Fault Tolerance in ZigBee Wireless Sensor Networks

Richard Alena, Ray Gilstrap, Jarren Baldwin, Thom Stone, Pete Wilson
NASA Ames Research Center
Moffett Field, CA 94035
650-604-0262
richard.l.alena@nasa.gov

Abstract—Wireless sensor networks (WSN) based on the IEEE 802.15.4 Personal Area Network standard are finding increasing use in the home automation and emerging smart energy markets. The network and application layers, based on the ZigBee 2007 PRO Standard, provide a convenient framework for component-based software that supports customer solutions from multiple vendors. This technology is supported by System-on-a-Chip solutions, resulting in extremely small and low-power nodes. The Wireless Connections in Space Project addresses the aerospace flight domain for both flight-critical and non-critical avionics. WSNs provide the inherent fault tolerance required for aerospace applications utilizing such technology. The team from Ames Research Center has developed techniques for assessing the fault tolerance of ZigBee WSNs challenged by radio frequency (RF) interference or WSN node failure.[1]

The ZigBee Network layer forms a mesh network capable of routing data around failed nodes. A two-tier ZigBee network is tested in the lab and various failures induced in sensor and router nodes, simulating realistic fault conditions. A ZigBee network analyzer is used to view the packet traffic and measure the response to these induced faults at the Network layer. Certain faults are induced using Radio Frequency (RF) interference or disruption of the Physical layer, so RF signal levels are monitored during the experiments. The speed at which an orphaned sensor node is detected and an alternative route formed is an important characteristic for fault-tolerant sensor networks. Our working definitions of metrics describing WSN fault tolerance are presented along with a summary of on-going test results from our development lab.

A brief overview of ZigBee technology is presented along with RF measurement techniques designed to gauge susceptibility to interference caused by other transmitters such as wireless networks. Since 802.11 and 802.15.4 technology share the 2.4 GHz ISM band, spectrum management is used to ensure every network has a reasonably clear channel for communications. Quantitative RF characterization of the WSN is performed under varying duty cycle conditions to understand the effect of wireless networks and other interference sources on its performance. Furthermore, multipath interference caused by delayed reflections of RF signals is a significant issue, given that the WSN must run in confined metallic spaces, which produce high levels of reflected multipath RF energy. The results of RF characterization and interference testing of our prototype WSN in the lab are presented and summarized. The architecture and technical feasibility of creating a single fault-tolerant WSN for aerospace applications is introduced, based on our experimental findings.

[1] U.S. Government work not protected by U.S. copyright.
IEEEAC Paper #1480, Version I, Updated December 9, 2010

TABLE OF CONTENTS

1. INTRODUCTION

The emergence of the IEEE 802.15.4 Personal Area Network (PAN) standard for wireless instrumentation has led to the creation of a number of different protocols for managing real-time data streams from multiple distributed sensors. Foremost among the emerging standards are the ZigBee protocol and the ISA.100 protocol. The ISA protocol is highly robust, delivering low bandwidth sensor data primarily for the factory automation market. The ZigBee Alliance defined the ZigBee protocol to address the low-cost home automation and smart energy markets among others, offering mid-range bandwidth and simplicity of PAN formation and maintenance. The Wireless Connections in Space Project collaboration between multiple NASA Centers addresses broad investigations into wireless data transfer methods for spacecraft, including primary avionics buses. This project was formed and funded by the Avionics Technical Discipline Team of the NASA Engineering and Safety Center and is intended to evaluate current and emerging technology relevant to aerospace applications for both critical and non-critical roles aboard human and robotic spacecraft and aircraft.

Wireless avionics technology has great potential for reducing mass and volume by eliminating cabling for many avionics applications,. However, its greatest strength may be producing new capabilities for the design of spacecraft and space missions. For example, wireless may produce a redundant layer for critical control that is insensitive to structural failure that would disrupt wired interconnects.

This coverage of certain common mode failure mechanisms is a key advantage for safety. Wireless can penetrate many materials without the use of actual physical penetrations, a key advantage for crossing pressure interfaces. Wireless components can be embedded in materials, and interrogated remotely using radio frequency (RF) energy, which could conceivably even power these embedded sensors or components. Finally, wireless data paths could result in unusual architectures for future avionics systems, with as yet unidentified properties and advantages.[1]

On the other hand, wireless signaling is always subject to interference from other RF sources. A wireless node that is transmitting generally cannot simultaneously receive signals resulting in half-duplex operation of all nodes – a distinct limitation. While frequency diversity can solve many of these problems, a high-energy broadband noise source (like the Sun) could prevent radio communication completely. Therefore it is very important to characterize the performance of a chosen technology against RF interference and for operation in the intended environment.

Ames Research Center (ARC) focused on the development of Wireless Sensor Networks (WSN) an area relevant to Developmental and Flight Instrumentation and ancillary sensing. Ames chose WSN technology as its focus given the emphasis on System Health Management in the Code TI Intelligent Systems Division.[2] WSN technology is a key enabler of health management for spacecraft, given its capability for providing real-time sensor data in situations where cabling may introduce undesirable complications. The ARC team chose to evaluate the ZigBee (ZB) Protocol due to certain potential advantages of this approach for WSN development:

- ZigBee provides mid-range bandwidth capability at a sustainable throughput rate of approximately 100 Kbps supporting higher performance sensor networks for more scalable sensor networks.
- ZB provides extremely low-cost solutions, with single “System on a Chip” components costs of under $5 featuring very small size and low mass.
- ZB provides a hierarchical mesh network architecture, where sensor nodes can be programmed to function at extremely low duty cycles, significantly reducing overall power consumption.

The reliability of WSNs is affected by in-band RF interference, multipath distortion due to reflection of the RF carrier waves and by WSN node anomalies such as sensor node, mesh routing or gateway functional failures. The methods used to test such fault modes include powering off nodes to simulate faults, introducing external RF sources using Wireless Local Area Network (WLAN) access points operating at the same frequency as the WSN and running the WSN in a closed metal environment to produce high levels of multipath interference. These methods are very similar to those used for WLAN evaluation.[3]

Reliability metrics used to quantify WSN behavior and fault tolerance include: measuring the reduction in throughput caused by external interference; changes in Received Signal Strength and Signal Quality indications; and the time required for the WSN to recover from an induced fault. RF signal levels are monitored using an ISM spectrum analyzer during all testing to understand ambient conditions and to confirm the level of interference being generated in our testbed. Packet capture and analysis tools, running at the 802.15.4 Media Access Control (MAC) layer and at the ZigBee Protocol layer allow measuring packet loss, monitoring packet retry operations and observing the behavior and timing of the Network (NWK) layer. The basic requirement that was levied upon our prototype WSN was one of single fault-tolerance, typical of non-critical functions, yet important for the overall success of the mission. This is consistent with the fault-tolerance necessary for Developmental and Flight Instrumentation or System Health Monitoring networks.

2. ZIGBEE AND IEEE 802.15.4 OVERVIEW

The ZigBee (ZB) protocol relies on the underlying IEEE 802.15.4 Physical (PHY) and MAC layers for packet data transport. The MAC layer uses Carrier Sense Multiple Access with Collision Avoidance (CSMA-CA) just like WLANs. ZB essentially resides at the same layer (L3) as the TCP/IP protocol used for Internet traffic, but is optimized for PANs. PANs are ad-hoc networks, created at run time using a generated PAN ID. Nodes join PANs based on their initialization parameters, either to join a specific PAN, or to join any PAN within range. This needs to be contrasted with TCP/IP networks where network administrators assign IP Addresses and routers connect various subnets into a larger Internet. A PAN self-organizes around a Coordinator, the primary ZB node that initiates PAN formation.

The sensor nodes, or End-Devices collect the sensor values and forward them to the PAN. They can connect directly to the Coodinator or through an intermediate Collector or Router. All nodes look for a Coordinator when they first power up. End-Devices can sleep for most of their duty cycle, waking only to read a sensor and forward the data to the PAN, significantly reducing power consumption. Coordinators and Routers are always active, as they are required to forward or collect data at any given time. End-Devices cannot route PAN data, they only originate data from their local sensors. A special type of node, the Network Capable Application Processor (NCAP) is used to bridge the ZB WSN to the conventional TCP/IP network. The NCAP can be either a coordinator or router and multiple gateways can be used for fault tolerance. A ZB PAN incorporates a hierarchical structure with roles defined in firmware that interacts with the ZigBee protocol stack. Many configurations are possible using this approach, from

long chains of repeaters for improving range to multiple parallel strings offering fault tolerance.

