International Workshop on Wireless and Industrial Automation
San Francisco, California, March 7, 2005

Improved Quality of Service in IEEE 802.15.4 Mesh Networks

Jay Werb
jwerb@sensicast.com

Michael Newman
mnewman@sensicast.com

Victor Berry
vberry@sensicast.com

Scott Lamb
slamb@sensicast.com

Sensicast Systems
220-3 Reservoir Street
Needham, MA 02494
Phone 781-453-2555

Daniel Sexton
Phone 518-387-4121
Email: sextonda@crd.ge.com

Michael Lapinski
Phone 518-387-6690
Email: lapinski@crd.ge.com

GE Global Research
One Research Circle
Niskayuna, NY 12309

In collaboration under a Department of Energy grant for the Industries of the Future, General Electric and Sensicast Systems have studied performance of 2.4 GHz IEEE 802.15.4 radio transceivers in factory environments, with particular attention to jamming from 802.11 and multipath fading. Temporal and frequency variations in link quality are explored. The implications for network reliability and protocol design are discussed.

I. INTRODUCTION

Wireless sensors for industrial applications are expected to open large opportunities for data collection where it has traditionally been considered technically impossible or cost prohibitive. To overcome installation and acceptance barriers a wide variety of requirements must be satisfied. Some of these barriers include cost and reliability. Short-range wireless technologies such as IEEE 802.15.4 [1] combined with mesh networking techniques are being widely considered as the answer to both cost and reliability in industrial settings. However RF communications, particularly indoors, is well known to be unpredictable.

Wireless Mesh Sensor Networks are being deployed today in various monitoring and control applications. Some radio network designs, such as ZigBee, presume that radio connectivity is reasonably consistent over time. Others take the opposite approach of presuming that links are entirely unreliable, and build large degrees of physical redundancy into the network in the hope that a collection of redundant but unreliable individual links will result in a reliable overall system. Surprisingly little work has been done in the middle ground, of endeavoring to understand the root cause of link failure in real-world factory environments and applying this knowledge in the design of protocols that adaptively detect and use workable radio channels.

Proper understanding of the channel characteristics is needed in order to determine adequate design margins to minimize the installation effort or the amount of physical network reconfiguration required as the environment around the network changes. One approach would simply be to over-configure the network by increasing the node density with additional mesh routing nodes. However this can cause issues with additional installation cost, network maintenance and decreased network capacity. A better approach would be to just slightly over-configure the network by understanding the appropriate required design margins. Of course success in this approach requires collection and analysis of a statistically relevant and representative set of Radio Frequency (RF) channels and environments.

In our work on the Department of Energy grant we have attempted to better understand what effects are present in the RF environment in industrial facilities. As we wish to employ standards when possible, and need an international solution in fielded products, we decided to focus our attention on the performance of 2.4 GHz IEEE 802.15.4 physical radios in this environment. We hope to gain insight into the characteristics required of a mesh network as well as the suitable design margins required.

II. BACKGROUND

In our project we needed to make several important technology and design decisions early so that we could launch critical product development activities. One of these decisions was the choice of a physical radio technology. There were several options available to us. To make an initial decision we first took measurements of a readily accessible and representative harsh indoor environment. We took an initial channel measurement using a network analyzer and two antennas spaced 25 feet apart. An image of the facility is shown in Figure 1 and the schematic of the test equipment in shown in Figure 2. Initial measurements of this environment as shown in Figure 3 illustrates that in the commonly used 2.40 – 2.83 GHz radio band there is both significant frequency selective fading as well as flat fading depending on the area of interest. The dark points represent the actual measurement whereas the lighter points represent the fading as averaged over a 2MHz bandwidth.

