

Multichannel Reliability Assessment in Real World WSNs

Jorge Ortiz



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2010-65

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-65.html>

May 12, 2010

Copyright © 2010, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

I dedicate this report to my loving wife Aileen Cruz, my parents Judy Zamora and Rudy Ortiz, my wonderful grandmother, Lilia Fajardo, and my sister Patty Rivera. Without their love and full support, I would not be where I am today.

I also want to thank my advisor David Culler, my thesis co-reader Randy Katz, and my lab-mates and colleagues Prabal Dutta, Jay Taneja, Xiafan Jiang, Jaein Jeong, Stephen Dawson-Haggerty, and Daekyeong Moon. All have provided useful guidance and motivation to pursue scientific truth and rigorous exploration that help advance the computer science field of knowledge. Thank you everyone.

Multichannel Reliability Assessment in Real World WSNs

by Jorge Ortiz

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

David E. Culler
Research Advisor

May 11, 2010

* * * * *

Randy H. Katz
Second Reader

May 11, 2010

Multichannel Reliability Assessment in Real World

WSNs

Jorge Ortiz

Computer Science Division

University of California, Berkeley

Berkeley, CA 94707

May 11, 2010

Abstract

We study the utility of dynamic frequency agility in real-world wireless sensor networks. Many view such agility as essential to obtaining adequate reliability in industrial environments. We introduce two facets of connectivity graphs – Multichannel Links (MCLs) and Multichannel Triangles (MCTs) – that identify instances in the network where switching channels may improve reliability. We study, empirically, how frequently MCLs and MCTs occur in live networks and determine whether multihop provides a comparable solution without the complexity of switching channels. We examine connectivity graphs of live networks over each 802.15.4 channel and find that MCLs and MCTs are extremely rare in practice. Almost no MCLs are found in any connectivity graph while MCTs occur between 0-200 parts per million (ppm). Furthermore, we show that MCLs are rarely im-

portant for routing while each MCT has a single-channel routing solution. We also find that there are channels that are always good for connectivity and offer comparable routing costs, with respect to transmission count, in comparison to multichannel communication. Thus, the justification for channel agility in industrial environments applies in the absence but not in the presence of multihop routing.

1 Introduction

Reliability is of great concern for wireless sensing in industrial settings – rooms with lots of metal surfaces and rotating machinery make it a harsh environment for radio-frequency (RF) communication. In [26], Sexton et al. show that RF signals in this environment have a wide dynamic range and vary substantially in stability. Because of the strict reliability requirements of industrial monitoring applications, three standards bodies have formed specifically to address this concern: IEEE 802.15.4e [1], ISA SP100.11a [2], and WirelessHART [9]. Largely based on [26] and studies like it [10, 18], these standards bodies have agreed that frequency diversity is absolutely necessary to provide high levels of reliability in wireless communication.

Beyond the standards groups, there has been much work in the research community to develop multichannel protocols. Many are evaluated in simulation [13, 31, 32] while others have been implemented and evaluated in practice [8, 21, 22]. Each study states various assumptions about the value of multiple channels, with some implicit validation. However, the validation is mostly with respect to network capacity. None of the studies closely examine the contribution of multiple channels with respect to the primary motivation – reliability.

Furthermore, these wireless devices form mesh networks and route over multiple hops. Routing enables communication among devices that cannot communicate directly. Thus, it is

an alternative to frequency diversity even for nodes that are in close proximity. It also provides receiver diversity. In this paper we find frequency diversity is not necessary for high reliability in the presence of routing. In practice, we show that even on a single channel, route diversity – multiple choices for routing at each hop – offers the same level of reliability as the multichannel solution.

To show this we distill the multichannel reliability assumption into two observable graph-theoretic objects that we can explicitly test for on live networks: Multichannel Links (MCLs) and Multichannel Triangles (MCTs). These objects capture locations in the network where channel-switching is necessary for the reliability of communication. After identifying instances of these objects we examine them with respect to routing on a *single channel* to determine if there is also a routing solution. Our results establish the following:

- Although there are many unidirectional links, MCLs – links that are unidirectional or nonexistent on some channel and bidirectional on another – are rare.
- Instances where a multichannel solution is necessary are extremely rare. The graphical facet, an MCT occurs only about 0-200 parts per million (ppm).
- In every instance where a multichannel solution is necessary for communication there is a single-channel routing solution available.

When reliability is a priority, our results suggest that the tradeoff to consider is between protocol complexity or resource use. One can design a multichannel protocol to find an opportunity for successful communication on a different channel or deploy extra sensors to enrich the deployment connectivity graph.

This study focuses on communication *reliability*. That is, the successful delivery of data

between nodes. Reliability is a separate concern from latency and throughput. One can obtain perfect reliability with high latency and low throughput by re-transmitting forever, but we want to achieve efficient reliability. Thus, we examine the trade-offs and associated costs between single and multichannel communication from a reliability stand-point. To the best of our knowledge this is the first study to systematically assess the utility of multichannel operation for reliable communication in the presence of routing. Many have looked at link behavior in isolation to evaluate the implications on reliability. However, we consider reliability in the context of the link and network layers by comparing the gains of both frequency and route (space) diversity.

2 Motivation

Sensornets are deployed in real-world environments, often use batteries as their main power source, and on mote-class devices, the radio consumes the most energy [17]. Communication cost is essential. As background, we describe general approaches to dealing with unpredictability of wireless signal propagation to obtain efficient packet delivery.

2.1 Wireless Propagation

Sensornets are deployed in various types of environments over extended periods of time and they are subject to unpredictable internal and external interference, collisions, physical obstructions, and multipath fading. These factors change over time and are non-uniform throughout the network. Different parts of the network may experience different phenomena that cause the connectivity to vary.

Internal interference may occur when nodes in the same network transmit simultaneously.

Collisions are a form of internal interference – multiple senders, within transmission distance of one another, transmitting simultaneously to the same receiver – and is avoided with Carrier Sense Multiple Access (CSMA). However, hidden terminal problems may occur and CSMA does not solve the problem entirely. Internal interference can be largely avoided by scheduling transmissions in the network (e.g. time-division multiple access (TDMA)).

External interference occurs when devices outside the network generate RF signals that prevent reception. For example, 802.11 shares the same RF frequency range as 802.15.4. When a mote and an 802.11 client transmit on any overlapping frequency, simultaneously, interference may occur. Microwaves, cordless phones and Bluetooth devices also transmit RF signals in the same frequency range and serve as external interferers to 802.15.4 networks. In industrial environments there may be many unintended source of RF interference.

Placement also contributes to loss. When there is Non-line-of-site (NLOS) communication, signals bounce off surfaces in the environment lengthening the propagation distance, weakening its strength by time the signal reaches the receiver. Moreover, reflections can cause destructive interference in certain locations and these locations change with changes in the environment. This is referred to as multipath-induced narrowband fading and has been demonstrated in various studies.

2.2 Diversity Helps

Spatial diversity transforms the propagation problem that causes multipath fading into a feature by using either multiple antennas or various choices for receivers in the network. Multiple-input multiple-output (MIMO) technology uses multiple antennas at the transmitter and receiver to increase the probability of packet reception. Various mechanisms at the sender and receiver,

such as spatial multiplexing and precoding, are used to reduce the effects of multipath fading [19]. Similarly, receiver diversity in the form of route choices at each hop is also used. The latter is examined in this paper.

Frequency diversity is used in various forms at different layers of the communication stack. Modulation techniques, such as direct-sequence spread spectrum (DSSS), are used at the physical layer in order to minimize the effects of noise on a given channel. Radios may also use frequency-hopping spread spectrum (FHSS) to communicate using multiple channels. Most multichannel MACs run on radios using wideband modulation techniques and provide the additional frequency diversity by explicit channel hopping at the link layer.

Diversity helps to hide the wireless communication errors. However, it is not clear how much each level of diversity improves communication reliability, as many of the mechanisms are redundant. Early versions of the 802.11 standard used both FHSS and DSSS. However, studies indicated 802.11 FHSS did not coexist well with other FHSS systems [29] and the standard eventually changed the modulation scheme to only DSSS [20].

