AFRL-IF-RS-TR-2007-62
Final Technical Report
March 2007

WIRELESS INTRUSION DETECTION

Johns Hopkins University

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory Rome Research Site Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-IF-RS-TR-2007-62 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/	/s/
E. PAUL RATAZZI Work Unit Manager	WARREN H. DEBANY JR, Technical Advisor Information Grid Division Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
MAR 2007	Final	Sep 04 – Oct 06

4. TITLE AND SUBTITLE

WIRELESS INTRUSION DETECTION

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA8750-04-1-0273

5c. PROGRAM ELEMENT NUMBER
62702F

6. AUTHOR(S)
Albert A. Tomko
Christian J. Rieser
Louis H. Buell
David R. Zaret
William M. Turner

5d. PROJECT NUMBER
4519

5e. TASK NUMBER
WI

5f. WORK UNIT NUMBER
D1

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Johns Hopkins University
3400 N. Charles St.
Baltimore MD 21218-2680

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL/IFGB
525 Brooks Rd
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-IF-RS-TR-2007-62

12. DISTRIBUTION AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 07-092

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This report describes a Wireless Intrusion Detection (WIND) system that utilizes physical layer features, derived from individual radio frequency packets, to identify network intrusions. The features considered include those intrinsic to the packet source, as well as those related to the propagation path between the source and a network access point. It is shown that the statistics of a set of packet features can be used to fingerprint each packet source in the network, thereby providing a mechanism for identifying rogue node activity, such as a spoofing attack. Empirical results are presented for IEEE 802.11b networks. Initial test results suggest WIND can achieve a 99% probability of detection with a 10% false alarm rate.

15. SUBJECT TERMS
Wireless networking, intrusion detection, 802.11g, physical layer, PHY, radio frequency, fingerprinting, specific emitter identification

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			E. Paul Ratazzi
U	U	U	UL	138	19b. TELEPHONE NUMBER (*Include area code*)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Section 1

OVERVIEW

1.1 BACKGROUND

Intrusion detection is a particularly important issue for wireless networks because of the inherent non-exclusivity of the transmission medium. Unlike coaxial or fiber optic cables, the wireless medium cannot be secured by physical means, so rogue transmitters may inject packets into the network from great distances and through obstructions. Because the delivery mechanism in packet networks is independent of the source, it is easy to spoof source identification [e.g., Medium Access Control (MAC) address or Internet Protocol (IP) address]. Traditional intrusion detection systems (IDSs) work at the data link layer and above, and thus are unable to distinguish between packets originating from legitimate nodes and packets from an intruder spoofing the identity of a legitimate node.

In 2001, using internal research and development (IRAD) funding, The Johns Hopkins University Applied Physics Laboratory (JHU/APL) began development of a new concept for wireless intrusion detection that utilizes physical layer features to construct an RF fingerprint of each network transmitter. The features are derived from measurements of individual radio frequency (RF) packets transmitted within the network. This approach exploits the transmitter and propagation channel characteristics that are inherently encoded in the electromagnetic wave of each transmitted packet. The packet features are highly dependent on both the individual transmitter's construction and the parameters of the environment between the transmitter and receiver node. The RF fingerprint is thus unique to a particular source and source location. Proof-of-concept testing conducted at JHU/APL demonstrated the viability of this approach, and it was subsequently patented [1]. The use of RF fingerprinting to identify network transmitters will make it much more difficult for an adversary to mimic a legitimate node.

1.2 OBJECTIVE

In response to Air Force Research Laboratory (AFRL) Solicitation #BAA 04-04-IFKA – Wireless Cyber Operations, JHU/APL proposed to develop a wireless intrusion detection (WIND) system using physical layer packet features [2]. The objectives of this research effort were to identifying an optimal set of features, develop fingerprinting and intrusion detection methodologies, and exploring approaches for real-time implementation of WIND, with a focus on IEEE 802.11b wireless local area networks (WLANs) [3]. The WIND Program was funded under a Grant from AFRL, with supplemental funding from the Naval Security Group (NSG). The NSG supplemental funding was directed toward extending WIND to IEEE 802.11g WLANs [4], system modularization, and interfacing with a conventional IDS.

1.3 SCOPE

This report describes research activities undertaken by the staff of JHU/APL to develop a physical-layer IDS for wireless networks, with a focus on the IEEE 802.11b/g Standards. The intrusion detection system utilizes RF features, derived from transmitted network packets, to identify network intrusions.

This report has four main sections. Section 2 describes the WIND architecture and implementation. Section 3 presents a summary of the temporal and statistical characteristics of RF features derived from the numerous experiments conducted in support of the algorithm development. Section 4 outlines a statistics-based WIND algorithm that operates on the features extracted from network RF packets, as well as a notional interface between the WIND system and a traditional IDS. Section 5 provides the major conclusions of this study, and recommendations for future work. Appendix A is a list of references. Appendix B and Appendix C provide a summary of IEEE 802.11b intrusion detection experiments conducted at JHU/APL as part of this effort. Appendix D describes an alternative WIND algorithm, based on clustering techniques, that was explored as part of this effort. Performance and implementations issues lead to the abandonment of this clustering method, in favor of the statistics-based approach.

Some of the material presented here was included in an earlier WIND Progress Report [5]. It is included here to provide a coherent comprehensive record of JHU/APL efforts under the WIND contract.

Section 2

WIND SYSTEM ARCHITECTURE AND IMPLEMENTATION

The purpose of the WIND system is to detect intrusions into a wireless network. Intrusion is initiated by forging the identity of a valid network user, and then transmitting packets containing the forged identity to a wireless access point (AP) in the network. The WIND system uses information collected at the network physical layer to distinguish between valid and forged packets. The following sections describe the WIND system architecture and implementation.

2.1 WIND SYSTEM ARCHITECTURE

The WIND architecture is comprised of several functional blocks. Figure 2-1 illustrates the relationship within the WIND architecture between a WIND sensor, the WIND algorithm, and a traditional IDS. The WIND system is envisioned to supplement traditional network intrusion detection systems, which utilize data collected at higher layers of the network protocol stack. The upper half of the figure shows one sensor, which consists of two or more receive subsystems and a packet processor. Multiple sensors would likely be used in an actual WIND deployment. For example, a WIND sensor might be collocated with each AP in the network.

Each receive subsystem consists of an antenna, a RF threshold detector, a RF down-converter, and a high-speed digitizer. In general, it is desirable to use antennas with different reception characteristics, to ensure that a diverse set of RF features is measured for each packet. For example, one antenna might be omni-directional and another might be directional, or two antennas might have different polarizations (e.g., one vertical and the other horizontal). The received power levels from the antennas will thus be inherently different. Given the signal from the antenna, the threshold detector identifies the presence of a packet, records the time-of-reception, and initiates packet capture through down conversion and digitization. The processor then demodulates and decodes each packet to recover the packet's embedded source identifier, which will be called the internal source identifier. The packet processor also records a set of RF features derived from each packet, such as the received power level, rise-time from the leading edge of the packet, the carrier frequency error and the propagation delay between the times of packet reception at each sensor.

The set of features measured using all sensors constitutes the basic input to the WIND algorithm and are referred to as the feature vector. Each feature vector has an associated time-of-reception and internal source identifier. For example, consider an experiment in which packets are captured by a single sensor with two receive subsystems (i.e., two channel operations), and suppose received power, P, rise-time, T, and frequency error, F, are the three

features measured from each antenna. The feature vector measured at time t_i with internal source identifier S_j is

$$\vec{\xi}_{ij} \equiv (P_1, T_1, F_1, P_2, T_2, F_2)_{ij} \tag{2-1}$$

If more features or more channels are used, the number of feature vector elements is increased accordingly. Moreover, if multiple sensors are employed, the length of the combined feature vector is a multiple of the length of the feature vector generated by a single sensor. It is certainly true that some feature vector elements will be highly correlated. For example, since frequency error is an intrinsic property of the node transceiver, there would seem to be little value in using both F_1 and F_2 in the feature vector. However, as will be shown subsequently, low signal strength can distort intrinsic packet features like frequency error or rise-time, leading to false intrusion alarms. By using multiple measures of each feature derived from different antennas or sensors, one can mitigate the effects of low signal strength and other propagation related impairments.

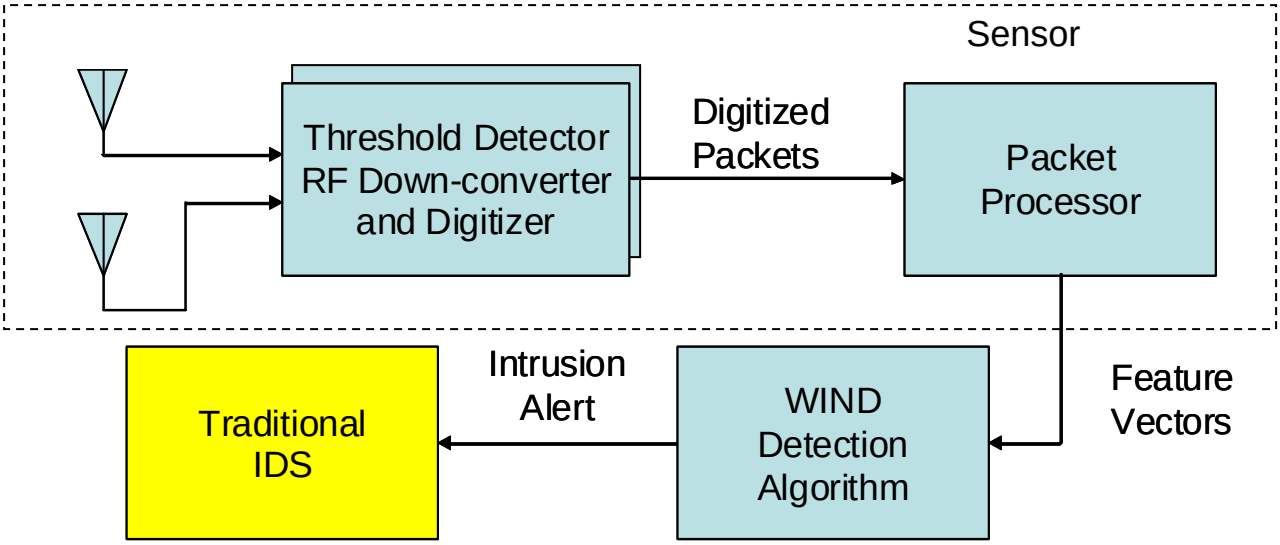

Figure 2-1 WIND Architecture

2.2 WIND SYSTEM IMPLEMENTATION

As part of a previous IRAD effort, JHU/APL developed a system that captures IEEE 802.11b wireless network packets. This initial WIND packet capture system consisted of a custom built two channel RF down-converter module with two pair of inphase (I) and quadrature (Q) analog outputs, a four input high speed digitizer (14 bit, 33 MS/s, Max 80 MS/s) with a controller and high-speed data storage system, two waveform generators, an in-house implementation of a IEEE 802.11b receiver (MATLAB and C) running on a separate PC controller, and a feature extraction engine implemented in MATLAB. This system was used to

conduct the set of experiments described in Appendix B. The data collected by this system were extremely valuable for characterizing the long-time scale (i.e., hours) stability of candidate packet RF features, and for investigations physical-layer intrusion detection methods. Nevertheless, the initial WIND packet capture system had serious limitations. These included a low packet capture rate (about 100 packets per hour per wireless source node), large disjointed components, limited data storage, and no easy migration path to other WLAN types (e.g., IEEE 802.11g), without a significant software development effort.

To address these limitations, JHU/APL developed a new high-speed packet capture system based on modular PXI components [6], and integrated it with a commercial IEEE 802.11b/g software package that provided packet demodulation and decoding, MAC identification and feature extraction. Figure 2-2 shows the new system and a mapping between the WIND functions and implementation.

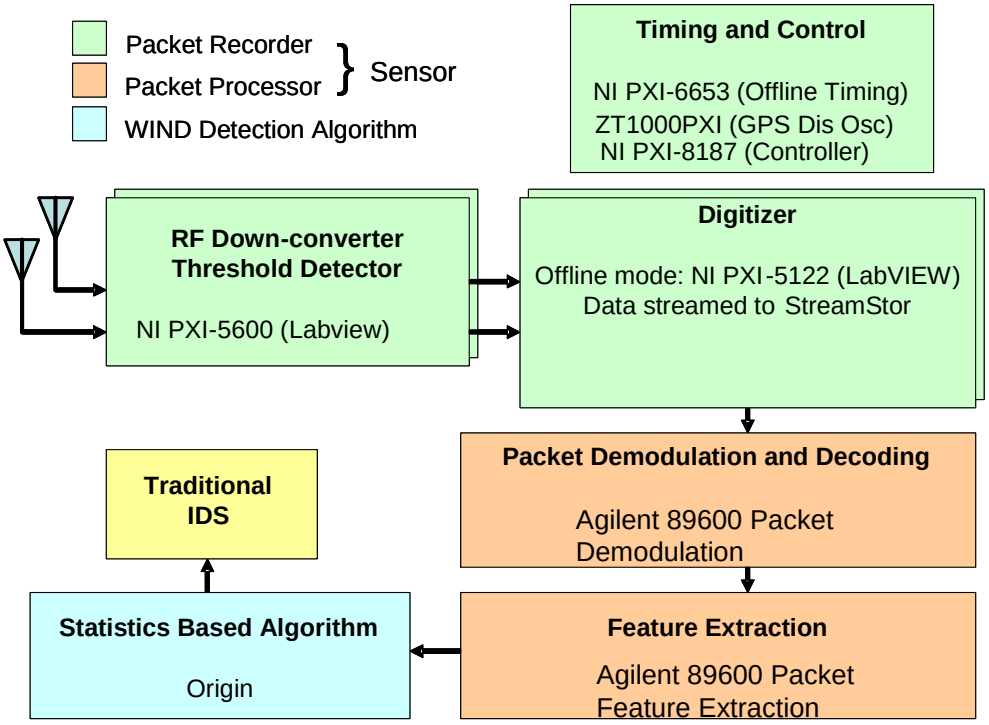

Figure 2-2 Current WIND Implementation

A modular National Instruments (NI) platform was chosen for its small form factor and to facilitate miniaturization and replication of the WIND system functionality. An eighteen slot 3U NI chassis (PXI-1045) houses standard off-the-shelf modular PXI hardware components. The system is small enough that it could be mounted in a field transportable case. The following components were selected to implement WIND functionality.

- 2.7 GHz PXI-5600 RF down-converter replaces the custom RF down converter box.
- PXI-5122 high speed digitizer card (14 bit, 55 MS/s, Max 100 MS/s) and PXI 808 StreamStor card [7] (110 MB/sec) provide the functionality of the large form factor signal digitizer and data archival chassis.
- Two 1U 1.6 Terabyte Conduant StreamStor disk arrays (Big River DM-4 [7]) permit continuous high speed packet capture streaming capability (110 MB/sec).
- A PXI-6653 Timing Module provides the functionality of the two discrete waveform generators. A PXI-5620 high speed digitizer card (14 bit, Max 64 Msps) and a PXI-6608 Timing Module are used for the online mode of the packet capture software.
- A PXI-8187 Embedded Controller provides the functionality of the separate PC controller used in the original JHU/APL system.
- A ZTEC ZT1000 PXI GPS Time, Frequency and Synchronization [8] module provides a frequency standard locked to GPS, and synchronized module timing via the PXI star trigger.

The ZTEC module was not part of the original design of the new packet capture system as described in the WIND Progress Report [5]. It was added to allow multiple packet capture systems (PXI chassis) to be very accurately synchronized (i.e., to within 25 ns). Measurements involving multiple NI packet capture systems, like Experiment 33 described in Appendix C, utilize this capability.

The NI system also has the capability to host select real time operations in a custom field programmable gate array (FPGA). This capability has not yet been utilized as the current WIND system does not operate in real time.

The NI LabVIEW programming environment [6] and programmable nature of the PXI hardware modules facilitate flexible prototyping. A customized LabVIEW application was developed to provide required packet recording functionality. Figure 2-3 shows the WIND User Interface implemented in LabVIEW. It provides for hardware parameter adjustment, packet display and capture, and data post-processing.

Initially, a commercial software package from SeaSolve [9] was selected to provided IEEE 802.11b/g packet demodulation and decoding, MAC identification and feature extraction. Serious deficiencies in this product were identified during system testing (see Section 2.3), and the SeaSolve software was subsequently replaced with a product from Agilent [10].

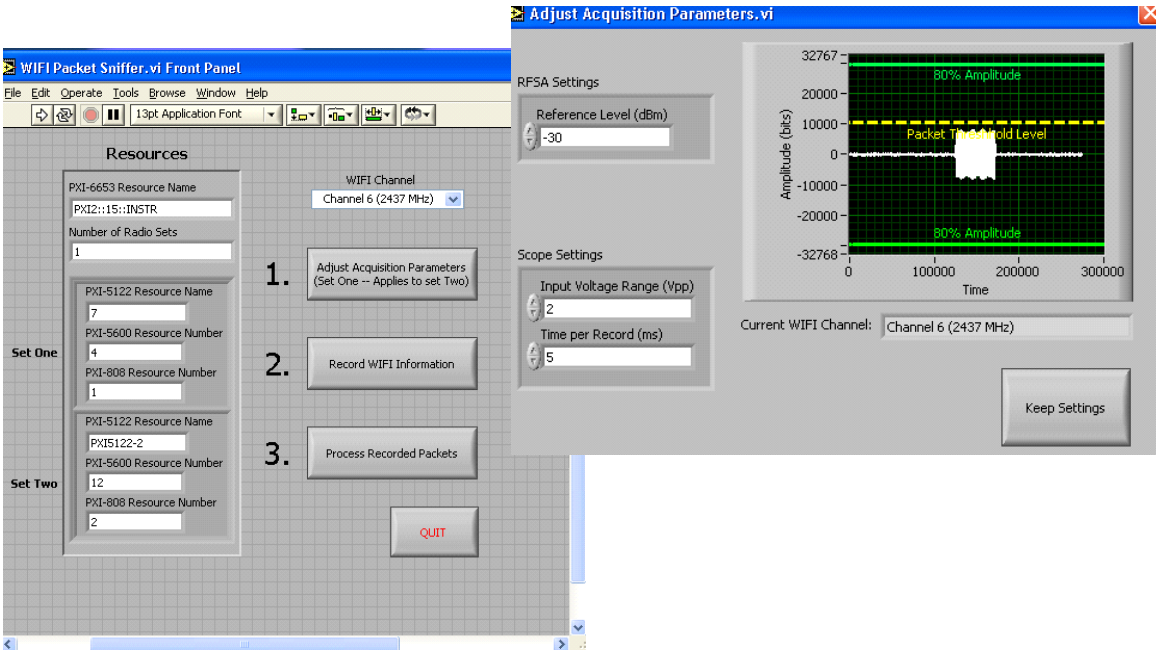

Figure 2-3 WIND Packet Recording Software

Later sections of this report discuss a prototype WIND intrusion detection algorithm developed using the OriginLab Scientific Graphing and Analysis software [11]. A full implementation of the detection algorithm would most likely be done in the C language, and would exploit the well defined API available with the Agilent software to seamlessly mesh the WIND feature extraction and intrusion detection functions.

2.3 SYSTEM EVALUATION

Numerous tests were conducted to evaluate the performance of the NI/SeaSolve system described in the WIND Progress Report [5]. These tests showed that the NI system significantly improved the packet capture rate. While the original JHU/APL system typically captured about 100 packets per hour per wireless network card, the new system typically acquires 50 or more packets per second per card (about one thousand times improvement).

During system testing serious deficiencies were identified in the SeaSolve software. For example, the frequency error (previously called frequency offset [5]) statistics generated by the SeaSolve software were dramatically different than those obtained using the original JHU/APL demodulator and decoder code. Figure 2-4and Figure 2-5 illustrate this discrepancy. They show frequency error time histories for three IEEE 802.11b cards derived for the same set of digitized packets using the APL software and the SeaSolve software, respectively. The mean and variance of the SeaSolve frequency errors are much larger than those obtained using the JHU/APL code. To resolve this discrepancy, the same digitized packet set was run through the Agilent software package [10], with the result shown in Figure 2-6. The Agilent results are in excellent agreement with the JHU/APL software. Another identified

SeaSolve deficiency is that the time-domain data output clipped the leading edge of many packets, leading to inaccurate rise-time estimates. Moreover, the SeaSolve packet processing time was prohibitively long when used in the off-line mode to process archived digitized packets. These deficiencies, and other factors, prompted a switch to the Agilent software. The Agilent software was found to be compatible with the NI output, provided the required 802.11b/g packet demodulation and decoding functionality, and generated the required data for accurate feature extraction. The Agilent software also provides additional outputs related to the intrinsic performance of the packet source. These parameters are being explored as potential new features for source fingerprinting (see Section 3). The Agilent software has a well defined API that can used to create customized feature outputs, and can facilitate integration of the demodulation, decoding and feature extraction functions with the WIND intrusion detection algorithm. Finally, similar Agilent software is available for IEEE 802.16 and most other existing or emerging wireless communication standards. This could aid in the migration of WIND technology to other wireless environments.

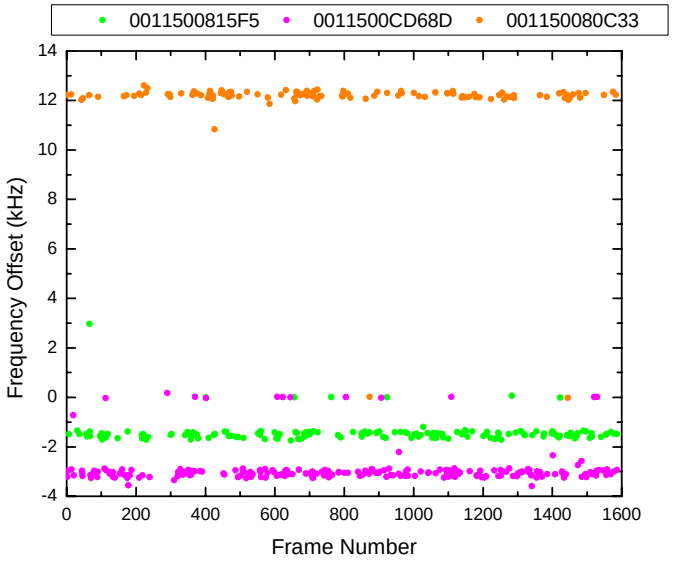

Figure 2-4 Frequency error from JHU/APL software

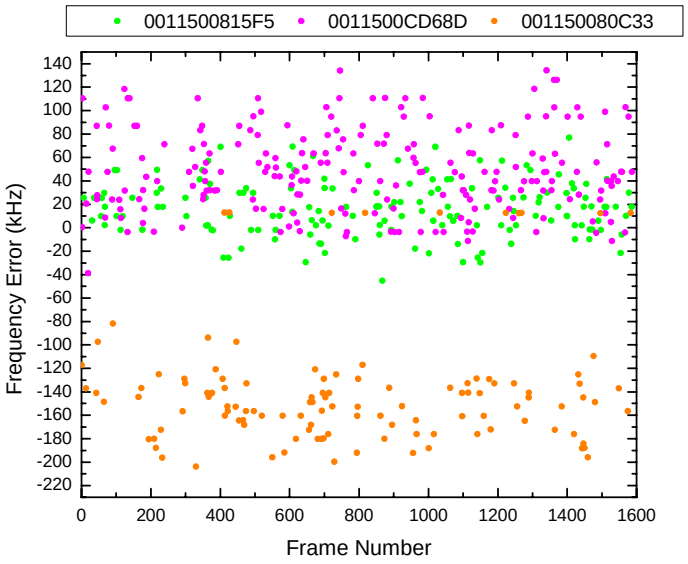

Figure 2-5 Frequency error from SeaSolve software

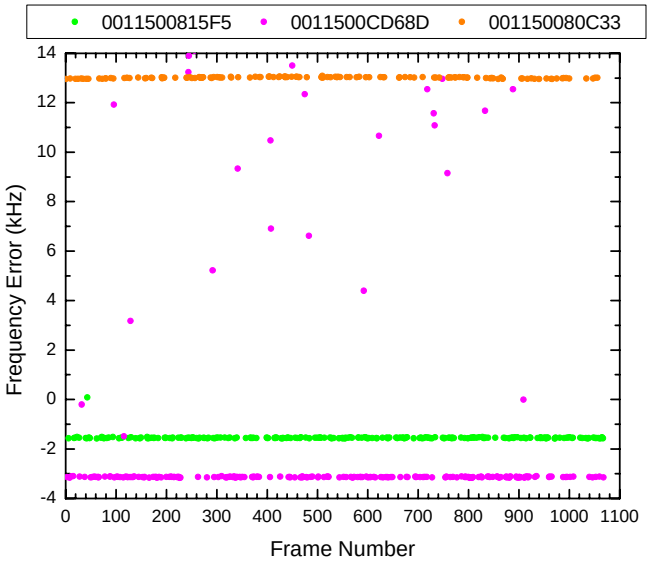

Figure 2-6 Frequency error from Agilent software

2.4 FEATURE EXTRACTION

In the new WIND system feature extraction is performed by importing 5 ms IF data segments into the Agilent 89600 VSA Software environment [10], and demodulating the first valid IEEE 802.11b/g packet detected in each segment. Figure 2-7 illustrates this process. It shows four Agilent display panels for a typical packet decode. These include the search window (upper left), which shows the imported 5 ms data segment, the detected leading edge of the packet (upper right), the packet power spectrum (lower left) and the symbol table (lower right). The lower section of the symbol table provides the Physical-Layer Service Data Unit (PSDU) payload symbol data bits for the demodulated packet, while the upper section gives the modulation quality error data results and Physical Layer Convergence Procedure (PLCP) header information for the packet. The MAC address of each detected packet was extracted from the PSDU payload symbol data bits. A VBScript macro [10] was used to automate the process of importing each data segment and writing out the feature data.

2.4.1 PACKET RISE-TIME

Figure 2-8 illustrates the process used to compute packet rise-time in the new WIND system. The blue curve on the lower graph of this figure is a sample of the IF voltage amplitude at the leading edge of a packet as detected by the Agilent 89600 VSA Software. These data were first low-pass filtered to identify the pulse envelope (red curve on the lower graph). Then the first and second time derivatives of the voltage envelope are computed (center and upper graphs). The packet rise-time taken as the time difference between the extrema of the second derivative. Observed rise-times were typically about 1 μs. The rise-time data provided in Appendix C were generated using this procedure. The original JHU/APL packet capture system used in our initial experiments (Appendix B) used a somewhat different procedure to determine rise-time, and the rise-time was recoded in terms of number of samples (i.e., at 33 MS/s sampling rate, 1 sample = 0.03 μs).

2.4.2 RECEIVED POWER

In the new WIND system, the received power level for a given packet was computed by direct integration of the packet power spectrum (lower left window on Figure 2-7) over a 22 MHz bandwidth about the band center frequency. The received power levels (in dBm) include the gain of a low-noise amplifier at the antenna output, which ranged from 10 to 25 dB, depending on the particular antenna used. The original JHU/APL packet capture system computed received power was derived from a subset of the time-domain I Q pairs following the turn-on transient, and power levels were recorded in dB relative to an arbitrary fixed reference level.

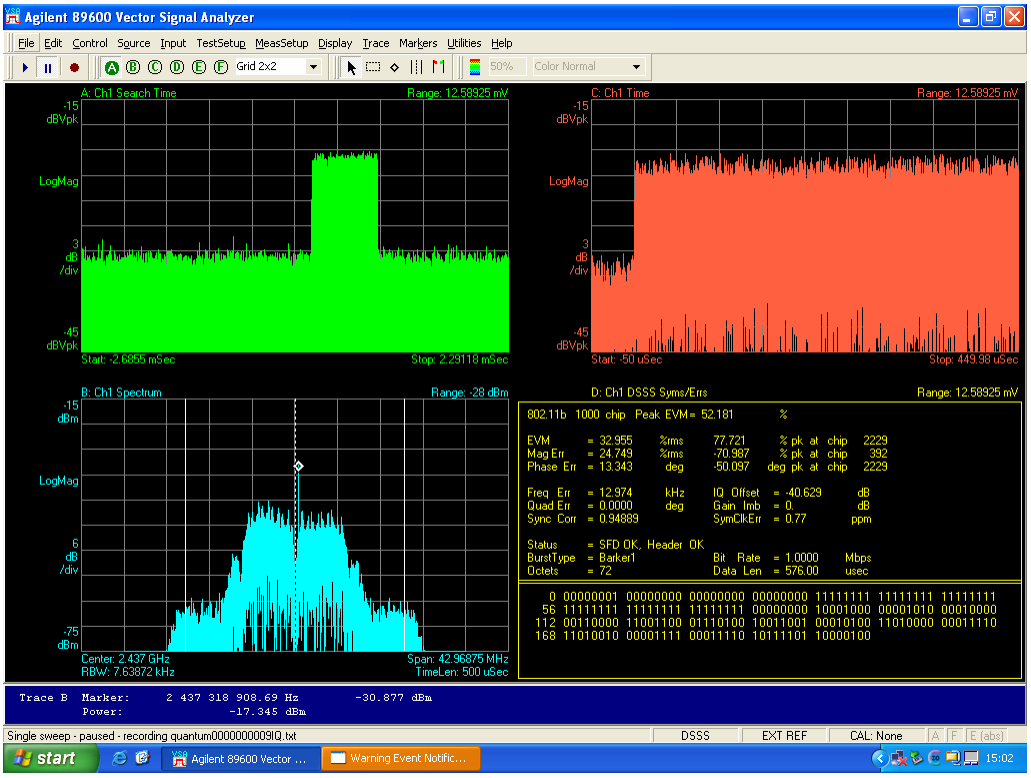

Figure 2-7 Agilent 89600 Packet Demodulation Displays

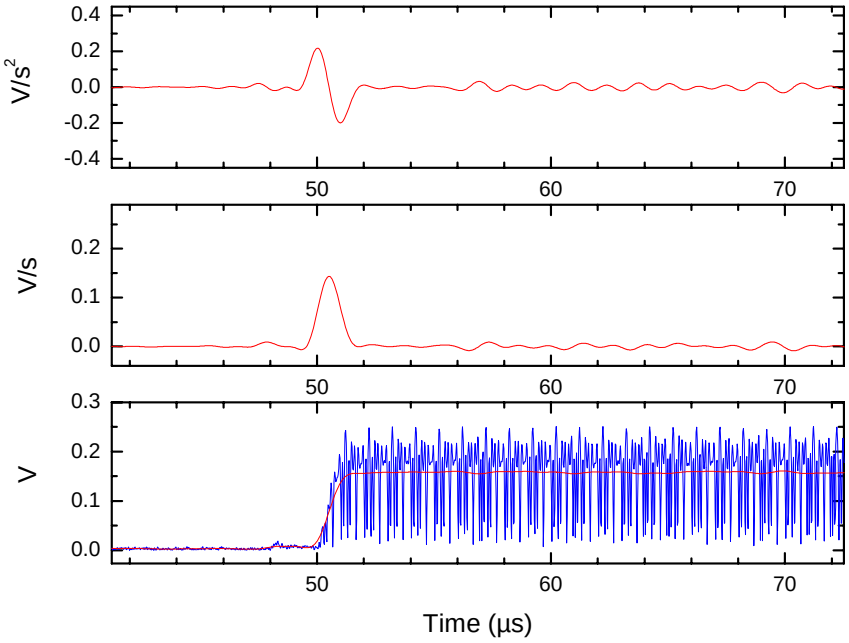

Figure 2-8 Rise-Time Calculation Procedure

2.4.3 FREQUENCY ERROR AND RELATED PACKET INFORMATION

The Agilent 89600 VSA Software [10] outputs frequency error directly. The following additional parameters are available for each detected packet:

I Q Offset - the magnitude of the carrier feed-through signal. When there is no carrier feed-through, IQ offset is zero (-infinity dB).

I Q Quadrature Error - indicates the orthogonality error between the I and Q signals. Ideally, I and Q are 90 degrees apart. A quadrature skew error of 3 degrees means I and Q are 87 degrees (or 93 degrees) apart.

I Q Gain Imbalance - compares the gain of the I signal with the gain of the Q signal.

S**ynchronization Correlation -** cross correlation of the preamble of the measured signal with an ideal Barker sequence. It is a figure of merit indicating the quality of the PLCP preamble synchronization data fields. A value of 1 indicates perfect correlation and a value of 0 no correlation.

Error Vector Magnitude (EVM) - a measure of the quality of the modulation on a signal derived by comparing the transmitted signal relative to a perfect theoretical signal. EVS is highly correlated with BER (bit-error-rate).

802.11b 1000 Chip Peak EVM - the normalized peak EVM over 1000 chips (see Section 18.4.7.8 “Transmit modulation accuracy” in Reference 2).

Header Status - PLCP Header status string indicating the status of the PLCP Header. It indicates if the header is valid or if an error condition exists, such as invalid header bits, a header CRC failure, or an invalid header phase shift.

MAC Status - status of the PSDU MAC frame: whether the frame check sum passes or fails.

Burst Type - the detected spread sequence code scheme/PLCP PPDU format type:

0	Barker 1	11 Chip Barker	1 Mbps	1 b/S	DBPSK
1	Barker 2	11 Chip Barker	2 Mbps	2 b/S	DQPSK
2	CCK 5.5	8 chip CCK	5.5 Mbps	4 b/S	DQPSK
3	CCK 11	8 chip CCK	11 Mbps	8 b/S	DQPSK

where CCK denotes complementary code keying.

Bit Rate - the modulation data rate as detected in the PLCP Header.

EVM Peak - the worst-case peak EVM expressed as a percentage of the square root of the mean power of the ideal signal.

Octets - the number of decoded Octets in the PSDU computed from the PLCP header data.

Data Time Length - the time length of the PSDU (payload data) in seconds as detected in the PLCP Header.

Symbol Clock Error - the difference between the ideal and actual symbol clock frequency in parts per million (ppm).

EVM Peak Location - the number of the chip with the worst-case peak EVM.

Magnitude Error - the vector magnitude difference between the I/Q measured signal and the I/Q reference signal at the measured chip time expressed as a percentage of the square root of the mean power of the ideal signal.

Magnitude Error Peak - the worst-case peak magnitude error value expressed as a percentage of the square root of the mean power of the ideal signal.

Mag Error Peak Location - the chip time with the worst-case peak magnitude error.

Phase Error - the phase difference between the I/Q reference signal and the I/Q measured signal measured at the chip time.

Phase Error Peak - the worst-case peak phase error value.

Phase Error Peak Location - the chip time with the worst-case peak phase error value.

All of these parameters are stored for each recorded packet, along with the main feature vector components: time-of-detection, MAC address, rise-time, received power and frequency error.

Section 3

EXPERIMENTAL FINDINGS

Many key design considerations must be understood to develop a wireless intrusion detection scheme. These considerations include the number of sensors required; the number, orientation, type (e.g., directional, omni-directional) and polarization of the monitoring antennas; the number and type of electromagnetic characteristics in the feature vector; the required length of the packet sequences; and the appropriate classification techniques. To address these many factors, JHU/APL conducted a series of experiments to collect samples of packets from a variety of physical environments with different network node densities (number of nodes) and spatial distributions. The resultant library of observations was used to study the temporal and statistical characteristics of candidate RF features, and to evaluate candidate intrusion detection algorithms.

3.1 EXPERIMENT OVERVIEW

Table 3-1 provides an overview of the thirty-four experiments conducted by JHU/APL in support of the WIND effort. All of the experiments used wireless network cards configured as IEEE 802.11b nodes [3]. For all experiments, all nodes operating in Ad Hoc mode [3], except Experiment 34, which employed infra-structure mode with several access points. The first twenty-four experiments were conducted with the original JHU/APL data collection system, and are described in Appendix B. The remaining ten experiments were collected with the new data collection system (see Figure 2-2) and are described in Appendix C.

The following sections discuss the variability observed in experimental data collected and key observations that help shape the algorithm development process. Section 3.2 discusses the temporal variability of feature elements, using Experiment 21 as an example. Section 3.3 describes a statistical characterization of the experimental data. Section 3.4 provides an example of a process for physical layer intrusion detection, utilizing data from Experiment 6. This process provides the foundation for the algorithm presented in Section 4.

3.2 TEMPORAL VARIABILITY

The values of the feature vector elements for a given internal source identifier will vary with time due to propagation effects, motion of the network nodes or other objects, and other environmental factors such as oscillator drift in the node electronics. Figure 3-1 through Figure 3-3 illustrate this point. They show the temporal variation of the six feature vector elements defined in Equation (2-1) for eight stationary nodes measured over 16 hours. Data plotted in these figures was collected in Experiment 21. Channel 1 of the sensor utilizes an

omni-directional antenna (vertical probe), while Channel 2 employs a directional antenna (parabolic). Each node is a laptop computer equipped with a PC Card WiFi (IEEE 802.11b) transceiver. The internal source identifier is the MAC address, which is shown at the top of each figure. WiFi cards from three different manufacturers are present in the network. These data show that, in general, each observed element of the feature vector is a non-stationary random process characterized by one or more central values.

Table 3-1 Overview of Experiments Conducted in Support of the Algorithm Development

Experiments	Antennas	Experiment Focus
1-8	Dual Polarization Horn (Channels 1 and 2)	Different numbers of intruders with staggered activity in a network made up of the same brand of cards
9-24	Polarized Parabolic (Channel 1) Omni-directional (Channel 2)	Larger mixed network consisting of cards from several manufacturers. Various scenarios including: -intruder nodes with vertically or horizontally polarized antennas, placed nearby the sensor as well as in adjacent rooms, -test cases to verify correct operation of the test bed, and -experiments to observe the long-term variation of RF features
25, 28, 29, 32	Dual Polarization Horn (Channels 1 and 2)	Large (24 node) compact and dispersed geometries; short-time scale feature stability and uniqueness
26, 27, 39, 31	Polarized Parabolic (Channel 1) Omni-directional (Channel 2)	Large (24 node) compact and dispersed geometries; short-time scale feature stability and uniqueness
33	Orthogonal Horizontally Polarized Horns (Channels 1 and 2) Omni-directional (Channel 3)	Multiple monitors, dispersed geometry; short-time scale feature stability and uniqueness
34	Omni-directional (Channel 1) Horizontally Polarized Horn (Channel 2)	Dispersed geometry; indoor and outdoor environments; short-time scale feature stability and uniqueness

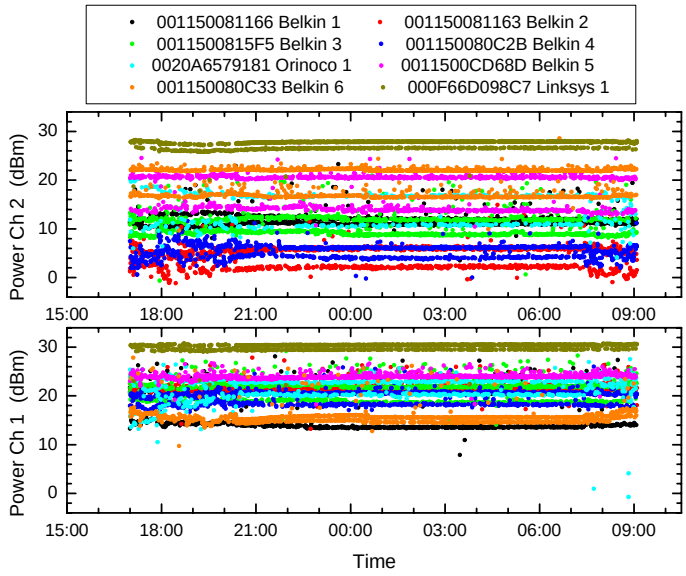

Figure 3-1 Temporal Variation of Received Power

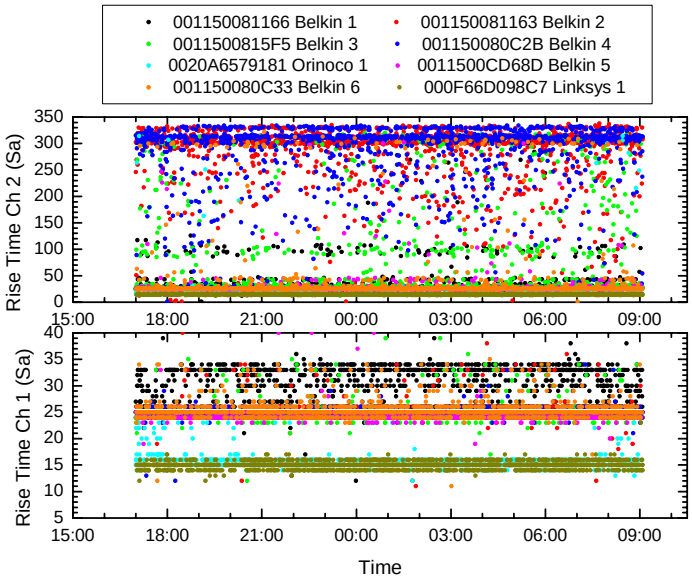

Figure 3-2 Temporal Variation of Rise Time

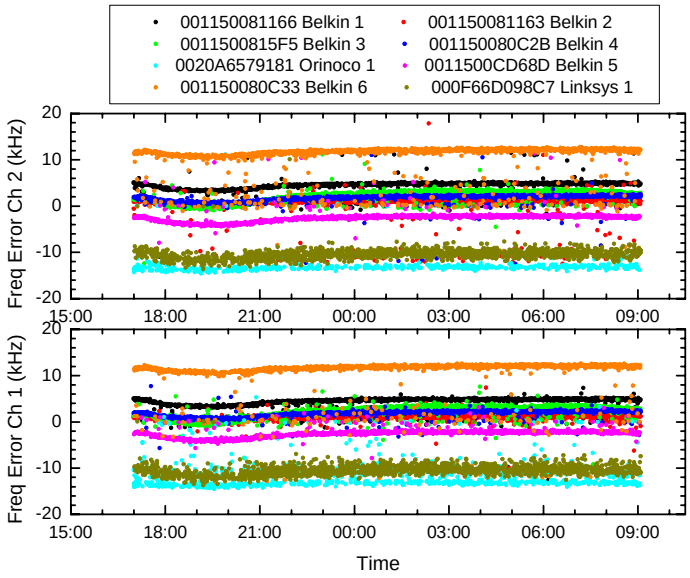

Figure 3-3 Temporal Variation of Frequency error

3.3 STATISTICAL CHARACTERIZATION

Consider a sequence of feature vectors of length $\Delta t = t_{i+L} - t_i$ for the J-th internal source identifier:

$$\{\vec{\xi}_{iJ}, \vec{\xi}_{(i+2)J}, ..., \vec{\xi}_{(i+L)J}\}. \quad (3\text{-}1)$$

The available empirical data suggest that, if Δt is small (e.g., L <100 packets), the variation of each feature vector element can be treated as a stationary distribution that drifts slowly with time. Figure 3-4, Figure 3-5 and Figure 3-6 show cumulative probability distributions (CPDs) constructed from the data presented in the previous section. The solid curves were each formed from a representative short sequence of packets. The symbols β and Ω denote the lower and upper limits of the CPD temporal variation over the 16 hour period. From these and the previous figures one can make the following observations:

1) over any small interval Δt each element of the feature vector has a unique statistical distribution characterized by one or more prominent central values (modes);
2) for a given feature and time interval, the feature distributions of some internal source identifiers may overlap (e.g., Belkin cards 2 and 4 in the WiFi example; see Figure 3-6); however,
3) if the number of feature vector elements is sufficiently large and diverse, one can uniquely identify each source using the measured mode characteristics (e.g., number of modes, central values, etc.); and

4) the presence of forged packets within a given Δt will be evident as a change in the mode characteristics relative to those observed at earlier times (e.g., a change in the number of modes observed for one or more feature vector elements).

