Channel-Specific Wireless Sensor Network Path Data

Lance Doherty, William Lindsay, Jonathan Simon
Dust Networks
Hayward, CA, USA
{ldoherty, blindsay, jsimon}@dustnetworks.com

***Abstract*—Channel-specific path data for a 44-node 2.4 GHz wireless sensor network deployed in an industrial setting is presented. Each node generates one data packet every 28 seconds with the number of transmissions, received acknowledgements, average RSSI, and other metrics for a path to a single neighbor on a single channel for every 15 minutes of operation. Twenty-six days of data were recorded, revealing the scale of time-variation of stability throughout the network and how this is a frequency-dependent quantity. Particularly on low-power paths, both RSSI and stability are observed to vary in unpredictable ways that differ from other paths in the same spatial vicinity. A time-varying model is proposed for simulation of networks in low-noise environments. Channel hopping and path diversity succeed in maintaining near-perfect reliability at a delivered rate of 1.0 kb/s despite this time- and frequency-variance.**

I. INTRODUCTION

Wireless sensor networks operating indoors face RF propagation challenges that lead to time-varying signal and interference strength at the receiver. These effects are difficult to predict during the provisioning stage of a network, and in the field of wireless sensor networks, signal strength effects have been studied mainly with the objective of localization [1]. Multi-path effects in particular pose problems as they can have different impacts on different communication channels and can change as humans and machinery alter the RF environment. An example of the severity of multi-path effects in the 900 MHz band over the size scales of interest is presented in [2]. While predictive strategies for determining the number and location of nodes in a network exist, the actual measured performance of a well-planned network can vary significantly from what is predicted in these models. A deployment in an industrial environment [3] was measured to assess the channel characteristics of a small number of paths with the same radio hardware used here and it showed large variation in path stability. Protocols such as ZigBee [4] allow for star-connected single-channel networks to be formed which can result in data loss if this variation is sufficiently large.

This paper details an experiment to measure the real time-varying effects on different channels in a monitoring sensor network carrying actual traffic. While behavior averaged over time and frequency channels is easier to monitor, a description of a time-varying model representative of the observed data is given. Simulating using this type of model can yield more realistic network performance estimates and identify the amount of resource over-provisioning required to ensure reliable functionality in the face of time-variance.

II. TEST LOCATION

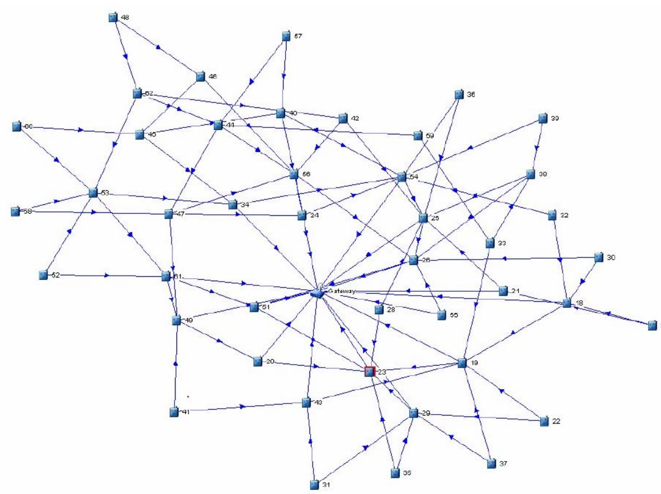

Figure 1. An example of the multi-hop network topology. Arrows represent used paths and point towards the gateway in the middle of the figure.

The test network was deployed in a printing factory in Berkeley, California, and data was collected over 26 days. The building has a rectangular footprint, measuring 250 feet x 225 feet and is three stories tall. The factory is divided into three distinct areas with different propagation obstacles. The south third contains numerous small job printing cells that process pamphlets, voting ballots, handouts and small catalogs. The central third houses the lithography and digital media center on the ground floor, with two floors of general office space above. The north third houses one large printing machine that takes 5-ton paper rolls in at one end, and pushes completed technical manuals out the other end. Also in this bay are the air-handling motors for the entire facility. There are many obstacles in the work area that could impede RF communication and cause multi-path reflections, but there are no major sources of interference.

