EXPLOITING THE PHYSICAL LAYER TO ENHANCE WIRELESS OPERATION WITH COGNITIVE RADIOS

BY ROBERT D. MILLER

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Professor Wade Trappe

and approved by

_____

_____

_____

_____

New Brunswick, New Jersey

January, 2011

ABSTRACT OF THE DISSERTATION

Exploiting the Physical Layer to Enhance Wireless Operation with Cognitive Radios

by Robert D. Miller

Dissertation Director: Professor Wade Trappe

Wireless communication systems have undergone considerable evolution in the past decade. This is in large part due to significant advancements made in underlying physical (PHY) layer technologies, resulting in substantial performance leaps in data rates and reliability. These strides have made wireless devices the platform of choice for communicating. Accordingly, considerable progress in the realm of Software Defined Radio (SDR) has made seamless cross-protocol communication not only plausible but a near-term certainty. This is primarily due to the openness of the physical layer to the system user and developer. Unfortunately, this same openness also provides an adversary with a powerful point of attack. It is therefore essential to consider physical layer exploits in order to enhance operation in future wireless communication networks. In this thesis, we consider using physical layer exploitation to: (1) enhance situational awareness, (2) act as an adversary, and (3) mitigate poor environments and adversarial conditions.

By enhancing its situational awareness, a device can make truly intelligent operational decisions. In order to efficiently and effectively utilize the RF spectrum a device should know what services are available and the level to which they are utilized — perhaps even down to the identity of neighboring devices. It is also advantageous for a device to know its physical location, characteristics of its environment (e.g. indoor vs. outdoor), and of course whether or not there is an adversary in the region. In this thesis, we explore new physical layer based techniques to acquire this valuable information.

Next, we acknowledge that sometimes the best defense is a good offense as we explore attack strategies focused on the physical layer. By thinking like an adversary, one can better anticipate possible attacks and determine the appropriate remedies. Since most 3G and 4G wireless standards and protocols incorporate some form of multi-input multi-output (MIMO) technology, we pay specific attention to MIMO operation. Whereas the majority of related research has assumed the presence of an unintelligent jammer, our focus will be on truly smart attacks.

In the final part of this thesis, we consider what to do once an accurate description of the operational scenario is achieved. Accurate knowledge of the environment plays a key role in Dynamic Spectrum Access (DSA), where devices adapt modulation schemes and protocols to both optimize communications and minimize interference with existing wireless infrastructures. Additionally, accurate situational awareness provides insight into potential communication hazards — from severe multipath conditions to adversarial attacks. In this thesis, we present unique physical layer methodologies that can be used to overcome channel degradations due to *both* natural phenomena and adversarial activity.

In each part of this thesis, we accompany theoretical results and findings with simulations and real-world experimentation in order to illustrate the feasibility and applicability of the proposed techniques. Real-world implementations were conducted using current SDR architectures.

Acknowledgements

The more that I learn, the more that I learn how much more there is to learn.

This thought rings true as I stare into the face of my son, Ryan, cheerfully wondering what more there is to learn. It seems that every step through life brings a new challenge, a new door, a new problem. And too often the solution brings forth yet another problem. And so it goes, for work, for school, and for life and in general... but if I have learned one thing, it is to always meet a challenge head on and with full focus, and win or lose, you can stand proud in the result.

Throughout my graduate studies I have been blessed with the guidance, support, and friendship of many. I would first like to express my sincere gratitude to my advisor, Dr. Wade Trappe, whom I consider both a mentor and a friend. His insight, encouragement, and direction have made the journey through graduate school arduous yet rewarding, and at times even enjoyable.

I would also like to acknowledge the great discussions with other WINLAB faculty and staff. To Dr. Rich Howard, for many enlightening technical discussions. To Dr. Marco Gruteser, for guidance and support particularly with regard to the tire pressure sensor work. To Ivan Seskar, for endless hours answering my many questions. And to my thesis committee, Dr. Dipankar Raychaudhuri, Dr. Yanyong Zhang, and Dr. Paul Prucnal. A heartfelt thankyou also extends to fellow WINLAB students and graduates, Dr. Wenyuan Xu, Dr. Suhas Mathur, and Shridatt "James" Sugrim, for the numerous technical discussions and debates.

A special thankyou is reserved for my corporate colleagues for their understanding and support of my graduate school journey. To Brian Abbe, a special

thankyou for always supporting this goal, and for the original introduction to Rutgers by way of Dr. David Daut. To Dr. Paul Zablocky, for always taking the time to discuss technical ideas and experiments. To Vince Simpson, Kevin Uher, and Bret Eddinger, for being so supportive of the work/school/family balance. And to Brian Hetsko, Ronald Li, and Yaakov Gorlin, for the numerous hours spent in technical discourse.

Finally, I extend my utmost gratefulness for the support of my family. To my wife, Michelle, for always being so patient and loving towards this goal. To my mother and father, Roberta and Robert, for raising me to search for truth and meaning, and to get up when you fall down. To my sister, for her friendship and encouragement. And last but first, to God, for His presence, comfort, insight, and strength through all of life's struggles and blessings.

Dedication

To my family.

Table of Contents

List of Tables

List of Figures

Chapter 1
Introduction

1.1 Motivation

Wireless communication systems have undergone considerable evolution in the past decade. In 1999, only 27% of US citizens used cell phones. In 2009, this number skyrocketed to 89% [1]. Wireless devices now permeate through society, and the general public has grown accustomed to high quality, dependable wireless links to both cellular and broadband services. As of 2010, 59% of American adults access the Internet wirelessly from either a laptop or cellular phone [2]. Society has grown accustomed to tweeting and texting, syncing with email and calendar events, and streaming multimedia content such as YouTube and Pandora at a moment's notice from any location — and all from a single wireless device.

The capabilities needed to support these demands have arisen in large part due to significant advancements in underlying physical (PHY) layer technologies. Advanced coding and modulation techniques [1] now come close to achieving the Shannon capacity of wireless channels. Moore's law has resulted in devices that are capable of employing these advanced techniques, and Eveready's law has helped to maximize device usability. Popular handheld devices such as the Apple iPhone last for days on a single charge and fit conveniently in a user's pocket.

In order to meet user demand, most wireless communication devices support

[1] Error corrective coding (ECC) techniques such as low-density parity check (LDPC) codes and Turbo-codes, coupled with modulation techniques such as Trellis-coded modulation and multi-input multi-output (MIMO) systems offer throughput increasingly closer to the capacity limit for Gaussian channels given by the Shannon-Hartley theorem [3].

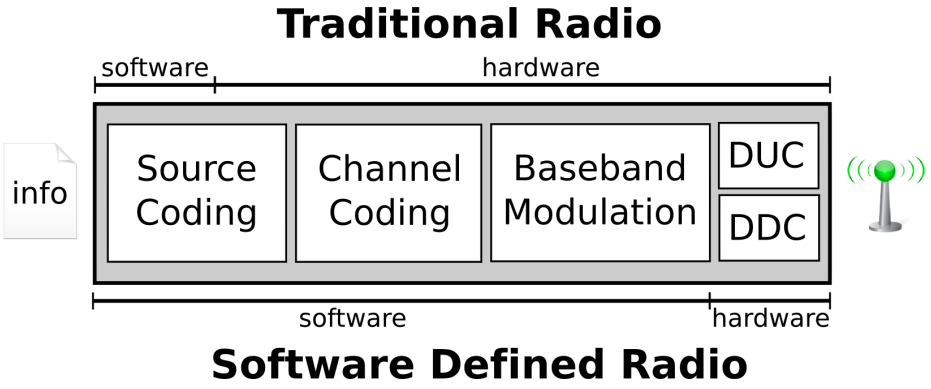

Figure 1.1: The general communication system task allocation between hardware and software is illustrated for a traditional radio and a software defined radio.

multiple protocols. Access to GSM, CDMA, WiFi, and Bluetooth networks are often offered in single platform solutions. This is presently done by including multiple, protocol-specific hardware units. But as protocols expand and grow, devices need to include more and more hardware units to meet user demand; these hardware units are often application specific integrated circuits (ASICs). Recent advancements in the realm of Software Defined Radios (SDRs) offer a more promising solution to this dilemma. By leveraging the increased processing power available on modern platforms, software can perform the jobs traditionally relegated to hardware. By pushing the software processing closer to the antenna, the wireless device becomes more readily adaptable — developers can wirelessly deliver code updates and new protocol stacks on the fly to facilitate new communication schemes. Further, protocols like WiFi and Bluetooth that share spectral resources can also share processing resources. This is essential in future wireless networks, where devices may choose to associate with existing protocols or implement their own. Figure 1.1 presents the functional mappings between hardware and software for traditional radios and Software Defined Radios.

We now introduce the concept of a Cognitive Radio (CR), which we define simply as a Software Defined Radio with an intelligence engine. A CR is capable

of intelligently deciding how best to use and adapt its resources in any given environment. Consider a relevant example, as FCC regulations now allow secondary usage of idle spectral resources in selected radio frequency (RF) bands [4]. A CR can use the capabilities inherent in an SDR to implement a wide array of spectral sensing algorithms in order to detect primary user occupancy. The CR can then implement a communication protocol catered specifically for the situation at hand. The Cognitive Radio concept was first introduced in [5], with much recent interest thanks to prominent advancements in SDR architectures [6, 7, 8, 9, 10].

It is clear that future wireless networks will be made up of more and more Cognitive Radios, giving software developers complete control over the majority of the transmit and receive path functions. Thus, developers have total access to information from all layers of the protocol stack — from application (APP) to physical (PHY). While traditionally, developers would have access to data bits (and perhaps rough signal quality estimates), now they have access to the actual received waveform. The data processing inequality tells us that such access inherently provides more information [11]. While this added information can greatly benefit wireless operation, it also provides new cogent attacks for an adversary. In this thesis, we present new and novel techniques to harness information available at the physical layer in order to: (1) enhance situational awareness, (2) act as an adversary, and (3) mitigate poor environments and adversarial conditions.

1.2 Situational Awareness

By enhancing its situational awareness, a device can make truly intelligent operational decisions. In order to efficiently and effectively utilize the RF spectrum a device should know what services are available and the level to which they are utilized — perhaps even down to the identity of neighboring devices. It is also

advantageous for a device to know its physical location, characteristics of its environment (e.g. indoor vs. outdoor), and of course whether or not there is an adversary in the region. In Chapter 2 of this thesis, we explore new physical layer based techniques to acquire this valuable information.

We begin by addressing the issue of service discovery and device identification in the 2.4 GHz ISM band. Our focus is on exploiting physical layer details of the WiFi and Bluetooth protocols to reliably detect networks and specific devices. We then introduce new signal processing techniques that leverage physical layer information from existing transmitters to discover device mobility and location. We demonstrate the effectiveness of these algorithms by using multiple protocols and services such as GSM, WiFi, FM radio, and broadcast television.

1.3 Attack

Next, we acknowledge that sometimes the best defense is a good offense, as Chapter 3 explores attack strategies focused on the physical layer. By thinking like an adversary, one can better anticipate possible attacks and determine the appropriate remedies. We begin with a survey of cogent, protocol-specific attacks that rely upon PHY layer exploitation and have been implemented with current SDR platforms. Acknowledging that most 3G and 4G wireless standards and protocols incorporate some form of multi-input multi-output (MIMO) technology, we perform an analysis on the physical layer weaknesses of such systems. Proper operation of MIMO systems demand accurate and timely knowledge of the wireless channel, whether at the transmitter, the receiver, or both. Our work explores the effects of inaccurate channel knowledge, and proposes methodology to force misestimation in MIMO systems. We target the two most popular MIMO schemes — the capacity achieving singular value decomposition (SVD) based scheme, and

the practical Alamouti space-time block code (STBC) scheme. The vulnerabilities that we develop are illustrated via real-world attacks. We note that the majority of related research has assumed the presence of unintelligent jammers rather than truly smart adversaries.

1.4 Mitigation

In Chapter 4 of this thesis, we consider what to do once an accurate description of the operational scenario is achieved. Accurate knowledge of the environment plays a key role in Dynamic Spectrum Access (DSA) [12, 13, 14], where devices adapt modulation schemes and protocols to both optimize communications and minimize interference with existing wireless infrastructures. Additionally, accurate situational awareness provides insight into potential communication hazards — from severe multipath conditions to adversarial attacks. In this Chapter, we present unique physical layer methodologies that can be used to overcome channel degradations due to *both* natural phenomena and adversarial activity.

We begin the Chapter with an overview of traditional mitigation strategies such as simply transmitting with more power. We continue by addressing the channel estimate vulnerabilities introduced in Chapter 3. Rather than focusing specifically on MIMO systems, however, we present a general physical layer based protection scheme to authenticate channel state information estimates. The techniques are then validated through SDR experimentation in real environments.

Next, we pay attention to the difficulties associated with reliably communicating in high multipath environments. With the added complexity of a high powered jammer, we develop strategies to overcome such daunting conditions by leveraging the multipath in our favor. Again, these techniques are verified in real world experiments using a software defined radio.

Figure 1.2: Pictured is the GNU Radio/USRP SDR platform used for experimentation in this thesis. The RF to baseband conversion is provided by the USRP, thus allowing GNU Radio processing blocks to be run on the host platform.

1.5 Experiment Methodology

In each part of this thesis, we accompany theoretical results and findings with simulations and real-world experimentation in order to illustrate the feasibility and applicability of the proposed techniques. Real-world implementations were conducted using a current SDR architecture — the Universal Software Radio Peripheral (USRP) [6] and GNU Radio [7]. GNU Radio is an open source, free software toolkit that provides a library of signal processing blocks that run on a host processing platform. Algorithms implemented using GNU radio send/receive baseband data directly to/from the USRP, which is the hardware that provides radio frequency (RF) access via an assortment of daughterboards. Many of these SDRs are deployed in the WINLAB Orbit grid at Rutgers University [15], and are publicly available for use. The USRP/GNU Radio SDR platform is depicted in Figure 1.2 and described in more detail in the following section.

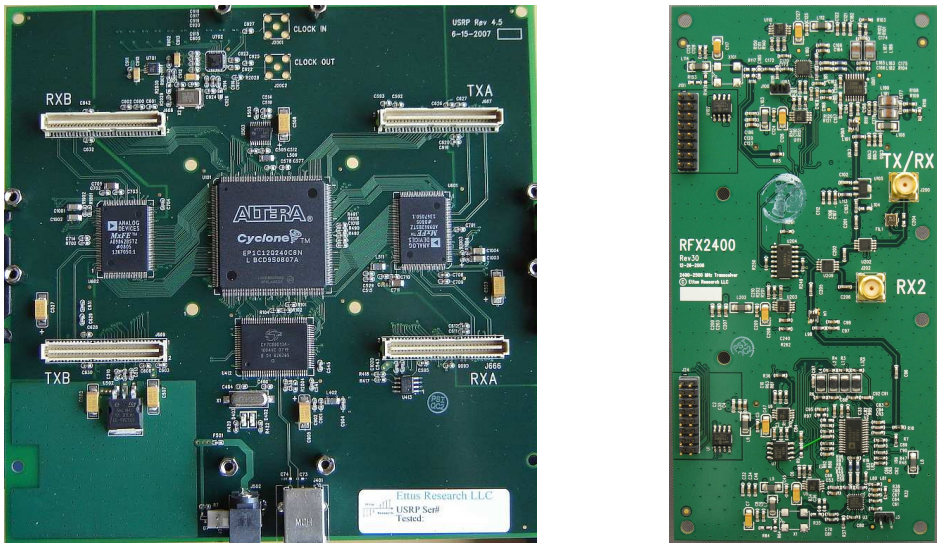

(a) USRP Motherboard

(b) RFX-2400 Daughterboard

Figure 1.3: Pictured above is the USRP motherboard (a) and the RFX-2400 daughterboard (b). The RFX-2400 provides RF to IF translation for the 2.4 GHz ISM band, and the USRP motherboard is responsible for the upconversion, downcoversion, and channelization. Baseband data is interchanged with the host processing platform over a USB 2.0 interface.

1.5.1 USRP

Motherboard

The USRP motherboard is the core hardware unit; it supports the simultaneous transmission and reception of four real or two complex channels in real-time. For reception it utilizes four 12-bit analog-to-digital converters (ADCs) operating at 64 MHz, and four digital-downconverters (DDCs) with programmable decimation rates. The transmit side of the USRP incorporates four 14-bit digital-to-analog converters (DACs) that operate at 128 MHz, and two digital-upconverters (DUCs) with programmable interpolation rates. The on-board Altera Cyclone FPGA is responsible for the channelization, down-conversion, and up-conversion. Data is transferred between the host computer and the USRP via a USB 2.0 interface, therefore limiting the sustainable data rate to 32 MBps half-duplex. Default

Table 1.1: RF Daughterboards

Name	**Frequency Range**	**TX**	**RX**
BasicRX	0.1 - 300 MHz		✓
BasicTX	0.1 - 300 MHz	✓	
LFRX	DC - 30 MHz		✓
LFTX	DC - 30 MHz	✓	
TVRX	50 - 80 MHz		✓
DBSRX	800 - 2400 MHz		✓
RFX-400	400 - 500 MHz	✓	✓
RFX-900	750 - 1050 MHz	✓	✓
RFX-1200	1150 - 1450 MHz	✓	✓
RFX-1800	1500 - 2100 MHz	✓	✓
RFX-2200	2000 - 2400 MHz	✓	✓
RFX-2400	2300 - 2900 MHz	✓	✓
XCVR2450	2400 - 2500 MHz	✓	✓
	4900 - 5900 MHz	✓	✓
WBX	50 - 2200 MHz	✓	✓

operation results in 16-bit complex sampling, thus delivering an effective total bandwidth of 8 MHz. Modification of the sample size to 8 and 4 bits is supported, thus sacrificing sampling accuracy for respective increased bandwidths of 16 and 32 MHz. For the experiments conducted in this thesis, we employed the default 16-bit complex sample size for both receive and transmit. The USRP motherboard is shown in Figure 1.3 (a).

RF Daughterboards

The USRP motherboard by itself does not provide direct RF access, but rather it utilizes an assortment of daughterboards to perform the RF to IF translation. These daughterboards plug directly into the USRP motherboard. Some daughterboards are receive-only, some are transmit-only, and some are transceivers — devices that provide both receive and transmit functionality. Table 1.1 lists the daughterboards associated with the USRP platform. Throughout this thesis, we refer to the daughterboards by their names listed in the table. Figure 1.3 (b) depicts the RFX-2400, which provides access to the 2.4 GHz ISM band.

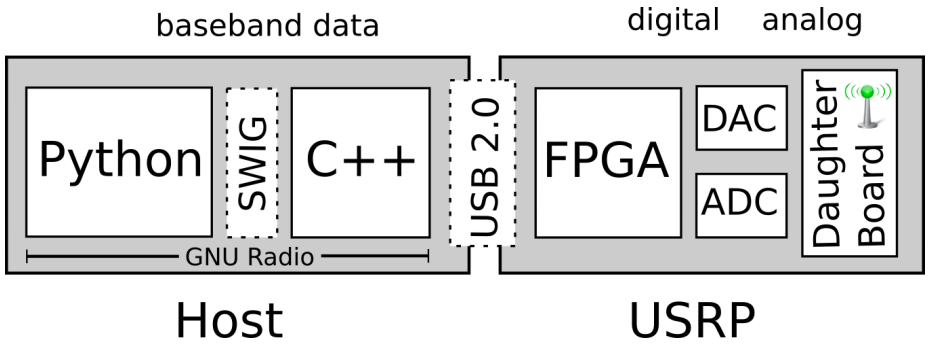

Figure 1.4: The general USRP architectural flow is depicted above. A Python flowgraph is created by connecting GNU Radio processing blocks. The blocks are implemented in C++ and made callable in Python via SWIG. Baseband data is transferred between the USRP and the Host over a USB 2.0 interface. The FPGA on the USRP is responsible for upconversion, downcoversion and channelization of the baseband data. The DACs and ADCs provide the analog to digital transformation, and daughterboards provide RF access.

1.5.2 GNU Radio

GNU Radio is an open source, free software toolkit that provides a library of signal processing blocks for developing communications systems and experiments. The processing blocks are written in C++, and made callable from Python via SWIG (Simplified Wrapper and Interface Generator). Standard practice is to create a Python flowgraph by connecting the appropriate signal processing blocks. If a custom signal processing block is desired, it can be created in C++ and integrated into the GNU Radio development environment via SWIG. The overall operational flow is depicted in Figure 1.4.

Chapter 2
Situational Awareness

2.1 Motivation

In future wireless communication networks, it will be important for devices to obtain accurate situational awareness in order to maintain efficient, effective, and secure communications. Two main categories of situational awareness emerge: (1) identifying available resources and active devices and (2) recognizing environmental conditions in the region. Many methods of acquiring these types of situational knowledge exist.

Addressing the issue of service discovery, researchers have proposed the introduction of a global control channel whereby devices can be informed as to the policies and services that are present in a region [16]. But if a control plane does not exist, devices may wish to perform their own discovery routines. One option would be for cognitive platforms to implement full protocols in order to engage and probe the existing services and devices. This however can prove timely, envelope ample resources, and in fact may be ineffective with uncooperative networks or devices. Additionally, many cognitive platforms may be incapable of implementing the full protocols of existing services. While the device would then not be able to associate with such networks, it is still advisable to know of the network's existence in order to select a communication method that will work in the region *without* ample interference. As an example, a device implementing the 802.11 protocol may wish to know about 802.16 devices in the region due to co-existence results in [17]. In the sections immediately following, we consider the physical

limitations of cognitive platforms while harnessing their lower layer access in order to effectively and efficiently perform service discovery and device identification by leveraging physical layer information and partial protocol knowledge.

In the later sections, we investigate methods of obtaining environmental knowledge. To better operate in future wireless networks, a device should know if the channel is line-of-sight or riddled with severe multipath, if it is mobile, or if it is under an adversarial attack. We will examine new methods of determining such knowledge based upon physical layer reconnaissance.

2.2 System Details

Experiments conducted in this Chapter leveraged the USRP SDR platform describe in detail in Section 1.5. For RF access in Section 2.3, we used the RFX-2400 daughterboard to conduct our experiments, as the focus was on the ISM band. In Section 2.4, we employed a variety of daughterboards in order to access cellular, FM, television, and WiFi signals. These daughterboards included the DBSRX, TVRX, RFX-1800, and RFX-2400.

An important note about our experiments relates to our usable spectral snapshot. Due to limitations of our host computer's USB 2.0 interface, we were restrained to a usable spectral bandwidth of 4 MHz (using complex 16-bit samples).

2.3 Resources and Devices

2.3.1 Service Discovery

Spectral awareness plays an essential role in Cognitive Radio (CR) operation. In order for a CR to make intelligent operational decisions it should be knowledgeable of available services. For example, if a CR desires to access the Internet via WiFi networks, it may decide to join one existing WiFi network, or perhaps set up

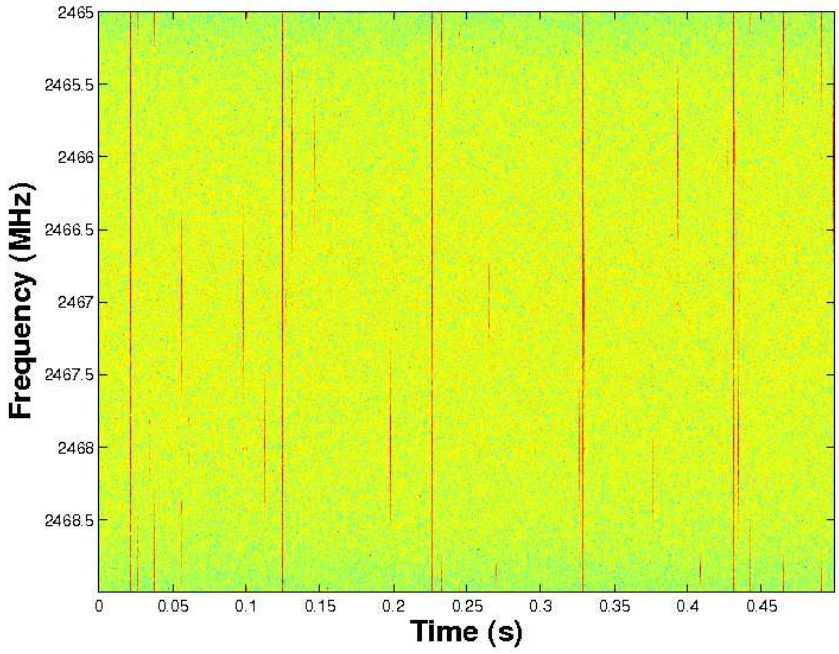

Figure 2.1: Spectral activity is displayed for the ISM band. A spectrogram is shown for 0.5 seconds of activity over 4 MHz of spectrum centered at 2467 MHz.

its own WiFi network. Here it is advantageous for the CR to monitor various portions of the ISM band in order to estimate the best channel to use. The mere presence of other WiFi networks in addition to existing Bluetooth piconets would be important information to be used to select an optimal channel.

Using GNU Radio, the USRP board, and the RFX-2400 daughterboard [6, 7], we illustrate it is possible to detect services and devices with relatively narrow-band spectral surveillance (i.e. 4 MHz), in spite of the fact that the underlying protocols themselves may employ a broader spectral range (e.g. 20 MHz in the case of WiFi). We note that, although we illustrate our techniques for Bluetooth and WiFi, our strategies may be applied to identify other wireless technologies.

Figure 2.1 shows the spectrogram of a snapshot from our platform that was centered at 2467 MHz. Note that the data spans half of a second and lies within the spectral range of WiFi devices. Upon examination, one immediately notices

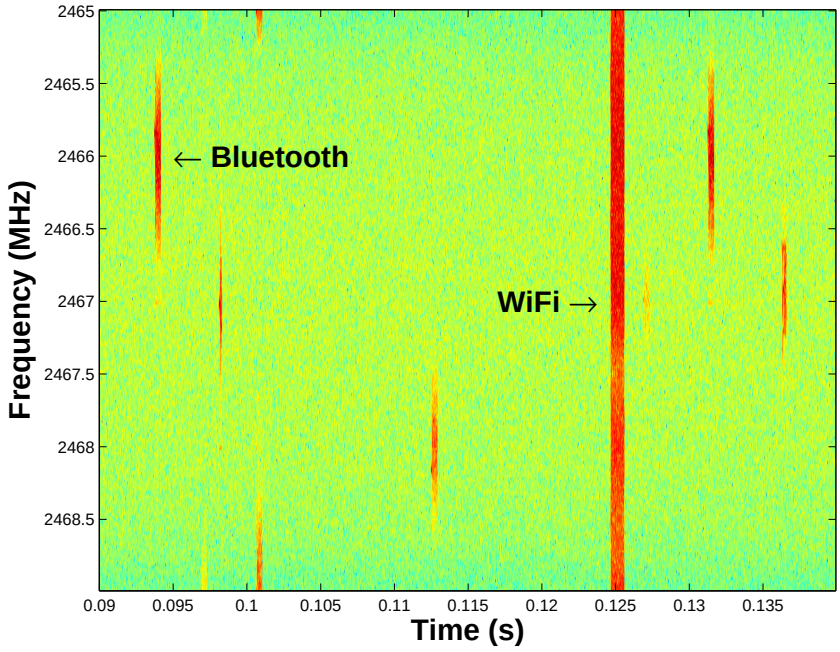

Figure 2.2: A closer look into the spectrogram of Figure 2.1. Note the Bluetooth bursts hopping through the band amidst the WiFi beacon.

various broadband, periodic bursts. These are in fact WiFi beacons being broadcast by an 802.11g Access Point (AP) in range of our receiver. WiFi (802.11g) beacon signals are 20 MHz in bandwidth and default to a pulse-repetition interval (PRI) of 102.4 ms on most APs.

Observe that there are other bursts hopping through the spectrogram. A zoomed-in view of the data collection from 0.09 to 0.14 seconds (Figure 2.2) reveals that the bursts span 1 MHz and vary in burst length. These sporadic narrow-band bursts are in fact Bluetooth bursts. The Bluetooth protocol mandates that transmissions frequency hop over 79 MHz of the ISM band while maintaining an instantaneous frequency of 1 MHz using Gaussian Frequency Shift Keying (GFSK) [18]. Bluetooth also exhibits a time-division multiplexed (TDM) nature where timeslots exist to guide network transmissions. A timeslot is 625 μs in length, and devices are allowed to transmit continuously for 1, 3, or 5 timeslots.

Detecting the presence of services is plausible by monitoring narrow-band

Algorithm 1: Service discovery in the ISM band

```
while (1) do
    /*** Data Collection ***/
    Collect new samples;
    /*** Time/Frequency Analysis - detect individual bursts ***/
    num_burst = 0;
    for ( each burst in time ) do
        for ( each distinct frequency ) do
            /*** extract physical properties ***/
            leading_edge[num_burst] = start_time;
            trailing_edge[num_burst] = end_time;
            ...
            bandwidth[num_burst] = burst_bw;
            num_burst++;
        end
    end
    /*** Service Discovery Phase ***/
    for ( each detected burst ) do
        if ( ( burst_width ≤ 5 timeslots ) && ( bandwidth ∼ 1 MHz ) ) then
            Bluetooth service found;
        end
        if ( ( bandwidth ∼ min (sample_rate, 20MHz) ) && ( periodic ) ) then
            WIFI AP;
        end
        /*** Checks for other available services ***/
        ...
    end
end
```

activity and leveraging protocol-specific features. As listed in Algorithm 1, CRs first collect data and identify bursts. Then, properties of individual bursts are extracted. Finally, CRs decide that a service is present if the burst patterns and properties match the intrinsic time and frequency properties of the corresponding protocol. For example, the presence of various 1 MHz GFSK bursts that do not exceed 5 Bluetooth timeslots in length indicates the presence of a Bluetooth piconet. Likewise, periodic, broadband bursts suggest the presence of a WiFi network.

Although using a small amount of spectrum to identify broader band services

is inherently desirable due to the associated reduction in sampling and computation, a natural questions that arises is whether there is any need for a bandwidth-limited CR device to know about services it cannot utilize. For example, with our experimental setup, we cannot join an 802.11g network due to the bandwidth limitation of the USRP and the USB 2.0 interface. The most obvious answer is that knowledge of services operating in the region plays an important role in choosing an optimal communication scheme. To illustrate, consider the co-existence issues inherent to WiFi and Bluetooth [17]; knowledge of only 802.11g services operating in a region may suggest that a CR should not choose to implement a frequency-hopping spread spectrum (FHSS) scheme.

Establishing knowledge of all existing services is therefore important. With this in mind, we present a general PHY/MAC classifier in Figure 2.3. The classifier expands upon the notion of service discovery, addressing prominent protocols outside of the ISM band such as cellular services (e.g. GSM and CDMA) and broadcast media (e.g. radio and television broadcasts). As an example, detection of a TDMA, constant-envelope signal in the 1800 MHz band with a burst width of 577 μs is a clear indicator of an active GSM network.

While knowledge of the presence of certain services can guide CR operations by giving a hint of spectral activity, it would be more advantageous for a CR to know more detailed information about the services present in the RF environment. For example, it would be very useful to know how many WiFi networks and Bluetooth piconets exist. It would be even more desirable to know how many distinct WiFi or Bluetooth devices are operating in the region. We now discuss how this information can be extracted from our narrowband data by leveraging protocol-specific properties.

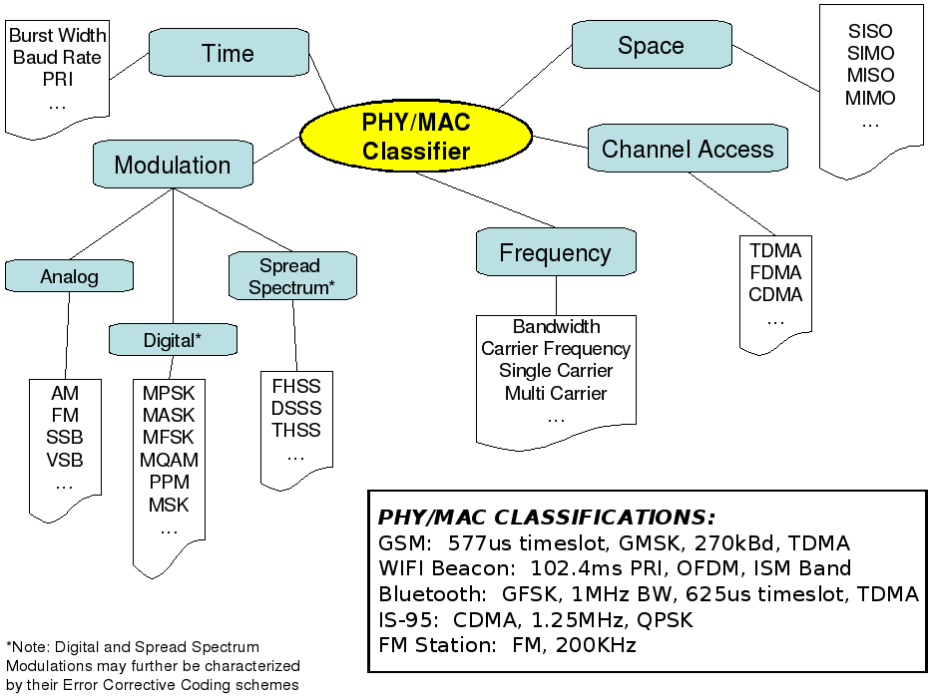

Figure 2.3: A general PHY/MAC classifier is illustrated. Using PHY and MAC based feature vectors provides reliable service discovery across the RF spectrum. For example, detection of TDMA, constant-envelope bursts using GMSK modulation and burst widths of 577 μs in the 1800 MHz band indicates GSM operation.

2.3.2 Device Identification

In this section, we investigate the problem of device identification in both Bluetooth and WiFi networks. For Bluetooth, we perform device identification on two levels. First, we detect distinct Bluetooth piconets; second, we identify individual Bluetooth devices. In a similar manner, we identify distinct WiFi Access Points and then illustrate how our methodology can be extended to identify individual WiFi devices.

Bluetooth Piconets

Every piconet has 1 device that acts as the Master, and up to 7 active devices operating as Slaves. A piconet is always synchronized to its Master's clock, and

timeslots are defined based on this reference. The Master is only allowed to begin transmitting its bursts at the beginning of even timeslots, while the Slaves may only begin their transmissions at the beginning of odd timeslots [18].

Time-binning approach: By analyzing the detected leading edge times, we can determine a lower bound on the number of piconets. Since piconets operate independently, we can view a particular piconet's timeslot structure as a uniformly distributed random variable between 0 and $625\mu s$ ($U(0, 625)\mu s$), where the random variable represents the relative starting point of a new timeslot. By dissecting an arbitrary timeslot into sub-intervals, we can collapse all of the starting times into bins. Starting times that fall into the same bins are then likely to belong to the same piconet. Since the specification allows for 20μs of time uncertainty ($\pm 10\mu s$ of jitter), it is wise to choose a time bin resolution δ slightly bigger than the maximum allowed uncertainty window (i.e. $20\mu s$).

In our case, we chose $\delta = 25\mu s$, resulting in 25 time bins. Analysis of our data revealed all Bluetooth bursts falling into the same bin, validating that they all came from the same piconet. As an example, consider the leading edge times of the bursts transmitted at 2466 MHz shown in Figure 2.2. The first burst is at $t_1 = 0.09378$ seconds while the second burst is detected at $t_2 = 0.13127975$ seconds. Thus, the second burst is $250ns$ shy of 60 timeslots away from the first burst ($(t_2 - t_1)/t_s \approx 60$; $(t_2 - t_1) \mod t_s = -250ns$). Since $250ns$ is much smaller than the time uncertainty window $20\mu s$, we conclude that these bursts belong to the same piconet. Using GNU Radio and the USRP, we developed a Bluetooth Access Point Detector using this approach. Figure 2.4 depicts a screenshot of our real-time histogram, where a single Bluetooth network was active.

