Countering the parking lot attack: Employing monopulse radar methods to detect and geo-locate RF targets operating in the 2.4 GHz ISM band

by

Daniel James Gieseman

A dissertation submitted to the graduate faculty

in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Major: Computer Engineering

Program of Study Committee:
Thomas Daniels, Major Professor
Yong Guan
Douglas Jacobson
Reginald Souleyrette
Joseph Zambreno

Iowa State University

Ames, Iowa

2015

DEDICATION

I dedicate this dissertation to my family and friends, who always believed I could accomplish this endeavor; to my wife Oksana, who repeatedly encouraged me to continue my research, even in the face of much adversity and time challenges; and to our son Leo, whose sense of marvel and amazement at the world around him provides wonderful inspiration for our future.

TABLE OF CONTENTS

ABSTRACT

Many attacks on information systems occur when an adversary exploits wireless networking technology to remotely gain access to sensitive or confidential data housed within a targeted facility. Where such attack vectors exist, even the most stringent physical security safeguards can fail in preventing an attacker from executing a stand-off attack aimed at compromising facility systems. This class of attack, where the attacker remains positioned outside the physical confines of a facility and instead penetrates a network using a wireless vulnerability, is commonly referred to by security researchers as the "Parking Lot Attack".

In this work, we present a scheme deploying an integrated network of sensors intended to detect and geo-locate any wireless emitter attempting the Parking Lot Attack. A novel feature of our system is the use of monopulse radar methods to assign Line-of-Bearing estimates to any detected RF target. Our design provides for a series of cooperative, mechanically steered, detection sensors each employing a remotely controlled monopulse antenna array. Combining the Line-of-Bearing estimates obtained from multiple detection sensors supports the processing functionality required to geo-locate any RF station actively transmitting within range of our system.

Our research program encompasses three facets, presented as separate chapters in this dissertation. We first describe our system design and architecture, and then we proceed into a quantitative study focused on analyzing the performance of a prototype detection sensor we developed to support field-operational experimentation. We also constructed a software model of our detection system, suitable for simulation studies. We describe how the simulation toolset can be utilized to answer questions about system deployment and operational strategy.

CHAPTER 1. BACKGROUND AND DISSERTATION OVERVIEW

The topic of localization using radio frequency methods interests us from an information systems security viewpoint due to the enhanced situational awareness obtained when one is provided precise, geo-spatial location data for any broadcasting RF station operating within an established zone of interest. The zone of interest may be broad in geographic scope, such as tracking the location of a mobile device subscriber as they move about the area served by regional or national cellular infrastructure, or as in our specific focus of study, the area may be more geographically constrained, pertaining only to the coverage area within range of the wireless LAN (WLAN) environment of a facility.

The WLAN environment is typically deployed to serve wireless stations operating internally, within the confines of the facility, but as we will show, the convenience of wireless accessibility has been repeatedly exploited by attackers as a means of gaining easy access to facility information systems. Security researchers use the term “Parking Lot Attack” to describe scenarios where an adversary recognizes wireless vulnerabilities that enable external remote access. The adversary then exploits those wireless vulnerabilities to gain access to the sensitive systems and data contained with the facility; this is often accomplished without the need to thwart any physical access controls setup to prevent unauthorized facility access.

In this work we focus to either prevent or reliably detect when a Parking Lot Attack is carried out against a facility. We present our research into a system we developed using RF localization methods to detect emitters actively transmitting within the environment which is external to the facility we aim to defend against an adversary employing the Parking Lot Attack.

We begin Chapter 2 with a review of relevant and contemporary RF localization work, emphasizing wherever possible those systems intended specifically for security domain applications. We seek to compare and contrast the underlying methods presented in each example of past work with that of our own system. Following Chapter 2, we present research originally prepared as a series of papers, each focused on a related aspect of our study into Parking Lot Attack detection and methods for RF target spatial-attribution. These papers are currently in the peer review process at several refereed journals and conferences.

In Chapter 3, we present the concept of operation, and architecture for our system designed to protect a facility from stand-off wireless attacks. We include a detailed threat model describing the facility attack surface and the capabilities and intentions of the adversary seeking to attack the facility protected by our system. Chapter 3 concludes with detailed discussion of a prototype detection sensor design which we constructed to perform real RF localization tests against wireless targets positioned on a surveyed field test range. A novel feature of our design is the use of monopulse radar methods to enhance the detection sensor Line-of-Bearing estimates collected by the sensor. Monopulse sensor readings support the positional spatial-attribution of detection targets using triangulation techniques; the more accurate the sensor reading, the more accurately the sensor position estimates become. The use of low-cost and commercially available hardware was also a design objective and something we touch on in Chapter 3.

Chapter 4 is focused on the design and evaluation of a series of field experiments we conducted using the prototype detection sensor hardware we presented in Chapter 3. In our experiments we quantify detection sensor performance in terms of the boresight error, which was calculated using data collected by lobing an RF target positioned at a surveyed, pre-

known location. We also present several software implemented digital signal processing (DSP) detectors and discuss the tradeoffs of these detectors in terms of the detection system error budget. Chapter 4 concludes with a detailed quantitative study of the sensor field performance and recommendations for future sensor design research and platform tools.

In Chapter 5 we choose a more theoretical path, exploring questions focused around determining the best strategy for effectively deploying and operating the network of detection sensors we designed for the protection of a facility from the Parking Lot Attack. To study these strategy questions, we developed a simulation tool capable of performing many iterations of randomized attack and defend scenarios. We first present several core strategy questions and then detail a series of simulation experiments which were performed to quantitatively assess our performance questions. The simulations permitted us to more fully explore our system in terms of scalability and true detection network performance. Since we only constructed a single instance of our prototype detection sensor for field research, simulations were a logical choice for exploring how a cooperative network deploying multiple sensors could realistically perform in a simulated hostile threat environment.

Chapter 6 summarizes the important findings of our research. In addition, we feel that we have uncovered many more questions than answers in our research, so we conclude this work by making recommendations for further study of the research problems we feel are still important topics for study in this domain.

CHAPTER 2. RELATED WORK

Prior art played an important role in laying the foundation for our research effort. The topic of localization using RF technologies is nearly as old as the subject of radio itself. Marconi described to the IEEE society in 1922 how, years earlier, he observed radio waves reflecting from objects in the environment and how he believed that such phenomena could aide in guiding ships at sea through foggy conditions [1]. It was not long after the advances of Marconi that the British "Chain Home" system of radio direction finding, led by Robert Watson-Watt, was fielded to protect United Kingdom airspace from hostile aircraft during World War II [2] [3].

While our research program borrows methods and techniques from the radar literature, the system we have developed fits better in the localization category of RF triangulation. This is primarily because our system lacks the target ranging capability inherent to any modern radar system. However, our system does employ monopulse radar methods to improve the accuracy of a Line-Of-Bearing (LOB) estimate calculated for any detected RF target. Excellent primers on the topic of monopulse radar can be found in [4], [5], [6].

It is the combination of multiple LOB estimates collected from spatially separated sensors in our system that enables triangulation to occur. For once two or more devices have estimated the LOB to a target, the LOB lines may be geometrically intersected to calculate a position point. In practice, LOB lines are better represented as polygons, due to estimation errors inherent to our detection system, and the position point is better characterized as a probable region of signal origination, falling within the geometric region where the LOB polygons intersect. We will explore this in much more detail in subsequent chapters.

Since the pioneering work of Marconi and the technologies inspired by Watson-Watt and others, there have been countless developments in the domain of RF localization, with topics ranging from those focused on military and defense, commercial applications of location-based services, and those serving utility functions such as personnel, asset, and inventory tracking. The entire field of study concerning wireless position estimation can best be described as extremely broad; for the purposes of providing the necessary backdrop contextualizing our research, we have elected to focus on those radio-location systems supporting an intrusion detection or information security feature set. A thorough technical treatment of the many types of localization schemes pertaining to security and electronic warfare models can be found in [7].

We can further subdivide security and intrusion detection oriented systems into those systems which operate using directional antennas, which are similar in architecture to our own detection and localization system, and those systems which employ other triangulation methods, such as Time Difference of Arrival (TDOA) and RF propagation path modeling based on Received Signal Strength (RSS). Our own system, using directional antennas would best be classified as an Angle of Arrival (AoA) system.

Localization Systems Employing RSS and TDOA Methods

There are excellent commercial and research examples of localization systems employing RSS mapping for device tracking [8], [9], [10], [11], [12]. These systems feature solid indoor performance with accuracy potential to within a few meters in indoor environments. While most systems in this category were designed for indoor operation, a University of Washington study did report attempting the method outdoors [10]. The drawback of these systems is the costly setup time and the periodic re-sampling required for

maintaining system accuracy. This is because most examples are not zero knowledge systems, instead the functionality relies on the creation of a database of signal propagation characteristics requiring potentially thousands of empirically collected sample points to accurately model the path traveled by wireless signals within the service set domain. Although, at least one system, RADAR from Microsoft Research reported an attempt to reduce the setup time required by these systems using sophisticated signal propagation modeling [8].

Localization using the RSS and propagation modeling approach has matured to the point where many product offerings are available which feature the technique. A typical example for personnel and inventory tracking can be found in Cisco's Wireless Location Appliance [9]. This system advertises the capability to track the location of thousands of wireless 802.11 devices within a service set domain. There are security features bundled into this product, a white-list of allowed devices can be maintained, so that unauthorized devices can be flagged and located based on a signal fingerprint lookup found in the signal propagation database. A changing propagation environment will necessitate the need for periodic resampling to occur in order to maintain the accuracy of the signal propagation model. Comparable commercial systems are presented in [11] and [12].

From a security context, a key drawback of systems employing RSS was the host-based orientation of the localization method. Hosts utilize a priori signal propagation data for the transmit/receive characteristics of WLAN Access Points operating in the environment to estimate their position. This is because it is simpler to collect samples for a small group of access points with known signal transmit power levels, than to collect propagation samples for each individual device operating within the WLAN environment. Propagation models

depend on knowing the transmit power, and the received signal strength to estimate device location. Ultimately, localization then depends on the subscribing hosts being cooperative in sharing their estimated position to the central network, something that is obviously not likely to occur in the context of an enemy perpetrating a wireless intrusion.

The most notable example of using TDOA for position estimation would be the Global Positioning System (GPS). Systems using TDOA use the known locations of a network of stationary reference stations, along with a precisely synchronized distributed clock, to measure the time it takes for a signal originating at a reference station to arrive at the receiving device. Multiple readings received from different reference stations can then be aggregated to triangulate position. The nature of the timing and measurement methods make these systems more complex to implement. They also suffer from the same host-based limitations we just described as being inherent to methods using RSS, making them of little use to a security or intrusion detection system. The system is again dependent on an enemy being cooperative in broadcasting her calculated position estimates. Nevertheless, TDOA is an important technology to understand when contextualizing the system we have developed in our research.

Most TDOA examples in the research literature attempt to use Access Point infrastructure that has been modified for accurate timing measurements to duplicate GPS functionality in indoor environments. A by no means exhaustive grouping of TDOA research studies using WLAN access points for both indoor and outdoor localization can be found in [13], [14], [15], [16].

Localization Systems Using Directional Antenna Methods for Intrusion Detection

Systems using directional antennas feature less synchronization and timing complexity than more sophisticated TDOA methods, and require much less setup time when compared to the empirical propagation modeling methods used in RSS location schemes. Directional antenna methods provide the capability to setup and operate in an environment with zero knowledge of the RF landscape. This capability is critical to the security context of our research, as these systems represent a class of functionality in direct contrast to the methods we have previously reviewed: host-based cooperation is not required for positional estimates to be made. This permits directional antenna systems to potentially detect and locate the uncooperative adversary. These systems are however much more susceptible to multipath fading and interference making them better suited for outdoor environments. However, at least one system, a University of Illinois study, employed outdoor directional antennas for location tracking indoors [17]. The system we present does indeed leverage this better suitability for outdoor or clear line-of-sight requirements, by deploying directional antenna-based detection sensors externally, on the facility perimeter, with outward facing orientations.

A paper published by researchers at the University of New Orleans is the standout example of a system with direct security and intrusion detection context [18]. In this paper the authors demonstrate the use of a high-gain antenna system for active transmitter localization. The authors combine an anomaly and signature-based Wireless Intrusion Detection System (WIDS) with the several types of high-gain antennas to locate a wireless intruder.

The University of New Orleans study concludes by repeating the utility of directional dish antennas as a means of accurately measuring the angular bearing of a receiver to a transmitter. The authors also mention their intended next step is motorizing the directional antenna for automated tracking of a transmitter. They mention that a drawback of this approach is that only active transmissions are detected. A passive eavesdropper cannot be located with directional-antenna based location methods.

Our conceptual design builds upon systems like that presented by University of New Orleans and utilizes a steerable directional antenna to attempt emitter triangulation using line of bearing estimation. Conceptually, our emitter spatial attribution system is composed of a distributed network of steerable high-gain antennas. Antenna angular direction is controlled automatically using a positioning motor. Directional heading is also sensed so that a scan bearing can be measured and reported.

One unique aspect of our system is that our design provides for the interoperation of multiple sensors in a cooperative manner. In this scheme, a spatially distributed network of independently scanning sensors would communicate and share collected LOB data using a wired distribution system, operated out-of-band from the wireless domain under protection. Sensors would communicate calculated LOB vectors along with GPS location coordinates to a command and control sub-system using the distribution system. Command and control would fuse disparate sensor LOB measurements to perform emitter triangulation calculations and direct future sensor scanning movements for dynamic tracking of emitters.

We found a scheme very similar to our design in use in a University of Greenwich led localization study [19], where a mobile rover operating a mechanically steered directional antenna was used to probe an environment for the location of operating WLAN Access

Points. Similar to our own system, the application was intended for outdoor deployment and operation. The location of the rover was tracked and maintained using GPS, enabling the mobile sensor to perform triangulations using a database created by the sensor as it moved through and about an environment. Unlike our system, this scheme did not utilize an antenna array, whereas our system seeks to enhance LOB angular estimates by applying monopulse radar techniques enabled by the use of an antenna array. Key finding of the University of Greenwich study were that higher gain antennas were needed to increase accuracy and the ability to detect targets operating more than 50 meters from the collection sensor. Our system employs 16dBi high-gain antennas to counter just such a limitation. A system nearly identical to the University of Greenwich system is described in [20].

Directly related to the prior research and again featuring similar directional antenna functionality was a Plymouth University study focused on discovering hidden WLAN enabled mobile devices [21]. This study featured a hand-held device with a high-gain antenna. The operator manually panned the antenna while watching signal strength readings filtered to only show reading for the network MAC address of the device being targeted. Again, this device only used a single antenna, and required the operator to actively engage in seeking the target. Our system instead focuses on passive monitoring methods, with operator notification only in the event of unauthorized detection. The Plymouth University system could be a useful tool to employ in combination with our own system, where the hand-held scanning system is used to very precisely sweep an area where our detection system estimates that a wireless attack is originating from.

We found very few examples employing monopulse radar antenna arrays in the 2.4 GHz spectrum band. One example was found using monopulse methods to track RFID tags

in the retail store setting [22]. No examples were found using monopulse methods for WLAN intrusion detection, which is the basis for our system of detecting unauthorized WLAN intrusions. We now turn to a detailed study of our detection system architecture, functionality, and theoretical performance.

References

[1] G. Marconi, "Radio Telegraphy," *Proceedings of the Institute of Radio Engineers,* vol. 10, no. 4, pp. 215-238, 1922.

[2] B. A. Austin, "Precursors to RADAR - The Watson-Watt Memorandum and the Daventry Experiment," *International Journal of Electrical Engineering Educators,* vol. 36, pp. 365-372, 1999.

[3] J. Gough, Watching the Skies: The History of Ground Radar in the Air Defense of the United Kingdom, Her Majesty's Stationary Office, 1993.

[4] D. R. Rhodes, Introduction to Monopulse, Artech House, 1980.

[5] S. M. Sherman and D. K. Barton, Monopulse Principles and Techniques, Artech House, 2011.

[6] A. I. Leonov, K. I. Fomichev, W. F. Barton and D. K. Barton, Monopulse Radar, Artech House, 1986.

[7] R. A. Poisel, Electronic Warfare Target Location Methods, Artech House, 2012.

[8] P. Bahl and P. V. N., "RADAR: An In-Building RF-based User Location and Tracking System," Microsoft Research, 2000.

[9] Cisco Systems, "The Cisco Wireless Location Appliance," [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/wireless/wireless-location-appliance/product_data_sheet0900aecd80293728.html. [Accessed 09 2015].

[10] J. Letchner, D. Fox and A. Lamarca, "Large-scale localization from wireless signal strength," in *Proc. of the National Conference on Artificial Intelligence*, 2005.

[11] Hewlett-Packard, "Accurate Indoor Positioning with HP Location Analytics," [Online]. Available: http://www8.hp.com/h20195/v2/GetPDF.aspx/4AA5-6244EEW.pdf. [Accessed 09 2015].

[12] Apple, Inc, "iOS: Understanding iBeacon," Apple, Inc, 02 2015. [Online]. Available: https://support.apple.com/en-gb/HT202880. [Accessed 09 2015].

[13] X. Luo, W. Li and J. Lin, "Geometric Location Based on TDOA for," *ISRN Applied Mathematics,* vol. 2012, 2012.

[14] Z. Nemec and P. Bezousek, "The Time Difference of Arrival Estimation of Wi-Fi Signals," *Radio Engineering,* vol. 17, no. 4, pp. 51-54, 2008.

[15] K. Keunecke and G. Scholl, "Wi-Fi-Based Performance Analysis of TOA/TDOA Estimators by Stochastic Channel Simulations," in *The Tenth International Symposius on Wireless Communications Systems*, Berlin, 2013.

[16] B. Li, I. J. Quader and A. G. Dempster, "On Outdoor Positioning with Wi-Fi," *Journal of Global Positioning Systems,* vol. 7, no. 1, pp. 18-26, 2009.

[17] T. Pongthawornkamol, "Zero-knowledge real-time indoor tracking via outdoor wireless directional antennas.," in *PERCOM: IEEE International Conference on Pervasive Computing and Communications*, Mannheim, 2010.

[18] F. Adelstein, A. Prasanth, R. Joyce and R. G. Golden, "Physically Locating Wireless Intruders," *Journal of Universal Computer Science*, vol. 11, no. 1, pp. 4-19, 2005.

[19] A. Ibrahim and D. Ibrahim, "Real-time GPS based outdoor WiFi localization system with map display," *Advances in Engineering Software*, vol. 41, no. 9, pp. 1080-1086, 2010.

[20] A. Subramanian, "Drive-by Localization of Roadside WiFi Networks," in *INFOCOM: The 27th Conference on Computer Communications*, Phoenix, AZ, 2008.

[21] N. C. M. Dagnall, "A Feasibility Study into Tracking Wi-Fi Enabled Mobile Devices," *Proceedings of the Ninth International Network Conference (INC2012)*, pp. 85-92, 2012.

[22] R. Parada, J. Melia-Segui, A. Carreras and R. Pous, "Study of a Monopulse System with RFID Antennas for Applications Oriented to Retail Industry," in *UbiComp Proceedings 2013*, Zurich, Switzerland, 2013.

CHAPTER 3. COUNTERING THE PARKING LOT ATTACK – DESIGN FOR A DETECTION SYSTEM EMPLOYING MONOPULSE RADAR METHODS TO DETECT AND SPATIALLY ATTRIBUTE RF TARGETS IN THE 2.4 GHZ ISM BAND

A paper submitted to *DEFCON Hacking Conference*

D. J. Gieseman[1,2] and T. E. Daniels[1]

Abstract

Many attacks on information systems occur when an adversary exploits wireless networking technology to remotely gain access to sensitive or confidential data within a targeted facility. Where such attack vectors exist, even the most stringent physical security safeguards can fail in preventing an attacker from executing a stand-off attack aimed at compromising facility systems. This class of attack, where the attacker remains positioned outside the physical confines of a facility and instead penetrates a network using a wireless vulnerability, is commonly referred to by security researchers as the "Parking Lot Attack." In this work, we present a scheme deploying an integrated network of sensors intended to detect and geo-locate any wireless emitter attempting the Parking Lot Attack. We first introduce the context for such a system by presenting a threat model describing the facility and data systems targeted for attack. Specific vulnerabilities in the attack surface of our model, which make the Parking Lot Attack a viable and preferred

[1] Graduate Student and Assistant Professor, respectively, Department of Electrical and Computer Engineering, Iowa State University.
[2] Primary researcher and author.

exploitation vector, are explained. We describe the motivations and capabilities of the adversary employing this attack, and provide a constrained, but realistic Parking Lot Attack scenario which drives our detection system design. A novel feature of our system is the use of monopulse radar methods to assign Line-of-Bearing estimates to a detected RF target. We cover a concept of operation and propose a deployment scheme for facility protection using our system. We then discuss in detail the design and architecture of a second generation sensor implementation; a device which we constructed to perform real operational experiments. We conclude with a brief demonstration of several key detection and sensing features of the device, saving more in-depth treatment for a planned follow on paper focused on quantitatively analyzing the performance of our system.

Background and Research Context

The home, business, appliance, and automobile are ever increasingly connected using radio frequency technologies. As one example, consider the broad range of devices that operate in just the license-free Industrial, Scientific, and Medical (ISM) RF bands. From a cyber-security perspective, these wireless communications platforms present an adversary an attack surface that is much more accessible – and hence much more vulnerable – than their hard-wired counterparts.

Vast amounts of time and resources are expended securing wireless protocol stacks from eavesdropping and unauthorized access. However, high-profile and costly attacks against wireless networks still persist [1] [2]. These attacks permit a stand-off adversary remote access to sensitive systems and data that would be more easily secured had system communications avoided the use of wireless technologies altogether. To counter this threat, monitoring agents serve alongside other underlying information

assurance mechanisms to provide tactical threat intelligence and situational awareness with the aim of ensuring that system integrity is maintained and that the system remains available. Monitoring of wireless – or any network communications traffic – is a mature discipline with many useful systems and tools supporting features such as load balancing and performance measurement, in addition to the detection of misuse and anomalous network communications.

Spatial Attribution in Wireless Network Monitoring Systems

Spatial attribution of wireless devices and RF activity within a secured environment is a more recent addition to these monitoring tools, and remains an interesting area of exploration for wireless security researchers. Monitoring systems capable of spatial attribution are being developed with the intent of extending device position information to traditional monitoring applications, thus augmenting traditional situational awareness capabilities with the locations of wireless devices interacting with the WLAN environment of a facility. Commercial and research systems for indoor spatial attribution of wireless devices are shown here [3] [4] [5].

Indoor environments in modern office structures are filled with the metallic materials that wreak havoc on radio frequency transmissions. This complicates the spatial attribution problem for systems seeking to track and geo-locate RF devices operating inside a structure. Studs internal to walls, building elevator shafts, heating and ventilation systems, and miles of cabling all contribute to signal multi-path reflections causing fading, cancellation, and other forms of RF propagation jitter. These non-linear effects markedly impact how systems for indoor wireless device spatial attribution are designed and function. Many of these systems, which behave like and resemble wireless LAN

access points, operate using pre-configured, static maps of the wireless environment inside a facility [3], [5]. These maps are developed while moving a target device with a known position and RF signature about a facility during a physical audit phase of the system deployment.

Statement of Research Objective

One limitation inherent to these systems is their indoor-oriented focus towards the RF environment internal to the facility where they are deployed. They do not monitor or attempt to spatially attribute devices connecting to internal wireless networks from outside the facility. However, one common attack vector found on the attack surface of wireless networks is the compromise of the sensitive systems and data of a facility by an adversary positioned **outside** the physical confines of the facility; an attack commonly referred to in literature as the "Parking Lot Attack" [6], or PLA for short.

- Our objective is to research systems and methods designed to detect an adversary employing the Parking Lot Attack as the primary penetration vector during a Wireless LAN incursion.

As previously stated, an adversary using the Parking Lot Attack vector is located off-premises, not actually physically present within a facility targeted for malicious penetration. If our aim is to detect the PLA, a critical facet of our research should then focus on the problem of detecting and spatially attributing RF devices operating external to the physical confines of a facility of interest. For if we know that a wireless device is located outside a facility, and that this device is attempting communications with a network internal to the facility, we can then alert security personnel who can investigate whether this is malicious or unauthorized behavior. The focus on attacks originating

external to a facility presents a different RF environment than that of the detection and monitoring schemes described above for indoor applications.

However, we feel that looking outward for an attack originating external to the facility perimeter provides us with advantages in terms of more well-behaved RF emissions, relative to the indoor environment. The outdoor environment surrounding a facility we wish to protect from external attack can be conditioned much more easily for RF sensing. Sensor locations can be selected such that clear lines-of-sight mitigating multipath reflections can be obtained about the perimeter, allowing more accurate angle sensing of RF targets operating nearby. In other words, the dense undergrowth of metallic obstructions present in the facility interior is removed from the equation. Due to this, the system of sensors we discuss and have deployed for our research is designed and functions differently than those systems previously mentioned for indoor spatial attribution. One benefit of this shift in design focus is that our system for monitoring the external environment of a facility does not require previous RF environmental mapping. In terms of form and function, our system resembles much more a radar installation than the wireless router add-on systems we described for interior deployments.

- The electronic toolset requirements for the detection and location of RF devices external to a facility differ from the toolset required for interior facility monitoring.
- This can be attributed to an external RF propagation environment featuring a more open electromagnetic landscape that can be conditioned for clear lines-of-sight. Clear lines-of-sight support accurate detection sensing.

Enhancing RF Situational Awareness to Counter the Parking Lot Attack

It is our premise that accurate and precise wireless device spatial attribution would enable virtual exclusion zones to be configured that monitor wireless activity external to a facility. We define an RF Exclusion Red Zone, or just Red Zone for short, as a zone where wireless activity, when detected within the zone perimeter, would be classified as potentially malicious, and in need of further investigation. As an example, a threat analyst could direct counter-measures towards an adversary attempting to execute the Parking Lot Attack against a facility defended by this capability. The system could detect an emitter transmitting in an unauthorized zone, and notify the threat analyst so that an attack disposition can be determined and a suitable defensive posture taken.

WIDAR – Wireless Intrusion Detection and Ranging

Our research goes beyond a purely theoretical design as we present a second-generation implementation of our system concept, in the form of a network of sensors which we call WIDAR – short for **W**ireless **I**ntrusion **D**etection **a**nd **R**anging. Before we delve into the details of our implementation design and architecture, we first cover the threat model which conceptually defines the theoretical facility we wish to protect from the Parking Lot Attack, as well as the capabilities of the ever-persistent and diabolically-evil Adversary; whose aim is to steal the sensitive data stored in the information systems maintained within this facility.

We then detail specific system architecture and device capabilities and at the same time offer our motivation for development decisions which were made during system design and implementation. Since the focus of the second part of this work is on a real implementation constructed for real system experiments, we also include many diagrams

and photos of our system detection sensor, as constructed, along with screen captures of system control and analysis software. Whenever appropriate, we also include presentations of data collected from ongoing preliminary bench testing of our device, saving a full analytical review for a forthcoming paper focused around quantitative experiments analyzing device spatial attribution outputs when operated against RF targets with known position.

Threat Model and Adversary Capabilities

The threat model sets the stage for experiments and analysis of scenarios where the Parking Lot Attack is the primary vector chosen by an adversary. The location, technology, and methods selected for the attack directly drive the design and implementation of our system for detecting and countering the PLA. It is important to discuss in detail the physical environment and technological circumstances that lead an attacker to choose the PLA as the best option for successful compromise. While defining the threat landscape we will, at the same time, place some presumptive constraints on our model that we feel are necessary to maintain focus on the wireless attack vector. By applying such rigid constraints, we can eliminate the need to consider parallel threats as part of this research, allowing for a more controlled study of the Park Lot vector and its detection.

We divide the system threat model into two categories:

- A description of the facility protected by our system, which we call the Facility Under Protection.

- A description of the motivation and technical capabilities of the theoretical Adversary, who seeks to carry out offensive information operations against the wireless attack surface of our model facility.

The Facility Under Protection

The Facility Under Protection (FUP), along with the threat model for it, both serve to contextualize our research focus, which is to detect and geo-locate an attacker carrying out the Parking Lot Attack. We use the term *Under Protection* to indicate that the facility is being guarded by our system; a system composed of a network of sensors designed to protect against external wireless attacks. A diagram illustrating the FUP model is shown in Figure 3.1.

In our model, size and physical layout of the facility and its internal floor plan are not critical parameters. Instead, the key takeaway is that this facility houses secured information systems hosting sensitive systems and data. For scenario realism, imagine that these systems stage and process information considered vital to the security of a nation, or databases storing the personally identifying and financial transactions of customers doing business with a large corporation.

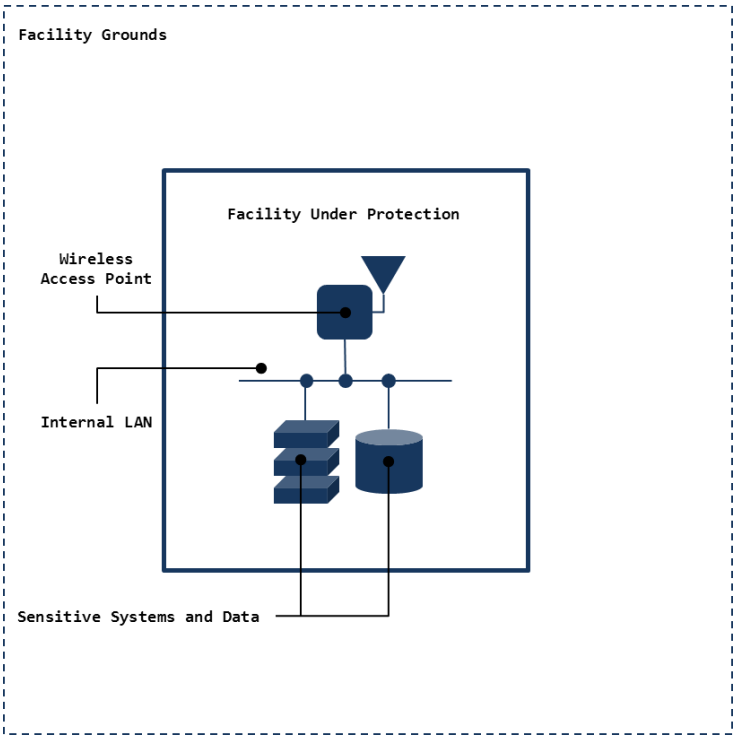

Figure 3.1. Model of Facility Under Protection. A wireless access point inside the facility presents a vulnerability permitting external attacks on the sensitive systems and data contained within the seemingly secure facility. The Access Point may or may not connect directly to the sensitive systems and data, instead our model assumes that, once internal LAN access is obtained by the Adversary, a mix of stepping stone and privilege escalation attacks can provide a path to the sensitive systems and data. In our model, an attack originating from the Facility Grounds is the preferred exploitation vector of choice for our Adversary.

We assume that the facility has been hardened against physical attack and penetration in some manner, thus making attack vectors executed against wireless infrastructure a more tempting and preferred avenue of compromise. While there are a multitude of wireless vulnerabilities that may be discussed in the context of the attack surface of our model facility, only one is germane to our research:

- Some segment of the data communications network operated within our model facility has – intentionally or un-intentionally – a remotely accessible Wireless Access Point. Given the multitude of reported real world PLA examples, this should not be an entirely unbelievable premise.

Lastly, although it is conceivable that wireless attacks could just as easily be executed from a location *internal* to the FUP, we wish to constrain our attack scenario even further by adding that an internally originating wireless attack is not a preferred attack vector. In our model, this is made true by the one reason we already mentioned: strict physical access controls. Physical access to the facility is authenticated, and perhaps entrants are also subject to search, or, as can be commonly seen in highly secure facilities, wireless devices are banned from the facility altogether. The main point we want to make is that for an inside attack to be carried out, the risk of an attacker being detected while executing an internally launched attack is orders of magnitude greater than the risk outwardly presented by the Parking Lot Attack vector. We say *outwardly presented*, because our model assumes that an adversary is not aware of the counter-measures our system deploys at the facility to detect and thwart PLA attempts.

- In our Threat Model, the FUP is MOST vulnerable to an adversary executing the Parking Lot Attack.

This leaves the remote or stand-off attack as the path of least resistance for an attacker wishing to penetrate the sensitive systems and data of the FUP. In Figure 3.1 we define the region external to the facility interior as the Facility Grounds. In most real world examples, the grounds will contain some type of parking area for employees and visitor access to the facility.

- The Facility Grounds serve as the launch pad for the Parking Lot Attack.

The Wireless Access Point and Theoretical Vulnerabilities

Even though it can be assumed that connections to the FUP wireless access point are authenticated and encrypted, in our threat model the use of authentication and

encryption mechanisms do not entirely mitigate the remotely executed threats directed at the FUP. There are many examples of attacks exploiting vulnerabilities in provably secure – at least provable in the mathematical sense – wireless networking systems [7] [8] [9] [10]. In, other words, our threat model contains vulnerabilities that mainly exist due to flaws in engineering and implementation, not necessarily for design reasons. These are the domain of buffer and heap-overflows, integer off-by-one's, race-conditions, and statistical attacks; exploitable implementation flaws employed by attackers as attack vectors against software targets.

Adversary Capabilities—Carrying Out the Parking Lot Attack

The Adversary in our threat model is both determined to gain access to the sensitive systems and data housed in the secured Facility, and she has the skills and expertise required to reach this objective. The Adversary is a cautious and capable planner, and we should assume that a combination of offsite and onsite pre-attack reconnaissance, obtained through careful surveillance, has led the Adversary to discover the Wireless Access Point in use at the FUP. Furthermore, the hardened nature of the physical security at the FUP was also discovered as part of the same comprehensive and sophisticated pre-attack preparation. Knowledge of these environmental conditions has led her to make the determination that the best avenue of attack is to overtly or covertly enter the facility grounds, either on foot or in a vehicle, in order to stealthily launch her wireless offensive from a location external to the facility structure itself.

- The determined and well prepared Adversary is going to discover and exploit the FUP vulnerability presented by the Wireless Access Point.

- The FUP is in need of a system that can detect the Parking Lot Attack so that counter-measures can be directed against the Adversary.

Fingerprinting of services and passive monitoring of RF activity near the FUP has provided her with a multitude of intelligence assisting her offensive. We should assume that this intelligence is sufficient to enable the Adversary to develop and dry-run a series of attack scenarios. The result of this careful planning culminates in several full-dress rehearsals, each resulting in full compromise of a simulation environment mimicking the targeted FUP Sensitive Systems and Data. The electronic battle plan for the attack is entirely predicated on the external access afforded by the FUP Wireless Access Point.

- The Wireless Access Point is the crack in the armor of the FUP, and our Adversary is prepared, and is confident that her attack will succeed.

We show our attacker executing the Parking Lot Attack in Figure 3.2. We should note that the Adversary is sophisticated and experienced enough to know that she cannot simply drive up, find a parking location, and pull out a laptop computer and begin working while sitting behind the wheel. She will most likely employ camouflaging methods that hide the fact that her attack is being executed. We mention this to make the point that any detection scheme we design to detect and thwart the Parking Lot Attack should not rely exclusively upon visual detection of the Adversary in the act of carrying out the Parking Lot Attack – using, for example, security or close-circuit television cameras. Instead a detection scheme should rely on sensors operating in the PHY and DATA LINK layers to detect any cyber-assault.

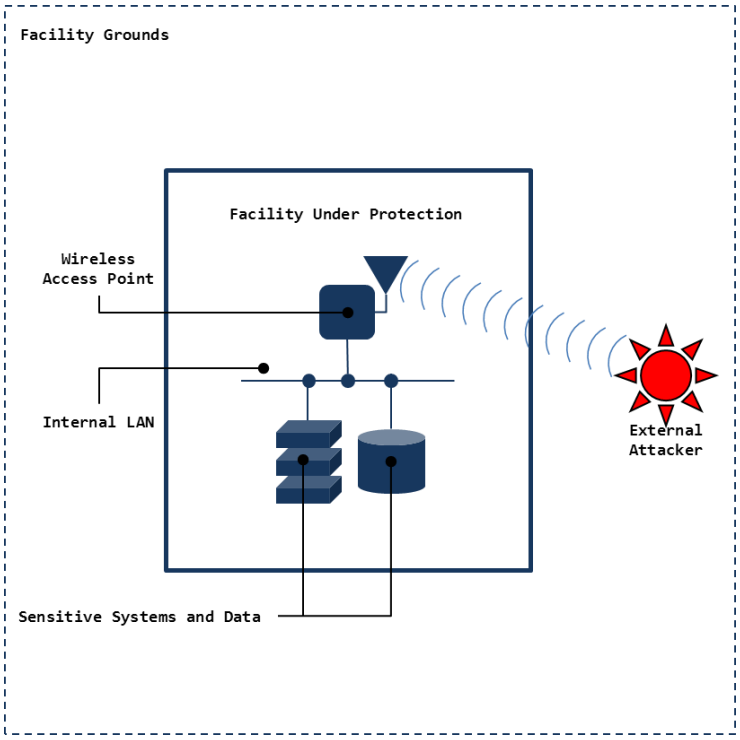

Figure 3.2. The Adversary executes the Parking Lot Attack. A wireless access point inside the facility is permitting an attacker, who is on the Facility Grounds but has not physically entered the hardened Facility interior, to gain access to Sensitive Systems and Data.

It is unimportant whether the sensitive systems and data are actually directly accessible via the wireless local area network upon a successful PLA launched wireless compromise. Rather, our threat model simply assumes that there exist other vulnerabilities permitting stepping stone and privilege escalation attacks using a variety of means, once wireless entry is achieved.

System Conceptual Overview and Operational Strategy

We next conceptually describe the system conceptual model, and the constituent components of this model. We then describe how components are deployed in and about the FUP perimeter to scan for and detect any RF emitters operating within the pre-designated boundary of a sensor scan zone. We also detail how sensor components are connected via a distribution system enabling communications, and the manner by which

sensor data are aggregated. All devices are administered via a centralized command and control system utilizing this same distribution system.

The RF Exclusion Red Zone

The components interoperate as part of an integrated system designed to detect and assign a spatial location to any RF target operating within a pre-designated RF exclusion area, which we call a “Red Zone”. The Red Zone must be pre-defined, following system deployment, by a Threat Analyst as a region on the facility grounds where RF activity is unauthorized and where RF activity, if detected, triggers an alert. The Threat Analyst can then make a determination whether or not an actual attack is occurring and whether to deploy countermeasures.

In our system conceptual model, we designate the entire facility grounds, external to the FUP as the RF Exclusion Red Zone. Figure 3.3 illustrates the RF Exclusion Red Zone, in the context of our FUP model.