The network manager is located on the top floor of the office section. 44 radio nodes were deployed throughout the facility, in the manufacturing areas near and around various printing machines, throughout the lithography areas, and in the office areas, in a relatively uniform distribution. The nodes report to the network manager all neighbors within RF range, and from those potential connections the network manager attempts to make the healthiest possible network. Many of these potential paths go unused until they are needed to repair failures. The least connected nodes reported only 4 potential neighbors, while the most connected node reported 26 potential

1-4244-1251-X/07/$25.00 2007 IEEE.

neighbors. The farthest nodes in the resulting self-assembled mesh network had a minimum depth of 3 hops. The average hop depth of all packets for all nodes was 2.48 hops per packet. A snapshot of the connectivity of the network showing only the used paths is given in Figure 1. This topology varied with time.

All nodes in the network use the same hardware and software and perform both data generation and routing functions. Nodes use the TI Chipcon CC2420 [5] radio with a power amplifier to increase output power. Output is nominal 15 dBm EIRP and shows device-to-device variation from +12 to +17 dBm at 25 °C.

III. Network Protocols

The network deployed was architected to provide high-reliability collection of periodic data from all the sensor nodes and follows link provisioning rules similar to those presented in [6]. The network is many-to-one: all multi-hop data is collected at a *gateway* node which relays the packets to the network manager and then to the user. The network operates using the Dust Networks Time-Synchronized Mesh Protocol (TSMP, [7]). The centrally-computed TSMP schedule dictates which of the 16 available channels as defined in the 802.15.4 PHY layer specification [8] (starting at 2.40 GHz and with 5 MHz spacing) should be used for each transaction. During a transmit slot, the transmitting node first performs Clear Channel Assessment (CCA) on the specified channel, and if it passes, transmits the packet to the waiting receiver. TSMP does not schedule colliding transmissions so CCA is used solely to avoid interference from outside the network. If the CRC of the packet passes at the receiver, it immediately sends an acknowledgement to the original transmitter on the same channel within the same 31.25 ms time slot. Every successful message pass consists of the successful transmission and reception of both the original packet and an acknowledgement on the same channel. The next transmission between the same two nodes is, in general, on a different channel and cycles through the list of channels approximately evenly in a pseudo-random manner.

In this experiment, instead of sending sensor data to the gateway, nodes recorded and reported meta-data on the quality of their neighbor paths. A *path* represents all transmissions between a pair of wireless nodes. At the bottom of the mesh, nodes keep time synchronization through messages once per minute and these messages form the basis for the link quality reports. Since these nodes have two parents, even the least busy among them should record statistics for at least 30 transmit attempts per 15 minute reporting period. Further up the mesh, nodes forward on the meta-data reports as well as send synchronization messages so their own reports are based on more transmission attempts per interval.

We allow for each node to have up to 8 neighbors and we communicate on all of the 16 channels. Each path is thus broken down into 16 *path-channels* representing the traffic on a particular channel for that path. A full description of all path-channels for a single node requires 128 entries of which we can fit 4 in a single packet. We stagger these reports evenly at each node to avoid congestion, so to report on all path-channels once per 15 minutes requires one packet every 28 seconds per node. We report on all path-channels even if there are zero transmissions during the interval. Each report contains the following statistics for a path-channel: path ID, channel ID, number of transmit attempts, number of transmit CCA fails, number of transmits without ACK reception, and the mean RSSI and LQI. Tallies are kept on both endpoints of a path.

A successful transmission sequence results in the transmitting node incrementing the number of transmissions and averaging in the new RSSI and LQI measurement based on the ACK. A transmission without receiving an ACK results in incrementing the “No ACK” counter at the transmitter and could result in a few possible outcomes at the receiver. For the results that follow, we define the *stability* of a path-channel, as measured at the transmit end of the path, to be:

$$stability = 1 - \frac{\#NoACK}{\#Transmissions} \quad (1)$$

It is possible for the original packet to be received and queued by the receiver, but for the ACK to not be heard by the original transmitter. A successful transmission with a missed ACK creates a duplicate packet and counts against stability.

IV. Time-Averaged Network Stability

The network self-configured and settled into a steady-state topology before the 26 days of path-channel data were collected. The 88 paths present at the beginning of the data collection were tracked for the duration of the experiment; some were pruned over the course of the collection as part of an ongoing topological optimization to improve performance. Averaging over all of these paths gives a coarse estimate of overall network behavior. Figure 2 shows the time-averaged network stability for each channel. Based on this plot, no channel appears significantly better or worse than any other when averaged over time and paths, and the network appears to function well on all frequencies.