It is possible that the time-binning approach will falsely declare two independent piconets as the same piconet when two piconets have overlapping uncertainty windows. One way to detect this would be to look for a scenario where two bursts are transmitted at the same time but on different frequencies. This would be a

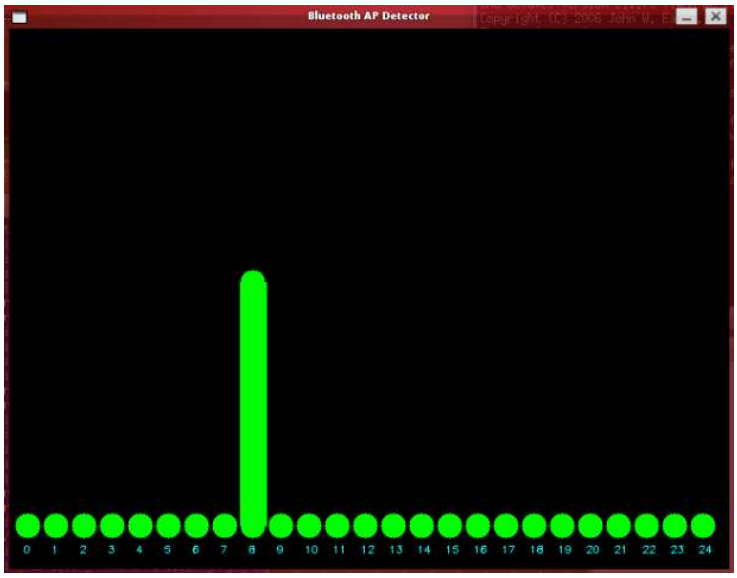

Figure 2.4: A PHY-based Bluetooth Piconet detector was implemented using GNU Radio and the USRP. Leading-edge burst times are mapped to an intra-timeslot based histogram as described in the time-binning approach of this Chapter in order to estimate the number of active piconets. The experiment was run during a call to voicemail using a Bluetooth enabled cellular phone and Bluetooth earpiece dongle. Results clearly show the presence of a single Bluetooth piconet. The y-access reflects the number of detected bursts in the processing window.

clear indicator of separate piconets. Additionally, real-world devices inherently exhibit local oscillator drift. This drift is independent from device to device and will cause a device's leading edge bin to drift over time. This phenomenon was verified during real-world experiment using GNU Radio and the USRP. Yet another alternative approach exists — examination of the actual bits transmitted.

Bit comparison approach: We emphasize that such an approach does not require a full implementation of the Bluetooth protocol, but merely requires some knowledge of basic modulation schemes.

The general Bluetooth packet structure, as depicted in Figure 2.5, is comprised of an access code, a header, and a data payload. All packets must have an access code, while the presence of the header and the data payload depends on the type of message conveyed. For actions such as paging and inquiry scans, the packet only

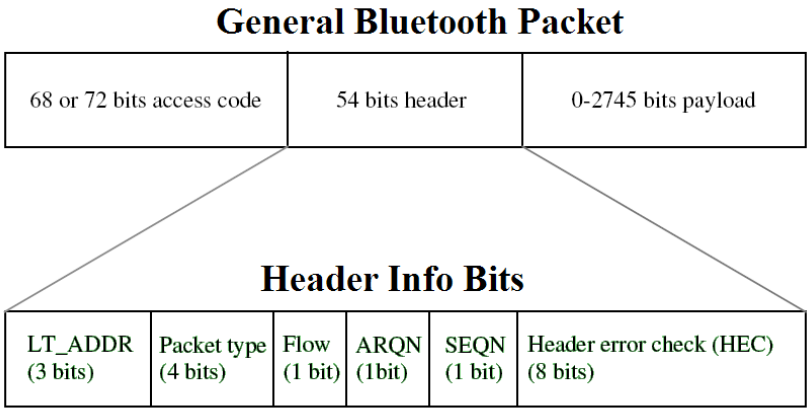

Figure 2.5: The Bluetooth packet structure is illustrated. The 18 header information bits are repetition encoded by 3 to generate the 54-bit packet header.

contains a 68-bit access code. During normal operation within a basic piconet, all packets begin with a 72-bit access code known as the Channel Access Code (CAC). This is followed by a header and when pertinent, payload data. The CAC is derived from the Master's unique device address and is therefore particular to a given piconet. By demodulating the bursts and comparing CACs, we can better estimate the number of distinct piconets, even when the piconets are synchronized to within the same time uncertainty window.

We now illustrate the CAC comparison procedure with the two Bluetooth bursts that we analyzed earlier. In order to obtain the access codes, we must properly demodulate the bursts. The second and third subplots in Figure 2.6 depict the demodulation, where the instantaneous frequency is plotted versus the sample number for each respective burst. The first subplot is the normalized burst power. Since high energy indicates the presence of a packet, the relative burst power plot is shown to illustrate the actual start and end point of the packets. During the packet transmissions, the reader can clearly see the fluctuation between two distinct frequencies as specified in the protocol (i.e. GFSK).

The bit-wise comparison of our two demodulated Bluetooth bursts is shown

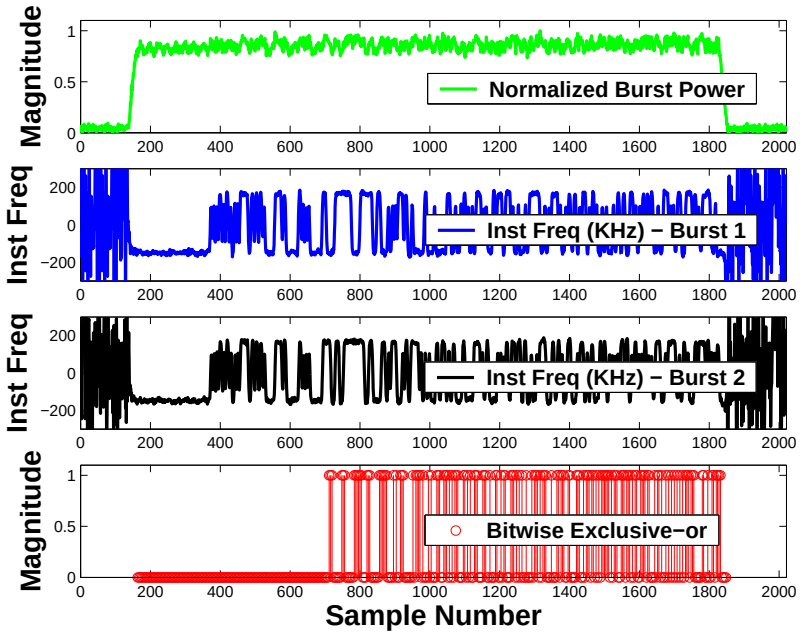

Figure 2.6: Bluetooth Piconet differentiation is performed by demodulating and comparing Channel Access Codes (CACs) from two Bluetooth bursts. The CAC is conveyed by 72-bits at the beginning of each Bluetooth packet. Bit differences are illustrated via an exclusive-or between two demodulated Bluetooth bursts. A 0 indicates bit-agreement, whereas a 1 indicates bit-disagreement. It is clear that the access codes are the same, therefore confirming that these bursts came from the same piconet. Note that the graph is organized to incorporate our sampling rate of 4 samples per bit.

in the fourth subplot in Figure 2.6. An exclusive-or was performed using the two bit sequences. A value of 1 indicates bit disagreement, whereas a value of 0 signifies commonality. The bit disagreement in the plot illustrates the location of the packet header and payload. The leading 0s indicated that the 72-bit CACs from the two Bluetooth bursts are identical. Therefore, these packets belong to the same piconet.

Bluetooth Devices

The upper bound on the number of distinct active Bluetooth device in the region is $8\times$ (*number of active piconets*), since only 7 active Slaves are allowed per piconet.

A better estimation can be achieved by further leveraging some protocol-specific information.

As there can be multiple Slaves communicating with the Master, each packet needs to contain the identity of the Slave involved in the communication. This information is found in the Logical Transport Address (LT_ADDR) in the packet header. The LT_ADDR is a 3-bit information field derived from the first 9 transmitted bits of the packet header. (The packet header uses a simple 3-bit repetition procedure for its Forward Error Correction (FEC), and techniques in [19] can be employed to unwhiten the data). The Master denotes the destination of a transmitted packet by specifying the intended Slave in the LT_ADDR. In a similar manner, Slaves include their own LT_ADDR when transmitting packets to the Master. The LT_ADDR of 0 (000) is reserved for broadcast packets, while $1-7$ (001-111) correspond to particular Slaves. By monitoring the LT_ADDRs of the basic piconet packets, we can identify distinct users of a particular piconet, and therefore obtain a better estimate of usage within the entire channel. The two Bluetooth bursts that we examined above had LT_ADDRs of '000', indicating that they were broadcast packets. Other bursts revealed an LT_ADDR of '001' for the same CAC, indicating that the specific piconet that we detected had 2 active users (i.e. the Master and 1 Slave).

We have shown that leveraging PHY layer information clearly results in a reliable estimate of Bluetooth services and unique devices. However, a CR should also account for the stationarity or transience of the networks being observed. As such, it makes sense for a CR to maintain a table of information that is updated with the most current information. In this manner it can phase out old data due to periods of inactivity. This is especially important in Bluetooth, since the protocol allows Slaves to switch functionality with the piconet Master. As a general process, we propose that every burst we detect have its CAC compared to a table of access codes for known piconets. If it does not match any of the

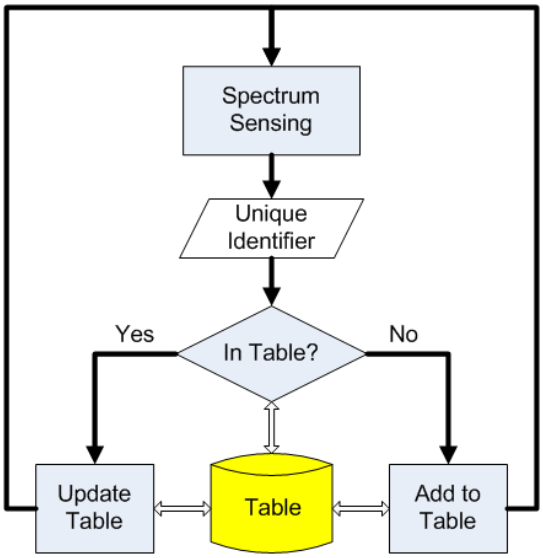

Figure 2.7: Unique identifier table population is illustrated. For Bluetooth, unique identifiers include the Channel Access Code (CAC) and the Logical Transport Address (LT_ADDR).

known piconets, then we insert the newly discovered piconet information into the table. Likewise, we can keep track of individual Slaves by their LT_ADDR *and* CAC. This general process is illustrated in Figure 2.7.

WiFi Access Points

WiFi beacons are periodic, broadband bursts that are broadcast by Access Points. These beacons contain useful identification information such as the service set identifier (SSID) and the AP name. Beacons are always present, as they function as the heartbeat of a given network (and are present even if the AP is hidden). Beacon information can be immediately extracted by properly demodulating the beacon signals. But, this is a bit more complicated given the limitations of our research hardware. Since the WiFi (802.11g) beacon is 20 MHz wide, it is out of scope of our equipment (i.e. 4MHz). However, alternative methods of AP identification can be pursued using our limited spectral snapshot. One method

performs an analysis on the periodic structure of the beacons, while another examines the physical properties of each beacon in more detail.

Periodicity: By default, WiFi APs broadcast beacons every 102.4 ms. Given a situation where every AP maintains the default periodic repetition interval (PRI), we can estimate the number of APs by monitoring the number of beacons that occur within a 102.4 ms period. In cases where APs do not use the default beacon PRI, an accurate estimation of the number of APs is also discernible by analyzing the start times of the observed beacons. Assuming AP PRIs are stationary, all unique PRIs can be determined using standard deinterleaving algorithms [20]. We are then back to the original task of identifying unique APs given knowledge of a beacon PRI.

Channel estimation: A particular scenario arises, however, where standard deinterleaving algorithms break down. Consider the case where a beacon frame is detected every 51.2 ms. There is an inherent uncertainty here in determining how many distinct APs exist. There could only be one with a PRI of 51.2 ms, or two with PRIs of 102.4 ms, or three with PRIs of 153.6 ms, etc. By leveraging the unique capabilities inherent to CRs, the dilemma can in fact be resolved. Since a CR by nature has direct access to the physically received waveform, it can exploit non-standard features to best determine a physical differentiator (e.g. other than just bits). Let us elaborate by discussing the 802.11g protocol a bit more in detail.

WiFi beacons operating in 802.11g only mode begin with an 8 μs training sequence followed by an 8 μs equalization sequence [21]. As previously stated, an 802.11g WiFi channel is 20 MHz wide. Being an OFDM signal, one channel contains 64 equally spaced sub-channels, where each sub-channel is 312.5 KHz wide. During the training sequence, every fourth sub-channel is active with a phase relationship such that the peak-to-average power ratio is minimized. Subsequently, the equalization sequence modulates every sub-channel with equal power [21].

It is therefore intuitive that the equalization sequence would make an excellent differentiator since it is in fact a channel sounding waveform between the AP and the CR. Figure 2.8 shows both the training sequence and the equalization sequence for one of our observed WiFi beacons. The reader can see the distinguishing spectral characteristics of the two sequences, even given our narrowband snapshot.

A CR may wish to take advantage of this broadcasted channel estimation signal by using it as a means of unique AP identification. Note that an estimate of the channel can be obtained by accounting for the spectrum of the transmitted waveform. This adjustment factor is known a priori and it will remain common for all channel estimates since the waveform is always the same. Let $s(t)$ be the transmitted equalization signal with $S(f)$ corresponding to its Fourier Transform. If we denote the channel between the AP and the CR as $h(t)$, and channel noise as $n(t)$, then a basic linear time invariant model for the waveform received by the CR is:

$$r(t) = s(t) * h(t) + n(t)$$
$$R(f) = S(f)H(f) + N(f)$$

Estimating the channel spectrum yields:

$$\hat{H}(f) = \frac{R(f)}{S(f)} = H(f) + \frac{N(f)}{S(f)}$$

Access Point Discrimination

With a valid channel estimate, it is proposed that unique APs can be differentiated within the time-coherence of the channel by performing cross-correlations between new channel estimates and known channel estimates. Given K known users and their corresponding channel estimates, $\hat{H}_i(f)$, where $i = 1, ..., K$, we can obtain a vector of cross-correlations, $\mathbf{v}$, between known channel estimates and

the newest channel estimate, $\hat{H}_{K+1}(f)$, via:

$$v(i) = \frac{< \hat{H}_i(f), \hat{H}^*_{K+1}(f) >}{\|\hat{H}_i(f)\|\|\hat{H}^*_{K+1}(f)\|} \tag{2.1}$$

If the largest correlation does not exceed a given threshold, then we can declare the presence of a new user. Similar to the case of Bluetooth devices discussed earlier, a table can be maintained to aid in AP discovery. And, just like the Bluetooth user table, the WiFi AP table can be updated over time in order to remove old users. Note that it is also possible to update the table to account for the time-coherence of the channel.

However, it should be noted that the aforementioned channel estimate, $\hat{H}(f)$, has one subtle flaw: the use of phase as a differentiator when considering bursty transmissions from a common source. This is because there is no deterministic way to estimate the initial phase of the transmitted waveform. Initial phase offsets can be attributed to various factors not limited to local oscillator drifts and even software. We must therefore consider the initial phase to be random and thus should not include it in our channel estimate. We can however base our channel estimate on the magnitude response of the channel. The following equation illustrates the relationship of this new channel estimate, $\hat{H}_M(f)$, to the old channel estimate, $\hat{H}(f)$. Utilizing solely the magnitude response in performing the correlations has a performance effect. The typical correlation range of [-1,1] is now collapsed to [0,1], as seen by bounding $\hat{H}_M(f)$ by $\|\hat{H}(f)\|$:

$$\begin{aligned}
\hat{H}_M(f) &= \frac{\|R(f)\|}{\|S(f)\|} \\
&= \frac{\|S(f)H(f) + N(f)\|}{\|S(f)\|} \\
&\leq \frac{\|S(f)H(f)\| + \|N(f)\|}{\|S(f)\|} \\
&\leq \|H(f)\| + \frac{\|N(f)\|}{\|S(f)\|} = \|\hat{H}(f)\|
\end{aligned}$$

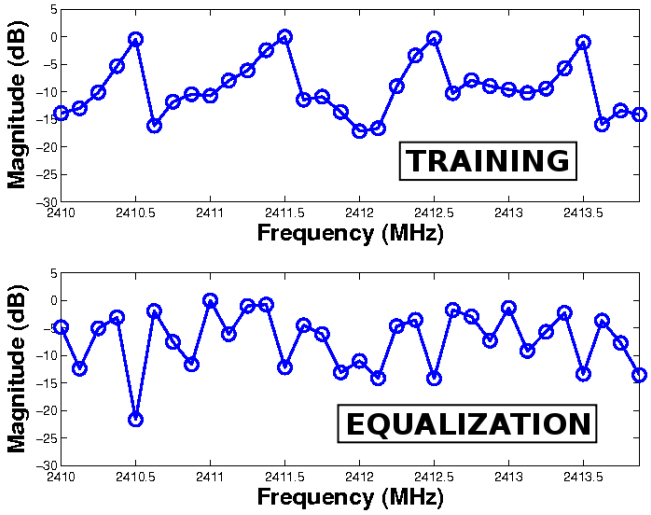

Figure 2.8: The first $8\mu s$ sequence of the WiFi beacon is used as a training sequence, while the second $8\mu s$ sequence is used to aid in equalization. The training sequence uses every fourth sub-channel, while the equalization sequence is modulated with equal power over every 312.5 KHz sub-channel.

Experimental Validation

To validate the feasibility of using channel estimation to identify APs, we conducted an experiment using the beacons of two APs at different locations. The devices were placed much further than a wavelength away from each other and the USRP. Figure 2.9 (a) shows the time-series magnitude of one of the data sequences. Two sets of periodic beacons with varying amplitudes are evident.

Since our narrowband snapshot limits us to 32 samples to represent the equalization sequence, our correlations were performed over a longer data sequence. In order to prevent this action from jeopardizing our experiment, both APs were set up with exactly the same parameters (e.g. Channel, SSID, PRI, name, etc.). This forced both APs to transmit identical bit-sequences. Our analysis used the appropriate waveform, $S(f)$, to calculate our channel estimate, $\hat{H}_M(f)$. With a wider bandwidth, however, it is desirable to only use the equalization sequence for channel estimation.

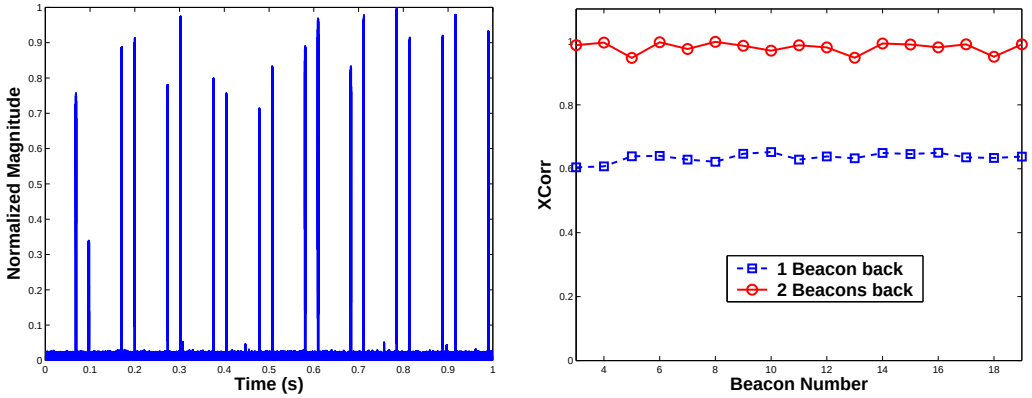

(a) Normalized Time-Series

(b) Channel Estimate Cross-Correlations

Figure 2.9: (a) Normalized time-series beacon data for two active 802.11g APs is shown. Note both the periodicity and the varying magnitudes. (b) Channel estimate cross-correlation results against the previous 2 beacons show strong correlations for the appropriate beacon signal.

Using our outlined methodology, every detected beacon resulted in a channel estimate. That channel estimate was then correlated against every channel estimate seen within the last 102.4 ms. In our experiment, this resulted in two correlations per detected beacon. Since our experiment consisted of two distinct APs alternately transmitting, one would expect to correlate best with the second to last beacon detected. Figure 2.9 (b) presents our results. As can be seen, the correlation routine was reliable in its ability to differentiate between APs based solely on the physical layer signature associated with the AP. The results were reproducible for various AP locations and over numerous days. It is quite clear that the proposed methodology therefore provides a plausible solution to the previously outlined deinterleaving dilemma.

WiFi Devices

Now that unique WiFi APs can be reliably detected, we propose that the same correlation methodology can be leveraged to estimate the number of distinct WiFi devices in the area. This includes both APs and clients. While data packets do

not exhibit the periodicity of the beacon frames, they do in fact contain known bit-sequences. These bit-sequences can be utilized to obtain unique channel estimates, and these channel estimates could be exploited to differentiate between distinct WiFi devices. One notable challenge, though, is that mobility might affect the temporal coherence of the channel and, consequently, necessitate more frequent comparisons.

2.4 Environment

We now investigate methods of obtaining environmental knowledge. To better operate in future wireless networks, a device should know if the channel is line-of-sight or riddled with severe multipath, if it is mobile, or if it is under an adversarial attack. Here, we examine new methods of determining such knowledge based upon physical layer reconnaissance. We refer to our suite of algorithms as PLATEAU: Physical Layer Techniques for Enhanced Situational Awareness.

2.4.1 PLATEAU

Wireless devices permeate our society, and with the influx of so many new devices and services, in order for future wireless devices to operate effectively, it is essential that they have awareness of their operating conditions so as to support appropriate adaptation of communication waveforms and protocols. At a minimum, this involves scanning spectrum for occupancy and fallow bands [22]. However, scanning spectrum is but a small component of the scenario describing a wireless device and its environment, and additional forms of information, such as whether the device (or its neighbors) are mobile, or whether the device is indoors/outdoors or moving between, are valuable and would help support protocol adaptation.

Although it may be possible to augment future radio platforms with additional

sensors in order to determine the isolated aspects of a device's operating conditions (for example, inertial sensors have been proposed for determining whether a device is mobile), such methods would ultimately require additional hardware and lead to increased costs and larger form-factors. Ultimately, though, it would be more desirable for a wireless device to use its native functionality, i.e. the radio itself, to assess its situation without having to resort to extraneous and costly methods for gaining situational awareness.

In this Section, we show that it is possible for a wireless device to solely use its ability to receive and analyze ambient signals in order to answer fundamental questions about a its operating conditions. In particular, we present a collection of algorithms, which we call PLATEAU (Physical Layer Techniques for Enhanced Situational Awareness), that infer coarse levels of mobility and location-oriented situational awareness. We show that it is possible to devise signal processing algorithms that extract physical layer information arising from the wealth of existing, ambient radio signals (such as FM, GSM Downlink, Advanced Television Standards Committee (ATSC) signals), to acquire accurate situational awareness. We describe two classes of PLATEAU's algorithms: mobility determination, and indoor versus outdoor discrimination. We support the validity of our algorithms by utilizing measurements from a software defined radio platform (SDR).

Related Work

Work has been done to utilize physical layer information for limited situational awareness. In [23], the authors leverage PHY layer information to discover services and detect devices. Our algorithms extend these ideas in order to extract even more information from the physical layer.

In [24], the authors investigate and implement a passive radar system by using GSM base stations. The principles used however are to implement a tradition passive radar system, where a line-of-sight (LOS) signal and a reflection are need to

obtain a directional estimate. Our algorithms reduce the operational complexity by only requiring a single arriving signal.

Researchers in [25] detect co-located transmitters based upon the observation of similar fading characteristics. This is the reverse view of the problem that we address. Further, the authors only use WiFi client transmissions, which means that they cannot guarantee constant transmission power.

PLATEAU Overview

PLATEAU provides situational awareness by exploiting physical layer signal properties that are already present in the environment. By applying machine learning and passive radar techniques, a device can accumulate accurate situational awareness to drive intelligent operational decisions. The best signals to leverage are broadcast signals from stationary transmitters across a wide array of frequencies, and prime candidates include ATSC, FM, GSM Downlink, and WIFI Access Point signals. In this paper, we focus on signals that have constant envelope and/or are broadcast with constant power. Such signals include GSM Downlink Control Channels [26], WIFI beacons [27], and the ATSC pilot tone [28]. In this section, we explore methods of exploiting the aforementioned signals to acquire two main types of coarse situational knowledge - *mobility* and *location*.

Mobility

If a device is mobile, it should choose a communication scheme that can mitigate fading effects and the non-stationarity of a time-varying channel. Hence, mobility information is very valuable for guiding communication protocol adaptation. We now propose three methods of mobility detection.

Spectrogram Feature Extraction

Figure 2.10 depicts the spectrograms of an ATSC signal as received from a stationary and mobile receiver. The ATSC signal was collected on Channel 43 (647 MHz) using a sample rate of 8 MHz. The spectrograms, $S(m,k)$, were computed from the data samples, $x(n)$, with $N = 2048$ point non-overlapping FFTs and a rectangular windowing function, $w(n)$, according to

$$S(m,k) \quad = \quad 20\log_{10}|\sum_{n=0}^{N-1} x(n+mN)w(n)e^{\frac{-j2\pi nk}{N}}|,$$

with no zero-padding, thus $k = 0, ..., N-1$. Mobility is indicated by the horizontal spectral gradients attributed to fading and Doppler shifts due to the movement. We therefore propose mobility detection using a decision statistic obtained from the two-dimensional convolution of the spectrogram with a modified Prewitt [29] horizontal edge-detection filter, P_M, where the modification incorporates a smoothing filter (for simplicity, we utilize a 20 by 20 moving average filter). Mobility is then indicated by edges across multiple frequencies. Thus, defining $V = S * P_M$, our decision statistic, $d(m)$, becomes:

$$d(m) = \sum_{k=0}^{W-1} V(m,k)/W,$$

where $W = N + b - 1$ due to b rows in the detection filter, P_M. Mobility is detected when the decision statistic exceeds a threshold, τ. We have used supervised learning techniques to determine τ, and Figure 2.10 illustrates our mobility detection results using a learned threshold of 0.025.

Power Statistics

Computation of the spectrogram is a costly procedure, and since wireless devices are typically compute and power-constrained, it is wise to consider more efficient techniques for mobility detection. One such method is to analyze the magnitudes of ambient signals. Two experiments were conducted using GNU Radio [7] and

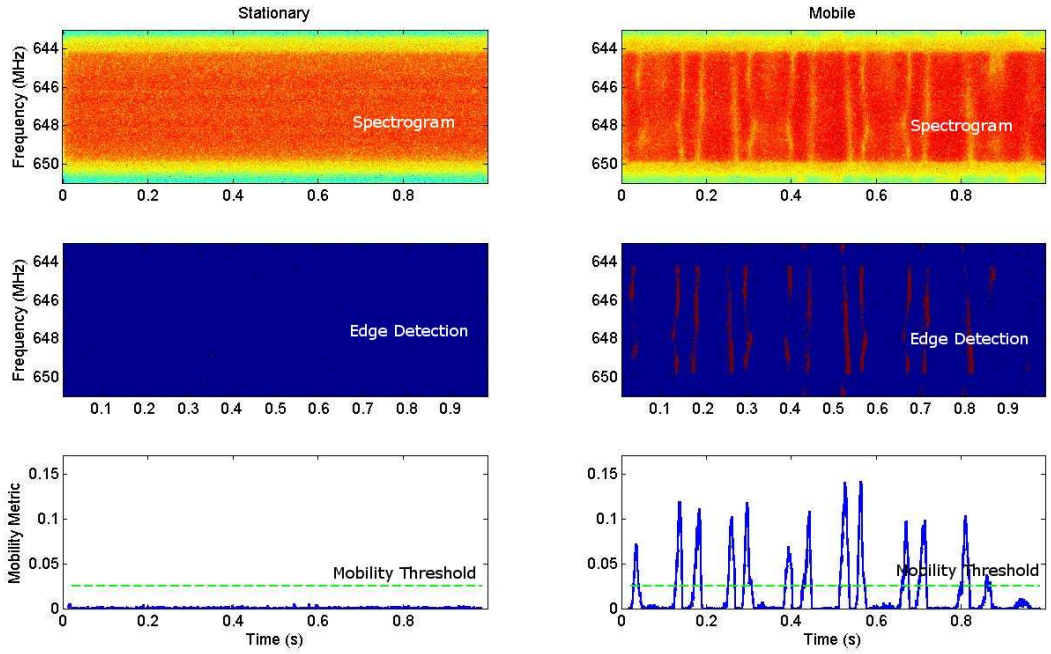

Figure 2.10: Mobility detection is demonstrated by a spectrogram edge detection algorithm. Stationary and mobile spectrograms are displayed. For the mobile spectrogram, an edge detection algorithm was implemented to show its applicability to mobility detection.

a Universal Software Radio Peripheral (USRP) [6] radio in a car monitoring FM, GSM, and ATSC signals. In both experiments, the car remained stationary for a few seconds before accelerating and maintaining a speed of 25 mph. At the end of each experiment, the car came to a complete stop. Both experiments lasted 30 seconds, and were conducted over the same linear stretch of road. The first experiment simultaneously monitored a GSM Downlink Control Channel in the 1900 MHz band and an ATSC signal from Channel 33 (587 MHz). The second experiment simultaneously monitored the same GSM channel and an FM signal at 99.1 MHz.

Mobility detection is possible by monitoring the differential unbiased sample standard deviations of received signal magnitudes. Using the time-series data,

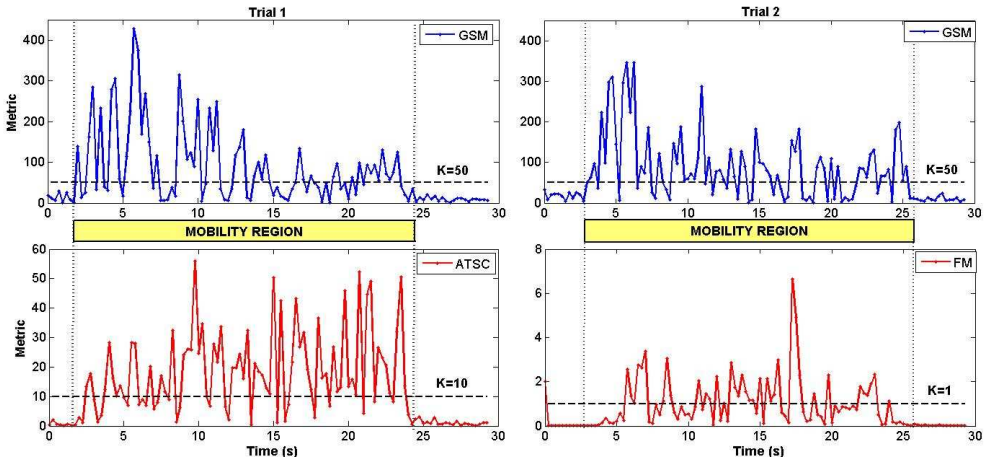

Figure 2.11: A car transitioned from a stationary position to a speed of 25 mph before coming to a complete stop. Mobility detection is demonstrated using GSM, FM, and ATSC broadcast signals using a power statistic processing algorithm .

$x(n)$, we calculate the mobility metric, $d(m)$, according to

$$\begin{aligned} u(m) &= \frac{1}{N}\sum_{n=0}^{N-1}|x(n)| \\ \sigma(m) &= \frac{1}{(N-1)}\sum_{n=0}^{N-1}||x(n)|-u(m)| \\ d(m) &= |\sigma(m)-\sigma(m-1)|. \end{aligned}$$

When the mobility metric rises above the threshold, K, mobility is inferred. As before, this threshold can be obtained from supervised learning techniques and is signal dependent due to modulation characteristics. For the mobile car experiments, we utilize $K = 80$ for the GSM signals, $K = 10$ for the ATSC signal, and $K = 1$ for the FM signal. Figure 2.11 plots the results for the two experiments. During the mobile portions of the experiments, our mobility metric for each signal correctly identifies the car as moving.

Doppler Shift Tracking

By tracking the Doppler shifts associated with the pilot tone of the ATSC signal [28], we can arrive at a simpler alternative to mobility detection that provides

a better quantification of the mobility (i.e. speed). The passive radar scenario that we are considering incorporates a transmitter that we know cannot move. Therefore the Doppler shift, f_d, can be calculated as $f_d = \frac{v}{\lambda} cos(\theta)$, where v is the velocity of our device, λ is the wavelength, and θ is the approach angle ($\theta = 0°$ corresponds to moving directly towards the transmitter). Without knowing device locations, we can bound the Doppler shift as

$$|f_d| \quad \leq \quad |\frac{v}{\lambda} cos(\theta)| \leq \frac{v}{\lambda} = \frac{vf}{c},$$

where c is the speed of light and f is the frequency of the broadcast signal. Note that higher transmission frequencies result in better Doppler shift resolutions. This is especially important if we want to detect mobility at reasonable speeds. For instance, ATSC Channel 33 transmits its pilot tone at 584.31 MHz resulting in a maximum Doppler shift of 22 Hz for a car moving at 25 mph. Figure 2.12 shows the ATSC pilot tone Doppler shifts incurred during the moving car experiment that we previously described. Mobility is again clearly evident. The velocities obtained from the measured Doppler shifts represent a minimum velocity because of the unknown approach angle. Thus, when monitoring multiple signals, the maximum velocity magnitude should be considered as the most accurate. In this experiment, we estimate a velocity of about -25 mph from the Doppler shift. Since this is the speed at which we were moving, we conclude that we were moving almost directly away from the broadcast tower.

Viewing the results from Figure 2.12, the astute reader may notice the slight Doppler drift associated with the stationary sections of the experiment (the beginning and the end). This drift is attributed to the local oscillator (LO) at the receiver, and therefore may cause some confusion and uncertainty with mobility estimation. We can mitigate this issue using a reference signal at a much lower frequency in order to factor out the effects of the LO. By using a very low frequency reference, any large shifts calculated from this reference signal can be immediately removed from the higher frequency estimation since it is most likely due to

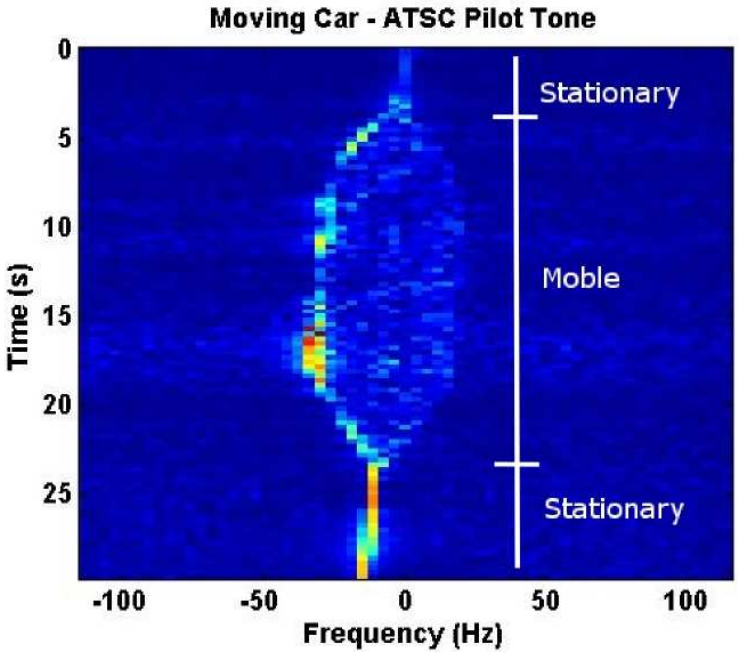

Figure 2.12: A car transitioned from a stationary position to a speed of 25 mph before coming to a complete stop. The spectrogram of an ATSC pilot tone region is displayed. The mobility region is indicated by Doppler deviations from nominal.

LO drift. An example of a prime reference signal would be the ATSC Channel 2 pilot (54.31 MHz), as the maximum Doppler shift seen by a car moving at 25 mph from this signal is 2 Hz.

Location

Coarse location information is useful in selecting communication schemes since indoor and outdoor environments have different channel characteristics (e.g. multipath levels and path loss exponents). Since cellular, FM, and ATSC towers are located outside, while WIFI access points are primarily located inside, a device may monitor measurements from these sources to deduce whether it is indoors or outdoors.