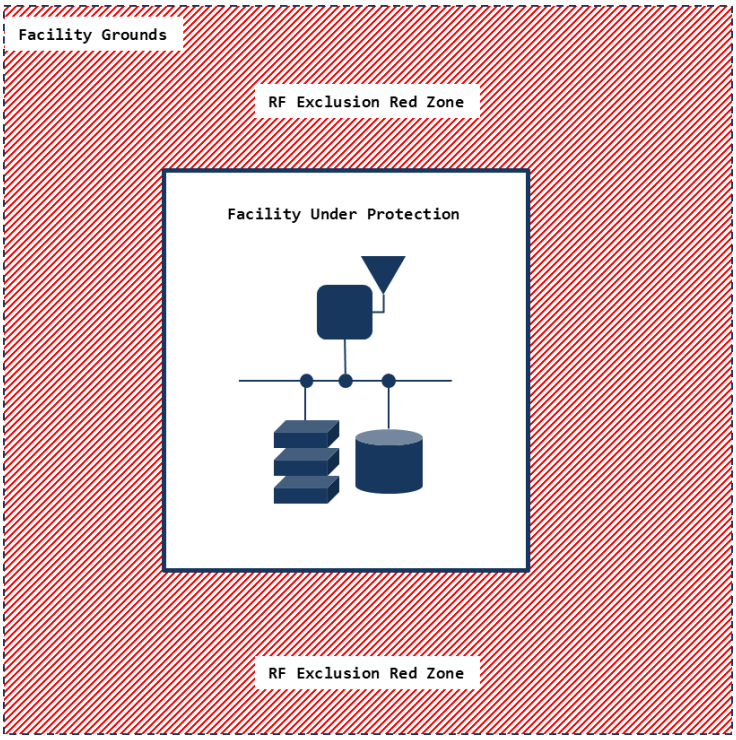

Figure 3.3. The RF Exclusion Red Zone about the Facility Under Protection Perimeter. In our model we designate All Space exterior to the Facility Under Protection as Red Zone. Any unauthorized RF emitter operating in this Red Zone will cause an alert to be issued to an on duty Threat Analyst.

The System Sensor

The system components which are most mission-critical to our model facility attack detection system are the devices for detecting and spatially tracking wireless intrusion attempts at the facility perimeter. Indeed the entire premise of our system, monitoring RF Exclusion Red Zones, is predicated upon the existence of some capability to accurately detect and spatially attribute RF targets operating in the vicinity of the FUP. Since our system is called WIDAR (Wireless Intrusion Detection and Ranging), the devices deployed about the facility edge to detect intrusions are called WIDAR devices, WIDAR sensors, or simply “detection sensors” throughout the remainder of this paper.

In our system design, it is critical that a sensor in our model be capable of producing the following detection outputs:

- Line-of-Bearing estimate (from detection sensor boresight) for any active targets falling within the detection array beam pattern.
- Azimuthal Heading estimate indicating the angular bearing of the detection sensor array.

These two detection outputs: the sensor being able to estimate angle of arrival, and the sensor tracking which direction it is scanning when an angle of arrival is estimated, drive the entire design of the system detection sensor. Two entire sub-systems of our device are dedicated to creating and managing these critical detection attributes. These sub-systems are described in detail in this work.

Sensor Deployment Scheme for the Facility Under Protection

The sensors in our system are deployed about the Facility Grounds such that the data from multiple sensor scans provide an overlapping coverage of RF activity. In practice, the extents and geometry of this overlapping coverage form the boundaries of the RF Exclusion Red Zone depicted in Figure 3.3. A hypothetical sensor deployment scheme is shown in Figure 3.4, where sensors are deployed at intervals on the exterior surface of the FUP. In a real implementation, sensors would be deployed on a building rooftop or on utility poles providing clear line-of-sight views of the grounds surrounding a facility. For an in-depth treatment of our recommended strategy for sensor deployment and system operation guidelines developed from extensive simulation runs, see our companion work [11].

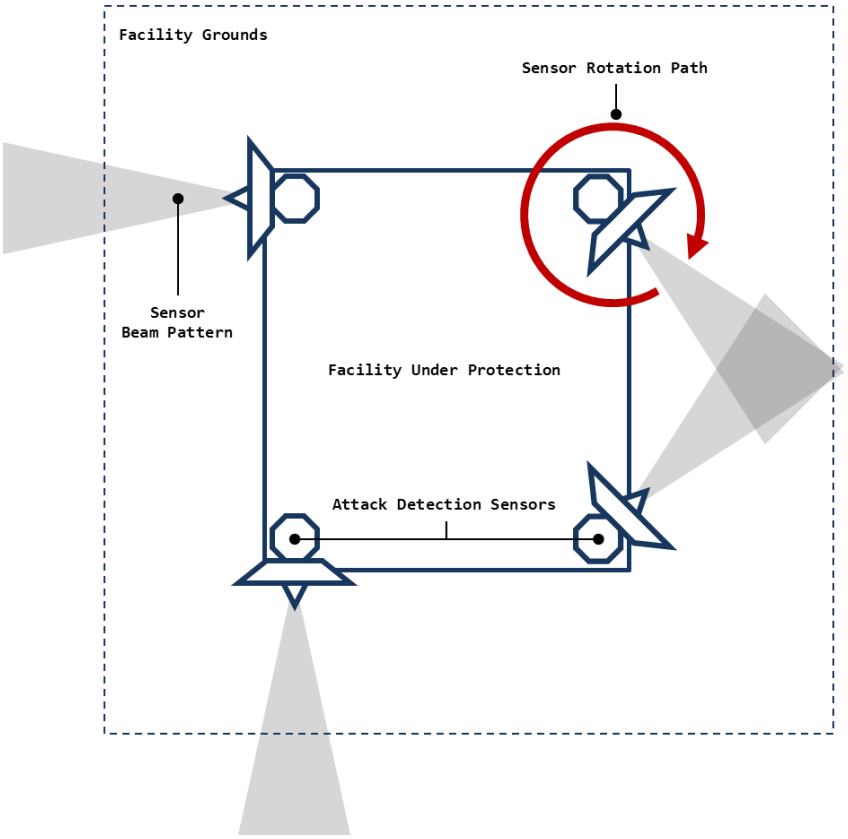

Figure 3.4. Detection Sensors Deployed on Facility Under Protection. The Detection Sensors can rotate 360 degrees to scan the RF environment external to the Facility Under Protection. The Scan Patterns of the Sensors have overlapping Fields of View. The locations of each sensor are known to the WIDAR Command and Control System.

Sensors are Mechanically Steered

The rightmost pair of sensors in Figure 3.4 illustrates how the sense patterns in our system should be arranged such that coverage overlaps to support the Red Zone designation capability of our system. Our present implementation features sensors operating mechanically steered antenna arrays. Each sensor is capable of continuous 360º rotation, with this rotational functionality also shown in Figure 3.4. In our prototype sensor implementation, continuous rotation is achieved through the use of a slip-ring connector joining the upper antenna chassis to the lower mechanical rotation assembly. Furthermore, a magnetic rotary encoder on each sensor is able to accurately track the

current azimuthal heading of the detection sensor, in the form of an angular bearing output. Each sensor periodically communicates current bearing data to a centralized Command and Control system – described in detail in a following section.

Sensor Movements are Camouflaged

In our design, deployed sensors are fitted with radome covers offering protection from the harshness of exposure to an outdoor environment, along with a degree of camouflage. By employing camouflage, we seek to obscure system capabilities and defensive posture by hiding whether a sensor is actively scanning, and therefore not revealing the exact directional bearing towards which the sensor array is currently scanning. By doing so, we seek to make using evasive physical counter-measures against our detection system more difficult for the adversary. An example of an evasive physical counter-measure would be an adversary conducting wireless communications only when she knows that a detection sensor is not actively oriented towards her position.

Detecting the Parking Lot Attack

Each detection sensor is capable of sensing RF activity within its area of control, and each device can sense and calculate an estimated Line-Of-Bearing (LOB) for any active RF emitters operating within range of the sensitivity limits of the device. A device accomplishes LOB detection using a mechanically steered antenna array, which is rotated through a scan pattern in search of anomalous RF signals. The scan capabilities and the RF sensing outputs from the antenna array attached to each device are described in detail later in this work.

Figure 3.5 illustrates two sensors in the process of detecting the Adversary as she executes the PLA. Unbeknownst to the Adversary, her active attack has been detected by

two independently scanning sensors tasked with protecting the facility. Each sensor is calculating and maintaining a lock on the LOB to the position of the RF device operated by the Adversary. The individual LOB estimates are being communicated to the Command and Control system where these data are integrated, resulting in a position fix estimate being made available to an on-duty Threat Analyst.

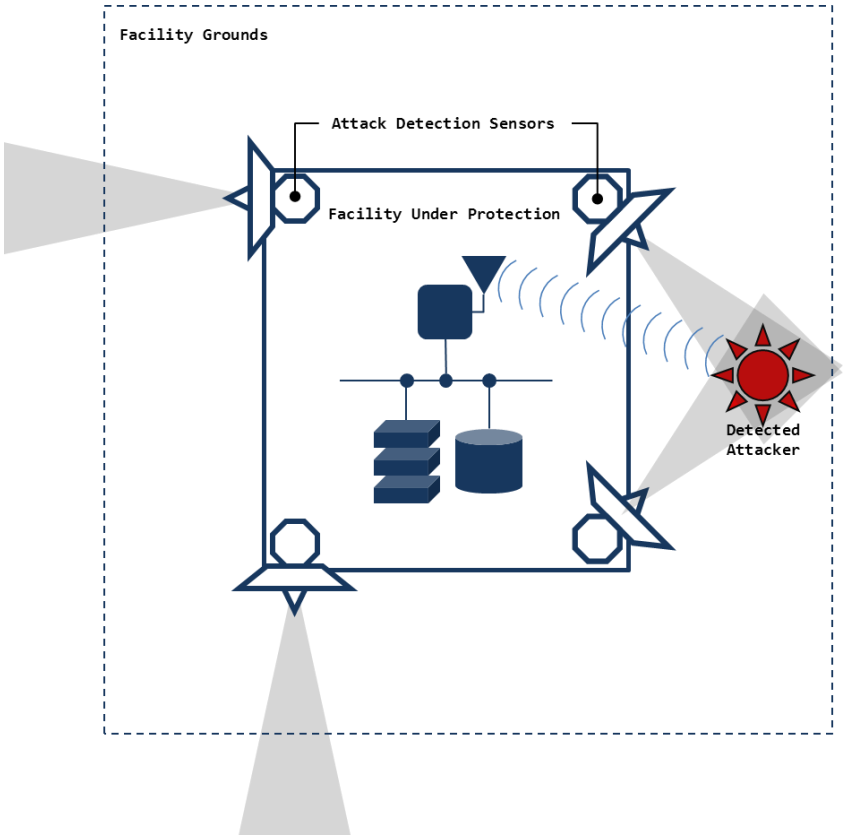

Figure 3.5. Sensors Detecting an Attack Attempt. Line-of-bearing data from two or more sensor angle estimates are integrated by a centralized Command and Control yielding a position estimate. System spatial attributes in turn feed functionality in higher-level Monitoring Systems such as a Red-Zone Intrusion Detection System which, when triggered, can direct a response team or activate other counter-measures.

Active Versus Passive Attack Detection

At this point, a distinction should be made regarding active versus passive attack detection. Our system of sensors detects only active RF targets operating within the beam pattern of the high gain antenna array deployed on each sensor, and as such, our system

does not have the capability to detect an Adversary who is only passively capturing RF traffic. For the purposes of this research, while many attacks originate using passive methods, we stipulate that the most damaging attacks will progress to some sort of active operations [12], and we assert that in order for our Adversary to really "get the goods" – reaching the sensitive systems and data of the FUP – she will need to switch to an active attack mode.

- Attacks which extract the most value from a target will require active RF emissions due to the need to utilize networking protocols with handshaking between target and attack hosts.

We feel this is a realistic requirement as many attack scenarios will require some usage of a connection oriented TCP/IP protocol based network service in order for an attack against a high value target to be carried out. For example, consider a NETBIOS/SMB connection to an LDAP/Active Directory File System Share, using TCP/IP over 802.11 WLAN frames.

Indeed, our threat model assumes that the Adversary utilizes some form of passive traffic analysis in order to discover the vulnerable Wireless Access Point of the FUP. However, it is also mentioned in our threat model description that passive pre-attack methods only yield reconnaissance intelligence, and do not provide the Adversary the desired access to the sensitive systems and data.

- Our system is not designed to detect or counter Passive Only Traffic Analysis Attacks.

- Our system assumes that an Attacker must utilize Active Transmission to perform any attacks against High Value FUP Sensitive System and Data Targets.

Sensor Command and Control System

Each detection sensor in our system communicates via a distribution system providing backplane connectivity to a centralized Command and Control (C2) segment. In our design, wireless communications from sensor to C2 system must be avoided, in order to remain out-of-band from the wireless communications network being monitored for protection; if the sensors utilized wireless communications they would interfere directly in the RF environment being monitored for threats and unnecessarily complicate our detection scheme. Therefore, it is a design requirement that communications with the C2 system be performed over a wired connection that is electronically separated from wireless activity of interest. As shown in Figure 3.4, there are multiple WIDAR devices deployed at strategic points along the perimeter of our model FUP, such that when more than one detection sensor triggers and collects LOB data for a would-be attacker these data can be integrated by C2 software to estimate a spatial location for an emitter, within some acceptable statistical error margin.

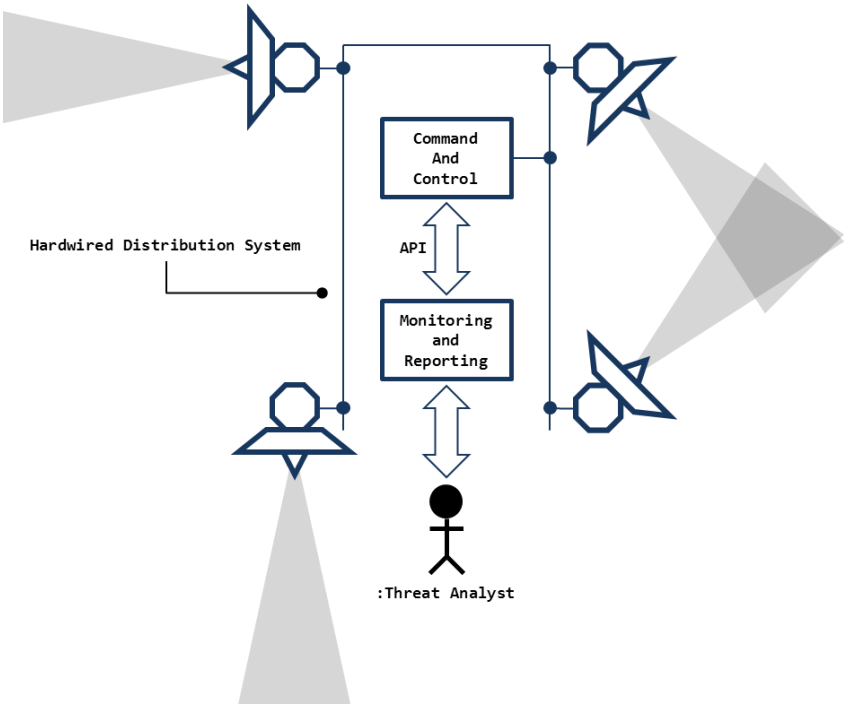

Figure 3.6. Command and Control System. Sensors are connected via a hardwired distribution system, such as Ethernet, to a centralized system featuring support for the higher-level Monitoring and Reporting System through an Application Programming Interface. A threat analyst is shown interacting with the Monitoring System directly, and with the C2 System via API calls built into the Monitor System. In this way, the threat analysts can receive Detection Alerts, as well as Configure Red Zone boundaries and other high-level system control parameters.

LOB estimates are integrated by C2 for input into a Target Location Processor (TLP) which outputs estimates of spatial attribution parameters. In our present system, the TLP utilizes a simple trigonometric position estimator, which triangulates a target location based on target LOB data, and a priori knowledge of system sensor deployment locations. Figure 3.7 depicts the simple triangulation scheme we are utilizing in the preliminary phases of our system testing. Poisel's Target Location Methods in Communications Electronic Warfare [13] details location estimators based around gradient descent least squares error minimization algorithms, which we plan to explore in a future paper centered on testing and improving the accuracy and performance of our WIDAR Sensor.

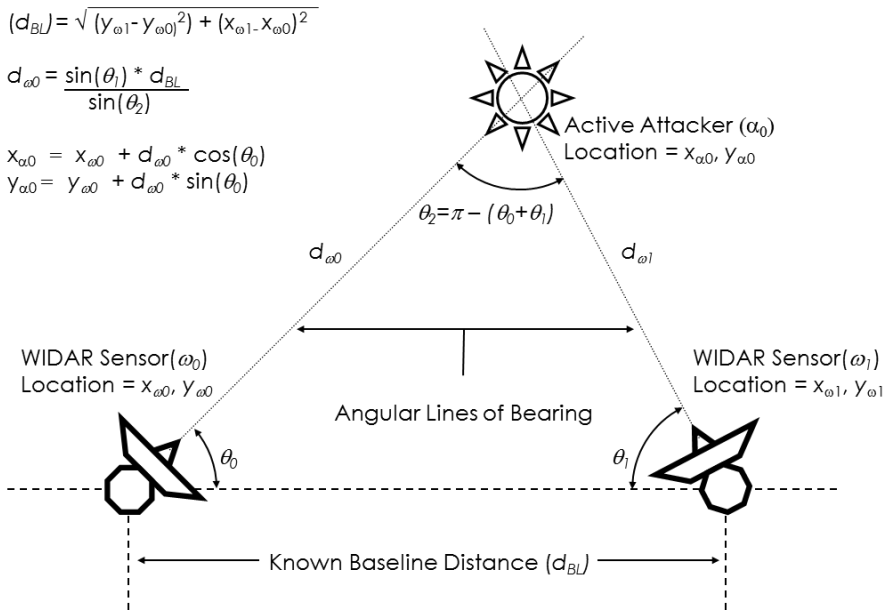

Figure 3.7. Simple Triangulation Model for Initial System Testing. This scheme operates on the assumption that the Command and Control system has a priori knowledge of the distance between sensors. Since sensors are statically deployed about the exterior of the Facility Under Protection, these metrics are easily obtained.

Monitoring and Reporting System

Potential detections are displayed to security analysts via visualization and reporting tools interfaced to the WIDAR C2 module via an extensible application programming interface (API). We show screen captures of two such Monitoring and Reporting System tools which were developed as prototypes during our research in Figure 3.8 and Figure 3.9.

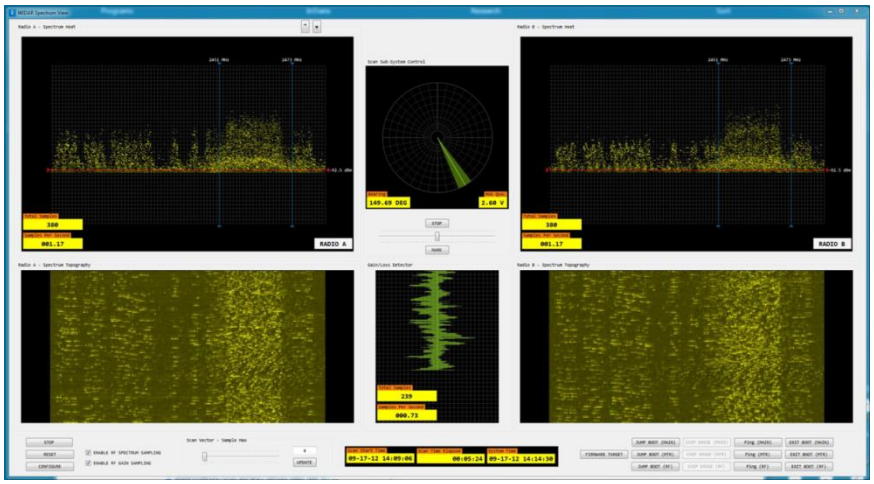

Figure 3.8. Monitoring and Reporting System – Prototype 1. In this configuration, RSSI data from the device antennas are displayed in separate visualization panels. The upper portion of each panel area displays spectrum usage as a function of channel power density, while the lower section displays the same density data, but with the change in density over time added on the abscissa axis. In the center of the display is the monopulse ratio of the antenna array.

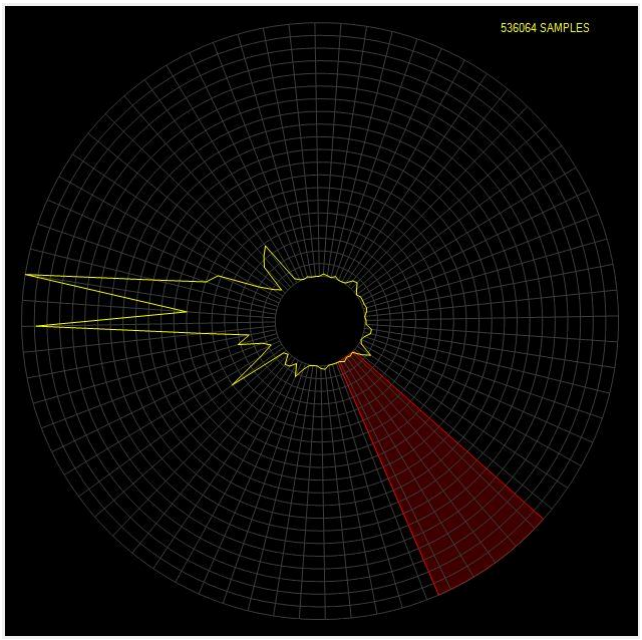

Figure 3.9. Monitoring and Reporting System – Prototype 2. In this configuration, RSSI data from both antennas of the device array are integrated into a common data model. The results are displayed as a polar plot of RSSI level for each bearing scanned. The screen capture above shows the current bearing of the device array using a red pattern – the device is currently scanning with a bearing of approximately 135° – with the plot lines in yellow also indicating strong RF activity detected when the device scanned near the 270° bearing.

WIDAR Detection Sensor System Architecture

At this point, we have presented the threat landscape for a hypothetical facility along with the capabilities and motivation of an adversary who is determined to launch an attack on a wireless network she had detected in operation at this facility. Our adversary has made a thorough assessment, and after much planning, she has concluded that for an attack to present the least amount of detection risk the attack should originate from a location on the facility grounds, but not actually inside the physical confines of the facility as it is too tightly secured for direct penetration. Tactically speaking, the Adversary has selected from her arsenal the Parking Lot Attack as the best weapon for accomplishing her mission objective: penetrating the sensitive systems and data of the facility.

We have also described our own secret weapon, which is a network of sensors deployed on and about the facility we are protecting, with the aim of detecting just such a class of attack. Due to clever camouflage techniques, the Adversary remains unaware that such a system is defending the FUP. We have also illustrated how the sensors are integrated, and communicate with a centralized Command and Control system. Each sensor has the capability to calculate a Line-of-Bearing for any RF Target detected within the sensor operational range. LOB's from multiple sensors can be combined to form a position fix by the Command and Control system. This system is capable of providing enhanced situational awareness to the security team tasked with protecting the sensitive systems and data of the facility, in terms of emitter detection and spatial attribution. A Threat Analyst can use situational decision support data to make an assessment if an attack is underway. The Threat Analyst may then choose whether to direct

countermeasures such as a security incident response team to further investigate and potentially interdict any attack.

Architectural Overview

We next focus on a detailed description of the design and capabilities of the sensor device forming the core of our WIDAR system, with the aim of explaining how a sensor implements Line-of-Bearing detection functionality. These devices constitute the backbone in our system for countering the Parking Lot Attack. Figure 3.10 shows an illustration of a sensor implementation we constructed to perform experiments during this research.

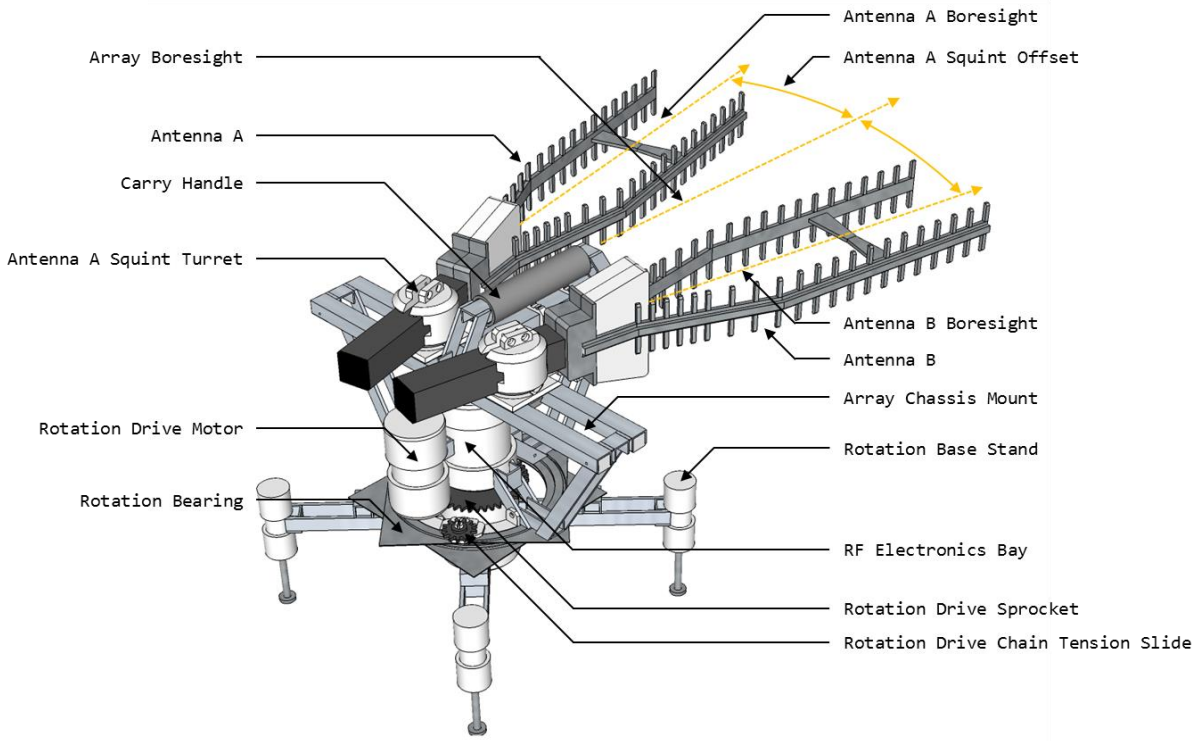

Figure 3.10. A second-generation implementation of our detection sensor. Visible are the two high-gain antennas of the array, and the rotation chassis. The rotation chassis is mounted to a fixed base, supported by four adjustable ground leveling pads.

Figure 3.11 presents the system block diagram for a typical WIDAR sensor. The two primary sub-systems responsible for the critical RF detection and sensor orientation detection attributes are highlighted on the right side of the block diagram. These two sub-systems communicate with a master microcontroller unit (MCU) via the slip-ring connector. The Master Control MCU, in turn, is interfaced to an Ethernet controller providing wired backbone communications between the sensor and the C2 system. There are two primary sub-systems responsible for RF Target Detection and Spatial Attribution:

- RF Sense Blocks – The sub-system tasked with sampling the device antenna array for RF Signals data.
- Array Orientation Blocks – The sub-system tasked with mechanically steering the antenna array during search. This includes functionality for sensing the angular bearing of the antenna array.

Both of these sub-systems appear as highlighted sections of Figure 3.11.

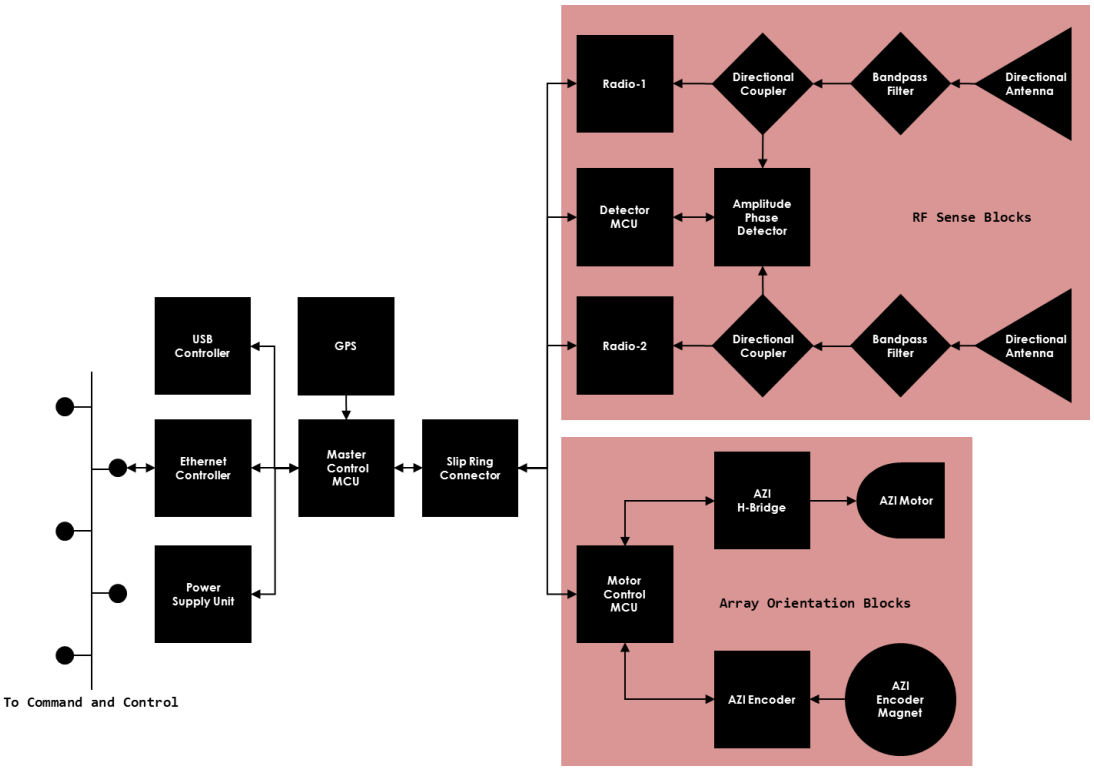

Figure 3.11. A Block Diagram of the Detection Sensor. The two sub-systems considered critical for RF detection and antenna array orientation tracking are background highlighted on the right side of the diagram.

Functionality for RF Target Detection and Position Estimation

Each detection sensor deployed as part of the WIDAR system contains a dual-antenna array for RF sensing. Two antennas are operated, borrowing simultaneous lobing techniques from Monopulse Radar systems. These antennas are operated as an antenna array with a combined beam pattern capable of yielding angular sense outputs that are more accurate and less susceptible to RF noise due to jitter than a system only operated using a single antenna. Detailed treatments on the theory and operation of monopulse radar systems can be found here [14] [15] [16]. We review the technique, from the context of our implementation in the following two sections.

The RF Sense-Array Sub-System and the Monopulse Concept

In a monopulse antenna array, power readings from multiple antennas are read by a comparator with ports fed by both antennas, to form a single ratio of the antenna gain present in each of the array elements. Phase ratios can be used as well, but more stringent array geometry is required to maintain phase coherence along the array boresight. The ratio, either from gain or phase, is called the *Monopulse Ratio* by many texts on the subject. While our system design and implementation is capable of measuring both the Gain and Phase monopulse ratios, it has not yet made use of phase, due to this more complex design requirement. A device measuring this ratio, called a monopulse processor, forms an output measurement that is purely a function of Angle-of-Arrival for an emitter detected in the two beams of the array antenna patterns.

Furthermore, the use of a monopulse array offers an advantage over less sophisticated emitter detection strategies. A naïve strategy sometimes encountered in systems attempting RF target tracking is where a single antenna is mechanically steered and the Received Signal Strength Indication (RSSI) is monitored during this sweep. A database of RSSI measurements is built during the sweep, which can be analyzed to estimate a peak RSSI following a complete sweep rotation, and when paired with a bearing measurement output from the mechanical steering sub-system, an LOB can be calculated. The monopulse literature calls this technique sequential lobing [15] [16].

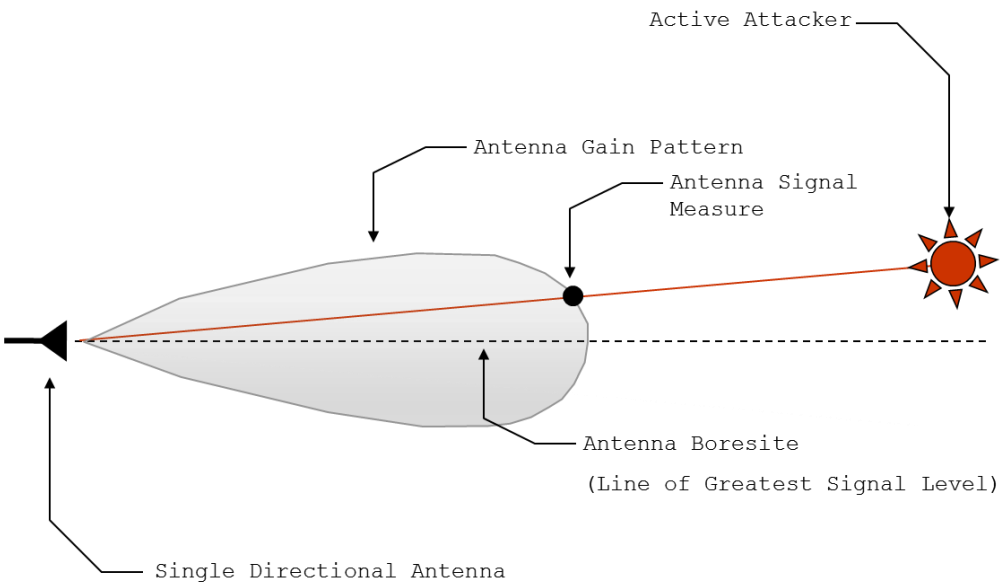

Figure 3.12. Sequential Lobing using a Single Directional Antenna.

The trouble with this method is that variations due to nonlinearities inherent to the way an RF signal propagates introduce too much unwanted noise in the RSSI measurements from sample to sample. Radar engineers use terms like scatter, scintillation, and glint to classify the sources of these nonlinearities [13] [16] and group them all into an effect called pulse-to-pulse jitter. Multipath reinforcement or cancellation of a signal also has a significant impact on the accuracy of a system using sequential lobing techniques. The end result of all of these variations is that an RSSI measurement could vary in a significant way from sample to sample obtained during a collection sequence. The effect of this noise in collection is that any database recording the peak measurement data becomes skewed by the uncertainty in whether a peak is the result of a real emitter RSSI measurement being detected or erroneous estimation.

In contrast, employing monopulse methods when lobing an emitter results in the effects of pulse-to-pulse jitter becoming common mode to both array antennas. By

forming the ratio of signal power entering each antenna, nonlinearities are assumed to affect the signal entering the beam pattern of each antenna in the same way, simultaneously. Thus, the ratio formed becomes purely a function of angle-of-arrival of the detected RF energy in the beam [15]. More importantly, each measurement of the gain or phase ratio, from the antenna array stands alone, yielding a line-of-bearing for an emitter detected in the array beam pattern. Indeed, the term monopulse derives from this fact, that a single (e.g. "mono") measurement of a radar pulse is sufficient to estimate an LOB for a detected target.

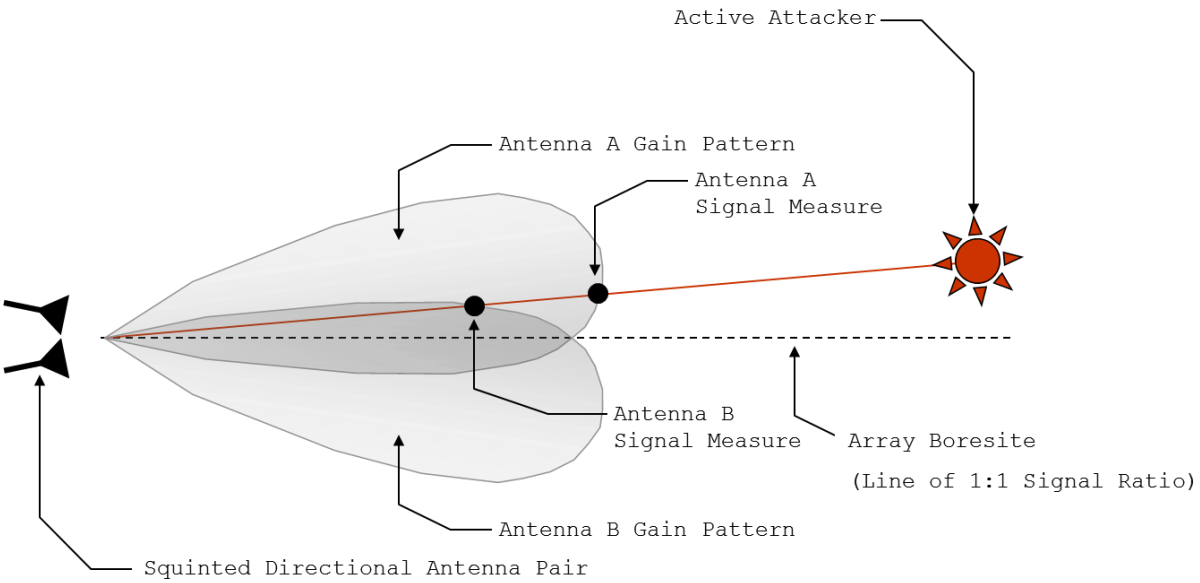

Figure 3.13. Simultaneous Lobing using the Monopulse Technique. A ratio is formed using gain data measured at Antenna A and Antenna B. Since both antennas are experiencing the same RF non-linearities due to measurement error, jitter, and multipath reflections, these errors become common mode, leaving the ratio formed purely the result of angle of the active transmitter from the array boresight. Each ratio measurement contains angle information, resulting in line-of-bearing estimate from only a single pulse of RF measurement, hence the name Monopulse.

A final note on monopulse: the ratio eliminates an ambiguity that arises during sequential lobing, namely that an increase or decrease in RSSI detected during sequential lobing does not yield any directionality parameter (left or right in azimuthal traverse)

since the RSSI parameter at that measurement point is a 1-dimensional data term. By contrast, since the monopulse lobing measurement is formed from a ratio of the gain difference between two separate antenna beam patterns, the ratio provides a directional component with each measurement. For example, if Antenna A has higher gain, and Antenna B has lower gain, the target must be more directly in the pattern of Antenna A, so the target should lie to the left of the antenna array boresight, and the monopulse ratio would be a negative result – assuming the ratio is normalized to zero when a target is directly on boresight of the monopulse array. In mathematical terms, the monopulse lobing ratio parameter is represented as an odd function of transmitter angle-of-arrival at the antenna array.