2.3 Standardization

Three standards bodies have formed to address some of the issues just discussed (mostly in the context of WSN deployments in industrial settings). In their proposals [1, 9] and standards [2], they introduce many forms of diversity. One is frequency diversity and the other is receiver diversity. Many reasons are given to justify each level of redundancy. With largely overlapping sets of goals, we examine the reasons more closely and analyze the underlying assumptions and related issues.

The goals of the standards efforts are include reliable packet delivery, long deployment

lifetime, adjustable quality-of-service (QoS), and fault tolerance. However, some goals are fundamentally at odds in their extreme, so mechanisms are proposed to provide a compromise. Several claims can be pulled directly from the standards documents and associated presentations.

First, it is stated that “channel hopping [is used] to provide a level of immunity against interference from other RF devices operating in the same band, as well as robustness to mitigate multipath interference effects” [2]. There are three assumptions made here:

1. There is interference on the current channel.
2. The width of the congested band is small.
3. The sender has a load to offer.

If these communication conditions hold then frequency-hopping (FH) *may* help. FH may be of no help if it switches to another congested channel or may even *hurt* performance by switching from a good channel to a congested one. The non-zero cost of switching is wasted when there is nothing to send.

Second, they explicitly state the use of TDMA “to allow a device to access the RF medium without having to wait for other devices” [2] and multihop mesh networking to “support end-to-end network reliability in the face of changing RF and environmental conditions” [2]. TDMA does reduce internal interference by explicitly causing devices to wait. Local scheduling is sufficient to prevent collisions to a common receiver but more global scheduling is required to avoid hidden terminals. Still external interference remains. FH is natural to include with TDMA but is orthogonal and does not come for free. The more sparse the channel usage, the more costly is the join operation. TDMA prevents collisions when two senders want to send

to the same receiver but is still susceptible to hidden and exposed terminal problems, as well as external interference. Multihop mesh networking allows for communication between nodes that are not directly connected for reasons of distance or interference. It is important to note that reliability, when routing over multiple hops is about reachability in the connectivity graph. This is important to consider, as FH fundamentally affects the available connectivity graph, which may also affect overall reliability. Section 5.4 examines the effects of FH on the connectivity graph and link quality.

Third, each standard's goals includes a network lifetime constraint. 802.15.4e states that they wish to obtain "long operational life for battery powered device (> 5 years)" [1]. This has deep implications on protocol efficiency. Given a fixed energy budget and lifetime constraints, we want to maximize the transmit efficiency. How does FH affect the communication efficiency and how does it compare with the single-channel case?

3 Guiding Study

An important study that played a role in the ISA standard, SP100.11a. The authors examine the behavior of wireless links in industrial environments and recommend various forms of diversity – FH being one of them. We revisit the results of the study and argue that the conclusions would have been different had they considered routing as an alternative to FH.

In [26], Sexton et al. measure the multipath delay spread and link characteristics in several industrial facilities and find that 802.15.4 radios may suffer in these types of environments. The CC2420 transmits at 250 kbps with a chip rate of 2000 kChips/sec [4]. Therefore each chip takes 500 nanoseconds to transmit. If the delay spread is greater than 50 nanoseconds it could present a problem for the CC2420, since it has no equalization. The RMS delay spread was

measured between 10-200 nanoseconds and similar results were obtained in [10].

The study also measures link quality. Six motes were placed in an industrial machine room measuring 43 ft x 66 ft. The observed loss rate for every pair of nodes on every channel in a machine room is reproduced in Figure 1. A path ij in the figure represents transmission from node i to node j and the loss rate is presented for each channel. A second study in a compressor house shows similar results.

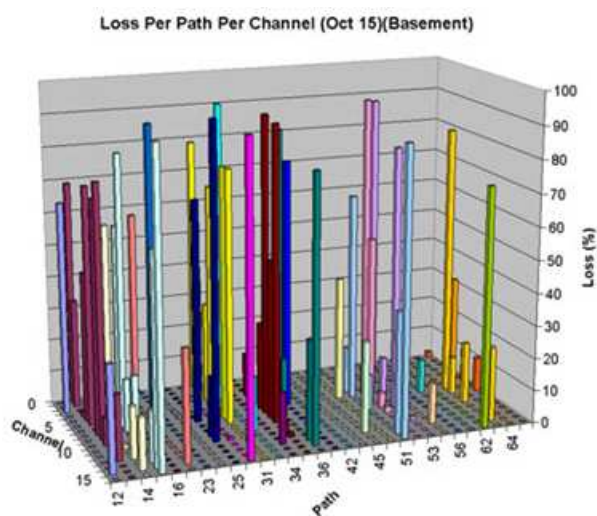


Figure 1: Sexton et al.'s [26] link connectivity measurements in an industrial environment. The Path axis is a link in our terminology (i.e. 12 is the directional link from node 1 to node 2). The height shows the loss rate and the y-axis is the channel.

We can observe that many of the paths are well connected on every channel but several have a high loss rate only on specific channels. Although it's difficult to see in the figure (directly pulled from their paper), this pattern is often asymmetric (i.e., the loss pattern from i to j is different than that from j to i).

3.1 Study's Conclusions

From this data the following conclusions are drawn:

“[In the first experiment] there was no channel that allowed for reliable communications over all paths for all units throughout the entire test period. [In the second experiment], only channel 15 was clear for all paths. None of the paths were very symmetric for all channels. The results of these experiments clearly show that a frequency agile approach might be more robust than a single channel approach...” [26]

We may observe that the study focuses on direct connectivity between all pairs of nodes in the network. In this context, the conclusions are, in fact, sound. However, we do not expect direct links between all nodes to be present in wireless meshes. Communication between widely separated pairs of nodes is accomplished by routing over multiple hops. Sometimes multiple hops are required even for nodes in close physical proximity when direct communication is not present. This leads us to examine the utility of frequency agile approaches in the context of routing, not just in the direct case. To do this we first formalize the basis for their observations and then study them in empirical settings that we can examine in depth.

3.1.1 Link Asymmetry

One observation made in this work is that many links are asymmetric and the conclusion states that “a frequency agile approach might be more robust...”. Observe Figure 2, which shows three links on distinct channels c_1 , c_2 , and c_α between nodes i and j . To formalize the notion of a useful asymmetric link we say a link (i, j) , with a known packet reception rate (PRR), is a *Multichannel Link (MCL)* if there exists distinct channels c_1 and c_2 and a link usability threshold T , such that $\text{PRR}((i, j))_{c_1} \geq T$, $\text{PRR}((j, i))_{c_1} < T$ and there is at least one channel c_α where $\text{PRR}((i, j))_{c_\alpha}$ and $\text{PRR}((j, i))_{c_\alpha}$ are greater than some usability threshold. Link (i, j)

is also an MCL if on some channel c_2 where $PRR((i, j))_{c_2}$ and $PRR((j, i))_{c_2}$ are both below threshold and there is at least one channel c_α where $PRR((i, j))_{c_\alpha}$ and $PRR((j, i))_{c_\alpha}$ are both greater than threshold. In other words, an MCL is a link that is unidirectional or non-existent on some channel and bi-directional on another channel.




Figure 2: An example of a Multichannel Link (MCL). An MCL is a link that is unidirectional or non-existent on some channel and bi-directional on another channel.

Multichannel communication allows the possibility of communicating from i to j on c_1 and j to i on c_2 , hence the two nodes have good bi-directional communication even though they are connected only by two asymmetric links. Multichannel communication also allows the possibility of finding a channel c_α where link (i, j) is bi-directional.

It is important to note that link-level acknowledgements in 802.15.4 utilize the same channel as the packet they cover, so bi-directionality on a single channel is essential for reliability through re-transmission. Therefore, the ability of multichannel communication to construct good bi-directional communication from two unidirectional links is not possible in practice. Only the option of finding a channel with good bi-directional connectivity for this link is a viable option.

3.1.2 Network Connectivity

In a mesh network, we are concerned with the connectivity of the entire network even where there is no connectivity between all pairs of nodes. Many pairs of nodes are connected bi-directionally on a single channel and still frequency agility is required for all nodes to communicate with one another in the absence of routing. An example is shown in Figure 3, where i can communicate bi-directionally with j on c_1 , i can communicate bi-directionally with k on c_2 , and j can communicate bi-directionally with k on c_3 but there is no channel where all three can communicate.