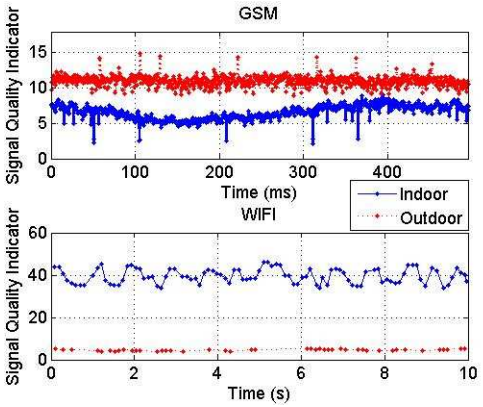

(a) SQI Measurements

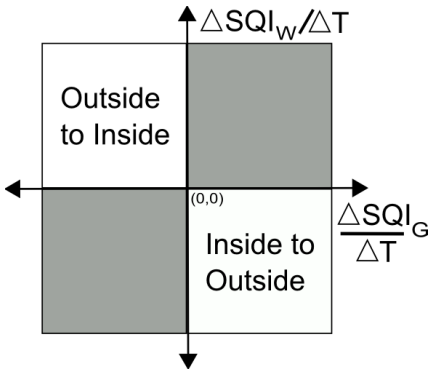

(b) SQI Decision Regions

Figure 2.13: (a) Signal Quality Indicators (SQI) are plotted for GSM and WIFI signals for indoor and outdoor locations. (b) SQI decision regions for coarse location discovery are depicted.

Signal Quality Comparison

In general, signals emanating from outside sources will degrade when the receiver moves inside. Contrarily, signals emanating from inside sources will improve when the receiver moves inside. By monitoring signal quality levels from inside and outside sources, a device can estimate when it has moved from an inside location to an outside location or vice-versa.

Experiments were conducted by defining and monitoring signal quality indicators (SQI) for GSM Downlink Control Channel bursts and WIFI beacon signals from a single access point. Because GSM Downlink Control bursts are transmitted with the same power, and GMSK modulation is constant-envelope, a perfectly received burst will have a constant magnitude [26] — thus variations in the magnitude can primarily be attributed to noise, multipath, and fading. To provide a fair comparator, we define the GSM SQI as the ratio of the sample standard deviation to sample mean of the signal magnitude for a given burst. Using N samples for the m^{th} GSM burst, the SQI, SQI_G, is calculated as $SQI_G(m) = \frac{u(m)}{\sigma(m)}$, where $u(m)$ and $\sigma(m)$ were previously defined.

Algorithm 2: Modified k-means clustering for location discovery

```
/*** Traditional k-means clustering ***/
cluster = kMeans(burstMags, k)
/*** Agglomerative Cluster Refinement ***/
while (clusterOverlap(cluster, tolerance) is TRUE) do
    cluster = RefineCluster(cluster, tolerance)
end
/*** Location Discovery ***/
if (NumClusters ≥ ThreshNum) AND (σ ≤ ThreshVar) then
    Device is outside
else
    Device is inside
end
```

For the WIFI signal, we choose to monitor the WIFI beacon because it is transmitted with constant power, and it is easily identifiable by its periodicity and signal structure [23]. However, the WIFI beacon is not constant-envelope, so we modify our SQI calculation to be the ratio of the average burst power to the noise floor. Thus we define the WIFI SQI for the m^{th} beacon as $SQI_W(m) = \frac{\mu(m)}{\eta(m)}$, where $\mu(m)$ is the sample mean of the beacon magnitude, and $\eta(m)$ is the noise floor magnitude associated with the m^{th} WIFI beacon. The noise floor estimate can be obtained from the mean magnitude of the data immediately before and after the leading and trailing edges of the beacon. Figure 2.13 (a) shows the results of the experiment, where the Signal Quality Indicators are plotted for the two signals versus time. The GSM SQI clearly increases when the receiver is transitioned to an outdoor location, while the WIFI SQI experiences a severe drop in quality level. By simultaneously monitoring these SQIs, a device can estimate when it has moved from inside to outside or vice-versa. The SQI Decision Region is illustrated in Figure 2.13 (b), where time-averaged differential SQI measurements are used to detect inside/outside transitions — false alarms due to fading are mitigated by the time-averaging process.

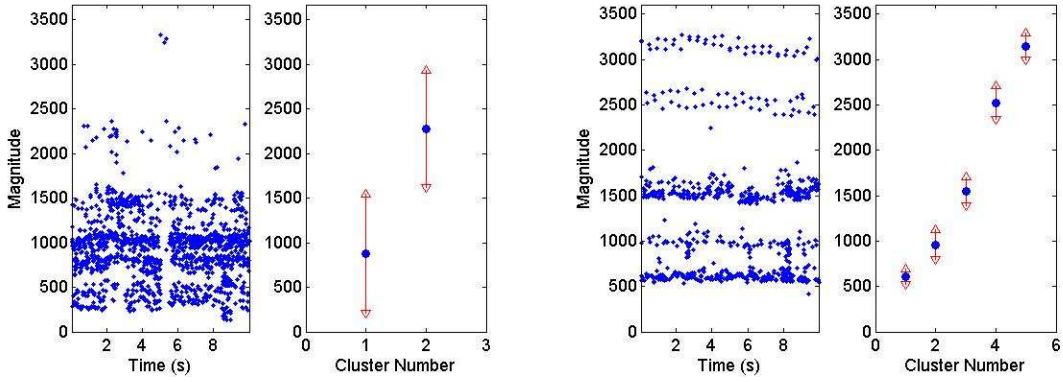

(a) Indoor (b) Outdoor

Figure 2.14: Clustering analysis of ISM band activity centered at WIFI Channel 6 is performed. Detected burst magnitudes for indoor and outdoor receiver locations are plotted versus time in addition to modified k-means clustering analysis. An outdoor receiver location is characterized by more clusters with higher densities and smaller variances.

Clustering Analysis

Alternatively, we may perform clustering analysis on *all* signals in the ISM band. To determine whether a device is inside or outside, we leverage the notion that a receiver located outside will have more relatively stationary paths from emitters. Whereas, when the receiver is inside, there will be more multipath signal arrivals due to an increased number of reflectors within close proximity — thus the arriving signal magnitudes will *smear* into each other. Therefore, we wish to analyze the magnitude clusters and their densities.

To estimate the number of clusters, we first considered the QT clustering algorithm [30], however its uniform cluster diameter was undesirable. More fitting for this application would be an agglomerative hierarchical clustering algorithm. We thus introduce a modified k-means algorithm [31], where we first perform traditional k-means clustering using a large k (e.g. 10). We then reduce the number of clusters by combining them in a manner related to overlapping variances. If two cluster means are within α standard deviations of each other, we combine the

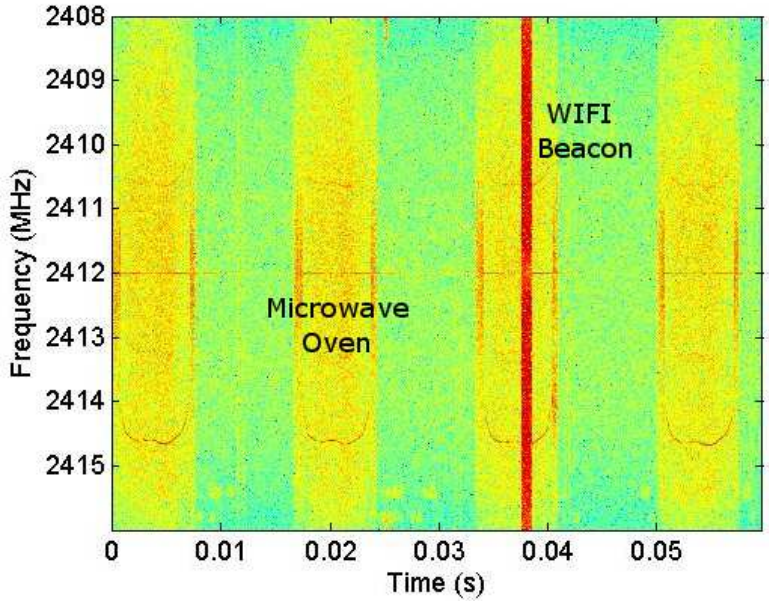

Figure 2.15: A microwave oven is interfering with WIFI Channel 1 activity.

clusters. This process is repeated until no more combining occurs. The number of final clusters and their densities can be used to estimate whether or not a device is inside or outside. Our modified k-means algorithm for location discovery is listed in Algorithm 2. Results from the implementation of this algorithm on burst magnitudes in the ISM band are depicted in Figure 2.14, where indoor and outdoor location is correctly identified based purely on a cluster number threshold of 3 using a cluster separation factor of $\alpha = 2$.

Unconventional Indicator Detection The detection non-communication based indicators can also aid in indoor/outdoor estimation. One example is the detection of microwave oven interference. Note that this is also very useful information to have when deciding what communication scheme to use, as microwave oven interference may be best mitigated by certain protocol selections (e.g FHSS). Figure 2.15 shows a spectrogram from the ISM band where a microwave oven is interfering with WiFi activity.

PLATEAU Summary

Accurate situational awareness is important for cognitive devices to optimize operational decisions. Through PLATEAU, we have presented and demonstrated techniques that detect device mobility and coarse location by leveraging physical layer properties from ambient radio signals.

2.5 Situational Awareness Summary

While all of the situational awareness techniques discussed in this chapter have been passive, we note that active techniques can also be employed — especially to discover services and devices in the area. For instance, an active CR spectral sensing routine may invoke the Bluetooth Service Discovery Protocol, whereby a CR may send a discovery request to Bluetooth devices in its area in order to obtain a list of services available. A similar discovery packet also exists in the 802.11 protocol in the form of Probe Request Frames. Note that this is of course dependent on the functional limitations of the CR itself.

Although only a limited number of protocols have been examined, the methodology proposed can be extended to other protocols. Furthermore, spectral sensing techniques can be developed to identify and avoid non-communication based interferers such as microwave ovens.

Chapter 3
Attack

3.1 Motivation

In most any competitive arena, a good defense is founded upon its understanding of possible attacks. This holds true for the wireless security world as well. By thoroughly understanding the vulnerabilities and risks associated with specific protocols, a wireless network can better defend itself against adversarial attacks.

Attacking a wireless network is an extremely broad area, with many degrees of freedom including: channel conditions, power constraints, hardware constraints, covertness constraints, etc. In this Chapter, we do not attempt to propose any general techniques that will work in *any* situation — rather we illustrate that attack strategies become increasingly effective and efficient when considering the specific protocol of interest. Further, access to the PHY layer provides new capabilities and advanced functionality to both launch and defend against attacks.

We begin this Chapter by presenting a survey of existing protocol specific attacks. These exploits target popular wireless protocols such as WiFi, Bluetooth, and GSM, in addition to lesser known yet equally pervasive sensor networks such as tire pressure monitoring systems (TPMS). Common to each exploit is its reliance upon physical layer access with a current SDR platform.

The second portion of this Chapter focuses on multi-input multi-output networks (MIMO), because it is incorporated in some fashion in many emerging wireless technologies, ranging from 802.11n to WiMAX to 4G cellular systems.

We present a thorough physical layer vulnerability study on two of the most popular MIMO techniques — the capacity achieving singular value decomposition (SVD) based scheme, and the practical Alamouti space-time block code (STBC) scheme. Analysis includes both theory and simulations, culminating in results from real-world experimentation with the GNU Radio/USRP SDR platform.

3.2 Exploit Survey

Numerous protocol exploits have been made possible by recent advancements in SDR technology. The low-cost, fully programmable GNU Radio/USRP SDR platform played a key role in each of the following protocol-specific exploits.

3.2.1 WiFi

The ADROIT project [32] was developed by BBN Technologies and funded by DARPA's ACERT program in an effort to create collaborative teams of CRs. As a result, the group was able to implement an 802.11 receiver using GNU Radio and the USRP. But due to the spectral limitations of this SDR, functionality was limited to receiving only 1 Mbps DBPSK signals. In [33], researchers moved despreading functionality into the FPGA in order to support full-bandwidth 802.11b signal reception. Using this code, full eavesdropping capability at the PHY layer can be achieved by simply implementing the higher layer protocol parsers (and of course traditional security cracks).

3.2.2 Zigbee

Another popular protocol with SDR support is IEEE 802.15.4 (Zigbee). In [34], researchers provide full transceiver functionality for the Zigbee protocol, thus opening the door for both passive and active attacks — from eavesdropping to rogue packet injection and device impersonation.

3.2.3 Bluetooth

GR-Bluetooth

In [19], a novel attack against the Bluetooth protocol is discussed using GNU Radio and the USRP. Leveraging physical layer access and intentional aliasing, the GR-Bluetooth code facilitates full eavesdropping capabilities against unencrypted Bluetooth devices. Minor modifications to the RFX-2400 daughterboard and FPGA firmware are needed in order to allow the aliasing to occur. Further enhancements to the work include a Wireshark plugin to decode higher layer Bluetooth messages, in addition to keystroke extraction code to reveal keystrokes from Bluetooth keyboards.

Car Whisperer

The researchers in [35] present software capable of associating with unsuspecting Bluetooth devices by exploiting passkey flaws — the majority of Bluetooth devices employ standard passkeys of 0000 or 1234. The software allows for full voice eavesdropping and is also capable of voice injection. Note that eavesdropping is possible even if the target device is not active. While the researchers used Bluetooth dongles rather than an SDR to conduct these attacks, we mention this work because the extension of the GR-Bluetooth work outlined above to include transmitter capabilities opens the door to utilize the Car Whisperer techniques.

3.2.4 GSM

GSSM

In [36], the Groupe Special (Software) Mobile (GSSM) project opened the door for GSM exploitation. The code used the GNU Radio/USRP SDR to provide demodulation capabilities for the GSM downlink control channel. A real time

tunneling interface to Wireshark was also implemented.

Airprobe

The exploitation of GSM continued in [37], which uses the GNU Radio/USRP SDR to passively decrypt and decode GSM calls and text messages. The Airprobe software utilizes cryptographic rainbow tables to crack the A5/1 cryptography used in most GSM networks. As expected, utilization of this technique comes with its proper share of legal concerns.

OpenBTS

The Open Base Transceiver Station (OpenBTS) [38] project provides a GNU Radio/USRP based GSM access point, thus providing GSM transceiver capabilities. The software allows standard GSM cellular phones to complete telephone calls without any existing telecommunication provider networks. While the creators emphasize the care that should be taken to avoid any legal trouble with operating such a system (e.g. not connecting antennas and ensuring the use of a test location area code), most adversaries will not be inclined to follow such measures. As such, rogue Base Stations can be readily implemented.

3.2.5 MBTA Charlie Card

Researchers in [39] drew much media and legal attention for their exploits related to the Massachusetts Bay Transportation Authority's (MBTA) Charlie Cards. These cards are MiFARE-based, contactless, stored value smart cards that are used to pay subway fares. By reverse engineering the protocols, the researchers demonstrate various available exploits such as altering the monetary value stored on the cards. The team incorporates the GNU Radio/USRP SDR to attack the RFID vulnerabilities of the system.

3.2.6 TPMS

In [40], we provide a thorough study on the physical layer vulnerabilities associated with modern tire-pressure monitoring systems (TPMS). [1] The purpose of TPMS is to monitor tire status (e.g. pressure and temperature) in order to provide timely alerts to drivers in an effort increase overall road safety. As a result of the Ford-Firestone tire failure controversy [41], TPMSs represent the first federally mandated in-vehicle wireless sensor network [42]. In our study, a variety of the most common TPMS sensors were backward engineered in order to discover the protocols used to send messages. We found that the TPMS sensors were typically awoken with continuous wave (CW) signals at 125 KHz and then responded with data bearing messages at 315/433 MHz using ASK or FSK signals. These signals were unencrypted, and contained *unique* identifiers that can therefore be used to identify vehicles and thus individuals. Specific vulnerabilities of such a system include: individual tracking/anonymity violation, hostile targeting based upon sensor ID, and spoofing the TPMS sensor to either attack the car's computer or to trick an individual to pull over. Real-world experimentation and attack demonstrations were conducted using the USRP/GNU Radio SDR to illustrate practicality and feasibility of the exploits against test equipment and researcher-owned sensors and vehicles.

3.2.7 Exploit Summary

The complete PHY layer access associated with SDRs provides a powerful point of attack for current and emerging systems. Most emerging wireless communication systems, such as 802.11n, WiMAX, and LTE, all incorporate some form of multi-input multi-output (MIMO) communication scheme in an effort to enhance

[1] This research study was conducted in collaboration with researchers and developers from Rutgers University and the University of South Carolina.

throughput and quality of service. Such systems depend upon reliable PHY layer channel estimates to operate effectively. In the next section of this thesis, we analyze the vulnerabilities of MIMO systems when subjected to PHY layer attacks from cognitive platforms.

3.3 MIMO

A major benefit of multi-input multi-output (MIMO) wireless communication systems [43, 44] is the ability to perform well in scenarios traditionally viewed as poor — such as richly scattering environments. Consequently, many emerging wireless technologies, ranging from 802.11n to WiMAX to 4G cellular incorporate some form of MIMO. Anticipating that systems might be deployed in adversarial situations, research has examined jamming MIMO wireless systems [45, 46, 47, 48, 49]. But rather than considering *smart* adversaries, investigations have focused on broadband jammers or incidental jamming as a result of co-channel interference. While these results are indeed important, they fall far short of capturing the shortcomings associated with an intelligent, capable adversary.

In this Section, we study how an intelligent adversary can disrupt MIMO communication by targeting the generally over-looked, but essential, channel estimation procedure. Since there are many MIMO schemes, we begin our analysis by first exploring attacks on (optimal) capacity-achieving MIMO systems employing the singular value decomposition (SVD) to create virtual, parallel channels. Since such systems require channel state information (CSI), we analyze the vulnerabilities associated with jamming the CSI estimation procedure. SVD-based waterfilling MIMO represents a theoretically ideal form of MIMO. On the other hand, practical MIMO implementations often incorporate space-time block coding (STBC), and a common form of STBC that appears across many standards (e.g. 802.11n and 802.16) is the Alamouti space-time block code. Consequently,

we also focus on the vulnerabilities associated with the baseline Alamouti 2-by-1 (2 transmit antennas and 1 receive antenna) STBC. Our results can be easily extended to other STBCs, and thus illustrate an underlying fragility in emerging wireless systems.

3.3.1 Related Work

An excellent overview of MIMO fundamentals, a survey of key research, and applicability to real-world implementations are provided in [43]. Findings that were key to initial MIMO developments are found in [44, 50, 51], and capacity for Gaussian and spatially-correlated Rayleigh fading channels are presented in [52] and [53]. Initial MIMO demonstrations can be found in [54], with improvements and extensions in [55, 56, 57].

Research has also analyzed performance degradation due to noise or other physical phenomenon (e.g. a time-varying channel). In [58], CSI errors due to noise in the channel are investigated for SVD-based MIMO systems. The authors use a Wiener filter to track the CSI to avoid stationarity issues in the SVD calculation. The authors in [59] look at SVD-based MIMO system performance for unencoded BPSK modulation in the presence of white noise. In [60], it is shown that the capacity of SVD-based MIMO systems degrade greatly with incorrect CSI as a result of a time-varying channel. They propose a remedy involving feedback to achieve near capacity performance. In [61], the authors present a model for time-varying channels without feedback when the receiver has an estimate or perfectly knows the CSI.

The work on jamming MIMO is predominantly concerned with adversaries that do not take advantage of system protocols, i.e. the jammers are unintelligent, with noise-like impact on the target receiver. In [45], the authors investigate uncorrelated jamming in non-coherent wideband fading channels (since the problem is understood well for the coherent regime). They show that naive energy-limited

jammers do not affect capacity in the wideband regime. In [46], the authors examine a network MIMO scenario where sensors collaborate to decode a message broadcast by a multi-antenna transmitter in the presence of a white-noise jammer. [48] looks at finding an optimal training sequence for a MIMO system under flat fading conditions with spatially colored interference due to thermal noise and multiple non-intelligent, independent jammers. The solution (which, frankly, may be attacked), requires information feedback. In [49], the researchers assume that the jammer produces a spatially correlated Gaussian interference signal and that the jammer has full CSI but does not have any knowledge of the users' signals. All of these papers do not involve an intelligent adversary.

On the other hand, there has been non-MIMO work involving a smart adversary [62, 47, 63, 64, 65]. In [62], the authors assume that a non-MIMO system transmits independent, identically distributed (i.i.d.) Gaussian random variables with an input power constraint. In the paper, the jammer taps the channel and perfectly feeds back a signal. Most closely related to our work is [47]. Here, the authors assume that the users and jammer have independently fading MIMO channels. They investigate scenarios involving different levels of CSI accuracy. For each situation, optimal transmit strategies are discussed. When no CSI is known, they show that the optimal jamming strategy is for the jammer to use equal power allocation. For perfect CSI, the jammer should use a matched-waterfilling strategy. Finally, for partial CSI, the jammer should beamform in the direction of the transmit correlation eigenvectors and perform proportional power allocation.

In our work, we consider a more intelligent (and capable) adversary, which is quite plausible given recent advancements in software defined radios (SDRs) [5, 10, 23, 66]. A capable SDR with protocol knowledge can achieve system synchronization therefore enabling increasingly smarter attacks. In this Section, we analyze *smart* MIMO attacks, where the CSI is targeted by jamming during the channel estimation procedure — for the majority of MIMO schemes, it is essential

to obtain accurate channel estimates in order to realize the gains associated with MIMO. By attacking only the CSI, the jammer remains fairly covert and power conservative as she only needs to operate during channel sounding (typically a small fraction of user transmission time). In particular, our strategy for a jammer is as effective and more efficient than a jammer that blasts throughout the data transmission period.

3.3.2 MIMO Overview

Consider a flat-fading MIMO channel, where Alice (who wishes to send a message to Bob) has n_t transmit antennas, while Bob has n_r receive antennas. The channel matrix between Alice and Bob is a complex $n_r \times n_t$ matrix, $\mathbf{H}$, describing the propagation effects between each of Alice's antennas and Bob's antennas. If Alice transmits the n_t-dimensional signal $\mathbf{x}$ to Bob, we can represent the n_r-dimensional received signal as $\mathbf{y} = \mathbf{Hx} + \mathbf{n}$, where $\mathbf{n}$ is n_r-dimensional additive noise. MIMO techniques can be classified into three main categories: *precoding*, *spatial multiplexing*, and *diversity coding*. A wide array of MIMO implementations exist, each of which exploits these principles to enhance communication rate and reliability.

3.3.3 SVD-based MIMO

One of the popular classes of MIMO systems, which can achieve optimal communication rates, involves precoding. Rather than transmitting $\mathbf{x}$, Alice precodes the signal using the SVD of the channel matrix, $\mathbf{H}$. Recall that the SVD decomposes a matrix as $\mathbf{H} = \mathbf{U}\Sigma\mathbf{V}^H$, where $\mathbf{U}$ represents the left singular vectors of $\mathbf{H}$, $\Sigma = diag\{\sigma_1, ..., \sigma_n\}$ provides the singular values of the channel matrix along its diagonal, $\mathbf{V}$ corresponds to the right singular vectors, and H indicates the conjugate transpose.

If Alice performs the SVD of $\mathbf{H}$ and pre-codes $\mathbf{x}$ with $\mathbf{V}$ by transmitting $\mathbf{Vx}$, then Bob receives $\mathbf{r} = \mathbf{HVx} + \mathbf{n}$. Bob can then decode the signal from Alice by operating on $\mathbf{r}$ with $\mathbf{U}^H$. Applying the unitarity of $\mathbf{U}$ and $\mathbf{H}$, we have

$$\begin{aligned} \mathbf{d} &= \mathbf{U}^H\mathbf{r} = \mathbf{U}^H\mathbf{HVx} + \mathbf{U}^H\mathbf{n} \\ &= \mathbf{U}^H\mathbf{U}\Sigma\mathbf{V}^H\mathbf{Vx} + \mathbf{U}^H\mathbf{n} = \Sigma\mathbf{x} + \mathbf{U}^H\mathbf{n}. \end{aligned}$$

In this manner, Alice and Bob can communicate over $\min(n_r, n_t)$ parallel single-input single-output (SISO) channels.

MIMO Capacity

Alice and Bob can achieve capacity over the MIMO channel, $\mathbf{H}$, by employing a waterfilling solution over the associated parallel SISO channels as computed above [67]. The mutual information between Alice and Bob is

$$I(\mathbf{H}, \mathbf{Q}) = \log_2(\det([\mathbf{I}_{n_r} + \rho\mathbf{HQH}^H]))$$

where $\rho = E_s/\gamma_n^2$ is the SNR and $\mathbf{Q}$ is the input covariance matrix. The optimal $\mathbf{Q}$, denoted $\mathbf{Q}^\star$, achieves capacity (by maximizing the mutual information) by allocating power optimally into the right singular vectors of $\mathbf{H}$. This optimal power allocation is $\{p_1^\star, ..., p_n^\star\}$ where $\mathbf{Q}^\star = \mathbf{V}diag\{p_1^\star, ..., p_n^\star\}\mathbf{V}^H$. Capacity for this channel is

$$C(\mathbf{H}) = \sum_{k=1}^{n} \log_2[1 + \rho p_k^\star \lambda_k],$$

where $\lambda_k = \sigma_k^2$, and the optimal power allocation $\{p_1^\star, ..., p_n^\star\}$ is obtained by water-filling using $p_k^\star = (\mu - \frac{1}{\rho\lambda_k})^+$, where $(z)^+ = \max(z, 0)$. Here μ is chosen according to the transmitter's power constraint (e.g. $\sum_{k=1}^{n} p_k^\star = 1$). Note that in the low SNR regime, waterfilling reduces to allocating all available power to the strongest eigenmode, while in the high SNR regime, waterfilling reduces to uniformly distributing the power over all non-zero eigenmodes.

3.3.4 Jamming SVD-based MIMO

Now suppose that Eve enters the picture to jam the communication between Alice and Bob, and (by Kirkhoff's Principle) that she knows the channel sounding waveforms. Since Alice and Bob are using a MIMO scheme, which requires CSI, Eve can jam (1) the data *and* the CSI, (2) only the data, or (3) only the CSI. It is the third case, where Eve attacks only the CSI estimation procedure, that we are interested in as it provides the most efficient and covert avenue of attack (i.e. a short burst against CSI can be more devastating than a short burst on the data).

To understand how Eve may attack the CSI, we must first understand how and where she can affect $\mathbf{H}$. Denoting $\hat{\mathbf{H}}_A$ as Alice's estimate of the channel, and $\hat{\mathbf{H}}_B$ as Bob's estimate, CSI estimations are obtained as follows:

- Bob transmits a channel sounding waveform to Alice, who estimates $\hat{\mathbf{H}}_A$ and uses this to precode her communication (e.g. standard beamforming).
- Alice transmits a channel sounding waveform so that Bob can estimate $\hat{\mathbf{H}}_B$ to enhance decoding. For the Alamouti STBC (used in 802.11n and WiMAX [68]), only the receiver needs an estimate of the channel.
- Alice transmits a channel sounding waveform so that Bob can estimate $\hat{\mathbf{H}}_B$. Bob then sends $\hat{\mathbf{H}}_B$ back to Alice in a message. Such feedback is common in cellular communications, where feedback on the downlink is used to refine coding on the uplink, and vice-versa.
- Alice transmits a channel sounding waveform so that Bob can estimate $\hat{\mathbf{H}}_B$. Bob then transmits his own sounding waveform so that Alice can obtain $\hat{\mathbf{H}}_A$. This procedure is allowed for in WiMAX [69].

Given this categorization, Eve will either interfere with channel sounding, or corrupt the feedback channel (we note that although authentication *should* be

explicit in CSI feedback, unfortunately this is generally not the case). Thus either Alice ($\hat{\mathbf{H}}_A$), Bob ($\hat{\mathbf{H}}_B$), or both may have an erroneous channel matrix. For SVD-based MIMO systems, both $\hat{\mathbf{H}}_A$ and $\hat{\mathbf{H}}_B$ are needed, so the jammed waveform is

$$\begin{aligned} \mathbf{s}_J &= \hat{\mathbf{U}}_B^H \mathbf{H} \hat{\mathbf{V}}_A \mathbf{x} + \hat{\mathbf{U}}_B^H \mathbf{n} \\ &= \hat{\mathbf{U}}_B^H \mathbf{U} \Sigma \mathbf{V}^H \hat{\mathbf{V}}_A \mathbf{x} + \hat{\mathbf{U}}_B^H \mathbf{n}, \end{aligned}$$

where $\hat{\mathbf{U}}_B$ are the left singular vectors of $\hat{\mathbf{H}}_B$, and $\hat{\mathbf{V}}_A$ are the right singular vectors of $\hat{\mathbf{H}}_A$, and $\mathbf{H}$ is the true channel.

Eve knows the channel

Let us assume for the moment that Eve knows $\mathbf{H}$ exactly, *and* can give both Alice and Bob an arbitrarily chosen matrix. Since Alice and Bob are seeking to waterfill over the eigenmodes of the channel, an excellent attack would be for Eve to convince them to perform opposite-waterfilling! Eve can accomplish this by computing the SVD of the channel, reversing the ordered singular values (i.e. for n singular values, $\hat{\sigma}_i = \sigma_{n-i}$, $\hat{\mathbf{H}} = \mathbf{U}\hat{\Sigma}\mathbf{V}^H$), and administering a corrupted $\hat{\mathbf{H}}$ to both Alice and Bob. Figure 3.1 illustrates this process. It is important to note that this result is in clear contrast to [47], where the recommendation is for Eve to add noise to the existing eigenmodes in a manner proportional to eigenvalue significance. The difference in attack methodology arises because we attack *only* the channel estimation procedure, while [47] assumes that jamming occurs during data transmission.

In the high SNR regime, however, this attack strategy will not have much effect since each eigenmode represents an excellent channel. In such a scenario, it would be better for Eve to give Alice and Bob a corrupted $\hat{\mathbf{H}}$ so that their singular vectors are orthogonal to the correct singular vectors (i.e. make $\hat{\mathbf{U}}_B^H \mathbf{U} = \mathbf{0}$, $\mathbf{V}^H \hat{\mathbf{V}}_A = \mathbf{0}$, or both)!

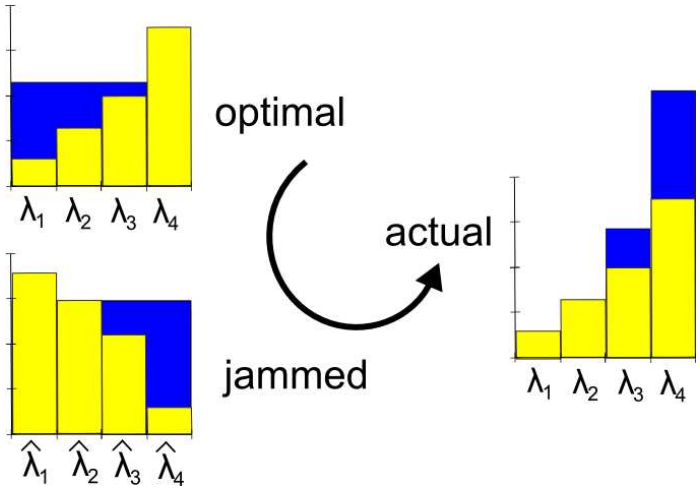

Figure 3.1: The opposite waterfilling attack is illustrated. *Optimal* waterfilling is not employed because of the *jammed* channel estimates resulting in *actual* distribution of power in an opposite-waterfilling manner.

If, on the other hand, Eve cannot arbitrarily give a corrupted $\hat{\mathbf{H}}$ to Alice/Bob, then she must attack the channel sounding process itself. In this case, there are added constraints to note. Often, the estimation procedure involves transmitting a channel sounding waveform from a single antenna and using it to estimate the channel to each receive antenna (and then cycling through the transmit antennas). For each of Alice's transmissions, n_r channel elements are calculated by Bob. Thus $\hat{\mathbf{H}}_B$ will be populated column-wise (h_{1i},h_{2i},...,$h_{n_r i}$). If messaging is not used to return the estimate to Alice, then Alice's channel estimate, $\hat{\mathbf{H}}_A$, will fill in row-wise (h_{i1},h_{i2},...,h_{in_t}). If messaging is used, we note that quantization levels must be considered (in WiMAX [69], 6-bit quantization is used to represent CSI). While these observations represent added constraints to Eve's operating procedure, it does not change the operating goals presented above, and (as we will see in Section 4.2) there is a general procedure that is effective for Eve to employ to corrupt the measured $\hat{\mathbf{H}}$.

In reality, Eve might only have an estimate of the channel, and she would be unlikely to *arbitrarily* perturb the channel estimation process. If Eve jams the

channel sounding process, then the estimation of $\mathbf{H}$ will be corrupted by an error matrix $\mathbf{E}$, and thus it would be valuable for Eve to quantify the impact she could have on the singular subspaces. Quantifying this impact is very important as the SVD is a non-continuous function, e.g. consider the following perturbation of $\mathbf{H}$ by an error matrix $\mathbf{E}$.

$$\mathbf{H} = \begin{pmatrix} 1 & 0 \\ 0 & 1+\epsilon \end{pmatrix} \qquad \mathbf{E} = \begin{pmatrix} 0 & \epsilon \\ \epsilon & -\epsilon \end{pmatrix}$$

$$\hat{\mathbf{H}} = \mathbf{H} + \mathbf{E} = \begin{pmatrix} 1 & \epsilon \\ \epsilon & 1 \end{pmatrix}$$

$$\mathbf{V} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad \hat{\mathbf{V}} = \tfrac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

It is clear that a 45° relative shift in the singular vector spaces occurs for arbitrarily small values of ϵ. So, should Alice and Bob be wary of small attacks drastically altering the singular vectors? The answer is generally "no" because such extreme perturbations are related to matrices with close singular values [70], and most real-world MIMO channels do not have close singular values. Consider a MIMO channel under Rayleigh fading with uncorrelated antenna elements. Here, the channel matrix contains zero-mean, complex, normal i.i.d. elements. The singular values of such a matrix are the square roots of the eigenvalues ($\sigma_i = \sqrt{\lambda_i}$) of the matrix, $\mathbf{W} = \mathbf{H}\mathbf{H}^H$ (for $n_t > n_r$, otherwise we use $\mathbf{W} = \mathbf{H}^H\mathbf{H}$). Here, $\mathbf{W}$, is called a central Wishart matrix, and its eigenvalues are characterized in [71]. Defining $n = min(n_t, n_r)$, $m = max(n_t, n_r)$, and noting the ordered eigenvalues of $\mathbf{W}$ as $(\lambda_1 \geq \lambda_2 ... \geq \lambda_n)$, the resulting joint p.d.f of eigenvalues is [53]

$$p(\lambda_1, \lambda_2, ..., \lambda_n) = K \prod_{i=1}^{n} e^{-\lambda_i} \lambda_i^{(m-n)} \cdot \prod_{i<j} (\lambda_i - \lambda_j)^2,$$

where K is a normalization constant

$$K = \frac{\pi^{(n(n-1))}}{\Gamma_n(m)\Gamma_n(n)},$$

and Γ_n is defined as $\Gamma_n(a) = \pi^{(n(n-1)/2)} \prod_{i=1}^{n} (a-i)!$.

Figure 3.2 depicts distance distributions related to the eigenvalues of $\mathbf{W}$ for a 3 by 4 Rayleigh fading MIMO channel. Since the channel matrix is 3 by 4 there are only 3 eigenvalues to consider. For illustrative purposes, the fading parameters were chosen so that the eigenvalues were generally less than 20. Plotted is the probability versus mean and minimum eigenvalue distances. In this case, the expected minimum eigenvalue distance is 1.8, and the expected average eigenvalue distance is 3.1. Hence, close singular values are not probable. Generally, the probability of encountering a channel of n eigenvalues where the minimum eigenvalue distance exceeds δ is given by:

$$P(\triangle\lambda_{min} \geq \delta) = \sum_{i=1}^{n-1} P((\lambda_i - \lambda_{i+1}) \geq \delta).$$

For the 3 by 4 channel considered, the probability that the minimum eigenvalue distance is greater than 1 is 86%. Moreover, 97% of the time the minimum eigenvalue distance encountered is greater than 0.5.

