but long links tend to have low reliability. On the other hand, we prefer links with high reliability and predictability, but these tend to be short. Some routing protocols attempt to capture the trade-offs between these two concerns by accounting for retransmissions in the cost function.

In our analysis, we find that with a lenient topology formation criteria and a stringent routing criteria, shortest path routing indeed tends to pick links that are on the hairy edge. It incorporates long links, even if they are weak, and allows few alternatives. With a more stringent topology formation criteria, the weak links are eliminated and the remaining set is of high quality. The mean link quality is far above the threshold. Considering slightly suboptimal routes dilutes the few long links, providing routing with a rich set of reasonable candidates. The two approaches together appear to be better than either alone.

4 Route Stability

As we have seen in the previous section, a simple PRR filter can dramatically increasing the quality and predictability of links. A natural question becomes: what happens if we increase $\mathbb{T}$ to 0.7 or even 0.9; and to what extend does topology formation criteria affect routing?

Because routes are established on top of topology graphs, the topology formation criteria inevitably plays a big role in the dynamics of routes. In this section, we step back from a microscopic view of individual links to see the behavior of routes over time.

4.1 Route lifetime

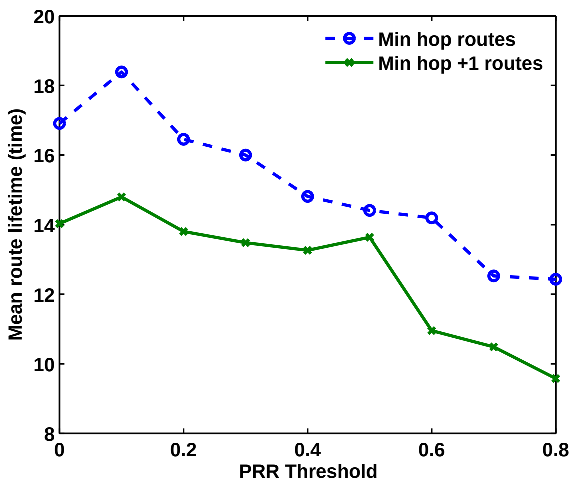

Figure 11. The average amount of time a route between a source and destination will last. "Broken" means having a link which falls under the PRR threshold. Increasing the stretch factor results in generally less reliable routes, since there are more opportunities for links in the path to fail. Also, increasing the the threshold decreases path longevity

Considering link dynamics over the entire day, we observed that even a modest topology formation criteria eliminated much of the link variation. Can we extrapolate that increasing the threshold results in more stable routes? To answer this question we look at the slice-by-slice dynamics. Surprisingly, the answer is no.

One metric of a route's performance is its lifetime. We define route lifetime as the amount of time between the initial formation of the route and the first time any link in the route is dropped from the topology graph. Once some link is no longer in the graph, the route is considered broken. One might expect that as the topology formation criteria gets more stringent, the route lifetime would increase as a result of higher quality links. But to the contrary, as we can see in Figure 11, the average route lifetime actually decreases as we increase the PRR threshold. This is an artifact of our binary topology formation criteria. Using a stringent binary formation criteria such as high PRR threshold, a link may be considered bad at times, and therefore drop out of the topology graph, when it's good a majority of the time. This tells us that sometimes it's better to have a more lenient topology formation criteria to avoid frequent route repair and recovery. Figure 11 also shows that routes using links from MIN_1 have longer lifetime than routes using links from MIN_0. This suggests that allowing more flexibility in route selection also improves route stability. This is understandable since MIN_1 has more links with better quality on average.