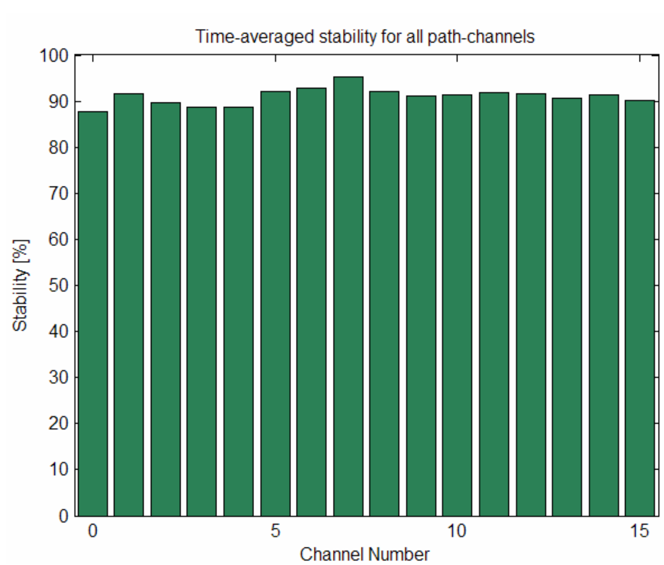

Figure 2. Stability as a function of channel averaged over all paths and time. A mean of 226,000 transmit attempts occurred on each channel.

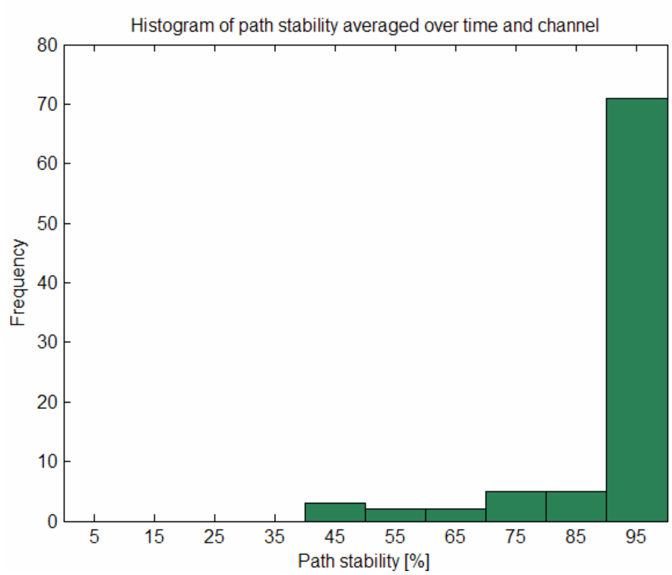

Figure 3. Histogram of 88 path stability values. Each path is averaged over all channels and the duration of the experiment.

Using CCA as a means of reducing energy costs by avoiding doomed transmissions does not succeed in this environment. The most prevalent channel for CCA fails results in 0.13% of transmissions being aborted, far below the ~10% of packets that eventually fail on this same channel. The small number of CCA fails indicates that stability loss is likely not due to in-band interference from external energy sources.

Mean stability analysis at the path-level shows that some paths underperform and appears in the histogram in Figure 3. While the majority of paths have better than 90% stability, the lowest averages 42%. Still, if this behavior did not vary over time and frequency, the needs of the network could be met by scheduling every node 2.4 times as much bandwidth as it would need if it had paths with perfect stability.