Since $\mathbf{H}$ is generally well-behaved under Rayleigh fading conditions, Alice and Bob need not worry about drastic changes in the singular vectors when subjected to general low powered attacks or even noise. But this does not mean that the singular values are safe. In fact, Eve can use perturbation theory to gauge her singular value attacks. In general, small random perturbations of the channel result in small perturbations of the singular values. Also, small singular values tend to increase under random perturbation on the order of square root of the number of transmitters [70]. Inherently, by injecting noise Eve makes bad channels appear better — which is Eve's goal. Eve may bound her effect on the singular values by [72, 70, 73]:

$$|\sigma_k(\mathbf{H}+\mathbf{E}) - \sigma_k(\mathbf{H})| \leq \sigma_1(\mathbf{E}) = \|\mathbf{E}\|_2,$$

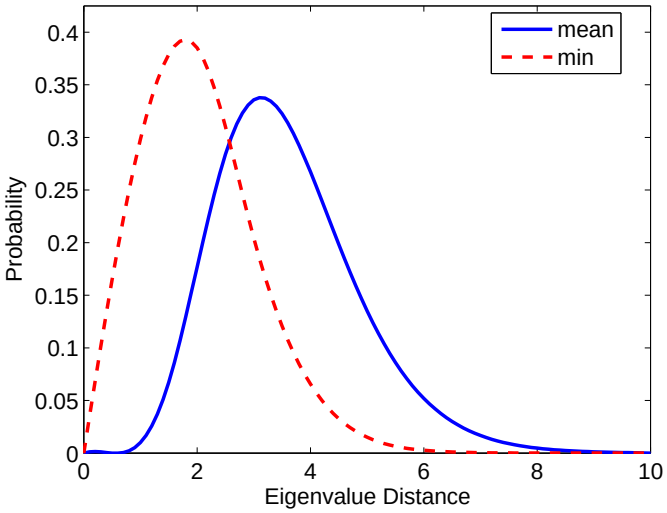

Figure 3.2: An eigenvalue distribution analysis of the Wishart matrix for a 3 by 4 Rayleigh fading MIMO channel is shown. Plotted is the probability distribution versus mean eigenvalue distance as well as minimum eigenvalue distance. Due to channel constraints, eigenvalues are less than 20. Close singular values are clearly not probable.

or more generally, by [73, 74, 70]:

$$\sum_{k=1}^{n}(\sigma_k(\mathbf{H}+\mathbf{E})-\sigma_k(\mathbf{H}))^2 \leq \|\mathbf{E}\|_F^2.$$

By utilizing these bounding equations, Eve may determine how effective she can be as a jammer when attacking the singular values of the channel. As noted earlier, in high SNR scenarios, Eve should *not* direct her attack against the singular values, but instead should attack the singular vectors. This is challenging to accomplish if Eve merely injects noise, however, as we will see shortly, is quite possible by exploiting the steps in the channel sounding process!

Eve doesn't know the channel

We now consider the most general case, where Eve uses no knowledge about $\mathbf{H}$. In this scenario, by using knowledge of the sounding process, Eve can still be very effective in subverting estimation of $\mathbf{H}$.

Eve can act as a jammer by injecting *random* eigenmodes. First consider the case where estimation messaging (feedback) is used. A special case of this attack would be Eve jamming during the channel sounding from Alice's first antenna element. By using enough power, Eve can cause the first physical channel to appear to be the best, and thus force Alice and Bob to waterfill primarily over this *physical* mode (since it appears optimal). This will have the effect, *on average*, of emptying power uniformly into all of the *actual* eigenmodes of the channel. In other words, water-filling is thwarted and we merely have a uniform power allocation scheme! Note that Eve may choose to attack any physical antenna (or random realization) and still achieve the same result. Additionally, attacking *all* of the channel sounding will make the channel appear singular, and deliver similar results. Utilizing this type of attack in reasonable SNR regimes will on average reduce the capacity from C to $C/min(n_t, n_r)$.

In the high SNR regime, the effect of this attack will be minimal, and Eve *could* revert to traditional jamming by jamming the data sequences. However, we must note that the feedback provides yet another point of attack for Eve (as authentication of CSI feedback is generally not employed). Rather than jamming only the Alice→Bob sounding procedure, Eve should also inject an appropriate (random and false) encoding of $\hat{\mathbf{H}}_B$. By forcing a dominant mode to appear in the Alice→Bob estimation, she forces emphasis on a *specific* vector of $\hat{\mathbf{U}}_B$. On the other hand, Eve can feedback a false, singular version of $\hat{\mathbf{H}}_B$ (e.g using the feedback scheme described on pg 449 of [69]), where she has emphasized one particular random singular vector. The net result is that Alice will waterfill into this false dominant mode, while Bob will attempt to decode on a false dominant mode, resulting in significant energy being lost in the transfer from the Alice→Bob dominant mode to the Bob→Alice dominant mode.

When feedback is not used, Eve can employ a comparable strategy by interfering separately in the Alice→Bob and Bob→Alice estimation processes. Here,

Eve jams enough energy into a specific transmitter to force Bob to believe in a particular false dominant mode. She then jams energy into the Bob→Alice estimation to force a false dominant Bob→Alice mode. Again, Alice will then waterfill into a false dominant mode, and Bob will decode on a false dominant mode not related to the true channel.

In both cases, for high SNR, Eve does not have to jam the data transmissions, and the strategies presented apply well to cases where Eve does know the channel.

3.3.5 Real World MIMO

Theoretical MIMO results, coupled with real-world feasibility demonstrations [54, 51], have led to MIMO (often coupled with orthogonal frequency-division multiplexing) being included in many emerging wireless systems. Although particular implementations vary, a major commonality between most of the standards is the inclusion of the Alamouti space-time block code (STBC). In fact, 802.11n, 802.16, and 3GPP (Release 7 and 8) all include this MIMO technique with varying antenna constellations[68]. 802.11n currently supports 2-by-2 (2 transmitters and 2 receivers), while 802.16 and 3GPP Release 7 support 2 by 1. All include options for higher order antenna constellations in the future (mainly 4-by-4). Since Alamouti is such a prevalent component of emerging systems, we now examine attacks specifically targeted at the Alamouti STBC scheme in an effort to determine any general vulnerabilities.

3.3.6 Alamouti STBC Overview

The Alamouti 2-by-1 scheme [75] is essentially a spatial repeater with a decoding trick used to minimize computation at the receiver. Assume that Alice has two transmit antennas and wishes to send two symbols, c_1 and c_2, to Bob, who has a single receive antenna. During the first symbol period, Alice simultaneously

sends c_1 from antenna 1 and c_2 from antenna 2. The next symbol period, Alice simultaneously sends $-c_2^*$ from antenna 1 and c_1^* from antenna 2, where $*$ denotes the complex conjugate. Let the channel coefficients h_1 and h_2 represent the paths between Alice's two antennas and Bob. The received signal for the first symbol period is $r_1 = c_1h_1 + c_2h_2 + n_1$, and for the second symbol period is $r_2 = -c_2^*h_1 + c_1^*h_2 + n_2$,where n_1 and n_2 are additive noise components. By conjugating the signal received during the second symbol period, the received signal over the two symbol periods is given by $\mathbf{r} = \mathbf{Gc} + \mathbf{n}$, where

$$\mathbf{r} = \begin{pmatrix} r_1 \\ r_2^* \end{pmatrix} \quad \mathbf{G} = \begin{pmatrix} h_1 & h_2 \\ h_2^* & -h_1^* \end{pmatrix}$$

$$\mathbf{c} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \quad \mathbf{n} = \begin{pmatrix} n_1 \\ n_2^* \end{pmatrix}.$$

Decoding is done at the receiver by taking advantage of the fact that $\mathbf{G}^H\mathbf{G} = \alpha\mathbf{I}_2$, where $\alpha = |h_1|^2 + |h_2|^2$. Including ambient receiver noise, decoding is achieved by selecting the symbol-tuple $\hat{c}$ that minimizes:

$$\begin{aligned} d &= |\mathbf{G}^H\mathbf{r} - \alpha\hat{\mathbf{c}}|^2 = |\mathbf{G}^H(\mathbf{Gc} + \mathbf{n}) - \alpha\hat{\mathbf{c}}|^2 \\ &= |\mathbf{G}^H\mathbf{Gc} + \mathbf{G}^H\mathbf{n} - \alpha\hat{\mathbf{c}}|^2 = |\alpha\mathbf{c} + \mathbf{G}^H\mathbf{n} - \alpha\hat{\mathbf{c}}|^2 \\ &= |\alpha(\mathbf{c} - \hat{\mathbf{c}}) + \mathbf{G}^H\mathbf{n}|^2. \end{aligned}$$

As an example, given binary signaling with symbols c_1 and c_2, there are 4 symbol-tuple's to utilize:

$$\hat{\mathbf{c}} \in \{(c_1, c_1), (c_1, c_2), (c_2, c_1), (c_2, c_2)\}.$$

Extension to the 2-by-2 scheme is straightforward, as the minimization is performed over both receive antennas.

3.3.7 Jamming the Alamouti STBC

During channel estimation jamming, the matrix $\mathbf{G}$ becomes a corrupted $\hat{\mathbf{G}}$, which affects decoding as:

$$\begin{aligned} d_J &= |\hat{\mathbf{G}}^H\mathbf{r} - \hat{\alpha}\hat{\mathbf{c}}|^2 = |\hat{\mathbf{G}}^H(\mathbf{G}\mathbf{c}+\mathbf{n}) - \hat{\alpha}\hat{\mathbf{c}}|^2 \\ &= |\hat{\mathbf{G}}^H(\mathbf{G}\mathbf{c}+\mathbf{n}-\hat{\mathbf{G}}\hat{\mathbf{c}})|^2 = |\hat{\mathbf{G}}^H(\mathbf{G}\mathbf{c}-\hat{\mathbf{G}}\hat{\mathbf{c}}) + \hat{\mathbf{G}}^H\mathbf{n}|^2. \end{aligned}$$

We call d_J the Alamouti metric for clarity, and note that $\hat{\alpha} = |\hat{h}_1|^2 + |\hat{h}_2|^2$. For an error to occur, the value of d_J for an incorrect symbol-tuple must be minimal. Ignoring noise, absolute minimization occurs when $(\mathbf{G}\mathbf{c} - \hat{\mathbf{G}}\hat{\mathbf{c}})$ lies in the null space of $\hat{\mathbf{G}}^H$. We now consider two special cases: $\hat{\mathbf{G}}^H = \mathbf{0}$ and $\mathbf{G}\mathbf{c} - \hat{\mathbf{G}}\hat{\mathbf{c}} = \mathbf{0}$. The first case implies that Eve has convinced Bob that the channel is $\mathbf{0}$, which is extremely unlikely and would immediately reveal Eve's presence. The second case, however, deserves further examination.

Selective Symbol Jamming

We begin our analysis with a powerful adversarial model where Eve is equipped with ample situational awareness and can force Bob to use a $\hat{\mathbf{G}}$ of her choosing. Given the signal constellation, the channel estimate, and the symbol-tuple to be transmitted, Eve can choose the symbol-tuple, $\hat{\mathbf{c}}$, that she wishes Bob to decode by giving Bob $\hat{\mathbf{G}}$, so that $\hat{\mathbf{G}} = \mathbf{G}\mathbf{c}\mathbf{V}_{\hat{c}}\Sigma_{\hat{c}}^{-1}\mathbf{U}_{\hat{c}}^H$, where $\hat{\mathbf{c}} = \mathbf{U}_{\hat{c}}\Sigma_{\hat{c}}\mathbf{V}_{\hat{c}}^H$. We note that $\Sigma_{\hat{c}}^{-1}$ is not the traditional matrix inverse of $\Sigma_{\hat{c}}$, but rather $\Sigma_{\hat{c}}^{-1}$ is computed by transposing $\Sigma_{\hat{c}}$ and inverting its non-zero diagonal elements (i.e. singular values).

Let us look at an example using the real-valued channel $\mathbf{h} = [7\ {-8}]$, and BPSK modulation with symbols '-1' and '1'. The resulting symbol-tuples are defined in Table 3.1. Table 3.2 depicts a $\hat{\mathbf{G}}$ (computed via the SVD method described above) that forces Bob to decode symbol-tuple $\mathbf{c}(1)$ rather than the transmitted symbol-tuple $\mathbf{c}(0)$, when there is no additional channel noise. The table also displays

Table 3.1: BPSK Symbol Tuples

BPSK Symbol Tuples			
$\mathbf{c}(0)$	$\mathbf{c}(1)$	$\mathbf{c}(2)$	$\mathbf{c}(3)$
$\begin{bmatrix} -1 \\ -1 \end{bmatrix}$	$\begin{bmatrix} -1 \\ 1 \end{bmatrix}$	$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$	$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$

Table 3.2: Selective Symbol Jamming

Selective Symbol Jamming			
Transmitted $\mathbf{c}(0)$	Decoded $\mathbf{c}(1)$	$\hat{\mathbf{G}}$	
$\begin{bmatrix} -1 \\ -1 \end{bmatrix}$	$\begin{bmatrix} -1 \\ 1 \end{bmatrix}$	$\begin{bmatrix} -0.5 & 0.5 \\ -7.5 & 7.5 \end{bmatrix}$	
Metrics, d_J			
$\mathbf{c}(0)$	$\mathbf{c}(1)$	$\mathbf{c}(2)$	$\mathbf{c}(3)$
160	0	320	160

the metrics that Bob computes for the other possible symbol-tuples when $\mathbf{c}(0)$ is transmitted. It is evident that minimization occurs with $\mathbf{c}(1)$.

But, a jamming attack that replaces $\mathbf{G}$ with $\hat{\mathbf{G}}$ affects *all* of the transmitted symbol-tuples. Table 3.3 illustrates Eve's impact across every possible transmitted symbol-tuple. For example, when $\mathbf{c}(0)$ is transmitted, $\mathbf{c}(1)$ is decoded with a metric of 0. Likewise, when $\mathbf{c}(3)$ is transmitted, $\mathbf{c}(2)$ is decoded. Note that transmitted symbol-tuples $\mathbf{c}(1)$ and $\mathbf{c}(2)$ have identical decoded metric values, and that they also exhibit two possible valid decoded symbol-tuples ($\mathbf{c}(0)$ and $\mathbf{c}(3)$), both of which are incorrect. The capability to strategically change certain symbol-tuples would allow the jammer Eve to disrupt *specific* data sequences within messages. Fortunately, for Alice and Bob, it is difficult for Eve to affect $\hat{\mathbf{G}}$ in the *exact* manner considered above. First, even letting Eve replace $\mathbf{G}$ must adhere to simple restrictions, as even a simple implementation of the Alamouti

Table 3.3: Jammed Metrics - Ideal

Alamouti Metrics, d_J				
TX \ RX	$\mathbf{c}(0)$	$\mathbf{c}(1)$	$\mathbf{c}(2)$	$\mathbf{c}(3)$
$\mathbf{c}(0)$	160	0	320	160
$\mathbf{c}(1)$	0	160	160	0
$\mathbf{c}(2)$	0	160	160	0
$\mathbf{c}(3)$	160	320	0	160

Symbol-tuple Transitions

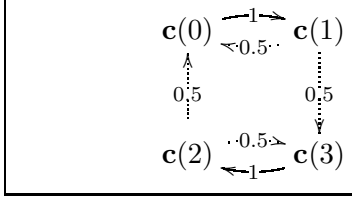

scheme would check that $\hat{\mathbf{G}}$ satisfies:

$$\hat{\mathbf{G}} = \begin{pmatrix} a & b \\ b^* & -a^* \end{pmatrix}.$$

We acknowledge, though, that our outlined general exploitation of $\hat{\mathbf{G}}$ may be achievable by way of an internal system-level attack. However, for this work, we are only focused on exploitations that are available at the physical RF level. So, we must consider perturbation matrices generated from $\hat{\mathbf{h}} = [\hat{h}_1 \ \hat{h}_2]$. Given this constraint, Eve is *still* capable of selective symbol jamming. Table 3.4 depicts the symbol-tuple transition table when Bob performs decoding with $\hat{\mathbf{h}} = [-7 \ -8]$. Note that while the desired symbol-tuple transition metric for $\mathbf{c}(0) \Rightarrow \mathbf{c}(1)$ is not 0, it does represent a global minimum (of 21).

At this point, let us take a moment to discuss error nomenclature— specifically, the differences between symbol-tuple errors, symbol errors, and bit errors. Symbol-tuple errors occur when Bob incorrectly decodes a symbol-tuple. This does not necessarily mean that *all* of the symbols that Bob decodes are in error. Rather, it ensures that at least one symbol within the symbol-tuple is in error. For the Alamouti 2-by-1 STBC, there are 2 symbols transmitted per symbol-tuple, so a symbol-tuple error implies at least one symbol was decoded in error.

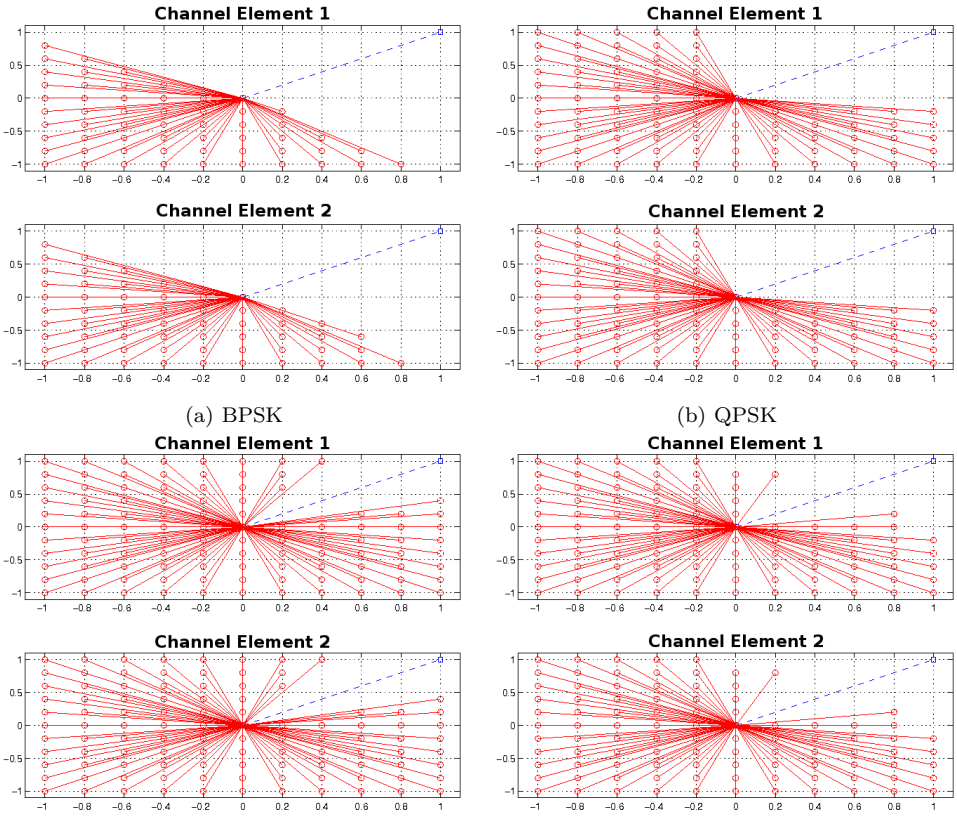

(a) BPSK (b) QPSK

(c) 8PSK (d) 16QAM

Figure 3.3: Jamming the Alamouti 2-by-1 space-time block code (STBC) scheme using MPSK and QAM constellations: (a) BPSK (b) QPSK (c) 8PSK (d) 16QAM. The actual channel is given by $\mathbf{h} = [(1+j)(1+j)]$

For the ideal selective symbol jamming case presented in Table 3.3, we see that symbol-tuple $c(0)$ is decoded incorrectly as symbol-tuple $c(1)$. Looking at the symbol-tuple mappings in Table 3.1, we see that this produces 1 symbol error. If, in fact, the symbol-tuple had been decoded as $c(3)$, then there would be 2 symbol errors. Bit errors then arise as a result of the symbol definitions. At this juncture, we do not consider bit-to-symbol mappings since they are user-defined and inconsequential to our current analysis.

Table 3.4: Jammed Metrics - Practical

Alamouti Metrics, d_J				
TX \ RX	$\mathbf{c}(0)$	$\mathbf{c}(1)$	$\mathbf{c}(2)$	$\mathbf{c}(3)$
$\mathbf{c}(0)$	210	21	319	241
$\mathbf{c}(1)$	319	210	241	21
$\mathbf{c}(2)$	21	241	210	319
$\mathbf{c}(3)$	241	319	21	210

Symbol-tuple Transitions

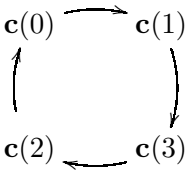

Simulation Results

We now present results from a simulation study. Under the Alamouti 2-by-1 scheme, two symbols are sent over two symbol periods from two antennas to a single receiver. The narrowband, time-invariant channel matrix for such a scheme is given by two complex coefficients, $\mathbf{h} = [h_1\ h_2]$. In jamming this channel, we look at perturbations of $\mathbf{h}$ in the complex plane given a "jammer power constraint." In our simulations, the true channel is set to $\mathbf{h} = [(1+j), (1+j)]$, and the jammer is empowered to alter the in-phase (I) and quadrature (Q) components of an h_j arbitrarily in the range of $[-1, 1]$ in step sizes of 0.2 (e.g. one jammed channel instance is $\mathbf{h}_J = [(-1 + 0.6j), (-0.2 - 0.4j)]$). For each possible jammed channel, the simulation iterates over every possible input codeword (or symbol-tuple) associated with the user specified signal constellation. (For a given constellation with M symbols, there are M^2 possible 2-dimensional codewords.) For each jammed channel instance, the simulation calculates the number of symbol errors that occur for every possible input codeword (for the Alamouti 2-by-1 scheme, the maximum number of symbol errors *per* input codeword is 2). The simulation repeats this process for every possible jammed channel case. At the end of the simulation, the jammed channels that resulted in the most symbol errors (which

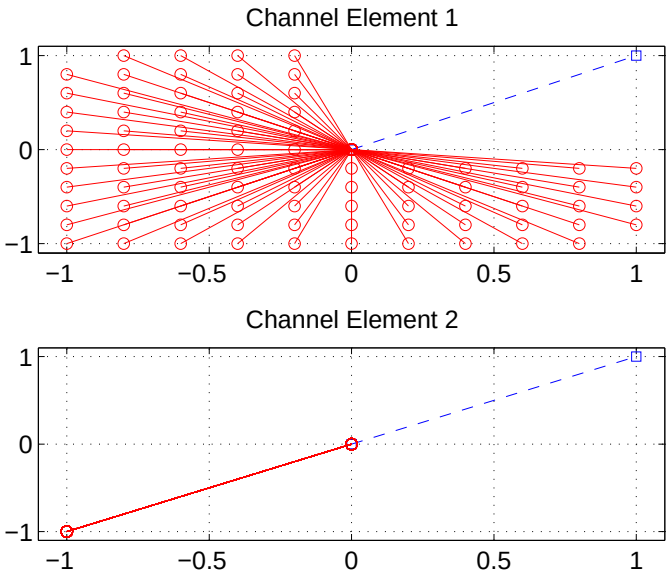

Figure 3.4: Jamming the Alamouti 2-by-1 space-time block code (STBC) for a QPSK constellation where we hold the jammed second channel coefficient constant at -1-j. The actual channel is given by $\mathbf{h} = [(1+j)(1+j)]$.

in all cases examined turned out to be the maximum errors possible, which is $2M^2$) are logged. Each simulation produces two subplots. The first relates to the first channel coefficient, h_1, while the second subplot relates to the second channel coefficient, h_2). The dotted line is the unjammed channel coefficient, while the solid lines are the jammed channel coefficients. While these plots do not directly illustrate which h_{J1} goes with which h_{J2}, it is interesting to note the structure of the jammed channel coefficients when compared to the actual signal constellation. M-PSK cases for $M = \{2, 4, 8\}$ and 16-QAM are shown in Figure 3.3, where a "perfect" channel of $\mathbf{h} = [(1+j), (1+j)]$ was used. Note how the symbol error regions mimic the signal constellation.

In order to illustrate the 1-to-1 nature of the jammed coefficients, Figure 3.4 shows results from a QPSK simulation where the second jammed coefficient is held constant ($h_{J2} = [-1-j]$). Each of the points displayed for the first jammed coefficient, h_{J1}, results in the maximum number of symbol errors. Again, the

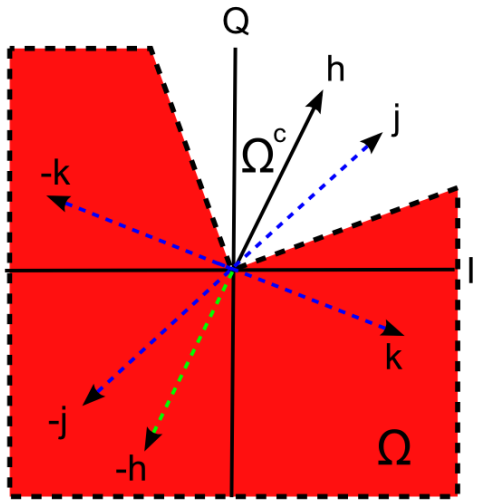

Figure 3.5: The maximally effective jamming region, Ω, is illustrated for a single channel coefficient, h, under the Alamouti 2-by-1 STBC using QPSK. Channel estimates -h, -j, k, or -k all lie in the desired jamming region.

error region clearly relates to the signal constellation. This is important because 64-QAM is mentioned for WiMAX and 3GPP MIMO standards when using Alamouti-based schemes (and extension of these results to 64-QAM is straightforward).

It is apparent that the maximal symbol error region related to the Alamouti STBC mimics the signal constellation. Figure 3.5 illustrates the maximally effective jamming region, Ω, associated with QPSK *for a single channel coefficient*. If Bob decodes using channel coefficients that lie within their respective Ω regions, then every transmitted symbol will be recovered in error. Let us examine the first channel element, so that $h = h_1$ from Figure 3.5. Bob will be effectively jammed on this antenna if he decodes with j, k, $-k$, or $-h$, whereas decoding with j (or of course h), will not incur any errors. The same geometrical region jamming phenomenon holds for any M-ary PSK or QAM modulation. Note that attacking either the I *or* the Q component of a given channel coefficient does *not* guarantee inclusion within Ω (the reader can verify this by considering QPSK modulation

Table 3.5: Jammed Metrics - Inverted

Alamouti Metrics, d_J				
TX \ RX	$\mathbf{c}(0)$	$\mathbf{c}(1)$	$\mathbf{c}(2)$	$\mathbf{c}(3)$
$\mathbf{c}(0)$	376	266	266	0
$\mathbf{c}(1)$	266	376	0	266
$\mathbf{c}(2)$	266	0	376	266
$\mathbf{c}(3)$	0	266	266	376
Symbol-tuple Transitions				
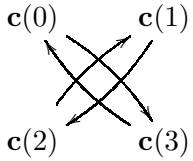				

with $h_1 = 1 + 0.2$j where the Q component is inverted). However, observe that for any of the examined schemes, it is sufficient to simply invert the channel coefficients (both I and Q) in order to maximally jam the communication. Thus, Eve need not even know the modulation as long as she can reliably *invert* Bob's estimate of the true channel.

Channel Inversion Attack

We now examine an attack where Eve successfully inverts the channel estimate so that Bob uses $\hat{\mathbf{h}} = -\mathbf{h}$ (i.e. $\hat{\mathbf{G}} = -\mathbf{G}$) to decode messages sent by Alice. To gain a better understanding of the attack, let us look at two examples using different modulations: (1) BPSK, and (2) QPSK. First, reconsider the previous BPSK example. Table 3.5 illustrates the decoded metrics given a noiseless channel and the resulting symbol-tuple transitions when Eve enacts a channel inversion attack. It is evident that each symbol-tuple is perfectly disturbed. In fact, the symbols themselves are all reflected in the I/Q plane (i.e. "$c_0 = 1''$ $\leftrightarrow$ "$c_2 = -1''$ and "$c_1 = j''$ $\leftrightarrow$ "$c_3 = -j''$).

Now consider an example with QPSK and a true channel of $\mathbf{h} = [(7-2j)(-8+4j)]$. Using $\{1, j, -1, -j\}$ as the QPSK symbols, we define symbol-tuple mappings

Table 3.6: QPSK Symbol Tuples

QPSK Symbol-tuples			
$\mathbf{c}(0)$	$\mathbf{c}(1)$	$\mathbf{c}(2)$	$\mathbf{c}(3)$
$\begin{bmatrix} 1 \\ 1 \end{bmatrix}$	$\begin{bmatrix} 1 \\ j \end{bmatrix}$	$\begin{bmatrix} 1 \\ -1 \end{bmatrix}$	$\begin{bmatrix} 1 \\ -j \end{bmatrix}$
$\mathbf{c}(4)$	$\mathbf{c}(5)$	$\mathbf{c}(6)$	$\mathbf{c}(7)$
$\begin{bmatrix} j \\ 1 \end{bmatrix}$	$\begin{bmatrix} j \\ j \end{bmatrix}$	$\begin{bmatrix} j \\ -1 \end{bmatrix}$	$\begin{bmatrix} j \\ -j \end{bmatrix}$
$\mathbf{c}(8)$	$\mathbf{c}(9)$	$\mathbf{c}(10)$	$\mathbf{c}(11)$
$\begin{bmatrix} -1 \\ 1 \end{bmatrix}$	$\begin{bmatrix} -1 \\ j \end{bmatrix}$	$\begin{bmatrix} -1 \\ -1 \end{bmatrix}$	$\begin{bmatrix} -1 \\ -j \end{bmatrix}$
$\mathbf{c}(12)$	$\mathbf{c}(13)$	$\mathbf{c}(14)$	$\mathbf{c}(15)$
$\begin{bmatrix} -j \\ 1 \end{bmatrix}$	$\begin{bmatrix} -j \\ j \end{bmatrix}$	$\begin{bmatrix} -j \\ -1 \end{bmatrix}$	$\begin{bmatrix} -j \\ -j \end{bmatrix}$

in Table 3.6. Under the channel inversion attack, the symbol-tuple transitions are $c_i \leftrightarrow c_{(i+6)\%16}$ for $i \in \{2, 3, 6, 7, 10, 11, 14, 15\}$.

Channel Inversion Attack Mitigation

Since the channel inversion attack causes symbols to be reflected in the I/Q-plane, higher-layer logic may mitigate this attack. If Bob suspects that he is under this attack, he can perform source decoding based upon the decoded symbols *and* their reflection in the I/Q-plane. It should be noted that to mitigate *any* stationary channel estimate attack, Bob can perform source decoding based upon every possible symbol-tuple mapping. However, complexity rises exponentially in order to gain such reliability.

An alternative mitigation to the channel inversion attack entails altering the modulation scheme. Since the attack does not necessarily focus upon the magnitude of the transmission, reasonable immune modulation schemes would include amplitude-based modulations. However, it is important to note that some amplitude-based modulations can be viewed as forms of phase-shift keying, or at least maintain some phase-related dependencies. For instance, consider binary pulse-amplitude modulation (2-PAM). This scheme is essentially equivalent

to BPSK, and is therefore affected by the channel inversion attack. Likewise, 4-PAM, is also susceptible to the attack. Under the 4-PAM scheme, amplitude levels of $-3A$, $-A$, A, and $3A$ are used to transmit data. Since the symbol associated with $-3A$ is equivalent to $3A$ with a 180-degree phase shift, it will suffer degradation under the channel inversion attack. We now propose a slight modification to the general PAM scheme that *is* immune to the channel inversion attack. By forcing any M-PAM scheme to be centralized about MA, the modulation becomes purely magnitude-based, where negative values become non-problematic . Spherical decoding based on pure signal magnitudes will remain impervious to the attack. This remains true for *any* channel-based attack that is stationary over the decoding process (which is true for any Alamouti STBC). While this modified PAM scheme is resistant to the channel inversion attack, there is an obvious bandwidth tradeoff.

802.11n Application

Let us take a moment to discuss applicability to 802.11n. To do this, we first provide some protocol background. 802.11n allows for three main modes of operation: Legacy Mode, Mixed Mode, and Green Field. Pertinent to this work are Mixed Mode and Green Field, where packet fields are designated for channel sounding. These fields are called High Throughput Long Training Fields (HT-LTF), and are used by the receiver to estimate the channel matrix, $\mathbf{H}$ [76]. While the beginning packet structure in Mixed Mode and Green Field differ, both end with HT-LTFs preceding the actual signal data. The two packet structures are illustrated in Figure 3.6. Under MIMO operation, Eve need only jam during HT-LTF transmission in order to attack the channel estimation procedure. Under 802.11n, each HT-LTF field lasts for 4us (with an optional extended HT-LTF lasting another 4us). This represents a relatively insignificant jamming time relative to data transmission time. (A 1500 byte data packet transmitted at 54 Mbps

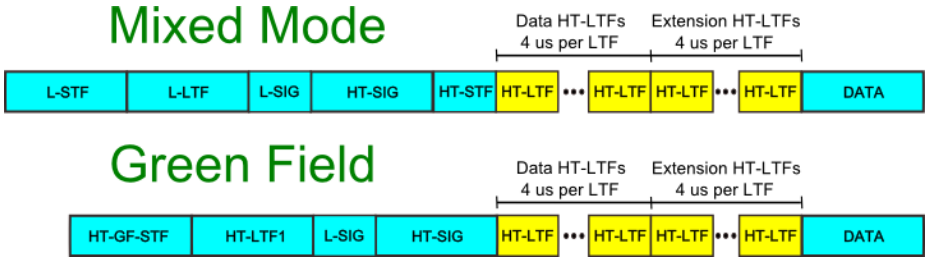

Figure 3.6: The 802.11n packet structure for Mixed Mode and Green Field operation are depicted. High Throughput Long Training Fields (HT-LTF) are used to estimate the channel matrix, **H**, between the transmitter and the receiver.

will last about 222us). While Eve may wish to implement a channel inversion attack on the Alamouti scheme, she still has options if it is not in use. Earlier, in Section 3.3.4 we illustrated the effects of jamming during SVD-based MIMO. Similar research and conclusions could be applied to any scheme utilizing channel estimation. Consider a simple scenario where Eve has a single antenna. Continual jamming during the HT-LTFs would result (given appropriate channel coherence time) in the same estimate for every channel. Thus, **H** would appear singular.

Real World Implementation

Implementing any of the aforementioned attacks in the real world raises some questions and concerns. For instance, Eve can never truly know the phase of the signals arriving at Bob (from Alice or Eve) due to channel unknowns (e.g. multipath, antenna separations, etc.). Thus, if Eve is conducting a channel inversion attack, she cannot guarantee that her signal is 180 degrees out of phase with Alice's when it arrives at Bob. However, given limited situational knowledge such as her *jammer to signal ratio* (J/S), Eve can reliably estimate her success by analyzing jamming regions. Let us revisit the Alamouti 2-by-1 STBC for QPSK. If h is the actual channel coefficient between Alice and Bob, and r is Eve's transmission as seen by Bob, then the probability of moving the channel estimate into the

maximally effective jamming region, $P(\Omega)$, can be calculated for a single antenna element by:

$$P(\Omega) = \begin{cases} 0 & \text{if } \frac{|r|}{|h|} \leq \frac{\sqrt{2}}{2}, \\ \frac{2}{\pi r^2}[\cos^{-1}(\frac{A}{|r|})r^2 - AB] & \text{if } \frac{\sqrt{2}}{2} < \frac{|r|}{|h|} \leq 1, \\ \frac{3}{4} - \frac{1}{\pi r^2}[AB + \sin^{-1}(\frac{A}{|r|})r^2 + h^2] & \text{if } \frac{|r|}{|h|} > 1, \end{cases}$$

where $A = \frac{|h|\sqrt{2}}{2}$, $B = \sqrt{r^2 - \frac{h^2}{2}}$, and we define $r^2 = rr^*$, and $h^2 = hh^*$. The J/S can be calculated in decibels (dB) by $10\log_{10}\frac{r^2}{h^2}$. Figure 3.7 illustrates two extreme J/S scenarios (J/S $\ll$ 0 dB and J/S $\gg$ 0 dB). When the J/S $\ll$ 0, Eve simply does not perturb the channel estimate enough to make a difference. However, when the J/S $\gg$ 0, Eve has a 3/4 chance of maximally interfering with a given antenna element. (For a general MPSK scheme, the jamming success probability of a single antenna is $(M-1)/M$ when the J/S $\gg$ 0.)