The Monopulse Processor

There are various accepted methods of monopulse processor implementation in the present state of practice. Most use differing forms of microwave signal combiners such as hybrid rings or junctions to produce simultaneous sum and difference outputs for a given antenna feed network. Most tracking radars implement the combiner network using waveguide technology [16], but there are research papers that describe transmission networks implemented in stripline and microstrip for low power microwave monopulse systems [17] [18].

The front-end feed network containing sum and difference ports is typically down-sampled by mixing with an intermediate frequency VCO, followed by the voltage (not power) being digitized by a pair of base-band ADC units. The ADC pair is a peripheral of an MCU or FPGA, which can then provide the sum and difference channel outputs to a signal processing computer for range and angle-of-arrival analysis.

The WIDAR Monopulse Processor Implementation

To minimize cost and implementation complexity, our system follows a slightly different approach to monopulse processor implementation than most conventional systems. Our feed network is composed of a bandpass filter fed PCB signal trace entering a discrete integrated circuit. The Analog Devices AD8302 Gain-Phase Comparator [19] features dual matched logarithmic amplifier networks. While not as commonplace as microwave combiner networks, *Monopulse Principles and Techniques* does present at least one design utilizing a logarithmic amplifier comparator as a method for implementing a monopulse processor [16].

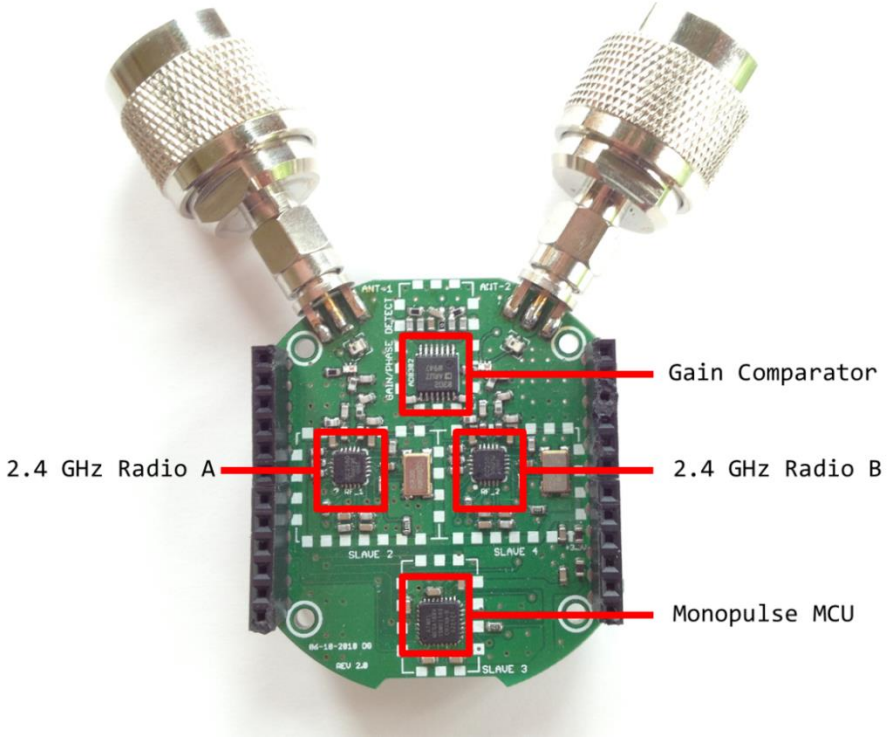

Figure 3.14. WIDAR Device Monopulse Processor Board. At the top is the AD8302 Gain/Phase Comparator. Dual 2.4 GHz radios are available for spectrum sampling and received signal strength applications. An Atmel ATmega 168 MCU is located near the bottom of the PCB. This MCU can configure the onboard radios, but is primarily tasked with sampling the AD8302 Gain Comparator using its built-in hardware ADC. Two BNC to RP-SMA connectors are connected to the PCB, allowing the board to be interfaced with the chassis 2.4 GHz antenna array.

Preliminary Monopulse Processor Data

Figure 3.15 depicts a command and control software tool we created to visualize the monopulse ratio data output of the detection sensor. The vertical white line in the center of the plot shows where the monopulse ratio is balanced between both the left and right array antennas. The figure shows four complete rotations of the antenna array chassis. Near the bottom of the plot, there was not an active RF emitter operating in the area, so the ratio remains closely aligned with the vertical central axis.

An active RF Target was positioned so that it would be lobed by the monopulse array. This can be seen in the top half of the signal plot, where monopulse ratio data are shown for two complete rotations of the sensor while the emitter target was active. When the emitter is in the array beamwidth, the monopulse ratio signal first peaks towards one antenna of the array, followed by peaking on the other array antenna.

Figure 3.15. A Gain Plot collected from two rotations of Device Antenna Array. Two full rotations of the device antenna array are shown. Two sinusoidal waveforms shown near the top half of the display indicate that an active emitter is detected in the array beam pattern. The RF Target is on boresight when the waveforms cross the center axis – which is the point when the antenna gains from each array element are detected as being equal. A 256 point moving average filter is implemented to assist waveform smoothing.

The Antenna Array Azimuth Orientation Control Sub-System

The Antenna Array Azimuth Orientation Control Sub-System is responsible for mechanically steering the antenna array to permit RF Sensing along a desired scan bearing. The sub-system accomplishes this task through two key functional blocks:

- Mechanical Rotation using DC Gearhead Motor Drive under software control.
- Position Sense using highly sensitive Magnetic Encoders.

Mechanical Rotation System for Array Azimuth Positioning

Physical positioning of the antenna array is controlled by the Motor Control Sub-System shown in Figure 3.11. This sub-system is implemented by a PCB featuring the STM Microelectronics L298 H-Bridge driver. This driver has the capability to drive up to two DC gearhead motors at up to 2amps continuous output, with 3A peak output capabilities. Our chassis integrates a single DC gearhead motor to drive the rotation of the antenna array. This motor typically draws only 60 mA of current during position change operations. The driver is also capable of directional control – forward and reverse – of the motor, along with an enable input pin, that can be switched on and off using PWM for precise speed control.

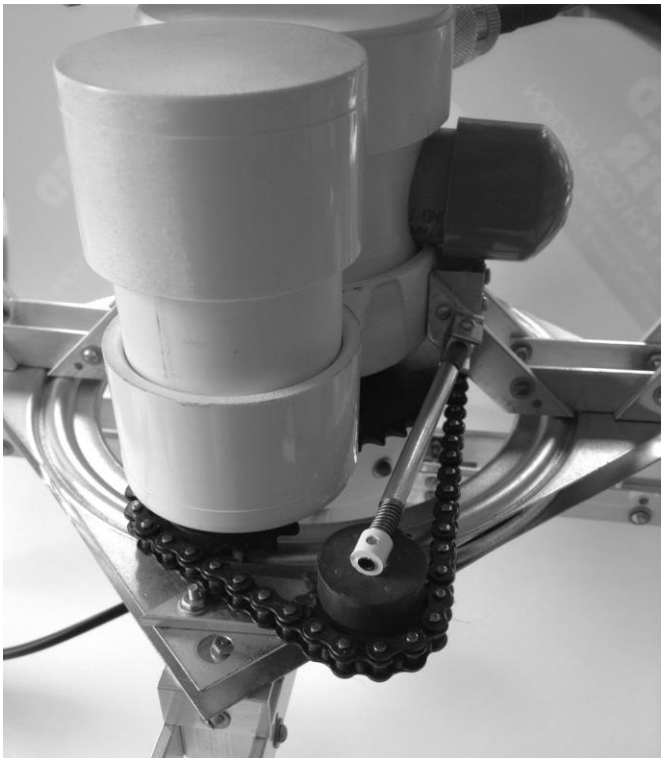

Figure 3.16. A close up shot of the motor and mechanical assembly on the sensor. The DC gearhead motor is housed in the PVC enclosure seen in the top left. A drive sprocket attached to the motor rotates a chain attached to the larger sprocket affixed to the chassis base. Also visible is a tension system added to maintain proper chain tension without over-stressing the drive train.

The PCB completes the mechanical positioning sub-system by integrating a control microprocessor – the Atmel ATmega 168 MCU - for sourcing PWM speed control signals, and responsible for direction and speed selection of the antenna array. This MCU receives control instructions from the Master Control MCU. The Master Control MCU can control the array position directly, or relay higher level array positioning commands from the C2 sub-system coming off the wired distribution system of the device.

Magnetic Rotational Encoder for Array Azimuth Position Sensing

The WIDAR device utilizes a chassis mounted magnetic rotational encoder to sense azimuthal rotation bearing. Our system implementation uses the Austrian Microsystems AS5306 mag encoder with quadrature A,B, and INDEX pulses typically found on most incremental encoder hardware [20]. Our design operates the AS5306 encoder in conjunction with the MR-12-72 ring magnet, featuring 72 magnets arrayed in a back-to-back pole configuration. The AS5306 IC is capable of sensing the magnetic pole change at each magnet junction, as well as providing interpolated encoder pulses based on reading the magnetic field change between junctions. This combination of pole change indexing and sensed interpolation yields 5760 encoder ticks per revolution or a positional accuracy of .0625 degrees. A photograph of the encoder mounted to the chassis of our sensor implementation is shown in Figure 3.17.

Figure 3.17. A close up shot of the encoder mounted just above a magnetic ring installed atop the main drive sprocket. The encoder package is mounted in a small custom PCB and the PCB is mounted on a spring loaded floating support cantilevered from a mount point located beneath the main chassis housing. The spring loaded support maintains the encoder sensor in proper position floating .5mm above the magnetic ring, and also allows for any rotational misalignment to be tolerated between the upper and lower chassis housings.

As shown in Figure 3.11, the encoder is interfaced to the Motor Control MCU of our WIDAR system, and each tick is processed by an interrupt service routine (ISR) run on the MCU. The ISR, when triggered, can read the quadrature values on the A and B encoder pins to determine the chassis azimuth rotation direction, and in turn, increment or decrement the internal position state counter block.

The INDEX pulse output by the encoder is also utilized to detect missed encoder pulses (missed due to slow interrupt handling or mag-sensor noise). Misses, if undetected introduce drift into the encoder state counter block. The INDEX pulse is triggered every 160 channel pulses, and when the INDEX pin is high on the MCU, the channel pulse delta from a previous INDEX pin reading can be calculated to ensure that the channel pulse total is 160. If the sum is not the expected value, a drift correction offset can be determined using the channel pulse total subtracted from the expected value.

To complement the INDEX pulse, the AS5306 encoder hardware outputs an analog magnetic field strength indication (MFI) on a pin that is interfaced to one of the Motor Control MCU 10-bit analog digital converters. Currently, no closed-loop control processing includes this reading, but C2 user interface software does display the voltage reading to an analyst, which proved useful to detect a problem encountered during bench testing when metal filings debris became attached to the ring magnet. The voltage dropped below the threshold recommended by the datasheet and INDEX pulses showed missing position channel pulses whenever the encoder head passed over a specific region of the magnet. This region contained metal filings, which when removed, solved the missing INDEX pulse problem.

Our system also includes provisions for the Motor Control MCU to send an error packet indicating the drift detection for processing by higher level C2 software. After some initial tuning and some debris removal from the ring magnet, our implementation encoder has not exhibited any operational drift, and tests using the encoder values to park the antenna array at a home position have shown the encoder to be highly accurate. Figure 3.18 shows a logic analyzer screen capture of the encoder INDEX and A/B channels.

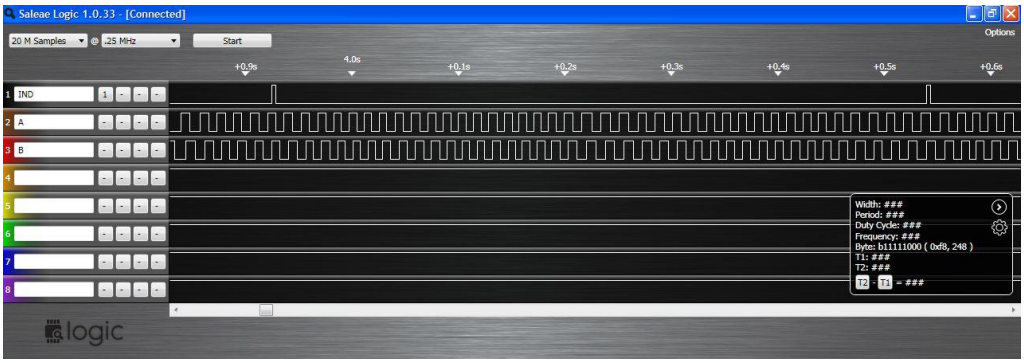

Figure 3.18. Encoder pins captured on logic analyzer. In this capture the A and B channels can be seen to be 90° out of phase with each other. This quadrature shift is used to determine the direction of rotation when the array is mechanically steered. The INDEX pulse can also be seen triggering after each succession of 40 A/B pulses. Counting encoder ticks and analyzing whether 40 ticks were received after each INDEX tick can be used to error check the encoder for missed detections, which would introduce drift errors.

In our present implementation, the main routine running on the Motor Control MCU reads the current position state counter block for packaging inside an encoder position packet for transmit to the Main Control MCU. Each encoder position packet contains the current azimuth position, and the current rotation direction. This packet sending operation is performed continuously on the Motor Control MCU, with a new packet being readied whenever the previous encoder position packet is finished transmitting.

Visualizing Array Position – A Screen from the Monitoring and Reporting System

Figure 3.19 shows the C2 software user interface which displays array traverse position to an analyst. In the prototype version shown, the user interface also has a slider providing the analyst manual forward and reverse control of the antenna array traverse position. An emergency stop button is also provided to halt all motor positioning activity if necessary.

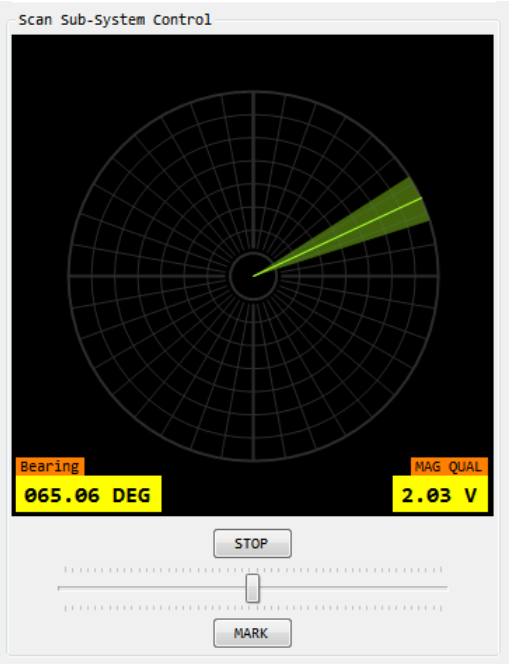

Figure 3.19. The Monitoring and Reporting System view of antenna array azimuthal bearing. In this screen capture the array has an orientation of 65.06° – indicated by the green shaded graphical sweep displayed on top the polar grid, and the bearing readout at the lower left. Also visible are a slider control enabling the Threat Analyst to manually steer the sensor antenna array – this feature is useful for sensor calibration and for fine grained control of the array when manually combing a Red Zone area. A reading showing the strength of the magnetic field in proximity to the encoder ring magnet is also displayed in the lower right of this GUI.

Recommendations for Improved Array Position Control on Future Implementations

Our present implementation operates using a simple open-loop position control routine. This routine is a simple counter which ramps the h-bridge controlling motor speed from zero to the peak speed value, and then back down again to zero. While this routine is effective in scanning the device position array, and we were able to tune it to achieve a desired step arc length, better scan capability could be achieved with a two stage closed loop PID controller. We recommend that any future implementation integrate PID control of velocity with a second PID stage controlling position.

This capability would enable precision scan tracking – for example the array could be instructed to scan a specific pattern. Also homing – the ability of the controller

to accurately position the array along a specific boresight vector – would be a useful capability enabled by these enhancements. One possible scenario using this behavior is the designation of scan zones for the device array, where the device sweeps a pre-selected sector instead of performing a complete rotation. Another scenario would be the precision tracking of a detected target that is moving.

Minor Sub-Systems

We briefly cover, for completeness, several of the minor sub-systems which provide operational support for the sensor.

The Slip Ring – Enabling Continuous Rotation Capability

As previously mentioned, the WIDAR sensor is capable of 360º continuous rotation scanning. This is mechanically accomplished through the use of the 12 channel Keyo KYC-12A slip ring connector, shown in Figure 3.20. The use of a slip ring facilitates a simpler control scheme, avoiding the use of mechanical limit switches that would be needed if our system only supported a partial rotation range. We incorporated continuous rotation into our sensor implementation following lessons learned from our first generation chassis. The prior implementation of the chassis lacked slip ring support and suffered from a limited range of scan, required mechanical and electronic limiting in the design to prevent chassis damage in the event of over-steer, and the wiring harness cabling the antenna array in this implementation was much more complex due to the increased range of travel and mechanical stress.

Design Flaws Inhibiting Slip Ring Communications Signals

In our present device implementation, the two sub-systems shown shaded in Figure 3.11 both communicate with the Master Control MCU and Power Service Units

via the 12 slip ring channels. This includes both control and 12V power signals for the chassis steering motor. The control signals are high speed SPI signals which are very susceptible to clock interference and cross talk. Exasperating this problem is the fact that the slip ring channels are not shielded from one another.

This choice of design has introduced many challenges in our operational testing. The SPI bus communications fails whenever the chassis position motor is engaged, even at the lowest data rate, and with the addition of filtering circuitry on all clock and bus lines. Failure occurs in the form of SPI calls which poll the RF Target Detection Sub-System. Calls to read registers on RF sub-system hardware become corrupted in transit. Since this hardware uses a communications protocol that does not implement foreword error correction or command packet checksum, corrupted commands are interpreted by the hardware as other configuration instructions, which lead to an invalid state of operation. This failure robs our implementation of a true track-while-scan capability, a key requirement if our tracking system is to have the capability to rapidly detect and spatially attribute attacking RF Targets.

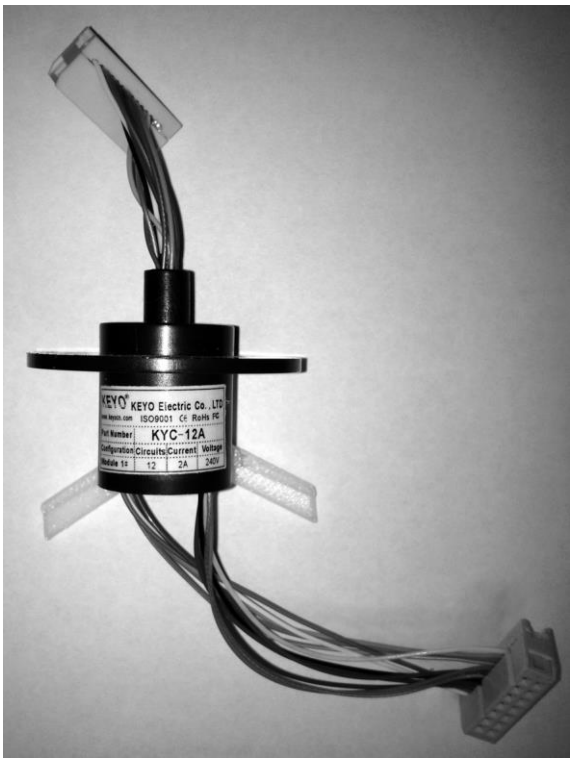

Figure 3.20. The Keyo KYC-12A slip ring connector supporting continuous 360° rotation of 12 control channels. This is an inexpensive connector sourced directly from a Chinese distributor. The ability to continuously rotate the sensor antenna array supports a simpler control scheme; one that does not require limit switches or a wiring harness with cabling supporting a large amount of lateral travel. However, noise is an issue when using this connector for high speed digital signals. Also, we attempted to send analog power signals side-by-side with these digital signals – with some isolation practices employed. This proved unworkable during use cases attempting continuous track-while-scan operations. To counter noise issues on this prototype, the mechanical steering sub-system is never operated in unison with the RF sense sub-system. This permitted research tests, but would not be practical in a real deployment situation.

For research data collection purposes, we settled on a fix, with the solution being to disable RF sense operations while the motor is positioning the antenna array. A solution we call the “Pan-then-Scan” mode of operation. This permits all device functionality to be tested in an operational environment, but at the expense of much reduced scanning speed. This solution is only a remedy for data collection purposes, enabling us to collect real field data in an operational setting.

As such, we did not find this speed reduction to be a barrier to our research testing since our RF Targets were static and in known positions. However an attacker with

knowledge of this reduction in scan speed could easily exploit such an implementation flaw, in effect creating blind spots in our Red Zone coverage. Any next-generation implementation will need to resolve this issue, allowing true Track-While-Scan capability. We recommend the following changes to better facilitate a Track-While-Scan capability on any future implementation:

- Eliminate any Analog Power Signals on the upper (rotating) part of the device chassis.
- It follows then that the entire Antenna Array Orientation Sub-System should then be relocated to the lower (fixed) section of the device frame.
- Do not attempt high-speed SPI communications over an unshielded slip ring network. Choose a scheme employing differential signals – such as RS-485 – which will yield better communications reliability and permit a higher sampling rate given adequate transceiver selection.
- Introduce hardware capable of verifying the integrity of the digital communications signals passing over the slip ring. Employ a simplex scheme capable of CRC or FEC.

Communications Interface to Wired Distribution System

A dedicated Serial-to-Ethernet Controller serves as the communications interface between the Command and Control system and the Master Control MCU. Figure 3.21 shows the Lantronix XPORT controller [21], which features buffered IO, and easy TCP/IP configuration via an onboard web server.

Figure 3.21. The XPORT 100mbps Serial-to-Ethernet Bridge. Limitations in the internal resonator clocking the master MCU keep serial communications speeds between MCU and bridge to a modest 38,400 BAUD.

GPS Sub-System

GPS hardware interfaced to the master MCU provides geo-location functionality enabling each sensor to accurately determine spatial location. The coordinates of the sensor location when combined with the coordinates of any other sensor enable the calculation of the Known Baseline Distance parameter, feeding the triangulation estimates outlined in the calculations presented in Figure 3.7.

Power Service Unit

All power for the device is provided via the Power Service Unit (PSU) shown in Figure 3.11. The PSU consists of a single PCB integrating discrete packaged switch mode power supplies sourcing 3.3V and 5.0 V for the digital electronics of the WIDAR device. The PSU also provides an analog pass through of the 12 input supply serving as the source input to the motor control sub-system.

Other Specialized Functionality of Interest in this Research

The device we constructed for our experiments is a second-generation platform developed with the intention of RF target spatial attribution in the 2.4 GHz band. In this section we touch on several implementation features which we have discovered to be useful in experiment design and platform configuration flexibility.

Generic Packet Communications Stack for All System Controllers

All MCUs in our design are AVR controllers from Atmel – ATmega 168 processors to be specific. The datasheet for this line of micros is available here [22]. Firmware for these systems is developed using the Atmel Studio 5 IDE with initial programming handled by the AVRISP-MKII in system programmer. The ISP is only used for an initial burn of a custom and very lightweight bootloader onto each MCU. All subsequent programming is performed over the Ethernet-to-Serial-to-SPI packet communications sub-system.

Making Firmware Changes – Using a Bootloader for In-System-Programming over Ethernet

As previously noted in our architecture description, there are three separate Atmel AVR ATmega 168 microcontrollers operating within each WIDAR sensor environment. Our initial implementation utilized an Atmel MK II AVR programmer to program each MCU. This required the programmer to be directly connected to the In System Programming (ISP) port on each system PCB board when firmware program changes were required (and there were many changes required and many still ongoing during experiments with our system). The device chassis design had access points intended to support direct connection of the ISP programmer, with clock connection routed to the target MCU by means of a rotary encoder. During our bench tests, the high speed clock would not operate properly over the mechanical rotary encoder, which resulted in failed EEPROM updates. Furthermore, sometimes, an MCU on the shared programming bus (and not an intended firmware update target) would cease operating due to EEPROM corruption following the programming of another target MCU on the shared bus. We

believe this was due to undesired crosstalk of the clock signal into other MCU clock lines.

The failure of the ISP ports accessible on the chassis exterior mandated that we then had to disassemble the device chassis in order to access the ISP port on each system PCB directly. This was an unacceptable and painstaking process which needed to be remedied with an engineering design change. The solution we arrived at was to pre-program each MCU with a small footprint (e.g. less than 1024 kB) bootloader capable of loading firmware object code into MCU EEPROM via serial or SPI communications, rather than the dedicated device programmer hardware (the device programmer was still required to burn in the initial bootloader image).

The use of bootloaders on each system MCU has enabled rapid reprogramming of any MCU firmware of the WIDAR device without the need for disassembly or connection of a dedicated firmware programmer. Figure 3.22 shows a screen shot of the device firmware update control software user interface. Firmware update control software addresses updates to each MCU bootloader as a node on the main system bus, with each MCU made reachable over either of the Ethernet to Serial, or Serial to SPI communications bridges.

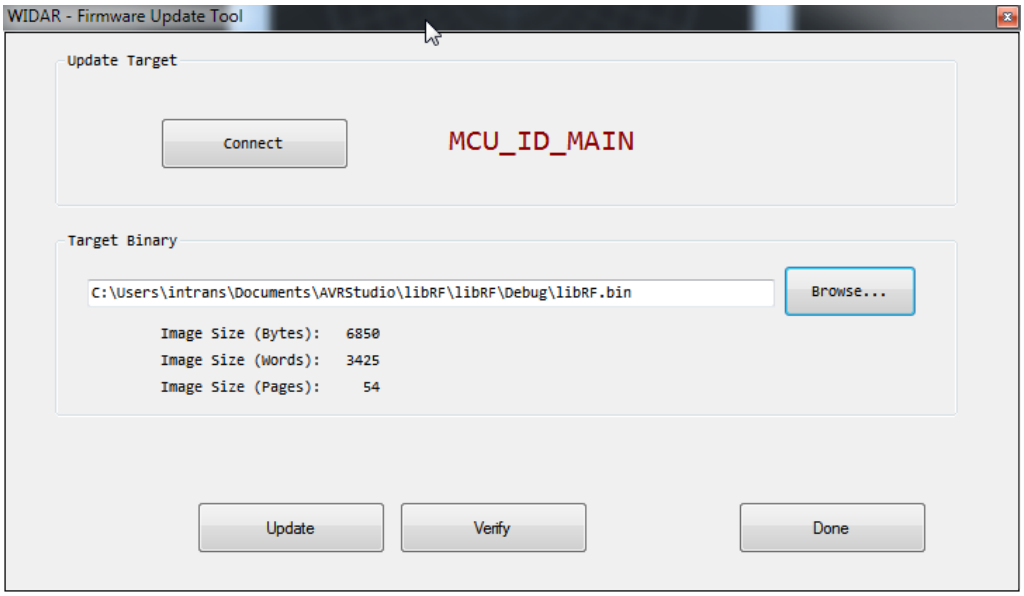

Figure 3.22. Device Firmware Update Tool. A binary file is produced by the MCU Development IDE, and once compiled, can be uploaded to any of the four MCUs on the device over an Ethernet connection between MCU and the device C2 system.

Spectrum Analysis using System on Chip Radios

Our device is also equipped with a pair of 2.4 GHz system on chip radios featuring programmable channel bandwidths. A PCB in the RF Sense sub-system hosts the two radios, monolithic CC2500 radio-on-chip packages from Texas Instruments [23]. These packages are surface mount, and require very little external hardware, except for an antenna matching network. Each of these radios is interfaced to the common SPI bus connecting all device peripherals to the Master Control MCU.

While the CC2500 is a full transceiver featuring several configurable modulation/demodulation schemes, we utilize only the RX mode of operation and the Receive Signal Strength Indicator (RSSI) estimation functionality onboard the radio, which serves our device as a low-cost spectrum analyzer. The two radios embedded in our system each support spectrum sampling from one antenna in the device array.

This spectrum sampling capability is useful in making determinations of the modulation scheme in use by a detected RF target – for example it is easy to distinguish a target using Bluetooth versus a target employing 802.11 protocols. Furthermore, when detecting WiFi traffic, we can easily make the determination if the target is in a channel band that is the same as that in use by the FUP and is a direct threat to Wireless Access Point attacks.

Visualization of Channel RSSI Outputs

The primary purpose of integrating system on chip radios into our system is to provide spectrum visualization services to further enhance the RF situation awareness capability offered by the device to Threat Analysts. Figure 3.23 shows a screen capture of our Monitoring and Reporting tool actively sensing 2.4 GHz activity and providing visual spectrum utilization data. Activity in 802.11 channel 11 is clearly visible in the spectrum diagram. Additional RF activity is also shown as a series of peaked narrowband channelized waveforms in the display.

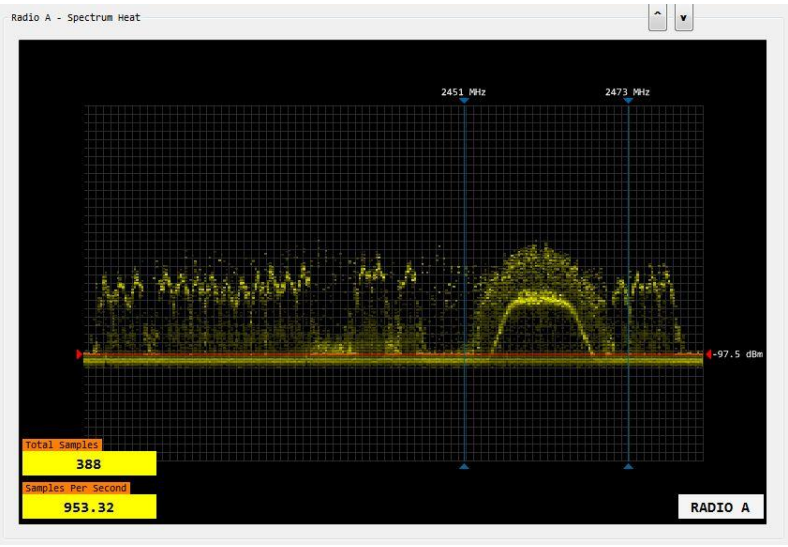

Figure 3.23. Spectrum Heat Plot from 2.4 GHz Radio Sub-System of WIDAR Device. The device was not actively rotating when data were collected for this Heat Plot.

To complement the visualization functionality shown in Figure 3.23, our system features one additional tool designed for visualizing spectrum samples over time. Called a “waterfall” or “topographic” plot by similar software tools [24], this tool depicts a top-down look at the spectrum peak values, with time plotted on the ordinal axis. A plot of the spectrum visualized in this manner is shown in Figure 3.24. This plot is visualizing the same spectrum samples as those collected and displayed in Figure 3.23. In this diagram, past heavy usage of the 802.11 channel 11 can clearly be seen on the right side. Intermittent activity in the other parts of the spectrum can be seen as small dashed patterns in the plot.

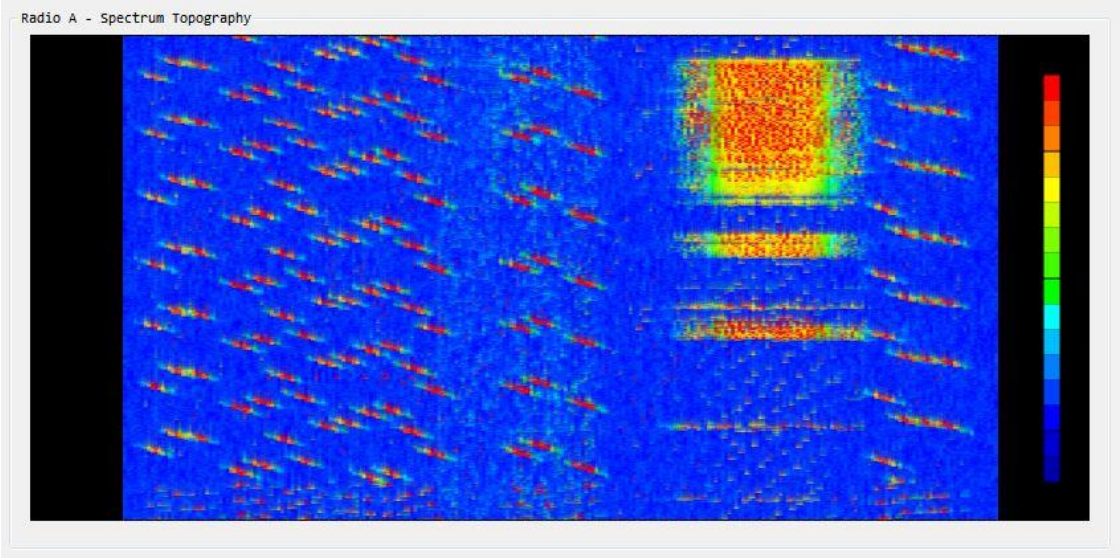

Figure 3.24. Spectrum Topographical Plot from 2.4 GHz Radio Sub-System of WIDAR device.

Spectrum Sample Metadata – Enhancing Visualization Using Duty Cycle Metrics

The channel duty cycle is defined as the number of times the sampled channel is above some predefined RSSI threshold, divided by the total number of channel samples (equ. 1). In an RF spectrum sampling context, channel duty can be also referred to as Channel Utilization. Typical duty thresholds are set 10 dB to 20 dB above the spectrum

noise floor [25]. In our system, the Channel Sample Configuration GUI provides a front-end permitting the selection of the Duty Threshold. Careful selection of this parameter can eliminate spurious detected RSSI values, while emphasizing higher power and more frequent utilization of a sampled RF channel.

For visualization applications, weighing the duty cycle typically involves displaying the channel power levels with higher measured Channel Duty (and hence, heavier channel utilization) as more pronounced or bolder color schemes.

$$Channel\ Duty = \frac{Channel\ Samples\ Exceeding\ RSSI\ Threshold}{Total\ Channel\ Samples} \quad (1)$$

It is also common to aggregate channel samples by either frequency or duty cycle into a history data structure typically called a heat map, which organizes the spectrum into channels and further sub-divides each channel into discrete cells, one for each value of power level quantization. Each of these cells maintains a count of the number of samples that are identified as belonging to that power level. Again, each cell can assign a weight based on Duty Cycle to these sample counts in order to draw further emphasis towards cells with higher channel utilization. In Figure 3.24 a WLAN operating in 802.11 Channel 11 (spanning 2451MHz to 2473MHz) can be seen to have a higher Channel Duty than other spectrum activity, thus giving the obtained spectrum power density samples more weight in those channels, making the WLAN density plot appear bolder and more pronounced in the visualization tool.

Concluding Remarks

We have presented a system designed to detect and counter the Parking Lot Attack. Our system utilizes a distributed network of sensors to secure a region surrounding a facility where RF activity is monitored with a spatial attribution capability. Our system design detects and locates any active attempts to connect with wireless infrastructure in use inside the facility. Countermeasures can be directed towards locations of interest by threat analyst receiving notifications from our system.

A key feature of our system is the use of monopulse radar methods to aide in RF transmitter geo-location. To test the capabilities of this feature, we constructed a real implementation of the sensor in our design. We provided an in-depth treatment of the sensor placement strategy for facility protection, along with hardware details of our implementation. We also provided recommendation for improvements in our system and device design, stemming directly from lessons learned during operational tests of our sensor.

To compliment this work, which is focused on an architectural description of our system, we intend to release a follow on performance-focused work, concentrating on analyzing the detection and spatial attribution accuracy of our device implementation. Preliminary bench tests show promise when line of sight measurements are taken from an RF Target actively transmitting and detected in the lobes of the device antenna array.

References

[1] SeatlePI, "Feds: Wi-Fi hacking burglars targeted dozens of Seattle-area businesses," 19 09 2011. [Online]. Available: http://www.seattlepi.com/local/article/Feds-Wi-Fi-hacking-burglars-targeted-dozens-of-2178421.php#ixzz1jwWj2LWH. [Accessed 08 2013].

[2] The Seattle Times, "High-tech hacker gets almost 8 years in $3M Seattle theft ring.," 13 07 2012. [Online]. Available: http://seattletimes.com/html/localnews/2018684238_hacker14m.html. [Accessed 08 2013].

[3] Cisco, "Cisco Mobility Services Engine," [Online]. Available: http://www.cisco.com/en/US/products/ps9742/index.html. [Accessed 08 2013].

[4] V. Bhargava and M. L. Sichitiu, "Physical Authentication through Localization in," in *IEEE GLOBECOM*, 2005.

[5] J. Werb and C. Lanzl, "Designing a positioning system for finding things and people indoors," *IEEE Spectrum*, vol. 35, no. 9, pp. 71-78, 1998.

[6] W. A. Arbaugh, N. Shankar, Y. Wan and K. Zhang, "Your 802.11 wireless network has no clothes.," *IEEE Wireless Communications*, vol. 9, no. 6, pp. 44-51, 2002.

[7] F. T. Sheldon, J. M. Weber, S.-M. Yoo and W. D. Pan, "The Insecurity of Wireless Networks," *IEEE Security & Privacy*, vol. 10, no. 4, pp. 54-61, 2012.

[8] S. Fluher, I. Mantin and A. Shamir, "Weaknesses in the Key Scheduling Algorithm of RC4," in *Proceedings of the 4th Annual Workshop on Selected Areas of Cryptography*, 2001.

[9] AirCrack-NG, "AirCrack-NG," [Online]. Available: http://www.aircrack-ng.org/. [Accessed 08 2013].

[10] E. Tews and M. Beck, "Practical attacks against WEP and WPA.," in *ACM Conference on Wireless Network Security (WiSec)* , 2009.

[11] D. J. Gieseman and T. E. Daniels, "A Strategy for Facility Wireless Attack Detection using Cooperative Mechanically-Steered RF Detection Sensors," Iowa State University, Ames, IA, 2015.

[12] US Department of Justice, "Retail Hacking Ring Charged for Stealing and Distributing Credit and Debit Card Numbers from Major U.S. Retailers.," 05 08 2008. [Online]. Available: http://www.justice.gov/opa/pr/2008/August/08-ag-689.html. [Accessed 08 2013].

[13] R. A. Poisel, Electronic Warfare Target Location Methods, Artech House, 2005.

[14] A. I. Leonov, K. I. Fomichev, W. F. Barton and D. K. Barton, Monopulse Radar, Artech House, 1986.

[15] D. R. Rhodes, Introduction to Monopulse, Artech House, 1980.

[16] S. M. Sherman and D. K. Barton, Monopulse Principles and Techniques, Artech House, 2011.

[17] L. H. Smit, "The analysis and design of a stripline monopulse comparator using equivalent circuits and mode matching techniques.," IEEE, 1996.

[18] U. Bulus, O. Kizilbey, H. Aniktar and A. Gunes, "Broadband Direction-Finding Antenna Using Suspended Microstrip-Line Hybrid Coupler for Handheld Devices," *IEEE ANTENNAS AND WIRELESS PROPAGATION LETTERS*, vol. 12, pp. 80-83, 2013.