The later two observations provide a basis for verifying the authenticity of the internal source identifier and of flagging network intrusions. The next section provides an example that further illustrates these points.

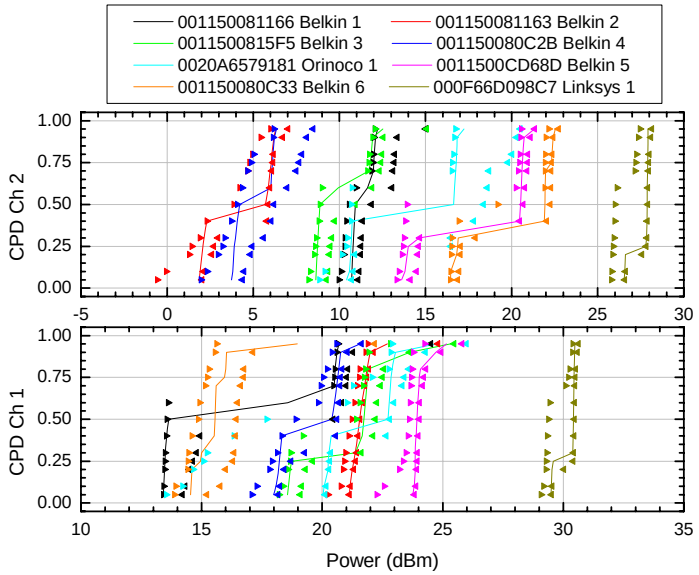

Figure 3-4 Received Power CPDs

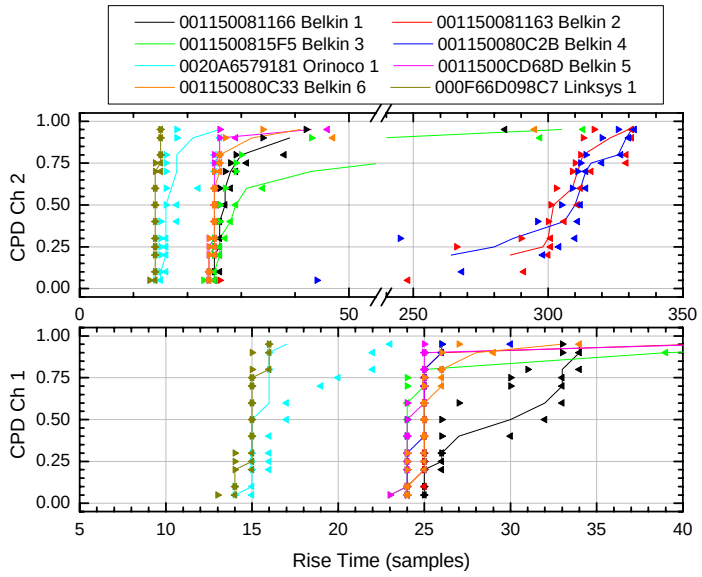

Figure 3-5 Rise Time CPDs

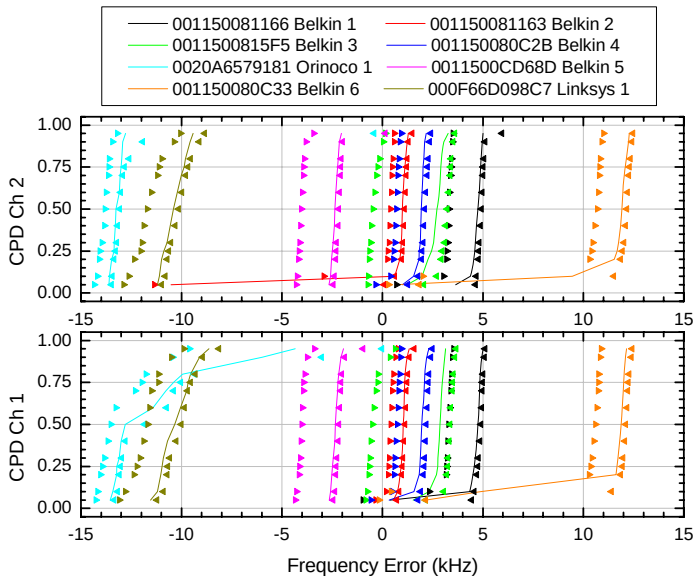

Figure 3-6 Frequency error CPDs

3.4 AN EXAMPLE OF PHYSICAL LAYER INTRUSION DETECTION

Figure 3-7 through Figure 3-12 show probability distributions constructed from a small number of packets ($L \approx 20$) during a simulated network intrusion event. Data was derived from Experiment 6. During the intrusion event a Belkin WiFi node attempts to spoof the MAC address of a Linksys WiFi node. This case is denoted by the red shaded distributions on each figure. Data from a second Linksys WiFi node (without spoofing) are shown for comparison (blue shading). Probability distributions are presented for three features: frequency error (Figure 3-7 and Figure 3-8), received power (Figure 3-9 and Figure 3-10) and packet rise-time (Figure 3-11 and Figure 3-12). For each pair of figures, the first figure shows the feature distributions prior to the spoofing attempt, and the second figure gives the feature distributions during the intrusion. The distributions of each feature display one or more principal peaks (modes), and frequently one or more secondary peaks. The characteristics of the principal modes for all measured features provide the unique identifier (external source identifier) for each node in the network. The measurable characteristics of each mode include the central (peak) value, the distribution spread about the central value, and the area under each peak distribution. The principal modes of a given feature may overlap (e.g., the rise-time distributions for the two Linksys cards on Figure 3-11), and it is plausible that in a dense network with many similar nodes two or more nodes may have a number of similar mode characteristics. However, based on the available empirical data, it seems highly unlikely that all of the characteristics of all of the principal modes for all the measured features will be nearly the same for any two nodes. Moreover, the length of the feature vector may be made arbitrarily large by combining the measurements of several WIND sensors. In the example shown here, the two Linksys cards are

easily distinguished by the principal modes of their frequency error and received power distributions (Figure 3-7 and Figure 3-9, respectively). The pairs of figures for the control node (blue shading) show that the distributions evolve with time, but the locations of the principal modes do not change significantly from one small sample of packets to the next (i.e., the modes are quasi-stationary; Figure 3-1, Figure 3-2 and Figure 3-3 also illustrate this point). The introduction of spoofed packets (included in the red shaded distributions on Figure 3-8, Figure 3-10 and Figure 3-12), significantly alters the number of principal modes observed for a given feature. Therefore, the WIND intrusion detection algorithm must track the temporal evolution of the characteristics of the principal modes of each element of the feature vector, looking for anomalies in the expected mode characteristics based on the immediate past (previous feature vector sequences). The secondary peaks in one or more feature probability distributions can lead to false alarms, so filtering and fitting techniques are used to map the raw multi-modal probability distributions into a better-defined set of Gaussian distributions.

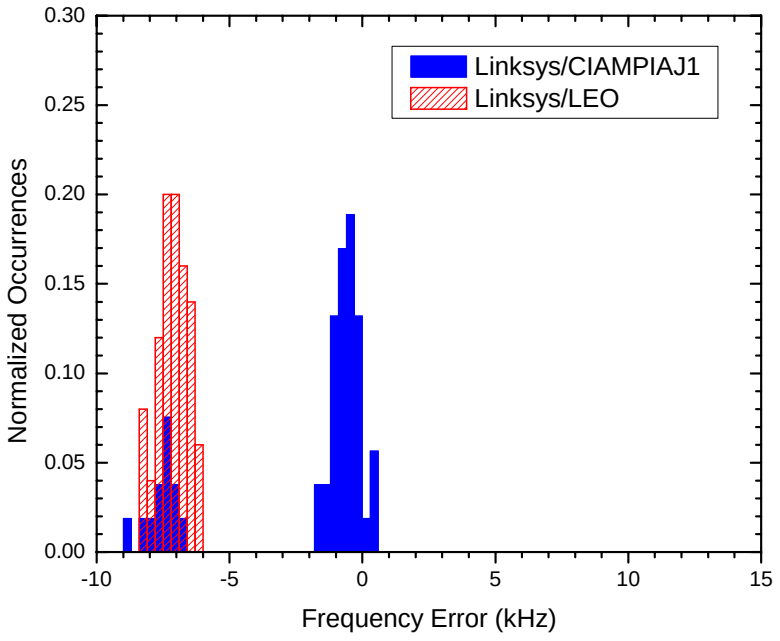

Figure 3-7 Frequency error Distributions Prior to Intrusion

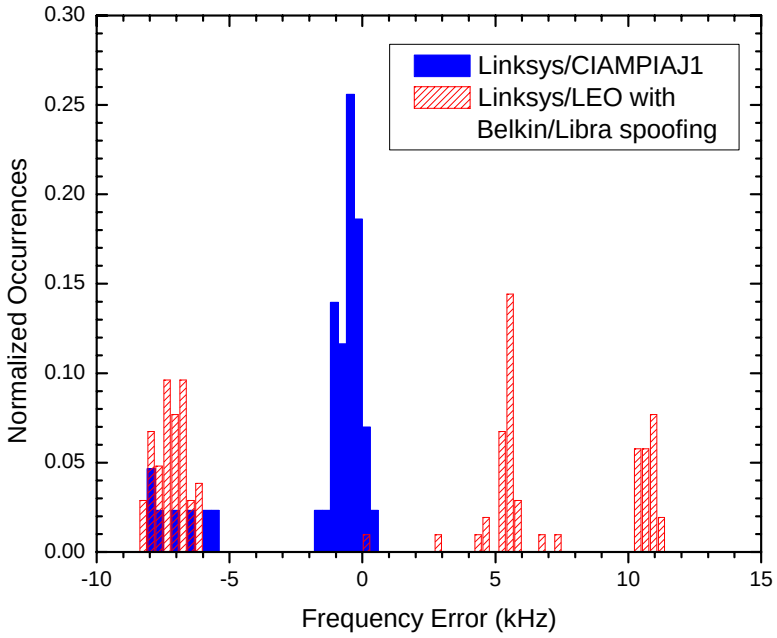

Figure 3-8 Frequency error Distributions During Intrusion

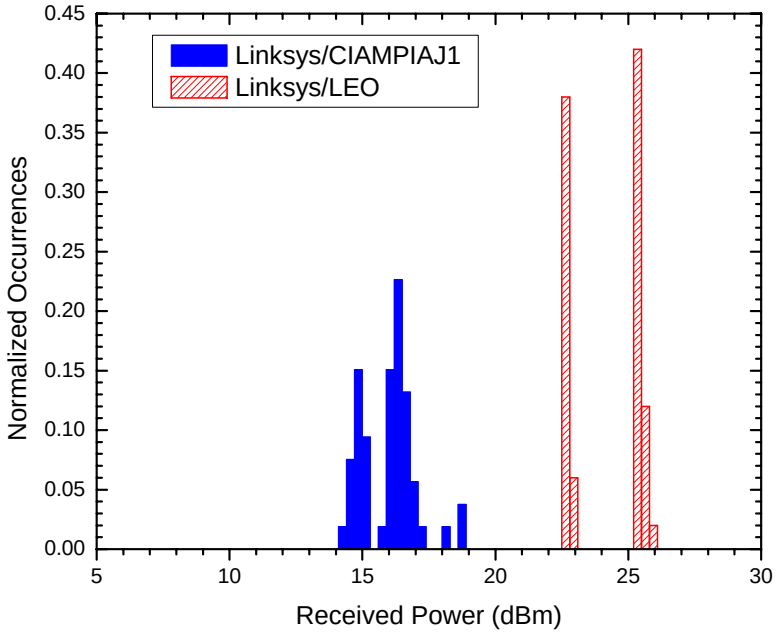

Figure 3-9 Received Power Distributions Prior to Intrusion

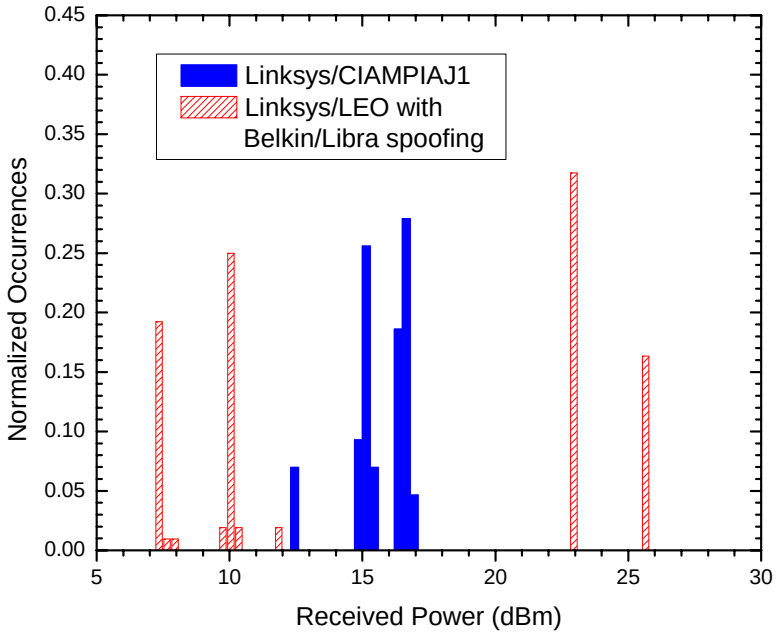

Figure 3-10 Received Power Distributions During Intrusion

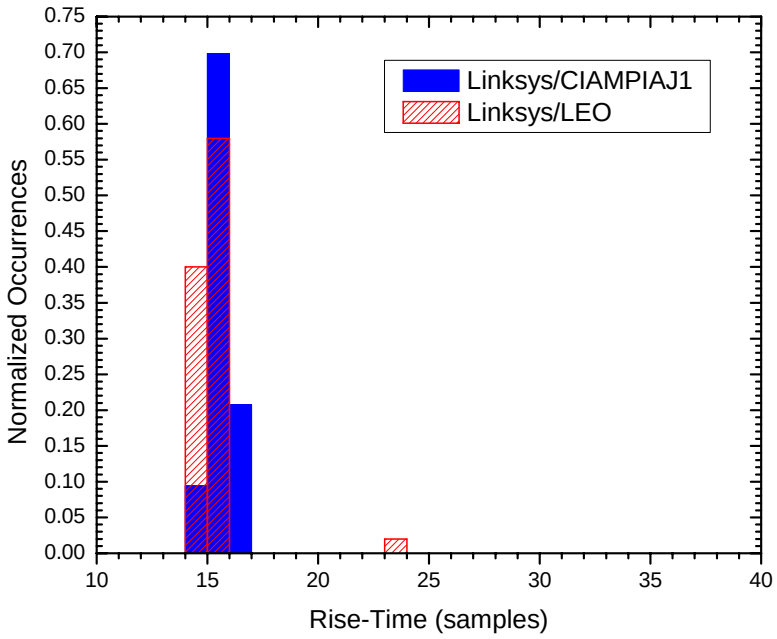

Figure 3-11 Packet Rise-Time Distributions Prior to Intrusion

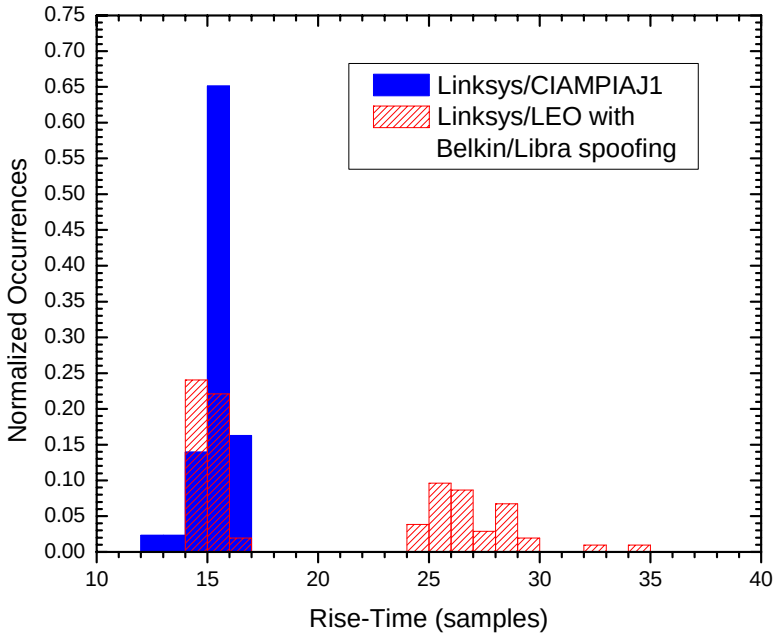

Figure 3-12 Packet Rise-Time Distributions During Intrusion

Section 4

DEVELOPING A WIND DETECTION ALGORITHM

Using insight gained from the observations detailed in the previous section, an algorithm was developed that tracks the statistically evolution of the feature space and identifies intruder activity within tens of packets[1]. The approach maps a sequence of feature vectors into a unique identifier, called an RF fingerprint. The RF fingerprint will evolve with time due to the non-stationary nature of the feature vector. A method for tracking and predicting the temporal evolution of the RF fingerprint is described and pseudo-code is presented that implements a WIND detection algorithm. A method to interface the WIND system with a traditional IDS is also proposed.

4.1 FEATURE MAPPING

A process for mapping a sequence of feature vectors into a unique identifier (RF fingerprint) is described in this section. The sequence of feature vectors is a set of sequences spanning all elements of the feature vector. For example, for the feature vector and feature vector sequence defined in Equations (2-1) and (3-1), there are six sequences of feature vector elements for each internal source identifier:

$$\begin{aligned} &P_{11}, P_{12}, \ldots, P_{1L} \\ &T_{11}, T_{12}, \ldots, T_{1L} \\ &F_{11}, F_{12}, \ldots, F_{1L} \\ &P_{21}, P_{22}, \ldots, P_{2L} \\ &T_{21}, T_{22}, \ldots, T_{2L} \\ &F_{21}, F_{22}, \ldots, F_{2L} \end{aligned} \tag{4-1}$$

Each sequence is mapped in the same way using a four-step process consisting of binning, filtering, peak detection and fitting. Figure 4-1 provides an illustration of this process for a sequence of *P* data. The process implementation used in this example is based on methods contained in OriginLab Peak Fitting Module (PFM) Version 7 [11].

[1] As part of this research, JHU/APL explored multiple approaches to WIND algorithm development. An algorithm based on clustering techniques was initially pursued. Based on the experimental findings, it became evident that the algorithm would not provide adequate performance without significant additional research. Results from this approach can be found in Appendix C.

First, a normalized histogram is constructed from the data sequence. The choice of bin width is fixed for each feature type. For WiFi data sets, the applicable bin widths are 0.1 dB for *P* data, 1 sample for *T* data and 150 Hz for *F* data. The number of counts in each bin is normalized through division by the sample size, *L*. The choice of sample size is a compromise. If *L* is too small, the histogram will be poorly defined; but if *L* is too large, false modes or mode broadening may appear due to the non-stationary nature of the data (i.e., if the modes drift significantly over the duration of the sample). $L \approx 20$ is used here. The equivalent Δt will be determined by the packet rate and the efficiency of the data collection system in acquiring packet feature vectors.

The second step is to smooth the histogram using a low pass Fast Fourier Transform (FFT) filter. This step helps remove small fluctuations in the histogram that might be interpreted as additional modes. The smoothing is accomplished by removing Fourier components with frequencies higher than $1/(N\,\Delta)$, where *N* is the number of bins considered at a time and Δ is the bin width. *N* values of 3 to 5 have been found to produce good results.

The third step is peak detection. There are numerous ways to implement this step. The PFM approach searches the smoothed second derivative of the filtered data looking for peaks that are above the standard deviation of the filtered data.

The fourth step is to fit each peak to a Gaussian model using the Levenberg-Marquardt method. The Gaussian model is of the form:

$$y = y_o + \frac{A}{w}\sqrt{\frac{4\ln(2)}{\pi}}e^{-4\ln(2)\frac{(x-x_c)}{w^2}}, \tag{4-2}$$

where

x_c – central value (mode),
A – area under the Gaussian curve,
y_o – baseline, and
w – full width of Gaussian curve at half maximum.

A maximum of 10 iterations are used to perform the fit. If the relative change of the reduced chi-square value between two successive iterations is less than the tolerance value (0.05), then no more iterations are performed. No weighting is used, and the baseline is fixed at $y_o = 0$.

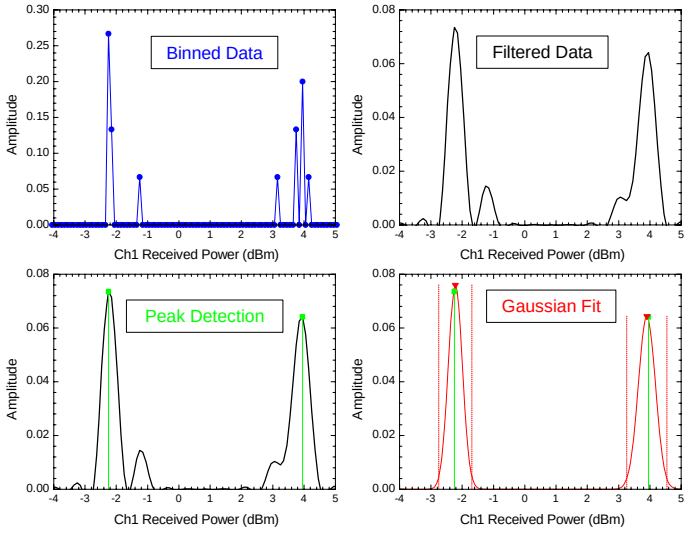

Figure 4-1 Illustration of the Feature Mapping Process

Each sequence of feature vector elements is thus mapped into one or more principal modes, and each mode is characterized by the set of fit parameters: (x_c, A, w). In the Figure 4-1 illustration, the red inverted triangles on the Gaussian Fit graph denote the values of x_c for the two observed modes, and the red dashed lines specify the effective mode boundaries $x_c \pm w$. The RF fingerprint is the superset of fit parameters (mode characteristics) encompassing all modes of all feature vector elements.

The RF fingerprint will evolve with time due to the non-stationary nature of the feature vector. A method for tracking and predicting the temporal evolution of the RF fingerprint is described in the next section.

4.2 RF FINGERPRINT EVOLUTION

Figure 4-2 shows the temporal evolution of one component of a feature vector: F_1, the frequency error on Channel 1. Measured feature vectors were grouped into fixed time blocks Δt with $L \approx 20$. The resultant probability distributions are plotted as a color-coded contour surface with independent variables of frequency error and time block index, k. The absolute time associated with each block is

$$t_k = t_o + (2k-1)\Delta t/2, \tag{4-3}$$

where t_o is the absolute time when the first feature vector is collected.

The non-stationary nature of this feature vector component is obvious. The data cluster around a single modal peak, whose central value changes with time. Some outliers (secondary peaks in the probability distribution) are also apparent (isolated cyan clusters against the blue background). The feature mapping process attempts to track the modal peak while removing the outliers.

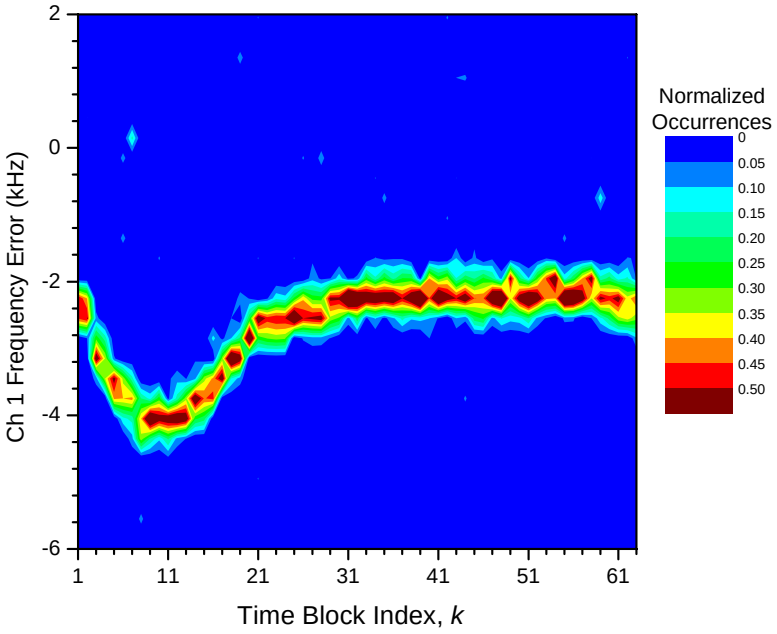

Figure 4-2 Temporal Evolution of Frequency Error

Figure 4-3 shows the result of the feature mapping process when applied to the data given in Figure 4-2. The central values $x_c(t_k)$ are plotted as black dots with error bars spanning $\pm w(t_k)/2$. A single mode is found in each time block, except at block $k = 7$, which is bimodal. Here, one of the secondary peaks is falsely interpreted as a principal mode. If only a single feature vector component were used for intrusion detection, this would lead to a false alarm condition. However, with a multi-element feature vector, one can prevent false alarms of this type by insisting that the modal characteristics of two or more feature vector elements be outside the expected range before an alarm is issued. Figure 4-4 illustrates this point. It shows the time evolution of all three mode characteristics of the Gaussian model for two elements of the feature vector: the frequency errors observed on Channels 1 and 2 of the WIND sensor. The Channel 1 data show outliers for all three mode characteristics at block $k = 7$, and for one of the mode characteristics ($w/2$) at block $k = 44$. By contrast, the Channel 2 data have no significant outliers. If one insists that both feature vector elements must have at least one outlier mode characteristic in the same time block, the likelihood of a false alarm is significantly reduced.

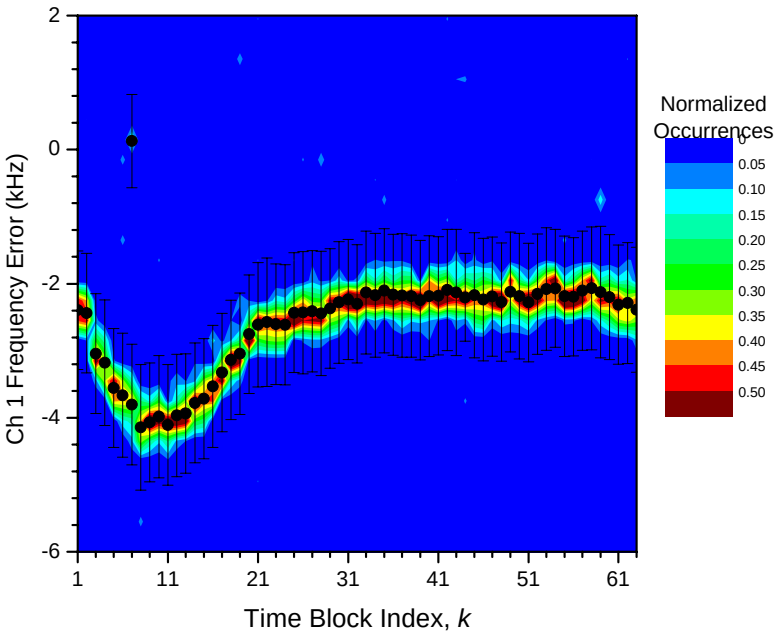

Figure 4-3 Feature Mapped Frequency Error

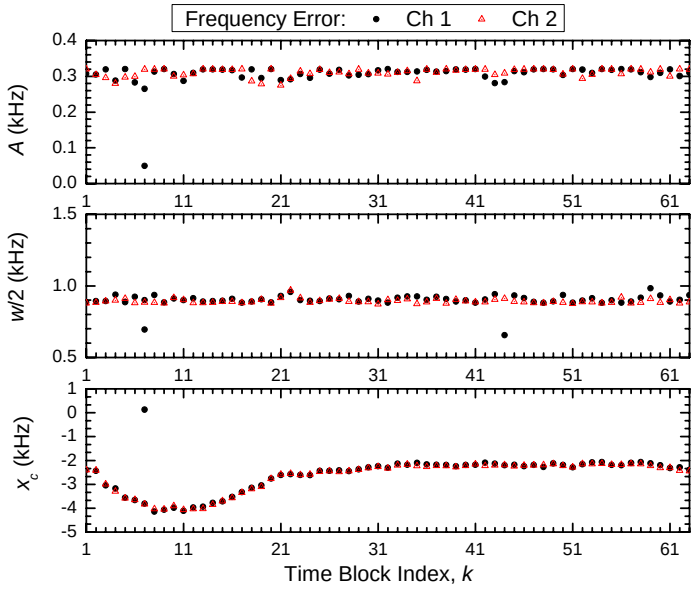

Figure 4-4 Modal Characteristics of Frequency Error

Figure 4-5 through Figure 4-7 show the temporal evolution of the RF fingerprint for an intrusion scenario in which WiFi node Belkin 3 (see Figure 3-1 through Figure 3-6) transmits forged packets with the internal source identifier (i.e., MAC address) of node Belkin 1. The intrusion begins during time block 21 and ends during time block 36. Since both WiFi cards are from Belkin, and rise-time is most sensitive to implementation differences between manufacturers, only frequency error and received power are considered. Figure 4-5 shows the temporal variation of $x_{cp}(t_k) \pm w_p(t_k)$ for all detected modes p of the frequency error on Channels 1 and 2. Figure 4-6 gives the corresponding variation of the area characteristic $A_p(t_k)$. Figure 4-7 shows the temporal variation of $x_{cp}(t_k) \pm w_p(t_k)$ for all detected modes p of the received power on Channels 1 and 2. The modes are color coded such that the mode with the smallest value of x_{cp} is colored black, the next lowest is red, and then green, blue and cyan (see Figure 4-7 for the entire coloring sequence). Although $A_p(t_k)$ for the frequency error gives a clear indication of the period of intrusion in this case, where only two modes are present, it is generally of limited utility, particularly when more than two modes are present. Consequently, the intrusion detection process only uses x_{cp} and w_p. Figure 4-5 and Figure 4-7 are used below to illustrate the process. The process assumes packet feature vectors are sorted by internal source identifier, accumulated over a fixed time block Δt, and transformed into a set of mode characteristics, as described above. The process is executed for each internal source identifier.

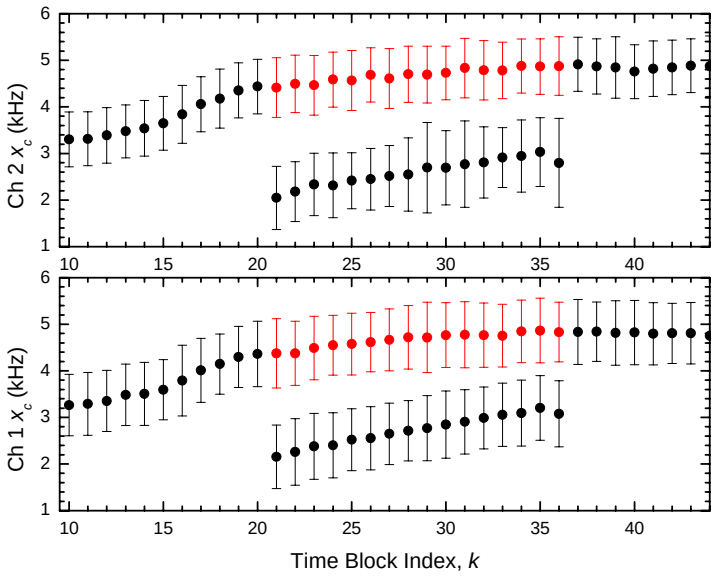

Figure 4-5 Intrusion Detection Using Frequency Error Data

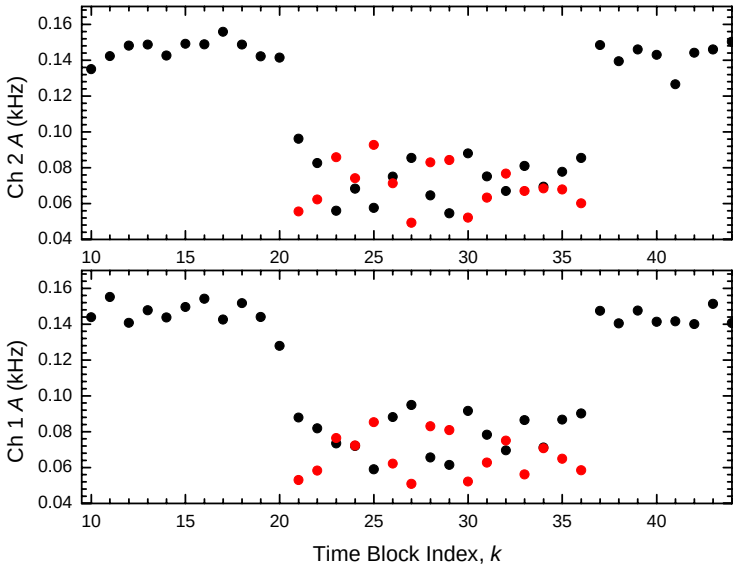

Figure 4-6 Intrusion Detection Using Frequency Error Data

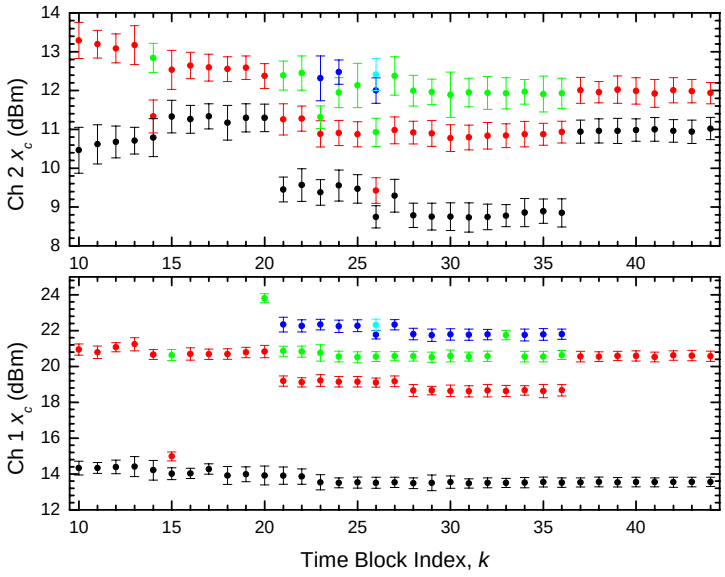

Figure 4-7 Intrusion Detection Using Received Power Data

Let $M_n(t_k)$ denote the number of modes detected in time block t_k for the n-th component of the feature vector. Define δ_{mp} as the absolute value of the difference between the central values of the m-th mode detected at t_{k-1} and the p-th mode detected at t_k:

$$\delta_{mp} = \left| x_{cm}(t_{k-1}) - x_{cp}(t_k) \right|. \tag{4-4}$$

Finally, let $Y_{pn}(t_k)$ denote the flag function for the p-th mode of the n-th component of the feature vector. $Y_{pn}(t_k)$ can have two values: 0 or 1. Initially $Y_{pn}(t_k) = 0$. If $Y_{pn}(t_k) = 1$, the p-th mode is said to be flagged.

The first step in the process is to check for any overlapping modes. For example, the black and red modes at time block t_{14} on the upper graph of Figure 4-7 are overlapping. If $p1$ and $p2$ are two modes at t_k for a given n, the modes overlap if

$$[x_{cp1} - w_{p1}, x_{cp1} + w_{p1}] \bigcap [x_{cp2} - w_{p2}, x_{cp2} + w_{p2}] \neq \varnothing, \tag{4-5}$$

where $\varnothing$ is the empty set. Two overlapping modes p_1 and p_2 are replaced by a single mode p_3 with $x_{cp3} = (x_{cp1}+x_{cp2})/2$ and $w_{cp3} = (w_{p1}+w_{p2})/2$. $M_n(t_k)$ is adjusted to reflect the new mode count. This step could be combined with the feature mapping (filtering) process, and it may not be necessary once higher packet acquisition rates are achieved by the sensor system.

If $M_n(t_k) > M_n(t_{k-1})$, then for each mode $1 \leq m \leq M_n(t_{k-1})$ compute δ_{mp}, and associate mode m with the particular choice of p for which δ_{mp} is a minimum. Also, set $Y_{pn}(t_k) = 1$ for any mode p at t_k that is not associated with a mode at t_{k-1}. Association determines which modes belong to a logical temporal sequence.

If $M_n(t_k) \leq M_n(t_{k-1})$, then for each mode $1 \leq p \leq M_n(t_k)$ compute δ_{mp}, and associate mode p with the particular choice of m for which δ_{mp} is a minimum. Also, set $Y_{pn}(t_k) = 1$ for any mode p at t_k that is associated with a mode m having $Y_{mn}(t_{k-1}) = 1$.

If $M_n(t_k) = M_n(t_{k-1})$, it is possible that the legitimate node (Belkin 1 in the intrusion scenario) did not transmit during the current time block, and that the intruder seized this opportunity to transmit forged packets. For example, suppose the red (legitimate) modes on Figure 4-5 were absent. Then the black (forged) packets occurring between t_{21} and t_{36} might be mistaken as legitimate. This case is addressed by comparing $x_{cp}(t_k)$ with an expected value derived form the previous time history. Since the time steps are assumed to be small relative to the time scale for temporal evolution of the mode characteristics, one requires that $x_{cm}(t_{k-1}) - \alpha w_m(t_{k-1}) < x_{cp}(t_k) < x_{cm}(t_{k-1}) + \alpha w_m(t_{k-1})$, where m and p are associated modes. If this condition is not met, disassociate p from m, and set $Y_{pn}(t_k) = 1$. The choice of α will depend on the stability of the feature vector component. For frequency error data of Figure 4-5, $\alpha = 1$ would be a reasonable choice.

After forming the above associations and setting the flags as appropriate for the current time block, compute the number of flags set for all modes of all feature vector components:

$$Y(t_k) = \sum_n \sum_p Y_{pn}(t_k). \tag{4-6}$$

If $Y(t_k) > Y_T$, where Y_T is a user specified flag count threshold, issue an intrusion alert message to the IDS. At a minimum, the alert message would contain the time t_k and the internal source identifier.

Finally, compute the number of legitimate modes detected at t_k,

$$Y_L(t_k) = Y(t_k) - \sum_n M_n(t_k), \tag{4-7}$$

and compare it to the previous value, $Y_L(t_{k-1})$. If $Y_T < Y_L(t_k) < Y_L(t_{k-1})$, create the missing legitimate modes using the previous values of $x_{cm}(t_{k-1})$ and $w_m(t_{k-1})$, and associate the new modes with the previous modes from which they were generated.

4.3 EXAMPLE APPLICATION OF WIND DETECTION ALGORITHM

Consider the application of this process to the $n = 4$ feature vector components shown in Figure 4-5 and Figure 4-7, and assume $Y_T = 2$. Prior to t_{14}, a single mode is present for F_1 and F_2 (lower and upper graphs of Figure 4-5, respectively), and two modes are present for P_1 and P_2 (lower and upper graphs of Figure 4-7, respectively). All flags are zero prior to t_{14}. At $t_{14,}$ a third mode appears for P_2. The black and red modes overlap, and are thus combined as a single mode. So there is no net change in the number of modes at this time block.