For ease of presentation, we continue by assuming that Eve is a very strong jammer (i.e. J/S $\gg$ 0 dB). Under the well accepted wide-sense stationary uncorrelated scattering model (WSSUS) [77], in a rich multipath environment Eve's effect on each antenna element will be independent with regard to phase. Thus, Eve will penetrate *both* effective jamming regions 9/16 of the time. Furthermore, Eve has 3/8 chance of interfering effectively with antenna 1 *or* antenna 2. So, Eve is only completely ineffective as a jammer 1/16 of the time. Additionally, these probabilities equate to $((M-1)/M)^2$, $2(M-1)/(M^2)$, and $1/(M^2)$, respectively for a general MPSK constellation. But, there is more that Eve can do to *ensure* her effectiveness.

To *guarantee* some level of jamming success, Eve may implement an oscillating channel inversion attack. Although Eve cannot be certain of the arriving phase, she can oscillate her transmit signal by 180 degrees. By doing so within the time coherence of the channel (and assuming some synchronicity between bursts), she ensures that *at least* one out of every two jamming attempts lies within the

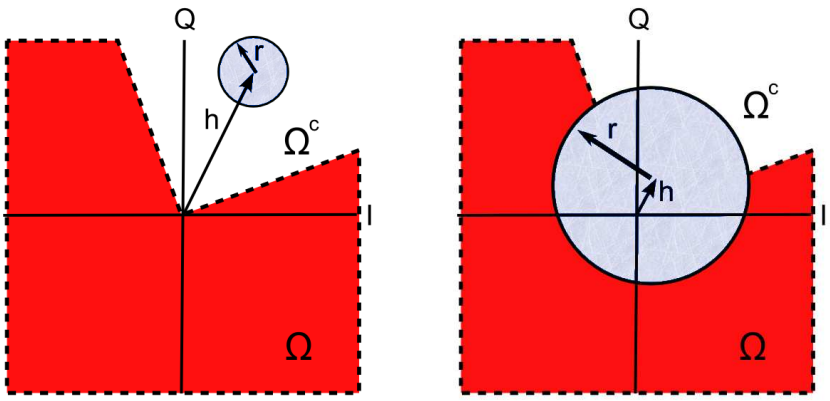

(a) J/S $\ll$ 0 dB (b) J/S $\gg$ 0 dB

Figure 3.7: Jammer-to-signal ratio (J/S) impact regions are investigated. (a) When the J/S $\ll$ 0 dB, the jammer has no real effect. (b) When the J/S $\gg$ 0 dB, the probability of randomly perturbing the channel estimate for a single antenna into the jamming region approaches 3/4.

optimal jamming region for a single antenna. Consider the jamming region finite state machine in Figure 3.8. Here, Ω_1 indicates the optimal jamming region for antenna 1, and Ω_2 indicates the jamming region for antenna 2. Four possible states arise as a result of Eve's oscillating attack. There is a 9/16 chance that Eve's initial attack falls within the optimal regions for both antenna 1 and 2. As can be seen, there is only 1/9 chance of leaving that state to the completely ineffective state of (Ω_1^c,Ω_2^c), while remaining in the totally effective state occurs with probability 4/9. Note that as a jammer Eve will be *totally* effective 25% of the time (9/16 $\times$ 4/9), and completely ineffective *never*. In fact, the worst performance, occurring only 6.25% of the time, will jam the desired regions half of the time. Note that effectively defeating signal data even half of the time is sufficient to defeat most inter-burst source decoding schemes.

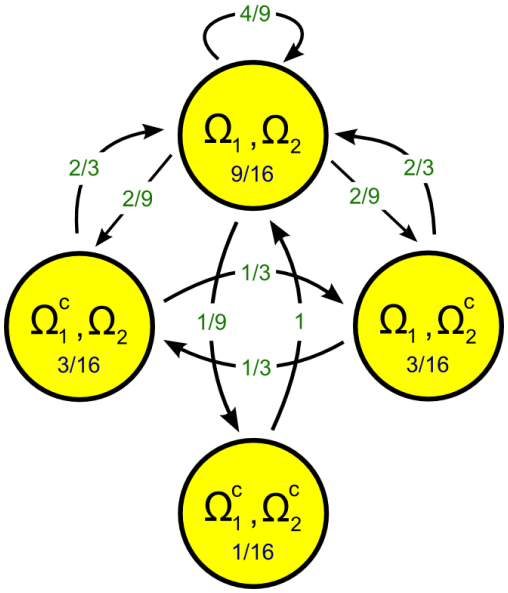

Figure 3.8: Above is the Jamming Region Finite State Machine for an oscillating channel inversion attack against the 2-by-1 Alamouti space-time block code (STBC) using QPSK when the jammer-to-signal ratio (J/S) is large (J/S $\gg$ 0dB).

Alamouti STBC Jamming Experiment

To illustrate the effectiveness of jamming the training sequence (e.g. the oscillating channel inversion attack) in the real-world, experiments were conducted using GNU Radio [7] and the Universal Software Radio Peripheral (USRP) [6]. An Alamouti 2-by-1 scheme was implemented using QPSK modulation in the 1800 MHz band. Additionally, an Agilent Vector Signal Analyzer (VSA) and an Agilent Arbitrary Waveform Generator (ARB) were used.

Alice was outfitted with 2 transmit antennas and sent bursts once every $4ms$ to Bob. In order for Bob to estimate the channel, Alice preceded her data transmissions with channel sounding waveforms, TS1 and TS2. TS1 was sent from transmit antenna 1, and was shortly followed by TS2 from antenna 2. Thus, Bob could obtain the channel estimates necessary for decoding under the Alamouti scheme. This procedure follows real world operation, such as found in Mixed and Green Field Modes in 802.11n [27]. For this experiment, Alice's data payload

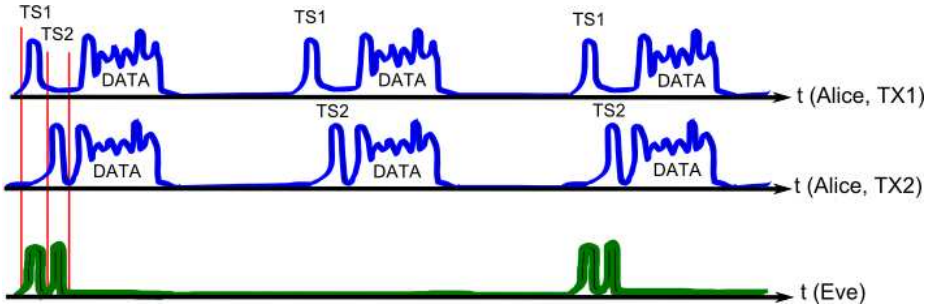

Figure 3.9: Alice transmits TS1 and TS2 from her antennas so that Bob can estimate the channel coefficients. Eve transmits her own channel sounding waveforms during these slots in an effort to perturb Bob's estimations. For our experiments, Eve only interferes with every other burst.

consisted of 12 symbols transmitted at 12.5 kBd. Eve then entered the picture by jamming with her own channel sounding waveforms, but only during every other burst (i.e. every $8ms$). This jamming procedure is illustrated in Figure 3.9 and the experimental setup is shown in Figure 3.10. Care was taken to produce a stationary, yet richly scattered environment (e.g. no line of site and many reflectors).

As mentioned earlier, we oscillate Eve's jamming waveform to *ensure* some level of jamming success. So, Eve's transmissions were incrementally phased-shifted by 90° for every burst (which was chosen rather than 180° degrees to provide an expanded sample set for analysis). Additionally, several J/S scenarios were tested (J/S $\in \{20, 10, 0\}$ dB), each of which resulted in 200 jammed and 200 unjammed bursts. We did not explore J/S $\ll 0$ because we showed earlier that it would not be effective (see Figure 3.7).

In analyzing the experimental results, we pay close attention to Bob's phase estimates for TS1 and TS2 for both the jammed and unjammed scenarios. This is because phase errors will mainly be responsible for symbol-tuple decoding errors in these J/S regions. Statistical results of the experiment are listed in Table 3.7. For each J/S scenario, the average phase offset and its variance are calculated for

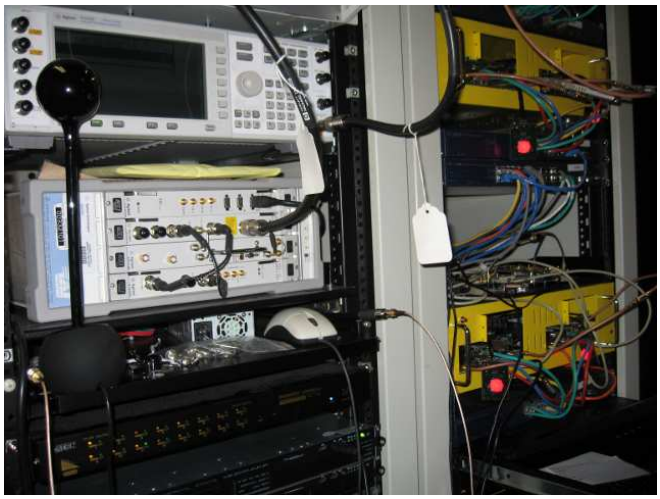

Figure 3.10: Equipment used in the Alamouti 2-by-1 space-time block code (STBC) jamming experiment is pictured. These platforms created and processed the waveforms; antennas were connected to these devices and positioned such that the environment was richly scattering.

the jammed and unjammed cases for each training sequence. Additionally, the symbol error rates are listed.

First, let us examine the large J/S cases (i.e. J/S $\in \{20, 10\}$). In these two experiments, Eve caused phase errors at both antennas reliably by an incremental 90 degrees for each jammed burst (and with a minimal variance). The reader should note, however, that there is another drift occurring on a burst-by-burst basis during unjammed operation. Here, Bob's phase estimates are drifting by a random yet constant offset from experiment to experiment (and the variances of the drift are also notably small). This drift is attributable to physical phenomena such as local oscillator (LO) drift and frequency offsets between tuners. (Note that this could also occur in mobile scenarios as a result of Doppler). So, it is apparent that Eve's actual effect at Bob is a function of her waveform *and* the unjammed phase drift (in fact it is the difference in the arriving phases between Alice's and Eve's transmissions). This is an important observation since in the real-world, Eve may not need to phase shift her waveform at all, therefore relying

on the effective phase shift incurred due to the aforementioned phenomena. A superior attack strategy would be for Eve to observe Alice's drift at Eve's own antenna, and then match her oscillating transmission accordingly. Note that this does not guarantee her initial phase arrival, only that the incremental drift from burst to burst is more controllable — which is her goal. In fact, it is noted that the initial phases for all J/S scenarios were random.

Now let us examine the low J/S experiment in more detail. It is clear that the unjammed phase drift is constant and with a low variance. However, while the average effect of the jammer produces the 90 degree incremental offset, the variance of the offset is very large (for both antennas). So, Eve effectively cannot control her impact.

Finally, we examine the symbol error rates from the experiments. For each J/S scenario, the symbol error rate is quite large (which was Eve's objective). The astute reader may wonder why the lower J/S scenario is showing a higher symbol error rate. This is because the higher J/S experiments were *not* accounting for Alice's phase drifts.

To summarize, in the high J/S arena, a real-world, oscillating channel inversion attack has been shown to be realizable and quite powerful. Additionally, as a direct result of our experimental analysis, an important feature of this type of attack would involve incorporating Alice's phase drift. In the low J/S regime, the oscillating channel inversion attack is still effective, however it is less controllable (as shown by the large variances listed in Table 3.7).

3.4 Attack Summary

In this Chapter, we illustrated various vulnerabilities associated with current and emerging wireless systems when subjected to a smart adversary with physical

Table 3.7: Experimental Results

Experimental Results (200 burst pairs/run)					
	Not Jammed		Jammed		
	TS1 (degs)	TS2 (degs)	TS1 (degs)	TS2 (degs)	J/S (dB)
Offset Average mean($\Delta\Theta$)	16	16	90	90	20
	121	121	90	90	10
	-130	-130	90	90	0
Offset Variance var($\Delta\Theta$)	2.85	2.73	0.15	0.18	20
	5.3	5.78	0.27	0.33	10
	2.56	2.57	38.86	69.74	0
Symbol Error Rate	<0.0001		0.6500		20
	<0.0001		0.7463		10
	<0.0001		0.7711		0

layer access. Given the recent advancements in SDR development, such an adversarial model is increasingly more pragmatic. Although much research has been done in the area of MIMO, little research has considered limitations due to a truly smart adversary. In this Chapter, we introduced and explored a strategy that can be used against MIMO systems that require accurate channel state information. By targeting the channel estimation procedure, an adversary may launch effective jamming attacks against unsuspecting users. We have presented jamming methodology for SVD-based MIMO systems in addition to recommending strategies to undermine the popular Alamouti STBC-based MIMO scheme. In addition, such attacks have been proven viable by way of analysis, simulations, and real-world experimentation. The attack strategies we present are general and may be applied to other MIMO systems.

Chapter 4
Mitigation

4.1 Motivation

The value of wireless communications arises from the ability to establish reliable communication between nodes that are not connected via an infrastructure. The quality of the communication between two entities is governed by several factors, including transmission power, modulation format, transmit and receive antenna capabilities, and environmental conditions.

Given reliable situational awareness, a device is more capable of choosing an effective and efficient communication scheme. Optimal modulation schemes and protocols can be determined, in addition to successful mitigation of adversarial operation. In this Chapter, we first consider traditional mitigation methodology, and then introduce new physical layer techniques that can be used to overcome channel degradations due to *both* natural phenomena and adversarial activity.

After our overview of classical mitigation techniques, we present novel new ways to address the MIMO channel estimate attack strategy developed in Chapter 3. Our methods provide general channel estimate authentication by exploiting physical layer properties of the wireless channel. The proposed solutions work for any system that relies upon channel estimate accuracy.

We then consider resolutions to another undesirable communications scenario, where we describe a unique mitigation strategy to an urban environment riddled with severe multipath in the presence of a powerful jammer. In this work, we

leverage the multipath effects of the channel in order to secure a reliable communication link.

4.2 Popular Techniques

Many techniques currently exist to mitigate poor channel environments or adversarial activity. Each technique aims to exercise an available degree of freedom. A listing of popular mitigation strategies follows:

Transmission Power

The most obvious mitigation strategy to address an adversary or poor environment is to simply transmit with more power [78]. However, this may be wasteful, or even futile if addressing a hidden node problem. Further, adding more power is often an ineffective strategy in an adversarial scenario since adversaries are commonly equipped with more resources.

Frequency Agility

Changing to a different frequency can be effective for fading environments or when retreating from adversaries that cannot track the retreat in an expedient manner [79, 80]. Care must be taken to not switch to a fading channel, or one occupied with other users. In this context, we comment that frequency hopping spread spectrum (FHSS) techniques provide an excellent solution, particularly if the hopping sequence is secret.

Modulation Agility

Another mitigation strategy is to dynamically adapt the modulation techniques [78, 81, 82]. One may liken this to 802.11 operation, where various modulation techniques are available for use depending upon the environmental conditions and user assets. An excellent adversarial mitigation strategy would be to switch to direct sequence spread spectrum (DSSS) modulation

when under attack, thus mitigating any narrow band attacks via the virtue of signal despreading.

Spatial Diversity

Changing device location (i.e. physically retreating) is a good mitigation technique for multipath rich environments and adversarial scenarios [78]. Falling under this classification of techniques would be classical beamforming and MIMO strategies.

Advanced Techniques

Advanced techniques exist such as those described in [83], where anti-jamming timing channels are used to communicate low-bit rate messages during active adversarial operation.

While each of the above techniques have their own merits, it is quite advantageous to employ mitigation strategies that combine aspects of each category. In fact, optimal mitigation decisions can be made by leveraging machine learning techniques in conjunction with the situational awareness achieved from methods described Chapter 2. We now introduce PHY layer techniques to deal with the channel estimation attacks presented in Chapter 3.

4.3 Channel Estimate Authentication

4.3.1 Motivation

Reliable communication over the wireless channel is often hampered by the effects of the channel itself. Transmitted waveforms interact with reflectors in the channel resulting in multiple distorted copies of the signal arriving at the intended receiver. Modern wireless systems mitigate these effects by using channel estimates to remove the distortion. Specific known waveforms are often transmitted

within data packets so that a receiver can obtain these channel estimates — such waveform segments are commonly referred to as *pilots*.

Pilots that are used in modern wireless communication systems are often simple waveforms (e.g. tones) and are unencrypted and unencoded — therefore they are unprotected and vulnerable to attack. Researchers have shown that current software defined radios (SDR) and limited protocol knowledge can be used to perform practical channel estimation attacks that greatly hamper system performance [84]. It is therefore vital that future wireless communication systems protect the channel estimation procedure. An obvious option would be to encode or encrypt the pilots themselves. However, such a resolution is impractical as it would overly complicate the detection process. Existing research has focused on transmitter authentication rather than channel estimate authentication, and solutions either do not address channel estimate authentication or require undesirable constraints such as operation within the channel coherence time or sacrificing system throughput for authentication [85, 86, 87, 88, 89]. Further, solutions are often restricted to operation in multipath-rich environments. In this paper, we propose methodology to embed channel estimate authentication messages into a physical (PHY) layer feature vector associated with the pilot waveform. The techniques work as a PHY Layer overlay to existing channel estimation procedures while also providing a level of channel estimation error inference. Further our pilot protection procedure works in any channel environment, without coherence time constraints, and without sacrificing any data throughput.

4.3.2 Channel Estimation Overview

Reliable wireless communication is complicated by the effects of the channel since a transmitted waveform experiences changes in amplitude and phase due to propagation loss, interaction with reflectors, and channel noise [3]. Over the timespan of a single packet, most channels are modeled as linear, time-invariant (LTI)

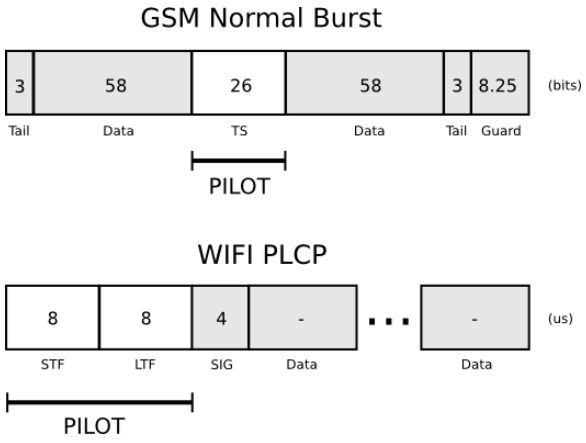

Figure 4.1: The GSM and WIFI protocols dedicate specific known signal segments, referred to as *pilots*, for the purpose of channel estimation.

systems, such that if Alice transmits $x(t)$, then Bob receives

$$y(t) = h(t) * x(t) + n(t),$$

where $n(t)$ is additive noise, $*$ denotes convolution, and $h(t)$ represents the channel response. Given an LTI channel with D distinct multipath components,

$$h(t) = \sum_{d=0}^{D-1} a_d e^{j\theta_d} \delta(t - \tau_d),$$

where a_d represents the amplitude attenuation, θ_d the phase shift, and τ_d the time-delay induced by the channel for the d_{th} signal path. Research has been done to characterize the distributions of a_d, θ_d, and t_d, which are highly dependent upon environment type (e.g. urban versus rural) [90, 91, 77]. For example, in a multipath-rich urban environment, a_d is often viewed as a Rayleigh random variable, and θ_d is assumed to be uniformly distributed [77].

In order to decode a signal properly, it is important for wireless communication systems to accurately estimate the effects of the channel. As mentioned in Section 4.1, protocols typically dedicate specific portions of the transmitted waveforms for the sole purpose of channel estimation. These waveform segments are known to both the transmitter and receiver, and are commonly referred to as *pilots*. In GSM, the pilot signal for normal bursts is a unique series of 26 transmitted bits

located in the middle of the waveform [92]. For 802.11 (WiFi) [27], the pilots are constant OFDM symbols transmitted at the beginning of the waveform. Pilot locations for GSM normal bursts and WiFi packets are depicted in Figure 4.1.

Real world channels are not stationary, however, as they change over time, frequency, and space. After a short amount of time, the channel decorrelates with itself, and pilots must be reissued. Knowledge of a channel at a particular time will not help in the future: a time window over which a channel induces a predictable phase is known as the coherence time. Similarly, the range of frequencies that experience comparable fading effects is referred to as the coherence bandwidth. Likewise, a wireless channel rapidly decorrelates with itself in space for distances larger than one half of the waveform's wavelength [93]. These are important observations because most new wireless communication systems operate over time, frequency, and space. We note that spatial diversity can result from mobility, but also by using multiple antennas, as is prominent in emerging wireless systems. For example, 802.11n, WiMAX, and LTE employ multi-input multi-output (MIMO) schemes coupled with OFDM, where messages are sent simultaneously from multiple antennas over numerous frequencies. Therefore, these protocols need to estimate the channel *reliably* in time, frequency, and space. Regardless of protocol, the goal remains the same — accurate estimation of channel state information (CSI).

Now consider that Alice wants to send a message over the wireless medium to Bob. To enable general channel estimation, Alice transmits a set of pilots $\mathbf{c} = \{c_1, c_2, c_3, ...c_k\}$ over time, frequency, and space (e.g. by using multiple antennas). Her goal is to choose $\mathbf{c}$ such that it properly covers the changes in the channel without being wasteful. In other words, the pilots should be transmitted sparsely enough so as not to waste resources but with high enough density such that the channel estimates obtained by Bob can be used to decode the data. This

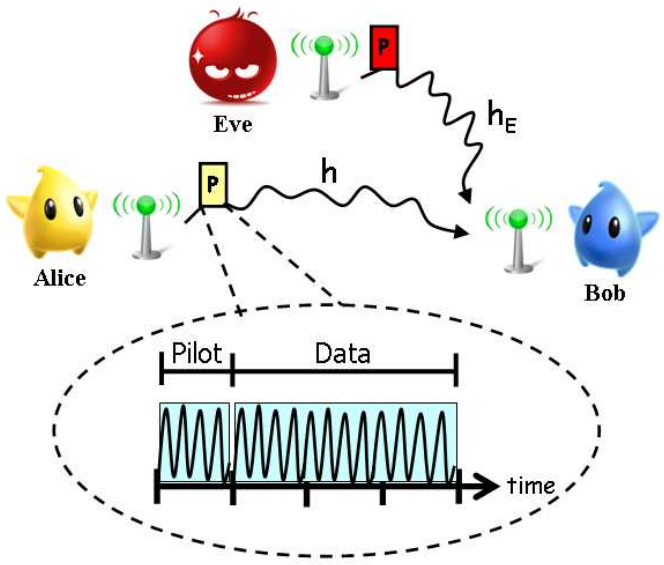

Figure 4.2: Alice transmits a packet to Bob over the wireless channel, h. The packet consists of a pilot and data. Evil adversary Eve attempts to interfere with the communication by transmitting her own rogue packet over the channel, h_E.

applies to both streaming and packetized data. For this paper, we consider packet-based protocols, however all principles and techniques can be directly applied to streaming data. There are thus many piloting strategies available. Classical results are listed in [94, 95], however the widely adopted "block" and "comb" schemes only consider time and frequency in a basic manner.

Assuming that the selected channel sounding sequence, $\mathbf{c}$, is properly chosen for the channel, inaccurate CSI estimates are still possible. Most commonly, this occurs due to channel noise, unintentional interferers (e.g. co-channel interference), and even malicious users. In fact, the authors in [84] show that targeted CSI attacks provide a covert, effective, and practical means of degrading system performance. Such an adversarial model is depicted in Figure 4.2, where the adversary Eve hampers communication between Alice and Bob by sending rogue transmissions. It is therefore important to be able to trust CSI estimates, which is why we propose a way to *authenticate* the estimates. Our work fills a major void because current and emerging wireless communications systems do not

provide a way to authenticate the channel estimates. Further, our methodology provides CSI authentication across coherence time intervals, and can be applied in any environment (not just richly-scattering). Moreover, our techniques also infer the *cause* of the invalid estimate, thus driving more intelligent mitigation decisions. For example, channel noise may be overcome simply with more power, while adversarial activity may require a more complex response [96, 97].

Our authentication techniques work in a manner similar to codeword transmission. The goal is to select $\mathbf{c}$ such that it encodes an authentication message. Our techniques leverage the degrees of freedom available in the channel while exploiting the fact that many channel sounding solutions exist. Properly decoding the message authenticates the CSI, while errors indicate an invalid channel estimate in addition to its probable cause (e.g. adversarial activity). Since most future wireless systems are multi-antenna (e.g. MIMO) and multi-carrier (e.g. OFDM), this is where we concentrate. We note that applying our techniques to single-antenna, single-frequency systems can be achieved by merely simplifying the methods.

4.3.3 Related Work

Work has been done in the area of physical layer authentication. In [85], the authors investigate and implement techniques that embed power-based authentication messages (referred to as tags) into the data transmissions. The technique requires accurate CSI estimates, which can be attacked because the pilot transmissions remain unaltered and unprotected. Further, the implementation lowers the overall system throughput because the tag is located within the data itself. Finally, the authentication scheme is susceptible to replay attacks.

In [86], the authors also use the entire transmission to send an authentication message, thus resulting in the same drawbacks with regard to throughput. The

scheme involves sending the authentication message by embedding a channel response into the message and relying on equalization techniques at the receiver to resolve the authentication message. Richly scattering environments may complicate processing, and the technique itself relies heavily upon successive messages within the coherence time of the channel.

In [87, 88, 89], the authors leverage statistically similar channel estimates over time to authenticate the transmitter. The techniques rely upon reliable successive CSI estimates in richly scattering environments, and [88] assumes a burst structure not emulative of current systems. Further, the computational complexity may not be practical.

Our work fills a major void in existing literature because it provides a practical method of CSI estimate authentication that is computationally simple and does not depend upon successive reliable estimates within the channel coherence time. Our techniques are also applicable to *any* channel, not just richly scattering environments. Further, the techniques that we present do not lower data throughput since the authentication messages are solely embedded within the pilots themselves. Data authentication is inherently part of our techniques because rogue data packets will not share the same CSI as the authenticated pilots, thus resulting in demodulation errors at the receiver. Finally, our techniques provide a level of error inference, which is important to properly drive mitigation decisions.

4.3.4 CSI Protection: The General Framework

We now present a general framework for protecting the CSI estimation procedure by embedding authentication messages within the pilot waveforms. Our framework assumes that Alice and Bob are equipped with a secret key, k. While such a key can be obtained by traditional key dissemination mechanisms, it is possible to procure a key passively and without direct communication via the techniques

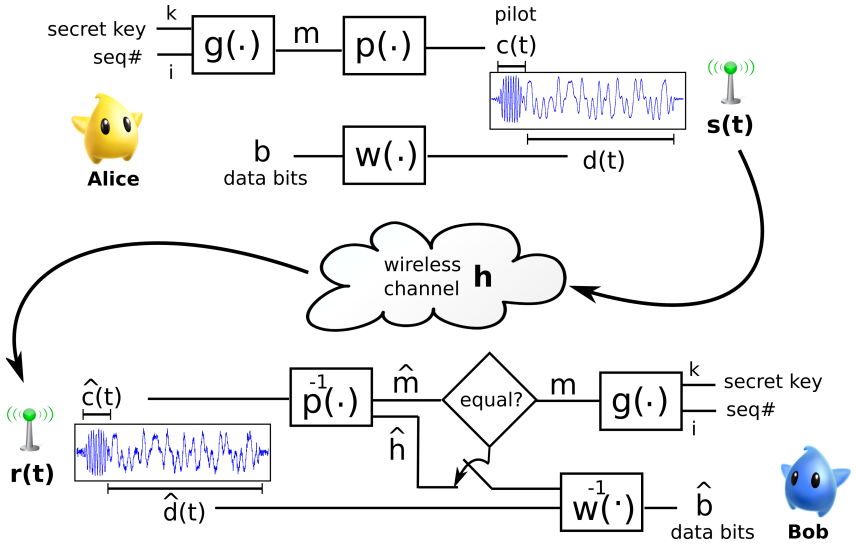

Figure 4.3: The general CSI protection procedure is depicted. Alice and Bob use a secret key, k, and a sequence number, i, to achieve CSI authentication by encoding messages into the waveform pilot. If the authentication message fails, then Bob should not trust the channel estimate to decode the data waveform.

described in [98, 99, 100, 101, 102]. Using k, Alice and Bob can embed authentication messages into the physical attributes of the channel sounding pilots themselves. A new authentication message, m_i, will be created by $m_i = g(k, i)$, where k is the secret key, i is the message sequence number, and $g(\cdot)$ is a one-way function. The sequence number prevents replay attacks, and the one-way function allows quick recovery from authentication errors. We incorporate the sequence number to prevent replay attacks, and use a one-way function to allow quick recovery from any authentication errors. To prevent replay attacks (e.g. Eve simply repeats Alice's transmissions back to Bob), Alice and Bob are required to use different keys for each directions of the communication. One may think of these authentication messages in a similar manner to the the pseudo-random hop sequence utilized in Bluetooth [103] or GSM [92].

The following general procedure, which is illustrated in Figure 4.3, describes how CSI estimate authentication can be achieved by using the pilot waveform to embed authentication messages:

1. Alice computes an authentication message, m_i, using the packet number, i, the secret key, k, and the one-way function, $g(\cdot)$.

$$m_i = g(k, i)$$

2. The authentication message is then mapped to a physical pilot waveform via $p(\cdot)$. Shortly, we will introduce a few of these pilot mapping functions.

$$c_i(t) = p(m_i)$$

3. Alice then generates the data waveform. Given L data bits for the i^{th} packet, $\underline{b}_i^L = \{b_0, b_1, \cdots b_{L-1}\}$, Alice generates the data waveform via the data waveform mapping function, $w(\cdot)$.

$$d_i(t) = w(\underline{b}_i^L).$$

4. Alice prepends the pilot to the data waveform[1], and transmits $s_i(t)$.

$$s_i(t) = [c_i(t)\ d_i(t)]$$

5. The transmission is affected by the wireless channel, so that Bob receives $r_i(t)$:

$$r_i(t) = \hat{s}_i(t) = [\hat{c}_i(t)\ \hat{d}_i(t)]$$

6. Bob uses the pilot waveform to decode the authentication message, $\hat{m}_i$, and estimate the CSI, $\hat{h}_i$.

$$(\hat{m}_i, \hat{h}_i) = p^{-1}(\hat{c}_i(t))$$

7. (a) If the authentication message is correct, then Bob can trust the CSI estimate and use it to recover the data bits.

$$\hat{\underline{b}}_i^L = w^{-1}(\hat{d}_i(t), \hat{h}_i), \text{ iff } \hat{m}_i = m_i$$

[1]The concatenation of the pilot and data waveforms should be done carefully to avoid phase discontinuities. The application of a Savitzky-Golay smoothing filter would achieve this while preserving relative signal magnitudes.

(b) If the message is incorrect, then Bob may sound an alert and/or demodulate the data *without* the CSI estimate.

$$\hat{\underline{b}}_i^L = w^{-1}(\hat{d}_i(t)), \text{ iff } \hat{m}_i \neq m_i$$

In our work, we focus on using *only* the pilot waveform to embed the authentication message. While the data waveform can also be used to convey our physical layer authentication message, we note that using only the pilot waveform allows us to select pilot mapping functions that prevent eavesdroppers from using the channel estimates obtained by the pilot to decode the data waveform.

The rest of this paper focuses on presenting various pilot waveform mapping functions, $p(\cdot)$. We begin by assuming single carrier MIMO capabilities and proceed by extending the schemes to incorporate multiple carriers. We then discuss practical extensions before demonstrating a selection of the techniques in real-world environments using a current cognitive platform. We now present two general classifications of pilot waveform mapping functions that we refer to as Frequency Quantization and Selective Usage.

4.3.5 Frequency Quantization

Signals transmitted within the same channel coherence bandwidth, Δ_{BW}, will experience comparable channel effects [104]. Therefore, if a pilot is transmitted at a nominal frequency of $\tilde{f}$, then the same pilot transmitted at frequency f_0 will experience similar fading effects if $|\tilde{f} - f_0| < \Delta_{BW}$. Thus, $a_d^{\tilde{f}} \approx a_d^{f_0}$, $\theta_d^{\tilde{f}} \approx \theta_d^{f_0}$, and $\tau_d^{\tilde{f}} \approx \tau_d^{f_0}$ for each of the D multipath components of $h(t)$. The coherence bandwidth of micro-cellular environments typical do not go below 80 KHz [105] — later we use this value in our real world experiments.

Using a single transmit antenna, Alice can embed an authentication message into the *frequency* of the pilot without sacrificing CSI validity so long as the transmission remains within Δ_{BW}. By quantizing Δ_{BW} into N distinct levels

surrounding the frequency that Alice would *normally* use to transmit the pilot, $\tilde{f}$, she can send an authentication message by transmitting her pilot in the appropriate frequency interval. Note that the pilot is transmitted at the offset frequency, f_0, while the data is still transmitted at the nominal frequency, $\tilde{f}$. Since Bob expects the authentication message at f_0, he can still properly decode the data transmitted at the nominal frequency, $\tilde{f}$, by accounting for the known frequency offset in his data recovery routines. And by definition of the coherence bandwidth, the CSI estimate for f_0 will approximate $\tilde{f}$.

We now use the frequency quantization concept to introduce three practical channel estimate authentication schemes: (1) Relative Frequency Codebook (RFC), (2) Binary Freqeuncy Codebook (BFC), and (3) Joint Frequency-Power Codebook (FPC). Figure 4.4 depicts the three methods.

Relative Frequency Codebook

Because there are instabilities and drifts associated with independent local oscillators [84], *absolute* frequency is not a suitable choice in the real world. For example, 802.11g requires oscillator accuracy within 25ppm [27]. For signals at 2.4 GHz, this equates to a 60 KHz frequency offset. Later in our experiments, we use a current SDR platform [6] with an oscillator rated at 20 ppm. Thus, for practical applications, we propose that Alice utilize *relative* frequency. Given a MIMO scenario with M transmit antennas, Alice should transmit a baseline pilot using her first physical transmit antenna. She can then use the remaining $M-1$ antennas to transmit authentication messages using the relative frequency offsets from the baseline pilot transmission. Given N frequency quantization levels and M antennas, Alice can achieve b authentication bits per packet, where $b = (M-1)\log_2 N$.

We describe our Relative Frequency Codebook construction using an example.