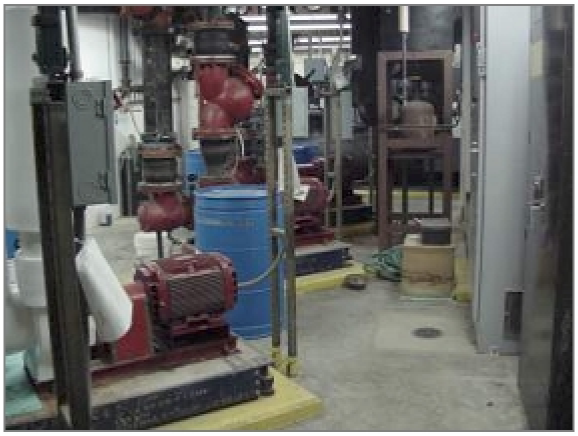

Figure 1: Initial Test Environment

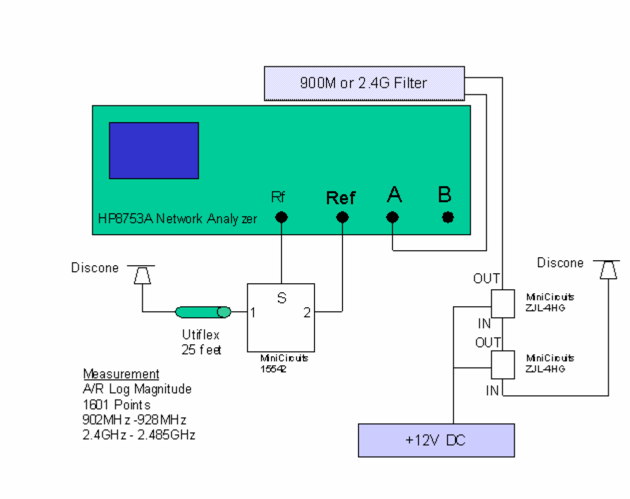

Figure 2: Initial Test Configuration

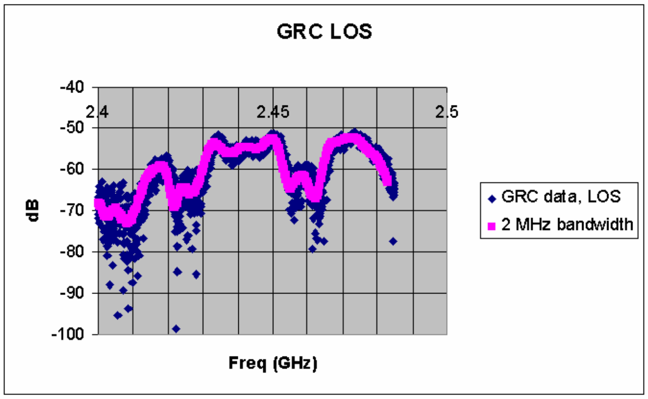

Figure 3: Channel Fading – Line of sight channel

Examination of Figure 3 shows deep fading at certain frequencies. The 2 MHz bandwidth plot is intended to approximate the bandwidth of IEEE 802.15.4, and clearly shows nulls on the order of 15 dB even with spreading. It is also interesting to note that the fades are frequency-specific. This and similar data led us to suspect that link quality could be substantially improved with a protocol that utilizes multiple frequencies. With this observation we took two parallel tracks. First, we tested the fundamental performance of IEEE 802.15.4 radios in factory environments to be sure we understand the root causes of link failure. Second, we implemented a first system based on 802.15.4 radios, with a design that incorporates channel diversity, path diversity, temporal diversity, and increased transmit power. Our strategy is summarized in Figure 4.

	Static Multipath	*Time Variant Multipath*	*Static Interference*	*Time Variant Interference*
Spatial Diversity	😐	😐	😐	😐
Frequency Diversity	☺	☺	☺	☺
Temporal Diversity	🛇	☹	☹	☺
Transmit Power	😐	😐	😐	😐
Overall Risk Coverage	☺	☺	☺	☺

Figure 4: Cost/Benefit of Risk Mitigation Strategies
Key: ☺=Good 😐=Fair ☹=Minimal 🛇=None

As shown in Figure 4, we identified four general causes of reduced link quality in a factory, and four strategies that together would help overcome these issues. Causes of reduced link quality include:

Static Multipath: Pairs of radios in particular positions may experience radio nulls due to destructive interference. In some static environments, these nulls may not change over time. However, even tiny changes in the environment can change the multipath profile, and we would expect to find "Static Multipath" only in sections of a facility where there is minimal human activity.

Time Variant Multipath: We and many others have observed stationary radios that seem to work for a while, then "mysteriously" stop working for a period of time. This is generally due to changing multipath conditions. Time variant multipath can make systems extremely difficult to install, as a link that seems to be good at one time can become unreliable seemingly at random. The multipath profile can change every few seconds in environments that are used by people, or every few hours with the movement of objects such as vehicles, equipment, doors, and chairs.

Static Interference: Interference in the environment, such as microwave ovens or RFID interrogators, may completely block one or all of the channels for a period of time.

Time Variant Interference: Interference from other wireless devices, such as WiFi and Bluetooth, are usually bursty in nature.

To address these risk factors, we identified four basic strategies that we are employing in our first systems.