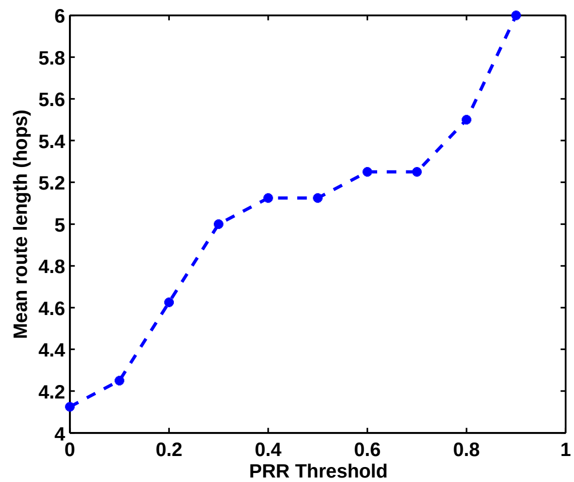

Figure 12. The average path length between a set of 8 senders sending to 8 destinations.

One consequence of the hairy edge hypothesis is that by using more stringent topology formation criteria, we trade off link length for link quality and predictability. Simple geometry indicates that as we become more stringent on link quality and predictability and link length decreases, it will take more hops to connect two points in space. We explore this effect in Figure 12. As we expect, the mean route length does increase as PRR threshold increase. This is another reason to be more lenient on topology formation.

4.2 Route choices over time

One outstanding question is whether it is necessary that topology formation be a dynamic process reacting to changing conditions, or if it is possible to develop a static set of routes which then remain unchanged. Figure 13 shows the size of each important set and how those sets change between each time-slice. The lowest section of the bars correspond to links in MIN_0 and the highest to links in MIN_4. We see strong evidence here for the statement that route choices *do* change over time and *are* affected by our topology formation criteria.

Increasing the topology formation threshold filters out the subset of the population of links in the route sets that fluctuate more often, so the route-sets do not change as frequently. This is seen in figure 13(b). To quantify the amount of variation observed we defined the churn of the route sets as shown in equation 2, where churn is Γ, MIN_ε is the link route-set, and t is time (or time-slice). Table 1 shows the difference in observed changes in each set of route choices.

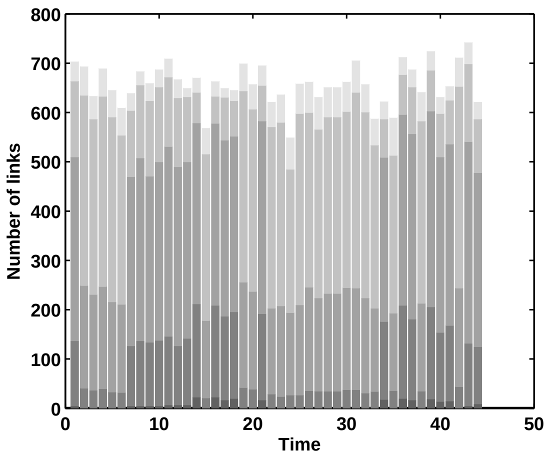

(a) $\mathbb{T} = 0.0$

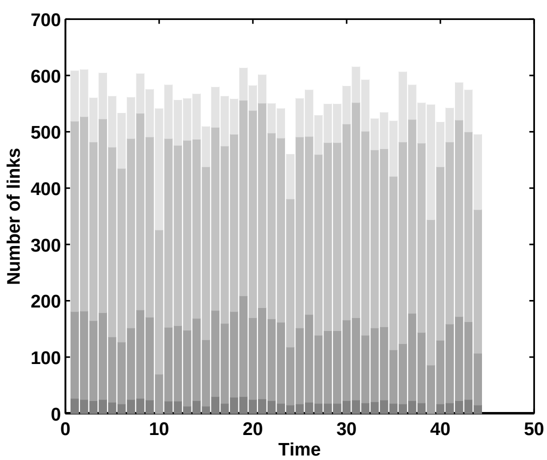

(b) $\mathbb{T} = 0.5$

Figure 13. Change in the size of important sets over time. The bars are stacked so that the lowest piece of bar is $|MIN_0|$, directly below $|MIN_1|$ and so on.

$$\Gamma = \frac{||\Delta MIN_\varepsilon||}{\Delta t} \qquad (2)$$

The key feature of figure 13 is the significant variation in time of the size of route sets. This indicates that even in the case of 13(b) where a threshold results in a significant damping of the dynamic activity, there is still a dynamic component to the network link state, and this variation occurs in a way that will affect routing protocols. This is a strong indication that there must be some dynamic component to route formation, but it is not sufficient.