At $t_{15,}$ a third mode appears for P_1. The black mode at t_{15} is associated with the black mode at t_{14}, the green mode at t_{15} is associated with the red mode at t_{14}, and the red mode at t_{15} is flagged. $Y(t_{15}) = 1 < Y_T$. The next change in the number of modes occurs at t_{20}, where the number of modes for P_1 increases from 2 to 3. The black mode at t_{20} is associated with the black mode at t_{19}, the red mode at t_{20} is associated with the red mode at t_{19}, and the green mode at t_{20} is flagged. Again, $Y(t_{20}) = 1$.

At t_{21}, the number of modes for F_1 and F_2 increases from 1 to 2, the number of modes for P_1 increases from 2 to 4, and the number of modes for P_2 increases from 2 to 3. For both F_1 and F_2, the red mode at t_{21} is associated with the black mode at t_{20}, and the black mode at t_{21} is flagged. For P_1, the black mode at t_{21} is associated with the black mode at t_{20}, the green mode at t_{21} is associated with the red mode at t_{20}, and the red and blue modes at t_{21} are flagged. For P_2, the red mode at t_{21} is associated with the black mode at t_{20}, the green mode at t_{21} is associated with the red mode at t_{20}, and the black mode at t_{21} is flagged. $Y(t_{21}) = 5$, and an intrusion alert is issued to the IDS.

At t_{22}, the number of modes does not change for any of the feature vector components. For both F_1 and F_2, the red mode at t_{22} is associated with the red mode at t_{21}, the black mode at t_{22} is associated with the black mode at t_{21}, and the black mode at t_{22} is flagged. For P_1, each colored mode at t_{21} is associated with the same color at t_{20}, and the red and blue modes at t_{22} are flagged. $Y(t_{22}) = 5$, and an intrusion alert is again issued to the IDS. With the exception of combined modes at t_{23} and t_{24} on P_2, there are no mode changes until t_{26}. Intrusion alerts are issued at t_{23} through t_{25}.

At t_{26}, P_1 has overlapping blue and cyan modes that are associated with the blue mode at t_{25}; the red and combined modes are flagged. P_2 has overlapping blue and cyan modes that are associated with the green mode at t_{25}. The red mode is associated with the green mode of the previous time block, and the red mode is associated with the black mode of the previous time block. The black mode is flagged, along with the red and combined modes. $Y(t_{26}) = 6$, and an intrusion alert is again issued to the IDS. There are no changes in the number of modes between t_{27} and t_{36}, except for a missing legitimate mode at t_{33}. $Y_L(t_{33}) = 5$ and $Y_L(t_{32}) = 6$, so the missing mode is created and assigned the values of the x_c and w values of the green mode at t_{32}. $Y(t_k) = 5$ at each time block, and intrusion alerts are issued.

At t_{36} and beyond, the number of modes for all feature vector components return to their pre-intrusion values, no flags are set, and no additional intrusion alerts are issued.

For this example, $Y(t_k) \geq 5$ throughout the intrusion event. This suggests that the intrusion would have been detectable if a smaller number of feature vector components were used, or if the modes of several feature vector components were nearly identical for legitimate and forged packets (i.e., overlapping modes).

Once an intrusion has been identified during a particular time block t_k, it is possible to compute the probability that a given packet captured during t_k originated from a rogue network node or a legitimate node. The intrusion detection process associates individual modes of each feature vector component with the intruder or a legitimate node. The corresponding set of mode parameters (x_{cp}, w_p, A_p) define one or more Gaussian probability distributions which, in sum, represent the probability that a given packet is a rogue packet, or a legitimate packet. The contents of each packet could thus be accompanied by an indicator of the authenticity of the packet, albeit with a latency of the order of Δt.

In the mode associate process, the expected value of x_c at t_k is approximated by the previous value at t_{k-1}. This is equivalent to using the first term in a Taylor series expansion of x_c near t_k. A better estimate of x_c at t_k might be possible using a higher order Taylor series expansion with backward difference estimates of x_c temporal derivatives, or by using a polynomial fit to the previous values of x_c.

4.4 ALGORITHM PSEUDO-CODE

Pseudo-code is described below that implements the core statistics based algorithm described in the previous sections. The algorithm steps include:

1. Get a new feature vector from the packet processor.
2. If the internal source identifier is a new one then
    a. Request source authentication from higher layers of protocol stack.
    b. Create new source vector list for this internal source identifier.
3. Add the feature vector to the appropriate source vector list based on the internal source identifier (i.e., sort the incoming feature vectors by MAC address).
4. If number of vectors in a given source vector list exceeds threshold value L (or elapsed time Δt since the last RF fingerprint update exceeds the user specified time threshold Δt_C) then, for each feature vector element.
    a. Construct a histogram from the source vector list.
    b. Extract and record the mode characteristics from each histogram.
    c. Compare the current mode characteristics to the expected values based on the previously recorded mode characteristics; flag any anomalies using the procedure described above.
5. If the number of flags for a given internal source identifier exceeds the alarm threshold, Y_T, issue an intrusion alert.
6. Return (get another feature vector).

4.5 WIND INTERFACE WITH A TRADITIONAL IDS

The WIND system monitors wireless packet activity and uses extracted physical and link layer information to generate a feature vector for each captured packet. This feature vector includes the claimed network source identifier (i.e. MAC address), the time, and a number of signal features. The WIND algorithm tracks the evolution of the feature vectors. When features from a set of packets with the same source identifier differ significantly from the existing history, the algorithm generates an alarm. The alarm denotes whether an intrusion is suspected, a watch should be initiated on packets from the network source identifier in question, or the suspected threat has passed. A notional IDS message format is provided in Figure 4-8.

Time	Network source identifier (MAC)	Intrusion flag

Figure 4-8 Notional IDS message format

As currently envisioned, the WIND algorithm will record suspected intrusion events to a log file. This log file will be scanned by an application that will send messages to a traditional IDS. By using a secondary application working from a standardized log file, the WIND system can be easily adapted to work with a wide range of IDS systems. For example, if

an IDS utilized the Java message server to exchange information between collection sensors and decision making sensors, a Java WIND reporting application would be implemented to provide similar functionality.

4.6 WIND PERFORMANCE AND OPTIMIZATION

A series of short duration (40 second) experiments were conducted with the new (NI/Agilent) WIND system to examine packet feature stability, uniqueness, and sensitivity to antenna choice and network configuration. Four features were considered: frequency error, rise-time, received power, and IQ offset. The details of each experiment, and the resultant observed feature statistical distributions are provided in Appendix C. This subsection uses the results of these experiments to identify an optimal feature set, and to estimate WIND performance. WIND performance is specified in terms of probability of successful detection of an intrusion event, and false alarm probability.

4.6.1 FREQUENCY ERROR

Table 4-1 and Table 4-2 provide a summary of the frequency error observations derived from Experiments 29 through 33. A spatially dispersed network was used for these experiments, and data were collected using three different antennas: a dual polarization horn, vertically polarized omni-directional antenna, and a horizontally polarized parabolic dish antenna. Table 4-1 provides a summary for the dual polarization horn measurements, while Table 4-2 presents the corresponding data for omni-directional and parabolic antennas. The data within each table are subdivided by card group (Belkin/Linksys or D-Link/Orinoco), and antenna. The first column in each table is the ID number assigned to a particular card within a given card group (see Appendix C for the specific card and MAC address associated with each ID and card group). The frequency error data for each card (ID and card group) were sorted into 200 Hz bins and histogrammed using the procedure outlined above. Appendix C provides the histograms. The resultant modal peak (bin center frequency at which the most probable value occurs) is identified in each table as Freq. Two percentage values are also listed: Raw % and Adj %. These values are the observed number of packets falling within the central frequency error distribution, expressed as a percentage of the total number of observed packets, before (Raw %) and after (Ajd %) the filtering described below. The central frequency error bounds are ± 300 Hz around the modal peak (i.e., the central bin and one bin on either side). In two cases, the horizontally polarized horn data from Experiment 32 and the omni-directional antenna data from Experiment 30, the modal peak values have been reduced by 6.2 kHz to correct for a systematic error that resulted when the Channel 1 PXI-5600 down-converted was not properly locked to the frequency standard (see Section 2).

The table values labeled ND indicate that either there were insufficient number of packets detected to form a meaningful frequency error distribution, or a distinct central value could not be found. Most of the data from the six Linksys cards (the even numbered ID values in

Experiments 32 of Table 4-1 and Experiment 31 of Table 4-2) fall in the latter category. No clear frequency error modal peaks were found for these cards in most of the experiments (see Appendix C). Further investigation showed that almost all of the Linksys packets detected by WIND employed high data rate CCK 11 modulation (Burst Type 3).

Table 4-1 Dual Polarization Horn With Dispersed Spatial Configuration

ID	Experiment 32 (Belkin/Linksys; dispersed)						Experiment 29 (D-Link/Orinoco; dispersed)					
	Horn Horizontal			Horn Vertical			Horn Horizontal			Horn Vertical		
	Freq# (kHz)	Raw (%)	Adj (%)	Freq (kHz)	Raw (%)	Adj (%)	Freq (kHz)	Raw (%)	Adj (%)	Freq (kHz)	Raw (%)	Adj (%)
01	-6.2	81	81	ND	ND	ND	-77.4	100	100	ND	ND	ND
02	ND	ND	ND	ND	ND	ND	-19.8	42*	92	-19.8	32*	100
03	-7.0	69	69	-7.0	75	99	-75.0	83	100	-75.0	63	100
04	ND	ND	ND	ND	ND	ND	-12.8	68	100	-12.8	63	87
05	5.2	69	79	5.2	83	83	ND	ND	ND	-132.6	36	89
06	ND	ND	ND	ND	ND	ND	-11.2	51	88	-11.2	65*	89
07	-3.4	73	76	-3.4	59	59	-91.0	53	78	-91.0	45	73
08	ND	ND	ND	ND	ND	ND	**-13.6**	54	100	**-13.4**	38	[4]
09	-2.4	73	82	-2.4	86	86	-47.6	51	83	-47.6	35	70
10	ND	ND	ND	ND	ND	ND	-14.8	61	100	-14.8	79	97
11	-8.8	79	81	-8.8	78	82	-51.4	36	92	-51.4	35	100
12	ND	ND	ND	ND	ND	ND	**-13.6**	41*	97	**-13.6**	31	93

 # Common -6.2 kHz offset removed; * Packets with low sync correlation.

Table 4-2 Omni-Directional and Parabolic With Dispersed Spatial Configuration

ID	Experiment 31 (Belkin/Linksys; dispersed)						Experiment 30 (D-Link/Orinoco; dispersed)					
	Omni-Directional			Parabolic			Omni-Directional			Parabolic		
	Freq (kHz)	Raw (%)	Adj (%)	Freq (kHz)	Raw (%)	Adj (%)	Freq# (kHz)	Raw (%)	Adj (%)	Freq (kHz)	Raw (%)	Adj (%)
01	-6.2	100	100	-6.2	100	100	-78.6	55	89	-78.4	92	92
02	ND	ND	ND	ND	ND	ND	-20.0	84	97	-20.0	94	99
03	-7.0	16*	[3]	-7.0	94	97	-76.6	91	92	-76.6	96	97
04	ND	ND	ND	ND	ND	ND	-13.4	80	98	-13.4	26*	96
05	5.0	59	100	ND	ND	ND	-133.4	36	90	-133.4	36	100
06	**-11.4**	15	[4]	ND	ND	ND	**-11.8**	85	98	-11.8	78	97
07	-3.6	60	97	-3.6	54	95	-91.6	55	84	-91.6	78	83
08	ND	ND	ND	ND	ND	ND	**-14.8**	36*	91	**-14.6**	43	89
09	-2.6	33*	100	-2.6	39*	95	-49.8	47	94	-49.8	47	96
10	ND	ND	ND	ND	ND	ND	-15.4	66	95	-15.4	90	94
11	-8.6	75	93	-8.6	69	94	-53.4	42	100	-53.2	71*	100
12	-7.8	7	[2]	ND	ND	ND	**-14.4**	29*	88	**-14.4**	59	93

 # Common -6.2 kHz offset removed; * Packets with low sync correlation.

Raw % is the percentage of packets falling within the frequency error central distribution. All packets with frequency errors outside the central distribution contribute to the false alarm rate. Using frequency error alone, the average probability of false alarm, P_{FA}, can be approximated as:

$$P_{FA} = 1 - \text{Avg [Raw \%]}. \quad (4\text{-}8)$$

From Tables 4-1 and 4-2 one finds $P_{FA} > 30\%$ for all four antennas (i.e., 36% for Horn Horizontal, 44% for Horn Vertical, 46% for Omni-Directional, and 31% for Parabolic). Such a high false alarm probability is clearly unacceptable.

An investigation was conducted determine why so many packets have frequency error values outside of the central distribution. It was found that for all cards, regardless of the manufacturer, the majority of packets falling outside the frequency error central distribution were CCK 11, while the majority of those within the central distribution were low data rate Barker 1 (Burst Type 0). WIND does not discriminate between packet modulation types in forming feature statistics. The only requirement for a packet to be accepted by WIND is that the decoded MAC address be valid. As will be shown below, it is essential that WIND consider the quality of the received packets. The disproportionate number of CCK 11 packets falling outside the frequency error central distribution is a symptom of poor packet quality.

One measure of packet quality is the error vector magnitude, EVM [12]. IEEE Standard 802.11b [3] requires that the worst-case peak EVM (expressed as a percentage of the square root of the mean power of an ideal signal) not exceed 35% for high data rate modes like CCK 11. Figure 4-9 illustrates the impact of exceeding this threshold on frequency error. It shows observed frequency error as a function of worst-case peak EVM for CCK 11 and Barker 1. When the 35% worst-case peak EVM limit is exceeded, the frequency error estimate from CCK 11 packets becomes unstable. WIND should exclude such packets from the RF fingerprint because these packets would be rejected by the AP (or an Ad Hoc node), and a re-transmission would be requested.

The prevalence of CCK 11 packets in the experimental data of Table 4-1 and Table 4-2 can be traced to the way in which the experiments were conducted. The wireless nodes were operated in Ad Hoc mode, so packets were exchanged between nodes without an intermediate AP. The link quality between two Ad Hoc nodes may be sufficient to exchange data at 11 Mbps (CCK 11), but the quality of the same CCK 11 packets as detected at WIND may be inadequate for RF fingerprint construction. To reduce false alarms due to poor packet quality, WIND RF fingerprints should be based only on those packets that are accepted by the AP and forwarded up the protocol stack. Future experiments should have WIND co-located with an AP (or Ad Hoc node) and sharing the same antenna.

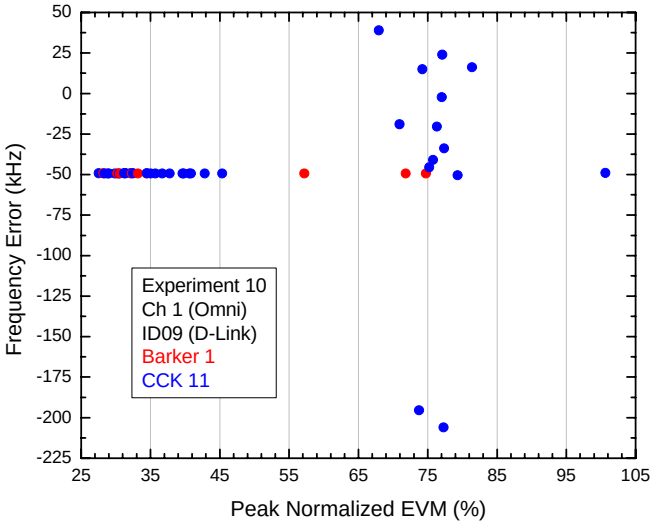

Figure 4-9 Frequency Error Variation with Worst-Case Peak EVM

The packets captured by WIND were filtered to exclude CCK 11 packets with worst-case peak EVM > 35%. In addition, it was observed that in some instances (marked with an asterisk in the tables) the frequency error for Barker 1 packets was adversely affected by low sync correlation. Consequently, Barker 1 packets with sync correlation less than 0.80 were also excluded. The resultant adjusted percentage of packets falling within the frequency error central distribution is given by the Adj % values in Table 4-1 and Table 4-2. The average false alarm rates derived from the Adj % values are a significant improvement: 13% for Horn Horizontal (12% with the low sync correlation values removed), 15% for Horn Vertical (5% with the low sync correlation values removed), 6% for Omni-Directional, and 5% for Parabolic.

Each of the experiments summarized in Table 4-1 and Table 4-2 was limited to twelve wireless cards because that was the number of available laptop PCs. Both experiments in each table were conducted under nearly identical conditions (i.e., same laptops, geometry and antennas), so the combined results of both experiments will be used here to estimate WIND performance for a large network.

Figure 4-10 shows the frequency error distributions from the combined experiments for the omni-directional antenna after filtering. The lower graph shows the frequency distributions for seventeen cards (5 Belkin, 6 D-Link and 6 Orinoco), while the upper graph is a blowup focused on the -30 to 0 kHz region. Only two of the distributions overlap: ID08 and ID12 from Experiment 30. The frequency error distributions span a range of more than 150 kHz, while the individual card central distributions span less than 0.6 kHz. This shows that the probability of detecting an intrusion event using frequency error alone is very high.

The experiments summarized in Tables 4-1 and 4-2 were conducted hours, or in some cases, days apart. Yet, for a given card, the frequency error modal peak identified by WIND varied very little from one experiment to the next. The typical variation was 0.7 kHz, with a maximum range of 2.2 kHz. This is consistent with earlier experiments (see Figure 4-3). The data also show that frequency error is independent of antenna type or polarization.

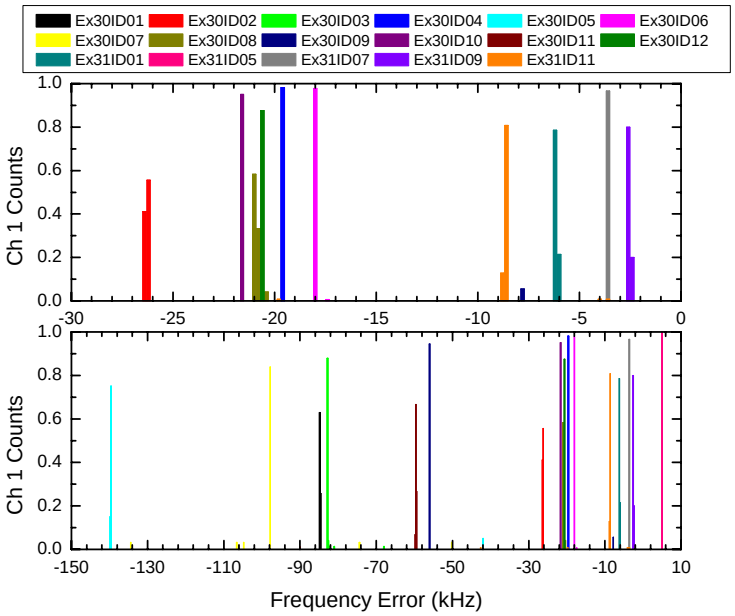

Figure 4-10 Frequency Error Distributions (Omni-Directional Antenna; Filtered)

Let N denote the number of cards observed during the joint experiment (excluding those cases with an insufficient number of filtered packets). Considering all possible scenarios where Card X attempts to spoof Card Y. The total number of such permutations is

$$P_2^N = \frac{N!}{(N-2)!}. \tag{4-9}$$

Let Q denote the number of card permutations with overlapping frequency error distributions:

$$[\text{Card X: FreqX}\pm 300\text{ Hz}] \cup [\text{Card Y: FreqY}\pm 300\text{ Hz}] \neq \varnothing. \tag{4-10}$$

Q represents the number of spoofing scenarios that are undetectable using frequency error. The probability of detection, P_D, can then be approximated as

$$P_D = 1 - \frac{Q}{P_2^N}. \tag{4-11}$$

For example, from Table 4-1, N = 16 for the Horn Vertical cases, and there are two overlapping distributions, which are identified by the bold face Freq values: -13.4 kHz and -13.6 kHz. In this case, Q =2, and the undetectable spoofing scenarios are:

"D-Link/Orinoco ID 08 spoofs D-Link/Orinoco ID 12"

and

"D-Link/Orinoco ID 12 spoofs D-Link/Orinoco ID 08".

The corresponding probability of detection is 0.99 or 99%. In fact, one finds $P_D \geq$ 99 % for all four joint experiments (antennas) formed from Table 4-1 and Table 4-2.

4.6.2 RISE-TIME

Figure 4-11 shows the rise-time error distributions from the combined experiments for the omni-directional antenna after filtering. These results are typical of those observed for the other experiments. The observed rise-times of the Belkin cards are significantly different from those of the D-Link/Orinoco group, but the rise-times of the D-Link and Orinoco cards are indistinguishable from each other. Rise-time also cannot distinguish between cards from a given manufacturer. Rise-time is not a viable feature for RF fingerprinting of IEEE 802.11b devices.

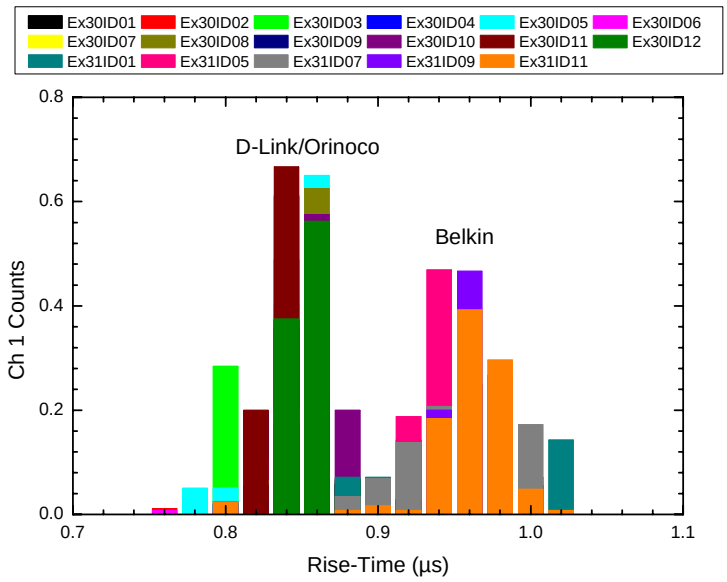

Figure 4-11 Rise-Time Distributions (Omni-Directional Antenna; Filtered)

4.6.3 RECEIVED POWER

Figure 4-12 and Figure 4-13 show some typical received power distributions from omni-directional and directional antennas. One each figure the upper graph is from the horizontally polarized parabolic antenna, while the low graph is from the vertically polarized omni-directional antenna. The data in Figure 4-12 are from the six D-Link cards, while those in Figure 4-13 are from the six Orinoco cards. The bin width is 0.5 dB.

The received power distributions for both card sets as observed on the omni-directional antenna sets are diffuse, with no clear central peak. They also display significant overlap from one card to the next. Using a directional antenna tends to make the individual receive power distributions more compact, with a central peak apparent in many cases. It also spreads out one card distribution from the next. Nevertheless, even with a directional antenna, many of the card distributions overlap. The received power distributions were observed to be reasonable stable over time, but this is largely because none of the card host laptop PCs were moved in any way during the measurements, and the PCs were not in active use by a human operator. Previous measurements of received power as a function of PC orientation are shown in Figure 4-14. The radiation patterns of wireless network cards are highly structured. Consequently, a small movement of the PC can significantly alter the received power. In a mobile environment, received power would be highly variable, and of no practical use for RF fingerprinting.

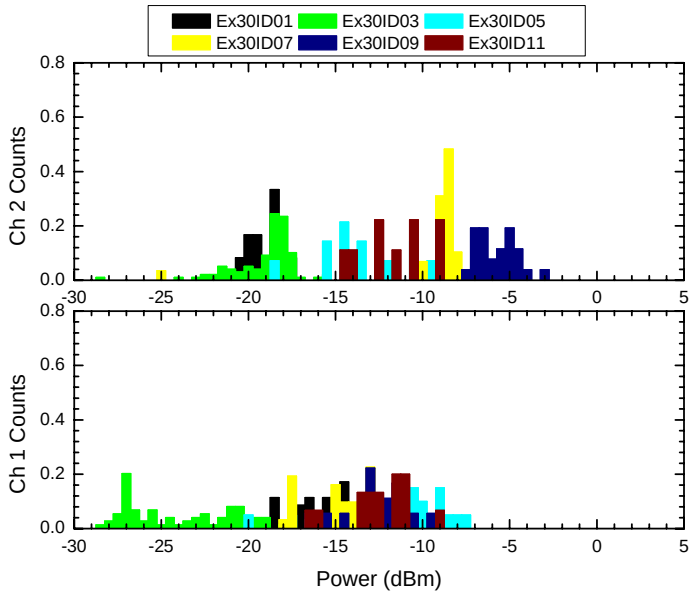

Figure 4-12 Received Power Distributions for D-Link Cards (Experiment 30)

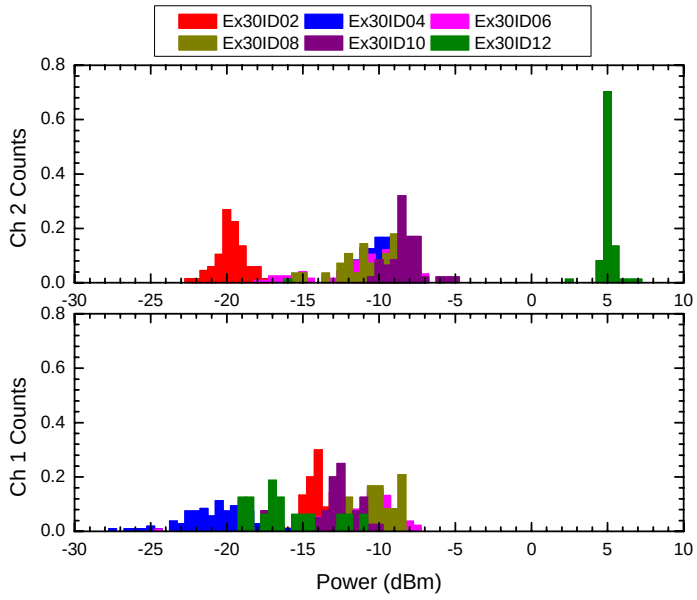

Figure 4-13 Received Power Distributions for Orinoco Cards (Experiment 30)

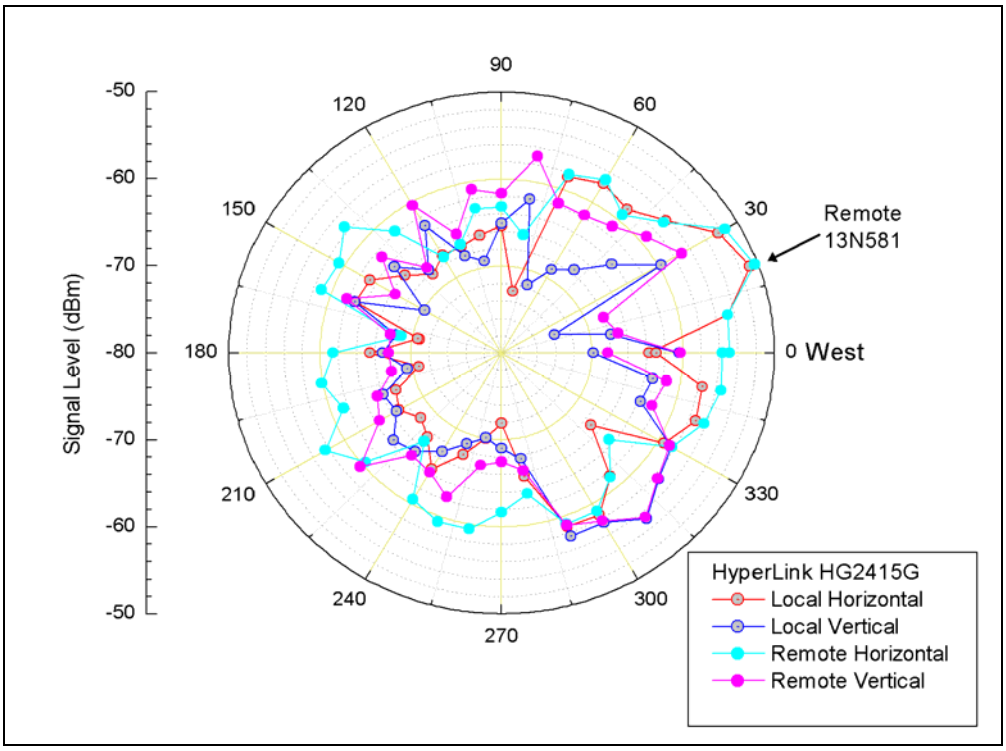

Figure 4-14 Received Power Variability

4.6.4 IQ OFFSET

Some initial measurements of IQ offset with a small number of cards suggested that this might be a useable feature. However, after testing a larger sample of cards, one finds that IQ offset is very unstable, even over relatively short time scales. Figure 4-15 shows the temporal evolution of IQ offset for a typical sample of eight cards (four D-Link and four Orinoco) from Experiment 30. The IQ offset values for a given card are neither unique nor stable.

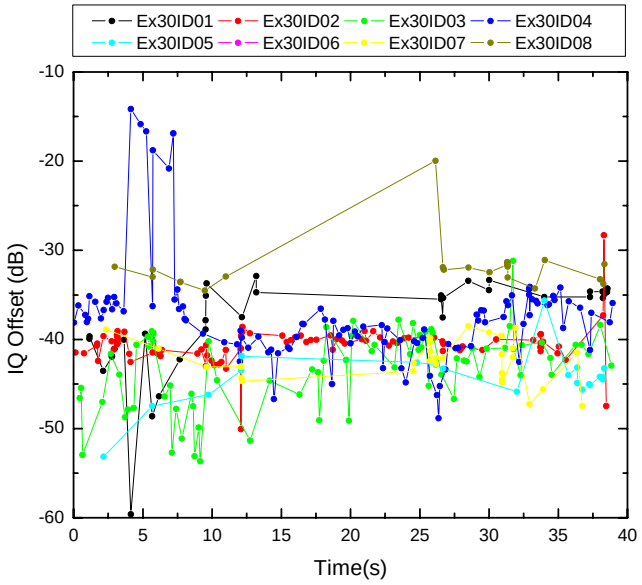

Figure 4-15 Observed Short-Term Temporal Variability of IQ Offset (Experiment 30)

Section 5

CLOSURE

5.1 SUMMARY

The objectives of the WIND research effort were to identify an optimal set of RF features, develop fingerprinting and intrusion detection methodologies, and exploring approaches for real-time implementation of WIND, with a focus on IEEE 802.11b wireless local area networks. A series of twenty-four experiments were conducted using a legacy packet capture system to identify candidate RF features, and to study their long-time scale (hours) stability and uniqueness (Appendix B). Three candidate features were identified: packet rise-time, received power level and frequency error (Section 3). These data were used to test two intrusion detection methodologies: one based on clustering techniques (Appendix D), and another based on tracking the temporal evolution of feature statistics (Section 4). Significant difficulties were encountered in implementing the clustering approach, and it was subsequently abandoned in favor of the more intuitive statistics method.

Concurrent with these activities, JHU/APL designed, implemented and tested a new modular, high speed packet capture and analysis system (Section 2). The packet capture system utilizes COTS PXI components, and allows continuous digitization and storage of RF packets at any frequency below 2.7 GHz in a 22 MHz bandwidth. Two packet capture systems were assembled, and a capability was developed to synchronize the two systems using GPS. Several commercial IEEE 802.11b/g demodulation and decoding packages were evaluated as a means of RF feature extraction, with the Agilent 89600 VSA software being eventually selected.

After the new packet capture and analysis system was tested, ten additional experiments were conducted to examine packet feature stability on short time scales (<1 minute), feature uniqueness from one card to the next, and feature sensitivity to antenna choice and network configuration (Appendix C). Twenty-four cards from four manufacturers (six from each) were measured using four different combinations of antenna and polarization for monitoring. The network geometries studied included compact and dispersed configurations indoors, and indoor/outdoor configuration in an “Internet Café” setting.

5.2 CONCLUSIONS

Only one of the four candidate RF features studied as part of this project is suitable for fingerprinting IEEE 802.11b devices: frequency error. The other three are too variable, or do not have statistics that are unique from one card to the next.

To reduce false alarms due to poor packet quality, WIND RF fingerprints should be based only on those packets that are accepted by the AP and forwarded up the protocol stack. WIND should be collocated with the AP and utilizing the same antenna(s).

The observed frequency error distributions are compact (<0.6 kHz spread), and, in the majority of observed cases, have little or no overlap. The frequency error central values vary significantly from one card manufacturer to the next, with ensemble of observed distributions spanning more than 160 kHz. Frequency error is very stable over short time scales (<1 minute) and drifts slowly over long-time scales (hours). The maximum observed long-term drift is 2.2 kHz. Frequency error is insensitive to antenna type or network geometry.

The available frequency error data suggest that a physical layer intrusion detection system based solely on the statistics of frequency error can achieve a 99% probability of detection, with an average false alarm rate of 10%.

5.3 RECOMMENDATIONS

1) Consider augmentation of AFRL Wireless Intrusion Detection System (WIDS) with RF fingerprints derived from frequency error.

2) Exploit the WIND packet capture and analysis capabilities developed as part of this effort to study fingerprinting and other network security techniques as applied to newer wireless technologies, particularly those employing orthogonal frequency division multiplexing (OFDM) modulation. The WIND measurement system would be ideally suited for studies of multiple-input multiple output (MIMO) wireless architectures like IEEE 801.11n, and could be easily adapted for use with IEEE 802.16 (WiMax).

3) Hall et al. [13] have developed a statistical RF fingerprinting technique based on the measurement of wireless device turn-on transients. This technique uses the entire transient (rise-time portion plus settling time afterwards). The authors claimed that this approach yields an average detection rate of 95% with a zero false alarm rate. This technique should be evaluated using the WIND packet library.

Appendix A

LIST OF REFERENCES

[1] Muaddi, A., and Tomko, A., "Wireless Network Physical Layer Intrusion Detection System", Patent Application 1807, The Johns Hopkins University Applied Physics Laboratory, April 2003.

[2] Phamdo, N, and Tomko, A., "Wireless Intrusion Detection (WIND)," Technical Proposal in Response to Air Force Research Laboratory (AFRL) Solicitation #BAA 04-04-IFKA – Wireless Cyber Operations, The Johns Hopkins University Applied Physics Laboratory, 13 May 2004.

[3] "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Higher-Speed Physical Layer Extension in the 2.4 GHz Band," *IEEE Std. 802.11b-1999*, http://standards.ieee.org/.

[4] "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications—Amendment 4: Further Higher-Speed Physical Layer Extension in the 2.4 GHz Band," *IEEE Std 802.11g-2003*, http://standards.ieee.org/.

[5] Buell, L. H., Rieser, C. J. , Tomko, A. A., and Zaret, D. R., "Wireless Intrusion Detection," AI-06-004, Applied Information Sciences Department, The Johns Hopkins University Applied Physics Laboratory, January, 2006.

[6] National Instruments PXI Products and Specifications, http://www.ni.com/pxi/.

[7] Conduant Corporation StreamStor and Big River Product Information, November 16, 2004, http://www.conduant.com/news/pressrelease2.html.

[8] ZTEC ZT1000 GPS PXI Specifications, http://www.amplicon.co.uk/data/ZT1000.pdf.

[9] SeaSolve WiLANTA Software, http://www.seasolve.com/products/wilanta_lvsa/index.html.

[10] "Agilent 89600 Series Vector Signal Analysis Software 89601A/89601AN/89601N12 Data Sheet," http://cp.literature.agilent.com/litweb/pdf/5989-1786EN.pdf, Agilent Technologies, Inc., August 31, 2006.

[11] OriginLab Scientific Graphing and Analysis Software, Version OriginPro 7.5, http://www.originlab.com/.

[12] Hassun, R., Flaherty, M., Matreci, R., and Taylor, M., "Effective Evaluation of Link Quality Using Error Vector Magnitude Techniques," Proceedings of the 1997 Wireless Communications Conference, 11-13 August 1997, Boulder, CO, pp. 89-94.

[13] Hall, J., Barbeau, M., and Kranakis, E., "Radio Frequency Fingerprinting for Intrusion Detection in Wireless Networks," Manuscript submitted to IEEE Trans.On Dependable and Secure Computing, July 2005, http://www.scs.carleton.ca/~jhall2/Publications/IEEETDSC.pdf.

Appendix B

LONG-TIME SCALE INTRUSION DETECTION EXPERIMENTS

This Appendix provides details of the initial set of twenty-four experiments conducted by JHU/APL to facilitate WIND algorithm development. The experiments were conducted in the JHU/APL Computer Network Operations (CNO) Laboratory. Laptop computers, equipped with PC Card format IEEE 802.11b wireless network cards from various manufacturers, were used as packet sources. The experiments utilized JHU/APL's original low-speed packet capture system, which only acquired about 10 packets per minute per card. The experiments were thus directed toward study the long-term behavior of packet RF features. The experiment durations ranged from two to sixteen hours. Experiments 1 through 8 used a dual polarization horn antenna for signal detection. For Experiments 9 through 24, Channel 1 was a polarized parabolic antenna and Channel 2 was an omni-directional antenna. Experiments 1 through 8 tested different numbers of intruders with staggered activity in a network made up of the same brand of cards. Experiments 9 through 24 were performed in a larger mixed network consisting of cards from several manufacturers. Some of these experiments utilized intruder nodes with vertically or horizontally polarized antennas, placed nearby the sensor as well as in adjacent rooms. Several experiments were performed as test cases to verify correct operation of the test bed. Later experiments were pursued to observe the stability of RF variations over a longer period of time. The arrows shown between nodes on figures in the appendix indicate the direction of pings between network elements.

B-1 EXPERIMENT 1

Purpose: To investigate different numbers of intruders with staggered activity in a network made up of the same brand of cards. The entire network (four authorized laptops) along with one intruder laptop was comprised entirely of Belkin PC Cards. Network traffic consists of pings sent continuously between the nodes. Elapsed Time: approximately 3 hours.

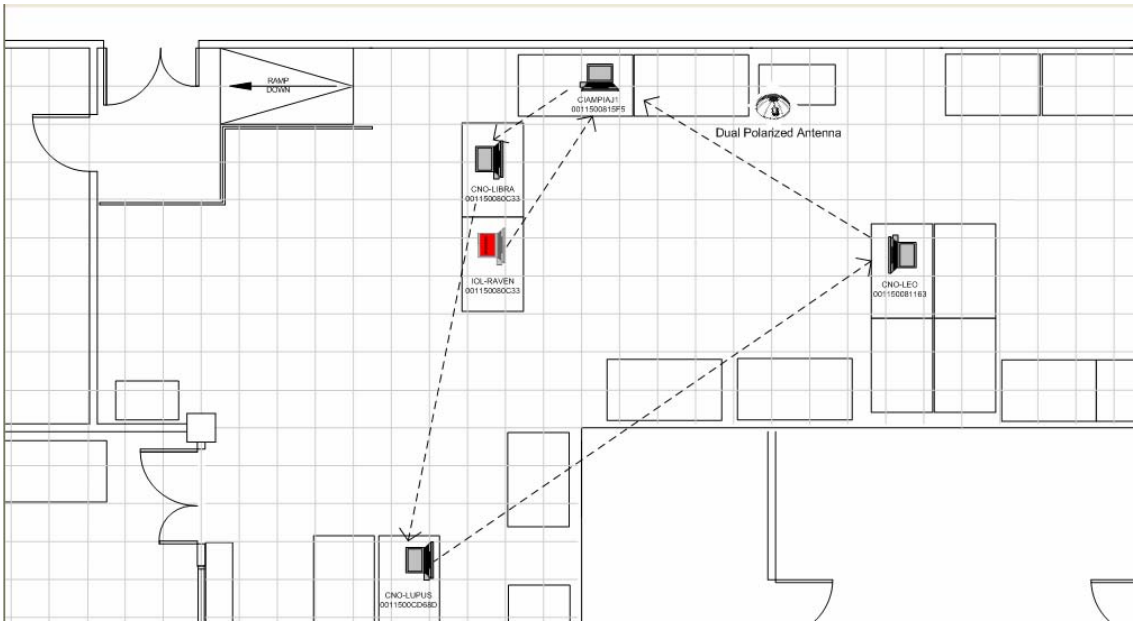

Figure B-1 Configuration for Experiment 1

B-2 EXPERIMENT 2

Purpose: To investigate different numbers of intruders with staggered activity in a network made up of the same brand of cards. The entire network (three authorized laptops) along with two intruder laptops was comprised entirely of Belkin PC Cards. Both intruders were turned on and off simultaneously. Network traffic consists of pings sent continuously between the nodes. Elapsed Time: approximately 2.5 hours.