Path diversity (mesh networking): Our systems do not rely on a single path. Installation and the protocol are arranged so that a small number of alternative paths are supported, providing some path diversity. We choose to support a small number of paths because path diversity involves extra hardware and extra power. Still, even a small number of alternate paths enables packets to adaptively find routes away from interferers, and also provides a degree of multipath immunity.

Frequency diversity: Frequency diversity is a little-appreciated but powerful way to improve the performance of a wireless sensor network. In our implementation, each packet requires an acknowledgement at each hop, and in the absence of an acknowledgement the sender tries to transmit on an alternative frequency and/or path. The adaptive use of the radio band can in theory provide strong multipath immunity. In our experience, links with frequency diversity have provided surprisingly consistent performance over long periods of time in commercial operation. In addition, frequency diversity can help a network operate in the presence of radio interference if such interference is channel limited.

Temporal diversity: Retries are of course a fundamental technique in acknowledged protocols to prevent like devices from interfering with each other. When alien devices are causing the interference, temporal diversity only helps if the channel is changing in the time scale of the retries, such as when the interferer is operating a radio protocol such as WiFi or Bluetooth.

Increased transmit power: More transmit power improves the link margin in general. Our research in factories so far indicates that increased power is necessary to achieve our objective of reliable connectivity at a range of 100 meters per hop. In addition, we have found in physical experiments and in simulation (to be published later this year) that IEEE 802.15.4 radios at very low power will experience heavy interference from WiFi and Bluetooth. To coexist with these other users of the channel, an IEEE 802.15.4 device needs to operate at power levels similar to the interferer. We are therefore targeting 15 dBm as the output power from our radios, rather than the 0 dBm that is commonly used. This power level is at an "inflection point" in terms of battery life. At about 15 dBm, about 35% of the battery's capacity is used by the transmitter, depending of course on the details of the protocol. Above 15 dBm, the transmitter percentage begins to rise precipitously.

As shown in Figure 4, a combination of these techniques can work together to provide a more reliable link.

III. MEASURING IEEE 802.15.4 LINK PERFORMANCE IN FACTORIES

In an earlier paper, we describe strategies that we employed for industrial channel measurement [2]. Here we focus on the measurements that most directly indicate the performance of IEEE 802.15.4 radios in an actual industrial environment. To perform this test we constructed six units each of which housed a radio module and a single board computer with local data storage – Figure 5. These units were placed at locations in industrial facilities where wireless sensors for equipment monitoring would typically be placed. The boxes recorded the transmission performance for every packet sent on every channel so that a history of path performance could be determined. The data was then extracted from the units and stored in a database where it could be processed and sorted as desired.

The indoor channels that we measured all had various amounts of interference from other radio systems. Although interference is not a topic for this paper it had a substantial effect on our measurements and had to be dealt with as the data was processed and analyzed.

Figure 5: Radio Performance Tester

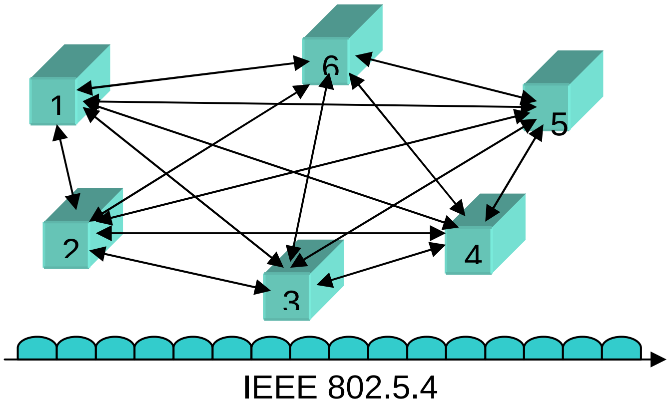

Figure 6: Testing Approach – 6 devices and 16 radio channels

Data collected from the actual radios using the units illustrated in Figure 5 has yielded interesting information. We have not completely harvested all the information in this data but we have gathered some interesting results. In one experiment the boxes were distributed in a machine room (Figure 1) with a floor plan illustrated in Figure 6. Each of the locations for the units is marked 1 through 6. The test was allowed to run for 4 hours. There is little to no motion in this installation as it is in an isolated area. During the experiment we collected lost packets, RSSI (received signal strength indicator) and LQI (link quality indicator) for each packet sent. Since the units are synchronized in time there can be no packet collisions. At any point in time only one unit was transmitting while the remainder were listening and recording. Transmission was then cycled to another unit until all had a time to transmit. This sequence occurred over and over again. The results of each transmission were recorded and is summarized in Figure 9.