$\mathbb{T}$	MIN_0	MIN_1	MIN_2	MIN_3
0	6.1	48.4	73.6	104.1
0.50	–	5.9	21.9	26.3

Table 1. Churn of route sets shown in figure 13.

4.3 Route redundancy

We continue our discussion of the need for dynamism in topology formation here by addressing a possible means of stabilizing the shifting topology present in 13. It is plausible that if each node in the network maintains a list of next hops for a path, rather then a single parent, much of the dynamics will be hidden by the redundancy. In order to evaluate this, we construct multitrees in our topology graph. For a given destination, we construct a k-tree with k parent pointers at each node pointing to possible next hops.

Figure 14 shows network connectivity over time using a tree topology. Network connectivity is the fraction of nodes which have some path to the destination within the routing multitree and topology graph. For the single parent case, we see that the higher the topology formation criteria, the higher the average network connectivity over time. A higher selection criteria improves the average links quality, so you would expect links to fluctuate less and routing paths to be up longer.

As expected, even a small amount of redundancy (two parents) shows a nearly 50% improvement in coverage over a single parent, and added redundancy in the form of more parents helps slightly more. This suggests that while there are dynamic events which require a routing response, they are relatively infrequent in a static network utilizing redundancy.

Redundancy is good for improving average network connectivity and robustness and is unaffected by our topology formation criteria. Furthermore, and perhaps more importantly, for all cases, the network coverage is always below 100%. This means that adaptive route repair is still necessary and should not be done offline.

5 Related Work

Routing protocol design in wireless sensor networks has been motivated by the lessons learned from network-layer design in internet research. Early on, OSI-model layering was incorporated to provide a strict layering abstraction between the network layer (layer 3) and the link layer (layer 2). However, unlike internet system, sensornets are heavily resource constrained. Energy consumption is of highest importance for optimization. As such, work in the area began to diverge for the standard networking approach with a rearrangement of the design goals.

Routing design is primarily interested in finding the "best" routes while minimizing energy consumption. Data transmission is an expensive activity [21, 13] so increasing message-delivery efficiency becomes a top priority. Many applications have stringent reliability constraints [21, 18, 26], so maximizing reliability is also of highest priority. Both