Suppose Alice is equipped with $M = 2$ transmit antennas and uses $N = 8$ quantization levels over 80 KHz to provide CSI authentication. By frequency shifting the pilot for the second antenna, Alice can send 3-bit ($=\log_2 8$) authentication messages using 10 KHz frequency offset increments. For simplicity, consider that Alice uses a nominal carrier frequency, $\tilde{f}$, of 2.4 GHz, and encodes her messages by using a frequency offset that is proportional to the two's complement form of her message, m. Denoting the two's complement operator as †, Alice transmits her message bearing tone at

$$f_0 = \tilde{f} + (\Delta f)(m^{\dagger}),$$

where $\Delta f = 10$ KHz. Now envision that Alice wishes to send the following four authentication messages

$$\underline{m} = \{010,\ 111,\ 011,\ 001\}$$
$$(\underline{m}^{\dagger}) = \{2,\ -1,\ 3,\ 1\}$$

Alice can accomplish this over four packets by transmitting the pilots from the first antenna at $\tilde{f} = 2.4$ GHz, and the pilots from the second antenna at

$$\begin{aligned}\underline{f}_0 &= 2400 \text{ MHz} + \{20,\ -10,\ 30,\ 10\} \text{ KHz} \\ &= \{2400.02,\ 2399.99,\ 2400.03,\ 2400.01\} \text{ MHz}\end{aligned}$$

In the real world, local oscillator variability (at both Alice and Bob) often results in frequency drifts. Suppose that inter-packet local oscillator variability results in the following frequency drifts as seen by Bob.

$$\underline{f}_d = \{10,\ -10,\ 20,\ 30\} \text{ KHz}$$

Under such conditions, Bob receives the nominal, $\underline{\tilde{f}}_B$, and message bearing, $\underline{f}_{0,B}$,

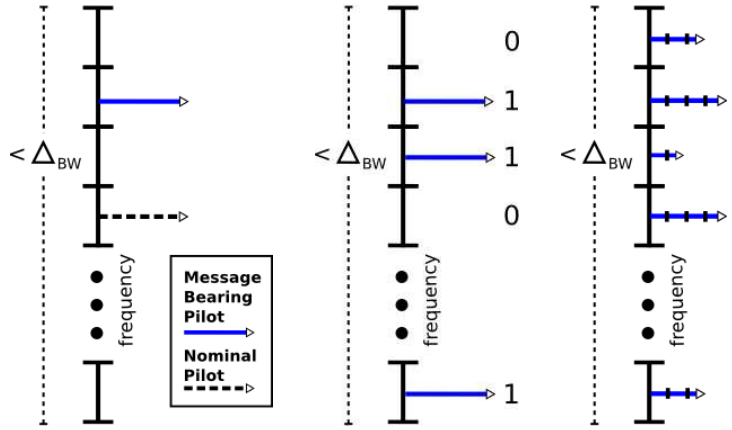

(a) RFC (b) BFC (c) FPC

Figure 4.4: Three Frequency Quantization schemes are depicted. The Relative Frequency Codebook (RFC) scheme in (a) utilizes relative frequency offsets from nominal to embed an authentication message. The Binary Frequency Codebook (BFC) scheme in (b) uses a binary codeword generated from multiple pilots. The Joint Frequency-Power Codebook (FPC) scheme in (c) transmits pilots at each sub-interval, using relative power levels to convey the authentication message.

pilots at

$$\begin{aligned}\underline{\tilde{f}}_B &= \underline{\tilde{f}} + \underline{f}_d \\ &= \{2400.01,\ 2399.99,\ 2400.02,\ 2400.03\}\ \text{MHz} \\ \underline{f}_{0,B} &= \underline{f}_0 + \underline{f}_d \\ &= \{2400.03,\ 2399.98,\ 2400.05,\ 2400.04\}\ \text{MHz}\end{aligned}$$

Bob decodes the authentication messages via

$$\begin{aligned}(\underline{\hat{m}}^\dagger) &= \|\frac{(\underline{f}_{0,B} - \underline{\tilde{f}}_B)}{\Delta_f}\| \\ &= \{2,\ -1,\ 3,\ 1\} \\ \therefore\ \underline{\hat{m}} &= \{010,\ 111,\ 011,\ 001\}\end{aligned}$$

Thus, despite oscillator instabilities, Bob recovers the authentication messages correctly. Figure 4.4 (a) illustrates the quantization of Δ_{BW} for the Relative Frequency Codebook authentication scheme.

As one would expect, authentication reliability grows with the number of receive antennas. The authors note that smart antenna techniques could be used by Bob (perhaps steered by previously authenticated CSI estimates) to increase the detection and resolvability of the authentication messages.

This authentication scheme offers error inference by physical layer pilot observation. For instance, channel noise simply increases the noise floor. Co-channel interference does the same, but in a sporadic manner since transmissions are asynchronous. For interfering devices, invalid authentication messages would be seen at unexpected times. While an adversary can mimic both channel noise and co-channel interference, we have noted that truly smart attacks will target the pilots in a synchronous fashion. But because Eve is not aware of the authentication message to send, she does not know where in frequency to transmit the rogue pilots.

A key point arises when we consider what happens if Eve is equipped with perfect knowledge of the the proper frequency intervals. Even with this seemingly great advantage, it is extremely probable that there are frequency offsets between Eve's pilots and Alice's pilots. These deviations are attributable to local oscillator differences. If this frequency error is resolvable, then Bob can detect Eve's presence. Later in the paper we will illustrate this principle in a real-world experiment.

The Relative Frequency Codebook authentication scheme is amenable to CSI estimation methods that are currently in use. For example, consider 802.11n where sequential transmissions from each antenna are used in the beginning of every packet to estimate the channel [76]. These pilots are called high throughput long training fields (HT-LTFs) and are illustrated in Figure 3.6 of Chapter 3.

Binary Frequency Codebook

We now consider a similar technique that increases the authentication bit-rate by further leveraging properties of the channel coherence bandwidth. As we have noted, a pilot will experience comparable channel effects if it is transmitted anywhere within the channel coherence bandwidth. Thus, if Alice transmits multiple pilots within a given channel coherence bandwidth using identical power, then either *all* or *none* of them will arrive at Bob. Alice may therefore send an authentication message by utilizing *all* of the N frequency quantization levels in a binary fashion (i.e. on/off). In this manner, she can send an N-bit authentication message per transmitter at a given frequency. The Binary Frequency Codebook scheme is illustrated in Figure 4.4 (b). Key to the success of this technique is pilot resolvability. When considering the use of channel sounding tones as pilots, protocol parameters such as the pilot length and frequency separation must be carefully selected so as to ensure frequency resolvability [106]. With this scheme, a single frequency reference is needed per packet, resulting in $b = (N-1)+(M-1)N$ authentication bits per packet.

Suppose that Alice wishes to use the Binary Frequency Codebook scheme to send the following 15-bit authentication message, $m = \{010110011100101\}$. By using the same $M = 2$ antennas and $N = 8$ quantization intervals from the RFC example, Alice can achieve this over a single packet transmission. If Alice uses the lowest frequency interval from the first antenna to send the reference pilot, then this leaves 7 message bits for the first antenna, and 8 for the second antenna. Thus, Alice transmits $\{10101100\}$ from antenna 1 and $\{11100101\}$ from antenna 2. Given a 1 is transmitted by the presence of a pilot, Alice transmits pilots at the following frequency intervals, where $\underline{f}_0$ indicates antenna 1 and $\underline{f}_1$ indicates

antenna 2:

$$\begin{aligned}
\underline{f}_0 &= 2400 \text{ MHz} + \\
&\{-40, -30, -20, -10, 0, 10, 20, 30\}. * \{1, 0, 1, 0, 1, 1, 0, 0\} \text{ KHz} \\
&= \{2399.96,\ 2399.98,\ 2400.00,\ 2400.01\} \text{ MHz} \\
\underline{f}_1 &= 2400 \text{ MHz} + \\
&\{-40, -30, -20, -10, 0, 10, 20, 30\}. * \{1, 1, 1, 0, 0, 1, 0, 1\} \text{ KHz} \\
&= \{2399.96, 2399.97,\ 2399.98,\ 2400.01,\ 2400.03\} \text{ MHz}
\end{aligned}$$

Error inference remains the same for this scheme as the Relative Frequency Codebook method, but with a major enhancement regarding smart adversarial detection. If Alice transmits pilots from a given antenna using identical power and within the same coherence bandwidth then the pilots will arrive at Bob with the same power. Hence, any pilots with different arriving power would immediately reveal an adversary's presence — note that it is virtually impossible for an adversary to match the power received at Bob. In Section 4.3.8, we illustrate this principle by way of a real-world experiment.

For MIMO operation, Bob will be equipped with multiple receive antennas. In such a scenario, smart adversary detection becomes more reliable, as each receive antenna will have independent power levels from each of Alice's transmissions in multipath-rich environments [77].

Joint Frequency-Power Codebook

Further leveraging properties of the coherence bandwidth, Alice can add *transmission power* to her arsenal for authentication. By transmitting relative power within Δ_{BW}, Alice increases the authentication bit rate by reducing her adversarial detection effectiveness. Under this scheme, Alice should transmit pilots at *every* frequency quantization level, but using Q power quantization levels. Because Alice is now using relative power, a pilot reference is needed for each

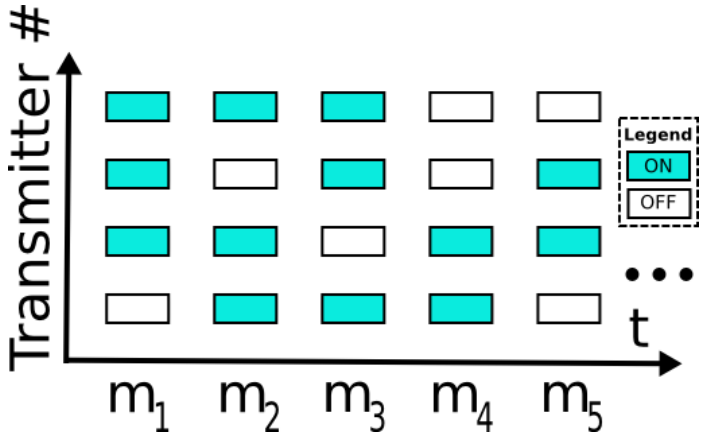

Figure 4.5: Under Selective Usage, Alice keeps a selection of transmitters idle during the channel sounding process. For packet i, Bob receives authentication message m_i. In the example above, if Alice's transmitter is active to send a 1, then Bob receives $m_1 = [0\ 1\ 1\ 1]$ and $m_2 = [1\ 1\ 0\ 1]$.

antenna. Thus, Alice can achieve $b = (N-1)\ \log_2(Q)$ authentication bits per transmitter.

Suppose that Alice wishes to use the Join Frequency-Power Codebook scheme to send the following 14-bit authentication message, $m = \{01011001110010\}$. Using the same $N = 8$ quantization intervals from the previous examples, we note that Alice can send the 14-bit message with a single transmit antenna by using $Q = 4$ power quantization intervals. Alice uses transmit powers of {0, -3, -6, -9} dBm to convey bits of {00, 01, 10, 11}, respectively, and transmits her reference pilot at 0 dBm again at the lowest frequency quantization interval. To send the message, Alice maps {01, 01, 10, 01, 11, 00, 10} onto the 7 message containing pilot frequencies, and therefore transmits pilots at {2399.96, 2399.97, 2399.98, 2399.99, 2400.00, 2400.01, 2400.02, 2400.03} MHz using power levels of {0, -3, -3, -6, -3, -9, 0, -6} dBm, respectively.

For the Joint Frequency-Power Codebook scheme, the CSI estimates again remain valid, however the detection of adversarial activity may be tougher to notice in the power domain as an adversary now has an increased range of power to match. The Joint Frequency-Power Codebook scheme is depicted in 4.4 (c).

4.3.6 Selective Usage

Some real world devices will not be able to implement any of the Frequency Quantization schemes due to transmitter limitations or OFDM implementations. In OFDM, carrier frequencies are selected carefully so as to minimize inter-carrier interference. Thus, we propose a lower bit-rate alternative that we refer to as Selective Usage. Without altering the actual frequency of the pilots, Alice may choose to *selectively* omit the pilot for a given transmit antenna. While CSI information cannot be estimated by Bob for the omitted antenna, he may utilize the last CSI estimate or an appropriate estimate based upon previous measurements [107]. Additionally, Bob can always use the data itself to perform timing, frequency, and phase recovery when the pilot is absent. While this is less efficient from a processing perspective, it does allow for CSI authentication without loss of accuracy due to CSI estimate interpolation. Selective omission of pilots will effectively transmit the authentication message to Bob in a binary fashion using each nominal pilot frequency. For instance, Alice can send a 1 to Bob by transmitting a pilot, and a 0 by remaining idle. Given M transmit antennas, Alice can send b authenticated bits per packet, where $b = \log_2 M$. Under the Selective Usage scheme, transmissions from antennas that should remain silent would indicate smart adversarial activity. The Selective Usage procedure is illustrated in Figure 4.5.

4.3.7 Extensions

In the previous sections, we assumed multiple antennas at a single frequency. We now explore extensions to our schemes that enhance the authentication bit rate further and/or deal with system constraints.

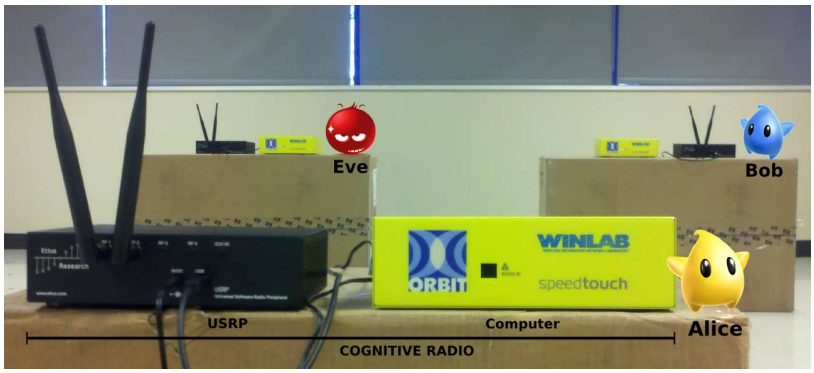

Figure 4.6: The experimental setup is shown. Alice, Bob, and Eve are comprised of cognitive platforms made up of a computer and a USRP.

Multiple Frequency Extensions

The extension to multiple frequency usage is straightforward. By using K carrier frequencies, the Frequency Quantization schemes can all achieve increased authentication bit rates. For the Relative Frequency Codebook (RFC) scheme, only a single reference is needed per packet. Thus, the authentication bit rate becomes

$$b_{RFC} = (M-1)\log_2 N + (K-1)M\log_2 N.$$

The Binary Frequency Codebook (BFC) scheme also only requires a single reference per packet, resulting in an authentication bit rate of

$$b_{BFC} = KMN - 1.$$

For the Joint Frequency-Power Codebook (FPC) scheme, a reference is needed at *each* carrier, therefore achieving an authentication bit rate of

$$b_{FPC} = KM(N-1)\log_2 Q.$$

For the Selective Usage (SU) scenario, the authentication bit rate per packet is similar to the BFC scheme in that each carrier operates independently, thus resulting in

$$b_{SU} = K\log_2 M.$$

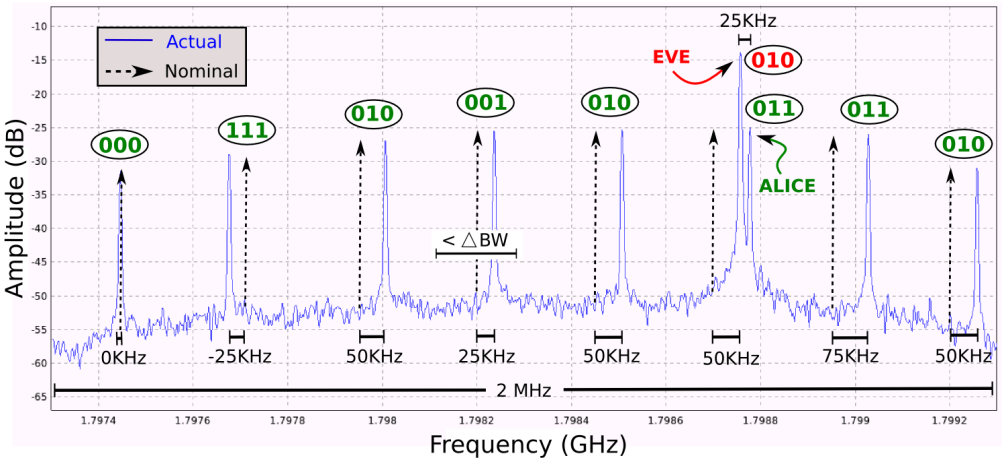

Figure 4.7: The Relative Frequency Codebook (RFC) scheme was implemented in a real-world channel estimate authentication experiment. Eve the adversary was present in order to inject false channel sounding pilots. Above, Alice transmits the authentication message {000, 111, 010, 001, 010, 011, 011, 010}, and Bob receives {000, 111, 010, 001, 010, 01$\underline{0}$, 011, 010}. The error (01$\underline{0}$) in the authentication message is due to Eve's inability to closely match Alice's pilot frequency and reveals Eve's presence.

Note that these multiple frequency extensions hold for the single antenna scenario, where the authentication bit rate can be calculated by using $M = 1$ in the above equations.

OFDM Extensions

Popular communication standards such as 802.11n and WiMAX utilize OFDM due to its ease of implementation and spectral robustness with regard to fading and inter-carrier interference suppression. Such implementations are typically Fast Fourier Transform (FFT) based, therefore restricting the use of user defined frequencies as called for in the Frequency Quantization scheme. However, slight modifications result in successful OFDM usage of Frequency Quantization CSI authentication. By incorporating a larger FFT, OFDM systems effectively can utilize more relative frequencies. In essence, the size of the FFT will govern the quantization levels per CSI estimate. As an example, 802.11 uses a 64-point FFT

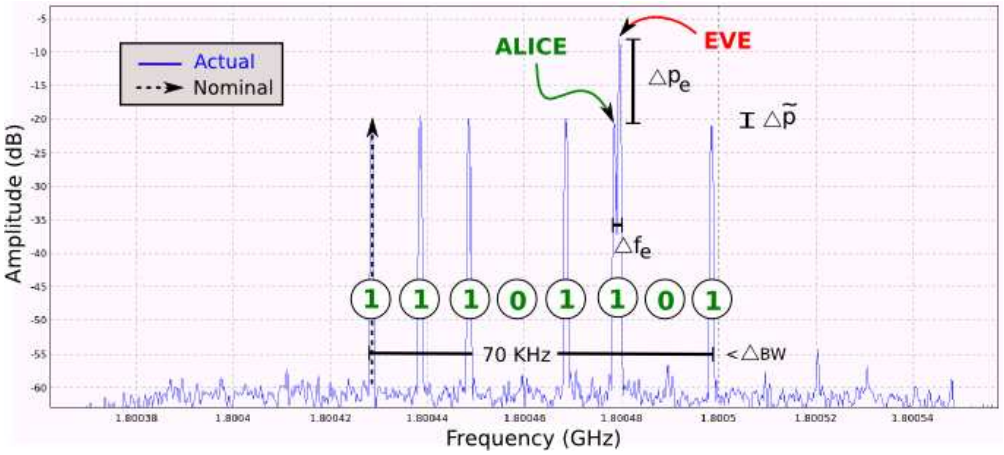

Figure 4.8: The Binary Frequency Codebook (BFC) scheme was implemented in a real-world channel estimate authentication experiment. Eve the adversary was present in order to inject false channel sounding pilots. Above, Alice transmits m ={11101101}, and Bob receives the correct authentication message because Eve is operating close enough to Alice's active pilot (i.e. Δf_e is not resolvable). But because the power deviation at the sixth pilot is much greater than the nominal power deviation expected ($\Delta p_e >> \Delta \tilde{p}$), Eve's presence is revealed.

during channel estimation [27]. By increasing the FFT to 256 points, Alice can achieve 4 authentication bins per pilot. With this minor modification, the data transmission may resume the legacy FFT size thus maintaining efficient use of the transmission device. Note that WiMAX implementations already utilize a 256-point FFT (with 8 carriers set aside as pilots) [69]. Further, carriers are often left idle during pilot transmission. In 802.11, the training sequence (TS) portion of the pilot only modulates on every 4^{th} carrier [27]. By altering this carrier allocation, Alice can incorporate CSI authentication. The astute reader may question the frequency offset incurred by shifting the pilot, however since Bob knows the message that he expects to receive, the offset is in fact correctable. If the Frequency Quantization schemes are too problematic to implement, then the OFDM systems may always revert to Selective Usage CSI authentication.

Single Frequency SISO Applications

While the schemes that we presented are geared towards multi-carrier MIMO operation, they can also be applied to single frequency SISO systems that utilize channel state information. For single frequency SISO application, CSI authentication messages can be sent either inter-packet (given a large enough channel coherence time) or by *shifting* the pilots in frequency within the transmission. Further extensions and their ramifications are left to the reader for consideration.

4.3.8 Experimental Validation

To illustrate the feasibility of our techniques, we conducted real-world experiments using a current cognitive radio architecture — the GNU Radio/USRP SDR platform that we describe in detail in Section 1.5 [6, 7]. Each player in the experiments that follow is a GNU Radio/USRP cognitive platform. The experimental setup is pictured in Figure 4.6.

Relative Frequency Codebook Experiment

In the Relative Frequency Codebook experiment, Alice uses two transmit antennas to send 8 authentication messages (by using 8 carriers) across 2 MHz in the 1800 MHz band. The carriers are each separated by 200 KHz. At a given carrier, Alice embeds an authentication message by transmitting the pilot at a frequency that is offset from nominal. She uses the same pilot message mapping function discussed in the example of Section 4.3.5, but with a frequency quantization interval (Δf) of 25 KHz. Each authentication message conveys 3-bits, hence valid frequency offsets are $\{-100, -75, -50, -25, 0, 25, 50, 75\}$ KHz, which equate to $\{100, 101, 110, 111, 000, 001, 010, 011\}$ in bits. In the experiment, Alice transmits an authentication bit sequence of $\{000, 111, 010, 001, 010, 011, 011, 010\}$, resulting in frequency offsets of $\{0, -25, 50, 25, 50, 75, 75, 50\}$ KHz. Figure 4.7 depicts the

spectrum received by Bob due to transmissions from Alice's second antenna during pilot activity. The nominal pilot frequencies (as transmitted by Alice's first antenna) are depicted with the dashed lines.

Also present in the experiment is an adversary, Eve, whose goal is to attack the CSI. For illustrative purposes, Eve is only active on a single carrier. In the experiment, ample jammer power allows Eve to trick Bob into decoding the wrong CSI pilot at the sixth carrier. But because Eve does not know the authentication message, she cannot predict exactly where to transmit the rogue CSI pilot tone. Since the rogue pilot is offset by -25 KHz from Alice's pilot, Bob decodes $\{000, 111, 010, 001, 010, 01\underline{0}, 011, 010\}$, where $\underline{0}$ indicates the bit received in-error. And since the authentication message is incorrect, Eve's presence is revealed.

Note that if Alice is only equipped with a single transmit antenna, she can still use the Relative Frequency Codebook methodology by using one of the pilot tones as the frequency reference. All other nominal frequencies can then be calculated from the reference. The authors state this to portray practical adaptability to SISO applications.

Binary Frequency Codebook Experiment

In the Binary Frequency Codebook experiment, we isolate activity at a single carrier frequency in order to focus on a particular authentication message. Using a single antenna, Alice transmits an 8-bit authentication message, $m = \{11101101\}$, in the 1800 MHz band over 80 KHz ($< \Delta_{BW}$). Figure 4.8 shows the spectrum as seen by Bob. Again, Eve is present, and solely attacks the sixth bit in the authentication message. Note that if Eve attacks an inactive bit, then the authentication message will fail and her presence will be revealed. However, if she is close enough in frequency to an active bit (i.e. Δf_e is not resolvable), then the authentication message will be correct despite an inaccurate CSI estimate. Such is the case for this experiment, but because Alice and Bob are using the Binary

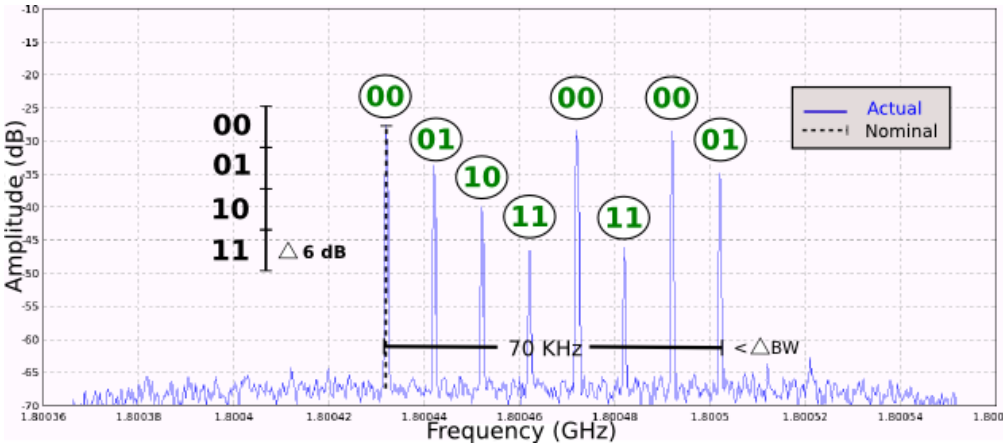

Figure 4.9: A real-world experiment was conducted using the Joint Frequency-Power Codebook scheme, where Alice sends an authentication message of m ={01,10,11,00,11,00,01}. Alice transmits a nominal pilot at the lowest frequency, and uses the remaining 7 frequencies to send messages by backing off transmission power. Using backoffs of {0,6,12,18} dB, Alice sends authentication bits of {00,01,10,11}.

Frequency Codebook Authentication scheme, Bob can still detect Eve's presence. Because the transmissions all occur within the coherence bandwidth of the channel, power deviations from nominal are a clear indicator of an attacker. We see that Eve's power deviation from Alice's is quite large ($\Delta p_e \approx 10\ dB >> \Delta\tilde{p}$).

Joint Frequency-Power Codebook Experiment

In the Joint Frequency-Power Codebook experiment, we isolate activity at a single carrier frequency again in order to focus in upon a specific authentication message. Using a single antenna, Alice transmits a 14-bit authentication message, m = {00101100110001}, in the 1800 MHz band over 80 KHz. Like the Binary Frequency Codebook scheme, it is essential that the authentication message be sent within the coherence bandwidth of the channel (i.e. $< \Delta_{BW}$). Using this scheme, a power reference is necessary for every message carrier. Alice chooses to use the lowest frequency carrier to transmit the nominal reference. In order to send authentication bits {00,01,10,11}, Alice uses relative power attenuations

of {0,6,12,18} dB from nominal. Therefore, to send the desired 14-bit authentication message, Alice transmits {00,01,10,11,00,11,00,01} by transmitting her pilot carriers at levels of {0, -6, -12, -18, 0, -18, 0, -6} dBm. Figure 4.9 shows the spectrum as seen by Bob for this authentication message. As can be seen, the message is fully recovered by analyzing the relative powers of the received pilot signals.

4.3.9 Channel Estimate Protection Summary

In this Section, we have provided physical layer approaches to authenticating channel state information estimates. Methodology has been proposed to support existing protocols such as 802.11n in addition to new cognitive protocols. Experimental verification was performed using USRP/GNU Radio SDR platforms.

Having proposed techniques to provide channel estimate authentication, we now proceed to address communication challenges associated with multi-path rich urban environments with the presence of a very strong adversary. Again, we focus on PHY layer techniques that can be utilized by emerging cognitive platforms.

4.4 Radio Teaming

Although theoretical results may suggest that it is possible to establish a reliable wireless link between any two parties, e.g. by increasing transmission power or adjusting the modulation format to operate at a lower data rate[108], there are practical limits to what is possible through such methods. In particular, commodity radios, such as an 802.11 radio, are limited in the types of modulation mechanisms (e.g. in 802.11b the lowest data rate is 1 Mbps, which corresponds to employing Direct Sequence Spread Spectrum (DSSS) modulation) that can be used or are constrained in terms of the amount of power they may employ (e.g. a 802.11 WLAN card can only transmit with a power level up to 100mW). These

device limitations, however, correspond to actual limits on the ability to establish communication links. For example, if we fix a modulation scheme and limit the amount of power, then there is a corresponding distance to which communications can reliably be established.

A consequence of this observation is that it may be impossible to establish reliable communication between devices that are beyond the communication range supported by transmitting at the lowest data rate with the highest allowed transmit power. Such a problem might arise in various realistic operational scenarios, such as the well-known urban canyon or in the presence of radio interference. In an urban canyon setting, the existence of a complicated multipath environment (e.g. such as midtown Manhattan, with its streets intertwined amongst tall buildings) can sufficiently degrade communications in spite of entities being in moderately close proximity of each other. On the other hand, in an interference setting, the existence of an interferer near the recipient could lead to a resulting signal-to-interference level that prevents successful demodulation of communication. In these challenging settings, given the limitations of commodity devices, a natural question that arises is whether it is still possible to establish communications between a sender and a receiver.

In this Section, we overcome the limits of a single commodity device through a form of cooperative communication that is amenable to commodity devices. Specifically, we propose that a device forms a *radio team* of similar devices within range of each other, and that this team acts synergistically to communicate with the receiver. Whereas the communication literature that has examined cooperative communication [109, 110, 111, 112, 113] requires stringent synchronization of devices at their physical layer, our approach works as an overlay on an existing wireless link-layer. Specifically, the radios work together to establish a common time line and to exchange a schedule to coordinate their synchronous transmissions. The resulting cacophony of simultaneous transmissions will arrive at the

receiver incoherently, producing perturbations in the receiver beam response pattern, that can be used to establish a reliable communication channel between the team and the receiver.

4.4.1 Related Work

Much research has been done in the area of cooperative communications [109, 110, 111, 112, 113]. However, the bulk of the focus has been on traditional beamforming and MIMO techniques that require a very high level of synchronization and intelligence amongst participating parties. For instance, standard beamforming requires that all transmissions arrive at the same time *and* in-phase with each other at the receiver. Our concept removes this level of synchronization, thereby making it an adoptable and reliable communication overlay for existing commodity wireless devices. We now overview some other non-traditional techniques.

In [114], the collision that occurs when two packets are transmitted simultaneously by two senders is actually used to the network's advantage. The signal after the collision is the sum of the colliding signals along with attenuation, phase and time shifts. If the receiver knows one of the packets, it can cancel the known packet's signal, and then proceed to decode the other packet. The receiver carries out a series of computations to calculate the channel attenuation and phase shift to decode the packets. Although they assume no synchronization between the transmitted signals, they assume that the receiver has knowledge of one of the packets. In our approach, we do not assume a priori knowledge of the packet. Rather, the receiver looks at perturbations in the beam-pattern response to decode the signal. In addition, [114] provides capacity benefit for high SNR operating regimes, while our approach is geared specifically towards low SNR/SINR scenarios.

In [115], a single antenna is used to achieve some of the benefits of a MIMO system. Along with transmitting their own messages, the radio nodes relay each

other's messages in a time slotted fashion. Hence, a form of spatial diversity is generated, and the receiver sees independently faded versions of the signal. This approach requires very high levels of synchronization and can be tough to achieve in the real world. Also, coordinating the cooperation is a challenging task as it involves a complex partner assignment scheme that is needed to achieve the desired data path diversity. In our approach, all the radio nodes in the team transmit the same message at roughly the same time.

Network MIMO is another diversity scheme that allows multiple concurrent transmissions using space-time coding techniques. It requires multiple transmit antennas at proper distances from each other to increase diversity in order to combat channel fading [109]. By sending signals that carry the same information through different paths, multiple independently faded replicas of the data is obtained at the receiver and hence, more reliable reception is achieved. The protocol that we propose is quite similar to MIMO, however it greatly reduces the level of synchronization needed. In addition, the location of each radio teaming transmitter is not restricted in any manner.

Another area of related work is anomaly detection (or deviation detection). In anomaly detection, some objects have attribute values that differ significantly from the expected or typical attribute values. The underlying cause of an anomaly provides important information since objects may be of different types or belong to different classes. There have been many approaches to anomaly detection including *model based techniques*, *proximity based techniques* and *density based techniques* [116]. In [117], a statistical signal processing technique is used to detect network anomalies. In our protocol, we seek to communicate with the receiver by causing variations in its beam-pattern response such that it deviates from its norm. In our experiments, we use a statistical distance-based technique to detect the anomaly. Our approach can also be compared to significance testing, where a hypothetical result is considered false if the observed result differs sufficiently

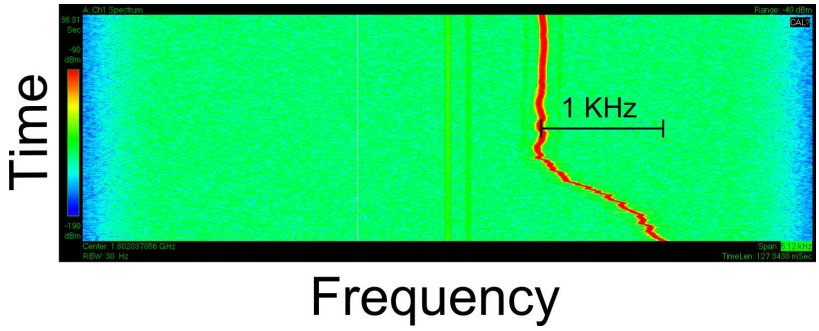

Figure 4.10: Local oscillator instability due to small temperature variation is demonstrated using a current Software Define Radio (SDR) architecture. Slight temperature variations produce a frequency deviation of over 1 KHz in the 1800 MHz band. The SDR's ovenized crystal oscillator (OCXO) is rated at 20 ppm.

from the hypothetical result. In our approach, we find the difference between the observed result and a base result, and then decode the signal based on the difference.

Finally, we also draw a comparison to CDMA techniques. In CDMA, a rake receiver is used to recombine multipath components emanating from a single transmitter [118]. While the rake receiver leverages the different time-delays associated with the multipaths, we look to take advantage of the different angle-of-arrivals (AOAs) associated with each multipath component at the intended receiver.

4.4.2 Strategy Overview

Let us now elaborate on the motivation behind radio teaming. It is clear that an urban canyon presents a difficult environment for reliable wireless communications. While cooperative communications may represent the most desirable solution, standard forms of beamforming and MIMO can be challenging to implement in certain scenarios. Consider the example where the "radio team" is a group of spatially diverse, independent transmitters, none of whom have a reliable

link to their intended receiver. It is perfectly feasible that units may synchronize finely in time based upon a shared clock (e.g. GPS). But, in order to properly beam form the units, the transmitters would need exact knowledge of their environment or user-specific feedback from the receiver so that each transmitter could properly coordinate its signal's arrival at the receiver. Otherwise, the signals will arrive with random phases despite being perfectly time-slotted. Furthermore, the random jitter and drifts associated with the fact that each unit has an independent local oscillator (and most likely a low-cost one) greatly reduces the chance of any meaningful analog gain associated with constructive interference. To illustrate this phenomenon, Figure 4.10 depicts a frequency drift of over 1 KHz that was incurred due to a temperature variation of only a few degrees Fahrenheit at the transmitter's local oscillator. The transmitter used was a USRP SDR platform with a RFX-1800 daughterboard serving as the RF frontend. The transmitter/receiver in this experiment is indicative of what is used in commodity communication platforms. Since heterogeneous units are not readily capable of synchronizing to the level needed to support MIMO or beamforming, is there something else that they can do to work cooperatively to overcome a bad link?

Overcoming the environment

The first motivating scenario that we will consider is one where our communications are hampered solely by the environment (i.e. no adversary is present). The major reason for the inability of the transmitters to reach the receiver in this scenario is the fact that the environment is multipath rich, resulting in significant attenuation and degradation of communication versus distance. Consequently, normal modulation methods may be unlikely to successfully decode at the receiver. A straight-forward approach to overcoming this problem might be to have a group of transmitters each *try* individually to communicate and hope that at least one of the individual links turns out to be a good one.

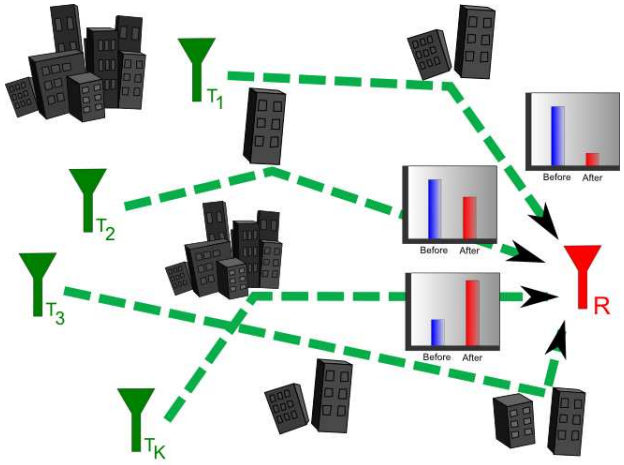

Figure 4.11: Multipath exploitation in the urban canyon is illustrated. Signals emanating from different transmitters travel along different paths to a common receiver. Signal statistics differ as a function of arriving look-angle at the receiver.