[19] Analog Devices, "AD8302.pdf," [Online]. Available: http://www.analog.com/static/imported-files/data_sheets/AD8302.pdf. [Accessed 08 2013].

[20] Austrian Microsystems, "AS5306 Linear Position Sensor," [Online]. Available: http://www.ams.com/eng/Products/Magnetic-Position-Sensors/Linear-Incremental-Magnetic-Position-Sensors/AS5306. [Accessed 08 2013].

[21] Lantronix, "XPORT Embedded Ethernet Device Server," [Online]. Available: http://www.lantronix.com/device-networking/embedded-device-servers/xport.html. [Accessed 08 2013].

[22] Atmel Corporation, "ATmega168," [Online]. Available: http://www.atmel.com/devices/atmega168.aspx. [Accessed 08 2013].

[23] Texas Instruments, "CC2500," [Online]. Available: http://www.ti.com/product/cc2500. [Accessed 08 2013].

[24] metageek, "Visualize and Troubleshoot Wi-Fi Interference with Chanalyzer Pro," [Online]. Available: http://www.metageek.net/products/chanalyzer/. [Accessed 08 2013].

[25] metageek, "What Happened to My Duty Cycle?," 01 07 2011. [Online]. Available: http://www.metageek.net/blog/2011/07/what-happened-to-my-duty-cycle/. [Accessed 08 2013].

CHAPTER 4. ANALYZING LINE OF BEARING ESTIMATES COLLECTED FROM A DEVICE EMPLOYING MONOPULSE RADAR METHODS TO TRACK RF TARGETS IN THE 2.4 GHZ ISM BAND

A paper submitted to *IEEE Transactions on Information Forensics and Security (TIFS)*

D. J. Gieseman[1,2] and T. E. Daniels[1]

Abstract

In prior work we presented the design and architecture for a device utilizing monopulse radar methods to detect and spatially attribute an RF target transmitting in the 2.4 GHz ISM band. The aim of our system was to detect and defend an environment against an adversary employing the Parking Lot Attack as a penetration and exfiltration vector. A tertiary implementation objective was to use low cost and off-the-shelf components in the construction of our device, while obtaining an acceptable performance with regard to spatial attribution accuracy. The ability of our device to accurately spatially attribute a wireless target is directly dependent on the capability of our device to both detect an active RF emitter and to subsequently assign a Line-of-Bearing (LOB) to any detected target. In this work we seek to use statistical methods to quantitatively assess the performance of our device for both detection and LOB estimation accuracy. A novel feature of our hardware is the application of monopulse radar principles and techniques to the 2.4 GHz ISM band. We begin by reviewing the concept of monopulse

[1] Graduate Student and Assistant Professor, respectively, Department of Electrical and Computer Engineering, Iowa State University.
[2] Primary researcher and author.

in radar applications and how the concept is implemented on our prototype sensor. We next describe the testing methodology we used to systematically evaluate and quantify the performance of the instrument. Our experiments first focus on the establishment of baseline performance metric using data collected from a sequentially lobed single antenna detection sensor configuration. We then conduct experiments using a monopulse enabled sensor configuration, presenting two different DSP detectors for target detection and LOB estimation. Our experiments reveal statistically significant performance gains over the baseline sequential lobing scheme. We conclude with recommendations for device improvements and topics for future research.

Introduction

The topic of spatial attribution in wireless local area networks interests us, from a security standpoint, as a way of detecting unwanted or malicious use of a wireless network. We use the term spatial attribution to describe the process of assigning a geo-location to a previously unidentified wireless emitter. The emitter, in our case, is typically the wireless local area network (WLAN) radio found on modern portable computers and handheld devices.

Expanding on this idea, consider the increased situational awareness that location intelligence for devices connected to a wireless network could provide a wireless network administrator for a given environment. Given sufficient spatial accuracy, location-enabled intelligence regarding authorized and unauthorized RF activity would enable administrators and security analysts to make decisions about the present wireless threat environment for a facility. For example, tools could be created that map out red-zones surrounding a facility perimeter, where any detected unauthorized wireless activity

triggers alerts that supplement the analytic and reporting capabilities of existing intrusion detection systems.

State-of-the-Practice in Spatial Attribution for Wireless Devices

Contemporary systems for identifying the location of wireless users are presented here [1], [2], [3]. These systems rely on a previously determined audit of signal propagation and signal strength estimates received from a distributed network of base stations within a wireless environment to perform host-based location estimation. A shortcoming of these systems is that constant change in the landscape of the wireless environment degrades positional accuracy, leading to the need for frequently recurring RF audits in order to maintain the accuracy of the location estimation system. Additionally, these systems are predominantly marketed towards personnel tracking and inventory management than for facility protection and security.

Furthermore, most systems are commonly deployed within the interior of a facility, where instead our research seeks to focus on the environment immediately external to a facility, and the security threats existing there. We focus on the external environment, as it is common for the outdoor areas in the immediate vicinity of a facility to be used as the launch pad for externally staged wireless attacks [4], [5], [6]. This is the so-called “Parking Lot Attack”, which our research concerns itself with detecting.

WIDAR: Our System for RF Target Spatial Attribution

In prior work [7], [8] we presented WIDAR (Wireless Intrusion Detection and Ranging), our system employing monopulse radar methods, for detecting and tracking an emitting RF target operating in the 2.4 GHz ISM band. One characteristic common to monopulse devices is that they utilize a multiple antenna array, in contrast to the single,

high gain antenna found on less sophisticated direction finding systems. The array antennas signals feed into specialized hardware that calculates the instantaneous ratio formed when measuring and comparing the individual array element signals. This hardware is called the monopulse processor, and the computed array ratio is called the monopulse ratio [9].

The term monopulse stems from the theoretical observation that each monopulse ratio measurement contains information sufficient to estimate the instantaneous Line of Bearing (LOB) to any emitter actively transmitting within the geometry of the array beam pattern. This differs from LOB estimates made by a singular antenna, where in those schemes an entire rotation of the antenna platform is required to be performed before enough data are gathered to calculate a LOB estimate.

There are many other benefits inherent to the monopulse method of LOB estimation, an important point being that any RF non-linearities, such as those stemming from multipath fading or destructive/constructive interference become common mode to all antenna elements in the array, nullifying the effects of jitter when the monopulse ratio is calculated to yield an angle estimate [9], [10]. Contrast this again to the less sophisticated, single antenna lobing scheme where readings are taken sequentially, and noise issues are integrated rather than cancelled.

Conceptually, our system is composed of a directional antenna array composed of dual antennas mounted on a mechanically steerable chassis. Figure 4.1 displays an illustration of our prototype WIDAR system detection sensor. There are two primary sub-systems: one tasked with RF sensing, and the other responsible for mechanically steering and sensing the azimuthal orientation of the antenna array.

The array antennas feed into low noise power amplifiers and band pass filters. Following the band pass filters, each signal path is split by a directional coupler. One component of each split signal is then input into a logarithmic gain detector [11].

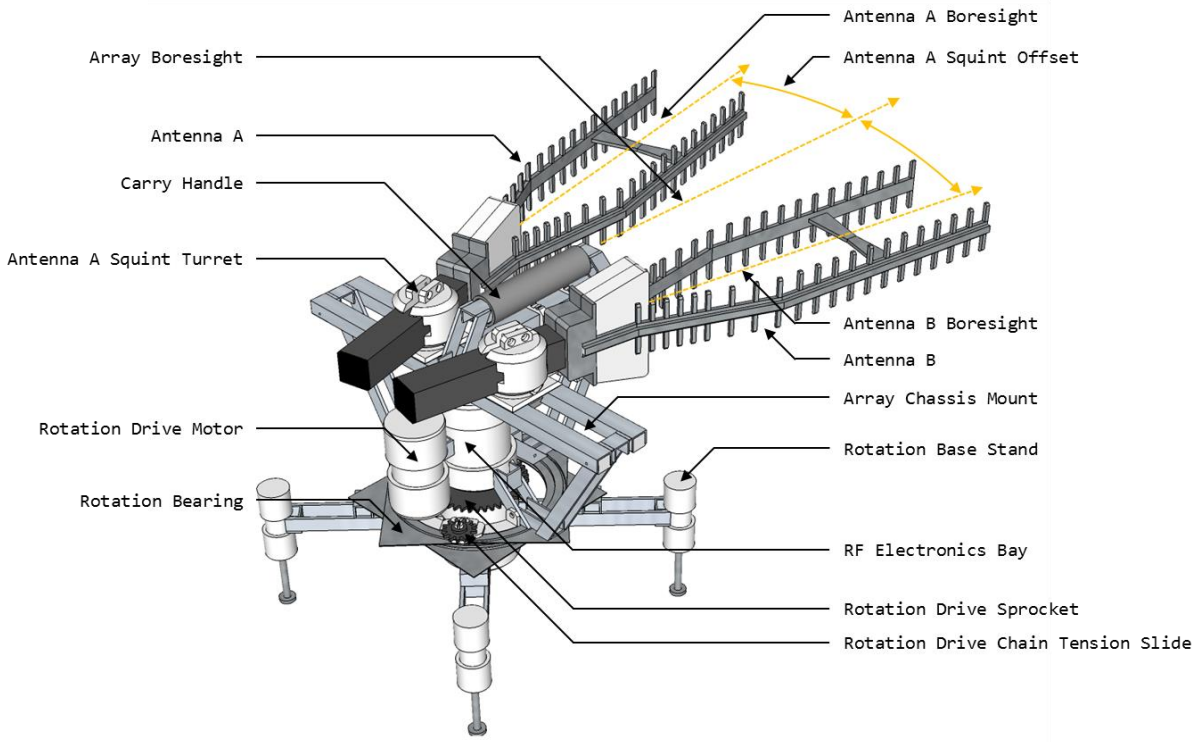

Figure 4.1. The WIDAR System Detection Sensor. Our device features a mechanically-steered, dual high-gain antenna array, and the onboard RF processing electronics capable of providing LOB estimates using monopulse radar principles and methods. The RF functionality is complemented by the rotational encoder hardware necessary to precisely report the azimuthal orientation of the array chassis. As shown in the figure, the array boresight is found at the mid-line, halfway between the boresights formed by each antenna in the array. A slip ring permits the upper chassis mounting the array to rotate continuously, at the same time allowing control and data signal communications between the lower and upper chassis assemblies.

The gain detector is an integrated circuit containing two RF input ports. These ports feed an integrated pair of tightly matched logarithmic detectors. Sampling the difference of these two detectors forms a gain/loss ratio, which is an acceptable source for monopulse ratio calculations [10]. The gain detector places the magnitude of the

detector gain ratio on an analog output pin of the detector IC which is, in turn, sampled by a 10 bit ADC [12] running on a microcontroller housed in the electronics bay of the detection sensor chassis.

The remaining component of each split signal feeds a 2.4 GHz radio [13] tasked with measuring spectrum channel power. Channel power samples are averaged and can be viewed by analysts, as a supplement to the monopulse ratio, providing additional situational awareness capabilities.

The steering sub-system is composed of an H-bridge controlling the direction and speed of a rotation drive motor, and a magnetic encoder [14] that can accurately sense the azimuthal rotation of the upper sensor chassis, which mounts the antenna array. The magnetic encoder on our prototype provides 5760 encoder pulses per revolution, providing 0.0625 degree azimuthal orientation accuracy. The encoder also features an index pulse every 160 encoder pulses, which we use to detect slips or skips in the microcontroller interrupt routines responsible for counting encoder ticks. This permits us to detect and correct for azimuthal sensor drift when and if it occurs. A conceptual block diagram detailing each of these sub-systems is shown in Figure 4.2.

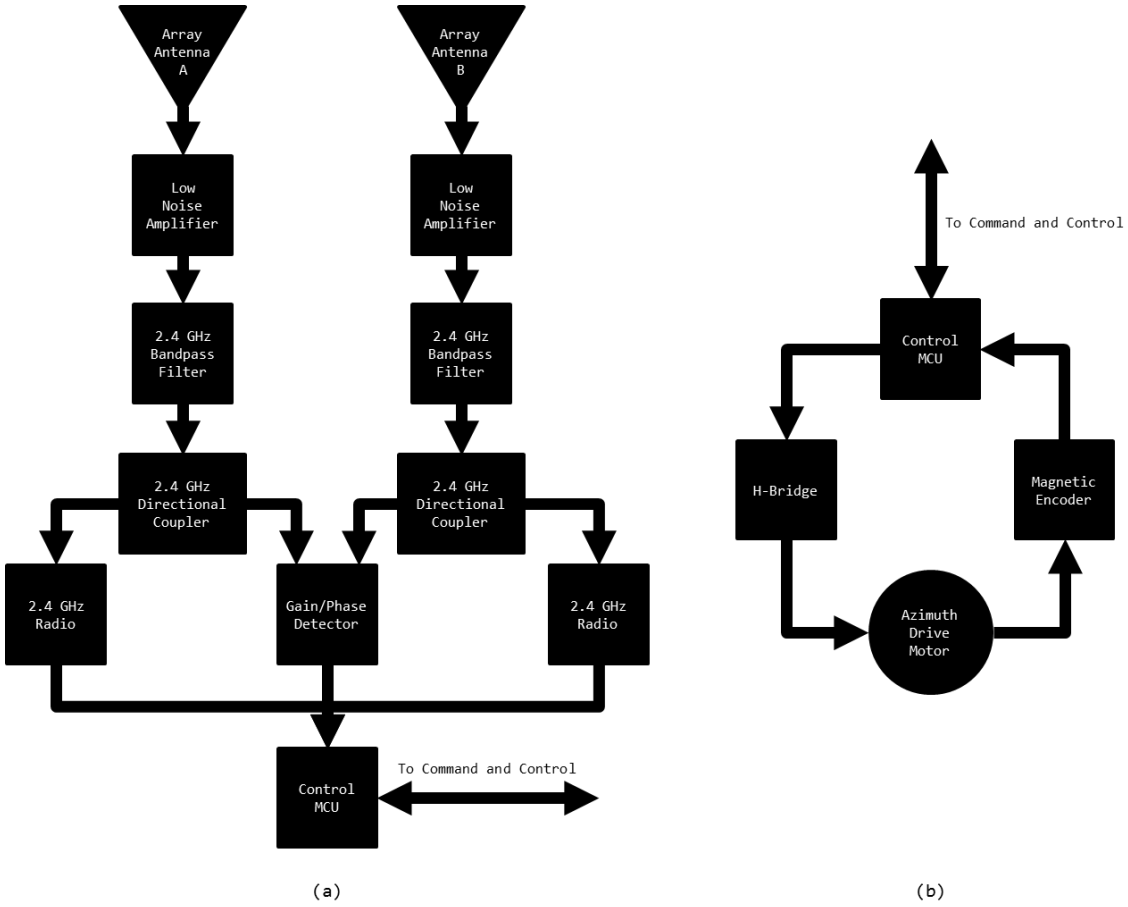

Figure 4.2. Block Diagram of Critical Detection Sensor Sub-Systems. Figure 4.2.a shows functional blocks composing the RF Sensing sub-system, and Figure 4.2.b shows functional blocks composing the Rotation Scan sub-system. Both sub-systems send and receive commands and data to a higher level Command and Control system.

Detection Performance Evaluation Experiments

We wish to quantitatively analyze the performance of the system detection sensor we just described. Theoretical models estimate the line-of-bearing accuracy for our device at 1/10 array element beamwidth or less [10]. Given that our device antenna array has a single-element beamwidth of 25°, we expect to see performance metrics supporting gross angular resolutions of 2.5°, or ±1.25° from true heading.

Using a Multi-Phase Detection Evaluation Approach

The experiments designed to evaluate detection sensor performance are divided into two distinctive research phases. In the first phase, we perform a battery of tests to define and quantify baseline performance metrics. These baseline metrics are intended to serve as a comparison gauge for the second phase of our experiments. The system performance baseline of Phase I is generated from data collected using an alternate configuration of our detection sensor, which features a singular antenna target detection and lobing scheme; a scheme referred to as sequential lobing in monopulse radar literature [9], [15], [10].

Phase II performs a similar battery of detection experiments, however, in this phase the detection sensor configuration is reset to the monopulse-enabled design. We hypothesize that the monopulse configuration of our prototype implementation will significantly outperform the baseline metrics collected using the single antenna sequential lobing scheme. Our two-phase, baseline-then-measure approach is summarized in Table 4.1.

Table 4.1. Summary of Detection Evaluation Experiment Phases. In Phase I, a baseline is developed using a single antenna only capturing peak RSSI. In Phase II, we capture metrics using a monopulse array to compare against the Phase I baseline data. The sensor configurations used in both Phase I and Phase II utilize the same mechanically steered base chassis for tracking and reporting antenna bearing.

Evaluation Phase	Detector Configuration	Mechanically Steered	Detector
I	Sequential Lobing Antenna	Yes	Peak RSSI
II	Monopulse Antenna Array	Yes	Monopulse Ratio

The Field Test Range

In both phases of our performance evaluation, we perform a battery of in-field detection experiments against a moveable RF target with a pre-known position and a

controlled RF transmission profile. To ensure repeatability and to permit the accurate positioning of the RF Detection Target used in our experiments, we designed and surveyed a Field Test Range with pre-measured distance offsets located at known locations from a base position. These distance offsets were measured along a known directional bearing. The detection Field Test Range is detailed in Figure 4.3.

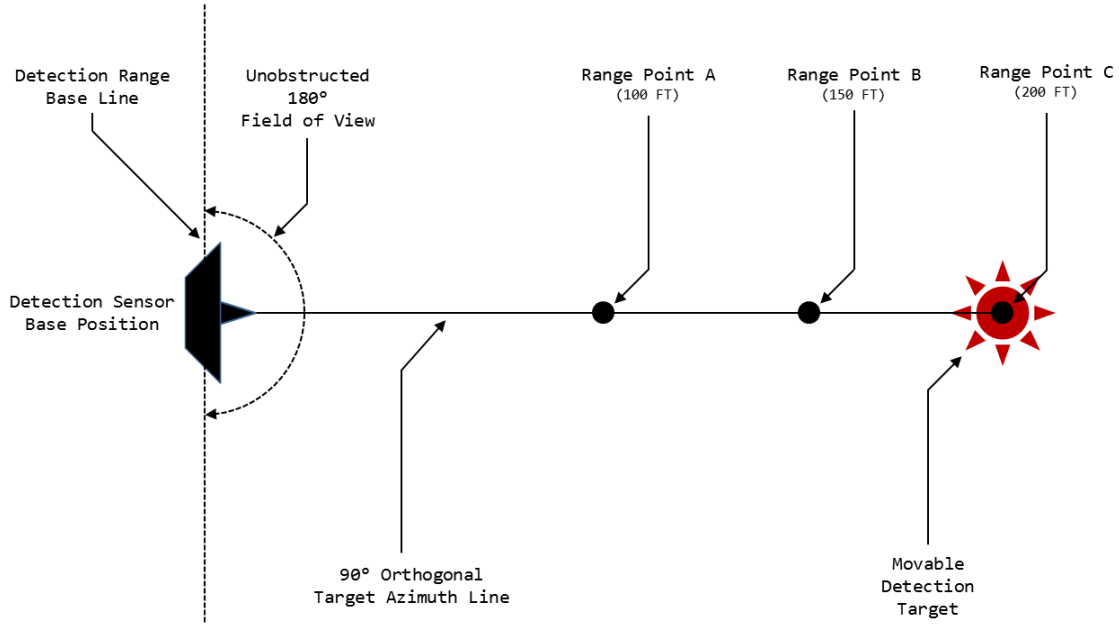

Figure 4.3. Layout Plan-View of the Field Testing Range. The Detection Sensor is shown positioned at the intersection of the Detection Range Baseline and the Target Position Line. Range Test Points (labeled A, B, and C) are shown positioned at known distances from the Detection Range Baseline.

We selected a field location situated on flat terrain and free from natural obstructions, with the intent of minimizing multipath reflections. We also desired to be far from man-made structures to avoid being in close proximity to potential sources of active WLAN activity to minimize direct sources of interference in the 2.4 GHz ISM band. The objective was to reduce the likelihood that detector readings logged during our experiments were affected by externally created signals and were more likely to include

only those signals produced by the RF Detection Target. Figure 4.4 shows an aerial map view depicting the Field Test Range location.

Figure 4.4. Aerial View of the Field Test Range. We selected a site with the intent of being far from man-made and natural structures.

We used a 200 foot field tape measure to survey the detection range baseline and then measured a series of three Range Points along a Target Position Line oriented 90° orthogonal to the Detection Range Baseline. In all of our experiments, we carefully positioned an actively transmitting RF Detection Target on one of the Detection Range Points and collected detection data by mechanically sweeping the antenna chassis of the detection sensor through a minimum 180 degree sweep arc. During the two evaluation

phases, sweep runs at each Range Point were repeated after we made changes to the detection sensor configuration used during each phase. Figure 4.5 shows photographs from the actual field testing range we surveyed and utilized during our testing phases. The photographs are from Phase I tests; a collection scheme where only a single antenna detection sensor configuration was used.

(a)

(b)

(c)

Figure 4.5. Photos of Detection Range During a Field Test. Figure 4.5.a shows a view of the Detection Sensor Base Position. A close-up detail of the platform and the Detection Sensor is shown in Figure 4.5.b. Also visible is the detection range measuring tape which was staked to the ground at the base of the Detection Sensor platform. Figure 4.5.c shows a view of the detection sensor scanning down range. The detection target can be seen at the far end of the range measuring tape.

Surveying the Ambient 2.4GHz Spectrum of the Field Test Range

We used the spectrum channel sampling capability we designed for the WIDAR command and control system [7] to survey RF activity in the testing range wireless environment, both before and after we began operating any RF detection targets on the Field Test Range. Figure 4.6 shows the results of spectrum channel sampling before and after we deployed RF detection targets. Given the crowded and busy spectrum channel samples that we were used to seeing when developing and testing our system in a lab environment, we were pleased to find only faint 2.4 GHz signals operating in the spectrum when we collected data on the field test range.

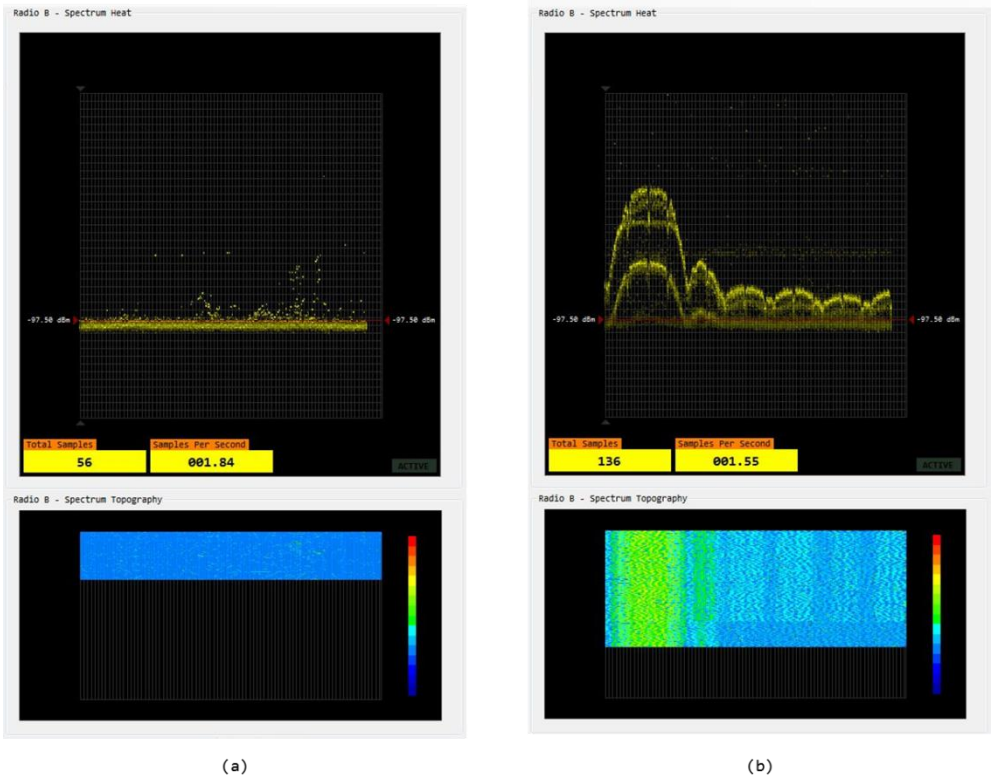

Figure 4.6. Comparing Environmental Channel Spectrum Power Samples. Shown are samples taken before the detection target was activated (Figure 4.6.a) and after (Figure 4.6.b) the RF detection target was made active. The RF detection target operated on 802.11g Channel 1, and can be seen clearly in the right image. The samples in the left image do show some weak activity, in what appears to be 802.11 Channel 11, although the activity is only just above the test range noise floor.

Phase I Experiments: Establishing Baseline Detection Sensor Metrics

During Phase I, we utilized an older generation laptop running Xubuntu Linux 14.02 LTS as the Detection Target on our RF Test Range. Figure 4.7 shows the Detection Target. The RF Detection Target ran LORCON [16], a C library supporting raw WLAN packet injection of 802.11 frames directly into the wireless environment of our RF test range. Since we could craft and inject raw 802.11 packets directly into the RF medium, no supporting infrastructure such as an Access Point or WLAN router were required to simulate active WLAN traffic. Using the LORCON library, we wrote a simple packet flooding tool to inject spoofed 802.11 Management Frame Beacon Packets [17].

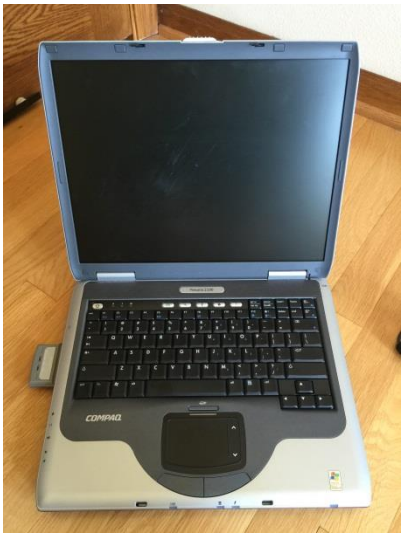

Figure 4.7. The Phase I detection target was an older laptop that we fitted with an 802.11g card supporting Linux raw wireless packet injection. The 802.11g card contained WLAN hardware built around the Atheros chipset.

Phase I Detection Sensor Configuration

In Phase I of our evaluation experiments, we fielded an alternate detection sensor configuration employing a less sophisticated method of LOB estimation. In this scheme

only a single antenna is mechanically rotated, while a Received Signal Strength Indication (RSSI) is continuously recorded. A windowed moving average filter is then run against the recorded RSSI data to calculate a peak value that is recorded following each completed antenna chassis rotation. The LOB pointing to the emitter was assigned to this peak value calculated by the filter.

The shortcomings of this method are well documented [9], [10] [16], but we chose to use metrics from this method as a gauge to determine whether our monopulse array detection prototype outperforms, underperforms, or compares similarly to an intentionally less sophisticated baseline method. Figure 4.8 shows the detection sensor in the single antenna lobing RSSI sensing configuration.

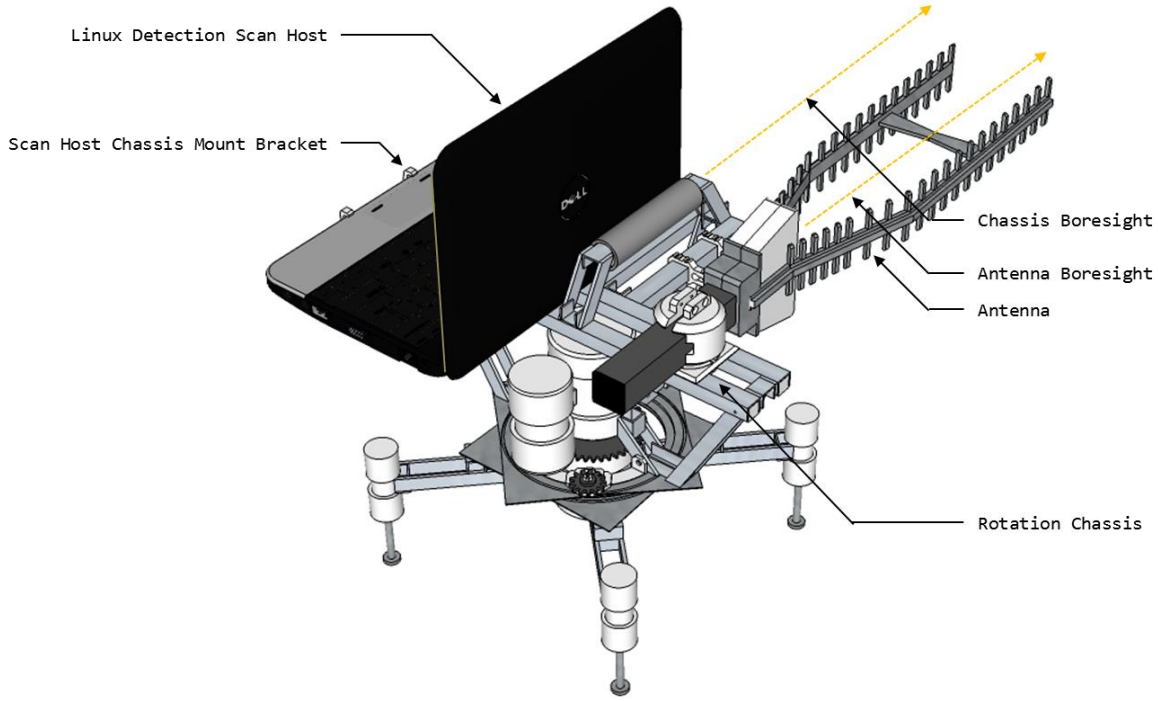

Figure 4.8. Chassis Configuration for Sequentially-Lobing Peak RSSI Detector. In this configuration, one monopulse array antenna is removed, and instead a small computer running Ubuntu Linux is mounted to the rotation chassis assembly. The single remaining antenna is boresighted in parallel to the chassis boresight, in contrast to the squinted boresight scheme used in the monopulse configuration of the detection sensor.

Phase I Detection Sensor Data

Data for Phase I were collected using two completely autonomous hardware systems, each with separate sense and data logging functionality. Azimuthal orientation data for the detection sensor chassis were logged using a laptop running the WIDAR Command and Control (C2) software [7]. The C2 software provided a user interface with the data logging functionality necessary to save rotation angle, encoder tick count, and timestamp to comma delimited ASCII files. These data were timestamped with a clock value set using NTP to ensure that the RSSI values, which were collected using different hardware, could be joined in a post-collection processing tool.

To log RSSI values, we used a small netbook physically mounted to the mechanically steered upper assembly of the detection sensor chassis. This computer ran the Ubuntu 11.02 LTS Linux operating system, which enabled a variety of security oriented tool chains to be deployed. We also configured the real time clock on this computer using the NTP protocol, and then we verified that this machine, along with the detection sensor C2 machine, were both running clocks that appeared tightly synchronized with each other.

Additionally and perhaps more notably, this netbook featured an after-market modification that we made which permitted an external omnidirectional antenna to be connected directly to the WLAN card of the device. The computer and the external antenna modification are shown in Figure 4.9. The modification enabled us to directly connect a single high-gain antenna to the netbook; the exact same antenna used in the monopulse configuration of the detection sensor.

(a) (b)

Figure 4.9. Photos of Linux Laptop Modified for Connection of External Antenna. The external RP-SMA connector is shown in Figure 4.9.a. Figure 4.9.b shows how an external antenna can then be mounted to the connector.

The WLAN interface on the netbook was placed in monitor mode, which supports promiscuous packet sniffing, along with 802.11 channel hopping. We ran the Wireshark protocol analzyer [18] on the netbook to enable raw packet capture, and most importantly, the logging of all wireless activity seen by the netbook on the monitor mode enabled WLAN interface.

This version of Linux also supported the mac80211 [19] driver which conveniently prepends additional WLAN metadata onto all received WLAN packets. These metadata were accessible in Wireshark under the radiotap headers [20], [21] protocol-decode section of the packet capture trace file. Figure 4.10 shows how

Wireshark presents these data. The radiotap metadata for Signal and Signal Noise were critical elements supporting our objectives for Phase I: the collection of baseline performance metrics.

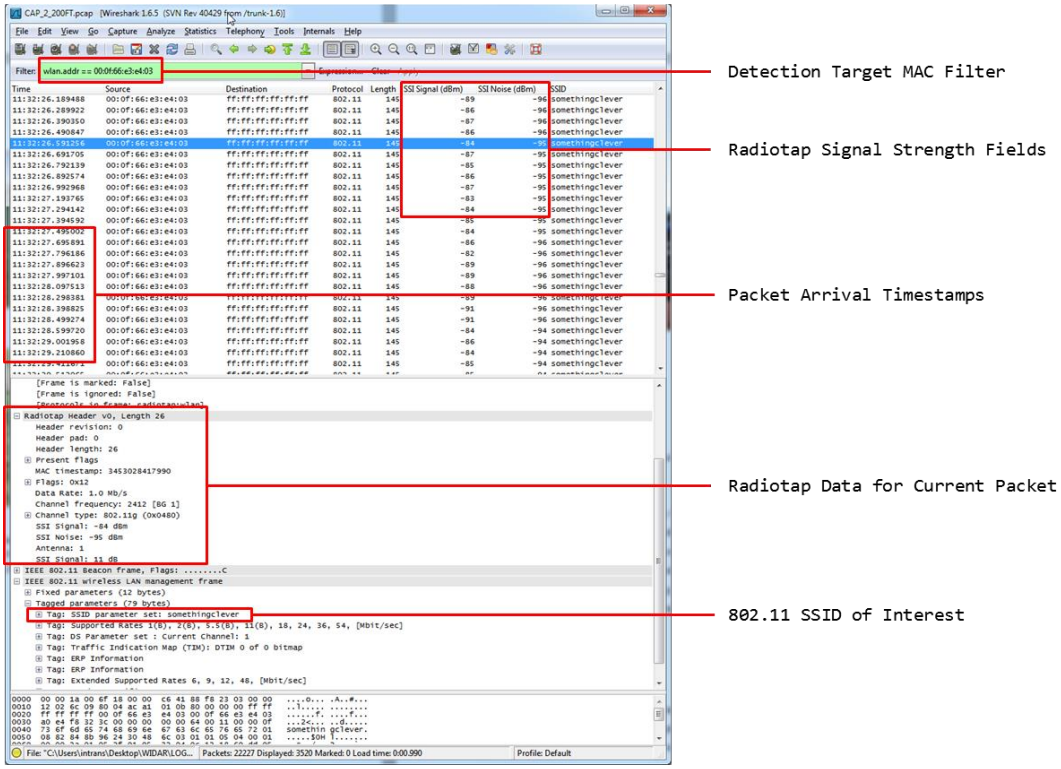

Figure 4.10. Wireshark View of Detection Sensor RSSI Packets. On Linux,Wireshark is able to use the libpcap driver to capture WLAN network traffic on any WLAN interface placed in Monitor mode. The Linux mac80211 driver system also prepends the radiotap header to wireless network packets. The radiotap extensions provide the signal strength and signal noise measurements necessary to create the RSSI logs for later analysis alongside the antenna bearing logs.

We were also able to use Wireshark to filter detection sensor data by MAC address to select only those packets that matched the RF detection target MAC address. Protocol stack browsing functions of Wireshark were also used to view those protocol attributes that we required for our analysis. For the final step, Wireshark data export functionality was then used to save only the attributes were interested in, to a comma-delimited text file.

Processing Phase I Data for Analysis

We just described how data collected from the sensor produced two different output files, one from the detection sensor C2 machine, and one from the netbook attached to the antenna rotation chassis. These two data files both contain timestamp information that was previously synchronized using the NTP.

We created a tool using Microsoft VB.NET 4.5 to join the two separate datasets using the common timestamp. The tool used a fuzzy matching algorithm to join the data, requiring only that any matched timestamps be within a preset threshold time, to be joined. If there was more than one candidate meeting the join threshold, additional proximity analysis was performed to determine the nearest timestamp match. All joins were one to one. Figure 4.11 contains a block diagram illustrating the data join process. Table 4.2 shows the data format of the resulting join process. A data file in this format is then fed to analysis logic tasked with target detection.

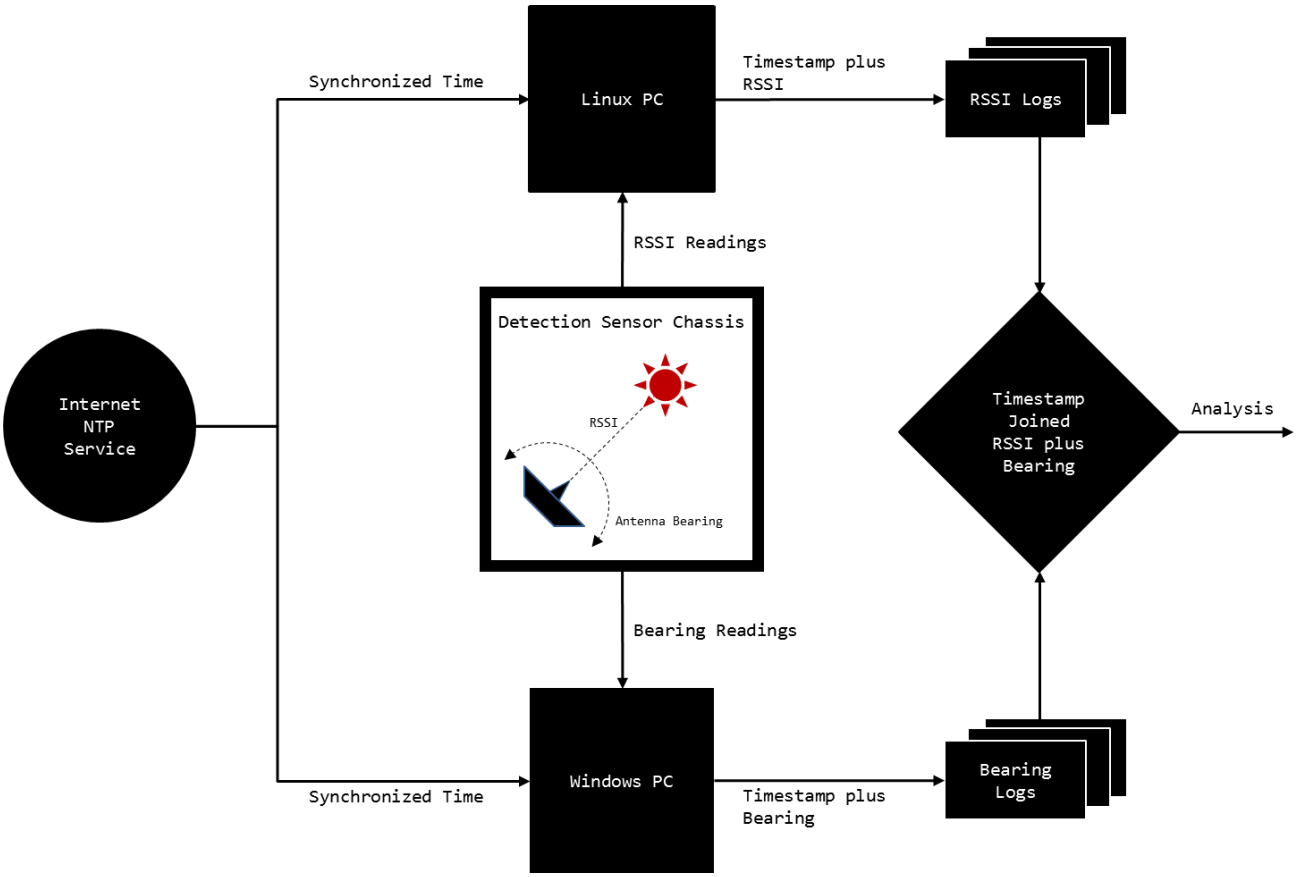

Figure 4.11. Data Fusion Process for the Peak RSSI Detector. In this detection mode of operation, two computers are required, one responsible for logging antenna bearing, and the other responsible for logging 802.11 RSSI and Signal to Noise readings. A network time protocol service is used to accurately synchronize the real time clocks which run independently on each machine. The common timestamp information permits logged data to be joined together during post collection analysis.