Figure 3: Multichannel Triangle (MCT): $c_1 \neq c_2$ AND $c_1 \neq c_3$ and there is no channel c_α where all three nodes can communication.

A *Multichannel Triangle (MCT)* consists of a 3-tuple of nodes (i, j, k) such that (i, j) share a bi-directional link on channel c_1 , (i, k) share a bi-directional link on channel c_2 , and (j, k) share a bi-directional link on channel c_3 , but $c_1 \neq c_2$ and $c_1 \neq c_3$ and there is no channel where all three can communicate.

Observe that for k to communicate with both i and j in our example requires two transmissions, one on channel c_1 and one on channel c_2 . Also observe that with routing this can also be accomplished by two transmission on c_1 first by k and then by i . Similarly, a routing solution exists on c_2 . Routing solutions are not considered in the analysis of [26].

Note that both MCLs and MCTs are necessary for completeness. MCLs are most important when partitions occur in the network on a single channel. If an MCL exists that bridges the partitions then we have identified a valuable instance where multichannel provides reliability. If the graph is disconnected, reliable packet delivery cannot be obtained. Moreover, MCTs are important for identifying opportunities for general multichannel connectivity, assumed common in real networks. However, MCTs do not cover all opportunities.

We can directly test the existence and prevalence of these graphical facets in real networks by examining their connectivity graphs on each channel. We can also assess their importance by testing whether there exists a single-channel routing solution everywhere an MCT occurs. Furthermore, we can see how often the connectivity graph becomes disconnected and whether there is a multichannel connected graph available in those instances – this identifies the most important MCLs. In the following section we describe our experimental setup and methodology and follow it up with our results.

4 Experimental Setup and Methodology

To examine RF characteristics and connectivity we placed a set of motes in three distinct environments: an industrial machine room environment with many metal and concrete surfaces and several active engines and pumps. We also deployed in a computer room environment with several computer racks, air conditioners, and storage boxes, and on a testbed with motes sitting amongst several 802.11 access points scattered throughout the ceiling of an office environment.

Each setting is subject to various forms of external interference and narrowband fading. The main source of loss in the industrial setting is due to NLOS communication and multipath-induced narrowband fading. We observed some external interference but activity was sporadic

and short-lived. Most of the communication between motes in the computer room and testbed is NLOS. Furthermore, both settings are subject to 802.11 interference from multiple access points and numerous active clients.

The motes used for the deployments in the machine and computer rooms ran with the b6lowpan [15] stack. Motes handle various experiment commands that are delivered over the routing tree constructed by the stack. Data is also collected over the routing tree after each experiment. For the testbed we used the Ethernet back-channel for command delivery and data collection.

4.1 Deployment Setup

For the machine room and computer room we placed motes in locations in the network where sensing might take place. In the machine room, we placed motes on top of moving engines and between pipes. The machine room deployment is separated into two separate rooms that are side by side. Both rooms have similar equipment and are separated by a wall. We placed motes in both rooms and tested connectivity among all the motes in the deployment. Similarly, we placed motes inside computer racks, next to active air conditioning units, and at varying heights inside the computer room.

On the testbed we had no choice in node placement. Nodes are scattered on the ceiling across an entire floor in an office environment. The motes sit amongst active 802.11 access points that see a varying number of clients and activity throughout the day. Since the floor is partitioned into several rooms, many of the links are NLOS.

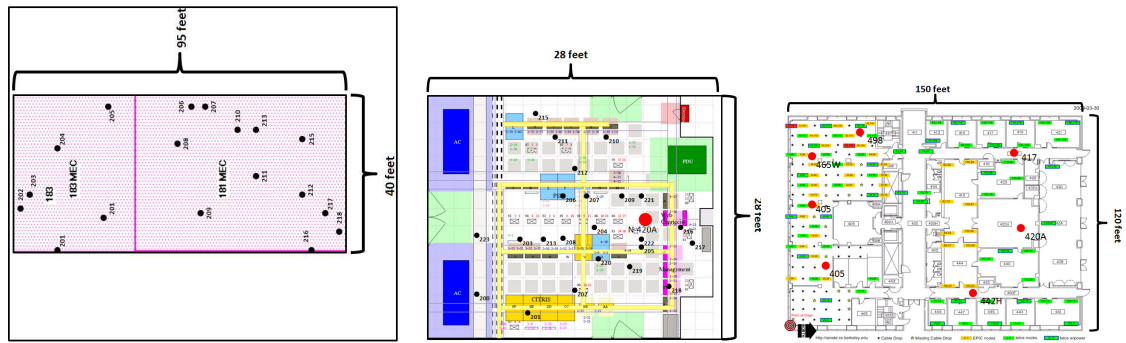
4.1.1 Industrial Machine Room Setup

Figure 4 shows a picture of the industrial setting and figure 5(a) shows the placement of nodes in the setting. This environment consists of two rooms on a single floor that cover a 95 ft x 40 ft area. It is similar to those described in Sexton's study. There are several moving engines and pipes with material flowing through them. There are also lots of metal surfaces and concrete floors; the type of environment where RF signals scatter and may be more prone multipath-induced narrowband fading. We used 20 TelosB [6] which uses the CC2420 radio.



Figure 4: Machine room setting.

We use the WiSpy spectrum analyzer [7] to characterize the RF environment and find some RF noise generated near engines in the room. RF activity is sporadic and spread throughout the frequency band however we do not examine this data very closely, since it does not seem to affect our results.



(a) Machine room node placement map. (b) Computer room node placement. The larger dot on the map denotes an 802.11 access point. (c) Testbed node placement. The larger dots on the map denote the placement of 802.11 access points.

Figure 5:

4.1.2 Computer Room Setup

Like the machine room, nodes were placed in locations where sensing might take place. For example, nodes were placed inside racks, near air conditioners, at the extremities of the room and at different heights. The room is 28 ft x 28 ft and we placed 23 telosb nodes throughout the room.

The computer room contains two rows of racks and three rows of metal storage shelves. It also contains a pair of running air conditioners and storage boxes filled with computer equipment and lots of wiring in and above the racks. Most of the communication in this environment is NLOS. Figure 5(b) shows the node placement overlaid on the map of the room.

There is an access point sitting on the ceiling in the middle of the computer room. We ran the experiments in the evening during the spring break session, so the RF activity was fairly low for each run.

4.1.3 Testbed Setup

On the testbed, motes are placed to allow for full network connectivity. This is the largest of the 3 deployments we examined, with 55-60 MicaZ [3] motes. Floor dimensions are 128 ft x 128 ft and it is partitioned into multiple rooms in an office environment. Since the testbed is inside the computer science building of the university it sees lots of human traffic and 802.11 activity. There are 7 access points sitting among the motes in the network and connectivity between the motes varies throughout the day.

802.11 interference is a potentially significant factor in the results of our experiment in this environment. We closely examined the RF environment using both the WiSpy Spectrum analyzer and the Cisco Wireless Control System (WCS) [5] which gives us direct measurements from the access points in the building. We found that the peak time of activity is during afternoon hours and that the radio utilization averages between 20-60% over all 7 access points.

4.2 802.11 interference characteristics

Although 802.11 interference affected our results our network remained connected in the majority of our snapshot samples. In this section we take a close look at the observed interference pattern in the office environment. 802.11 played a minimal role in the machine room and the computer room. The machine room is located in a remote area of the building with almost no human traffic or offices. The computer room experiments were run during a low-traffic period (spring break) at the university.

Figure 6 shows the worst 15-minute period on the worst channel throughout our experiments in the office environment. The red line is set at -77 dBm, the maximum noise level considered by the CSMA protocol on the CC2420 for clear channel assessment. In the figure,




Figure 6: This graphs shows the nosiest 15-minute period in the testbed environment. The channel with the worst noise occurred on 802.15.4 channel 17.

the noise is above this level only 21% of time. This implies that there is an opportunity to successfully receive 79% of the time. Even during the worst interference period on the worst channel there is ample opportunity for successful transmission.