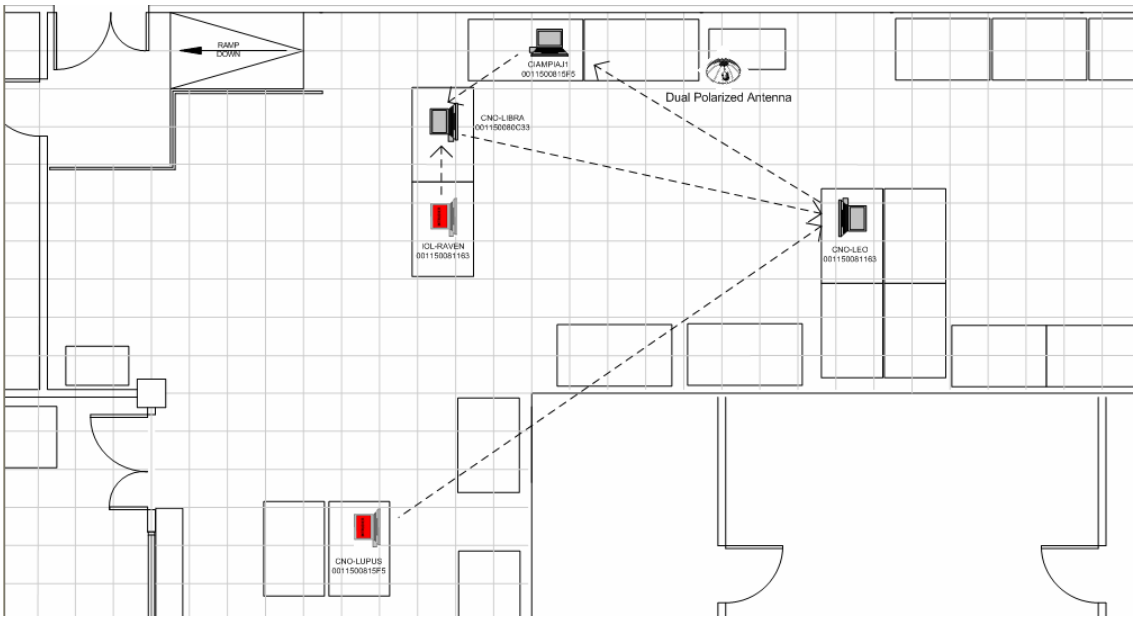

Figure B-2 Configuration for Experiments 2 through 4

B-3 EXPERIMENT 3

Purpose: To investigate different numbers of intruders with staggered activity in a network made up of the same brand of cards. The entire network (three authorized laptops) along with two intruder laptops was comprised entirely of Belkin PC Cards. Intruders are turned on separately, each for 45 minutes, and never exist simultaneously on the network. Network traffic consists of pings sent continuously between the nodes. Elapsed Time: approximately 2.5 hours.

B-4 EXPERIMENT 4

Purpose: To investigate different numbers of intruders with staggered activity in a network made up of the same brand of cards. The entire network (three authorized laptops)

along with two intruder laptops was comprised entirely of Belkin PC Cards. Intruders were turned on separately, each for an hour. After the first 30 minutes of the first intruder's activity, intruder #2 was turned on. After another 30 minutes, intruder #1 was turned off. Another 30 minutes later, intruder #2 was turned off and the experiment ended after a final 30 minutes of regular network activity (no intruders). The result was a 30 minute overlap of intruder activity in a 1.5 hour intruder activity period. Elapsed Time: approximately 2.5 hours.

B-5 EXPERIMENT 5

Purpose: To investigate different numbers of intruders with staggered activity in a network made up of the same brand of cards. The network consists of four authorized laptops, each using a Linksys PC Card, along with one intruder laptop using a Belkin PC Card. Network traffic consists of pings sent continuously between the nodes. Elapsed Time: approximately 2.5 hours.

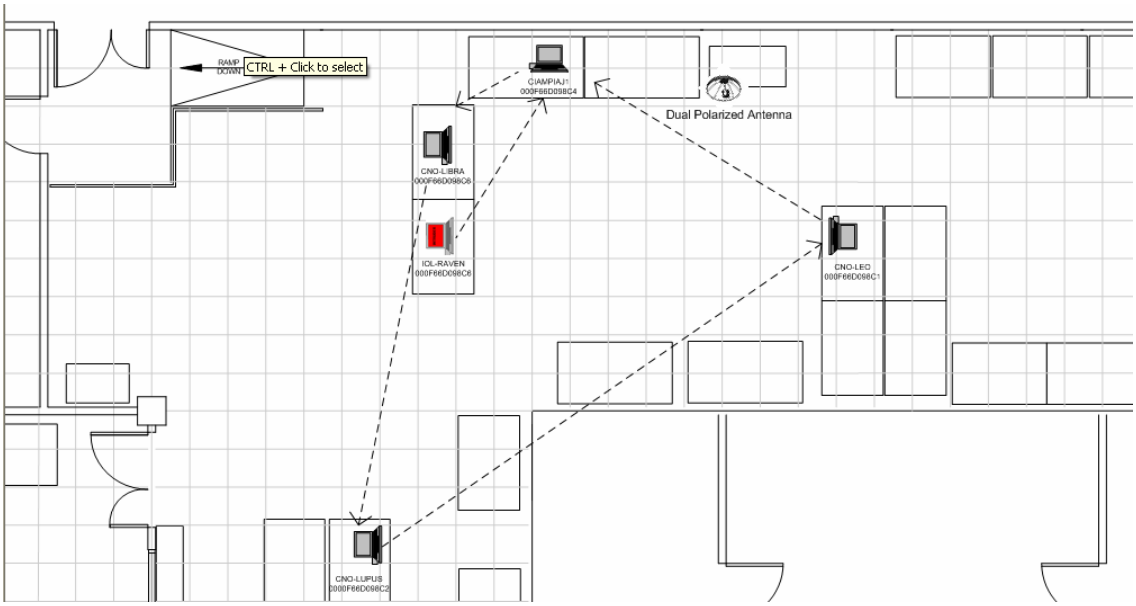

Figure B-3 Configuration for Experiment 5

B-6 EXPERIMENT 6

Purpose: To investigate different numbers of intruders with staggered activity in a network made up of the same brand of cards. The network consists of three authorized laptops, each using a Linksys PC Card, along with two intruder laptops using Belkin PC Cards. Both intruders were turned on and off at the same time. Network traffic consists of pings sent continuously between the nodes. Elapsed Time: approximately 2.5 hours.

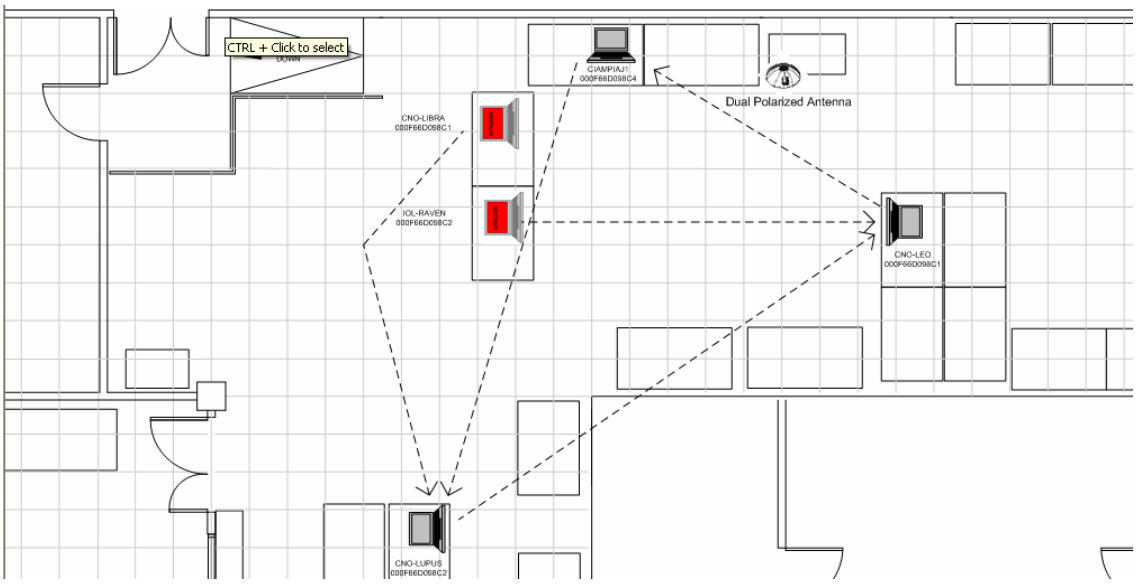

Figure B-4 Configuration for Experiment 6

B-7 EXPERIMENT 7

Purpose: To investigate different numbers of intruders with staggered activity in a network made up of the same brand of cards. The network consists of three authorized laptops, each using a Linksys PC Card, along with two intruder laptops using Belkin PC Cards. Each intruder was turned on separately (each for 45 minutes) and never existed on the network simultaneously. Network traffic consists of pings sent continuously between the nodes. Elapsed Time: approximately 2.5 hours.

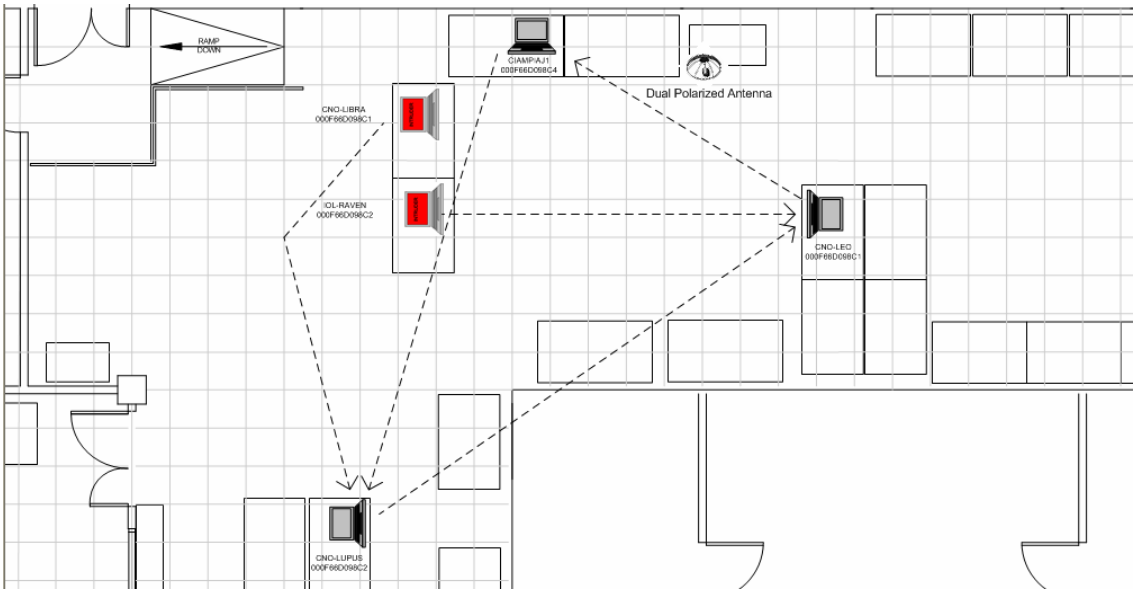

Figure B-5 Configuration for Experiment 7

B-8 EXPERIMENT 8

Purpose: To investigate different numbers of intruders with staggered activity in a network made up of the same brand of cards. The network consists of three authorized laptops utilizing Linksys PC Cards. The two intruder laptops, however, are comprised of Belkin PC Cards. The intruders were turned on separately, each for an hour. After the first 30 minutes of the first intruder's activity, intruder #2 was turned on. After another 30 minutes, intruder #1 was turned off. Another 30 minutes later, intruder #2 was turned off and the experiment ended after a final 30 minutes of regular network activity (no intruders). The result was a 30 minute overlap of intruder activity in a 1.5 hour intruder activity period. Elapsed Time: approximately 2.5 hours.

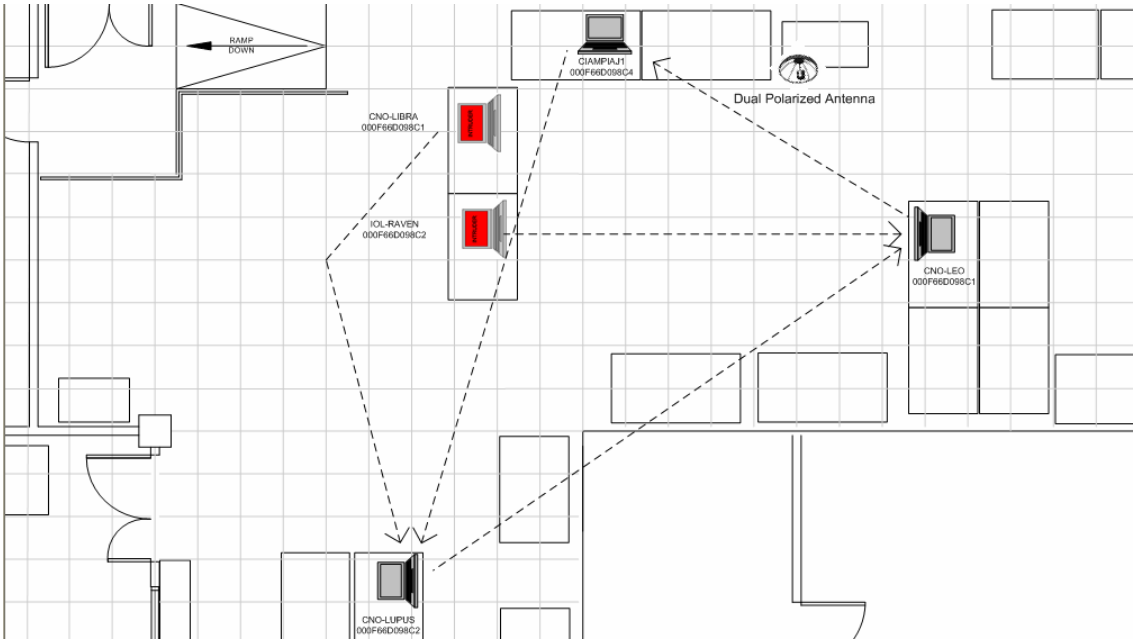

Figure B-6 Configuration for Experiment 8

B-9 EXPERIMENT 9

Purpose: To verify correct operation of the test bed. This experiment was used to test whether the new antenna configuration had its full functionality. Four regular laptops (with Linksys PC cards) comprise the network. One intruder laptop (equipped with a Belkin PC card) penetrates the network. The intruder was turned on after the first half hour and remained on for 1.5 hours, and was then turned off. Network traffic was small (pings). Elapsed time: approximately 2.5 hours.

B-10 EXPERIMENT 10

Purpose: To show that a unique signature can be obtained from each PC card, regardless of intrusion. Within this experiment, there were six regular laptops with no intruders. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. Elapsed time: approximately 3 hours.

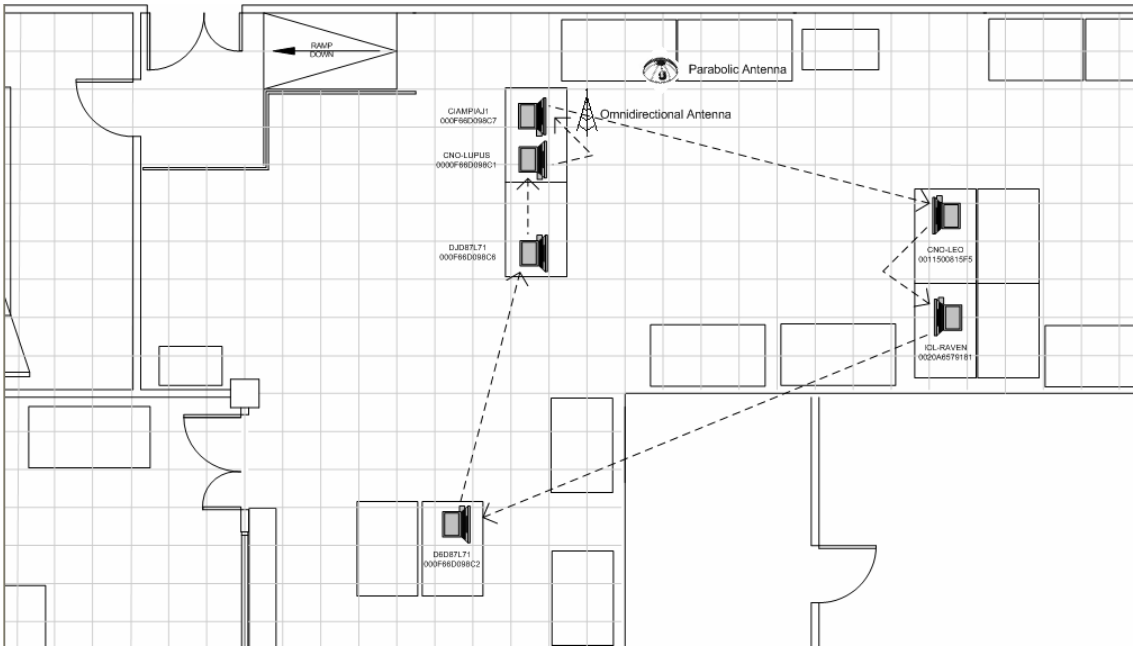

Figure B-7 Configuration for Experiment 10

B-11 EXPERIMENT 11

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes with horizontally polarized antennas, placed near the sensor. Within this experiment, there were five regular laptops with one intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. The orientation of the parabolic antenna was horizontal. The intruder was turned on after 30 minutes of non-intrusion and was then turned off after another 1.5 hours. Elapsed time: approximately 2.5 hours.

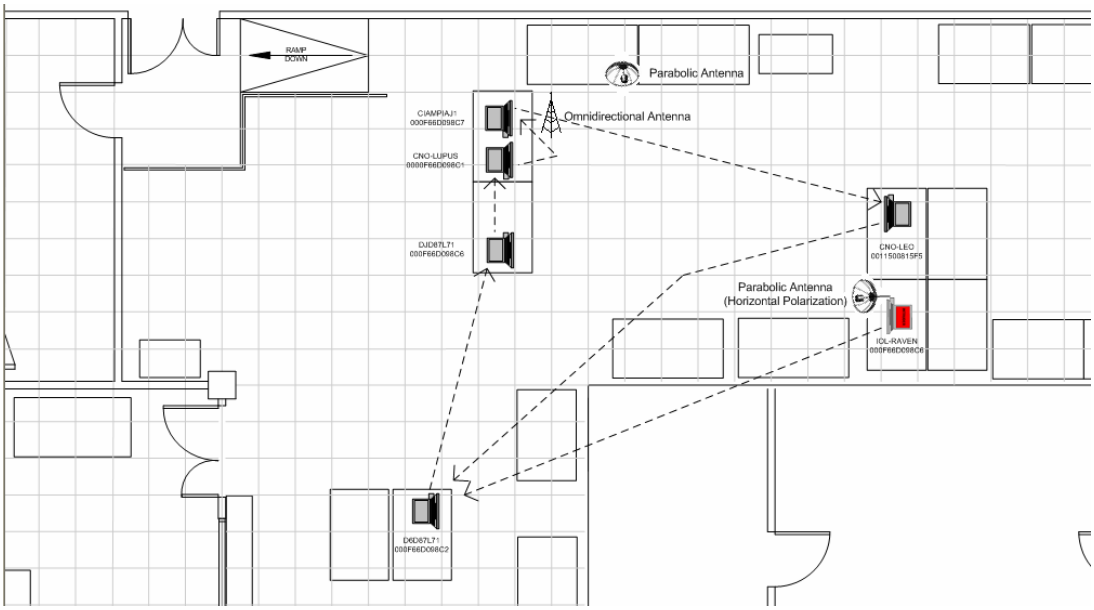

Figure B-8 Configuration for Experiment 11

B-12 EXPERIMENT 12

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes with vertically polarized antennas, placed near to the sensor. Within this experiment, there were seven regular laptops with one intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. The orientation of the parabolic antenna was vertical. The intruder was turned on after 30 minutes of non-intrusion and was then turned off after another hour. Elapsed time: approximately 2 hours.

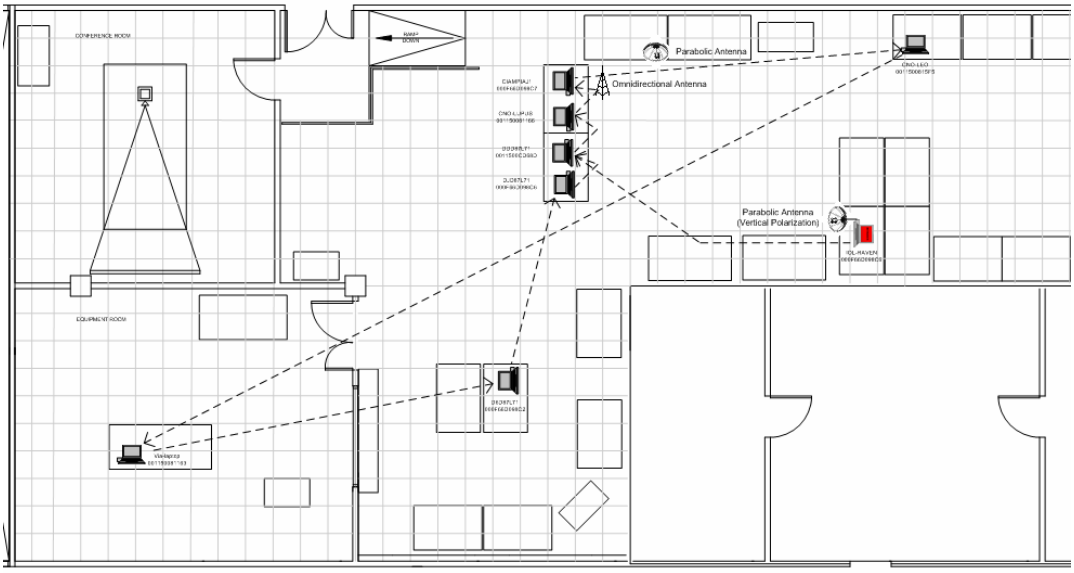

Figure B-9 Configuration for Experiment 12

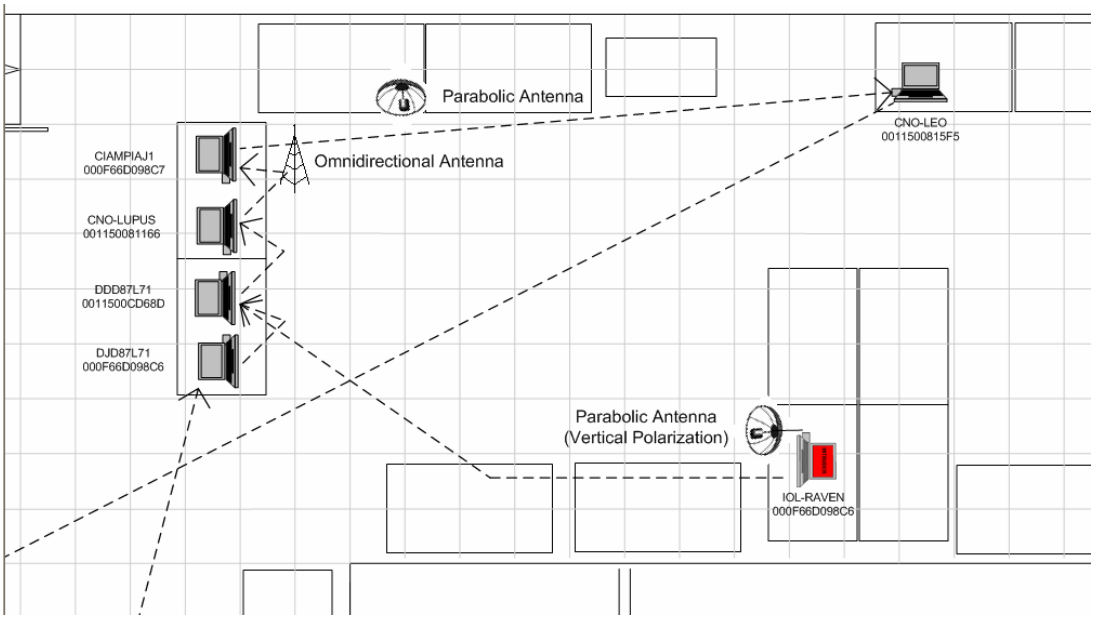

Figure B-10 Experiment 12 Detail 1

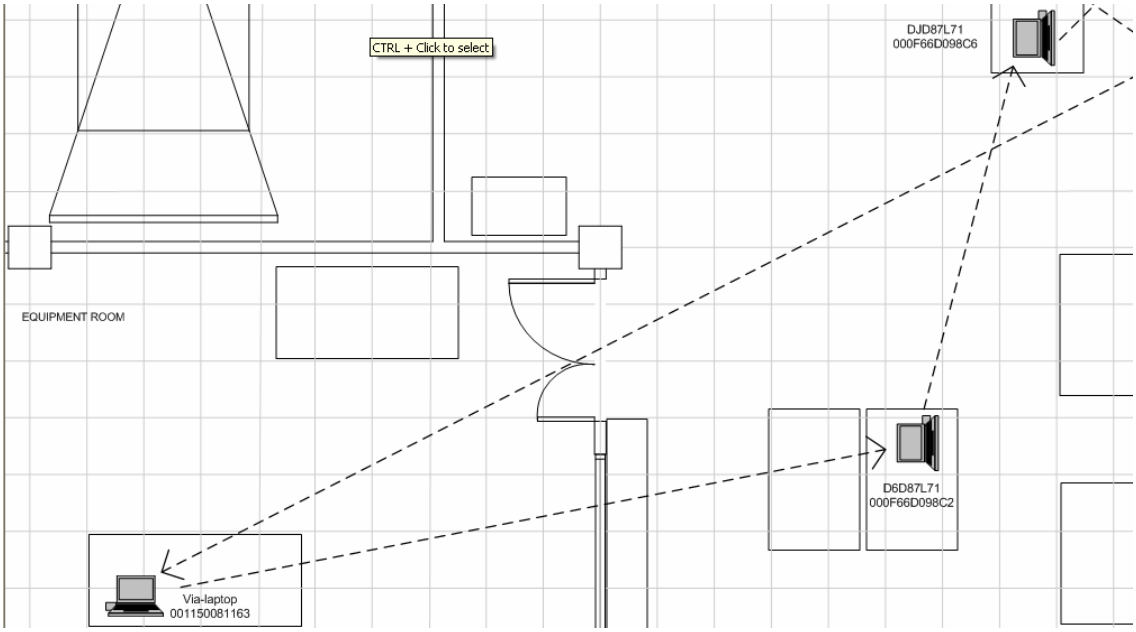

Figure B-11 Experiment 12 Detail 2

B-13 EXPERIMENT 13

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes placed in adjacent rooms. Within this experiment, there were seven regular laptops with one intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. The parabolic antenna was not used. The intruder was placed inside the storage room and was turned on after an hour of non-intrusion and was then turned off after another hour. Elapsed time: approximately 3 hours.

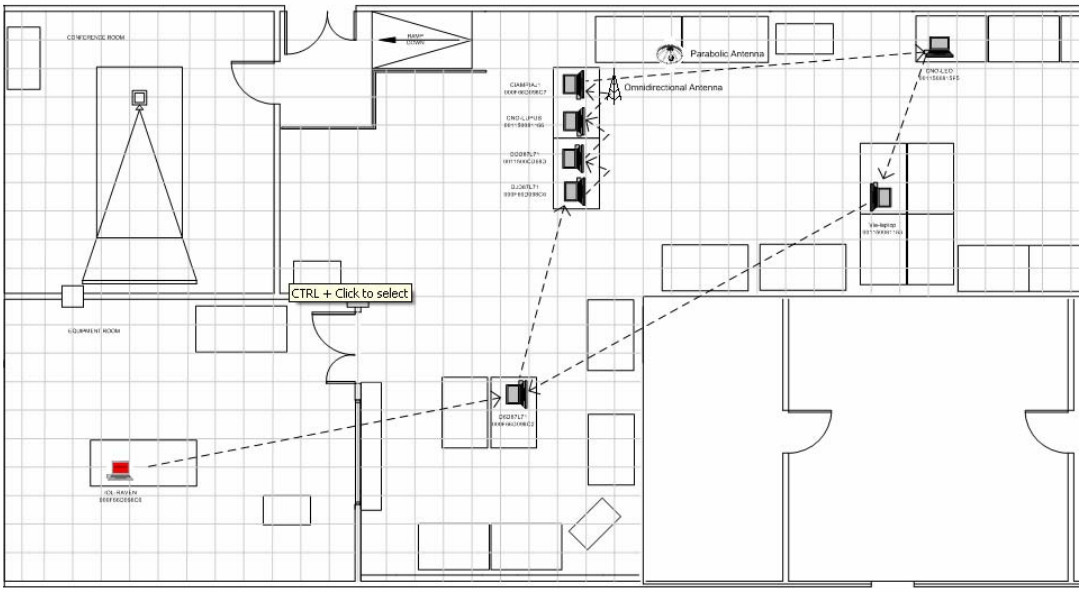

Figure B-12 Configuration for Experiment 13

B-14 EXPERIMENT 14

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes with vertically polarized antennas, placed in adjacent rooms. Within this experiment, there were seven regular laptops with one intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. The parabolic antenna was vertically polarized. The intruder was placed inside the storage room and was turned on after an hour of non-intrusion and was then turned off after another hour. Elapsed time: approximately 3 hours.

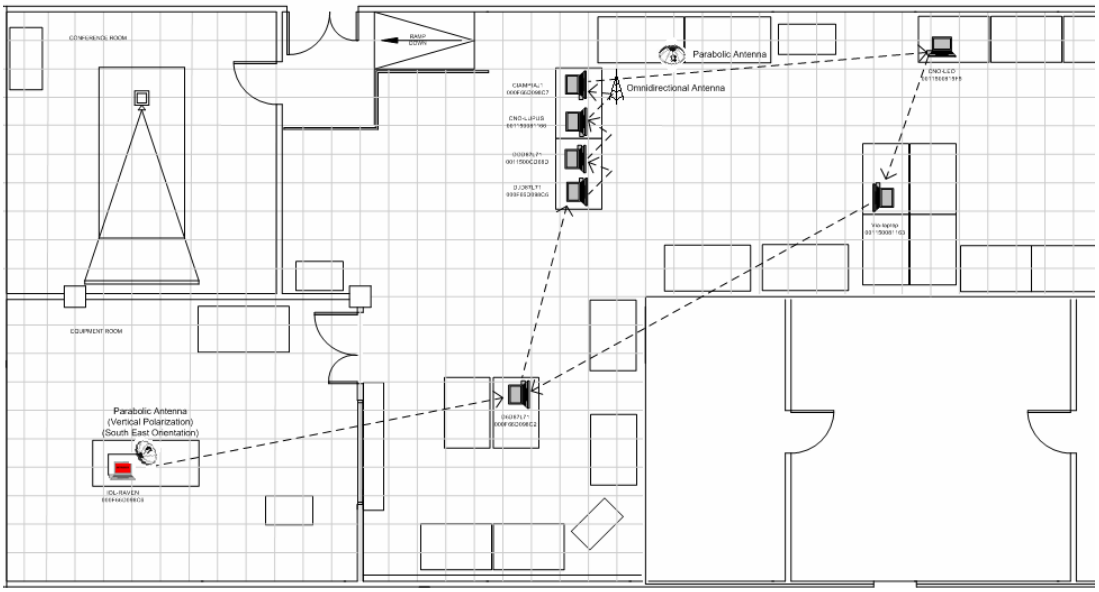

Figure B-13 Configuration for Experiment 14

B-15 EXPERIMENT 15

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes with horizontally polarized antennas, placed in adjacent rooms. Within this experiment, there were seven regular laptops with one intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. The parabolic antenna was horizontally polarized. The intruder was placed inside the storage room and was turned on after an hour of non-intrusion and was then turned off two hours later. Elapsed time: approximately 4 hours.

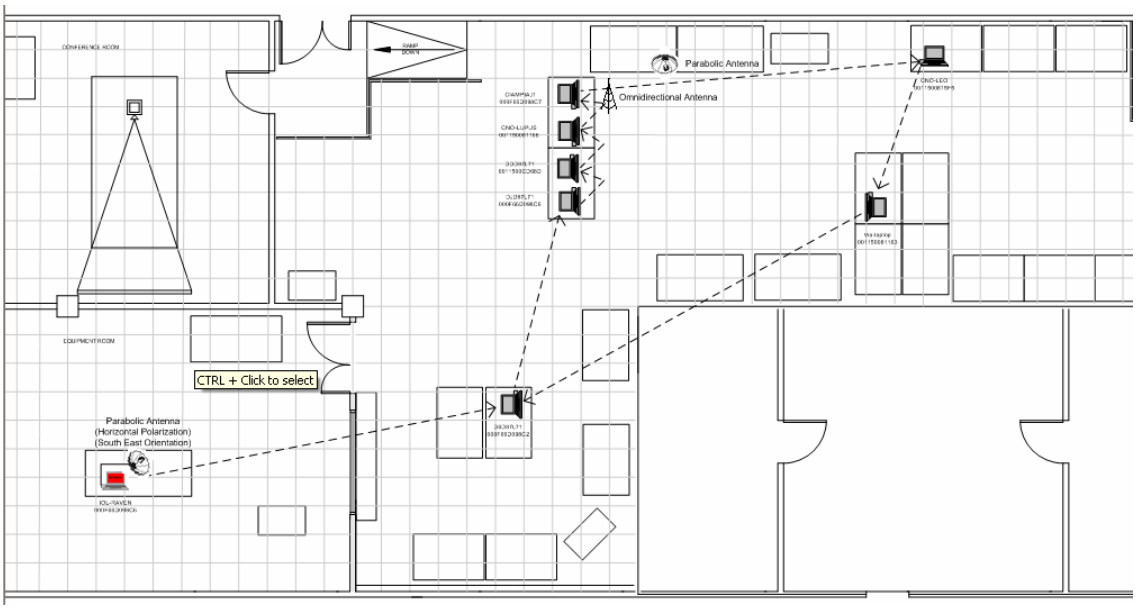

Figure B-14 Configuration for Experiment 15

B-16 EXPERIMENT 16

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes placed in adjacent rooms. Within this experiment, there were seven regular laptops with one intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. No antenna was used. The intruder was placed inside room 17-N430 and was turned on after a period of non-intrusion. Elapsed time: approximately 2.5 hours.

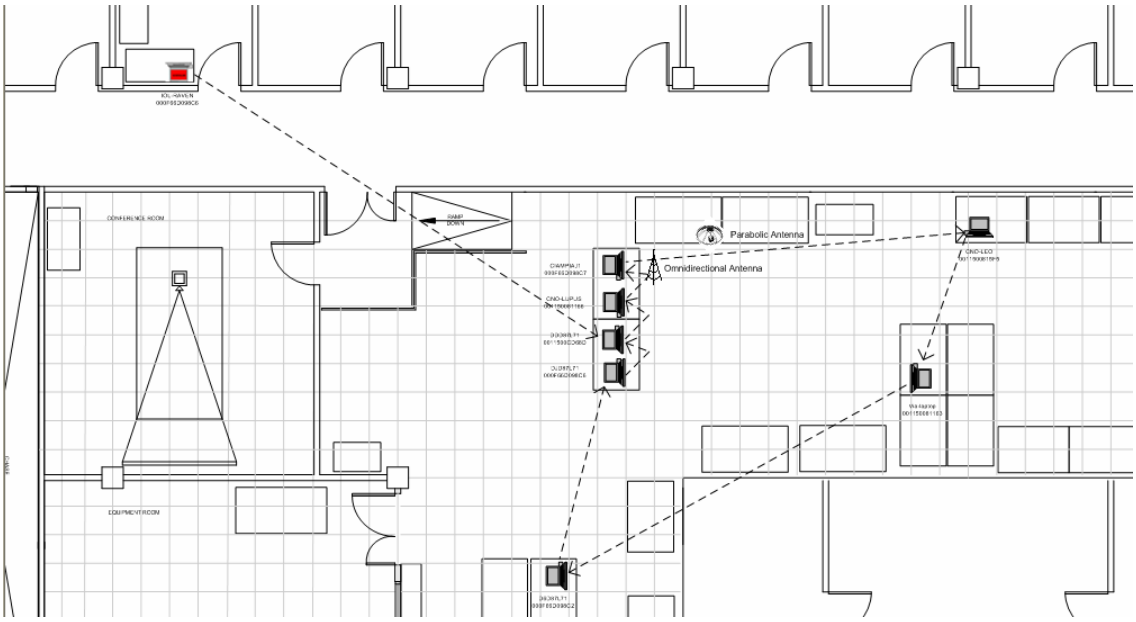

Figure B-15 Configuration for Experiment 16

B-17 EXPERIMENT 17

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes with horizontally polarized antennas, placed in adjacent rooms. Within this experiment, there were seven regular laptops with one intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. The parabolic antenna was horizontally polarized. The intruder was placed inside room 17-N430 and was turned on after a period of non-intrusion. Elapsed time: approximately 2.5 hours.

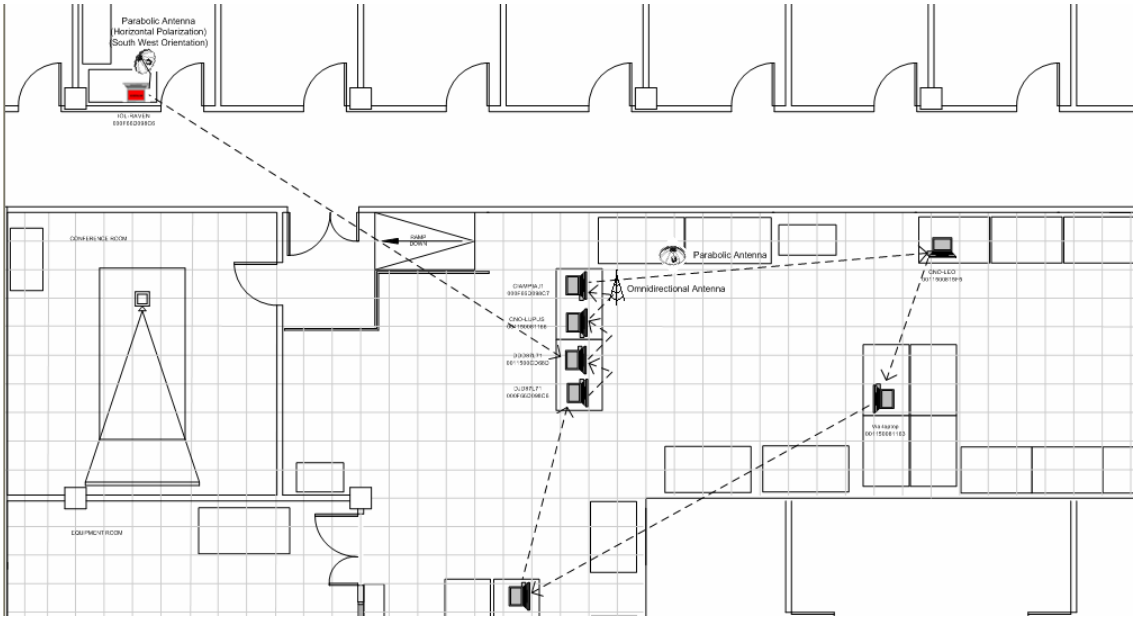

Figure B-16 Configuration for Experiment 17

B-18 EXPERIMENT 18

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes with vertically polarized antennas, placed in adjacent rooms. Within this experiment, there were seven regular laptops with one intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. The parabolic antenna was vertically polarized. The intruder was placed inside room 17-N430 and was turned on after a period of non-intrusion. Elapsed time: approximately 2.5 hours.

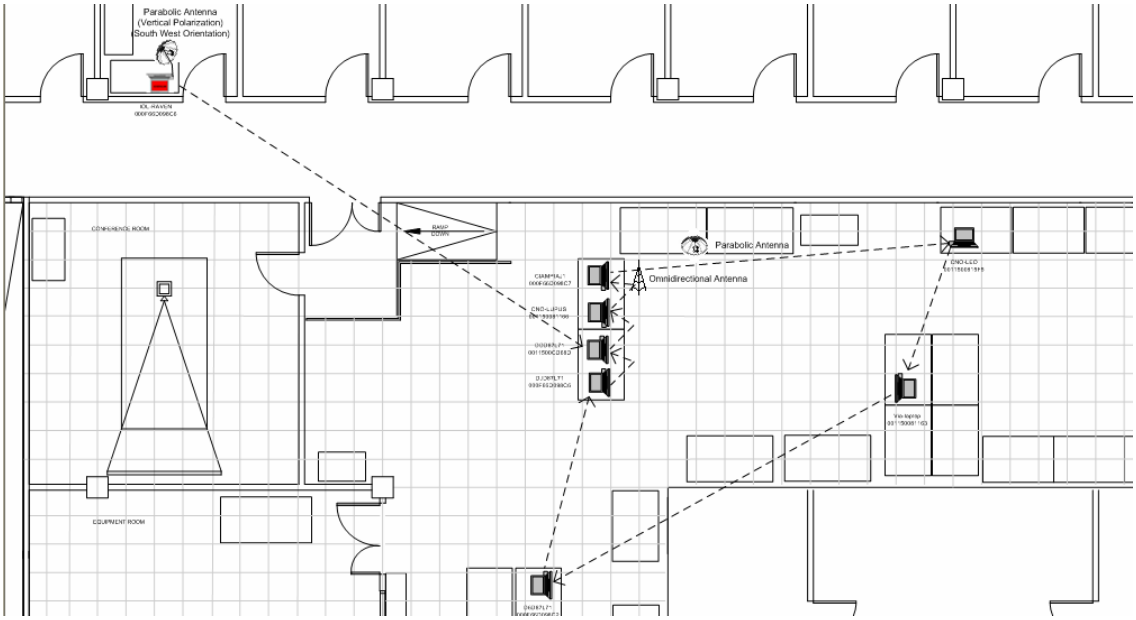

Figure B-17 Configuration for Experiment 18

B-19 EXPERIMENT 19

Purpose: To determine whether or not packets emitted by the Orinoco card are being detected by the receiver antennas. Within this experiment, there were eight regular laptops with no intruder. The network was mixed and consists of Linksys, Belkin, and Orinoco PC cards. The antenna that was connected to the Orinoco card was vertically polarized and both were placed inside room 17-N430. Elapsed time: approximately 3 hours.