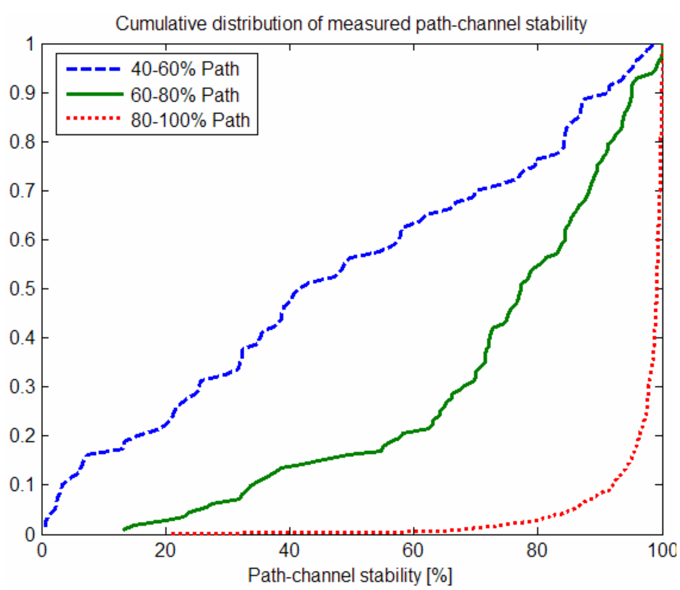

Figure 4. Measured CDFs for path-channel stability based on the overall stability of the path. Each curve represents two bars from Figure 3. Even some 80-100% paths have path-channels with stability below 42%.

As shown in Figure 4, lower stability paths tend to be composed of lower stability path-channels and the overall stability can change the shape of this distribution. Mistakenly choosing a low-stability path-channel in a single-frequency non-redundant network would result in data loss since individual path-channel stability can be arbitrarily small even on otherwise strong paths.

V. Time-Averaged Single Path Stability

In contrast to the relatively flat by-channel distribution of the entire network, a given path can have significantly different stability values for different channels. Furthermore, other paths in the same vicinity may not have the same behavior. The stability of three different paths, again averaged over time and broken down into path-channel statistics, is given in Figure 5. A high-traffic low-stability path (44→56) is shown to have a high path-channel stability near 100% and a low near 20%. The two other plotted paths each involve one of the endpoints of the 44→56 path. Both of these paths are high-stability, but both have low stability on channel 5 where the original 44→56 path is relatively stable. Since a full roundtrip transaction is required for the 44→56 path to succeed, localized interference would impact all three paths in this case. Neither of the other paths exhibits low stability on channels 2 or 12 where the original performs the worst, and both have at least one low-stability channel where another path performed relatively well.

The path-dependent behavior illustrates that path-channel stability is not exclusively a function of a single endpoint; it is dependent on the physical nature of the path between the two nodes. This is likely due to obstacles with different reflective and permissive properties as seen by the different wavelengths that interact in complex ways. Measurement of noise in the environment around a node, such as in a pre-provisioning site survey, is not sufficient to predict the frequency behavior of associated paths. The variability across channels on a given path emphasizes the benefit of frequency hopping on a per-packet basis.

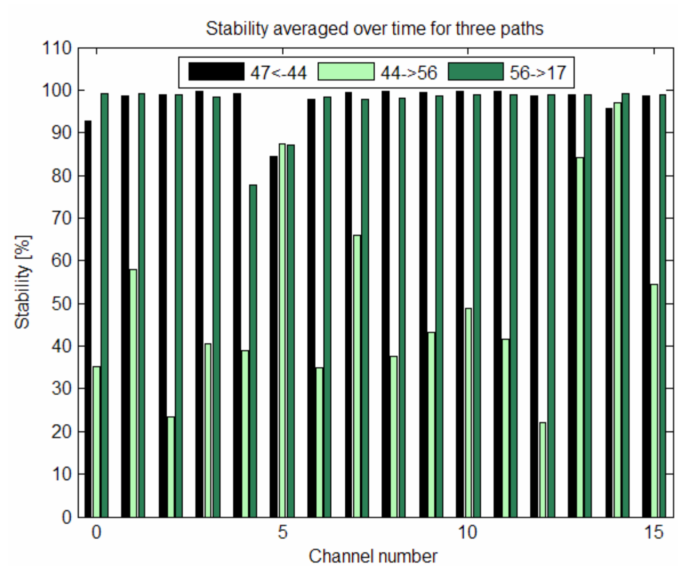

Figure 5. Time-averaged stability as observed on three different paths sharing endpoints. Channel 5 has the closest stability for all three paths.

VI. Time Series Network Stability

Averaging over time, as done in the previous sections, obscures the temporal behavior which can result in networks being under-provisioned. The time series data for the entire network, weighted by the number of transmission attempts on each path-channel, appears in Figure 6. Each data point represents stability in the previous 15 minute interval reconstructed from the individual path-channel reports which arrive at the gateway spread out in time. Stability tends to change as much as 5% between intervals. A 5-hour moving average is also plotted in Figure 6. With less variation, the longer-term trends and periodicity are more apparent. The stability appears to oscillate with a period of a day which could correspond to large machinery moving in similar patterns during each work day. There is also a lower overall stability between days 15 and 24. No changes were done to the experiment during this period, this type of behavior results as the factory environment changes and the system must be robust enough to meet specifications even when stability values are at their lowest.