Although the environment is 802.11-rich, overall channel utilization, even on the worst channel, is relatively low. If the application is highly duty-cycled and real-time delivery constraints are not a concern, buffering data during busy periods may improve reliability even further. In our experiments we did not buffer data for longer than the default CSMA-backoff period, therefore our results reflect an un-optimized approach to handling 802.11 interference.

Furthermore, a vast majority of interference periods lasted for 10 seconds or less. Figure 7 shows the distribution of transmission times on the worst channel over an 1-hour period. The distribution is similar over larger timescales. This further corroborates the data observed in figure 6. Although there are 4 different access points within range of the our WiSpy receiver, utilization of even the most congested channel is relatively low and sporadic. To see the effect this has on a specific workload, lets examine the case where we have a sense-and-send application running at 1% duty cycle. If the transmitter wish to send during the beginning of an heavy interference period, it would accumulate approximately 48, 60-byte packets in its buffer before there is an opportunity to send. Moreover, the average inter-transmission period (excluding

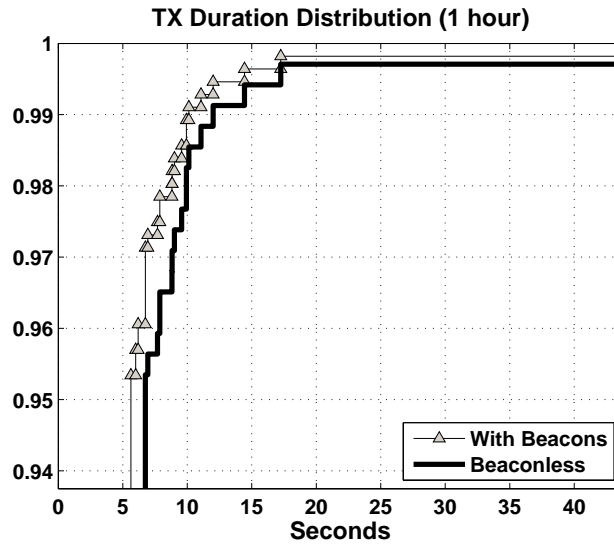
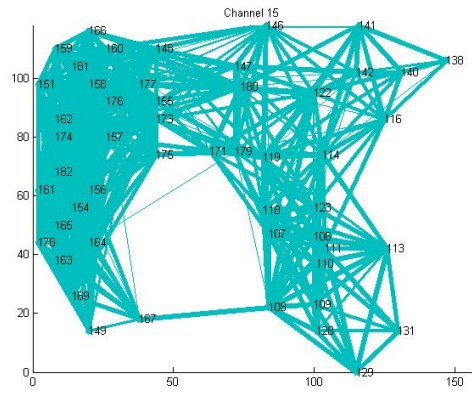


Figure 7: Inter-transmission time on worst channel (17) over 1 hour period. The channel with the worst noise occurred on 802.15.4 channel 17.

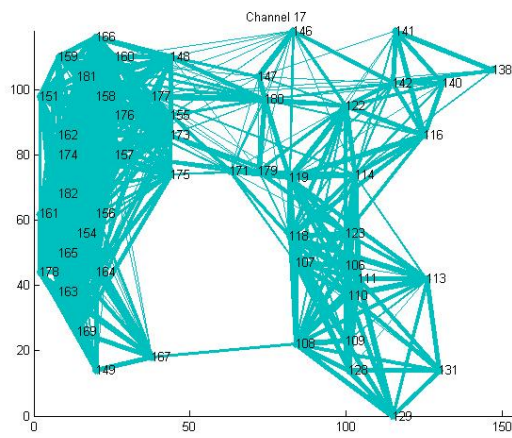
beacons) is about 6 seconds, more than enough time to drain the 48-packet buffer as 250 kbps. Although the workload dictates the resource consumption, this analysis presents a real-world near worst-case example of the resource in a realistic interference environment.

The data presented here only begins to address some of the assumptions made by the standards bodies. We can see that there is sometimes interference on a channel, but that the interference does not last for extended periods of time. This ultimately had little effect on the connectivity graphs we observed. Figure 8 shows two snapshots of the connectivity graph on channel 15, which does not overlap with any 802.11 channels and channel 17, which lies directly within the range of 802.11 channel 6. These snapshot was taken in the afternoon – the highest network-traffic period. Observe that although there are differences in link quality, as represented by the thickness of the lines that represent links, the overall connectivity remains rich.

Since most transmissions last for short time interval, interference did not play a major role in



(a) Connectivity graph on channel 15.



(b) Connectivity graph on channel 17.

Figure 8: Connectivity graphs observed on channel 15, which does not overlap with any 802.11 frequency and channel 17, which sit directly within the 802.11 channel 6 frequency band.

overall connectivity and reliability. Although this is only a single interval in time in a particular environment, our data suggests that this is the actually the common case.

4.3 Channel Probing

In each experiment we use a set of motes, each with CC2420 low-power 802.15.4 radio. During each run, each mote sends 100 broadcast packets with a 20 millisecond inter-packet interval. Only a single node sends at a time while the rest of the motes log all received packets to local flash memory. Once a mote is done sending, another mote starts and this continues until all

motes send 100 packets. Then each mote switches to the next channel and the process repeats. When every mote sends 100 packets on every channel (11-26), the data is collected from each mote for processing. In each deployment we run this experiment several times. We ran the experiment in the machine room twice, the computer room three times and the testbed 17 times continuously over the span of a week.

Wireless link dynamics vary widely over time due to various factors. Broadly, link behavior affects the structure of the underlying connectivity graph and our methodology allows us to capture many snapshots of the graph to examine communication opportunities under varying degrees of duress. In addition, link behavior is fundamentally statistical and taking many samples increases our confidence on our observations.

4.3.1 Analytical Approach

Each broadcast packet contains a sender ID and a local sequence number. When a node receives a broadcast packet it extracts both of these values and logs them along with its own ID and current channel. The testbed experiment also logs timestamp and received signal-strength indicator (RSSI) values. Using this information we separate the data set into bins separated by channel and use each subset to study the connectivity graph on each channel.

For each directional link in the connectivity graph we calculate the packet reception rate (PRR) and set a threshold on link quality to construct the connectivity graph. We then run the MCL and MCT locaters on the traces as well as Dijkstra's shortest-path algorithm. In calculating the cost of transmission over a link, we use the expected transmission count (ETX) metric and compute the sum for all links along a path to determine its cost. Equation 1 shows how to calculate ETX; l_f is the forward link PRR and l_b is the backward link PRR.

$$ETX = \frac{1}{l_f * l_b} \quad (1)$$

After the construction of a graph for each channel we count the number MCLs. For each link on a particular channel, we search for the same link that exists in only one direction on any channel. If the link is either not found or it exists unidirectionally on some channel, it is added to the set of MCLs. This set enumerates the number of opportunities there are for multichannel to enable communication between a pair of nodes. In addition, we ran connectivity tests for all the observed connectivity graphs using Tarjan's connectivity algorithm.

In the search for MCTs we create various sets that consist of 3-tuples of unique nodes in the network that share bi-directional links between each other. The first set, the *single-channel set*, takes every set of three nodes in the network that are bi-directionally connected on a *single* channel. For example, if there exists bi-directional links (i, j) , (i, k) , and (j, k) for unique nodes i , j , and k and each link is on the same channel, then it is included in the single-channel triangle set, also referred to as set S .

We also construct another set, similar to set S , except that the constraints on the links are loosened to include triangles that occur across channels. Therefore, if there exists bi-directional links (i, j) , (i, k) , and (j, k) for unique nodes i , j , and k on *any* channel, then it is included in the global multichannel triangle set, hereafter referred to as set M .

Finally we construct the set of interest, the *MCT set*. This set includes all element in set M that are not in set S . In other words, it includes all sets of 3 nodes that are not connected on a single channel but are connected on multiple channels – the definition of an MCT. We refer to this set as set M_u (unique multichannel triangles).


4.3.2 How representative are samples?

Commands from the experimental driver are sometimes lost on the path to a mote, so retries are used. During the testbed experiments, retries sometimes fail, at which point the node is removed from the experiment and not considered in the analysis. This is why the testbed data contained 55-60 motes per experiment. Furthermore, the untethered versions of this experiment has to wait for the underlying routing structure to converge before sending commands on each channel. The warm-up time varies from channel to channel and can last between 0.5 and 15 minutes.