Due to some "skipping" by the system, the data collected from experiment 19 was not able to be fully demodulated. However, the small amount of data that was able to be demodulated showed that the omni directional antenna was detecting packets emitted by the Orinoco PC card. However, the amount of packets triggered were still quite few in number (ten packets were triggered by the Orinoco card in the first 12 minutes).

B-20 EXPERIMENT 20

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes placed nearby the sensor. This experiment consists of seven regular computers and one intruder. The laptops were constantly pinging each other and the network consists of a mix of Linksys, Belkin, and Orinoco cards. No antenna was used, and the intruder was placed inside the main CNO laboratory. The non-intrusive laptops were scattered throughout the laboratory and some were placed in adjacent rooms as well. Elapsed time: approximately 3 hours.

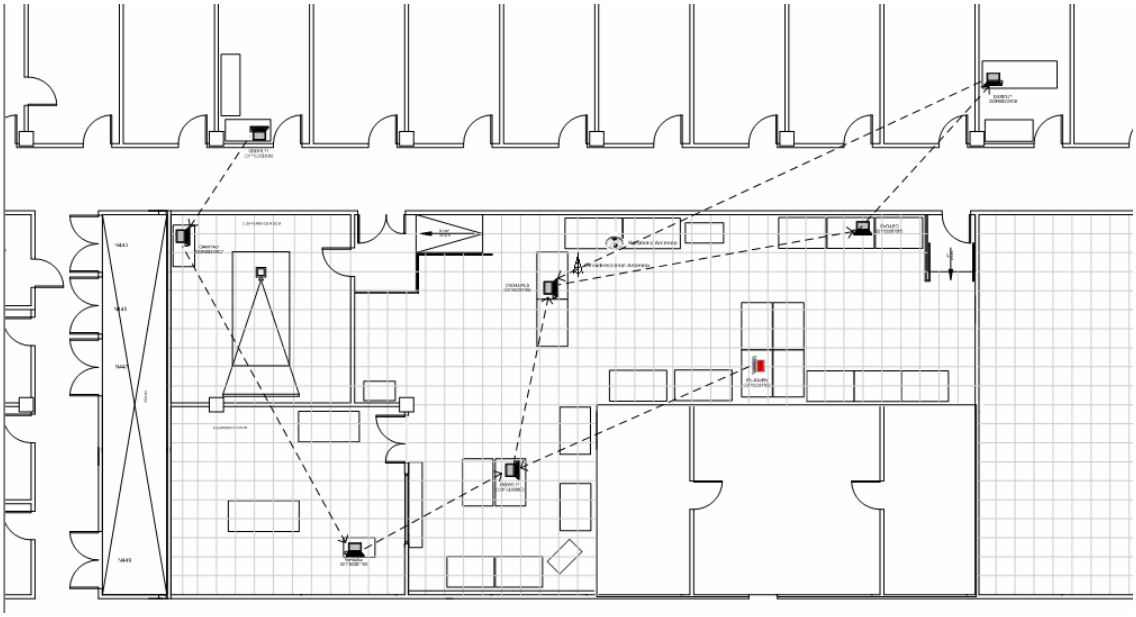

Figure B-18 Configuration for Experiment 20

B-21 EXPERIMENT 21

Purpose: To investigate how the radio frequency characteristics deviate over long periods of time for each PC card. This experiment consists of eight regular laptops. All laptops were placed in close proximity to the omni directional antenna and were constantly pinging each other. This was an overnight experiment. Temperature readings were saved for this experiment. Elapsed time: approximately 16 hours.

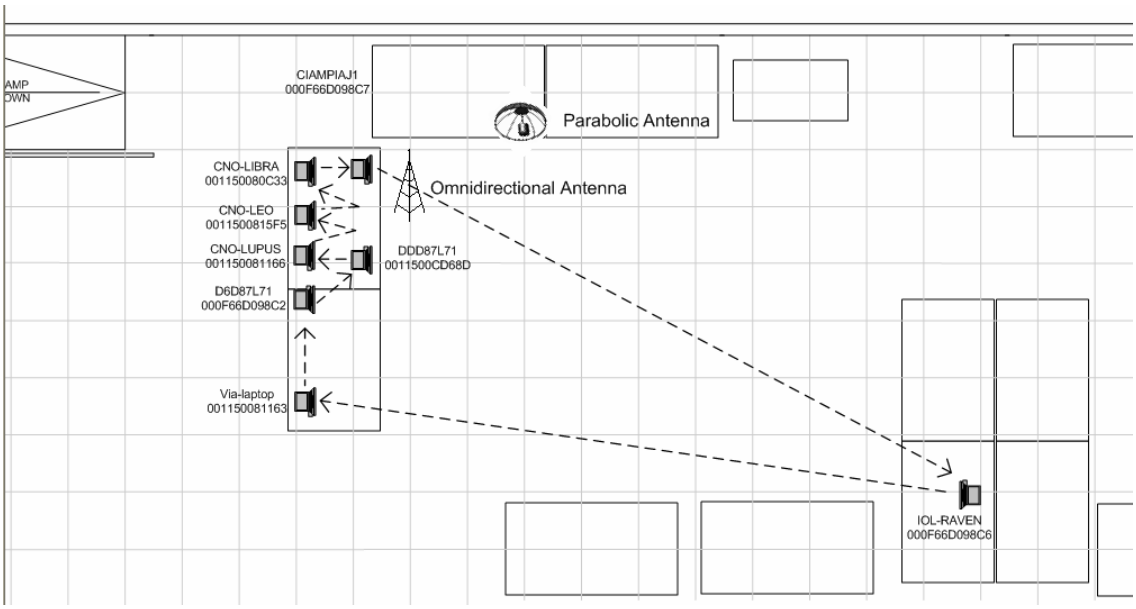

Figure B-19 Configuration for Experiment 21

B-22 EXPERIMENT 22

Purpose: To investigate a larger mixed network consisting of cards from several manufacturers utilizing intruder nodes with horizontally polarized antennas, placed nearby the sensor. This experiment consists of seven regular computers and one intruder. The laptops were constantly pinging each other and the network was made up of a mix of Linksys, Belkin, and Orinoco cards. The antenna was horizontally polarized and the intruder was placed inside the main CNO laboratory. The non-intrusive laptops were scattered throughout the laboratory and some were placed in adjacent rooms as well. Elapsed time: approximately 3 hours.

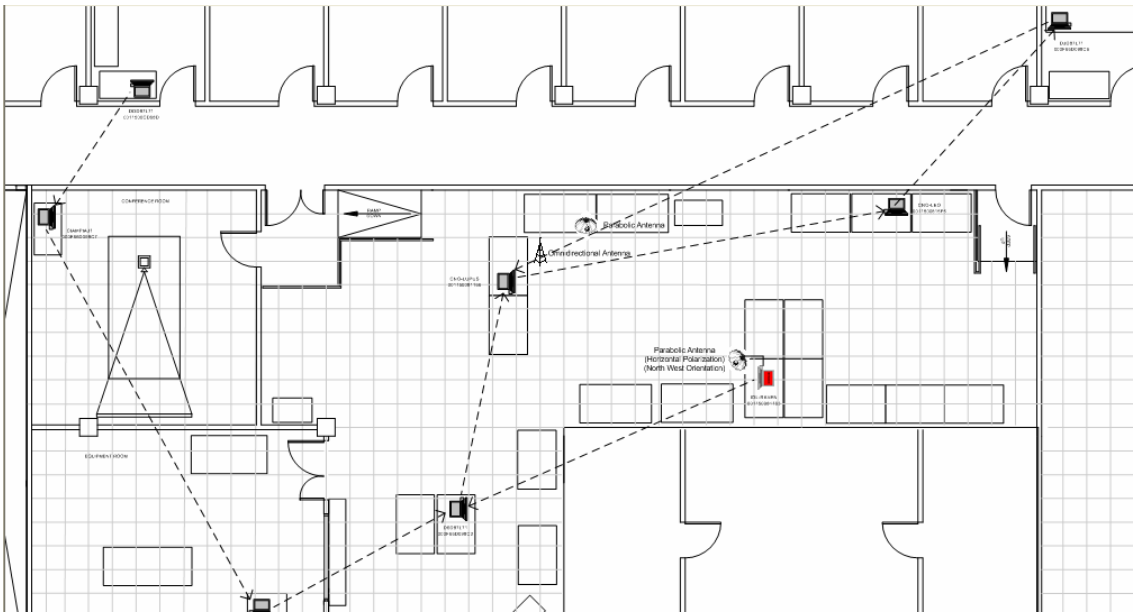

Figure B-20 Configuration for Experiment 22

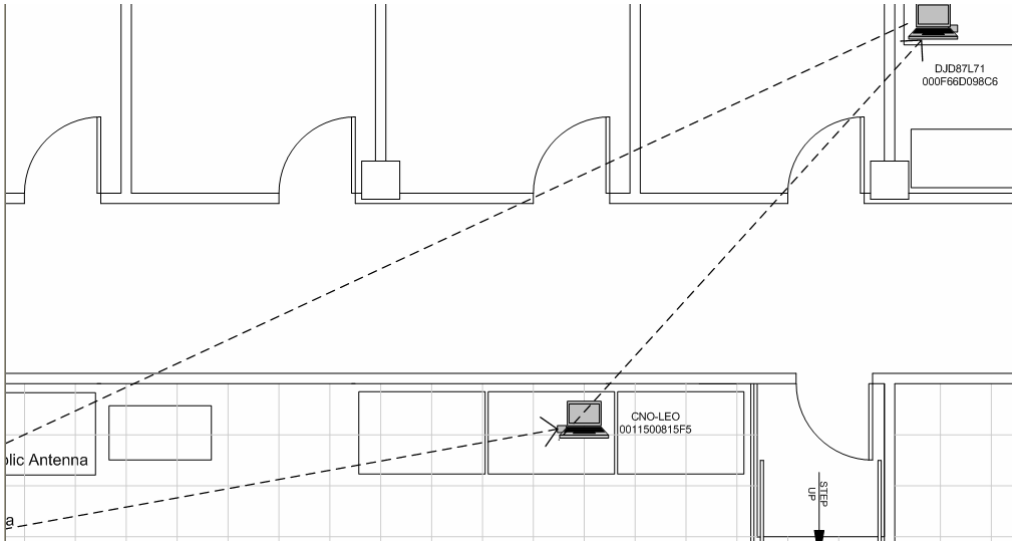

Figure B-21 Experiment 22 Detail 1

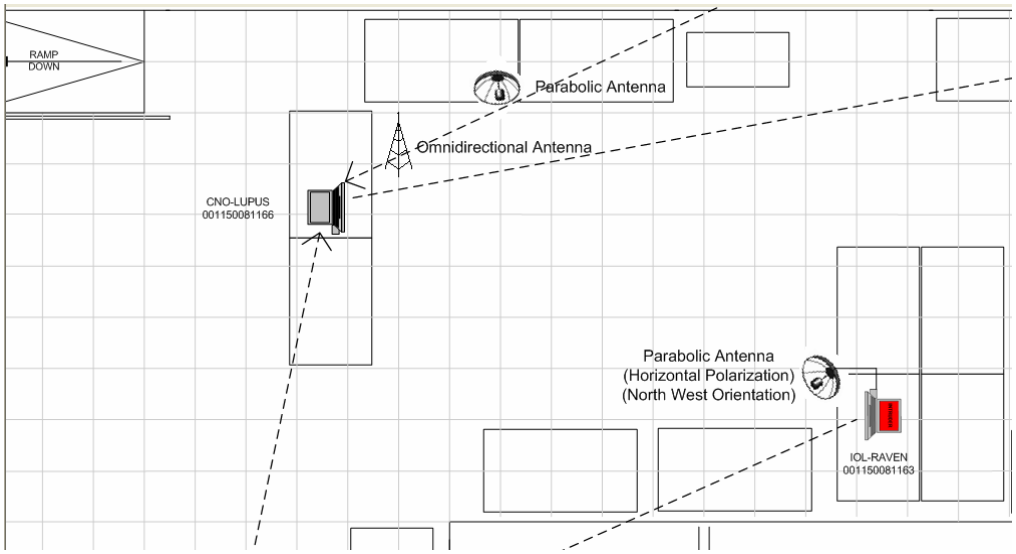

Figure B-22 Experiment 22 Detail 2

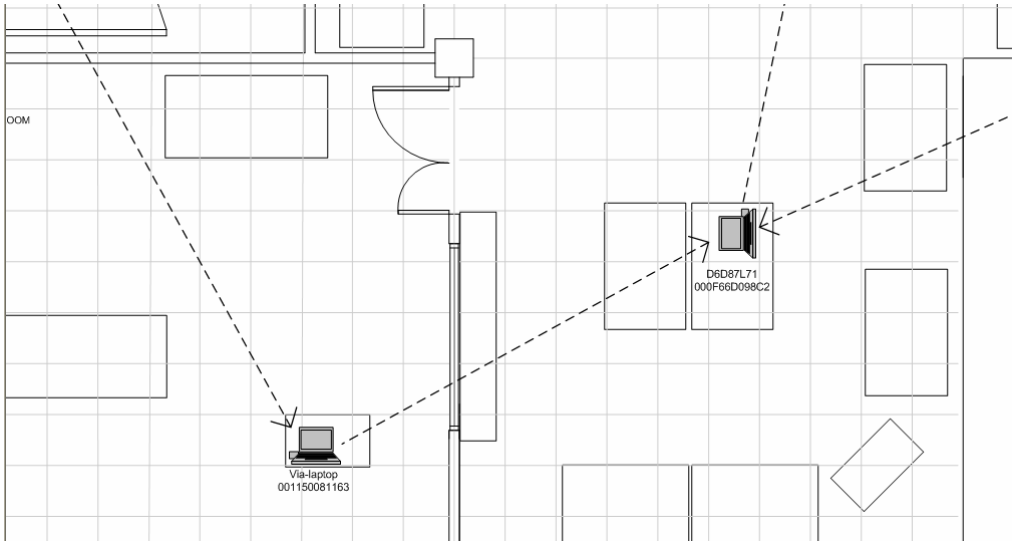

Figure B-23 Experiment 22 Detail 3

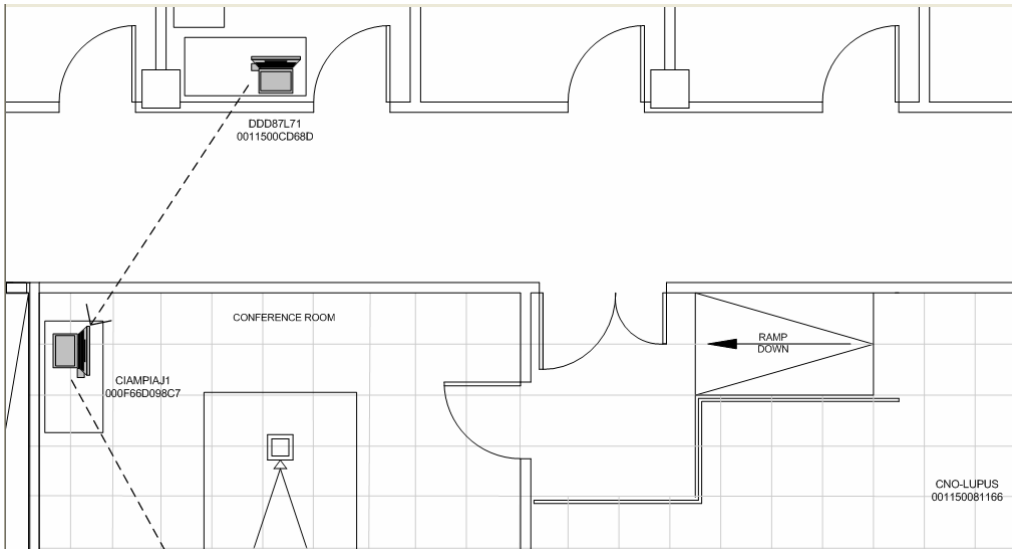

Figure B-24 Experiment 22 Detail 4

B-23 EXPERIMENT 23

Purpose: To investigate how the radio frequency characteristics deviate over long periods of time for each PC card. This experiment consists of seven regular laptops. All laptops were placed in close proximity to the omni directional antenna and were constantly pinging each other. This was an overnight experiment. Temperature readings were saved for this experiment, but were lost somehow during file transfer. Elapsed time: approximately 16 hours.

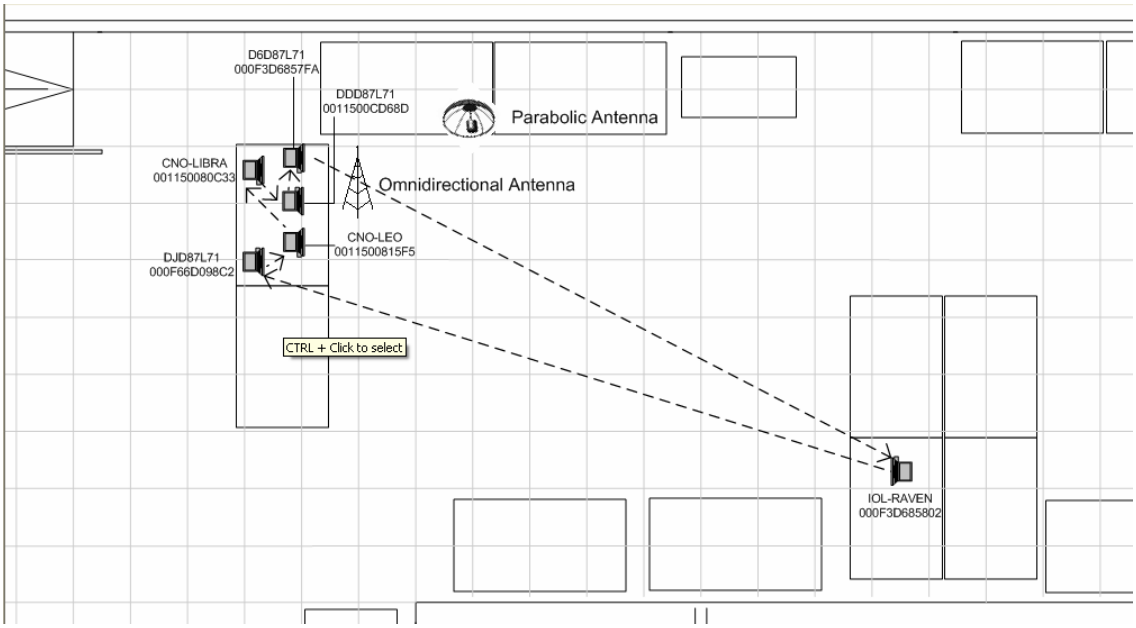

Figure B-25 Configuration for Experiment 23

B-24 EXPERIMENT 24

Purpose: To investigate how the radio frequency characteristics deviate over long periods of time for each PC card. This experiment consists of four regular laptops. All laptops were placed in close proximity to the omni directional antenna and were constantly pinging each other. This was an overnight experiment. Temperature readings were saved for this experiment. Elapsed time: approximately 16 hours.

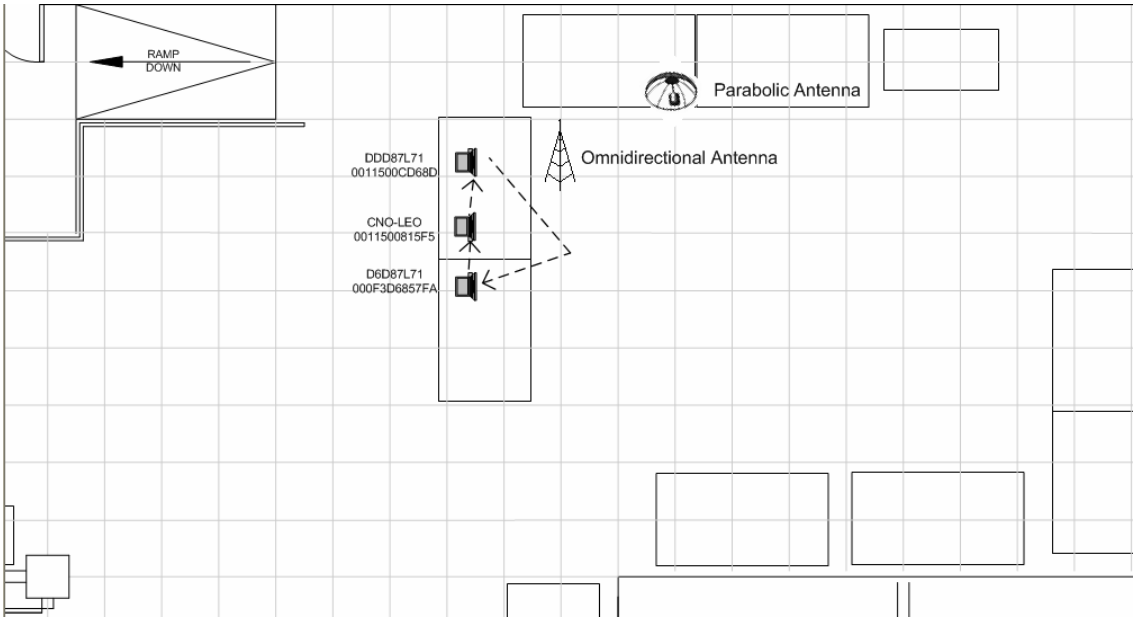

Figure B-26 Configuration for Experiment 24

Appendix C

SHORT-TIME SCALE INTRUSION DETECTION EXPERIMENTS

This appendix describes a series of ten experiments that were conducted to examine packet feature stability, uniqueness, and sensitivity to antenna choice and network configuration. These experiments utilized the new JHU/APL high-speed packet capture system described in Section 2 (see Figure 2-2). The experiments focused on packet RF features measured over short-time scales (<1 minute). Four features were considered: frequency error, rise-time, received power, and IQ offset. In each of the experiments packet features were measured for a set of PC Card format WiFi cards configured to operate as IEEE 802.11b devices. Channel 6 (2.437 GHz) and Ad Hoc mode were used for all experiments except Experiment 34, which employed Channel 11 (2.462 GHz) and Infrastructure mode.

In each experiment, the NI PXI system was used to capture a small sample (typically about 40 seconds) of RF data from several antennas. The RF data were down-converted to a 22 MHz bandwidth centered on an intermediate frequency (IF) of 15 MHz, digitized at 55 MS/s using a 14 bit A/D, and streamed continuously to a disk array. The digitized data were subsequently broken into 5 ms segments for analysis. Agilent 89600 VSA Software was used to detect and demodulate the first valid packet in each segment, and to extract the packet features. A feature vector was constructed for each detected packet from each antenna. It consists of the segment number, time of detection (to the nearest 5 ms), MAC address, the four features, and the additional Agilent 89600 VSA Software parameter outputs listed in Section 2. The features vectors from each antenna were sorted by MAC address, and a histogram was constructed for each combination of feature, MAC address and antenna. The bin widths used to construct the histograms were 0.2 kHz for frequency error, 0.02 μs for rise-time, 0.25 dB for received power, and 0.5 dB for IQ offset. Each histogram was normalized to the total number of packets detected. A histogram was not constructed if the total number of packets detected was less than 5**.**

In Experiments 25 through 28, the host laptop PCs were configured in a compact spatial arrangement (Figure C-1) within the CNO laboratory. In Experiments 29 through 33, the host laptop PCs were dispersed about the CNO laboratory (Figures C-2 and C-3). Experiment 34 was conducted in an "Internet Café" setting on the JHU/APL campus, with some of the laptops on an outdoor patio, and others inside a large open cafeteria (Figure C-4).

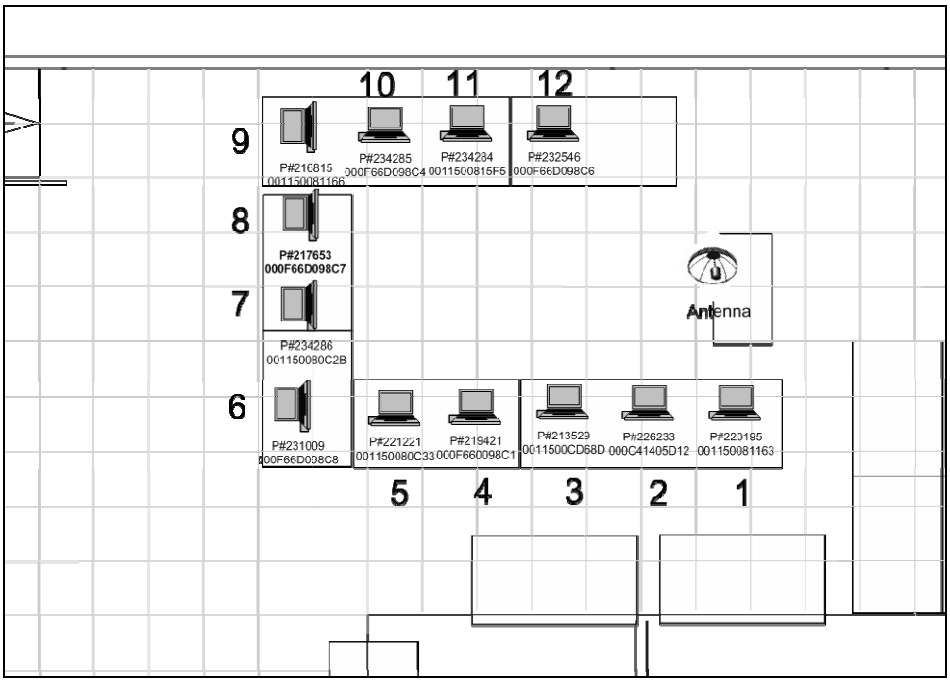

Figure C-1 Network Configuration for Experiments 25 through 28

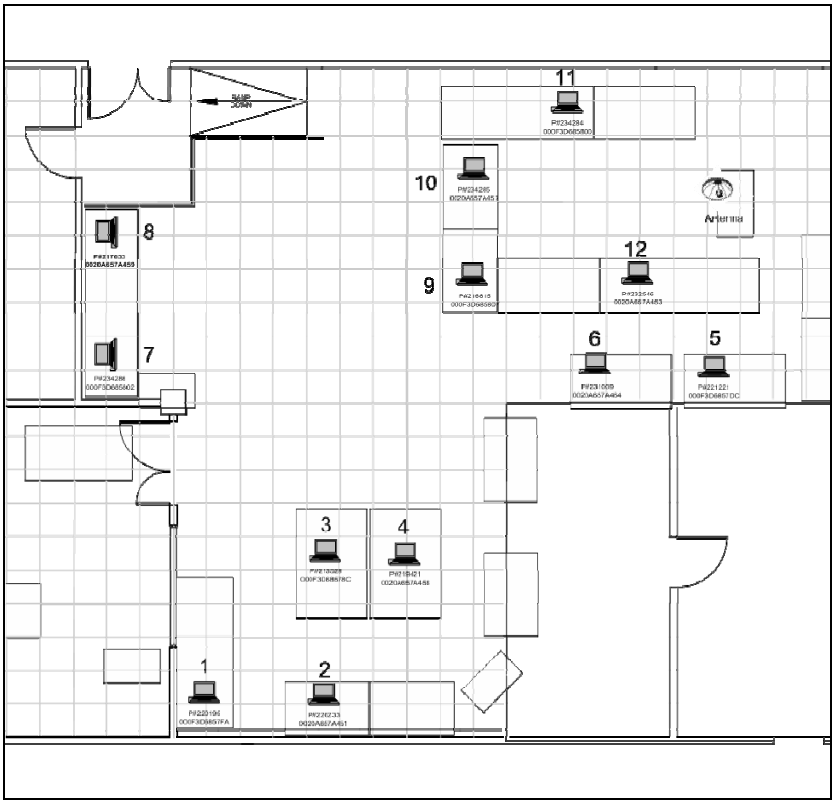

Figure C-2 Network Configuration for Experiments 29 through 32

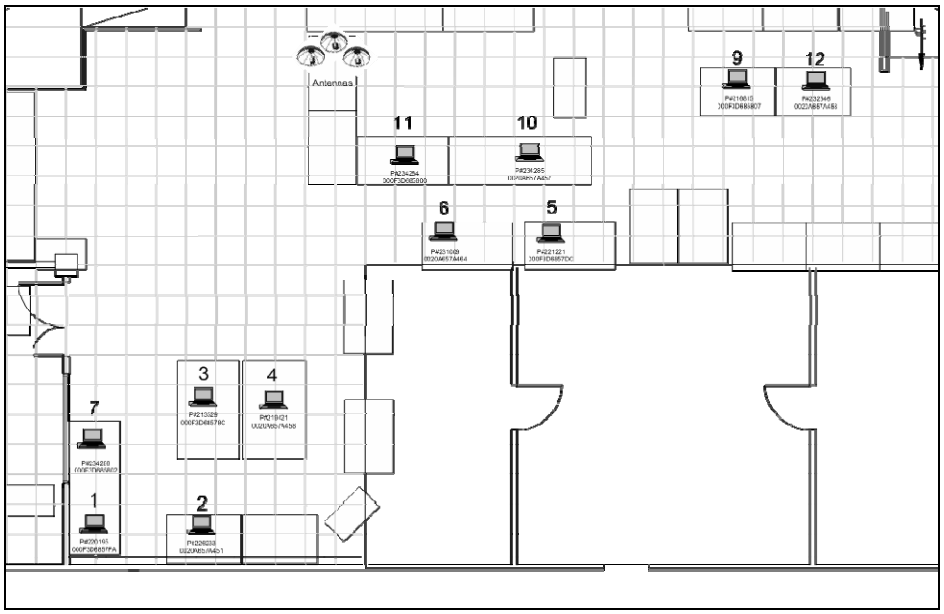

Figure C-3 Network Configuration for Experiment 33

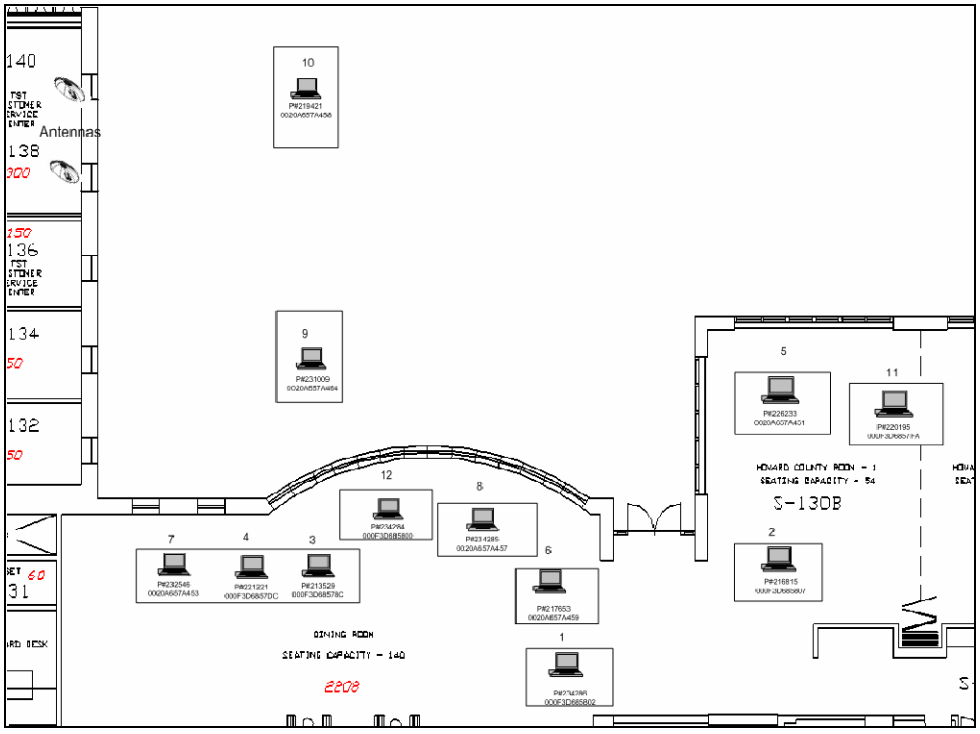

Figure C-4 Network Configuration for Experiment 34

C-1 EXPERIMENT 25

Experiment 25 was a two channel measurement using a dual polarization horn antenna. Table C-1 provides a summary of the packets captured for each card. It lists the host laptop property number, card MAC address, card type, and the number of packets detected on each channel. Channel 1 (Ch1) is the horizontal polarization, and Channel 2 (Ch2) is the vertical polarization. The network configuration is the compact arrangement shown in Figure C-1. The detected Linksys packets were all Burst Type 3 (CCK 11), while only about 11% of the Belkin packets were CCK11, with the remainder being Burst Type 0 (Barker 1). The lower data rate Barker 1 packets were obviously more easily detected, as reflected in the much higher packet counts for the Belkin cards compared to the Linksys cards.

Table C-1 Experiment 25 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	001150081163	Belkin	Ch1ID01:379	Ch2ID01:379
2	226233	000C41405D12	Linksys'	Ch1ID02:31	Ch2ID02:19
3	213529	0011500CD68D	Belkin	Ch1ID03:188	Ch2ID03:181
4	219421	000F66D098C1	Linksys	Ch1ID04:18	Ch2ID04:35
5	221221	001150080C33	Belkin	Ch1ID05:143	Ch2ID05:79
6	231009	000F66D098C8	Linksys	Ch1ID06:30	Ch2ID06:19
7	234286	001150080C2B	Belkin	Ch1ID07:109	Ch2ID07:155
8	217653	000F66D098C7	Linksys	Ch1ID08:27	Ch2ID08:26
9	216815	001150081166	Belkin	Ch1ID09:302	Ch2ID09:339
10	234285	000F66D098C4	Linksys	Ch1ID10:36	Ch2ID10:27
11	234284	0011500815F5	Belkin	Ch1ID11:102	Ch2ID11:46
12	232546	000F66D098C6	Linksys	Ch1ID12:36	Ch2ID12:31

Belkin = Belkin Wireless Notebook Network Card IEEE 802.11g/54 Mbps Model F5D7010 Version 1315
Linksys' = Linksys Wireless-B Notebook Adapter Model No: WPC11 Version 4
Linksys = Linksys Wireless-G Notebook Adapter Model No: WPC54G Version 2

Figures C-5 through C-12 are pairs of figures showing measured feature distributions for four features: frequency error, rise-time, received power and IQ offset. The first figure of each pair is for the Belkin cards, while the second figure is for the Linksys cards. There are two graphs on each figure. The upper graph (Ch 2) is from the horn vertical polarization, while the lower graph (Ch 1) is from the horn horizontal polarization.

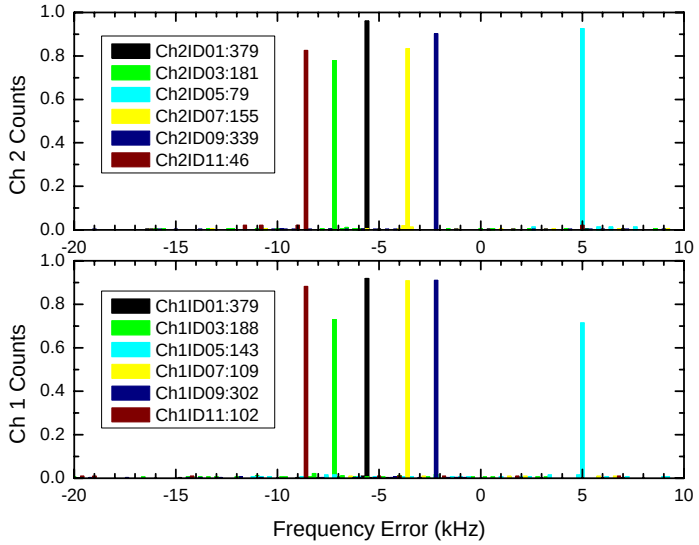

Figure C-5 Frequency Error for Belkin Cards (Experiment 25)

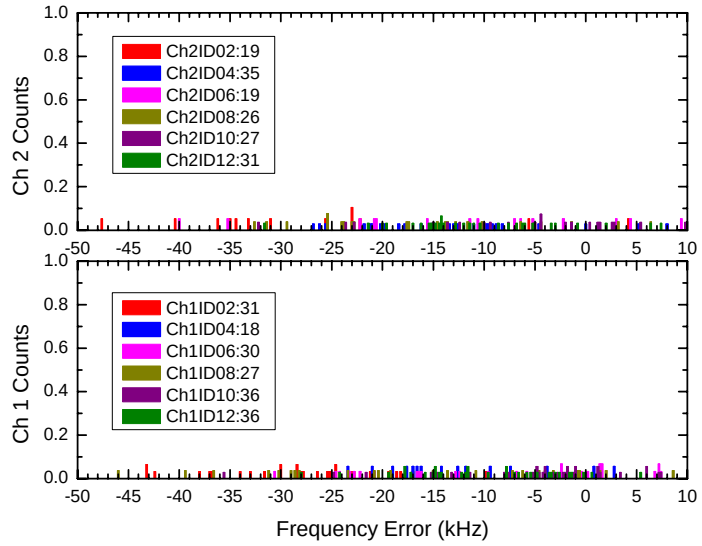

Figure C-6 Frequency Error for Linksys Cards (Experiment 25)

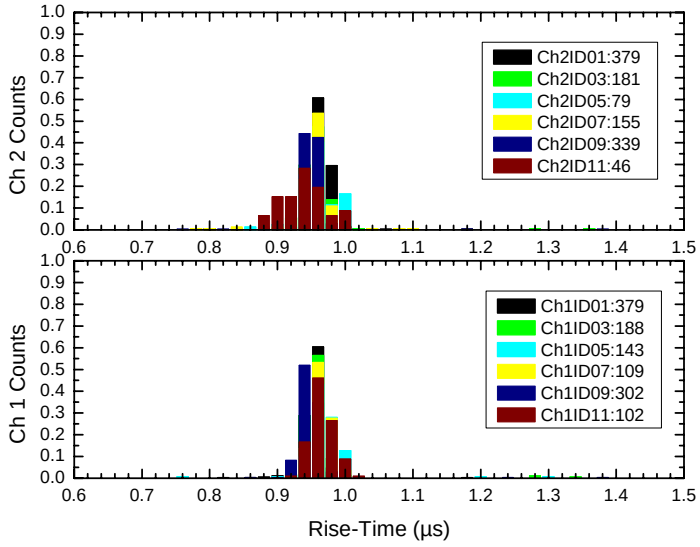

Figure C-7 Rise-Time for Belkin Cards (Experiment 25)

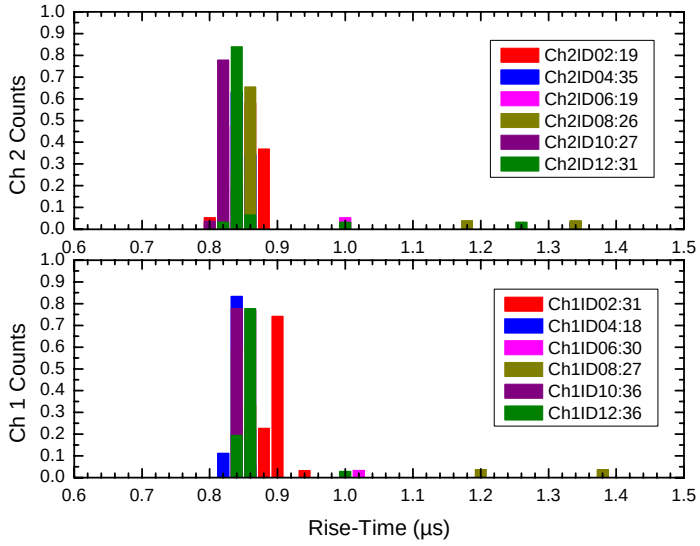

Figure C-8 Rise-Time for Linksys Cards (Experiment 25)

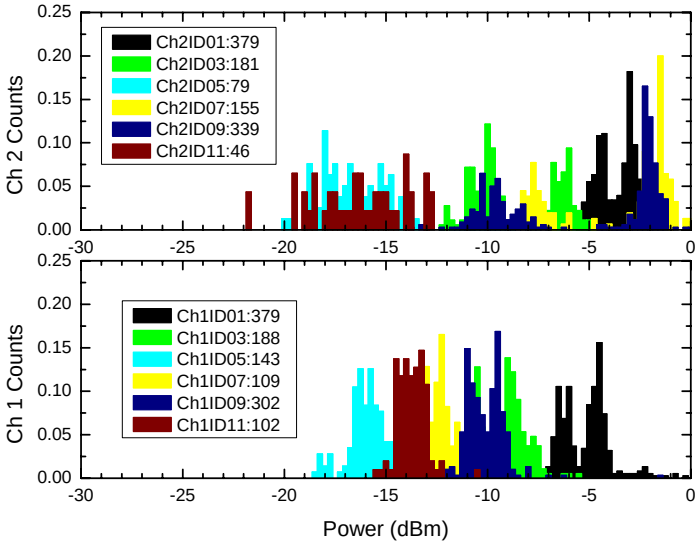

Figure C-9 Received Power for Belkin Cards (Experiment 25)

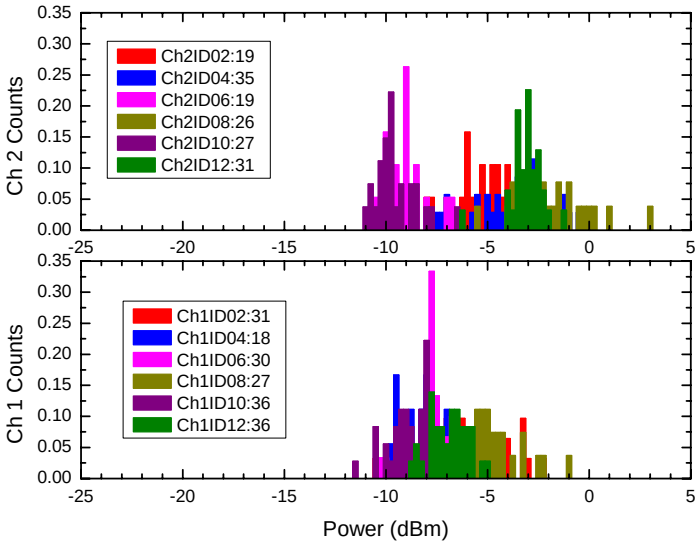

Figure C-10 Received Power for Linksys Cards (Experiment 25)

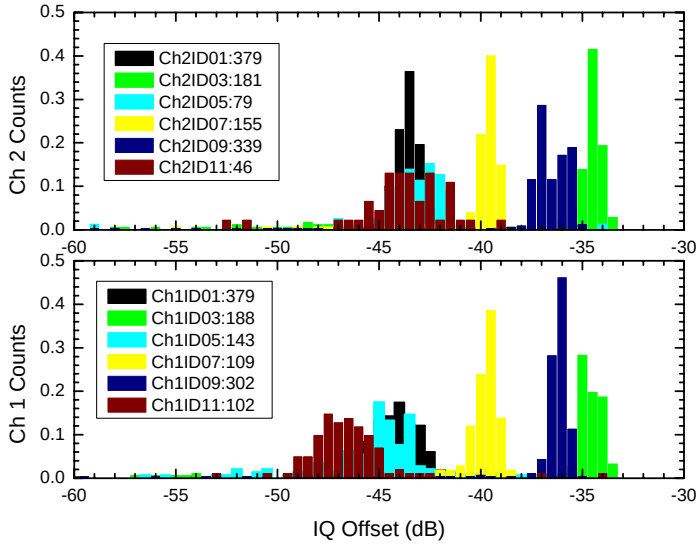

Figure C-11 IQ Offset for Belkin Cards (Experiment 25)

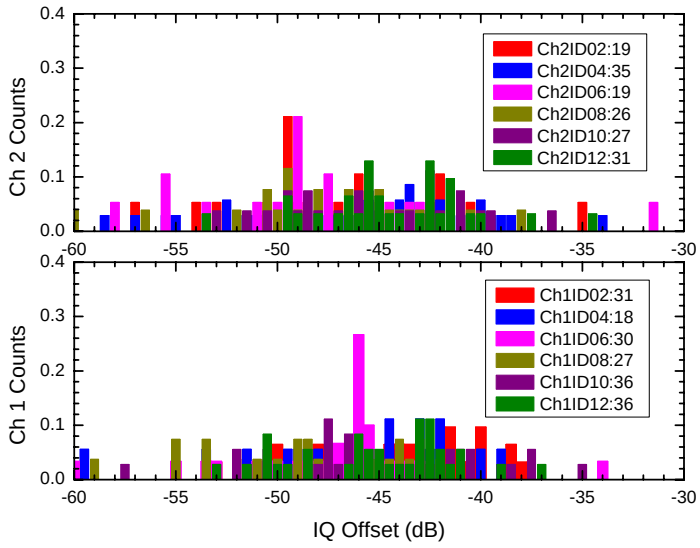

Figure C-12 IQ Offset for Linksys Cards (Experiment 25)

C-2 EXPERIMENT 26

Experiment 26 was a two channel measurement of packet features from the same twelve WiFi cards as in Experiment 25, and in the same compact spatial arrangement. However, for this experiment, Channel 1 was connected to a vertically polarized omni-directional antenna, while Channel 2 was attached to a horizontally polarized parabolic antenna. Table C-2 provides a summary of the packets captured for each card.