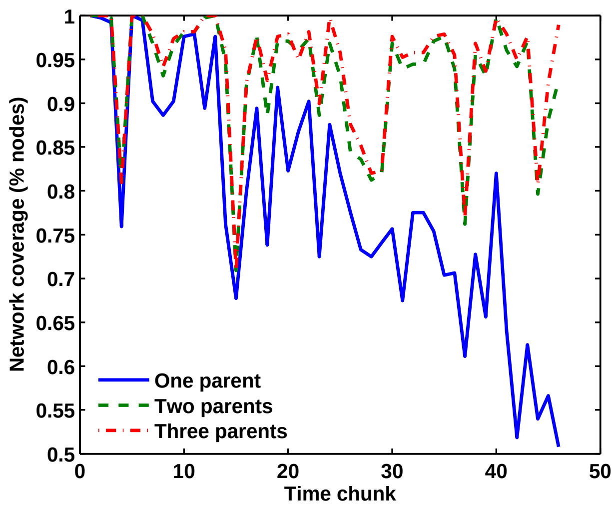

(a) $\mathbb{T} = 0.0$

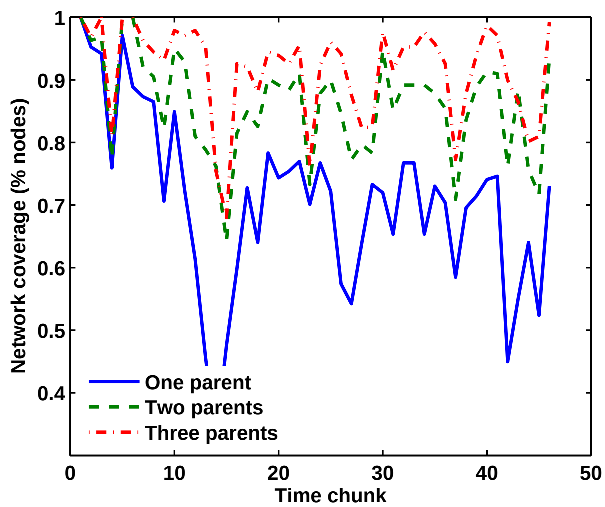

(b) $\mathbb{T} = 0.75$

Figure 14. Local redundancy improves the degree of connectivity and thus reliability in a tree-based protocol.

incur a latency and throughput tradeoff, however, most applications operate at low duty cycle with low sensing and data-send frequency [21, 28], so low latency and high throughput are not as important.

There have been many wireless routing protocols introduced in the literature [7, 8, 15, 17, 22, 24, 2, 20, 23]. Each of them is interested in minimizing the hop-count to the destination in comparison to the minimum hop-count (stretch) and maximizing the stability of route selection. Consequently, each maintains a list of available neighbors and uses a metric for expelling neighbors that do not meet the minimum threshold. The main metrics used are per-link packet reception rate (PRR) or the expected transmission count metric (ETX) [12] along a path. These metrics are agnostic to link dynamics and look to adaptively adjust their link ratings so that the selections can adapt the behavior of the links in the current environment. Our work generalizes the selection criteria by comparing the impact of local decisions (neighbor selection through link filtering) on route behavior on traces taken from a real wireless testbed setting that experiences a wide range in link dynamics.

Several studies have focused on characterizing the dynamics of wireless links in the context of low-power radios [34, 27, 33, 29, 30, 10]. In [34] Zuniga et al. examine the root causes of link asymmetry and formulate an analytical model of link behavior. They also highlight the causes for the transition region, where link behavior varies substantially, and incorporate this observation into the formulation of their model. They compare their model to experiments using the `mica2` [4] mote that uses a `CC1000` [1] radio and show how the model closely fits the results of the data. In our work, we also examine the extent of the transition region, however we used `micaZ` [5] motes which use the `CC2420` radio. We also examine the transition region in a real testbed environment. Prior studies use a linear topology to observe the behavior of the transition region.

The authors of [33] formulate a radio model using both the `CC1000` and `CC2420` called the radio irregularity model (RIM). However, their experimental setting was an empty parking lot with a pair of `mica2` and `micaZ` motes whereas we measured the transition region on a real wireless sensor-net testbed in an office building. The paper also shows how radio irregularity affects the link layer and network layer through experiments using a CSMA-based MACs and several routing protocols. They show that radio irregularity has almost no effect on the link-layer, but that some routing protocols – mainly proactive protocols that try to do routing through localization – are more more susceptible to radio irregularity than reactive protocols like AODV. We are agnostic to the link-layer protocol and distill routing mechanisms to examine their impact on the goals for wireless routing.

This study also includes recommendations to improve routing, which include a link-asymmetry detection component that filters out links that do not have two-way communication. They propose a learning function approach to link estimation that tracks the behavior of links over time and filters out links that are not well-behaved. We also make recommendations to improve routing, however, our recommendations general and can be applied to the parameters in most routing schemes, whereas theirs are more focused on new mechanisms that should exist within the family of protocols they examined.

In the interest of maximizing message-delivery efficiency, routing protocols want to minimize the number of retransmissions. Several studies examine the decoupling of the link-selection mechanism into its own component called the link estimator [25, 14]. In [14] Fonseca et al. propose the use of a cross-layer link estimator that uses four bits of information for 3 layers (physical layer, link layer, and the network layer) to estimate the quality of a link over time. Again, our work is agnostic to the mechanisms in the link-layer and instead focuses on making decisions based on the dynamics observed by the network layer. Links are inherently dynamic and can be very unpredictable regardless of the link-estimation metric, as such, link estimation cannot hide the fundamental behavior of link and our study still stands to improve routing protocol design decisions.