The basic operation of the ZigBee PAN is driven by the Coordinator, which advertises the number of the PAN it is creating by broadcasting a Beacon. Other router and end-point devices then issue an Associate request to join the PAN, which is acknowledged by the Coordinator. The ZB Network (NWK) layer then assigns IEEE Short Addresses to the devices to use as members of the PAN. From then on, all devices use the Short Address assigned to them to communicate with other nodes. All nodes provide a basic heartbeat throughout the network, which is used to detect when a node is orphaned. An end-device that loses its connection to the PAN transitions to Orphan state. An orphaned node issues a Rejoin request, which is acknowledged by the Router. This is the mechanism used for node failover – rejoin the PAN within a couple of heartbeats of link failure.

ZigBee uses the Ad-hoc On-demand Distance Vector (AODV) routing algorithm, which records the logical distance to the next router for path optimization. Routing is a function of the ZB Network Layer. No global routing table is ever present in the PAN; routing is done by hopping from one router to the next. Each router maintains its own routing table for its local neighbors. As nodes associate with, or drop out from contact with a given router, this table is updated. Routes are established using route discovery in which the originating device broadcasts a Route Request and the destination devices send back the Route Reply.

The diagram below describes the complex ZB Protocol Stack, which consists of the Application (APL) Layer on top of the Application Support Layer (APS), and the Network (NWK) Layer, which provides routing and network management.[4] These layers sit on top of the IEEE 803.15.4 MAC and PHY layers, completing the full stack. The ZB APL Layer contains the Application Framework, which contains the specific application code defining the role and function of each node within the PAN. Application Objects interact with End-Points (specific software entity used primarily to bind an application to a specific ZB device) within the APL, and interact with the ZB Protocol Stack through the ZigBee Device Object (ZDO) Layer.

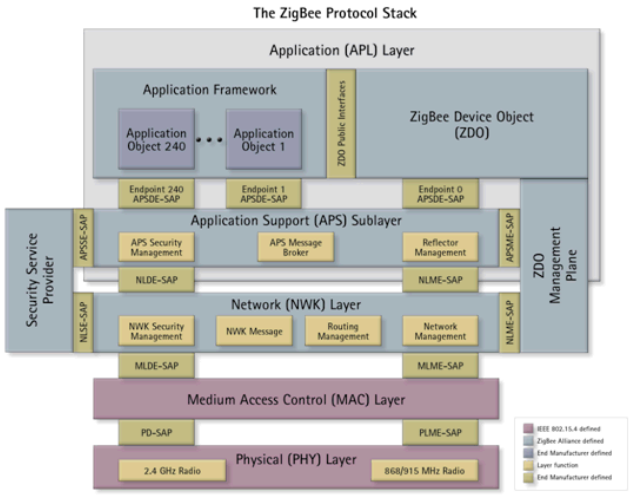

Figure 1. ZigBee Protocol Stack

The ZB protocol has been evolving, with the current version named ZigBee 2007. An extension called ZigBee PRO incorporates security functions, such as key management, authentication and encryption. Additional functions include ZB Channel management to help identify and avoid RF interference. Commissioning is the setup and installation of a complete ZB-based system. Commissioning software interacts with ZB devices to define the PAN to which each device associates by default, their role in the PAN, default routing and configuration and other customizations needed to create a logical network supporting the desired functions. The ZB protocol is evolving rapidly, as are many of the alternative protocols, in response to market needs and technological advances.

3. ZIGBEE FAULT TOLERANCE TESTING

The WSN testbed, built using ZigBee Evaluation Kit components, was used to evaluate key aspects of ZB protocols and failover behaviors. This section defines the metrics, test procedures and test results from the Fault Tolerance Assessment of the WSN testbed addressing the following goals:

- Measure mesh properties of complex ZigBee configurations
- Determine current technology performance parameters
- Determine best way to characterize fault-tolerance behavior
- Determine optimum configurations for fault tolerance and performance

The method was to setup various configurations, (1-hop, 2-hop) with variable numbers of ZB End-Devices (1, 5, 10) and define and measure parameters relevant to redundant, extensible and scalable wireless mesh networks:

- Define test topologies and methods for fault injection
- Identify 802.15.4 and Zigbee protocol packets and handshakes
- Determine timing for formation of PAN, data transfer and failover
- Measure total latency and variability for each network operation

The following table defines the various measurements associated with ZigBee WSN fault tolerance measurements.[5] Most are timings based on certain ZB protocol events, signaled by packets of certain types.

Table 1. Test Measurement Parameter Definition

End Device	A ZigBee reduced function device that is unable to serve as a gateway or coordinator on the network.
End Device Association Time	The time period between a ZigBee End Device sending an initial PAN Association Request to a Gateway or Coordinator and the device sending an End Device Announcement.
IEEE Address Time	The time period between a ZigBee node sending an initial IEEE Address Request and sending an APS IEEE Address Acknowledgement.
Orphan Transition Time	The time period between a ZigBee device discovering a link is broken and declaring orphan status.
PAN Reconstruction Time	The time period between a ZigBee device discovering a link is broken and the orphaned device sending an End Device Announcement
Date Transfer Cycle Rate	A Zigbee Data report is sent from each end device at this rate and the coordinator replies with a Zigbee Data Acknowledgement.

The basic ZigBee testbed components are shown in the photo as Figure 2. The active topology of the testbed can be monitored with the Texas Instruments (TI) SensorMonitor application, included as part of the development kit. The Daintree Packet Analysis Tool provides a separate IEEE 802.15.4 receiver and software that enables monitoring of all PAN packet traffic, without affecting the network under test. Changes in topology are immediately captured in both the Daintree console and in the TI SensorMonitor. The Daintree Packet Analysis Tool's packet capture and decoding capability was used to evaluate items such as latency, latency variability, and failover behavior. A packet capture screenshot is shown in Figure 3.

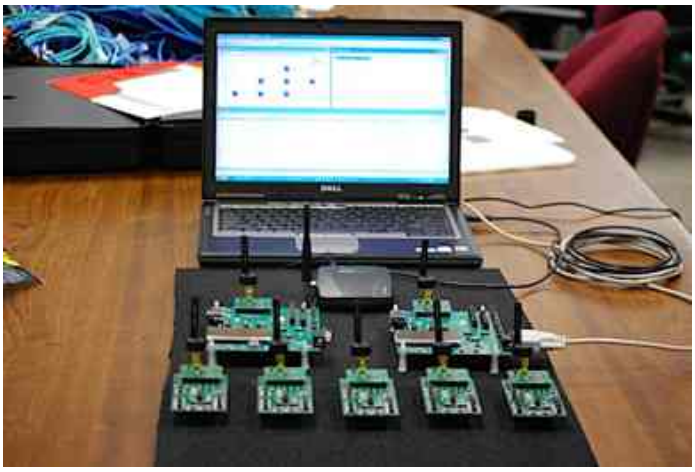

Figure 2. ZigBee Testbed Components

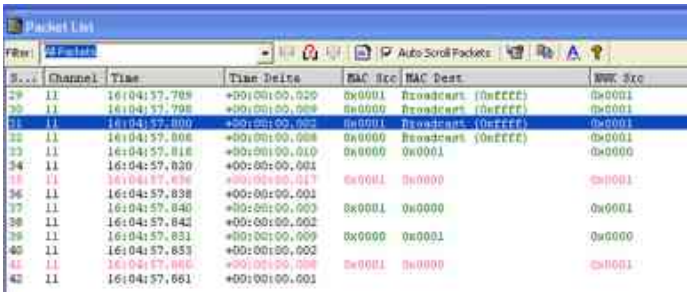

Figure 3. Daintree ZB and 802.15.4 Packet Analyzer

The IEEE 802.15.4 standard is used as the PHY and MAC layers for the ZigBee protocol.[6] Each packet produces an acknowledgement at the MAC level, allowing detection of broken connections. The beacon request and response handshake is used to initiate a Personal Area Network (PAN), playing the key role in node discovery. The association process creates the PAN, assigning IEEE short addresses for convenience using the Zigbee Network (NWK) layer. The End-Device association creates the PAN and the End-Point association connects the software cluster together. The Zigbee Application Support (APS) layer allows data transfer at a periodic rate.

Example parameters to be measured would be the response time for any 802.15.4 packet acknowledgement, the beacon response time and parameters, the PAN formation handshakes, the routing messages, the node association (and re-association) mechanisms and the data transfer methods. Each test parameter is measured using the Packet Analyzer by determining the timing between the packets signaling key events for each test case, injecting a fault into the PAN

and determining the sequence and timing of the resulting reconfiguration. Multiple runs for each test case were performed and summary results reported.