Table C-2 Experiment 26 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	001150081163	Belkin	Ch1ID01:91	Ch2ID01:141
2	226233	000C41405D12	Linksys'	Ch1ID02:70	Ch2ID02:55
3	213529	0011500CD68D	Belkin	Ch1ID03:38	Ch2ID03:68
4	219421	000F66D098C1	Linksys	Ch1ID04:10	Ch2ID04:42
5	221221	001150080C33	Belkin	Ch1ID05:310	Ch2ID05:203
6	231009	000F66D098C8	Linksys	Ch1ID06:8	Ch2ID06:32
7	234286	001150080C2B	Belkin	Ch1ID07:69	Ch2ID07:74
8	217653	000F66D098C7	Linksys	Ch1ID08:22	Ch2ID08:9
9	216815	001150081166	Belkin	Ch1ID09:236	Ch2ID09:258
10	234285	000F66D098C4	Linksys	Ch1ID10:29	Ch2ID10:36
11	234284	0011500815F5	Belkin	Ch1ID11:55	Ch2ID11:16
12	232546	000F66D098C6	Linksys	Ch1ID12:13	Ch2ID12:3

Belkin = Belkin Wireless Notebook Network Card IEEE 802.11g/54 Mbps Model F5D7010 Version 1315
Linksys' = Linksys Wireless-B Notebook Adapter Model No: WPC11 Version 4
Linksys = Linksys Wireless-G Notebook Adapter Model No: WPC54G Version 2

Figures C-13 through C-20 are pairs of figures showing measured feature distributions for four features. The first figure of each pair is for the Belkin cards, while the second figure is for the Linksys cards. There are two graphs on each figure. The upper graph (Ch2) is from the horizontally polarized parabolic antenna, while the lower graph (Ch 1) is from the omni-directional antenna.

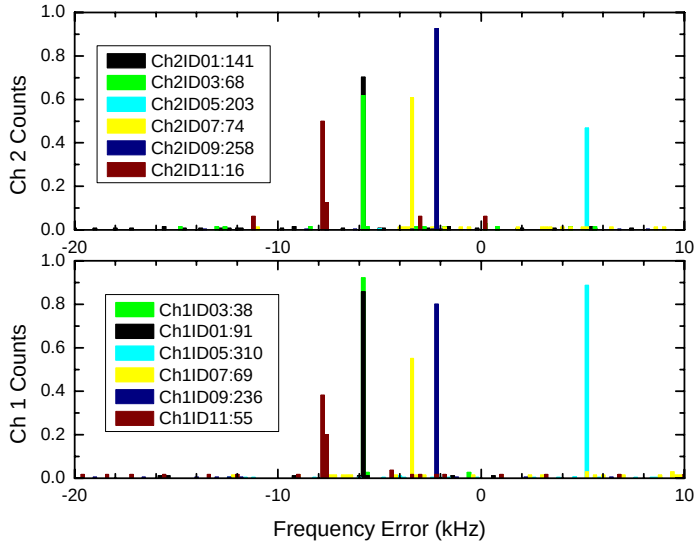

Figure C-13 Frequency Error for Belkin Cards (Experiment 26)

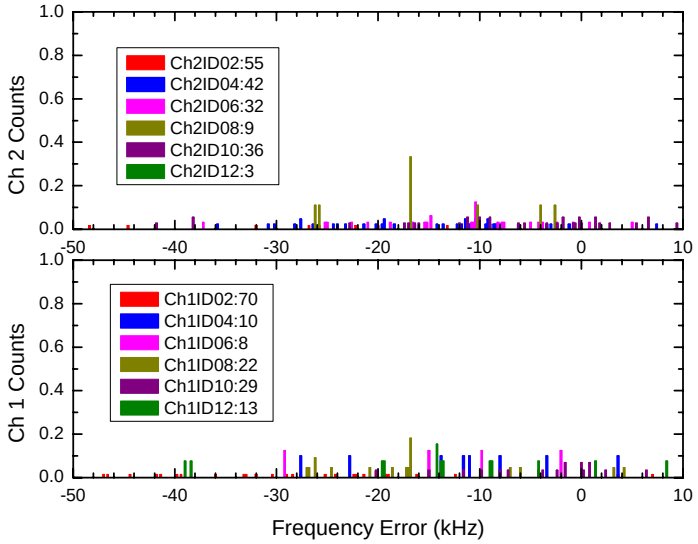

Figure C-14 Frequency Error for Linksys Cards (Experiment 26)

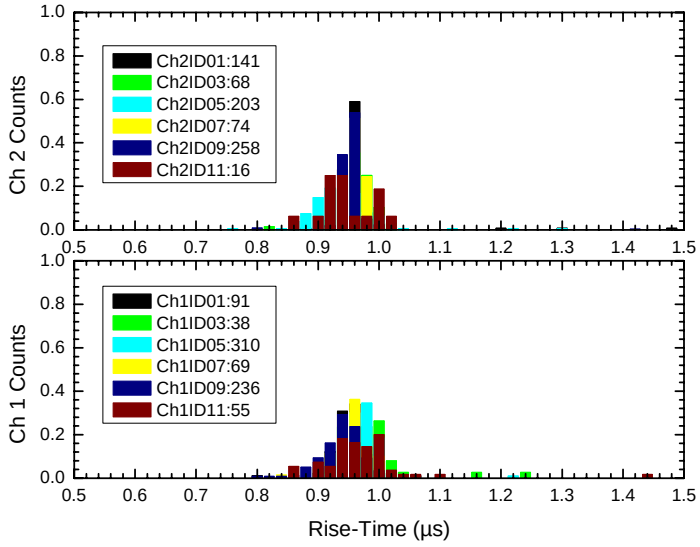

Figure C-15 Rise-Time for Belkin Cards (Experiment 26)

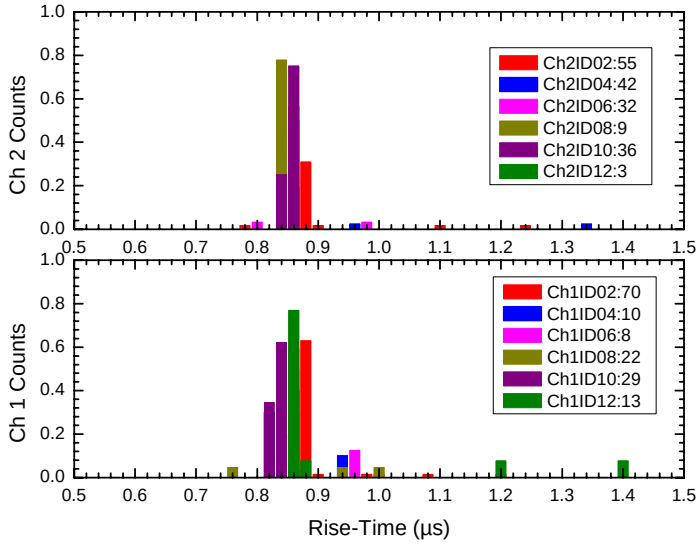

Figure C-16 Rise-Time for Linksys Cards (Experiment 26)

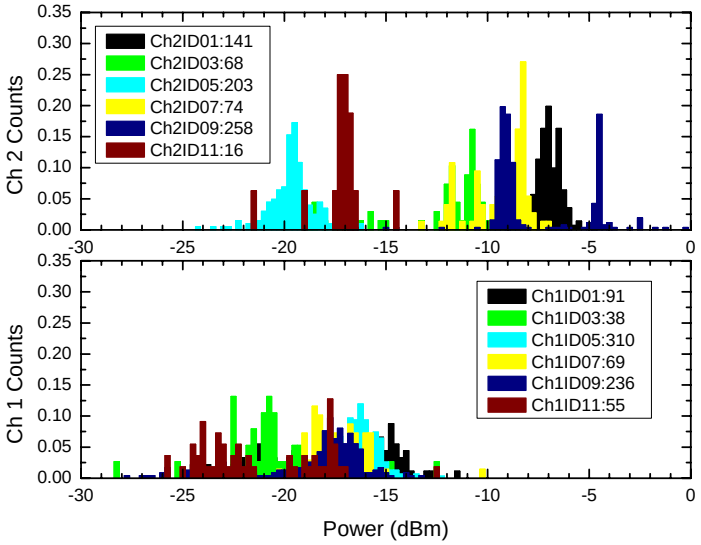

Figure C-17 Received Power for Belkin Cards (Experiment 26)

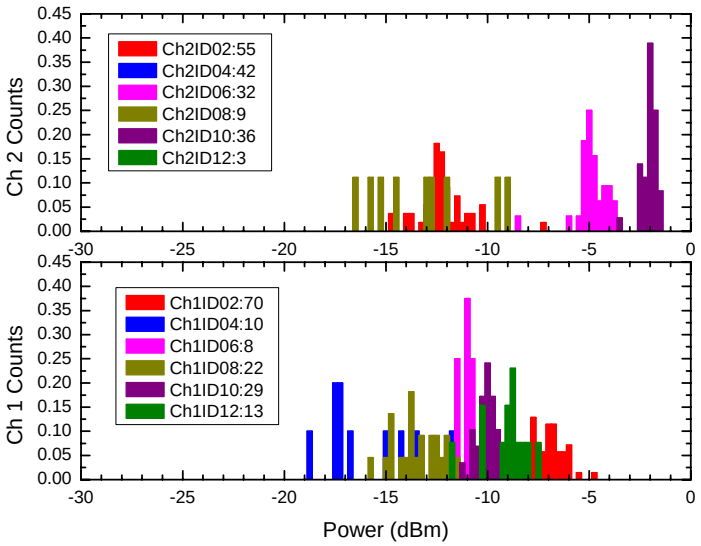

Figure C-18 Received Power for Linksys (Experiment 26)

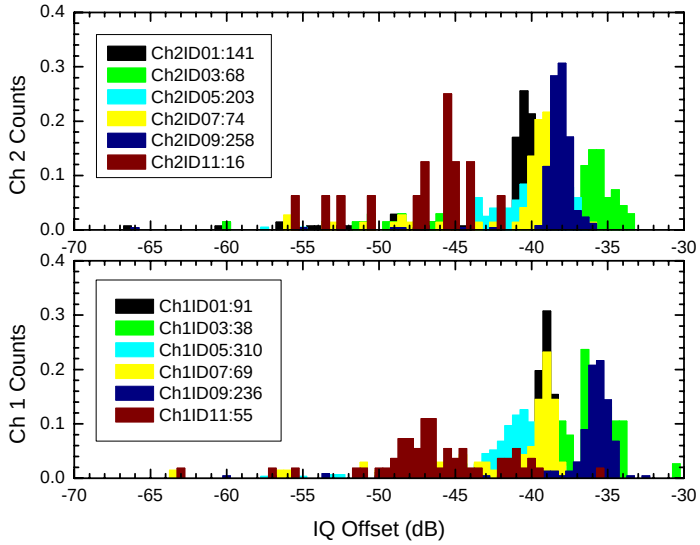

Figure C-19 IQ Offset for Belkin Cards (Experiment 26)

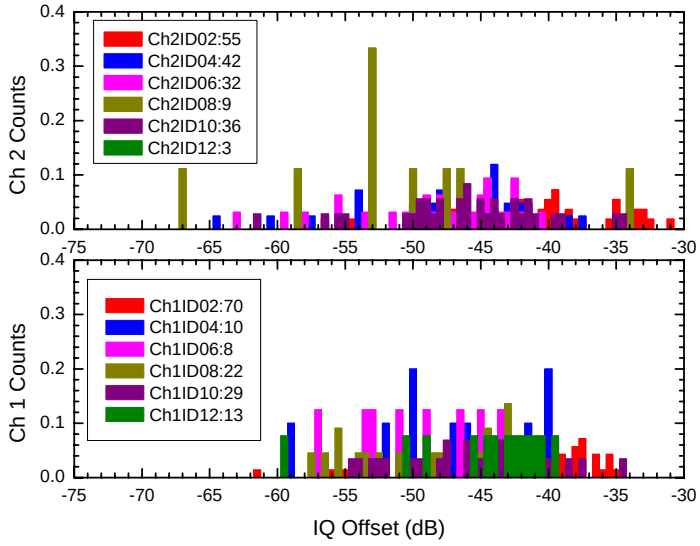

Figure C-20 IQ Offset for Linksys Cards (Experiment 26)

C-3 EXPERIMENT 27

Experiment 27 was a two channel measurement of packet features from a second group of twelve WiFi cards, with six from D-Link and six from Orinoco. The network configuration is that as shown in Figure C-1. Channel 1 was connected to a vertically polarized omni-directional antenna, while Channel 2 was attached to a horizontally polarized parabolic antenna.

Table C-3 Experiment 27 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	000F3D6857FA	D-Link	Ch1ID01:50	Ch2ID01:58
2	226233	0020A657A451	Orinico	Ch1ID02:67	Ch2ID02:44
3	213529	000F3D68578C	D-Link	Ch1ID03:65	Ch2ID03:72
4	219421	0020A657A458	Orinico	Ch1ID04:15	Ch2ID04:66
5	221221	000F3D6857DC	D-Link	Ch1ID05:27	Ch2ID05:31
6	231009	0020A657A464	Orinico	Ch1ID06:58	Ch2ID06:81
7	234286	000F3D685802	D-Link	Ch1ID07:44	Ch2ID07:40
8	217653	0020A657A459	Orinico	Ch1ID08:70	Ch2ID08:72
9	216815	000F3D685807	D-Link	Ch1ID09:41	Ch2ID09:46
10	234285	0020A657A457	Orinico	Ch1ID10:15	Ch2ID10:54
11	234284	000F3D685800	D-Link	Ch1ID11:53	Ch2ID11:42
12	232546	0020A657A453	Orinico	Ch1ID12:26	Ch2ID12:80

D-Link = D-Link AirPlusG DWL-G630 P/N: BWLG630NA.C1 H/W Version C1 F/W Version 3.00
Orinico = Proxim Gold ORiNOCO 11b/g PC Card Model: 8470-FC PN: 67877/1

Figures C-21 through C-28 are pairs of figures showing measured feature distributions for four features. The first figure of each pair is for the D-Link cards, while the second figure is for the Orinoco cards. There are two graphs on each figure. The upper graph (Ch2) is from the horizontally polarized parabolic antenna, while the lower graph (Ch 1) is from the omni-directional antenna.

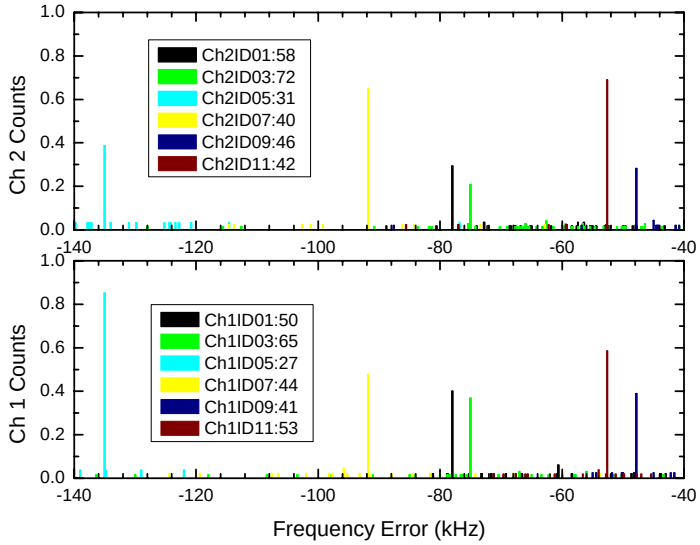

Figure C-21 Frequency Error for D-Link Cards (Experiment 27)

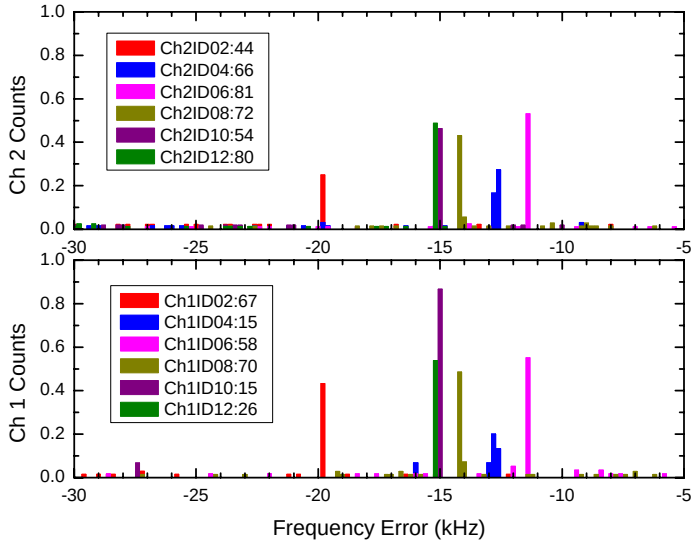

Figure C-22 Frequency Error for Orinoco Cards (Experiment 27)

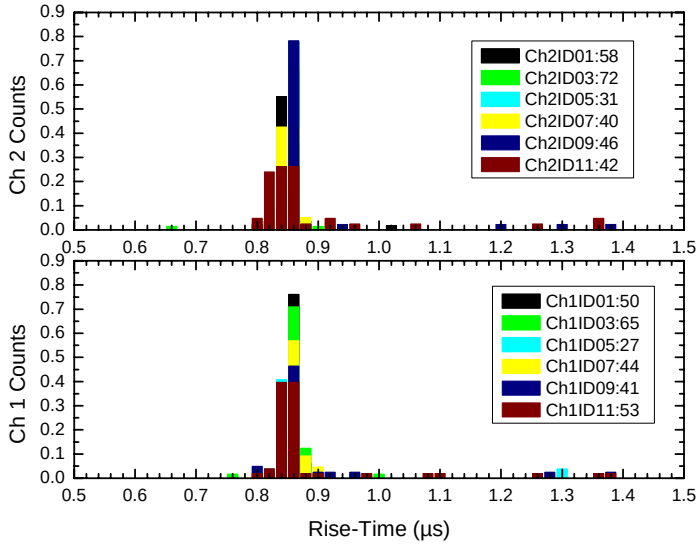

Figure C-23 Rise-Time for D-Link Cards (Experiment 27)

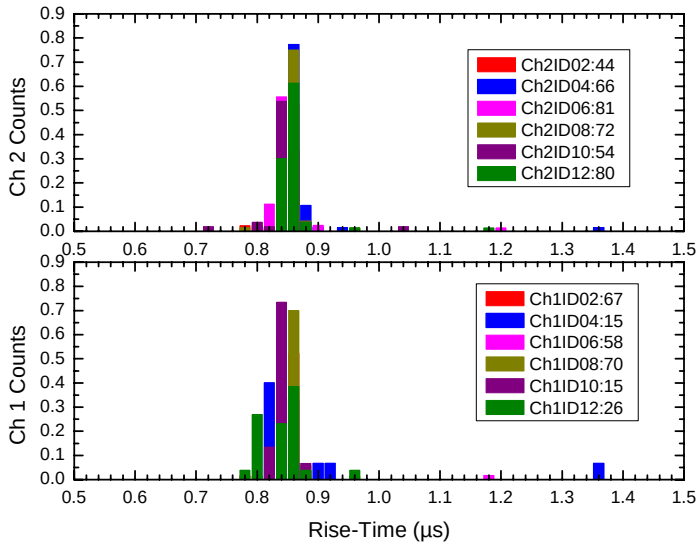

Figure C-24 Rise-Time for Orinoco Cards (Experiment 27)

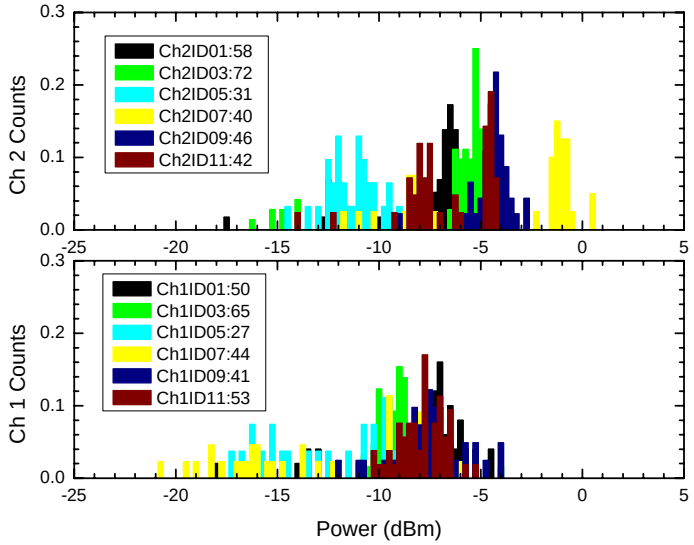

Figure C-25 Received Power for D-Link Cards (Experiment 27)

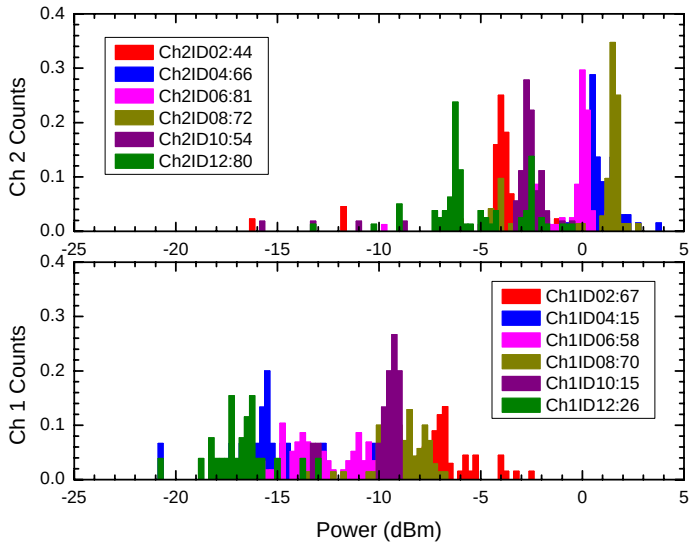

Figure C-26 Received Power for Orinoco Cards (Experiment 27)

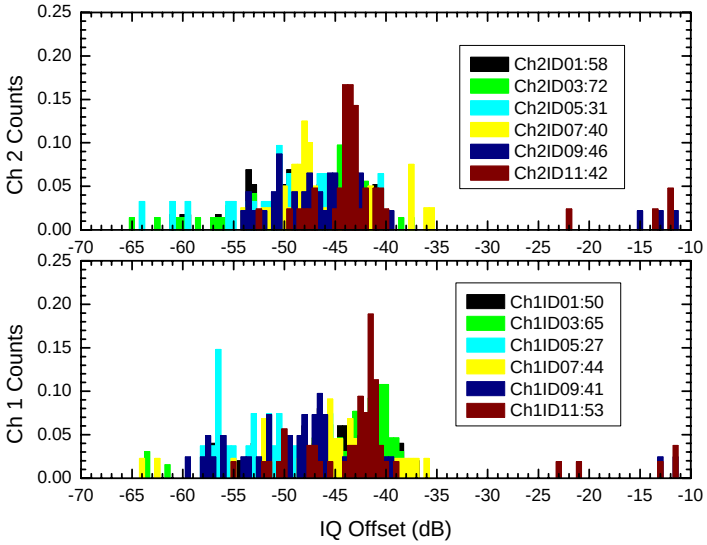

Figure C-27 IQ Offset for D-Link Cards (Experiment 27)

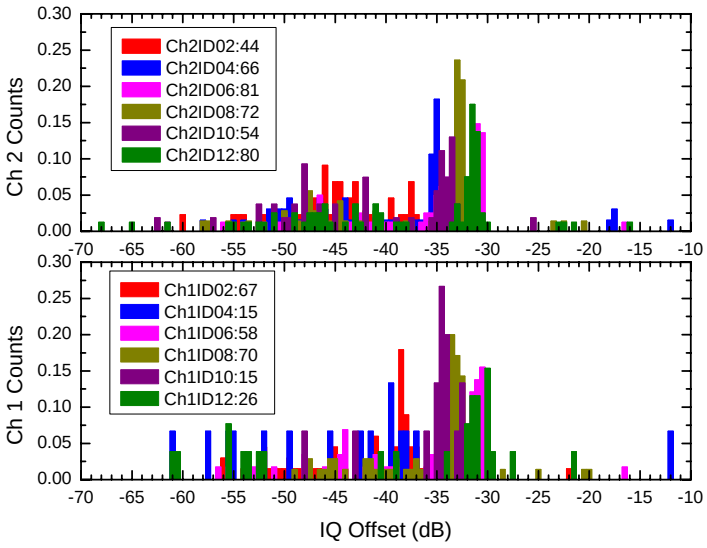

Figure C-28 IQ Offset for Orinoco Cards (Experiment 27)

C-4 EXPERIMENT 28

Experiment 28 was a two channel measurement of packet features from the second group of the twelve WiFi cards using the dual polarization horn antenna. The network configuration is that as shown in Figure C-1. Channel 1 (Ch1) is the horizontal polarization, and Channel 2 (Ch2) is the vertical polarization.

Table C-4 Experiment 28 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	000F3D6857FA	D-Link	Ch1ID01:25	Ch2ID01:24
2	226233	0020A657A451	Orinico	Ch1ID02:21	Ch2ID02:61
3	213529	000F3D68578C	D-Link	Ch1ID03:45	Ch2ID03:77
4	219421	0020A657A458	Orinico	Ch1ID04:49	Ch2ID04:49
5	221221	000F3D6857DC	D-Link	Ch1ID05:58	Ch2ID05:48
6	231009	0020A657A464	Orinico	Ch1ID06:75	Ch2ID06:82
7	234286	000F3D685802	D-Link	Ch1ID07:61	Ch2ID07:24
8	217653	0020A657A459	Orinico	Ch1ID08:86	Ch2ID08:87
9	216815	000F3D685807	D-Link	Ch1ID09:61	Ch2ID09:57
10	234285	0020A657A457	Orinico	Ch1ID10:64	Ch2ID10:60
11	234284	000F3D685800	D-Link	Ch1ID11:23	Ch2ID11:44
12	232546	0020A657A453	Orinico	Ch1ID12:84	Ch2ID12:86

D-Link = D-Link AirPlusG DWL-G630 P/N: BWLG630NA.C1 H/W Version C1 F/W Version 3.00
Orinico = Proxim Gold ORiNOCO 11b/g PC Card Model: 8470-FC PN: 67877/1

Figures C-29 through C-36 are pairs of figures showing measured feature distributions for the four features. The first figure of each pair is for the D-Link cards, while the second figure is for the Orinoco cards. There are two graphs on each figure. The upper graph (Ch 2) is from the horn vertical polarization, while the lower graph (Ch 1) is from the horn horizontal polarization.

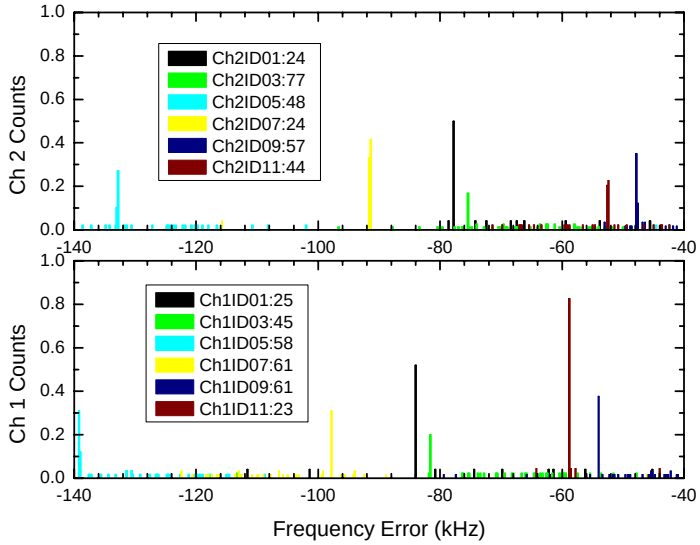

Figure C-29 Frequency Error for D-Link Cards (Experiment 28)

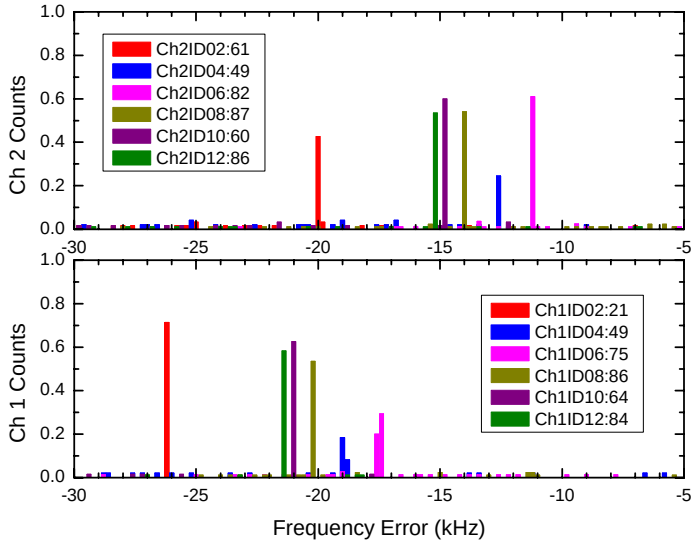

Figure C-30 Frequency Error for Orinoco Cards (Experiment 28)

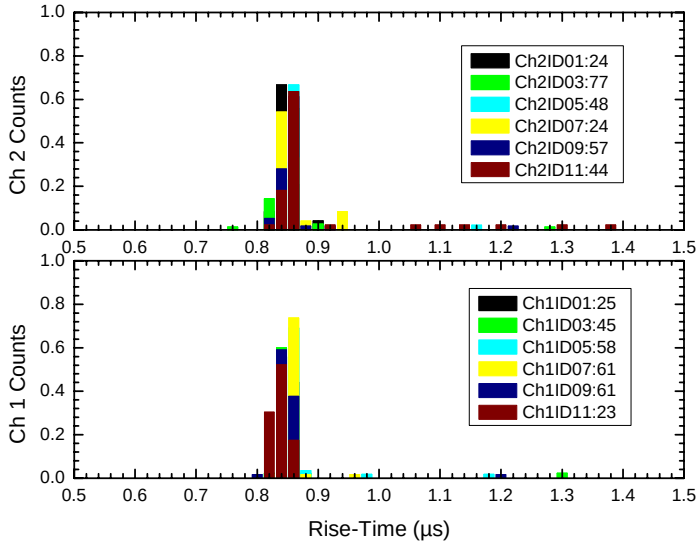

Figure C-31 Rise-Time for D-Link Cards (Experiment 28)

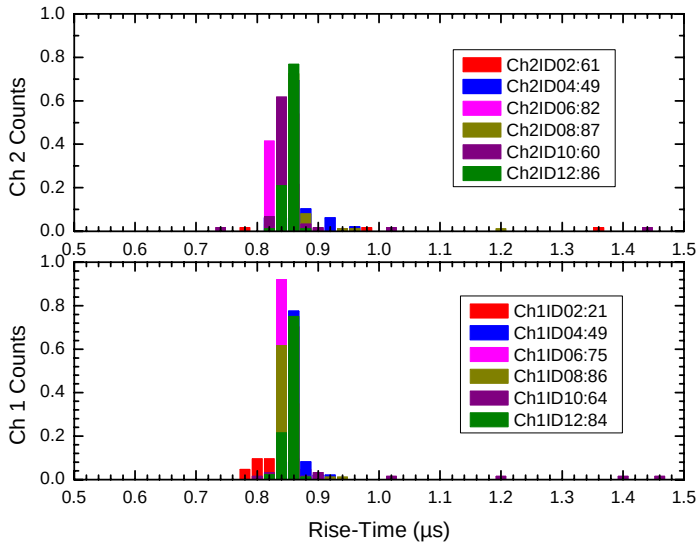

Figure C-32 Rise-Time for Orinoco Cards (Experiment 28)

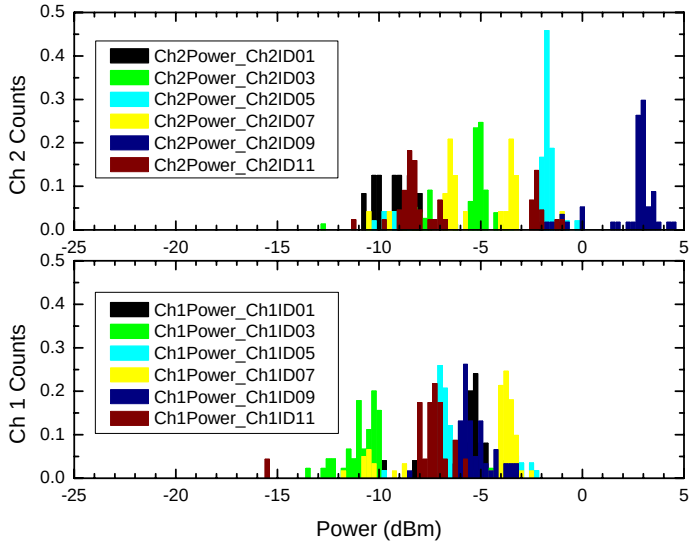

Figure C-33 Received Power for D-Link Cards (Experiment 28)

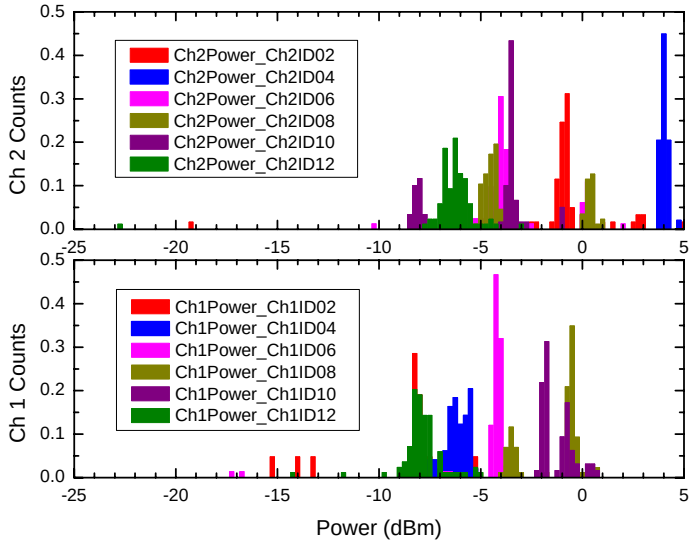

Figure C-34 Received Power for Orinoco Cards (Experiment 28)

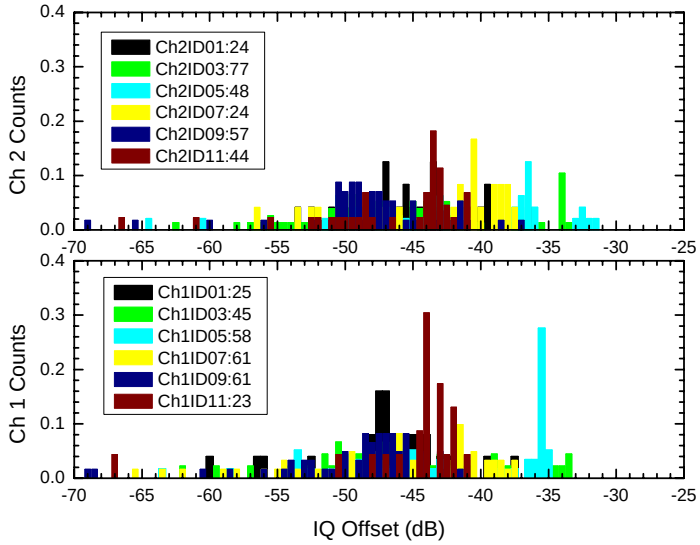

Figure C-35 IQ Offset for D-Link Cards (Experiment 28)

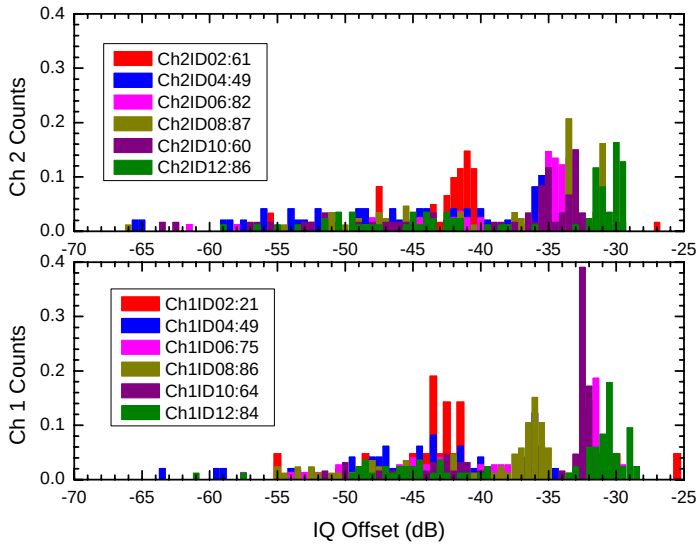

Figure C-36 IQ Offset for Orinoco Cards (Experiment 28)

C-5 EXPERIMENT 29

Experiment 29 was a two channel measurement of packet features from the second group of twelve WiFi cards. The dual polarization horn antenna was used for these measurements. The network configuration was the dispersed arrangement shown in Figure C-2. Channel 1 (Ch1) is the horizontal polarization, and Channel 2 (Ch2) is the vertical polarization.