By sampling with a constant inter-packet interval we may raise the concern that our data is statistically biased because of aliasing. However, during each experiment, there are various random stalls and retries by the experimental driver that de-synchronizes the samples. Furthermore, we obtained many samples in each environment and note the similar trends in each run. Therefore, this sampling methodology is sufficient to capture the underlying connectivity of the network.

In capturing directional properties of links it may be worrisome that each link sample (100 packets) may be separated by at $2N + \epsilon$ sample times, where N is the number of nodes in the network and ϵ is the a small random wait time caused by stalls and retries. This kind of separation between samples may statistically de-correlate the measurements in both link directions. However, by taking many samples we can bound of the error for the PRR measured in each direction of a bi-directional link. Furthermore, the main sources of uncertainty – external interference and changes in the environment – are effectively removed from the two experiments with the smallest sample sets. The industrial environment and machine room had no random movement or significant RF interference and there were many samples collected from


the testbed.



(a) Links found in the industrial setting.



(b) Links found in the computer room.



(c) Links found on the tested.

Figure 9: Properties found in these links match those found in each industrial setting examined in [26]. Note the similarities between these graphs and figure 1. We observe much wider bands of fading links.

5 Experimental Results

In this section, we examine the prevalence and importance of MCLs and find that although bi-directional links are common, MCLs are not. Neither sets are proportionally large enough to affect multihop reachability between the nodes in the network. Furthermore, we observe the existence and prevalence of MCTs and find that they are extremely rare in practice and that

each instance of an MCT has a routing solution.

Figure 9 shows the equivalent of the connectivity measurements in [26] for our test environments. Since we have many more nodes in our networks, in the figure we take a random sample of 6 nodes and show the loss rate between all pairs in that set on each of the channels. We do see that on some links in a particular direction, connectivity is nearly perfect on some channels while there is almost no connectivity on others. The band of this fading is much less 'narrow' than in the Sexton study. In general, there is less connectivity between our nodes. This does imply that the topology of connectivity seen on one channel may be very different from that on other channels, which is likely to have a serious impact of routing protocols that use multiple channels. In the next few sections we study the observed connectivity of these graphs in greater detail.

5.1 Multichannel Links in Practice

Asymmetric links are indeed common in our networks. The number of them also varies substantially by node placement, channel, link-quality threshold, and time. We define a link between a pair of nodes as unidirectional if the PRR is greater than some threshold, T , in one direction and less than T in the other. Although stronger criteria would be to require a difference of ϵ around threshold T , we allow even a small difference to be most generous to the prevalence of situations where FH provides benefit. The vast majority of potential links lie far from the threshold regardless.

In examining our connectivity graphs in all environments and thresholds between $T = 1\%$ and $T = 90\%$ with a step of 10, we observed that 32-36% of the links in the machine room are unidirectional, 18-34% of the links in the computer room are unidirectional, and 10-46% of the

links on the testbed are unidirectional.

The distribution of unidirectional links varies substantially by channel. In the machine room, the fraction of unidirectional links, across channels, varies between 8-77%. In the computer room this population varies from 5-42% and on the testbed it is between 10-70%. Still, we were able to maintain connectivity between all the nodes in the network, through routing, for the lifetime of the experiment on every channel in the machine and computer rooms and more than 98% of the time on the testbed.

On any network, the minimum number of directional links needed for routing in a network of N nodes is $2(N - 1)$. So it is not surprising that we are able to route to every node in the network even with fraction of the nodes being unidirectional. The minimum percentage of links that connects the network in the machine room and the computer room is about 5%, while only about 2% of the links need to be present in the testbed.

MCLs are thought to be common in the industrial setting and it is believed that they exist because of multipath interference that causes narrow band fading. In our environment, only a small fraction of the links are considered MCLs. For each trace collected in each environment, we found that the population of MCLs ranges from about 2-6% of all links in the network. This small fraction suggests that MCLs are actually quite rare and may only play a role in sparse networks.

This may also suggest a small gray region in our deployments [30, 33] – locations in the network where connectivity between the sender and receiver are at the edge of radio connectivity. Generally, the population of links in the gray region is small, since reliable communication is desirable and gray-region links have more unpredictable link quality [16]. Sparser networks have more links at the edge of network connectivity and thus there may be some links that are

above the goodness threshold on some channels and not on others, raising the likelihood that they are MCLs. In practice, sparse network deployments should be avoided. The Wireless HART deployment guides specifically suggest that each node have at least 3 neighbors [9]. If these are followed, the network should be very well connected, decreasing the likelihood of MCLs.



5.1.1 Multichannel Link Importance

While the machine and computer rooms were connected for all experiments (all threshold and channels), the testbed became disconnected several times. However, each time there was a partition in the network, there was no channel available to connect the remaining connected components. We explored this further, by examining all the connectivity graphs, for all experimental runs and for all thresholds. On the testbed, only 1.8% of the connectivity graphs (2720 graphs examined) had a corresponding connected multichannel graph. A vast majority of the time, when the graph became disconnected, it was disconnected on *every channel*, not just the one that the network is currently operating on.



This may indicate much wider noise correlation across channels in the operating frequency band. This also directly addresses the assumption made in the standards about FH's ability to avoid interference. It can only improve reliability if there is an opportunity to transmit on *some* channel that is interference free. According to our data, that opportunity is quite rare. Furthermore, this number is optimistic since there must not only be a free channel, but the protocol must find it to be successful. The wider the interference is across channels, the smaller the probability of finding a free channel for transmission.

5.2 Multichannel Triangles in Practice



MCTs are also important in identifying instances in the network where multichannel offers a communication solution where single-channel communication does not. We examined the connectivity graphs for each environment and all thresholds. As a working example, when the threshold is set to 50%, the diameter of the machine room network is between 5-6 hops, the diameter of the computer room is between 2-3 hops and the diameter of the testbed is between 3-4 hops.



(a) Machine room multichannel triangle-set count. (b) Machine room triangle-set count.



(c) Computer room multichannel triangle-set count. (d) Computer room MCT-set count.



(e) Testbed multichannel triangles. (f) MCTs on with N hop solution on testbed.

Figure 10:

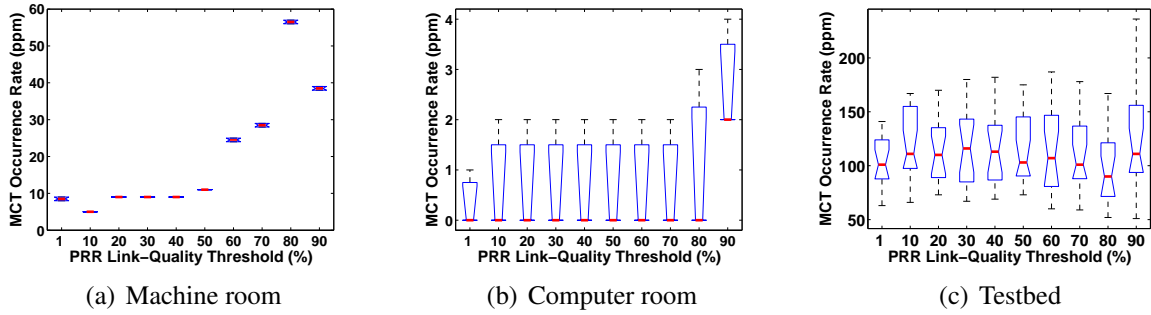


Figure 11: MCT occurrence rate distribution for all experimental runs and link-quality thresholds.

Figure 10(a) and figure 10(b) show the set of non-unique multichannel triangles and corresponding MCTs in the network, respectively, for the industrial machine room facility. The number of MCTs in the network for both runs was extremely small. For example, the figure 10(a) contains 554,578 non-unique multichannel triangles while the figure on the right shows 11 MCTs – an occurrence rate of about 20 parts per million (ppm). The second run has similar results with over 500,000 multichannel triangles and 10 MCTs.