Table 4.2. Sample Timestamp-Joined Signal, Noise, and Bearing Data output by the detection sensor.

TimeStamp	BEARING	SSI_SIGNAL	SSI_NOISE
10:46:11.444	0.19	-37	-94
10:46:11.444	0.19	-36	-94
10:46:11.444	0.19	-36	-94

All collected detection sensor output samples feature the following attributes:

- **Timestamp** – the time of the sample – with a resolution of 1 millisecond.
- **Bearing** – the orientation of the array azimuthal heading – with a resolution of .0625° [14].

- **Signal Strength Indication** – the signal strength at the antenna, in dBm, with a resolution of 1dBm [20], [21].
- **Signal Strength Indication Noise Floor** – The noise floor at the antenna, in dBm, with a resolution of 1 dBm.

Phase I Target Detection Scheme

With RSSI based signal and noise data it is a simple matter to calculate the signal to noise (SNR) ratio and then to filter the resulting SNR signal searching for peak sample values. Since both the signal and the noise values of the detection sensor data are logarithmically scaled, the SNR is calculated by subtracting the noise from the signal. We performed a window moving average filter against the resulting SNR signal.

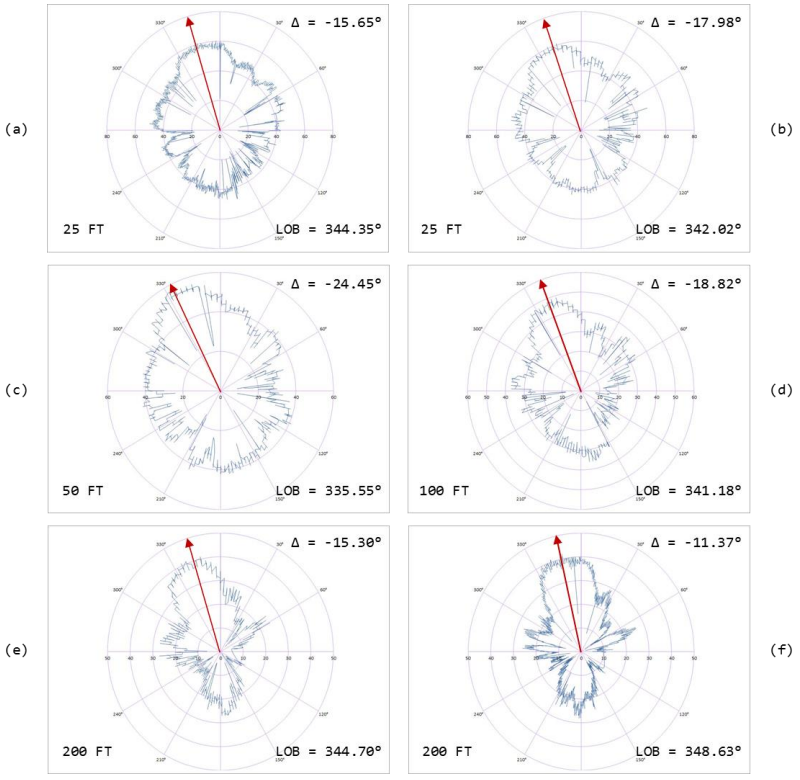

Figure 4.12. Polar Plots showing the LOB assigned to 6 different detection sensor runs using targets at varied distances and sampling frequency.

We also desired to know whether the choice of window size used in windowed moving average calculations would show any statistically significant impact on detector performance. We scripted the moving average filtering so that selected window sizes ranging from 1 (no averaging) to 255 were reported, using a window size step that double the window size of the previous iteration step. Following the filter step, we selected the windowed moving average RSSI value with the highest peak and assigned that value as the Line of Bearing to the detection target. Figure 4.12 features polar plots showing the LOB results from performing the windowed moving average filter on RSSI values from 6 separate detection sensor rotations.

Phase I Statistical Analysis of Baseline Metrics

We calculated average, variance, standard deviation, and confidence interval values for the treatment of Moving Average Window Size, and for Target Distance from Detection Sensor. The statistics in the lower rows of Table 4.3 show values for the different Window Size treatment, with boresight mean error values shown between -16.30° and -19.84° with a wide range of variance values, yielding 95% confidence intervals ranging from ±3.53° to ±6.42° from mean boresight error values. The polar plots in Figure 4.12 taken in conjunction with these results led us to believe that our antenna was either grossly misaligned during the field tests or that the antenna has a distinctive gain geometry skew that differs significantly from the product data sheet.

Table 4.3. Phase I Statistical Results.

	1	15	31	63	127	255	MEAN	VAR	STDEV	CI (95%)
25 FT	-16.94	-15.28	-16.94	-21.72	-24.72	-22.56	-19.69	14.44	3.80	3.99
25 FT	-15.50	-16.10	-14.88	-15.97	-15.97	-10.66	-14.84	4.41	2.10	2.20
50 FT	-28.75	-22.31	-23.94	-27.09	-20.53	-24.53	-24.52	9.13	3.02	3.17
100 FT	-15.69	-17.28	-18.82	-20.38	-23.37	-21.16	-19.45	7.68	2.77	2.91
200 FT	-12.69	-14.28	-14.28	-17.35	-20.60	-21.50	-16.78	13.27	3.64	3.82
200 FT	-11.87	-12.53	-11.38	-10.19	-13.85	-11.94	-11.96	1.47	1.21	1.27
MEAN	-16.91	-16.30	-16.70	-18.78	-19.84	-18.72				
VAR	37.37	11.29	18.88	32.80	17.64	34.64				
STDEV	6.11	3.36	4.35	5.73	4.20	5.89				
CI (95%)	6.42	3.53	4.56	6.01	4.41	6.18				

ANOVA tests conducted on the mean boresight error across different Moving Average Window sizes, indicate no significant difference in means collected using different moving average window sizes. Table 4.4 shows these ANOVA results.

Table 4.4. ANOVA Results Comparing Mean Boresight Error. The boresight means were collected using different moving average window sizes. The null hypothesis is accepted; there is no significant difference in LOB mean error when using different moving average window sizes.

ANOVA – Windowed Moving Average

SS	**df**	**MS**	**F**	**P-value**	**F crit**
61.16984	5	12.23397	0.480919	0.787618	2.533555

The statistics in the right columns of Table 4.3 show values for boresight errors collected using four different Range Point distances of 25, 50, 100, and 200 feet. The boresight mean error values shown range between -11.96° and -24.52°, again with a wide range of variance values, yielding 95% confidence intervals ranging from ±1.21° to ±3.99° from mean boresight error values. ANOVA results on the data range do indicate a significant difference in the boresight error means when the RF detection target is placed at different range points using the RSSI method.

Table 4.5. ANOVA results when comparing mean LOB collected at different distances from the detection sensor. The null hypothesis is accepted in this case, indicating that there is significant different in the boresight error means.

ANOVA – Distance to Detection Sensor

SS	df	MS	F	P-value	F crit
572.33674	5	114.4673483	13.62732	5.6491E-07	2.533554548

Pairwise analysis of the means show significant differences between the 25 foot and 200 foot tests, as well as the 50 foot tests, and 200 foot tests. We looked more closely at the polar plots in Figure 4.12 and concluded that the RF detection target may have been too close to the detection sensor for the 25 and 50 foot tests. The dynamic range appears to be reduced for these plots, while the dynamic range for the 100 and 200 foot plots appears much wider; this is particularly easy to distinguish in the 200 foot plots, where a large main gain lobe is distinctly visible. We hypothesize that the significant mean boresight error may be due to the differences in dynamic range at different distances. It is a problem for future research.

Phase I Conclusions

The results of the LOB estimation performed using the RSSI detection methods showed gross errors of between 3% and 6% from true boresight angle. The results of polar plotting our data visually show a distinctive skew which we surmise may be due to antenna misalignment during data collection or a geometric defect in the antenna gain pattern. Since we only desired to establish baseline metrics to compare the performance of the monopulse enabled design, we did not devote additional effort to following up on these interesting questions. Our primary takeaway from the analysis of Phase I results

was that placing the RF detection target too close to the detection sensor may saturate the detector resulting in suppressed signal dynamic range.

Phase II Experiments: Evaluating Monopulse Detection Sensor Performance

In Phase II, we returned the detection sensor operating mode to the monopulse configuration, by reinstalling the dual antenna array on the device chassis, and removing the Linux computer from the chassis frame. Given the low dynamic range attained in the data collected using the close-in distances of 25 and 50 feet, we adjusted the field test range points to instead use range point distances of 100, 150, and 200ft in Phase II.

The Phase II RF Detection Target

The RF Detection Target was also swapped out in Phase II. We utilized the same Linux netbook that was used for RSSI logging in Phase I, due to the ability to connect an external antenna. In Phase II, we connected a 4dBi omnidirectional ¼ wave 2.4 GHz antenna to the device, as depicted by the photograph shown is Figure 4.9.b. The netbook was primarily selected due to the increased battery life versus the older computer that was used as the detection target in Phase I. This enabled us to stay fielded longer and collect a much larger number of detection sensor samples.

Similar to Phase I, we ran LORCON, the C library supporting direct 802.11 WLAN packet injection on Linux platforms. We again utilized the beacon flood tool we created for Phase I, which injected a continuous stream of 802.11 Management Frame Beacons into the 2.4 GHz environment of the field test range, avoiding the need to set up and deploy any additional wireless network infrastructure to the field test range.

Settings the Squint Angle of the Detection Sensor Antenna Array

The process of overlapping antenna beam patterns is called squinting in the monopulse literature [9], [15], [10]. Theory states that optimal detection LOB resolution is obtained when the antennas in a monopulse array are squinted with an overlap of ½ beamwidth each [10]. Given that the antennas used on our array operated with 25° H-Plane beamwidth, we configured our array to have each antenna squinted at 6.25° each, to arrive at the recommended 12.5° overlap.

The rotating portion of the detection sensor chassis features antenna mounting assemblies with easy-to-adjust squint angles. These assemblies are called “squint turrets” in our design. Figure 4.13 shows the squint turrets. We used a machinist’s protractor to align the arrays on the detection sensor chassis.

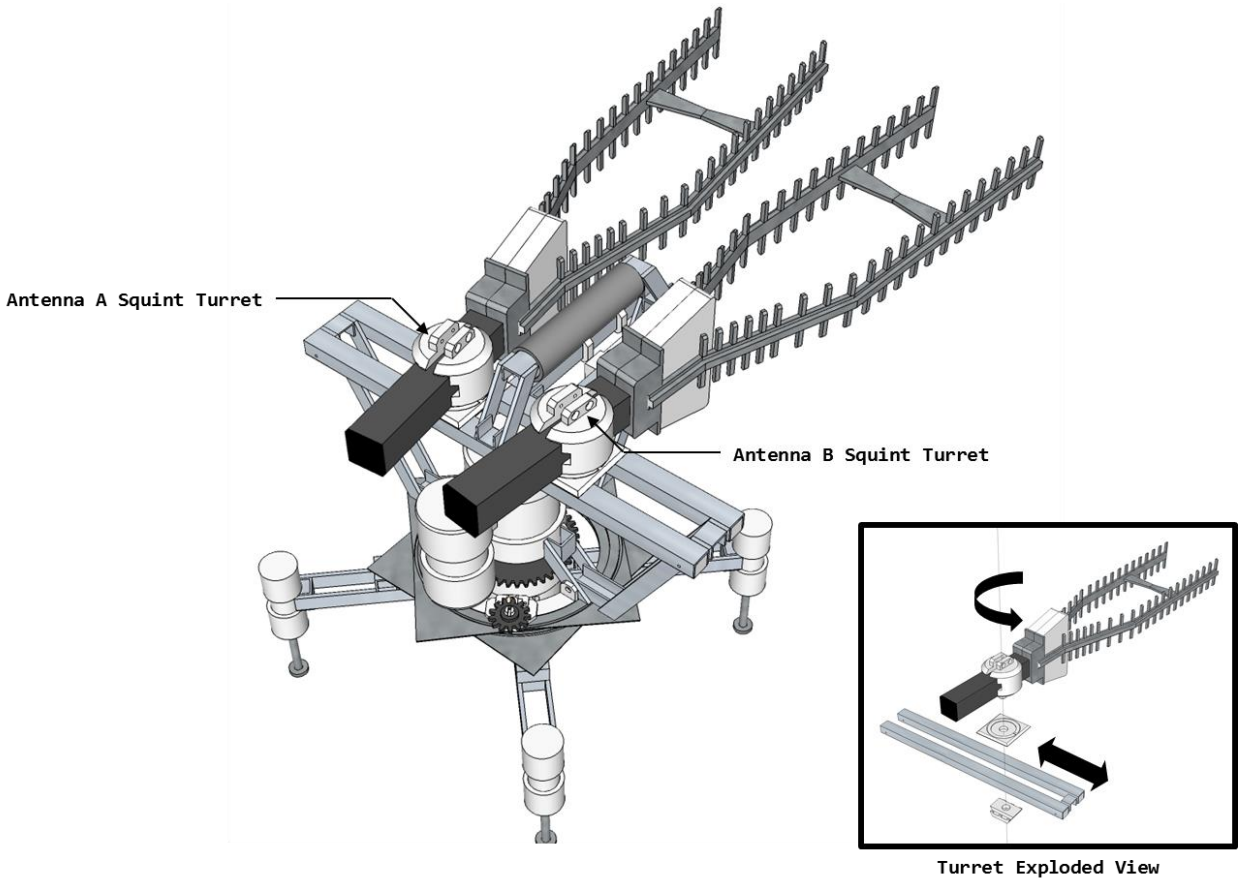

Figure 4.13. Detection Sensor Showing Both Antenna A and Antenna B Squint Turrets. The inset box, pictured in the lower right of the figure, illustrates how the squint turrets permit the lateral positioning of each antenna, in addition to allowing adjustments to be made to the boresight angle orientation of each antenna to any azimuth bearing between 0° and 45°. The entire assembly is fastened together using a machine screw and wing-nut, permitting quick, in-field adjustments of antenna lateral offset and squint settings.

The overlapping beam patterns form a composite beam pattern when presented to the ratio-detection circuitry housed within the RF electronics bay of our device. A schematic diagram detailing the individual beam patterns and the overlaying composite beam pattern is shown in Figure 4.14. According to the datasheet provided by the manufacturer of the array antennas, the composite beam pattern should have a main lobe measuring 37.5°, in practice we observed gain sensitivity beginning to peak at 25° to 30°

off the array boresight, yielding an actual operational beamwidth closer to 50° or 60°.

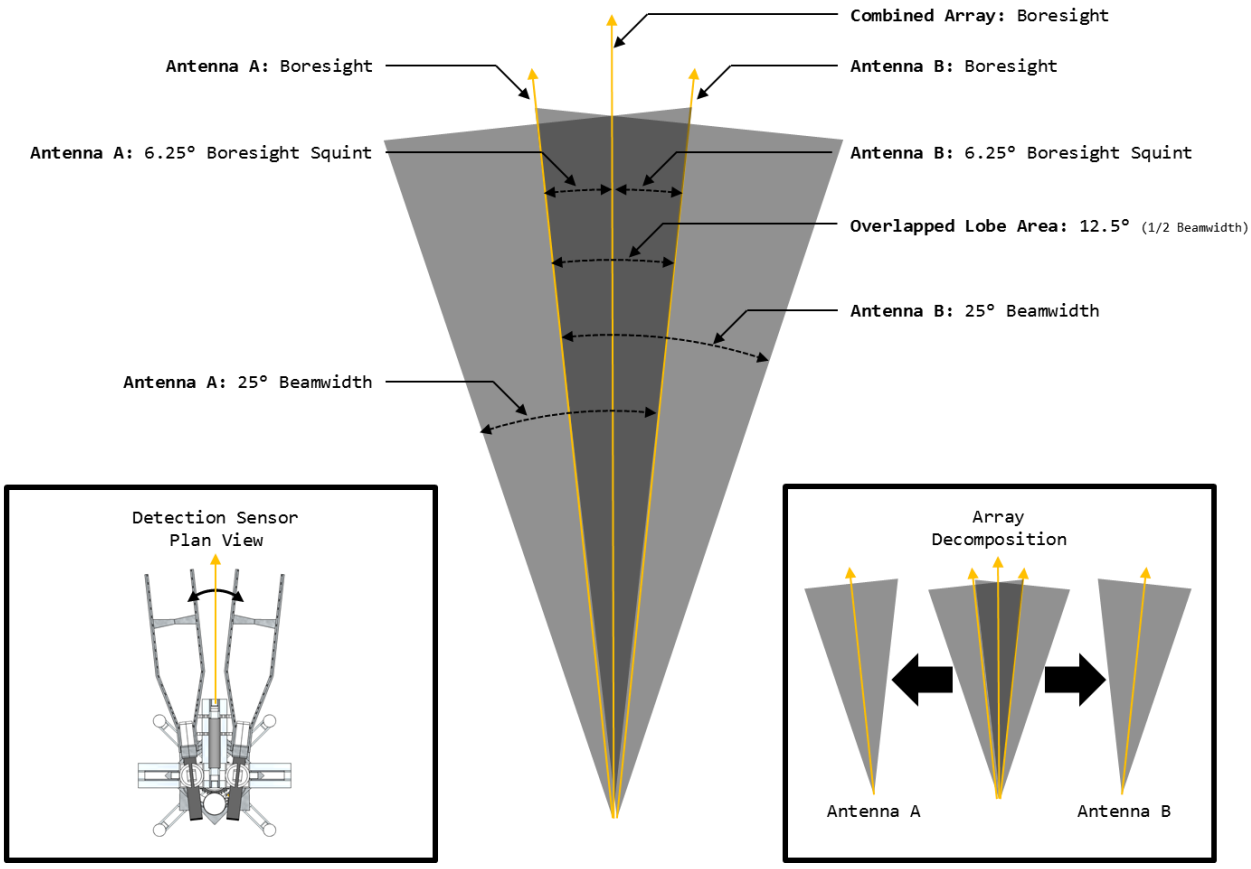

Figure 4.14. Schematic Diagram Detailing Antenna Beam Pattern Overlaps for Array Antennas A and B. The dark gray area is the ½ beamwidth overlap created by the squinted antennas. The left inset area shows the physical squint geometry of the arrays on the actual detection sensor. The right inset area shows the decomposed beam patterns of the individual antennas.

Phase II Detection Sensor Data

In Phase II collected data do not require the complex post-collection processing procedures that arose during Phase I, where it was necessary for data from disparate collection sensors to be joined using carefully synchronized timestamp attribution. Rather, the detection sensor in the Phase II monopulse configuration conveniently outputs

a single data stream featuring integrated timestamp, array bearing, and monopulse ratio measures. This results from the more tightly integrated and purpose built nature of the monopulse collection sensor hardware. Three samples of data output from the detection sensor are shown in Table 4.6 below.

Table 4.6. Sample Monopulse Ratio Data output by the detection sensor.

timestamp	bearing	avg	min	max	window
10:07:47:978	224.81	474.33	467	551	4
10:07:48:010	224.81	496.33	484	548	4
10:07:48:025	224.81	508.00	480	546	4

All collected detection sensor output samples feature the following attributes:

- Timestamp – the time of the sample – accurate to the nearest millisecond.
- Bearing – the orientation of the array azimuthal heading – accurate to .0625° [14].
- Average Monopulse Ratio – the monopulse ratio that is the result of performing a windowed moving average on continuously sampled values resulting from microcontroller analog digital conversions of detector circuit analog values. The monopulse ratio is expressed in unit-less decibels, and has a measurement resolution of .059 dB [11].
- Minimum and Maximum Monopulse Ratio – the min and max monopulse ratio values that were recording during the present sample period.
- The Moving Average Window Size – the moving average window size was configurable using the Command and Control software of the detection sensor.

As we previously alluded to, all data output from the detection sensor in the monopulse operating mode are spatially coherent; meaning every azimuthal-heading value in the output data are accompanied by monopulse data collected simultaneously while the detection sensor was oriented at the precise heading recorded in the data stream. No interpolation or time skewing is required when analyzing these data. Figure 4.15.a shows a plot of 101,355 average monopulse ratio samples captured during a field test against an RF detection target located 150 Feet from the detection sensor baseline. This signal was normalized to unit amplitude prior to plotting. As can be seen by counting the peaks in Figure 4.15.a, the samples shown result from 17 complete rotations of the detection sensor. Figure 4.15.b shows a magnified view of the same sample data, plotting the signal detail for 3000 monopulse ratio samples; the data for just one complete detection sensor rotation.

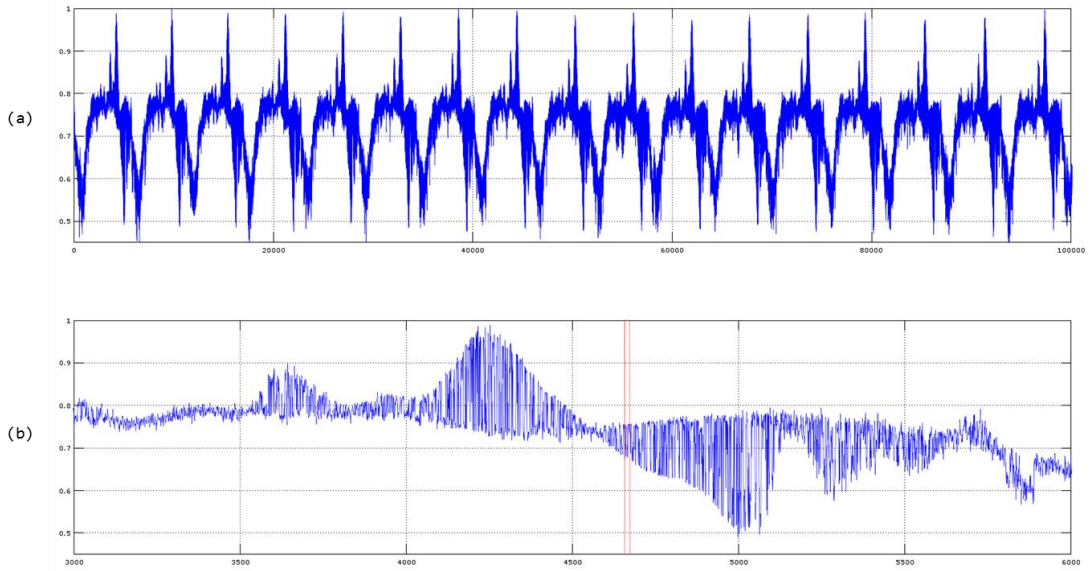

Figure 4.15. Monopulse Ratio Data for 17 Complete Detection Sensor Rotations. For this plot, the RF detection target was placed at Range Point B (150 feet from detection sensor). These data were normalized to unit amplitude prior to plotting. Figure 4.15.a shows all collected samples, while Figure 4.15.b shows a close up detail for only a single detection sensor rotation. The LOB for the emitter is shown as a red line in Figure 4.15.b.

Phase II Target Detection Schemes

We attempted many signal processing techniques in an effort to discover a repeatable and reliable indicator of RF activity occurring within the detection sensor array beamwidth. Notable methods we employed include the Discrete Fourier Transform (DFT) and a windowed DFT technique using Welch's method. After many attempts, we found no telling indicators using frequency domain representations and visualizations of the monopulse ratio signal.

Focusing on time domain plots similar to those in Figure 4.15 we were able to see that the shape of the signal peaks showed distinct visual correlation corresponding to those times when an emitter was active in the detection sensor beamwidth. However, the signal amplitude also varied with rotation, to the point that some samples had amplitudes above or below those of these peak patterns, even though there was not any emitter active in any array antenna beam during these sample times. We looked for additional signal processing methods which could normalize the signal in a way that eliminated these amplitude discrepancies.

Several signal statistics did produce telling results; most notably when we plotted the variance of the signal average amplitude, calculated using a sliding window of signal samples. The windowed signal variance of the monopulse ratio clearly indicated peaks closely corresponding to when an active transmitter was being lobed by the monopulse antenna array, while conveniently attenuating the problematic peaking that was sometimes present in the monopulse ratio signal, even when no emitter was active.

We calculated the windowed monopulse ratio variance using Octave [22] and the Octave signal processing package extensions. We time-shifted the resulting variance

signal so that each variance value was time-aligned with the index position corresponding to the halfway point of the variance calculation window. Finally, we normalized the variance signal to unit amplitude, so we could directly overlay monopulse ratio and monopulse variance signals, using Octave to create comparison plots.

In Figure 4.16.a the blue-colored signal represents the averaged monopulse ratio samples from two completed detection sensor rotation cycles. The red-colored signal in the variance plot shown in Figure 4.16.b shows the characteristic ‘twin peak’ signal pattern corresponding directly to the high and low amplitude peaks seen in the monopulse ratio data. These two patterns also correspond with the timing of the detection RF target entering the approximate 50° peak-to-peak beamwidth of the detection sensor antenna array. Figure 4.16 also displays green lines for reference, indicating where the actual detection target LOB is found in the context of these signal traces.

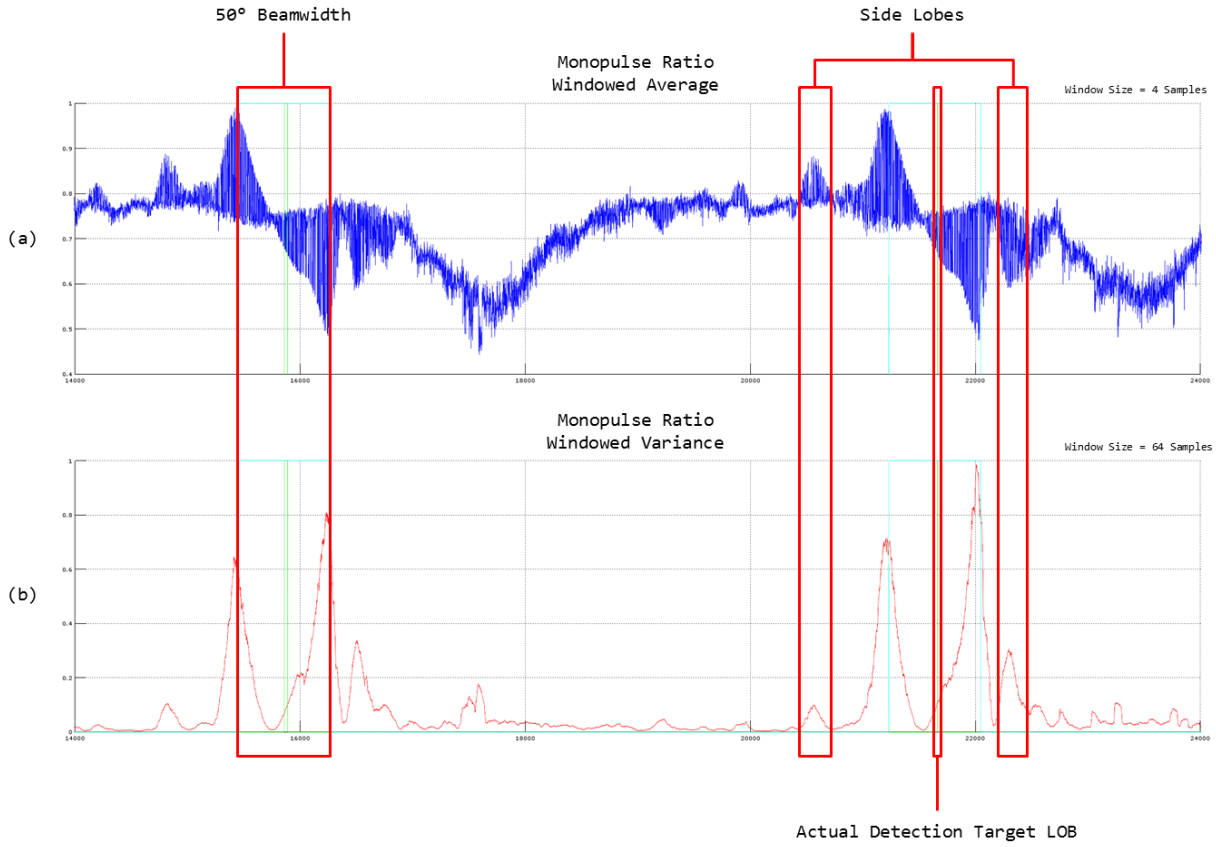

Figure 4.16. Monopulse Ratio Windowed Variance for Two Complete Detection Sensor Rotations. The monopulse ratio was collecting using a 4 point moving average, while the Monopulse Variance was collected using a 64 point moving average. In both figures, the points where the emitter target enters the beam are clearly visible. The variance statistic was found to be a reliable detector for determining when an emitter is operating within the array beamwidth. The actual detection target LOB is shown by the green line on the left side detection, while the LOB is shown highlighted in red on the right side detection.

It is interesting to note that the actual LOB does not fall exactly where the two main lobes of the monopulse ratio signal switch from a ratio dominated by Array Antenna A to a ratio dominated by Array Antenna B. Rather, a distinctive right-shifted bias can be observed, as seen in the figure. We attribute this to mismatches in the stripline microwave transmission channels on the PCB we implemented for the detection sensor. One array antenna experiences more transmission channel loss than the other. Figure 4.17 depicts a detail plot illustrating this observed bias.

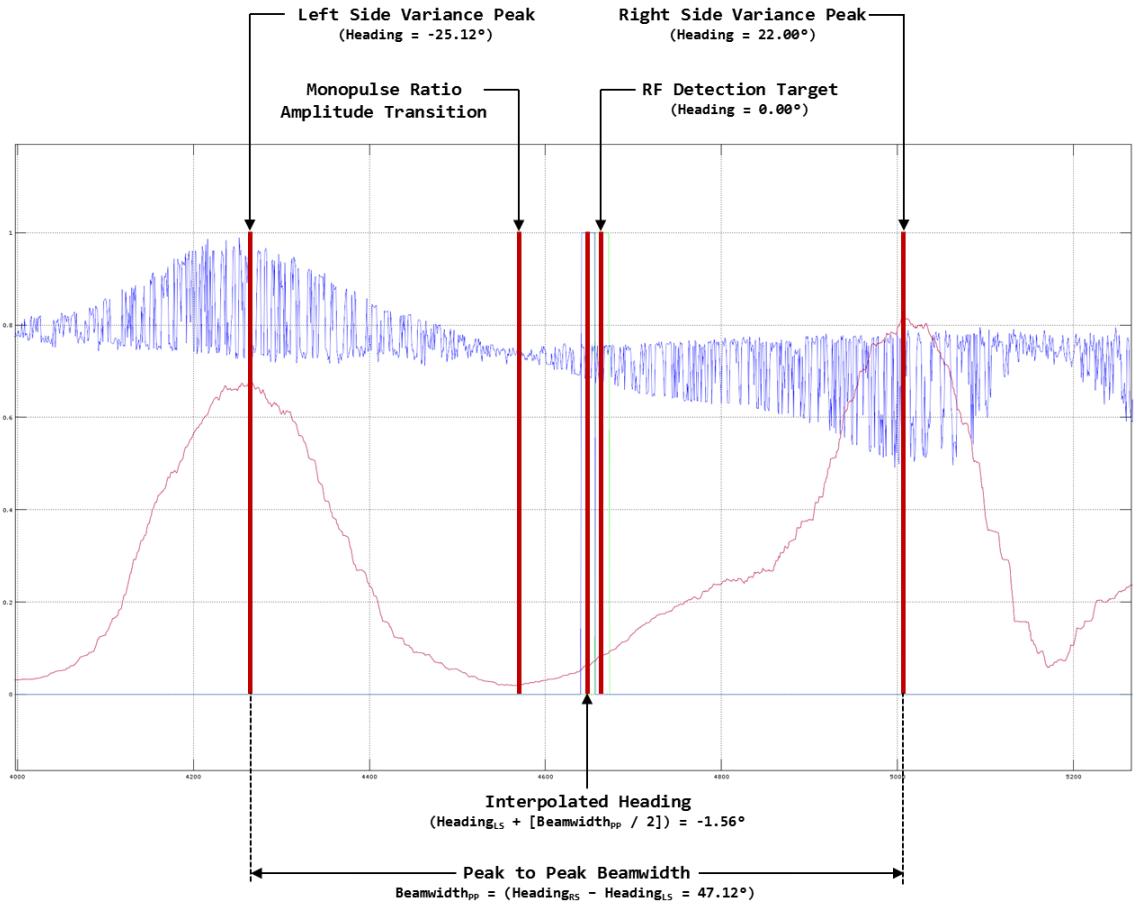

Figure 4.17. Detail Plot Illustrating Bias Observed in Monopulse Ratio Mean. In a properly matched and tuned detection sensor, the monopulse ratio amplitude transition and the bearing to the RF Detection Target should closely align. In our system there was a distinctive offset biased towards Antenna A. We attributed the bias to flaws in our implementation of microwave transmission lines on the detection sensor PCB. We did not explore this bias in greater detail, since we also observed that the Left and Right monopulse variance peaks were still highly correlated with the actual bearing to the RF Detection Target.

We did not fully explore more test cases that could assist in further attributing the detection sensor bias, since we also observed that, despite this bias, the actual target emitter LOB aligns very closely with the point lying half the distance between the peaks of the main lobes of the monopulse windowed variance signal. Having made this observation, we next set out to calculate the LOB estimate using the indices of the variance signal peaks, which we manually identified using the plots we made using

Octave. These indices directly correspond to index values in a synchronous heading signal that is also an output of the detection sensor.

Heading Interpolation Using a Windowed Variance Peak-to-Peak Detector

Given that we can visually identify peaks in the monopulse windowed variance signal, and given that the signal array indices denoting the peak values also index directly into a time-synchronous heading signal which stores the antenna chassis azimuthal bearing, it is a simple operation to lookup the headings for each peak, and then to interpolate a heading for the point located half way between these bracketing bearings. This point is what we call the Line of Bearing to the Detection Target. Equation 1 describes the simple calculation that is necessary to calculate the detection target LOB estimate. In the equation, θ_{RS} is the Right Side Peak Heading Angle and θ_{LS} is the Left Side Peak Heading Angle, in respect to the position of each peak on the signal plot horizontal axis.

$$LOB = \frac{\theta_{RS} - \theta_{LS}}{2} \quad (1)$$

Obtaining the resulting LOB estimate then permits us to compare the estimate directly against the known LOB to the target emitter, the position of which we surveyed when we set up the field test range. We can calculate the difference between these two bearings to produce an error signal, on which we can perform a statistical analysis that indicates the effective performance of the WIDAR system monopulse detection sensor.

The manual procedure relies on the signals analyst to visually examine the processed monopulse variance signal to identify the characteristic twin peaks produced when an active emitter is operating in the detector array beamwidth. Figure 4.18 shows

polar plots for the results of performing the manual procedure on 6 separate sensor data collection runs.

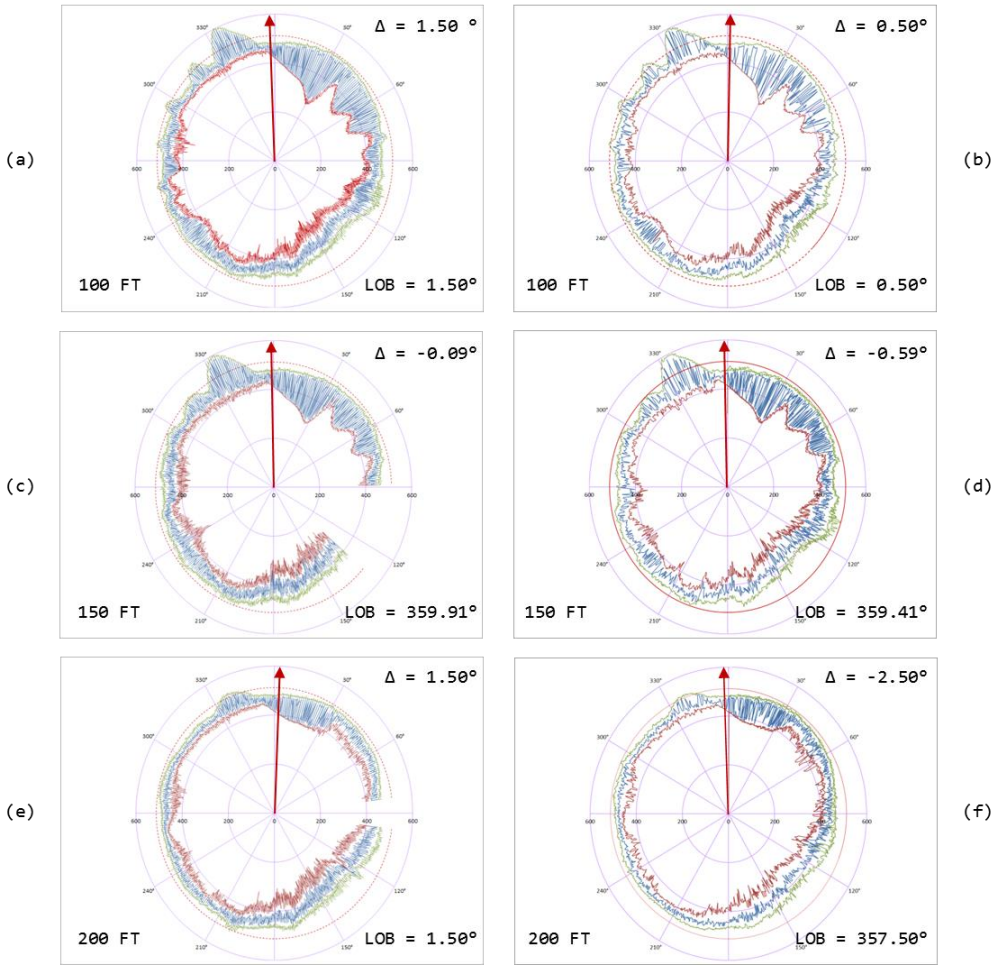

Figure 4.18. Polar Plots of Sampled Monopulse Array Measurements. The plots were created by manually performing interpolation using the Windowed Variance Peak to Peak Method. The Octave command line was used to perform the calculations, with visual inspection performed on Octave plots.