Table C-5 Experiment 29 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	000F3D6857FA	D-Link	Ch1ID01:5	Ch2ID01:3
2	226233	0020A657A451	Orinico	Ch1ID02:26	Ch2ID02:28
3	213529	000F3D68578C	D-Link	Ch1ID03:6	Ch2ID03:8
4	219421	0020A657A458	Orinico	Ch1ID04:44	Ch2ID04:32
5	221221	000F3D6857DC	D-Link	Ch1ID05:3	Ch2ID05:22
6	231009	0020A657A464	Orinico	Ch1ID06:43	Ch2ID06:37
7	234286	000F3D685802	D-Link	Ch1ID07:34	Ch2ID07:51
8	217653	0020A657A459	Orinico	Ch1ID08:35	Ch2ID08:8
9	216815	000F3D685807	D-Link	Ch1ID09:49	Ch2ID09:40
10	234285	0020A657A457	Orinico	Ch1ID10:67	Ch2ID10:81
11	234284	000F3D685800	D-Link	Ch1ID11:25	Ch2ID11:26
12	232546	0020A657A453	Orinico	Ch1ID12:58	Ch2ID12:59

D-Link = D-Link AirPlusG DWL-G630 P/N: BWLG630NA.C1 H/W Version C1 F/W Version 3.00
Orinico = Proxim Gold ORiNOCO 11b/g PC Card Model: 8470-FC PN: 67877/1

Figures C-37 through C-44 are pairs of figures showing measured feature distributions for the four features. The first figure of each pair is for the D-Link cards, while the second figure is for the Orinoco cards. There are two graphs on each figure. The upper graph (Ch 2) is from the horn vertical polarization, while the lower graph (Ch 1) is from the horn horizontal polarization.

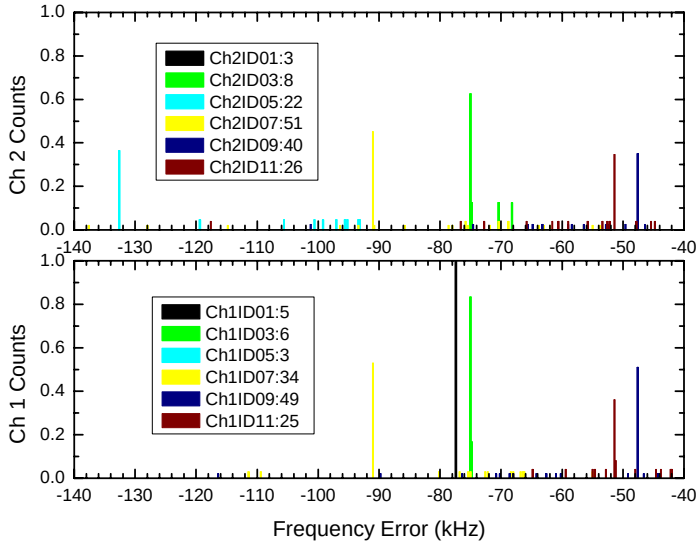

Figure C-37 Frequency Error for D-Link Cards (Experiment 29)

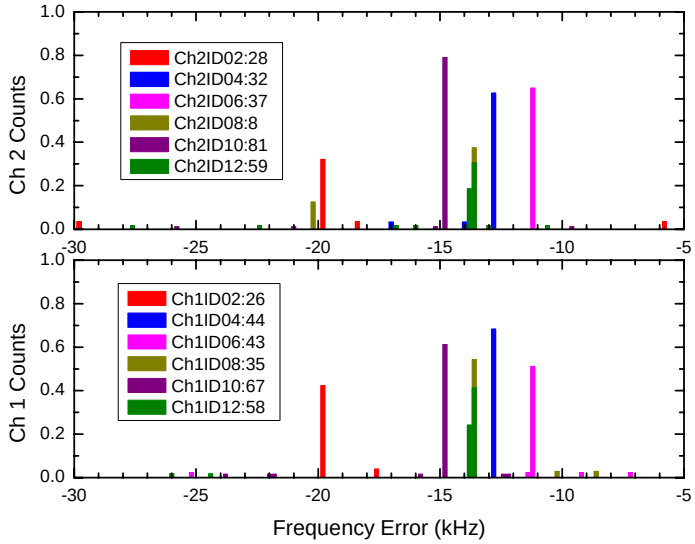

Figure C-38 Frequency Error for Orinoco Cards (Experiment 29)

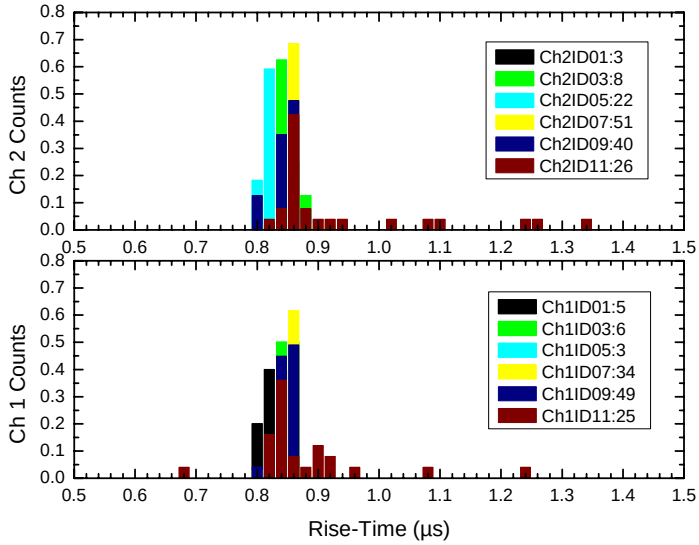

Figure C-39 Rise-Time for D-Link Cards (Experiment 29)

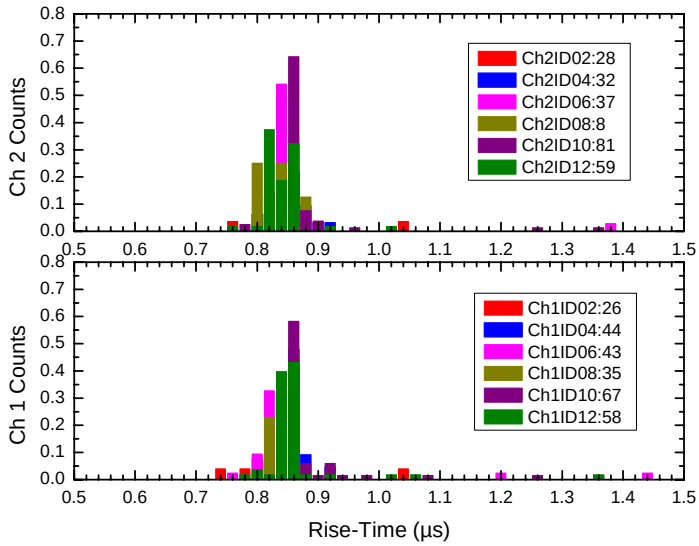

Figure C-40 Rise-Time for Orinoco Cards (Experiment 29)

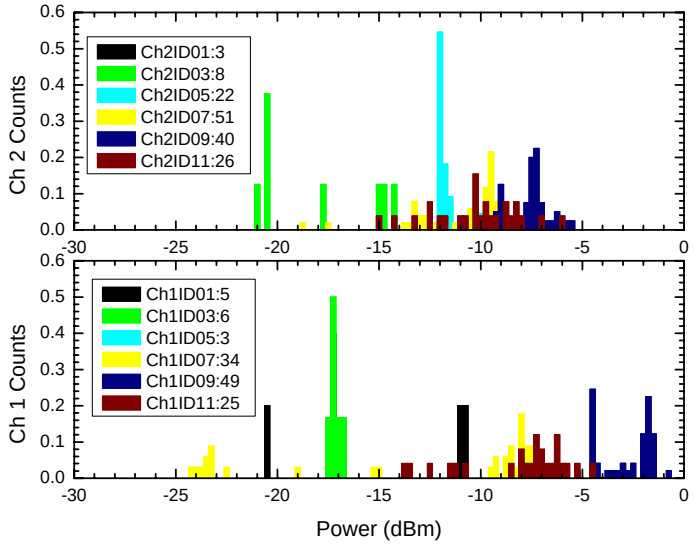

Figure C-41 Received Power for D-Link Cards (Experiment 29)

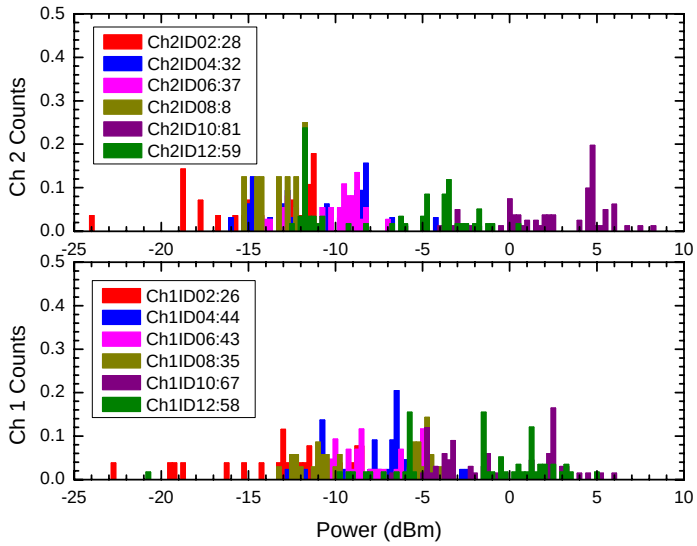

Figure C-42 Received Power for Orinoco Cards (Experiment 29)

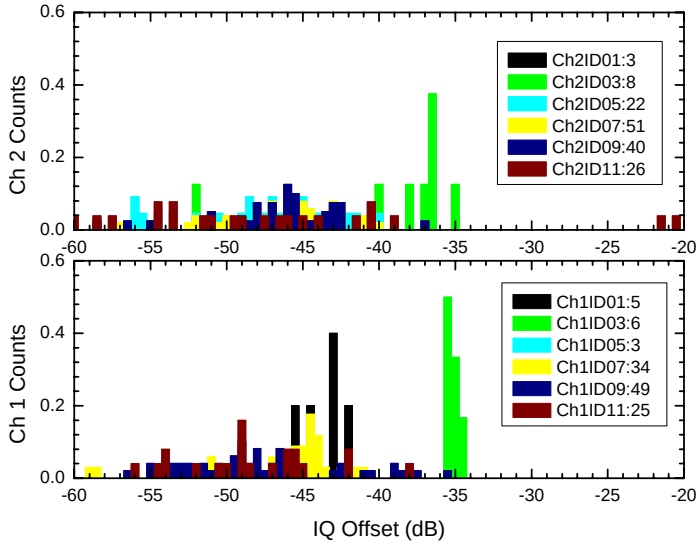

Figure C-43 IQ Offset for D-Link Cards (Experiment 29)

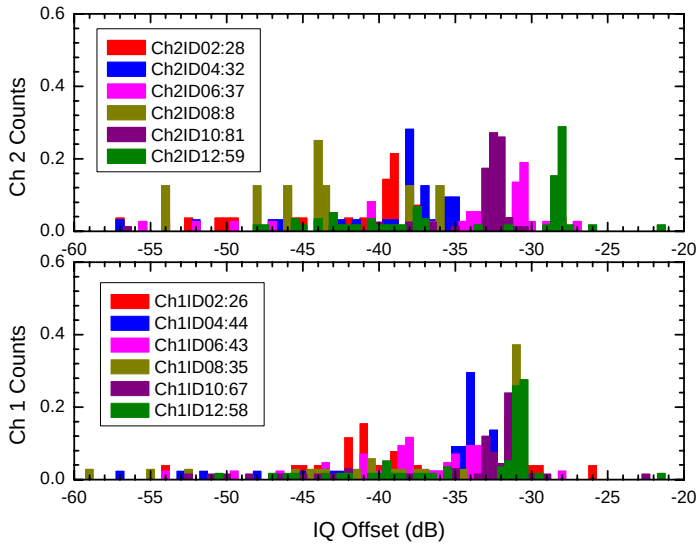

Figure C-44 IQ Offset for Orinoco Cards (Experiment 29)

C-6 EXPERIMENT 30

Experiment 30 was a two channel measurement of packet features from the second group of twelve WiFi cards, configured as shown in Figure C-2. Channel 1 was connected to a vertically polarized omni-directional antenna, while Channel 2 was attached to a horizontally polarized parabolic antenna.

Table C-6 Experiment 30 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	000F3D6857FA	D-Link	Ch1ID01:56	Ch2ID01:12
2	226233	0020A657A451	Orinico	Ch1ID02:103	Ch2ID02:70
3	213529	000F3D68578C	D-Link	Ch1ID03:75	Ch2ID03:99
4	219421	0020A657A458	Orinico	Ch1ID04:132	Ch2ID04:88
5	221221	000F3D6857DC	D-Link	Ch1ID05:50	Ch2ID05:39
6	231009	0020A657A464	Orinico	Ch1ID06:156	Ch2ID06:154
7	234286	000F3D685802	D-Link	Ch1ID07:47	Ch2ID07:32
8	217653	0020A657A459	Orinico	Ch1ID08:60	Ch2ID08:58
9	216815	000F3D685807	D-Link	Ch1ID09:36	Ch2ID09:55
10	234285	0020A657A457	Orinico	Ch1ID10:58	Ch2ID10:49
11	234284	000F3D685800	D-Link	Ch1ID11:36	Ch2ID11:17
12	232546	0020A657A453	Orinico	Ch1ID12:49	Ch2ID12:117

D-Link = D-Link AirPlusG DWL-G630 P/N: BWLG630NA.C1 H/W Version C1 F/W Version 3.00
Orinico = Proxim Gold ORiNOCO 11b/g PC Card Model: 8470-FC PN: 67877/1

Figures C-45 through C-52 are pairs of figures showing measured feature distributions for four features. The first figure of each pair is for the D-Link cards, while the second figure is for the Orinoco cards. There are two graphs on each figure. The upper graph (Ch2) is from the horizontally polarized parabolic antenna, while the lower graph (Ch 1) is from the omni-directional antenna.

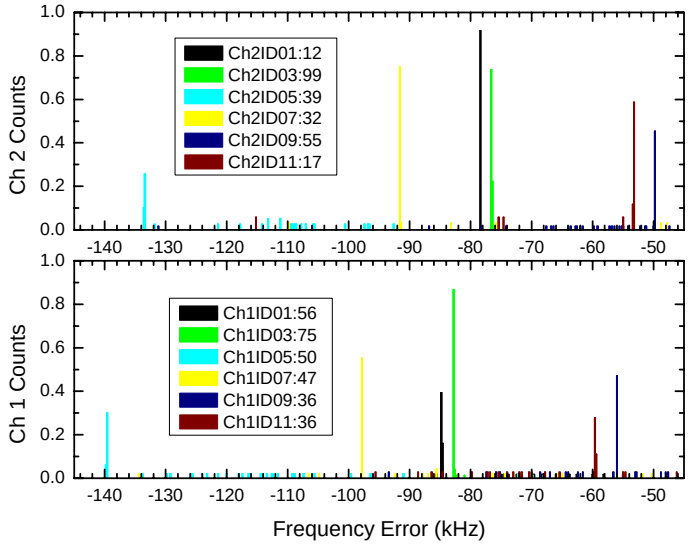

Figure C-45 Frequency Error for D-Link Cards (Experiment 30)

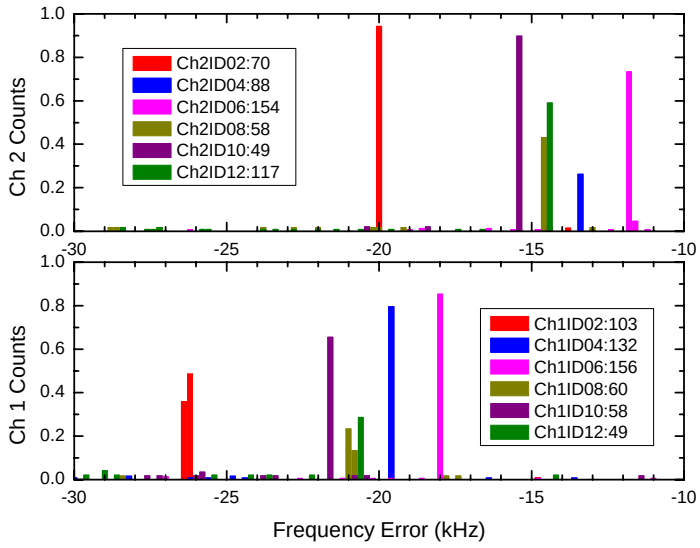

Figure C-46 Frequency Error for Orinoco Cards (Experiment 30)

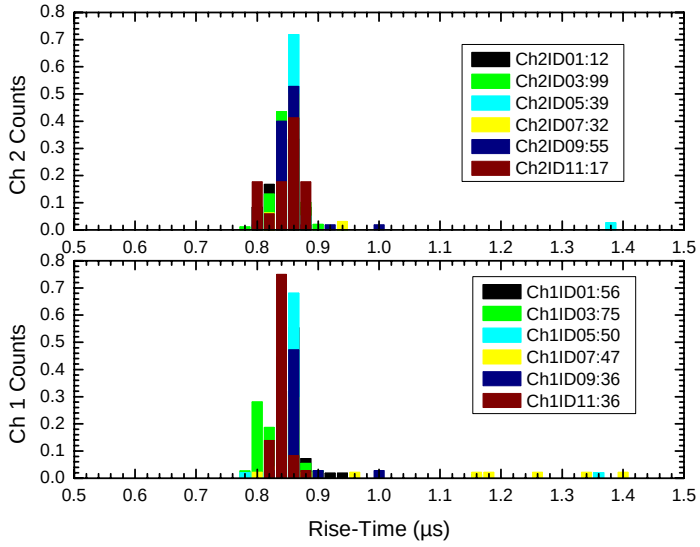

Figure C-47 Rise-Time for D-Link Cards (Experiment 30)

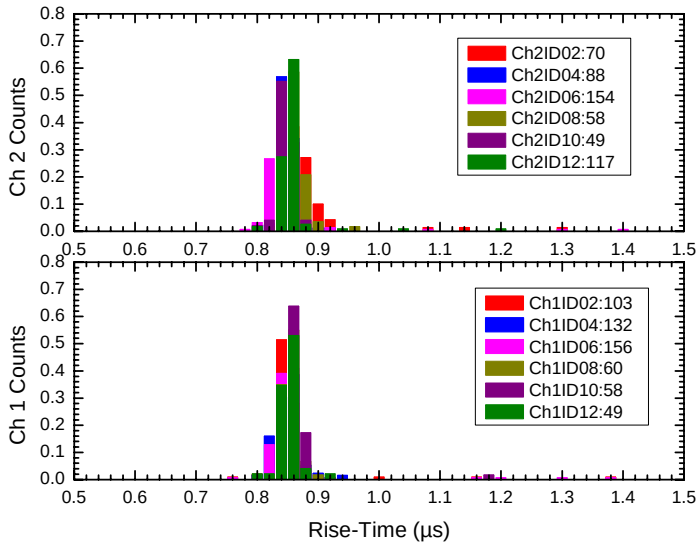

Figure C-48 Rise-Time for Orinoco Cards (Experiment 30)

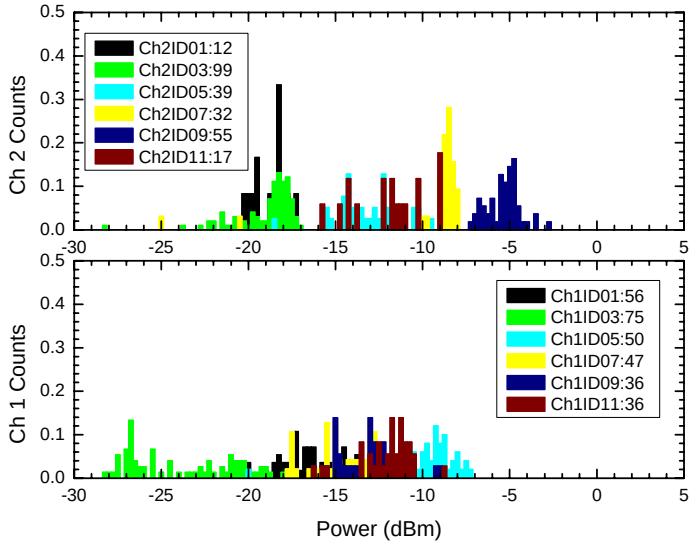

Figure C-49 Received Power for D-Link Cards (Experiment 30)

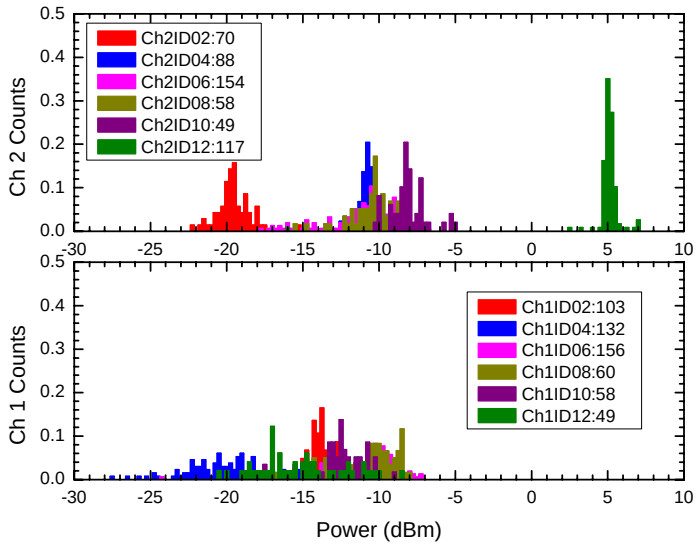

Figure C-50 Received Power for Orinoco Cards (Experiment 30)

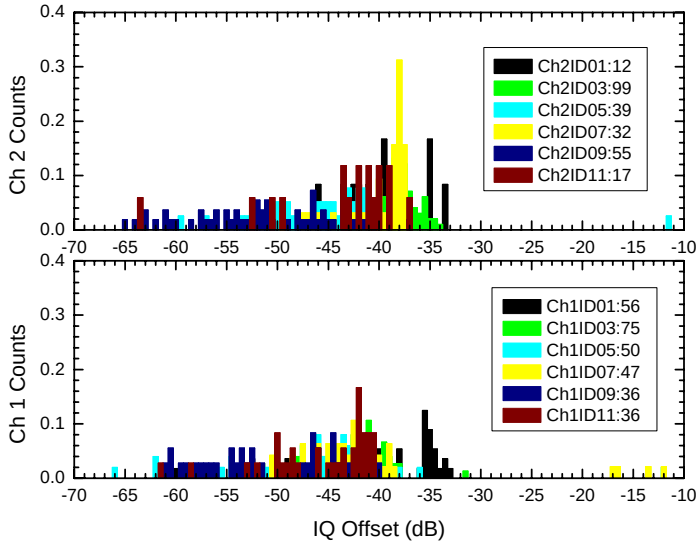

Figure C-51 IQ Offset for D-Link Cards (Experiment 30)

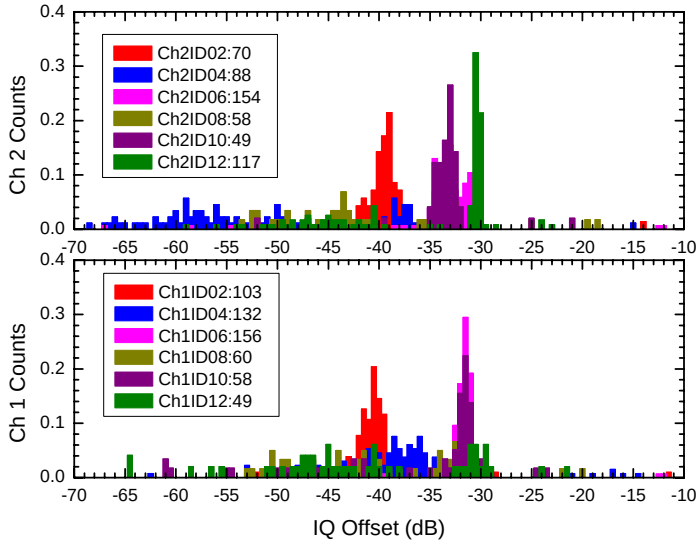

Figure C-52 IQ Offset for Orinoco Cards (Experiment 30)

C-7 EXPERIMENT 31

Experiment 31 was a two channel measurement of packet features from the first group of twelve WiFi cards, configured as shown in Figure C-2. Channel 1 was connected to a vertically polarized omni-directional antenna, while Channel 2 was attached to a horizontally polarized parabolic antenna.

Table C-7 Experiment 31 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	001150081163	Belkin	Ch1ID01:14	Ch2ID01:32
2	226233	000C41405D12	Linksys'	Ch1ID02:2	Ch2ID02:121
3	213529	0011500CD68D	Belkin	Ch1ID03:19	Ch2ID03:33
4	219421	000F66D098C1	Linksys	Ch1ID04:1	Ch2ID04:4
5	221221	001150080C33	Belkin	Ch1ID05:58	Ch2ID05:1
6	231009	000F66D098C8	Linksys	Ch1ID06:27	Ch2ID06:4
7	234286	001150080C2B	Belkin	Ch1ID07:47	Ch2ID07:69
8	217653	000F66D098C7	Linksys	Ch1ID08:32	Ch2ID08:69
9	216815	001150081166	Belkin	Ch1ID09:39	Ch2ID09:44
10	234285	000F66D098C4	Linksys	Ch1ID10:29	Ch2ID10:18
11	234284	0011500815F5	Belkin	Ch1ID11:157	Ch2ID11:123
12	232546	000F66D098C6	Linksys	Ch1ID12:27	Ch2ID12:35

Belkin = Belkin Wireless Notebook Network Card IEEE 802.11g/54 Mbps Model F5D7010 Version 1315
Linksys' = Linksys Wireless-B Notebook Adapter Model No: WPC11 Version 4
Linksys = Linksys Wireless-G Notebook Adapter Model No: WPC54G Version 2

Figures C-53 through C-60 are pairs of figures showing measured feature distributions for four features. The first figure of each pair is for the Belkin cards, while the second figure is for the Linksys cards. There are two graphs on each figure. The upper graph (Ch2) is from the horizontally polarized parabolic antenna, while the lower graph (Ch 1) is from the omni-directional antenna.

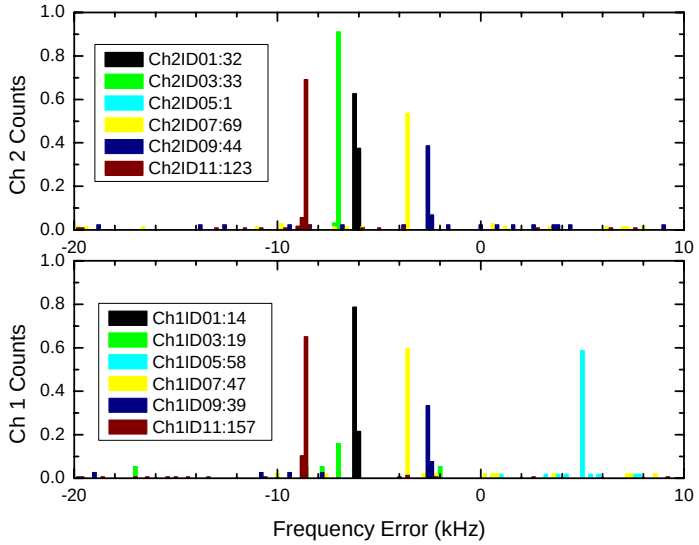

Figure C-53 Frequency Error for Belkin Cards (Experiment 31)

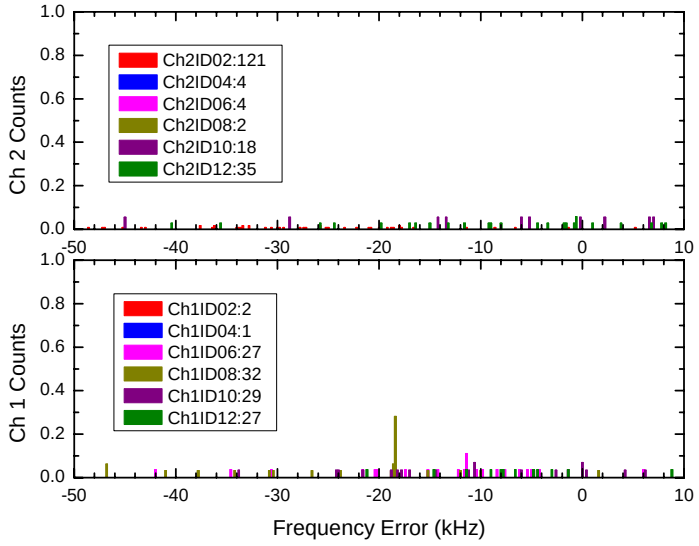

Figure C-54 Frequency Error for Linksys Cards (Experiment 31)

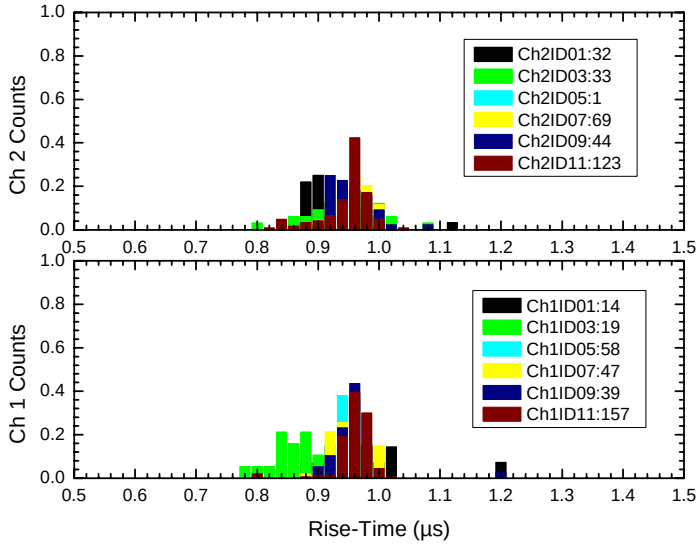

Figure C-55 Rise-Time for Belkin Cards (Experiment 31)

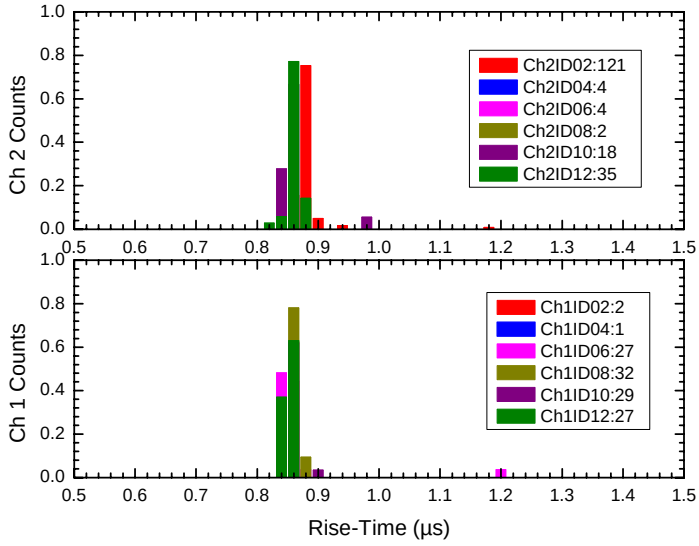

Figure C-56 Rise-Time for Linksys Cards (Experiment 31)

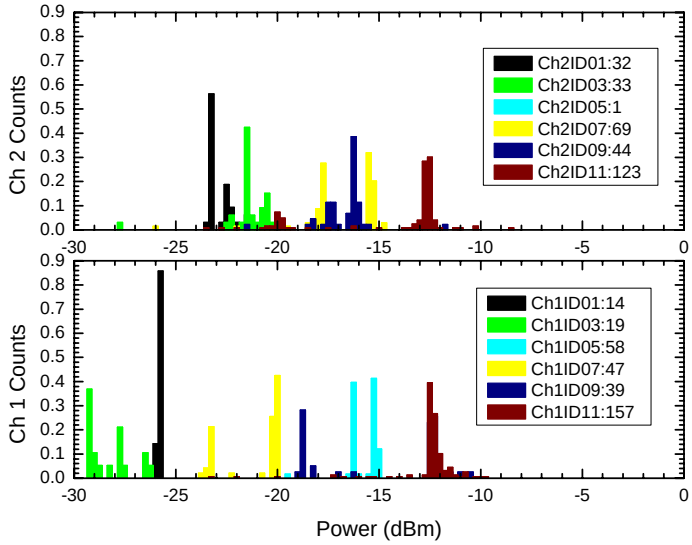

Figure C-57 Received Power for Belkin Cards (Experiment 31)

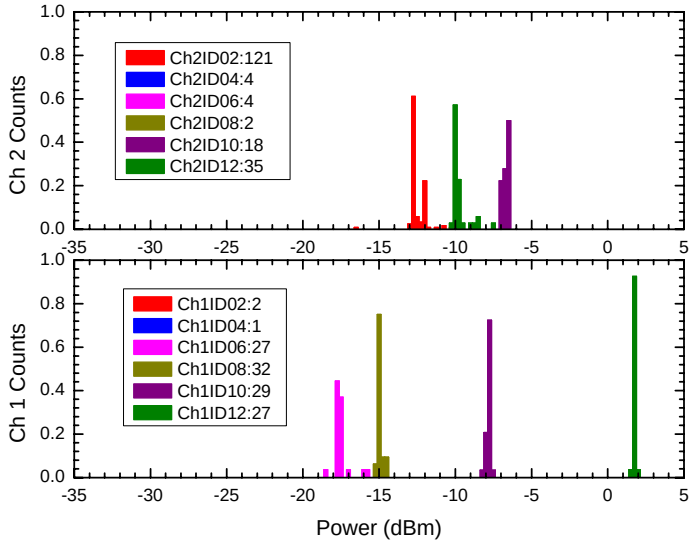

Figure C-58 Received Power for Linksys Cards (Experiment 31)

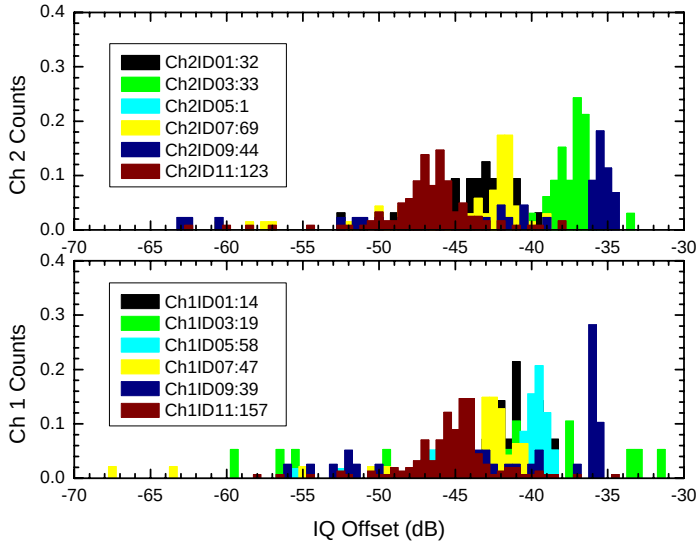

Figure C-59 IQ Offset for Belkin Cards (Experiment 31)

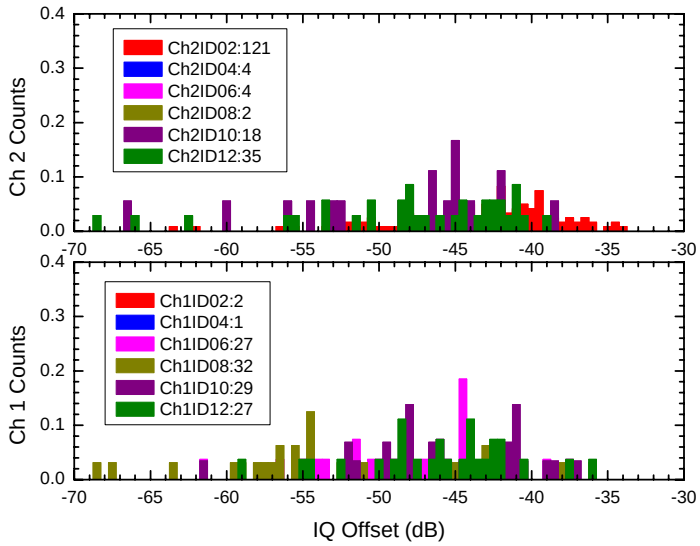

Figure C-60 IQ Offset for Linksys Cards (Experiment 31)

C-8 EXPERIMENT 32

Experiment 32 was a two channel measurement using a dual polarization horn antenna. Table C-8 lists the host laptop property number, card MAC address, card type, and the number of packets detected on each channel. Channel 1 (Ch1) is the horizontal polarization, and Channel 2 (Ch2) is the vertical polarization. The network configuration is the dispersed arrangement as shown in Figure C-2. The first group of twelve WiFi cards was used for these measurements.

Table C-8 Experiment 32 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	001150081163	Belkin	Ch1ID01:42	Ch2ID01:1
2	226233	000C41405D12	Linksys'	Ch1ID02:85	Ch2ID02:15
3	213529	0011500CD68D	Belkin	Ch1ID03:261	Ch2ID03:332
4	219421	000F66D098C1	Linksys	Ch1ID04:22	Ch2ID04:25
5	221221	001150080C33	Belkin	Ch1ID05:150	Ch2ID05:167
6	231009	000F66D098C8	Linksys	Ch1ID06:14	Ch2ID06:30
7	234286	001150080C2B	Belkin	Ch1ID07:262	Ch2ID07:39
8	217653	000F66D098C7	Linksys	Ch1ID08:14	Ch2ID08:29
9	216815	001150081166	Belkin	Ch1ID09:147	Ch2ID09:187
10	234285	000F66D098C4	Linksys	Ch1ID10:27	Ch2ID10:30
11	234284	0011500815F5	Belkin	Ch1ID11:113	Ch2ID11:174
12	232546	000F66D098C6	Linksys	Ch1ID12:1	Ch2ID12:30

Belkin = Belkin Wireless Notebook Network Card IEEE 802.11g/54 Mbps Model F5D7010 Version 1315
Linksys' = Linksys Wireless-B Notebook Adapter Model No: WPC11 Version 4
Linksys = Linksys Wireless-G Notebook Adapter Model No: WPC54G Version 2

Figures C-61 through C-68 are pairs of figures showing measured feature distributions for four features. The first figure in each pair is for the Belkin cards, while the second figure is for the Linksys cards. There are two graphs on each figure. The upper graph (Ch 2) is from the horn vertical polarization, while the lower graph (Ch 1) is from the horn horizontal polarization.