Figure 11 summarizes our results. The boxplot shows the distribution of MCT occurrence rates for each run and threshold. Observe that for all three environments, the MCT occurrence rate is extremely small. The rates range from 0-200 ppm. This indicates that MCTs are actually not that important to consider if your deployment is provisioned to be connected on a single channel – since they occur so infrequently. Also observe the effects of thresholding on the MCT occurrence rate. In the machine room, the occurrence rate increases by almost a factor of 6 from $T = 1\%$ to $T = 80\%$. Only 2 MCTs were found in the computer room and the testbed did not significant variations.

Initially, one might expect to observe more MCTs as the link population admits more gray-region links. The reason it rises so sharply is related to the link population distribution. A vast majority of the link population is either very high quality or very poor (nonexistent), i.e. a

Run ID	MCTs	2-Hop	N-hop
1	145	127	18
2	173	143	30
3	180	156	24
4	155	143	12
5	152	134	18
6	145	97	48
7	87	81	6
8	105	87	18
9	72	67	5
10	74	56	18
11	97	97	0
12	115	115	0
13	99	99	0
14	87	87	0
15	105	99	6
16	107	89	18
17	81	75	6

Table 1: Routing solutions on testbed with threshold set at 50.

bimodal distribution. As the threshold increases, we exclude more links from the population at a faster rate (for all the links on the “good” portion of the bimodal distribution). On the testbed, this is not as pronounced.

Another interesting observation is the differences in the MCT occurrence rates across the three environments. The testbed environment had a *higher* occurrence rate than either the computer room or the machine room. The difference ranged from a factor of less than 2 to almost a factor of 5. This implies that the effects of 802.11 are more serious to consider than the multipath fading. The testbed surely suffered some loss due to multipath, but it was also exposed to the most dynamic environment of the 3 settings. The testbed was susceptible to RF interference, human movement, and other factors.

5.2.1 Channel Distribution

Table 1 shows the associated routing solution count for each MCT found on testbed ($T = 50\%$). Notice, *every MCT has a single-channel routing solution*. The vast majority of routing solutions are two hops in length. Similar results are seen for all thresholds in each all environments tested. This demonstrates that there is *some channel* channel that provides a routing solution to the MCT. However, it does not directly address whether there is a *single channel* that is good for the entire network to use.

We observe that the channel distribution graph shows that every channel where an MCT has an edge there is also a route solution. Furthermore, the machine room and computer rooms connectivity graphs were connected for all experimental runs and the testbed was connected over 98% of the time. On the testbed, channels 25 and 26 were free, for all runs and all experiments. Finding the best communication channel was easier than expected. With a spectrum analyzer and an engineered network deployment (i.e. Wireless HART's recommendation to have 3 neighbors per node), one can likely choose a single channel over which to route over, reliably, for the lifetime of the network.

5.3 A Closer Look At Routing Solutions

We take a closer examination of the use of routing in cases where an MCT is present. In our evaluation, we make two important assumptions. First, we assume expected transmission count (ETX) can serve as a proxy for energy consumption. Second, we assume that a vertex in the triangle wants to send the same data to the other two vertices. The multichannel solution only considers sending the packet to either vertex by switching channels. Therefore we add together the ETX of each link from a source to the other two destinations over links with the lowest




Figure 12: Two-hop single-channel solutions in a Multichannel Triangle instance.

ETX on any channel in the MCT-instance. For the single-channel solution we pick channel 26, one of the better channels throughout our experiments, and calculate the ETX route cost from that corresponding source and destination, with the third node serving as a forwarder along the path.

In considering each solution we also cycle through each vertex in the MCT and set the other two nodes as destinations. After we consider all 6 source-forwarder-destination paths we compare each solution by source node and rank them according to the ratio of the cost for the single-channel route solution to the cost of the multichannel solution.

Figure 12 shows an example of a pair of two-hop single-channel routing solutions for an MCT. Let $cost(ij)_{c_\alpha}$ be the cost of link ij on channel c_α . In the figure, the triangle on the left is an example of an MCT where i is the source, k is the destination, and j is the forwarder while the second example swaps the destination and forwarder roles. Let's assume the "best"(lowest ETX) route solution is (i, j, k) (figure on the left). Then the "best" cost ratio calculation is shown in Equation 4.

$$cost(MC) = cost(ij)_{c_1} + cost(ik)_{c_2} \quad (2)$$

$$cost(SC) = cost(ij)_{c_1} + cost(kj)_{c_1} \quad (3)$$


$$R_{cost} = \frac{cost(SC)}{cost(MC)} \quad (4)$$

We compare the “best” (lowest ETX) and “worst” (highest ETX) route solutions with the corresponding multichannel solution, and calculate their ratio. When the ratio is below 1, the single-channel route solution is cheaper than the multichannel solution. In other words, remaining on channel 26 is a better option than switching channels. When the ratio is above 1, switching channels is cheaper. Finally, when the ratio is equal to 1, their costs were the same. The routing cost ratio calculation is equivalent to *transmission stretch*.


5.3.1 Routing Solutions and Costs

For every MCT found in each run of each experiment, there exists a single-channel routing solution. Table 1, shows the breakdown of the number of MCTs found on the testbed and how many corresponding two and N-hop routing solutions were also found. Every MCT found in every environment has a routing solution on channel 26 (and probably others). Channel 26 is amongst the set of non-overlapping channels with 802.11 on 802.15.4. We chose this channel after a quick examination of our WiSpy data.


Figure 13 shows the routings costs for all routing solutions in the each of our environments. These were calculated with the link-quality threshold set at 50%. Observe that each ratio sample is around 1, showing that the routing solution on channel 26 has a comparable cost to the multichannel solution. Figures 13(b) and 13(d) show the distribution for the machine room and the testbed. The average best-case ratio is 0.97. This indicates that if the protocol remains on channel 26 for the entire experiment it can save 3% in energy-cost. The worst-case ratio is 1.22. Therefore, if the routing protocol chooses poor paths on channel 26 then it consumes




(a) Best(top) and worst(bottom) cost ratio for the machine room setting.



(b) Cost ratio CDF the Best/Worst ratios in the machine room setting.



(c) Best(top) and worst(bottom) cost ratio for the testbed room setting.



(d) Cost ratio CDF the Best/Worst ratios in the testbed room setting.

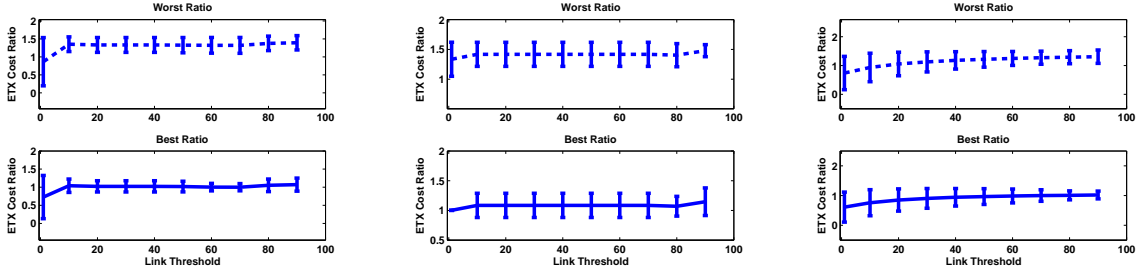
Figure 13: Single-channel to multichannel communication cost ratio comparison. The connectivity graph constructed with threshold $T = 50\%$ and the routing channel is 26.

approximately 22% more energy than using multichannel for transmission.

Of course, this is only part of the cost comparison as there are many other protocol-dependent factors that may lead to inefficiencies in energy consumption. However, this sets a basis for comparison based on transmission links.

Run ID	Best	Worst
1	1	1.5
2	1	1
3	1	1.5
4	1	1.5
5	1	1.5
6	1.5	1.5

Table 2: Cost ratio comparison for the machine room. Threshold set to 50%, routing solutions channel set to 26.



(a) Basement cost ratios as a function of threshold. (b) Machine room cost ratios as a function of threshold. (c) Testbed cost ratios as a function of threshold.

Figure 14: Cost ratios in each environment as a function of link-quality threshold on the connectivity graph.

5.3.2 Link Threshold Effects On Ratio

Recall that to construct a connectivity graph, one must set an initial threshold on the population of links in the traces. For the results that have thus far been shown, the link threshold was set at 50%.