As shown by the LOB values listed in Figure 4.18, the results of using the manual peak identification show promise, the mean error for the 6 plots was .05°; well under the predicted theoretical performance of the monopulse array. In all of our experiments the true LOB was at heading 0.00°.

This method was then automated using the `FindPeaks()` command found in the Octave signal processing package [22]. A block diagram of the automated detection process is shown in Figure 4.19. Before inputting signal data into the peak detector shown in the diagram, it is necessary to first segment the monopulse ratio signal into discrete blocks of signal samples grouped by 360 ° rotations. Signal data are then run through the detector blocks, where the windowed variance calculation is performed, along with time shifting and signal amplitude normalization.

For our research we simply ran the `FindPeaks()` function against the normalized signal, after tuning the function to select peaks above a tunable amplitude parameter. In our testing, we selected .82 as the peak height parameter, and 200 samples for the minimum inter-peak spacing. This process found 100% of the 17 peak pairs in our data sample, with zero type I or type II errors. We did include logic in our detector to report any rotation segments where the peak detector located more or less than the required 2 windowed variance peak values.

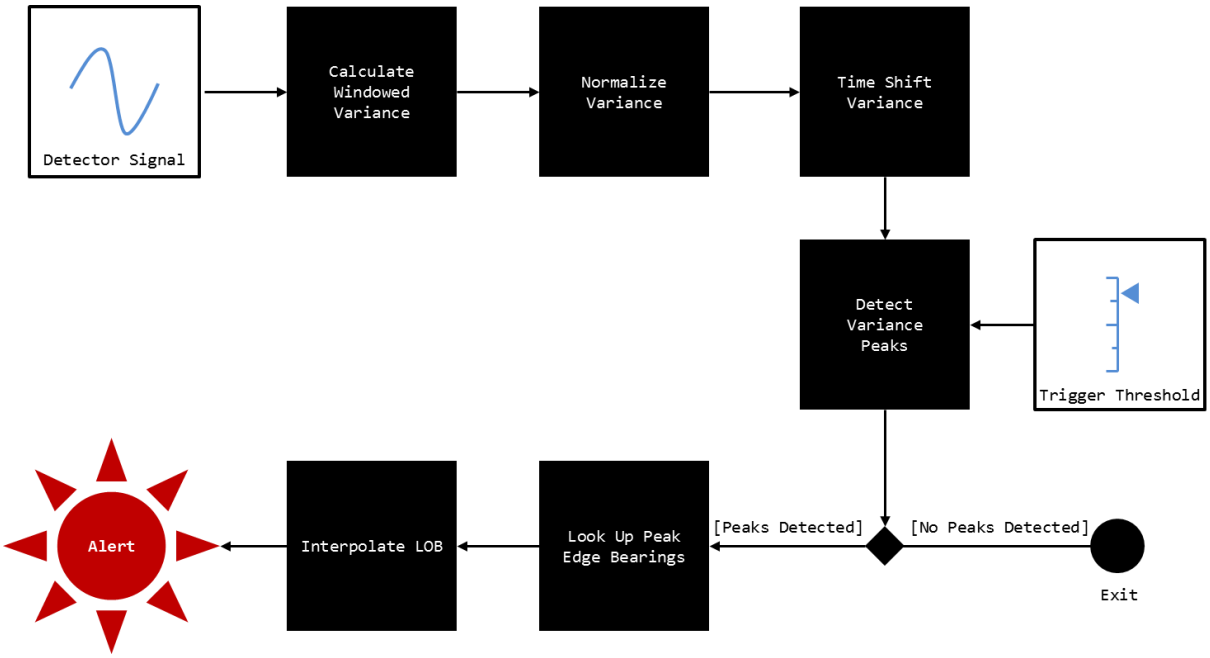

Figure 4.19. Block Diagram Detailing the Windowed Variance Peak to Peak Detector.

Statistical Analysis of the Windowed Variance Peak to Peak Detector

We processed monopulse ratio data collected from our sensor through the detector to test the LOB interpolation performance. We also desired to know whether the choice of window size used in windowed variance calculations would show any statistically significant impact on detector performance. We selected window sizes ranging from 32 to 256 (in steps of 32 samples each), and then created an Octave script that would run the detector interpolation algorithm described in Figure 4.19, incrementing the variance window size with each iteration.

Table 4.7 shows mean boresight error values ranging from -1.40° to -1.80° with variance values tightly grouped about the mean, yielding 95% confidence intervals ranging from ±.26° to ±.49 from mean boresight error values. Portions of the confidence interval overlap with the boresight error of ±1.25° predicted by monopulse theory.

Furthermore, the low mean boresight error of the monopulse array stands in stark contrast to the much larger errors observed in data collected using the sequential lobing scheme of Phase I.

Table 4.7. Statistical Analysis of Interpolated LOB Estimated Using Windowed Variance Peak to Peak Detector.

	32	64	96	128	160	192	224	256
MEAN	-1.41	-1.40	-1.45	-1.50	-1.80	-1.66	-1.64	-1.75
VAR	0.65	0.55	0.52	0.26	0.19	0.32	0.76	0.91
STDEV	0.81	0.74	0.72	0.51	0.43	0.56	0.87	0.95
CI (95%)	0.41	0.38	0.37	0.26	0.22	0.29	0.45	0.49
MEAN - CI	-1.83	-1.78	-1.82	-1.76	-2.02	-1.95	-2.09	-2.24
MEAN + CI	-1.00	-1.02	-1.08	-1.24	-1.57	-1.37	-1.19	-1.26
MIN	-2.53	-2.94	-3.09	-2.53	-2.91	-2.75	-3.35	-3.38
MAX	0.35	0.04	-0.16	-0.56	-1.10	-0.56	-0.25	0.44
RANGE	2.88	2.97	2.94	1.97	1.81	2.19	3.10	3.82
SKEW	0.90	0.41	-0.58	0.06	-0.83	0.02	-0.04	0.53
Avg Error Arc (Inches)	-44.42	-43.96	-45.45	-47.02	-56.39	-52.21	-51.61	-54.93
Avg Error/Tgt Dist (%)	-2.47%	-2.44%	-2.53%	-2.61%	-3.13%	-2.90%	-2.87%	-3.05%

Table 4.8. ANOVA Results Comparing Boresight Error Means. The null hypothesis is accepted – no significant mean inequality exists.

ANOVA – Peak-to-Peak Windowed Variance

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	2.889797059	7	0.412828	0.797871	0.590459	2.081872

It is interesting to note that similar to Phase I, ANOVA tests on the choice of window size parameter indicate that no significant mean difference is attributable to the variance calculation window size. This is an important discovery since any operational system would benefit from the performance increase achieved by selecting a smaller variance calculation window, with the aim of reducing computation complexity.

This analysis was performed on the 101,355 sample dataset we collected when the RF detection target was placed at the 150 foot Range Point. For reference we have converted the boresight angle errors shown in Table 4.7 into distance measures relative to the 150 foot target offset. These values are shown in the bottom 4 rows of Table 4.7. For example, consider the 32 point variance window size having a mean boresight error of.-1.41°. This rotational error translates to an arc distance of 44.42 inches, given a radius of 150 feet from detection sensor to the RF detection target. This value expressed as a percentage of distance to target is also shown in Table 4.7. For the 32 point window, the percentage shown is -2.47%.

Heading Interpolation Using a Matched Filter Detector

We next modified the Windowed Variance Peak to Peak detector to include a matched filter. The idea behind this change was that in a real world situation multiple targets and increased environmental noise would hinder the performance of any detector based purely on the windowed variance peak to peak method. Instead we chose to test whether a matched filter, which brings the capability to detect signals buried in noise, would also provide similar LOB interpolation performance. The modified detector block diagram is shown in Figure 4.20.

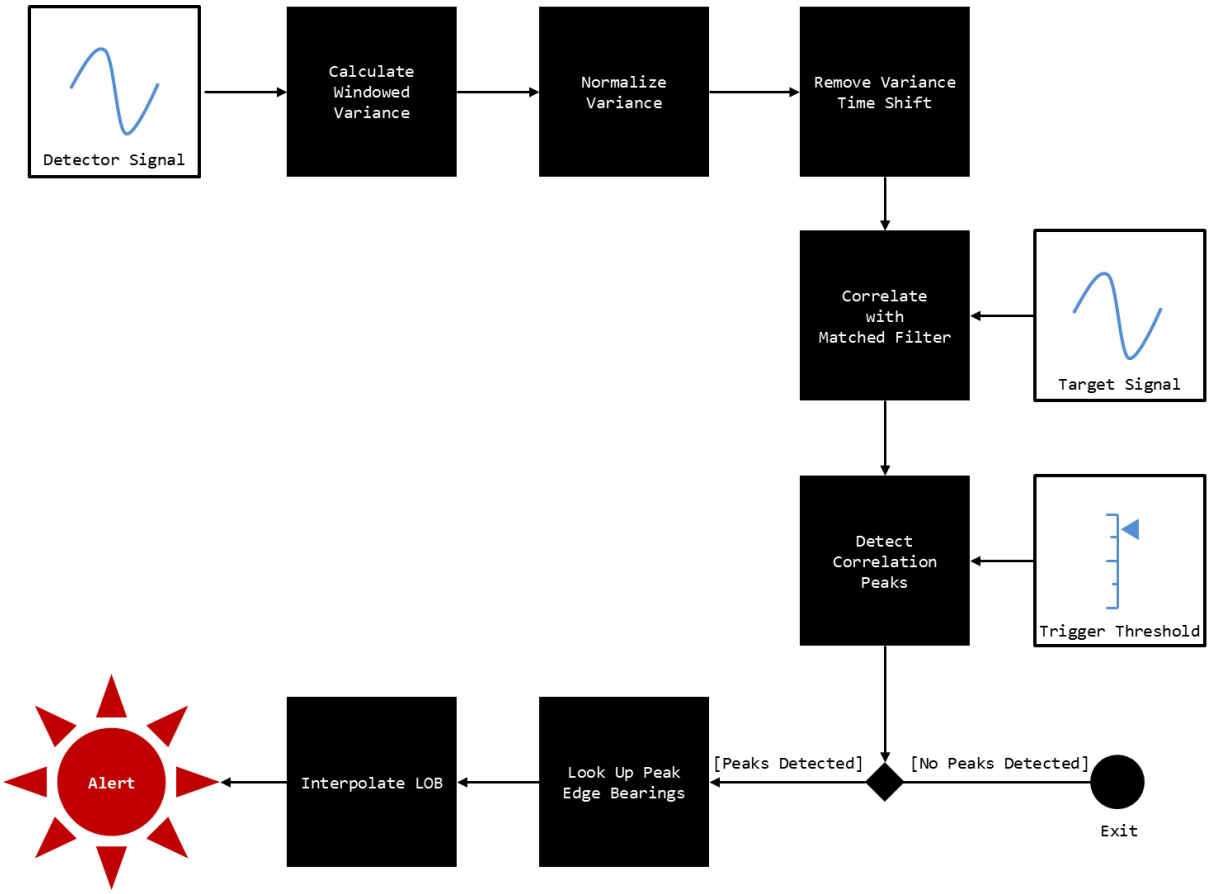

Figure 4.20. Block Diagram of the Matched Filter Detector. This is a modified version of the Windowed Variance Peak-to-Peak detector, where a Matched Filter is used to correlate signal data with a target image of the Peak-to-Peak signal pattern. Peak detection is then used to identify highly correlated regions of the signal which are then run through Bearing Lookup and Interpolation blocks to estimate a Target LOB.

In this detector we input the characteristic twin peaks, seen in the windowed variance signal when a target is active, as the target signal we want the matched filter to correlate with. There are screen shots of the matched filter being identified and selected using Octave shown in Figure 4.21. Due to how the matched filter correlation is performed in Octave, the actual target signal used in filtering was a time-reversed copy of the originally selected target signal.

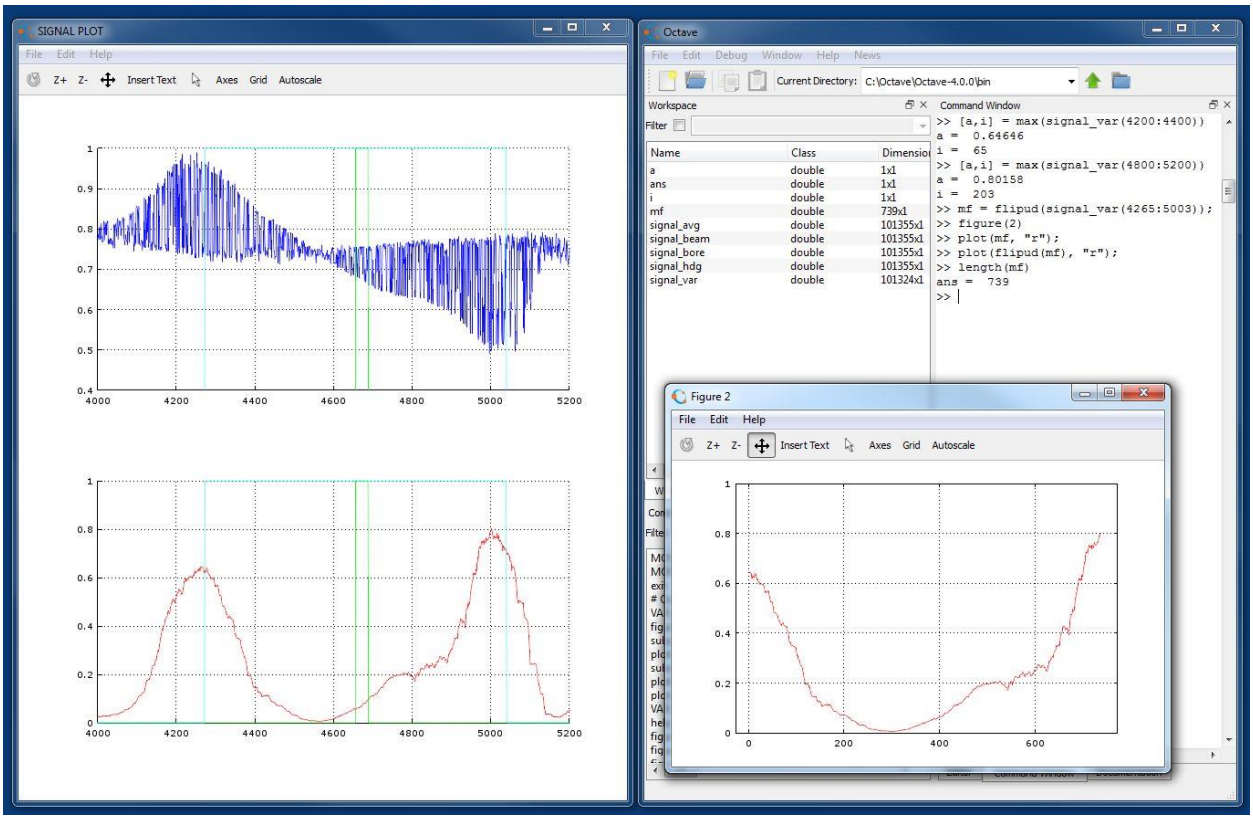

Figure 4.21. Using Octave to Identify and Isolate the Matched Filter Correlation Target Pattern.

We used the Octave `Filter()` command to run the correlation filter against the same 101,355 sample dataset we used when assessing the previous Peak to Peak detector treatment. A notable difference being that there is no longer a need to perform the rotation segmenting that was utilized in the last detector. This method could also be performed in real time on detection sensor data, assuming the proper optimizations were implemented.

The matched filter produces a separate signal containing amplitude peaks at sample points where the signal is strongly correlated with the matched filter target signal. The matched filter signal then enters a peak detector, once again implemented using the `FindPeaks()` function in the Octave signal processing package. Figure 4.22.a shows a

plot of monopulse ratio data corresponding to 3 complete detection sensor rotations. In Figure 4.22.b we show the signals output from the various stages of the matched filter detector blocks. The windowed variance of the monopulse ratio signal is plotted using the red lines in Figure 4.22.b, while the matched filter correlation signal is plotted using light-green lines.

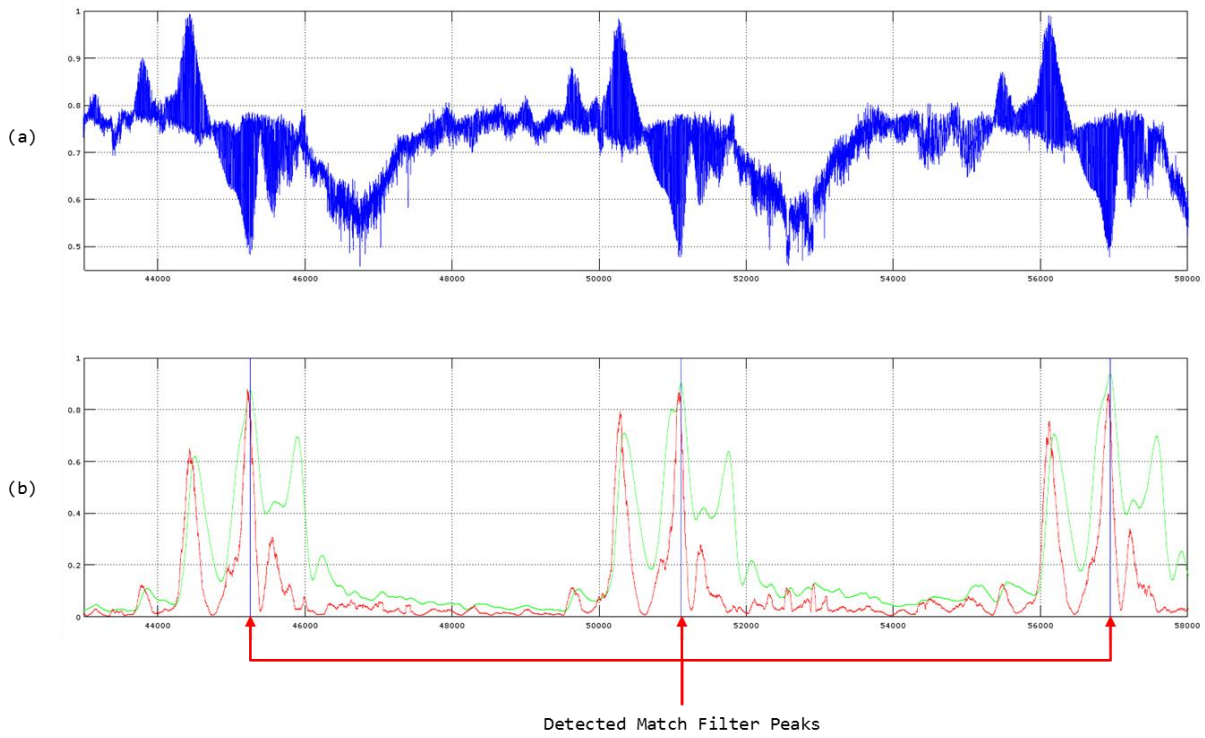

Figure 4.22. The matched filter output signal is shown in green in Figure 4.22.b. The peak detection algorithm uses a cutoff amplitude of .82 and produces a detected peaks signal. Detected peaks are marked by blue line segments in Figure 4.22.b. The index of the detected peaks are passed to the LOB calculation blocks, where the peaks are used to lookup bearing information for each peak and the corresponding bearing information for the opposite peak located at the beginning of the matched filter.

The matched filter signal also appeared to strongly correlate with any side lobes located near the variance signal main lobes, hence we had greater difficulty tuning the peak detector in this stage of our detector design. We selected filter parameters of .82 for minimum peak height, and again used 200 as the minimum sample spacing between

peaks. After careful tuning we had zero false positives, but did track several false negatives.

Matched Filter Method Statistical Analysis

We processed monopulse ratio data collected from our sensor through the detector to test the LOB interpolation performance. Since this method still integrated a windowed variance signal feeding the matched detector, we followed steps similar to the prior detector analysis, selecting window sizes ranging from 32 to 256 (in steps of 32 samples each), and then creating an Octave script that would run the detector interpolation algorithm using all these window sizes.

Table 4.9 shows mean boresight error values ranging from 0.98° to 5.86° with variance values tightly grouped about the mean except for the 256 point windowed variance treatment. The boresight error values have 95% confidence intervals ranging from ±.50° to ±1.33 from mean boresight error values. Most average values are much larger than those mean error values estimated using the Windowed Variance Peak-to-Peak detector, and larger variances are also observed.

Table 4.9. Statistical Analysis of Interpolated LOB Estimated Using Windowed Variance Matched Filter Detector.

	32	64	96	128	160	192	224	256
MEAN	0.98	2.56	1.62	2.32	3.45	3.95	5.86	4.36
VAR	1.49	1.26	0.96	1.20	1.29	2.05	1.86	6.69
STDEV	1.22	1.12	0.98	1.10	1.13	1.43	1.36	2.59
CI (95%)	0.63	0.58	0.50	0.56	0.58	0.74	0.70	1.33
MEAN - CI	0.35	1.99	1.12	1.75	2.86	3.21	5.16	3.03
MEAN + CI	1.61	3.14	2.12	2.88	4.03	4.68	6.56	5.69
MIN	-0.84	0.13	0.19	0.69	1.13	1.13	3.03	-0.06
MAX	3.72	4.63	4.16	5.54	6.38	7.35	9.19	9.19
RANGE	4.56	4.50	3.97	4.85	5.25	6.22	6.16	9.25
SKEW	0.58	-0.24	0.77	1.32	0.60	0.35	0.30	-0.13

An ANOVA test performed against the boresight mean error data indicates significant mean error differences exist between boresight errors. The ANOVA test on boresight mean error is shown in Table 4.10.

Table 4.10. ANOVA Results Comparing Boresight Error Means. The null hypothesis is rejected – at least one mean inequality exists.

ANOVA – Matched Filter on Windowed Variance

Source of Variation	***SS***	***df***	***MS***	***F***	***P-value***	***F crit***
Between Groups	458.95	7.00	65.56	6.87	0.00	2.08

Pairwise analysis using Tukey Kramer multiple comparisons indicates that the larger window sizes of 224 and 256 are significantly different than boresight errors collected using smaller window sizes. Since smaller window sizes yield boresight errors that are smaller with tighter variances we conclude that a detector employing smaller detection windows offers statistically significant improvements over larger window sizes. Figure 4.23 plots the boresight error mean and standard deviation for all windowed

variance window sizes. The results of Tukey Kramer ranged q test for multiple mean comparisons are shown in Table 4.11.

.

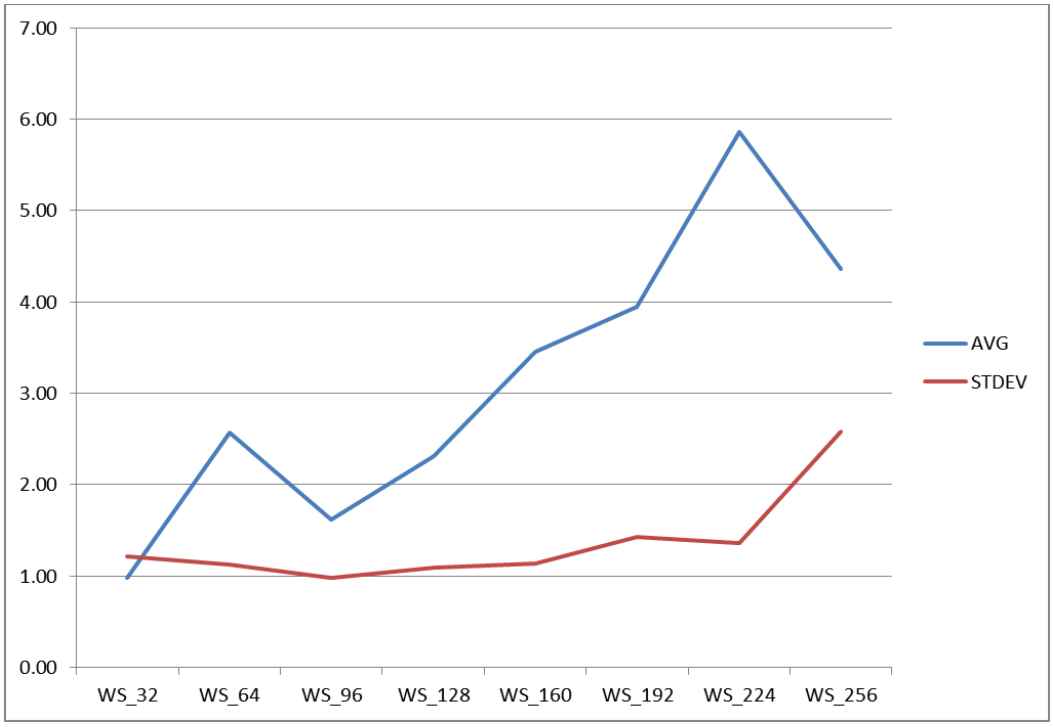

Figure 4.23. Mean and Standard Deviation Differences for Windowed Variance Window Sizes.

Table 4.11. Tukey Kramer Multiple Comparisons on Pairwise Mean Differences (Ranged q Value = 4.363; df = 128).

Window Pair	Mean Diff	Critical Range (q = 4.363)	Results
32 to 64	1.58	3.27	FALSE
32 to 96	0.64	3.27	FALSE
32 to 128	1.34	3.27	FALSE
32 to 160	2.47	3.27	FALSE
32 to 192	2.97	3.27	FALSE
32 to 224	4.88	3.27	TRUE
32 to 256	3.38	3.27	TRUE
64 to 96	0.94	3.27	FALSE
64 to 128	0.25	3.27	FALSE
64 to 160	0.88	3.27	FALSE
64 to 192	1.38	3.27	FALSE
64 to 224	3.30	3.27	TRUE
64 to 256	1.79	3.27	FALSE
96 to 128	0.70	3.27	FALSE
96 to 160	1.83	3.27	FALSE
96 to 192	2.33	3.27	FALSE
96 to 224	4.24	3.27	TRUE
96 to 256	2.74	3.27	FALSE
128 to 160	1.13	3.27	FALSE
128 to 192	1.63	3.27	FALSE
128 to 224	3.54	3.27	TRUE
128 to 256	2.04	3.27	FALSE
160 to 192	0.50	3.27	FALSE
160 to 224	2.41	3.27	FALSE
160 to 256	0.91	3.27	FALSE
192 to 224	1.91	3.27	FALSE
192 to 256	0.41	3.27	FALSE
224 to 256	1.50	3.27	FALSE

Phase II Conclusions

Both LOB Detectors we analyzed in Phase II out-performed our baseline metrics collected from the sequential lobing scheme used in Phase I. The Windowed Variance Peak-to-Peak detector provided consistent mean boresight errors, and standard deviations

of 1° or less for all variance window sizes. The performance of this detector very nearly matched that of the ±1.25° that textbook theory estimated for a device with the beamwidth of the antennas that were used.

In our experiments with the matched filter detector we observed similar boresight errors, at least when those estimates were performed using smaller variance calculation windows. The overall mean and overall variance for boresight error for this detector were higher than the detector using only the window variance peaks for LOB estimation. However, the addition of a matched filter correlator in this detector should offer superior detector performance in noisy environments, or in multiple target environments. This is left as a question for future research.

Recommendations for Continued Exploration

We explored the capabilities of a detection sensor designed to detect and estimate a Line of Bearing for targets operating in the 2.4 GHz ISM band. While the results of testing the monopulse array concept do show promising LOB estimation capabilities having accuracies that could permit a system to detect and spatially attribute a wireless attacker, many questions remain unanswered. We have the following recommendations for those wishing to conduct further research in this area:

- Explore Software Defined Radio (SDR) tools for next generation of detection sensor. Devices like the HackRF [23] and the USRP-B210 [24] offer technical advances that far exceed our current generation detection sensor in terms of sensitivity and detection bandwidth. These products can operate in RF bands from baseband to 6GHz with 70MHz sampling bandwidths. An SDR toolchain would support true Sum/Difference on complex signal samples to be

performed, enabling the more commonly implemented forms of monopulse processors to be developed, purely in software. Furthermore, when couple with tools such as GNU Radio [25], the capability to sample spectrum using monopulse methods *and* demodulate wireless network traffic becomes available, presenting many advancement opportunities.

- Experiment with Phased Array detection sensors. Utilizing the SDR tools just mentioned, an interesting research area would be to implement a phased array radar system using beam/null shaping principles and techniques. Research such as this could lead to far more sensitive detectors than those we presented, and could potentially eliminate the need for a mechanically steered chassis.
- Research the potential for miniaturizing an ISM band monopulse sensor to be a drone payload. RF environmental surveys could be performed using a single drone, where our system currently would depend on a network of fixed position cooperative sensors to estimate angular LOB to any detected target. A mobile sensor could integrate readings from multiple spatial locations which could dramatically increase location estimation accuracy.
- Many other avenues of exploration exist, such as multiple target detection and tracking, and an entire treatment of counter measures an attacker could employ to defeat detection, such as deployment of decoys and randomized slow-speed wireless traffic patterning.

References

[1] Cisco, "Cisco Mobility Services Engine," [Online]. Available: http://www.cisco.com/en/US/products/ps9742/index.html. [Accessed 08 2013].

[2] J. Werb and C. Lanzl, "Designing a positioning system for finding things and people indoors," *IEEE Spectrum*, vol. 35, no. 9, pp. 71-78, 1998.

[3] V. Bhargava and M. L. Sichitiu, "Physical Authentication through Localization in," in *IEEE GLOBECOM*, 2005.

[4] SeatlePI, "Feds: Wi-Fi hacking burglars targeted dozens of Seattle-area businesses," 19 09 2011. [Online]. Available: http://www.seattlepi.com/local/article/Feds-Wi-Fi-hacking-burglars-targeted-dozens-of-2178421.php#ixzz1jwWj2LWH. [Accessed 08 2013].

[5] The Seattle Times, "High-tech hacker gets almost 8 years in $3M Seattle theft ring.," 13 07 2012. [Online]. Available: http://seattletimes.com/html/localnews/2018684238_hacker14m.html. [Accessed 08 2013].

[6] US Department of Justice, "Retail Hacking Ring Charged for Stealing and Distributing Credit and Debit Card Numbers from Major U.S. Retailers.," 05 08 2008. [Online]. Available: http://www.justice.gov/opa/pr/2008/August/08-ag-689.html. [Accessed 08 2013].

[7] D. J. Gieseman and T. E. Daniels, "Countering the Parking Lot Attack – Design for a Detection System Employing Monopulse Radar Methods to Detect and Spatially Attribute RF Targets in the 2.4 GHz ISM Band," Iowa State University, Ames, IA, 2015.

[8] D. J. Gieseman and T. E. Daniels, "A Strategy for Facility Wireless Attack Detection using Cooperative Mechanically-Steered RF Detection Sensors," Iowa State University, Ames, IA, 2015.

[9] D. R. Rhodes, Introduction to Monopulse, Artech House, 1980.

[10] S. M. Sherman and D. K. Barton, Monopulse Principles and Techniques, Artech House, 2011.

[11] Analog Devices, "AD8302.pdf," [Online]. Available: http://www.analog.com/static/imported-files/data_sheets/AD8302.pdf. [Accessed 08 2013].

[12] Atmel Corporation, "ATmega168," [Online]. Available: http://www.atmel.com/devices/atmega168.aspx. [Accessed 08 2013].

[13] Texas Instruments, "CC2500," [Online]. Available: http://www.ti.com/product/cc2500. [Accessed 08 2013].

[14] Austrian Microsystems, "AS5306 Linear Position Sensor," [Online]. Available: http://www.ams.com/eng/Products/Magnetic-Position-Sensors/Linear-Incremental-Magnetic-Position-Sensors/AS5306. [Accessed 08 2013].

[15] A. I. Leonov, K. I. Fomichev, W. F. Barton and D. K. Barton, Monopulse Radar, Artech House, 1986.

[16] "LORCON - A sane interface for crafting and transmitting packets on wireless interfaces across multiple platforms.," [Online]. Available: https://code.google.com/p/lorcon/. [Accessed 08 2015].

[17] M. S. Gast, 802.11 Wireless Networks: The Definitive Guide, 2nd Edition, O'Reilly Media, 2005.

[18] Wireshark Foundation, "Wireshark," [Online]. Available: https://www.wireshark.org/. [Accessed 08 2015].

[19] wireless.kernel.org, "About mac80211," [Online]. Available: https://wireless.wiki.kernel.org/en/developers/documentation/mac80211. [Accessed 08 2015].

[20] "Radiotap," [Online]. Available: http://www.radiotap.org/. [Accessed 06 2015].

[21] Wireshark Foundation, "Display Filter Reference: IEEE 802.11 Radiotap Capture header," [Online]. Available: https://www.wireshark.org/docs/dfref/r/radiotap.html. [Accessed 08 2015].

[22] GNU Octave, "Octave-Forge," [Online]. Available: http://octave.sourceforge.net/. [Accessed 08 2015].

[23] Great Scott Gadgets, "The HackRF One," [Online]. Available: https://greatscottgadgets.com/hackrf/. [Accessed 08 2015].

[24] Ettus Research, "USRP B210 Board," [Online]. Available: http://www.ettus.com/product/details/UB210-KIT. [Accessed 08 2015].

[25] J.-P. Lang, "Welcome to GNU Radio," [Online]. Available: http://gnuradio.org/redmine/projects/gnuradio/wiki. [Accessed 08 2015].

CHAPTER 5. A STRATEGY FOR FACILITY WIRELESS ATTACK DETECTION USING COOPERATIVE MECHANICALLY-STEERED RF DETECTION SENSORS

A paper submitted to *9th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2016)*

D. J. Gieseman[1,2] and T. E. Daniels[1]

Abstract

Our research is focused on the development of a system of sensor devices employing monopulse radar methods to detect and spatially attribute RF targets transmitting in the 2.4 GHz ISM band. The intent is to deploy this type of system on or about a facility to protect the premises from overt or covert cyber-attacks exploiting wireless vectors. In this paper, we seek to focus on deployment and operational questions we have regarding our detection system, with the aim of arriving at a coherent strategy for system operation. To answer such questions, we present a simulation tool with the purpose of exploring candidate sensor deployment and operational detection schemes. We present simulation scenarios which vary sensor placement, modes of sensor operation, and sensor mechanical capabilities. Quantitative analytics accompany each scenario. We then present a study of position estimation error, where we program detection sensors to simulate the mean Line-of-Bearing estimation errors we obtained

[1] Graduate Student and Assistant Professor, respectively, Department of Electrical and Computer Engineering, Iowa State University.
[2] Primary researcher and author.

from a prototype sensor during real field tests. We conclude with several recommendations for an effective protection strategy, based on results obtained from our simulations and scenario analysis.

Introduction

In previous work, we presented our design for a system of detection devices employing monopulse radar methods to detect and spatially attribute RF targets transmitting in the 2.4 GHz ISM band [1], [2]. We designed the detection sensor of that system with the capability to assign a Line-of-Bearing (LOB) to any RF targets actively emitting within the operational range of the sensor [3]. In conjunction with the system design and concept of operation detailed in that study, we also presented an actual implementation of a research prototype.

The sensor prototype featured a mechanically steered antenna array alongside integrated monopulse processing hardware and software. The intent was to deploy a cooperative network of these sensors about the external vicinity of any facility we desired to protect from stand-off wireless attacks. The sensors would scan the external grounds of the facility while looking for active wireless transmissions. They were intended to provide the increased situational awareness necessary to support decision tools that could make a determination whether any detected activity was part of an unauthorized network intrusion. While we described a conceptual deployment scheme indicating the rudimentary placement of detection sensors, we did not provide an in-depth treatment or recommend any strategy for effective device deployment and operation on or around the Facility Under Protection (FUP).

In a follow up paper, we discussed the results of performance tests conducted using the research prototype in a tightly controlled field setting [2]. During field tests, the prototype sensor showed the promise of producing reasonable target LOB estimates; however, we also concluded that many technical challenges still lay ahead before such devices were deemed effective for real world facility protection. The data obtained during our system field trials provided insight into how well an actual device might perform in a real facility protection context, in terms of performance and device accuracy.

In this paper, we put these data to valuable use; feeding them as inputs into a parametric simulation tool we developed to model our detection system, enabling us to research questions about how best to deploy and operate a system with the capabilities we designed. We show how the simulation environment can be a useful tool, providing us the liberty to set aside many of the technical challenges arising from the physics of real world RF detection; instead permitting us to focus on the development of a strategy for facility protection using our system. Any strategy research can become very broad if not constrained. To maintain focus in our work, we chose to constrain our research to questions aimed at guiding operational and deployment aspects of our system:

- How well would a network of our prototype sensors (as implemented) actually perform?
- What are the marginal performance gains achieved when mechanical enhancements are made to our detection sensors?
- How many sensors should be deployed about a facility to ensure adequate protection?

- Should sensors be mast mounted a distance away from the FUP, or should they be concealed on the rooftop of the FUP?
- Can modulating the WLAN data rate play a useful role in ensuring full attacker geo-location and spatial attribution?
- What is the mean position estimate error we can expect when multiple detection sensor LOB estimates are integrated to triangulate a detected attack?

By using our simulation tool to explore these questions we begin to gain important knowledge about how best to protect a facility from attack using our system. We next present the simulation architecture in more detail, prior to embarking on an examination of the actual simulation experiments we performed to explore the strategy questions we outlined above.

Simulation Model Architecture

Our simulation model is a custom written tool implemented using the Microsoft Visual Studio .NET Framework 4.5 toolchain. The software features a reusable object model that is both scalable and configurable, implementing simulation constructs for all of the pieces at play in our wireless attack research. Abstractions exist to model and simulate the facility targeted by wireless attacks, the detection sensors that defend the facility, and the adversary who carries out wireless attacks. In addition to those primary model objects, many other classes were implemented for use in the development of realistic wireless attack and defend simulation environments: network adapters, servers, workstations, and laptops, and most importantly, the data files residing on the files systems of these simulated devices. In all of our simulation scenarios, these data files are

the primary target of the adversary, who seeks to steal the files without being detected by our sensor system.

The simulation model integrates a custom geographic information system (GIS) engine that is capable of spatially referencing and visually displaying any of the sensor placement and adversary attack scenarios we wish to simulate in this research. Figure 5.1 shows a screen capture taken from a running simulation run.