Figures 9 and 10 plot the packet loss rate versus path and channel. The nomenclature for path is as follows: Path_{12} represents the packet loss information for the path from unit 1 as the transmitter to unit 2 as the receiver and Path_{21} represents the reverse path. From this experiment we not only studied the performance of the radios but also the symmetry of the channel. In this experiment paths were both line of sight (LOS) as well as non line of sight (NLOS).

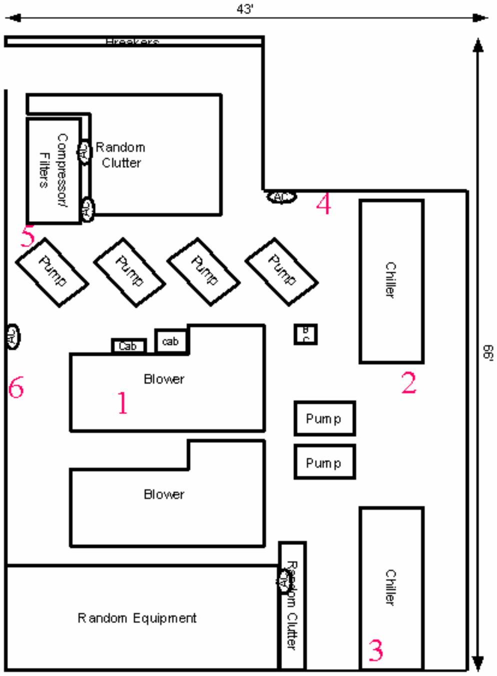

Figure 7: Machine room floor plan

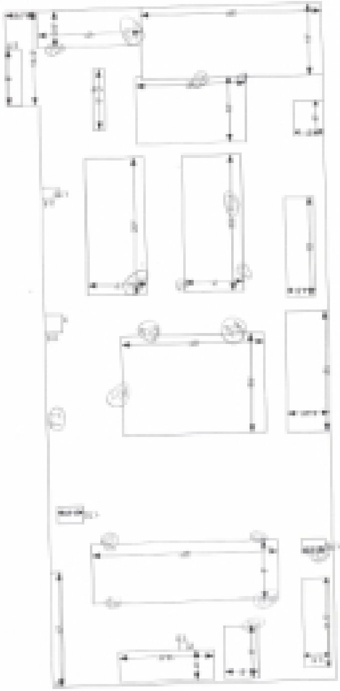

Figure 8: Compressor house floor plan

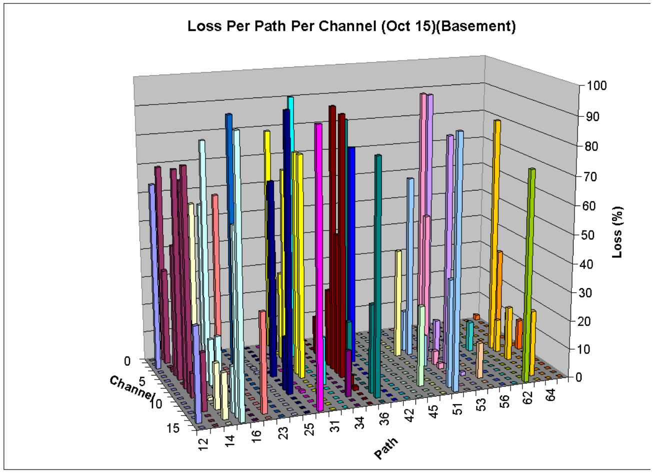

Figure 9: Packet loss versus path and channel – Machine Room

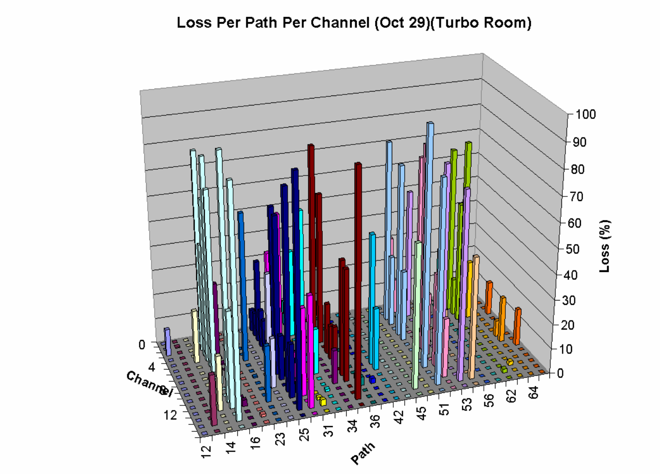

Figure 10: Packet loss versus path and channel – Compressor House

Figure 10 plots the same experiment in a different environment – this data was collected in an industrial gas compression facility. Both data sets show rather vividly the effect of frequency selective fading. In this case of Figure 9 there was no channel that allowed for reliable communications over all paths for all units throughout the entire test period. In Figure 10 only channel 15 was clear for all paths. The results of these experiments clearly show that a frequency agile approach or a multi-channel approach might be more robust than a single channel approach. This finding is consistent with the initial channel measurements – Figure 3.