Fundamentally, we want to observe how changing the connectivity graph affects the population of MCTs population observed and the corresponding routing solutions. In varying our link threshold, we noticed changes in the size of the MCT population, especially as the threshold grew closer to 90%.

The authors of [16] show that increasing link threshold on a connectivity graph lengthens the number of hops to a destination and reduces the links on the fringes of connectivity. If routes do indeed get longer then the routing cost should increase and the ratio measurements should also increase. However, this is not what we observed. Figure 14 shows the ETX ratio calculation as a function of threshold. Note the slight increase in cost for the testbed routing solution. As the threshold increases, the link-population choice becomes shorter in length and the choices available are essentially the same for both the multichannel and single-channel solutions, so the ratios do not change much.

Protocol	ROM	RAM
PracMac	9544	761
B-MAC	3046	166
B-MAC w/ACK	3040	168
B-MAC w/LPL + ACK	4386	277

Table 3: Example code size comparison in bytes between single and multichannel MAC protocols [22, 24].

5.4 Other Associated Costs

A good indication of code complexity is code size. Reducing the complexity of the protocol reduces the state and the likelihood of race conditions [24]. Therefore it is desirable to keep code size small. Table 3 shows the code sizes of the original implementation of the default, single-channel MAC in TinyOS, B-MAC, and compares the size against an implementation of a multichannel MAC protocol also written in TinyOS, PracMac [22]. PracMac is almost 5 times larger in RAM and more than twice as large in ROM. Furthermore, motes are resource constrained, and the larger the stack, the smaller the space for applications.

We also simulated FH on the testbed data set, run ID 1, using the first of five pre-set hop sequences in the SP100.11a standard [2]. The links formed by this process were of very poor quality (about 25% or below PRR). With these links, we constructed a sparse, poorly connected graph. Although the simulation is not realistic, it may indicate a problem with FH in an 802.11-rich environment. As mentioned earlier, there are 7 access points sitting amongst the nodes on the testbed. With the access point center frequencies set to 802.11 channels 1, 6, and 11, there are only 4 of the possible 16 channels that do not overlap 802.11 transmissions. Therefore, the hop sequence chooses a “bad” channel 75% of the time. SP100.11a and Wireless Hart are certainly aware of this and make explicit recommendations to blacklist the 802.11 channels a priori. We suspect that if this is done in the testbed environment, the links qualities will not differ from that of remaining on a single, good channel (since the FH protocol is left to choose

from only “good” channels).

Osterlind et al. [28] show the maximum CC2420 radio send-rate is approximately 225 kbps. Therefore, a 60-byte packet takes about 2.1 ms to transmit. The channel-switch time is approximately 1.4 ms – 67% of the time it takes to send a packet. The RSSI read time is almost 500 microseconds, bringing the total-switch plus RSSI-reading time to 1.9 milliseconds, or over 90% of the time it takes to send a 60-byte packet. If a multichannel scheme ranks each channel based on RSSI, the protocol must scan the RSSI on all 16 channels. Scanning takes the same amount of time to send over 14 60-byte packets. If the scheme includes hop-by-hop communication (control-channel scheme), there is an exchange of at least 2 more packets to negotiate a channel at each hop. Even if the transmission cost is amortized over a large burst of packets on the data-exchange channel, the scanning and negotiation overhead is non-negligible.

In the FH case, the overhead is still non-negligible. SP100.11a uses a send frame of 10 ms. In this frame, approximately 5 packets can be sent. If the offered load is greater than 5 packets (> 300 bytes) the sender and receiver must both switch channels. Switching channels takes 1.4 ms, or 14% overhead. Single-channel communication would not incur any extra overhead if a good channel is chosen.

Of course, single-channel communication is not free. Some provisioning and planning is necessary. One must survey the deployment environment, choose the right channel, and test the connectivity over time. One must also set up the network. However, you have to do this anyway, according to SP100.11a and WirelessHART, as both make recommendations about topology properties and blacklisting. Still, FH may offer higher network capacity, as multiple senders in the same space can transmit simultaneously without interfering with one another.

6 Related Work

There has been a gold-rush effect in networking research community to explore the multi-channel protocol design space. In the 802.11 research community there have been numerous publications in theory [12, 14], simulation [11, 27], and practice [25, 23]

In sensor networks, energy consumption is of highest priority and communication consumes the most energy. Therefore reliability and efficiency is of utmost importance. The focus of multichannel work has been to increase packet-delivery reliability. Throughput is a secondary goal as sensor networks mostly transmit at very low data rates and operate at low duty cycles. Several multichannel MACs have been built and studied for sensornets [13, 21, 22, 31, 32]. Y-MAC [21] and the Time Synchronized Mesh Protocol (TSMP) [8] – the state-of-the-art TDMA-based multichannel protocols – specifically cite reliability as a design goal. Although PracMac’s main priority is to improve throughput, the authors are motivated by the same scenario as Y-MAC – a congested cell of simultaneous senders.

Y-MAC partitions time into slots and allocates a fixed interval for broadcast messages and another for unicast messages. During the unicast transmission interval, nodes compete using CSMA on a dedicated control channel. Each node has a frequency hopping schedule and given this information both the sender and receiver can hop to the same frequency to continue the transmission of data.

Y-MAC uses two common approaches to deal with *the possibility* of narrow-band congestion. TDMA is used to prevent interference amongst the nodes in the same network. In addition, CSMA is used to resolve contention events. Finally, frequency-hopping is used to prevent correlated losses across any single channel. The authors cite the need for multichannel communication to improve performance in areas of high node density. The underlying assump-

tion is that high density implies high channel contention. Contention occurs when nodes have a load to offer simultaneously. This assumption is highly workload dependent. For example, uncoordinated event detection may trigger many nodes to fire when an event is detected. Since the event is observed by all nodes simultaneously the channel quickly becomes congested. If a routing tree is used, nodes near the root have more load to offer to the root, creating a condition where multichannel communication up the tree may help. However, throughput is still limited by the parent of any sub tree and it's not clear that even a multichannel protocol approach would improve throughput substantially. The underlying assumption that motivates this work is neither explicitly stated nor qualified.

Although these assumption are not examined in this paper, our results show that congestion events are rare. Even during the worst contention period on the noisiest channel the medium was free almost 80% of the time. Furthermore, this is not an uncommon result for 802.11 traffic patterns. Achieving high utilization of network resources is difficult and uncommon in real-world networks.

PracMac [22] uses a control-theoretic approach to dynamically allocate channels for each node. Multiple home frequencies are used and an algorithm that clusters nodes that communicate frequently into the same home-frequency group is used. The algorithm tuned to try to maximize throughput. The authors of the protocol are also motivated by the congested cell scenario without explicitly examining the frequency of its occurrence. The results do not examine the efficiency of the protocol and neither Y-MAC nor PracMac examine the effects of their protocol on the performance of the network layer. My study suggests that route diversity is an effective alternative to frequency diversity in real world scenarios and deployments and the data suggests that instances when multichannel may improve reliable are rare.

TSMP defines a grid where each row is a frequency and each column is a time-slice. There are 16 total frequencies and 60 time-slots per second. Each node generates a pseudo-random hop-sequence and shares this information with its neighbors. Once a schedule is established, each pair-wise transmission takes place in a unique slot. To maintain accurate link-quality measurements each node transmits small amounts of data every few minutes. Both of these protocols aim to improve reliability and pay the added synchronization overhead in order to maintain transmission schedules. Our data show that the added overhead and protocol complexity may be unnecessary to achieve better reliability.

The added complexity in these protocols is apparent in their design. Our data suggests that instances where added complexity actually helps are rare in practice and an engineering decision must be made about whether it is worth introducing more complexity. The designer must also have a good understanding of the deployment environment. None of the studies looked closely at the motivating assumptions for their work. If the data corroborates the assumptions the added complexity is justifiable. Our data suggests that the motivation for multichannel for increased reliability is rare in real-world environments and that single-channel routing is sufficient and highly reliable.

7 Conclusion

Radio frequency (RF) communication is unpredictable and link quality varies over time. Wireless routing protocols manage link uncertainty by periodically assessing the quality of available links and constructing a snapshot of the underlying connectivity graph over which routing is done. In most large sensor network deployments, multihop routing is used to allow nodes out of communication range to communicate. Therefore, we contend that RF link characteristics

should not be evaluated in isolation to assess the reliability of the network. Instead, one needs to consider the overall set of choices that routing protocols have when routing. If the number of routing choices is small and links become unstable then reliability decreases. Conversely, lots of choices and high-quality links increases reliability and delivery efficiency.