One-Hop PAN Configuration

The initial network configuration tested was a simple 1-Hop topology used to gain foundational knowledge of how the Zigbee end-device nodes interact within a PAN. The time associated with PAN setup, data exchange, and PAN reconfiguration was measured by capturing the packet traffic during association and node failure. The SensorDemo application, embedded in the ZB modules and used for these tests, produces a reading of the ZB chip's internal temperature sensor every two seconds from each End-Device. Collector modules can be used as intermediate routers, simply forwarding the packets to the Coordinator, which passes the data stream to the PC via a serial interface. The PC runs the SensorMonitor application, displaying PAN configuration, module addresses and temperature values on a GUI screen.

For this simple test, a single End-Device is connected to the Coordinator, acting as the gateway to the application running on the PC. To perform a test run, the gateway is initiated and a senor node is bound to the network. The PAN Association time was determined for each End-Device using the Daintree Packet Analyzer, which produces time stamps on every packet accurate to 1 msec. The number of End-Devices was varied from one to five to ten nodes. Five test runs were performed for each configuration. The gateway node was then shut down causing the sensor node to transition to an orphan state and the time required for the End-Device to detect its orphan status is measured using the Packet Analyzer. The following graph in Figure 4. contains the data summary for each configuration, averaging all the runs.

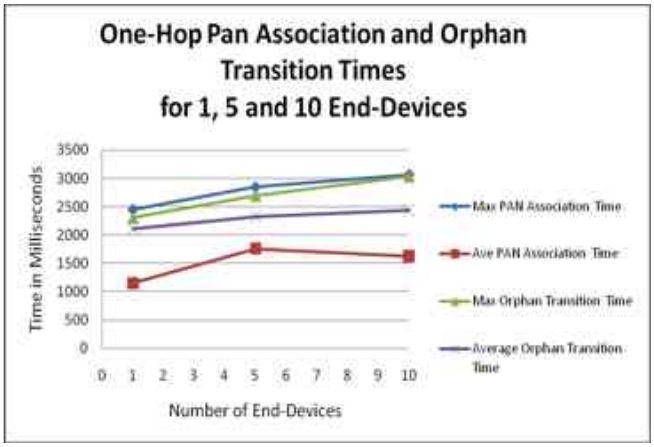

Figure 4. One-Hop PAN Association and Orphan Transition Times for 1, 5 and 10 End-Devices

The PAN Association Time is the addition of PAN End-Device Association, IEEE Short Address assignment and End-Point Association Times. Once a PAN was created the sensor node maintained the PAN connection by sending periodic data reports every two seconds. MAC layer acknowledgments were still maintained at around 1 ms. If the sensor node failed to receive an 802.15.4 MAC layer acknowledgment from the gateway, it entered a "frantic" state in which it rapidly resent the data packet in order to re-establish a connection with the gateway. After a short period of time with no responding MAC acknowledgement from the gateway, the sensor node sent out a ZB NWK layer orphan notification, which was used to measure the Orphan Transition Time.

The test data shows a large variation in the PAN association times for the nodes, varying from 0.8 to 3 seconds. The time it takes for the nodes to declare themselves orphans is dependent upon the data cycle time and varies less than the PAN Association time. These times do not change significantly with the number of End-Devices. These test results are to be expected – the ZB protocol is executed at the data cycle rate, so Orphan transition is declared after one missed data cycle. Scaling does not change these values because the network throughput and ZB protocol stack execution time in the microcontroller are not limiting factors for this number of sensors and at this data transfer cycle rate.

For the next test case, the data transfer cycle rate was increased to understand how incremental increases in network traffic can affect Zigbee PAN reconfiguration times. We modified the TI Sensor Demo software to adjust the ZigBee Data Report cycle time based on input from the joystick. The tests preformed in the previous 1-hop test case were repeated with 1 and 5 sensor nodes to determine the quantitative effect of increased PAN traffic. The new rates used were 1000 msec, 500 msec and 250 msec data transfer cycle times, corresponding to doubling, quadrupling and octopling the network traffic.

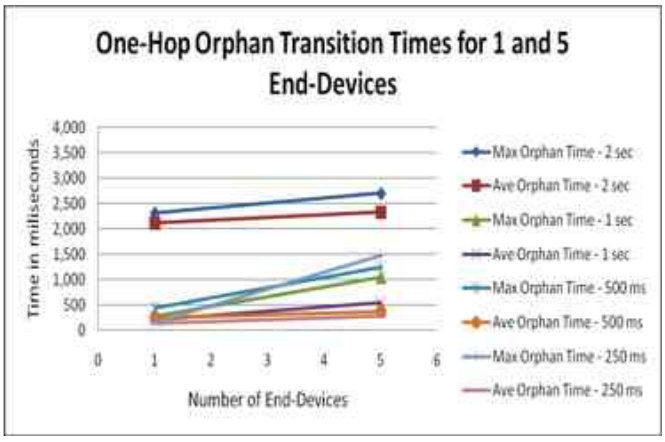

Figure 5. One-Hop Orphan Transition Times for Different Data Cycle Rates

From comparing the data for each test case, the PAN Association Times and the Orphan Transition Times are comparable. The Association Time is not affected by data

transfer cycle time, since data is not flowing at the time of PAN formation. The Orphan Transition Time was affected, since orphan state was declared after missing only ONE data transfer cycle. Decreasing the data cycle time reduces all average Orphan Transition Times. However, after reducing data cycle time below 500 msec, the transition time appears to be limited by the stack execution time and further reductions are not seen.

Two-hop PAN Configuration

This network configuration was tested to study how routing data through a router node would affect the PAN association and orphan transition times. A Zigbee PAN was established that contained the gateway and one end-device sensor node routed through a collector node. Turning the collector power off induced the router fault and resulted in the sensor node re-connecting directly to the gateway. Figure 6 shows the original PAN topology and a visualization of the self-reconfigured PAN after router failure.

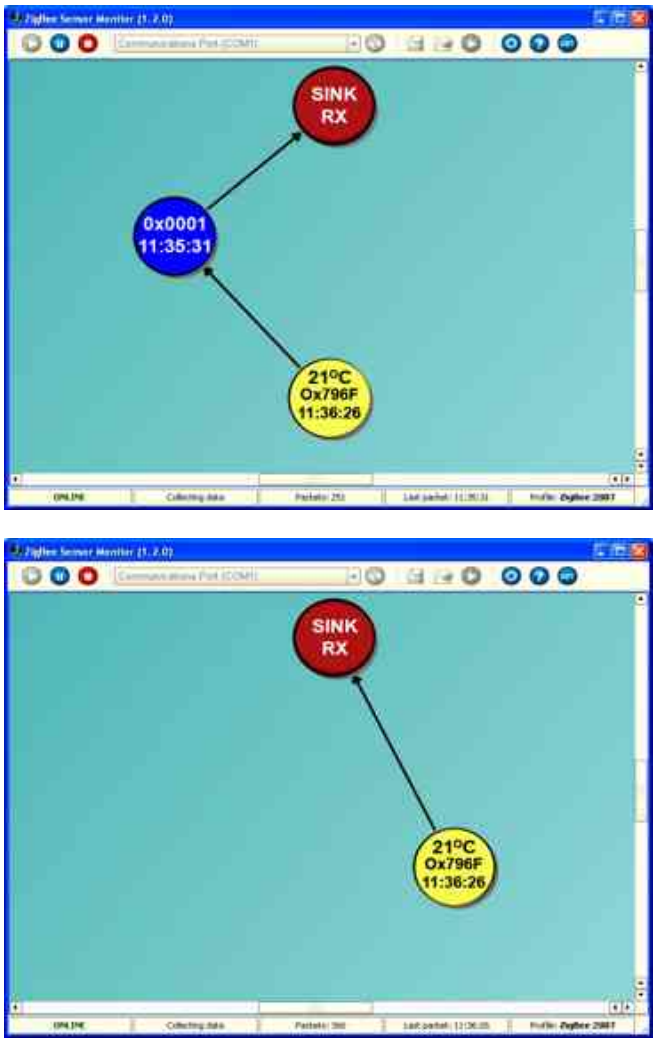

Figure 6. Initial and Failover State for Two-Hop Configuration

Five runs of each configuration (1, 5 and 10 sensor nodes) were done and the results averaged and presented in Figure 7. PAN Reconstruction time was measured, which is the sum of the Orphan Transition Time plus the PAN Re-association Time. Due to the use of the router, the PAN Association time now includes the additional time required to perform the one additional association through the router for the entire chain.