In 802.11 wireless networking link behavior [6] and link estimation has also been examined [31, 19], however since 802.11 devices are usually plugged in to a power source, there has not been as much work in the area of explicit link estimation.

6 Conclusion

We began our study with the goal to characterize the underlying impact of link dynamics on wireless routing and to answer a set of questions, including the validity of the "hairy edge" hypothesis.

Using a connectivity graph built from emperical link reception data on a relatively deep testbed, we are able to obtain a rich set of topology graphs using various formation criterias; and from which we base our analysis of routing algorithm design trade-offs. This methodology allows us to draw conclusions that are relevant to a wide range of wireless routing protocols and are not specific to any particular mechanism.

Using shortest path as a metric for importance and variance and entropy as metrics for unpredicatability, we observe that the "hairy edge" phenomenon is in fact pronounced when the topology formation criteria is lenient. When combined with a stringent routing criteria, shortest path routing indeed tend to pick links that are on the hairy edge. On the other hand, a more stringent topology formation criteria weeds out poor and unpredictable links, leaving a set of high quality links.

However, link churn is not the same as route churn. When we look at the stability of routes (in the form of route lifetimes), a more stringent topology formation criteria actually reduces the route lifetime. A link is often considered bad at times by the stringent selection criteria when it is actually usable. A more strigent topology criteria also results in deletion of long links. This leads to routes with longer hop counts, decreasing the stability of the overall routes. If we relax our route selection criteria, we gain significantly more links. While the inidividual links may not be as well directed, we find that having the extra links available for immediate local repair significantly increase the overall network coverage and robustness.

Many existing routing protocols incorporate mechanisms that perform the topology formation and route selection discussed in this study. Our findings can directly guide the design choices of existing and future protocols. For example, CTP, like many protocols, decouples link estimation from routing. Our results suggest that a lenient link estimator should be used to maintain stable routes. CTP's use of multiple parents is highly recommended, however eight candidates is excessive. Two is sufficient in most environments. CTP's adaptive beaconing will help maintain a fresh topology, which is important when links fluctuate. Currently CTP only uses links on the minhop route, but it may benefit from relaxing this constraint and allowing routing through links in the minhop+1 routes. AODV, on the other hand, performs topology formation by putting all received nodes in the neighor set. It is likely to benefit from using a low threshold filter that takes out the obvious bad links. AODV might also benefit from allowing suboptimal entries in its routing table, such as maintaining multiple entries for a particular destination.

Starting from a comprehensive trace of the dynamics on all links in a large, realistic testbed, we are able to isolate the fundamental interactions between topology formation and routing algorithm. This provides a means of assessing design trade-offs and optimizations for specific algorithms or for entire classes of algorithms in graph theoretic terms.

7 References

[1] Chipcon cc1000. chipcon cc1000 low power radio transceiver. `http://www.chipcon.com`.

[2] Collection tree protocol. `http://www.tinyos.net/tinyos-2.x/doc/html/tep119.html`.

[3] Distributed System Design from a Senor Net Perspective.

[4] Mica2. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/MICA2_Datasheet.pdf`.

[5] Micaz. `http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/MICAz_Datasheet.pdf`.

[6] 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.

[7] S. Basagni, I. Chlamtac, V. R. Syrotiuk, and B. A. Woodward. A distance routing effect algorithm for mobility (dream). In *MOBICOM*, pages 76–84, 1998.

[8] M. Caesar, M. Castro, E. B. Nightingale, G. O'Shea, and A. Rowstron. Virtual ring routing: network routing inspired by dhts. In *SIGCOMM '06: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 351–362, New York, NY, USA, 2006. ACM Press.

[9] A. Cerpa and D. Estrin. Ascent: Adaptive self-configuring sensor networks topologies. *IEEE Trans. Mob. Comput.*, 3(3):272–285, 2004.

[10] 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*, pages 414–425, 2005.