This same time data is plotted for each channel separately in Figure 7. Each of the 16 channels is represented by one vertically offset line and each data point represents the network-wide stability average over 15 minutes: each line is a vertically-compressed by-channel version of Figure 6. The channels are numbered sequentially from 0 at the bottom to 15 at the top of the plot. No single channel is significantly better than any other and decreases in stability happen on different channels at different times. This suggests that there is no single channel that could be selected to ensure good network performance over all time on all paths. It is not clear what causes the long periods of instability on different channels (for example the second from the top), but the overall decrease in network stability is more the result of the superposition of several channels exhibiting variation rather than any particular channel showing exceptionally large variation.

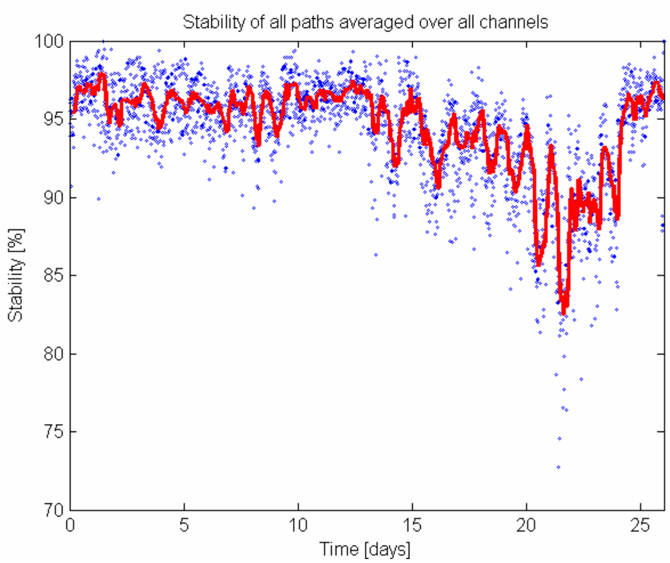

Figure 6. Overall network stability as a function of time throughout the experiment with each point representing one 15 minute interval. The solid line is a 5-hour moving average.

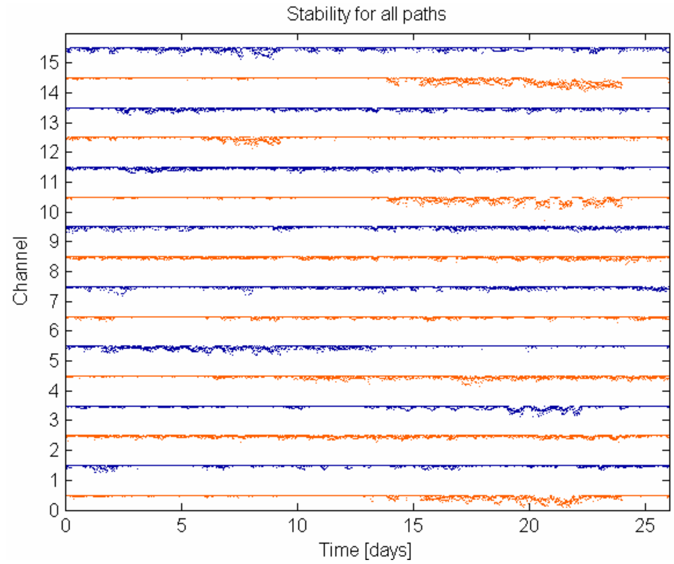

Figure 7. Network stability separated out by channel throughout the experiment. Each horizontal line represents one frequency channel with stability of 0% at the corresponding axis tick. These channels 0 through 15 correspond sequentially to 802.15.4 channels 11 through 26.