However, such a strategy would not work if the multipath environment is sufficiently complex. However, in radio teaming, the strategy is to create a team with multiple transmitters at different locations, and properly coordinate how communications are transmitted by the team and decoded by the intended receiver so that we increase our probability of reaching the desired receiver with our message. To accomplish this, we need to use the multipath nature of the environment to our advantage. Under the wide-sense stationary-uncorrelated scatterer (WSSUS) model [77], we can assume independence of arriving signals not only in time and phase, but also in *angle*. The key to our proposed idea is to utilize the angle-of-arrival (AOA) as an added degree of freedom in the multi path-rich environment. In this manner, rather than relying on one solid multipath link from the radio team making it to the receiver, we leverage the fact that many multipath components arrive at different AOAs. The key benefit here is that each individual link does *not* need to be of high enough quality to support conventional communication. The cumulative effect of the radio team will simply *distort* the beam-pattern at the receiver in a manner that deviates from the ambient background received

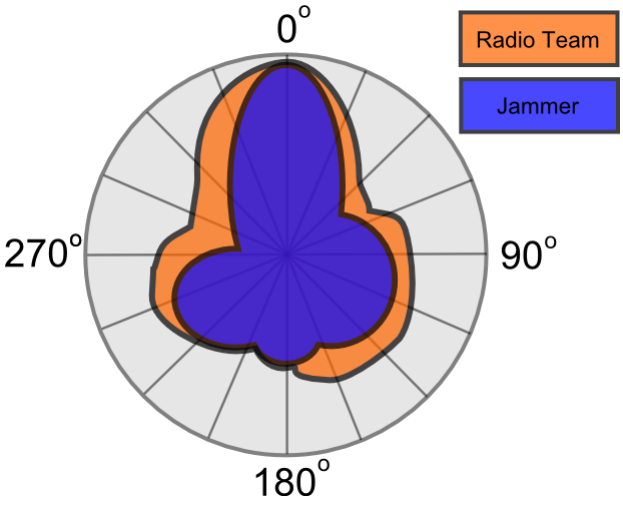

Figure 4.12: Beam pattern distortion is illustrated. Jammer activity is distorted by the team of radio transmitters. The extra energy arriving at angles where jammer energy is low provides detectable distortion for the receiver.

beam-pattern. By appropriately coordinating the teams' transmissions, they will be able to convey the message by modulating the beam-pattern, alternating from ambient background beam-pattern to a distorted beam-pattern response. Figures 4.11 and 4.12 illustrate the idea. Each transmitter $(T_1, T_2, ..., T_K)$ in Figure 4.11 has its own set of multipath links to the receiver, R, resulting in various different angles of arrival. The cumulative effect of the disparate arrival angles results in the distortion seen in Figure 4.12. Two beam patterns are shown. One is indicative of the normal ambient background beam-pattern response, while the other illustrates the *added* distortion based upon radio teaming activity.

We reiterate that the the goal of the radio teaming protocol is to merely distort the beam-pattern at the receiver *enough* to be perceived at the receiver as a distortion, and use this effect to convey the team's message.

Overcoming interference

The second scenario that we consider is the case where there is an adversary present in the multipath-rich environment. A key point to note here is that we

must assume that the adversary is *not* in control of the multipath, and so it cannot control its paths to the intended receiver. For our discussion, we shall assume that the adversary is emitting a constant, background signal with the objective of disrupting all communications to the receiver. Consequently, the multipath components corresponding to the adversary's signal will be constrained in terms of AOAs because it is physically at a single location. We do not consider more complex scenarios where the adversary conducts adaptive jamming games, and note that this is a topic of ongoing investigation.

In the radio teaming problem with the constantly emitting jammer in the background, the radio team geographic dispersal allows for some of its many multipath signals to arrive on angles other than the angles that the jammer's signal arrives at. As a result, the radio teaming transmitters will therefore "waterfill" the jammed beam-pattern at the receiver. Just as before, radio teaming leverages the multipath (usually considered a detrimental effect) to the advantage of communications.

4.4.3 Protocol

We now present the radio teaming protocol in a detailed manner. Assume transmitter T_1 wants to send a message to receiver R using a conventional communication method (e.g. BPSK). Now suppose that the multipath effects of the channel and/or the presence of a jammer prevent their communication - either the SNR or the SINR is too low. At this juncture, an alternative communication solution is needed.

The protocol that we now describe makes two main assumptions: (1) the receiver has a directional antenna and is capable of taking measurements over various look-angles, such an antenna maybe electronically steerable, (2) the transmitter has a team within a communicable distance that is willing to help convey its message.

When T_1 realizes that it cannot reliably communicate with R, for e.g. due to the failure to receive an acknowledgement from R, it will ask for help from its radio team. Coarse timing synchronization will be carried out amongst the team, where timeslot boundaries will be coordinated. The team will then decide upon the signaling specifics (modulation, frequency, symbol mapping, and symbol timing). Finally, T_1 will tell its team the intended message and the time at which to start transmission.

Meanwhile, receiver R initiates the radio teaming reception. This might occur for e.g. on a time out condition at the receiver which indicates that the receiver needs to use teaming in order to receive communication. The receiver R will utilize its directional antenna to monitor the beam-pattern response over N look-angles. Using a baseline beam-pattern derived from non-radio teaming activity, the receiver looks for beam-pattern deviation to perform its symbol recovery. Note that the receiver may perform deviation detection using a variety of methods (e.g. statistical analysis, clustering algorithms, etc.). Figure 4.12 illustrates a general form of beam-pattern distortion.

It is important to note at this juncture that the radio teaming protocol is focused upon sending reliable data from the radio team to the receiver, and not vice-versa. In this work, we only analyze the protocol involved in getting data to the receiver in situations where conventional communication is not reliable. We assume that the receiver is more capable than any individual in the radio team, and may in fact have the option of resolutions such as increasing its signal power. We justify this assumption due to the increased capabilities already inherent to the receiver (i.e. directional antenna capable of "rotational" operation).

Simple binary communications

Now consider a simple binary communications example. Radio T_1 first synchronizes with its team using Network Time Protocol (NTP) in order to achieve

Algorithm 3: Radio Teaming for Simple Binary Communication using OOK.

```
/****** Transmitter Processing ******/
/*** Establish Timeslots ***/
NTPSynchronization( );
/*** Coordinate message and start time ***/
(message, start time, frequency) = MessageCoordination( );
/*** Tune to appropriate frequency ***/
TuneTransmitter(frequency);
/*** Transmit the message ***/
if current time is greater than start time then
    for (each new symbol in message) do
        if symbol is 1 then
            Transmit;
        end
        if symbol is 0 then
            Remain Idle;
        end
    end
end
/****** Receiver Processing ******/
/*** Compute baseline feature vector ***/
baseline = ComputeBaseline( );
/*** Decode the message ***/
for (each new symbol k) do
    /*** Collect data for each of the N look-angles ***/
    for (each collection angle n in N) do
        metric[n] = power_statistics;
    end
    /*** Perform symbol decoding ***/
    symbol[k] = decision(metric, baseline);
end
```

time-slotting accuracy on the order of milliseconds [119, 120]. The team then decides to use a particular carrier frequency. In our experiments we use 2402 MHz. On-off keying (OOK) is subsequently selected as the modulation, with a symbol period of 1 second. The symbol, 1, will be conveyed by all the team members transmitting and the symbol, 0, will be conveyed by all of them remaining idle. Finally, radio T_1 informs the team to send the message, m=10110, and to start transmitting the message at timeslot M. Figure 4.13 illustrates this process and its anticipated effects at the receiver. We note that the signaling of the radio

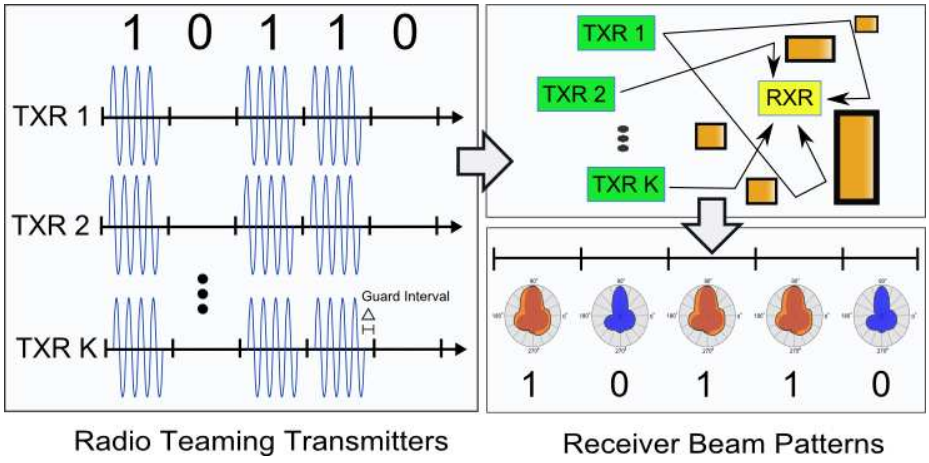

Figure 4.13: An overview of the radio teaming protocol is provided. The radio team leverages the variability in the arriving angles to distort the beam pattern at the receiver. By *modulating* this distortion, the team sends a message to the receiver In the example above, K transmitters send an emergency message of 10110 to the receiver despite the presence of a strong jammer.

team should incorporate proper guard time so as to avoid inter-symbol interference (ISI) at the receiver. To elaborate, the guard time should be sufficiently long to account for synchronization accuracy and the delay spread of the channel.

Now it is up to the receiver to decode the symbols by monitoring the beam pattern. The baseline beam pattern is defined by a feature vector that is created using power statistics gathered over each of the N look-angles during non-radio teaming activity. Let receiver R use the mean and standard deviation to define its baseline feature vector.

$$\mathbf{y} = (\mu_1, \mu_2, ..., \mu_N, \sigma_1, \sigma_2, ..., \sigma_N)$$

For every revolution of the directional antenna, a new feature vector is calculated. For the k^{th} revolution,

$$\mathbf{x}^k = (\mu_1^k, \mu_2^k, ..., \mu_N^k, \sigma_1^k, \sigma_2^k, ..., \sigma_N^k)$$

The Euclidian distance is then used to calculate the difference between the two

feature vectors. It is given by

$$d^k = ||\mathbf{x}^k - \mathbf{y}||^2$$

Finally, symbols are decoded at the receiver using

$$s^k = \begin{cases} 1, & \mathtt{d^k} >= \tau \\ 0, & \mathtt{d^k} < \tau \end{cases}$$

where s^k is the symbol for the k^{th} revolution and τ is the symbol decision threshold. Under this protocol, receiver R is able to reliably decode the message being sent by T_1 with the help of its radio team. The processing flow for this binary schema is further described in Algorithm 3.

Advanced communications

While the previous binary example may suffer from a low-bit rate (1 bps), there exist many immediate extensions to bolster both bit-rate and reliability. One example would be an OFDM extension, where the radio team selects M frequencies as OFDM carriers. Each carrier then modulates an OOK symbol on each carrier. In this manner, the capacity of the radio team garners a 2^M improvement over the single carrier OOK case given the same symbol timing. It is important here to select the carriers such that LO drifts associated with multiple transmitters do not overlap in frequency.

If the goal of the radio team is reliability rather than capacity, the team may take a spread spectrum approach. Each team member may utilize direct sequence spread spectrum (DSSS) by using a common chip-sequence on each symbol. This scenario further leverages the multipath of the environment. And, since the radio teaming methodology aims to work at low SINR and SNR, the spreading of the transmission power has a manageable effect on the performance of the system.

While the above cases utilize frequency as an added degree of freedom, we now discuss further ways to leverage time diversity in conjunction with the radio teaming methodology. It is clear that the receiver needs to perform its beam-pattern observations over a given interval, which in the previous cases defined the symbol period. However, borrowing from pulse-width modulation (PWM) techniques, we can successfully transmit data from the radio team by distorting the environment for time-lengths that are multiples of the receiver rotational period. In this manner differential and relative schemes can be utilized to achieve more methods of reliable, covert signaling. In [83], similar extensions are made to salvage a communications link in the presence of an adversary.

It is important to note at this point that advanced equipment such as Software Defined Radios (SDR), particularly at the receiver, provide a myriad of options in terms of exploiting the radio teaming methodology. Consider an SDR with a phased antenna array acting as a receiver. In this scenario, switching through N look angles can be performed digitally and therefore in a very fast and efficient manner.

4.4.4 Simulations

Before discussing the details of the simulations, let us take a moment to justify the radio-teaming methodology. We are investigating a scenario where members of a radio team lie within a communicable distance of one another, and various radio teaming transmitters may have a path— albeit a degraded or indirect one — to the receiver. A natural question that one might ask is: if one of the radios in the team has a path to the receiver, why can't that particular radio just relay the appropriate messages alone? The answer is because in this situation the SNR/SINR is too low to support conventional modulation schemes. To understand this, let us consider BPSK since it represents the conventional communication method with the lowest data-rate and greatest reliability. Figure 4.14 illustrates bit-error rate

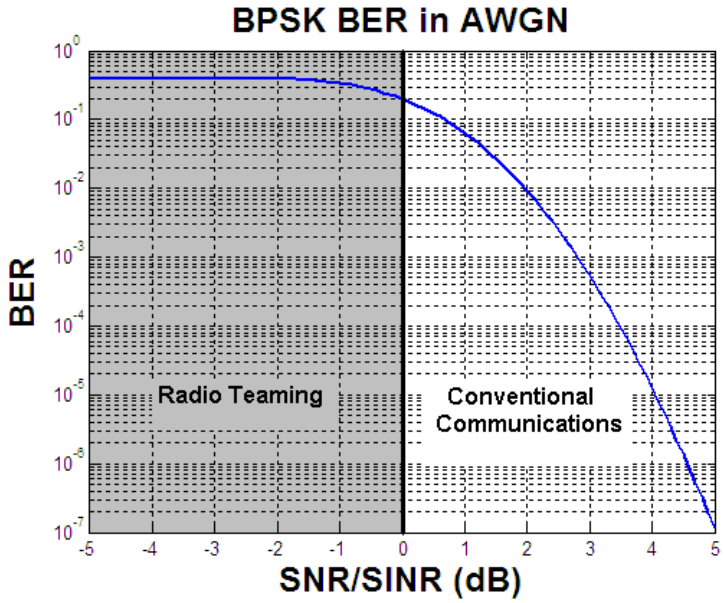

Figure 4.14: Bit-error rate curve for conventional BPSK communication in AWGN. Radio teaming is useful when conventional communications schemes break down.

versus SNR for conventional BPSK communication in the presence of additive white Gaussian noise (AWGN) [121]. In SNR/SINR regimes where we propose to use radio teaming (below 0dB), conventional communication methods break down. In fact, in signaling regions where we will demonstrate radio teaming, the BPSK is basically equivalent to the flipping of a coin (i.e. the BER is about 0.5).

Another question one might have is: what advantage is there in having a radio *team* as opposed to using a single radio transmitter on its own to disrupt the ambient beam-pattern response at the receiver? As we will see, by using multiple transmitters that concurrently transmit, there will be many independent multipaths arriving at different angles at the receiver, and hence using more transmitters will increase the probability of producing a beam-pattern distortion at the receiver that is sufficiently different from the ambient background levels, therefore increasing detectability. We now share some results from our simulations to illustrate this principle.

First let us describe the simulation setup. In our simulations, we consider a radio team consisting of K separate transmitters that are using the OOK version of our radio teaming protocol. We assume that the multipath environment can be modeled using a wide-sense stationary uniform scattering model [77], which is an accepted model for multipath-rich environments, such as in urban canyons or indoor environments. We choose OOK for various reasons — (1) it is easy to implement (which supports our real-world experiments in Section 4.4.5), (2) it provides a good mechanism to induce beam-pattern distortion, and (3) it is easy to model, simulate, and analyze. Extensions to other forms of radio teaming are straight-forward.

To add to the realism of the simulation, we conducted an experiment where we gathered data to act as a trace in support of our simulations. Our experiment involved a continuous-wave (CW) jammer. The data trace was collected in the WINLAB ORBIT Grid [122], where the jammer was created using a standard signal generator. The data collects were taken using a vector signal analyzer (VSA) in conjunction with a MAXRAD 18 dBi directional panel antenna over eight uniform receiver look-angles. The simulation incorporates H arriving signals from the radio team, where the SINR is uniformly distributed between -20 and -10 dB as seen at the receiver. Various other levels of control are also parametrized in the simulation. For instance, the angle-of-arrival (AOA) of an incoming signal is uniformly distributed between the eight angles corresponding to the jammer collects. The simulation also allows for variability in the phases of the arriving signals, in addition to frequency offsets.

Now that we have described the simulation, let us introduce our first simulation where we wish to investigate the effect of the size of a radio team. Specifically, we analyzed several cases where we had H=$\{1,4,8,12\}$ arriving signals arrive at the receiver at the previously stated SINRs. Figure 4.15 shows the beam-pattern response as we progressively increased the number of radio teaming transmitters.

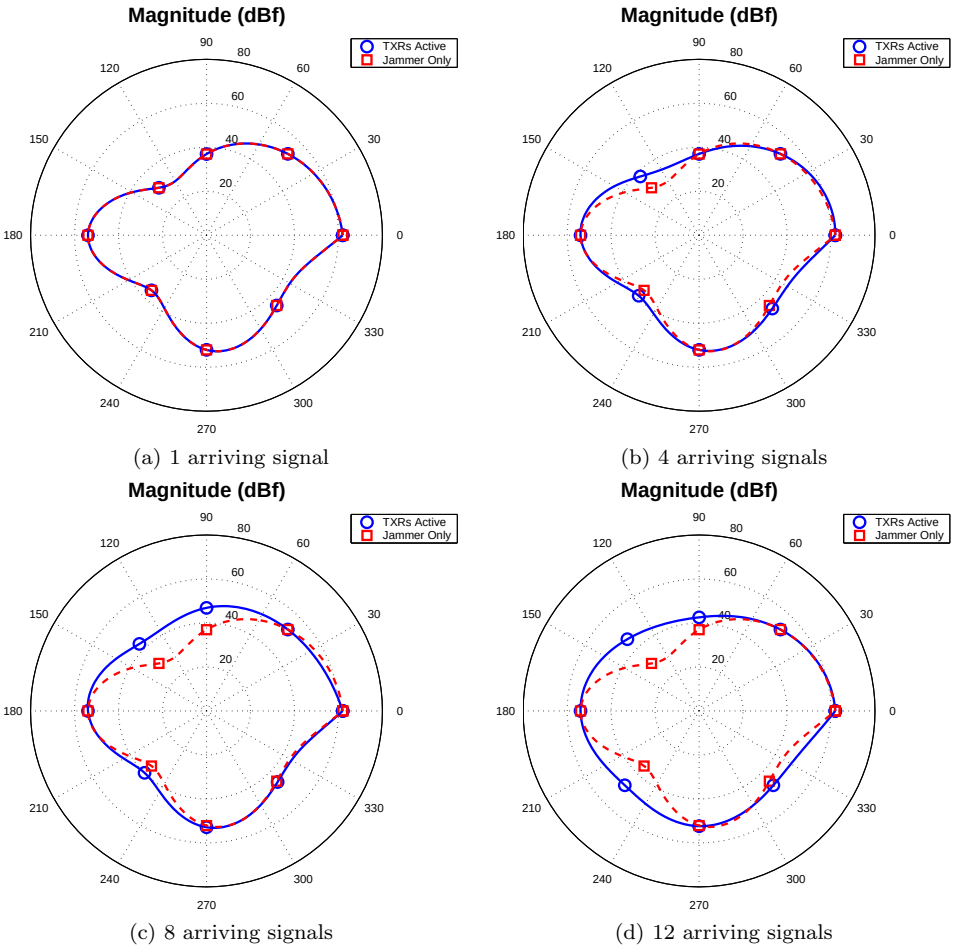

(a) 1 arriving signal

(b) 4 arriving signals

(c) 8 arriving signals

(d) 12 arriving signals

Figure 4.15: Beam-pattern magnitude distortion increases for a progressive number of arriving signals. Results from Matlab simulations using H ={1,4,8,12} arriving signals are illustrated.

The plots clearly demonstrate the escalating beam-pattern distortion that results from increasing the number of radio teaming transmitters. It is apparent from the results that increasing the size of the radio team increases the probability of having more multipath components reach the intended receiver, thus adding to the beam-pattern distortion at the receiver. The more distortion that the radio team can create at the receiver, the easier it is for the receiver to reliably recover the transmitted message.

We now continue with our radio teaming feasibility investigation by taking

a more in-depth look at scenarios involving an ample number of radio teaming transmitters. As such, four distinct simulations with H=12 arriving signals were conducted. The goal of these simulations was to investigate the necessity — or lack thereof — of synchronization between arriving signals at the receiver specifically with regard to phase, frequency, and angle (we assume a low enough data rate and proper guard time relative to delay spreads such that time delays do not play a significant role). Just as in the previous simulation, arriving SINRs were uniformly distributed between -20 and -10 dB.

We now describe the setup mathematically. Let $x(t)$ denote the baseband representation of the transmitted waveform at a given radio teaming transmitter. Given OOK modulation, we can represent this signal as,

$$x(t) = [\sum_{n=0}^{N-1} As_n\delta(t - nT)] * [u(t) - u(t - T)],$$

where s_n is one of the N binary symbols to be transmitted in the symbol period, T, A is the amplitude of the waveform, δ is the dirac-delta function, $u(t)$ is the unit-step function, and $*$ represents convolution. At the receiver, any given multipath arrival of the transmitted signal can be represented at baseband as,

$$y(t) = G(t)x(t - \tau)e^{j2\pi f(t)t}e^{j2\pi\theta(t)} + n(t),$$

where $G(t)$ is the gain of the channel, $f(t)$ is the frequency offset (e.g. associated with LO drifts or Doppler effects), $\theta(t)$ represents the phase shift, τ is the time delay, and $n(t)$ is additive channel noise. While each of these parameters are time-varying, we assume stationarity and justify it by selecting a small enough investigation interval. Since we are investigating scenarios where the SINR is exceptionally large, we relax the arriving radio-teaming signals to have a common, time-invariant gain, G, and also drop the noise component, $n(t)$. Finally, given a relatively low symbol period we can assume negligible time delays with respect

Table 4.1: Simulation Cases

Parameter	Case 1	Case 2	Case 3	Case 4
AOA	Jammer	Jammer	Random	Random
Phase/Freq	Sync	Async	Sync	Async

to the different signal arrivals, and represent each of the H arriving signals as:

$$y_h(t) = Gx(t)e^{j2\pi f_h t}e^{j2\pi\theta_h}.$$

Adding in our jammer, $b(t)$, we therefore model the arriving signals at the receiver as,

$$r(t) = \sum_{h=1}^{H} y_h(t) + b(t)$$

As stated, the simulations that we now discuss deal with four distinct cases where $H = 12$ arriving signals are modeled in order to investigate the effect of synchronization on the radio teaming protocol. The first simulation, Case 1, constrains all of the multipath arrivals of the radio teaming transmitters to coincide with that of the jammer. In addition, the arriving radio team transmissions are fully synchronized (i.e. frequency, and phase) at the receiver. This simulation represents the best possible scenario with regard to constructive interference at the receiver when all AOAs are constrained to be co-linear with the jammer.

The second simulation, Case 2, involves removing the synchronization of the arriving signals at the receiver. This is the real-world scenario where all arriving signals are co-linear. Here we have the frequency offsets independent and uniformly distributed between +/-500 Hz from the jammer's center frequency. We also distribute the phase offset uniformly between 0 and 2π. Note that the energy received for Case 2 is upper-bounded by Case 1 (i.e. $|r_2(t)|| \leq |r_1(t)|$). This is an important observation since it tells us that we should expect more beam-pattern distortion for Case 1 than Case 2.

Table 4.2: Simulation Parameters

SINR (dB)	AOA (Degrees)	Phase (Degrees)	Frequency Offset (Hz)
-18.75	90	104	-316
-13.18	90	359.00	177
-12.73	225	0	-441
-16.77	315	216.00	287
-10.71	135	58.00	164
-11.97	45	262.00	-299
-15.29	315	292.00	-374
-11.72	90	64.00	268
-18.70	180	119.00	-292
-12.44	180	44.00	-173
-18.43	225	165.00	-426
-15.58	180	216.00	-231

The third simulation, Case 3, adds AOA diversity while re-instating the synchronization of Case 1. For our model, we choose to uniformly distribute the AOAs over the eight jammer angles (i.e. 0, 45, 90, ... 315 degrees). This is the best case scenario where transmitters are all constructively interfering at the receiver, but from different angles of arrival.

The final simulation, Case 4, is the most realistic, as it combines AOA diversity with asynchronous signal arrivals. Table 4.1 provides a brief description of the four distinct cases.

Detailed results for the four cases are found in Figure 4.16, Figure 4.17, and Table 4.2. Figure 4.16 and Figure 4.17 depict the beam-pattern distortion graphically by plotting beam-pattern magnitudes in both Polar and Cartesian systems, in addition to illustrating the beam-pattern variance for the four simulated cases. Figure 4.16 illustrates Cases 1-2, where the AOA is constrained, while Figure 4.17 depicts Cases 3-4, where the AOA is random. Table 4.2 provides the characteristics of the arriving radio teaming signals related to Figure 4.16 and Figure 4.17. The full table represents Case 4, while all of the other cases are represented by constraining the appropriate column parameters.

We note that adding synchronization at the receiver only provides minimal

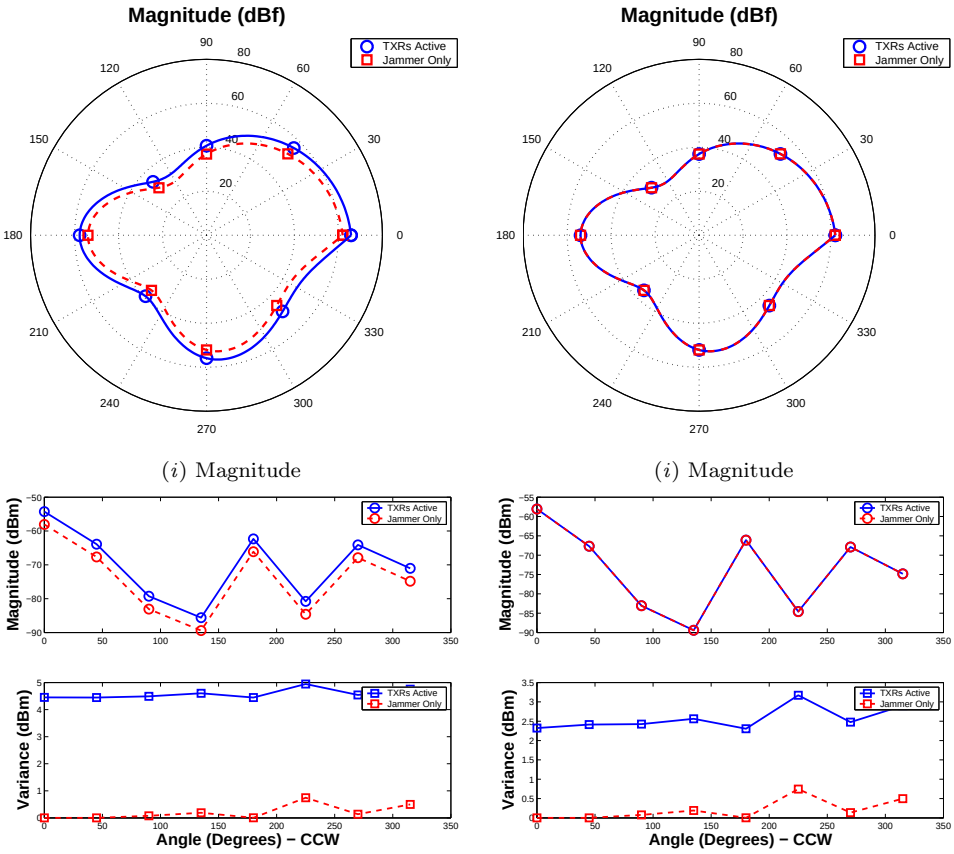

(*i*) Magnitude

(*ii*) Magnitude and Variance

(a) Case 1

(*i*) Magnitude

(*ii*) Magnitude and Variance

(b) Case 2

Figure 4.16: Beam-pattern statistics are depicted for Matlab simulation cases 1 and 2. (a) Case 1: AOA constrained and synchronized. (b) Case 2: AOA constrained, not synchronized.

signal gain (by way of constructive interference) and also limited distortion of the beam-pattern (in fact it is linear in Case 1). As expected, randomizing the AOAs provides much more distortion of the beam-pattern. What is exciting, is the fact that the beam-pattern exhibits ample distortion with random AOAs whether or not the arriving signals are synchronized. Because the random AOA, asynchronous simulation (i.e. Case 4) most represents our problem scenario in the real world, we conclude that our simulations show that the radio teaming protocol provides a viable communication mechanism when conventional communication is just not an option — even when in the presence of a strong jammer. We now

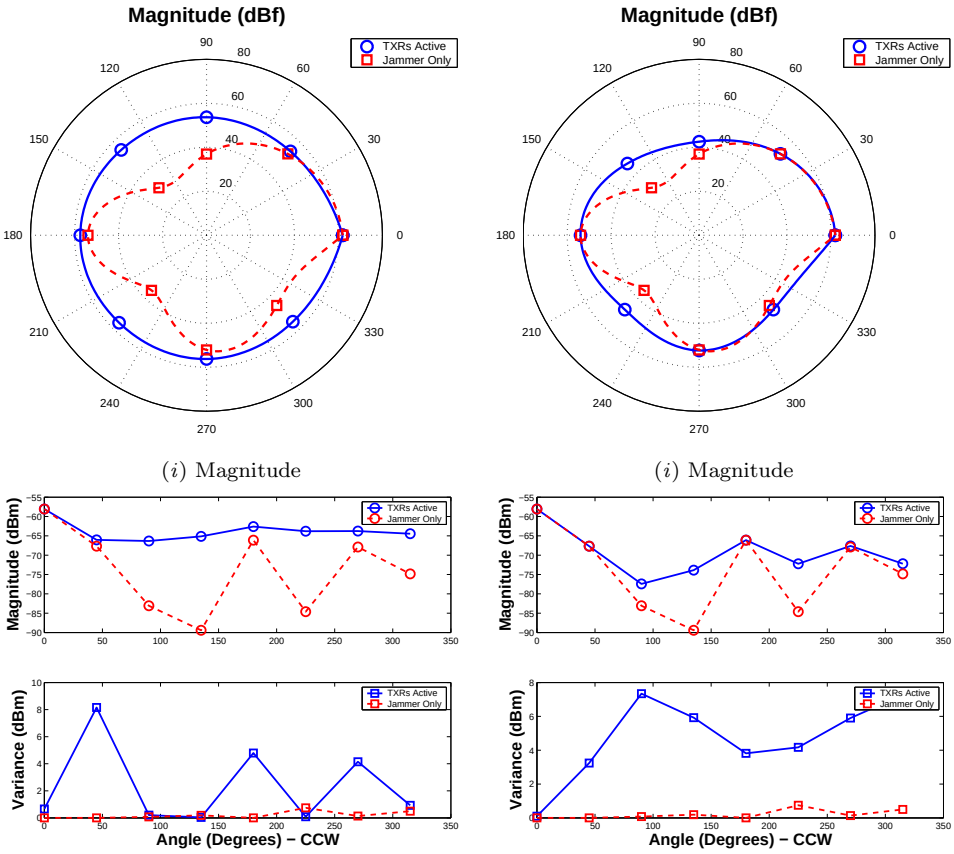

(*i*) Magnitude

(*ii*) Magnitude and Variance

(a) Case 3

(*i*) Magnitude

(*ii*) Magnitude and Variance

(b) Case 4

Figure 4.17: Beam-pattern statistics are depicted for Matlab simulation cases 3 and 4. (a) Case 3: Random AOA, synchronized. (b) Case 4: Random AOA, not synchronized. Case 4 is most representative of a real-world scenario; the promising results indicate real-world feasibility of the radio teaming protocol

continue by presenting real-world experimental results.

4.4.5 Experimental Results

Since the first step in the radio teaming protocol is time synchronization, an experiment was conducted in the WINLAB ORBIT Grid where 12 ORBIT nodes were used to synchronize with each other using Network Time Protocol (NTP). Each node was allowed to update its clock once every 64 seconds, and the experiment was conducted over the course of an hour. At any given time during

Table 4.3: NTP results. The average offset and its variance are calculated for synchronization tests conducted with 12 radio teaming nodes. The calculations are relative to Node 1.

Time (minutes)	Offset Mean (milliseconds)	Offset Variance (milliseconds)
0	320.598	449.183
5	374.129	452.190
10	-5.806	0.494
15	-7.614	0.429
20	-.364	0.334
25	7.340	0.200
30	1.605	0.154
35	8.599	0.114
40	4.360	0.113
45	2.456	0.101
50	5.560	0.119
55	3.282	0.115
60	-0.135	0.197

the experiment, timing offsets existed between Node 1 and its eleven other team members. Table 4.3 provides the measurements for these timing offsets relative to Node 1. Listed in this table is the average offset (in milliseconds) at every five minute interval, as well as the variance of the offsets. After only ten minutes (or about ten updates), we see that Node 1's eleven neighbors were on average 5.806 msec adrift, with a variance of only 494 microseconds.

Table 4.4 provides a more in-depth look at the synchronization level between the radio teaming nodes at precisely thirty minutes into the experiment. Displayed is the timing offset in milliseconds between each of the twelve nodes used in the experiment. Results corresponding to the row with Node 1's data was used to generate the data for the "30 minute" entry of Table 4.3. It is clear that this level of synchronization is more than enough to establish timeslotting in order to conduct further radio teaming experiments.

After demonstrating the feasibility of timeslotting our radio teaming transmitters, we conducted another experiment in an indoor laboratory facility to

Table 4.4: Synchronization offsets in milliseconds at 30 minutes into the 12 node NTP experiment.

Node	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
1	0	-5.0	-9.8	5.8	14.1	-3.4	11.2	-2.3	21.5	2.8	-23.7	6.4
2	5.0	0	-4.8	10.7	19.0	1.6	16.2	2.7	26.5	7.7	-18.7	11.4
3	9.8	4.8	0	15.5	23.8	6.4	21.0	7.5	31.3	12.5	-13.9	16.2
4	-5.8	-10.7	-15.5	0	8.3	-9.2	5.5	-8.0	15.8	-3.0	-29.4	0.6
5	-14.1	-19.0	-23.8	-8.3	0	-17.5	-2.8	-16.3	7.5	-11.3	-37.7	-7.7
6	3.4	-1.6	-6.4	9.2	17.5	0	14.7	1.2	25.0	6.2	-20.3	9.8
7	-11.2	-16.2	-21.0	-5.5	2.8	-14.7	0	-13.5	10.3	-8.5	-34.9	-4.9
8	2.3	-2.7	-7.5	8.0	16.3	-1.2	13.5	0	23.8	5.0	-21.4	8.6
9	-21.5	-26.5	-31.3	-15.8	-7.5	-25.0	-10.3	-23.8	0	-18.8	-45.2	-15.2
10	-2.8	-7.7	-12.5	3.0	11.3	-6.2	8.5	-5.0	18.8	0	-26.4	3.6
11	23.7	18.7	13.9	29.4	37.7	20.3	34.9	21.4	45.2	26.4	0	30.1
12	-6.4	-11.4	-16.2	-0.6	7.7	-9.8	4.9	-8.6	15.2	-3.6	-30.1	0

illustrate the feasibility of the radio teaming idea in a multipath-rich yet quasi-stationary environment (i.e. no mobile reflectors were present). A single CW jammer (generated by an Agilent 83620B swept signal generator) was used in the presence of two radio teaming transmitters for this experiment. A current SDR architecture was used to construct our radio teaming transmitters. Specifically, we used the GNU Radio/USRP SDR platform with two RFX-1800 daughterboards [6, 7]. The technical details of this SDR platform are described in detail in Section 1.5. The radio team is depicted in Figure 4.18, where the separation between the two transmitters is 1.5m, which is sufficiently far apart to provide independent multipath from the two transmitters relative to the receiver.

The modulation employed by the radio team was OOK with a symbol period of 8 seconds. Together, the team broadcasted a 20-bit message, $\mathbf{m} = 11101101000100110100$, to the receiver, which again was an 18 dBi MAXRAD directional panel antenna that rotated through 8 look angles separated by 45 degrees. The receiver gathered its beam-pattern statistics over eight uniform look-angles, and arriving signals from the radio team resulted in SINRs of approximately -20 dB.

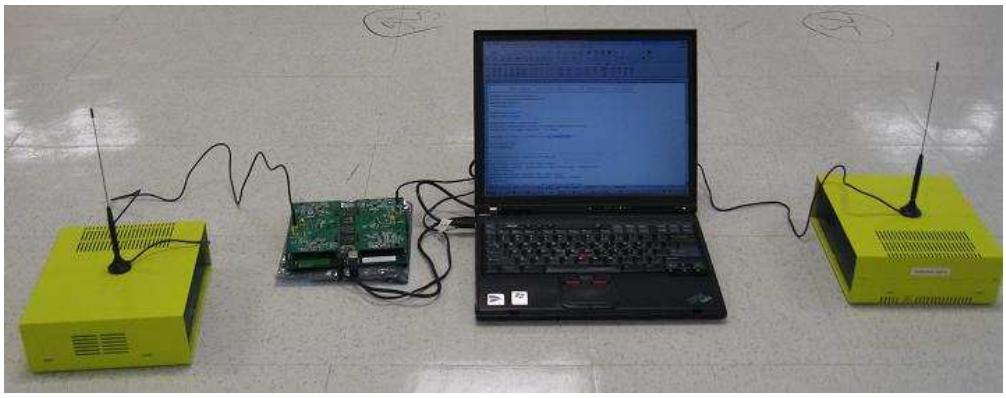

Figure 4.18: The radio team consisted of 2 OOK transmitters created using the USRP/GNU Radio SDR platform.