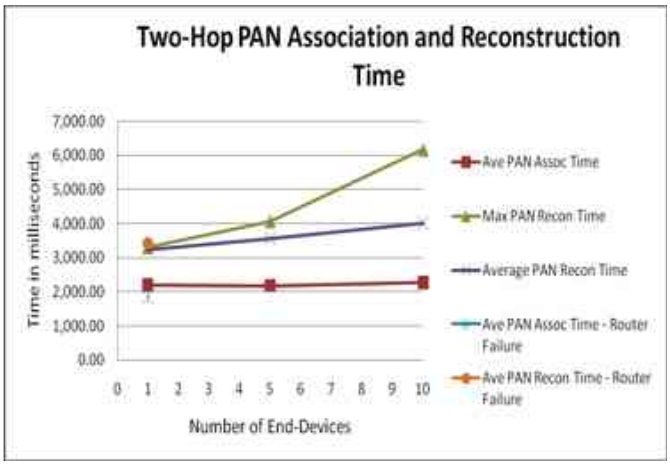

Figure 7. Two-Hop PAN Association and Recontruction Times

The PAN Association Time measures the time required to establish connections and relay data via the Collector acting as a router. This routing increased the total PAN association time due to the latency caused by the routing. This was a 1,150 ms increase for the 2-hop case compared to the average PAN Association Time recorded for the 1-hop case. The PAN Reconstruction time was a bit over 3000 msec, which is the addition of the average Orphan Transition Time (2000 ms) and the average Association Time (1000 ms) for the one End-device case.

Two-Hop with Alternate Router

This network configuration was tested to understand how ZigBee End-Devices re-route using an alternate router after the original router fails, producing a test of mesh routing and recovery. The gateway was initiated and a sensor node bound to the network through a router. A second router was connected to the network to serve as a backup router for the sensor data. The initial router was then shut down in order to observe how the PAN deals with the sudden failure of a data path. Figure 8. shows a pictorial view of the network before and after router failure.

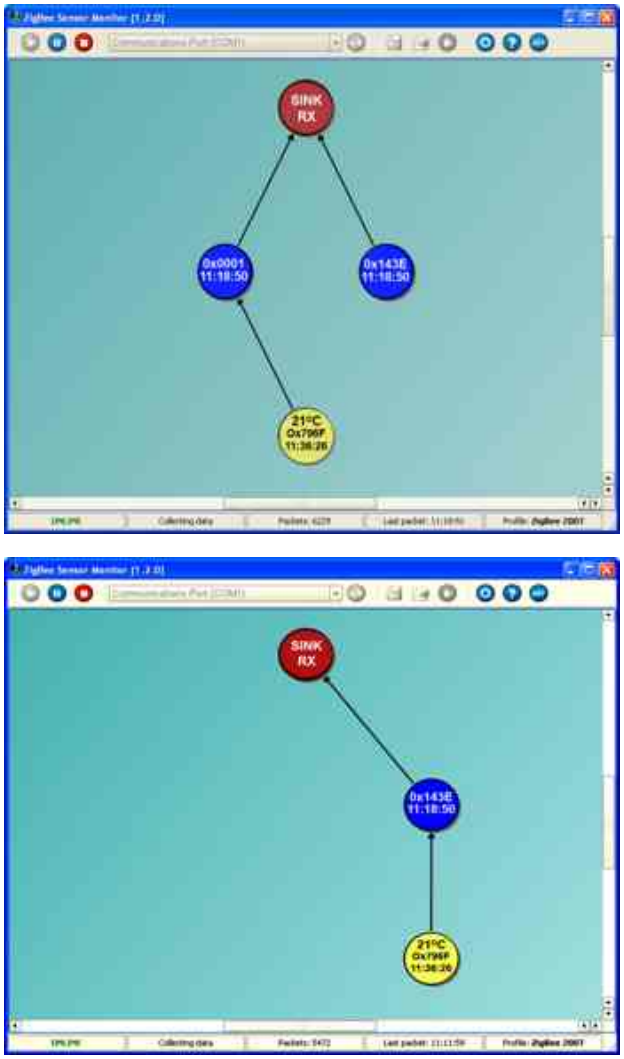

Figure 8. Failover to Alternate Router for Two-Hop Configuration

The data obtained is shown on the previous graph, as two points as only one end-device was tested. Once the initial router failed, the sensor nodes transitioned to orphan state and used the Zigbee Protocol to search for alternate routes to the Coordinator. Since only the alternate router allowed associations, the orphaned node had to connect through it to reach the Coordinator. PAN association times are comparable to the previous test case, as were PAN Reconstruction Times. This is to be expected, since failover to the Coordinator, or failover to an alternate router use similar mechanisms.

Fault Tolerance Discussion

The objective of the ZigBee Fault Tolerance test series was to determine the factors affecting PAN construction and reconstruction times by varying the number of End-Devices, the type of re-routing required for reconstruction and its dependence upon data cycle time. The results show that PAN Association Time is determined by ZigBee Protocol Stack (Z-stack) execution time for performing the needed handshakes and is generally less than 3 seconds. The Association Time is slightly affected by the number of End-Devices, but our test protocol brought each one on-line sequentially. If all End-Devices were brought on line simultaneously, then we would expect PAN association time to increase significantly due to the flood of network traffic and the resultant increased processing load of the coordinator.

As for PAN Reconstruction time, this is also less affected by number of End-Devices and data cycle rate than had been surmised. Minimum Reconstruction times were on the order of 100-200 ms, and Maximum Reconstruction Times scaled with the number of End-Devices. Below 500 ms data cycle time, the Reconstruction times were not affected by shorter data cycle times. However, for data transfer cycle times longer than 1 second, the Orphan Transition Times would be lengthened in direct proportion, resulting in delayed re-routing, confirming the key role played by Data Report cycles in orphan detection and transition. The PAN can form and re-form within 1-4 seconds, with a maximum of 6 seconds for 10 End-Devices. This is fast enough to ensure data loss for only a short time period. While not adequate for flight critical sensors, this failover time is adequate for ancillary data collection from aerospace vehicles.

5. ZIGBEE RF INTERFERENCE TESTS

The RF Interference metrics and test protocols address the following goals:

- Measure relevant parameters of the RF Physical layer
    - 2.4 GHz ISM band Spectrum
    - Radiated output power and received signal strength indication (RSSI) for each ZB transmitter
    - Radiated output power and received power for each WLAN transceiver
- Measure relevant parameters at the MAC layer nominally and in the presence of multipath interference
    - Packet Loss Rate
- Measure WSN RF compatibility at the MAC and Protocol layers with active WLANs operating within WSN channel allocation.
    - Packet Loss Rate
    - Data Throughput Rate

ZigBee RF Characteristics

The diagram below shows the channel allocations for IEEE 802.15.4 (PAN) IEEE 802.11b/g (WLAN) and 802.15.1 (Bluetooth) devices, which share the 2.4 GHz ISM band.[7] The ISM band is an unlicensed public band shared by many devices using multiple standards and all devices within the ISM band are required by the FCC to use some form of spread-spectrum modulation to minimize the potential interference between devices. Spread spectrum uses a spreading code to widen the frequency range of the modulated signal, also reducing its overall power level, allowing multiple ISM devices to co-exist within the same geographical area. As the diagram shows, 802.15.4 uses 16 channels within the 2.4 GHz ISM band, with each channel occupying about 5 MHz. By contrast 802.11b uses about one-third of the ISM band (25 MHz) to carry 11 Mbps, 802.11g uses one-third of the ISM band to carry 54 Mbps and 802.11n looks like multiple 802.11g WLANs. Bluetooth by contrast, uses the entire 84 MHz spectrum to carry its low-power signal. Each standard uses a different form of spread-spectrum modulation: DSSS (direct-sequence spread-spectrum) or FHSS (Frequency-hopping).