VII. Time Series Single-Path Stability

A time-series stability plot for a single path between nodes 24 and 17 is shown in Figure 8. At the path-channel level, the system looks much more tractable than when the network is considered as a whole. Each channel seems to have three distinct phases: the first with near-perfect stability, the second with zero stability, and the third with uniformly distributed stability between the two extremes. Furthermore, the per-channel stability varies with a daily period supporting the previously suggested workday model. The weekly periodicity is also apparent as effects seen Monday-Friday are not repeated during the two weekend days. Some channels transition upwards during these intervals while others transition downwards: the periodically changing environment can have both positive and negative effects on different channels but tend to repeat consistently over time. The complex dynamics of Figure 7 are the superposition of the simpler states occupied by individual path-channels.

As another example, a single channel from the 53→61 path is plotted in Figure 9 along with the RSSI averaged during the same 15 minute intervals. This path-channel is high stability compared to those seen in Figure 8 and the periods of lowest stability match up with the lowest measured RSSI. The signal strength is measured independently by the two path endpoints on the same bit sequence of the CC2420 preamble. The parent node 61 measures the RSSI of the data packets and the child node 53 measures the RSSI of the ACKs. As reported, the parent consistently measures RSSI about 3 dB lower than the child. This is likely due to differences in transmission power or RSSI measurement which both lie within specification and are systematic. That the two values track together suggests channel properties are unchanged on the time scale between the original packet and the ACK, on the order of 10 ms.

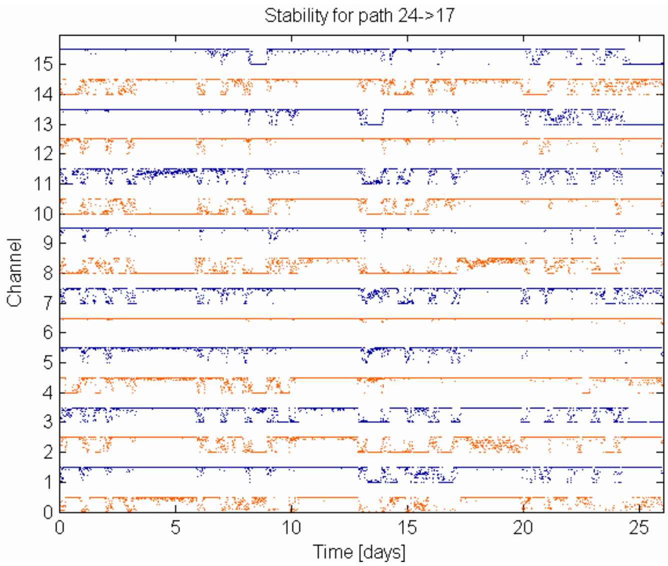

Figure 8. The stability of the path from node 24 to 17 shows daily periodicity corresponding to operational days in the factory.

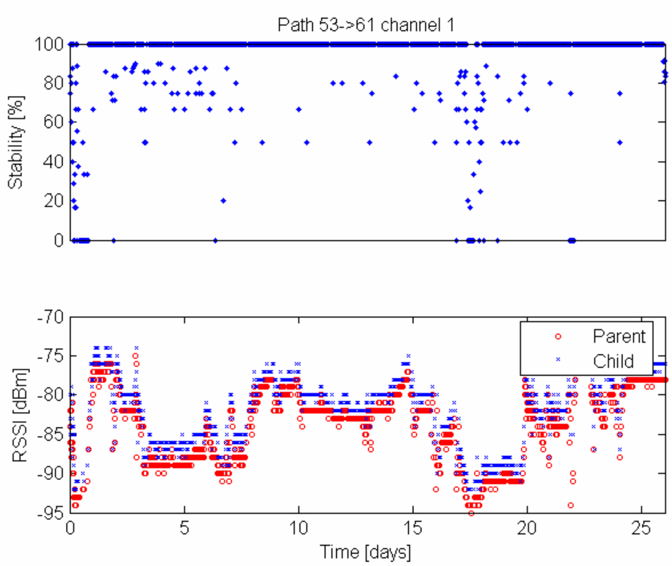

Figure 9. Stability and RSSI of a single path-channel. The top plot shows periods of decreased stability that correlate to lower RSSI in the bottom plot.

The three paths compared earlier in Figure 5 are shown in a time series for channel 5 in Figure 10. This was the channel that had the closest overall stability across the three paths. The failures are spread out differently over the three paths, but all seem to have the same three states described earlier. This evidence further supports that geographic proximity is not sufficient for path similarity. Though not shown, each path-channel correlates to RSSI readings which fluctuate over the same approximate range as in Figure 9 confirming that the changing physical media between path endpoints is responsible for the variation in stability, not outside interference.