In this work we examined the set of assumptions used to motivate the inclusion of multichannel communication in industrial environments for more reliable communication. We looked at the frequency of unidirectional-link occurrence and the impact they have on the connectivity of the network. We also transformed the assumptions to a set of graphical facets that can be directly observed, called Multichannel Link (MCLs) and Multichannel Triangles (MCTs). MCLs and MCTs capture instances in the connectivity graph where switching channels offers an opportunity for reliable communication where single channel does not.

Our data show that unidirectional links are common and the set of unidirectional links varies quite substantially over time. The population of unidirectional links varied by time, channel, and deployment. Still, in most instances, each of our network deployments contained a path from any source to any destination. Rarely do unidirectional links play a role in overall network reliability. Furthermore, although MCLs are thought to be common in industrial settings, only 2-6% of the links throughout all of our connectivity graphs were classified as MCL link. In addition, although MCLs were sometimes very important for overall connectivity (network bridge link), a *vast majority of the time when the graph became disconnected on some channel it was disconnected on every channel*. There was no opportunity, in these case, for routing or frequency-switching to allow nodes to communication reliably.

In our search for MCTs, we found that MCTs are very rare in practice. The population of MCTs varies between 0-200 ppm. In addition, although it is believed that the industrial set-

ting is the worst for wireless communication, our data show that the office setting was much harsher. The MCTs occurrence rate was 5 times higher than the occurrence rate in the industrial environment. Finally, in the context of routing, MCTs are even less important for reliability. The routing and transmission stretch for communicating was comparable to the multichannel solution. Using transmission stretch as a proxy for energy consumption, we can see that single-channel routing provides comparable reliability for the same cost as the multichannel alternative. This is significant to consider in a typical deployment since single-channel protocols are less complex, require less memory, and provide comparable performance as frequency-hopping protocols.

8 Future Work

For future work we will examine various instances of multichannel and single channel communication protocols and compare the effects that each have on the diversity of observable connectivity graphs. Reliability in multihop networks is entirely dependent on the the difference between the underlying connectivity graph and the snapshot of the connectivity graph that the nodes in the network have. Snapshots of the connectivity graph are taken from the network layer through the MAC layer. It is important to understand how different MAC protocols affect the construction of the connectivity graph snapshot and ultimately, how well the MAC layer stabilizes this view. We will tease apart the various factors that affect this view, such as external interference and workload, by continuously monitoring the environment on an orthogonal data-collection band. We will also explore other relative benefits of multichannel communication, such as increased bandwidth and examine protocols that could take advantage of the spatial-reuse through simultaneous transmissions on orthogonal channels.

References

- [1] <http://www.ieee802.org/15/pub/TG4e.html>.
- [2] <http://tinyos.stanford.edu/ttx/2007/viewgraphs/standards-sp100.pdf>.
- [3] www.xbow.com.
- [4] Chipcon cc2420. cc2420 2.4 ghz ieee 802.15.4/zigbee-ready rf transceiver. <http://www.chipcon.com>.
- [5] Cisco wireless control system – product overview. <http://www.cisco.com/en/US/products/ps6305/>.
- [6] Telosb. http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/TelosB_Datasheet.pdf.
- [7] Wispy spectrum analyzer v2.4. http://www.metageek.net/products/wi-spy_24x.
- [8] Technical overview of time synchronized mesh protocol (tsmp). White paper, Dust Networks, June 2006.
- [9] Wirelesshart technical data sheet. White paper, Dust Networks, September 2007.
- [10] J. B. Anderson, T. S. Rappaport, and S. Yoshida. Propagation measurements and models for wireless communications channels. In *IEEE Communications Magazine*, January 1995.
- [11] P. Bahl. Ssch: slotted seeded channel hopping for capacity improvement. pages 216–230, 2004.
- [12] N. B. Chang and M. Liu. Optimal channel probing and transmission scheduling for opportunistic spectrum access. In *MobiCom '07: Proceedings of the 13th annual ACM international conference on Mobile computing and networking*, pages 27–38, New York, NY, USA, 2007. ACM.
- [13] X. Chen, P. Han, Q.-S. He, S. liang Tu, and Z.-L. Chen. A multi-channel mac protocol for wireless sensor networks. *Computer and Information Technology, 2006. CIT '06. The Sixth IEEE International Conference on*, pages 224–224, Sept. 2006.
- [14] X. Chen, P. Han, Q.-S. He, S. liang Tu, and Z.-L. Chen. A multi-channel mac protocol for wireless sensor networks. *Computer and Information Technology, 2006. CIT '06. The Sixth IEEE International Conference on*, pages 224–224, Sept. 2006.
- [15] S. Dawson-Haggerty. Berkeley ipv6. <http://smote.cs.berkeley.edu:8000/tracenv/wiki/blip>.
- [16] S. Dawson-Haggerty, J. Ortiz, X. F. Jiang, and D. E. Culler. The effect of link churn on wireless routing. Technical Report UCB/Eecs-2008-109, Eecs Department, University of California, Berkeley, Aug 2008.
- [17] 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.
- [18] R. Ganesh and K. Pahlavan. On the modeling of fading multipath indoor radio channels. pages 1346–1350 vol.3, Nov 1989.
- [19] D. Gesbert, M. Shafi, D. shan Shiu, P. Smith, and A. Naguib. From theory to practice: an overview of mimo space-time coded wireless systems. *Selected Areas in Communications, IEEE Journal on*, 21(3):281–302, Apr 2003.
- [20] IEEE. Ieee 802.11 wireless lan medium access control (mac) and physical layer (phy) specifications.
- [21] Y. Kim, H. Shin, and H. Cha. Y-mac: An energy-efficient multi-channel mac protocol for dense wireless sensor networks. In *IPSN*, pages 53–63, 2008.
- [22] H. K. Le, D. Henriksson, and T. F. Abdelzaher. A practical multi-channel media access control protocol for wireless sensor networks. In *IPSN*, pages 70–81, 2008.
- [23] A. Mishra, S. Banerjee, and W. Arbaugh. Weighted coloring based channel assignment for wlans. *SIGMOBILE Mob. Comput. Commun. Rev.*, 9(3):19–31, 2005.

- [24] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 95–107, New York, NY, USA, 2004. ACM.
- [25] E. Rozner, Y. A. Mehta, A. Akella, and L. Qiu. Traffic-Aware Channel Assignment in Enterprise Wireless LANs. In *Proceedings of the IEEE International Conference on Network Protocols (ICNP '07)*, Beijing, China, Oct. 2007.
- [26] D. Sexton, M. Mahony, M. Lapinski, and J. Werb. Radio channel quality in industrial wireless sensor networks. In *SICON '05 Sensors for Industry Conference*, 2005.
- [27] J. So and N. Vaidya. A multi-channel mac protocol for ad hoc wireless networks, 2003.
- [28] F. sterlind and A. Dunkels. Approaching the maximum 802.15.4 multi-hop throughput. In *Proceedings of the Fifth ACM Workshop on Embedded Networked Sensors (HotEmNets 2008)*, June 2008.
- [29] A. Stranne, F. Floren, O. Edfors, and B.-A. Molin. Throughput of ieee 802.11 fhss networks in the presence of strongly interfering bluetooth networks. pages 161–165 vol.1, Sept. 2002.
- [30] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *SenSys*, pages 1–13, 2003.
- [31] G. Zhou, C. Huang, T. Yan, T. He, J. A. Stankovic, and T. F. Abdelzaher. Mmsn: Multi-frequency media access control for wireless sensor networks. In *INFOCOM*, 2006.
- [32] G. Zhou, J. A. Stankovic, and S. H. Son. Versatile low power media access for wireless sensor networks. In *Third Workshop on Embedded Networked Sensors (EmNets 2006)*, 2006.
- [33] M. Zuniga and B. Krishnamachari. An analysis of unreliability and asymmetry in low-power wireless links. *TOSN*, 3(2):7, 2007.