[11] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. *Wireless Networks*, 8(5):481–494, 2002.

[12] D. S. J. D. Couto, D. Aguayo, J. C. Bicket, and R. Mor-

ris. A high-throughput path metric for multi-hop wireless routing. In *MOBICOM*, pages 134–146, 2003.

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

[14] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis. Four-bit wireless link estimation. In *Sixth Workshop on Hot Topics in Networks (HotNets VI)*, 2007.

[15] R. Fonseca, S. Ratnasamy, J. Zhao, , C. T. Ee, D. Culler, S. Shenker, and I. Stoica. Beacon vector routing: Scalable point-to-point routing in wireless sensornets. In *Proceedings of the Second USENIX/ACM Symposium on Network Systems Design and Implementation (NSDI 2005)*, 2005.

[16] 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. *Technical Report UCLA/CSD-TR, 02-0013, Department of Computer Science, University of California at Los Angeles*, 2003.

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

[18] S. Kim, S. Pakzad, D. E. Culler, J. Demmel, G. Fenves, S. Glaser, and M. Turon. Health monitoring of civil infrastructures using wireless sensor networks. In *IPSN*, pages 254–263, 2007.

[19] S. Lee, B. Bhattacharjee, and S. Banerjee. Efficient geographic routing in multihop wireless networks. In *MobiHoc*, pages 230–241, 2005.

[20] P. Levis, N. Patel, D. Culler, and S. Shenker. Trickle: A self-regulating algorithm for code maintenance and propagation in wireless sensor networks. In *First USENIX/ACM Symposium on Network Systems Design and Implementation (NSDI)*, 2004.

[21] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson. Wireless Sensor Networks for Habitat Monitoring. In *Proceedings of the ACM International Workshop on Wireless Sensor Networks and Applications*, Sept. 2002.

[22] Y. Mao, F. Wang, L. Qiu, S. S. Lam, and J. M. Smith. S4: Small state and small stretch routing protocol for large wireless sensor networks. In *NSDI*, 2007.

[23] S. Nath, P. B. Gibbons, S. Seshan, and Z. R. Anderson. Synopsis diffusion for robust aggregation in sensor networks. *TOSN*, 4(2), 2008.

[24] C. E. Perkins and E. M. Belding-Royer. Ad-hoc on-demand distance vector routing. In *WMCSA*, pages 90–100, 1999.

[25] J. Polastre, J. Hui, P. Levis, J. Zhao, D. E. Culler, S. Shenker, and I. Stoica. A unifying link abstraction for wireless sensor networks. In *SenSys*, pages 76–89, 2005.

[26] G. Simon, M. Maroti, A. Ledeczi, G. Balogh, B. Kusy, A. Nadas, G. Pap, J. Sallai, and K. Frampton. Sensor network-based countersniper system. *ACM Sensys*, 2004.

[27] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. Understanding the causes of packet delivery success and failure in dense wireless sensor networks. In *SenSys*, pages 419–420, 2006.

[28] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh. Fidelity and yield in a volcano monitoring sensor network. In *OSDI*, pages 381–396, 2006.

[29] K. Whitehouse, C. Karlof, A. Woo, F. Jiang, and D. E. Culler. The effects of ranging noise on multihop localization: an empirical study. In *IPSN*, pages 73–80, 2005.

[30] A. Woo, T. Tong, and D. E. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *SenSys*, pages 14–27, 2003.

[31] H. Zhang, A. Arora, and P. Sinha. Learn on the fly: Data-driven link estimation and routing in sensor network backbones. In *INFOCOM*, 2006.

[32] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *SenSys*, pages 1–13, 2003.

[33] G. Zhou, T. He, S. Krishnamurthy, and J. A. Stankovic. Models and solutions for radio irregularity in wireless sensor networks. *TOSN*, 2(2):221–262, 2006.

[34] M. Zuniga and B. Krishnamachari. An analysis of unreliability and asymmetry in low-power wireless links. *TOSN*, 3(2):7, 2007.