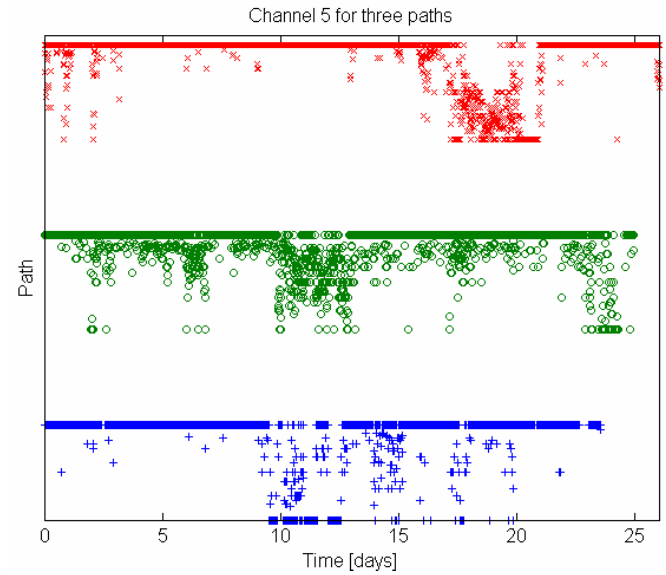

Figure 10. Time series for the same three paths looked at earlier. The top section is 56→17, the middle is 44→56 and the bottom is 44→47.

VIII. Modeling Networks

A multi-hop network is a complex system that is not completely determined by its mean behavior. For the performance quantities of general interest (*i.e.* latency, reliability, power) a simulated network where all paths have the mean path stability will appear better than one where the stability varies over time and channel in the ways seen in this experiment. This section describes a method to generate a time-varying per-channel map of stability in a network that matches the collected data from a simpler set of measurements. The test factory environment was practically noise-free, and the proposed model is for simulation of this type of system. For noisy environments, more involved methods are required such as those proposed in [9]. This model does not capture the large-scale network-wide variations seen in Figure 6 between days 15 and 24, but should serve as a first attempt to simulate an industrial environment with one scale of periodicity. It is not necessarily representative of all such indoor factory wireless networks. The process is to construct time variation at the path-channel level with a simple model which, when superimposed, form an aggregate system with complex network behavior.

The first step is to determine the distribution of average path stability for the network as in Figure 3. This distribution is dependent on the distance between wireless devices and the specifics of transceiver power and sensitivity. For a simulated network, each node pair within communication range should be randomly assigned an overall path stability based on this distribution - call this path stability s_p. Next, each of the 16 channels can be modeled independently by randomly sampling 16 times from the appropriate CDF in Figure 4, and each assigned a path-channel stability $s_{p,c}$. For example, a path with s_p = 85% stability might have several path-channels near 100% stability, a couple in the 40-80% range, and the rest in between.

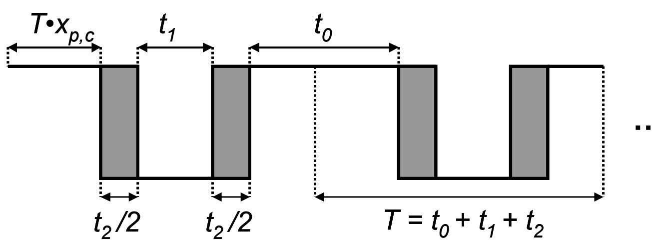

Figure 11. Time-varying periodic channel stability model. Two full periods are depicted starting with a random offset.

In our deployment, the channels vary periodically with the workday. Call this period $T = 1$ day in this case. Inspection of all the collected time-series of path-channel stability suggests that imperfect stability comes during two phases: the first (t_1) is when the channel fails completely; the second (t_2) is when the stability is approximately uniformly distributed between 0 and 100%. For almost all path-channels, there are times of perfect stability (t_0). Motivated by traces like those in Figures 8 and 10, we make the simplifying assumption that $t_1 = t_2$ and that the t_2 is broken into equal halves on either side of t_1. Figure 11 shows a model that captures these phenomena in a periodic manner. Since the mean stability during time t_2 is 50%, we can calculate these three times as follows:

$$\begin{aligned} t_0 &= T(4s_{p,c} - 1)/3 \\ t_1 = t_2 &= 2T(1 - s_{p,c})/3 \end{aligned} \quad (2)$$

The offset $x_{p,c}$ is chosen uniformly randomly in [0, 1] to ensure that not all channels fail during the same period. While simultaneous failures would be useful for simulating a worst case, in the experiment we saw some path-channels improve at the same time that others got worse. For $s_{p,c} < 25\%$, we recommend that no full 100% stability period be modeled ($t_0 = 0$) and the length of t_2 shrunk accordingly. Only 28 of 1408 observed path-channels had stability this low but those that did rarely exhibited extended periods of perfect stability.