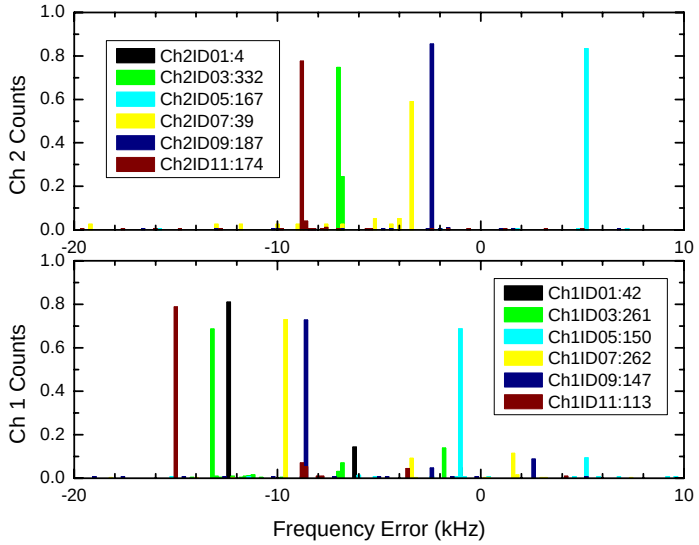

Figure C-61 Frequency Error for Belkin Cards (Experiment 32)

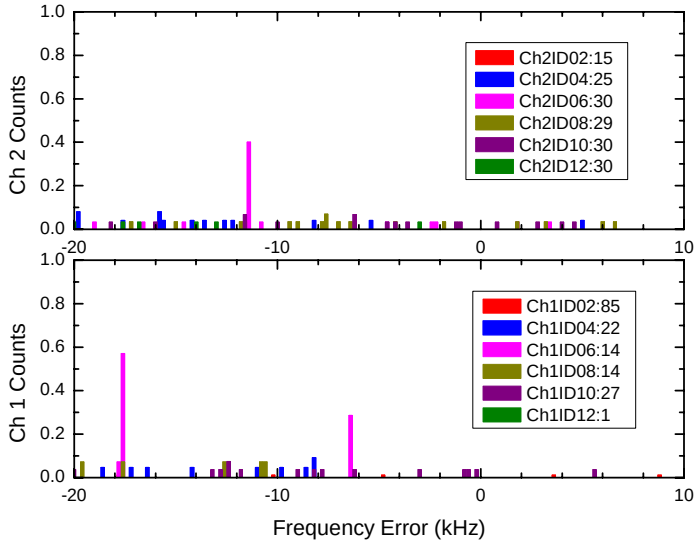

Figure C-62 Frequency Error for Linksys Cards (Experiment 32)

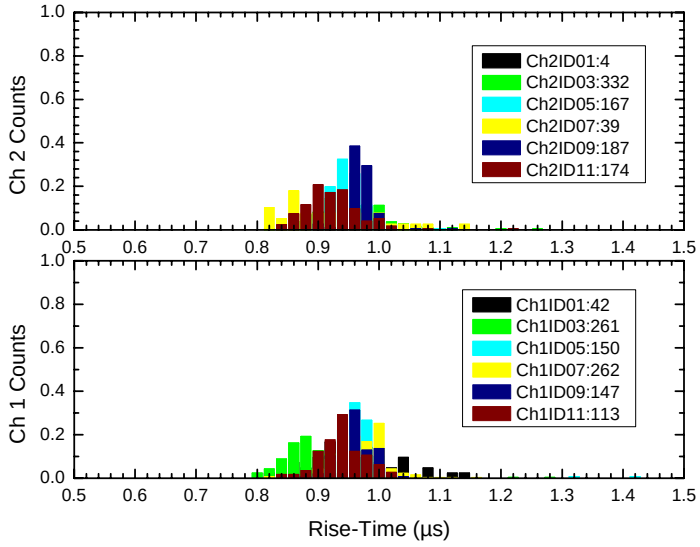

Figure C-63 Rise-Time for Belkin Cards (Experiment 32)

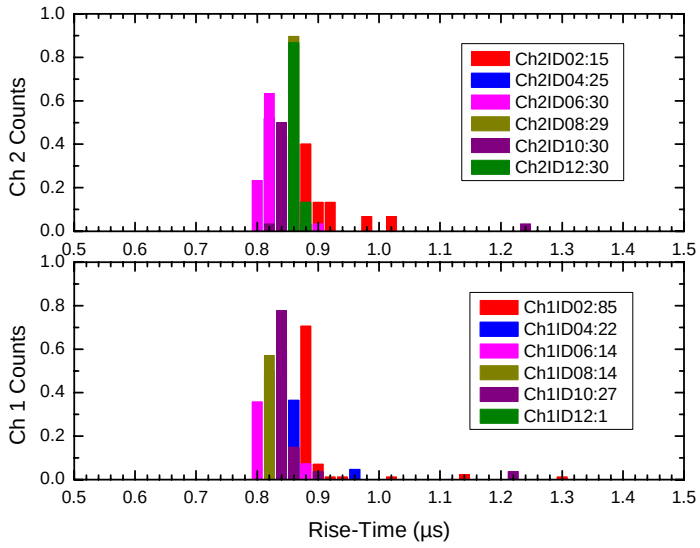

Figure C-64 Rise-Time for Linksys Cards (Experiment 32)

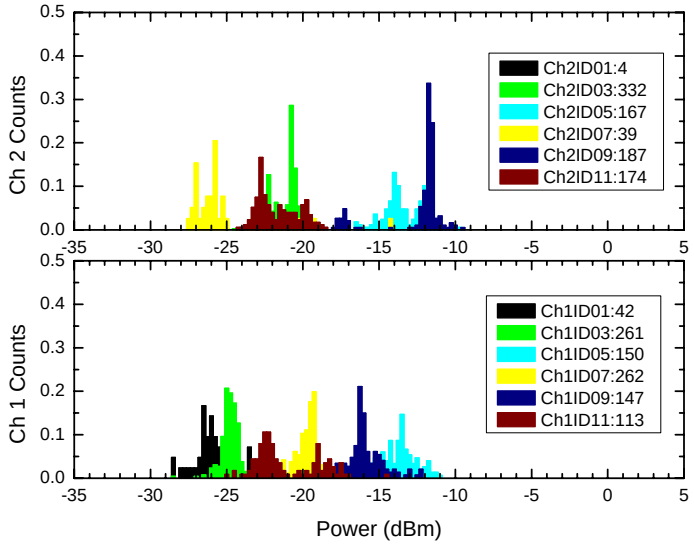

Figure C-65 Received Power for Belkin Cards (Experiment 32)

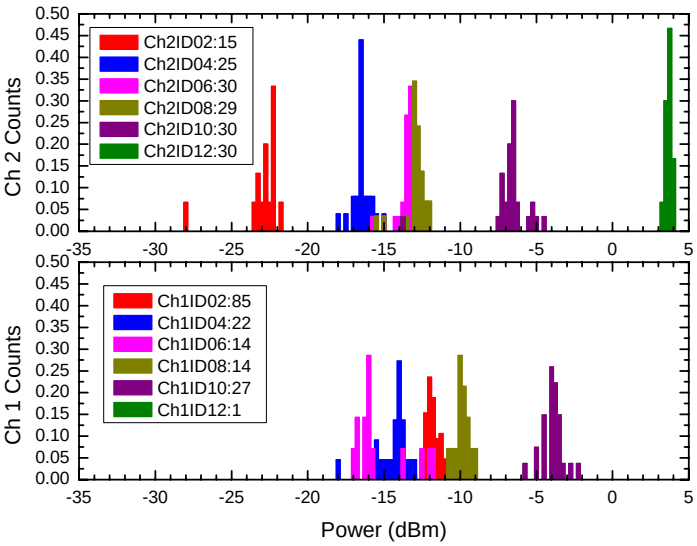

Figure C-66 Received Power for Linksys Cards (Experiment 32)

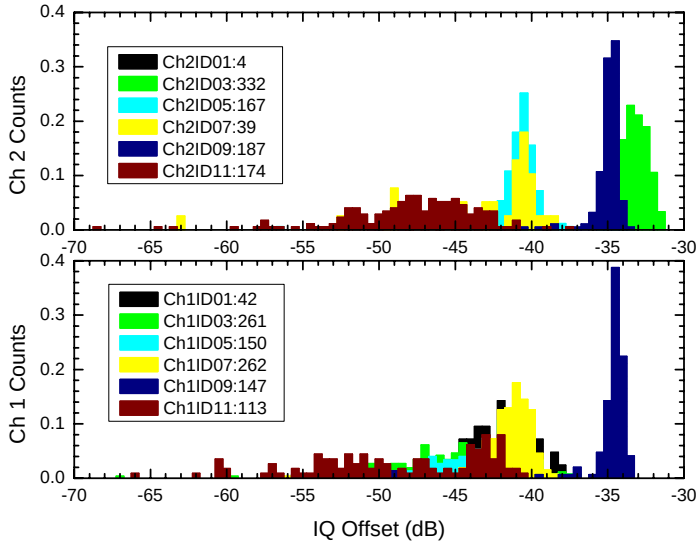

Figure C-67 IQ Offset for Belkin Cards (Experiment 32)

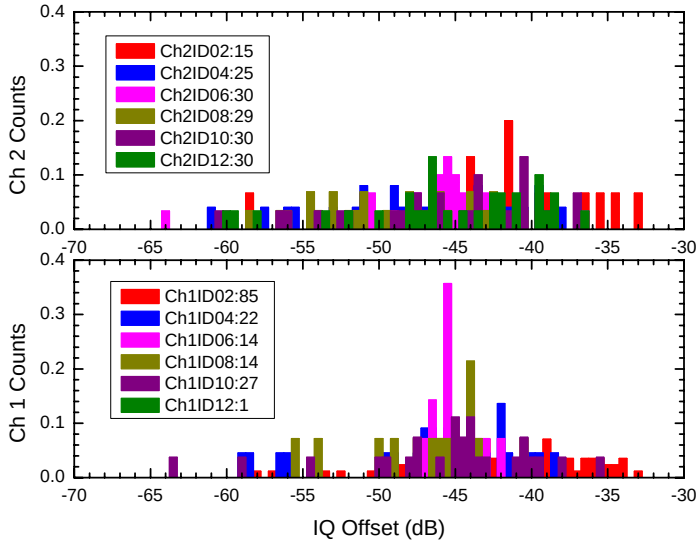

Figure C-68 IQ Offset for Linksys Cards (Experiment 32)

C-9 EXPERIMENT 33

Experiment 33 was a dual monitor experiment using the second group of twelve WiFi cards. The first monitor employed a pair of orthogonal horizontally polarized horn antennas for reception (Channels 1 and 2), while the second monitor used a vertically polarized omni-directional antenna for reception (Channel 3). The network configuration was the dispersed arrangement as shown in Figure C-3. Table C-9 lists the host laptop property number, card MAC address, card type, and the number of packets detected on each channel.

Table C-9 Experiment 33 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets	Ch 3 Packets
1	220195	000F3D6857FA	D-Link	Ch1ID01:58	Ch2ID01:70	Ch3ID01:102
2	226233	0020A657A451	Orinico	Ch1ID02:55	Ch2ID02:56	Ch3ID02:81
3	213529	000F3D68578C	D-Link	Ch1ID03:63	Ch2ID03:59	Ch3ID03:82
4	219421	0020A657A458	Orinico	Ch1ID04:48	Ch2ID04:26	Ch3ID04:129
5	221221	000F3D6857DC	D-Link	Ch1ID05:46	Ch2ID05:67	Ch3ID05:93
6	231009	0020A657A464	Orinico	Ch1ID06:88	Ch2ID06:145	Ch3ID06:154
7	234286	000F3D685802	D-Link	Ch1ID07:15	Ch2ID07:50	Ch3ID07:74
8	217653	0020A657A459	Orinico	Ch1ID08:0	Ch2ID08:0	Ch3ID08:0
9	216815	000F3D685807	D-Link	Ch1ID09:47	Ch2ID09:11	Ch3ID09:61
10	234285	0020A657A457	Orinico	Ch1ID10:16	Ch2ID10:54	Ch3ID10:67
11	234284	000F3D685800	D-Link	Ch1ID11:46	Ch2ID11:38	Ch3ID11:87
12	232546	0020A657A453	Orinico	Ch1ID12:24	Ch2ID12:18	Ch3ID12:107

D-Link = D-Link AirPlusG DWL-G630 P/N: BWLG630NA.C1 H/W Version C1 F/W Version 3.00
Orinico = Proxim Gold ORiNOCO 11b/g PC Card Model: 8470-FC PN: 67877/1

Figures C-69 through C-76 are pairs of figures showing measured feature distributions for four features. The first figure of each pair is for the D-Link cards, while the second figure is for the Orinoco cards. There are three graphs on each figure. The upper graph (Ch3) is from the vertically polarized omni-directional antenna, the middle graph (Ch2) is from one of horizontally polarized horn antennas, and the lower graph (Ch 1) is from the other orthogonal horizontally polarized horn antenna.

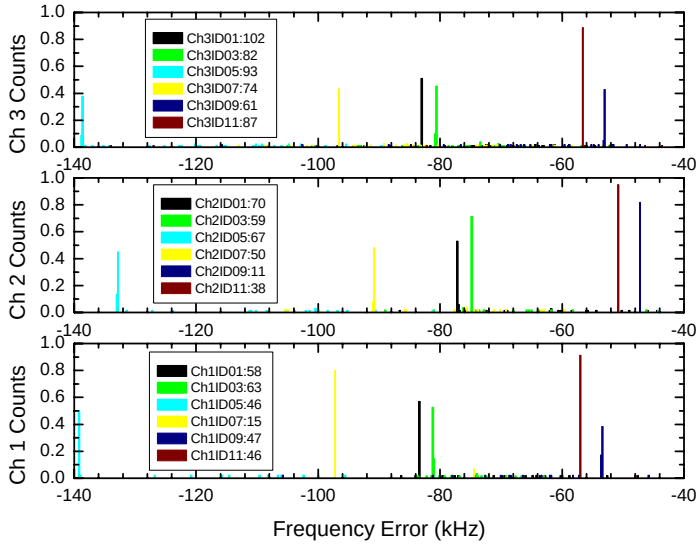

Figure C-69 Frequency Error for D-Link Cards (Experiment 33)

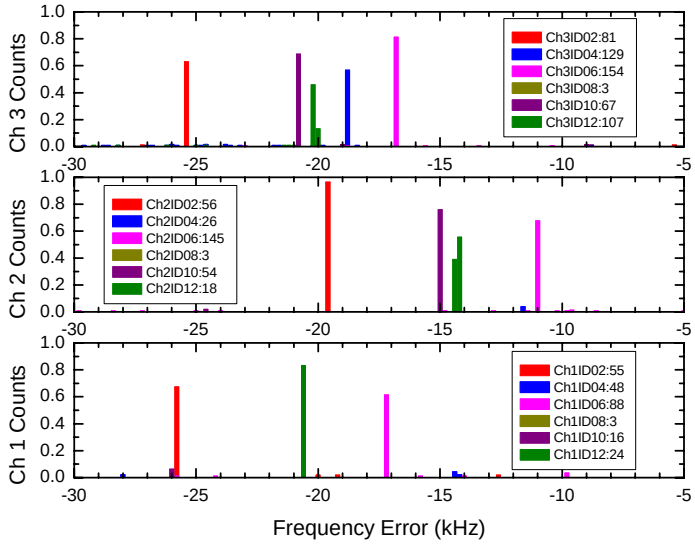

Figure C-70 Frequency Error for Orinoco Cards (Experiment 33)

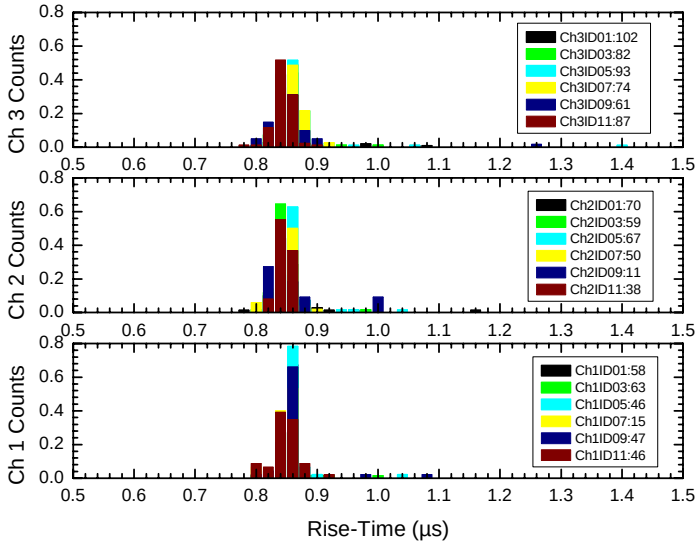

Figure C-71 Rise-Time for D-Link Cards (Experiment 33)

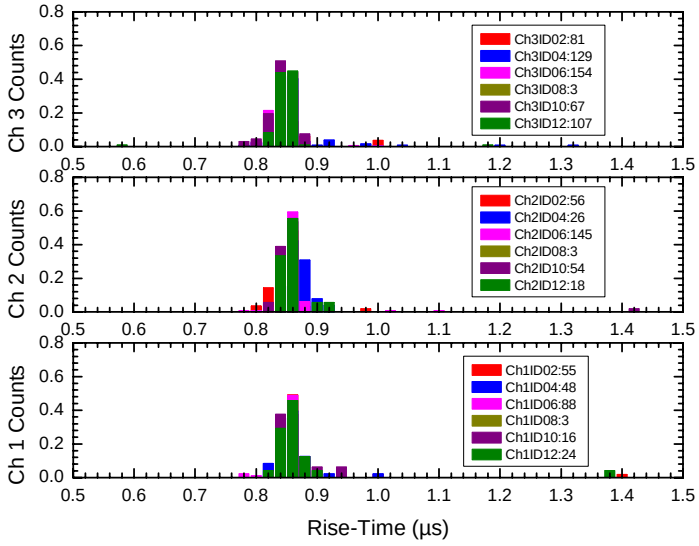

Figure C-72 Rise-Time for Orinoco Cards (Experiment 33)

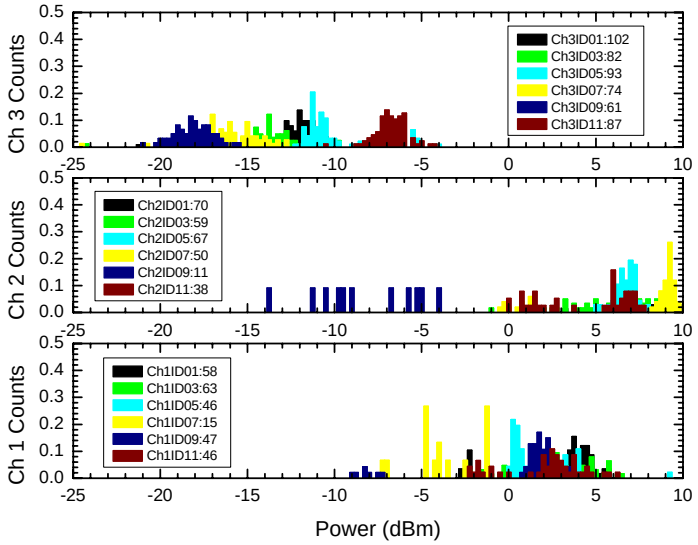

Figure C-73 Received Power for D-Link Cards (Experiment 33)

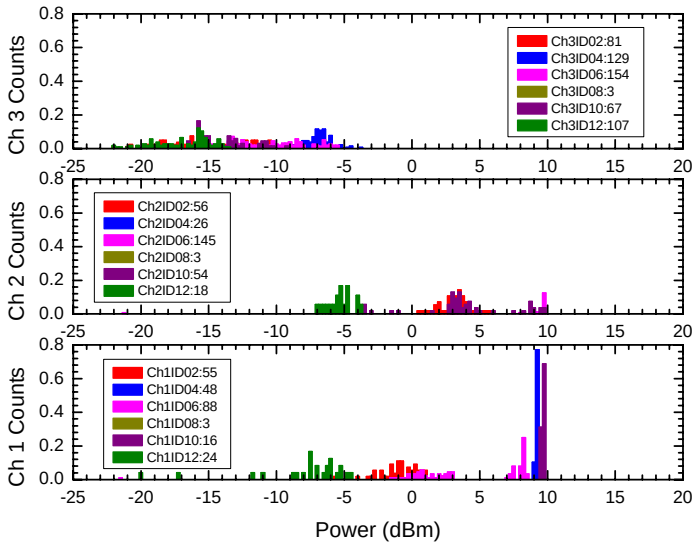

Figure C-74 Received Power for Orinoco Cards (Experiment 33)

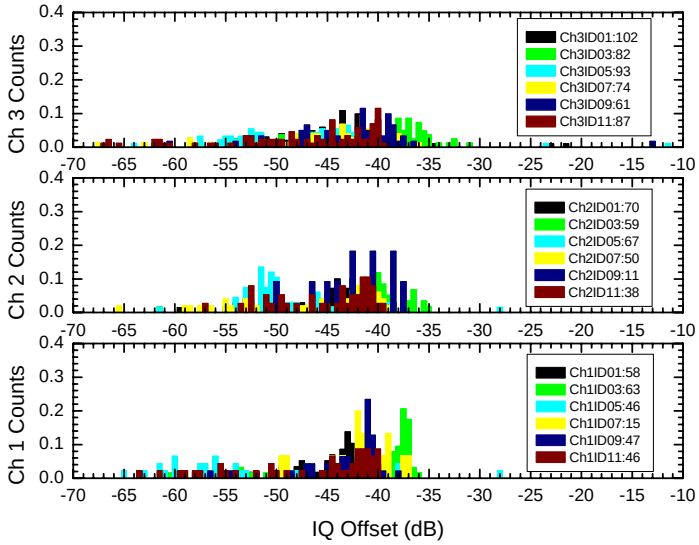

Figure C-75 IQ Offset for D-Link Cards (Experiment 33)

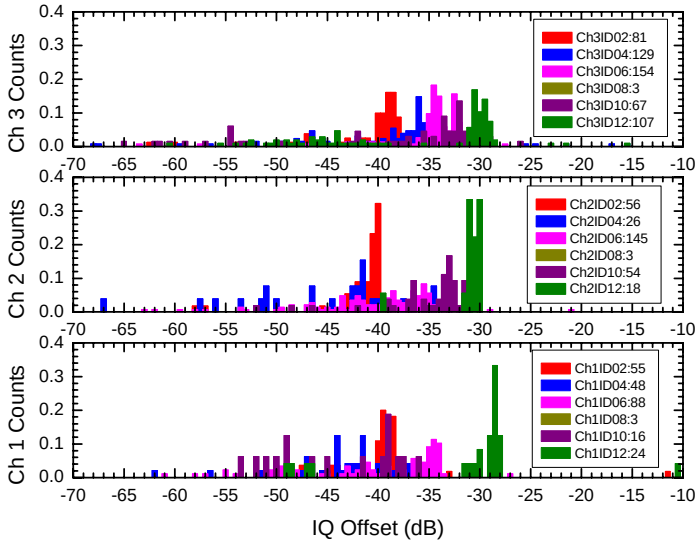

Figure C-76 IQ Offset for Orinoco Cards (Experiment 33)

C-10 EXPERIMENT 34

Experiment 34 was a single monitor measurement with a vertically polarized omni-directional antenna on Channel 1, and a horizontally polarized horn antenna on Channel 2. The data were collected in an "Internet Café" setting (Figure C-4) using the second group of twelve WiFi cards. Table C-10 provides a summary of the detected valid packets for each card. This experiment was a first attempt at packet capture in for an infrastructure configuration. Connectivity to local AP was good, but the monitor location was far removed from the AP, with and obstructed view of the indoor WiFi cards, and the duration of the experiment was short (<90 s). Consequently, very few valid packets were collected for many of the cards. Additional experiments of this type will need to be conducted to determine optimal strategies for packet capture in infrastructure configurations.

Table C-10 Experiment 34 Packet Captures

Laptop	Property Number	Card MAC Address	Type	Ch 1 Packets	Ch 2 Packets
1	220195	000F3D6857FA	D-Link	Ch1ID01:0	Ch2ID01:0
2	226233	0020A657A451	Orinico	Ch1ID02:0	Ch2ID02:0
3	213529	000F3D68578C	D-Link	Ch1ID03:24	Ch2ID03:0
4	219421	0020A657A458	Orinico	Ch1ID04:33	Ch2ID04:50
5	221221	000F3D6857DC	D-Link	Ch1ID05:4	Ch2ID05:0
6	231009	0020A657A464	Orinico	Ch1ID06:12	Ch2ID06:8
7	234286	000F3D685802	D-Link	Ch1ID07:2	Ch2ID07:2
8	217653	0020A657A459	Orinico	Ch1ID08:3	Ch2ID08:34
9	216815	000F3D685807	D-Link	Ch1ID09:46	Ch2ID09:2
10	234285	0020A657A457	Orinico	Ch1ID10:25	Ch2ID10:8
11	234284	000F3D685800	D-Link	Ch1ID11:5	Ch2ID11:4
12	232546	0020A657A453	Orinico	Ch1ID12:14	Ch2ID12:40

D-Link = D-Link AirPlusG DWL-G630 P/N: BWLG630NA.C1 H/W Version C1 F/W Version 3.00
Orinico = Proxim Gold ORiNOCO 11b/g PC Card Model: 8470-FC PN: 67877/1

Figures C-77 through C-84 are pairs of figures showing measured feature distributions for four features: frequency error, rise-time, received power and IQ offset. The first figure of each pair is for the D-Link cards, while the second figure is for the Orinoco cards. There are two graphs on each figure. The upper graph (Ch 2) is from the horn horizontal polarization, while the lower graph (Ch 1) is from the vertically polarized omni-directional antenna.

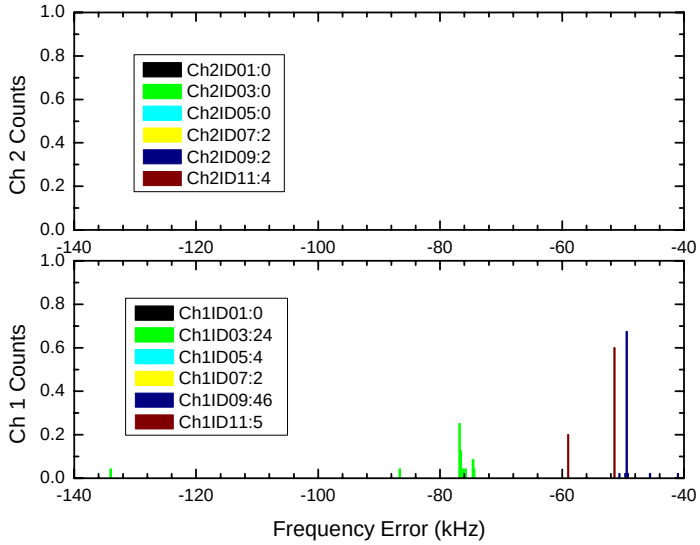

Figure C-77 Frequency Error for D-Link Cards (Experiment 34)

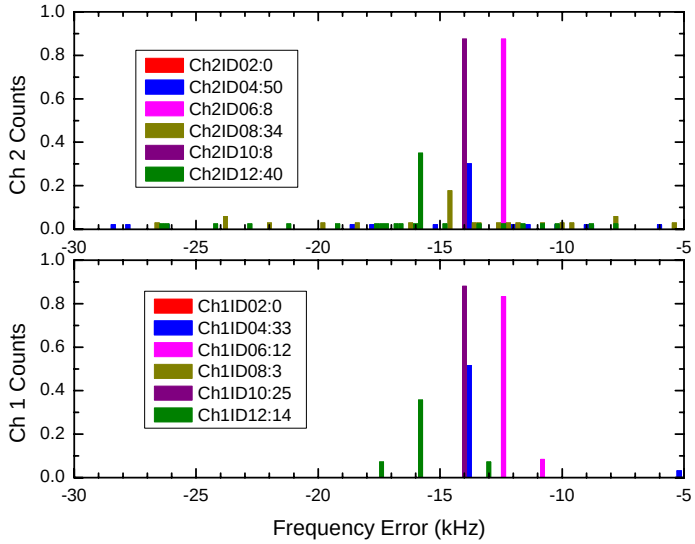

Figure C-78 Frequency Error for Orinoco Cards (Experiment 34)

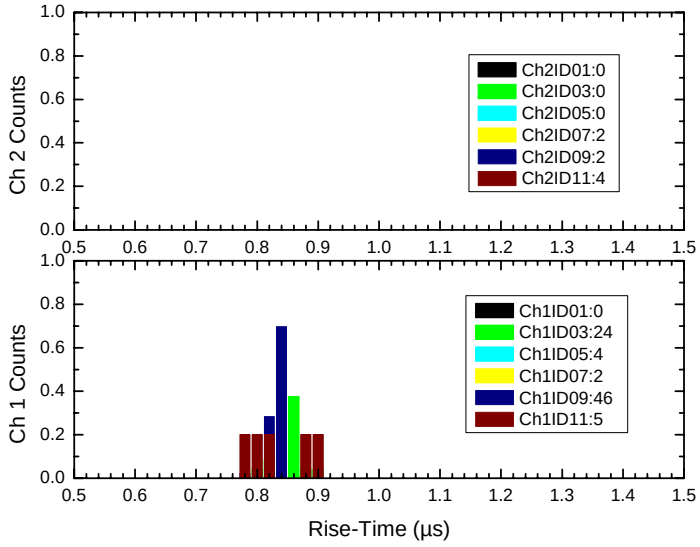

Figure C-79 Rise-Time for D-Link Cards (Experiment 34)

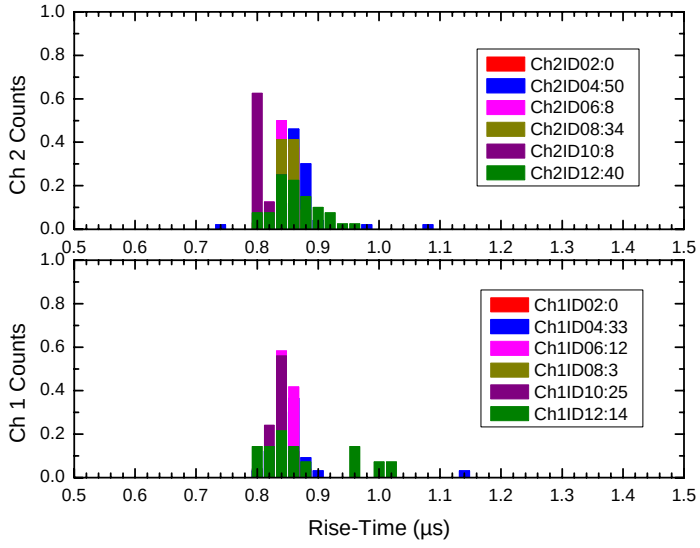

Figure C-80 Rise-Time for Orinoco Cards (Experiment 34)

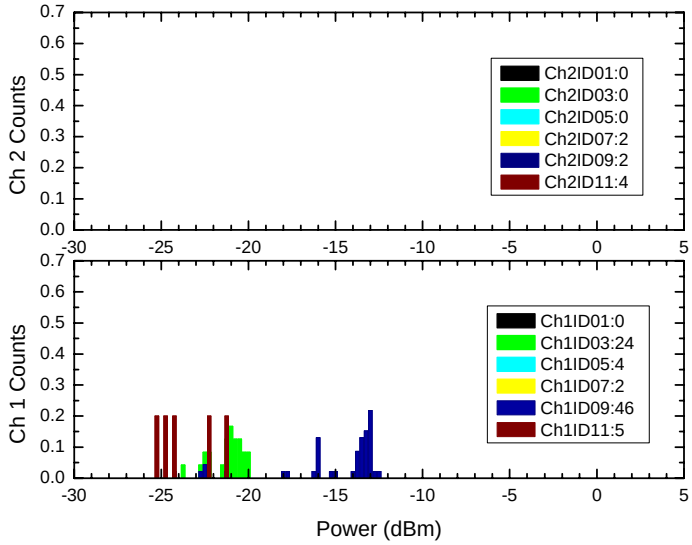

Figure C-81 Received Power for D-Link Cards (Experiment 34)

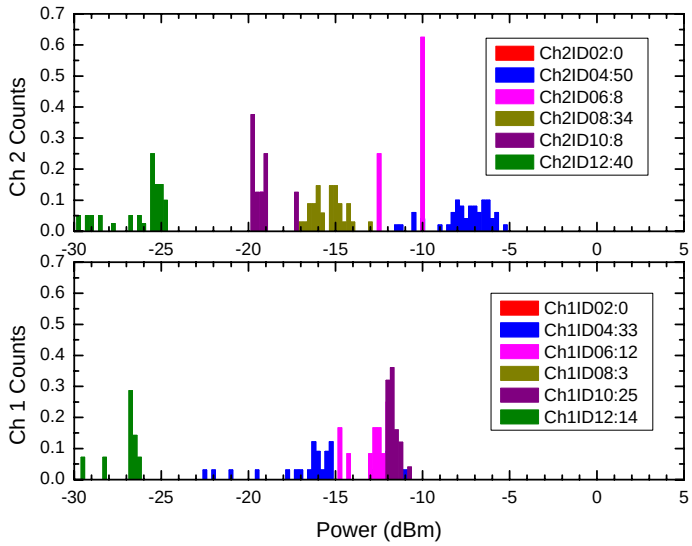

Figure C-82 Received Power for Orinoco Cards (Experiment 34)

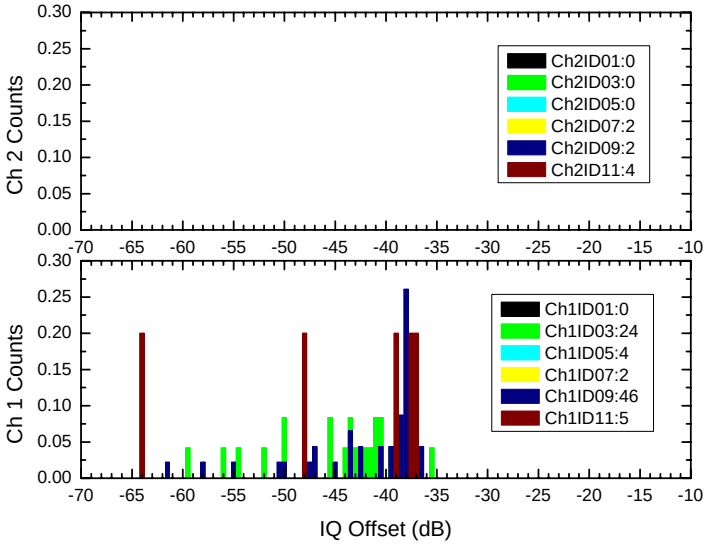

Figure C-83 IQ Offset for D-Link Cards (Experiment 34)

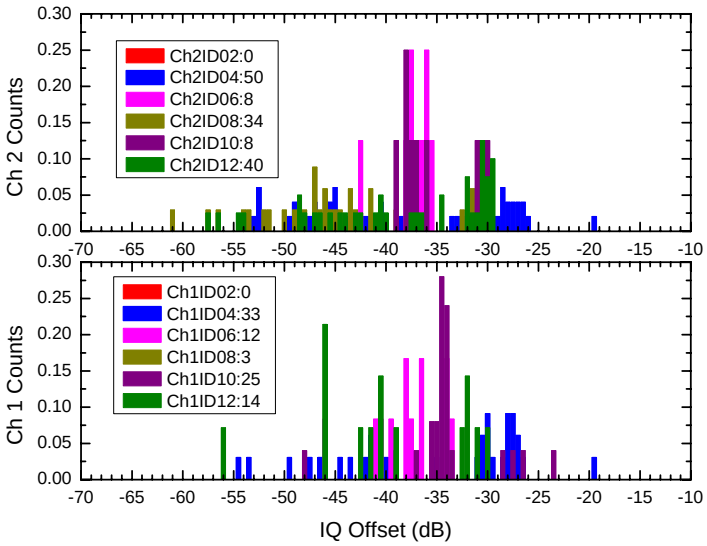

Figure C-84 IQ Offset for Orinoco Cards (Experiment 34)

Appendix D

A CLUSTERING ALGORITHM FOR ANOMALY DETECTION

This section provides an overview of research pursued on clustering algorithms for the WIND project. There are two principal approaches to the development of anomaly detection algorithms: supervised and unsupervised. In supervised anomaly detection, one builds a model of normal system behavior, and then flags anomalies by looking for deviations from that model. A necessary condition for building a supervised anomaly detection algorithm is that adequate collections of purely normal data be available for training the model. In unsupervised anomaly detection, one does not build a model from purely normal training data. Instead, an unsupervised anomaly detection algorithm takes as input a set of unlabeled data, and attempts to find anomalies buried within the data. Thus unsupervised anomaly detection is reminiscent of a classical outlier detection problem. The usual motivation for considering unsupervised anomaly detection is that purely normal training data simply is not available. However, there is a second possible reason. It may happen that the parameters that characterize normal behavior are not constant from one application to the next. In other words, the characterization of normal behavior one derives from a training set, may not be adequate for characterizing normal behavior in the future. Under these conditions, supervised anomaly detection is unlikely to be successful.

The WIND problem is one in which the very notion of "normal behavior" is problematic. That is, the signature of a device in feature space is sensitive to geometry, temperature, and other environmental conditions; and so the "normal behavior" one learns from one application may not apply to a different application. It was this understanding of the problem that led to the investigation of unsupervised anomaly detection algorithms. The specific approach to anomaly detection that was investigated involves geometric clustering in feature space. Under this approach, for any application, the first packets received are all "normal". As the packets arrive, they form different clusters in feature space corresponding to different, "legal" devices. Once this picture of normal behavior has been developed, an intruder can be detected because its packets will form a new, "outlying" cluster. This approach will work, in principle, even if there is no notion of "normal behavior" that remains constant from one application to the next. In particular, the approach described does not require that the clusters associated with a particular device be consistent from one application to the next. It requires only that one have an adequate supply of normal packets for constructing initial, normal clusters in any given application.

As a first attempt, a very simple clustering algorithm in six-dimensional feature space was implemented. The six dimensions are defined by two channels each for frequency error, power, and rise time. The distance metric is defined by Euclidean distance in feature space, supplemented by weights that compensate for the different units used to characterize the different dimensions. The algorithm uses a parameter *maxDist*: a packet is assigned to a cluster only if its distance from the centroid of the cluster is less than or equal to *maxDist*. Simple trial

and error was used to find acceptable values for *maxDist*. Not surprisingly, the algorithm is sensitive to the value of this parameter; different values may work better for different data sets. A single fixed value of *maxDist* was used in obtaining all of the results reported below. The algorithm uses a second parameter *k*: it recomputes the centroids for each cluster after every *k*-th iteration.

A first-cut clustering algorithm in 6-dimensional feature space is presented below:

Geometric Clustering
j = 0;
while there are unprocessed packets
j += 1;
p = current packet;
if *p* is within *maxDist* of some cluster
then add *p* to its closest cluster;
else
form new cluster with centroid *p*;
if *j* mod *k* = 0
then recompute centroids for all clusters;

For purposes of an initial assessment of the clustering approach, it is sufficient to focus exclusively on the devices in each experiment for which there was no intruder. Clearly, a necessary condition for the success of this approach is that the clustering algorithm be able to separate different devices in feature space. This prerequisite is necessary for the algorithm to recognize an intruder cluster.

Initial results for various experiments are tabulated below in Table D-1 through D-6. For each experiment, only those devices for which there was no intruder, and for which at least 50 packets were collected were included. In each table, **ci** notation in the first row denotes cluster **i**. The cluster numbering is generated internally by the algorithm, and has no other significance. Only those clusters that contain at least 50 packets are displayed. Table D-1 provides an example of the kind of results one would hope to see, with each card having a distinct central cluster.

Table D-1 Results for Experiment 16

MAC address	c1	c2	c3
1150081166	179	0	17
000F66D098C7	0	156	3
0011500CD68D	2	2	372

Different devices are almost perfectly associated with unique clusters. Not surprisingly, however, most of the results were more ambiguous. The following example (Table D-2) is more typical.

Table D-2 Results for Experiment 21

MAC address	c1	c2	c3	c4	c5	c6	c7
1150081166	1431	11	69	3	0	1	0
1150081163	2	1438	100	7	0	4	0
0011500815F5	388	115	914	3	0	18	0
001150080C2B	4	1152	53	0	0	4	0
0020A6579181	0	0	13	365	526	3	0
0011500CD68D	0	9	31	0	0	1311	0
001150080C33	2	1	13	0	0	3	1391
000F66D098C7	0	0	7	1512	43	0	0

In this example, the algorithm was unable to distinguish between 1150081163 and 001150080C2B. As Table D-3 through Table D-6 show, most of the results look like Table D-2, where most, but not all, of the devices are unambiguously separable in feature space. It is believed that these initial results indicate that this clustering approach has some promise. But considerably more research is needed in order to address algorithm limitations and transform the approach into an algorithm that is useful in practice.

Table D-3 Results for Experiment 10

MAC address	c1	c2
0011500815F5	1	50
000F66D098C6	184	0

Table D-4 Results for Experiment 13

MAC address	c1	c2	c3	c4
000F66D098C2	7	5	1	159
000F66D098C7	0	4	157	0
0011500CD68D	1	31	83	0
1150081166	314	91	0	2

Table D-5 Results for Experiment 15

MAC address	c1	c2	c3	c4
000F66D098C2	0	135	3	0
000F66D098C7	0	52	0	142
0011500CD68D	3	1	233	0
1150081166	210	1	7	1

Table D-6 Results for Experiment 20

MAC address	c1	c2	c3	c4	c5
000F66D098C2	4	120	0	0	0
0011500815F5	21	1	1	162	21
001150080C33	9	0	280	9	3
1150081166	58	2	1	20	265

LIST OF ACRONYMS

A/D	Analog to Digital Converter
AFRL	Air Force Research Laboratory
AP	Access Point
BAA	Broad Agency Announcement
BER	Bit Error Rate
CCK	Complementary Code Keying
CNO	Computer Network Operations
COTS	Commercial Off the Shelf
CPD	Cumulative Probability Distribution
DBPSK	Differential Binary Phase Shift Keying
DQPSK	Differential Quadrature Phase Shift Keying
EVM	Error Vector Magnitude
FFT	Fast Fourier Transform
FPGA	Field Programmable Gate Array
GPS	Global Positioning System
ID	Identification
IDS	Intrusion Detection System
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
IQ	In-phase/Quadrature
IRAD	Internal Research and Development
JHU/APL	Johns Hopkins University Applied Physics Laboratory
MAC	Medium Access Control
ND	No Data
NI	National Instruments
NSG	Naval Security Group
OFDM	Orthogonal Frequency Division Multiplexing
PC	Personal Computer
PFM	Peak Fitting Module
PLCP	Physical Layer Conversion Protocol
PPDU	PLCP Protocol Data Unit
PSDU	Physical Layer Service Data Unit
PXI	PCI Extensions for Instrumentation
RF	Radio Frequency
VSA	Vector Signal Analyzer
WIDS	Wireless Intrusion Detection System

WIND	Wireless Intrusion Detection
WLAN	Wireless Local Area Network