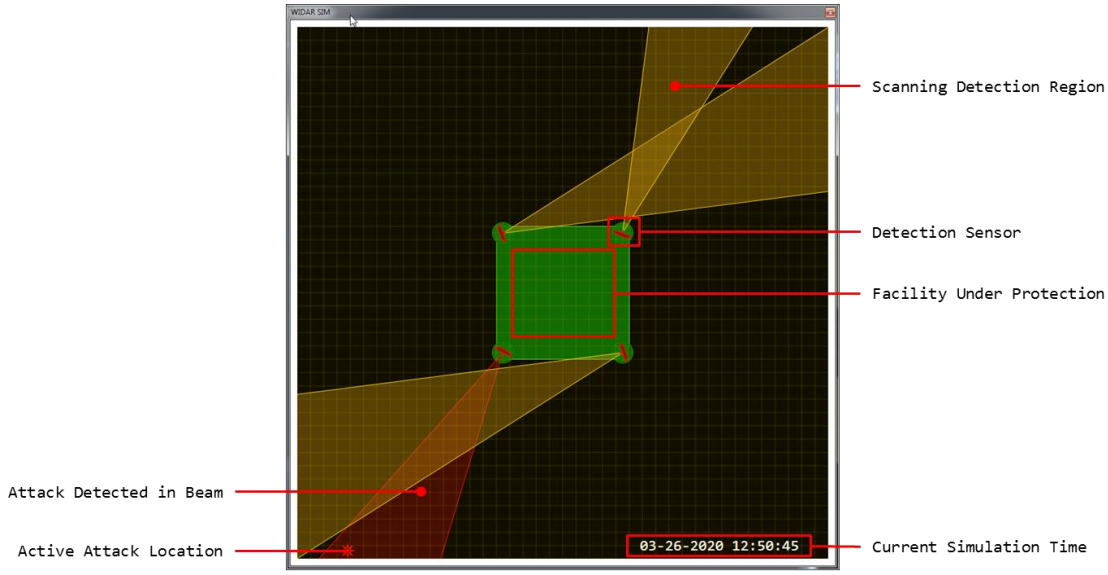

Figure 5.1. Screen Capture from Running Simulation Software. In this scenario, four wireless detection sensors (shown as red directional antennas) are actively scanning for wireless attacks against a simulated facility, represented by the green rectangle. The sensor on the lower left corner of the facility has detected an active attacker within the scanning beam of the detection sensor. In our simulation, detected attacks are indicated by changing the scanning beam color from gold to red. The attacker is shown represented as a flashing star, in the lower left hand corner of the simulation window. The current simulation timestamp is displayed in the lower right corner of the simulation window.

Simulating Random Arrivals using a Discrete Poisson Probability Model

A core requirement for our simulation tool is the capability to generate randomly occurring attacks against the wireless access point of the facility. These attacks must be randomly distributed, both temporally and spatially. Our simulation tool integrates a Poisson Arrival Generator to support randomized discrete arrivals. An example of a

random discrete arrival that we would generate for the simulation would be the number of attacks per year against a simulated facility, or the number of computers operated on a simulated facility network. We based the randomized arrival engine on prior work we developed as part of external research [4]. The simulation engine generates random arrivals using a pseudo random number generator (PRNG) as the input to a Poisson distributed Cumulative Distribution Function (CDF), as shown in Figure 5.2.

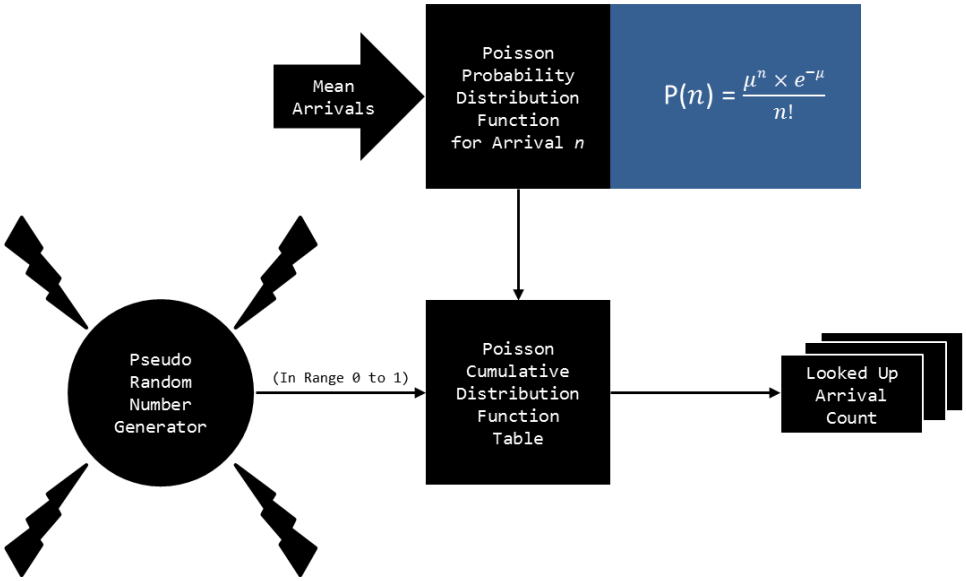

Figure 5.2. Conceptual Diagram of the Simulation Tool Poisson Arrival Engine. The PRNG generates a pseudo random number which is used to choose a discrete value from a Poisson CDF. The CDF table is generated at simulation startup based on a Mean Arrivals parameter dependent upon the average number of arrivals per time period for some simulation artifact.

In the figure, the CDF depicted is generated during the initialization of the simulation by first iterating through potential discrete arrival possibilities and calculating the Poisson distributed probability for each discrete arrival. The only input to the Poisson probability function is the mean arrivals we desire to see for some aspect of our simulation. Equation 1 shows the Probability Distribution Function (PDF) for the Discrete Poisson Distribution.

$$P(n) = \frac{\mu^n \times e^{-\mu}}{n!} \quad (1)$$

The probability that is generated for each discrete possibility during each iteration step is added to a summary table maintaining a running sum of the cumulative discrete arrival probability. Each row in this table is indexed by the discrete arrival count value corresponding to the probability that was just added to the running sum. This process continues until the CDF sum saturates near a very close single precision floating-point approximation of cumulative probability equal to one.

At this point, any probability value we generate, for example a random probability generated by the PRNG, can be used to reverse-lookup a discrete arrival value from the CDF summary table. We utilize this method; randomly generating a probability value, followed by looking up the discrete arrival value for that random probability, to generate many types of random arrival events in our simulation. The PDF and CDF distributions are plotted in Figure 5.3. Figure 5.3.a and Figure 5.3.b display the PDF and the CDF for a Mean Arrivals value of 5. Figure 5.3.c shows a plot of 10000 arrivals looked up from the CDF Summary Table using PRNG values as index. After only 10000 samples, the PDF can clearly be seen emerging from the data.

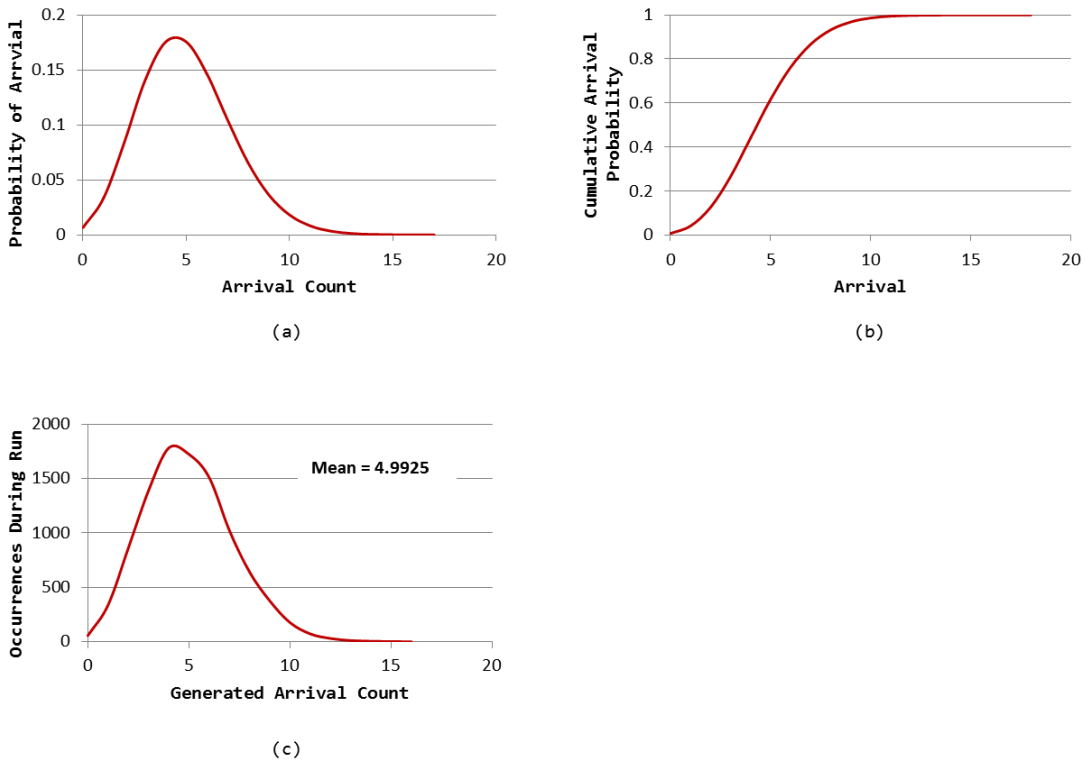

Figure 5.3. Discrete Poisson Probability Distribution Function and Cumulative Distribution Function. Figure 5.3.a shows the Probability Distribution Function (PDF) for Mean Arrival ($\mu = 5$). An iterative procedure is used to generate probabilities for discrete arrival values. Figure 5.3.b shows the Discrete Poisson Cumulative Distribution Function (CDF) for Mean Arrival ($\mu = 5$). A value generated between 0.0 and 1.0 by the PRNG is used to perform a reverse lookup onto the curve. The lookup value determines the number of discrete arrivals that are randomly generated for a given simulation state. Figure 5.3.c shows a comparison plot depicting Mean Distribution from 10000 Runs of the Random Arrival Generator. The Run Generator was configured with a Mean Arrival Rate of 5. The results should closely align with the Probability of Arrival Distribution shown in Figure 5.3.a, and they do, with an observed Mean Arrival Rate of 4.9925 after 10000 samples. In the Discrete Poisson Distribution, the population mean and population variance are always equal to the input mean arrival rate parameter.

Simulation Logical Objects

We next turn to describing the simulation tool software. The software architecture is logically divided into object classes responsible for implementing specific detection system functionality. Our objective was to equip the software with model elements for each of the important classes of objects in our detection system. Important components are the Facility we wish to protect with the detection system, the Detection Sensors we will deploy on the facility, and the Adversary, who carries out facility wireless attacks.

Simulating the Facility Under Protection

The Facility Under Protection (FUP) or just simply "facility" serves as the target of attack during our simulation runs. The facility is also the logical container for a variety of simulation sub-system components as shown in Figure 5.4. At simulation startup, a single facility object is created. The Poisson Arrival Generator is then used to generate a random number of internal servers, based on a mean number of computers the simulated facility is programmed to contain. The mean facility computer count is input by the user before the simulation is run. Each server, in turn, generates a random internal file system, containing randomly sized data files, each with a randomly assigned target value.

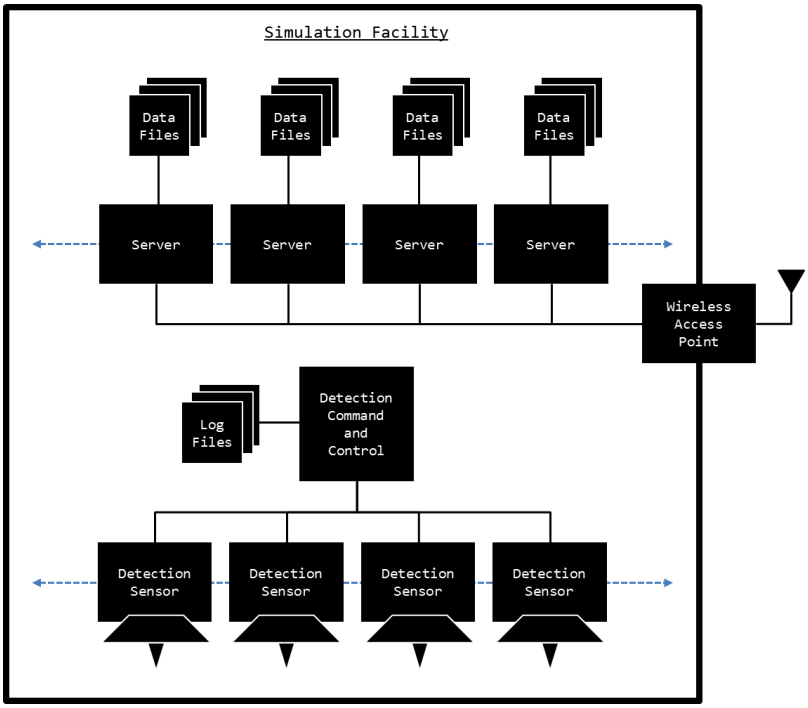

Figure 5.4. Simulation Facility Architecture. The facility logically contains the detection system – with detection sensors and command and control functionality – as well as a wireless access point vulnerable to external wireless attack, and a network of servers housing sensitive systems and data which are targeted for attack by the simulation adversary.

We assigned classification values to the data files in the simulation model to simulate the variety of files that are commonly found on computer networks, from mundane system files, to highly sensitive commercial databases, to the classified secrets maintained by nation states. Enabling a data classification system in the simulation permits the running of scoring schemes against simulation results to compute, for example, a post-simulation severity score for each successful attack against the FUP. Our model currently contains the following target file classifications:

- Spam = 0
- Unclassified_Non_Critical = 1
- Unclassified_Critical = 2
- Classified_Secret = 3
- Classified_Top_Secret = 4
- Classified_Ultra_Secret = 5
- Classified_Apocalyptic_Event = 6

To expose the necessary facility vulnerabilities described by our threat model [1], we equip the target facility with a Wireless Access Point (WAP) and provide for the WAP to be topologically linked with the common network resources shared by the facility file system servers.

Additionally, and most importantly, the FUP in our simulation is initialized with a set of detection sensors. The number of detection sensors and their geometric placement on the actual facility are configurable and initialized at simulation startup. Each detection sensor reports to a Detection Command and Control (DCC) object in our simulation model, as shown in Figure 5.4. The DCC serves as the central point for attack detection

data collection, and is also a convenient location in the software model to place data logging functionality that is valuable for assessing the performance of simulation scenarios.

Simulating Detection System Sensors

The simulation detection sensors are designed to model the mechanical and physical behavior of our real world prototype, while at the same time implementing the monopulse lobing detection capabilities our simulation is premised upon. Figure 5.5 displays key details for a typical simulation detection sensor. For our simulation, each detection sensor has parametrically programmable rotation RPM and scanning beamwidth. The initial boresight heading and the detection scan rotation direction (clockwise or counter-clockwise) can be programmed independently for each simulated sensor.

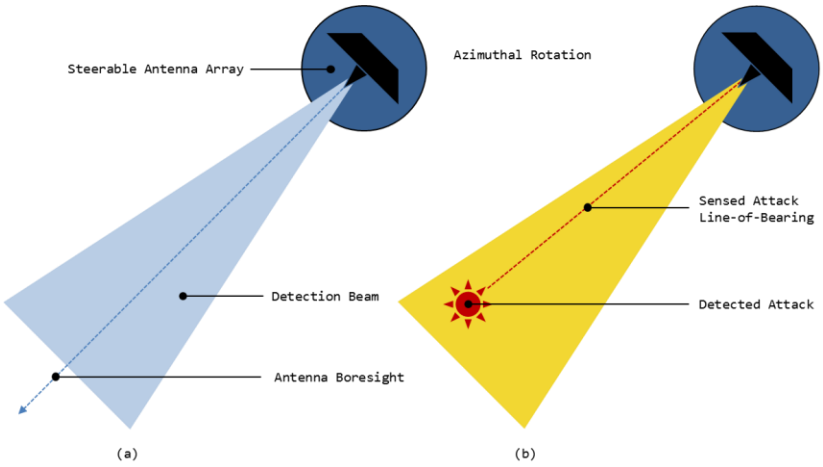

Figure 5.5. Conceptual Diagram of Simulation Detection Sensor and Beam Pattern. One or more detection sensors are actively rotated during the simulation. Each sensor has configuration parameters controlling the rate of azimuthal rotation, the detection beamwidth, starting boresight bearing, and operating mode (PAN_THEN_SCAN or CONTINUOUS). In Figure 5.5.b, the detection sensor has detected an active attack. Every detection sensor can sense any wireless transmissions active within the detection beam boundary. The detection sensor returns a Line-of-Bearing estimate for any detected target. The estimate is subject to random estimation error about a programmable mean boresight error. The LOB estimate is relayed to the detection command and control system.

In order to accurately model our functional prototype, we included several behaviors in the simulation sensors that control how detection scanning is performed. Our prototype suffered from electrical motor noise issues on the control data bus when the device was actually rotating [1]. This physically prevented our real world device from actually performing a detection scan while at the same time actively rotating the array. We termed this mode of operation "Pan Then Scan", and included functionality to model Pan Then Scan in our simulated detection sensors. The Pan Then Scan mode featured several additional configuration parameters, allowing the user to specify the simulation time duration of the pan operation, the detection scan operation, as well as the number of degrees to step the boresight heading during the pan portion of the Pan Then Scan mode.

Of course, since simulated detection sensors do not physically suffer the actual electrical noise issues inherent to our test implementation, we also included a Continuous Scan mode of operation in our sensor model. We designed our model to include both behaviors with a comparison scenario in mind. The scenarios section of this paper details the performance comparison results observed during our simulated comparison of the Pan Then Scan Mode and Continuous Scan Mode of operation.

Simulating the Adversary

The attacking adversary in our simulation directs wireless attacks against the FUP. Attacks are conducted using a simulated laptop, connecting to the target facility using a Wireless LAN Adapter. The attack laptop WLAN Adapter has a configurable mean file download data rate. Figure 5.6 depicts the conceptual model describing the Simulation Adversary. At simulation startup, the simulation runtime randomly generates the number of attacks that will occur during the course of the simulation run time.

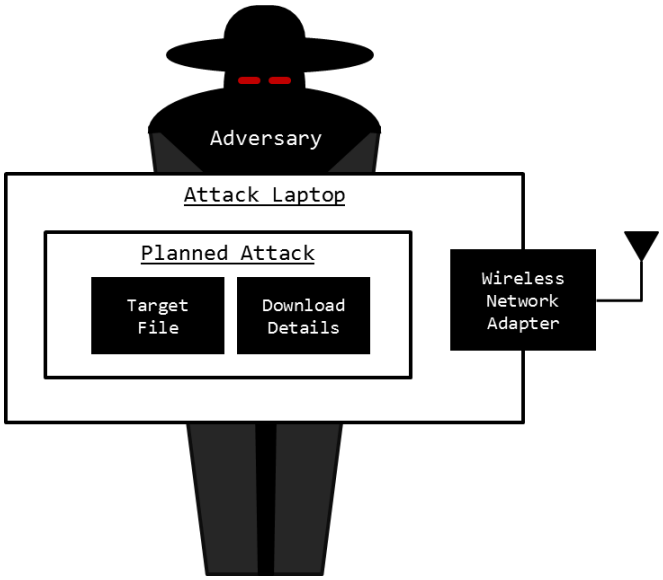

Figure 5.6. The Simulation Adversary

The randomized attack times generated by the simulation runtime then drive the creation of a single adversary simulation object for each random attack time. The adversary is initialized with the time of attack and the facility to target. The adversary randomly generates a skill level for herself, and then proceeds to plan an attack, based on the internally generated skill level, against the target facility. We have included the following Adversary Skill Levels in our simulation model:

- LookingForWiFiToCheckEmail = 0
- ScriptKiddie = 1
- DisgruntledFormerEmployee = 2
- SysAdmin = 3
- PenTester = 4
- Ninja = 5
- NationState = 6

Skill level drives the attack planning strategy executed by the adversary. The higher the skill level assigned by the simulation system, the more likely the adversary is to target sensitive data and execute increasingly sophisticated attacks. For example, if an attacker is generated with the relatively low level of Script Kiddie (level 1 attacker), she may be more inclined to simply randomly select files, and may spend excessive amounts of time actively connected to the FUP WLAN system. Continuing the example, an attacker assigned the Nation State (level 6 attacker) may carefully select the highest value target files, and may also choose target files which are downloaded the fastest, to minimize detection risk. Our adversary generator can be configured to preset the skill level of the attacker, or to parametrically generate a skill level from an Average Attacker Skill Level value set at system startup. This mean value will then be input into the simulation system Poisson Arrival Generator, and a random skill level is then assigned to the adversary.

The Simulation Clock

The entire simulation progresses from a start time to an end time by means of a simulation clock commonly referenced by all simulation entities. The simulation clock is the engine that causes the simulation objects to change state from one time instant to the next. The clock functions by raising a software tick event. Each element holding a reference pointer to the simulation clock is notified of the clock tick event, and takes action dependent upon the programmatic behavior setup for that object. Figure 5.7 illustrates conceptually how the simulation clock relays clock tick events to simulation objects.

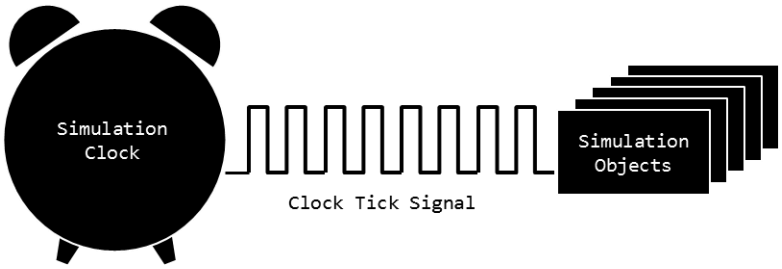

Figure 5.7. Simulation Clock – All Simulation Objects maintain an internal reference to the Master Simulation Clock. Upon receiving a Clock Tick Signal, each simulation object behaves according to the internal programming logic of the simulation object. For example, a detection sensor might rotate an amount based on the time since the last clock signal, or an adversary might launch a pre-planned attack.

The user can configure the clock to raise tick events with any level of granularity desired for simulation purposes. We have selected a 20ms time base for our research runs. Therefore, each clock tick increments the simulation state by 20ms per clock tick. The interval between clock ticks is also configurable, allowing for simulation time to be sped up or slowed down as appropriate. Also, the direction of simulation time flow is user changeable, permitting the simulation to be run in reverse or forward time change directions.

The simulation clock is initialized at simulation startup with a start time and end time in system date format. The simulation clock can be positioned programmatically to any time point, in between the start end time span. We have somewhat arbitrarily chosen January 01, 2020 to December 31, 2020 as the start and end times for all of our simulation runs.

Positioning the simulation time will result in all simulation child objects taking on the state reflective of that time point. For example, the rotating detection sensors will calculate the current detection boresight heading for any time point provided by the simulation clock. This is possible since each detection sensor is aware of the initial

startup boresight heading, and our simulation makes the assumption that rotation rates are constant, therefore the future boresight bearing can be deterministically computed for any simulation time point.

Simulation Scenarios

Scenario 1 – Simulating the Prototype Detection Sensor Implementation

The objective for Scenario 1 was the establishment of baseline performance parameters for the detection system. For the baseline, we chose to model the prototype sensor we created for field testing our system. This was driven by an interest in quantifying how that sensor would actually perform in a real facility protection environment. This experiment has two treatments (Treatment I and Treatment II), each simulating a single detection sensor operating in two different scanning modes.

For Treatment I, we configured the simulation software to operate a single sensor with performance characteristics matching that of our prototype implementation. We previously described how the prototype sensor could only operate in the Pan Then Scan mode of operation, due to electrical noise issues. The array steering sub-system on that sensor supported rotation scan speeds of 6RPM.

For Treatment II, we configured the simulation software to operate the same detection sensor, keeping all parameters the same, except for enabling a continuous scanning capability. We hypothesized that there would be a significant performance gain, in terms of the number of attacks detected by the system. Key parameters for Treatment I and Treatment II are listed in Table 5.1.

Table 5.1. Treatment Parameters for Scenario 1 Comparison Study. In Scenario 1, we wish to identify whether there is a performance gain achieved when a single detection sensor is upgraded from Pan Then Scan mode of operation to a Continuous detection scanning capability. The different Scan Modes are highlighted in yellow.

Scenario Parameter	Treatment I	Treatment II
Simulation Duration	1 Year	1 Year
Simulation Runs	538	538
Facility Average Attacks Per Year	5	5
Detection Sensors Deployed	1	1
Detection Sensor RPM	6	6
Detection Sensor Beamwidth	25	25
Detection Sensor Scan Mode	PAN_THEN_SCAN	CONTINUOUS
Detection Capability	LINE_OF_BEARING	LINE_OF_BEARING

Results and Analysis

We performed 538 trials of the simulation baseline. We then modified the baseline parameters to cause the simulated detection sensor to operate in continuous scanning mode of operation. This parameter change is shown in the highlighted section of Table 5.1. Results comparing the two treatments are shown in Table 5.2. Table 3 shows two-sample t-test statistics we ran to compare sample means from the two treatments. Our expectations are that both sets of trials should have statistically equal attack arrival means, but that we would observe a statistically significant performance gain when changing sensor scan modes from Pan Then Scan to Continuous.

Table 5.2. Comparison of Scenario 1 Treatment Results (n = 538).

Comparison Result	Treatment I	Treatment II	Delta
Total Attacks Conducted	2736	2711	-25
Total Attacks Detected	394	858	464
Observed Attack Mean	5.09	5.04	
Observed Detection Mean	0.73	1.59	0.86
Detection Percentage	14.40%	31.65%	17.25%

Table 5.3. Two Sample t-Test Results for Observed Attack and Observed Detection Means.

	Observed Attack Mean	Observed Detections
t-Stat	0.34	10.54
t-Critical	1.96	1.96

Our results show that between the two trials of 538 runs each, there were 25 fewer attacks in the second trial. Even though there were marginally fewer attacks, there was an increase of 17.25% in the rate of detection for the CONTINUOUS mode of detection scanning, versus the PAN_THEN_SCAN mode of detection scanning. The t-statistic far exceeds the critical value set for the detection mean comparison study, allowing us to safely conclude that the CONTINOUS mode of operation represents a statistically meaningful performance improvement. This means that our real world sensor would benefit from engineering improvements decreasing motor noise on the power bus sufficient to allow for continuous scan operation.

Still even after enabling continuous detection scanning, a single sensor detecting only 31.65% of the total attacks conducted against the protected facility represents poor performance. Nearly 7 in 10 attacks still succeed against our hypothetical facility! We next turn to increasing the amount of sensors deployed about the facility, to assess whether there are significant performance gains when sensors become cooperative.

Scenario 2 – Simulating a Network of Continuously Scanning Detection Sensors

Continuing our previous track of research, we modified the initial conditions of our simulation model, adding four detection sensors to our facility. Adding four sensors increases the simultaneous coverage area for facility protection detection scheme, along with adding an additional powerful capability: if an attack is now detected by two or

more sensors, the system can then employ triangulation methods to precisely geo-locate the position of the attack origin. To reflect this capability enhancement, we add metrics recording geo-location events to all simulation analysis from this point forward.

As in Scenario 1, this experiment has two treatments (Treatment I and Treatment II). The sensor configuration for each of the treatments is identical in this scenario; only the number of simulation sensors deployed is varied. In Treatment I, we deploy the single sensor, operating with continuous scan capabilities. In Treatment II we deploy 4 identical sensors, spatially arranging them to be located on the 4 corners of the of the simulated FUP rooftop. The simulation parameters for the experiment treatments are shown in Table 5.4.

Table 5.4. Treatment Parameters for Scenario 2 Comparison Study. In Scenario 2 we wish to quantify whether there is a performance gain achieved when, instead of a single detection sensor, the detection system employs a cooperative network of detection sensor. The different scan modes are highlighted in yellow.

Scenario Parameter	Treatment I	Treatment II
Simulation Duration	1 Year	1 Year
Simulation Runs	538	538
Facility Mean Attacks Per Year	5	5
Detection Sensors Deployed	1	4
Detection Sensor RPM	6	6
Detection Sensor Beamwidth	25	25
Detection Sensor Scan Mode	CONTINUOUS	CONTINUOUS
Detection Capability	LINE_OF_BEARING	LINE_OF_BEARING

Results and Analysis

We performed 538 trials with a 4 sensor network, operated using the Treatment II simulation parameters. For Treatment I, we borrowed the results from the single continuous sensor we simulated in Scenario 1. Results comparing the two sets of trials

are shown in Table 5.5. Table 5.6 shows two-sample t-test statistics we ran to compare sample means from the two treatments.

Table 5.5. Comparison of Scenario 2 Treatment Results (n = 538).

Comparison Result	Treatment I	Treatment II	Delta
Total Attacks Conducted	2711	2698	-13
Total Attacks Detection Only	858	1682	824
Total Attacks Geo-Located	0	878	878
Total Attacks Undetected	1853	1016	-837
Observed Attack Mean	5.04	5.01	-0.03
Observed Detection Mean	1.59	3.13	1.54
Detection Percentage	31.65%	62.34%	30.69%

Table 5.6. Two Sample t-Test Results for Observed Attack and Detection Means.

	Observed Attack Mean	Observed Detections
t-Stat	0.18	12.99
t-Critical	1.96	1.96

Scenario 3 – Simulating Line-of-Sight Detection Behavior

At this point in the simulations, the detection sensors deployed atop the facility operated without any line-of-sight (LOS) considerations. For example, if an attacker were located near the east side of the facility, any sensor operating on the west side of the facility was permitted to detect this attack. In Scenario 3, we seek to add further realism to the sensor behaviors by placing line-of-sight constrains on the simulation attack detection logic. This is shown in Figure 5.8. In Figure 5.8.a and Figure 5.8.b we show how the facility itself creates effective blind-spots for detection sensors. These blind spots result from the nature of the building materials used in the typical office or laboratory complex, which are known to severely reflect, impede, or attenuate RF transmissions.

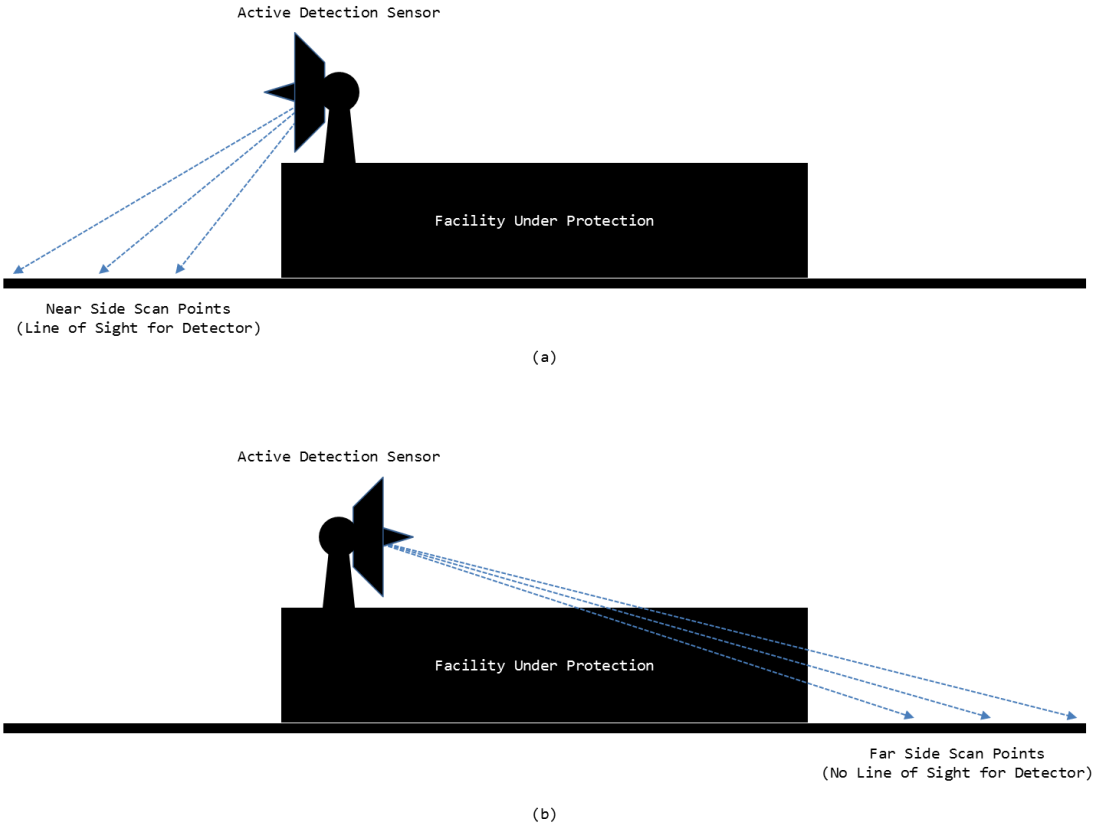

Figure 5.8. Elevation Model for the Facility Under Protection. In Figure 5.8.a, a detection sensor has clear line-of-sight to detect and scan for targets even in close proximity to the FUP. In Figure 5.8.b, a detection sensor has its scanning field of view obstructed by the facility structure. This serves as a dead-zone for our detection sensor scanning path.

To model line-of-sight detection requirements, we implemented geometry processing rules in the simulation software requiring that any line segment, drawn between the origin point of a wireless attack and any sensor detecting the attack, may not have any common intersection with the polygon representing the facility footprint. If an intersection with the facility perimeter is found, the detection is declared invalid, since the sensor does not have clear line-of-sight to the emitter target. These rules are direct result of our field testing; our prototype did not perform well without direct line-of-sight to an RF target.

Detection Sensor Performance with Line-of-Sight Rules

We first sought to quantify the impact that modeling line-of-sight performance would have on system detection capabilities. We designed a simple set of treatments, Treatment I featured a sensor operating without LOS constraints, while Treatment II operated with LOS constraints. The treatments and the simulation parameters are shown in Table 5.7.

Table 5.7. Treatment Parameters for Testing Non-LOS and LOS Simulation Rules.

Scenario Parameter	Treatment I	Treatment II
Simulation Duration	1 Year	1 Year
Simulation Runs	538	538
Facility Mean Attacks Per Year	5	5
Detection Sensors Deployed	4	4
Detection Sensor RPM	6	6
Detection Sensor Beamwidth	25	25
Detection Sensor Uses LOS	NO	YES
Deployment Position	FACILITY	FACILITY

We ran 538 trials of each treatment configuration. The results are shown in Table 5.8. Table 5.9 shows the t-test results performed on the treatment means.

Table 5.8. Comparison of Line-Of-Sight Treatment Results (n = 538).

Comparison Result	Treatment I	Treatment II
Total Attacks Conducted	2698	2819
Total Attacks Detection Only	1682	1491
Total Attacks Geo-Located	878	584
Total Attacks Undetected	1016	1328
Observed Attack Mean	5.01	5.24
Observed Detection Mean	3.13	2.77
Observed Geo-Location Mean	1.63	1.09
Detection Percentage	62.34%	52.86%

Table 5.9. Pairwise t-test Results Between Treatment Means.

	Observed Attack Mean	Observed Detections	Observed Geo-Locations
t-Stat	1.40	2.62	4.8703
t-Critical	1.96	1.96	1.9622

Our study revealed that both the mean number of attacks detected and the mean number of attacks geo-located were sharply impacted by the implementation of line-of-sight rules in the simulation. We hypothesized that attack detections would be affected by this change.

Selecting a Detection Scanning Scheme Under Line-of-Sight Rules.

Given that our simulation model now featured the ability to include FUP geometry impacts on LOS in system trials, we next turned our focus towards designing experiments which could test various scan pattern and sensor placement schemes. This gave rise to a series of research questions: Since vast areas of the facility grounds were now invisible to a detection sensor – due to detections being blocked by the facility itself – was it better to place the sensors a distance away from the facility, e.g. by mounting them high atop mast towers? Or, could a performance increase be obtained if sensors were programmed to scan a limited rotation sweep arc instead, to avoid scanning areas where the sensors lacks LOS?

To explore these questions we selected three treatments for trialing in our simulation tool (Treatment I, Treatment II, and Treatment III). Treatment I is the same sensor scheme we have been operating in prior scenarios – 4 sensors concealed on the rooftop of the FUP – only these sensors now integrated line-of-sight detection rules. We added this Treatment for experiment control purposes. We wanted a baseline to compare

other treatment results. We were interested in relative differences from the baseline control value.

For Treatment II, the sensors remained deployed on the facility rooftop; however we programmed the simulation such that the rotation path of each detection sensor was constrained to only sweep out 270° of a full circular arc, as shown in Figure 5.9. We hypothesized that programming the detection sensor to avoid scanning along azimuth lines where no line-of-sight existed would be a more efficient use of the detection sensor, when compared with the continuous scanning sensor scheme used in Treatment I.

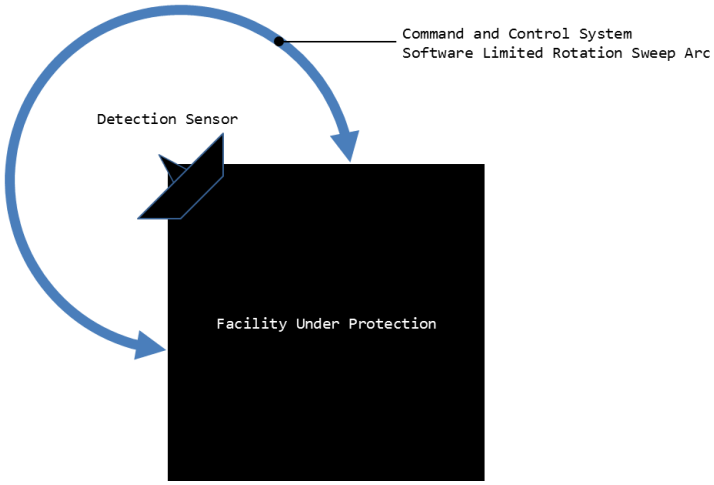

Figure 5.9. Constraining Detection Sensor Sweep Patterns to Avoid Dead-Zone Areas having Zero Line-of-Sight.

In Treatment III, we re-located the detection sensors to be mounted on mast towers located a distance away from the FUP. We selected 50 meters as the location offset. The sensors were programmed for continuous rotation, similar to Treatment I. Line-of-sight rules still applied to these sensors; however, we hypothesized that sensors deployed away from the facility would have a broader area of detection coverage,

providing better detection performance. A screen capture of a simulation running with sensors deployed on masts is shown in Figure 5.10. The parameters for each treatment are summarized in Table 5.10.