For each path-channel, a simulator is required to store only the sampled value $s_{p,c}$ and the offset $x_{p,c}$. During simulation, if a transmission occurs on path-channel (p,c) and the time is during a t_0 period, the transmission is successful. If it occurs during a t_1 period, it fails. If it occurs during a t_2 period, it succeeds with probability one-half. Empirically, this will result in the stability values shown in Figure 11 wherein t_2 periods will in general exhibit all values of stability between 0 and 1.

IX. Conclusions

The purpose of this paper is twofold. First, a long-term study of frequency and path specific behavior was presented. Second, a model to generate some of the more important behaviors was proposed. In general, we found that network performance is the superposition of relatively simple dynamics at the path-channel level but that these dynamics are difficult to predict without an empirical study of the system. For example, while RSSI is strongly correlated to stability of path-channels in our experiment, RSSI itself varies significantly (±10 dB) over time. We found it consistent over millisecond timescales but prone to precipitous decreases over minutes. There is no site survey that can be done to adequately predict the behavior of a network installed at that location, and paths in similar locations often have different patterns of stress.

Typically, network statistics are averaged over time and frequency and mask the underlying physics that cause complex variations at the network level. Using mean behavior to plan network architecture can lead to networks with under-provisioned bandwidth as it is on low stability path-channels during their worst times that packets are lost. The simple model that we propose maintains the same mean path properties while introducing the types of time-varying stability that contribute to decreased performance in real networks. Simulating a network with varying path-channel stability values allows a planner to judge whether enough bandwidth has been allocated for typical interactions among independent paths.

In the face of these challenges, our approach remains to employ path and frequency diversity in all data traffic. By hopping over all channels equally, we avoid the chance of having a single path-channel failure result in data loss. We overprovision bandwidth, at some cost in power, to ensure that the vast majority of potential time variations do not push the network beyond its limits. By encouraging each node to have two parents, the resulting mesh is less susceptible to unpredictable effects on paths. In the 26 days of operation of this network, only 17 packets were lost out of a total of 3.6 million generated even though some individual paths failed or were pruned. This represents a delivered data reliability of 99.9995%. Data packets in our test network contained 80B of payload (not including packet overhead and ACK traffic) which resulted in an overall consistent data collection rate of 1.0 kb/s at the gateway.

References

[1] K. Whitehouse, C. Karlof, D. Culler. "A Practical Evaluation of Radio Signal Strength for Ranging-based Localization". ACM MC2R, Special Issue on Localization Technologies and Algorithms, 2007.

[2] K. Sohrabi, B. Manriquez, and G. Pottie, “Near-ground wideband channel measurements”. Proceedings of the 49th Vehicular Technology Conference, Houston, May 1999.

[3] D. Sexton, M. Mahony, M. Lapinski and J. Werb, “Radio Channel Quality in Industrial Wireless Sensor Networks”, Proceedings of Sicon 2005, Houston, TX, February 2005.

[4] The full 802.15.4/ZigBee specification can be obtained for non-commercial purposes at zigbee.org.

[5] TI Chipcon CC2420 Datasheet

[6] L. Doherty and D.A. Teasdale, “Towards 100% Reliability in Wireless Monitoring Networks”, Proceedings of PE-WASUN 2006, Torremolinos, Spain, October 2006.

[7] A technical description of TSMP is available at: http://www.dustnetworks.com/docs/TSMP_Whitepaper.pdf

[8] The IEEE 802.15.4 specification can be downloaded at: http://standards.ieee.org/getieee802/download/802.15.4-2003.pdf

[9] H. Lee, A. Cerpa, P. Levis, “Improving wireless simulation through noise modeling”, Proceedings of IPSN 2007, Cambridge, MA, April 2007.