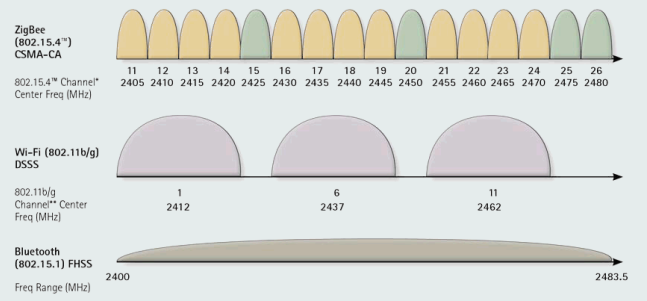

Figure 9. ISM Spectrum Diagram

The different standards provide different RF power output levels geared toward their particular market and function. WLANs can output up to 1 W of power, but typically provide just 100 mW of RF power, capable of covering about 1 Km range. 802.11g is limited to 30 mW to limit interference with 802.11b. 802.15.4 PAN radios can output up to 10 mW of RF power, but typically run at 1 mW per node to reduce overall power consumption. This yields a range of about 200 to 800 m. Bluetooth defines two classes of devices – one has WLAN type power and range, but the class used for headsets output just 2.5 mW. Therefore, the potential for interference between each type of ISM device is a highly complex interaction dependent upon power, distance, channel used and modulation scheme.

RF Spectrum Measurements

The physical test configuration is shown below in Figure 8. The WLAN and ZigBee networks were located in close proximity to maximize the effects of mutual RF interference. The WiSpy RF spectrum monitor was placed in the middle to measure the total RF power spectrum. A separate WLAN interference source was setup to run within the ZigBee channel allocation and created maximum interference for the ZigBee WSN. The *iperf* application was run on two PCs, one connected to the wireless access point (WAP) and the other a WLAN client, and each transmitted TCP/IP packets as quickly as possible to fully load the WLAN and produce maximum RF transmission during these interference tests. Two Zigbee Collectors were programmed with the *Transmit Application* firmware for measuring WSN throughput and were connected to two PCs running *SmartRF Studio* for measuring WSN packet loss rates. The WiSpy spectrum analyzer is used to capture the entire ISM (2400 – 2500 MHz) band and display the results in peak and average RF energy level in dBm, the duty factor of transmitted energy together with a waterfall diagram showing emissions over time. The WiSpy *Channalizer* application is used to capture ambient RF characteristics of the WSN Testbed without any Zigbee devices active. Next, an active packet transfer operation is conducted between two ZigBee nodes and the RF spectrum captured again, to show the spectrum produced by the ZigBee devices. Finally, a spectrum with WLAN interference concurrent with ZigBee operation is captured to show maximum RF signal levels from all sources.

The RF Spectrum measurement procedure was designed to provide characterization of the ambient RF emissions in the lab, to measure Received Signal Strength Indicator (RSSI) values from WSN receivers and to characterize the RF interference spectrum including energy levels. The RF test environment had a WLAN 802.11g access point about 5 meters from the ZigBee devices. The channel assignment for the local WLAN Access Point (WAP) changed without notice and had been detected on WLAN Channel 1, 4 and 11. Therefore, baseline RF spectrum measurement was part of the nominal test procedure to account for ambient RF in the analysis of RF characteristics and compatibility.

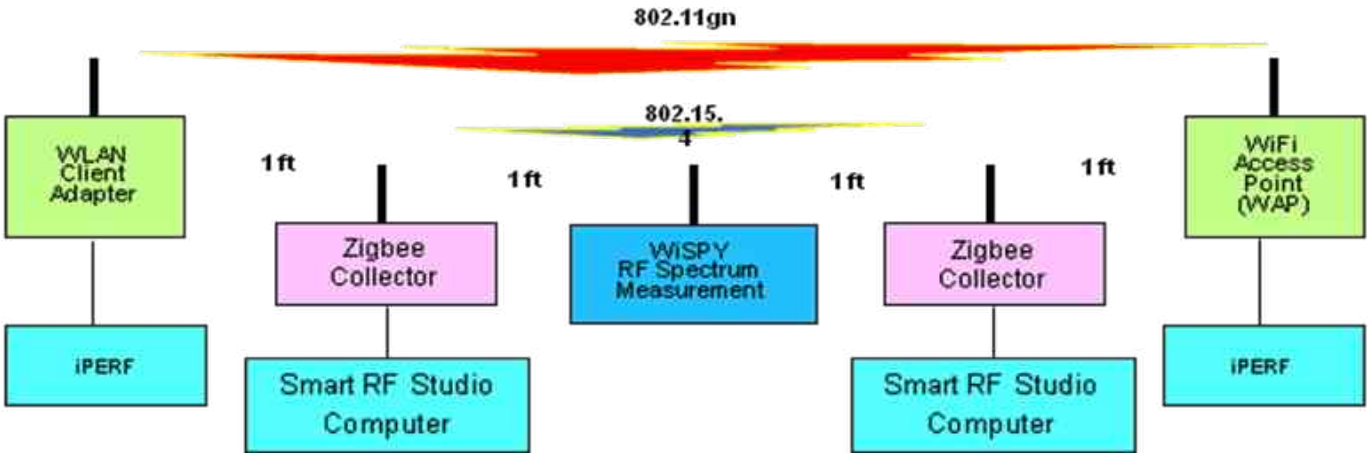

Figure 10. RF Interference Configuration Diagram

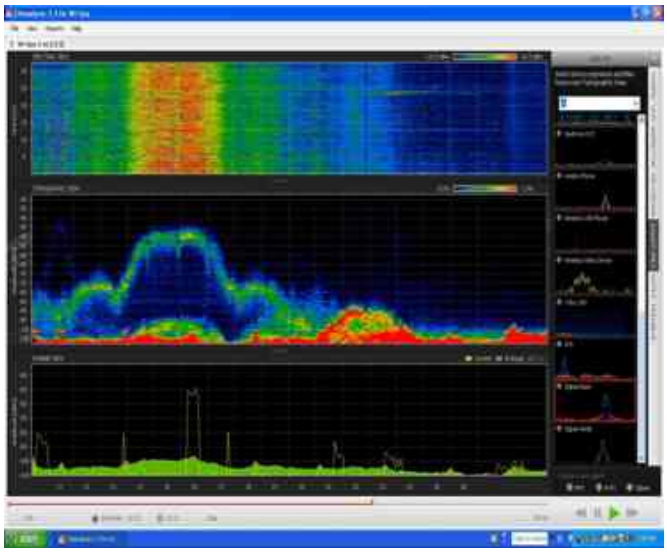

Figure 11. RF Test Lab Baseline Spectrum

The baseline RF spectrum shown in Figure 11 above shows the local WAP operating on WLAN Channel 4 with ZigBee operating on ZB Channel 11. The amount of RF energy is significant at -45 dBm within this frequency band. The distance from the WAP to the WiSPY spectrum analyzer was 3.5 m, with an expected signal level of -40 dB. Generally, this same level of energy is emitted by our Zigbee nodes running at 1 dBm at a distance of about 0.8 m. Therefore the WiFi signal is at least as strong as the ZigBee signal. Apparently there was plenty of traffic on the WAP, since the duty factor for the WLAN signal appears to be high, as evidenced by the red/orange area in the waterfall diagram.

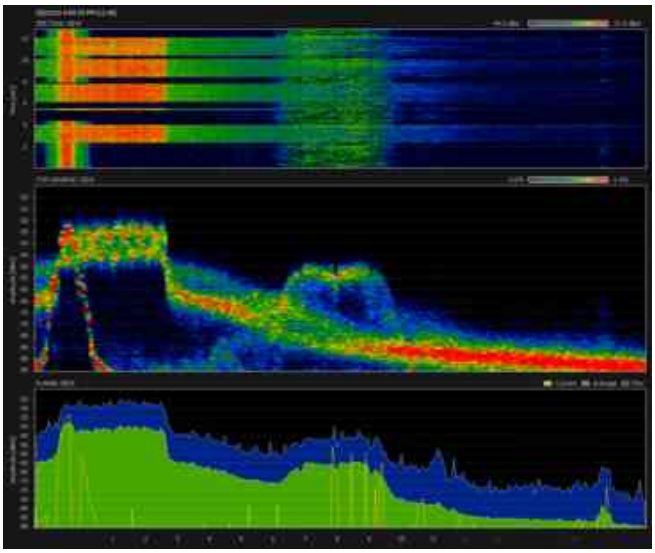

Figure 12. WLAN G mode with ZigBee Spectrum

The WiSpy graph above shows the RF spectrum in the lab during the WLAN RF Interference test sequences. The WAP was set to 802.11g mode on Channel 1 and the ZB WSN was set on ZB Channel 11, overlapping at the lower end of the WLAN spectrum. Note the high throughput on both the WSN and WLAN as indicated by the red duty cycle indication in the top waterfall diagram representing four runs. Note that the WLAN and WSN signal levels measured by the WiSpy located between the two networks are roughly the same at about -32 dBm. This setup was designed to ensure both networks interfered with each other without overloading any receivers.