In our system we require that a message is acknowledged and that the acknowledgement (ACK) is received to complete a transaction. (In a typical scenario, lost ACKs cause unnecessary retries, wasting power and bandwidth but not compromising reliability.) Currently the ACK must occur on the same channel and the reciprocal path as the original message. If there is minimal reciprocity in the paths then the diversity gain could be compromised, potentially requiring an ACK on multiple channels. Our test data does not collect reciprocal path data within the coherence times we measured but we ran a correlation between the data to see if there was reciprocity anyway. For the machine room data in Figure 9 the correlation for reciprocity was 0.91 and for the compressor house data in Figure 10 the correlation for reciprocity was 0.87 both with a high confidence factor. With this level of correlation it appears that we can rely on a frequency agile protocol that ACKs on a single symmetrical channel.

IV. IMPLEMENTATION ISSUES

Path and temporal diversity have been widely used in mesh networking, but frequency diversity is rarely discussed. This section provides an overview of the issues we have found relevant in our commercial implementations of frequency diversity with 802.15.4 transceivers.

The IEEE 802.15.4 MAC does not support frequency diversity per se, but it does allow higher layers to change the frequency of operation. One way to address this is with a thin layer between the MAC and network layers that handles frequency diversity issues. This layer needs to buffer transmissions so that packets are not sent during the period when the system is hopping from one radio channel to another. A beaconing strategy can be used to allow new nodes to find a network and gain time synchronization. Once time synchronization is achieved, the network protocol can operate more or less blind to the fact that the network is frequency hopping. The main caveat is that retry strategy is cognizant of the channels used. When a packet is not ACKed, an enhanced 802.15.4 MAC might employ non-compliant retries by first attempting to use another path, and then waiting for the next hop.

IEEE 802.15.4 radios support 16 channels in the 2400-2483 MHz band, each channel separated by 5 MHz. These channels are largely independent from each other; IEEE 802.15.4 specifies 30 dB of alternate channel rejection. As these channels are independent, it is productive to utilize all 16 channels with sequences that favor large frequency shifts at each hop (10 Mhz or more). However, it is not productive to transmit network “advertisements” on all 16 channels, and use of a smaller number of “control channels” is allowed and desirable.

While balanced use of the available channels is not required in the North America under Ch. 15.247, it does appear to be helpful under European rules where 100 mW of output power is allowed for a frequency hopping implementation involving balanced use of at least 15 channels [4]. Without frequency hopping, the IEEE 802.15.4 radios in Europe are limited to about 10 mW due to the relatively narrow band emissions of the devices. Thus, a frequency hopping IEEE 802.15.4 implementation in Europe is expected to have about a 30 dB edge over a single-frequency alternative, about 10 dB due to higher allowable power, and about 20 dB[1] due to process gains from frequency diversity.

V. CONCLUSIONS

IEEE 802.15.4 at 2.4 GHz appears to be a suitable physical layer protocol for use in industrial environments; however much more testing and experience is needed. Diversity schemes such as spatial diversity through mesh routing, and frequency diversity will significantly help to increase the reliability of the network.

REFERENCES

[1] IEEE Standard 802.15.4 – 2003, Standards for Telecommunications and Information Exchange Between Systems – Local Area and Metropolitan Area Networks - Specific Requirements Part 15.4: Wireless Media Access Control (MAC) and Physical Layer (Phy) Specifications for Low Rate Wireless Personal Area Networks (WPAN), *Standard*
[2] D. Sexton, M. Mahoney, M. Lapinski, J. Werb, "Radio Quality in Industrrial Wirless Sensor Networks", Sensors for Industry Conference, Houston TX USA, 8 February 2005.
[3] FCC rules, Section 15.247
[4] Final Draft ETSI EN 300 328 V1.6.1 (2004-07)

---

[1] Assumes 99% link availability, a NLOS channel with Rayleigh statistcs