Figure 4.19 illustrates the beam-pattern response at the receiver for two symbols — one symbol corresponds to a 1 since it occurred during radio teaming activity, while the other symbol corresponds to a 0 since it resulted from jammer activity only. One immediately notices the "waterfilling" effect on the received power per angle statistics in addition to added variability to the variance of the statistics that occurs during radio teaming activity. It is also interesting to observe the similarity between the beam-pattern distortions of these experimental results and Case 4 of the simulations from Figure 4.17 (b). One may wonder why only 2 radio teaming transmitters appears to mimic the simulated case of 12 arriving signals, but the reader should keep in mind that 2 radio teaming transmitters does not equate to 2 arriving signals. In fact, given the multipath-rich experimental environment, the 2 transmitters were responsible for a larger number of multipath arrivals as seen by the receiver (our own experience suggests that there are 4 to 5 significant multipaths in this room for any transmission). It is quite interesting to note the behavior of the beam-pattern magnitude at 270 degrees in Figure 4.19 — the radio team appears to actually decrease the energy received when they are active. We attribute this phenomenon most likely to transient jammer behavior or channel variation, but it may also be a result of destructive

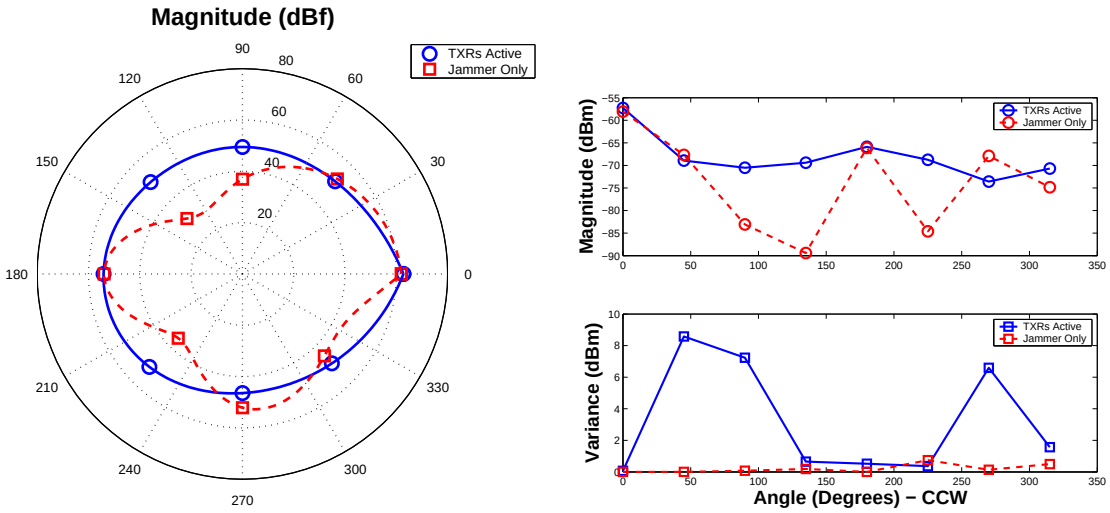

(a) Magnitude

(b) Magnitude and Variance

Figure 4.19: Experimental results are shown for a single symbol. (a) Polar beam-pattern statistics for a single symbol are plotted. (b) X-Y beam-pattern statistics for a single symbol are illustrated.

interference with the jammer in this symbol period at this specific angle.

Given ample beam-pattern distortion, is was then up to the receiver to decode the radio teaming message. Figure 4.20 depicts the symbol decoding process as the feature vector distances are plotted versus symbol number. The decision threshold, τ, was computed as a function of the mean and variance statistics during non-radio teaming activity. Despite the low signaling levels, the complete message, $\mathbf{m} = 11101101000100110100$, was successfully decoded by the receiver.

It is worthwhile to note that the effective bit-rate of this experiment was 0.125 bps at a low SINR of roughly -20dB. Although this might seem to be a very low data rate, we note that this is a limitation of the experimental setup that we employed, which was intended to support the feasibility of radio teaming. A significant improvement in communication rates is possible by employing more precise equipment, such as a phased array at the receiver, which would have vastly improved the effective look-angle switch-rate. Another avenue for improving the data rate would be to employ better timeslotting mechanisms to allow for tighter

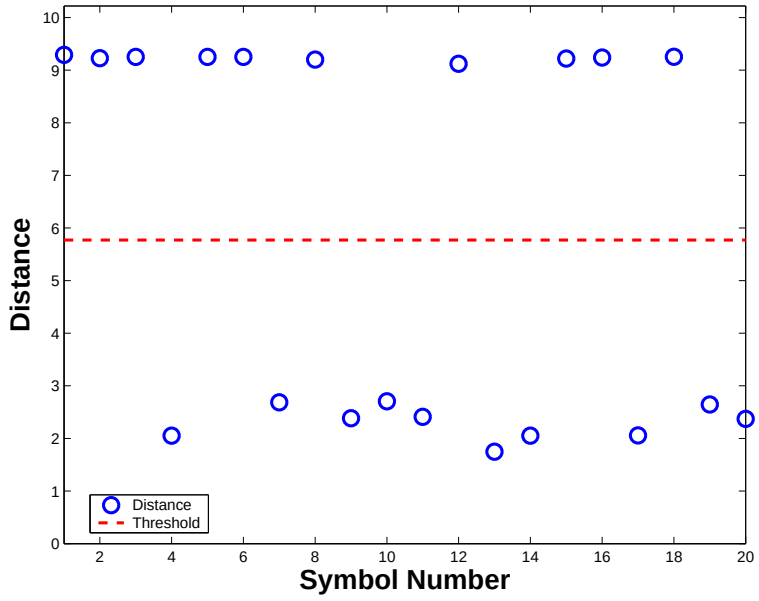

Figure 4.20: Experimental results are shown for all symbols. Symbol recovery is shown for all symbols, resulting in a perfectly recovered message of $\mathbf{m} = 11101101000100110100$.

alignment of transmitters.

4.5 Mitigation Summary

In this Chapter, we have explored various popular mitigation strategies and have proposed methodology to leverage machine learning techniques in adaptively selecting a path forward. We have also introduced a new, practical form of cooperative communications that is able to operate in harsh communication environments, such as an urban canyon or a scenario involving a jammer. Although most cooperative communications research has mainly focused on traditional MIMO and beamforming methodologies, where units share a high level of synchronization in time, frequency, and space, our radio teaming approach relaxes these stringent synchronization constraints. In this manner, radio teaming becomes a viable option for heterogeneous commodity devices in need of an alternative

communication scheme— in fact, radio teaming can work as an overlay for existing wireless networks, relying upon simple network-level time synchronization between transmitters that is possible through services, such as NTP. We have outlined the radio teaming protocol, providing a baseline variation that provides the equivalent of an on-off-keying (OOK) cooperative physical layer, and have outlined more general extensions that can allow for faster data rates (such as by employing OFDM). We have supported our radio teaming method through simulations, where we have demonstrated each of the building blocks of the radio teaming protocol in order to illustrate the factors related to the feasibility of radio teaming. We then validated radio teaming using real measurements in an indoor multipath environment involving an RF jammer and a receiver that employed a rotational directional antenna. Our experimental results showed that radio teaming can provide a reliable data link in a multipath-rich environment even in the presence of a strong adversary.

Chapter 5
Conclusions

5.1 Thesis Summary

It is clear that advancements in the realm of software defined radio promise near-term, large-scale usage of cognitive radio devices, where Dynamic Spectrum Access and cross-protocol wireless communication will be the norm. Since SDR platforms provide full access of the physical layer to its users and developers, it is essential to understand the benefits and drawbacks associated with PHY layer access in order to enhance operation in future wireless communication networks. In this thesis, we have investigated the usage of PHY layer access to: (1) enhance situational awareness, (2) act as an adversary, and (3) mitigate poor environments and adversarial conditions.

In Chapter 2, we began by investigating the use of physical layer information to discover services and identify devices. Specific techniques were developed to detect Bluetooth and WiFi services and devices by leveraging a PHY/MAC classification approach. Additionally, it was shown that differentiating between devices was plausible by way of channel estimate based feature vector correlations. We then took advantage of physical layer features from existing broadcast signals, such as cellular and broadcast television, to obtain coarse device mobility and location.

In Chapter 3, we started with a survey of cogent, protocol-specific attacks that relied upon PHY layer exploitation and were implemented with current SDR platforms. Attacks incorporated well known protocols from Bluetooth to GSM, to

lesser known yet equally pervasive tire pressure monitoring systems. Acknowledging that most 3G and 4G wireless standards and protocols incorporate some form of multi-input multi-output (MIMO) technology, we performed a thorough analysis on the physical layer weaknesses associated with two of the most popular MIMO techniques — singular value decomposition (SVD) based MIMO, and the Alamouti space-time block code. Our attacks focused on the efficacy of the channel estimates obtained during MIMO operation, as proper operation of MIMO systems demand accurate and timely knowledge of the wireless channel.

In Chapter 4, we provided an overview of traditional mitigation strategies such as transmitting with more power, changing modulation scheme, and changing carrier frequency. We then addressed the specific channel estimate vulnerabilities from Chapter 3 by proposing a general channel estimate authentication scheme. By embedding messages into a physical feature vector of the channel sounding pilot, we showed it was possible to authenticate channel state estimates and transmitters. Further, our techniques can be applied to any environment, operate over channel coherence times, and do not impinge upon protocol throughput. Finally, we introduced a radio teaming mitigation strategy to deal with richly scattering environments coupled with the presence of a strong jammer. By leveraging the multipath in favor of the radio team, messages can be conveyed to the receiver by modulating beam pattern distortion.

In each Chapter of this thesis, theoretical results were accompanied by simulations and real-world experimentation. We implemented our protocols and techniques using the USRP/GNU Radio SDR platform, thus illustrating implementation feasibility and the applicability of the proposed techniques.

5.2 Future Work

Leveraging the physical layer to obtain situational awareness, attack techniques, and mitigation strategies is a continually changing area of research. SDR platforms will continue to expand in capability, and new protocols will emerge. Further, future wireless networks will allow dynamically changing protocols to be used by CRs in order to optimize secondary usage of primary bands.

Immediate extensions exist to the device identification ideas presented in Section 2.3.2. Continuing the channel estimate based device differentiation, more experiments can be conducted using WiFi clients rather than just Access Points. This problem is inherently different and more difficult because these devices do not transmit specific bursts in periodic fashions and often may be mobile.

To extend the PLATEAU work in Section 2.4.1, investigations are warranted to incorporate channel model driven statistical clustering analysis to determine device location. Multipath delay profiles of general indoor and outdoor channels can be used to further leverage stationary broadcast services and perhaps even non-stationary transmitters. Statistically, the multipath delay profile of inside to inside, inside to outside/outside to inside, and outside to outside channels differ, and therefore should be detectable.

Continuing with situational awareness extensions, we consider that a major focus of this thesis is to obtain accurate situational awareness given the physical constraint of a partial bandwidth observation (Sections 2.3.1-2.3.2). Maintaining this constraint, partially observable modulation recognition algorithms should be investigated. A CR is never guaranteed to have a complete spectral snapshot of an existing device, and therefore it is advantageous to determine modulation even if missing some of the signal. Most modulation recognition algorithms focus on partial observability in time (and even space), but not frequency.

Relevant extensions of the radio teaming work from Section 4.4 include investigations and experimentation using a single radio team member, where rotation of the effective transmissions may effectively alter the beam-pattern distortion at the receiver. This can be accomplished by way of physically rotating the antennas, or via beamforming with a phased antenna array. Effective modulation may be done via the mechanisms discussed in Section 4.4.3 or by the using the rotational speed of the distortion to send the message.

Further work is also needed to address the vulnerabilities of tire pressure monitoring systems that were described in Chapter 3.2.6. A notable PHY layer based authentication approach is to incorporate expected Doppler profiles from the sensor transmissions to assist input validation.

The research ideas presented in this Section are by no means meant to provide a definitive summary of future physical layer exploitation research. Rather, each idea is presented in hopes of providing new avenues for exploration.

References

[1] Are You Obsessed With Your Cell Phone? http://cellphones.org/blog/are-you-obsessed-with-your-cell-phone/.

[2] A. Smith. Mobile Access 2010. July 2010.

[3] A. Goldsmith. *Wireless Communications*. Cambridge University Press, New York, NY, 2005.

[4] Federal Communications Commission. Unlicensed operation in the TV broadcast bands. *ET Docket No. 04-113*, May 2004.

[5] J. Mitola III. *Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio*. PhD thesis, Royal Institute of Technology (KTH), May 2000.

[6] USRP. http://www.ettus.com.

[7] GNU Radio. http://www.gnu.org/software/gnuradio/.

[8] The Rice WARP platform. http://warp.rice.edu/news.php.

[9] Texas instruments. http://focus.ti.com/lit/ml/sprt406/sprt406.pdf.

[10] B. Ackland (PI), M. Bushnell, D. Raychaudhuri, C. Rose, and T. Sizer. NeTs-ProWin: High Performance Cognitive Radio Platform with Integrated Physical and Network Layer Capabilities. *National Science Foundation NeTS-0435370*.

[11] T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., New York, NY, 1991.

[12] E. Hossain, D. Niyato, and Z. Han. *Dynamic Spectrum Access and Management in Cognitive Radio Networks*. Cambridge University Press, Cambridge, UK, 2009.

[13] T. C. Clancy III. *Dynamic Spectrum Access in Cognitive Radio Networks*. PhD thesis, University of Maryland, College Park, April 2006.

[14] T. Clancy and N. Goergen. Security in Cognitive Radio Networks: Threats and Mitigation. *Third International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CrownCom)*, May 2008.

[15] ORBIT. http://www.orbit-lab.org.

[16] X Jing and D. Raychaudhuri. Global Control Plane Architecture for Cognitive Radio Networks. *IEEE International Conference on Communications*, pages 6466–6470, June 2007.

[17] M. Shoemake. Wi-Fi (IEEE 802.11b) and Bluetooth: Coexistence Issues and Solutions for the 2.4 GHz ISM Band. *White Paper*.

[18] Wireless Personal Area Networks Working Group. IEEE Std 802.15.1-2005 Part 15.1: Wireless medium access control (MAC) and physical layer (PHY) specifications for wireless personal area networks (WPANs). *White Paper*, 2005.

[19] D. Spill and A. Bittau. BlueSniff: Eve meets Alice and Bluetooth. *In Proceedings of USENIX Workshop on Offensive Technologies (WOOT)*, 2007.

[20] H. K. Mardia. New techniques for the deinterleaving of repetitive sequences. *IEE PROCEEDINGS*, 136:149–154, August 1989.

[21] Inc. Agilent Technologies. MIMO Wireless LAN PHY Layer [RF] Operation and Measurement. September 2005.

[22] T. Yucek and H. Arslan. A Survey of Spectrum Sensing Algorithms for Cognitive Radio Applications. *IEEE Communications Surveys and Tutorial*, 11(1):116–130, 2009.

[23] R. Miller and W. Xu and P. Kamat and W. Trappe. Service discovery and device identification in cognitive radio networks. In *Proceedings of the IEEE Workshop on Networking Technologies for Software Defined Radio (SDR) Networks (Held in Conjunction with IEEE SECON)*, pages 40–47, 2007.

[24] D. K. P. Tan, H. Sun, Y. Lu, M. Lesturgie, and H. L. Chan. Passive radar using Global System for Mobile communication signal: theory, implementation and measurements. *IEE Proceedings on Radar, Sonar and Navigation*, 152:116–123, June 2005.

[25] G. Chandrasakearan, M. Ergin, R. Martin, M. Gruteser, J. Yang, and Y. Chen. DECODE: Detecting Co-Moving Wireless Devices. In *Proceedings of the 5th IEEE International Conference on Mobile Ad-hoc and Sensor Systems (IEEE MASS 2008)*, pages 315–320, September 2008.

[26] ETSI. Digital cellular telecommunications system (Phase 2+); Modulation in (GSM 05.04), February 2002.

[27] 802.11 Working Group. IEEE Standard 802.11: Wireless LAN Medium Access Control and Physical Layer Specifications, June 2007.

[28] Advanced Television Systems Committee. ATSC Digital Television Standard Part 2 - RF/Transmission System Characteristics (A/53, Part 2:2007), January 2007.

[29] A. K. Jain. *Fundamentals of Digital Image Processing*. Prentice Hall, Upper Saddle River, NJ, 1989.

[30] L. J. Heyer, S. Kruglyak, and S. Yooseph. Exploring expression data: identification and analysis of coexpressed genes. *Genome Res.*, 9:1106–1115, 1999.

[31] J. MacQueen. Some methods for classification and analysis of multivariate observations. *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297, 1967.

[32] http://acert.ir.bbn.com/projects/adroit/.

[33] H. Firooz. Implementation of Full-Bandwidth 802.11b Receiver. http://span.ece.utah.edu/pmwiki/pmwiki.php?n=Main.80211bReceiver.

[34] S. Knauth. Implementation of an IEEE 802.15.4 Transceiver with a Software-defined Radio setup. Technical report, Lucerne University of Applied Sciences, 2008.

[35] trifinite group. Car Whisperer. http://trifinite.org/trifinite_stuff_carwhisperer.html.

[36] J. Lackey. Group Special (Software) Mobile. http://thre.at.gsm, 2007.

[37] K. Nohl. Airprobe. https://svn.berlin.ccc.de/projects/airprobe/, 2010.

[38] H. Samra and D. A. Burgess. OpenBTS. http://openbts.sourceforge.net/.

[39] R. Ryan, Z. Anderson, and A. Chiesa. Anatomy of a Subway Hack. http://tech.mit.edu/V128/N30/subway/Defcon_Presentation.pdf.

[40] I. Rouf, R. Miller, H. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe, and I. Seskar. Security and Privacy Vulnerabilities of In-Car Wireless Networks: A Tire Pressure Monitoring System Case Study. *in Proceedings of the 19th USENIX Security Symposium*, August 2010.

[41] S. Govindjee. Firestone Tire Failure Analysis. 2001.

[42] Department of Transportation National Highway and Traffic Safety Administration. 49 CFR Parts 571 and 585 Federal Motor Vehicle Safety Standards; Tire Pressure Monitoring Systems; Controls and Displays; Final Rule. *http://www.tireindustry.org/pdf/TPMS_FinalRule_v3.pdf*.

[43] D. Gesbert, M. Shafi, D. Shiu, and P. Smith. From theory to practice: An overview of space-time coded MIMO wireless systems. *IEEE Journal on Selected Areas on Communications*, 21(3):281–302, April 2003.

[44] G. J. Foschini. Layered space-time architecture for wireless communication in a fading environment using multi-element antennas. *Bell-Labs Technical Journal*, pages 41–59, 1996.

[45] S. Ray, P. Moulin, and M. Medard. On Jamming in the Wideband Regime. *IEEE ISIT 2006*, pages 2574–2577, July 2006.

[46] S. Farahmand, A. Cano, and G. B. Giannakis. Anti-jam distributed MIMO decoding using wireless sensor networks. *IEEE ICASSP 2008*, pages 2257–2260, April 2008.

[47] E. A. Jorswieck, H. Boche, and M. Weckerle. Optimal transmitter and jamming strategies in Gaussian MIMO channels. *IEEE VTC 2005*, 2:978–982, June 2005.

[48] B. Park and T. F Wong. Optimal training sequence in MIMO systems with multiple interference sources. *IEEE GLOBECOM 2004*, 1:86–90, December 2004.

[49] M. H. Brady, M. Mohseni, and J. M. Cioffi. Spatially-Correlated Jamming in Gaussian Multiple Access and Broadcast Channels. *2006 40th Annual Conference on Information Sciences and Systems*, pages 1635–1639, March 2006.

[50] G. J. Foschini and M. J. Gans. On limits of wireless communication in a fading environment when using multiple antennas. *Wireless Personal Communications*, 6(3):311–335, March 1998.

[51] BLAST: Bell Labs Layered Space-Time - An Architecture for Realizing Very High Data Rates over Fading Wireless Channels. http://www1.bell-labs.com/project/blast/.

[52] I. E. Telatar. Capacity of multi-antenna gaussian channels. *European Transactions on Telecommunications*, 10(6):585–595, 1999.

[53] M. Chiani, M. Z. Win, and A. Zanella. On the capacity of spatially correlated mimo rayleigh-fading channels. *IEEE Transactions on Information Theory*, 49(10):2363–2371, October 2003.

[54] P. W. Wolniansky, G. J. Foschini, G. D. Golden, and R. A. Valenzuela. V-BLAST: An Architecture for Realizing Very High Data Rates Over the Rich-Scattering Wireless Channel.

[55] T. Mao and M. Motani. STBC-VBLAST for MIMO Wireless Communication Systems. *IEEE*, pages 2266–2270, 2005.

[56] R. A. Valenzuela, G. D. Golden, C. J. Foschini, and P. W. Wolniansky. Detection algorithm and initial laboratory results using V-BLAST space-time communication architecture. *Electronics Lett.*, 35(1), January 1999.

[57] S. H. Nam, O. Shin, and K. B. Lee. Transmit Power Allocation for a Modified V-BLAST System. *IEEE Transactions on Communications*, 52(7):1074–1079, July 2004.

[58] G. Lebrun, S. Spiteri, and M. Faulkner. Channel Estimation for an SVD-MIMO System. *EEE Communications Society*, pages 3025–3029, 2004.

[59] L. M. Garth, P. J. Smith, and M. Shafi. Exact Symbol Error Probabilities for SVD Transmission of BPSK Data over Fading Channels. *IEEE*, pages 2271–2275, 2005.

[60] G. Lebrun, J. Gao, and M. Faulkner. MIMO Transmission Over a Time-Varying Channel Using SVD. *IEEE Transactions on Wireless Communications*, 4(2):757–764, March 2005.

[61] M. Medard. The Effect upon Channel Capacity in Wireless Communications of Perfect and Imperfect Knowledge of the Channel. *IEEE Transactions on Information Theory*, 46(3):933–946, May 2000.

[62] T. Basar. The Gaussian Test Channel with an Intelligent Jammer. *IEEE Transactions on Information Theory*, 29(1):152–157, January 1983.

[63] W. Xu, W. Trappe, Y. Zhang, and T. Wood. The Feasibility of Launching and Detecting Jamming Attacks in Wireless Networks. *in Proceedings of ACM International Symposium on Mobile Ad Hoc Networking and Computing (Mobihoc)*, pages 46–57, 2005.

[64] G. Noubir and G. Lin. Low-power DoS attacks in data wireless lans and countermeasures. *SIGMOBILE Mobile Computing and Communications Review*, pages 29–30, 2003.

[65] Y. Hu, A. Perrig, and D. B. Johnson. Wormhole Attacks in Wireless Networks. *IEEE Journal on Selected Areas in Communications*, 24(2):370–380, 2006.

[66] Z. Li, W. Xu, R. Miller, and W. Trappe. Securing wireless systems via lower layer enforcements. In *WiSe '06: Proceedings of the 5th ACM workshop on Wireless security*, pages 33–42, New York, NY, USA, 2006. ACM Press.

[67] C. Oestges and B. Clerckx. *MIMO Wireless Communications*. Academic Press, Oxford, UK, 2007.

[68] Introduction To MIMO Systems. http://www.rohde-schwarz.com.

[69] 802.16 Working Group. IEEE Standard 806.16e-2005: Air Interface for Fixed and Mobile Broadband Wireless Access Systems, February 2006.

[70] G. W. Stewart. Perturbation Theory for the Singular Value Decomposition. September 1990.

[71] A. T. James. Distributions of matrix variates and latent roots derived from normal samples. *Ann. Math. Statist.*, 35(1):475–501, 1964.

[72] H. Weyl. Das asymptotische Verteilungsgestez der Eigenwert linearer partieller Differentialgleichungen (mit einer Anwendung auf der Theorie der Hohlraumstrahlung). *Mathematische Annalen*, 71:441–479, 1912.

[73] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, MD, 1996.

[74] L. Mirksy. Symmetric Gage Functions and Unitarily Invariant Norms. *Quarterly Journal of Mathematics*, 11:50–59, 1960.

[75] S. M. Alamouti. A simple transmit diversity technique for wireless communications. *IEEE Journal on Selected Areas on Communications*, 16(8):1451–1458, October 1998.

[76] N. Venkatesh. Wireless Handheld Devices - The 802.11n Advantage. *Mobile Handset DesignLine*, July 2008.

[77] P. A. Bello. Characterization of Randomly Time-Variant Linear Channels. *IEEE Transactions on Communications Systems*, CS-11(4):360–393, December 1963.

[78] W. Xu, K. Ma, W. Trappe, and Y. Zhang. Jamming Sensor Networks: Attack and Defense Strategies. *IEEE Networks Special Issue on Sensor Networks*, 20(3):41–47, June 2006.

[79] W. Xu, W. Trappe, and Y. Zhang. Channel Surfing: Defending Wireless Sensor Networks from Jamming and Interference. *Proceedings of the 6th International Conference on Information Processing in Sensor Networks*, pages 499–508, 2007.

[80] J.L. Dalmau-Royo, J.A. Delgado-Penin, J. Serrat-Fernandez, and R. Valle-Alarcon. An in-band frequency agility modem implemented with a dsp device. In *Military Communications Conference, 1989. MILCOM '89. Conference Record. Bridging the Gap. Interoperability, Survivability, Security., 1989 IEEE*, pages 80 –84 vol.1, October 1989.

[81] A. Misra, V. Krishnamurthy, and R. Schober. Stochastic learning algorithms for adaptive modulation. In *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, volume 4, page IV, May 2006.

[82] Rakesh Rajbanshi, Qi Chen, Alexander M. Wyglinski, Gary J. Minden, and Joseph B. Evans. Quantitative comparison of agile modulation techniques for cognitive radio transceivers. In *Consumer Communications and Networking Conference, 2007. CCNC 2007. 4th IEEE*, pages 1144 –1148, 2007.

[83] W. Xu, W. Trappe, and Y. Zhang. Anti-Jamming Timing Channels for Wireless Networks. In *ACM Conference on Wireless Security (WiSec)*, pages 203–213, 2008.

[84] R. Miller and W. Trappe. Subverting MIMO wireless systems by jamming the channel estimation procedure. *Proceedings of the third ACM conference on Wireless network security*, pages 19–24, March 2010.

[85] P. Yu, J. Baras, and B. Sadler. An Implementation of Physical Layer Authentication Using Software Radios. *Report No. ARL-TR-4888*, July 2009.

[86] N. Goergen, W. S. Lin, K. J. R. Liu, and T. C. Clancy. Authenticating MIMO Transmissions Using Channel-Like Fingerprinting. *IEEE Global Communications Conference (GLOBECOM)*, December 2010.

[87] L. Xiao, L.J. Greenstein, N. Mandayam, and W. Trappe. Using the Physical Layer for Wireless Authentication in Time-Variant Channels. *IEEE Transactions on Wireless Communications*, pages 2571–2579, July 2008.

[88] L. Xiao, L. J. Greenstein, N. B. Mandayam, and W. Trappe. A Physical-Layer Technique to Enhance Authentication for Mobile Terminals. *IEEE International Conference on Communications (ICC)*, pages 1520–1524, May 2008.

[89] L. Xiao, L. J. Greenstein, N. B. Mandayam, and W. Trappe. MIMO-Assisted Channel-Based Authentication in Wireless Networks. *IEEE Conference on Information Sciences and Systems (CISS)*, pages 642–646, March 2008.

[90] ITU-R Recommendation M.1225. Guidelines for evaluation of radio transmission technologies for IMT-2000, 1997.

[91] V. Erceg, et al. Channel Models for Fixed Wireless Applications. *IEEE 802.16.3c-01/29r4*, July 2001.

[92] ETSI. Digital cellular telecommunications system (Phase 2+); Multiplexing and multiple access on the radio path (GSM 05.02), 1999.

[93] W. C. Jakes Jr. *Microwave Mobile Communications*. Wiley-IEEE Press, Piscataway, NJ, 1994.

[94] M. Hsieh and C. Wei. Channel Estimation for OFDM Systems Based on Comb-Type Pilot Arrangement in Frequency Selective Fading Channels. *IEEE Transactions on Consumer Electronics*, pages 217–225, January 1998.

[95] R. Tesi, M. Hamalainen, and J. Iinatti. Channel Estimation Algorithms Comparison for Multiband-OFDM. *The 17th Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, September 2006.

[96] W. Xu and W. Trappe and Y. Zhang. Channel surfing: Defending wireless sensor networks from jamming and interference. In *in Proceedings of the 6th International Conference on Information Processing in Sensor Networks (IPSN07)*, pages 499–508, 2007.

[97] M. Han, T. Yu, J. Kim, K. Kwak, S. Lee, S. Han, and D. Hong. OFDM Channel Estimation With Jammed Pilot Detector Under Narrow-Band Jamming. *IEEE Transactions on Vehicular Technology*, 57(3):1934–1939, May 2008.

[98] S. Mathur, W. Trappe, N. Mandayam, C. Ye, and A. Reznik. Radio-telepathy: extracting a secret key from an unauthenticated wireless channel. *Proceedings of the 14th ACM International conference on Mobile computing and networking*, pages 128–139, 2008.

[99] H. Iwai and H. Sasaoka. Secret information and sharing techniques based on radio wave propagation. *IEICE Transactions B (Japanese Edition)*, J90-B(9):770–783, September 2007.

[100] S. Jana, S.P. Nandha, M. Clark, S.K. Kasera, N. Patwari, and S. Krishnamurty. On the Effectiveness of Secret Key Extraction Using Wireless Signal Strength in Real Environments. *In the ACM Sigmobile International Conference on Mobile Computing and Networking (MOBICOM)*, September 2009.

[101] N. Patwari, J. Croft, S. Jana, and S. K. Kasera. High-Rate Uncorrelated Bit Extraction for Shared Secret Key Generation from Channel Measurements. *IEEE Transactions on Mobile Computing*, 9(1):17–30, January 2010.

[102] R. Liu and W. Trappe. *Securing Wireless Communications at the Physical Layer*. Springer, New York, NY, 2010.

[103] Wireless Personal Area Networks Working Group. IEEE Standard 802.15.1-2005 Part 15.1: Wireless medium access control (MAC) and physical layer (PHY) specifications for wireless personal area networks (WPANS), 2005.

[104] T. S. Rappaport. *Wireless Communications: Principles and practice*. Prentice Hall, Upper Saddle River, NJ, 2002.

[105] S. Y. Seidel, T. S. Rappaport, S. Jain, M. L. Lord, and R. Singh. Path Loss, Scattering, and Multipath Delay Statistics in Four European Cities for Digital Cellular and Microcellular Radiotelephone. *IEEE Transactions on Vehicular Technology*, 40(4):721–730, November 1991.

[106] S. J. Orfanidis. *Introduction to Signal Processing*. Prentice Hall, Upper Saddle River, NJ, 1996.

[107] Y. Tsai and C. Rose. MIMO power strategies for limited transmitter CSI. *IEEE Conference on Information Sciences and Systems*, March 2010.

[108] J. G. Proakis. *Digital Communications.* McGraw Hill, New York, NY, 2001.

[109] Introduction to mimo systems. http://www.rohde-schwarz.com.

[110] A. Nosratinia, T. E. Hunter, and A. Hedayat. Cooperative Communication In Wireless Networks. *IEEE Communications Magazine*, October 2004.

[111] A. Scaglione, D. L. Goeckel, and J. N. Laneman. Cooperative Communication in Mobile Ad-Hoc Networks. *IEEE Signal Processing Magazine*, September 2006.

[112] Y. Hong, W. Huang, F. Chiu, and J. Kuo. Cooperative Communication in Resource Constrained Wireless Networks. *IEEE Signal Processing Magazine*, May 2007.

[113] B. J. Sepko and W. Lee. A Study on Effect of Interference in Cooperative Communication. *IEEE, Radio and Wireless Symposium*, pages 651–654, January 2008.

[114] S. Katti, S. Gollakota, and D. Katabi. Embracing Wireless Interference: Analog Network Coding. In *SIGCOM '07*, 2007.

[115] A. Sendonaris, E. Erkip, and B. Aazhang. User Cooperation Diversity - Part 1 System Description. *IEEE Transactions on Communications*, 51(11), November 2003.

[116] P. Tan, M. Steinbach, and V. Kumar. *Introduction To Data Mining.* Pearson Addison Wesley, Boston, 2006.

[117] M. Thottan and C. Ji. Anomaly Detection in IP Networks. *IEEE Transactions on Signal Processing*, 51(8), August 2003.

[118] M. Chugh, D. Bhatia, and P. T. Balsara. Design and Implementation of configurable W-CDMA Rake Receiver Architectures on FPGA. *Proceedings of IPDPS05*, April 2005.

[119] D. Deeths and G. Brunette. Using NTP to Control and Synchonize System Clocks. *Sun MicroSystems Online Blueprint*, July 2001.

[120] S. E. Butner and S. Vahey. Nanosecond-scale Event Synchronization over Local-area Networks. *IEEE Proceedings of Local Computer Networks*, (261-269), November 2002.

[121] A. J. Coulson. Bit Error Rate Performance of BPSK Modulated OFDM Synchronized using a Pilot Signal. *IEEE Symposium on Personal, Indoor and Mobile Radio Communications*, 2(F86-F89), October 2001.

[122] M. Ott, I. Seskar, R. Siraccusa, and M. Singh. ORBIT Testbed Software Architecture: Supporting Experiments as a Service. 2005.

Curriculum Vitae

Robert D. Miller

Education

2002-2011 Ph.D., Electrical and Computer Engineering, Rutgers University

2000-2001 M.S., Electrical Engineering, Princeton University

1996-2000 B.S. Computer Engineering, Boston University

Employment

2010 Lead Associate, Booz Allen Hamilton, Red Bank, NJ

2006-2010 Associate, Booz Allen Hamilton, Eatontown, NJ

2004-2006 Senior Consultant, Booz Allen Hamilton, Eatontown, NJ

2002-2004 Consultant, Booz Allen Hamilton, Eatontown, NJ

2001-2002 Software Engineer, Aware, Inc., Bedford, MA

2000 Computer Engineer, Stanford Linear Accelerator Center, Menlo Park, CA

Publications

2010 R. Miller and W. Trappe, "Physical Layer Techniques for Advanced Situational Awareness", ICASSP, 2010.

I. Rouf, R. Miller, H. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe, and I. Seskar, "Security and Privacy Vulnerabilities of In-Car Wireless Networks: A Tire Pressure Monitoring System Case Study", USENIX, 2010.

R. Miller and W. Trappe, "Subverting MIMO Wireless Systems by Jamming the Channel Estimation Procedure", WISEC, 2010.

2009 W. Trappe, N. Mandayam, M. Littman, R. Howard, D. Raychaudhuri, C. Rose, I. Seskar, and R. Miller. "Exploiting Cognitive Radio Technologies to Assure Communication Superiority in Tactical Networked Scenarios." DARPA White Paper - DARPA SN-09-60 Machine Learning for Behavioral Control of Cognitive Radios (ML BCCR). Fort Monmouth, NJ, September, 2009.

R. Miller, "Fundamentals of Radar Signal Processing (Richards, M.A.; 2005) [Book review]" IEEE Signal Processing Magazine. May 2009, vol. 26, no. 3, pp. 100-101

2008 R. Miller, S. Jain, and W. Trappe, "Radio Teaming: Establishing Communication When Communication is Not Possible", IEEE MASS, 2008.

2007 R. Miller, W. Xu, P. Kamat, and W. Trappe, "Service Discovery and Device Identification in Cognitive Radio Networks," in Proceedings of the IEEE Workshop on Networking Technologies for Software Defined Radio (SDR) Networks (Held in Conjunction with IEEE SECON) 2007.

2006 Z. Li, W. Xu, R. Miller and W. Trappe,"Securing Wireless Systems via Lower Layer Enforcements", in Proceedings of the 2006 ACM workshop on Wireless security(WiSe), pg. 33-42, 2006.