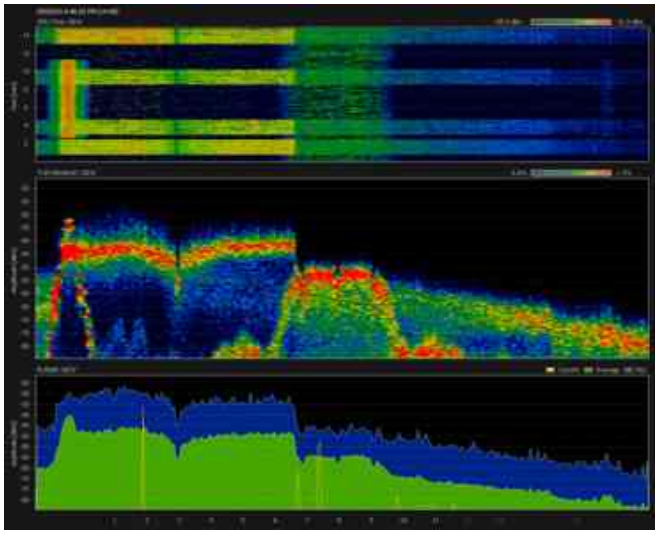

Figure 13. WLAN 802.11n Interference Spectrum

The WLAN WAP was next set to 802.11n mode Channel 3 at a rate of 65 Mbps. The ZB WSN was left at its original setting. The spectrum produced during the packet transfer test and the throughput tests is shown in Figure 13. Note that the n-mode WLAN looks like a g-mode WLAN of about double the bandwidth. The average power of the WLAN was less than the prior test, given that similar data throughput was now spread over double the bandwidth.

Packet Loss from RF Interference and Multipath

The next test series measured MAC layer packet loss when the WSN was challenged by multipath interference and by a controlled source of in-band WLAN interference. The ZigBee development kit comes with the SmartRF Studio application that can directly measure packet transfer loss rate at the 802.15.4 layer. The application can set the active ZigBee channel (11-26) and output power level (-30 dBm to +6 dBm) and other parameters of the radio portion of the chipset. The application allows setting up packet transfers and measuring the packet loss rate under various conditions. Packet errors are usually due to RF interference disrupting the ability to interpret the 802.15.4 packets. For these tests, the ZigBee protocol stack is NOT active. Two collectors are used to measure packet loss rate with the WSN running at a moderate data rate under varying conditions of RF interference from multipath or WLAN transmissions.

The first test determined the extent that multipath interference, caused by reflections of the carrier wave by conductive surfaces could affect IEEE 802.15.4 packet transfer. The exact mechanism is corruption of the modulated data by the overlay of delayed signals from the same transmitter that traveled by a different and somewhat longer path. A baseline was obtained for the two ZigBee collectors placed on the testbench located a very short distance (12") apart to measure packet loss rate without multipath interference.

For testing multipath interference effects, the ZigBee collectors were moved into a metal desk drawer of dimensions (12"W X 20"L X 6"H) as shown below. The node antennas were placed about 12" apart, and about 2" from the metal sides. The WiSpy was placed between the two in the same drawer. The doors were closed so four-way reflection occurred at each wall of the drawer, producing nearly 100% multi-path RF energy. Two different output power levels (1 and -19 dBm) were used for the Collectors and packet loss rates measured for each case. The RF energy from the ZigBee nodes was concentrated in the drawer, producing a standing wave pattern, although this did not significantly affect measured RSSI. The WSN RF level was as expected given their close proximity. Also, the metal shielding of the drawer significantly attenuated (-15 dB) the external WAP signal.

Figure 14. Multipath Test Configuration in Drawer

One thousand packets were transferred using SmartRF Studio and the packet loss rate measured. Error rate remained at the same level as it was before the Collectors were placed in the drawer. The first run (Multipath 1) used 0 dBm (1mW) as the ZB transmit power and the second run (Multipath 2) used -19 dBm as the transmit power. The lower power resulted in somewhat higher, but still nominal, packet loss rates. This data is shown in Figure 15. below.

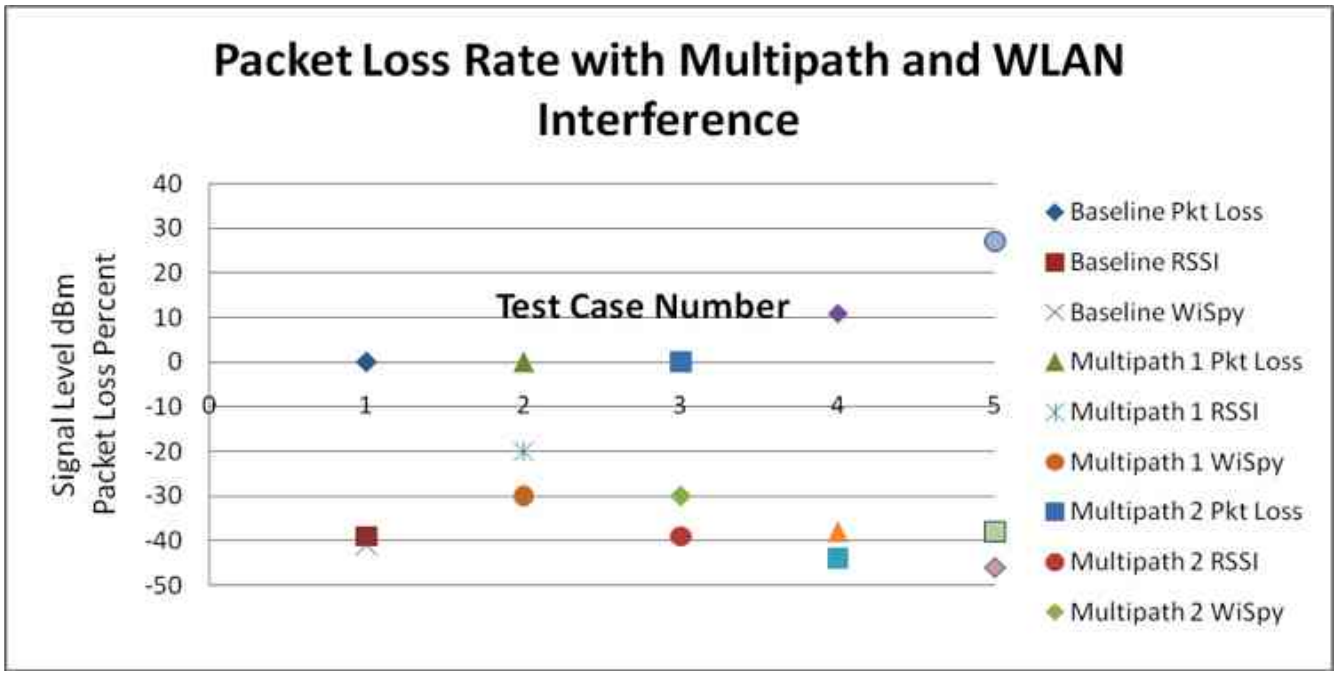

Figure 15. Packet Loss Rate with Multipath and WLAN Interference

The packet loss test procedure was repeated using a controlled RF WLAN source, where the mode (802.11g or n) and operating channel could be set and which could be fully loaded with TCP/IP traffic for stimulating maximum RF interference within the ISM band. The main objective was to quantify the 802.15.4 packet loss rate vs WLAN duty cycle for each of the WLAN modes. The packet loss test was repeated, this time using the WLAN in 802.11g mode at 54 Mbps raw data rate with a total throughput of 9.3 Mbps (5.1 Mbps to client concurrent with 4.2 Mbps from client to WAP). The WLAN throughput without ZigBee running was 9.7 Mbps (4.8+4.9Mbps), indicating some minor interference from the ZigBee WSN. This test showed significantly higher packet loss rates than the baseline test, indicating that the WLAN (transmitting within the ZB channel) was now directly interfering with ZB and increasing packet loss from negligible to about 10 percent. Prior packet loss rates never exceeded 1% under any conditions other than direct WLAN interference. The graph in Figure 15 above summarizes the packet loss tests. Please note, the negative numbers on the x-axis are Signal Level, while the positive numbers indicate Packet Loss.

The 802.11g mode interference test produced two significant results. The packet loss from in-band WLAN g-mode was significant at 10%. This means when a packet is lost it will trigger complex packet retry mechanisms at the MAC and Protocol layers, a response relevant to the throughput tests described in the next section.