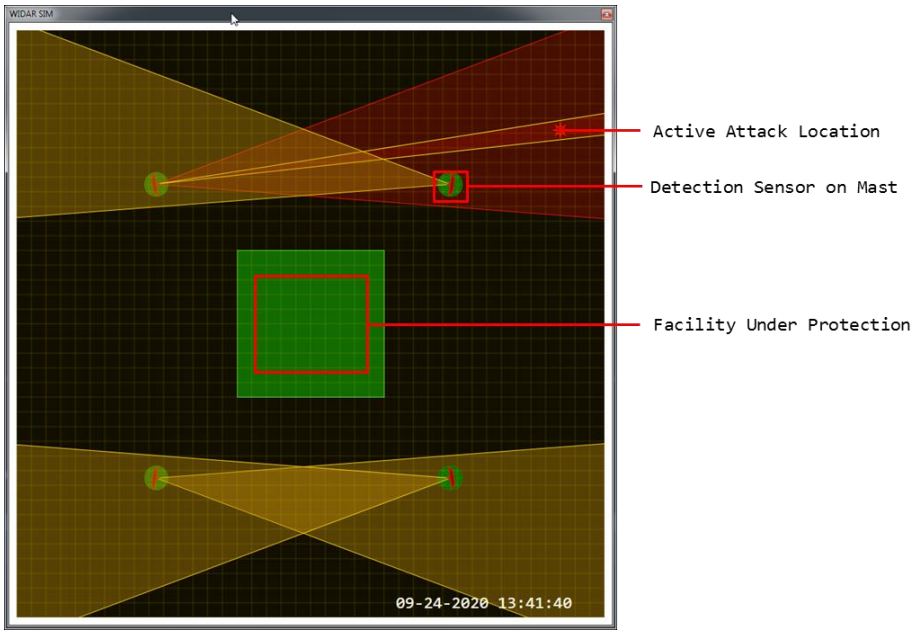

Figure 5.10. A Simulation Trial Running with Detection Sensors Deployed on Masts about the FUP.

Table 5.10. Treatment Parameters for Testing Three Different Sensor Detection Scanning Schemes.

Scenario Parameter	Treatment I	Treatment II	Treatment III
Simulation Duration	1 Year	1 Year	1 Year
Simulation Runs	538	538	538
Facility Mean Attacks Per Year	5	5	5
Detection Sensors Deployed	4	4	4
Detection Sensor RPM	6	6	6
Detection Sensor Beamwidth	25	25	25
Detection Sensor Scan Mode	CONTINUOUS	ZONE SWEEP	CONTINUOUS
Deployment Position	FACILITY	FACILITY	MAST

Results and Analysis

We ran simulation runs with 538 trials for each of the treatments. Table 5.11 shows results comparing the three separate treatments. The results from pairwise t-tests between all treatment combinations are shown in Table 5.12.

Table 5.11. Comparison Results from Simulations using Three Different Detection Scanning Treatments.

Comparison Result	Treatment I	Treatment II	Treatment III
Total Attacks Conducted	2819	2646	2759
Total Attacks Detection Only	1491	1700	1519
Total Attacks Geo-Located	584	584	759
Total Attacks Undetected	1328	946	1240
Observed Attack Mean	5.24	4.92	5.13
Observed Detection Mean	2.77	3.16	2.82
Observed Geo-Location Mean	1.09	1.09	1.41
Detection Percentage	52.86%	64.23%	54.97%

Table 5.12. Pairwise t-test results between treatment means.

	Attacks		Detections		Geo Locations	
Treatment Pair	t-Stat	t-Critical	t-Stat	t-Critical	t-Stat	t-Critical
I with II	1.98	1.96	2.88	1.96	0.00	1.96
I with III	0.65	1.96	0.38	1.96	2.93	1.96
II with III	1.39	1.96	2.57	1.96	3.04	1.96

In this scenario context, our results are somewhat inconclusive. Data did indicate that Treatment II, which employed a detection scanning scheme avoiding unnecessary movements through areas where no LOS existed, did yield a significant increase in detections (14% more detections, t = 2.88/2.57; t-critical = 1.96) versus the Treatment I scheme employing continuous scans following a complete rotation arc.

However, results also indicated that the Treatment III scheme, which deployed sensors atop masts at a distance away from the FUP, also yielded significantly higher geo-located attack detections (30%, t = 2.93/3/04; t-critical = 1.96) when compared to

both Treatments I and II. We should emphasize that, while the pairwise analysis results are mixed, the Treatment III results should be given more important consideration due to the design objective of our system to not only detect, but actually interdict attacks. The significant increase in spatial geo-location percentage provided by Treatment III offers FUP security personnel the ability to not only detect attacks, but also launch search and seizure countermeasures against the adversary.

Faced with these mixed results and given the importance we assign to attack spatial attribution, we conclude that a scheme employing a mast mounted sensor deployment strategy provides the most impact in terms of system performance. The 30% gain in geo-locations of Treatment III offsets the smaller 14% gain in detections achieved with the Treatment II scheme, which follows the constrained sweep pattern.

Still, given the data presented so far, geo-location rates remain far too low for our system to be practical in a real setting. Geo-locating 759 attacks out of the 2759 conducted in the Treatment III simulation trials indicates that our system is only spatially attributing a low percentage of all attacks conducted against the FUP (only 27.5% of attacks were geo-located in the Treatment III trials).

We next explore a scheme designed to significantly increase the spatial attribution capabilities, even when adhering to the low speed mechanical steering constraints imposed by simulations mimicking the low (6 RPM) rotation scan capabilities of our real world prototype design.

Scenario 4 – Simulating WLAN Data Rate Modulation to Slow Attack Progression

We modeled our next strategy consideration off of those older movies and television shows where the police detective attempts to run a telephone call trace on a

suspect and must keep the target talking on the line long enough for the trace to complete. Following a similar approach, we modified the simulated command and control system of the detection network to have the capability to lower the overall data rate of the WLAN environment, whenever an active attack is detected. For example, an attack might be launched using WLAN hardware supporting 100Mbps, but once any sensor in our system detects such an attack in progress, the command and control system can begin modulating the data rate of the system Wireless Access Point (WAP). This is depicted in Figure 5.11

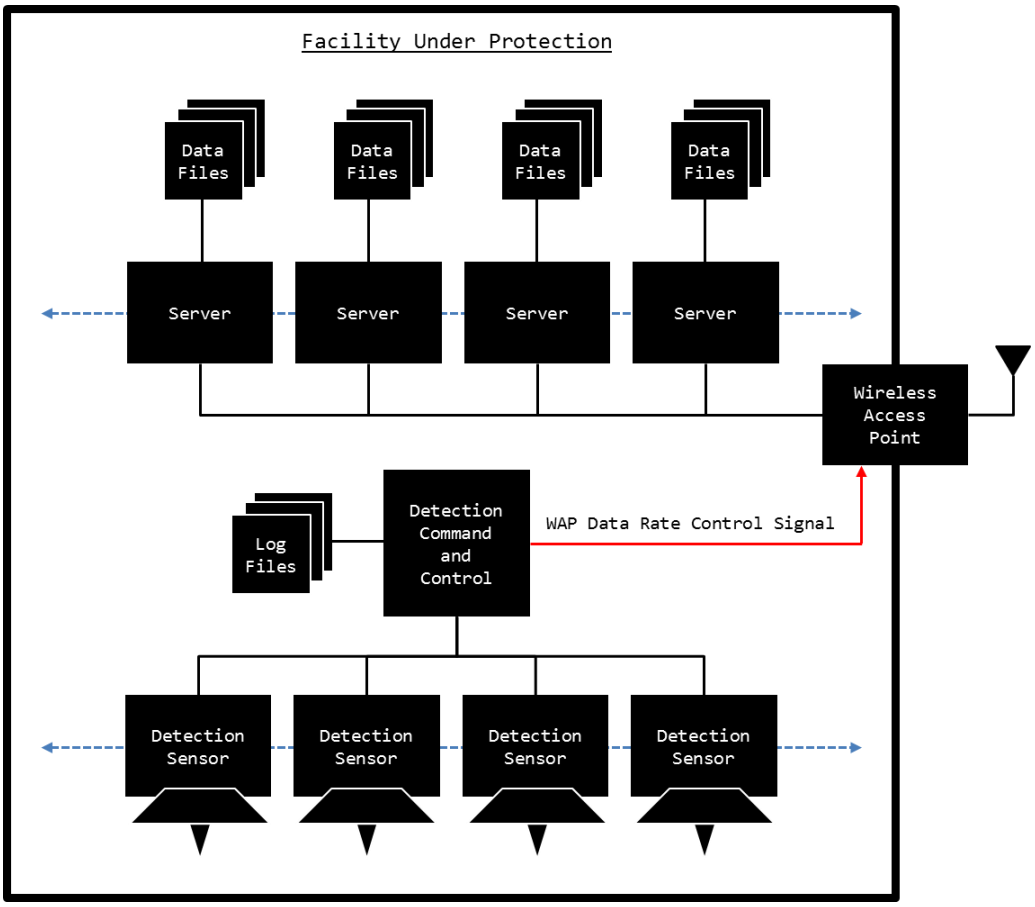

Figure 5.11. Placing the Wireless Access Point Data Rate Under Control of the Detection Command and Control Sub-System. In this scenario, the Command and Control system can throttle the facility Wireless Access Point data rate. The scenario assumes a honey-pot or honey-net situation, where data are made available for download, but when an active download attempt is detected, the Command and Control System dynamically modifies the data rate to be slower, thus stretching the attack duration with the aim of increasing the chances of successfully geo-locating the adversary.

In our simulation trials, we designed two treatments (Treatment I, and Treatment II). In Treatment I we operated the WLAN environment as we have been in all prior scenarios, with an average attacker WLAN data rate of 100Mbps, and an average FUP WAP data rate of 50Mbps. In Treatment II, we programmed the command and control system to drop WLAN data rates to 5Mbps whenever adversary activity was detected, from whatever the current high speed rate was before any attack was detected.

We hypothesized that a lowered the data rate will allow even a low-speed mechanically steered system more effective scan time, increasing the odds that more than one detection system sensor lobes the active emitter, enabling geo-location methods to be performed. In this way our system can mimic the "keep them talking longer" methods we discussed earlier. Table 5.13. shows the treatment parameters we used during simulation trials.

Table 5.13. Treatment Parameters for Normal Data Rate and Modulated Low Speed Data Rate.

Scenario Parameter	Treatment I	Treatment II
Simulation Duration	1 Year	1 Year
Simulation Runs	538	247
Facility Mean Attacks Per Year	5	5
Mean Adversary Skill Level	NATION_STATE	NATION_STATE
Detection Sensors Deployed	4	4
Detection Sensor RPM	6	6
Detection Sensor Beamwidth	25	25
Detection Sensor Scan Mode	CONTINUOUS_NORM_DATA_RATE	CONTINUOUS_LOW_DATA_RATE
Deployment Position	FACILITY	FACILITY

Data Rate Modulation: Keen Strategy or Pure Folly?

Before presenting our simulation results we feel it is important to discuss the merits of modulating the data rate when an active attack against the FUP is detected by our system. When consideration is being given whether or not to provide a data rate

modulation capability to the WLAN environment operated within proximity to the FUP, one should also consider whether is it more advantageous to instead provide the capability to simply halt any WLAN operations whenever unauthorized WLAN activity is detected by the sensor network. We call this the Halting mode of operation. Under the Halting mode of operation, sensors would sacrifice any ability to geo-locate an active emitter, since WLAN operations are immediately disabled upon attack detection. Indeed when the prevention of data exfiltration of any sort attains a higher priority than capturing or revealing the location of an attacker, then the Halting mode of operation should be preferred.

That being stated, one real set of circumstances where lowering the data rate of the WLAN would be desirable, even a key part of a highly effective detection system strategy, would be in a honey-netting environment [5] [6]. In those situations, we have deliberately set out to lure an adversary to conduct wireless attacks against an intentionally made vulnerable WAP with the objective of monitoring the tradecraft being employed by the adversary or, as in the context of our spatially enabled detection system, we seek to capture the adversary when she is in the act of perpetrating her attack. We feel the honey-netting context provides the most supportive rationale for employing the lowered data rate strategy, in place of the simpler and less risky Halting Strategy.

Results and Analysis

We performed simulation runs of 538, and 247 trials each using the respective treatment parameters listed in Table 5.13. We ran fewer trials for Treatment II only because a software design peculiarity inherent to the simulation software implementation meant that simulating the extremely low data rates in the modulated data rate trials

caused our simulations to run significantly slower than the high data rate trials. A comparison of treatment results is shown in Table 5.14. Table 5.15 shows the results of pairwise t-tests we performed against mean observed attacks, observed detections, and observed geo-location simulation results.

Table 5.14. Comparison of Simulation Results using Two Different WLAN Data Rate Treatments.

Comparison Result	Treatment I	Treatment II
Total Attacks Conducted	2819	1253
Total Attacks Detection Only	1491	1224
Total Attacks Geo-Located	584	1175
Total Attacks Undetected	1328	29
Observed Attack Mean	5.24	5.03
Observed Detection Mean	2.77	4.92
Observed Geo-Location Mean	1.09	4.72
Detection Percentage	52.86%	97.69%

Table 5.15. Pairwise t-test Results Between Treatment Means.

	Observed Attack Mean	Observed Detections	Observed Geo-Locations
t-Stat	1.08	12.10	21.21
t-Critical	1.96	1.96	1.97

Significantly increased detection and geo-location rates (97.6% and 93.8%) were observed in the Treatment II trials, where the detection system command and control was able to reduce WLAN data rate at the instant an attack was first detected. We concluded that this behavior is a powerful weapon to include in the protection strategy of our detection system. Indeed enabling this capability yielded the most desirable results observed in simulations so far, towards an effective strategy capable of detecting and defeating the vast majority, in terms of raw numbers, of attacks launched against the FUP.

We next turn to a study of geo-location performance and the capabilities of the detection system to accurately and precisely estimate the position coordinates of any attack lobed by more than one detection sensor.

Scenario 5 - Simulating Detection System Position Estimation Error

Having explored various strategies for sensor placement on and about the facility, as well as examining sweep-pattern control, and honey-netting schemes, we lastly turn to an analysis of sensor position error estimates. We first describe the concept of error in the position estimates output by our detection system. We show that the position estimate is dependent upon the boresight error of each simulated detection sensor, which impacts the accuracy of the LOB calculated by each sensor. We then present a randomized, time dependent process which dynamically models the boresight error of each simulated detection sensor. Finally, we apply three different boresight error treatments to simulation runs integrating this model, and discuss the detection performance impact that randomized sensor boresight error has on position estimation outputs of the detection system.

The Concept of Position Estimation Error

When a sensor in our system detects an attack, the sensor calculates a LOB for the detected emitter. Logically, the LOB can be represented as a line segment, with origin at the detection sensor, and termination at the emitter. When two or more sensors detect the same emitter, triangulation techniques permit us to estimate the sensor position. Think of this as the finding the coordinates of the single point where the two equations describing the LOB line segments intersect. In practice a real detection sensor is not going to possess laser-like, straight line accuracy; there is a boresight error impacting the detection sensor

LOB calculation. This results in the LOB becoming a pie-shaped wedge, instead of a linear vector. Furthermore, where two wedges representing separate device LOB estimates intersect, the position estimate for the triangulated emitter becomes a region as well. We define this region as the Position Estimate, i.e. the region where the detection sensor network estimates that the attack is originating from. This is illustrated in Figure 5.12

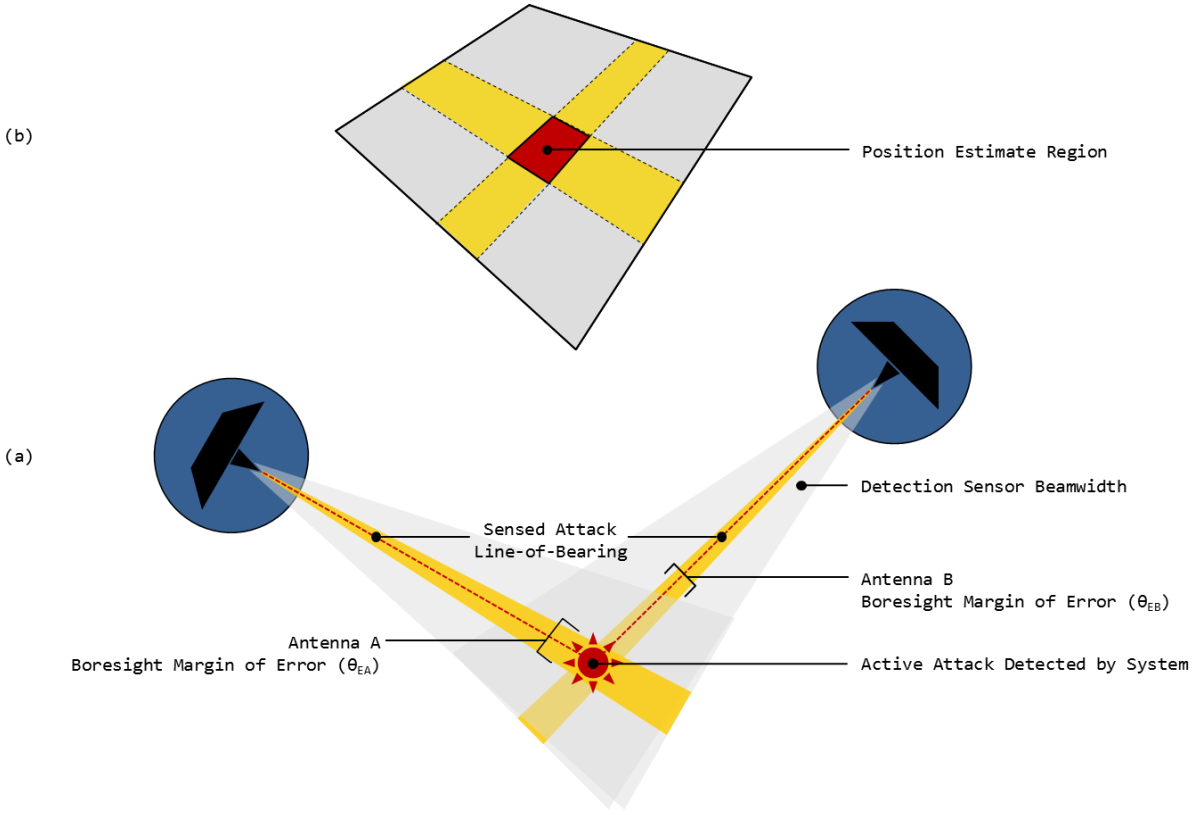

Figure 5.12. Conceptual Model of Detection Sensor Position Estimate. In Figure 5.12.a, two sensors are shown lobing an active attack. The yellow colored bands in the center of each detection beam correspond to the simulated mean boresight error – the actual emitter LOB error is somewhere in this band. Antenna A and Antenna B both have independent and time-varying boresight errors (θ_{EA} and θ_{EB}). Figure 5.12.b depicts a close-in representation showing the region where the two LOB error band polygons intersect. This is the Position Estimate for the detection system.

Modeling Boresight Estimation Error Using a Randomized Process

As shown in Figure 5.12, simulation detection sensors do not operate with constant boresight errors. This is due to non-linearities inherent to RF propagation. To model this behavior, we implemented a random process within the simulation which outputs a new boresight error for each detection sensor at each simulation time increment. Figure 5.13 features a block diagram describing the randomized boresight error process.

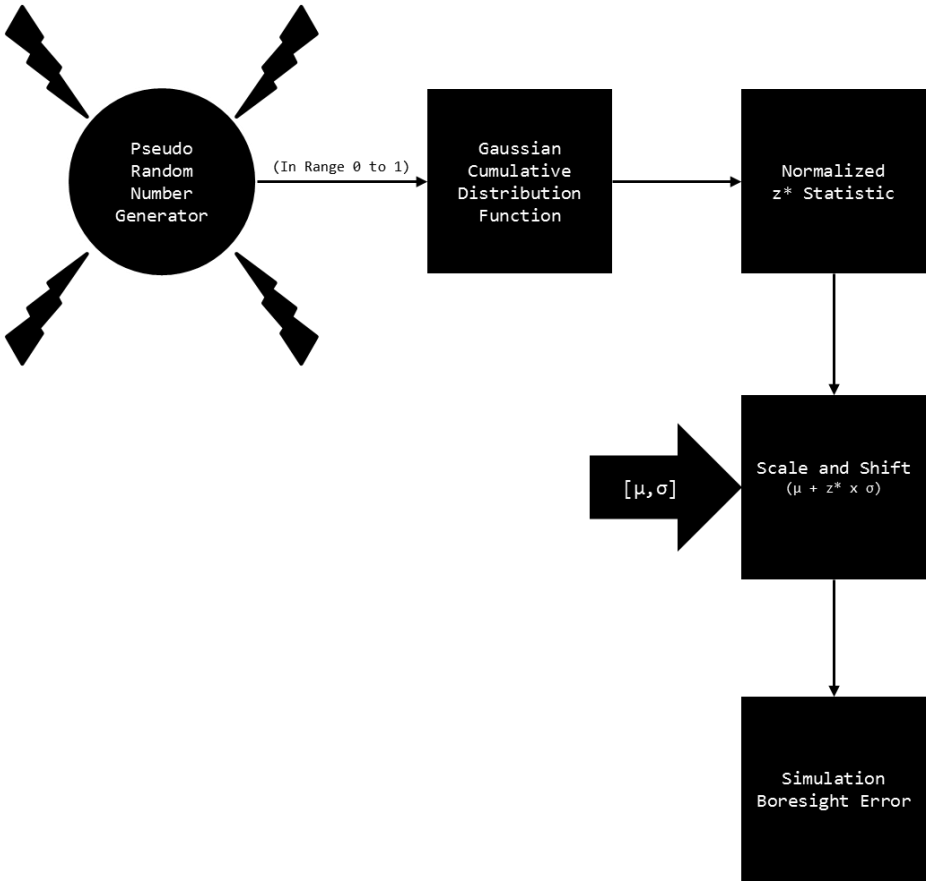

Figure 5.13. Randomized Boresight Error Estimation Process. A PRNG is shown generating a probability value which is input to the inverse Gaussian CDF. The result is a z* statistic representing the normalized Gaussian mean for that probability. The normalized mean is then scaled and shifted by the model parameters for Mean Boresight Error and Mean Boresight Standard Deviation. This is the simulation Current Boresight Error which is assigned to a detection sensor for one simulation time slot.

As is shown in the figure, a probability value is generated by a pseudo-random number generator (PRNG). In our implementation, this is a single-precision, floating-

point value between 0 and 1. We input the probability into a routine performing an inverse-lookup of a corresponding mean value from the Cumulative Distribution Function of the Gaussian Normal Distribution. The routine returns a normalized z* value (μ=0; σ=1), which is then scaled and shifted to align with the mean and standard deviation parameters input into the simulation model at startup. We call this result the Mean Boresight Error (MBE). The MBE for each detection sensor is updated at each simulation time instant – the value varies constantly with simulation time about the preset mean, and is independently generated for each sensor.

The equation for the Cumulative Distribution Function (CDF) of the Gaussian Normal Distribution appears in Equation 2. No closed form solution exists for the CDF of the Gaussian Normal Distribution; instead we selected a numerical approximation from here [7] and integrated source-code for the algorithm directly into the simulation module.

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-t^2/2}\, dt \tag{2}$$

We sanity checked the algorithm output using the `Norm.Inv()` function in Microsoft Excel to verify the numerical approximation from the algorithm implementation exactly matched the approximations output by Excel. Results from this function, and results from our algorithm implementation were numerically identical. This provided us confidence in our implementation, but more importantly we no longer needed to rely on a table of z-values to generate randomized means, something we wished to avoid in order to provide greater precision to our simulation process.

Applying the Randomized Model to Simulation Detection Sensors

We next programmed each simulated detection sensor to have a time dependent boresight error which randomly varied the accuracy of the LOB estimate calculated by the sensor. This resulted in the LOB estimates becoming the wedge shaped polygons we described in Figure 5.12. The point of the wedge corresponded to the position of the sensor, and the angle spanned between the left and right edges of the wedge corresponded to the random error mean generated by the boresight error model. Field tests using our sensor prototype against RF targets with known headings behaved exactly in the same manner; LOB values for these targets fell in a range having a mean, and standard error, with the true mean falling within a confidence interval based on the standard error and number of samples in our estimate [2].

Any attack that was sensed by a single sensor could then be spatially attributed to the wedge-shaped polygonal region determined by the random boresight error that the sensor was operating with at the time of attack detection. When two or more sensors were able to lobe an attack, the resulting wedges from each sensor LOB could then be intersected to form the Position Estimate. To perform polygon intersection and polygon area calculations in a computationally efficient manner in our simulation runs, we implemented 2-D polygon intersection algorithms based on the computational geometry tools described here [8]. The calculated area of the position estimate region is a key performance metric in our experiments – allowing us to compare simulation runs with different detection sensor Mean Boresight Error treatments.

Selecting Mean Boresight Error Treatments for Simulation Runs

The MBE for each sensor was then normally distributed and model driven using our random process. This permitted each detection sensor to simulate position estimation errors that were parametrically driven by configuration values for mean (μ) and standard deviation (σ). We only needed to initialize our simulation detection sensors with boresight error mean and boresight error standard deviation values that we were interested in testing in our simulation experiments. These configuration values were input at simulation startup. Each simulation run was performed using differing treatments, with respect to MBE μ and σ configuration values. We arrived at the individual MBE treatments by averaging boresight errors measured during actual field tests of our prototype detection sensor. Four treatments using different mean boresight error statistics are shown in Table 5.16.

Table 5.16. Four Different Mean Boresight Error Treatments Obtained from Detection Sensor Field Tests.

Treatment	Boresight Mean Error	Boresight Std Dev
I - Sequential Lobing	-16.70°	4.35°
II - Monopulse I	-1.41°	0.81°
III - Monopulse II (Best)	0.98°	1.22°
IV - Monopulse II (Worst)	2.56°	1.12°

Treatment I configures the detection sensors with the large mean boresight error we observed when using our field prototype in a sequential lobing single antenna configuration. All of the remaining treatments are configured with values taken from the monopulse configuration of the detection sensor we fielded. They differ only in the software monopulse processing detector that was used to detect emitter activity, and perform monopulse LOB estimate calculations. Treatment II uses a boresight error mean

value resulting from a Peak Variance detector. Treatments III and IV both use a Matched Filter detector, but since this detector featured a significantly wider boresight error variance, relative to the Treatment II detector, we chose to simulate the best and worst case boresight error values, so that we could ascertain how sensitive position estimate was to each error measurement. For an in-depth technical discussion of the how we arrived at these treatment parameters, see [2].

Discussion of Simulation Run Results

We ran simulation runs using each of the treatment parameters, stopping each simulation when 1000 geo-located detections were logged. A side-by-side comparison of results is shown in Table 5.17. The Position Estimate Error for each treatment is highlighted in the table. ANOVA and pairwise analysis indicate that all of the position error means are significantly different from the other. The first two rows of the table show the sample boresight error mean and standard deviation, which closely matched that of the parameters we input to the randomized lookup module of the simulation. This was expected and served as a sanity check on the data we collected from the simulation.

Table 5.17. Summary of Simulation Results using Different Mean Boresight Error Treatments.

	Treatment I	Treatment II	Treatment III	Treatment IV
Sample Boresight Mean Error (deg)	-16.59	-1.41	0.95	2.56
Sample Boresight Std Dev (deg)	4.35	0.80	1.20	1.13
Position Estimate Area Mean (sq ft)	7714.56	67.38	38.76	205.09
Postion Estimage Area Std Dev (sq ft)	8529.25	246.21	132.99	603.57
Position Estimate CI	528.64	15.26	8.24	37.41
Mean + CI	8243.20	82.64	47.01	242.50
Mean - CI	7185.92	52.13	30.52	167.68
Attack Distance Mean	132.29	128.77	134.34	130.61
Attack Distance Std Dev	54.64	54.75	54.78	52.94
Attack Distance Min	283.23	281.06	284.71	281.059
Attack Distance Max	6.43	10.01	11.12	10.94

Does the Detection System Produce Actionable Threat Intelligence?

In [1] we detailed the threat model and environmental landscape that would make attacks staged and launched externally to the FUP viable attack vectors. We described the Parking Lot Attack and commented that an adversary would most likely camouflage her attack such that it would be visually difficult to detect using surveillance methods alone. Our system was designed to counter this attack vector, by providing actionable threat intelligence that would permit facility security personnel to narrow the geographic source of any detected threats to a spatial area that was rapidly searchable. For example, we think that narrowing a suspected attack source to only 1 or 2 target vehicles, instead of searching all the vehicles that could potentially be located in a research lab or office parking lot would represent reasonable actionable threat intelligence.

To assess whether our detection system could meet the challenge presented by such a narrow search area size, we selected a target vehicle footprint of 52 square feet, based on a U.S. government vehicle size report found here [9]. We then plotted the

Position Estimate regions shown in Table 5.17 so that we could visually inspect each area relative to one another. Figure 5.14 shows the plotted results.

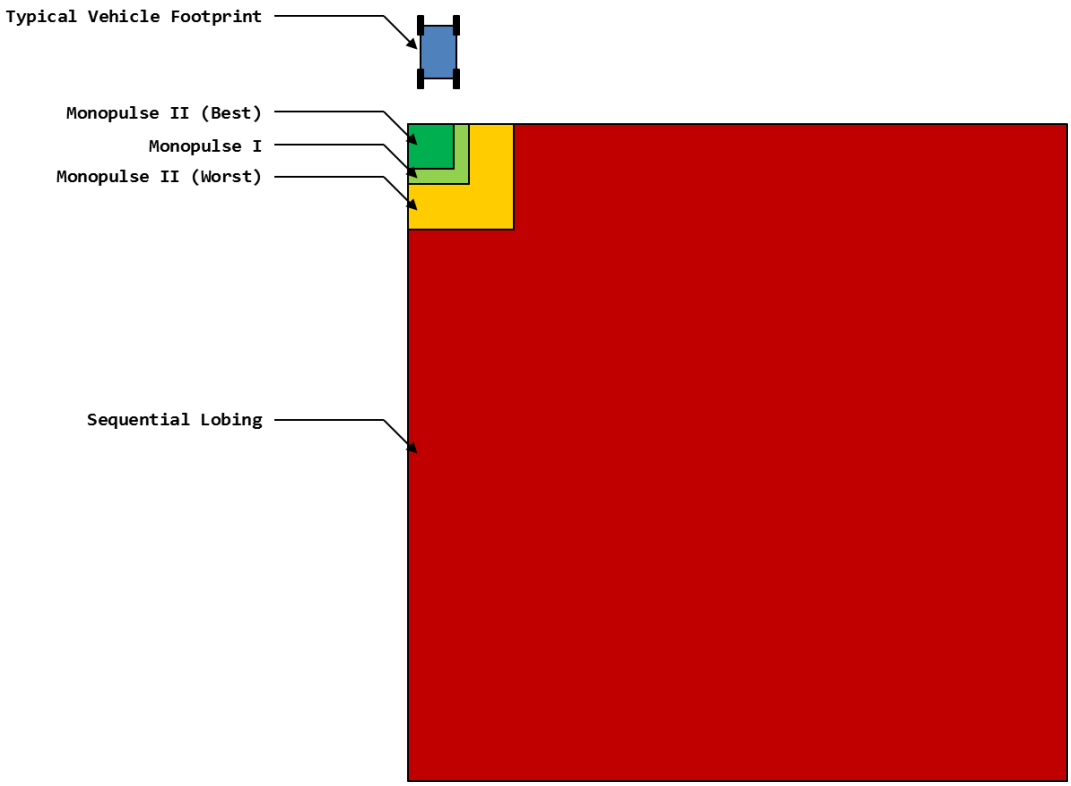

Figure 5.14.Relative Footprint Sizes Derived from Study Position Estimation Errors.

In the figure, it can be clearly seen that the Position Estimate for results obtained from the sequential lobing (Treatment I) target tracking scheme – about 7700 square feet – would most likely be unacceptable in terms of actionable threat intelligence, although this region is still smaller than the typical sub-urban house lot, and is several orders of magnitude smaller than the entire external vicinity of the FUP we used in our simulation, which comprised 150,000 square feet.

Both Treatment II and Treatment III produced Position Estimate results well within the bounds of our one or two vehicle search region requirement we imposed as a reasonable search area expectation. Even the 205 square foot zone produced by the

Treatment IV Position Estimate – about 4 vehicles – could be arguably called reasonable, depending on the nature of the threat environment and the risk imposed by any successful attack.

We conclude that each of the estimation treatments do result in a substantial narrowing of the zone that would need to be searched to find and locate any unauthorized transmitter detected by our system. The monopulse methods operate very near the performance threshold we set for our system, showing promising results that encourage us to continue further research into our detection system.

Research Conclusions

We presented a simulation tool capable of answering questions about how a network of cooperative mechanically-steered, monopulse-enabled detection sensors might protect a hypothetical facility from standoff externally launched wireless attacks. Our intention was to show that there are many types of questions that can arise during the course of determining what an effective strategy would be for device spatial deployment and operational considerations. It was not our intention for this work to be an exhaustive study of every minute operational detail. Instead we wanted to emphasize the utility that simulations modeling can afford a designer of such a system.

Using the simulation tools we developed we first were able to conclude that our prototype sensor, which employed a pan then scan mechanical steering scheme, would not be an effective detection sensor in a real environment. We also concluded that, even when the chassis was modified for continuous rotation, the sensor only achieved a simulated detection success rate of 65% and only a 13% geo-location rate, even after our

simulations increased the number of sensors operating in the detection network from 1 to 4.

We increased the simulation realism by incorporating line-of-sight detection rules into the simulation model, and then ran scenarios where we tested the effectiveness of a zone sensor sweep patterns versus simple continuous sensor rotation and deployment of sensor on mast towers instead of the facility roof top. We concluded that sensor scan sweeps increases operational effectiveness, but overlapping scan zones on masts situated away from the actual facility offer even better performance in terms of effective spatial attribution. Furthermore, maintaining the continuous rotation mode of operation, instead of an oscillating back and forth scan arc is a simpler scheme requiring less complex mechanical control.

We showed that permitting the detection system command and control logic to intentionally degrade the WLAN data rate during an active attack offers significant benefits in terms of both detection percentage and geo-spatial attribution. This technique would be especially relevant in a honey-netting context, where permitting the sensor to throttle WAP data rate to maintain longer sessions with a lured adversary would be highly desirable.

Finally, we concluded that simulation trials configured with mean boresight errors obtained from actual device field tests indicate that position estimate error regions provide significant reductions in the search space, and that the monopulse boresight error estimates featured regions with areas comparable to the footprints we presented for the average vehicle. We found these results promising and we intend to further pursue research questions in the domain of wireless intrusion detection and spatial attribution.

References

[1] D. J. Gieseman and T. E. Daniels, "Countering the Parking Lot Attack – Design for a Detection System Employing Monopulse Radar Methods to Detect and Spatially Attribute RF Targets in the 2.4 GHz ISM Band," Iowa State University, Ames, IA, 2015.

[2] D. J. Gieseman and T. E. Daniels, "Analyzing Line of Bearing Estimates Collected from a Device Employing Monopulse Radar Methods to Track RF Targets in the 2.4 GHz ISM Band," Iowa State University, Ames, IA, 2015.

[3] R. A. Poisel, Electronic Warfare Target Location Methods, Artech House, 2012.

[4] D. J. Gieseman and T. H. Maze, "Evaluating Capacity and Delay Given the Implementation of ITS Technology at Truck Weight and Safety Inspection Stations," *IET Intelligent Transportation Systems*, vol. 1, no. 2, pp. 124-130, 2007.

[5] L. Spitzer, "Honeypots: Catching the Insider Threat," in *Computer Security Applications Conference, 19th Annual*, 2003.

[6] L. Spitzner, "The Honeynet Project: Trapping the Hackers," *IEEE Security and Privacy*, vol. 1, no. 2, pp. 15-23, 2003.

[7] P. J. Acklam, "An algorithm for computing the inverse normal cumulative distribution function," [Online]. Available: home.online.no/~pjacklam/notes/invnorm. [Accessed 08 2015].

[8] J. O'Rourke, Computational Geometry in C, 2nd Edition, Cambridge University Press, 1998.

[9] U.S. Department of Energy, "Average Vehicle Footprint for Cars and Light Trucks," 09 2011. [Online]. Available: http://energy.gov/eere/vehicles/fact-693-september-19-2011-average-vehicle-footprint-cars-and-light-trucks. [Accessed 08 2015].

[10] S. M. Sherman and D. K. Barton, Monopulse Principles and Techniques, Artech House, 2011.

[11] D. R. Rhodes, Introduction to Monopulse, Artech House, 1980.

CHAPTER 6. CONCLUSIONS AND RECOMMENDATIONS FOR FURTHER RESEARCH

Results and Conclusions

We presented the design and architecture for an integrated network of cooperative mechanically steered sensors designed to detect and spatially attribute RF targets operating in the 2.4 GHz ISM band. It was our intention to deploy this sensor platform to defend a facility against an adversary employing the Parking Lot Attack as a vector for information system compromise and data exfiltration.

The primary finding of this work is that the utility of such a detection system does indeed show promise, provided that the environment about the facility under protection can be conditioned for the proper line-of-sight characteristics that we were able to demonstrate during our own controlled field tests. Results we obtained from tests using both a prototype detection sensor and simulation studies show that the region of estimated position error calculated by this type of detection system yields a perimeter that can be easily searched and policed by a counter intrusion team, at least for the relatively close proximity detection zones surveyed and simulated for our research. Additionally, monopulse radar methods, which led to us to incorporate an antenna array into our sensor architecture, enabled us to show that these techniques can outperform less sophisticated schemes, where only a sequentially-lobing single directional antenna is employed for detection.

Recommendations for the Direction of Additional Research

Developments in software defined radio (SDR) have made possible exciting new possibilities in this research domain. Where our present research prototype could only

sample a single 333kHz 2.4 GHz ISM band channel, platforms such as the HackRF or the Ettus Research USRP device permit wide band monopulse detectors to be created entirely within software. These tools, when coupled with software radio development environments such as GNU Radio, would permit much more flexibility in terms of prototype design, capabilities, and device sensitivity. We recommend that any future studies explore the capabilities of these platforms and tools. Furthermore, these technologies permit detection sensor designs that are much smaller in terms of physical size, to the point where it is entirely feasible to deploy a monopulse sensor on a mobile and autonomous drone. We also recommend continued research involving studies where a single or cooperative network of drone mounted detection sensors would be employed for facility wireless attack detection and prevention. Whether these platforms would augment a fixed network of mechanically steered – sensors similar to those presented in this research – or whether they provide the situational awareness capabilities to necessary to independently defend a facility remains an interesting research question from a theoretical, simulation, and real hardware prototyping perspective.