The WLAN interference test was repeated using the same WAP, but now configured in 802.11n mode at 65 Mbps raw data rate. The WLAN produced a throughput of 8.2 Mbps with the WSN running, but over 16.6 Mbps when running alone. While statistical error (variations in throughput from run to run) can account for some of this variation, most of it must be attibuted to the WSN signal interfering with the WLAN, halving the overall throughput. This was an unexpected result – the WSN signal interferes much more with the N mode WLAN under nearly the same operating conditions as compared to G mode.

When the packet test was repeated using 802.11n mode WLAN interference source, the packet loss rate was significantly higher at 30% compared to 10% for the G mode case. This was not expected since it seemed like the duty cycle of the WLAN was considerably lower than in the G mode case. Signal levels were consistent with prior cases after RSSI values and WiSpy Spectrum Plots were compared.

ZigBee Throughput versus WLAN G and N

The third test suite measured the effect of WLAN interference on maximum data throughput at the ZigBee Protocol layer. The purpose was to measure the effect of increased packet loss rates on the ZigBee Protocol by observing the reduction in WSN maximum throughput rate. Throughput testing is a very sensitive measurement, because throughput can be greatly affected by packet loss rate. Packet loss is greatly magnified because packet retry (in response to a lost acknowledgement at the MAC layer) takes much longer than nominal packet transfer and therefore significantly reduces throughput even though packet loss

rate is low. Therefore, we expected throughput loss to be a multiple of packet loss rate for these tests.

The test method used to measure ZigBee network maximum throughput was the Transmit Application provided by TI as part of their evaluation kit. The Transmit Application was used to send ZigBee packets between two collector nodes that had the Transmit Application loaded into the CC2530 firmware. The Transmit application simply sends data packets at the maximum rate from one ZigBee node to another. The kilobytes per second (Kbps) transfer rate is displayed for both transmit and receive sides respectively along with a cumulative byte count. The effect of RF interference from WLAN operating in the 802.11g and 802.11n modes could then be determined with the entire ZigBee stack operating thereby making measurements at the Protocol layer.

The first run tested the ZB WSN without the WLAN running, resulting in around 100 Kbps, the nominal full throughput rate as shown in the first column in Figure 16 below. The next column shows the WSN throughput results while running the WLAN concurrently in the same way as the packet transfer test. The third column is the WLAN throughput with WSN running and the last column shows WLAN throughput without the WSN. Therefore, all combinations of RF interference on WSN and WLAN throughput are measured and shown in the graph below.

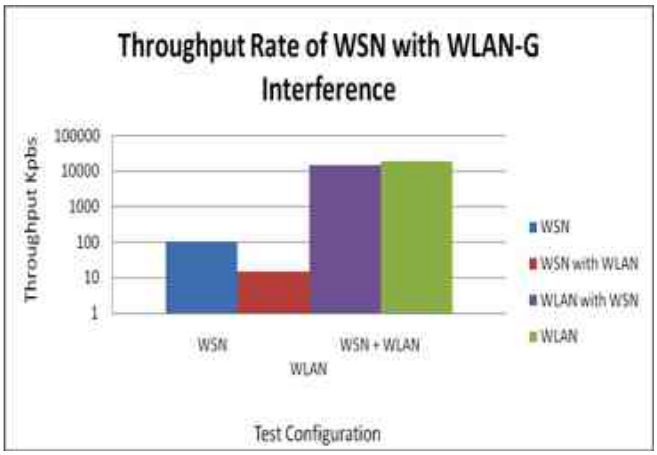

Figure 16. Throughput Rate with WLAN-G Interference

We actually measured a 90% loss of ZB throughput. Also, in certain cases, the WLAN interference shut down all ZB traffic for a period of 1-2 seconds. The ZB WSN interfered with the WLAN, but only reduced WLAN throughput by about 25%.

The test was repeated for the N-mode WLAN and the results shown in Figure 17. in the same way as for the G mode. The differences between the G and N mode are significant.

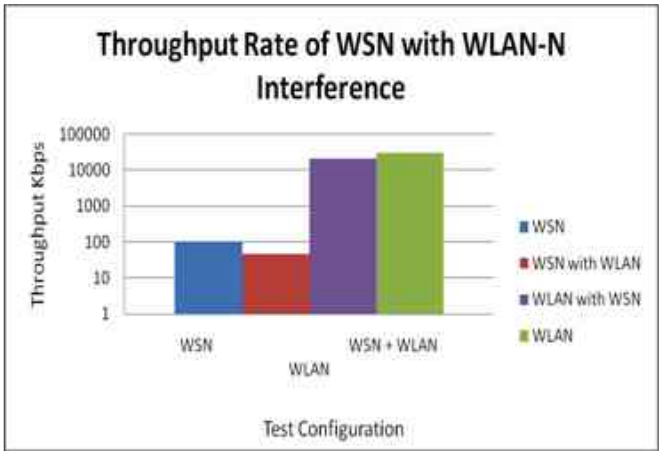

Figure 17. Throughput Rate with WLAN-N Interference

The N mode WLAN interferes much less with ZB throughput than the G mode WLAN. However, the reduction of ZB throughput is much less with this type of interference, and the shutdown phenomenon does not occur. The ZB WSN significantly interferes with the WLAN, reducing WLAN throughput by about 50%.

Interference Discussion

The multipath testing was the most important, since highly multipathing environments enclosed within wing and other metallic structures require sensors. Very little effect on packet loss rate was seen, although many more tests will be performed in a variety of volumes of differing material, size and shape. Throughput testing in the presence of multipath will help verify these results.

The WLAN interference test results were surprising and did not confirm our initial hypothesis–that MAC packet loss rate would be a good indicator of data loss rate at the Protocol level. While packet loss was lower for the G-mode case, the throughput was reduced by a factor of 10. The N-mode packet loss tests indicated that the effect on WSN throughput would have been expected to be greater due to higher packet loss, but the WSN throughput was much less affected than in the G-mode interference case. Moreover, the WSN affected the WLAN throughput much more for the N-mode case, (30 Mbps reduced to 20 Mbps).

The *Smart RF Studio* application generates 5 pkts/sec while the *Transmit Application* sends 20,000 pkts/sec. This major difference in WSN duty cycle changes the probability of intercept between WSN and WLAN, which means the packet loss results cannot be compared directly with the throughput loss results. These measurements must be made concurrently, with the WSN running in the same way for each test, resulting in a controlled experiment.

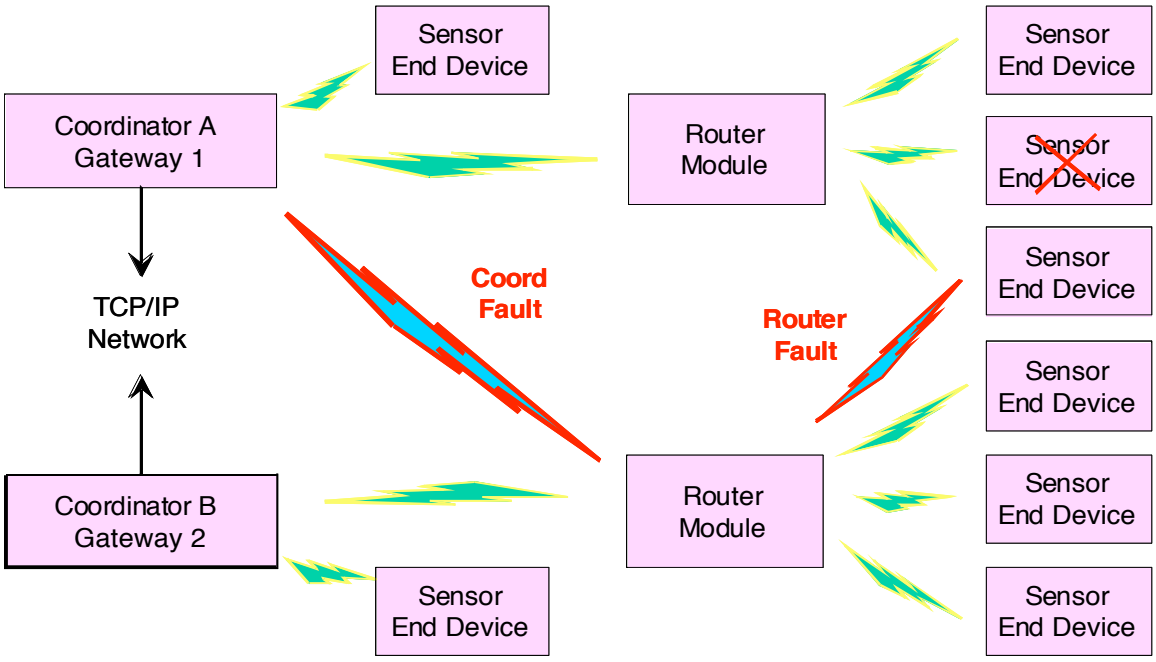

Figure 18. Proposed Fault Tolerant WSN Architecture

6. Fault Tolerant Architecture

The results of fault tolerance analysis and testing help define a single fault-tolerant architecture for WSNs applied to Developmental and Flight Instrumentation aboard future aerospace vehicles. Wireless methods allow miniaturization of sensor systems, allow them to be installed in locations where running cables is prohibitive and significantly reduce interconnect mass and volume. Fault tolerance testing has shown that the failover mechanisms employed in the ZigBee 2007 protocol are robust and work quickly. A sensor node never failed to recognize its orphan state nor failed to re-configure with the PAN. However, these tests were not performed if the WSN had marginal RF links or if the WSN was scaled up to hundreds of nodes.

The conceptual WSN architecture is shown in Figure 18. above. The architecture shown is a balanced mesh with more parallel pathways than required for nominal operation, reflecting the use of redundancy for implementing fault tolerance. The failover mechanisms may not necessarily result in NO loss of data, there may simply be a short loss of streaming data during reconfiguration. It is not known how much internal buffering of data streams is performed within the ZigBee protocol.

Multiple Coordinators are used for fault tolerance, with each Coordinator acting as a Gateway to the avionics system. The exact method of setting up parallel coordinators has not been identified, but seems to be the only outstanding design issue. The routers are also redundant, allowing sensor nodes to reconfigure to the alternate router in case of primary router failure. The alternate paths for coordinator or router failure are shown in red. If a sensor node fails, there is no alternative. However, the design must include enough sensors that one or more can fail without loss of critical flight test data. Therefore, redundant sensors are employed, of which one can fail and the other supplies the needed data. Interestingly enough, such redundancy will give you cross-checking of data accuracy if the WNS is operating nominally, so sensor redundancy provides additional benefits overall.

Certain rules must be observed for this architecture to work. ALL redundant nodes must be within RF range of each other, since the alternative paths must be supported by the PHY layer. Certain logic must be built into the WSN mode firmware and in the WSN commissioning mechanism to setup the default nominal configuration upon startup. This capability is built into the ZigBee protocol. Other factors will improve scalability. Limiting the length of a chain will speed up reconfiguration. Balancing routers will also speed up reconfiguration. An analysis of overall WSN throughput for each link, for both nominal and off-nominal conditions will support scalability studies and can be performed during the design phase.

7. CONCLUSIONS

The test results were generally favorable for the use of ZigBee technology for WSN applied to non-critical ancillary data collection aboard aerospace vehicles. The quantified results for PAN formation showed completion of ad-hoc configuration within about two to three seconds that held up for double hop networks through a router. The fault tolerance testing showed that failover occurred both with high-reliability (we did not see any failed PAN re-association) and within a short interval. PAN reformation generally occurred within 2-3 seconds, but could take as long as 6 seconds for router failover with five sensor nodes. Analysis showed that PAN reformation time was paced by data cycle rates and network traffic load. PAN reformation would be rather slow if many routers were supporting many sensor nodes. A router failure would require each sensor node to be re-associated with the coordinator, resulting in many ZigBee protocol layer exchanges.

Preliminary multipath testing has shown minimal effect on ZB WSN performance, which means that WSNs can be deployed within enclosed metallic volumes aboard spacecraft and aircraft. One must ensure a proper RF propagation path, but multipath interference is not expected to be a problem. Further testing will be performed.

WLANs interfere with ZB WSNs only when operated within the same area of the ISM spectrum. The interference is mutual, that is ZB also affects WLAN throughput, but to a lesser degree. If operated within the same area of the RF spectrum, WLAN 802.11g can shutdown ZB WSNs based on the periodic loss of all throughput. The exact response of the ZigBee Protocol to this interference source which results in this behavior will be investigated in future tests, where packet loss rate is measured concurrently with throughput.

Generally an interference source must be equal or higher in power at the receiver than the intended signal before function is compromised. Therefore, physical separation, shielding, power limitation and spectrum management can all be used to help alleviate RF interference concerns Effective spectrum management and electromagnetic compatibility testing can prevent such problems.

REFERENCES

[1] Richard Alena, Steven R. Ellis, Jim Hieronymus, Dougal Maclise "Wireless Avionics and Human Interfaces for Inflatable Spacecraft" IEEE Aerospace Conference 2008, Big Sky MT.

[2] Fernando Figueroa, Randy Holland, John Schmalzel, Dan Duncavage, Rick Alena, Alan Crocker, "ISHM Implementation for Constellation Systems," AIAA 2006-4410, 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, July 9-12, 2006, Sacramento Convention Center, Sacramento, CA.

[3] Alena R., Ossenfort J., Lee C., Walker E., Stone T., "Design of Hybrid Mobile Communication Networks for Planetary Exploration" IEEE Aerospace Conference 2004, Big Sky, MT.

[4] Getting Started with ZigBee and IEEE 802.15.4, Daintree Networks Inc. 2008

[5] Richard L. Alena, John P. Ossenfort IV, Kenneth I. Laws, Andre Goforth "Communications for Integrated Modular Avionics" IEEE Aerospace Conference 2007, Big Sky MT.

[6] Getting Started with ZigBee and IEEE 802.15.4, Daintree Networks Inc. 2008

[7] IEEE Personal Area Network Standard 802.15.4

BIOGRAPHY

Richard L. Alena is a Computer Engineer in the Intelligent Systems Division at NASA Ames. Mr. Alena worked on the Ground Data System and performed Communications Analysis during operations for the LCROSS Lunar Mission. He was the co-lead for the Advanced Diagnostic Systems for International Space Station (ISS) Project developing model-based diagnostic tools for space operations. He was the chief architect of a flight experiment conducted aboard Shuttle and Mir using laptop computers, personal digital assistants and servers in a wireless network for the ISS. He was also the technical lead for the Databus Analysis Tool for on-orbit diagnosis of ISS avionics. He was group lead for Intelligent Mobile Technologies developing planetary exploration systems for field simulations. Mr. Alena holds an M.S. in Electrical Engineering and Computer Science from the University of California, Berkeley. He is the winner of an Ames Honor Award for Engineering in 2010, the NASA Silver Snoopy Award in 2002, a NASA Group Achievement Award for his work on the ISS Phase 1 Program Team and a Space Flight Awareness Award in 1997.

Jarren A. Baldwin is an Electrical Engineer currently pursuing graduate degrees at Stanford University. He has been a NASA Ames Intern in the Special Projects Division: Information Technology Directorate since 2007. His primary interest is the hardware and software technologies associated with wireless communication systems. Jarren developed the failover methodology and metrics for the project described in this paper. Mr. Baldwin holds a B.S. in Electrical Engineering from the University of Illinois at Chicago.

Ray Gilstrap is a network engineer at NASA Ames Research Center. He holds a B.S. in Electrical Engineering from Florida Agricultural and Mechanical University and an M.S. in Electrical Engineering from the University of California Berkeley. He is been involved in numerous projects in the areas of network architecture design, space communications, satellite and wireless networking for field operations, network security, and multimedia.

Thom Stone is a Senior Computer Scientist with Computer Sciences Corp. He is attached to the NASA Research and Engineering Network project at Ames Research Center (ARC). Mr. Stone has been at NASA ARC employed by various contractors since 1989. He was an engineer with the NASA Science Internet project office where he led the project that bought reliable Internet connections to remote locations including US bases in Antarctica including McMurdo Station and Amundson Scott South Pole Station. He was principal engineer for communications for the NASA Search for Extraterrestrial Intelligence (SETI) project and was a senior engineer for the Space Station Biological Research Project. Before his involvement with NASA, Stone was employed in the computer and communications industry and taught telecommunications at the undergraduate level.

Pete Wilson P.E. is a Consulting Engineer and Business Developer at Radiokinetics. He has lived and worked as an Electrical Engineer in the San Francisco Bay Area for over 20 years. Pete has a broad background in electronics R&D and business development. With positions at Texas Instruments, Burr-Brown, GreenSpring Computers, InVision Technologies and various other enterprises he has developed expertise in ISM wireless technologies, analog signal-chain and data acquisition. Pete received a BS in Electrical Engineering from University of California San Diego in 1991 and has been a Licensed Professional Engineer since 1994. Pete is a master sailor but spends most of his free time with his family.